IMPROVED ACTIVE SHAPE MODELS FOR SEGMENTATION OF THE PROSTATE ON MR IMAGERY

BY ROBERT JAMES TOTH

A thesis submitted to the

Graduate School—New Brunswick

Rutgers, The State University of New Jersey

in conjunction with

The Graduate School of Biomedical Sciences

the University of Medicine and Dentistry of New Jersey

in partial fulfillment of the requirements for the

Joint Degree of Master of Science

Graduate Program in Biomedical Engineering

Written under the direction of

Anant Madabhushi

and approved by

New Brunswick, New Jersey

Oct, 2010

ABSTRACT OF THE THESIS

Improved Active Shape Models for Segmentation of the Prostate on MR Imagery

by Robert James Toth

Thesis Director: Anant Madabhushi

Segmentation aims to determine which locations within an image contain the object of interest. Segmentation of the prostate boundary on clinical images is useful in a large number of applications including calculation of prostate volume pre- and post-treatment, detection of extra-capsular spread, and creation of patient-specific anatomical models. Manual segmentation of the prostate boundary is, however, time consuming and subject to inter- and intra-reader variability. T2-weighted (T2-w) Magnetic Resonance (MR) structural imaging (MRI) and MR Spectroscopy (MRS) have recently emerged as promising modalities for detection of prostate cancer *in vivo*. With the recent advance in prostate imagery, we have generated an accurate prostate segmentation system for MR imagery. Our system builds upon the popular Active Shape Model (ASM) framework, in which a statistical description of the shape is first generated, after which an appearance of the object of interest is modeled.

In our system, the shape model can be generated in either 2D or 3D, and is defined by a set of anatomical landmarks. For the appearance model, we offer several improvements. We generated statistical texture features of the prostate images, and use those features to overcome limitations of solely using intensities. In addition, we use intelligent feature selection algorithms including forward feature selection and adaboost to

determine which features to include in our segmentation system. The statistical appearance models are not modeled as a simple Gaussian distribution, but rather as a sum of Gaussians, resulting in more accurate models. In 2D, a local appearance model is generated for each landmark location on the prostate border. However, in 3D this is infeasible, so we generate a global appearance model describing the voxels within the prostate. The 2D ASM resulted in a Dice similarity coefficient (DSC) of 0.85, while our 3D system resulted in a DSC of 0.89 (over 56 and 37 studies respectively). This is comparable to other state of the art prostate MR segmentation schemes. Finally, we have shown that in the specific application of prostate volume estimation, our system performs more accurate volume estimations than currently employed clinical models. Our system achieved a correlation (R^2 value) with the ground truth volume of 0.82 while the clinical model achieved an R^2 value of 0.70. Our system had a volume fraction of 1.05 in comparison to the ground truth volume, while the clinical model achieved a volume fraction of 1.14. Overall, we have developed an efficient, accurate, and useful prostate segmentation scheme for MR imagery.

Acknowledgements

I would first and foremost like to acknowledge my advisor and mentor, Dr. Anant Madabhushi. He believed in me from the start, and has since pushed me to achieve more than I thought was possible. His relentless dedication to all members of his lab, from undergraduates to post-docs, has inspired me and showed me what true leadership means. On that note I would like to acknowledge my other lab members. There are too many to list every single colleague individually, but needless to say I have learned more from all of you than from any book or research paper, especially Satish Viswanath, Jon Chappelow, Pallavi Tiwari, and Shannon Agner. You all have inspired me to continue my research towards a PhD. I would also like to mention my collaborators, especially, but not limited to, Dr. Rosen and Dr. Bloch, without you this research would never have beeen possible. Finally, I would like to acknowledge my family for their support. They always have my best interests in mind and have showed me what it takes to succeed and to be a good person.

This work was made possible via grants from the Wallace H. Coulter Foundation, Aresty Foundation for Undergraduate Research, New Jersey Commission on Cancer Research, National Cancer Institute (Grant Nos. R01CA136535-01, ARRA-NCl-3 R21CA127186, R21CA127186, R03CA128081-01, and R03CA143991-01), and The Cancer Institute of New Jersey.

Dedication

This is dedicated to my sister, Leigha Marie Toth. It is said that you can choose your friends but not your family, and fortunately to me you're both.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Background and Motivation

Prostatic adenocarcinoma (CaP) is the second leading cause of cancer related deaths among men in the United States, with an estimated 192,280 new cases in 2009 (Source: *American Cancer Society*). The current standard for detection of CaP is transrectal ultrasound (TRUS) guided symmetrical needle biopsy, which can have a high false negative rate [1]. Recently, multi-modal Magnetic Resonance (MR) Imaging (MRI) comprising both structural T2-weighted (T2-w) MRI [2, 3] and MR Spectroscopy (MRS) [4, 5, 6, 7, 8, 9, 10, 11, 12] has emerged as promising multi-modal image data for early detection of CaP [5, 6].

Several computer-aided diagnosis (CAD) schemes have emerged in recent years for automated CaP detection from prostate T2-w MRI [2, 13, 14] and MRS [8, 15, 16]. MRS measures the relative concentrations of different biochemicals and metabolites in the prostate, and changes in relative concentrations of choline, creatine, and citrate are highly indicative of the presence of CaP. It's important to note that MRS acquisition has a lower resolution than MRI acquisition, and thus each MRS metavoxel (containing a spectral signal) is approximately 13 times the radius of an MRI voxel (containing a single intensity value). CAD systems for prostate cancer utilize advanced machine learning algorithms, using the data at hand (whether intensities or spectral signals) to determine which locations within the prostate have a high probability of being cancerous.

With the increasing use of structural, functional, and metabolic MR imaging of the prostate [13, 15, 17, 18, 19, 20], MR imaging is used to evaluate efficacy of CaP therapy [21, 22, 23]. This is typically done by estimating prostate volume, pre-and post-treatment for CaP on MRI. Prostate volume has been shown to be a strong predictor

of treatment outcome for patients with prostate cancer [24], especially when combined with a baseline prostate-specific antigen (PSA) level [25]. Prostate volume is important in determining PSA density [26] and is an important factor in diagnosing and managing both benign and malignant conditions of the prostate [27]. The most common method for estimating prostate volume is by manually determining the anterior-posterior, transverse, and rostral-caudal measurements of the prostate and estimating the prostate as a simple ellipsoid [22, 26, 27, 28, 29, 30, 31, 32, 33, 23, 34]. While TRUS is cheaper and easier to perform than endorectal *in vivo* prostate MRI, it lacks the high structural information (such as a clear view of the prostate border) associated with T2-weighted MRI (T2-w), and has a lower signal to noise ratio. Most prostate volume estimations are done using TRUS imagery, yet it has been shown that there is a strong correlation ($R^2 = 0.93$) between volume estimations obtained using TRUS, and volume estimations obtained from MR imagery [27]. In addition, the ellipsoidal model can yield accurate volume estimations for T2-w MR imagery of the prostate, even when an endorectal coil is used (as in our data) [33].

Segmentation of the prostate can be useful in both CAD systems and in volume estimations of the prostate. Segmentation is the process of defining where an object is in an image. In CAD systems, segmentation of the prostate is an important first step for determining a region of interest (ROI) on which to perform the algorithm. Once the prostate surface is delineated, the CAD algorithm can be implemented on just the part of the image in which the prostate is seen. In addition, prostate segmentation could be useful in using planimetry to determine prostate volume. Given a set of segmentations, prostate volume estimation is done by aggregating the areas associated with the gland segmentations on a series of 2D slices [27]. Overall, accurate prostate segmentation is useful both as a first step for CAD systems to aid in detecting prostate cancer, and for accurately determining prostate volume to evaluate therapies of prostate cancer. Manual segmentation of the prostate, however, is not only laborious, but is also subject to a high degree of inter-, and intra-observer variability [35].

1.2 Previous Work

1.2.1 Active Shape Models

The Active Shape Model (ASM) [36] and Active Appearance Model (AAM) [37] are both popular methods for segmenting known anatomical structures. Active Shape Models (ASM's) [36] provide a statistical framework for automatically segmenting objects with a known shape from images and are particularly relevant in the context of medical imagery. ASM's involve the use of a series of manually landmarked training images to generate a point distribution model. An expert segments the prostate on a series of M training images, and manually aligns a set of N landmarks around the border of the prostate. Once a set of anatomical landmark points have been placed, principal component analysis (PCA) is performed to evaluate known variations in the shape. PCA essentially calculates the eigenvectors of the entire set of landmark points, and retains the top few eigenvectors which contain most of the shape variation. For N landmark points in d dimensions, the prostate shape in an image is defined by an $N \cdot d$ vector of Cartesian coordinates. PCA can then be performed on all training images to develop a statistical shape model of the prostate shape.

The appearance model for ASM's is usually defined by calculating the direction normal to the shape at each landmark point. Then, κ intensity values are sampled along this normal on each training image. A κ-dimensional Gaussian is defined for each landmark point, thus defining the statistical appearance model. To locate the border on a new image, the Mahalanobis distance to this appearance model is minimized, and the location in the image with the minimum Mahalanobis distance is chosen as the border [36]. The Mahalanobis distance measures the distance of a vector to a Gaussian distribution. This is repeated for all landmark points, and the statistical shape model is optimally fit to the border.

1.2.2 3D Active Shape Models

While the original proposed framework was mainly used for 2D images, it is easily extendable to 3D. 3D ASM's have been used for a multitude of tasks, such as bone

segmentation [38], liver segmentation [39], and artery segmentation [40]. A thorough review of 3D ASM's from recent years can be found in [41]. In 3D, a large set of landmarks can be used to define a triangulated surface. The variations captured by performing PCA on these landmarks will contain variations in all 3 dimensions. For N landmarks, each containing an x, y, and z coordinate, a maximum of $3 \cdot N$ principle components can be calculated. However, most (98%) of the shape variations seen in the training set can usually be captured using only the first few principle components. One advantage to generating 3D statistical shape models over 2D models is that unique variations at different z locations can be captured. In addition, a second advantage to 3D models is that interslice variations can be captured, which 2D models would completely ignore. This is more important if the interslice spacing is small, but is still important even when the voxels are anisotropic. However, most 3D ASM's suffer from distinct limitations which are described below.

1. In the traditional ASM methodology, a unique appearance model is created at each landmark point. In 3D, however, there can exist hundreds or thousands of landmarks used to define the surface shape, and a unique statistical appearance model at each landmark can quickly become infeasible.

2. A given landmark point must correspond to a specific anatomical location, and any slight misalignment could result in a severe degradation of the quality of the local appearance models.

3. Since each local appearance model is created by sampling along the normal direction, slightly inaccurate triangulations could drastically change the calculated normal directions, which would result in inaccurate local appearance models.

4. If one were to use κ neighborhood voxels to define the local appearance model, a κ dimensional Gaussian would be created at each landmark point, requiring a large number of training volumes.

5. The 1st dimension of the Gaussian representing the local appearance model must represent the intensity value *exactly* κ voxels away from the landmark point.

However, neighboring objects locations may be different across training images, and theres nothing to suggest that the intensity value *exactly* κ voxels away on each different training image would provide any meaningful appearance model.

6. The distribution of intensity values may not be accurately described by a Gaussian, which using the Mahalanobis distance assumes.

7. The image intensity may not be the optimal representation of the object's texture.

8. The fact that a region is searched near each landmark point in segmenting a new image requires an accurate initial location of the shape model.

1.2.3 Extensions to Active Shape Models

While ASM's set the groundwork for a very efficient and accurate segmentation system, there are some inherent limitations. The first is the requirement for a proper initialization. If the system is initialized too far from the ground truth, the system won't be able to converge on the correct object border. Over the last few years, some researchers have been exploring schemes for accurate and reproducible initialization of ASM's [42, 43, 44]. Seghers et al. [45] presented a segmentation scheme where the entire image is searched for landmarks. They however concede that accurately initialized regions of interest (ROIs) would greatly improve their algorithm's efficiency and accuracy. Ginneken et al. [46] pointed out that without *a-priori* spatial knowledge of the ROI, very computationally expensive searches would be required for ASM initialization, contributing to a slow overall convergence time. Multi-resolution ASMs have also been proposed, wherein the model searches for the ROI in the entire scene at progressively higher image resolutions [44]. Brejl et al. [43] presented a shape-variant Hough transform to initialize an ASM, but the scheme can be very computationally expensive. Cosio [42] presented an ASM initialization method based on pixel classification which was applied to segmenting TRUS prostate imagery. The method employs a Bayesian classifier to discriminate between prostate and non-prostate pixels in ultrasound imagery. A trained prostate shape is then fit to the edge of the prostate, identified via

the Bayesian classifier. A Genetic Algorithm [47] is employed to minimize the distance between the trained prostate shape and the edge of the prostate.

