

Fat Segmentation in Abdominal MR-scans

Thomas Hammershaimb Mosbech

Kongens Lyngby 2008

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

Abstract

This thesis describes a method for automatically segmenting abdominal adipose tissue from 3-dimensional magnetic resonance images. The segmentation distinguishes between three types of adipose tissue; visceral adipose tissue, deep subcutaneous adipose tissue, and superficial subcutaneous adipose tissue.

Prior to the segmentation, the image data is preprocessed to remove within-class image intensity inhomogeneities caused by the so-called bias field effect. The field is sampled as two classes of intensity points and the effect is estimated using an extension of thin plate splines.

The adipose tissue is labelled across the abdomen by unsupervised classification using fuzzy c-means clustering and locally determined thresholds.

The abdomen boundary is segmented, and the visceral adipose tissue is separated from the subcutaneous adipose tissue by means of active contours; incorporating intensity information derived through the unsupervised classification.

The subcutaneous adipose tissue layer is subdivided into a deep and superficial part by dynamic programming and a polar transformation of the image data.

In the absence of ground truth segmentations, the results are subject to a visual validation; good results are obtained across the broad spectrum of images present in the data set.

Resumé

Dette eksamensprojekt omhandler en metode til automatisk segmentering af fedtvæv i maveregionen ud fra tredimensionelle magnet resonans scanninger. Segmenteringen skelner mellem tre typer fedtvæv: Visceralt fedtvæv, dybt subkutant fedtvæv og overfladisk subkutant fedtvæv.

Inden segmenteringen gennemgår billederne en forbehandling for at fjerne uensartethed i intensiteterne inden for vævstyper. Dette felt af skævhed indsamles som to klasser af punkter med tilhørende billedintensiteter. Effekten af feltet estimeres efterfølgende ved brug af en udvidelse de såkaldte 'thin plate splines'.

Fedtvævet i maveregionen identificeres med en automatisk klassificering ved hjælp af såkaldt 'fuzzy c-means clustering' og lokalt bestemte tærskelværdier.

Den ydre kant af maveregionen, samt opdelingen mellem visceralt og subkutant fedtvæv, findes med en deformerbar model, der inkluderer information om billedintensiteterne udledt i den automatiske klassificering.

Det subkutane fedtlag opdeles i dybt fedtvæv og overfladisk fedtvæv ved anvendelse af dynamisk programmering og en polær transformation af billedet.

Grundet manglende reference-resultater foregår valideringen af resultaterne visuelt. Der opnås gode resultater for datasættet, der indeholder meget varierende billeder.

Preface

This thesis was prepared at the Department of Informatics Mathematical Modelling (IMM), the Technical University of Denmark in partial fulfillment of the requirements for acquiring the Master of Science degree in engineering, M.Sc.Eng. The thesis was supervised by Professor Rasmus Larsen, IMM. The work was carried out over a period of 5 months, corresponding to 30 ECTS credits.

Kongens. Lyngby, February 2008

Thomas Hammershaimb Mosbech

Acknowledgements

I would like to thank Kasper Pilgaard, MD, Steno Diabetes Center, for his assistance in evaluating the method and for giving valuable feedback throughout the project period.

I thank Peter Stanley Jørgensen, IMM, for lending me the source code from his master's thesis project.

A very special thank you goes to my two friends Casper Willestofte Berg and Hans Gregers Petersen with whom I have been sharing the office through the 5 months of the project period. Thank you for some very interesting conversations about football – providing nice diversions from the abdominal fat. I would also like to thank my friend Peter Eriksen for dropping by the office to be the fourth player in our exciting table football matches in the basement.

Finally I would like to thank Tove Pedersen, Gregers, Casper and my father Holger Mosbech for valued assistance in proof reading the thesis. John also deserves to be mentioned here for his many enthusiastic visits to our office throughout the project period.

Contents

Abstract	i
Resumé	iii
Preface	v
Acknowledgements	vii
1 Introduction	1
1.1 Background	1
1.2 The Task	2
1.3 Thesis Overview	4
2 Data	5
2.1 The Image Series	5
2.2 Image Variation	8

3	Preprocessing	11
3.1	Rectangular Region of Interest	11
3.2	Bias Field Correction	17
3.3	Results	49
4	Identifying Adipose Tissue	53
4.1	Fuzzy C-Means Clustering	54
4.2	Classification	64
4.3	Results	70
5	Identifying Regions	75
5.1	Active Contours	76
5.2	Dynamic Programming	90
5.3	Results	98
6	Combining Labels and Boundaries	103
6.1	Binary Masks and Set Operations	103
6.2	Results	108
7	Final Results	111
8	Conclusion	113
8.1	Future Work	114
A	Volume Results	115

CONTENTS

xi

B Segmentation Results

117

C Software Overview

139

D Data CD

141

Introduction

This thesis has been done in collaboration with Ph.D. student Kasper Pilgaard, MD, from Steno Diabetes Center, Copenhagen. In this chapter a short introduction to his medical research project is presented to provide background and perspective regarding the work of the thesis. After this, the project goal and task are introduced, and finally the chapter is concluded by a short overview of the thesis structure.

1.1 Background

Kasper Pilgaard is working on a research project entitled '*Influence of Fetal Growth Velocity on Metabolism in Young Adult Twins and Singletons*'. The study investigates the association between low birth weight and type 2 diabetes and associated defects of metabolism including insulin resistance (IR).

The association between low birth weight and type 2 diabetes is well established. It has been proposed, that this association is a result of a common genotype leading to both IR and/or lower insulin secretion and reduced birth weight. One of the primary objectives of the study is to distinguish between a genetic and non-genetic origin of metabolic changes predisposing type 2 diabetes.

The gold standard for measuring IR is the hyperinsulinemic euglycemic clamp, but this method is invasive, expensive and labor intensive – therefore impractical for use in clinical practice. Obesity, particularly abdominal obesity, is associated with an increased risk of IR and type 2 diabetes. In the study it is investigated whether an association between low birth weight and abdominal obesity storage can explain the negative correlation between birth weight and IR.

For measuring abdominal obesity, a relevant approach is to distinguish between subcutaneous adipose tissue (SAT) and visceral adipose tissue (VAT), since a strong association between the quantity of VAT and IR is known. SAT is defined as the layer of adipose tissue located just below the skin around the outer rim of the abdomen. VAT is located underneath the SAT layer – occurring as cushions and flakes around the organs.

Methods for measuring adipose tissue includes calculation of body mass index (BMI), waist-to-hip ratio (WHR), dual energy X-ray absorptiometry (DXA), computed tomography (CT) and magnetic resonance imaging (MRI). BMI provides a simple numeric measure, but it is very crude and does not provide any information on the distribution of the adipose tissue. WHR is also a crude measure but gives some information on distribution – but not the particular relation between SAT and VAT. Abdominal DXA scans provide a more accurate quantity measure – but not information on the distribution in the region. From CT it is possible to retrieve accurate information on both quantity and distribution by means of good visualisations, but due to radiation issues related to the image acquisition, limited use is generally advised. MRI does not pose the same kind of radiation issues while still providing useful visualisations.

A drawback of MRI compared to CT is the lack of standard intensity units. CT images are measured in Hounsfield units providing a known correspondence between tissue types and image intensity. This is not the case for MRI, and therefore manual adipose tissue classification for assessing quantity and distribution can be a demanding task requiring a large amount of interaction from trained technicians.

1.2 The Task

The goal of the work described in this thesis is to develop a method to automatically determine the quantity and distribution of abdominal adipose tissue from T1-weighted 3-dimensional MRI data. The results should be presented as percentages of the total abdomen volume.

T1-weighted MRI is a suitable choice for the task, as adipose tissue gives rise to high image intensities for this modality. The quantity is sought to be determined from image intensities, and the distribution is assessed by dividing the abdomen into anatomically defined regions – separating the adipose tissue into different types.

In [12] Kelly et al. show the relevance of considering a subdivision of the SAT layer. In the study it was proposed to assess SAT by partitioning it into the plane superficial to Scarpa’s fascia (superficial SAT) and that below the fascia (deep SAT). The fascia is a thin layer of connective tissue in the abdominal wall. The conclusion drawn was that the quantities of both deep SAT and VAT had strong independent associations to IR, whereas the superficial SAT had a very weak relation. This proposed strategy of distinguishing between three types of adipose tissue is adopted in the present work.

Figure 1.1 shows an example of a 2-dimensional MR scan of the abdominal region, manually partitioned to identify the regions containing the three types of adipose tissue considered; VAT, deep SAT and superficial SAT.




Figure 1.1: **(a)** A zoomed view of a 2-dimensional MR image from the data set available for this study – the adipose tissue appears as high image intensities. The image is slice 16 from patient 7. **(b)** A rough manual segmentation of the regions performed by the author with a consumer-level image editing software. The three colored areas illustrate the regions containing the three types adipose tissue considered in this work; red holds superficial SAT, blue holds deep SAT, and green contains VAT.

1.3 Thesis Overview

This section presents the overall structure and a brief overview of the chapters and appendices in the thesis.

The general structure chosen, is to evaluate the method by examples of intermediate results after presenting the theory of each step. These examples are chosen to illustrate the performance on images with different selected characteristics. These intermediate results are available for the entire data set on the CD-ROM enclosed in appendix D. Furthermore, chapters 3, 4, 5 and 6 are concluded by showing the results on a recurring set of randomly selected images.

- **Chapter 2** introduces the image data used in this project. The image variation both within and between patients is illustrated by examples.
- **Chapter 3** describes the two steps undergone to prepare the image data for abdominal adipose tissue segmentation; first the arms are removed, then the bias field effect is estimated and removed.
- **Chapter 4** presents the method used for labelling adipose tissue voxels in the preprocessed image data.
- **Chapter 5** covers the method applied for identification of the anatomically defined regions containing the different types of adipose tissue.
- **Chapter 6** describes how the tissue labels and identified regions are combined to form the final segmentation.
- **Chapter 7** presents a short evaluation of the final segmentation results.
- **Chapter 8** holds the conclusion.
- **Appendix A** contains a table with the quantities of the three types of adipose tissue segmented.
- **Appendix B** shows the final adipose tissue segmentation results by means of samples from all patients in the data set.
- **Appendix C** gives an overview of the structure of the software developed.
- **Appendix D** holds a CD-ROM with the image series, intermediate and final segmentation results for the entire data set.

This chapter introduces the image data used in this project. The data set consists of 3-dimensional T1-weighted MRI from 21 patients, acquired throughout the 5 months of the project period by the Department of Diagnostic Radiology, Glostrup Hospital. The medical study includes a higher number of patients, but at the project deadline, MRI data was available for 21 patients only. The patients are identified by the reference numbers used in the medical study. Section 2.1 introduces the image data structure, and section 2.2 shows examples of the variability present within the data set.

2.1 The Image Series

For each patient a series of 2-dimensional MR-scans has been acquired to form 3-dimensional image data. The in-plane resolution of each scan is 512×512 pixels of 0.90×0.90 mm, and the spacing between slices is 7 mm. The data is delivered as series of slices stored in the DICOM format, where the image data is accompanied by a header of meta data with information on the patient (e.g. date of birth, name, sex, weight) and the image (modality, position of the slice in the image series, settings used in the acquisition etc.). In this work the pixels of each image slice will be denoted *voxels*, since the context of the image

series allows them to be regarded as 3-dimensional *volume* elements instead of 2-dimensional *picture* elements.

By means of reference scans acquired prior to the image acquisition each image series is positioned to cover the volume delimited by the first and fourth lumbar vertebra of the patient. The spacing between slices is equal for all patients, so because of the variability in patient height, the number of slices forming the image series varies between 16 and 23 – the data set comprising a total of more than 400 slices. The series positioning is important in relation to usefulness of the segmentation results in the medical study, since the anatomically bounded unit enables comparison between patients. It also allows for incorporating other measurements performed on the same bounded volume unit – e.g. DXA scans for measuring the total quantity of adipose tissue.

As an example of the image data, a grid structure with an overview of all 16 slices from patient 7 is shown on figure 2.1. The example also illustrates an important property of the modality when segmenting adipose tissue: As noted in section 1.2, T1-weighted MRI data exhibits high intensities for adipose tissue, whereas air and tissue with high water content exhibit low intensity.

It should be noted, that the method used for showing images in this work scales the displayed intensities to fit the range of each individual image. Thus the displayed intensities should not be compared between image slices. In figure 2.1 intensities are displayed in a greyscale color map, where black corresponds to the lowest intensity and white corresponds to the highest.




Figure 2.1: Overview of the image series of patient 7; 16 slices forming the 3-dimensional image data. Moving from left to right, top to bottom; the top left is the highest slice i.e. the slice closest to the head of the patient (slice 1), the rightmost in the bottom row is the lowest slice of the series (slice 16).

2.2 Image Variation

The patients of the medical study are same-sex young monozygotic and dizygotic twin pairs and a control group of singletons. The patients are both male and female, and characterised as ranging from very lean to slightly obese.

Figure 2.2 shows two examples of the anatomical variation between slices within a patient. In comparison to the upper slices, the two lower slices exhibit more SAT in the posterior part of the abdomen, as these slices are located closer to the hips. The SAT layer appears as a continuous rim on most images, but some slices contain the patient’s umbilicus appearing as a low intensity *hole* in the anterior part (figure 2.2(b)). Furthermore the layer generally seems to be thinner in the area right below the spine and on both sides in the anterior part – on lean patients this results in the SAT layer being almost invisible here. The slice shown in figure 2.2(a) holds very little VAT as the patient is very lean, and furthermore the left side of the image covers the liver. Figure 2.2(b) contains more, but the two slices in figure 2.2(c) and 2.2(d) contain even more VAT.

In this work, depending on the purpose of the visualisation, image slices are also shown using the *jet* color map to provide a better display of intensity details compared to the greyscale color map. As with the grayscale image visualisation, the displayed intensities are scaled to fit the intensity range of each individual image; dark blue corresponds to the lowest intensities and dark red to the highest – passing through cyan, yellow and orange in between (cf. the color bars).

Comparing the patients, it is evident that the quantity of VAT varies greatly, but also the contrast exhibited by the tissue is different. The VAT intensities in figure 2.2(b) seem far less well-defined compared to the intensities in both figure 2.2(c) and 2.2(d). A general observation made on the properties of the VAT is that the smaller flakes are less *well-defined* – showing less contrast to the surrounding non-adipose tissue – compared to the larger VAT cushions and the SAT layer.

It should also be noted, that the images of figure 2.1 contain the patients arms, whereas neither of the two displayed on figure 2.2 do, this variation is handled by a preprocessing step in section 3.1 as the adipose tissue segmentation only concerns the abdomen.

The image data features variation both in terms of the quantity and distribution of adipose tissue, but also with respect to image quality and the presence of artifacts. Generally a small amount of low intensity noise is present in the intensities just outside the abdomen, but on other slices the noise is more severe inside the abdomen as well. Figure 2.3 shows an example of such strong artifacts




Figure 2.2: (a) and (b) Top and bottom slices of a female patient. Slice 1 exhibits very little SAT and VAT, the left side of the image covers the liver resulting in low intensities. Slice 20 features more VAT, and since the slice is located closer to the hips, a thicker SAT layer is also present. (c) and (d) Top and bottom slice of a male patient. Both slice 1 and slice 20 exhibit a larger quantity of VAT with better contrast compared to (a) and (b).

in the SAT layer of a bottom slice. After consulting Kasper Pilgaard on this matter, it was decided to exclude such faulty looking slices from the image series, as they hold no valid information. These corrupted slices are all located in the end of the series – thus the exclusion does not ruin the 3-dimensional image structure. Since the medical study involves comparisons within each pair of twins, the same slice is also removed from the image series of the corresponding twin. The patients IDs are noted such that this can be taken into account if incorporating other measurements in the the medical study.




Figure 2.3: A bottom slices featuring image artifacts in the SAT layer; the slice is removed from both the patient and the corresponding twin.

Overall the variety of the image data is large, and thus the method developed for the adipose tissue segmentation, should be able to handle this.

A last note on the image data concerns the observed spatially varying inhomogeneity of image intensities, e.g the general intensity level appears to be higher in the right side of the images shown in figures 2.2(b) and 2.2(c) compared to the left side. A correction of this *bias field* to make the adipose tissue intensities more homogeneous across the images is described in section 3.2.

Preprocessing

Before adipose tissue segmentation, the image data requires two steps of preprocessing. The first step (section 3.1) is to identify a rectangular region of interest around the abdomen for each slice – removing the arms, if present. The second step (section 3.2) is a correction of the bias field – a spatial non-uniformity of same-tissue voxel intensities, causing problems for the intensity based tissue classification applied in chapter 4. Section 3.3 finishes the chapter with concluding remarks on the methods applied and a few randomly chosen example results of the preprocessing.

3.1 Rectangular Region of Interest

As noted in chapter 1, some slices contain both the abdomen and the patient's arms. The arms are of no interest in abdominal adipose tissue segmentation, so the first task of the preprocessing is to remove them from the image series. This is done by identifying a rectangular region of interest (ROI) only containing the abdomen. The method introduced below is performed on all slices one by one. Section 3.1.1 evaluates the method by means of result examples to illustrate the performance.

Both the boundary of the abdomen and the arms are characterised by a shift in intensity from a homogeneous area of low intensities (the surrounding air) to the high intensities (skin and fat). These shifts form the base of the rectangular region identification. To locate the shifts, sums of image intensities \mathbf{I}_{ij} are computed for all rows and columns; vectors \mathbf{R} and \mathbf{C} with elements:

$$R_i = \sum_j \mathbf{I}_{ij} \quad (3.1)$$

$$C_j = \sum_i \mathbf{I}_{ij} \quad (3.2)$$

Figure 3.1 shows an example of the sums computed for an image slice; the intensity shifts related to the boundary of the abdomen are clearly visible. Furthermore the column sums also exhibit shifts related to the arms. The image is used throughout the section.

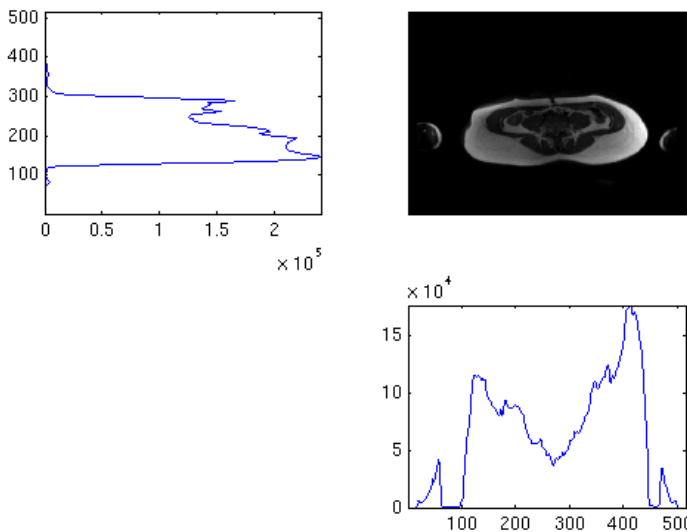


Figure 3.1: An example of row and column sums. The shifts related to the boundaries between background and the patient are clearly distinguishable in both directions. The image is slice 16 of patient 7.

The idea is to locate the intensity shifts by means of limits defined as percentages

of the maximum row and column sum of intensities:

$$R_{limit} = p \cdot \max(\mathbf{R}) \quad (3.3)$$

$$C_{limit} = p \cdot \max(\mathbf{C}) \quad (3.4)$$

where p is a percentage limit chosen to indicate the intensity shifts. If the percentage is set too low, the method could miss the gap of low intensity between the abdomen and the arms, but if too high the ROI could end up erroneously excluding parts of the abdomen. A value of $p = 0.05$ has been found to give good results for both the horizontal and vertical boundary identification.

First the vertical boundaries are located by regarding the vector of column sums \mathbf{C} . The vector is split in two equally sized halves. For each half the maximum value and limit is computed as in (3.4). The two boundaries are then located as the first column j with $C_j < C_{limit}$ – searching from the center and outwards to the left and right respectively. This method assumes that some extent of air is present between the abdomen and the arms enabling a separation by vertical boundaries. Figure 3.2 illustrates the method.

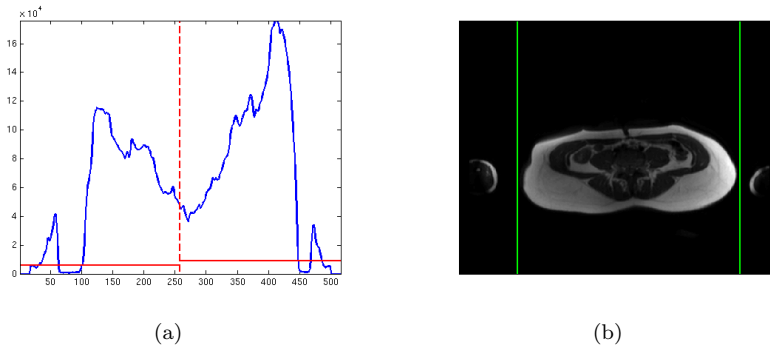


Figure 3.2: **(a)** The column sums with the limits for each half marked by the red lines – the vertical boundaries are located as the intersections closest to the center. **(b)** The image with the located vertical boundaries.

The same approach is used to determine the horizontal boundaries, only using the row sums \mathbf{R} searching each of the two halves for row i with $R_i < R_{limit}$. If the abdomen is elevated compared to the arms, the method would catch these high intensities and locate the lower boundary unnecessary far away from the abdomen. To prevent this, the row sums are only calculated for the area inside the vertical boundaries. Figure 3.3 illustrates the method.

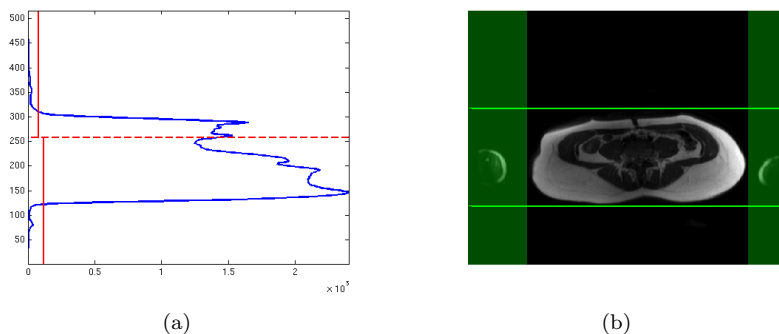


Figure 3.3: **(a)** The row sums with limits marked by the red lines – the two horizontal boundaries are located at the intersections with the row sum curve. **(b)** The image with the located horizontal boundaries. The green area indicates voxels outside the vertical boundaries, they are not included in the row sums.

The four boundaries are used to create a binary mask of same dimensions as the image – containing ones inside and zeros outside the identified rectangle. The mask is applied by means of an element-wise matrix multiplication with the image; all elements outside the rectangular region are set to 0 – approximating the low image intensity observed in the background area outside the abdomen (figure 3.4). The horizontal boundaries have no effect on removing the arms, but are applied to remove the small amount of image noise and artifacts in the air around the abdomen.

Another solution could have been to crop each slice by means of the boundaries instead of applying them as a binary mask. This solution has not been chosen, since it would have created image series with different sized slices, making the data less convenient to handle in the further processing. However the boundaries are used to provide zoomed views for better visualisation of results later. The next section shows some example results; slices before and after identifying and applying the rectangular ROI mask.




Figure 3.4: **(a)** The binary mask constructed from the located vertical and horizontal boundaries. **(b)** The image after the masking – the arms have successfully been removed.

3.1.1 Evaluation

Figure 3.5 shows an example of the method applied on an image with far less SAT than the previous example, the image is slice 1 of the same patient. The method performs equally well. Note the change in displayed image intensities caused by the change in intensity range. Because the arms contain some very high intensity voxels, the abdomen voxels appear brighter after the arms have been removed by the masking – the color bars of the two images show that the upper intensity limit is reduced from just above 3500 to around 1200.

Figure 3.6 shows an example of the method applied on a slice without arms present; the abdomen remains intact, and only very little change is visible in the intensities of the background area.

The data set holds one image series (patient 103) where some slices feature one arm and the abdomen located extremely close together. This causes the method to fail by including the arm in the rectangular ROIs found. Figure 3.7 shows an example of a slice leading to such an error. However this is only a problem for this single patient in the entire data set, and because the particular image data was received late in the project period, no automatic solution to the problem has been developed. After consulting Kasper Pilgaard on the matter, a decision has been made to exclude this image series from the data set, and remove the problematic arms manually at a later point before submitting the image series to further processing. A more detailed instruction to the technicians performing the scanning (and the patient) could eliminate the problem.




Figure 3.5: **(a)** The image before removing the arms. **(b)** The image after successfully removing the arms. Note the change in the image intensities displayed caused by the automatic intensity scaling. The image is slice 1 of patient 7




Figure 3.6: The method performed on a slice without arms present. **(a)** The image before identifying the rectangular ROI. **(b)** The image after applying the mask; very little change is observed as only low intensity background voxels are removed by the masking. The image is slice 1 of patient 21.




Figure 3.7: The method performed on a slice without a sufficient amount of space between the arm and the abdomen. **(a)** The image before attempting to identify the rectangular ROI to remove the arms. **(b)** The image after applying the masking technique – the arm in the right side of the image is not removed. The image is slice 3 of patient 103.

Apart from this single image series the method performs well on all image series. After this masking, the image data is ready for the second step of the preprocessing.

3.2 Bias Field Correction

This section describes the method applied to perform a correction of the intensities to remove the bias field effect. In section 3.2.1 the problem and the proposed solution are introduced. Sections 3.2.2-3.2.8 cover the estimation and subsequent removal of the effect. After the method has been presented, some remarks and examples of corrected images conclude the chapter in section 3.2.9.

3.2.1 Introduction

A known obstacle when working with MRI data is the presence of a *bias field*; a non-anatomic variability within same-tissue intensity values over the image domain. The effect is usually spatially smooth, and a level of the variation reaching up to 20% has been reported by Sled et al. in [13].

Apart from the slowly varying bias field another type of intensity non-uniformity is observed in the data; wide intensity peaks in the adipose tissue. These *bias peaks* occur in the left and right side of the slices, and only seem to occur for voxels of the adipose tissue class – magnifying the already high intensities.

Both the bias field and peaks are reported to be caused by imaging instrument calibration, patient anatomy and patient movement during the image acquisition period. The effect is therefore unique for each patient and scan session.

Figure 3.8(a) shows a slice with intensity peaks in both the left and right side. Figure 3.8(b) is a higher positioned slice from the same patient where far less SAT is present – here only the right peak is present and furthermore it exhibits smaller magnitude. Figure 3.8(c) shows a slice, where the bias field effect is clearly visible on both the SAT and VAT in the right side of the image. Figure 3.8(d) shows a slice, where the bias field effect is present on the outer part of the SAT layer more or less all the way around the abdomen.

Figure 3.9 shows the intensity histogram computed for all voxels located inside the rectangular ROIs (defined slice-by-slice in section 3.1). The image series is patient 16. The lower intensity bins exhibits a little of the background air voxels left after the rectangular region masking, but the main part of the histogram covers intensities from voxels of the abdomen. The range of the high intensity voxels of adipose tissue is stretched by the presence of the intensity peaks caused by the bias field.

These intensity inhomogeneities can often be overcome by human observers during manual segmentation tasks, but the effect can cause great difficulties for automatic intensity based tissue classification methods, where some extend of intra-class homogeneity is required. To provide ideal conditions for the adipose tissue classification, the histogram should exhibit a distinct peak for the adipose tissue intensities. This is not the case in this image data as the adipose tissue only gives rise to the flat bump to the right. As this is caused by the bias field effect, removing it from the image series is essential for a correct automatic classification of adipose tissue by means of image intensity.

