Atlas-Based Segmentation of Medical Images

Akinola Olanrewaju Akinyemi, MEng

A themed portfolio submitted to
the Universities of

Edinburgh

Glasgow

Heriot-Watt

Strathclyde

For the Degree of

Doctor of Engineering in System Level Integration

To Akintunde Akinyemi

1986-2005

My Dear Brother

Abstract

Atlas-Based Segmentation of medical images is an image analysis task which involves labelling a desired anatomy or set of anatomy from images generated by medical imaging modalities. The overall goal of atlas-based segmentation is to assist radiologists in the detection and diagnosis of diseases. By extracting the relevant anatomy from medical images and presenting it in an appropriate view, their work-flow can be optimised.

This portfolio-style thesis discusses the research projects carried out in order to evaluate the applicability of atlas-based methods to a variety of medical imaging problems. The thesis describes how atlas-based methods have been applied to heart segmentation, to extract the heart for further cardiac analysis from cardiac CT images, to kidney segmentation, to prepare the kidney for automated perfusion measurements, and to coronary vessel tracking, in order to improve on the quality of tracking algorithms.

This thesis demonstrates how state of the art atlas-based segmentation techniques can be applied successfully to a range of clinical problems in different imaging modalities. Each application has been tested using not only standard experimentation principles, but also by clinically-trained personnel to evaluate its efficacy. The success of these methods is such that some of the described applications have since been deployed in commercial products.

While exploring these applications, several techniques based on published literature were explored and tailored to suit each individual application. This thesis describes in detail the methods used for each application in turn, recognising the state of the art, and outlines the author's contribution in every application.

Contents

List of Tables

List of Figures

Acknowledgment

I would like to thank my supervisors, Dr. Ian Poole and Professor Yvan Petillot for their infinite display of knowledge and support throughout the Engineering Doctorate programme.

Many thanks to the Image Analysis team at Toshiba for their patience and tutorship everyday of the four years I spent learning image analysis, software engineering and a host of other useful things. I would also like to thank the rendering team for their hard-work in producing such brilliant visualisation, making my thesis all the more readable.

Without clinical knowledge, extensive validation and ground-truth, none of the methods in this thesis would have materialised. Many thanks to all the product analysts at Toshiba for putting in the gruelling hours to help my research.

Lots of love and gratitude to my family and friends for their inspiration and for making this process seem a lot easier than it actually was.

My thanks to God for granting me the good health and perseverance to pursue such a great achievement.

Author's Declaration

I, Akinola Akinyemi, declare that this thesis was composed entirely by myself and that the work contained herein is my own except where acknowledged in the text. A list of references has been given in the bibliography. I also declare that this thesis has not been submitted for any other degrees or professional qualifications at any university.

The material contained in this thesis is my own original work produced under the supervision of Ian Poole and Yvan Petillot. Parts of the work have appeared in other publications and these are listed in section 1.2.

(Akinola Akinyemi)

Chapter 1

Portfolio Overview

1.1 Introduction

1.1.1 Medical Imaging: A brief introduction

Medical imaging is defined as: "the techniques and processes used to create images of the human body (or parts thereof) for clinical purposes (i.e. medical procedures seeking to reveal, diagnose or examine disease) or medical science (including the study of normal anatomy and physiology)".

Over the years, several modalities have been invented and adopted for various tasks in medical imaging, however the most prevalent of these in use today are Projection Radiography (X-rays), Computed Tomography (CT), Magnetic Resonance Imaging (MRI), Ultrasound (US) and Positron Emission Tomography (PET).

Projection Radiography X-rays are transmitted through the body, understanding that different parts of the anatomy have different absorption properties, with bones being the highest. This makes projection radiography most suitable for determining the type or extent of fractures, dislocations, and arthritis. It is also used to visualise both benign and malignant tumours, and finds some use in chest pathology.

While conventional use of projection radiography is in planar x-ray, x-rays are also used in fluoroscopy and CT. Fluoroscopy makes use of low-dose x-rays to provide real-time imagery of moving structures usually during surgery. The risks associated with frequent use of this method come from repeated exposure to ionising radiation from the x-rays.

Computed Tomography The CT scanner consists of a patient table; a rotating x-ray source and rotating detectors enclosed in a gantry. The emitted x-ray beams are attenuated by the internal organs and the level of attenuation measured by the ring of detectors. Several mathematical methods, such as filtered back projection and iterative reconstruction techniques are used to

reconstruct the image from these attenuated signals. Modifications of this basic underlying technology have led to the development of helical CT, in which the patient table is moved through the gantry as the x-ray tubes rotate, thus producing a volume of contiguous slice data. Multi-slice/multi-detector CT (MDCT) scanners are now common; these utilise the same principles of the helical scanner, but contain multiple rows of detector rings. This allows the scanner to capture multiple slices per rotation; thus increasing the anatomical area imaged in a fixed time. Toshiba's Aquilion ONE 320-slice CT scanner can cover a 16cm vertical range, capable of imaging an entire organ, in one gantry rotation.