In addition to ASM initialization, the use of the Mahalanobis distance leads to some limitations. First of all, the Mahalanobis distance assumes that the distribution of intensity values is Gaussian, which need not necessarily be the case. ASM's normally find the location with the minimum Mahalanobis distance, and assume that is close to the object border. But outliers, local minima, and intensity artifacts often prevent accurate segmentations. Secondly, image intensities might not necessarily be the optimal texture to use, as intensities are prone to noise and artifacts, which detracts from an accurate segmentation. In addition, with limited training data, the Mahalanobis distance will be undefined if too many pixels are sampled.

Several improvements to traditional Mahalanobis distance have been proposed. The first is the popular Active Appearance Model (AAM), which creates a global appearance model of the object, and combines that model with the shape information [37]. In addition, the AAM model was improved to be more robust to occlusion and outliers (called Robust AAM) [48], and the ASM model was also independently improved to be more robust to outliers [49]. A major improvement to the traditional ASM is ASM with Optimal Features (ASMOF) [46], which was shown to offer significant improvements. ASMOF steers clear entirely of using the Mahalanobis distance, and instead creates a classifier as to whether a pixel is considered inside or outside of the object. Then, whichever features best classify pixels are used in the segmentation algorithm. The significant improvements offered by this approach show the usefulness of using features other than just image intensities, although it is unclear whether the improvements come from the features or from using a classifier instead of the Mahalanobis distance. A second segmentation system which builds upon the traditional ASM scheme is Minimal Shape and Intensity Cost Path Segmentation (MISCP) [45]. While this system contains many differences to the traditional ASM, we focus on two major improvements. The first improvement is the idea of sampling a neighborhood around each landmark point instead of just pixels along a straight line. This allows more information to be gleaned from the training data. The second improvement is the combination of multiple

statistical texture features to improve segmentation, over simply using intensities.

1.3 Proposed Prostate Segmentation Scheme

With the recent advancements of prostate imaging, several prostate segmentation schemes have been developed [3, 42, 50, 51, 52, 53, 54, 55]. In the past 3 years, prostate segmentation schemes specifically for MR imagery of the prostate have been developed by Klein et al. [56], Martin et al. [57, 58], Pasquier et al. [59], and Makni et al. [60]. Klein et al. [56] performed a registration between an MR image of the prostate and an atlas of training data to achieve a segmentation of the prostate. Martin et al. [58] also used an atlas of training images, but constrained the segmentation model through the use of a statistical shape model. Pasquier et al. [59] used an Active Shape Model [36] method for extracting a statistical shape model of the prostate and searched for strong gradients to detect the prostate edge. Finally, Makni et al. [60] used a statistical shape model of the prostate, and clustered the intensities in a manually placed region of interest into 3 clusters: surrounding tissues and fat, central prostate zone, peripheral prostate zone. Any voxels within the latter 2 zones were determined to be part of the prostate. In these studies, the number of prostate volume studies tested range from 12 to 50 studies and there are varying degrees of manual intervention, ranging from completely automated to fully interactive initialization of the segmentation.

For our system, the shape model was first developed in 2D, but has been since extended to be fully 3D. The main contributions are regarding the appearance model, and the initialization of the segmentation system. For the appearance model, we have explored the use of statistical texture features over simply using image intensities [61, 62] and alternatives to using the Mahalanobis distance to detect the prostate border [61]. The ASM appearance model typically uses the intensities of the image to learn a statistical appearance model. However, relying solely on the intensity information may not be sufficient for accurately detecting the correct boundary, especially if different regions of the image, or different regions within the desired object, have similar intensity values. This is particularly true of MR imagery where strong bias field inhomogeneity artifacts can significantly obfuscate object boundaries [2]. We calculate the gray level statistics

of each image by convolving a set of kernels with the intensity image. These include the Kirsch [63] and Sobel [64] kernels to better quantify the edges of the prostate border, and Gaussian and mean kernels to intelligently incorporate neighboring intensity information. While traditional ASM's use neighboring intensity information, they are dependent on the normal to the shape at any given landmark point. By contrast, the Gaussian and mean kernels take neighboring information into account and yet do not depend on the normal of the shape. In conjunction with the texture features, we found that using the Cartesian x and y coordinates of landmark point as additional "features" greatly improved our results. This was inspired by the work of Cosio [42], who using the Cartesian coordinates as additional dimensions of his distribution. Finally, we discovered that the texture features of the prostate boundary are not always optimally modeled as a Gaussian, as the Mahalanobis distance assumes. Thus, in our system we define the distribution as the sum of multiple Gaussians, otherwise known as Gaussian Mixture Models (GMM's) [65] to define the distribution of the prostate border at each landmark. A multi-dimensional GMM comprised of different kernel responses is thus created for each landmark point. As a final step in our training, we performed experiments testing various algorithms for feature selection to determine which features offered the highest segmentation accuracy. In summary, our contributions to the ASM appearance model are 3-fold:

1. Use of statistical texture features to better identify the prostate.

2. Use of Gaussian Mixture Models instead of the Mahalanobis distance for generating more accurate appearance models.

3. Use of feature selection algorithms to determine which features to use for segmentation.

In addition, we present a fully automated ASM initialization scheme for segmentation of the prostate on multi-protocol *in vivo* MR imagery by exploiting the MR spectral data. Note that for the studies considered in this work, the MRS data was acquired as part of routine multi-protocol prostate MR imaging and not specifically for the purposes of this project. While the resolution of MR and MRS data are different,

the identification of prostate spectra by eliminating non-informative spectra outside the prostate provides an initial accurate ROI for the prostate ASM. We leverage the idea first introduced in [8, 15], in which spectral clustering was employed to distinguish between prostatic and extra-prostatic spectra. We achieve this through replicated k-means clustering of the MR spectra in the midgland. Replicated k-means clustering aims to overcome limitations associated with the traditional k-means algorithm (sensitivity to choice of initial cluster centers) by randomly initializing the cluster centers, and repeating the process until a global minima is found. For each slice, the largest cluster (identified as the non-informative cluster) is eliminated. The mean shape of the prostate is then transformed to fit inside the remaining spectra, which serves to provide the initial landmark points for a 2D ASM. In addition, since the spectral data severely degrades away from the midgland of the prostate, we perform clustering in the midgland. The resulting initialized shape is then rescaled to account for the change in size of the gland towards the base and the apex.

Chapter 2

Thesis Overview

2.1 Organization

The rest of the thesis will be organized as follows. Section 2.2 contains a list of notation to be used throughout the thesis, and Section 2.3 contains a description of our prostate MR data. Chapter 3 will contain a full description of using MRS data to initialize the prostate segmentation scheme, and quantitative results for our ASM initialization. Chapter 4 will contain a full description of the traditional ASM segmentation method and our novel contributions to the ASM, as well as the results on prostate MR images. An application to prostate volume estimation is presented in Chapter 5, and finally concluding remarks and future directions are presented in Chapter 6.

2.2 Notation

We define a T2-w MR intensity image scene as $\mathcal{C} = (C, f)$, where C represents a set of spatial locations (voxels), $f(c)$ is the MR image intensity function associated with every $c \in C$. Throughout the thesis, we will use c, d, or e as placeholders for spatial locations within an image, and i, j, or k as placeholder for scalars. We define the associated spectral scene $\hat{\mathcal{C}} = (\hat{C}, \hat{F})$ where $\hat{C}$ is a 2D grid of metavoxels. Note that a metavoxel is a voxel at the lower spectral resolution. For each spatial location $\hat{c} \in \hat{C}$, there is an associated 256-dimensional valued spectral vector $\hat{F}(\hat{c}) = \left[\hat{f}_j(\hat{c}) \mid j \in \{1, ..., 256\}\right]$, where $\hat{f}_j(\hat{c})$ represents the concentration of different biochemicals (such as creatine, citrate, and choline). The distance between any two adjacent metavoxels, $\hat{c}, \hat{d} \in \hat{C}, \| \hat{c} - \hat{d} \|_2$, (where $\| \cdot \|_2$ denotes the L_2 norm) is roughly 13 times the distance between any two adjacent spatial voxels $c, d \in C$. We define $\mathbf{X} = \{c_1 \dots c_N\} \subset C$ as a set

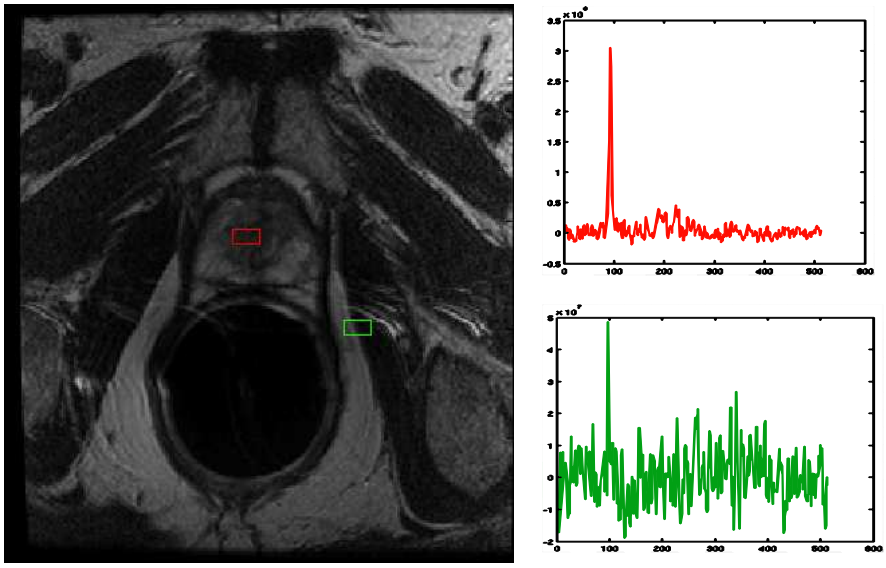

Figure 2.1: 2D section of T2-weighted MR image intensity scene. The spectral signal in red corresponds to the voxel shown in red (within the prostate), and the spectral signal in green corresponds to the voxel shown in green (outside the prostate).

of N landmarks used to define a given prostate shape. The κ-neighborhood of pixels surrounding each $c \in C$ is denoted as $\mathcal{N}_\kappa(c)$, where for $\forall d \in \mathcal{N}_\kappa(c), \| d - c \|_2 \leqslant \kappa, c \notin \mathcal{N}_\kappa(c)$. If we are referring to only those locations in $\mathcal{N}_\kappa(c)$ along the normal direction to the given shape, we denote those voxels are $\hat{\mathcal{N}}_\kappa(c)$.

2.3 Data Description

Our data comprises multi-protocol clinical prostate MR datasets including both MRI and MRS endorectal *in vivo* data. These were collected during the American College of Radiology Imaging Network (ACRIN) multi-site trial [66] (19 studies), from the University of California, San Fransisco (14 studies), and from the Beth Israel medical center in Boston (37 studies). The MRS and MRI studies were obtained on either 1.5 Tesla (1.5T) or 3T MRI scanners, and all MRI studies were axial T2-w images. Each 2D slices had a spatial XY resolution of 256×256 pixels, or 140×140 mm, for the 1.5T data and 512×512 pixels, or 140×140 mm, for the 3T data. An example of a T2-w MR image of the prostate along with its associated spectra is shown in Figure 2.1.

