

Well believe me, Mike, I calculated the odds of this succeeding versus the odds I was doing something incredibly stupid... and I went ahead anyway.

– Crow T. Robot (played by Trace Beaulieu)
Mystery Science Theatre 3000: The Movie, 1996.

University of Alberta

**FOCUSING ON THE MEDIUM: AN ALTERNATIVE APPROACH TO ULTRASOUND
IMAGE SEGMENTATION**

by

Brian Booth



A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of **Master of Science**.

in

Department of Computing Science

Edmonton, Alberta
Spring 2008



Library and
Archives Canada

Bibliothèque et
Archives Canada

Published Heritage
Branch

Direction du
Patrimoine de l'édition

395 Wellington Street
Ottawa ON K1A 0N4
Canada

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file *Votre référence*
ISBN: 978-0-494-45778-8
Our file *Notre référence*
ISBN: 978-0-494-45778-8

NOTICE:

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

AVIS:

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protègent cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.


Canada

Abstract

Ultrasound image segmentation has grown as a field in the hope of aiding busy radiologists in diagnosing patients, and due to high noise and low contrast in the resulting images, prior knowledge has been used extensively. However, given the strength of the prior knowledge often used, resulting algorithms are usually limited in their applicability and feasibility. In this thesis, I look at what are appropriate priors for ultrasound image segmentation. In doing so, I determine that ultrasound is a unique medium and that understanding of the image acquisition process is required.

Presented as a result of this prior knowledge review is a new approach to ultrasound image segmentation. A segmentation system is presented that relies solely on prior knowledge obtained from dissecting the image acquisition process. The image is decomposed into Synthetic A-Mode ultrasound scans and the resulting echo patterns are modeled to determine the location of likely tissue boundaries. These likely boundary locations are used to initialize a contour for a force-based segmentation scheme.

Results for the segmentation system are on par or better than other current methods, including a 73% segmentation accuracy on a set of 304 pork loin ultrasound images. A direct comparison with a similar existing algorithm in the field also shows the ability of my system to handle large variations in image content and noise. These results, obtained without the use of many limiting priors, show the value of both acknowledging the image creation process and of treating ultrasound as a unique medium in addressing the problem of segmentation.

*To the universe, or whoever decided that someone else would be Leonardo
and I would be... less remarkable.*

Acknowledgements

I gratefully acknowledge Dr. Alan Tong of Lacombe Research Centre, Agriculture and Agri-Food Canada for providing the pork loin ultrasound images used in this thesis. Similarly, I sincerely thank Dr. Vimal Patel of the Department of Radiology and Diagnostic Imaging at the University of Alberta for providing the liver and kidney ultrasound images used in this thesis. This thanks also extends to Dr. Lawrence Le of the Department of Radiology and Diagnostic Imaging at the University of Alberta, Dr. Edmond Lou of Capital Health, and the Capital Health Ethics Review Board for their work in setting up this project and for dealing with the ethical red tape surrounding the use of medical ultrasound images.

I extend a great deal of thanks to my Supervisor, Dr. Xiaobo Li, for his skepticism and for giving me the freedom to try whatever crazy idea I came up with. Also, I would like to thank him, as well as the Natural Sciences and Engineering Research Council of Canada (NSERC), the Informatics Circle of Research Excellence (iCORE), and the University of Alberta, for funding me throughout my graduate studies. A thank you also goes to Dr. Martin Jägersand for welcoming me into his lab and the computer vision research group, as well as Dr. Rong-Qing Jia for aiding in my thesis examination.

I would like to thank my sister, Susan, and my friends for giving reasons to procrastinate and for keeping my spirits up throughout my studies. This goes double for my office-mate Brendan Tansey; how we ever got any work done is still a mystery.

Finally, I would like to thank my parents for always holding me to a higher standard than the rest of the world does. Any success I achieve is as much due to their hard work as it is due to mine.

Table of Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 1.1 | Background | 2 |
| 1.2 | Purpose of this Research | 8 |
| 1.2.1 | Objective of the Thesis | 9 |
| 1.2.2 | Contributions of the Thesis | 9 |
| 1.3 | Thesis Outline | 10 |
| 2 | Ultrasound Basics | 12 |
| 2.1 | Signal Acquisition | 12 |
| 2.2 | The Ultrasonic Scanner & Image Types | 14 |
| 2.2.1 | Ultrasound Probes | 14 |
| 2.2.2 | The Ultrasound Digital Signal Processor (DSP) | 16 |
| 2.3 | Image Content & Noise | 17 |
| 2.4 | Imaging Artifacts | 18 |
| 2.5 | Related Work in Context | 21 |
| 2.6 | Assumptions | 26 |
| 3 | Segmentation System Overview | 30 |
| 3.1 | Synthetic A-Mode Ultrasound | 30 |
| 3.2 | Modeling Echo Patterns | 35 |
| 3.3 | Contour Initialization | 40 |
| 3.4 | Force-Based Segmentation | 47 |
| 3.4.1 | External Forces | 49 |
| 3.4.2 | Internal Forces | 51 |
| 3.5 | Conclusions | 54 |
| 4 | Experimental Results | 56 |
| 4.1 | Validation Issues in this Field | 57 |
| 4.2 | Agricultural Images Test Case | 59 |
| 4.3 | Medical Images Test Cases | 65 |
| 4.3.1 | Kidney Image Set | 66 |
| 4.3.2 | Liver Image Set | 71 |
| 4.3.3 | Parameter Settings and Their Effect | 76 |
| 4.4 | Conclusions | 79 |
| 5 | Algorithmic Comparison | 80 |
| 5.1 | Algorithm Description | 80 |
| 5.2 | Experimental Results | 84 |
| 5.3 | Conclusions | 92 |

| | | |
|----------|-----------------------|-----------|
| 6 | Conclusion | 93 |
| 6.1 | Summary | 93 |
| 6.2 | Future Work | 94 |
| | Bibliography | 96 |

List of Tables

| | | |
|-----|---|----|
| 2.1 | The pros and cons of various uses of knowledge in the field of ultrasound image segmentation. | 25 |
| 3.1 | Ultrasound Image Assumptions and how they are incorporated into The Algorithm. | 55 |
| 4.1 | Parameter settings used for the segmentation of 304 pork loin ultrasound images | 59 |
| 4.2 | Statistical results for the algorithm on 304 ultrasound images of pork loins | 60 |
| 4.3 | Parameter settings used for the segmentation of 30 human kidney ultrasound images | 66 |
| 4.4 | Statistical results for the algorithm on 30 ultrasound images of Human Kidneys | 67 |
| 4.5 | Parameter settings used for the segmentation of 20 human liver ultrasound images | 71 |
| 4.6 | Statistical results for the algorithm on 20 ultrasound images of Human Livers. | 72 |
| 5.1 | Ultrasound Image Assumptions and how they are incorporated into The Earlier Version of the Segmentation System. | 84 |
| 5.2 | Parameter Settings used for the Segmentation of each Data Set with the Earlier Segmentation System. | 85 |
| 5.3 | Contour Distance Results for our Early Segmentation Algorithm on all three Data Sets. | 86 |
| 5.4 | Overlap Score Results for our Early Segmentation Algorithm on all three Data Sets. | 86 |

List of Figures

| | | |
|------|--|----|
| 2.1 | The Reflection and Refraction of Sound Waves from an Ultrasonic Scanner. As in [37] | 13 |
| 2.2 | The Linear Probe of an Ultrasonic Scanner. From [40] | 14 |
| 2.3 | Ultrasound Imaging Types | 15 |
| 2.4 | The Effect of Boundary Orientation on Ultrasound Signal Acquisition | 18 |
| 2.5 | The Reverberation Artifact and how it Affects Ultrasound Images | 19 |
| 2.6 | Acoustical Shadowing and its Effects on Ultrasound Images | 20 |
| 3.1 | Overview of the Ultrasound Image Segmentation System | 31 |
| 3.2 | Pulse Trajectories from various Ultrasound Images | 32 |
| 3.3 | The Algorithm for Detecting Points Along a Curved Array Ultrasound Image Boundary | 33 |
| 3.4 | RANSAC Algorithm used to Obtain Image Boundary Lines for a Curved Array Ultrasound Image | 34 |
| 3.5 | A Sample Synthetic A-Mode Ultrasound Scan | 35 |
| 3.6 | A Sample Gumbel Distribution ($\mu = 0, \beta = 1$). | 36 |
| 3.7 | A Conceptual Description of the Echo Pattern Modeling Process | 37 |
| 3.8 | The Expectation-Maximization Algorithm used to fit Gumbel Distributions to Synthetic A-Mode Scans. | 38 |
| 3.9 | Algorithm used to determine the number of Tissue Boundaries in a Synthetic A-Mode Ultrasound Scan. | 39 |
| 3.10 | A Sequential View of the Model Fitting Process | 40 |
| 3.11 | RANSAC Algorithm used to Obtain The Initial Elliptical Contour for the Region of Interest. | 41 |
| 3.12 | Visual Interpretation of the Calculations used for Boundary and Separation Support Measures. | 44 |
| 3.13 | Visual Interpretation of the Calculations used for Region and Depth Support Measures. | 45 |
| 3.14 | Algorithm to obtain Likely Tissue Boundary Depths Expressed using Gaussian Distributions. | 46 |
| 3.15 | A Visual Description of Obtaining the Most Likely Depths in an Image. | 46 |
| 3.16 | The Calculation of Ellipse Support Values. | 48 |
| 3.17 | The Boundary Point Filtering Algorithm Expressed using Gaussian Distributions. | 51 |
| 3.18 | The External Boundary Point Force and how and where it interacts with the Contour | 52 |
| 3.19 | The Propagation of External Forces expressed in relation to the Contour Points | 53 |
| 4.1 | The Distribution of Scores for the Pork Loin Ultrasound Image Set. | 61 |

| | | |
|------|---|----|
| 4.2 | An Example of a Poor Result seen from using the Segmentation System on an Ultrasound Image of a pork loin (Contour Distance: 36.16 pixels, Overlap Score: 36.77%). | 62 |
| 4.3 | An Average Example of the Results seen from using the Segmentation System on an Ultrasound Image of a pork loin (Contour Distance: 14.92 pixels, Overlap Score: 72.36%). | 63 |
| 4.4 | An Example of a Good Result seen from using the Segmentation System on an Ultrasound Image of a pork loin (Contour Distance: 5.98 pixels, Overlap Score: 88.95%). | 64 |
| 4.5 | The Distribution of Scores for the Kidney Ultrasound Image Set. . . | 68 |
| 4.6 | An Example of a Poor Result seen from using the Segmentation System on an Ultrasound Image of a human Kidney (Contour Distance: 55.15 pixels, Overlap Score: 18.68%). | 69 |
| 4.7 | An Example of a Good Result seen from using the Segmentation System on an Ultrasound Image of a human Kidney (Contour Distance: 22.29 pixels, Overlap Score: 74.79%). | 70 |
| 4.8 | The Distribution of Scores for the Liver Ultrasound Image Set. . . | 73 |
| 4.9 | An Example of a Poor Result seen from using the Segmentation System on an Ultrasound Image of a human Liver (Contour Distance: 51.31 pixels, Overlap Score: 39.64%). | 74 |
| 4.10 | An Example of a Good Result seen from using the Segmentation System on an Ultrasound Image of a human Liver (Contour Distance: 13.88 pixels, Overlap Score: 82.29%). | 75 |
| 4.11 | Comparison of two kidney segmentation results, one using a class-based paramter setting and one using an image-tuned parameter setting. | 77 |
| 4.12 | Comparison of two liver segmentation results, one using a class-based paramter setting and one using an image-tuned parameter setting. | 78 |
| 5.1 | The Region-Based Pixel Force and its affect on the Contour | 82 |
| 5.2 | Overview of the Earlier Version of the Ultrasound Image Segmentation System | 83 |
| 5.3 | The Distribution of Scores for the All Three Ultrasound Image Sets. | 87 |
| 5.4 | An Example Result from using the Earlier Segmentation System on an Ultrasound Image of a pork loin (Contour Distance: 7.84 pixels, Overlap Score: 80.70%). | 88 |
| 5.5 | An Example of a Poor Result from using the Earlier Segmentation System on an Ultrasound Image of a Human Liver (Contour Distance: 81.30 pixels, Overlap Score: 38.48%). | 89 |
| 5.6 | An Example of a poor Initialization Result from using the Earlier Segmentation System on an Ultrasound Image of a human kidney (Contour Distance: 16.60 pixels, Overlap Score: 25.82%). | 90 |
| 5.7 | An Good Example Result from using the Earlier Segmentation System on an Ultrasound Image of a human liver (Contour Distance: 15.26 pixels, Overlap Score: 82.53%). | 91 |

Chapter 1

Introduction

Ultrasound imaging has been a popular diagnostic tool in the medical field since the early 1960's. Its main concept is to use the sound medium to gather information about the state of a patient internal anatomy and is most commonly known for its use in monitoring a fetus' growth and development.

Despite mainly being known for its use during pregnancy, ultrasound imaging is used almost everywhere in diagnostic medicine, whether it's checking for heart conditions by monitoring a person's heartbeat (known as echocardiography), to checking for cancerous masses, to diagnosing glaucoma by imaging a patient's eye [35]. Lately, the use of ultrasound has extended beyond the field of medicine. In the agricultural field, ultrasound is not only used to treat animals, but also to monitor and estimate the size of various cuts of meat from live animals [10].

The popularity of ultrasound is due to many factors, the main one being that it is a non-invasive imaging technique. It provides images of the internal state of a subject's body without having any foreign object enter that body. Of course, there are many non-invasive imaging techniques, including MRI (magnetic resonance imaging), PET (positron emission tomography), X-ray, and CT (computed tomography). Unlike those latter three imaging techniques, ultrasound is free of ionizing radiation, which is a key point especially in the field of obstetrics where a fetus cannot handle the levels of radiation used in X-ray, CT, or PET scans. Also, ultrasound is significantly cheaper than these imaging techniques and has a much faster acquisition time.

However, these advantages do not come without a cost. Ultrasound imaging

is known for its poor image quality. High levels of noise, low contrast between regions, and signal dropout during image acquisition are all major issues when it comes to analyzing the content of the resulting images. In fact, training is required in order to properly analyze the content of ultrasound images.

In the medical field, analyzing ultrasound images is one of the main tasks of a radiologist. As with other doctors, specifically in North America, radiologists are in high demand and are often overworked. Over the past two decades, efforts have been made to use computers to aid radiologists in their analysis and diagnosing tasks. Ideally, automating some of the preliminary tasks involved in ultrasound image analysis would lighten the workload on radiologists.

The shape or size of an imaged anatomical structure are often key to diagnosing a patient. To obtain this information, it is first important to detect the anatomical structure in the ultrasound image. This is essentially the problem of image segmentation: distinguishing objects in an image from their background. Proper ultrasound image segmentation solutions would lead to good measurements of the size - and proper classification of the shape - of various anatomical structures, making it faster and easier for a radiologist to analyze the results and successfully diagnose a patient.

1.1 Background

Current work in ultrasound image segmentation is extensive, and though many techniques are being proposed to solve this problem, certain key approaches have become commonplace. While these algorithms differ most notably in the type of image cues and prior knowledge incorporated into the segmentation process, it has been widely accepted, or at least implied, that segmentation algorithms for ultrasound images should include any and all information available in order to overcome the poor quality of the images. With this mindset however, there is a concern that the applicability of an algorithm may be limited by the type of image cues and priors used.

The features used to segment ultrasound images depend strongly on the ap-

proach taken to the problem. Contour-based segmentation algorithms search for the region of interest by attempting to locate its boundary, and therefore rely on features that will highlight said boundary. On the other hand, region-based segmentation methods attempt to identify a portion of the image where all points inside that part of the image are similar in appearance, according to a certain measure.

Though many of the current approaches in ultrasound image segmentation combine approaches or rely on the use of multiple features or priors, the main cues driving these segmentations can be placed into one of five different categories:

- *Gradients*. Contour-based segmentation feature. Look for areas of high gradient magnitude.
- *Gray-level Distributions and Histograms*. Region-based segmentation feature. Region similarity based on intensity model.
- *Texture*. Region-based segmentation feature. Region similarity based on texture measures.
- *Shape*. Prior often used as a constraint in contour-based methods.
- *Application-specific Knowledge*. Ad-hoc approaches based on specific image content.

This section focuses on these five forms of knowledge and how they are often incorporated into ultrasound image segmentation algorithms. It should be noted that the algorithms presented here are not an exhaustive covering of the field of ultrasound image segmentation, but is instead a brief overview of the most popular methods in the field. For further approaches, please refer to [38] and references therein.

Gradients

A common assumption in image processing is that a strong gradient magnitude is a good indicator of a region boundary. While the accuracy of this assumption is

debatable in the context of ultrasound images, it nevertheless remains a popular approach in the field.

The most common use of gradient knowledge is through the use of Active Contours. First presented by Kass *et al.*, the active contour approach takes a given initial contour and, while keeping the contour smooth, moves it to nearby areas of high gradient magnitude. This process is done by maximizing the gradient along the curve while minimizing the length and curvature of the curve as shown in the energy functional given in Equation 1.1 [31]. In this formulation, the curve is represented parametrically by $\mathbf{v}(s) = (x(s), y(s))$.

$$E_{snake} = \int_0^1 |\nabla(I(\mathbf{v}(s)))| - \gamma|\mathbf{v}_s(s)| - \lambda|\mathbf{v}_{ss}(s)| \quad (1.1)$$

In ultrasound image segmentation, Active Contours is used in combination with other techniques in order to automate the segmentation process. In the work of Tong *et al.* on pork loin images, intensity-based region growing is performed first and the obtained region boundary is used as an initial contour [42]. Mathematical morphology is then applied to the Active Contour result to smooth the contour to its final result. A similar approach is taken by Holmes and Robb [26] on ultrasound images of the prostate.

Gradient information has also been relied upon outside of Active Contour methods. In the work of Abolmaesumi and Sirospour, candidate edge points are detected using gradient magnitude along fan lines through a user-given point within the region of interest [1]. The candidate edge points on each fan line are related to those of neighbouring fan lines through the use of a probabilistic data association filter (PDAF). An interacting multiple model PDAF (IMM/PDAF) is included in this case due to the multiple candidate edge points on each fan line. A similar candidate edge point detection scheme is used by He and Jheng in combination with an Active Shape Model (ASM) to segment the tibia bone [24].

Gray-level Distributions and Histograms

Unlike gradient-based methods which attempt to detect a region boundary, histogram-based methods are used most frequently to characterize a region. Given the specific

nature of the noise in ultrasound images, the pixel intensities within and outside a region of interest can be characterized as coming from separate probability distributions.

Rayleigh distributions used are most frequently to model region intensity. The Rayleigh probability distribution is known to describe ultrasound image noise well and its distribution function is given in Equation 1.2 [27].

$$p(x; \theta) = \frac{x}{\theta} e^{-\frac{x^2}{2\theta}} \quad (1.2)$$

Cardinal *et al.* use Rayleigh distributions to describe the four different portions of a vessel wall in intravascular ultrasound images [12]. The distributions are fitted to the histogram using Expectation-Maximization and the segmentation performed using a level-set curve evolution with maximum likelihood as the energy functional. Jardim and Figueiredo perform a similar level-set segmentation, however the variance, θ , for the Rayleigh distribution of each region's intensity is updated iteratively instead of calculated ahead of time [29].

Other distributions have also been used to model region intensity in ultrasound images. Algolini *et al.* use a Gaussian distribution to describe the regions in echocardiographic images while maintaining a similar level-set segmentation approach as Jardim and Figueiredo [3]. August and Kanade also use a maximum likelihood level-set approach but with interior and exterior regions described by learnt histograms [4].

Texture

Certain ultrasound images, particularly prostate images, display noticeable texture differences between regions to the point that texture-based segmentation techniques have become quite popular. While these approaches have been both region and boundary based, co-occurrence matrix features have shown the most success with various forms of images. A co-occurrence matrix P is often defined as follows:

$$P(i, j) = \frac{1}{R} \sum_{m=1}^{M-\Delta x} \sum_{n=1}^{N-\Delta y} \delta(I(m, n) = i \text{ and } I(m + \Delta x, n + \Delta y) = j) \quad (1.3)$$

where the image I is of size $M \times N$ and $R = (M - \Delta x)(N - \Delta y)$. The matrix records the relationships between pixel intensities within a region of size $\Delta x \times \Delta y$. From this matrix, various features describing an image's texture can be calculated [37].