In [8] Hou presents a commonly used model for describing the bias field effect for an acquired voxel intensity $y_{observed}$ as the multiplicative link:

$$y_{observed} = y_{true} \cdot y_{bias} + \xi \quad (3.5)$$

Where y_{true} is the true intensity, y_{bias} is the bias field effect, and ξ is a noise term. To simplify computations the noise term is often ignored, and the natural




Figure 3.8: Four slices illustrating examples of the bias field effect. **(a)** An example of an image with intensity peaks in both left and right side. **(b)** An example of an image only containing an intensity peak in the right side – this slice is from the same patient as **(a)**. **(c)** An example of an image where the bias field effect is clearly visible for voxels of VAT as well as SAT in the right side of the image. **(d)** An example slice where the bias field effect is more or less equally high on the entire outer part of the SAT layer.




Figure 3.9: Intensity histogram of voxels inside the identified rectangular ROIs of the whole image series of patient 16. The few high leftmost bins hold the low intensity of the background air. The tail on the right is caused by the intensity peaks in the adipose tissue. The peak just below 200 holds intensities of non-adipose tissue voxels.

logarithmic transform is applied:

$$\begin{aligned} \log(y_{observed}) &= \log(y_{true}) + \log(y_{bias}) \Leftrightarrow \\ \log(y_{true}) &= \log(y_{observed}) - \log(y_{bias}) \end{aligned} \quad (3.6)$$

This formulation enables a removal of the bias field effect from an image by subtraction – if the effect is known.

In [9] Ji et al. states that one way of obtaining an estimate of the effect is by means of extra series of scans of uniform phantom performed in connection with each scan session. This increases the total scanning time, requires additional hardware, and furthermore phantom measurements cannot account for the part of the effect related to anatomy and patient movement during the acquisition.

As noted in chapter 2, only T1-weighted MRI data is available for this work, thus it is not possible to make any predictions about the bias field. A method for estimating the effect retrospectively is needed in order to perform the correction for each individual image series by (3.6).

The technique chosen for removing the bias field effect, is to exploit the expected

slowly spatially varying behavior (except for the areas of the bias peaks). The effect is sought to be estimated by a smooth function with parameters obtained through regression. The regression is done on a set of spatially dense reference points with intensities assumed to hold information about the bias field effect.

Such a technique has successfully been applied on abdominal 3-dimensional MRI by Jørgensen in [10]; fitting a 3-dimensional smooth function to a set of automatically sampled points from adipose tissue voxels of the entire abdominal region covered by the image series. The assumption is, that the image intensities of adipose tissue voxels are the same across the image domain if not corrupted by the bias field – the source of the variance in the sampled intensities should only be the bias field.

Some patients in the data set of this work feature very little VAT. Therefore estimating the bias field effect from adipose tissue intensities alone could lead to large areas inside the abdomen with no points available for the estimation. The method proposed in the following extends the fitting of a smooth function to using two classes of reference points; voxels of adipose tissue and voxels corresponding to tissue with high water content. Thus assuming that the spatial variation of the bias field effect is the same for voxel intensities within any of the two classes.

3.2.2 Sampling Voxel Points

In T1-weighted MRI data adipose tissue gives rise to high image intensities while high water content and air give rise to low image intensities. These intensity properties, along with the assumption of a slowly varying bias field, form the basis for the automatic sampling of the two classes of voxel points described in the following section. The method is applied on each slice separately.

3.2.2.1 Methodology

In section 3.1 a rectangular mask was applied to the image to remove the arms – padding with zero intensity voxels. This still leaves some low intensity background voxels outside the abdomen containing no useful information for estimating the bias field. Thus the first task is to identify a rough estimate of the abdomen boundary – forming a ROI where the sampling of voxel points can be performed.

The shape of the abdomen boundary is generally smooth and continuous, and as noted in section 3.1, it is characterised by an abrupt change in intensity from the homogeneous low intensity of the background to the high intensity of the skin and SAT. These characteristics are used to find a rough outline estimate by means of an active contour approach. The theory of active contours is covered in section 5.1, and the method will therefore not be elaborated further here. Figure 3.10(a) shows an example of a boundary retrieved by the method. From the enclosed voxels a binary mask for the ROI can be generated (figure 3.10(b)). The image is slice 16 from patient 7, and will be used as an example to illustrate the point sampling technique throughout this section.






Figure 3.10: **(a)** The abdomen boundary found using active contours. **(b)** The ROI mask generated from the voxels enclosed by the boundary curve.

To make use of the relative intensity properties of the two classes, all local maxima and minima are identified inside the ROI. The two sets of points created are denoted I_{max} and I_{min} respectively – corresponding to classes of adipose tissue points and non-adipose tissue (high water content) points. Figure 3.11 shows the image marking all these local extrema inside the abdomen ROI – a zoom has been applied for for a better view of the details.

A large part of the local maxima are located on true adipose tissue and many of the local minima are located in areas with high water content – but clearly some points from both classes are erroneously positioned. This is especially the case for the points of I_{min} located on the SAT layer; no high water content or air is present, and the local minima here are solely due to the intensity variation



(a)



(b)

Figure 3.11: Zoomed views of the abdomen with points marking the local intensity extrema inside the ROI. **(a)** The red points mark all local maxima, I_{max} . **(b)** The green points mark all local minima, I_{min} .

within the adipose tissue. Because the ROI is only roughly estimated, a few points from each class are also placed outside the abdomen.

For the two classes of points to provide valid information on the bias field effect, they must be trimmed, such that only true points are left in both classes. The trimming seeks to make use of the assumption of a slow spatial variation of the bias field; within small areas of the image the effect of the bias field is close to constant. The solution proposed is to divide the image into subregions and only select points exhibiting extreme intensities compared to the the other points of the class inside the subregion.

The subdivision is performed by means of overlapping rectangles covering the ROI. First the smallest rectangle covering the whole ROI is identified from the binary mask. This rectangle is then subdivided into smaller rectangular regions; n_r subdivisions vertically and n_c subdivisions horizontally are created with overlaps of o_r and o_c voxels respectively. The principle of the subdivision is illustrated in figure 3.12.




Figure 3.12: Illustration of the subdivision scheme applied on the ROI from figure 3.10(b). The area inside the rectangle containing the ROI is displayed – showing only the three overlapping rectangles covering the top left corner. For illustrative purposes $n_r = 5$, $n_c = 6$ and $o_r = o_c = 15$ are chosen.

Ideally the number of subdivisions n_r and n_c should be chosen as high as possible, such that no subdivision can be performed positioning a rectangle to cover an area without at least one true voxel point of each class. E.g. if a subdivision only contains voxels of the SAT layer, the local minima are caused by within-class intensity variation, and do not belong to high water content tissue as intended.

The use of many small rectangles is not possible though, as seen from figure

3.11(a), a trimming with too small rectangles on slices containing a large quantity of SAT could lead to falsely positioned points of I_{min} . The same issue is present for I_{max} , as some image series in the data set contain slices with large areas solely containing non-adipose tissue – e.g on patients with very little VAT, and on slices where part of the liver is present (such an example is presented later in this section). On the other hand, the rectangles should not be chosen too large as this would ruin the local scope of the regions.

Values of $n_r = 12$ and $n_c = 9$ with overlaps of $o_r = o_c = 10$ pixels are selected giving reasonable results in areas not suffering from these problems, the erroneously placed points caused by the single-class-tissue covering regions is dealt with later. Different values have been tested, and it was found, that a number of subdivisions around the chosen would yield more or less the same sets of points. The dependency is also reduced by further measures against sampling spurious points – as described later.


For each subdivision, R_{sub} , the points $I_{max,R} = I_{max} \in R_{sub}$ and $I_{min,R} = I_{min} \in R_{sub}$ are considered. To account for the assumed slow spatial variation of the bias field effect, the points are trimmed with respect to the extreme values of each class within the subdivision and a threshold percentage p_t . Points satisfying

$$\begin{aligned} I_{max,R} &< p_t \cdot \max(I_{max,R}) \\ I_{min,R} &> (1 + (1 - p_t)) \cdot \min(I_{min,R}) \end{aligned} \quad (3.7)$$


are discarded. A percentage of $p_t = 0.85$ was found to give good results for trimming points in both classes. Figure 3.13 shows images with the two sets of points after the trimming.

Regarding the points of I_{max} shown in figure 3.13(a) the trimming has removed the erroneously placed points, since a sufficient quantity of VAT is present. Figure 3.13(b) shows that the points of I_{min} left after the trimming suffer from the issue of large areas of SAT. Furthermore some of the points located outside the abdomen are also still present. For both classes a tendency for the points to huddle together is observed. The advantage of more information in these very densely sampled regions is outweighed by the computational load and the desired level of accuracy of the bias field estimation, thus it is favourable to spatially trim the point.

This spatial trim is done by another subdivision, this time the ROI is divided into equally sized non-overlapping rectangles. For each rectangle only the single point of I_{max} with the maximum intensity is kept, and only the single point with minimum intensity of I_{min} is kept. A choice of 10 rows and 10 columns of



(a)



(b)

Figure 3.13: **(a)** Points of I_{max} after subdividing and trimming with respect to intensity thresholds defined in relation to each the region maximum. **(b)** Points of I_{min} after subdividing and trimming with respect to intensity thresholds defined in relation to each region minimum.

rectangles is made, as this seems to provide a good trade-off, but as above the exact number of subdivisions used is not of great importance. Furthermore to remove points in I_{min} located outside the abdomen, no points of this class are kept in the subdividing rectangles located on the edge of the ROI. This also has the positive effect of removing some of these false points inside the SAT. Not all are removed since the SAT layer is thicker than the range of a single rectangle. Figure 3.14 shows the two classes of points after this spatial trimming – the trimming has removed very few points in the sparse sampled areas and many points in the densely sampled areas.




Figure 3.14: Points of both classes after the spatial trim.

The points of I_{max} now appear to be more equally distributed across the abdomen. This is also the case for points in I_{min} located in the VAT region – but some points in this class are still erroneously located on the SAT layer. The SAT layer generally exhibits high intensities with a high contrast on all image series, so a solution to this problem is to apply a hard *global* threshold; only points of I_{min} with an intensity below 400 are accepted across the entire image. This value is found to give good overall results – not removing true points from I_{min} from other parts of the image, but lower threshold values have also been found to remove the spurious points of I_{min} . Figure 3.15 shows the result of the point sampling method.

As noted earlier, similar problems occur in slices with very little VAT and/or artifacts causing poor contrast between intensities of the two classes. Figure 3.16(a) shows an example; slice 2 of patient 16, containing very little VAT. The dark area in the left side of the image is the liver. Points of I_{max} are erroneously still placed here after the trimming. To solve this a lower intensity threshold is applied to the points. By looking at intensity histograms like the one shown in figure 3.9, and inspecting intensities of points selected at this step for various patients, a value of 500 is determined. This value also removes true points from I_{max} in some slices, but the goal of removing spurious points is achieved, and






Figure 3.15: The final set of points: Points of both classes; I_{max} and I_{min} , after applying a hard intensity threshold to the latter.

the value is determined to provide the best overall trade-off between these two effects. Figure 3.16(b) shows the example slice and points after applying this threshold.



(a)



(b)

Figure 3.16: The identified points shown before **(a)** and after **(b)** applying the hard thresholds on both classes. The set of points in I_{min} remains unchanged, as no points are located on the SAT layer. A lot of spurious points in I_{max} are removed by applying the lower threshold; excluding all points with intensity below 500 – the effect is especially visible in the left side where the liver is present. The image is slice 2 from patient 16.

3.2.2.2 Evaluation

Figures 3.17 to 3.20 show examples of slices with corresponding classes of voxel points sampled using the method described in the previous. For each example the intensity range of the two classes are displayed along with the number of points. Note that the points of I_{min} all have been thresholded with a lower intensity limit of 100, the reason for this is explained in section 3.2.6. The examples are chosen to illustrate the performance of the method. The method generally seems to perform well by sampling true voxel points distributed across the abdomen as intended, but the performance can only be assessed after the bias field estimation and correction by means of these points.




Figure 3.17: Illustrative example of the effect of trimming the points of I_{max} using overlapping regions with locally determined intensity thresholds; no voxel points are sampled from the relatively high image intensities exhibited by the spine, since the surrounding adipose tissue voxels exhibit higher intensities. Regarding the intensities of I_{min} across several slices, almost no sampled points exhibit intensities very close to the threshold value of 400. This large margin could support the assumption of a clearer difference between intensities of the SAT layer and the high water content voxels of the VAT region.




Figure 3.18: Illustrative example of the effect of trimming the points of I_{max} using a hard globally defined lower threshold; voxel points of this class are only sampled from what seems like true adipose tissue – even though the slice contains a lot of noise and very little contrast. Areas not densely covered by I_{max} are covered by I_{min} instead.




Figure 3.19: Image slice illustrating the drawback of the hard lower threshold imposed on voxel points of I_{max} ; no voxel points are sampled from what appears to be small low intensity flakes of VAT near the spine (arrows). However the VAT region is still densely covered by points of I_{min} .




Figure 3.20: Example to support the choice of a relatively high threshold for I_{max} ; even though the lowest intensity is 523, a voxel point is still erroneously sampled from the spine (arrow). Selecting a lower threshold could lead to more erroneously sampled voxels in this area on other patients, and a choice has been made rather to include too few true than too many points with a larger chance of sampling erroneously. Furthermore, the other class I_{min} is assumed to hold information on the bias field in the regions, where true points are removed.

3.2.3 Estimating the Bias Field

As stated in section 3.2.1, the bias field effect is unique for each patient and scan session and must therefore be estimated for each of the image series in order to remove it. The bias field effect model adopted in the estimation is presented in (3.6), and repeated below. This means, that the intensities of the voxel points sampled as described in the previous section are logarithmically transformed prior to estimating the bias field effect as a smooth function.

$$\begin{aligned} \log(y_{observed}) &= \log(y_{true}) + \log(y_{bias}) \Leftrightarrow \\ \log(y_{true}) &= \log(y_{observed}) - \log(y_{bias}) \end{aligned}$$

For each patient the set of points is used as a whole, estimating the bias field for the entire image series as one volume, since the bias field is expected to vary smoothly between slices as well as within slices. The model chosen for estimating a smooth function from the two classes of sampled points is the *thin plate spline* model. The theory is presented in the following section.

3.2.3.1 Thin Plate Splines

Thin plate splines (TPS) were first introduced by Duchon in 1976 [4]. This implementation of the TPS is based on the theory presented by Green and Silverman in [6] and Hastie et al. in [7].

Consider a set of observations sampled from a single slice with coordinates $[s_1 \ s_2]^T$ and intensity y – forming a 3-dimensional space. The principle of fitting an interpolating TPS to these data points can be thought of, as an infinitely thin metal plate being forced through the set of points – creating an interpolating fit with minimal bending energy. This principle can be carried to any dimensional space, so voxel points sampled over all slices can be considered as a whole, extending the function from a *plate* to a hypersurface.

For the sampled voxel point data the problem consists of N observations \mathbf{s} in \mathbb{R}^3 , with coordinates $[s_1 \ s_2 \ s_3]^T$, value (intensity) y and a class-indicator c

$$c(\mathbf{s}) = \begin{cases} 0 & \text{if } \mathbf{s} \in I_{min} \\ 1 & \text{if } \mathbf{s} \in I_{max} \end{cases} \quad (3.8)$$

For these two-class observations, an interpolating TPS can be formulated as

$$f(\mathbf{s}) = \beta_0 + \beta_1^T \mathbf{s} + \gamma c(\mathbf{s}) + \sum_j^n \delta_j h_j(\mathbf{s}) \quad (3.9)$$

The coefficient γ can be interpreted as a constant difference between values of the two classes of observations. The basis-functions $h_j(\mathbf{s})$ of the TPS are defined by means of n knots \mathbf{t} with coordinates $[t_1 \ t_2 \ t_3]^T$ located on a regular grid covering the extent of the observation points. That is, for the i th observation point \mathbf{s}_i the j th basis-function h_j is defined as the cubed distance to the j th knot $\mathbf{t}_j = [t_{j1} \ t_{j2} \ t_{j3}]$:

$$h_j(\mathbf{s}_i) = \|\mathbf{s}_i - \mathbf{t}_j\|^3 \quad (3.10)$$

A set of linear constraints are added to the basis function coefficients δ_j :

$$\sum_{j=1}^n \delta_j = \sum_{j=1}^n \delta_j t_{j1} = \sum_{j=1}^n \delta_j t_{j2} = \sum_{j=1}^n \delta_j t_{j3} = 0 \quad (3.11)$$

The bending energy (curvature) of f can be written as the function $J(f)$:

$$J(f) = \int \int \int_{\mathbb{R}^3} \sum_i^3 \sum_j^3 \left(\frac{\partial^2 f}{\partial s_i \partial s_j} \right)^2 ds_1 ds_2 ds_3 \quad (3.12)$$

The disregard of the image noise in the bias field model adopted, is a crude assumption. Often the sampled intensities are corrupted by image noise, thus fitting an interpolating TPS to the data would lead to an over-fitting model. Instead a smooth hypersurface approximating the intensities is sought – the smoothing TPS. This can be formulated as minimising the function S subject to f (a penalised residual sum of squares):

$$S(f) = \sum_i^N \{y_i - f(\mathbf{s}_i)\}^2 + \alpha J(f) \quad (3.13)$$

With α controlling the penalty with respect to the smoothness – a tradeoff between less curvature of the fit and closeness to the data (a tradeoff between bias and variance of the model). A low α will allow f to have much curvature since less penalty is enforced, with $\alpha \rightarrow 0$ leading towards an interpolation of the data points. On the other hand a high α will allow less curvature, and $\alpha \rightarrow \infty$ will result in f being the least squares fitting hyperplane. Choosing an appropriate α is investigated in section 3.2.4.

The solution to (3.13) is obtained by solving a set of linear equations. For the knots and data points two coordinate matrices are defined:

$$\mathbf{T}_k = \begin{bmatrix} 1 & \cdots & 1 \\ \mathbf{t}_1 & \cdots & \mathbf{t}_n \end{bmatrix}_{[4 \times n]} \quad (3.14)$$

$$\mathbf{T}_d = \begin{bmatrix} 1 & \cdots & 1 \\ \mathbf{s}_1 & \cdots & \mathbf{s}_N \end{bmatrix}_{[4 \times N]} \quad (3.15)$$

The corresponding data values and class-indicators are gathered in two $N \times 1$ vectors \mathbf{Y} and \mathbf{C} .

Matrices with pairwise cubed distance measures forming the basis functions of f in (3.9) are arranged; \mathbf{E}_k and \mathbf{E}_d with elements computed as

$$\{\mathbf{E}_k\}_{ij} = h_j(\mathbf{t}_i), \text{ with } i, j = 1, \dots, n \quad (3.16)$$

$$\{\mathbf{E}_d\}_{ij} = h_j(\mathbf{s}_i), \text{ with } i = 1, \dots, N \text{ and } j = 1, \dots, n \quad (3.17)$$

(3.9) can then be written in matrix form

$$\mathbf{F} = \mathbf{E}_d \delta + \mathbf{T}_d^T \beta + \mathbf{C} \gamma = [\mathbf{E}_d \quad \mathbf{T}_d^T \quad \mathbf{C}] \begin{bmatrix} \delta \\ \beta \\ \gamma \end{bmatrix} \quad (3.18)$$

where the vector \mathbf{F} consists of N elements $\{\mathbf{F}\}_i = f(\mathbf{s}_i)$, and $\beta = [\beta_0\beta_1]_{[4 \times 1]}$.

Gathering the coefficients δ_j in a $n \times 1$ vector δ , [6] states that the curvature function $J(f)$ can be written as

$$J(f) = \delta^T \mathbf{E}_k \delta \quad (3.19)$$

The linear constraints on the coefficients can be incorporated by means of $\mathbf{T}_k \delta = \mathbf{0}$ and a 4×1 Lagrange multiplier vector λ .

With this setup (3.13) can be written as:

$$S(f) = [\mathbf{Y} - \mathbf{E}_d \delta - \mathbf{T}_d^T \beta - \mathbf{C} \gamma]^T [\mathbf{Y} - \mathbf{E}_d \delta - \mathbf{T}_d^T \beta - \mathbf{C} \gamma] + \alpha \delta^T \mathbf{E}_k \delta + \lambda^T \mathbf{T}_k \delta \quad (3.20)$$

Computing partial derivatives of (3.20) with respect to the four variables (δ , β , γ and λ), and setting $\frac{\partial S}{\partial \delta} = \frac{\partial S}{\partial \beta} = \frac{\partial S}{\partial \gamma} = \frac{\partial S}{\partial \lambda} = 0$, the following system of linear equations is composed:

$$\begin{bmatrix} \mathbf{E}_d^T \mathbf{E}_d + \alpha \mathbf{E}_k & \mathbf{E}_d^T \mathbf{T}_d^T & \mathbf{E}_d^T \mathbf{C}^T & \mathbf{T}_k^T \\ \mathbf{T}_d \mathbf{E}_d & \mathbf{T}_d \mathbf{T}_d^T & \mathbf{T}_d \mathbf{C}^T & \mathbf{0} \\ \mathbf{C} \mathbf{E}_d & \mathbf{C} \mathbf{T}_d^T & \mathbf{C} \mathbf{C}^T & \mathbf{0} \\ \mathbf{T}_k & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \delta \\ \beta \\ \gamma \\ \lambda \end{bmatrix} = \begin{bmatrix} \mathbf{E}_d^T \mathbf{Y}^T \\ \mathbf{T}_d \mathbf{Y}^T \\ \mathbf{C} \mathbf{Y}^T \\ \mathbf{0} \end{bmatrix} \quad (3.21)$$

The system is solved to estimate the parameters $\hat{\delta}$, $\hat{\beta}$, $\hat{\gamma}$ and $\hat{\lambda}$:

$$\begin{bmatrix} \hat{\delta} \\ \hat{\beta} \\ \hat{\gamma} \\ \hat{\lambda} \end{bmatrix} = \begin{bmatrix} \mathbf{E}_d^T \mathbf{E}_d + \alpha \mathbf{E}_k & \mathbf{E}_d^T \mathbf{T}_d^T & \mathbf{E}_d^T \mathbf{C}^T & \mathbf{T}_k^T \\ \mathbf{T}_d \mathbf{E}_d & \mathbf{T}_d \mathbf{T}_d^T & \mathbf{T}_d \mathbf{C}^T & \mathbf{0} \\ \mathbf{C} \mathbf{E}_d & \mathbf{C} \mathbf{T}_d^T & \mathbf{C} \mathbf{C}^T & \mathbf{0} \\ \mathbf{T}_k & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{E}_d^T \mathbf{Y}^T \\ \mathbf{T}_d \mathbf{Y}^T \\ \mathbf{C} \mathbf{Y}^T \\ \mathbf{0} \end{bmatrix} \quad (3.22)$$

From $\hat{\delta}$ and $\hat{\beta}$ the effect of the bias field can be estimated for all voxel positions \mathbf{x} . The class-indicator coefficient $\hat{\gamma}$ is only included to estimate the parameters enabling the use of two classes of observations, and is not included when using the parameters to compute values of the estimated TPS \hat{f} in image positions \mathbf{x} . These coordinates are gathered in a matrix \mathbf{T}_x similar to (3.15), and a basis function matrix of pairwise cubed distances to the knots \mathbf{E}_x is set up as in (3.17). A vector with the estimated bias field effect $\hat{\mathbf{Y}}_x$ with elements $\{\hat{\mathbf{Y}}_x\}_i = \hat{f}(\mathbf{x}_i)$ can then be computed as

$$\hat{\mathbf{Y}}_x = [\mathbf{E}_x \quad \mathbf{T}_x^T] \begin{bmatrix} \hat{\delta} \\ \hat{\beta} \end{bmatrix} \quad (3.23)$$

3.2.4 Effective Degrees of Freedom

Selecting a suitable α – the degree of smoothness of the TPS fit – is important in order to get a good approximation of the bias field effect. The estimated hypersurface should be flexible enough to enable a correction of the variation in the areas of the bias peaks. But at the same time the allowed flexibility should not be too high, since this could over-fit the data, increasing the risk of incorporating noise in the parameter estimation.

The smoothing effect of one α -value varies when fitting to different sets of observation – thus a fixed α cannot be used to obtain similar smoothness between images series. Since it is undesirable to select an appropriate value of smoothness for estimating the bias field of each patient, another way of representing this factor – independent of the observation set – is needed.

Hastie et al. [7] describe an intuitive way of specifying the amount of smoothing by the term *effective degrees of freedom*, df_α . The term can be interpreted as the effective dimensionality of the fit; specifying $df_\alpha = 5$ would lead to the least squares fitting hyperplane, while $df_\alpha = n$ would lead to an interpolating fit. The following section derives this term with respect to the TPS model formulation presented in section 3.2.3.1.

Consider $\hat{\mathbf{Y}}$; a vector of the fitted N values $\hat{f}(\mathbf{s}_i)$ in the data points \mathbf{s}_i with observed values \mathbf{Y} .

$$\hat{\mathbf{Y}}^T = [\mathbf{E}_d \quad \mathbf{T}_d^T \quad \mathbf{C}^T] \begin{bmatrix} \hat{\delta} \\ \hat{\beta} \\ \hat{\gamma} \end{bmatrix} \quad (3.24)$$

The fitted values $\hat{\mathbf{Y}}$ can be obtained through a projection of \mathbf{Y} :

$$\hat{\mathbf{Y}} = \mathbf{H}_\alpha \mathbf{Y} \quad (3.25)$$

where the *hat matrix* \mathbf{H}_α is the projection matrix 'putting the hat on \mathbf{Y} '. From this, df_α is defined as the trace of the hat matrix

$$df_\alpha = \text{trace}(\mathbf{H}_\alpha) \quad (3.26)$$

The hat matrix corresponding to (3.24) is now derived. First (3.22) is written in a more compact form:

$$\begin{bmatrix} \hat{\delta} \\ \hat{\beta} \\ \hat{\gamma} \\ \hat{\lambda} \end{bmatrix} = \begin{bmatrix} \mathbf{S} & \mathbf{Q}^T \\ \mathbf{Q} & \mathbf{0} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{D} \\ \mathbf{0} \end{bmatrix} = \mathbf{\Gamma}^{-1} \begin{bmatrix} \mathbf{D} \\ \mathbf{0} \end{bmatrix} \quad (3.27)$$

The matrix $\mathbf{\Gamma}^{-1}$ is split into four parts:

$$\mathbf{\Gamma}^{-1} = \begin{bmatrix} \mathbf{\Gamma}^{11} & \mathbf{\Gamma}^{12} \\ \mathbf{\Gamma}^{21} & \mathbf{\Gamma}^{22} \end{bmatrix} \quad (3.28)$$

$\mathbf{\Gamma}^{11}$ is the $(n + 4 + 1) \times (n + 4 + 1)$ top left corner:

$$\mathbf{\Gamma}^{11} = \mathbf{S}^{-1} - \mathbf{S}^{-1} \mathbf{Q} (\mathbf{Q}^T \mathbf{S}^{-1} \mathbf{Q})^{-1} \mathbf{Q}^T \mathbf{S}^{-1} \quad (3.29)$$

By means of this split $\hat{\lambda}$ can be disregarded in (3.27) – enabling the following conversion:

$$\begin{bmatrix} \hat{\delta} \\ \hat{\beta} \\ \hat{\gamma} \end{bmatrix} = \mathbf{\Gamma}^{11} \mathbf{D} = \mathbf{\Gamma}^{11} \begin{bmatrix} \mathbf{E}_d^T \\ \mathbf{T}_d \\ \mathbf{C} \end{bmatrix} \mathbf{Y}^T \quad (3.30)$$

By replacing the vector of parameters in (3.24) with this expression, $\hat{\lambda}$ can be found:

$$\begin{aligned} \hat{\mathbf{Y}}^T &= [\mathbf{E}_d \quad \mathbf{T}_d^T \quad \mathbf{C}^T] \mathbf{\Gamma}^{11} \begin{bmatrix} \mathbf{E}_d^T \\ \mathbf{T}_d \\ \mathbf{C} \end{bmatrix} \mathbf{Y}^T \\ &= \mathbf{H}_\alpha \mathbf{Y}^T \end{aligned} \quad (3.31)$$

Using this, the rigidity of the TPS is controlled by specifying a more intuitive df_α instead of α . In practice this is implemented by a bisection algorithm solving $df_\alpha = \text{trace}(\mathbf{H}_\alpha)$ starting with a fixed initial value of α and adjusting it accordingly until the solution is within a given precision of the desired df_α .