2.4 Quantitative Metrics for Segmentation Accuracy

The ground truth for the prostate boundary on the T2-w images was obtained by manual outlining of the prostate border on each 2D section by a solitary expert radiologist, one with over 10 years of experience in prostate MR imagery. For each image to be segmented, the expert segmentation the prostate region is denoted as a set of ground truth landmarks $\mathbf{X}_E$. For a given shape $\mathbf{X}$, the set of pixels contained within the shape is denoted as $S_{\mathbf{X}}$. The most commonly reported metrics for evaluating prostate segmentation performance are the Dice Similarity Coefficient (DSC) [67] and the mean absolute distance (MAD). These are defined as follows:

$$DSC(\mathbf{X}) = 2 \cdot \frac{|S_{\mathbf{X}} \cap S_{\mathbf{X}_E}|}{|S_{\mathbf{X}}| + |S_{\mathbf{X}_E}|} \tag{2.1}$$

$$MAD(\mathbf{X}) = \frac{1}{N} \sum_{n=1}^{N} (\|c_n - d_n\|_2, c_n \in \mathbf{X}, d_n \in \mathbf{X}_E) \tag{2.2}$$

DSC is a region based performance measures, in which a higher value indicates a more accurate segmentation (maximum value of 1.00). MAD is an edge based performance measures which evaluate proximity of the ASM extracted boundary compared to the manually delineated boundary, in which a lower value reflects a more accurate segmentation (minimum value of 0.00). In general the DSC is used to compare our method to other state of the art methods (as most prostate segmentation papers report the DSC value [56, 58, 59, 60]), and the MAD value is given for completeness.

Chapter 3

Novel ASM Initialization Methodology

Section 3.1 describes an overview of our initialization procedure. Section 3.2 describes our clustering algorithm in detail. Section 3.3 describes fitting the prostate shape to the resulting clusters. Section 3.4 discusses extending the results to the base and apex slices (since the clustering is performed in the midgland), and finally Section 3.5 discusses the quantitative results.

3.1 Overview of ASM Initialization Procedure

Figure 3.1 illustrates the modules and the pathways comprising our automated initialization system. First, replicated k-means clustering is performed on the spectra in the midgland. The largest cluster obtained is identified as the non-informative cluster corresponding to the extra-prostatic spectra and removed. The remaining spatial locations corresponding to the resulting spectra are denoted by S_{MRS}. The prostate shape is fit to the region corresponding to these informative spectra. The clustering results are then extended to the base and apex slices, ersulting in an initial segmentation $\mathbf{X}^0$.

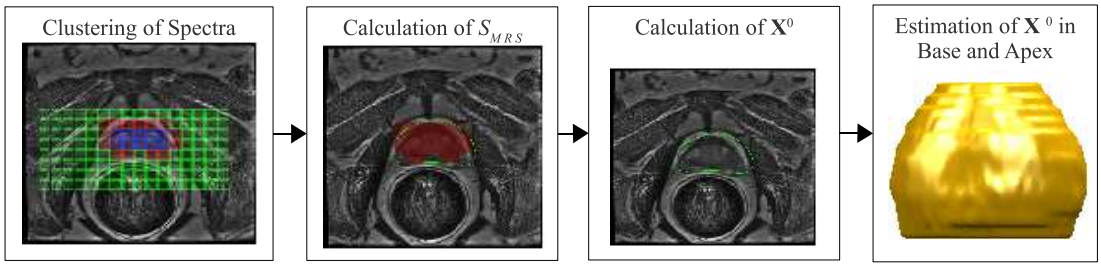

Figure 3.1: Pathways and modules involving in the MRS-based ASM initialization scheme for prostate segmentation on multi-protocol *in vivo* MRI.

3.2 Clustering of Spectra (Calculation of S_{MRS})

The crux of the methodology is to determine a set of prostate voxels (S_{MRS}) based on a clustering of the spectroscopic data. This algorithm is described in the form of a sequence of steps below.

1. For a given 2D MRS slice $\hat{\mathcal{C}} = (\hat{C}, \hat{F})$, we first obtain the MR spectra $\hat{F}(\hat{c}) = \left[\hat{f}(\hat{c}) \mid j \in \{1, \ldots, 256\}\right]$.

2. The metavoxels $\hat{c} \in \hat{C}$, are aggregated into k clusters $V_a \subset \hat{C}, a \in \{1, \ldots, k\}$, by applying k-means clustering to all $\hat{F}(\hat{c}), \forall \hat{c} \in \hat{C}$. k-means clustering aims to minimize the sum of distances to the clusters' centroids, for all clusters. Formally, it iteratively estimates

$$\underset{V_1,\ldots,V_k}{\operatorname{argmin}} \sum_{a=1}^{k} \sum_{\hat{c} \in V_a} \left\| \hat{F}(\hat{c}) - \frac{1}{|V_a|} \sum_{\hat{c} \in V_a} \hat{F}(\hat{c}) \right\|_2 . \tag{3.1}$$

3. Since the k-means algorithm is dependent on the starting locations of the centroids (i.e. which V_a each $\hat{c}$ initially belongs to), the result is sometimes a local minima instead of a global minima. To overcome this limitation, the clustering was repeated 25 times with random initial locations of the centroids, and the resulting clustering which yields the minimum value from Equation 3.1 is selected. Repeating the clustering more than 25 times did not significantly change the results.

4. The dominant cluster is identified as being extra-prostatic (non-informative), and the metavoxels in this cluster are removed. The set of remaining metavoxels is then defined as,

$$S_{\hat{MRS}} = \left\{ \hat{c} \mid \hat{c} \in V_a, a \neq \underset{a}{\operatorname{argmax}} |V_a| \right\}. \tag{3.2}$$

   The set of MRI voxels corresponding to metavoxels in $S_{\hat{MRS}}$ are then idenfitied. For our data, we found that $k = 3$ clusters yielded the best results. Figure 3.2a shows the 3 clusters $V_1 - V_3$ as colored metavoxels. The largest cluster is shown in

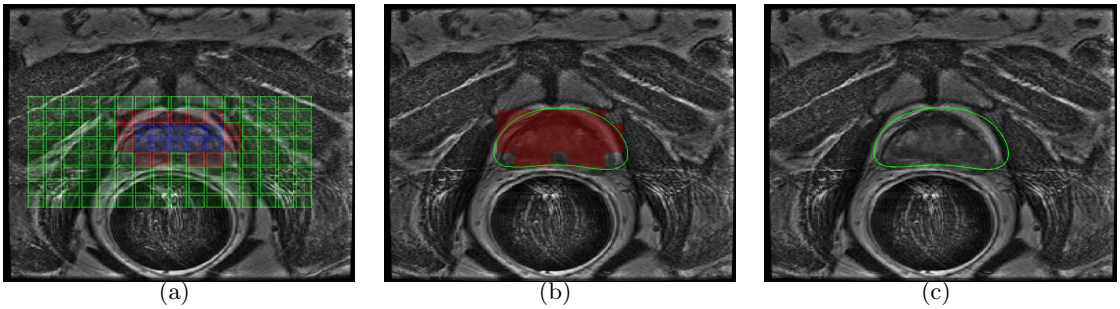

(a) (b) (c)

Figure 3.2: (a) MRS metavoxels $\hat{C}$ (shown as colored boxes) overlaid onto the T2-w image $\mathcal{C}$ where each color represents a different class resulting from the replicated k-means clustering scheme. The voxels associated with the informative metavoxels (S_{MRS}) are shown as a red overlay in (b), with the resulting shape initialization ($\mathbf{X}^0_{MRS}$) shown as a green line in (b) and (c).

green, and would be eliminated, yielding $S_{M\hat{R}S}$ as the cluster comprising blue and red metavoxels. The voxels associated with $S_{M\hat{R}S}$, denoted as S_{MRS} are shown in red in Figure 3.2b.

3.3 Fitting the Prostate Shape (Calculation of $\mathbf{X}^0$)

In [42], Cosio employed the Genetic Algorithm [47] to optimize the pose parameters of the prostate shape to fit a given binary mask. We found that using the objective function described below (Equation 3.3) yielded an accurate initialization for a given set of prostate pixels S_{MRS}. In addition, we found that a direct pattern search as presented in [68] yielded more accurate results than using the Genetic Algorithm. The mean shape $\overline{\mathbf{X}}$ constitutes a polygon, and the set of voxels inside this polygon is denoted as $S_{\overline{\mathbf{X}}}$. More generally, for a given set of affine transformations T, which represent scaling rotation and translation, the set of voxels within that polygon are denoted as $S_{T(\overline{\mathbf{X}})}$. The objective function we use aims to maximize the true positive ratio, so that the initialization is given as

$$\mathbf{X}^0 = T(\overline{\mathbf{X}}), \text{where } T = \underset{T}{\operatorname{argmax}} \left[\frac{S_{MRS} \cap S_{T(\overline{\mathbf{X}})}}{S_{MRS} \cup S_{T(\overline{\mathbf{X}})}} \right], \tag{3.3}$$

where the argmax calculation is performed using a direct pattern search as presented

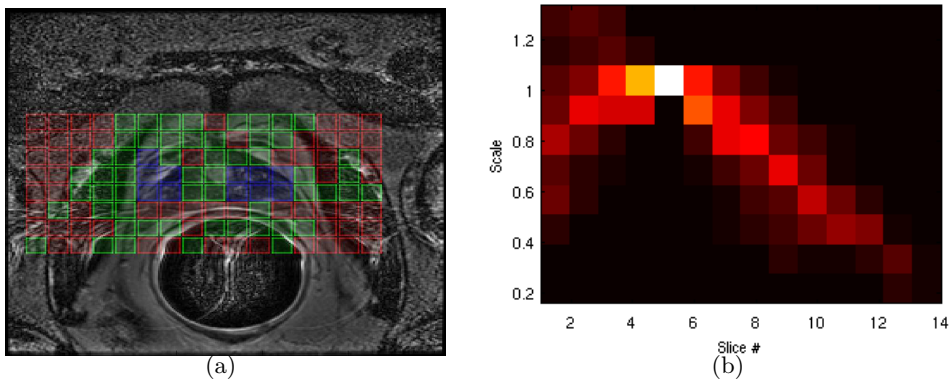

(a) (b)

Figure 3.3: (a) Example of a clustering result in the base of the prostate, in which the 3 colors represent $V_a, a \in \{1, 2, 3\}$. The results demonstrate the degradation of quality of MR spectra in the prostate near the base, a phenomenon which also occurs near the apex. (b) A 2D histogram of the relative size of the prostate as a function of the slice index from base to apex reveals that the prostate is largest in the center and tapers off towards the extrema.

in [68]. Essentially, Equation 3.3 determines which affine transformation T will yield the best initialization, and applies that transformation to the mean shape $\overline{\mathbf{X}}$, resulting in an initialization $\mathbf{X}^0$.

3.4 Initializing the Base and Apex Slices

The MRS spectra lose their fidelity towards the base and the apex of the prostate. This is demonstrated in Figure 3.3a, in which the 3 resulting clusters $(V_1, \ldots, V_3)$ are shown via red, green, and blue metavoxels respectively. It was found that the spectra in the midgland of the prostate yielded accurate estimations of S_{MRS}. For this reason, we perform our clustering algorithm in the midgland of the prostate, the results of which are then extended to the base and apex. Note that since the 2D T2-w MRI slices tend to cover the prostate from base to apex, the middle slice invariably corresponds to the midgland. We observed that the area of the prostate decreases to 80% its size in the base, and 30% its size in the apex. Figure 3.3b demonstrates this tapering off of the gland in the base and apex. Figure 3.3b is a histogram showing the size of the prostate relative to the central slice for all ground truth segmentations. Hence, to calculate $\mathbf{X}^0$

Table 3.1: Quantitative results of the various clustering algorithms used on the MRS data to initialize the segmentation. The values are given in terms of mean $\pm$ standard deviation, over 35 studies.

Metric	Hierarchical	Meanshift	Replicated k-means
DSC	0.28 ± 0.30	0.52 ± 0.33	0.75 ± 0.29
MAD (mm)	5.87 ± 4.35	5.55 ± 3.17	5.60 ± 4.19

for the remaining slices, $\mathbf{X}^0$ is first calculated for slice 5, and is linearly scaled down to 80% its size for the base, and to 30% its size for the apex of the bland.