In Muzzolini's seminal work, co-occurrence matrix features are used along with statistical classifiers to classify voxels in 3D ultrasound images. Once the voxels are classified, they are grouped into blocks up to a user-defined size [37]. Hao *et al.* also use co-occurrence matrix features along with other quantities such as local variance and wavelet energy signatures to create a feature space in which to perform region growing [21, 22].

Aside from texture features taken from co-occurrence matrices, Gabor filters are also popular for detecting locations of texture inhomogeneity. Chen *et al.* use responses from forty-two Gabor filters (six filters, each separated by thirty degrees, at seven different resolutions) to create a new image on which a modified Active Contour segmentation is performed [14]. The norm of the responses from each of the Gabor filters is used as the intensity in this new image.

Shape

Shape is possibly the most commonly used prior in ultrasound image segmentation and its incorporation into algorithms ranges from the use of smoothness terms in Active Contour approaches to the fitting of deformable models. Active Shape Models (ASM) and its variants are the most popular of the more explicit methods. Under the ASM approach, a region is represented by a fixed number of landmark points, (x_i, y_i) , selected, usually manually, along the boundary of the region to create a point distribution model (PDM) similar to that in Equation 1.4.

$$\mathbf{x} = (x_1, x_2, \dots, x_n, y_1, y_2, \dots, y_n) \quad (1.4)$$

The point distribution models from each training image are then combined in a matrix and principal component analysis (PCA) is performed to obtain a set of basic shapes, Φ , of the region of interest. At this point, the segmentation problem is reduced to determining the linear combination of the base shapes obtained from

PCA that best represent the region of interest as described in Equation 1.5.

$$\mathbf{x} \approx \mathbf{b} \cdot \Phi \quad (1.5)$$

Many variants of this type of Active Models have been applied to ultrasound image segmentation. Glasbey used the basic ASM described above in combination with a Bayesian framework to segment ultrasound images of sheep loins [19]. Bosch *et al.* added the intensity of the landmark points to their PDMs, creating what is known as an Active Appearance Model (AAM), and used the resulting model to segment heart chambers in echocardiographs [9]. In fact, the work of Bosch *et al.* is unique in that landmark points are selected from a sequence of images to create a 2D+time model for entire cardiac cycles. Jiang *et al.* also use the Active Model approach, but instead of adding the intensity of the landmark points to their PDMs like Bosch, Jiang uses gradient orientation and magnitude instead [30].

Chen *et al.* also use an explicit model in their level-set approach. In this case, the level-set energy functional includes a term describing the distance between the current contour location given by the level-set function and a registered version of the model [15].

However, the use of shape information is not limited to the fitting of deformable models. Knowing that the vessel wall in an intravascular ultrasound image is roughly circular, Zhu *et al.* convert their images to the radial coordinate system and use Active Contours on a filtered version of the transformed image to find the now-roughly-horizontal region boundaries [46]. Cancela *et al.* also shift their segmentation problem into a different coordinate system. In this case, the new coordinate system is defined using a model. The x-axis in the new coordinate system is the model contour with a selected point on the contour labeled as the origin [11]. All image points are translated according to the above description and like Zhu *et al.*, a roughly horizontal region boundary is detected.

Application-specific Knowledge

Ad-hoc methods are also common in ultrasound image segmentation due to the poor image quality and their ability to be designed around a particular context or a

particular data set. These algorithms are denoted as relying on multiple low-level image processing techniques to enhance the region of interest in order to make them easier to detect. The actual detection of the region of interest is either done manually or by comparison with a manually given shape.

For example, Youmaran *et al.* combine Gaussian and Wiener smoothing filters with thresholding and mathematical morphology to enhance ultrasound images of the eye. Detection of the anterior chamber and sclera regions in the eye are then done by comparing each resulting region with a template shape [45]. Hiransakolwong *et al.* also use comparisons with a series of template shapes for object detection following peak-and-valley filtering, smoothing through interpolation, and adaptive thresholding that reduce the image to a set of candidate regions [25].

Other algorithms rely on manual object detection. Hamou and El-Sakka enhance ultrasound images of the carotid artery using histogram equalization, Gaussian filtering, Canny edge detection, and mathematical morphology, before leaving the region selection task to an expert [20]. Flores *et al.* use a median filter, intensity-based region growing, and dilation to detect cancerous masses in breast tissue. In this case, the manual region selection is done as part of the initialization of the region growing method [2].

1.2 Purpose of this Research

Up until now, ultrasound image segmentation, as a problem, has been analysed in the same fashion as other segmentation problems on real world images. Gradients are used to represent region boundaries, texture or intensity to describe the regions, and shape constrains to overcome noise. However, all of these techniques try to avoid, or attempt to ignore, one key issue, specifically:

Ultrasound is a unique image modality with its own specific types of noise and artifacts due to its medium of acquisition.

This point may seem obvious, but its implications are far-reaching. Specifically, all assumptions made in image processing about the appearance of features in real world images cannot be automatically applied in ultrasound image segmentation.

For example, the assumption that edges have a high gradient magnitude is highly debatable in the context of ultrasound images. In fact, the mere recognition that the analysis of ultrasound images requires considerable training suggests that at least some, if not many, of the assumptions that our visual system makes - and that are often incorporated into image processing algorithms - do not hold for ultrasound images.

1.2.1 Objective of the Thesis

The objective of the thesis is twofold, the first objective being the detection of features unique to ultrasound images that could be useful to recognize image regions and their boundaries. In particular, I will look at the image acquisition process for ultrasound images to determine how to differentiate the noise and artifacts inherent in ultrasound images from the actual content of the image.

Secondly, I will show that these particular features can be incorporated into an ultrasound image segmentation algorithm and that key anatomical structures can be reasonably well segmented as a result. Furthermore, this ultrasound image segmentation algorithm will be fully automated and require no training data. These two requirements are imposed by the goal of the ultrasound image segmentation problem itself.

1.2.2 Contributions of the Thesis

Many of the contributions made by this thesis surround the approach taken to the problem, specifically the use of image acquisition information in the segmentation of ultrasound images. Others, as expected, are more implementation-specific. Regardless, great care was taken to ensure that the contributions herein can be generally applicable to the problem of ultrasound image segmentation. In particular, these contributions are:

- The decomposition of ultrasound images into synthesized acoustical signals known as Synthetic A-Modes Ultrasound Scans.
- The creation of an acoustical model for the detection of tissue boundaries.

- The detection of likely boundary locations in ultrasound images through the use of an acoustical model.
- An automated contour initialization approach that does not require explicit size, shape, location, or orientation information.
- A robust force-based segmentation framework that is immune to local noise.

Over and above these contributions is the fact that the proposed segmentation system can provide good results given a variety of images. As an example, results on a set of pork loin ultrasound images achieve 73% accuracy, better than most current methods.

Ultimately though, the greatest contribution of this thesis may in fact be the analysis of the ultrasound image acquisition process and the image processing assumptions obtained from that analysis. While the segmentation framework presented herein will be shown to be useful, the assumptions themselves could eventually lead to more sophisticated segmentation algorithms and even better results.

1.3 Thesis Outline

The remainder of this thesis is organized as follows. In Chapter 2, the medium of ultrasound will be analysed and the acquisition of ultrasound images discussed. In the new context of this information, related work in this field will be reconsidered in Section 2.5. Section 2.6 will give a description of the assumptions upon which my ultrasound image segmentation approach will be based.

My segmentation approach will be presented in Chapter 3. In particular, the creation of Synthetic A-Mode Ultrasound will be presented in Section 3.1 while the acoustical modeling of tissue boundaries will be discussed in Section 3.2. The contour initialization and force-based segmentation will follow in Sections 3.3 and 3.4 respectively.

Results for this segmentation algorithm on various sets of ultrasound images will be presented in Chapter 4. Also, this algorithm will be compared to an earlier version of this work in Chapter 5.

Finally, I will conclude and consider possible future work in Chapter 6.

Chapter 2

Ultrasound Basics

An understanding of the processes under which ultrasound images are formed is often necessary, or at least helpful, in order to understand the content of ultrasound images, and in turn properly process these images. Therefore, it is imperative that a brief overview is given of how an ultrasonic scanner is used, how it functions, how the images are created from the sound data, and the resulting complications that often arise throughout this process. For this chapter, it is assumed that the reader is familiar with longitudinal waves - of which sound waves are an example - as well as the concepts of frequency, amplitude, and propagation.

2.1 Signal Acquisition

The ultrasound signal acquisition process is essentially performed by emitting an ultrasonic pulse and recording the echoes created by the pulse as it travels through the subject. A simplified example of this process is shown in Figure 2.1. This process begins with a transducer: a small piece of ceramic that, when exposed to an electrical current, expands and contracts slightly, creating an ultrasonic pulse. The opposite is also true, that when hit with an ultrasonic pulse, the ceramic piece expands and contracts, emitting an electrical current. The frequency of the ultrasonic pulse is controlled by the electrical potential applied to the ceramic and is usually kept in the range of 1-10Mhz, with lower frequencies used to penetrate deeper into the subject. [41].

As the ultrasound pulse proceeds into the subject, it will encounter what are

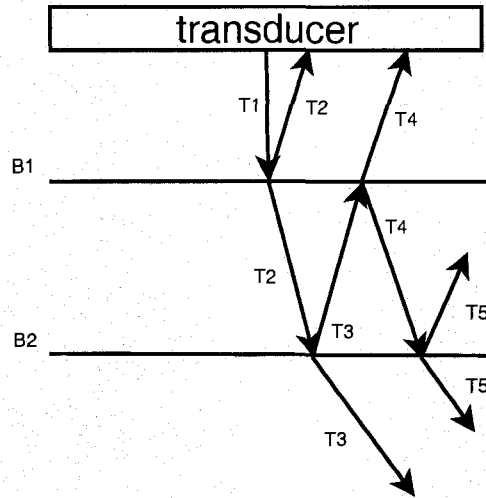


Figure 2.1: The Reflection and Refraction of Sound Waves from an Ultrasonic Scanner. As in [37]

known as acoustical boundaries. These are locations where the density of the medium through which the pulse is travelling changes significantly. At these locations, a portion of the ultrasound pulse will reflect back towards the transducer, while a weaker ultrasound pulse will refract through the boundary and continue deeper into the subject. The reflected pulse is known as an echo, and the greater the magnitude of the density difference at the boundary, the stronger the produced echo will be [37].

When the echoes return to the transducer, their strength is recorded based on the electrical potential they induce on the transducer. Also, the time between the emission of the ultrasound pulse and the recording of the echoes is recorded. Using the formula $velocity = \frac{2 \text{ distance}^1}{time}$, and knowing that the speed of sound through soft tissues is $1540 \frac{m}{s}$, the depth at which the echo was generated can be obtained.

The entire signal acquisition process takes on the order of 0.001 seconds. The ultrasound pulse takes one microsecond to be generated while the remaining 999 microseconds are used to record the echo of echoes emanating from the original pulse. This quick acquisition time allows ultrasound to be a real-time imaging technique; one of its main strengths [41].

¹The 2 used in this equation stems from the fact that the round-trip time was recorded by the transducer.



Figure 2.2: The Linear Probe of an Ultrasonic Scanner. From [40]

Note that as the ultrasound pulse travels through the subject, its amplitude will continuously weaken, as will the echoes being sent back to the transducer. This side effect is known as *attenuation* and is dependant on many variables including the frequency of the emitted ultrasound pulse and the density of the medium being traversed [36]. As a result of attenuation, there are finite depths at which an ultrasound pulse will no longer be able to return an echo strong enough to be recorded.

2.2 The Ultrasonic Scanner & Image Types

As important as the transducer is in ultrasound signal acquisition, it is only a small part of a larger machine known as an ultrasonic scanner. Conceptually, the scanner is divided into three key components: the ultrasound probe which acquires the data, the on-board digital signal processor (DSP) which cleans up and integrates the obtained data, and a screen for displaying the resulting ultrasound image. For the purposes of this research, it is the details of the first two components that are of greatest interest.

2.2.1 Ultrasound Probes

There are many different types of ultrasound probes, all of which are known to be small devices. The largest ultrasound probes are still small enough to be handheld, like the one seen in Figure 2.2. The small size of the probes allow for increased mobility and flexibility when imaging a patient. As a result, images of the same anatomical structure can often look very different based on the location and angle

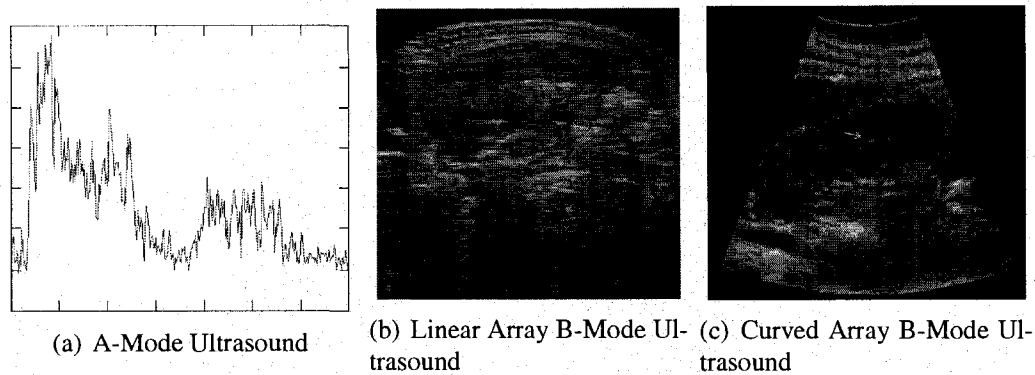


Figure 2.3: Ultrasound Imaging Types

at which the ultrasound probe is placed.

The main components of the ultrasound probe are the transducers, with probes differing in the number of transducers used and how they are setup. The most basic ultrasound probe consists of one fixed transducer. In this case, the ultrasound scan produced is one dimensional and measures the echo response with respect to depth. This is known as an A-Mode (amplitude mode) ultrasound scan. An example A-Mode scan can be seen in Figure 2.3(a). A-Mode ultrasound is rarely used for diagnostic purposes as the scans are difficult to visualize and therefore manually analyze.

The most commonly used ultrasound probes have either an array of fixed transducers or a curved array of transducers that rotate to receive signal responses from a fanned out area. These probes create B-Mode (brightness mode) ultrasound scans: the commonly known ultrasound image or slice. Figures 2.3(b) and 2.3(c) show examples of a linear array B-Mode ultrasound scan curved array B-Mode scan respectively. In a B-Mode scan, the array of transducers can be pictured at the top of the image with depth shown roughly along the vertical axis. The amplitude of the echoes received from a particular location is displayed as the brightness of the pixels. The brighter a location in the image is, the stronger the recorded echo for that location. Conceptually, a B-Mode ultrasound scan can be considered as an array of A-Mode ultrasound scans with each transducer in the ultrasound probe generating its own signal records.

Other probes also exist, including intraluminal array ultrasound probes that can

be inserted into the body. The resulting ultrasound scans are still B-Mode images, but the conceptual location of the transducer is no longer at the top of the resulting image. For example, in intravascular ultrasound scans (IVUS), the transducer is located in the middle of the image with the blood vessel wall imaged around it. In this case, the probe consists of a single transducer that rotates 360 degrees to create the resulting image. Also common is Trans-rectal Ultrasound scans (TRUS), which are similar to intravascular ultrasound scans and are used to image the prostate.

There are also additional ultrasound imaging modes that are not directly used in this thesis, including M-Mode (motion mode) Ultrasound, Doppler Ultrasound, and 3D Ultrasound. With the exception of Doppler Ultrasound which is a little more involved, the data acquisition process for these imaging modes is similar to that described herein. For further information on these ultrasound imaging modes, please refer to [37, 36].

2.2.2 The Ultrasound Digital Signal Processor (DSP)

Once the ultrasound data has been gathered, it is processed by an on-board digital signal processor. Due to competition among various vendors of ultrasound scanners, little is known what actions these DSP's perform, yet there are two known tasks performed by the processor that are worth noting. One of these tasks performed is a log-compression of the signal data. This process shrinks the range of the data for better visualization [38]. The other known task performed by the DSP is *time-gain compensation*.

Time-gain compensation is a tool used to overcome the attenuation effects of the ultrasound signal as it travels deeper into a subject by amplifying more distant echoes. This amplification is proportional to the round-trip time of the ultrasound pulse and, like the location and the angle of the ultrasound probe, is controlled by the user. Generally, a higher time-gain compensation setting will lead to a brighter image [37].

It is generally accepted that further processing is done by the DSP to improve the quality of the resulting ultrasound image. However, what other analysis a DSP performs is kept as proprietary company information, and therefore no assumptions

can be made about what that further processing involves. I simply note that the resulting ultrasound image is not a direct representation of the original ultrasound signals.

2.3 Image Content & Noise

As can be seen in Figure 2.3, ultrasound images can be difficult to interpret in spite of knowing how the images were acquired. One would expect strong echoes from tissue boundaries, where a change in medium density is generally large, and little to no echoes originating from within tissues. An ultrasound image, therefore, would be expected to show tissue boundaries as being brighter than the rest of the image while tissue regions should appear almost completely black.

In reality, this is not the case. Soft tissues do not appear with a uniform intensity within an ultrasound image. Instead, there appears to be a significant amount of variation in the appearance of the tissue itself. This non-uniformity is due to what has been termed *speckle noise*, an observable and random interference pattern.

Speckle noise is caused by many factors, the most obvious being that the tissues themselves are not of uniform density. Each tissue has its own micro-structure that scatters ultrasound echoes and produces particular echo patterns. As a result, a tissue region can have a noticeable granularity within an ultrasound image, which has led to the use of texture as a key image feature used in segmentation [38]. Also, the distribution of specularities in the image has been shown to follow a Rayleigh distribution, which has led to the distribution's inclusion in many probabilistic segmentation frameworks [27].

However, these is more to the textured appearance than simply tissue inhomogeneity. The superposition of echoes generated by the ultrasound pulses of neighbouring transducers generates small spots of relatively high intensity noise throughout the image. Also, general acoustical noise may be picked up by the probe and inserted into the image, especially in areas where the original signal has greatly attenuated. Both of these causes ultimately affect the reliability of speckle noise as a valuable feature for image segmentation.

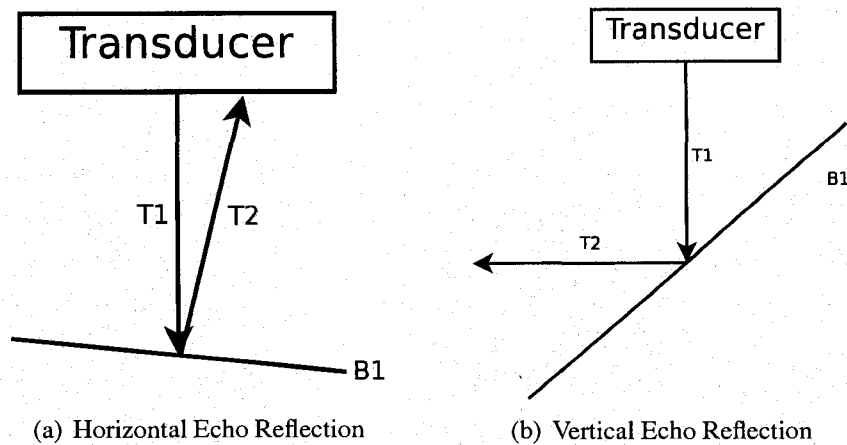


Figure 2.4: The Effect of Boundary Orientation on Ultrasound Signal Acquisition

Ultimately, while different tissues can best be detected based on their general intensity, there is no known connection between the density of the tissue and its average intensity. One special case is fluids. Water, blood, and other fluids are the easiest to observe in ultrasound images as they appear pitch black and void of any speckle noise. This phenomenon is due primarily to the more homogeneous nature of the fluid itself.

2.4 Imaging Artifacts

As much as speckle noise can plague ultrasound images, so too can various artifacts that develop through this imaging process. These artifacts can have the effect of blurring the image, missing key image information, or even adding misleading content to the image. As a result, it is important to discuss the four main ultrasound imaging artifacts and how they arise.

Orientation Dependency

Given that a transducer has a pre-defined size, the echoes it can record are limited to the ones that can propagate to the transducer's surface. Since these echoes are created from the original ultrasound pulse as it reflects off a tissue boundary, the angle of that reflection plays a key role in determining which echoes will be returned in the direction of the transducer and which echoes will miss the transducer




Figure 2.5: The Reverberation Artifact and how it Affects Ultrasound Images


completely. Consider Figure 2.4. The more perpendicular the tissue boundary is to the incident ultrasound pulse, the more likely the resulting echo will return to the transducer (as shown in Figure 2.4(a)). However, as the tissue boundary and the incident beam become closer to being parallel with each other, the more likely the reflected echo will miss the transducer completely (as in Figure 2.4(b)).