Selecting an appropriate df_α for the bias field correction is dealt with in a parametric study described in section 3.2.8.

3.2.5 Removing the Bias Field

After having estimated the bias field effect for the entire image series from the sampled voxel points with logarithmic transformed intensities, the estimate corresponds to $\log(y_{bias})$ in the bias field effect model of (3.6). The corrected image intensities can then be obtained by a subtraction followed by an exponential transformation of each voxel intensity:

$$\begin{aligned} \log(y_{true}) &= \log(y_{observed}) - \log(y_{bias}) \Leftrightarrow \\ y_{true} &= \exp(\log(y_{observed}) - \log(y_{bias})) \end{aligned} \quad (3.32)$$

3.2.6 1-Dimensional Example

In the following a 1-dimensional example *re-enacting* the bias field correction method is presented to study the behavior and properties in a low dimension for better comprehension. A vector of image intensities, $\mathbf{I}_{observed}$ forms the basis for the example (a single row of voxels in an image slice), $\mathbf{I}_{observed}$ is plotted in figure 3.21 – the high intensities on each side correspond to the high intensities in the SAT layer.

Since it is very unlikely, that the voxel point sampling technique described in section 3.2.2 will place points on the exact same row of the image, the points used for this example are placed manually on local extrema distributed across the vector – only approximating the behavior of the sampling technique. The points of the two classes are marked on $\mathbf{I}_{observed}$ in figure 3.21 – a total of 12 points are placed (6 of each class).




Figure 3.21: The curve shows the intensity vector $\mathbf{I}_{observed}$ used for this example. The two classes of points are marked in green (I_{min}) and red (I_{max}).

Before the bias field estimation $\mathbf{I}_{observed}$ is logarithmically transformed; $\mathbf{I}_{log} = \log(\mathbf{I}_{observed})$. A set of 9 knots for the basis functions of the TPS are positioned equally spaced in the range of the sampled points. \mathbf{I}_{log} , the points, and the knots are shown on figure 3.22(a). It is worth noting, that the logarithm clearly reduces the magnitude of the bias peaks. The contrary is the case for the two low intensity areas – their relative *depth* is enhanced.

A TPS, \mathbf{I}_{bias} , with $df_{\alpha} = 7$ is fitted to the points to estimate the bias field. The red curve in figure 3.22(a) shows \mathbf{I}_{bias} . The role of the estimated coefficient, $\hat{\gamma}$, of the class indicator variable from (3.9) is illustrated by the vertical lines below the points of class I_{max} . \mathbf{I}_{bias} is generally smooth but the closeness of the fit is affected by the two points located in the extremely low intensity area.

The corrected intensity vector $\mathbf{I}_{corrected} = \exp(\mathbf{I}_{log} - \mathbf{I}_{bias})$ is shown on figure 3.22(b). The effect of the fit's lack of closeness around the two points is clear; the areas next to them are clearly erroneously corrected to a much too high intensity.






Figure 3.22: **(a)** The blue curve is the logarithmic transformed intensities \mathbf{I}_{log} . The 12 points of the two classes (red and green) are shown along with 9 knots for the basis functions of the TPS (blue). The red curve is the estimated bias field with $df_{\alpha} = 7$. The vertical red lines below the points illustrate the class indicator coefficient γ . **(b)** The intensity vector after the correction. The areas around the two extremely low intensity points are clearly erroneously corrected.

By inspection of the image data set subject to the study, it is observed that voxels inside the abdomen with intensity values below 100 often are caused by void areas – e.g. air inside the stomach and intestines. It is observed that there generally is a significant difference between this and intensities of the non-adipose surrounding tissue, and that these voxels corresponding to air hold no information on the bias field effect. These observations make it possible to assist the TPS by narrowing the range of intensities to be fitted. Before the

voxel point trimming described in section 3.2.2, all voxel points of I_{min} with an intensity below 100 are excluded. Figure 3.23 shows how this step significantly reduces the error caused by voxels from the low intensity areas.



(a)



(b)

Figure 3.23: (a) The logarithmic transformed intensities displayed along with the estimated bias field ($df_\alpha = 7$); the two voxel points in the extreme low intensity areas are removed, improving the closeness of the fit compared to figure 3.22(b). (b) The corrected intensity vector; the error around the *holes* is significantly reduced.

3.2.7 Method Overview

This section briefly sums up the pipeline of the bias field correction described in the previous sections:

- Local high and low intensity voxel points are located inside the abdomen on all slices to form two classes of points.
- Local low intensity points with intensities below 100 are removed.
- The points are trimmed to be equally distributed across each slice and to avoid spurious points.
- All intensities are logarithmically transformed, and the bias field is estimated across the entire image volume using thin plate splines extended to fit 3-dimensional points of two classes.
- The correction is performed by subtracting the estimated bias field from the transformed intensities, and exponentially transforming this difference.

3.2.8 Selecting Effective degrees of Freedom

For the TPS estimation a regular grid of $11 \times 7 \times 5$ knots are used for each image series – covering the range of the corresponding set of sampled points. The relation between the number of knots in each direction is chosen in correspondence with the overall height and width characteristics of the ROIs and the number of slices together with the in-plane resolution and slice thickness – looking at all image series to find one common grid size. Furthermore the total number of knots is also considered, since adding only a few knots in each dimension has a large impact on the computational time used.

To determine a number of effective degrees of freedom, df_λ , suitable for all image series in the data set, a parametric study of df_λ is performed. The principle of the study is illustrated by an example slice from an image series corrected using different values of df_λ . The image displayed is slice 16 of patient 7, where a significant bias field effect is observed (figure 3.24). This slice was also used to illustrate the voxel point sampling technique in section 3.2.2. The histograms presented only contain intensities of voxels located inside the abdomen ROIs, hence the amount of voxels in the few bins holding the very low intensities is reduced compared the histogram presented in section 3.2.1. Furthermore only the bias field inside the ROIs is considered, since care should be taken not to

use estimated bias field values obtained from knots positioned outside the area covered by the sampled points.

The objective of the correction is to make the high intensities of adipose tissue voxels distinguishable from the low intensity voxels. For an evaluation of this, the shape of the intensity histogram can provide a good overview of the performance; a distinct peak in the high end of the intensity range could indicate that adipose tissue intensities are well-separable from the rest.




Figure 3.24: (a) The original intensities of the example slice used to illustrate the parametric study to determine a suitable df_λ for the bias field correction. (b) Histogram of original intensities of voxels inside the abdomen volume.

Figures 3.25 and 3.26 show the estimated bias field for the slice, the corrected image slice and the intensity histogram of the entire corrected image series.

It is seen that $df_\lambda = 5$ gives a much too rigid bias field estimation, and the corrected intensities of the adipose tissue voxels still appear very inhomogeneous. Increasing df_λ produces better and better results – the separability of two peaks in the intensity histogram improves. Moving from $df_\lambda = 50$ to 80 there is still a noticeable improvement, but increasing the value further does not seem to improve the bias field correction noticeably. Using $df_\lambda = 230$ seems to worsen the correction, this could be due to the fact that the estimation is allowed to be too flexible, fitting too close to the data and thus allowing noise to affect the estimation. $df_\lambda = 80$ is chosen – as this value seems to provide a good overall trade off; fitting the intensity variation caused by the bias field while remaining robust towards noise.

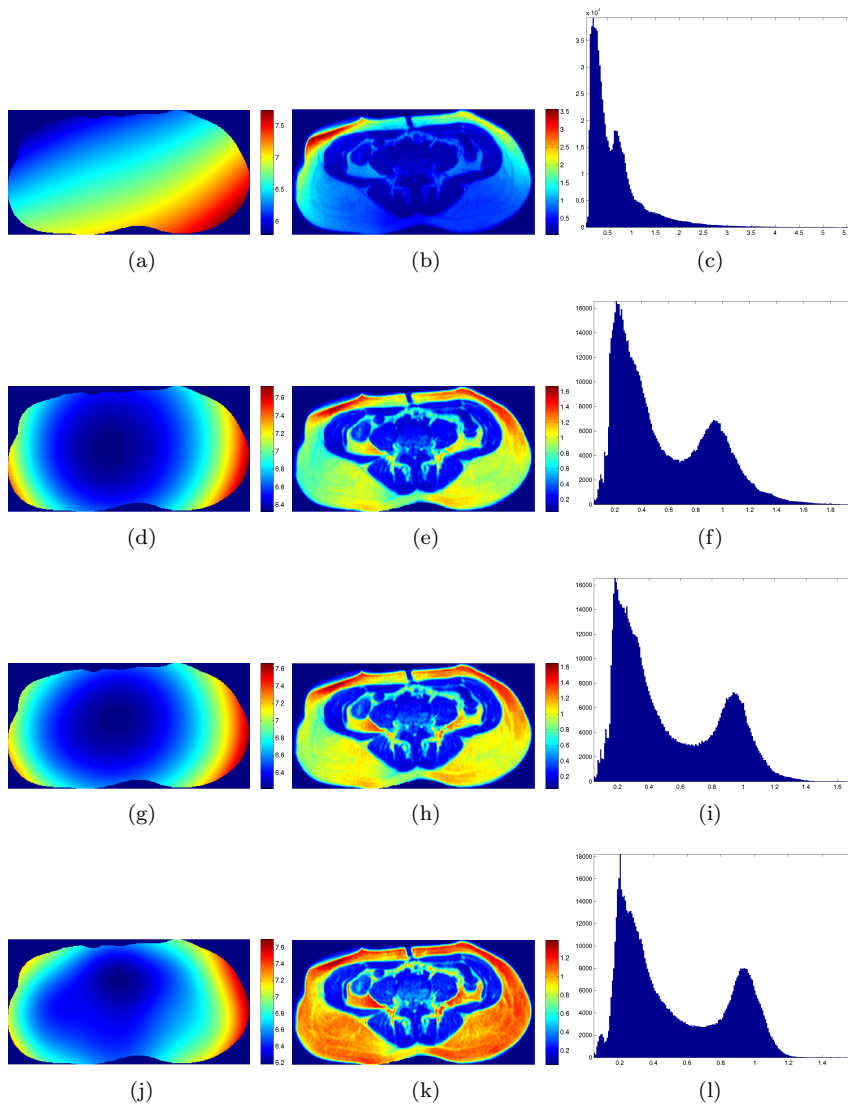


Figure 3.25: Illustrative example of the parametric study performed to determine a suitable df_λ for the bias field correction. The left column displays the estimated bias fields, The center column holds the corrected images. The right column shows intensity histograms of the entire corrected image series. The rows correspond to df_λ values 5, 10, 20 and 50 from top to bottom.

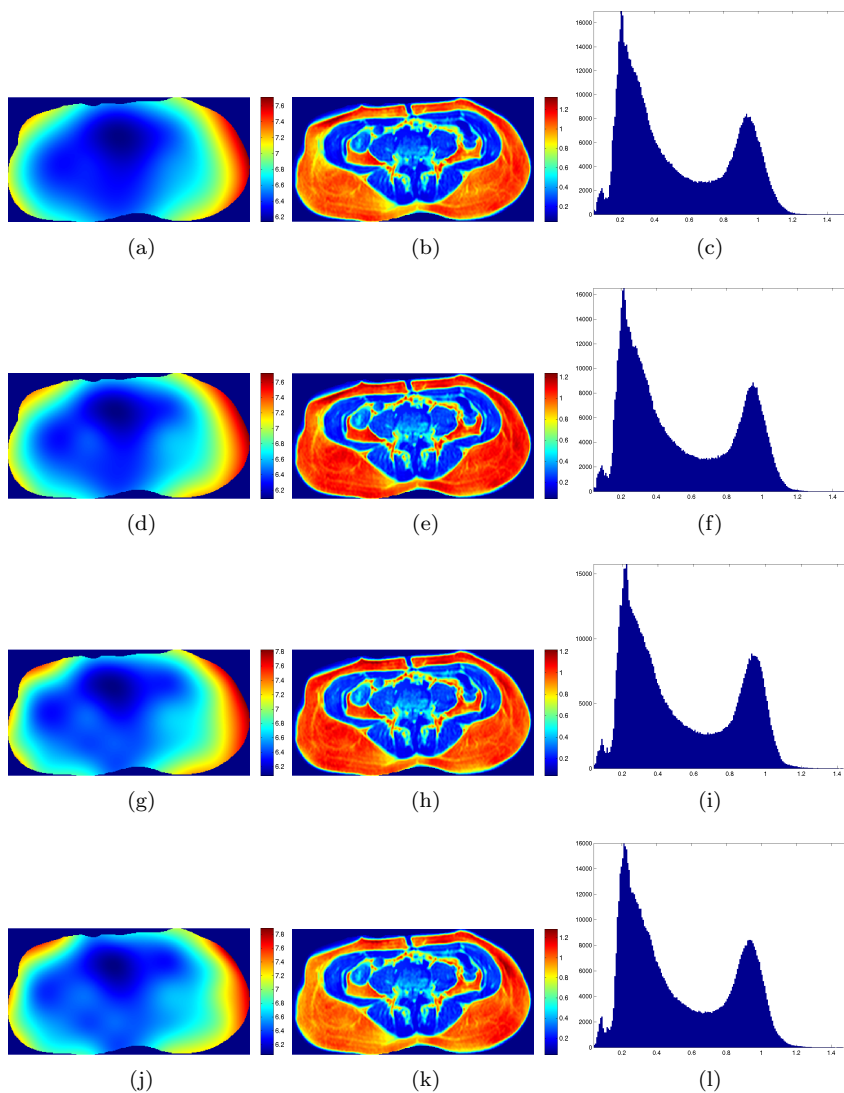


Figure 3.26: Illustrative example of the parametric study performed to determine a suitable df_λ for the bias field correction. The left column displays the estimated bias fields, The center column holds the corrected images. The right column shows intensity histograms of the entire corrected image series. The rows correspond to df_λ values 80, 120, 180 and 230 from top to bottom.

3.2.9 Evaluation

In this section examples of the bias correction results are presented and the method is evaluated by a visual inspection comparing selected slices and intensity histograms before and after the correction. The scale of intensities is different before and after the bias field correction, thus displayed intensities should not be compared – only the within-tissue homogeneity.

The general impression of the correction results is that, by relying on two classes of points, the method can produce valid estimates of the bias field effect in areas with little VAT. However for these images the method does require some extent of image quality. The presence of large noise contributions and/or low contrast can lead to a violation of the assumption of the bias field, being the only source of intensity variation with in each of the two classes, and thus add inaccuracies to the estimated field. It is also observed from the original images, that the flakes and cushions of VAT generally exhibit a less distinct transition to the non-adipose tissue compared to the SAT layer. This is carried to the correction results, making the VAT harder to delimit visually compared to the more well-defined boundary of the SAT layer.

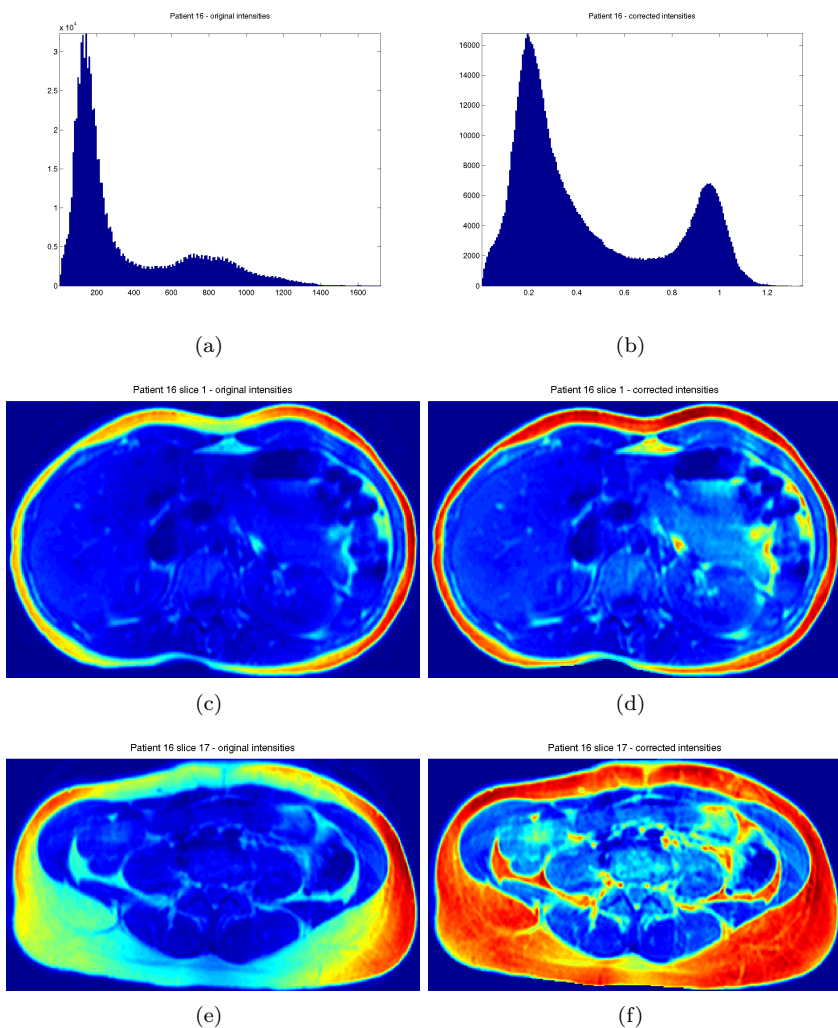


Figure 3.27: Illustration of the result of the bias field correction on patient 16. **(a)** and **(b)** Comparing the histograms the corrected intensities exhibits a much more separable high intensity peak, since the relative range of these adipose tissue intensities has been narrowed by the correction. **(c)** and **(d)** A slice with little SAT and VAT. The correction produces good results, not erroneously over-correcting the low intensities in the large parts of non-adipose tissue. **(e)** and **(f)** The contrast has been enhanced for adipose tissue voxels across the slice; the SAT layer intensities appear homogeneous and also much more similar to the intensities of VAT originally exhibiting lower contrast.

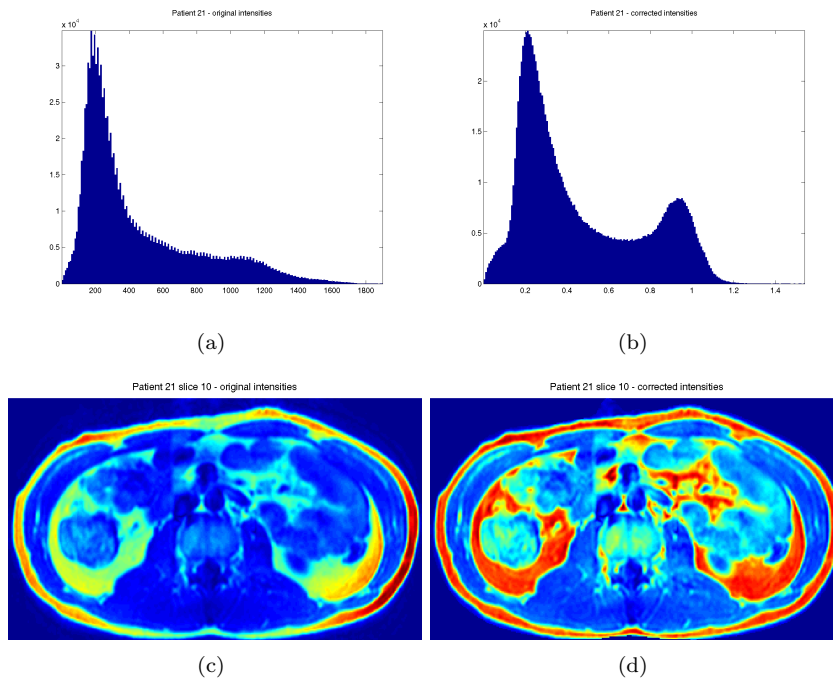


Figure 3.28: Illustration of the result of the bias field correction on patient 21. (a) and (b) The high intensity peak is evident on the corrected intensities. The slice example shows how the bias field effect is successfully removed from both the affected SAT and VAT. For (c) and (d) the corresponding points were shown on figure 3.17; no points were sampled from the spine area. Therefore these intensities still appear relatively low after the correction.




Figure 3.29: Two examples illustrating the weakness in using the same hard threshold value on I_{max} across all image series. The points used for the correction were shown on figures 3.19 and 3.20. Automatically determining an individual hard threshold value for each image series/slice would add more robustness to the method towards handling noise and contrast variations between image series. (a) and (b) The low contrast in the VAT flakes around the spine is not corrected sufficiently since no points of I_{max} are sampled here for the estimation. (c) and (d) The true VAT seems to be corrected sufficiently, but spurious points imposes inaccuracies in the correction of non-adipose tissue intensities.

3.3 Results

This section finishes the preprocessing described in this chapter with an overview of the method and some concluding remarks. 6 randomly selected slice examples of the final preprocessing results are displayed on figures 3.30 and 3.31.

The arms are removed from the image series such that only the abdomen is present. The method is successful on all but one of the 21 patients. For this single patient it was not possible to automatically remove both arms, since one of them was located very close to the abdomen. The image data for this particular patient was acquired late in the project period, and therefore the method was not modified to handle this particular case.

The within-class inhomogeneities caused by the bias field effect are removed to enable an intensity based tissue classification. The effect is estimated from intensities of points sampled in both adipose and non-adipose tissue across the entire abdomen. These points are sampled automatically as local intensity extrema, and the sampling method generally performs well. However for some image slices the intensities feature a high degree of noise and low contrast, causing the method to sample local extrema points with intensities not representing the bias field effect, but the error. Some measures are taken against this, but in order to gain overall performance across all image slices in all image series, a trade-off is made forcing some erroneous points to be sampled and some true points to be excluded – introducing inaccuracies to the estimated bias field and resulting correction.

From the sampled intensities the effect is estimated by fitting a smooth function approximating the assumed spatially smooth behaviour of the effect. The function consists of cubic spline basis functions exhibiting minimum bending energy, under a curvature constraint. The curvature constraint is imposed by means of specifying the effective degrees of freedom of the estimate – a measure to ensure a similar behaviour of the bias field estimate between images.

Overall the intensity histograms and image slices show that the corrected intensities are suitable for the automatic tissue classification presented in chapter 4.

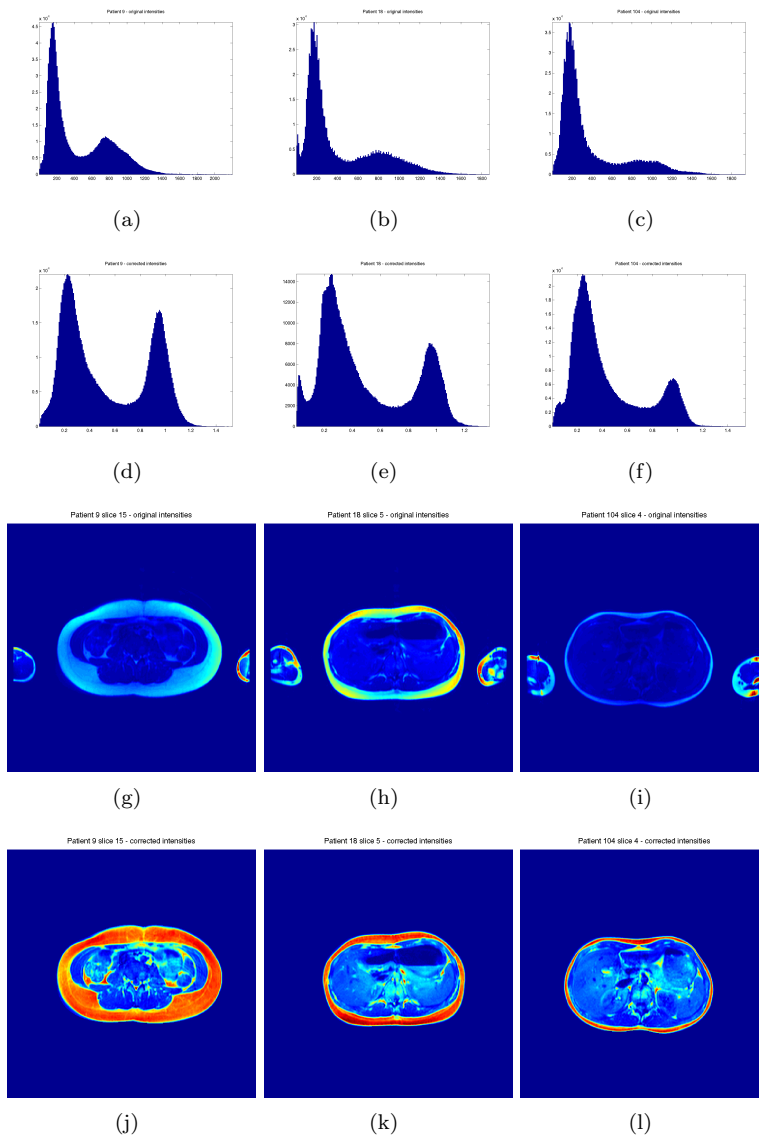


Figure 3.30: First row; histograms of original intensities. Second row; histograms of corrected intensities. Third row; image slices with original biased image intensities. Fourth row; image slices with corrected intensities. The three columns shown are; patient 9 slice 15, patient 18 slice 5 and patient 104 slice 4.

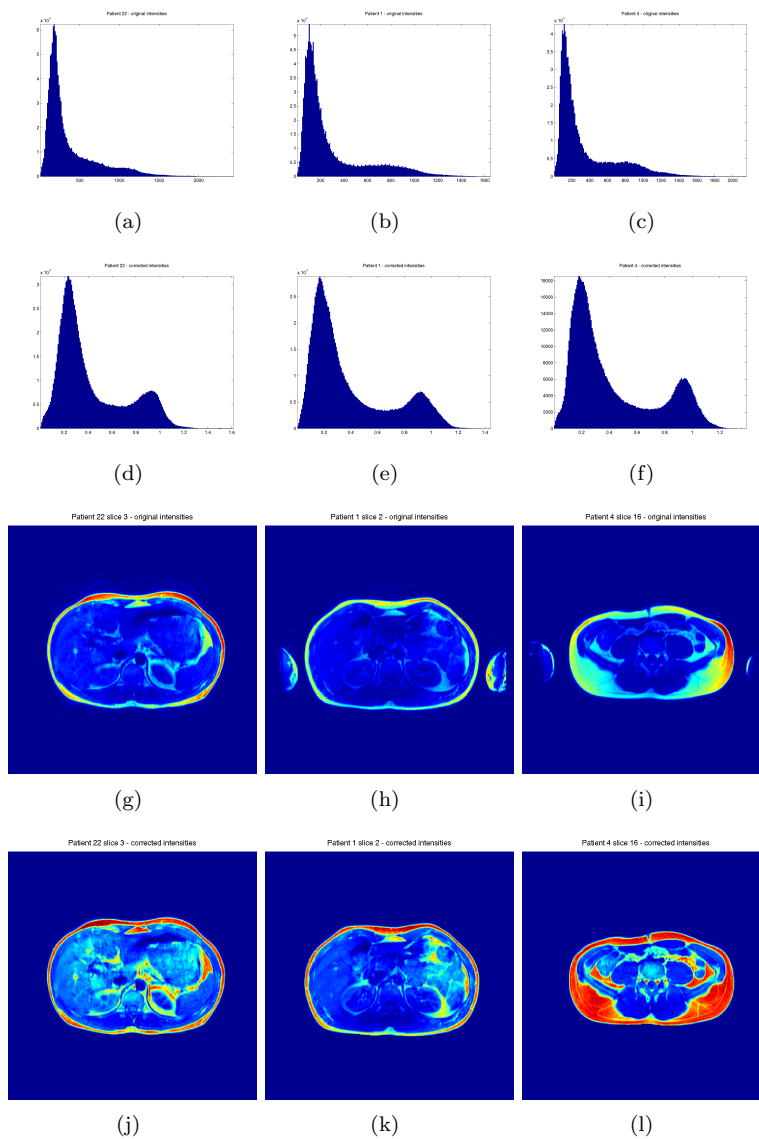


Figure 3.31: First row; histograms of original intensities. Second row; histograms of corrected intensities. Third row; image slices with original biased image intensities. Fourth row; image slices with corrected intensities. The three columns shown are; patient 22 slice 3, patient 1 slice 2 and patient 4 slice 16.