CT imaging is readily available in most hospitals due to its high image resolution (isotropic pixel sizes high-end CT scanners are typically less than $0.5mm^3$), which makes it applicable to a wide range of clinical procedures such as head and neck, cardiac, functional, pulmonary, abdominal and pelvic imaging. The use of short wavelength x-rays cause ionisation in body tissues, damaging cells over constant exposure. A single abdominal CT scan delivers up to $10mSv$ to the patient, compared with $\approx 3mSv$ average yearly background radiation dose in the United States [1].

Magnetic Resonance Imaging Magnetic Resonance Imaging (MRI) applies a magnetic field to force hydrogen nuclei in the body into alignment. Hydrogen is chosen because it is abundant in the human body in form of water and fat. The hydrogen nucleus possesses the 'spin' property which causes it to precess at a known frequency (termed the Larmor frequency) when an external magnetic field is applied, and all the nuclei produce a net magnetisation vector pointing in the direction of this external magnetic field. The Larmor frequency depends on the strength of the applied magnetic field, and is $\approx 42.58MHz/Tesla$ [2].

An RF pulse is applied at the Larmor frequency, following magnetisation until no component of the net magnetisation vector exists parallel to the magnetic field (saturation). Once this pulse is switched off, relaxation occurs emitting waves which are picked up by an RF receiver. Hydrogen protons in different tissues relax at different rates. The relaxation process is divided into T1 (longitudinal) and T2 (transverse); the point at which detection of the emitted wave occurs, along with the strength of the magnetic field is what distinguishes several MR imaging sequences that exist. An image signal is obtained from this RF wave by localisation to a particular slice and 2D co-ordinates using three gradient fields.

The main advantage of MRI over x-rays and CT imaging is that the RF pulses applied to the body are not of high enough frequency to cause ionisation, and hence the risks of tissue damage associated with the other two methods are not applicable. MRI is used for both structural and functional imaging, it provides better contrast between soft tissues than CT. It is used in functional brain imaging, muscle imaging, liver, kidney, and lung imaging. The main downside is that it is a very expensive modality to acquire.

Ultrasound The underlying principle of ultrasound is the piezo-electric effect, which transforms an electrical voltage to high frequency sound waves and vice-versa. An ultrasound scanner consists of a transducer, normally on a hand-held probe. The sound frequencies produced by this transducer are transmitted into the body; the same transducer (or a separate one on the same probe) is used to detect the reflected sound, or echo. The time taken for the echo to be received provides information about the depth of the reflecting anatomy.

Frequencies used in medical ultrasound are in the range of 1 - 10 MHz [3]. High frequency ultrasound has the advantage of producing better resolution due to the magnitude of the reflected signal, but more of the signal energy is absorbed by tissues, reducing its penetration potential. On the other hand, low frequencies produce relatively coarse resolutions, but are capable of imaging deeper tissues within the body. Varying the frequency of the ultrasound waves allows this modality to be used for a wide range of imaging applications such as obstetrics, vascular, abdominal and testicular imaging [4, 5, 6, 7]. Application of the doppler effect to ultrasound also allows it to be used for arterial imaging.

Ultrasound is the cheapest of the cross-sectional imaging modalities around at the moment, making it very efficient as a technology for initial evaluation. Ultrasound scanners are also portable, allowing them to be used in cases where the patient cannot be prepared for a CT or MR scan. Unlike CT, the signals emitted into the body in ultrasound imaging are non-ionising, and so there are no known risks attached to this modality. The low resolution of images compared with other cross-sectional modalities, like CT and MR, is the main disadvantage of ultrasound.

Positron Emission Tomography PET uses positron emitters, attached to pharmaceuticals in order to detect the uptake sites within the body. In PET, the positrons emitted combine with the electrons in the surrounding tissue briefly; and then annihilation occurs, releasing two photons in opposite directions. These photons are measured to reconstruct the image. The most common use of PET is in oncology, and so the most commonly used pharmaceutical is fluorodeoxyglucose (FDG). FDG is a glucose analogue, and the theory here is that cancerous cells are highly energetic, requiring a lot of glucose for metastases. FDG is labelled with fluorine-18 to produce 18F-FDG. This radionuclide has a half-life of 110 minutes [8], therefore it needs to be produced on-site to be clinically useful. Cancers currently investigated by PET are head & neck tumours, thyroid carcinomas, pulmonary nodules, lung cancer, breast cancer, pancreatic cancer, colorectal cancer, ovarian cancer, testicular tumours, Hodgkin's disease, and brain tumours [8].