The ultimate result of this limitation is that echoes from certain tissue boundaries will be missed and in turn, those tissue boundaries will not appear in the resulting ultrasound image. This is particularly a problem for vertical tissue boundaries as they are more parallel to the incoming ultrasound pulse.

Reverberation

The entire ultrasound imaging process is based on the recording of echoes created through reflection. However, the ultrasonic scanner, or the technician running the scanner, have little to no control over this echo creation process. Speckle noise is one result of this problem, reverberation is another. Frequently, when an echo returns to the transducer, it is not simply absorbed and recorded. Instead, a portion of that echo also reflects off the surface of the transducer and returns back towards the subject to interact once again with tissue boundaries. This process repeats itself until the offending ultrasound echoes become too weak to be recorded. The repetition of this process is known as reverberation.

The most likely location for reverberation is between the transducer's surface



(a) The Process of Shadow Creation



(b) The Effect of Shadowing on the Right Side of an Ultrasound Image

Figure 2.6: Acoustical Shadowing and its Effects on Ultrasound Images

and the skin of the subject, as shown in Figure 2.5(a). This is due to the fact that the transducer and the skin are strong acoustical boundaries. However, reverberation can also occur between two strong tissue boundaries. Usually, the resulting echoes from reverberation blend into the speckle noise and do not cause a noticeable difference in the image quality, but occasionally, these reverberated echoes have a noticeable impact. This impact is most problematic when imaging fluid-filled anatomical structures such as the one shown in Figure 2.5(b). Since fluid does not create any speckle noise, the reverberated echoes have nothing to blend into, and instead stand out quite obviously.

Shadowing

When an ultrasound pulse reaches an acoustical boundary, it generally reflects and refracts. Occasionally, the ultrasound pulse will reach a boundary so strong that the majority of the pulse is reflected and the refracted pulse ends up being so weak that no recordable echoes can be obtained from it. As a result, any anatomical information below such a boundary is never recorded, leading to ultrasound images that have portions of the image area completely void of content. This process, which is shown in Figure 2.6(a), is known as shadowing.

Shadowing is one of the more dramatic imaging artifacts in ultrasound images. As can be seen in Figure 2.6(b), acoustical shadowing can lead to significant portions of the resulting ultrasound image containing no anatomical information at all.

Blurring

Unlike the earlier imaging artifacts, there are multiple sources for the blurring in an ultrasound image. In particular, the movement of the subject, whether voluntary or involuntary, can shift tissue boundaries and lead to distortions and blurring in the resulting image.

Blurring can also occur as a result of the fine discretization that is common to ultrasound images. It is important to note that the ultrasound pulse emitted from a transducer is not infinitely thin. Instead, the emitted pulse occupies roughly the same area as the transducer's surface. The result of this observation is that an ultrasound pulse interacts with tissue boundaries over a particular area. This effect, combined with pixels that are fractions of a millimetre in size, results in echoes from a tissue boundary being recorded over various pixel depths, thereby creating an additional blurring effect.

2.5 Related Work in Context

Given a brief introduction to how ultrasound images are formed, more can now be said about current work in the area of ultrasound image segmentation, specifically how the use of prior knowledge can limit the applicability of current methods. As seen, there are various complications surrounding the analysis of ultrasound images. Speckle noise and image artifacts are perhaps the most obvious of issues, but there is also the concern of dealing with images of an anatomical structure taken of different probe locations or angles. Changes in probe position can cause significant differences in the shape, size, orientation, or even appearance, of a particular anatomical structure. Also, the task of analysing images from both healthy and ill patients can add additional variability to a data set.

Earlier, the current work in the area of ultrasound image segmentation was classified by the main types of image cues and priors used in the field. Each of these knowledge classes impose a particular set of assumptions on the data being analysed. How valid these assumptions are in the area of ultrasound is of great importance in developing a successful segmentation system. Therefore, I will consider

each class in detail.

Gradients

The main assumption behind gradient-based segmentation methods is that a gradient of high magnitude is a good indicator of a region boundary. While this assumption is reasonable for many real world images, it is rarely valid with ultrasound data for various reasons. First, the speckle noise that is inherent in the imaging process adds significant gradients within homogeneous regions that can be misleading. Also, the orientation dependency of the ultrasound signal causes the more vertical boundaries in an image to be either less pronounced or even non-existent. This combination of misleading gradients and missing boundaries can lead to unpredictable and often invalid results.

Gradients may be a reasonable choice for fluid-filled structures, such as heart chambers or cysts, as these anatomical structures are often void of speckle noise and, even though certain portions of the region boundary may be missing, the speckle noise outside the region of interest may be strong enough to substitute as a reliable boundary.

Even so, it is no surprise then that the gradient-based algorithms used in this field often incorporate additional priors into their segmentation frameworks, be it shape [1, 24] or intensity [26, 42] based priors.

Gray-level Distributions and Histograms

Algorithms in this class rely on the assumption that a region can be well described by a distribution of pixel intensities. As speckle noise is known to be Rayleigh distributed [27], the use of said distribution is both well supported and well used [12, 29]. However, one must be concerned when imaging a structure from different probe locations. The appearance, particularly the brightness, of the pixels in a region can be affected by both the depth of the region in the image as well as the time gain compensation setting used by the ultrasound technician to overcome the attenuation effect.

The other major concern in this situation is the possibility that a subject's ill-

ness may add too much variability in intensity within the region of interest that the Rayleigh distribution model for the region will not be appropriate. Unfortunately, little has been done in this field to test current methods on images of ill subjects [38].

The use of other gray-level distribution models, like the Gaussian distribution, is less well supported and may only work in particular settings or with a specific data set. In the case of the Gaussian distribution, fluid-filled structures are once again the most likely candidate for success as they contain little-to-no speckle noise.

Texture

There is strong support for the use of texture to define the appearance of region in ultrasound images as the texture in the image relates well to the micro-structure of the represented tissue. In particular, the prostate, breast cancer masses, and vessel walls contain strong texture cues that are relatively easy to separate from their surrounding regions [35]. In these cases, the speckle noise can be easily modeled by texture information obtained from co-occurrence matrices or Gabor filters.

However, there are often occasions where a region's texture is not visually apparent and cannot be well detected or described by either co-occurrence matrices or Gabor filters. Take for example the ultrasound images in Figure 2.3. As a result, there is no guarantee in general that a given ultrasound image will have reliable texture cues.

Also, as with algorithms that rely on gray-level distributions, texture-based methods are also susceptible to variation from the texture model due to the illness of a subject.

Shape

The use of shape priors is possibly the most common incorporation of knowledge in ultrasound image segmentation techniques as it is robust to speckle noise and imaging artifacts. However, the assumption that the region of interest maintains a reasonably consistent shape over multiple images is a tenuous one. The imaging of multiple subjects, both ill and healthy, can lead to a significant variation in shape of a particular tissue.

Even so, the imaging of the same anatomical structure from various probe positions may be the most difficult problem facing shape-based methods. Scanning a subject from different locations, or at different angles, can lead to ultrasound images of a structure that are significantly different in terms of the structure's shape, size, and orientation. Current methods generally avoid this problem by limiting their algorithms to the most frequently used view of a structure [9, 15, 19]. While certain views of an anatomical structure are more likely to be diagnostically useful, radiologists, in practise, will use whatever image or images they find to be the most helpful for diagnosing the patient. Any automated method should be able to do the same.

Shape priors would be most reasonably used in situations where the position of the ultrasound probe is relatively fixed with respect to what is being imaged, such as intravascular ultrasound images or trans-rectal ultrasound images. Still, caution must still be exercised to ensure that any shape-based method can handle the amount of variation in size and shape that can be obtained from multiple subjects in varying levels of health.

Application-specific Knowledge

The ad-hoc methods common to this class often make multiple assumptions about the content and structure of the images being segmented. In fact, these algorithms are often built around the images themselves. Specificity is their strength, but it is also their weakness as these approaches may have difficulty adapting to many possible variations in appearance. While it can be argued that the ability to handle variability is not a concern for algorithms in this class, it should be noted that different ultrasound technicians can obtain significantly different images of a tissue despite using the same probe position on the same subject. These differences are due to the technician's personal preference for various settings on an ultrasound scanner, of which time gain compensation is one.

What is also common for many of methods in this class is the use of knowledge via manual interaction. While approaches of this kind provide significant knowledge for image segmentation, they do little to reduce the time pressures imposed on

| | Pros | Cons |
|--------------------------------|--|--|
| Gradients | <ul style="list-style-type: none"> • Can identify boundaries of speckle-free fluid-filled structures. | <ul style="list-style-type: none"> • Speckle noise adds significant non-boundary related gradients. • Absent vertical boundaries go undetected. |
| Gray-level Distributions | <ul style="list-style-type: none"> • Rayleigh distributions accurately model region intensities. • Robust to speckle noise and missing boundaries. | <ul style="list-style-type: none"> • The use of non-Rayleigh distributions is not well-founded. • Unknown ability to deal with images from ill subjects. • Adjacent regions need visible intensity differences. |
| Texture | <ul style="list-style-type: none"> • Models tissue microstructure. • Robust to missing boundaries. | <ul style="list-style-type: none"> • Unknown ability to deal with images from ill subjects. • Requires identifiable texture cues that vary among adjacent regions. |
| Shape | <ul style="list-style-type: none"> • Can compensate for missing boundaries. • Adds anatomical knowledge. | <ul style="list-style-type: none"> • Limits the algorithm to one view of a structure. • Poor at handling images from ill subjects, particularly masses. |
| Application-specific Knowledge | <ul style="list-style-type: none"> • Caters to the specific needs of the application, thereby improving the quality of the segmentation. | <ul style="list-style-type: none"> • Often includes manual intervention. • Has the possibility of being too closely tied to one data set and not others. |

Table 2.1: The pros and cons of various uses of knowledge in the field of ultrasound image segmentation.

a radiologist.

The pros and cons of each prior knowledge class are summarized in Table 2.1.

Other Points on the State of the Field

It is important to note that current approaches in this field appear to be focused on five particular areas of application, in particular echocardiographs, obstetrics/gynecology, breast cancer detection, prostate cancer detection, and intravascular ultrasound [38]. It is no coincidence that these five areas of application have strong image features.

Texture is quite evident in both breast and prostate images as well as intravascular ultrasound images. Echocardiographs as well as ultrasound images from obstetrics/gynecology are aided by the extra contrast provided by fluids within the imaged area.

While there is some research being done on ultrasound image segmentation that is not related to the categories mentioned above, the number of successful results are far fewer due to the decreased reliability of various image features.

The specificity of algorithms to a particular area of application is also interesting and it could show two possible mindsets. First, there seems to be a strong desire to obtain practical solutions for specific diagnostic problems over attempting to address the ultrasound image segmentation problem as a whole. It may also show the inability to generalize image features, or even a class of image features, from one area of application to another. The sheer variety of appearances of tissues in ultrasound images may be difficult to cover with a single set of image features.

Finally, I must acknowledge the role of ethics constraints, at least in North America, on the development of ultrasound image segmentation systems. Obtaining ultrasound images, especially medical images, for research of this kind is a very time consuming and bureaucratic process. On top of this, radiologists are very busy and it is difficult for them to find time to release images for research purposes. As a result, the data sets obtained are often quite small, to the extent that it makes it difficult to justify learning-based methods as a practical approach.

2.6 Assumptions

Up until this point, there has been a strong emphasis in the area of ultrasound image segmentation on image features and shape as the sole guiding knowledge for segmentation algorithms. While the results from these methods have been good, the algorithms themselves have had limited areas of application and often ignore some of the more challenging aspects of how ultrasound is used in general.

In contrast, far less work has been done at analyzing the image acquisition process and obtaining information from there that may be useful for segmenting ul-

trasound images. It is here that this research hope to make a significant impact. However, to do so, I must first make the connection between the image features I will use and how those feature come to be. This is done by relying solely on the following observations which, for formality's sake, I will present as assumptions.

1. Bright pixels are likely to represent tissue boundaries.

As different tissues have different densities, it is expected that the echoes created from the tissue boundary will be stronger than those created within the tissue. In practice, this is indeed the case and as a result, can be reliably assumed.

2. Pixels within a region of interest are likely to be darker than surrounding pixels at the same depth.

Consider two ultrasound pulses that have traveled the same depth into the subject. The first pulse reaches the boundary of the tissue of interest and refracts through the boundary and into the tissue. Meanwhile, the second pulse misses the tissue of interest and, due to its proximity to the tissue of interest, is unlikely to have refracted through an unrelated tissue boundary. How will the echoes originating from these pulses differ?

As expected, the echoes from the first pulse will be weaker than those from the second. This phenomenon is due to the fact that the pulse travelling through the tissue of interest has been weakened by being partially reflected at the tissue surface. Conversely, the second pulse did not cross this boundary and as a result is the stronger of the two ultrasound pulses. Therefore, the echoes originating from the second pulse will be stronger and, in turn, the area surrounding the tissue of interest will appear brighter in the image than the region of interest.

3. Noise is more likely at deeper parts of the image than in more shallow portions of the image.

This assumption is based on both attenuation as well as the propagation of echoes in multiple directions within the subject being scanned. As ultrasound pulses progress deeper into the subject, the strength of the pulse will weaken, making it

more difficult to differentiate the original signal from noise. In fact, at a certain depth, it is impossible to differentiate the two.

Meanwhile, the deeper an ultrasound pulse travels, the more echoes are generated and scattered within the subject, leading to more noise that may corrupt the signal. The result of these two effects leads to noise playing a more prominent role as you go deeper into the image.

4. Horizontal tissue boundaries will appear brighter than vertical tissue boundaries.

Due to the aforementioned orientation dependency of the signal, this image effect is to be expected. The more parallel a tissue boundary is to an incoming ultrasound pulse, the less likely the echoes generated off that boundary will return to the transducer.

5. Tissue boundaries will appear fuzzy in the image.

Unfortunately, the interaction between an ultrasound pulse and a tissue's surface is not entirely simplistic. As mentioned, this interaction does not occur at an infinitely tiny point, nor is either quantity entirely static. As a result, the tissue boundaries in ultrasound images are not as sharp or clear as boundaries can appear in real world images.

6. Pixel intensity will decay following a tissue boundary.

Reverberation of echoes between the transducer and the skin of the subject is inevitable, and while this imaging artifact often does not create any visual difference within the resulting images, its presence is nonetheless undeniable and cannot be ignored in the modelling process.

7. Ultrasound images can be conceptually considered as an array of A-Mode ultrasound scans.

This assumption originates from how an ultrasound image is formed. While A-Mode ultrasound shows the ultrasound signal's effect from one transducer, an

ultrasound image shows the same for an array of transducers. Clearly, there is a strong connection between these two visualization modes.

Ideally, a one-dimensional line in an ultrasound image that passes through the represented location of a transducer could be looked at as being similar to an original A-Mode ultrasound scan. The main difference between this line and the real scan is the overlap of echoes that occurs through the use of multiple transducers in the creation of ultrasound images. That said, the effect of this difference is unknown and given the presence of a DSP on the ultrasound scanner and the fact that spatial information is relied upon in these images for diagnostic purposes, it is enticing to assume that the effect of echo superposition has been minimized prior to image formation.

8. The region of interest has a smooth, closed contour.

While this assumption does not relate to the ultrasound imaging process, it is still important to note for the upcoming segmentation method. The rationale that the region of interest have a smooth contour is reasonable considering that virtually all anatomical structures are generally smooth in shape. The assumption that this contour is also closed should not be a surprise either in that both the image and any anatomical structure is of finite size.

Chapter 3

Segmentation System Overview

Due to poor image quality, most ultrasound image segmentation systems are quite complex. The approach presented herein is similar in that regard and includes a significant amount of pre-processing. The amount of pre-processing is not surprising as I do not use any image feature directly. Instead, the image is interpreted based on models of how the image was formed and from this process, boundary information is obtained for use in the actual segmentation.

The complete segmentation system is shown in the block diagram in Figure 3.1. The approach begins by decomposing the image into one-dimensional ultrasound scans. Echo patterns in these Synthetic A-Mode ultrasound scans are then modeled in order to determine likely tissue boundary locations. Finally, the detected boundary locations are used to initialize a contour that is stretched and pulled by various forces to obtain a final segmentation. Throughout, I rely on the existing assumptions obtained from studying the ultrasound image acquisition process.

3.1 Synthetic A-Mode Ultrasound

In order to be able to use knowledge about how ultrasound data is acquired, we must first decompose the image into its separate ultrasound signals: one for each emitted pulse. Admittedly, by splitting the image into one-dimensional scans, we lose spatial information, but the hope is that we can gain an advantage by analyzing the original signals as closely as possible. Also, despite the obvious appearance of echo superposition in ultrasound images, I will assume that the ultrasound scans

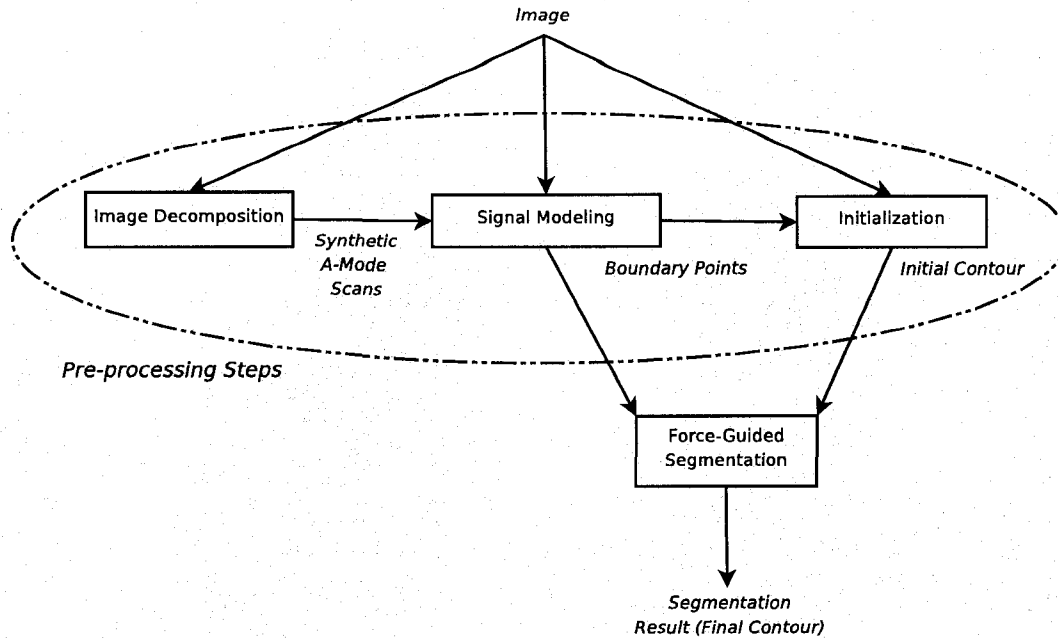



Figure 3.1: Overview of the Ultrasound Image Segmentation System

from each pulse are independent for the sake of simplicity.

Different ultrasound images will require different approaches depending on the construction of the ultrasound probe used to obtain the images. For the purposes of this thesis, we will restrict ourselves to linear and curved array ultrasound images similar to the ones seen earlier in Figure 2.3. For the linear array ultrasound images, the fixed array of transducers produce signals that correspond directly to the columns of the image. Therefore, each column of a linear array ultrasound image can be interpreted similarly to an A-Mode ultrasound scan, as described in Figure 3.2(a). Henceforth, we will consider these one-dimensional scans as Synthetic A-Mode Ultrasound scans since echo superposition, discretization, and processing done by the ultrasonic scanner have altered what would be an A-Mode ultrasound scan into what we are observing here.

The curved array ultrasound case, however, is not as straightforward. Here, the image area is shaped like a hand-held fan with the ultrasound pulses originating from the curved array positioned near the pivot of the fan. In fact, if we project the pulse trajectories backwards, they will all intersect at this pivot point. Therefore, the Synthetic A-Mode can be generated by lines that pass through this pivot point



(a) Pulse Trajectory from a Linear Array Ultrasound Image



(b) Pulse Trajectory from a Curved Array Ultrasound Image. Image Boundary lines shown in Blue.

Figure 3.2: Pulse Trajectories from various Ultrasound Images

and fan out through the image area. An example can be seen in Figure 3.2(b).

Generating synthetic A-Mode ultrasound scans from a curved array ultrasound image first requires the detection of this pivot point. To detect the pivot point, we first detect the left and right boundaries lines of the image area - shown in blue in Figure 3.2(b) - and obtain their intersection. Determining these boundary lines, in turn, involves first obtaining points along the image boundaries and fitting boundary lines to those points.