Identifying Adipose Tissue

After undergoing the preprocessing described in the previous chapter, the image data is ready for the adipose tissue segmentation. This chapter covers the first step of the segmentation: Identifying voxels corresponding to adipose tissue. Identifying the regions defining the three types of adipose tissue will be handled in the second step, described in chapter 5.

In this work the data set solely consists of T1-weighted MRI – no multi-modal image data is available, and the tissue identification is therefore only based on the corrected T1-weighted intensities. Furthermore no ground truth segmentations are available. The presence of such a training set of voxels labelled with respect to tissue type would enable the application of *supervised* classification methods – using the labels to obtain some prior knowledge of the relationship between the intensities and class membership. Since this information is not available, *unsupervised* classification is adopted for the solution – a technique for deriving the natural structure of data from the data itself to form the basis of the classification.

The intensity histograms of the abdomen in the corrected image data generally exhibit two peaks; one for low intensities (non-adipose tissue voxels), and one for high intensities (adipose tissue voxels). This indicates the possibility of identifying the adipose tissue using an intensity based thresholding – separating

the histogram peaks.

In intensity based thresholding voxels are assigned labels by comparing intensity values to one or more intensity thresholds. A single threshold partitions the set of voxels into two regions, multiple thresholds can create multiple regions. The thresholds can either be locally defined; different threshold values for subsections of the image, or globally defined; one value used across the entire image domain. The large spatial intensity variation caused by the bias field effect was dealt with by the preprocessing in chapter 3, but as noted some of the VAT appears less well-defined in comparison to SAT, thus locally defined thresholds are adopted to identify these two types of adipose tissue – this choice is elaborated in the end of the chapter.

The following sections, 4.1 and 4.2, will present a method for identifying adipose tissue voxels by automatically deriving optimal threshold values from the image data. Section 4.3 finished the chapter with concluding remarks and example results of the tissue classification applied to the corrected image data shown in section 3.3.

4.1 Fuzzy C-Means Clustering

As noted in the chapter introduction, no ground truth data is available to help determining the thresholds for the tissue classification, thus the identification of regions, or *clusters*, is unsupervised and solely relies on the structure of the image intensities.

A technique for solving this problem is fuzzy *c*-means clustering (FCM) described by Dawant and Zijdenbos in [3]. Jørgensen [10] and Engholm et al. [5] have achieved good results using FCM for segmenting adipose tissue from T1-weighted MRI without ground truth labelled voxels.

The FCM algorithm differs from other clustering algorithms (such as *k*-means clustering) as each voxel is associated with more than one cluster. This is done by means of continuous, *fuzzy*, membership values reflecting a certain degree of membership to all clusters – instead of assigning voxels to only a single cluster each. These membership values form a common reference frame of the connection between intensity and tissue types comparable between images, since they are derived to resemble the individual structure of the data.

The algorithm is formulated around the minimisation of a criterion function, J_{FCM} , expressing the quality of the clusters identified by a weighted sum of

squared distances:

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

The subscript j denotes a voxel location in the image domain Ω , k is the class index, C is the total number of classes. The parameter u_{jk} is the fuzzy membership value of class k for voxel j , and q is a weighting exponent greater than 1 defining the fuzziness of the classifications. y_j is the intensity of voxel j and v_k is the centroid of class k . The membership values must conform to the following constraints:

$$\begin{aligned} 0 &\leq u_{jk} \leq 1 \\ \sum_{k=1}^C u_{jk} &= 1 \end{aligned} \quad (4.2)$$

The criterion function J_{FCM} is minimised with respect to u_{jk} and v_k . This is done through iteratively updating them by evaluating the following two equations:

$$v_k = \frac{\sum_{j \in \Omega} u_{jk}^q y_j}{\sum_{j \in \Omega} u_{jk}^q} \quad (4.3)$$

and

$$u_{jk} = \frac{|y_j - v_k|^{\frac{-2}{q-1}}}{\sum_{k=1}^C |y_j - v_k|^{\frac{-2}{q-1}}} \quad (4.4)$$

The algorithm is initialised by assigning random values to u_{jk} (conforming to (4.2)), and then iteratively alternating between evaluating (4.3) and (4.4), until reaching convergence – when the resulting change in J_{FCM} between subsequent iterations is suitably small.

The entire abdominal volume is clustered as one; thus the image domain Ω is defined as voxels inside the ROIs (figure 3.10) of all slices in the image series. This seems plausible, since the bias field correction was performed for the entire 3-dimensional image series, such that the within-class intensities should be spatially homogeneous within slices as well as between slices.

Figure 4.1 shows an example of the resulting membership values as a function of voxel intensity. The values are obtained by the FCM algorithm with $C = 2$ and $q = 2$, the image data is patient 7. The membership curves of the two clusters are superimposed on the corresponding intensity histogram; the units on the y-axis are number of voxels, so the membership values ($0 \leq u_{jk} \leq 1$) have been magnified to a suitable scale for the plot. The vertical dashed lines mark the cluster centroids.




Figure 4.1: FCM algorithm membership values for $C = 2$ plotted as a function of voxel intensity. The plot is superimposed on the intensity histogram of the corresponding image series – the values of the membership function have been magnified for the plot. The vertical dashed lines mark the cluster centroids. The image data subject to the clustering is patient 7.

Voxels close to a cluster centroid are assigned high a membership value to that cluster – thus low to the other. It is observed that the centroids are located close to the two histogram peaks. This indicates that membership values obtained with $C = 2$ are useful for distinguishing the high intensity adipose tissue voxels from the low intensity non-adipose tissue voxels – the two distinct peaks generally observed in the intensity histograms of the corrected image series.

Figure 4.2(a) shows the values of the criterion function J_{FCM} for each iteration of the FCM algorithm to achieve the membership values shown on figure 4.1 . Figures 4.2(b) to 4.2(e) show membership values as a function of image intensity through the iterative procedure. 16 iterations are spent to reach convergence, but already after 8 iterations the criterion function reaches a steady level, and not much change is observed for the membership values through the last 8 iterations.

Figure 4.3 shows the effect of varying the q parameter of (4.1). The relation between the value of q and the amount of fuzziness in the membership values is evident. A value of q close to 1 results in almost no fuzziness – a very high membership to one of the classes is assigned to almost all intensities. A higher values of q results in only assigning high membership values to intensities very similar to either of the cluster centroids. It should be noted that while the shape of the two curves change with q , the two centroids remain almost constant at

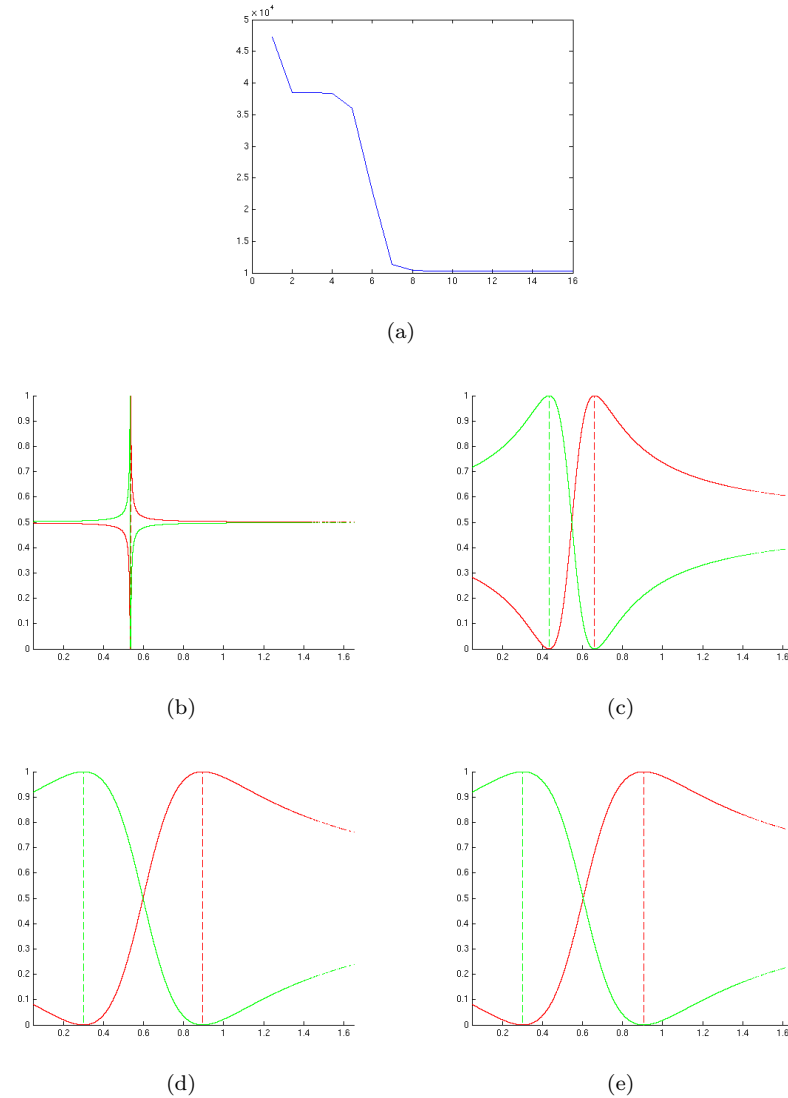


Figure 4.2: **(a)** The values of the criterion function J_{FCM} plotted for each of the 16 iterations used by the FCM algorithm to reach convergence. **(b)** After 2 iterations still no real structure is visible – the centroids values are equal and most intensities are assigned membership values of 0.5 for both clusters. **(c)** After 6 iterations the two clusters are becoming more separable. **(d)** After 8 iterations the centroids have moved a little more apart and the curves appear more round. **(e)** The final result; after 16 iterations. Not much change can be seen compared to the result after 8 iterations, but this is also observed in the little change in the corresponding values of the criterion function.

approximately 0.3 and 0.9. In this work $q = 2$ is chosen, as it appears to give

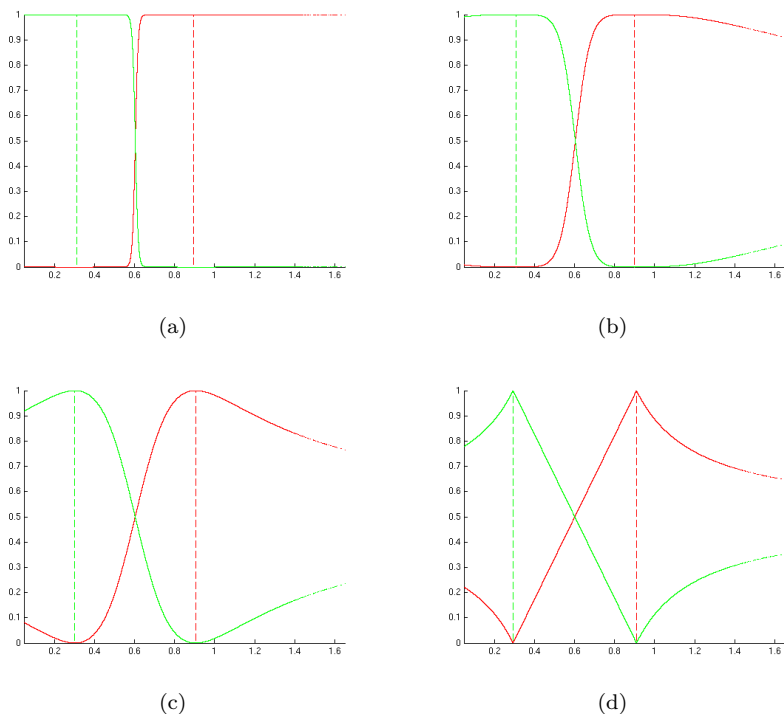


Figure 4.3: Examining the q parameter effect by means of the shape of the membership function plots. **(a)** $q = 1.1$; no fuzziness, intensities are almost assigned full memberships to either one of the two classes. **(b)** $q = 1.5$; a little more fuzziness is observed. **(c)** $q = 2.0$; smooth membership curves, exhibiting both round peaks but also a high degree of fuzziness. **(d)** $q = 3.0$; very pointy curves, almost no high membership values – only for voxel with intensities around the cluster centroids (very fuzzy).

a good tradeoff between high membership values close to the centroids and the degree of fuzziness assigned to intensities less similar to the centroids.

The assumed correspondence between the two identified clusters and classes of non-adipose and adipose tissue intensities is confirmed, by comparing intensities and assigned membership value in image slices; figure 4.4 shows an example of such a comparison. Figure 4.4(b) exhibits high membership for low intensity voxels while figure 4.4(c) for high intensity.

For this example it is furthermore observed that the membership values are well defined as either high or low – only a little fuzziness is present. This is because the high membership values are assigned to intensities around the two histogram peaks – that is, fewer voxels are present in the intensity range corresponding to the more fuzzy membership values (see figure 4.1).




Figure 4.4: Example of the membership values in figure 4.1 displayed for each voxel in an image slice. Note that only voxels within the ROI are considered for the classification. Membership values for voxels outside the ROI are shown as 0 for both clusters. **(a)** The image slice (after the bias field correction). **(b)** Membership values for the low intensity cluster. **(c)** Membership values for the high intensity cluster.

The cluster corresponding to adipose tissue is identified as the one with the highest centroid value v_a , its membership values are denoted u_{ja} .

4.1.1 Evaluation

Before presenting a way of translating the membership values into classification of non-adipose and adipose tissue voxels in section 4.2, some example results of the FCM algorithm are shown and evaluated.

The method is evaluated by means of selected histograms of the corrected intensities with membership values superimposed, and corresponding slice examples

of corrected intensities and membership values. For the examples only the values of u_{ja} are shown – since with $C = 2$, the membership values of the other cluster are just $1 - u_{ja}$.

For the data set, the assigned membership values generally seem to resemble the two classes of tissue defined, though with a little more fuzziness present in the VAT region. This is especially observed in images with a higher degree of noise and/or very little adipose tissue – where the high intensity peak appears less well-defined on the histogram. Overall the FCM algorithm seems to be a suitable method for automatically transforming intensities to values corresponding to the two tissue classes. As noted, these membership values provide a common reference-frame of intensity similarity to tissue classes comparable between all image series. This is exploited in the next section, where the membership values are used to automatically classify voxels as either non-adipose or adipose tissue.

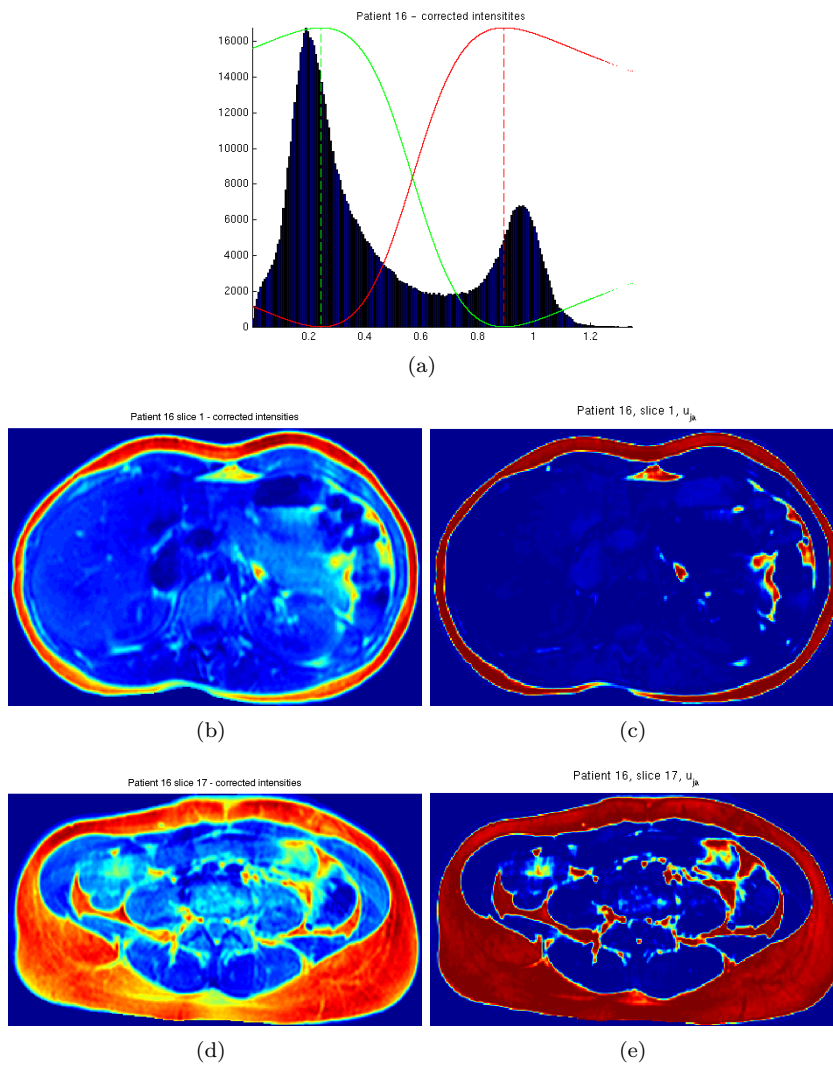


Figure 4.5: (a) The two histogram peaks are well separated. (b)-(e) The histogram shape is carried to the membership values; high values of u_{ja} are exhibited for the high intensity adipose tissue voxels in the corrected image slice. Only a little fuzziness is observed around the smaller and less distinct flakes of VAT.

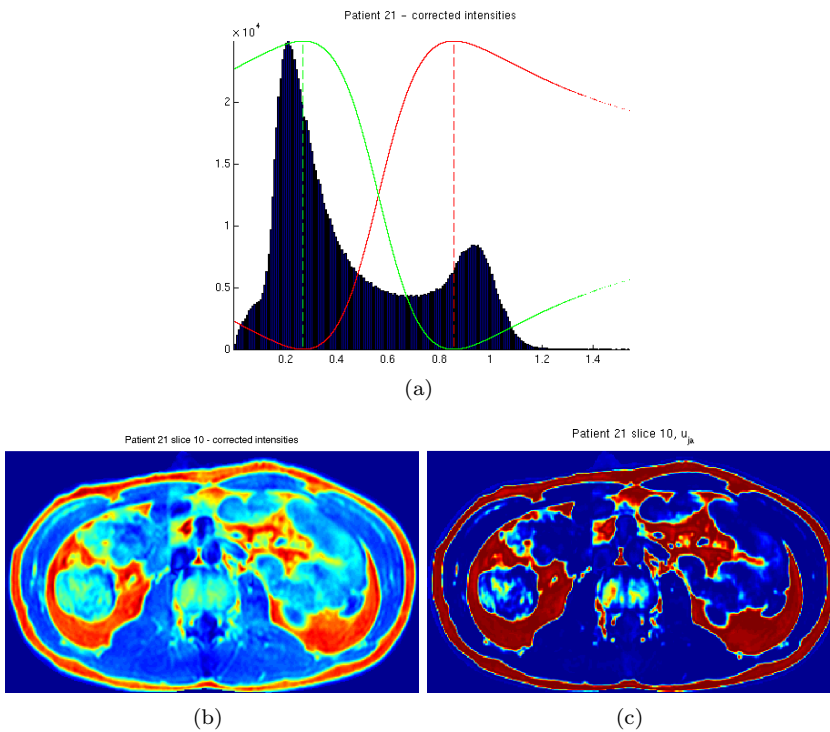


Figure 4.6: The correspondence between adipose tissue voxels and the membership values seems clear, although some fuzziness is observed in the spine area, where voxels with intensities of the histogram *valley* are present.

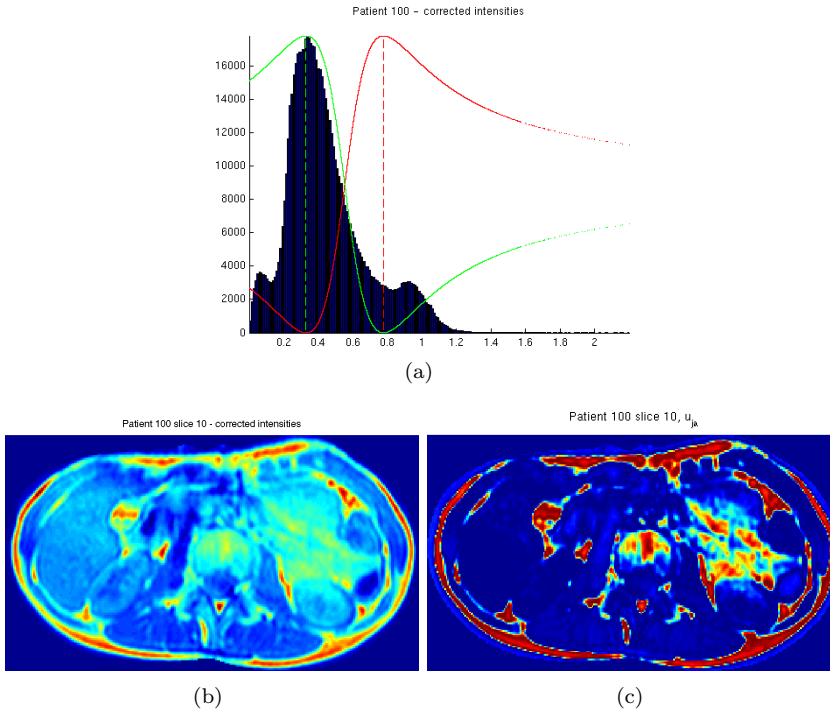


Figure 4.7: The corrected intensities of patient 100 exhibit low image quality. The noise has not been removed by the preprocessing, even though the points used in the bias field correction seemed to be correctly sampled (figure 3.18). This is carried to the derived membership values. **(a)** The noise in the image gives rise to a large skew peak in the lower range of the corrected intensities, making the peak of the few high intensity voxels indistinct. This is carried to the membership values; the centroid, v_a , is not located on the very faint high intensity peak, and a higher amount of intensities are assigned the fuzzy membership values. **(b)** and **(c)** Poor contrast voxels in the spine area and the right side of the image receive high membership values.

4.2 Classification

The classification applied in this work is a *discrete* scheme, labelling voxels as only one class each, considering two classes for the tissue labelling; adipose tissue and non-adipose tissue. Such a binary label f of voxel j is obtained by means of a threshold ε_m applied on the membership value of the cluster corresponding to the adipose tissue u_{ja} . The labelling function is formulated as

$$f(j) = \begin{cases} 1 & \text{if } u_{ja} \geq \varepsilon_m \\ 0 & \text{if } u_{ja} < \varepsilon_m \end{cases} \quad (4.5)$$

where $f(j) = 1$ means labelling voxel j as adipose tissue. This way a high value of ε_m can be interpreted as making the voxels labelled as adipose tissue more probable of being true adipose tissue. Using two clusters, $\varepsilon_m = 0.5$ corresponds to placing a threshold at the intersection of the membership curves – classifying as adipose tissue when this membership value is the highest of the two.

From the binary labelling the total quantity of adipose tissue inside the abdomen M_a can be obtained by simply counting the number of voxels classified as adipose tissue:

$$M_a = \sum_j f(j) \quad (4.6)$$

Figure 4.8 shows a discrete classification of adipose tissue done with $\varepsilon_m = 0.5$ from the membership values also shown on figure 4.4(c).




Figure 4.8: (a) The membership values for the adipose tissue class. (b) The labels from a discrete tissue classification with a threshold of $\varepsilon_m = 0.5$. The voxels shown in red are labelled as adipose tissue. The image is slice 16 from patient 7.

Another classification method using the membership values is the *fuzzy* scheme – this is not used in this work and will therefore just be touched briefly here.

In fuzzy classification the membership value of a voxel is used as a measure of the partial volume effect; the effect on the voxel intensity caused by different tissue types being present inside the volume covered by the voxel. This way an adipose tissue voxel is considered to contain a larger fraction of non-adipose tissue if exhibiting an intensity value far from the cluster centroid. The total quantity of adipose tissue inside the abdomen can be obtained by summing the membership values of the cluster corresponding to the adipose tissue class:

$$M_a = \sum_j u_{ja} \quad (4.7)$$

This means that the fuzzy classification corresponding to figure 4.8(b) will just be the membership values displayed in figure 4.8(a). As noted, the membership values for this particular slice contain very little fuzziness, and therefore the two classifications resemble each other well. But on the other examples shown in section 4.1.1 more fuzziness can be present due to low contrast and image noise in the VAT area. This way the fuzziness is not only reflecting the partial volume effect as intended. Furthermore as the results in this work are subject to visual inspection, this scheme seem less appropriate, as the fuzzy classifications are more difficult to grasp by the human eye, compared to the discrete.

Figure 4.9 shows the effect of a discrete classification using different values of ε_m in (4.5). As noted, the image displayed in figure 4.4 contains only a little amount of voxels with fuzzy membership values, therefore the changes observed between using $\varepsilon_m = 0.7$ and 0.85 are very small compared to $\varepsilon_m = 0.5$; the *certain* adipose tissue voxels are still labelled correctly. Not before raising the threshold to $\varepsilon_m = 0.95$, a significant change in the tissue classification becomes visible.

As seen in section 4.1.1 the FCM algorithm assigns a higher amount of fuzziness to the voxels in the VAT region on slices where less contrast is present and the adipose tissue voxels seem less well-defined. This means that a threshold of $\varepsilon_m = 0.5$ can cause the classification to be too *generous* – labelling too many voxels as VAT. From figure 4.9 it was observed that raising ε_m had little effect on labelling an image exhibiting good contrast in the VAT region, but for images with more fuzzy membership values it can make a greater difference.

By visual inspection of several slices, their membership values and the resulting tissue labels, a value of $\varepsilon_m = 0.85$ has been found suitable for labelling the VAT region. While $\varepsilon_m = 0.5$ gives good results for classifying the more well-defined SAT layer.

The example in figure 4.10 shows the classification performed with both $\varepsilon_m = 0.5$ and $\varepsilon_m = 0.85$. The two set of labels are superimposed for better comparison;




Figure 4.9: Adipose tissue labelled using various threshold values, ε_m , for the tissue classification done by means of the membership values shown in figure 4.4(c). (a) $\varepsilon_m = 0.5$, (b) $\varepsilon_m = 0.7$, (c) $\varepsilon_m = 0.85$ and (d) $\varepsilon_m = 0.95$.

that is, the red labels mark voxels labelled as adipose tissue with $\varepsilon_m = 0.85$, the green voxels show the extra voxels labelled by using $\varepsilon_m = 0.5$. Some difference between the two sets of labels is observed; the labelling becomes more tight around the smaller flakes of VAT. Furthermore the high intensity voxels in the spine area of the example slice are successfully excluded from the adipose tissue class by applying the higher threshold. For the two big VAT cushions on the two sides little effect is observed from raising the threshold, since their boundaries are well-defined.




Figure 4.10: An illustrative example of raising ε_m to improve the VAT classification. The two sets of labels ($\varepsilon_m = 0.5$ and $\varepsilon_m = 0.85$) are superimposed for a better view of the difference. The green voxels show the extra adipose tissue labels assigned by using the lower of the two thresholds.

This leads to the choice of applying locally defined thresholds for the classification distinguishing between the VAT region and the SAT layer. The boundary separating the regions is determined in chapter 5.