3.5 Results for Segmentation Initialization

We had 35 datasets with MRS signals available. We compared the efficacy of the replicated k-means clustering scheme with hierarchical clustering and mean-shift clustering. Hierarchical clustering generates a dendrogram based off the Euclidean distance between spectra, and hierarchically combines spectra which have a low distance between them into a single cluster. The process is repeated until a pre-specified number of clusters remain [69]. Mean-shift clustering attempts to iteratively learn the density of the feature space and yields a clustering result based of the manifold instead of a pre-specified number of clusters [70]. Each methodology was used to calculate S_{MRS} in the midgland, and extended the results to the apex and base. These estimations of $\mathbf{X}^0$ were compared to the ground truth segmentations $\mathbf{X}_E$ for all studies. Table 3.1 summarizes the results from clustering experiments performed. The results show that using replicated k-means clustering yielded a more accurate initialization.

Chapter 4

Novel ASM Segmentation Methodology

An overview of our methodology is presented in Section 4.1. Section 4.2 presents our description of how to generate a statistical shape model and Section 4.3 describes our improved appearance model in detail. Finally our methodology for segmenting an image of the prostate is presented in Section 4.4 and quantitative results are given in Section 4.5.

4.1 Overview of Segmentation Methodology

We begin by generating a statistical shape model from a set of manually placed landmark points. This shape model is defined by the mean Cartesian coordinates, and the first few eigenvectors of the Cartesian coordinates. Then, an appearance model is created. We have developed a methodology for creating both a local appearance model near each landmark point, and a global appearance model of the entire prostate. These trained shape and appearance model are then used to segment a new image of the prostate.

4.2 Generation of a Statistical Shape Model

The N Cartesian coordinates of a given shape is defined as $\mathbf{X} = \{x_1, y_1, \ldots, x_N, y_N\}$ in 2D and $\mathbf{X} = \{x_1, y_1, z_1, \ldots, x_N, y_N, z_N\}$ in 3D. The first step to defining a shape model is to take the mean Cartesian coordinates over all M training images as $\overline{\mathbf{X}} = \{\overline{x}_1, \overline{y}_1, \ldots, \overline{x}_N, \overline{y}_N\}$. Then, an eigen-analysis of the set of shapes is performed, where a matrix of the top β eigenvectors is given as $\mathbf{P} \in \mathbb{R}^{|\overline{\mathbf{X}}| \times \beta}$ with corresponding eigenvalues $\boldsymbol{\lambda} = \{\lambda_1, \ldots, \lambda_\beta\}$. β is chosen to represent 98% of the variance seen in the training data,

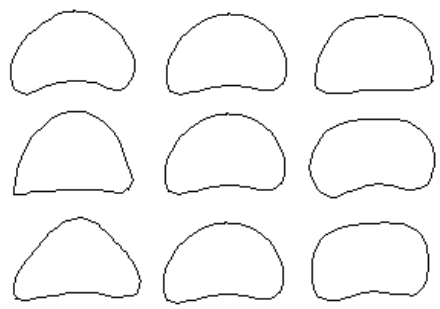

Figure 4.1: Column 1 shows $-3 \cdot \sqrt{\lambda_i}$, column 2 shows $\overline{\mathbf{X}}$, and column 3 shows $+3 \cdot \sqrt{\lambda_i}$. Row 1 signifies $i = 1$, row 2 signifies $i = 2$, and row 3 signifies $i = 3$.

or stated explicitely, β is set as large as possible such that $\sum_{i=1}^{\beta} \lambda_i \leqslant 0.98 \cdot \sum_{i=1}^{|\overline{\mathbf{X}}|} \lambda_i$ is still satisfied.

We can now define an ASM by the equation,

$$\mathbf{X} = \overline{\mathbf{X}} + \mathbf{P} \cdot \mathbf{b}, \tag{4.1}$$

where $\mathbf{b} \in \mathbb{R}^{\beta}$ is a vector defining the shape. Each individual element of $\mathbf{b} \in \{b_i, \ldots, b_\beta\}$ can range between ± 3 standard deviations from the mean shape $\overline{\mathbf{X}}$. This is accomplished by setting bounds on each $\mathbf{b}$, so that $-3 \cdot \sqrt{\lambda_i} < b_i < +3 \cdot \sqrt{\lambda_i}$ for $\forall b_i \in \mathbf{b}$. Therefore, our statistical shape model is defined by $\overline{\mathbf{X}}, \mathbf{P}$, and $\boldsymbol{\lambda}$, and by changing the variable $\mathbf{b}$, we can define an infinite number of prostate shapes, all within ± 3 standard deviations from the mean shape. An example of a 2D statistical shape model of the prostate is shown in Figure 4.1 in which column 1 shows $-3 \cdot \sqrt{\lambda_i}$, column 2 shows $\overline{\mathbf{X}}$, and column 3 shows $+3 \cdot \sqrt{\lambda_i}$. Row 1 signifies $i = 1$, row 2 signifies $i = 2$, and row 3 signifies $i = 3$.

4.3 Generation of a Statistical Appearance Model

4.3.1 Traditional ASM Appearance Model

We begin with a discussion of the traditional method for generating an appearance model (defined in [36]) and then discuss our improved appearance model. In the traditional ASM, for each landmark point c an appearance model is generated. The set of

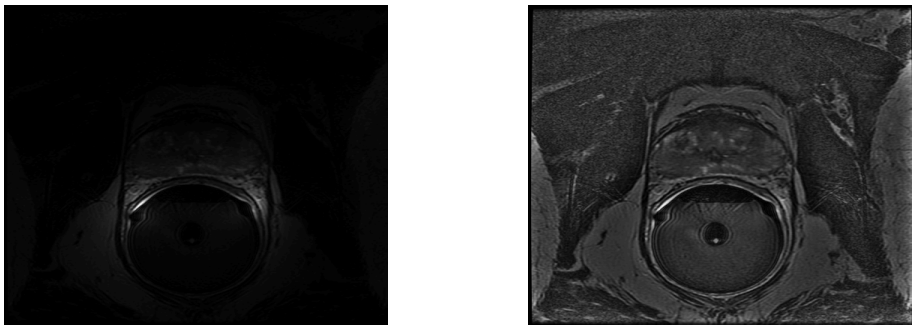

(a) (b)

Figure 4.2: (a) shows a 3 Tesla, T2-w MR image of the prostate. It can be seen that the strong bias field near the center of the image makes the rest of the image too dark to see. (b) shows the same image after preprocessing. The prostate can be seen quite visibly in the center of the image.

intensities in the normal direction to c is given as $F_{\hat{\mathcal{N}}}(c) = \{f(d)|d \in \hat{\mathcal{N}}(c)\}$. For the given anatomical landmark, the mean F over all training images is denoted as μ with a covariance matrix of Σ, which defines a Gaussian distribution and hence constitutes the appearance model for a given landmark point.

4.3.2 Preprocessing and Feature Extraction

In our appearance model, we begin by performing a bias field correction on each MR image. The endorectal coil used in our data creates excessively high intensity values near the prostate. The algorithm in [71] is used to perform the bias field correction on our images. An example of an MR image before and after the bias field correction is shown in Figure 4.2. Once the images have all been corrected, we convolve each image with a Gaussian kernel to remove noise in the images, and finally we normalize each image to have the same range of intensities.

Each image is then convolved with K kernels to derive statistical texture descriptors of the object. The set of kernels is denoted as $S_{\mathbb{k}} = \{\mathbb{k}_1, \ldots, \mathbb{k}_K\}$, where we use the 4 directional Kirsch kernels [63], the 4 directional Sobel kernels [64], Gaussian kernels with standard deviations $\sigma \in \{0.5, 1, 5\}$, and mean kernels with various window sizes.

Examples of a Kirsch and Sobel kernel are,

$$\Bbbk_1 = \begin{bmatrix} 5 & 5 & 5 \\ -3 & 0 & -3 \\ -3 & -3 & -3 \end{bmatrix}, \Bbbk_5 = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}.$$

Denoting the intensities of the neighborhood surrounding c as $F_{\mathcal{N}}(c) = \{f(d)|d \in \mathcal{N}_\kappa(c)\}$ and the convolution operator as $*$, the feature vector $\mathbf{G}(c)$ associated with each $c \in C$ is defined as,

$$\mathbf{G}(c) = \{G_k(c)|k \in \{1, \ldots, K\}\}, \text{ where } \; G_k(c) = F_{\mathcal{N}}(c) * \Bbbk_k. \tag{4.2}$$

4.3.3 Local Appearance Model

Once texture features have been extracted, we can either create a local appearance model near each landmark point c, or we can create a global appearance model for all voxels inside the prostate. In this section we discuss creating a local appearance model, and in Section 4.3.4 we discuss creating a global appearance model.

The most straightforward extension to the traditional ASM appearance model would be to model the distribution of features near each c as a Gaussian. Except in this case, we use each statistical texture feature as a dimension of the Gaussian, whereas the traditional ASM uses nearby intensities to define each dimension of the Gaussian. Hence, we extract μ and Σ as before, but use $\mathbf{G}(c)$ instead of $F_{\hat{\mathcal{N}}}(c)$.

However, this assumes that the distribution of features can be modeled as a Gaussian distribution, which may not necessarily be the case. To overcome this limitation, we use Gaussian Mixture Models (GMM's) [65] to define the distribution. We denote the set of feature vectors from M training images as $S_{\mathbf{G}} = \{\mathbf{G}_1, \ldots, \mathbf{G}_M\}$. We therefore have M datapoints, each with dimensionality $|\mathbf{G}|$. We model this distribution as a linear combination of Q Gaussians by maximizing the log-likelihood of the model. If η denotes a normal (Gaussian) distribution with mean $\mu \in \mathbb{R}^{|\mathbf{G}|}$ and covariance $\Sigma \in \mathbb{R}^{|\mathbf{G}| \times |\mathbf{G}|}$, and p denotes the probability returned by that distribution, then our set of Q distributions is defined by $\boldsymbol{\mu} = \{\mu_1, \ldots, \mu_Q\}, \boldsymbol{\Sigma} = \{\Sigma_1, \ldots, \Sigma_Q\}$ and $\mathbf{w} = \{w_1, \ldots, w_Q\}$ where w denotes

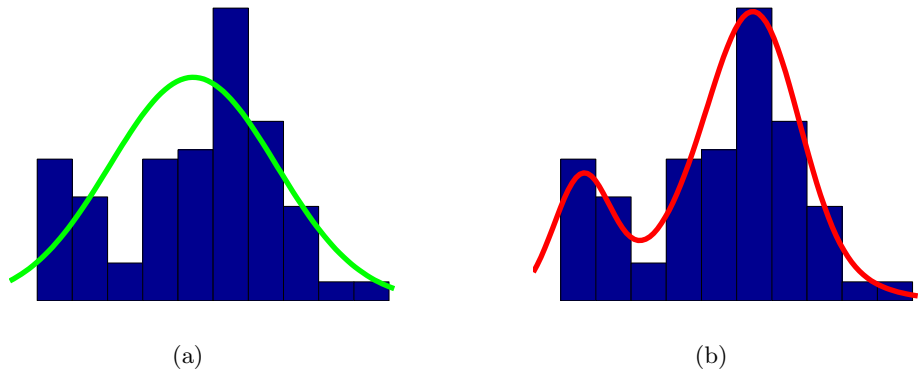

Figure 4.3: A histogram of a texture feature around a landmark, in which (a) shows a Gaussian fit in green, and (b) shows a GMM fit in red. The X axis refers to the value of each texture feature, and the Y axis refers to the number of voxels containing that value. It can be seen that the GMM more accurately models the histogram.

a weight parameter. We therefore use the Expectation Maximization algorithm [65] to maximize the log likelihood L, thereby defining our appearance model by the variables $\boldsymbol{\mu}, \boldsymbol{\Sigma}$, and $\mathbf{w}$, where

$$\boldsymbol{\mu}, \boldsymbol{\Sigma}, \mathbf{w} = \underset{\boldsymbol{\mu}, \boldsymbol{\Sigma}, \mathbf{w}}{\operatorname{argmax}} L, \text{where} \quad L = \sum_{m=1}^{M} \left(\ln \sum_{q=1}^{Q} w_q \cdot p\left(\mathbf{G}_m | \mu_q, \Sigma_q\right) \right). \tag{4.3}$$

To determine the optimal number of Gaussians (Q) to use, the Bayes Information Criteria (BIC) [72] is minimized, which aims to maximize the accuracy of the model while minimizing the number of parameters (the number of Gaussians) to avoid overfitting the training data. The BIC is defined as $BIC = -2 \cdot L + Q \cdot \ln(|\mathbf{G}|)$ where L is the log likelihood value calculated from the EM algorithm in Equation 4.3. Therefore, the Q which minimizes the BIC is selected. To explain the advantage of using GMM's over the Mahalanobis distance, we refer the reader to Figure 4.3, in which a histogram of a texture feature for a given landmark point is shown in blue. In Figure 4.3a, a Gaussian is fit to the data in green. In Figure 4.3b, a GMM is fit to the data in red, which clearly models the data more accurately.