The image boundary points are detected using the algorithm shown in Figure 3.3. The pixel in the upper, left-hand corner of the image is used to obtain the intensity of the uniform background area within the image that is outside of the viewing area. Using this as a cutoff value, the first and last points above this background intensity in each row of the image are obtained. For ideal ultrasound images, fitting a line to these points is all that would need to be done to obtain the image boundary points.

Unfortunately, it is not uncommon for curved array ultrasound image to have shadowing effects near the viewing boundaries. Consider the right image boundary in Figure 3.2(b). Simply obtaining the first inward pixels above the uniform background intensity will lead to the majority of the points chosen being within the image and not along its boundary. Even the left image boundary in Figure 3.2(b) shows slight shadowing in the middle section of the image, despite being visibly

```

DETECTIMAGEBOUNDARYPOINTS(I)
1  pbdy ← I(1, 1)
2
3  forall rows r in I
4    leftPts(r) ← MININDEX(I(r, :) > pbdy)
5    rightPts(r) ← MAXINDEX(I(r, :) > pbdy)
6  endfor
7
8  leftPts ← MEDFILT1(leftPts, 9)
9  rightPts ← MEDFILT1(rightPts, 9)
10
11 forall i in leftPts & rightPts
12   leftPts(i) ← MIN(leftPts(i - 1), leftPts(i))
13   rightPts(i) ← MAX(rightPts(i - 1), rightPts(i))
14 endfor
15
16 leftPts ← UNIQUE(leftPts)
17 rightPts ← UNIQUE(rightPts)
18
19 return [leftPts, rightPts]

```

Figure 3.3: The Algorithm for Detecting Points Along a Curved Array Ultrasound Image Boundary

prominent image boundary.

To account for these shadowing effects, the obtained sets of data points for both the left and right image boundaries are arranged based on depth and filtered with a one-dimensional median filter with a window size of nine pixels to remove obvious outliers. This process is followed by removing any points from a boundary set that is clearly within the viewing area. Note that for curved array ultrasound images, the viewing boundaries move outward as they descend in the image. As a result, any pixel in a boundary point set that is closer to the center-line of the image than an earlier pixel in the same set should be removed. For each boundary point set, the pixels are traversed depth-wise and replaced with the most outward value previously seen in the set, up to and including the old value of the pixel, as seen starting at line 11 of Figure 3.3. Duplicates are then removed to obtain the final image boundary point sets.

Using the image boundary point sets, we must find the most fitting boundary lines for the left and right sides of the image. Due to shadowing effects, it cannot be assumed that the point sets obtained using the above method are free of outliers.

```

LINERANSAC(pts)
1  for i ← 1 to N do
2    point1 ← RANDOM(pts)
3    point2 ← RANDOM(pts)
4
5    sampleLines(i) ← CROSS(point1, point2)
6    lineScore(i) ← SCORELINE(sampleLines(i))
7  endfor
8
9  return sampleLines(MAXINDEX(lineScore))

```

Figure 3.4: RANSAC Algorithm used to Obtain Image Boundary Lines for a Curved Array Ultrasound Image

As a result, we cannot simply determine the line of best fit using the least-mean-squared algorithm. Instead, the RANSAC algorithm derived from [23], is used.

Figure 3.4 displays the RANSAC algorithm used for this task. Essentially, the algorithm picks two points from the boundary point set at random and creates a line through these two points. This line is then scored based on a given criteria. In this particular application, the following criteria is used:

$$Score(l) = |p_{on}| - \lambda \sum_{i=1}^{|p_{out}|} d(p_{out}, l) \quad (3.1)$$

where p_{on} represents the set of image boundary points on the line, p_{out} represents the set of image boundary points on the outside side of the line (*i.e.*, points to the right or the right image boundary line, or points to the left of the left image boundary line), and d is the euclidean distance function. Basically, the score attempts to maximize the number of points that support the line while minimizing the number of image boundary points that are outsiders. The variable λ in Equation 3.1 weights the two criteria. As the possibility of a boundary point being outside the imaging area is very low, a relatively large value for λ will likely produce the best results. In practice, a value of $\lambda = 10$ was found to be appropriate.

The process of creating and scoring lines is repeated a fixed number of times with the highest scoring line chosen as the best representation of the boundary. The number of iterations RANSAC executes depends on a few variables and is usually

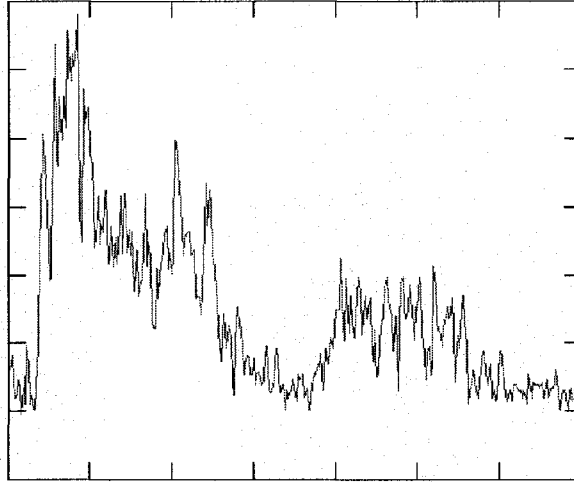


Figure 3.5: A Sample Synthetic A-Mode Ultrasound Scan

chosen using the following formula (from [23]):

$$N = \frac{\log(1 - p)}{\log(1 - (1 - \epsilon)^S)} \quad (3.2)$$

where p is the likelihood of getting a successful fit (usually set to 0.99), ϵ is the fraction of outliers in the boundary point set, and S is the number of points used for the data fitting (in this case, two). Ideally, it is best to choose a value for ϵ that is as large as is practical. $\epsilon = 0.8$ has been found to work rather well and results in $N = 113$ RANSAC iterations.

With the image boundary lines detected, we can determine the pivot point of the fanned-out region and use that point as the point of origin for the ultrasound pulses. Using this point of origin, as well as all the points between the boundary lines at the bottom of the image, we can easily interpolate at a fine scale between the two points to obtain all the pixels on each Synthetic A-Mode Ultrasound scan for a curved array ultrasound image.

3.2 Modeling Echo Patterns

With the creation of the Synthetic A-Mode scans complete, the goal of the system becomes the analysis of these one-dimensional scans for information on possible regions and their boundaries.




Figure 3.6: A Sample Gumbel Distribution ($\mu = 0, \beta = 1$).

A sample Synthetic A-Mode ultrasound scan is shown in Figure 3.5. For the sake of reminder, note that the scan displays echo amplitude versus depth in the subject, or more precisely, pixel intensity versus depth in the image. Also note the amount of noise that is still present in this mode. To identify a tissue and its boundaries will require an understanding of what type of echo pattern a tissue boundary generates.

Earlier, two particular imaging artifacts were mentioned that will aid in understanding these echo patterns: reverberation and blurring. The latter would result in relatively strong echoes over a small depth window while the former would cause a gradual decrease in echo amplitude following a tissue boundary as echoes reverberate between the transducer and the subject's skin. The result of these two artifacts would be a rise and gradual decay of echo amplitude for each tissue boundary crossed.

The Gumbel probability distribution follows a similar shape, as can be seen in Figure 3.6. The shape of the Gumbel distribution as well as its simple mathematical description make it a good candidate for modelling the echo patterns of tissue boundaries in this context. The tail of the Gumbel Distribution mimics the reverberation effect seen in the Synthetic A-Mode scans while the mode of the distribution can be taken as the most likely boundary location. The Gumbel distribution is given

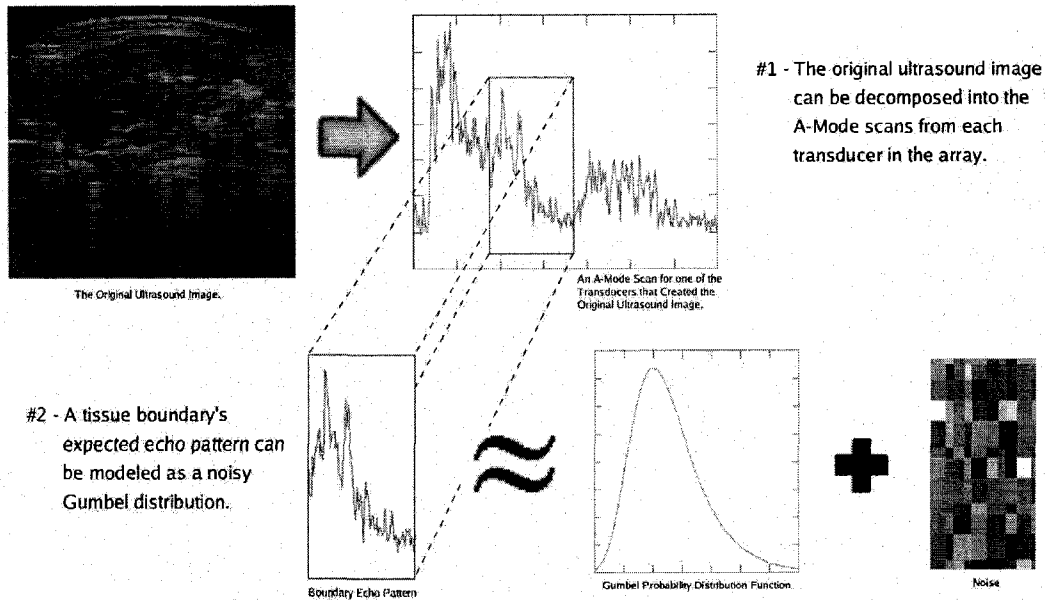


Figure 3.7: A Conceptual Description of the Echo Pattern Modeling Process

by the following probability distribution function:

$$f(x; \mu, \beta) = \frac{e^{-\frac{x-\mu}{\beta}} e^{-e^{-\frac{x-\mu}{\beta}}}}{\beta} \quad (3.3)$$

where μ is the distribution's mode and β is its spread. Using the Gumbel distribution as a model, I propose that a tissue boundary's echo pattern appears similar to a Gumbel distribution corrupted by noise. Figure 3.7 displays the conceptual steps of the segmentation system thus far.

With the Gumbel distribution as an appropriate echo pattern model, the problem of analyzing the Synthetic A-Mode scans is reduced to fitting the appropriate number of Gumbel distributions to each scan. The boundary locations given by the fitted Gumbel distributions can then be easily converted back to locations in the image.

The fitting of the Gumbel distributions to the scans can be achieved by many minimization techniques. Here, the task is performed using the Expectation-Maximization algorithm derived from [17] and given in Figure 3.8. First the algorithm converts the Synthetic A-Mode scan into a set of samples by using the echo intensity at each depth as the number of samples for that depth. These samples

```

EXPMAXGUMBEL(scan, numToFit)
1  scanSamples ← CONVERTINTENSITYTOSAMPLES(scan)
2  modes ← Linspace(MIN(scan), MAX(scan), numToFit)
3  spreads(i) ←  $\frac{\text{Var}[\textit{scan}]}{\sqrt{\textit{numToFit}}}$ 
4
5  while numChangedClass > tstop
6
7    # The Expectation Step #
8    foreach x in scanSamples
9      for i = 1 to numToFit
10        $p(x, i) = \frac{e^{-\frac{x-\mu(i)}{\beta(i)}} e^{-e^{-\frac{x-\mu(i)}{\beta(i)}}}}{\beta(i)}$ 
11     endfor
12   endfor
13   classLabels ← MAXINDEXOVERDIM(p, 2)
14
15   # The Maximization Step #
16   for i = 1 to numToFit
17     valsInClass ← scanSamples(FIND(scanSamples, classLabels == i))
18     classMean ←  $E[\textit{valsInClass}]$ 
19     classStdDev ←  $\sqrt{\text{Var}[\textit{valsInClass}]}$ 
20     spreads(i) ←  $\frac{\sqrt{6} \textit{classStdDev}}{\pi}$ 
21     modes(i) ← classMean -  $\gamma \textit{spreads}(i)$ 
22   endfor
23 endwhile
24
25 return [modes, spreads]

```

Figure 3.8: The Expectation-Maximization Algorithm used to fit Gumbel Distributions to Synthetic A-Mode Scans.

are then classified in the expectation step according to the most likely Gumbel distribution to have generated that sample. Once classified, the mean and standard deviation of each class are calculated and used to update the mode and spread of each Gumbel distribution according to the equations given on lines 20 and 21 (From [43], where γ is the Euler-Mascheroni constant), thereby maximizing the class' parameters. This process is repeated until the number of samples that change from one class to another goes below a given threshold.

As can be seen in Figure 3.8, the Expectation-Maximization algorithm requires that the user specify the number of Gumbel distributions to fit to each Synthetic A-Mode scan. Given that there may be a varying number of anatomical structures within any portion of an ultrasound image, the number of Gumbel distributions

```

NUMTOFITPOLY(scan)
1  for i ← 1 to 20 do
2    polyCoeff ← POLYFIT(scan)
3    resNorm(i) ← RESIDUALNORM(scan, polyCoeff)
4  endfor
5
6  bestFitDegree ← MININDEX(resNorm)
7  return CEIL(bestFitDegree/2)

```

Figure 3.9: Algorithm used to determine the number of Tissue Boundaries in a Synthetic A-Mode Ultrasound Scan.

to fit to each scan is unclear and must be determined before using Expectation-Maximization. However, a reasonable lower and upper bound on the number of tissue boundaries in a Synthetic A-Mode scan can be easily determined from viewing various ultrasound images. Clearly, zero is a reasonable lower bound while a reasonable upper bound on the tissue boundary number would be on the order of 10 given the limited depth achievable with ultrasound imaging.

In theory, one could simply run the Expectation-Maximization algorithm on each scan over the range of likely tissue boundary numbers and choose the result that is the best fit according to a given error measure. Unfortunately, each ultrasound image contains on the order of hundreds of Synthetic A-Mode ultrasound scans, so running the EM algorithm on each ten times would be quite time consuming.

Instead, I use the method provided in Figure 3.9 to estimate the number of tissue boundaries in a given Synthetic A-Mode ultrasound scan. Polynomials of degrees ranging from zero to twenty are fitted to the Synthetic A-Mode scan and the residual norm calculated for each case to determine the quality of the fit. While the fitted polynomials do not provide us with any information about the location of possible tissue boundaries, the degree of the polynomial of best fit does provide us with information on the complexity of the scan. Therefore, the degree of the polynomial with the lowest residual norm is determined and - due to the roughly parabolic shape of the Gumbel distribution - is divided by two to obtain an estimate for the number of Gumbel distributions to fit to the scan. This algorithm has the benefit of being

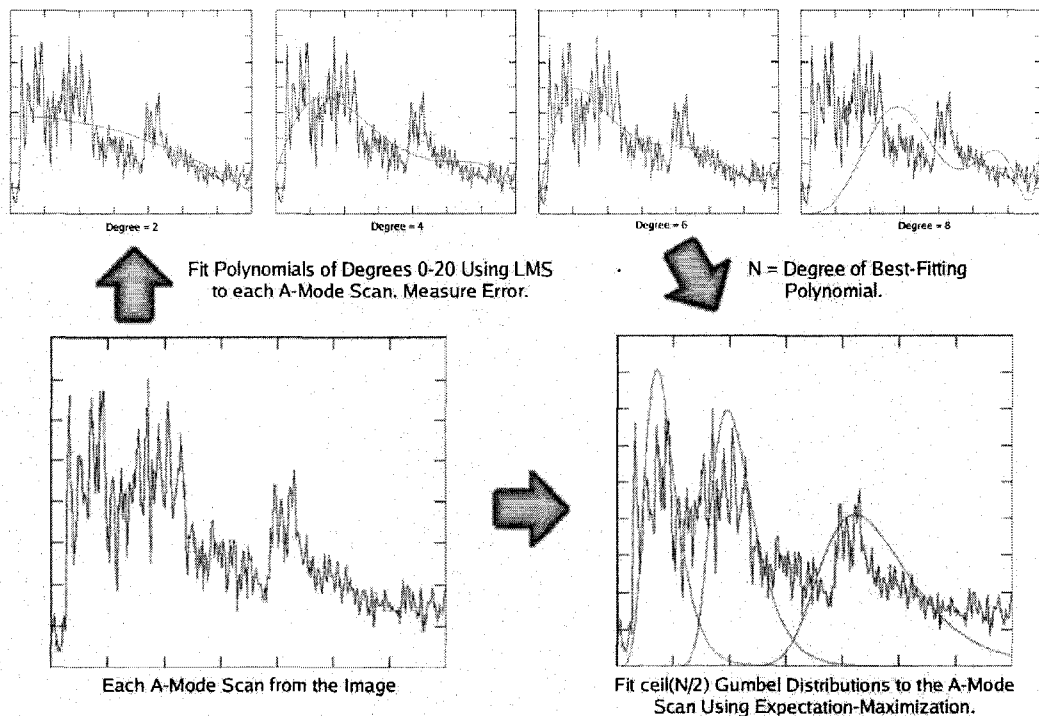


Figure 3.10: A Sequential View of the Model Fitting Process

much faster than multiple EM passes as the polynomial fitting can be done rather efficiently.

With the estimate of the number of tissue boundaries within a Synthetic A-Mode ultrasound scan, we can now fit the proper number of Gumbel distributions to each scan and, using the modes of each distribution, obtain likely tissue boundary locations within the ultrasound image. This entire process can be seen in Figure 3.10

3.3 Contour Initialization

With the knowledge of possible boundary locations now at hand, it is reasonable to progress with a contour-based segmentation approach. In order to do so, an initial contour is required. Given that most anatomical structures have a smooth closed contour, it only seems appropriate to use an ellipse as an initialization for the contour-based segmentation. Yet, at what location in the image should this ellipse be placed? At what orientation? How long should its major and minor axes be? Despite limiting the initial contour to an ellipse, there are still many options in

```

ELLIPSERANSAC(pts)
1  for i ← 1 to N do
2    points ← RANDOM(pts, 5)
3
4    ellipses(i) ← FITELLIPSE(points)
5    ellipseScore(i) ← SCOREELLIPSE(ellipses(i))
6  endfor
7
8  return ellipses(MAXINDEX(ellipseScore))

```

Figure 3.11: RANSAC Algorithm used to Obtain The Initial Elliptical Contour for the Region of Interest.

determining the ultimate shape, size, and orientation of this contour. This freedom, however, is not without complexity.

The ultimate goal in initializing the contour is to use the set of boundary points already obtained from the analysis of the Synthetic A-Mode scans to place and size an ellipse that best represents the region of interest. Unfortunately, this set of likely boundary points includes not only boundary points for the region of interest, but also boundary points for other regions in the image and possibly even some erroneous points. As a result, we cannot simply fit an ellipse to all the points in the boundary point set. Instead, a more intelligent algorithm is required.

In order to deal with the likely possibility of outsiders in the boundary point set, I once again rely on a RANSAC approach developed from [23]. As with the line fitting RANSAC algorithm presented in Section 3.1, the algorithm used here selects five points at random from the boundary point set, fits an ellipse to these five points, and scores the ellipse based on a given criteria. Again, the ellipse with the highest score is chosen as the initial contour. This ellipse fitting RANSAC algorithm is presented in Figure 3.11.

The ellipse fitting process is performed directly using the least squares approach derived by Fitzgibbon *et al.* [18]. Consider the formula for a general conic, given as:

$$F(\mathbf{a}, \mathbf{x}) = \mathbf{a} \cdot \mathbf{x} = ax^2 + bxy + cy^2 + dx + ey + f = 0 \quad (3.4)$$

where $\mathbf{a} = [a \ b \ c \ d \ e \ f]^T$ and $\mathbf{x} = [x^2 \ xy \ y^2 \ x \ y \ 1]^T$. Given the scale invariance

of the conic equation, the ellipse constraint can be written as $4ac - b^2 = 1$, or in matrix form as:

$$\mathbf{a}^T \begin{bmatrix} 0 & 0 & 2 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 \\ 2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \mathbf{a} = 1 \quad (3.5)$$

Given a set of n points \mathbf{D} , where $\mathbf{D} = [\mathbf{x}_1 \ \mathbf{x}_2 \ \cdots \ \mathbf{x}_n]^T$, the ellipse fitting problem is reduced to minimizing $E = \|\mathbf{D}\mathbf{a}\|^2$ according to the constraint in Equation 3.5. By introducing a Lagrange Multiplier λ and differentiating, the following generalized eigenvalue system can be obtained:

$$\mathbf{D}^T \mathbf{D} \mathbf{a} = \lambda \mathbf{C} \mathbf{a} \quad (3.6)$$

where \mathbf{C} is the constraint matrix from Equation 3.5. Fitzgibbon *et al.* showed that the generalized eigenvector associated with the only positive generalized eigenvalue of the above system is an ellipse that best fits the points given according to the algebraic error given above [18].