Before presenting the labelling of some selected slices for evaluating the method performance, a note should be made regarding the importance of performing the bias field correction prior to the tissue labelling. Figure 4.11 shows labels obtained by the FCM algorithm with $C = 2$, $q = 2$ and a threshold of $\varepsilon_m = 0.85$. A classification is performed on both the original and corrected intensities of patient 7. The effect is evident on the example slice shown; from the original image data almost no VAT is labelled as opposed to what appears to be a good tissue classification achieved from the corrected intensities.




Figure 4.11: Example to illustrate the effect of the bias field correction. **(a)** and **(b)** Original and corrected intensities. **(c)** and **(d)** Tissue classifications of original and corrected intensities. The classifications are done using $\varepsilon_m = 0.85$. The image is slice 10 of patient 7.

4.2.1 Evaluation

The evaluation of the discrete classification scheme is done by following up on the membership value examples from section 4.1.1, showing the corrected intensities and resulting tissue labels. The boundary defining the local classification thresholds is still to be determined at this point, thus both adipose tissue voxels classified with $\varepsilon_m = 0.5$ and $\varepsilon_m = 0.85$ are displayed for the whole slice.




Figure 4.12: These two examples illustrate that the use of two thresholds performs well on slices both with a small amount of adipose tissue and a larger amount, only labelling the most *certain* voxels as VAT. (a) and (b) The adipose tissue labels are assigned to what looks like true adipose tissue from the corrected intensities. On this image slice the posterior SAT layer is so thin, that a hole appears in the rim of labelled voxels – using either of the thresholds. (c) and (d) 0.85 seems to do much better than 0.5 on labelling VAT – as the latter value includes too many voxels.

The general impression of the resulting tissue classifications is that labels match the visual perception. Using the membership values from the FCM algorithm to determine the thresholds enables the classification to adapt to the natural structure of the intensities for each individual patient without any user-interaction.




Figure 4.13: The voxels containing noise in the right side of the image is classified as VAT using $\varepsilon_m = 0.5$, while $\varepsilon_m = 0.85$ reduces this amount significantly. This higher threshold also reduces the erroneous labelling in the spine area, but voxels are still falsely labelled as both non-adipose and adipose tissue.

The choice of using two locally defined thresholds seems reasonable for handling the difference in properties of the exhibited contrast between VAT and SAT.

4.3 Results

This section concludes the tissue classification with a short overview and comments on the method applied. Figures 4.14 and 4.15 show the corrected intensities and corresponding tissue labels for the 6 selected slices from section 3.3.

The labelling works unsupervised without ground-truth information – deriving the natural structure from the data itself to divide into optimal classes. The corrected intensity data generally feature two well-separable classes, this is the structure automatically exploited by the fuzzy c-mean clustering applied.

The derived structure is used to define a common reference-frame of membership values comparable between images. This makes the tissue classification more robust, as no data-dependent parameters are needed for the class separation. Inaccuracies in the corrected intensities and differences in adipose tissue characteristics are sought to be handled by using a locally defined classifier for the tissue labelling. This improves the overall performance, and a correspondence between the assigned labels and the expected results is generally present – although images suffering from intensity inaccuracies from more severe noise and bad contrast conditions can result in falsely labelled voxels.

The tissue labels of the data set were presented to Kasper Pilgaard, who verified that the classification method performed well across the variety of the image data set.

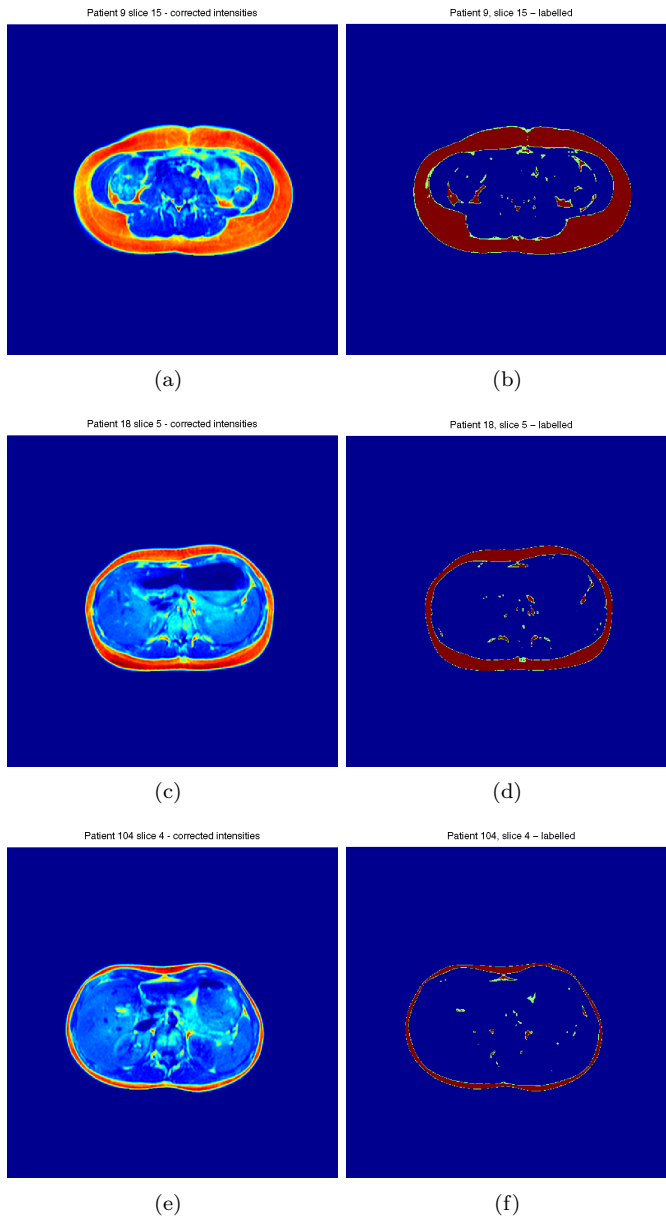


Figure 4.14: Each row holds corresponding corrected intensities and tissue labels for an image slice. The images are; patient 9 slice 15, patient 18 slice 5, and patient 104 slice 4.

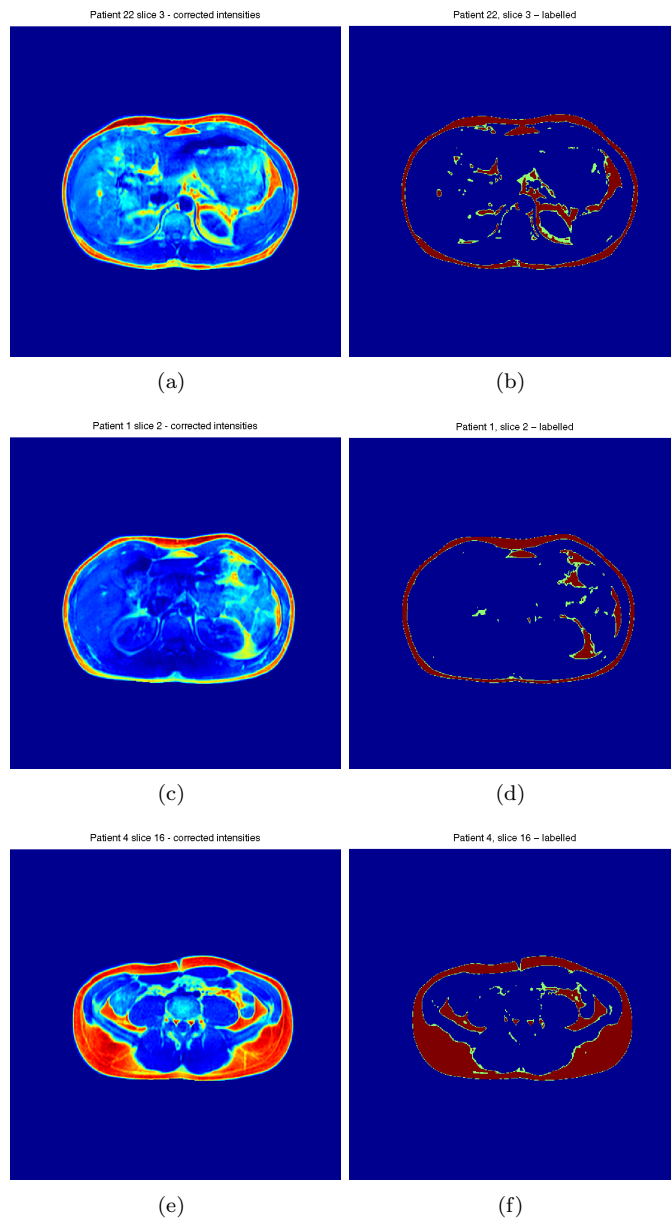


Figure 4.15: Each row holds corresponding corrected intensities and tissue labels for an image slice. The three images are; patient 22 slice 3, patient 1 slice 2, and patient 4 slice 16.

Identifying Regions

This chapter covers the second step of the segmentation process; to automatically divide the abdomen into the anatomically defined regions containing the three types of adipose tissue (as illustrated in figure 1.1).

Section 5.1 describes how the boundary of the abdomen and the interior boundary of the SAT layer are determined. In the image data the boundary of the abdomen is formed by the exterior SAT boundary located right underneath the thin layer of skin. The SAT layer is generally characterised as a rim of high image intensities with a smooth exterior boundary – easy to distinguish from the low intensity voxels present outside the abdomen. Furthermore the corrected intensities of the SAT layer appear homogeneous, and a connection between intensities and tissue type was derived through the tissue classification. These properties are exploited by means of applying a deformable model approach – *active contours*.

Section 5.2 describes the SAT subdivision done by means of locating Scarpa's fascia; a thin layer of connective tissue in the abdominal wall, dividing the SAT layer into a deep and a superficial part. On the images the fascia appears as a thin line of low intensity voxels following the shape of the abdomen, this enables the usage of *dynamic programming* for the subdivision.

5.1 Active Contours

The theory of the deformable model approach chosen in this work is covered in this section. In sections 5.1.1 to 5.1.3 the application is described; segmenting the abdomen boundary and the interior SAT boundary. Section 5.1.4 gives an evaluation of the method performance.

Deformable models is a widely applied technique for image segmentation. The term refers to closed curves or surfaces evolving through time in the image domain from an initial state to a final state. The variant of deformable model used in this work is *active contours*, introduced in 1987 by Kass et al. [11]. The theory behind the implementation is based on Xu et al. [14], extended by applying a method for incorporating image features similar to the formulation used by Engholm et al. in [5] for segmenting SAT in MRI.

The approach of active contours uses a parametric formulation, representing the model explicitly during the deformation by a predefined representation; storing point coordinates and information on the connections between points. For this work, the segmentation is performed slice-by-slice, thus the model used is a contour; a closed curve in 2-dimensions $\mathbf{X}(s) = (X(s), Y(s))$, with the natural parametrisation $s \in [0, 1]$. This parametrisation is implemented using curves of equally distributed points. The evolution is performed with the goal of minimising the curve energy \mathcal{E}_{curve} . This energy is composed of an internal and an external term:

$$\mathcal{E}_{curve}(\mathbf{X}(s)) = \mathcal{E}_{int}(\mathbf{X}(s)) + \mathcal{E}_{ext}(\mathbf{X}(s)) \quad (5.1)$$

The internal energy component, \mathcal{E}_{int} , is defined within the curve to govern the smoothness during the evolution. The external energy, \mathcal{E}_{ext} , is designed to incorporate the image features of interest in the segmentation task – here the SAT layer. By using a curve with a smoothness constraint, the segmentation gains robustness with regards to handling image noise and boundary gaps, compared to regular edge detection. For locating the SAT layer boundaries, this is an important property, since the rim is not perfectly continuous on all images. As noted in section 2.2, the anterior part of the rim can be pierced by the umbilicus, and on lean patients the layer can appear almost invisible in other parts of the image.

In [14] a dynamic formulation is given for minimising the energy formulation of (5.1), treating the curve $\mathbf{X}(s)$ as a function of time t deforming under the influence of forces relating to the two energy components, s is omitted from the

notation for simplicity:

$$\gamma \frac{\partial \mathbf{X}}{\partial t} = \mathbf{F}_{int}(\mathbf{X}) + \mathbf{F}_{ext}(\mathbf{X}) \quad (5.2)$$

The coefficient γ is added to make the units on both sides consistent. The idea is that the two forces should cancel out and thereby stop the deformation, when reaching the energy minimum. Formulations for these forces are described in the following.

With s being the natural curve parametrisation, the internal force $\mathbf{F}_{int}(\mathbf{X})$ of (5.2) is given by:

$$\mathbf{F}_{int}(\mathbf{X}) = \frac{\partial}{\partial s} \left(\alpha \frac{\partial \mathbf{X}}{\partial s} \right) - \frac{\partial^2}{\partial s^2} \left(\beta \frac{\partial^2 \mathbf{X}}{\partial s^2} \right) \quad (5.3)$$

The first-order derivative of the curve with respect to the curve parametrisation, $\frac{\partial \mathbf{X}}{\partial s}$, represents the degree of stretching, and the second-order derivative $\frac{\partial^2 \mathbf{X}}{\partial s^2}$ represents the curvature. The weighting parameters, α and β , are used for controlling the contribution of the two terms; a high α prevents the curve from stretching, while a large β makes the curve behave more rigid during the deformation.

A common choice for formulating the *image driven* force, $\mathbf{F}_{ext}(\mathbf{X})$, is to derive it from a *Gaussian potential energy* function, P , defined for the entire image domain. For point (x, y) in image \mathbf{I} this energy can be defined by means of the image gradient magnitude:

$$P(x, y) = -w_e |\nabla_{\sigma} \mathbf{I}(x, y)|^2 \quad (5.4)$$

where w_e is a positive weighting parameter. $\nabla_{\sigma} \mathbf{I}$ is the gradient operator applied on \mathbf{I} filtered with a Gaussian kernel of standard deviation σ – chosen to match the scale of the features of interest, and to smoothen out the effect of possible image artifacts. From (5.4) the corresponding external force component for a point \mathbf{X} on the curve can be written as the derivative:

$$\mathbf{F}_{ext}(\mathbf{X}) = -\nabla P(\mathbf{X}) \quad (5.5)$$

Figure 5.1 illustrates the behavior of \mathbf{F}_{ext} , by a vector field superimposed on the corresponding part of an image slice; a corner of the SAT layer on a corrected image. The vectors clearly point towards the interior and exterior boundary of the SAT layer. The figure also illustrates an issue related to this force formulation. Since both intensities outside the abdomen and on the SAT layer generally appear homogeneous, the magnitude of the gradient is very low in these areas.

This means, that no gradient information is available to guide the deformation, and thus the curve can get stuck if passing through here during the evolution. Using a larger Gaussian kernel with a larger σ in (5.4) can resolve this, by increasing the attraction rate of \mathbf{F}_{ext} , but this will also give a less accurate result, as the edge location becomes less distinct by the blurring.




Figure 5.1: The external force of (5.5) illustrated as a vector field superimposed on the corresponding part of a corrected image. To avoid clustering, the vectors are only shown for a grid of every fifth voxel. The image is a zoom of the top right corner of the abdominal region showing the high intensity SAT layer. It is observed, that \mathbf{F}_{ext} points towards both the interior and exterior boundary, but also exhibits close to zero magnitude outside the abdomen and inside the thicker part of the adipose tissue.

To solve this problem [14] presents the use of *multi scaled Gaussian potential forces*. The idea is to segment the image filtered with a kernel with a large σ to begin with. Then, when the model reaches its energy minimum, the image is segmented again using the previous result as initialisation and filtering with a smaller σ – repeating this procedure until the wanted accuracy of the boundary is reached.

The segmentation process should be fully automatic – also including the initialisation of the curve. For segmenting the boundary of the abdomen, this can be handled by placing a circular curve of equally spaced points positioned outside the abdomen, then allowing it to shrink to fit the boundary using the sufficiently blurred images.

However, using the Gaussian potential force formulation of (5.5) together with an automatic initialisation poses problems when segmenting the interior boundary of the SAT layer. Placing the model inside the abdomen to expand during

the deformation is not feasible, since the VAT can exhibit distinct edges, which can hold back the deformation from reaching the true boundary. Another solution could be to initialise the curve by means of the abdomen boundary, and then shrinking to fit the interior boundary. This requires some additional scheme to make the curve to *jump* past the exterior boundary, while still enabling (5.5) to stop the deformation at the interior boundary. Methods for handling this are not straightforward, and therefore the Gaussian potential force formulation is abandoned in this work. Instead another external force formulation is adopted for segmenting both the interior and exterior boundaries of the SAT layer.

The adopted external force formulation utilises the prior knowledge of intensity homogeneity; within the high intensity voxels of the SAT layer in the corrected images, and within the low intensity voxels outside the abdomen – both present in the original and corrected images. Furthermore the formulation can exploit information about tissue related intensities by means of the membership values gained from the classification process described in chapter 4 – this is described in the sections covering the application of the model. The new external force formulation is formed by two terms:

$$\mathbf{F}_{ext}(\mathbf{X}) = \mathbf{F}_d(\mathbf{X}) + \mathbf{F}_i(\mathbf{X}) \quad (5.6)$$

where \mathbf{F}_d is a *deflation* force, and \mathbf{F}_i is an *impurity* force, the two terms are presented in the following.

In 1991 Cohen [1] proposed the *balloon* model; adding a deflation force as an extension to the Gaussian potential force – to make the deformation less dependent on an initialisation within the gradient’s attraction range. In this work the deflation force will be combined with another image driven force to segment both the abdomen boundary and the interior SAT layer boundary. The deflation force for a point \mathbf{X} on the curve is defined as

$$\mathbf{F}_d(\mathbf{X}) = w_d \mathbf{N}(\mathbf{X}) \quad (5.7)$$

where w_d is a positive weighting parameter and $\mathbf{N}(\mathbf{X})$ is the inward unit normal for the curve at point \mathbf{X} . A deformation using this force alone will cause the curve to contract like deflating a balloon – hence the name. Selecting $w_d < 0$ will expand the curve instead. Figure 5.2 illustrates the nature of the deflation force as vectors drawn for every fifth point of a curve initialised for segmenting the abdomen boundary. The segmentation corresponds to the one shown earlier in figure 3.10 for the point sampling of the bias field correction (slice 16 from patient 7). This image will be used as an illustrative example throughout the section – the details of segmenting the abdomen boundary are elaborated in section 5.1.1.




Figure 5.2: The deflation force shown as vectors superimposed on the initialisation curve used for segmenting the abdomen boundary in an uncorrected image – here the image has been filtered with a Gaussian kernel to reduce the effect of artifacts.

The idea behind (5.7) is that the second force term of (5.6), \mathbf{F}_i , should be designed to neutralise the deflation by means of an appropriate image feature. The formulation adopted exploits the assumption of homogeneous regions to define \mathbf{F}_i . The force is directed opposite of \mathbf{F}_d , and the magnitude is defined by the intensities of the area enclosed by the initial position and the current position of the curve; the force magnitude is zero for a *pure* area, and grows proportional to the amount of *impure* voxels covered. The impurity is defined as an intensity threshold. The impurity force \mathbf{F}_i for point \mathbf{X} on the curve is formulated as:

$$\mathbf{F}_i(\mathbf{X}) = -w_i \mathbf{N}(\mathbf{X}) \int_{z=0}^{z=1} H\{k(\Theta - I(\mathbf{X}_{between}))\} dz \quad (5.8)$$

The force is directed against the inward normal, and weighted by the positive parameter w_i . H is the Heaviside function measuring the impurity based on an intensity threshold Θ and the image intensity I at the location $\mathbf{X}_{between}$ defined as

$$\mathbf{X}_{between} = (1 - z)\mathbf{X}_{outline} + z\mathbf{X} \quad (5.9)$$

For a point, the integral is performed over the line between the position on the initial curve ($z = 0$) and the current position ($z = 1$). In (5.8) $H\{n\} = 1$ for positive n , and 0 for negative. This way $k > 0$ results in Θ functioning as a lower intensity threshold defining the impurity, while $k < 0$ imposes an upper threshold. The magnitude of k is not important, only the sign.

Figure 5.3 shows the principle of calculating the impurity force for a point on the curve deforming from the initialisation shown on figure 5.2. By (5.8) and (5.9) the magnitude of the impurity force for a point is proportional to the percentage of impure voxels encountered when tracking a straight line between the point on the two curves ($z = 0$ and $z = 1$). With an appropriate value of Θ the Heaviside function only triggers, when the curve approaches voxels near the abdomen boundary – due to the homogeneity of the low intensity voxels – causing the deformation to stop.




Figure 5.3: Segmenting the abdomen boundary in an uncorrected slice; an illustrative example to explain the impurity force calculation. **(a)** The blurred image subject to the segmentation. The blue curve shows the initialisation ($z = 0$). The red curve is the final segmentation ($z = 1$). The green line is drawn between the same point on the two curves to illustrate how $\mathbf{X}_{between}$ is calculated. **(b)** The intensities forming $\mathbf{X}_{between}$ for $z \in [0, 1]$. **(c)** The corresponding values of the Heaviside function with $k < 0$; white indicates $H = 1$, black indicates $H = 0$.

Sections 5.1.1 to 5.1.3 describe how the active contour model is applied to segment the boundary of the abdomen and the interior boundary of the SAT layer. In section 5.1.4 some resulting boundary segmentations are presented and the method is evaluated.

5.1.1 Rough Abdomen Boundary

The first application of the active contour model is for achieving a rough estimate of the abdomen boundary, used to limit the the point sampling of the bias field correction, as described in section 3.2.2.

For this task the image data consists of original intensities with the rectangular masking applied (section 3.1). This rectangle is used for automatically initialising the model; the curve is initialised as a circle of equally distributed points, with center at the center of the rectangular ROI. The radius is chosen as the distance from the center to the corners – this way the entire abdomen is ensured to be enclosed by the initial curve. Figure 5.4 shows the principle.




Figure 5.4: Example of the automatic initialisation (blue). The green box and cross is the rectangular ROI and center used for determining the size and position of the initialisation.

During the evolution, the model must deform from the initial circular shape to the more oval boundary of the abdomen. The surface of the abdomen is generally smooth, but a high degree of flexibility is required to enable the deformation from the smoother circle. This means, that low values of α and β must be used for the internal force in (5.3).

For the two external force components in (5.6) the size of the weights w_d and w_i should generally be chosen, such that the curve reaches the force equilibrium at the wanted position. The magnitudes of \mathbf{F}_d and \mathbf{F}_i should be equal, when enclosing a sufficient amount of impure intensity voxels – providing the strength to stop the deflation. In practice the exact equilibrium of the forces is seldom

reached, and instead the deformation is terminated, when the difference between the forces is below a suitable tolerance.

For segmenting the abdomen boundary, the impurity measure of (5.8) is defined as an upper threshold on the dark intensities enclosed by the curves ($k < 0$). With regards to selecting a good value of Θ the uncorrected image data alone provides no information on the exact correspondence between intensities and tissue type (adipose or non-adipose). Therefore a choice has been made to determine one common value by means of a visual inspection of the entire image data set. On one hand, this value should not be set too low. In spite of the rectangular ROI masking, some slices feature bright artifacts in voxels just outside the abdomen – potentially stopping the deformation. On the other hand, if using a too high value, the model could end up cropping some of the SAT where low intensities are present. Furthermore, to reduce the dependency of an accurate Θ , and as a measure to make the deformation more robust with regards to image artifacts, the image is filtered with a Gaussian kernel. As noted earlier this makes the boundary location more inaccurate, but at this point it is not of great importance, since spurious points placed outside the abdomen will be removed as described in section 3.2.2.

Figure 5.5 shows steps of the curve deformation to segment the abdomen boundary. The green vectors illustrate the behavior of \mathbf{F}_{ext} ; \mathbf{F}_d contracts the curve until \mathbf{F}_i stops the evolution at the boundary.

To determine a set of parameter values giving good overall results, a parametric study has been performed by means of visual inspection of the resulting boundaries. The goal was, to find values resulting in abdomen boundaries not cropping the SAT, but still being accurate enough for the point sampling method to be able to trim the spurious points. The values $w_d = 1$, $w_i = 15$, $\alpha = 0.01$, $\beta = 0.01$ and $\Theta = 75$ have been found to result in good boundary estimates. Figures 5.6 to 5.8 illustrate the principle of the study, by showing segmentations resulting from varying the parameters one by one. It should be noted, that since the parameter space is high dimensional, and the parameter effects are correlated (the segmentation relies on reaching an equilibrium of the sum of forces), a certain parameter value, shown to corrupt the segmentation, can give good results – if the other parameters are adjusted accordingly. That is, the results shown should be regarded as illustrations of the general effect of each parameter.




Figure 5.5: An example of the curve deformation to achieve a rough estimate of the abdomen boundary in an image slice blurred with a Gaussian kernel. (a) to (e) show the deformation by means of *snapshots* taken at different stages of the evolution. The curve was initialised as shown on figure 5.4. The red curve shows the curve position. The green vectors illustrate the external force effect for every fifth point on the curve. The vectors are scaled within each image, and the magnitudes should not be compared between images. (f) The final segmentation result shown on the image without the Gaussian blur.




Figure 5.6: Example segmentations to illustrate the effect of varying the weights of the deflation and the impurity force. **(a)** Red; $w_d = 0.5$, green; $w_d = 1$ and magenta; $w_d = 5$. For $w_d = 0.5$ the segmentation fits more loosely around the abdomen, as \mathbf{F}_i needs less impure voxels to neutralise \mathbf{F}_d . $w_d = 5$ results in a clearly erroneous segmentation; \mathbf{F}_i can only stop the deflation, where a thick layer of SAT is present to trigger the Heaviside function. **(b)** Red; $w_i = 5$, green; $w_i = 15$ and magenta; $w_i = 30$. Since the effect of w_d and w_i are intertwined in (5.6), using $w_i = 5$ results in an erroneous fit similar to the one caused by $w_d = 5$. The same similarity is observed between $w_i = 30$ and $w_d = 5$.




Figure 5.7: Example segmentations to illustrate the effect of varying the elasticity and rigidity constrained in the internal force. **(a)** Red; $\alpha = 0.001$, green; $\alpha = 0.01$ and magenta; $\alpha = 0.1$. Very little difference is observed between using $\alpha = 0.001$ and $\alpha = 0.01$. Using $\alpha = 0.1$, the elasticity constraint drowns out \mathbf{F}_{ext} and produces a very smooth but highly erroneous segmentation. **(b)** Red; $\beta = 0.001$, green; $\beta = 0.01$ and magenta; $\beta = 0.05$. Applying $\beta = 0.001$ and $\beta = 0.01$ has effects similar to the observed for α . The reason for not choosing the lower parameter values is that they make the model less robust towards segmenting discontinuities in the boundary. $\beta = 0.05$ results in the curve cropping some of the SAT in the left side, where the abdomen boundary is less smooth.




Figure 5.8: Example to illustrate the effect of varying the intensity threshold of the impurity force. Red; $\Theta = 25$, green $\Theta = 75$ and magenta; $\Theta = 250$. Selecting $\Theta = 25$ results in a loose fit, especially in the anterior part of the abdomen, where a small degree of high intensity artifacts are present outside the abdomen, even after the rectangular ROI masking. $\Theta = 250$ creates a very tight fit cropping some SAT in the anterior part of the abdomen, where the SAT layer exhibits lower image intensities.

5.1.2 Abdomen Boundary

The quantity of adipose tissue should be presented as percentages of the total abdominal volume. To get a more accurate value of the volume, the rough boundary estimate found in the previous section is refined. The choice of $\Theta = 75$ was a crude approximation made to fit all images, thus the improvement of the boundary estimate is sought by determining individual values of Θ for each image series.

After the bias field correction and the tissue classification, information on the correspondence between adipose tissue voxels and image intensities is present by means of the membership values u_{jk} . Compared to intensities (both original and corrected) the values, $0 \leq u_{jk} \leq 1$, provide an intuitive correspondence between voxels and tissue class memberships – comparable between image series regardless of within-class intensity ranges.