While a GMM defining the distribution of statistical texture features of the border is useful, we can extend this even further by implementing an intelligent method for

selecting which features to include in our multi-dimensional distribution. Our segmentation algorithm was run through a cross-validation scheme using a forward feature selection [73] to choose which dimensions of $\mathbf{G}$ (i.e. which features) to include. This was done by running a 5-fold cross validation on a subset of the training data for each feature individually. The best performing feature was then selected using the mean Dice similarity coefficient (DSC) [67]. Then, each other feature was subsequently tested in a 5-fold cross validation scheme. The feature which improved the DSC the most was then selected. This was repeated until a maximum DSC value was achieved. In the experiments we performed, the DSC value was maximized after 4 features were selected. For the first feature selected, $\mu_q \in \mathbb{R}^1$, and $\Sigma_q \in \mathbb{R}^1$, for the second feature, $\mu_q \in \mathbb{R}^2$ and $\Sigma_q \in \mathbb{R}^{2\times 2}$, and for our data, $\mu_q \in \mathbb{R}^4$ and $\Sigma_q \in \mathbb{R}^{4\times 4}$.

4.3.4 Global Appearance Model

Having a unique appearance model for each landmark point is reasonable for 2D, but in 3D a substantial number of landmark points must be used to define the surface of the object, especially for regions of high curvature. In our system, we used $N = 100$ landmark points to define the shape for each slice in 2D, but required almost $N = 2000$ landmark points to define the surface for each volumetric image. Therefore, a unique appearance model for each landmark point can quickly become infeasible. In addition, for each patient, there is only 1 3D volume, while there are a multitude of 2D slices within that volume. Therefore, for a given number of patients, there are significantly more 2D slices than there are 3D volumetric images. So to accurately define a distribution, there may not be enough 3D volumetric images. For example, if a $|\mathbf{G}| = 15$ dimensional Gaussian is being modeled, several hundred 2D slices could be sufficient to define that 15-dimensional distribution if the ASM is being generated in 2D. However, those several hundred slices could perhaps only be contained within 50 or so 3D volumetric images, which would not be enough to accurately model the 15-dimensional distributions if local appearance models were to be used.

We therefore aim to create a global appearance model of the voxels within the prostate, instead of local appearance models used to define the prostate border, and

aim to minimize the dimensionality of this distribution. The algorithm in [42] defined the distribution of prostate voxels as a GMM, and defined the distribution of non-prostate voxels as a second GMM. A Bayesian classifier was used to distinguish between the 2 distributions when segmenting a new image. In [42], each distribution was 3 dimensional, defined by $\{f(c), x_c, y_c\}$. We did not want to exceed 3 dimensions for our distribution, yet wanted to use multiple statistical texture features. The solution we decided upon was to use the adaboost [74] algorithm to generate a weighted sum of GMM's (one for each texture feature).

The final outputs are the distributions ($\mathcal{D}$) for voxels inside and outside the prostate, and a weight (α_k) for each feature. We denote ω_+ as the class of voxels within the prostate, and ω_- as the class of voxels outside the prostate, with S_ω denoting the set of voxels in class ω. We also denote η as a normal (Gaussian) distribution, and use the method from Section 4.3.3 to define a GMM. $B(c)$ denotes a binary classification, in which $B_k(c) = 1$ if c is correctly determined to be within ω_+ using feature k and $B_k(c) = 0$ otherwise. Finally, the adaboost [74] algorithm (denoted as adaboost) returns a set of weights $\alpha_k, k \in \{1, \ldots, K\}$, for each feature given a set of binary classification results.

Algorithm *GlobalAppearance*:

Input $S_{\omega_+}, S_{w_-}, P(\omega_+), P(\omega_-), t$

for $k := 1$ to K **do**

  for $\omega := \{\omega_-, \omega_+\}$ **do**

    Calculate $[x_c, y_c, F_k(c)]\ \forall c \in S_w$

    Let $Q = \text{argmin}_Q\, BIC$

    Let $\mathcal{D}_{\omega,k} \sim \sum_{q=1}^{Q} \eta\left(\mu_q, \Sigma_q\right)$

  end for

  Let $B_k(c) = \left(\frac{P(\omega_+|c,\mathcal{D}_{\omega_+,k})\cdot P(\omega_+)}{P(\omega_-|c,\mathcal{D}_{\omega_-,k})\cdot P(\omega_-)} > t\right) \leftrightarrow \left(c \in S_{\omega_+}\right) \forall c \in S_{\omega_+} \cup S_{\omega_-}$

end for

Let $\mathbf{B} = \{B_k | k \in \{1, \ldots, K\}\}$

Let $\alpha_k = \text{adaboost}(\mathbf{B}), k \in \{1, \ldots, K\}$

return $\alpha_k, \mathcal{D}_{\omega_+,k}, \mathcal{D}_{\omega_-,k}, k \in \{1, \ldots, K\}$

In the algorithm above, t does not necessarily need to be 0.5, but this is what we found worked well. In addition, adaboost [74] is unique in that it chooses different features which correctly classify different sets of voxels, thereby not redundently selecting similar features. A weight of $\alpha_k = 0$ would indicate that the feature was not chosen by the adaboost algorithm. Therefore, our final global appearance model for the prostate consists of a GMM for the prostate ($\mathcal{D}_{\omega_+,k}$) voxels, the non-prostate voxels ($\mathcal{D}_{\omega_-,k}$), and a weight α_k for each feature $k \in \{1, \ldots, K\}$. It should be noted that for implementation purposes, if $\alpha_k = 0$, then $\mathcal{D}_{\omega,k}$ need not be defined.

4.4 Segmention of an Image of the Prostate

4.4.1 Traditional ASM Segmentation

The process of segmenting an image in the traditional ASM methodology is as follows: detect the border, fit the shape, and repeat until convergence. Given a set of landmark points $\mathbf{X}^i$ for iteration i, the goal is to find landmark points $\hat{\mathbf{X}}^i$ closest to the object border. The shape is then updated using Equation 4.1 where

$$\mathbf{b} = \mathbf{P}^T \cdot \left(\hat{\mathbf{X}}^i - \mathbf{X}^i\right), \tag{4.4}$$

and where each element of $\mathbf{b}$ can only be within ± 3 standard deviations of the mean shape. To determine $\hat{\mathbf{X}}$, the set of pixels along the normal to c (denoted as $\hat{\mathcal{N}}_\kappa(c)$) are checked as potential locations for the prostate border. The standard cost function for a given pixel to the training set is the Mahalanobis distance. Therefore, $\hat{\mathbf{X}}^i$ is defined as

$$\hat{\mathbf{X}}^i = \{d_n \mid n \in \{1, \ldots N\}\}, \text{ where } d_n = \underset{e \in \hat{\mathcal{N}}(c_n)}{\operatorname{argmin}} \left[\left(F_{\hat{\mathcal{N}}}(e) - \mu_n\right)^T \cdot \mathbf{\Sigma}_n^{-1} \cdot \left(F_{\hat{\mathcal{N}}}(e) - \mu_n\right)\right], \tag{4.5}$$

where μ_n and Σ_n are the mean and covariance matrix for landmark n defined in Section 4.3.1.

4.4.2 Segmentation Using a Local Appearance Model

If a local appearance model near each landmark point is created as described in Section 4.3.3, then Equation 4.4 can also be used to define the shape for iteration i. However, instead of minimizing Mahalanobis distance, as in Equation 4.5, we aim to maximize the probability given the GMM. In addition, we chose to not only search in the normal direction, but in a neighborhood surrounding each current landmark point. We therefore define our segmentation cost function as,

$$\hat{\mathbf{X}}^i = \{d_n \mid n \in \{1, \dots N\}\}, \text{ where } d_n = \underset{e \in \mathcal{N}_\kappa(c_n)}{\operatorname{argmax}} P(\mathbf{G}(e) | \boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n, \mathbf{w}_n), \tag{4.6}$$

where $\boldsymbol{\mu}_n, \boldsymbol{\Sigma}_n$, and $\mathbf{w}_n$ are the appearance model for landmark n, as defined in Section 4.3.3.

4.4.3 Segmentation Using a Global Appearance Model

If a global appearance model is created, as in our 3D methodology, a new image can be segmented using the trained distributions $\mathcal{D}$ along with the set of features selected by the adaboost algorithm. The set of voxels calculated to lie within the object are denoted as S_{ω_+}. The set of voxels calculated to lie within the current shape is denoted as $S_{\mathbf{b}}$. The shape is fit to the resulting set of voxels presumed to lie within the object of interest by optimizing $\mathbf{b}$ to maximize the Dice overlap between S_{ω_+} and $S_{\mathbf{b}}$, resulting in a final segmentation. This is summarized in the algorithm below.

Algorithm *ShapeFit*:

Input $\mathcal{C}, \mathcal{D}_{w_+,k}, \mathcal{D}_{w_-,k}, \alpha_k, k \in \{1, \dots K\}$

for $\forall c \in C, \forall k \in \{1, \dots, K\}$ **do**

  Calculate $[x_c, y_c, G_k(c)]$

end for

Let $S_{\omega_+} \leftarrow \left\{ c \middle| \left(\sum_{k=1}^{K} \alpha_k \cdot \frac{P(w_+|c,\mathcal{D}_{w_+,k}) \cdot P(w_+)}{P(w_-|c,\mathcal{D}_{w_-,k}) \cdot P(w_-)} \right) > t \right\}$

Let $\mathbf{b} \leftarrow \operatorname{argmax}_{\mathbf{b}} \frac{|S_{\mathbf{b}} \cap S_{\omega_+}|}{|S_{\mathbf{b}}| + |S_{\omega_+}|}$

Let $\mathbf{X} \leftarrow \overline{\mathbf{X}} + \mathbf{P} \cdot \mathbf{b}$

return $\mathbf{X}$

Table 4.1: DSC values from other state of the art prostate MR segmentation systems, the number of volumes used in the study, the efficiency (in seconds per volume), and the level of user interaction required.

Reference	Volumes	DSC	Seconds	Interaction
Local Appearance Model (2D)	56	0.85	55	minimal
Global Appearance Model (3D)	37	0.89	120	minimal
Makni et al. [60]	12	0.91	76 [60]	none
Klein et al. [56]	50	0.88	900 [60]	none
Pasquier et al. [59]	24	0.88	1200 [60]	medium
Martin et al. [58]	36	0.84	240 [58]	unknown

To perform the argmax calculation, we used a direct pattern search as presented in [68], in which each element in $\mathbf{b}$ was constrained to be within ± 3 standard deviations from the mean shape. Our final segmentation is thus the output $\mathbf{X}$.

4.5 Prostate Segmentation Results

Both the local appearance model (for the 2D ASM) and global appearance model (for the 3D ASM) were compared to other state of the art prostate segmentation schemes for MR imagery. We tested our local appearance model ASM over 56 studies in terms of DSC and MAD. The local appearance model in 2D achieved a mean DSC of 0.8483 with a standard deviation of 0.0448 (standard error of 0.0060 mm). Our mean MAD (over 56 studies) was 2.49 mm with a standard deviation of 0.73 mm (standard error of 0.10 mm). The 3D global appearance model was tested over 37 studies in terms of the DSC, which resulted in a mean DSC of 0.887 with a standard deviation of 0.0329 (standard error of 0.0055). Our system took under a minute for the 2D system and approximately 2 minutes for the 3D system.