PET scanners are the most expensive of all the mentioned modalities, and require a cyclotron on site, which is also costly.

The wide range of medical imaging applications targetted by these modalities is good reason for their popularity. A more detailed description of these and other modalities is provided in section 7.3.

Visualisation Cross-sectional modalities like CT, X-Ray, MRI, and PET are all capable of producing multiple 2D image slices along the axis of acquisition; the latest CT scanners can produce 320 slices in one gantry rotation. Rendering techniques have made it common practice to provide orthogonal and oblique multi-planar reconstructions (MPR), making it convenient for clinicians to read scans (figure 1.3). 3D volume rendering is also commonplace in most visualisation software, making it possible to perform non-invasive surgical planning. Advanced visualisation software provide user-interfaces that target the clinicians' workflow, making common protocols more efficient (figure 1.4).

The Digital Imaging and Communication in Medicine (DICOM) format for medical image storage and transmission has been widely adopted as the standard. It defines the way images are stored, allowing JPEG, JPEG 2000 and lossless JPEG codecs. DICOM contains predefined tags specifying demographic information such as age, sex, and the name of the attending health institution. It also defines a TCP/IP protocol for communicating images between systems.

The current state of the art in radiology labs is by far more technologically advanced than it was 20 years ago. Today, the presence of Picture Archival and Communications Systems (PACS) networks in hospitals means that scans in various modalities can be performed in the imaging centre and transferred to a database as DICOM images, where they can be pre-processed by a server running image analysis software to extract useful information for display and analysis and made accessible to radiologists, radiographers and other authorised medical personnel (see figure 1.1).

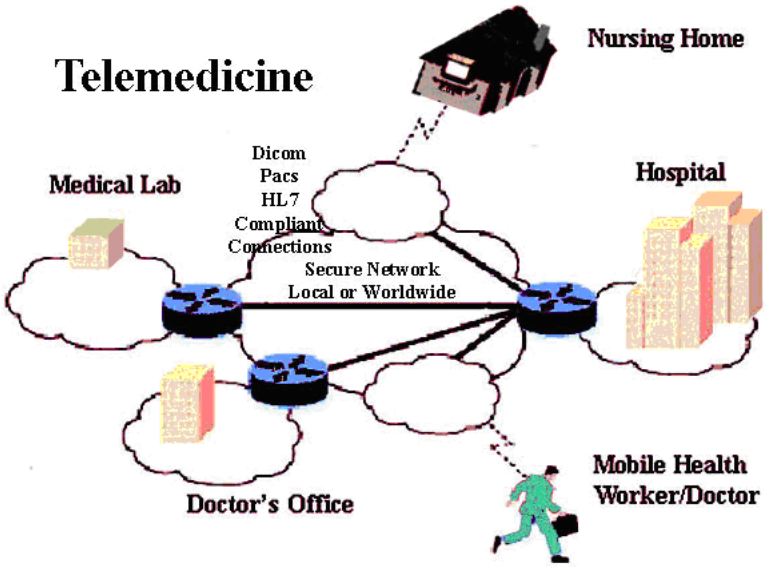

Figure 1.1: A typical PACS Network
The scans can be performed in a lab and the relevant information accessed at any location connected to the network.

1.1.2 Motivation

The increasing importance of large 3D and 4D medical datasets from ultrasound, MRI, and particularly MDCT are placing an increasing load on clinicians. There is therefore a growing need for automation to arrive quickly at the particular view or measurements they need to make their diagnosis.

1.1.3 Automatic Segmentation

Automatic segmentation is the process of delineating known structures of interest from these images. Such structures could be gross anatomy, such as the heart, kidneys, liver, and the brain, or sub-anatomical structures like the myocardium, coronary arteries and tumours. The aim of automatic segmentation is to be 1) robust, capable of dealing with varying image quality, 2) accurate, such that automatic measurements can be carried out based on the results, 3) fast enough to ease clinicians' workflow, and 4) reproducible, such that segmentations from an image can be re-generated precisely.