Prior to scoring the sample ellipses, those whose eccentricity is greater than 0.95 are discarded for efficiency's sake. Clearly, any ellipse with a large eccentricity will be too flat and too elongated to properly represent most anatomical structures. The eccentricity of each ellipse is calculated from the obtained conic equation using the method presented in [5].

Once the remaining sample ellipses are obtained, they are scored based on the following five criteria:

- Boundary Support N_{bdy} . The number of points within a distance t_{near} of the ellipse.
- Homogeneous Region Support S_{reg} . The largest percentage of the ellipse that is not split by points in the boundary point set.
- Depth Support S_{depth} . The amount of the ellipse that is at a depth in the image where there are a lot of boundary points.

- Separation Support N_{diff} . The portion of the ellipse whose nearest boundary point, for both the upper and lower portions of the ellipse, is not the same point.
- Image Inclusion Support N_{out} . The portion of the ellipse that is within the image area.


A weighted combination of these five support measures is used to obtain an ellipse's score:

$$Score(\mathbf{a}) = w_{bdy} N_{bdy} + w_{reg} S_{reg} + w_{depth} S_{depth} + w_{diff} N_{diff} + w_{out} N_{out} \quad (3.7)$$


The first four support criteria used are calculated using the Synthetic A-Mode ultrasound scans. The points of intersection between the scan and the ellipse are calculated and these points are then compared to the boundary points obtained from modeling the scan. The distance between the ellipse points and the nearest boundary points are measured and compared to t_{near} for boundary support, as seen in Figure 3.12(a). Separation support simply finds the closest boundary point to each of the two ellipse intersection points and determines whether the same boundary point is detected in both cases, as shown in Figure 3.12(b). If this is frequently the case, then the sample ellipse may be focused on one tissue boundary and not two.

The homogeneous region support requires two measures seen in Figure 3.13(a): the distance between the two ellipse intersection points (E_{size}), and the largest interval between the two ellipse intersection points that is not interrupted by a point in the boundary point set (R_{size}). These two measures refer to the ellipse size as well as the homogeneous region size. Both are used to obtain the percentage of the ellipse that can be considered as a homogeneous region.

The depth support is a little more involved. Naturally, if there are a large number of boundary points detected at a particular depth in the image, then it is likely that there exists a tissue boundary at that depth. Given the boundary point locations, likely boundary depths are calculated using the Algorithm presented in Figure 3.14. For each row in the image area, the fraction of that image row that contains boundary points is calculated. This creates a form of histogram over image depth. Once




(a) Boundary Support. $p1$ Supports the Boundary, $p2$ does not.




(b) Separation Support. The Left Scan Line Shows Good Separation Support, the Right Scan Line does not.

Figure 3.12: Visual Interpretation of the Calculations used for Boundary and Separation Support Measures.



(a) Homogeneous Region Support. R_{size} Shows the Largest Homogeneous Region Inside the Ellipse.



(b) Depth Support. d is Compared to the Mean and Standard Deviation of the Nearest Likely Depth.

Figure 3.13: Visual Interpretation of the Calculations used for Region and Depth Support Measures.

```

GETLIKELYDEPTHS( $S_{bndy}, I, L_{left}, L_{right}, N_{scans}$ )
1  foreach row  $r$  in  $I$  do
2     $imgSize \leftarrow$  GETIMAGEBOUNDARY( $L_{right}, r$ ) - GETIMAGEBOUNDARY( $L_{left}, r$ )
3     $bndyPtsInRow \leftarrow \sum_{\forall p \in S_{bndy}} \delta(p(y) == r)$ 
4     $pctBoundary(r) \leftarrow 100 \frac{bndyPtsInRow}{imgSize}$ 
5  endfor
6
7   $depthSamples \leftarrow$  CONVERTVALUETOSAMPLES(ROUND( $pctBoundary$ ))
8   $numLikelyDepths \leftarrow$  ROUND( $|S_{bndy}| / N_{scans}$ )
9  return EXPMAXGAUSSIAN( $depthSamples, numLikelyDepths$ )

```

Figure 3.14: Algorithm to obtain Likely Tissue Boundary Depths Expressed using Gaussian Distributions.

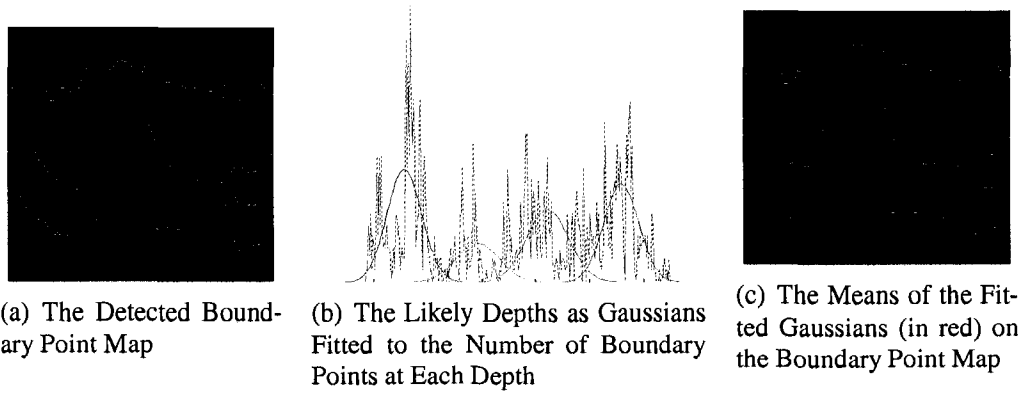


Figure 3.15: A Visual Description of Obtaining the Most Likely Depths in an Image.

complete, the Expectation-Maximization algorithm is used to fit N Gaussian distributions to the depth histogram, where N is the average number of boundary points from one of the image's Synthetic A-Mode scans (The EM algorithm used here is identical to the one in Figure 3.8 used for fitting the Gumbel distributions, except lines 20 and 21 are removed for obvious reasons). This process can be seen in Figure 3.15.

With the likely depths described using Gaussian functions, the Z-score for each ellipse intersection points can be obtained for each of the Gaussians. The lowest Z-score for each ellipse point is then chosen and negated to represent depth support. This process is illustrated in Figure 3.13(b)

Finally, the image inclusion support merely counts and negates the number of

ellipse points that are along the boundary of the image area. This score is usually weighted highly to ensure that the ellipse chosen as the initial contour is entirely within the image area.

The entire ellipse scoring algorithm is presented in Figure 3.16. For each Synthetic A-Mode scan i , its line equation $lines(i)$ and boundary points $p_{bndy}(i)$ are maintained.

3.4 Force-Based Segmentation

With the pre-processing now complete and the initial contour obtained, the segmentation process can now commence in shaping and moving the contour towards the proper region boundaries. This deformation is performed by applying forces of various types to the contour. The contour, in turn, will act much like an elastic band, stretching and flexing until an equilibrium is reached between the forces acting on the contour and the elasticity of the contour itself. Once this equilibrium has been reached, the segmentation process will be considered complete.

For the purposes of this algorithm, the contour S is represented as a set of contour points p . In order to maintain the elastic quality of the contour, the distance between each contour point is kept within an interval $[d_{min}, d_{max}]$, where d_{min} and d_{max} are chosen ahead of time to ensure that the elasticity of the contour is around the same order of magnitude as the forces acting on the contour. In practice, the values $d_{min} = 5 \text{ pixels}$ and $d_{max} = 40 \text{ pixels}$ work well for ultrasound images of various sizes and content. If the distance between two contour points falls out of this given interval, a contour point is either added or removed as appropriate. With these distance thresholds in mind, the initial contour point set is generated by selecting points on the initial ellipse that are 20 pixels apart.

The forces used in the segmentation are applied directly to the contour points in an iterative fashion with the contour points being moved in each iteration to a neighboring pixel location in the direction of the prevailing force. In a similar vein as Active Contours, the forces acting on each contour point can be classified into

```

CALCELLIPSE SUPPORT(ellipse, lines, pbndy,  $\mu_{depth}$ ,  $\sigma_{depth}$ )
1  foreach l in lines do
2    ellipsePts  $\leftarrow$  INTERSECT(ellipse, lines(l))
3
4    foreach x in ellipsePts do
5      foreach y in pbndy(l) do
6         $dist(x, y) \leftarrow \sqrt{(x - y)^2}$ 
7      endfor
8      if MIN(dist(x, :)) <  $t_{near}$  then
9         $N_{bndy} ++$ 
10     endif
11   endfor
12
13   if MININDEX(dist(ellipsePts(1), :))  $\neq$ 
14     MININDEX(dist(ellipsePts(2), :)) then
15      $N_{diff} ++$ 
16   endif
17
18    $ellipseSize(l) \leftarrow \sqrt{(ellipsePts(1) - ellipsePts(2))^2}$ 
19    $ptsOnInEllipse \leftarrow \{x | (x \in ellipsePts \vee x \in p_{bndy})$ 
20      $\wedge ellipsePts(1) \leq x \leq ellipsePts(2)\}$ 
21   foreach x in ptsOnInEllipse do
22     foreach y in ptsOnInEllipse do
23        $distInEllipse(x, y) \leftarrow \sqrt{(x - y)^2}$ 
24     endfor
25   endfor
26    $regionSize(l) \leftarrow MAX(distInEllipse)$ 
27
28   foreach x in ellipsePts
29     foreach d in  $\mu_{depth}$ 
30        $Z(x, d) = \frac{x - \mu_{depth}(d)}{\sigma_{depth}(d)}$ 
31     endfor
32      $S_{depth} \leftarrow S_{depth} - MIN(Z(x, :))$ 
33   endfor
34
35    $S_{reg} \leftarrow 100 \frac{\sum_{l \in lines} regionSize}{\sum_{l \in lines} ellipseSize}$ 
36
37   imgBoundaries = [Lleft, Lright, Rtop, Rbtm]
38   foreach l in imgBoundaries do
39     ellipsePts  $\leftarrow$  INTERSECT(ellipse, imgBoundaries(l))
40     if ellipsePts  $\neq \emptyset$  then
41        $N_{out} \leftarrow N_{out} - \sqrt{(ellipsePts(1) - ellipsePts(2))^2}$ 
42     endif
43   endfor
44
45   supports  $\leftarrow [N_{bndy}, S_{reg}, S_{depth}, N_{diff}, N_{out}]^T$ 
46   return weights  $\cdot$  supports

```

Figure 3.16: The Calculation of Ellipse Support Values.

one of two types: internal or external [31].

$$F(\mathbf{p}) = F_{ext}(\mathbf{p}) + F_{int}(\mathbf{p}) \quad (3.8)$$

External forces can be considered as those forces which are applied to the contour from an external source while internal forces represent more intrinsic properties of the contour, such as elasticity and smoothness. Both force classes themselves are also made up of multiple forces, the combination of which is a robust yet flexible framework through which a segmentation result can be obtained.

3.4.1 External Forces

The role of external forces is to drive the contour towards the boundary of the region of interest. For this task, knowledge of the region boundary is necessary. Once again, this knowledge comes in the form of the boundary points obtained from the Synthetic A-Mode ultrasound scans. These detected boundary points are used to generate two separate external forces:

$$F_{ext}(\mathbf{p}) = \mathbf{w}_{inS} F_{inS}(\mathbf{p}) + \mathbf{w}_{outS} F_{outS}(\mathbf{p}) \quad (3.9)$$

Intuitively, F_{inS} and F_{outS} represent the forces applied to the contour by boundary points within and outside the closed contour respectively. As these boundary points may represent the true tissue boundary, the force generated by these boundary points should pull the contour towards them. As a result, the outside boundary point force is defined as follows:

$$F_{outS}(\mathbf{p}) = \sum_{\forall \mathbf{v} \in BP} \left(1 - \frac{\mathbf{v}(y)}{\max(y)}\right) g(\mathbf{v}, \mathbf{p}) \quad (3.10)$$

where

$$g(\mathbf{v}, \mathbf{p}) = \begin{cases} \mathbf{v} - \mathbf{p} & \text{if } \forall \mathbf{p}' \in S, \mathbf{p} = \min(\mathbf{v} - \mathbf{p}'); \\ 0 & \text{otherwise.} \end{cases} \quad (3.11)$$

Essentially, each boundary point outside the closed contour pulls on its nearest contour point. The strength of this force is equivalent to the distance between the boundary point and the contour point. This force is then weighted in Equation 3.10 inversely to the boundary point's depth in the image. This weighting is done to

reflect the increased presence of noise deeper in ultrasound images due to signal attenuation. The inside boundary point force is defined almost identically.

As I noted earlier however, not all of the detected boundary points truly relate to the region of interest. Equation 3.10 does not take this into account. The key difference between the two boundary point forces is how they deal with noise in the boundary point set. The inside boundary point force simply imposes an envelope of width ρ around the inside of the contour. Any boundary point that is further than ρ pixels within the closed contour is disregarded as noise. Aside from this additional criteria, the inside boundary point force is identical to that defined in Equation 3.10.

The outside boundary point force uses a more stringent approach in dealing with likely erroneous boundary points. Reconsider once again the likely tissue boundary depths calculated using the algorithm in Figure 3.14. The Gaussian distributions obtained from this algorithm conceptually represent separate tissue boundaries. Comparing these Gaussian distributions with the initial elliptical contour, we can determine which boundary points are most likely referring to the region of interest.

Consider the set of detected boundary points. These points can be classified according to the Gaussian distributions obtained from the algorithm in Figure 3.14. Conceptually, these classes contain the boundary points for separate tissue boundaries.

The initial contour points can be classified in the same way. Assuming a reasonable initial contour, any class in which at least one contour point has been classified may contain boundary points related to the region of interest. Conversely, any class that does not contain a contour point can be assumed to represent a tissue boundary that does not belong to the region of interest. Therefore, any boundary points within that class are simply ignored. This filtering of erroneous boundary points is presented in Figure 3.17. Note that since this filtering process relies on the location of the initial contour, it is important for the contour to be initialized at the appropriate depth.

With erroneous boundary points filtered away, it is assumed that the remaining boundary points relate to the region of interest. As much as this filtering makes the

```

FILTERBOUNDARYPTS( $p_{bdy}$ ,  $\mu_{depth}$ ,  $\sigma_{depth}$ ,  $p_{contour}$ )
1  foreach  $x$  in  $p_{bdy}$  do
2    foreach  $u$  in  $\mu_{depth}$  do
3       $probDepth(x, u) \leftarrow \frac{e^{-\frac{(x-\mu_{depth}(u))^2}{2\sigma_{depth}(u)^2}}}{\sqrt{2\pi}\sigma_{depth}(u)}$ 
4    endfor
5     $depthClass(x) \leftarrow \text{MAXINDEX}(probDepth(x, :))$ 
6  endfor
7
8  foreach  $x$  in  $p_{contour}$  do
9    foreach  $u$  in  $\mu_{depth}$  do
10      $probContour(x, u) \leftarrow \frac{e^{-\frac{(x-\mu_{depth}(u))^2}{2\sigma_{depth}(u)^2}}}{\sqrt{2\pi}\sigma_{depth}(u)}$ 
11   endfor
12    $depthContour(x) \leftarrow \text{MAXINDEX}(probContour(x, :))$ 
13 endfor
14
15  $unsupportedClasses \leftarrow depthClass \setminus depthContour$ 
16  $unsupportedPoints \leftarrow \text{FIND}(depthClass == unsupportedClasses)$ 
17  $filteredPoints \leftarrow p_{bdy} \setminus unsupportedPoints$ 
18 return  $filteredPoints$ 

```

Figure 3.17: The Boundary Point Filtering Algorithm Expressed using Gaussian Distributions.

algorithm robust to noise, it also has the added benefit of allowing the external force to be applied globally to the entire image space. By using all remaining boundary points, the contour is less likely to get caught by localized noise.

Finally, note that both boundary point forces are weighted differently in the horizontal direction as they are in the vertical direction through the use of weight vectors. This differential weighting is a result of more pronounced vertical boundaries compared to horizontal boundaries.

A visual representation of the effects of the external forces on a contour is shown in Figure 3.18.

3.4.2 Internal Forces

Internal forces are used generally to keep the contour consistent and smooth and as such are often referred to as regularization forces. This application is no different in that regard. What is unique in this context is the possibility of significant noise




Figure 3.18: The External Boundary Point Force and how and where it interacts with the Contour

as well as large gaps along the region boundary, particularly vertically. As a result, a significant effort must be done to ensure the contour keeps its integrity. This goal is achieved through the use of two internal forces:

$$F_{int}(\mathbf{p}) = w_{curv} F_{curv}(\mathbf{p}) + w_{prop} F_{prop}(\mathbf{p}) \quad (3.12)$$

The first internal force is known as a curvature force and is similar to the one used by Kass *et al.* in their Active Contour framework [31]. In fact, the curvature force is simply the finite difference approximation of the contour's second derivative:

$$F_{curv}(\mathbf{p}_i) = \mathbf{p}_{i-1} - 2\mathbf{p}_i + \mathbf{p}_{i+1} \quad (3.13)$$

It is this curvature force that provides the contour with its elastic properties as well as enforcing a certain amount of smoothness. This force is also used while maintaining the contour itself. As mentioned, when the contour grows, new contour points have to be added to maintain a consistent curvature force. To ensure a greater smoothness to the contour, these contour points are only added to the contour if that portion of the contour has a low curvature. Heuristically, areas of low curvature are chosen to be portions of the contour where a contour points satisfies:

$$F_{curv}(\mathbf{p}) < \left[\frac{d_{max}}{4} \quad \frac{d_{max}}{4} \right] \quad (3.14)$$




Figure 3.19: The Propagation of External Forces expressed in relation to the Contour Points

While the curvature force provides some smoothing effects to the contour, its impact may not be enough. As a result, an additional internal force is added. The effect of the propagation force is to distribute the external forces more smoothly over the contour points. This is achieved by adding a portion of the external forces from neighboring contour points to a given contour point:

$$F_{prop}(p_i) = \frac{F_{ext}(p_{i-1}) + F_{ext}(p_{i+1})}{2} \quad (3.15)$$

Conceptually, the propagation force has an impact similar to applying a low-pass filter to the contour point's external forces. Note that the addition in Equation 3.15 is a vector addition as shown in Figure 3.19.

These two forces counterbalance the external forces and help shape the resulting contour by ultimately providing an equilibrium between the forces.

Halting Criteria

Given that there is no threshold on the force required to move a contour point, there is little chance that any contour point will reach a fixed location. Instead, when an equilibrium state is achieved, contour points usually end up oscillating forward and backward within a small region. Detecting this equilibrium state can be difficult. However, it is known that once in an equilibrium state, the number of contour points defining the region boundary will not change. Also known is that contour points are separated by a distance in the interval of $[d_{min}, d_{max}]$. As each iteration moves a

contour point to a adjacent location, it is safe to assume that any addition or removal of contour points must be done within $2 \cdot d_{max}$ iterations.

As a result, if $2 \cdot d_{max}$ iterations go by without any contour points being added or removed from the contour, it is assumed that the contour has reached the equilibrium state. Therefore, once this criteria is met, the algorithm has halted. Cubic B-Splines are then used to interpolate the final contour.

3.5 Conclusions

The ultrasound image segmentation system presented in this chapter can be best described as having four main components:

- *Image decomposer*: The creation of Synthetic A-Mode ultrasound scans that mimic the original ultrasound signals.
- *Echo Pattern Modeler*: The modelling of the echo patterns of a tissue boundary in a Synthetic A-Mode ultrasound scan with Gumbel distributions. Boundary points locations are obtained by fitting Gumbel distributions to the Synthetic A-Mode scans.
- *Contour Initializer*: The creation of an initial estimate of the boundary of the region of interest by fitting an ellipse to a sample set of boundary points. The best ellipse is chosen according to a detailed scoring criteria to act as the initial contour.
- *Force-Based Segmenter*: A robust set of external forces is used to drive the contour to a proper final segmentation. Strong internal forces are used to maintain the contour in an smooth and elastic fashion.