Makni et al. [60] achieved a mean DSC of 0.91, with approximately 76 seconds per volume, yet it was only over 12 studies. Klein et al. [56] achieved a median (not mean) DSC of 0.88 over 50 studies. However, it was noted in [60] that Klein et al.'s system took approximately 15 minutes per volume on a standard PC, while our ASM took under 1 minute per volume. Pasquier et al. [59] achieved a mean overlap ratio of 0.784 over 24 studies which corresponds to a DSC of 0.879 as per the equation in [56]. In addition, it

was noted in [60] that it takes approximately 20 minutes to segment an entire volume on a standard PC for the method by Pasquier et al., and it was claimed in [59] that the user must interactively place the segmentation model over the target. Finally, one of the most recent papers on prostate MR segmentation is by Martin et al. [58], in which a mean DSC of 0.84 was achieved over 36 studies, which took approximately 4 minutes per volume. Overall, given the efficiency of our segmentation system, our accuracy is comparable to many state of the art systems. The systems which outperformed our system were tested on less studies and took longer to run. The entire set of comparisons is given in Table 4.1.

Chapter 5

Application to Prostate Volume Estimation

One of the benefits for accurate prostate segmentation is in accurate estimation of the prostate volume. Hence we have employed our 2D prostate segmentation scheme for use in estimating the prostate volume. Of the 56 studies we used in our evaluation of our 2D ASM, we had clinical volume estimations for 45 of them. Section 5.1.1 describes the calculation of a ground truth volume estimation. Section 5.1.2 describes several volume estimations currently performed in the clinic. Section 5.1.3 describes our volume estimation methodology, and Section 5.2 gives the quantitative results for our volume estimations.

5.1 Methods for Estimating Prostate Volume

5.1.1 Ground Truth Volume Estimation

The ground truth volume (V_{Ex}) for the prostate in each of the 45 studies was determined as follows. For each study $\mathcal{C}$, an expert radiologist provided a manual segmentation of the prostate for all slices in which the prostate was visible. Figures 5.1(b),(f),(j) show the expert segmentations in red for a base, midgland, and apex MRI slice respectively. The set comprising the area estimates of the prostate from all M slices within a single 3D study $\mathcal{C}$ is denoted as $S_\Phi = \{A_m(\Phi) \mid m \in \{1, \ldots, M\}\}, \Phi = Ex$, where $A_m(\Phi)$ denotes the segmented area of 2D slice m. The estimated prostate areas (areas within manual delineations) on all slices are integrated and multiplied by the slice interval I (the spacing between adjacent slices). The prostate volume (V_{Ex}) in $\mathcal{C}$ is then calculated as

$$V_{\Phi} = I \cdot \sum_{m=1}^{M} A_m(\Phi). \tag{5.1}$$

5.1.2 Clinical Volume Estimation

Three common techniques currently in use in the clinic for prostate volume estimation (ellipsoid, Myschetzky, and prolate spheroid) depend on the transverse, cranio-caudal, and anterior-posterior distances, all determined from the 3D acquisitions. An expert manually determined the transverse (D_1), cranio-caudal (D_2), and anterior-posterior (D_3) lengths of the prostate for each of the 45 studies.

Ellipsoid

This estimation technique is based on the ellipsoid formula [26] and expressed as,

$$V_{Ell} = D_1 \cdot D_2 \cdot D_3 \cdot \pi/6. \tag{5.2}$$

Myschetzky

This formula [75] aims to slightly increase the volume estimate compared to the ellipsoidal model (Equation 5.2) and is expressed as,

$$V_{Mys} = D_1 \cdot D_2 \cdot D_3 \cdot 0.7. \tag{5.3}$$

Prolate Spheroid

This formulation models the prostate as a prolated spheroid instead of an ellipsoid and is expressed as [26],

$$V_{Sph} = (D_1)^2 \cdot D_2 \cdot \pi/6. \tag{5.4}$$

5.1.3 ASM Based Volume Estimation

We denote V_{ASM} as the volume estimation from the traditional ASM, and V_{MFA} as the 2D multi-feature ASM using a local appearance model, described in Section 4.3.3. V_{ASM} and V_{MFA} require no assumptions about the shape of the gland. This method

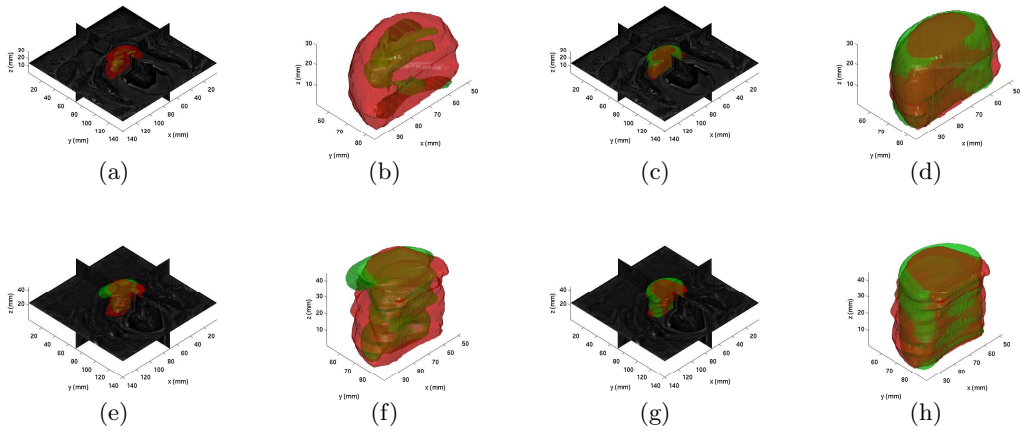

Figure 5.1: Images of the resulting volume from 2 studies. For each result, we show the volumes overlaid on the T2-w MR image (columns 1 and 3), and just the volumes themselves (columns 2 and 4). (a), (b), (e), (f) show $\Phi = ASM$ while (c), (d), (g), (h) show $\Phi = MFA$. In each image, the red represents the ground truth and the green represents the aggregation of segmentations for each slice.

involves an automated ASM segmentation of the prostate in each 2D slice, thereby yielding a set of area estimates for a given study $\mathcal{C}$. These slice areas are then used to determine the prostate volume (as described in Equation 5.1) by considering the inter-slice interval I.

For a given 3D MR acquisition, the 2D ASM algorithm is run on each of the M slices. The set of all M slice areas for a given 3D acquisition is denoted as $S_\Phi = \{A_m(\Phi) \mid m \in \{1, \ldots, M\}\}$. V_{ASM} and V_{MFA} are thus calculated using Equation 5.1.

5.2 Results of Volume Estimations

5.2.1 Pearson's Correlation Coefficient Between V_Φ and V_{Ex}

We first compared the ASM estimated volumes ($\Phi \in \{ASM, MFA\}$) with the clinical models ($\Phi \in \{Ell, Mys, Sph\}$). This was done by calculating the Pearson correlation coefficient [76] (the R^2 value) between V_Φ, for $\Phi \in \{Ell, Mys, Sph, ASM, MFA\}$ and the expert volume estimation V_{Ex}, over all testing studies. The hypothesis was that both V_{ASM} and V_{MFA} should have at least as high of an R^2 value as the commonly

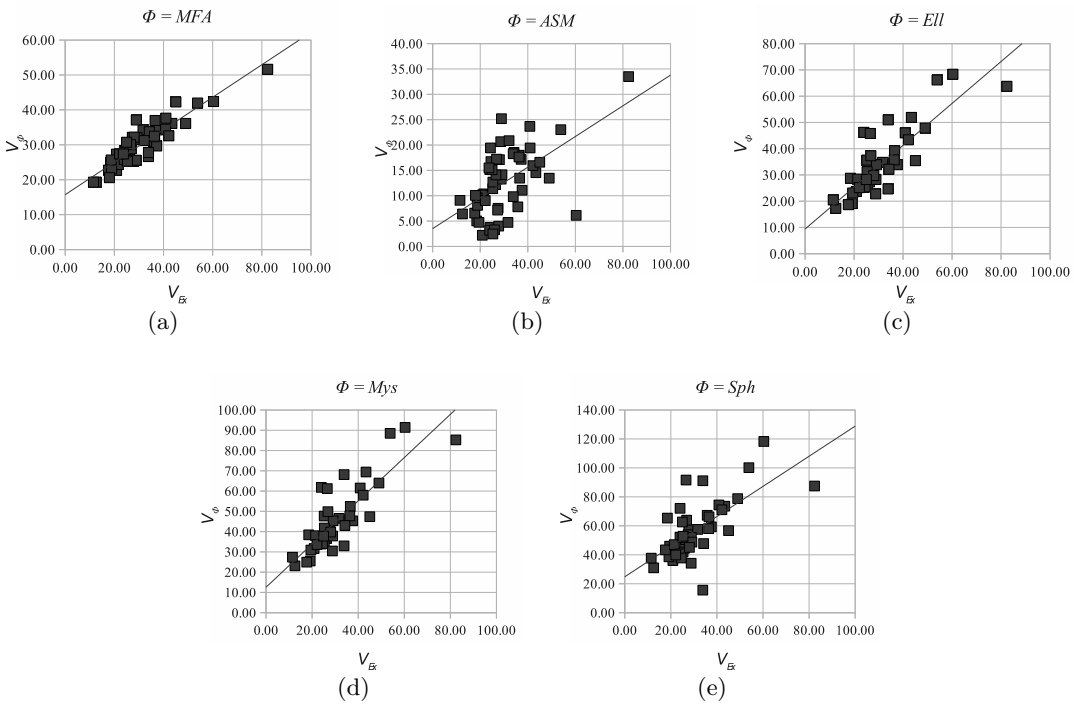

Figure 5.2: Each datapoint represents a patient. V_{Ex} is shown in the X-axis and V_Φ is shown on the Y-axis for $\Phi \in \{MFA, ASM, Ell, Mys, Sph\}$

employed clinical models.

The corresponding R^2 values between V_Φ, for $\Phi \in \{Ell, Mys, Sph, ASM, MFA\}$ and V_{Ex} were computed. The data is shown in Figure 5.2, and the results are shown in Table 5.1. The number of studies used to compute each R^2 value was 45 for the clinical volume estimations, and 56 for the ASM based volume estimations. The lowest R^2 value (0.31) was obtained by using the traditional ASM methodology, and comparatively the highest value (0.82) was obtained from the 2D MFA. It was to be expected that the ellipsoidal and Myschetzky have the same R^2 value, as they are simply scaled variants of each other. In addition, the prolate spheroid expectedly performed the worst of the clinical estimation techniques, as it only used 2 axes in its volume estimation while the ellipsoidal and Myschetzky used measurements from 3 axes.

Table 5.1: Pearson's correlation coefficient (R^2) between V_Φ and V_{Ex} for $\Phi \in \{Ell, Mys, Sph, ASM, MFA\}$.

Φ	Ell	Mys	Sph	ASM	MFA
R^2	.700	.700	.454	.314	.823

Table 5.2: Comparison of V_Φ/V_{Ex} for $\Phi \in \{Ell, Mys, Sph, ASM, MFA\}$ in terms of mean, standard deviation (std.), number of studies (#), and standard error (ste.)

Φ	Mean	Std.	#	Ste.
Ell	1.143	0.252	45	0.0376
Mys	1.528	0.337	45	0.0502
Sph	1.958	0.587	45	0.0875
ASM	0.433	0.189	56	0.0252
MFA	1.053	0.207	56	0.0277

5.2.2 Comparison of Volume Fractions (V_Φ/V_{Ex})

The volume fraction between V_Φ and V_{Ex} was calculated for each study as V_Φ/V_{Ex}, in which a value of 1.00 indicates that the estimated volume is exactly equal to the ground truth volume. The results from these calculations are shown in Table 5.2. The ellipsoid, Myschetzky, and prolate spheroid methodologies had mean volume fractions of $1.14, 1.53$, and 1.96 respectively, with standard deviations of $0.25, 0.34$, and 0.59 respectively. The fact that the ellipsoidal estimation performed better than the Myschetzky estimation was not surprising, as the Myschetzky correction aims to increase the ellipsoidal's estimation.