1.1.4 What is an Atlas?

The term atlas conventionally refers to a labelled map, used as a navigation aid to find the location of places relative to other points or within a geographical region. These atlases can be precise, providing one to one correspondence between points on the map and the actual locations. An atlas can therefore be generalised as a set of descriptors for a particular structure in relation to a set of structures. In medicine, the atlas of human anatomy is a fully labelled diagram of a human being, allowing its students to locate specific anatomical or sub-anatomical areas.

In the literature associated with this thesis, an atlas is considered as a means of uniquely describing specific structures within the body. It may be as simple as a conventional human atlas, consisting of a modality-specific medical image with regions labelled by clinical experts [9, 10, 11, 12, 13]. This is extended by aligning several labelled images, representing the atlas by the mean and variance images [14]. Other definitions maintain a single image as reference, and propagate the labelled structures from other images to obtain a probabilistic estimate of the true anatomy [15, 16]. Another concept of an atlas represents an anatomical structure by a set of parameters governing its shape, hence making such an atlas modality-agnostic. An example of this concept is in active shape models [17, 18, 19, 20], whereby anatomical variability is encoded by modelling its mean shape and principal component modes collected from several examples. An atlas could also be less physically comprehensible, existing as a set of semantic rules defining the characteristics of certain structures within the body [21]. Figure 1.2 shows some examples of how atlases can be represented.