This enables determining Θ in a generic fashion – similar to assigning labels by means of ε_m in (4.5): Instead of specifying Θ as an image intensity, it is specified as a membership value, which can then be transformed into the corresponding image intensity for each image series individually. The procedure can be viewed as segmenting images consisting of membership values instead of image intensities – this is the angle of approach for the following.

Since the abdomen boundary is characterised as a nearly continuous ring of SAT, it is regarded as the exterior SAT boundary. In section 4.2 a value of

$\varepsilon_m = 0.5$ was found to be a suitable membership value threshold for classifying voxels as SAT – thus choosing $\Theta = 0.5$ for the more accurate segmentation of the abdomen boundary seems rational.

For this segmentation the corresponding rough boundary estimate is used as initialisation. The image data subject to the segmentation now consists of membership values, and thus Θ is specified accordingly – as noted above. All other parameters remain unchanged, since the shape of the abdomen does not differ significantly from the one segmented of the blurred image slice. Figure 5.9 shows an example of the resulting segmentation – the curve is superimposed on two representations of the same image slice; the membership values and the corrected intensities.




Figure 5.9: Example of the more accurate abdomen boundary segmentation; the red curve is the initialisation, the magenta curve is the result. **(a)** The curves are superimposed on the membership values of the adipose tissue class. **(b)** The two curves superimposed on the corrected image.

5.1.3 Interior SAT Boundary

The last application of active contours is to separate the region containing VAT from the SAT layer. Once again within-class intensity homogeneity is exploited, as voxels of the SAT layer exhibit homogeneous (high) intensities after the bias field correction – and thus also similar membership values.

The procedure is to initialise the curve on the exterior boundary and from there deforming to fit the interior boundary by shrinking the curve. This behavior is achieved by defining the impurity measure using $\Theta = 0.5$ as a lower threshold ($k > 0$).

Figure 5.10 shows the resulting segmentation on the recurring example slice.



Figure 5.10: Segmentation of the interior SAT boundary. The red curve is the abdomen boundary used as initialisation. The magenta curve is the resulting interior SAT boundary estimate.

5.1.4 Evaluation

Before proceeding with the SAT layer subdivision a short evaluation of the method is presented.

The method is completely automatic, and the formulation of internal forces makes the segmentation more robust towards boundary discontinuities compared to regular edge detection schemes. Furthermore the method adopts information on the classified tissue in each individual image series by using the membership values.

The resulting segmentations seem to be dependent on the parameters selected to some extent. Thus to achieve overall good results across the large variety of boundary shapes present in the image series, a trade-off on the individual slice accuracy has been necessary.




Figure 5.11: Examples of the boundary segmentation. **(a)** Segmenting the interior SAT boundary, some errors occur when the VAT appears to be closely connected to the anterior SAT layer. A choice could have been to add further constraints on the flexibility of the deformation. However, this would prevent a close fit to the curvy shaped interior boundary featured in the posterior on slices with more deep SAT. The parameters chosen for the model should therefore be regarded as a trade-off to provide overall good results. **(b)** The discontinuity in the SAT layer at the umbilicus is handled in the abdomen boundary segmentation. The curvy fashion of the interior SAT boundary in the anterior part is also fitted accurately.

5.2 Dynamic Programming

In this section dynamic programming (DP) is presented as a method for subdividing the SAT layer into deep SAT and superficial SAT by means of locating Scarpa's fascia (SF) anatomically partitioning the layer. The general theory of DP is presented below, in section 5.2.1 the application is introduced, and in section 5.2.2 the method performance is evaluated.

Dynamic programming (DP) is a technique for finding optimal paths in graphs. A thorough general presentation of graphs and the theory of DP is done by Cormen et al. in [2]. In [3] Dawant et. al present DP as a method for boundary detection in medical images.

A graph consists of a set of points, *nodes*, and a set of *links* connecting the nodes. A path through the graph is defined as a set of links, that connects a *start* node to an *end* node.

DP works on graphs that are directed (the links are one-directional) and acyclic (no loops are present; a node can only be visited once on a path through the graph). Figure 5.12(a) shows an example of a directed acyclic graph. On the figure the nodes N are aligned in a grid; node $N(i, j)$ is located at row i and column j . Note that the links of the graph are constructed such that a node $N(i, j)$ only is connected to nodes in row $i + 1$ positioned at columns $j - 1$, j and $j + 1$. In this type of graph the set of possible start and end nodes are defined as the top and bottom row respectively; a path is a set of links leading from top to bottom

Assigning a static cost value to each node, an optimal path is defined as the path traversing links connecting nodes with the lowest overall sum of cost values (for other applications of DP an optimal path with maximum cost can be desirable, but here paths of minimum cost are regarded). These associated static cost values can be represented in a matrix structure N_s as is shown in figure 5.12(b), with element $N_s(i, j)$ corresponding to node $N(i, j)$.

In general, DP solves a problem by splitting it into several subproblems, and solving them recursively to form the overall solution. For the optimal path-problem this is done is by recursively defining an optimal path by means of a cumulative cost matrix N_c of the same size as N_s . The top row of the cumulative cost matrix is defined as the first row of the static cost matrix

$$N_c(1, j) = N_s(1, j) \quad (5.10)$$




Figure 5.12: **(a)** An example of a directed acyclic graph. The two nodes are labelled to illustrate the indexing; I is $N(1, 1)$, and II is $N(2, 3)$. **(b)** The static cost matrix corresponding to the graph.

For the rest of the rows the cumulative cost of a node $N(i, j)$ is calculated as

$$N_c(i, j) = \min(N_c(i-1, j-1), N_c(i-1, j), N_c(i-1, j+1)) + N_s(i, j) \quad (5.11)$$

The values are calculated row by row from top to bottom. This way the sub-problems are defined as recursively determining an optimal path through node $N(i, j)$ in row i for all columns j – using the previously computed optimal path costs N_c in row $i-1$ and the static cost of the particular node $N_s(i, j)$. Figure 5.13(a) shows the cumulative cost matrix computed for the graph and static cost matrix from figure 5.12. The overall cost of the optimal path can be read off as the lowest value of the bottom row in N_c . To enable a reconstruction of the path resulting in this cost, each calculation of a value N_c for a node in row i is accompanied by a pointer to the predecessor j with the minimum $N_c(i-1, j)$. Figure 5.13(b) shows this map of predecessors by means of arrows pointing either left, straight or right. The optimal path can then be found by backtracking from the lowest value in the bottom row following the predecessor map to the first row. On figure 5.13(b) this path is marked by shaded fields.

If regarding pixels as nodes, DP can be applied to 2-dimensional image data. By using the image intensities as cost values, an optimal path traverses the lowest overall intensities going from top to bottom of the image. Furthermore the link-structure illustrated in figure 5.12(a) corresponds to paths conforming to 8-connectivity, a suitable choice for finding smooth paths through images.

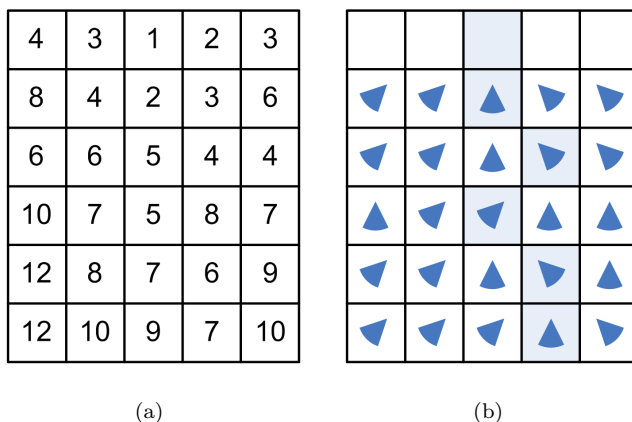


Figure 5.13: (a) The cumulative cost matrix computed for the graph and static cost matrix of figure 5.12. (b) The predecessor map corresponding to the cumulative cost matrix displayed as arrows. The optimal path is reconstructed by backtracking through the predecessors from the node with the lowest value in the bottom row of the cumulative cost matrix (the shaded fields of the matrix).

The DP scheme presented above finds a path by searching through each row, and the course of the path is constrained – by the selection of start and end node sets – to run from the top row to the bottom row of the image. This makes the algorithm work like a 1-dimensional search, and is thus not applicable for tracing the round shape of SF. To make the scheme usable, a spatial transformation must be applied to the image data. The following section presents such a transformation and further details on the application of DP to subdivide the SAT layer.

5.2.1 Subdividing the SAT layer

On the image data SF appears as a low intensity cleft in the high intensity SAT, and the smooth round shape generally follows the shape of the exterior SAT boundary. When the SAT layer is thin, SF coincides with the interior boundary, and thus can be impossible to make out. In these areas all SAT should be classified as superficial SAT.

In the spatial transformation of the image data, the shape of the fascia is exploited by means of applying a polar transformation. This is done by sampling

image intensities along spokes of equal length, evenly distributed on a circle. Since the coordinates of points located on the spokes generally are located between the voxel centers of the original image, intensities of the transformed image are sampled using interpolation. This way each spoke of intensities makes a row in the transformed image. The spokes are sampled inwards, such that the first columns of the transformed image corresponds to the part of the spokes furthest away from the center. Figure 5.14 shows the principle of the polar transformation, the example image will be used throughout this section (patient 7 slice 12). The first row of the transformed image corresponds to the red spoke on figure 5.14(a) – the rest are sampled by moving anticlockwise from here. The




Figure 5.14: **(a)** Illustration of the polar transformation; sampling voxels on spokes equally distributed on a circle. For illustrative purposes only 20 spokes are drawn on the image. The first spoke sampled is marked in red, the rest are sampled proceeding anticlockwise from here. **(b)** The resulting spatially transformed image; each row corresponds to a spoke. Note that the size of the two images can not be compared. The image is slice 12 of patient 7.

transformation is controlled by 3 parameters; the circle center (x_c, y_c) , the circle radius r and the *angular resolution* n_s (the number of spokes sampled). (x_c, y_c) and r are determined in the same way as the automatic initialisation of the active contour model in section 5.1.1 – ensuring that the entire abdomen is inside the circle. On one hand, n_s should not be set too low, as this will prevent the method from retrieving an accurate estimate of the fascia. On the other hand, if using too many spokes the path becomes more vulnerable with regards to image noise and smaller discontinuities. A value of $n_s = 200$ was found to be sufficient. In the example shown in figure 5.14 the length of the spokes are

automatically determined as $r = 183$.

After the transformation SF runs from top to bottom through the SAT layer in the unfolded image. However, applying the DP scheme at this point, would just result in finding an optimal path through the zero intensities outside the abdomen – only the SAT layer is of interest.

By combining the boundaries found in sections 5.1.2 and 5.1.3 with the labels determined in chapter 4, a binary mask can be created for extracting the SAT layer (the details of this method are covered in chapter 6, and thus only the result is presented here). Figures 5.15(a) and 5.15(b) show the image only containing the SAT layer before and after the spatial transformation. To remove the zero-intensity voxels outside the abdomen, the image is *flattened* by removing all these voxels, figure 5.15(c) shows the result. To prevent the path from entering the area of removed voxels inside the abdomen, the intensities here are set to high values (e.g. 100, compared to the corrected image intensities reaching just above 1.5), this way making it very unfavorable for the path to enter. An example of this high intensity padding is not shown, as the scaling with respect to intensity range prohibits a good visualisation.

Since SF is characterised as a closed curve on the images, it is desirable for the path to be connected at the start and end nodes (the same position in the top and bottom row of the matrix). This is enforced by means of an overlap in the image; padding with repetitions of rows in both top and bottom. Figure 5.16(a) illustrates this padding.

The image is now ready for applying the DP scheme described in the previous section. Figure 5.16(b) shows the resulting curve superimposed on the transformed image with the repetitions. On this particular slice the added overlap has little effect, since the SAT layer is very thin around the start and end nodes, but this could play a role for slices with more SAT. It is observed that the path generally follows the fascia; only in the bottom part the path appears to diverge a little from the expected, since the dark cleft is a bit faint.


The overlapping points are removed from the path, and the remaining are transformed back to the original image space. Figure 5.17 shows a zoomed view of the corrected image slice and the resulting curve; subdividing the SAT layer into deep SAT and superficial SAT. The curve resembles SF well, and in the anterior part – where the fascia coincides with the interior SAT boundary – almost no deep SAT is identified as intended.



(a)



(b)



(c)

Figure 5.15: (a) The masked image only containing the SAT layer of image of figure 5.14(a). (b) The spatially transformed image. (c) The transformed image after flattening with respect to the abdomen boundary.




Figure 5.16: (a) The transformed image with 20 rows of repetition added at both the top and bottom, the horizontal blue lines mark the extent of the original transformed image without the overlap. (b) The optimal path found by applying DP shown on the same image.




Figure 5.17: (a) A zoomed view of the corrected slice subject to the SAT layer subdivision. (b) The same image with the resulting curve superimposed.

5.2.2 Evaluation

This section finishes the SAT layer subdivision by means of some comments on the method and an example.

Applying DP, using the restrictive link structure and the polar image transformation, appears to be a suitable choice to match the generally smooth shape of SF and the abdomen. The method results in overall accurate segmentations without the use of data-dependent parameters, proving a high degree of robustness.

In the image data of this study, there seems to be a connection between the SAT thickness and the visibility of SF: In the thin parts of the layer, SF coincides with the interior SAT boundary and is not distinguishable. This is especially observed in the anterior part of the abdomen, since most of the patients feature little SAT here. Even on the patients with the largest quantity of anterior SAT, SF can be hard to determine visually.

In areas of the SAT layer without SF visible, the subdivision should track the interior SAT boundary. This is also achieved on most of the slices, but on some slices, the shape of the SAT layer hinders a correct segmentation. With the polar unfolding and the flattening with respect to the abdomen boundary, the right boundary of the transformed image reflects the shape of both boundaries. This means, that the structure of the links (restricting the change in columns to one position on each side between neighbouring rows) becomes insufficient for tracking the true path. Figure 5.18 illustrates the issue in the anterior part of the slice.

After consulting Kasper Pilgaard, a decision was made not to change the segmentation method, as it generally performed very well on the rest of the SAT layer. Instead the decision was to disregard the subdivision in the anterior SAT, since very little true deep SAT would be overlooked this way in the image data of this project. This is implemented in the next chapter, where the region boundaries are combined with the tissue labels.

Showing the results to Kasper Pilgaard gave a positive response, as he considered the concept of SAT layer subdivision to be in the *front line* of the research field.




Figure 5.18: Illustration of an erroneous subdivision in the anterior SAT layer. **(a)** The corrected intensities with superimposed curves for the interior SAT boundary (magenta) and the subdivision (red). **(b)** The transformed SAT layer, the VAT region has been partially cropped for the visualisation.

5.3 Results

This section concludes the region identification with a short overview and comments on the methods applied. Figures 5.19 and 5.20 show the region boundaries segmented on the recurring 6 slices.

In this chapter two methods were developed to identify the three regions containing the different types of adipose tissue considered.

For identifying the abdomen boundary and the interior SAT layer boundary a variant of active contours was developed: A deformable model incorporating a measure of robustness towards segmenting discontinuous boundaries. The specific formulation applied in this work incorporated prior knowledge of the homogeneous low intensities caused by the air outside the abdomen, and the homogeneous high values exhibited by the corrected intensities on the SAT layer. Furthermore the patient-specific relationship between intensities and tissue types was taken into account by means of the membership values derived in the tissue classification. Segmenting the interior SAT boundary, the method appeared to be relatively dependent of the parameter values specified, and therefore, in order to obtain overall good segmentations of the the variety in shape and thickness of the SAT layer, a trade-off was made. If the interior SAT layer boundaries had been more homogeneous (both between patients and within patients) more accurate segmentations would have been possible. To add more robustness to the method one could also have considered incorporating the 3-dimensional structure present in the image data, by extending the

active contours to active surfaces. However, the implementation of this is less straightforward in comparison to the 2-dimensional approach, and thus such a solution was not attempted in this work.

The subdivision of the SAT layer into deep and superficial SAT was performed using the concept of dynamic programming and a polar transformation of the image data. The polar transformation was developed to incorporate knowledge on the general shape of the abdomen and the smooth behaviour of Scarpa's fascia. This method provided good subdivisions across the entire image data set, using only the image transformation and link structure as a data-dependent choice – proving to be suiting across the large variety in SAT layer appearances.

Considering that both quantities of VAT and deep SAT have a reported connection to IR [12], the measure of deep SAT has an advantage towards measuring VAT on the data in this project. On the present image data, the VAT generally exhibits lower contrast and more indistinct transitions to the surrounding non-adipose tissue, making an accurate visual validation difficult – in comparison to the deep SAT, where the appearance of SF generally is well-defined; either clearly visible on the SAT layer or coinciding with the interior SAT boundary.

Presenting the segmented boundaries to Kasper Pilgaard confirmed the authors perception of overall good results.

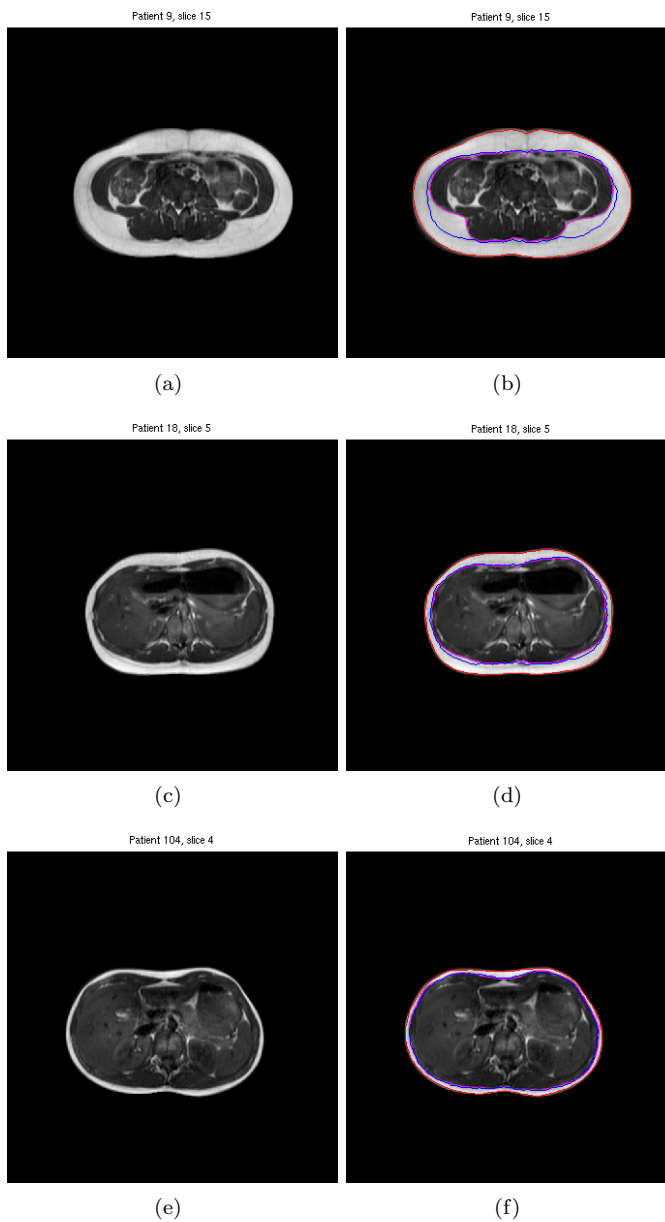


Figure 5.19: The first column holds images with corrected intensities, the second holds the same images, with the segmented region boundaries superimposed. The rows are; patient 9 slice 15, patient 18 slice 5, and patient 104 slice 4.




Figure 5.20: The first column holds images with corrected intensities, the second holds the same images, with the segmented region boundaries superimposed. The rows are; patient 22 slice 3, patient 1 slice 2, and patient 4 slice 16.

CHAPTER 6

Combining Labels and Boundaries

This chapter describes the last step of the segmentation process, where the results of the tissue classification of chapter 4 and the region identification of chapter 5 are combined to form the final result; segmenting images by dividing voxels into four classes, superficial SAT, deep SAT, VAT, and non-adipose tissue. To merge the labels and boundaries boolean set operations are applied. The procedure is explained in the following.

6.1 Binary Masks and Set Operations

For distinguishing between the three types of adipose tissue, three boundaries have been identified using active contours and dynamic programming; the abdomen boundary, the boundary subdividing the SAT layer and the interior SAT boundary. Here the abdomen boundary is also regarded as the exterior boundary of the SAT layer.

6.1.1 Binary Masks

The located boundaries are represented by sets of points forming closed curves. For each boundary, a binary mask is created. The mask is of the same size as the corresponding image, and is constructed by assigning the value 1 to all voxels enclosed by the respective curve, and the value 0 to voxels located outside. The three resulting masks are denoted M_{ext} , M_{sub} and M_{int} respectively. Figure 6.1 illustrates the principle on a single image slice, the slice will be used as an example through out the section.

As noted in section 5.2.2, the anterior part of the SAT layer should not be subdivided, instead all adipose tissue voxels should be classified as superficial SAT. This is implemented by locating the left and right edge of M_{int} , and setting all voxels above the abdomen between these vertical boundaries to 1, creating a new mask, M_{ant} . Figure 6.2 shows the resulting mask.

6.1.2 Set Operations

Masks for the three regions containing the three types of adipose tissue can now be created by means of boolean operators applied to M_{ext} , M_{sub} , M_{int} and M_{ant} .

As an intermediate calculation, a mask for the entire SAT layer, M_{SAT} , is computed by means of the exclusive operator, xor

$$M_{SAT} = \text{xor}(M_{ext}, M_{int}) \quad (6.1)$$

Another intermediate mask M_{antSAT} is created to hold the anterior part of the SAT layer by means of the set intersection operator, and

$$M_{antSAT} = \text{and}(M_{ant}, M_{SAT}) \quad (6.2)$$

The deep SAT M_{dSAT} mask is computed, such that none of the anterior SAT is included

$$M_{dSAT} = \text{and}(M_{sub}, \text{xor}(M_{antSAT}, M_{SAT})) \quad (6.3)$$

The superficial SAT layer M_{sSAT} mask can then be computed as another exclusive set

$$M_{sSAT} = \text{xor}(M_{SAT}, M_{dSAT}) \quad (6.4)$$

The mask containing VAT M_{VAT} is simply defined as

$$M_{VAT} = M_{int} \quad (6.5)$$




Figure 6.1: (a) The corrected image with the three identified region boundaries superimposed. (b) The exterior SAT mask, M_{ext} . (c) The subdividing SAT mask, M_{sub} . (d) The interior SAT mask, M_{int} . The image is slice 12 of patient 7.




Figure 6.2: The mask, M_{ant} , created to cancel the subdivision of the anterior SAT layer.




Figure 6.3: (a) The corrected image. (b) The three region masks created by the set operators. Light blue is M_{sSAT} , yellow is M_{dSAT} , and red is M_{VAT} .

Figure 6.3 shows the three regions identified for the masks corresponding to the previous example.

To achieve the resulting segmentation, the three masks are combined with the two sets of binary labelled voxels from the FCM classification, $M_{0.5}$ and $M_{0.85}$, corresponding to the membership value thresholds $\varepsilon_m = 0.5$ and $\varepsilon_m = 0.85$. As noted in chapter 4 applying these locally defined thresholds seem to be an appropriate measure to avoid labelling non-adipose tissue as VAT. Therefore the labels from $M_{0.85}$ are used for classifying VAT and $M_{0.5}$ are used for classifying the two types of SAT. By means of set intersections with M_{sSAT} , M_{dSAT} and M_{VAT} the segmentation is achieved. Figure 6.4(b) shows $M_{0.5}$ and $M_{0.85}$ along with the resulting segmentation, dividing the adipose tissue into the three different types.




Figure 6.4: **(a)** The two sets of labels superimposed; red is $M_{0.85}$ and green is the extra voxels labelled using $M_{0.5}$. **(b)** The resulting segmentation, combining the labels of **(a)** with the region masks of figure 6.3(b); Dark blue is non-adipose tissue and background, light blue is superficial SAT, yellow is deep SAT, and red is VAT.

6.1.3 Evaluation

In this section a short evaluation of the method for combining the boundaries and tissue labels is done. The slice used to illustrate the binary masks and set operations in this chapter, did not feature errors in the anterior SAT layer

subdivision. To illustrate the effect of applying M_{ant} to cancel such errors, the example from last chapter (figure 5.18) is reviewed in figure 6.5.




Figure 6.5: The example from figure 5.18 reviewed to illustrate the effect of applying M_{ant} . **(a)** The corrected image with curves for the interior SAT layer boundary and the SAT layer subdivision superimposed. **(b)** The resulting tissue labelling. The error in the anterior SAT layer subdivision is removed as intended, and only a very small amount of the SAT layer is classified as deep SAT – this corresponds to the visual perception of the corrected image, where Scarpa’s fascia is not visible.

6.2 Results

In this chapter a method for merging the tissue classifications and segmented boundaries was presented. The method works as intended, taking the locally defined thresholding into account – differentiating between SAT and VAT classification to lessen the effect of image noise and low contrast in the VAT region. Furthermore a measure was taken to avoid errors in the subdivision of the anterior SAT layer and worked as intended. Figures 6.6 and 6.7 show the results on the recurring 6 example slices. The final segmentation results are evaluated in the following chapter.

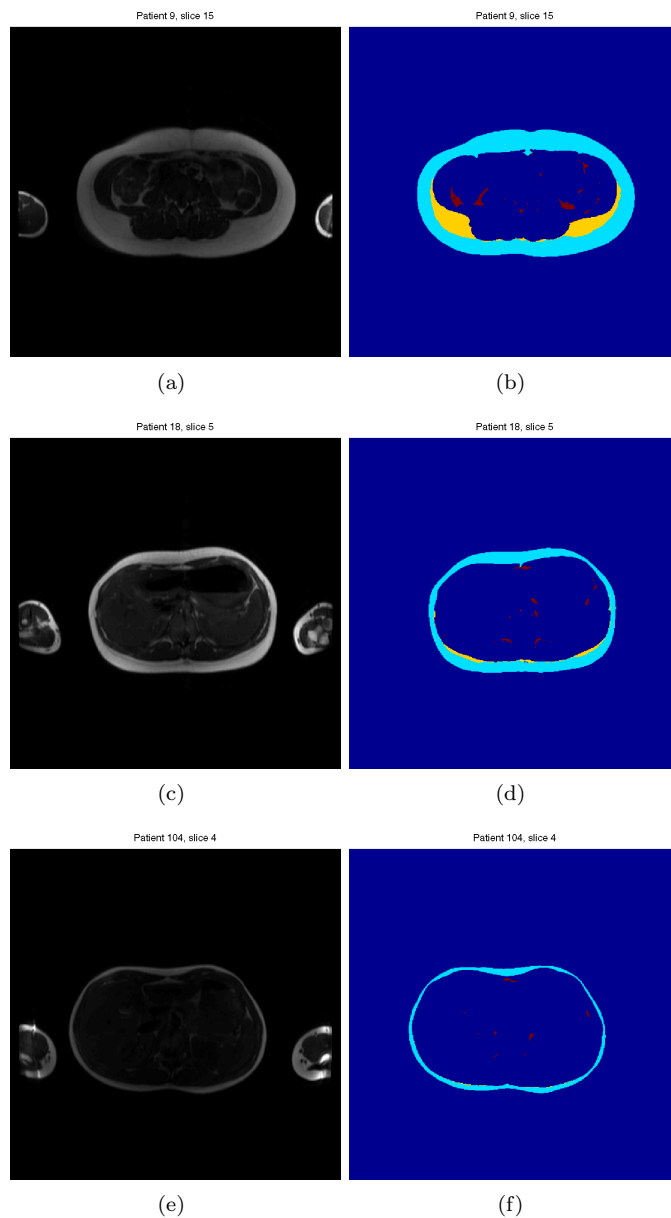


Figure 6.6: Each row holds corresponding original intensities and final segmentation for an image slice. The images are; patient 9 slice 15, patient 18 slice 5, and patient 104 slice 4.