The traditional ASM had a volume fraction of 0.43 with a standard deviation of 0.18. Finally, the MFA had a volume fraction of 1.05 with a standard deviation of 0.21. The fact that the mean volume fraction from the traditional ASM was significantly less than 1.00 suggests that the ASM consistently undersegmented the gland. This could have occured if the traditional ASM detected edges within the capsule as the capsule boundary, as reflected in Figure 5.1(b),(f). In Figure 5.1, the aggregation of segmentations from all slices are shown in green, with the aggregation of ground truth segmentations shown in red. As stated previously, the ASM consistently detected edges

within the prostate, and thus severely undersegmented the gland. However, the MFA as compared to the traditional intensity-driven ASM was able to adequately determine the correct boundary of the prostate. Note that almost any segmentation scheme (not just ASM's) would find it difficult to correctly identify the prostate boundary on the extreme base and apical sections [77].

Chapter 6

Concluding Remarks and Future Work

In segmenting new images, we have presented a fully automated and accurate ASM initialization scheme for prostate segmentation from multi-protocol *in vivo* MRI/MRS data. With the increasing use of MR imaging of the prostate, several institutions are beginning to acquire multi-modal MR prostate data, including MR spectroscopy [4, 5, 6, 7, 8]. The primary novel contribution of our work on ASM initialization is in leveraging information from one imaging protocol (spectroscopy) to drive the segmentation of the prostate on a different protocol (T2-weighted structural MRI). To the best of our knowledge this is the first instance of multi-modal information being used in this fashion for ASM initialization. Our method uses replicated k-means clustering to cluster the MRS spectra in the midgland, from which we eliminate the background spectra, fit the shape to the remaining spectra, and extend our initializations to the base and apex of the prostate.

We have developed a system for accurately segmenting prostate MR images. In both our 2D and 3D segmentation systems, multiple statistical texture features are extracted from the image, and used to better quantify the prostate. In addition, we have used intelligent feature selection algorithms (either forward feature selection in 2D or adaboost in 3D) to only select those features which offer accurate segmentations. Moreover, we have modeled the distribution of our feature vectors by using Gaussian Mixture Models (GMM's) with an adaptive number of Gaussians. The fact that we sum multiple Gaussians to determine the model for the distribution means that the underlying distribution does not necessarily have to be a Gaussian. Local appearance models around each landmark point are used when sufficient training images exist (such as in the 2D case), and a global appearance model of the prostate is generated when

limited training data prevents accurate models of the distribution (such as in the 3D case).

Our system was compared to other state of the art segmentation systems and performed admirably. It took either a similar amount of time, or less time, than other methods, was tested over more studies, and achieved very high accuracies (0.85 DSC for 2D and 0.89 DSC for 3D). Moreover, our model was used to determine the prostate volume and had a higher correlation with the ground truth ($R^2 = 0.82$) compared to such traditional schemes as the ellipsoid volume estimation ($R^2 = 0.70$), the Myschetzky volume estimation ($R^2 = 0.70$), the prolate spheroid volume estimation ($R^2 = 0.45$) and the traditional ASM volume estimation ($R^2 = 0.31$). In summary, our ASM volume estimation method can save valuable time for clinicians and can yield a consistently accurate prostate volume estimation which is extremely useful for evaluating post-therapeutic response to cancer therapy. Future work will entail testing our system on a large cohort of data and exploring other applications for our accurate segmentation system, such as in guiding biopsies and in treatment planning.

References

[1] W. Catalona et al. Measurement of prostate-specific antigen in serum as a screening test for prostate cancer. *J.Med*, 324(17):1156–1161, 1991.

[2] A. Madabhushi and J. Udupa. Interplay of inhomogeneity correction and intensity standardization in MR image analysis. *Medical Imaging, IEEE Transactions on*, 24(5):561–576, 2005.

[3] Y. Zhu, R. Zwiggelaar, and S. Williams. Prostate segmentation: a comparative study. In *Medical Image Understanding and Analysis*, pages 129–132, 2003.

[4] J. Kurhanewicz, D.B. Vigneron, H. Hricak, P. Narayan, P. Carroll, and S.J. Nelson. Three-dimensional h-1 mr spectroscopic imaging of the in situ human prostate with high (0.24-0.7-cm3) spatial resolution. *Radiology*, 198(3):795–805, Mar 1996.

[5] R. Kumar, R. Nayyar, V. Kumar, N.P. Gupta, A.K. Hemal, N.R. Jagannathan, S. Dattagupta, and S. Thulkar. Potential of magnetic resonance spectroscopy imaging in predicting absence of prostate cancer in men with serum prostate-specific antigen between 4 and 10 ng/ml: A follow-up study. *Technology and Engineering*.

[6] J.C. Vilanova and J. Barcelo. Prostate cancer detection: magnetic resonance (mr) spectroscopic imaging. *Abdominal Imaging*, pages 253 – 262, 2007.

[7] J.J. Hom, F.V. Coakley, J.P. Simko, Y. Lu, A. Qayyum, A.C.A. Westphalen, L.D. Schmitt, P.R. Carroll, and J. Kurhanewicz. High-grade prostatic intraepithelial neoplasia in patients with prostate cancer: Mr and mr spectroscopic imaging features–initial experience. *Radiology*, 242(2):483–489, 2006.

[8] P. Tiwari, M. Rosen, and A. Madabhushi. A hierarchical spectral clustering and non-linear dimensionality reduction scheme for detection of prostate cancer from magnetic resonance spectroscopy. *Medical Physics*, 36(9):3927–3939, 2009.

[9] J. Kurhanewicz, M.G. Swanson, S.J. Nelson, and D.B. Vigneron. Combined magnetic resonance imaging and spectroscopic imaging approach to molecular imaging of prostate cancer. *Journal of Magnetic Resonance Imaging*, 16(4):451–463, 2002.

[10] M. Zaider, M.J. Zelefsk, E.K. Lee, K.L. Zakian, H.I. Amols, J. Dyke, G. Cohen, Y. Hu, A.K. Endi, C. Chui, and J.A. Koutcher. Treatment planning for prostate implants using magnetic-resonance spectroscopy imaging. *nt J Radiat Oncol Biol Phys.*, 47(2):1085–1096, 2000.

[11] H. Kim, D.L. Buckley, D.M. Peterson, G.R. Duensing, J. Caserta, J. Fitzsimmons, and S.J. Blackband. In vivo prostate magnetic resonance imaging and magnetic resonance spectroscopy at 3 tesla using a transceive pelvic phased array coil: Preliminary results. *Investigative Radiology*, 38(7):443–451, 2003.

[12] F.V. Coakley, A. Qayyum, and J. Kurhanewicz. Magnetic resonance imaging and spectroscopic imaging of prostate cancer. *Journal of Urology*, 170(6):S69–S76, 2003.

[13] S Viswanath, B.N. Bloch, M Rosen, J Chappelow, N. Rofsky, R. Lenkinski, E. Genega, A. Kalyanpur, and A. Madabhushi. Integrating structural and functional imaging for computer assisted detection of prostate cancer on multi-protocol in vivo 3 tesla MRI. *SPIE Medical Imaging*, 7260, 2009.

[14] I. Chan, Wells I., William, Robert V. Mulkern, Steven Haker, Jianqing Zhang, Kelly H. Zou, Stephan E. Maier, and Clare M. C. Tempany. Detection of prostate cancer by integration of line-scan diffusion, t2-mapping and t2-weighted magnetic resonance imaging; a multichannel statistical classifier. *Medical Physics*, 30(9):2390–2398, 2003.

[15] P. Tiwari, M. Rosen, and A. Madabhushi. A hierarchical spectral clustering and nonlinear dimensionality reduction scheme for detection of prostate cancer from magnetic resonance spectroscopy (MRS). *Medical Physics*, 36(9):3927–3940, Sep 2009.

[16] P. Tiwari, A. Madabhushi, and M. Rosen. A hierarchical unsupervised spectral clustering scheme for detection of prostate cancer from magnetic resonance spectroscopy (mrs). In *MICCAI*, Lecture Notes in Computer Science, pages 278–286, 2007.

[17] B.N. Bloch, E. Furman-Haran, T.H. Helbich, R.E. Lenkinski, H. Degani, C. Kratzik, M. Susani, A. Haitel, S. Jaromi, L. Ngo, and N.M. Rofsky. Prostate cancer: Accurate determination of extracapsular extension with high-spatial-resolution dynamic contrast-enhanced and T2-weighted MR imaginginitial results. *Radiology*, 245:176–185, 2007.

[18] B.N. Bloch, N.M. Rofsky, R.H. Baroni, R.P. Marquis, I. Pedrosa, and R.E. Lenkinski. 3 tesla magnetic resonance imaging of the prostate with combined pelvic phased-array and endorectal coils: Initial experience. *Academic Radiology*, 11(8):863–867, Aug 2004.

[19] S. Viswanath, B.N. Bloch, E. Genega, N. Rofsky, R. Lenkinski, J. Chappelow, R. Toth, and A. Madabhushi. A comprehensive segmentation registration and cancer detection scheme on 3 tesla in vivo prostate DCE MRI. *MICCAI*, 1:662–669, 2008.

[20] A. Madabhushi, M. Feldman, D. Metaxas, J. Tomasezweski, and D Chute. Automated detection of prostatic adenocarcinoma from high resolution ex vivo MRI. *IEE Trans. Med. Imag.*, 24(12):1611–1625, 2005.

[21] J.M. Kaminski, A.L. Hanlon, E.M. Horwitz, W.H. Pinover, R.K. Mitra, and G.E. Hanks. Relationship between prostate volume, prostate-specific antigen nadir, and biochemical control. *International Journal of Radiation Oncology and Biological Physics*, 52(4):888–892, Mar 2002.

[22] A. Tewari, R. Indudhara, K. Shinohara, E. Schalow, M. Woods, R. Lee, C. Anderson, and P. Narayan. Comparison of transrectal ultrasound prostatic volume estimation with magnetic resonance imaging volume estimation and surgical specimen weight in patients with benign prostatic hyperplasia. *J. Clin. Ultrasound*, 24:169–174, May 1996.

[23] J. Sosna, N.M. Rofsky, S.M. Gaston, W.C. DeWolf, and R.E. Lenkinski. Determinations of prostate volume at 3-tesla using an external phased array coil : Comparison to pathologic specimens. *Academic Radiology*, 10(8):846–853, 2003.

[24] P.M. Pierorazio, M.D. Kinnaman, M.S Wosnitzer, M.C. Benson, J.M. McKiernan, and E.T. Goluboff. Prostate volume and pathologic prostate cancer outcomes after radical prostatectomy. *Urology*, 70(4):696–701, Oct. 2007.

[25] C.G. Roehrborn, P. Boyle, D. Bergner, T. Gray, M. Gittleman, T. Shown, A. Melman, R.B. Bracken, R.V. White, A. Taylor, D. Wang, and J. Waldstreicher. Serum prostate-specific antigen and prostate volume predict long-term changes in symptoms and flow rate: Results of a four-year, randomized trial comparing finasteride versus placebo. *Adult Urology*, pages 662–670, 1999.

[26] C.H. Bangma, A.Q.H.J. Niemer, D.E. Grobbee, and F.H. Schroder. Transrectal ultrasonic volumetry of the prostate: In vivo comparison of different methods. *The Prostate*, 28:107–110, 1996.

[27] S.C. Hoffelt, L.M. Marshall, M. Garzotto, A. Hung, J. Holland, and T.M. Beer. A comparison of ct scan to transrectal ultrasound measured prostate volume in untreated prostate cancer. *Int. J. Radiation Oncology Biol. Phys.*, 57(1):29–32, 2003.

[28] L.M. Eri, H. Thomassen, B. Brennhovd, and L.L. Hheim. Accuracy and repeatability of prostate volume measurements by transrectal ultrasound. *Nature*, 5(4):273–278, Dec. 2002.

[29] P.J. Littrup, C.R. Williams, T.K. Egglin, and R.A. Kane. Determination of prostate volume with transrectal us for cancer screening. part 11. accuracy of in vitro and in vivo techniques. *Radiology*, 52:49–53, 1991.

[30] G.J. Matthews, J. Motta, and J.A. Fracchia. The accuracy of transrectal ultrasound prostate volume estimation: Clinical correlations. *J. Clin. Ultrasound*, 24:501–505, 1996.