The main strength of this system is the knowledge used to guide the segmentation process. The algorithms presented herein rely only on the observations made on the ultrasound imaging process. Table 3.1 displays the key assumptions relied upon and how they are incorporated into the my segmentation approach.

| Assumption | Use In Algorithm |
|--|--|
| 3. Noise is more likely at deeper parts of the image than in more shallow portions of the image. | The segmentation forces applied by the boundary points are weighted inversly by their depth in the image. |
| 4. Horizontal tissue boundaries will appear brighter than vertical tissue boundaries. | The segmentation forces applied by the boundary points have separate weights for the horizontal and vertical directions. |
| 5. Tissue boundaries will appear fuzzy in the image. | The fuzzy tissue boundary is appropriately modeled by the smooth peak of a Gumbel distribution. |
| 6. Pixel intensity will decay following a tissue boundary. | This decay in intensity is captured by the tail of the Gumbel distribution that models the tissue boundary. |
| 7. Ultrasound images can be conceptually considered as an array of A-Mode ultrasound scans. | The main rationale used to split up the ultrasound image into Synthetic A-Mode ultrasound scans. |
| 8. The region of interest has a smooth, closed contour. | An ellipse is used to initialize the contour of the region of interest. |

Table 3.1: Ultrasound Image Assumptions and how they are incorporated into The Algorithm.

Another strength of the proposed method is its ability to overcome noise in the images. This strength stems from the fact that the image data is analysed using a model of the image acquisition process. Significant pre-processing is used to ensure that the most reliable region boundaries were detected and incorporated into the segmentation process. Even so, additional noise reduction is performed within the segmenter through the use of strong internal forces and boundary point filtering to lessen the effect of misleading boundary information on the end result.

Chapter 4

Experimental Results

The proposed ultrasound image segmentation system was created and tested on multiple data sets from both the medical and agricultural fields. The resulting tests were run using Octave version 2.9.9 under Scientific Linux, release 4.4 on a Intel® Pentium® 3.40 GHz HT machine.

Three data sets were constructed from ultrasound images of pork loins, human kidneys, and human livers. Great care was taken to ensure that the images in each data set contained significant variation not only in the subjects scanned, but also in terms of ultrasound probe location, probe type, scanner settings, the health of the subjects scanned, and even the manufacturer of the ultrasound scanners used to generate the images. As a result, significant differences in image quality are apparent in each data set as well as wide variation in the shape, size, location, and orientation of the region of interest.

The inclusion of such a variety of ultrasound images within the test sets is deliberate. The goal of these tests is to determine how the algorithm performs in a practical setting. In order to achieve this goal, the ultrasound images used to test the proposed segmentation system were selected by an independent radiologist who had no knowledge of the algorithm or of the the image features used therein.

In this chapter, I present the results of these tests. While obtaining the resulting segments of each image is straightforward, determining the quality of the segment's contours is not as clear of a procedure. Regardless, an attempt is made to quantify the segmentation results in a viable manner.

4.1 Validation Issues in this Field

On the whole, the issue of validating results in the field of ultrasound image segmentation has been a problematic one. There are various aspects of this problem that not only make it difficult to properly compare results within the field, but can also make it a dubious task to interpret the quality of any obtained result.

The majority of quantitative results in this field stem from a comparison between the segmentation result and a segmentation manually performed by an expert [38]. In essence, the expert's segmentation is taken as a ground truth. Unfortunately, there is significant variation between a region of interest traced by one expert as opposed to one traced by another. Intra-observer variation is also often quite high as tracing out an image region by hand is a task rarely performed by radiologists. With no ethical way of determining the accuracy of the expert's manual segmentation, we can only safely assume that the manual trace is, at best, an approximation.

Occasionally, inter-observer and intra-observer variance values will be used as a standard upon which segmentation systems are compared [9, 14, 21]. This approach would be the most appropriate way of assessing the quality of a segmentation algorithm, however its use in the field is often limited by the availability of experts to perform the manual segmentation.

An additional concern in interpreting segmentation results is the number of images used in the testing phase. Due to ethics constraints and the busy schedules of radiologists, the size of most test sets is often less than fifty images [38]. With such a small set of samples, it is difficult to generalize the results obtained by the segmentation algorithm for general use.

What is more concerning though is the role ethics constraints play when it comes to comparing results within the field. As of yet, no publicly available data sets exist for ultrasound image processing. For ethical reasons, no research group has been able to present their ultrasound image for public use. Therefore, each ultrasound image segmentation algorithm presented in the field has been tested on separate data sets. Surely, the quality of, and variation within, these data sets may vary significantly from one group to the next. Any direct quantitative comparison

between two approaches has limited analytical value as the difficulties encountered in one test set may not exist in another.

Not only do the differences in test sets affect the value of a direct quantitative comparison, it also limits the ability to duplicate the methods and the results obtained by others. It is not unusual to have algorithms be built specifically around a particular data set. As all ultrasound image segmentation algorithms include a particular amount of prior knowledge, each method is, by definition, limited in their ability to be generally applied. The strength of gradients, the quality of shape and texture cues, and the value of intensity variation between regions all vary from one data set to the next. Recreating the algorithm of another and applying it to the ultrasound images you have may not provide you with an appropriate comparison as that other person's algorithm may not be as well suited to your problem.

Finally, different segmentation approaches often use different measures to determine the quality of their results. Some researchers use boundary correlation [9, 19] while others measure the distance between contours [14, 1], while still others provide area-related scores [22, 11]. In some cases, only qualitative results are presented [16, 24], or segmentation quality is scored using an application-specific criteria that cannot be generalized [2, 29]. This variation in quality measures simply adds to the difficulty of assessing the results of various approaches in the field.

For the purposes of assessing the quality of the segmentation system presented herein, two of the most common qualitative measures will be used: the distance between the expert's contour and the resulting contour, and the overlap of the resulting segment and the expert's traced segment. The contour distance is measured as the average shortest distance between the two contours, in pixels. Meanwhile, the overlap score is defined as,

$$Score(A) = \frac{A \cap M}{A \cup M} \times 100 \quad (4.1)$$

where A is the set of pixels in the region segmented by the algorithm and M is the set of pixels in the region segmented manually by the expert. Note that this score provides a percentage area accuracy measure. Despite the limited value in doing so, these quantitative results will be directly compared to those of similar

| Ellipse Fitting Parameters | | | Segmentation Force Parameters | | |
|----------------------------|-------------|-----|--------------------------------------|---------------|------|
| Fraction of Outliers | ϵ | 0.7 | Bndy. Pixel Weight (Inside Contour) | $w_{inS}(x)$ | 0.00 |
| Boundary Sup. Weight | w_{bndy} | 1.0 | Bndy. Pixel Weight (Outside Contour) | $w_{inS}(y)$ | 0.95 |
| Bndy. Sup. Dist. Thresh. | t_{near} | 10 | Curvature Weight | $w_{outS}(x)$ | 0.85 |
| Region Sup. Weight | w_{reg} | 1.0 | Propagation Weight | $w_{outS}(y)$ | 0.25 |
| Depth Support Weight | w_{depth} | 1.0 | Contour Env. Width | w_{curv} | 0.10 |
| Separation Sup. Weight | w_{diff} | 1.0 | | w_{prop} | 0.10 |
| Img. Incl. Sup. Weight | w_{out} | 5.0 | | ρ | 20 |

Table 4.1: Parameter settings used for the segmentation of 304 pork loin ultrasound images

studies. Keep in mind though that the data on which these algorithms were tested is different. Also keep in mind the amount of prior knowledge used in each approach and how that may affect the ultimate results.

Due to time constraints, I was only able to obtain one manually traced contour for each of the images in my data sets. As a result, inter-observer or intra-observer results unfortunately cannot be presented.

4.2 Agricultural Images Test Case

With the help of the Lacombe Research Centre of Agriculture and Agri-Food Canada, a collection of 304 ultrasound images of *in vivo* pork loins were obtained for testing purposes. The ultrasound images were recorded with an Aloka Flexus Model SSD-1100 equipped with a 3.5MHz/127mm transducer Ultrasound system. The images are from between the 3rd and 4th ribs from the last rib and 7cm off the mid-line of the pigs. Due to the positioning of the loin with respect to the pig's ribcage, variations in probe position could not be achieved. Variation in probe type and scanner manufacturer was also unable to be achieved due to financial constraints. However, this data set still contains variation due to the large number of subjects included as well as the settings of the ultrasonic scanner. In many ways, this type of data set is similar to what is normally used in the field for testing purposes.

The images were segmented using the the proposed segmentation system tuned with the parameter settings given in Table 4.1. Parameter settings were chosen

| Statistical Measure | Contour Distance (in pixels) | Overlap Score (%) |
|----------------------------|-------------------------------------|--------------------------|
| Mean | 15.915 | 72.703 |
| Std. Deviation | 7.598 | 11.252 |
| Minimum | 5.852 | 36.771 |
| Maximum | 44.400 | 88.949 |


Table 4.2: Statistical results for the algorithm on 304 ultrasound images of pork loins

empirically as to obtain the best results. Contour distances and overlap scores were recorded of each segmentation result and are summarized in Table 4.2. On average, the pork loins were segmented with 73% overlap accuracy and the contours differed by approximately 16 pixels.

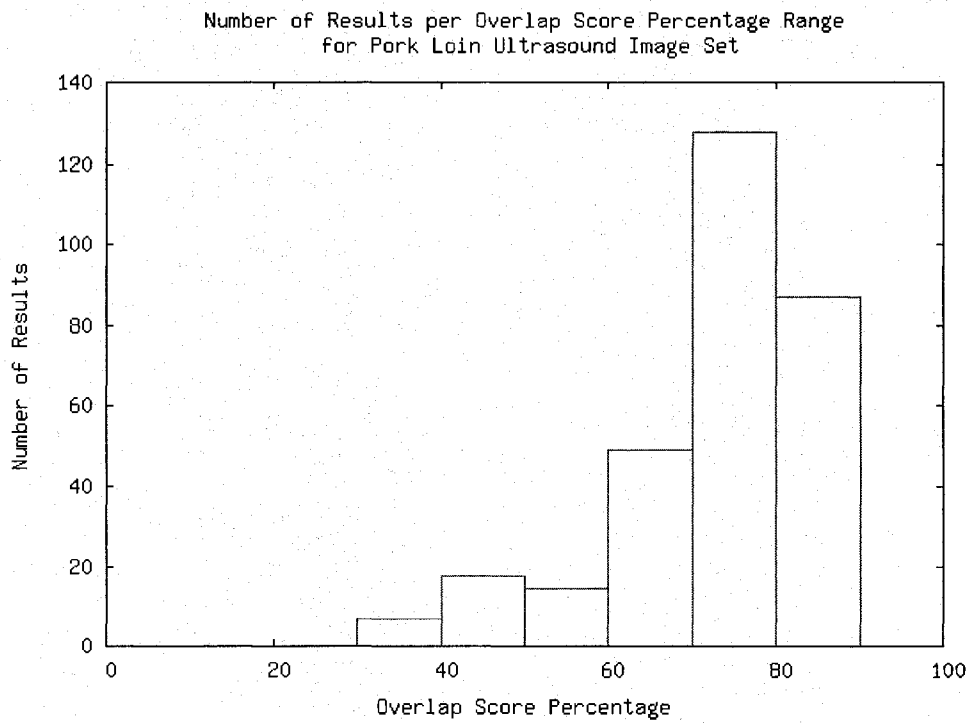
What is particularly of note though is the relatively large standard deviation in both cases. Upon further inspection, we see that these scores have the distributions shown in Figure 4.1, suggesting that for the majority of cases, the results are as good or better than the averages given in Table 4.2. These graphs also show a certain number of extreme cases that did not score well. An example of one of these extreme cases is shown in Figure 4.2.

The extreme poor cases similar to the one shown in Figure 4.2 are mainly the result of a poor initial contour. In these cases, reverberation effects above the skin layer distort the location of boundary points on the left side of the image, leading to a poor contour initialization. While the force-based segmentation is capable of overcoming this reverberation effect, the filtering of boundary points based on the initial ellipse location limits the boundary points that can apply a force on the contour. In the case shown in Figure 4.2, the boundary points relating to the interest region's bottom contour were mistakenly removed from the segmentation due to the poor location of the initial ellipse. It is possible that alternative parameter settings within the initial contour ellipse fitting procedure could have avoided this issue. In fact, allowing the parameter settings to somehow vary with respect to the image content would likely lead to better results.

Figure 4.3 shows an the results for an average case. Note that in general, the initial contour is reasonably well placed and the force-based segmentation is able




(a) Contour Distance Distribution




(b) Overlap Score Distribution


Figure 4.1: The Distribution of Scores for the Pork Loin Ultrasound Image Set.




(a) Original Image with Manually Traced Expert Contour in Green.



(b) The Detected Boundary Points (in white) and their relation to the Most Likely Boundary Depths (in red).




(c) The Fitted Ellipse Initial Contour.



(d) The Boundary Points Filtered by Depth and Ellipse Location.




(e) The Final Countour - The Segmentation Result.



(f) The Manually Traced Region Overlaid on the Detected Segment.


Figure 4.2: An Example of a Poor Result seen from using the Segmentation System on an Ultrasound Image of a pork loin (Contour Distance: 36.16 pixels, Overlap Score: 36.77%).




(a) Original Image with Manually Traced Expert Contour in Green.




(b) The Detected Boundary Points (in white) and their relation to the Most Likely Boundary Depths (in red).




(c) The Fitted Ellipse Initial Contour.



(d) The Boundary Points Filtered by Depth and Ellipse Location.




(e) The Final Countour - The Segmentation Result.




(f) The Manually Traced Region Overlaid on the Detected Segment.


Figure 4.3: An Average Example of the Results seen from using the Segmentation System on an Ultrasound Image of a pork loin (Contour Distance: 14.92 pixels, Overlap Score: 72.36%).




(a) Original Image with Manually Traced Expert Contour in Green.




(b) The Detected Boundary Points (in white) and their relation to the Most Likely Boundary Depths (in red).




(c) The Fitted Ellipse Initial Contour.



(d) The Boundary Points Filtered by Depth and Ellipse Location.



(e) The Final Contour - The Segmentation Result.



(f) The Manually Traced Region Overlaid on the Detected Segment.

Figure 4.4: An Example of a Good Result seen from using the Segmentation System on an Ultrasound Image of a pork loin (Contour Distance: 5.98 pixels, Overlap Score: 88.95%).

to move the contour with reasonable ease despite the presence of noise. These observances also hold true for better results as seen in Figure 4.4.

For comparison's sake, consider additional works in the area of agricultural ultrasound image segmentation. The most closely related work is that of Tong *et al.* who also segmented pork loin images. Their results on thirty pork loin images were used to estimate the muscle volume. The estimated muscle volume was then compared to the volume of the loin measured from the carcass of the pig and their method was found to predict 79% of the loin's variation in size [42]. Unfortunately, this method is difficult to compare to mine given the quantitative measure used.

Hao *et al.* present a method for segmenting infarcted, ischemic, and normal regions of pig's hearts from echocardiographs. Their results, on average show an overlap score of 62% [22]. Cancela *et al.* segment the rib eye area in cattle in sixty ultrasound images and obtain a score above 85% for 80% of their images [11]. Meanwhile, Glasbey segments 144 ultrasound images of sheep loins with a correlation of muscle depths of 52% [19].

On average, the results presented herein fall within the same range as those presented in the field of agricultural ultrasound image segmentation while using a much larger data set for testing and unlike Cancela *et al.* and Glasbey, my approach does not rely on learnt shape information. While this comparison is dubious in the sense that separate test sets are used, these results still show that my proposed approach is appropriate and does not include knowledge that may limit its practical application.

Preliminary results from this method were shown in [6].

4.3 Medical Images Test Cases

With the cooperation of Capital Health and the University of Alberta hospital, two sets of medical ultrasound images were obtained for testing purposes: a set of 30 kidney images and one of 20 liver images. Unlike the agricultural ultrasound images, the images within these test sets were obtained from various curved array ultrasound probes with a myriad of settings. The liver and kidney images were also

| Ellipse Fitting Parameters | | | Segmentation Force Parameters | | |
|----------------------------|-------------|------|-------------------------------|---------------|------|
| Fraction of Outliers | ϵ | 0.7 | Bndy. Pixel Weight | $w_{inS}(x)$ | 0.10 |
| Boundary Sup. Weight | w_{bndy} | 9.0 | (Inside Contour) | $w_{inS}(y)$ | 0.75 |
| Bndy. Sup. Dist. Thresh. | t_{near} | 10 | Bndy. Pixel Weight | $w_{outS}(x)$ | 0.10 |
| Region Sup. Weight | w_{reg} | 2.0 | (Outside Contour) | $w_{outS}(y)$ | 0.03 |
| Depth Support Weight | w_{depth} | 0.0 | Curvature Weight | w_{curv} | 0.05 |
| Separation Sup. Weight | w_{diff} | 2.0 | Propagation Weight | w_{prop} | 1.00 |
| Img. Incl. Sup. Weight | w_{out} | 50.0 | Contour Env. Width | ρ | 50 |

Table 4.3: Parameter settings used for the segmentation of 30 human kidney ultrasound images

obtained from multiple probe locations and cover a group of subjects in various stages of health. The amount of variation in image content in these data sets is far beyond what is normally found in the test sets of similar studies, in particular the variation due to probe position.

4.3.1 Kidney Image Set

The parameter settings for the kidney data set are presented in Table 4.3. Note that the region support criteria is weighted much less than the boundary support criteria. This particular weighting is due to the structure of the kidney itself. The kidney has a dense center known as its sinus which is surrounded by soft tissue and fluid. The kidney sinus is naturally bright in the ultrasound image with the rest of the kidney being relatively dark. As a result of this structure, the kidney sinus is often recognized as a tissue boundary itself, leading to less reliability in the region support score.

Also note that the depth support score is not used in this ellipse fitting procedure. This omission is due to the fact that the kidney often takes up only a small portion of the ultrasound image, with other anatomical structures comprising the rest of the image. Therefore, the detected likely boundary depths are often quite noisy and unreliable.

Results for the kidney data set are summarized in Table 4.4 and displayed in Figure 4.5. As expected, the numerical results are much lower than those obtained from the pork loin data set. The variation in size, shape, and location of the imaged

| Statistical Measure | Contour Distance (in pixels) | Overlap Score (%) |
|----------------------------|-------------------------------------|--------------------------|
| Mean | 34.102 | 48.999 |
| Std. Deviation | 12.480 | 15.504 |
| Minimum | 12.909 | 18.679 |
| Maximum | 62.041 | 74.786 |


Table 4.4: Statistical results for the algorithm on 30 ultrasound images of Human Kidneys

kidneys is simply too large to be properly assessed using a single set of parameters.


Figure 4.6 shows an example of the issues met with this data set. The presence of the kidney sinus within the images led to the ellipse fitting procedure to favor boundary support over region support. As a side effect of this choice, large initial ellipses were favored in many situations where the kidney itself was not that large. For the case in Figure 4.6, the initial contour was poorly chosen, leading to the filtering out of the boundary points associated with the bottom of the kidney. Attempts to use a strong internal force and interior boundary point force could not overcome this poor initialization.

However, there were a few images for which the segmentation results were encouraging. Figure 4.7 shows one of these good kidney segmentation results. In this particular case, the kidney was much larger in the image and better suited to the parameter settings used.

Comparatively, it is difficult to assess the relative quality of the proposed algorithm. Other works in the area of kidney ultrasound image segmentation are very limited due to the difficult qualities of the images. Those proposals that do exist simply do not use data sets with the same variability as the one used for testing here. Xie *et al.* have obtained overlap scores of 96.49% and contour separation distances around 3.25 pixels over their six image data set [44]. However, Xie *et al.* rely on a learnt deformable template in their algorithm and their data set simply features images of the left kidney scanned from the same probe position. Martín-Fernández and Alberola-López also rely on shape through the use of a manually traced initial contour, and while they do not provide any comparable numerical results, their seg-




(a) Contour Distance Distribution



(b) Overlap Score Distribution


Figure 4.5: The Distribution of Scores for the Kidney Ultrasound Image Set.



(a) Original Image with Manually Traced Expert Contour in Green.




(b) The Detected Boundary Points (in white) and their relation to the Most Likely Boundary Depths (in red).



(c) The Fitted Ellipse Initial Contour.



(d) The Boundary Points Filtered by Depth and Ellipse Location.




(e) The Final Countour - The Segmentation Result.



(f) The Manually Traced Region Overlayed on the Detected Segment.


Figure 4.6: An Example of a Poor Result seen from using the Segmentation System on an Ultrasound Image of a human Kidney (Contour Distance: 55.15 pixels, Overlap Score: 18.68%).



(a) Original Image with Manually Traced Expert Contour in Green.




(b) The Detected Boundary Points (in white) and their relation to the Most Likely Boundary Depths (in red).



(c) The Fitted Ellipse Initial Contour.