Figure 6.7: Each row holds corresponding original intensities and final segmentation for an image slice. The three images are; patient 22 slice 3, patient 1 slice 2, and patient 4 slice 16.

Final Results

This chapter presents the final results of the abdominal adipose tissue segmentation.

For each patient the quantities of the three types of adipose tissue segmented were calculated as percentages of the total abdomen volume, these results are included in appendix A. Since no ground truth is available to give a measure of the accuracy of the results, the evaluation relies on visual inspection only. As the data set comprises almost 400 slices, it would take up too much space to display them all in the report. Instead they are included on the CD-ROM in appendix D.

To give an overview of the results, segmentations of the top, middle, and bottom slices from all patients are included in appendix B. These three slices allow for a good assessment of the model's performance across the broad spectrum of images in the data set – both when comparing images from different patients and from the same patient. Each segmentation is presented together with the corresponding original image slice, as this allows for an evaluation of the results as a whole – including the preprocessing. The CD-ROM enclosed in appendix D also includes intermediate results from both the preprocessing and the steps of the adipose tissue classification and region boundary segmentation for a closer inspection of the method performance.

The overall impression is, that the quality of the results is quite good, and to a large extent the automatic segmentations resemble the authors perception of the true adipose tissue. The results were presented to Kasper Pilgaard, who confirmed this evaluation.

Generally the FCM algorithm succeeds in deriving the correspondence between intensities and tissue types from the structure of the corrected intensities. A remark could be directed towards the amount of adipose tissue classified in the VAT region, as the method can seem to label too little adipose tissue on lean patients, where the VAT flakes are thinner and more indistinct. However, due to the small amount of contrast present on these slices, the possible inaccuracies are difficult to assess visually.

On the image series from two pairs of twins (23 and 24, and 100 and 101) there is a particular large amount of image noise (possibly caused by the patient breathing during the image acquisition), and the area of the spine exhibits very high intensities compared to the surrounding tissue. This is carried from the bias field correction to the tissue classification and results in some erroneous labelling. However, classifying the spine as adipose tissue, does not appear to be a general issue across the data set, since attention was given to the problem in the bias field correction. After presenting the results to Kasper Pilgaard, he proposed to have these particular patients scanned again with greater attention payed to the image quality. And to see, if the relative high intensities of the spine could have been avoided in the first place. If this is not the case, a possible extension to the method would be to identify the boundary of the spine area in order to perform a correct tissue classification of the area independent of the high intensities.

The results generally exhibit an accurate partitioning between VAT and SAT. The trade-off, for segmenting both the posterior and anterior of the interior SAT boundary equally well, appears to work as intended. The subdivision of the SAT layer also appears to be successful in the majority of the images; correctly classifying close to no voxels as deep SAT on images with thin SAT layers, and locating Scarpa's fascia accurately when it is visible on the images.

Conclusion

The goal of this thesis was to develop an automatic method for segmenting abdominal adipose tissue from 3-dimensional magnetic resonance images. The segmentation should distinguish between three types of adipose tissue; visceral adipose tissue, deep subcutaneous adipose tissue, and superficial subcutaneous adipose tissue.

Before segmenting the adipose tissue, the image data was preprocessed. The arms were removed, such that the images only contained the abdominal region. The within-class inhomogeneities, caused by the bias field effect, were removed, to enable an intensity based tissue classification. The correction was done by estimating the effect retrospectively by means of a *thin plate spline* extended to fit 3-dimensional voxel points of two classes. The points were automatically sampled to provide information from the entire image volume – regardless of the quantity of adipose tissue present. To obtain similar smoothness of the estimates between image series, the rigidity of the fit was controlled by *effective degrees of freedom*. Overall the preprocessing was successful; removing the arms, and making the intensities of adipose tissue appear uniform across the image.

The adipose tissue voxels were classified with respect to intensity. As no ground-truth knowledge was available, the concept of unsupervised classification was adopted by applying *fuzzy c-means clustering*; deriving the natural structure of the data itself to form the basis of the classification. This resulted in a common

reference-frame of membership values comparable between image series. From the membership values the voxels were classified using locally defined thresholds. This was done to account for the slight difference observed between visceral and subcutaneous adipose tissue. The derived membership values made the method completely automatic and it performed well by adapting the labelling to the intensity structure of each individual image series.

To compute the total abdomen volume and to separate the visceral from the subcutaneous adipose tissue, the borders of the subcutaneous adipose tissue layer were located by applying *active contours*. The approach combined a *balloon force* with an image driven *impurity force*, to exploit the intensity homogeneities present after the preprocessing. Image information was generically incorporated by means of the membership values from the tissue classification. The method generally performed well across the variety of boundary shapes in the image data.

The subcutaneous adipose tissue was partitioned into a deep and a superficial part. This approach used the concept of *dynamic programming* and a polar transformation of the image data to exploit the shape of the abdomen and the behaviour of Scarpa's fascia. The method performed well, however some inaccuracies were encountered in the anterior abdomen, where no deep subcutaneous adipose tissue was generally present on the patients in the study. A decision was made to only consider the subdivision in the posterior part of the abdomen.

The assessment of the accuracy of the adipose tissue segmentations was done by visual inspection, and the overall impression was good. The results were presented to Kasper Pilgaard, MD, who confirmed this evaluation.

8.1 Future Work

This section lists some possible areas of future work for the developed method.

- Improving the robustness of the point sampling for the bias field estimation, by determining the thresholds uniquely for each slice or image series.
- Extending the active contours approach to 3-dimensional active surfaces, since volumetric information is available in the image data – adding some robustness to the boundary segmentation.
- Locating the spine area to avoid erroneous labelling of visceral adipose tissue.

APPENDIX A

Volume Results

This appendix contains the distributions of abdominal adipose tissue in each patient volume. The total volume is computed as the sum of voxels inside the abdomen boundary multiplied by the voxel dimensions.

The adipose tissue quantities of the three classes considered are stated as volume percentages, computed as the sum of voxels classified as each class divided by the total sum of voxels in the abdomen. The segmentations used for the computations are the same as the ones shown in B (only using the entire image series, and not just the three slices).

Patient	Total volume (litre)	VAT	deep SAT	superficial SAT
1	7.20	5.34	3.17	15.29
4	5.11	6.08	3.95	16.11
7	4.29	5.77	7.29	21.97
8	3.56	3.23	6.01	19.56
9	6.78	3.18	5.97	30.96
10	5.69	2.19	3.26	24.87
13	6.32	2.89	1.14	12.10
14	5.59	4.48	1.26	14.43
16	4.53	2.28	5.86	20.59
17	5.51	2.86	7.02	33.39
18	4.58	2.33	4.33	26.18
19	5.19	4.30	2.38	11.70
20	5.36	2.88	2.70	12.10
21	6.09	14.31	2.80	12.04
22	7.15	11.83	0.91	11.67
23	4.34	6.40	1.72	12.62
24	5.26	5.03	6.80	27.11
100	3.50	7.50	1.90	9.62
101	4.09	6.25	1.42	10.22
104	5.26	3.99	3.05	16.77

Table A.1: The volume results of the adipose tissue segmentation stated for each of the patients. The first column is the patient ID. The second is the total volume of the abdomen. The last three columns hold the quantity of the three types of adipose tissue considered; VAT, deep SAT and superficial SAT. These values are given as percentages of the total abdominal volume.

APPENDIX B

Segmentation Results

This appendix includes the final segmentation results. For each patient the top middle and bottom slice of the original MRI are displayed along with the corresponding segmentation results. The color coding of the segmentation is: Dark blue; background and non-adipose tissue. Light blue; superficial SAT. Yellow; deep SAT. Red; VAT.