[31] M.S. Nathan, K. Seenivasagam, Q. Mei, J.E.A. Wickham, and R.A. Miller. Transrectal ultrasonography: Why are estimates of prostate volume and dimension so inaccurate? *British J. of Urology*, 77(3):401–407, Mar 1996.

[32] S.B. Park, J.K. Kim, S.H. Choi, H.N. Noh, E.K. Ji, and K.S. Cho. Prostate volume measurements by trus using heights obtained by transaxial and midsagittal scanning: Comparison with specimen volume following radical prostatectomy. *Korean J. Radiol.*, pages 110–113, 2000.

[33] Chang Wook J, Hyoung Keun P, Sung Kyu H, Seok-Soo B, Oak Jong L, and Sang Eun L. Comparison of prostate volume measured by transrectal ultrasonography and MRI with the actual prostate volume measured after radical prostatectomy. *Urologia Internationalis*, 81(2):179–185, Aug. 2008.

[34] P.J. MacMahon, A. Kennedy, D.T. Murphy, M. Maher, and M.M. McNicholas. Modified prostate volume algorithm improves transrectal us volume estimation in men presenting for prostate brachytherapy. *Radiology*, 250:273–280, Jan. 2009.

[35] S.K. Warfield, K.H. Zou, and W.M. Wells. *Validation of Image Segmentation and Expert Quality with an Expectation-Maximization Algorithm*, volume 2488, pages 298–306. Springer Berlin Heidelberg, 2002.

[36] T.F. Cootes, C.J. Taylor, D.H. Cooper, and J. Graham. Active shape models - their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, Jan 1995.

[37] T.F. Cootes, G.J. Edwards, and C.J. Taylor. Active appearance models. In *ECCV '98*, pages 484–498, 1998.

[38] J. Fripp, P. Bourgeat, S. Crozier, and S. Ourselin. Segmentation of the bones in mris of the knee using phase, magnitude, and shape information. *Acad Radiol.*, 14(10):1201–1208, Oct 2007.

[39] T. Heimann, I. Wolf, and H.P. Meinzer. Active shape models for a fully automated 3d segmentation of the liver–an evaluation on clinical data. *Med Image Comput Comput Assist Interv.*, 9(2):41–48, 2006.

[40] K. Lekadir and G.Z. Yang. Carotid artery segmentation using an outlier immune 3d active shape models framework. *Med Image Comput Comput Assist Interv.*, 9(1):620–627, 2006.

[41] T. Heimann and H.P. Meinzer. Statistical shape models for 3d medical image segmentation: a review. *Med Image Anal.*, 13(4):543–563, Aug 2009.

[42] F.A. Cosio. Automatic initialization of an active shape model of the prostate. *Medical Image Analysis*, 12(4):469–483, Aug 2008.

[43] M. Brejl and M. Sonka. Object localization and border detection criteria design in edge-based image segmentation: automated learning from examples. *Medical Imaging, IEE Transactions on*, 19(10):973–985, Oct 2000.

[44] T.F. Cootes, C.J. Taylor, and A. Lanitis. Multi-resolution search with active shape models. *Computer Vision and Image Processing*, 1:610–612 vol.1, Oct 1994.

[45] D. Seghers, D. Loeckx, F. Maes, D. Vandermeulen, and P. Suetens. Minimal shape and intensity cost path segmentation. *Medical Imaging, IEEE Transactions on*, 26(8):1115–1129, Aug 2007.

[46] B. van Ginneken, A.F. Frangi, J.J. Staal, B. Romeny, and M.A. Viergever. Active shape model segmentation with optimal features. *medical Imaging, IEEE Transactions on*, 21(8):924–933, Aug 2002.

[47] M. Mitechel. *An Introduction to Genetic Algorithms*. MIT Press, 1998.

[48] Reinhard Beichel, Horst Bischof, Franz Leberl, and Milan Sonka. Robust active appearance models and their application to medical image analysis. *IEEE Trans Med Imag*, 24(9):1151–1170, Sep. 2005.

[49] M. Robers and J. Graham. Robust active shape model search. In *European Conference on Computer Vision*, pages 517–530, 2002.

[50] B. Chiu, G.H. Freeman, M.M.A. Salama, and A. Fenster. Prostate segmentation algorithm using dyadic wavelet transform and discrete dynamic contour. *Physics of Medical Biology*, 49(21):4943–4960, Nov 2004.

[51] J. Costa, H. Delingette, S. Novellas, and N. Ayache. Automatic segmentation of bladder and prostate using coupled 3d deformable models. In *MICCAI*, pages 252–260, 2007.

[52] H. M. Ladak, F. Mao, Y. Wang, D.B. Downey, D.A. Steinman, and A. Fenster. Prostate boundary segmentation from 2d ultrasound images. *Medical Physics*, 27(8):1777–1788, Aug 2000.

[53] Ning Hu, Dónal B. Downey, Aaron Fenster, and Hanif M. Ladak. Prostate boundary segmentation from 3d ultrasound images. *Medical Physics*, 30(7):1648–1659, 2003.

[54] S.D. Pathak, V. Chalana, D.R. Haynor, and Y. Kim. Prostate segmentation algorithm using dyadic wavelet transform and discrete dynamic contour. *IEEE Trans Med Imag*, 19(12):1211–1219, December 2000.

[55] L. Gong, S.D. Pathak, D.R. Haynor, P.S. Cho, and Y. Kim. Parametric shape modeling using deformable superellipses for prostate segmentation. *IEEE Trans Med Imaging*, 23(3):340–349, Mar 2004.

[56] S. Klein, U.A. van der Heide, I.M. Lips, M. van Vulpen, M. Staring, and J.P.W. Pluim. Automatic segmentation of the prostate in 3d MR images by atlas matching using localized mutual information. *Med. Phys.*, 35(4):1407–1417, Apr. 2008.

[57] S. Martin, V. Daanen, and J. Troccaz. Atlas-based prostate segmentation using an hybrid registration. *Int. J. CARS*, 3:485–492, 2008.

[58] S. Margin, V. Daanen, and J. Troccaz. Automated segmentation of the prostate in 3d MR images using a probabilistic atlas and a spatially constrained deformable model. *Med. Phys.*, Mar. 2010.

[59] D. Pasquier, T. Lacornerie, M. Vermandel, J. Rousseau, E. Lartigau, and N. Betrouni. Automatic segmentation of pelvic structures from magnetic resonance images for prostate cancer radiotherapy. *International Journal of Radiation Oncology*Biology*Physics*, 68(2):592–600, 2007.

[60] N. Makni, P. Puech, R. Lopes, and A.S. Dewalle. Combining a deformable model and a probabilistic framework for an automatic 3d segmentation of prostate on MRI. *Int J. CARS*, 4:181–188, 2009.

[61] R. Toth, J. Chappelow, M. Rosen, S. Pungavkar, A. Kalyanpur, and A. Madabhushi. Multi-attribute, non-initializing, texture reconstruction based asm (mantra). In *MICCAI*, volume 1 of *Lecture Notes in Computer Science*, pages 653–661, 2008.

[62] R. Toth, S. Doyle, S. Pungavkar, A. Kalyanpur, and A. Madabhushi. Weritas: Weighted ensemble of regional image textures for asm segmentation. In *SPIE Medical Imaging*, volume 7260, 2009.

[63] R.A. Kirsch. Computer determination of the constituent structure of biological images. *Comput Biomed Res.*, 4(3):315–328, Jun 1971.

[64] R. Duda and P. Hart. *Pattern Classification and Scene Analysis.* John Wiley and Sons, 1973.

[65] B.S. Everitt and D.J. Hand. *Finite mixture distributions.* Chapman and Hall, 1981.

[66] MR imaging and MR spectroscopic imaging of prostate cancer prior to radical prostatectomy: a prospective multi-institutional clinicopathological study. *http://www.acrin.org/Portals/0/Protocols/6659/ A6659partial_summary.pdf*.

[67] L.R. Dice. Measures of the amount of ecologic association between species. *Ecology*, 263:297–302, 1945.

[68] R.M. Lewis and V. Torczon. Pattern search algorithms for bound constrained minimization. *SIAM Journal on Optimization*, 9(4):1082–1099, 1999.

[69] T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning.* Springer, 2 edition, 2009.

[70] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5):603–619, 2002.

[71] M.S. Cohen, R.M. Dubois, and M.M. Zeineh. Rapid and effective correction of rf inhomogeneity for high field magnetic resonance imaging. *Human Brain Mapping*, 10(4):204211, 2000.

[72] S. Theodoridis and K. Koutroumbas. *Pattern Recognition*, page 253. China Machine Press, 4 edition, 2009.

[73] I. Guyon and A. Elisseeff. An introduction to variable and feature selection. *J. Mach. Learn. Res.*, 3:1157–1182, 2003.

[74] Y. Freund and R. Schapire. Experiments with a new boosting algorithm. *Proc. of Natural Conf. on Machine Learning*, pages 148–156, 1996.

[75] P.S. Myschetzky, R.E. Suburu, B.S. Jr. Kelly, M.L. Wilson, S.C. Chen, and F. Lee. Determination of prostate gland volume by transrectal ultrasound: correlation with radical prostatectomy specimens. *Scand J Urol Nephrol Suppl.*, 137:107 – 111, 1991.

[76] A.L. Edwards. *An Introduction to Linear Regression and Correlation*, chapter 4, pages 33–46. W.H. Freeman and Co Ltd, 1976.

[77] S.S. Mahdavi, W.J. Morris, I. Spadinger, N. Chng, O. Goksel, and S.E. Salcudean. 3d prostate segmentation in ultrasound images based on tapered and deformed ellipsoids. In *MICCAI*, volume 1 of *Lecture Notes in Computer Science*, pages 960–967, 2009.

Curriculum Vita

Robert James Toth

Education

2009-2010 M.S. in Biomedical Engineering, Rutgers University

2005-2009 B.S. in Biomedical Engineering (Summa Cum Laude), Rutgers University

Patents

- "Computer Assisted Diagnosis (CAD) of Cancer Using Multi-Functional in vivo Magnetic Resonance Spectroscopy (MRS) and Imaging (MRI)" by Anant Madabhushi, Satish Viswanath, Pallavi Tiwari, Robert Toth, Mark Rosen, John Tomaszeweski, Michael Feldman, PCT/US08/81656, Oct. 2008.

Publications

- Toth, R., Doyle, S., Rosen, M., Pungavkar, S., Kalyanpur, A., Madabhushi, A., WERITAS: Weighted Ensemble of Regional Image Textures for ASM Segmentation, In: SPIE 7258, pp. 725905-725905-11, 2009. (Best Paper Award)
- Toth, R., Tiwari, P., Rosen, M., Pungavkar, S., Kalyanpur, A., Madabhushi, A., An Integrated Multi-Modal Prostate Segmentation Scheme Combining Spectral Clustering and Active Shape Models, In: SPIE 6914, 2008.
- Toth, R., Chappellow, J., Rosen, M., Pungavkar, S., Kalyanpur, A., Madabhushi, A., A Multi-Attribute, Non-Initializing, Texture Reconstruction Based Active Shape Model (MANTRA), In: MICCAI 5241, pp. 653-661, 2008.
- Viswanath, S, Chappelow, J, Toth, R, Linkinski, R, Bloch, B, Madabhushi, A., A Comprehensive Segmentation Registration and Cancer Detection Scheme on 3 Tesla in vivo Prostate DCE MRI, In: MICCAI 11, pp. 662-669, 2008.
- Viswanath, S., Bloch, B.N., Rosen, M., Chappelow, J., Toth,, R., Rofsky, N., Lenkinski, R., Genega, E., Kalyanpur, A., Madabhushi, A., Integrating Structural and Functional Imaging for Computer Assisted Detection of Prostate Cancer on Multi-Protocol in vivo 3 Tesla MRI, In: SPIE 7260, pp. 72603I-72603I-12, 2009.
- Sparks, R., Toth, R, Chappelow, J, Xiao, G, Madabhushi, A. An Integrated Framework for Analyzing Three-Dimensional Shape Differences: Evaluating Prostate Morphometry. In: IEEE International Symposium on Biomedical Imaging, 2010.