(d) The Boundary Points Filtered by Depth and Ellipse Location.



(e) The Final Countour - The Segmentation Result.



(f) The Manually Traced Region Overlaid on the Detected Segment.

Figure 4.7: An Example of a Good Result seen from using the Segmentation System on an Ultrasound Image of a human Kidney (Contour Distance: 22.29 pixels, Overlap Score: 74.79%).

| Ellipse Fitting Parameters | | | | Segmentation Force Parameters | | | | |
|----------------------------|-------------|------|--------------------|-------------------------------|------|--------------------|---------------|------|
| Fraction of Outliers | ϵ | 0.7 | Bndy. Pixel Weight | $w_{inS}(x)$ | 0.00 | (Inside Contour) | $w_{inS}(y)$ | 0.01 |
| Boundary Sup. Weight | w_{bndy} | 4.0 | Bndy. Pixel Weight | $w_{outS}(x)$ | 0.50 | (Outside Contour) | $w_{outS}(y)$ | 0.25 |
| Bndy. Sup. Dist. Thresh. | t_{near} | 10 | Curvature Weight | w_{curv} | 0.01 | Propagation Weight | w_{prop} | 1.00 |
| Region Sup. Weight | w_{reg} | 1.0 | Contour Env. Width | ρ | 20 | | | |
| Depth Support Weight | w_{depth} | 0.0 | | | | | | |
| Separation Sup. Weight | w_{diff} | 5.0 | | | | | | |
| Img. Incl. Sup. Weight | w_{out} | 50.0 | | | | | | |

Table 4.5: Parameter settings used for the segmentation of 20 human liver ultrasound images

mentation results do appear to be similar to my better results, such as the one shown in Figure 4.7 [34]. Of course, the framework presented by Martín-Fernández and Alberola-López is not fully automated and therefore differs significantly from the approach presented herein.

Ultimately, no other attempt has been made to address the variability of the data set used herein.

4.3.2 Liver Image Set

The parameter settings for the human liver data set are presented in Table 4.5. The presence of strong speckle noise, as well as occasional masses, within the liver, leads to a significant number of additional boundary points being detected. As a result, both the region support score and the depth support scores used in the ellipse fitting procedure are less reliable for this data set. This degradation in reliability is shown in the weighting of the respective scores in the ellipse fitting procedure. The depth support is once again ignored while the region support is weighted much lower than the others.

Results for the liver data set are summarized in Table 4.6 and displayed in Figure 4.8. The results for the liver data set were better than those seen in the kidney data set, which is to be expected. The liver image set does not have the same level of variability as the kidney image set simply due to the much larger size of the liver, making the liver the dominant feature in many of the images.


However, there was still variation due to the combination of lengthwise and

| Statistical Measure | Contour Distance (in pixels) | Overlap Score (%) |
|----------------------------|-------------------------------------|--------------------------|
| Mean | 34.665 | 60.106 |
| Std. Deviation | 14.286 | 17.109 |
| Minimum | 13.885 | 26.862 |
| Maximum | 56.598 | 82.289 |


Table 4.6: Statistical results for the algorithm on 20 ultrasound images of Human Livers.

widthwise slices of the liver, and the differing sizes of the subjects. Again, the parameters of the algorithm were set to provide large initial contours, which worked well for cases where the liver took up a large portion of the image image area, such as with the case in Figure 4.10. However in certain cases, the widthwise slice of a smaller patient led to a much smaller portion of the image being represented by the liver. Figure 4.9 shows such a case. Note that the initial elliptical contour is too large for the region of interest due to the parameter settings. Also, the settings for the force-based segmentation led to that initial contour being expanded, resulting in a much poorer segmentation. Once again, the variability of the data set could not be covered by a single parameter setting.

While there have been many efforts to detect and classify diseased liver tissues, particularly the detection of masses [14, 33], little has been done in detecting the liver itself. The work of Kotropoulos and Pitas was used to segment the liver, however no quantitative results were given [32]. It appears that few formal attempts have been made at segmenting the liver as a whole. This result is not surprising as most segmentation approaches involving the liver rely on texture cues, particularly the work of Chen *et al.* [14] (also see references in [38]). The variability of texture cues within the liver images from ill subjects is known to be quite large. Determining the boundary of the liver while accounting for this variability would be a difficult task. As a result, even the simplest of comparisons could not be obtained for this test.



(a) Contour Distance Distribution



(b) Overlap Score Distribution

Figure 4.8: The Distribution of Scores for the Liver Ultrasound Image Set.



(a) Original Image with Manually Traced Expert Contour in Green.




(b) The Detected Boundary Points (in white) and their relation to the Most Likely Boundary Depths (in red).




(c) The Fitted Ellipse Initial Contour.



(d) The Boundary Points Filtered by Depth and Ellipse Location.




(e) The Final Countour - The Segmentation Result.




(f) The Manually Traced Region Overlayed on the Detected Segment.


Figure 4.9: An Example of a Poor Result seen from using the Segmentation System on an Ultrasound Image of a human Liver (Contour Distance: 51.31 pixels, Overlap Score: 39.64%).



(a) Original Image with Manually Traced Expert Contour in Green.




(b) The Detected Boundary Points (in white) and their relation to the Most Likely Boundary Depths (in red).



(c) The Fitted Ellipse Initial Contour.



(d) The Boundary Points Filtered by Depth and Ellipse Location.



(e) The Final Countour - The Segmentation Result.



(f) The Manually Traced Region Overlaid on the Detected Segment.

Figure 4.10: An Example of a Good Result seen from using the Segmentation System on an Ultrasound Image of a human Liver (Contour Distance: 13.88 pixels, Overlap Score: 82.29%).

4.3.3 Parameter Settings and Their Effect

Note that the majority of the variability seen in the medical data sets are due to the presence of images taken from multiple probe positions, scanners, and scanner settings. This level of variation in image content is not often seen in test sets in this field. The lack of variation in most test sets has often allowed other algorithms to succeed through the use of fixed parameter values. This situation is clearly not the case for the medical data sets presented herein.

It is clear that for the last two image sets, a single parameter setting simply could not account for the variability of the image content. This does not discount the algorithm as a whole. Instead, it is believed that by somehow tuning the parameters to the image content, better results could be obtained. While determining a proper method to tuning these parameters is beyond the scope of this thesis, it is still important to show that the algorithm can perform well given appropriate parameter settings.

In order to show that my proposed algorithm is indeed well-founded, I take my test cases which saw poor results earlier in Figures 4.6 and 4.9 and determine whether alternative parameter settings would lead to improved performance. If good results could be obtained by tuning the parameters to the images, it would show that the proposed segmentation framework is not only capable of handling the variability seen in medical radiology in practice, but that the system is conceptually appropriate.

Figure 4.11 compares the poor result obtained earlier on an ultrasound image of a kidney with a segmentation result obtained from the same algorithm with better tuned parameter settings. The tuned parameter settings were selected manually based on the analysis of the boundary point locations. Given more appropriate parameter settings, we see a vast improvement. The overlap score improves by almost 57% while contour distance decreases by over 42 pixels on average. Similar results can be seen for the liver test case in Figure 4.12. In this case, overlap score improves by over 37% and contour distance decreases on average by over 37 pixels.


Clearly, given appropriate parameter settings, this approach can produce good




(a) The initial contour (in red) and filtered boundary points (in white) using class parameters.




(b) The initial contour (in red) and filtered boundary points (in white) using tuned parameters.




(c) The Final Contour using Class Parameters



(d) The Final Contour using tuned parameters



(e) Contour Distance: 55.15 pixels, Overlap Score: 18.68%



(f) Contour Distance: 12.60 pixels, Overlap Score: 75.45%


Figure 4.11: Comparison of two kidney segmentation results, one using a class-based parameter setting and one using an image-tuned parameter setting.




(a) The initial contour (in red) and filtered boundary points (in white) using class parameters.



(b) The initial contour (in red) and filtered boundary points (in white) using tuned parameters.



(c) The Final Contour using Class Parameters



(d) The Final Contour using tuned parameters



(e) Contour Distance: 51.31 pixels, Overlap Score: 39.64%



(f) Contour Distance: 14.20 pixels, Overlap Score: 77.14%

Figure 4.12: Comparison of two liver segmentation results, one using a class-based parameter setting and one using an image-tuned parameter setting.

results. However, determining appropriate parameter settings given the image content is left as future work.

4.4 Conclusions

The proposed ultrasound image segmentation system was implemented and test on three different data sets. A set of 304 ultrasound images of pork loins was compiled with a similar degree of difficulty as test sets common to the field of ultrasound image segmentation. Two additional sets of medical ultrasound images, one of kidneys and one of livers, were also used to measure the quality of the algorithm. These two medical data sets contained a much larger degree of variability than what is normally seen in the field.

On average, the segmentation system showed an accuracy of 73% on the pork loin data set. While direct numerical comparisons are a dubious way of comparing comparing algorithms for various reasons, I do note that these results are better than many approaches presented in similar studies in agricultural ultrasound image segmentation.

Results on the medical images were significantly lower with the segmentation system obtaining an accuracy of 49% on the kidney image set and 60% on the liver image. It became clear that a single set of parameter settings could not encompass the entire variability in image content of these data sets. However, it is possible that by somehow tuning the parameters to the image content, I can obtain significantly improved results.

Chapter 5

Algorithmic Comparison

In order to provide further context for the results of the algorithm proposed herein, I present an earlier version of this segmentation system for the purposes of comparison. In doing so, we are able to directly compare numerical results I've obtained with those gathered from an algorithm that has been accepted in the field [8, 7]. Furthermore, by comparing these two segmentation approaches, we can gain an understanding of the benefits the final algorithm provides.

With the final segmentation system already presented, discussion will focus on the differences between it and this earlier approach. Results will be shown for the same data sets used earlier for testing purposes. As both algorithms were designed with the same goal in mind - to segment general ultrasound images void of limiting priors - the comparison in quality of these two algorithms has an appropriate context.

5.1 Algorithm Description

In order to best accentuate the differences between the proposed algorithm and its earlier version, it is best to separate the system into three separate components: the boundary point detection scheme, the contour initialization, and the force-based segmentation.

Boundary Point Detection

Recall that in the finalized version of the segmentation system, boundary points were detected from the Synthetic A-Mode ultrasound scans. In particular, the modes of the Gumbel distributions used to model the tissue boundary echo patterns in the scans were the boundary points used in the final approach.

The earlier segmentation system used a much simpler criteria for determining the presence of likely boundary points. Knowing that brighter pixels relate to stronger echoes, which in turn relate to areas of significant changes in density, a basic threshold on pixel intensity was used to classify likely boundary points. Any pixels over an intensity threshold t_{bright} were considered likely boundary points.

Also, unlike the final segmentation system which later filters out boundary points based on their location to the initial contour, no filtering of these boundary points is performed.

Contour Initialization

As with the final segmentation system, the initial contour for the region of interest is an ellipse that is generated by being fitted to a collection of boundary points. The final segmentation system uses a RANSAC method to create multiple candidate ellipses which are then scored to determine the best quality ellipse to use as an initial contour.

My earlier approach was again more simplistic. Instead of scoring candidate ellipses, one single ellipse is fitted to all the points in the likely boundary point set using the same algorithm described in Section 3.3 and taken from [18]. In using all the likely boundary points for the ellipse fitting, an implicit assumption is made that any noise in the boundary point set would have equally been generated within and outside the region of interest. If this assumption is true, then the effect on the ellipse of the noise in the boundary point set would be small enough for the ellipse to remain a reasonable initial contour.




Figure 5.1: The Region-Based Pixel Force and its affect on the Contour

Force-Based Segmentation

Unlike the pre-processing steps which have seen significant changes, the force-based segmentation framework has remained relatively the same. All internal and external forces that existed in the final version of the segmentation system existed in this earlier form as well. What is different however is the inclusion of an additional external force in the earlier segmentation system: a region based pixel force.

Earlier, it was noted that darker pixels are more likely to be within the region of interest as the ultrasound pulses imaging the interest region would be weaker than the pulses in the surrounding area. Therefore, we can recognize pixels that are likely to be within the region of interest based on their intensity. Once again, thresholding is used to detect likely region pixels. Any pixels below a threshold t_{dark} are classified as likely region pixels.

As with the likely boundary points detected earlier, these likely region pixels also apply a force to the contour. However, unlike the boundary pixel force which works in a global fashion, the region pixel force is applied locally to the contour. Essentially, the goal is to move the contour to absorb, in a way, the detected region pixels. This goal is accomplished using the following force formulation:

$$F_{reg}(\mathbf{p}) = \sum_{\forall \mathbf{v} \in DP} h(\mathbf{v}, \mathbf{p}) \quad (5.1)$$

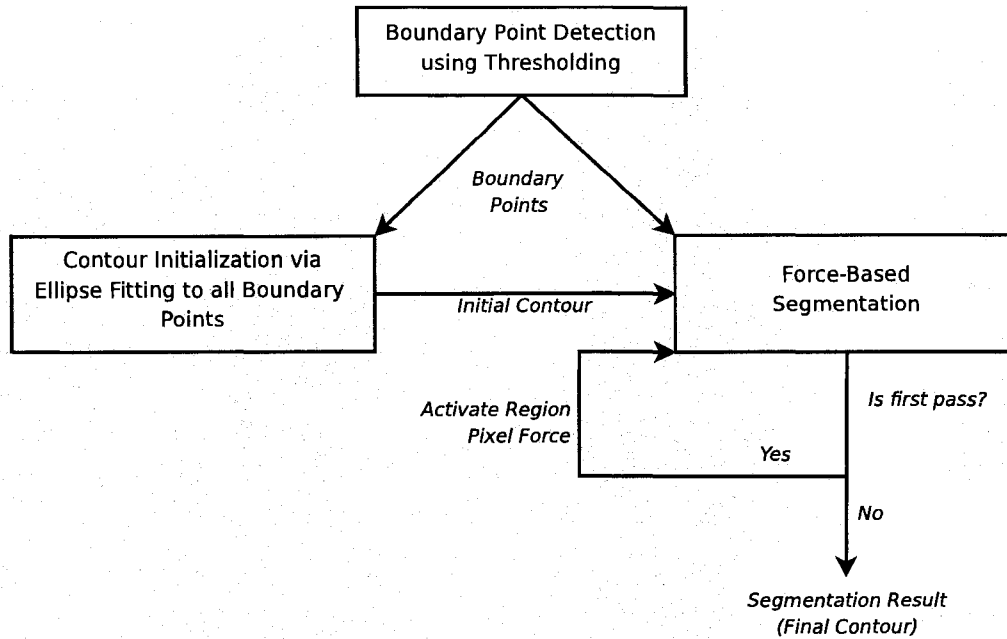


Figure 5.2: Overview of the Earlier Version of the Ultrasound Image Segmentation System

where

$$h(\mathbf{v}, \mathbf{p}) = \begin{cases} \mathbf{v} - \mathbf{p} & \text{if } \mathbf{v} \in 8\text{-Neighbourhood}(\mathbf{p}); \\ 0 & \text{otherwise.} \end{cases} \quad (5.2)$$

Effectively, the region pixel force only comes into play when there are likely region pixels along the outside edge of the contour. In this situation, the likely region pixels pull the contour in their direction, thereby including themselves within the closed contour and the region of interest. This process is shown visually in Figure 5.1.

The region pixel force works in parallel to the boundary pixel forces and the internal curvature and propagation forces described in Section 3.4 to obtain the final segmentation result. In order to be robust to poor initial contours, the region pixel force is not applied to the contour until the halting criteria is met. At that moment, the region force is added to the segmentation framework and the segmentation continues until the halting criteria is met a second time.

An overview of the earlier version of the segmentation system can be seen in Figure 5.2. Note that once again, this segmentation system also relies solely on

| Assumption | Use In Algorithm |
|--|--|
| 1. Bright pixels are likely to represent tissue boundaries. | Likely boundary pixels are detected based on their intensity. |
| 2. Pixels within a region of interest are likely to be darker than surrounding pixels at the same depth. | Darker pixels are detected as likely region pixels and apply an inclusion force to the contour. |
| 3. Noise is more likely at deeper parts of the image than in more shallow portions of the image. | The segmentation forces applied by the boundary points are weighted inversely by their depth in the image. |
| 4. Horizontal tissue boundaries will appear brighter than vertical tissue boundaries. | The segmentation forces applied by the boundary points have separate weights for the horizontal and vertical directions. |
| 8. The region of interest has a smooth, closed contour. | An ellipse is used to initialize the contour of the region of interest. |

Table 5.1: Ultrasound Image Assumptions and how they are incorporated into The Earlier Version of the Segmentation System.

prior knowledge obtained from the assumptions listed in Section 2.6. How these assumptions affect the working of this early version of the ultrasound segmentation system is summarized in Table 5.1.

5.2 Experimental Results

As mentioned, the pork loin, kidney, and liver data sets are once again used to measure the quality of the earlier ultrasound image segmentation system. Overlap scores and contour distances are once again the measures used to determine the quality of each segmentation result. Parameter settings for each data set are presented in Table 5.2. Intensity thresholds are given in percentage of pixels in the image in order to handle images of varying contrasts.

Results for the earlier segmentation system are summarized in Tables 5.3 and 5.4, with the range of the results displayed in Figure 5.3. Note that the results for the pork loin ultrasound image data set is very similar to those obtained by the final segmentation system. This is not surprising as there is a strong relationship in this image set between the boundary pixels selected by intensity and the boundary pix-

| Segmentation Parameters | | Pork Loins | Kidneys | Livers |
|--------------------------|---------------|------------|---------|--------|
| Boundary Pixel Threshold | t_{bright} | 13% | 2.5% | 7% |
| Region Pixel Threshold | t_{dark} | 26% | 20% | 10% |
| Boundary Pixel Weight | $w_{inS}(x)$ | 0.00 | 0.00 | 0.10 |
| (Inside Contour) | $w_{inS}(y)$ | 1.00 | 1.00 | 0.30 |
| Boundary Pixel Weight | $w_{outS}(x)$ | 0.18 | 0.30 | 0.50 |
| (Outside Contour) | $w_{outS}(y)$ | 0.04 | 0.04 | 0.30 |
| Region Pixel Weight | w_{dark} | 2.00 | 1.00 | 0.80 |
| Curvature Weight | w_{curv} | 0.08 | 0.10 | 0.10 |
| Propagation Weight | w_{prop} | 0.20 | 0.20 | 0.20 |
| Contour Envelope Width | ρ | 20 | 100 | 25 |

Table 5.2: Parameter Settings used for the Segmentation of each Data Set with the Earlier Segmentation System.

els selected from the modeled Synthetic A-Mode scans. Furthermore, the pork loin region is the only noticeable anatomical structure present in this image set. These two phenomenon can be seen from the example result in Figure 5.4.

However, what is also noticeable in Figure 5.4 is the effect of noise in the detected boundary point set. Despite strong internal forces attempting to keep the contour smooth, large pockets of erroneously selected boundary pixels in the center of, below, and to the left of the pork loin result in a rather jagged final contour. Although the parameters used were chosen to give the best numerical results, the contours obtained by this earlier segmentation system are not as visually appealing as those obtained by my final segmentation system.

The intensity threshold approach to determining likely boundary points also had difficulty recognizing particular boundaries. Consider the results for the liver image in Figure 5.5. In this case, the bottom-left liver boundary is not significantly brighter than its surroundings. Other than a slight decrease in intensity, it is hardly noticeable. As a result, no boundary points are detected along this tissue border, leading to a poor final segmentation. While well founded, selecting boundary points solely on intensity leads to noisy or incomplete boundary point sets in practice. As a result, the final segmentation approach uses an alternative boundary detection approach.