Figure B.1: Final segmentation of patient 1

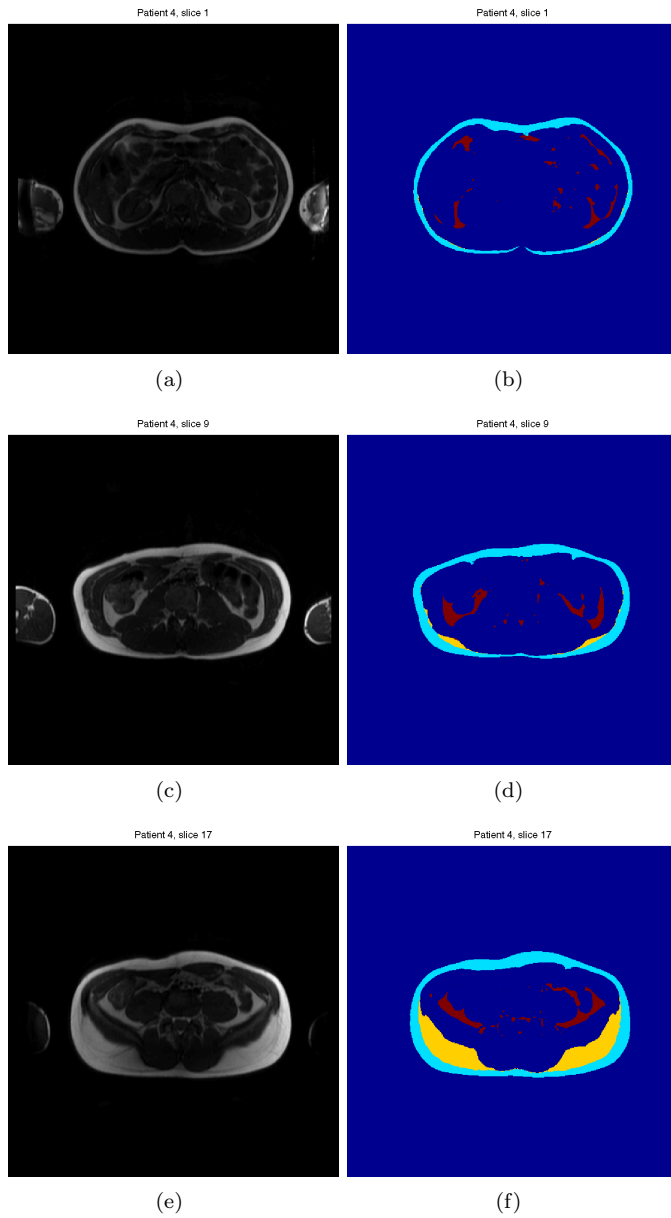


Figure B.2: Final segmentation of patient 4

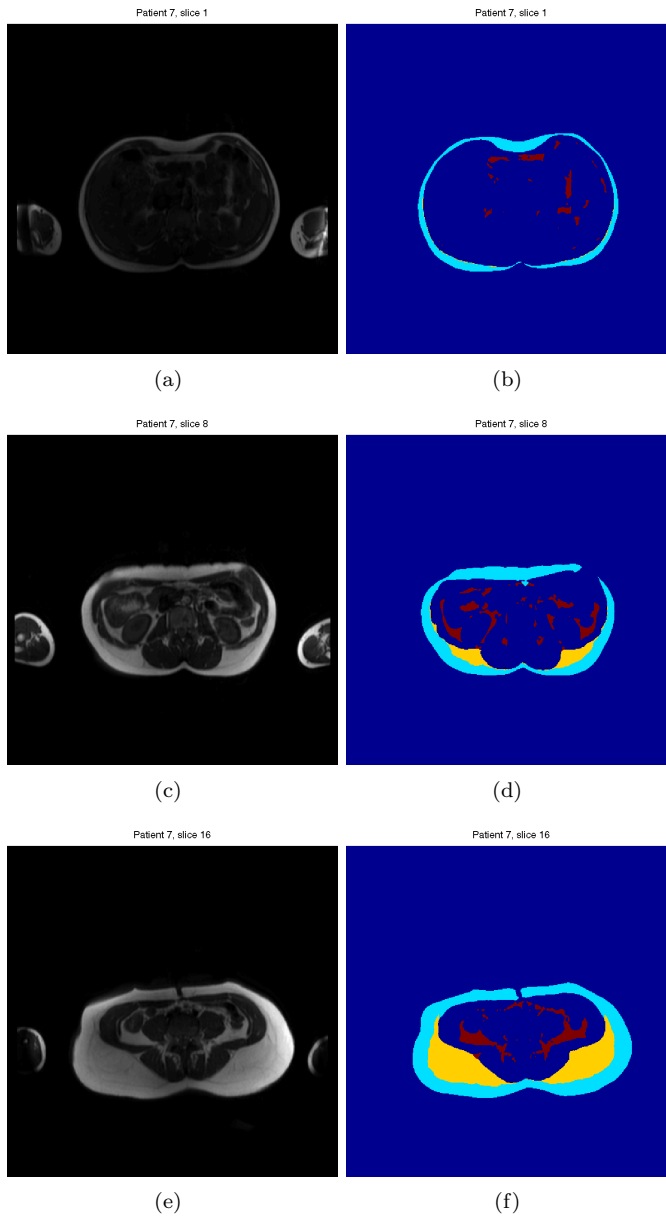


Figure B.3: Final segmentation of patient 7

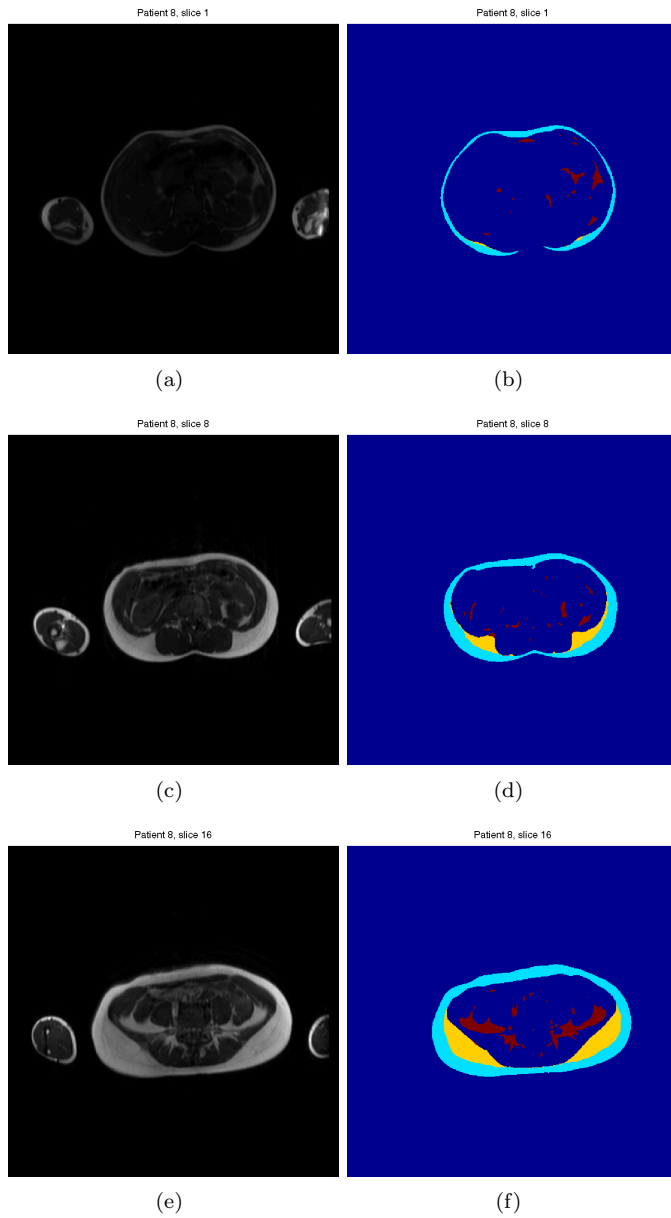


Figure B.4: Final segmentation of patient 8

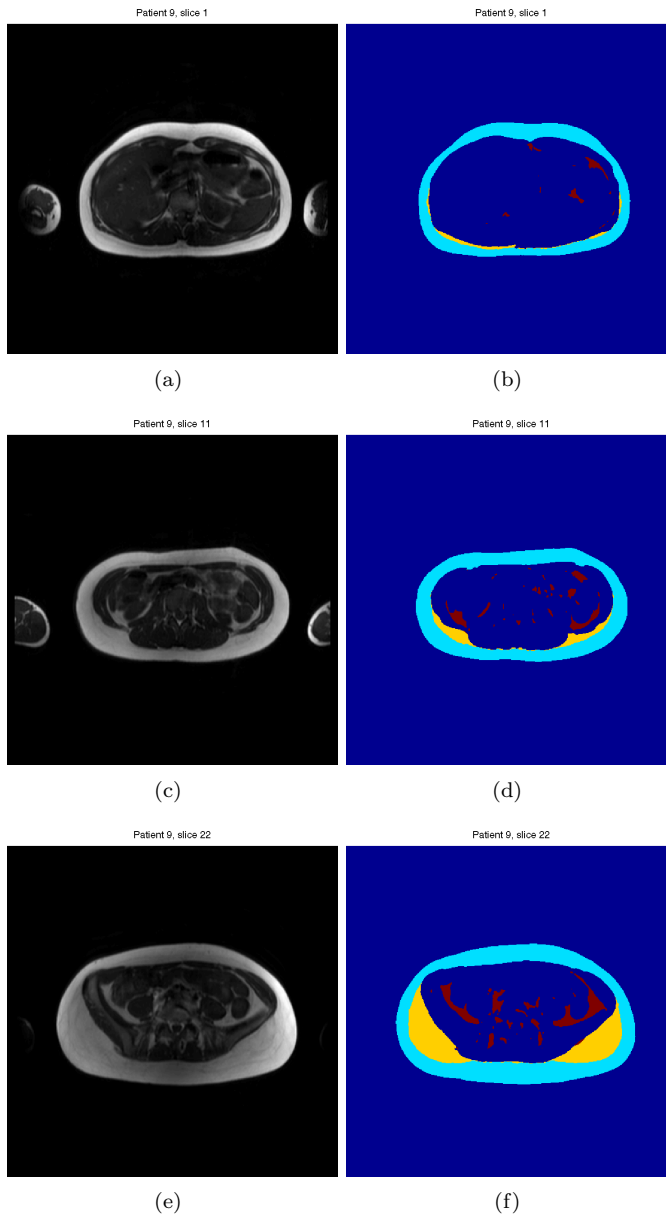


Figure B.5: Final segmentation of patient 9




Figure B.6: Final segmentation of patient 10

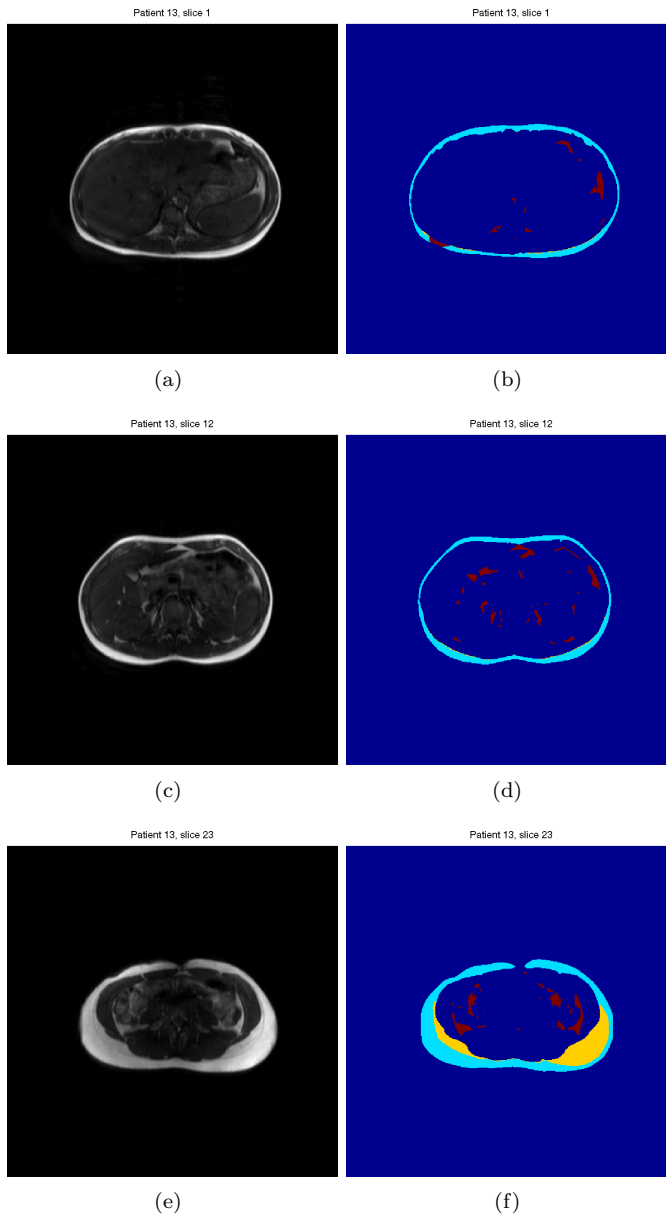


Figure B.7: Final segmentation of patient 13

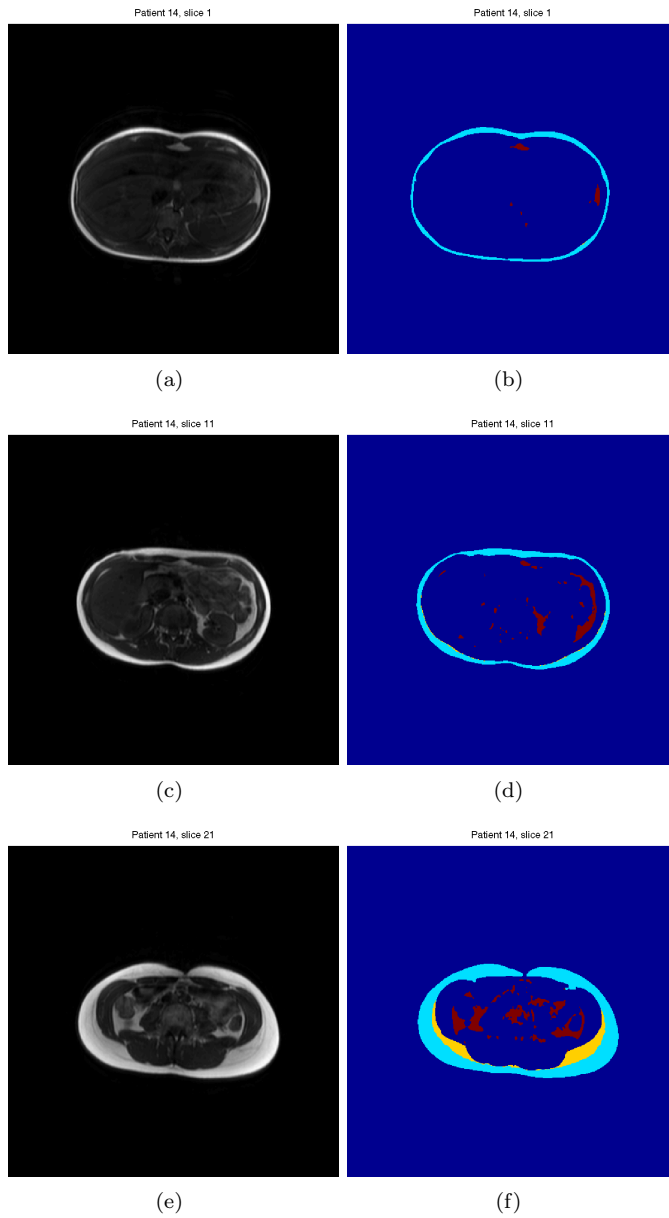


Figure B.8: Final segmentation of patient 14

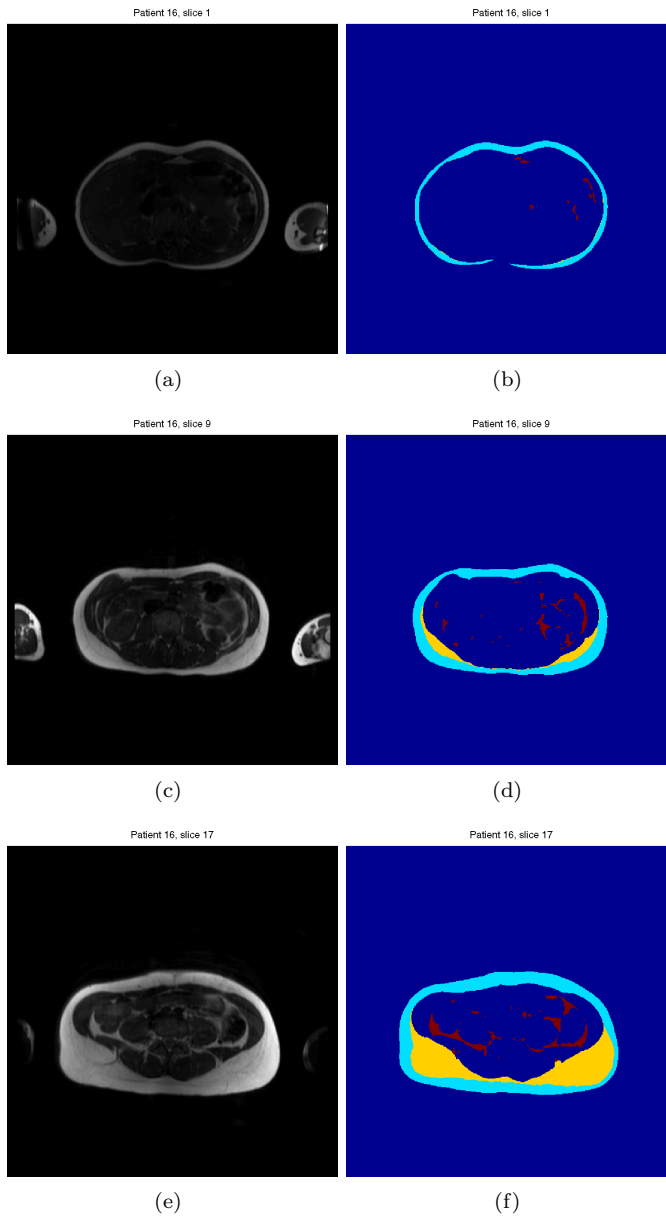


Figure B.9: Final segmentation of patient 16

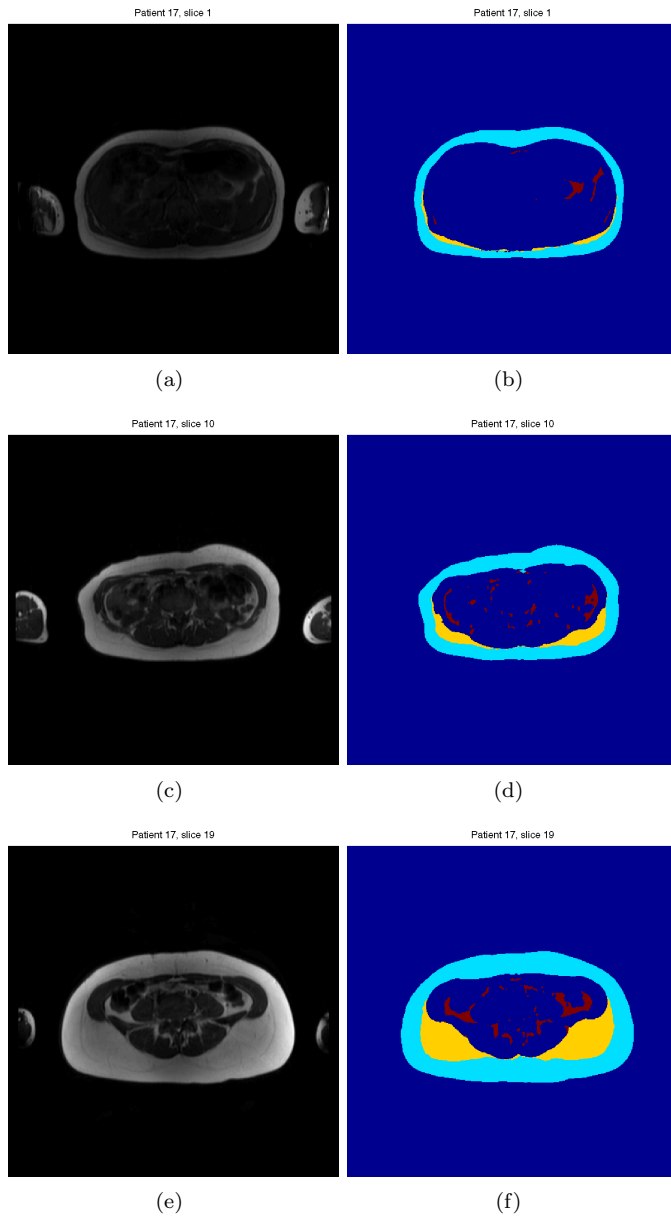


Figure B.10: Final segmentation of patient 17

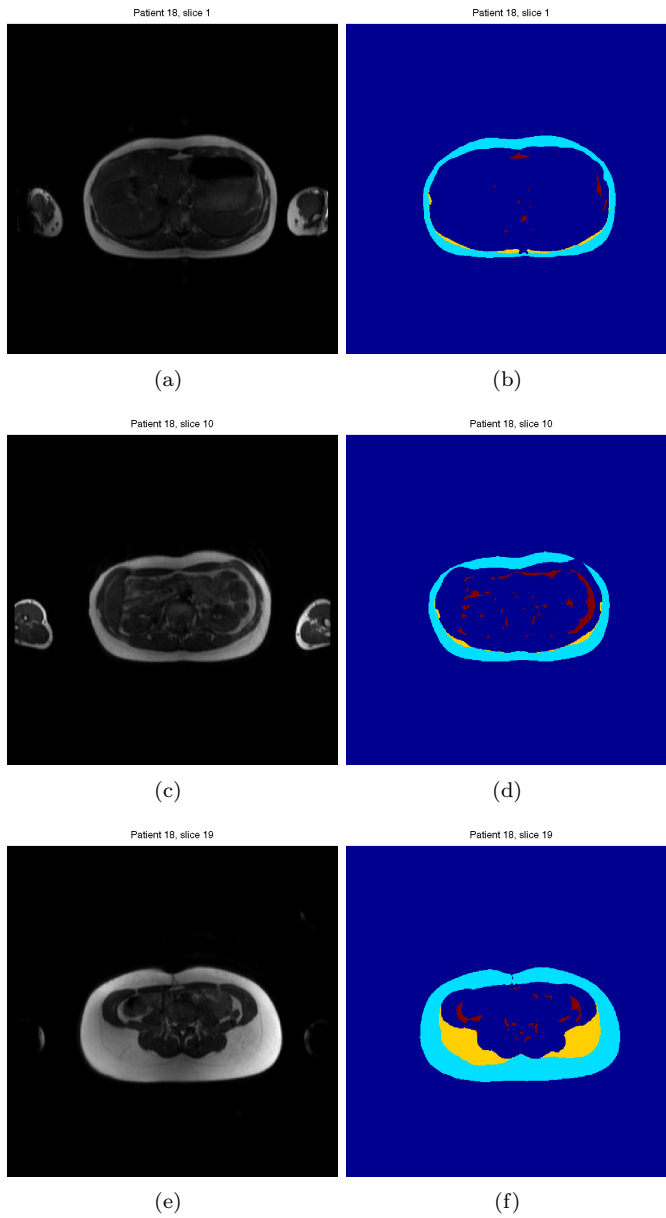


Figure B.11: Final segmentation of patient 18

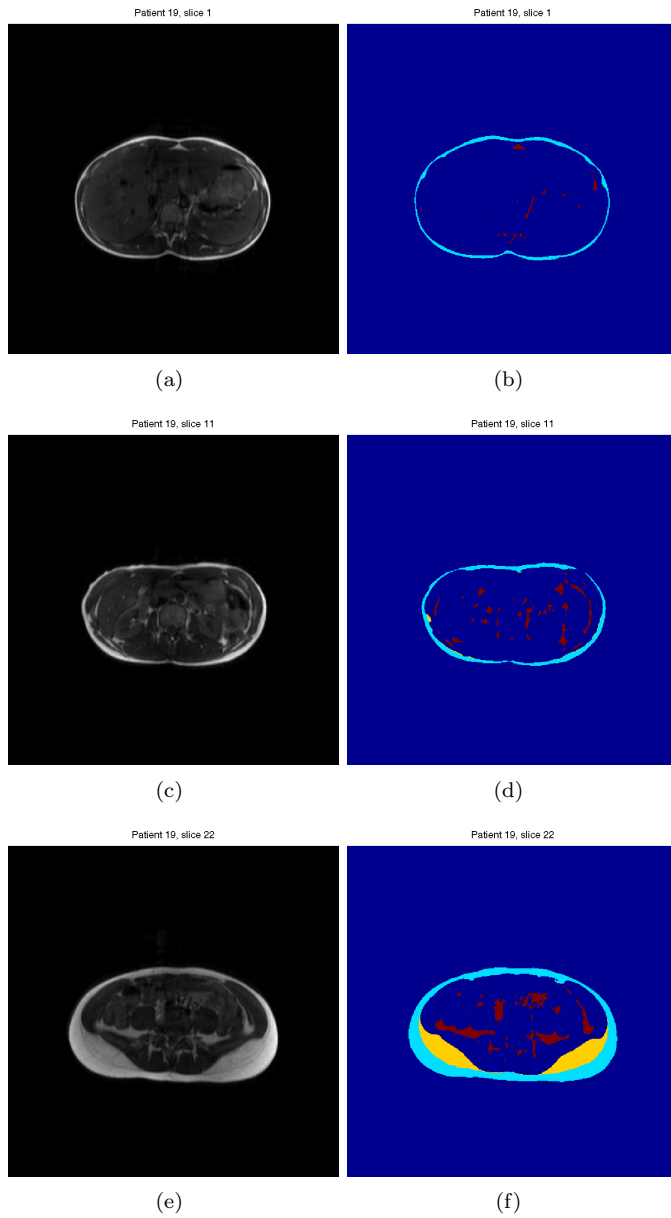


Figure B.12: Final segmentation of patient 19

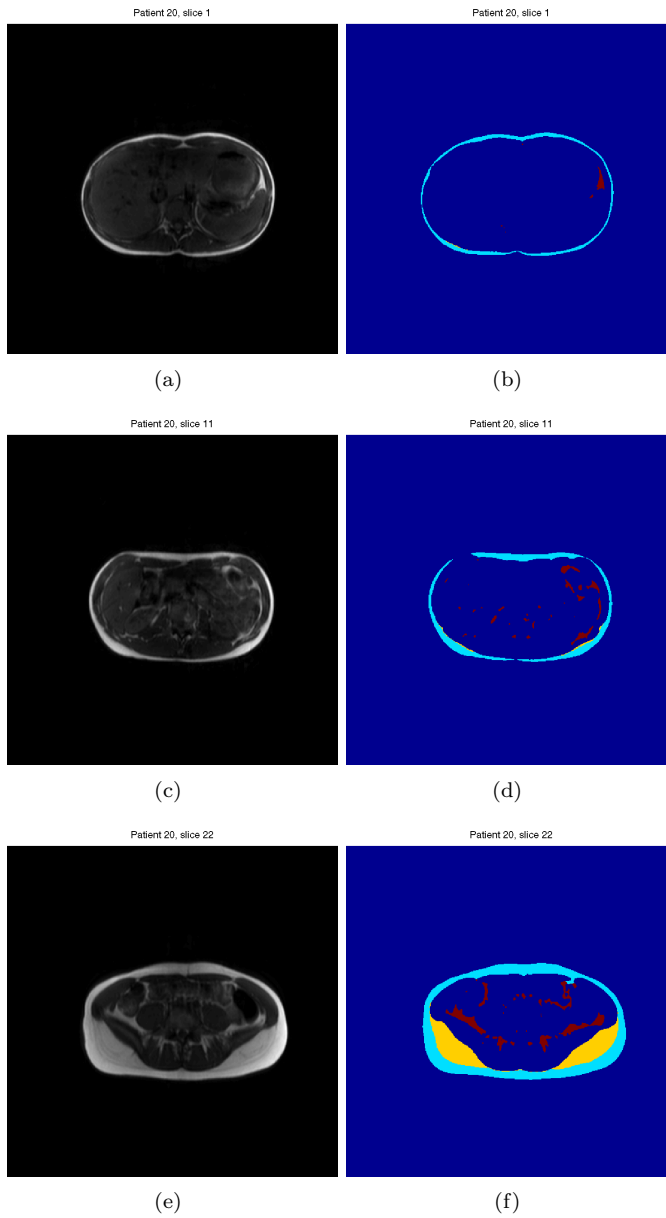


Figure B.13: Final segmentation of patient 20

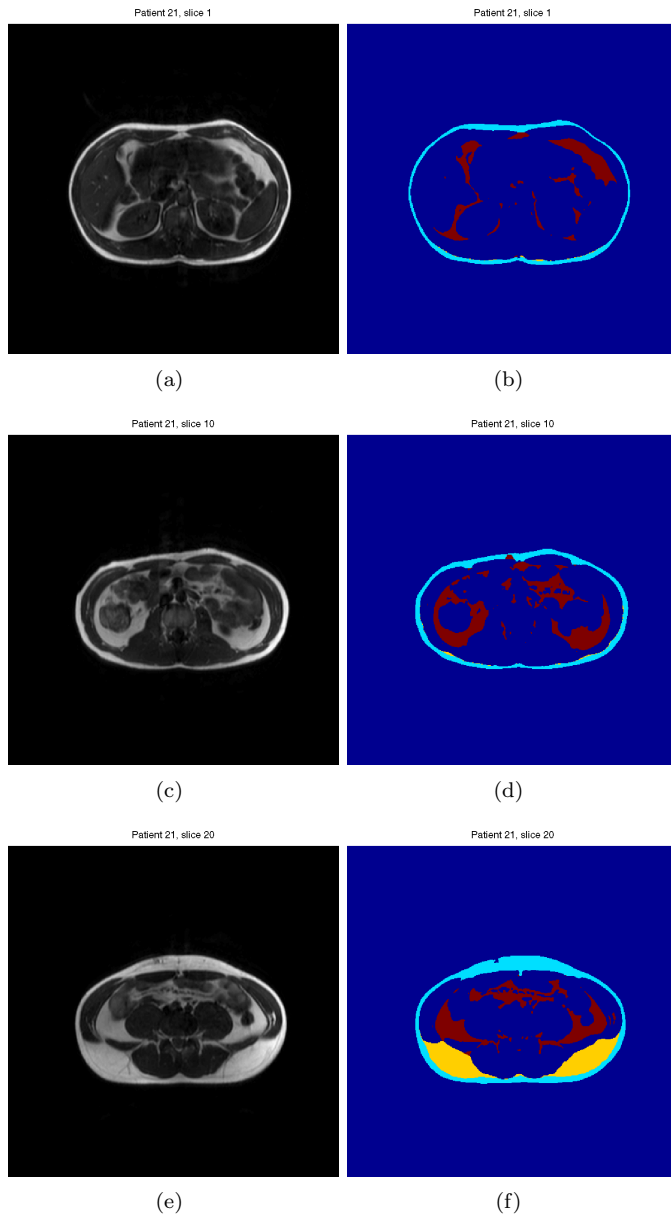


Figure B.14: Final segmentation of patient 21

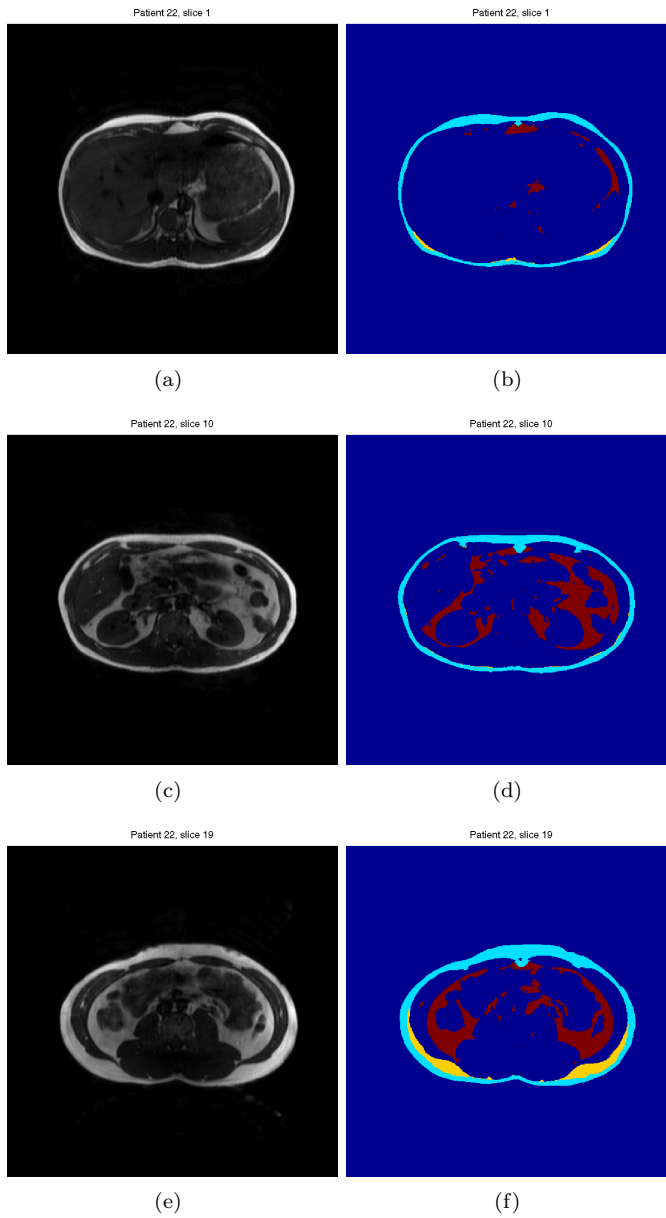


Figure B.15: Final segmentation of patient 22




Figure B.16: Final segmentation of patient 23

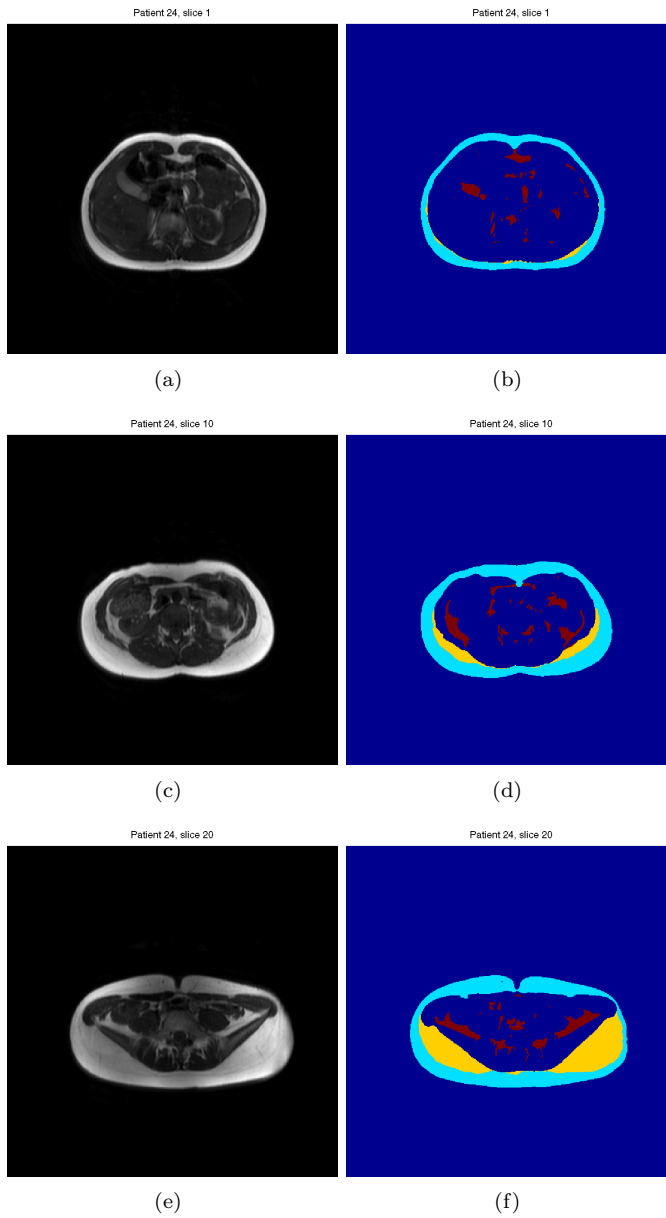


Figure B.17: Final segmentation of patient 24

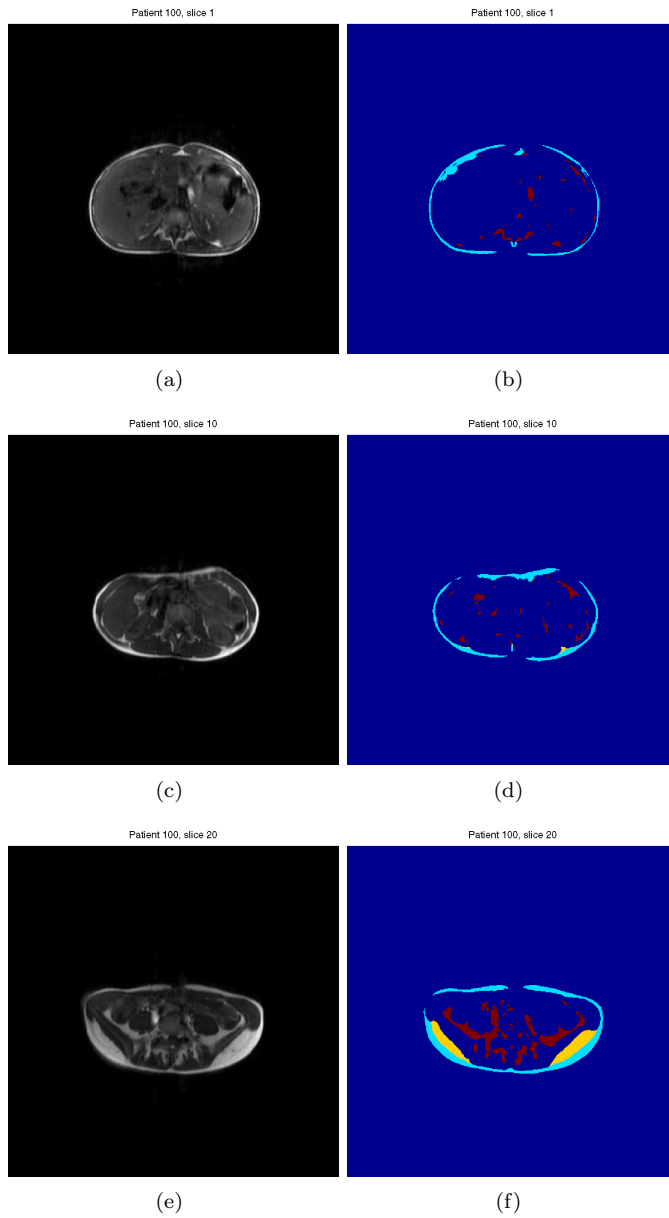


Figure B.18: Final segmentation of patient 100

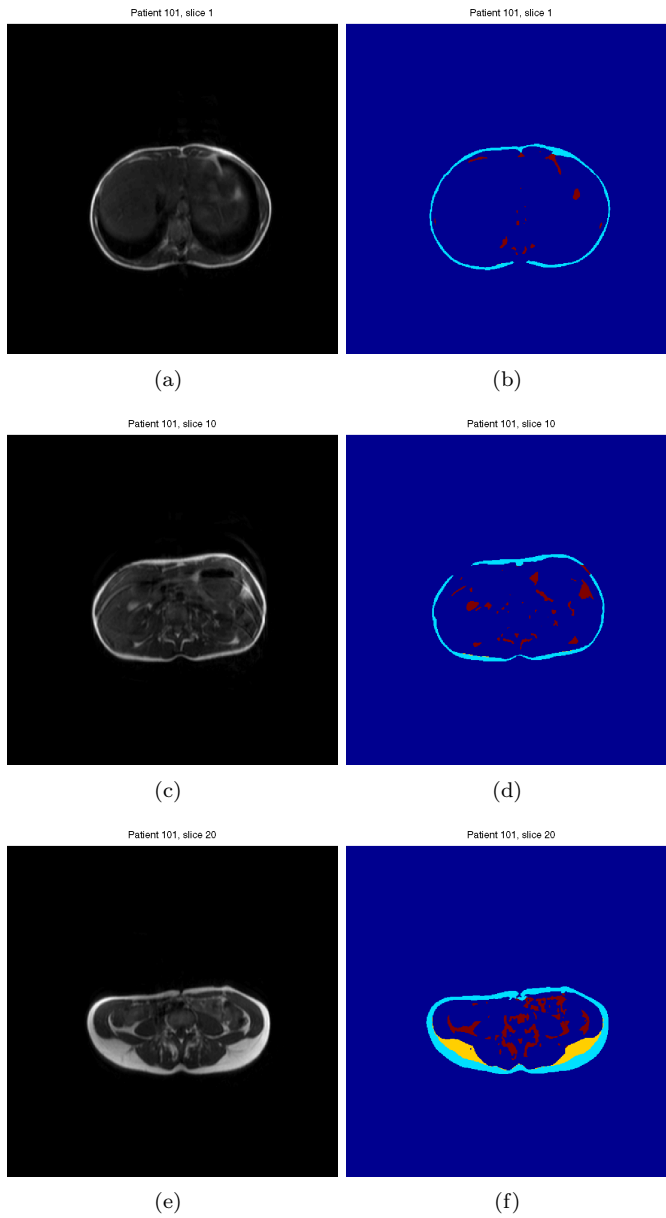


Figure B.19: Final segmentation of patient 101

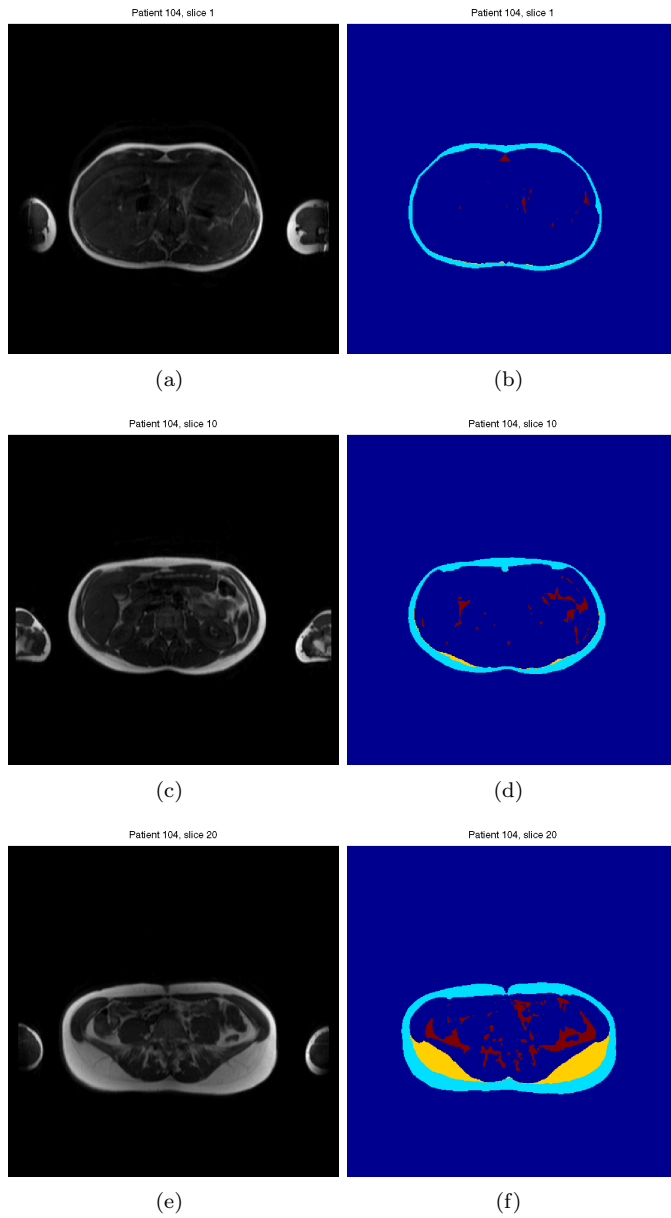


Figure B.20: Final segmentation of patient 104

Software Overview

This appendix presents a brief overview of the software developed for the segmentation. The software has been implemented using Matlab 7.4. Below is a list of the most important files and a short description of their functionality. Function hierarchy is denoted by the indentation. All source code is available on the CD-ROM enclosed in appendix D.

- **mainPreprocessing.m** loads the image data and header information from the DICOM files and gathers it in a single `.mat` file.
 - **correctbias3d.m** performs the bias field correction.
 - * **findBiasPoints.m** samples the two classes of points used for estimating the bias field.
 - * **tpsSmooth.m** estimates the bias field.
- **mainSegmentation.m** performs the segmentation
 - **doCluster.m** handles the tissue classification.
 - **doSegmentation.m** does the boundary segmentation to distinguish between the three types of adipose tissue.
 - **doProcessing.m** processes the segmentation results by some illustrative plots, and calculating volume percentages corresponding to the segmented adipose tissue.

APPENDIX D

Data CD

The enclosed CD-ROM holds the following for all patients:

- Original images
- Original images with sampled points used for the bias field correction
- Corrected images
- Tissue classified images with labels corresponding to both thresholds shown
- Images with the identified boundaries
- Final segmentation results

Furthermore the Matlab source code of the software developed is included.

Bibliography

- [1] L. D. Cohen. On active contour models and balloons. *CVGIP: Image Understanding*, 53(2):211–218, March 1991.
- [2] T.H. Cormen, C.E. Leiserson, R.L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2nd edition edition, 2001.
- [3] B.M. Dawant and A.P. Zijdenbos. *Handbook of Medical Imaging. Volume 2. Medical Image Processing and Analysis*, chapter 2. SPIE Press, 2000.
- [4] J. Duchon. Interpolation of functions of two variables following the principle of the bending of thin plates. *Revue Francaise d'Automatique Informatique Recherche Operationnelle*, 10:5–12, 1976.
- [5] 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.
- [6] P.J. Green and B.W. Silverman. *Nonparametric Regression and Generalized Linear Models: A Roughness Penalty Approach*. Chapman and Hall, 1994.
- [7] T. Hastie, R. Tibshirani, and J. H. Friedman. *The Elements of Statistical Learning*. Springer, 2001.
- [8] Zujun Hou. A review on MR image intensity inhomogeneity correction. *International Journal of Biomedical Imaging*, 2006:1–11, 2006.

-
- [9] Qing Ji, John O Glass, and Wilburn E. Reddick. A novel, fast entropy-minimization algorithm for bias field correction in mr images. *Magnetic Resonance Imaging*, 25:259–264, 2007.
- [10] P. S. Jørgensen. Segmentation of male abdominal fat using MRI. Master’s thesis, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, 2006. Supervised by Assoc. Prof. Rasmus Larsen, IMM.
- [11] M. Kass, A Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1(4):321–331, January 1988.
- [12] David E. Kelley, F. Leland Thaete, Fred Troost, Trina Huwe, and Bret H. Goodpaster. Subdivisions of subcutaneous abdominal adipose tissue and insulin resistance. *The American Journal of Physiology - Endocrinology and Metabolism*, 278:E941–E948, May 2000.
- [13] J.G. Sled, A.P. Zijdenbos, and A.C Evans. A nonparametric method for automatic correction of intensity nonuniformity in MRI data. *IEEE Transactions on Medical Imaging*, 17:87–97, 1998.
- [14] C. Xu, D.L. Pham, and J.L. Prince. *Handbook of Medical Imaging. Volume 2. Medical Image Processing and Analysis*, chapter 3. SPIE Press, 2000.