Another concern with the earlier segmentation system is the initial contour. As

| Stat. Measure | Pork Loins | Kidneys | Livers |
|----------------------|-------------------|----------------|---------------|
| Mean | 13.348 | 28.407 | 32.275 |
| Std. Deviation | 5.129 | 10.720 | 20.643 |
| Minimum | 3.814 | 10.791 | 10.004 |
| Maximum | 35.784 | 51.051 | 81.305 |

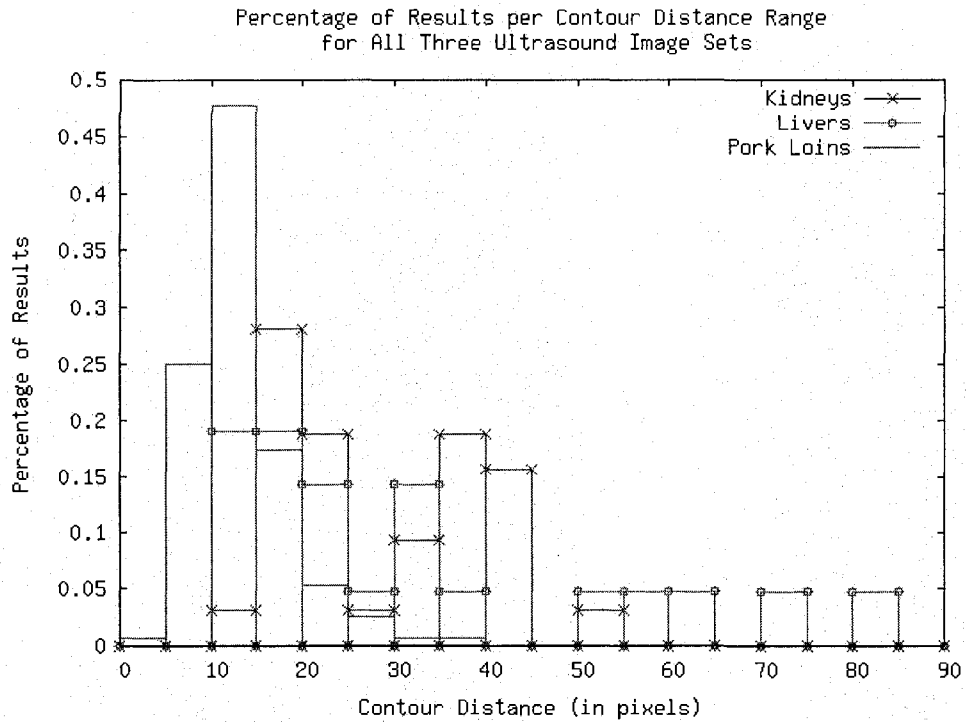
Table 5.3: Contour Distance Results for our Early Segmentation Algorithm on all three Data Sets.

| Stat. Measure | Pork Loins | Kidneys | Livers |
|----------------------|-------------------|----------------|---------------|
| Mean | 72.553 | 47.090 | 64.592 |
| Std. Deviation | 10.465 | 19.652 | 15.924 |
| Minimum | 8.244 | 7.030 | 38.479 |
| Maximum | 91.765 | 82.211 | 89.392 |


Table 5.4: Overlap Score Results for our Early Segmentation Algorithm on all three Data Sets.

just mentioned, the detected boundary point sets often have a significant amount of noise. It was assumed that any noise in the detected boundary points would be distributed in such a way as to have little impact in the initialization process. While this assumption held well for the pork loin images, the erroneous boundary points had a significant impact on the other two data sets, particularly the kidney ultrasound images. Figure 5.6 shows an example of a poor contour initialization on a kidney image. In this case, the skin and the liver boundaries added erroneous boundary points above and to the left of the kidney, leading to an unbalanced effect on the ellipse fitting. The end result is a significantly larger contour than the one desired. Clearly, not all boundary points should be used to initialize the contour. Therefore, the contour initialization method in the final segmentation system was altered to be more robust to outliers in the boundary point set.

Even though this earlier segmentation system has its faults, it is able to generate reasonable, and even good, results in many cases. Figure 5.7 shows one of these good segmentation results on an ultrasound image of a human liver. However, the limitations in the boundary point detection and contour initialization show a lack of flexibility in the overall method that needed to be addressed. While poor results obtained by my final segmentation system can be avoided by using tuned parameter




(a) Contour Distance Distribution




(b) Overlap Score Distribution


Figure 5.3: The Distribution of Scores for the All Three Ultrasound Image Sets.




(a) Original Image with Manually Traced Expert Contour in Green.



(b) The Detected Boundary Points (in white) and Region Points (in black).



(c) The Fitted Ellipse Initial Contour.




(d) The Final Countour - The Segmentation Result.




(e) The Manually Traced Region Overlaid on the Detected Segment.


Figure 5.4: An Example Result from using the Earlier Segmentation System on an Ultrasound Image of a pork loin (Contour Distance: 7.84 pixels, Overlap Score: 80.70%).




(a) Original Image with Manually Traced Expert Contour in Green.



(b) The Detected Boundary Points (in white) and Region Points (in black).



(c) The Fitted Ellipse Initial Contour.




(d) The Final Countour - The Segmentation Result.




(e) The Manually Traced Region Overlayed on the Detected Segment.

Figure 5.5: An Example of a Poor Result from using the Earlier Segmentation System on an Ultrasound Image of a Human Liver (Contour Distance: 81.30 pixels, Overlap Score: 38.48%).




(a) Original Image with Manually Traced Expert Contour in Green.




(b) The Detected Boundary Points (in white) and Region Points (in black).



(c) The Fitted Ellipse Initial Contour.



(d) The Final Countour - The Segmentation Result.




(e) The Manually Traced Region Overlayed on the Detected Segment.


Figure 5.6: An Example of a poor Initialization Result from using the Earlier Segmentation System on an Ultrasound Image of a human kidney (Contour Distance: 16.60 pixels, Overlap Score: 25.82%).




(a) Original Image with Manually Traced Expert Contour in Green.



(b) The Detected Boundary Points (in white) and Region Points (in black).



(c) The Fitted Ellipse Initial Contour.



(d) The Final Countour - The Segmentation Result.



(e) The Manually Traced Region Overlaid on the Detected Segment.

Figure 5.7: An Good Example Result from using the Earlier Segmentation System on an Ultrasound Image of a human liver (Contour Distance: 15.26 pixels, Overlap Score: 82.53%).

settings, here the poorer results are due to fundamental problems in the framework itself.

5.3 Conclusions

To provide further context to the results from my proposed segmentation system, an earlier version of my segmentation algorithm, which has generated reported results [7, 8], was used as a comparison. As both algorithms were made for exactly the same purpose, a reasonable comparison could be done with a method that already existed in the field.

The earlier segmentation approach differed from the final segmentation system in three ways:

- Boundary points were detected by intensity thresholds instead of from Synthetic A-Mode ultrasound scans.
- Contour initialization was done by fitting an ellipse to all detected boundary points instead of a subset of the detected boundary points.
- A region-based pixel force was added to the segmentation framework. Likely region pixels were detected as those pixels below a given threshold and these pixels push the contour to encapsulate them.

Results were obtained over the same data sets and using the same quality measures for both segmentation systems. While similar results were obtained with both algorithms, the earlier system has some fundamental flaws. The boundary point detection scheme turned out to incorporate a significant amount of noise while missing more subtle boundaries. Meanwhile, the contour initialization procedure was less susceptible to the noise in the boundary point set. Despite the best chosen parameter settings, these flaws could not be overcome. Both of these issues were addressed in the final segmentation framework, which given appropriately tuned parameters, can generate better results and handle more variability.

Ultimately, the final segmentation system proposed herein compares favorably to an existing algorithm in the field.

Chapter 6

Conclusion

6.1 Summary

As radiologists become more and more in demand, the desire for automated ultrasound image segmentation grows. Given the poor quality of ultrasound images, any and all image cues and priors have been relied upon to obtain reasonable results. Often, relying on these features has limited the applicability of certain proposed algorithms.

In this thesis, I concern myself directly with the question of what type of knowledge is appropriate for the task of ultrasound image segmentation. Unlike other approaches which treat ultrasound images in a similar fashion to real world images, I acknowledge instead the uniqueness of the medium itself. The ultrasound imaging process is discussed in detail and in doing so, I am able to infer various assumptions about the content of ultrasound images.

From these assumptions, I present an alternative approach to ultrasound image segmentation. A segmentation framework based solely on these obtained assumptions is presented. In this approach, the original ultrasound image is decomposed into separate one-dimensional signals, dubbed as Synthetic A-Mode Ultrasound Scans, in the hope of mimicking the originally obtained ultrasonic data as closely as possible. These Synthetic A-Mode scans are then analyzed as echo patterns. Likely tissue boundary echo patterns are detected and fitted Gumbel distributions are used to obtain likely tissue boundary locations. Using these obtained tissue boundary locations, an initial contour is created and segmentation is performed by

applying robust forces to this contour.

My segmentation system was implemented and tested on three separate data sets, one data set being similar to those in the field, and two data sets of a much higher degree of difficulty. Results of 73% accuracy on the pork loin data set seem to suggest that the segmentation system is as good or better than other algorithms in the area while not being as limited by the priors it uses. Further results on two medical image data sets also show good results assuming appropriate parameter settings are used.

Finally, a direct comparison was performed between this proposed segmentation system and an earlier version of my system that currently exists in the field. Numerical results were similar between both methods with the final segmentation system being more capable of handling noise and variability.

The combination of these results show that there is value in approaching ultrasound images as being a unique case in image processing and that understanding the ultrasonic imaging process will be important in obtaining feasible and reliable ultrasound image segmentation systems.

6.2 Future Work

There are various directions for future work based on what is presented herein. Both the algorithm and approach are new and can be improved upon. In terms of the algorithm itself, a key component moving forward is, as mentioned earlier, linking the parameter settings to features in the image, essentially automating the setting of various parameters. For example, by looking at the standard distributions of the most likely depths, we could determine how spread out the detected boundary points are in the image. If the boundary points are spread out (*i.e.*, the standard distributions are high), then it may be less likely for the best elliptical contour to have strong boundary support. Therefore, it may be reasonable to tie the boundary support weight to the standard distributions of the likely boundary depths.

On top of automating parameter settings, a more elegant way of determining the number of Gumbel distributions to fit to a Synthetic A-Mode scan would be useful.

A probabilistic approach similar to the infinite Gaussian mixture model by Rasmussen comes to mind [39]. Extending this work to 3D ultrasound data sets is also a possibility, as is reformulating the segmentation using a level-set approach similar to that of Caselles *et al.* to detect multiple regions of interest [13]. Improving the efficiency of the ellipse fitting procedure would also help to make the algorithm more usable. One idea of particular interest would be to incorporate texture and intensity information into the boundary detection process, thereby creating a boundary map similar to the saliency map of Itti and Koch for real world images [28]. This boundary map could ultimately improve the robustness of this algorithm even further.

As far as using knowledge of the image acquisition process for ultrasound image segmentation, there are many ways in which to move forward. Recognizing more imaging artifacts and determining better echo models for tissue boundaries are both open problems. By decomposing ultrasound images into Synthetic A-Mode scans, we have opened the problem of ultrasound image segmentation to the field of one-dimensional digital signal processing, a domain which includes a myriad of analysis techniques. Further study of the ultrasound imaging process could also lead to further recognition of important patterns within Synthetic A-Mode scans that could lead to a better overall understanding of image content.

In the end, by recognizing the uniqueness of ultrasound as a medium, we open up the door to a new way of assessing ultrasound signal content. Possibilities abound as to what can be useful in this approach to the problem of ultrasound image segmentation.

Bibliography

- [1] P. Abolmaesumi and M. R. Sirouspour. An interacting multiple model probabilistic data association filter for cavity boundary extraction from ultrasound images. *IEEE Transactions on Medical Imaging*, 23(6):772–784, June 2004.
- [2] Miguel Alemán-Flores, Patricia Alemán-Flores, Luis Álvarez-León, José Manuel Santana-Montesdeoca, Rafael Fuentes-Pavón, and Agustín Trujillo-Pino. Computer-aided measurement of solid breast tumor features on ultrasound images. In *Proceedings of Computer Vision and Mathematical Methods in Medical and Biomedical Image Analysis*, volume 3117, pages 353–364, October 2004.
- [3] Elsa Angelini, Jeffrey Holmes, Andrew Laine, and Shunichi Homma. Segmentation of r3d ultrasound with implicit deformable models without gradients. In *Proceedings of International Conference on Image and Signal Processing and Analysis*, volume 2, pages 711–716, September 2003.
- [4] Jonas August and Takeo Kanade. Weakly-supervised segmentation of non-gaussian images via histogram adaptation. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention*, pages 992–993, November 2003.
- [5] Ayoub B. Ayoub. The eccentricity of a conic section. *The College Mathematics Journal*, 34(2):116–121, March 2003.
- [6] Brian Booth and Xiaobo Li. Boundary point detection for ultrasound image segmentation using gumbel distributions. In *Proceedings of Signal Processing and Multimedia Applications (SIGMAP)*, pages 179–183, July 2007.
- [7] Brian Booth, Ryan Neighbour, and Xiaobo Li. On agricultural ultrasound image segmentation. In *Proceedings of IEEE International Conference on Signal Processing*, volume 2, pages 915–920, November 2006.
- [8] Brian Booth, Vimal Patel, Edmond Lou, Lawrence Le, and Xiaobo Li. Towards medical ultrasound image segmentation with limited prior knowledge. In *Proceedings of IEEE Digital Signal Processing Workshop*, pages 488–493, September 2006.
- [9] Johan G. Bosch, Steven C. Mitchell, Boudewijn P. F. Lelieveldt, Francisca Nijland, Otto Kamp, Milan Sonka, and Johan H. C. Reiber. Automatic segmentation of echocardiographic sequences by active appearance motion models. *IEEE Transactions on Medical Imaging*, 21(11):1374–1383, November 2002.
- [10] Dennis Burson and Eric Burg. Procedures for estimating pork carcass composition. National Pork Producers Council, 2001.

- [11] Pablo Cancela, Fernando Reyes, Pablo Rodríguez, Gregory Randall, and Alicia Fernández. Automatic object detection using shape information in ultrasound images. In *Proceedings of IEEE International Conference on Image Processing*, volume 3, pages 417–420, September 2003.
- [12] Marie-Hélène Roy Cardinal, Jean Meunier, Gilles Soulez, Roch L. Maurice, Éric Therasse, and Guy Cloutier. Intravascular ultrasound image segmentation: A three-dimensional fast-marching method based on gray level distributions. *IEEE Transactions on Medical Imaging*, 25(5):590–601, May 2006.
- [13] Vicent Caselles, Ron Kimmel, and Guillermo Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22(1):61–79, 1997.
- [14] Chung-Ming Chen, Henry Horng-Shing Lu, and Yu-Chen Lin. An early vision-based snake model for ultrasound image segmentation. *Ultrasound in Medicine and Biology*, 26(2):273–285, 2000.
- [15] Yunmei Chen, Feng Huang, Hemant D. Tagare, Murali Rao, David Wilson, and Edward A. Geiser. Using prior shape and intensity profile in medical image segmentation. In *Proceedings of IEEE International Conference on Computer Vision*, pages 1117–1124, 2003.
- [16] Franck Davignon, Olivier Basset, Gérard Gimenez, and Wilfried Mai. A multi-parameteric segmentation tool dedicated to ultrasonic data. In *Proceedings of IEEE Ultrasonics Symposium*, volume 2, pages 1761–1764, October 2002.
- [17] Arthur Dempster, Nan Laird, and Donald Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society*, 39(1):1–38, 1977.
- [18] Andrew Fitzgibbon, Maurizio Pilu, and Robert B. Fisher. Direct least squares fitting of ellipses. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(5):476–480, May 1999.
- [19] C. A. Glasbey. Ultrasound image segmentation using stochastic templates. *Journal of Computing and Information Technology*, pages 107–116, 1998.
- [20] Ali K. Hamou and Mahmoud R. El-Sakka. A novel segmentation technique for carotid ultrasound images. In *Proceedings of ICASSP*, pages 521–524, 2004.
- [21] Xiaohui Hao, Charles Bruce, Cristina Pislaru, and James F. Greenleaf. A novel region growing method for segmenting ultrasound images. In *Proceedings of IEEE Ultrasonics Symposium*, volume 2, pages 1717–1720, October 2000.
- [22] Xiaohui Hao, Charles Bruce, Cristina Pislaru, and James F. Greenleaf. Segmenting high-frequency intracardiac ultrasound images of myocardium into infarcted, ischemic, and normal regions. *IEEE Transactions on Medical Imaging*, 20(12):1373–1383, December 2001.
- [23] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in computer vision*. Cambridge University Press, 2 edition, 2003.
- [24] Ping He and Jun Zhang. Segmentation of tibia bone in ultrasound images using active shape models. In *Proceedings of IEEE Engineering in Medicine and Biology Conference*, volume 3, pages 2712–2715, October 2001.

- [25] Nuallsawat Hiransakolwong, Kien A. Hua, Khanh Vu, and Piotr S. Windyga. Segmentation of ultrasound liver images: An automatic approach. In *Proceedings of International Conference on Multimedia and Expo*, volume 1, pages 573–576, July 2003.
- [26] David R. Holmes III and Richard A. Robb. Real-time segmentation of trans-urethral ultrasound images for prostate brachytherapy. In *Proceedings of Medical Image Computing and Computer-Assisted Intervention*, pages 983–984, November 2003.
- [27] M. F. Insana, R. F. Wagner, D. G. Brown, and B. S. Garra. On the information content of diagnostic ultrasound. In *Proceedings of Information Processing in Medical Imaging*, pages 437–455, 1987.
- [28] Laurent Itti and Christof Koch. A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Research*, 40:1489–1506, 2000.
- [29] Sandra V. B. Jardim and Mario A. T. Figueiredo. Automatic contour estimation in fetal ultrasound images. In *Proceedings of IEEE International Conference on Image Processing*, pages 1065–1068, 2003.
- [30] Yifeng Jiang, Zhijun Zhang, Feng Cen, Hung Tat Tsui, and Tze Kin Lau. An enhanced appearance model for ultrasound image segmentation. In *Proceedings of International Conference on Pattern Recognition*, volume 3, pages 802–805, 2004.
- [31] Michael Kass, Andrew Witkin, and Demetri Terzopoulos. Snakes: Active contour models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 259–268, 1987.
- [32] Constantine Kotropoulos and Ioannis Pitas. Segmentation of ultrasonic images using support vector machines. *Pattern Recognition Letters*, 24(4-5):715–727, February 2003.
- [33] Wen li Lee, Yung-Chang Chen, Ying-Cheng Chen, and Kai-Sheng Hsieh. Un-supervised segmentation of ultrasonic liver images by multiresolution fractal feature vector. *Information Sciences*, 175(3):177–199, October 2005.
- [34] Marcos Martín-Fernández and Carlos Alberola-López. An approach for contour detection of human kidneys from ultrasound images using markov random fields and active contours. *Medical Image Analysis*, 9:1–23, 2005.
- [35] William D. Middleton, Alfred B. Kurtz, and Barbara S. Hertzberg. *Ultrasound, The Requisites*. Mosby, Inc., 2 edition, 2004.
- [36] William D. Middleton, Alfred B. Kurtz, and Barbara S. Hertzberg. *Ultrasound, The Requisites*, chapter Practical Physics, pages 3–26. Mosby, Inc., 2 edition, 2004.
- [37] Russell E. Muzzolini. *A Volumetric Approach to Segmentation and Texture Characterisation of Ultrasound Images*. PhD thesis, University of Saskatchewan, 1996.
- [38] J. Alison Noble and Djamel Boukerroui. Ultrasound image segmentation: A survey. *IEEE Transactions on Medical Imaging*, 25(8):987–1010, August 2006.

- [39] Carl Edward Rasmussen. The infinite gaussian mixture model. *Advances in Neural Information Processing Systems*, MIT Press, 12:554–560, 2000.
- [40] Daniel W. Rickey. Medical ultrasound linear array probe/ scan head/ transducer. <http://en.wikipedia.org/wiki/Image:UltrasoundProbe2006a.jpg>, February 2006.
- [41] K. J. W. Taylor. *Medical Imaging Techniques: A Comparison*, chapter Basic Principles of Diagnostic Ultrasound, pages 231–259. Plenum Press, 1979.
- [42] A.W.K. Tong, R. Qureshi, X. Li, and A.P. Sather. A system for ultrasound image segmentation for loin eye measurements in swine. *Canadian Agricultural Engineering*, 40(1):47–53, 1998.
- [43] Eric W. Weisstein. Gumbel distribution. From MathWorld—A Wolfram Web Resource. <http://mathworld.wolfram.com/GumbelDistribution.html>, March 2006.
- [44] Jun Xie, Yifeng Jiang, and Hung tat Tsui. Segmentation of kidney from ultrasound images based on texture and shape priors. *IEEE Transactions on Medical Imaging*, 24(1):45–57, January 2005.
- [45] R. Youmaran, P. Dicorato, R. Munger, T. Hall, and A. Adler. Automatic detection of features in ultrasound images of the eye. In *Proceedings of Instrumentation and Measurement Technology Conference*, pages 1829–1834, May 2005.
- [46] Hui Zhu, Yun Liang, and Morton H. Freidman. Ivus image segmentation based on contrast. In *Proceedings of SPIE Medical Imaging 2002: Image Processing*, volume 4684, pages 1727–1733, 2002.