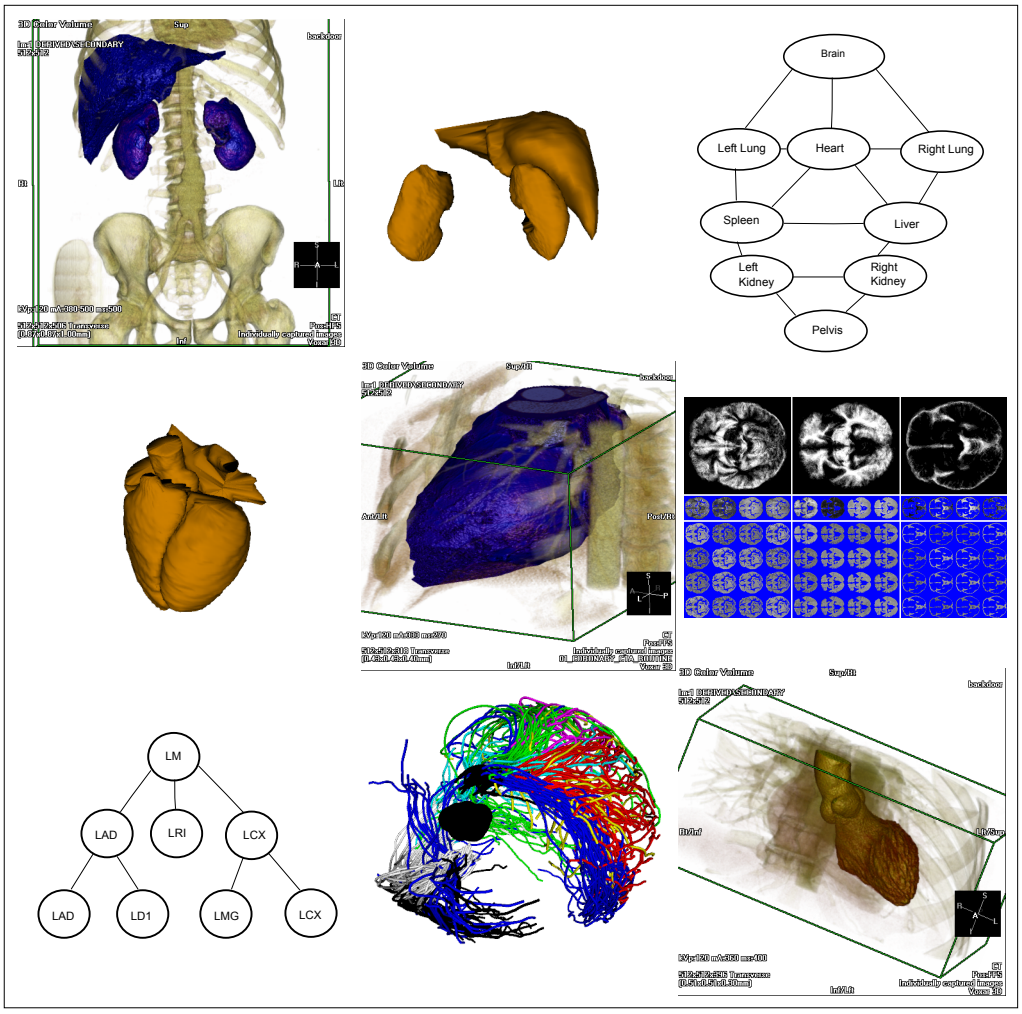

Figure 1.2: Atlases in Medical Image Analysis

Top Row (Left to right): An abdominal 3D CT scan with manually labelled kidneys and liver, a modality-agnostic mesh representation of the kidneys and liver from a single image, graph representing relative positions of anatomical structures. Middle Row: A mesh representation of sub-structures within the heart(visible are the ventricles and descending aorta), volume of interest from a 3D thoracic CT scan with a manually labelled heart, Mean and covariance matrices of MR intensity images for grey matter, white matter and CSF. Bottom Row: Tree representation of the topology of the coronary arteries, coronary artery centrelines from 42 images aligned to a reference space, labelled aorta and left ventricle blood pool from a single 3D thoracic CT scan.

1.1.5 Why Atlas-Based Segmentation?

The concept of using manually labelled medical images to segment relevant anatomical structures from a novel image becomes feasible due to advances in alignment technology, particularly registration. Registration is the process of aligning images on a voxel-based level, with advanced methods capable of providing sub-voxel accuracy through interpolation. Once an atlas image is accurately aligned to a novel image, segmentation is simply the process of propagating the relevant labelled voxels onto the novel image. Labelled anatomy within a scan is a source of positional, topological and shape information about the structures of interest. It provides this information in a physically conceivable framework, making it easy for clinicians to construct atlases as required. The extra benefit of this framework is that subtle features pertaining to anatomy, which may have been overlooked by conventional feature detection methods are implicitly described.

Anatomical variations naturally exist, and atlases can capture these variations in a number of ways provided enough data is available in the training/atlas-creation stages. Statistical atlases, modelling population variance are a good example of such means. Atlas-based segmentation therefore provides a fast, simple, elegant framework that the clinician as the end-user can place confidence in.

1.1.6 Chronology

The projects carried out over the past four years investigate the efficacy of atlas-based methods in solving clinical problems presented to Toshiba Medical Visualisation Systems (TMVS) for incorporation in advanced visualisation software. Five major pieces of work were carried out during the period of research, each focusing on a specific problem.

Coronary Artery Labelling

The first piece of work was aimed at automatically assigning the correct anatomical labels to pre-segmented coronary arteries. This work was targetted at an existing cardiac analysis software suite, capable of evaluating stenoses in coronary arteries as a means of assessing the risk of coronary artery disease (CAD). This work is somewhat disconnected from the other pieces of work, in that it does not apply the same definition of an atlas as the rest due to the nature of the task. The concept of an atlas is presented as a set of rules governing the physical, spatial and topological properties of each coronary artery. A feature-based Maximum a-Posteriori (MAP)classifier is applied in this and subsequent chapters.

Coronary Artery Atlas Creation

The second piece of work was aimed at detecting for removal, structures incorrectly identified as coronary arteries following automatic coronary artery centreline tracking. This work was also targetted at the cardiac analysis software suite, as a means of increasing the specificity of the existing artery centreline tracking algorithm. A coronary artery centreline atlas is created from a set of labelled cardiac CT angiography (CCTA) images, and used to detect non-arterial structures in the input centrelines. This is the first piece of work in the portfolio that addresses the concept of an atlas as it is used throughout the rest of the research.

Whole Heart Segmentation

The third piece of work was aimed at automatically segmenting the whole heart from 3D CCTA images to provide uninterrupted volume-rendered views for clinicians and to seed the other segmentation algorithms in the cardiac analysis package, such as aortic root segmentation and coronary ostia finding for arterial centreline tracking. Variation in cardiac anatomy proved to be an issue when using an atlas image pre-selected from a database of images, particularly when faced with patient images from different continents. This work therefore presents a novel method of selecting the most similar atlas image to the patient image following registration in order to increase the accuracy of the segmentation. The method for optimal atlas selection is used in subsequent pieces of work.

Kidney and Renal Cortex Segmentation

The fourth piece of work was aimed at automatically segmenting the kidneys and their renal cortices from 3D dynamic contrast-enhanced abdominal CT (DCE-CT) images in order to facilitate automatic measurements of perfusion within the kidneys. DCE-CT images are normally noisy because radiographers aim to reduce the overall radiation dose exposed to the patient by limiting the x-ray beam current. To add to this challenge, kidney structure varies widely among individuals, not only by shape, but also by their positioning within the abdomen. This work relied heavily on the optimal atlas selection procedure introduced previously, and also incorporated intensity-based classification as a refinement stage by using an semi-supervised MAP classifier. The concept of an atlas containing multiple labelled structures was introduced in this work.

Multi-Compartment Heart Segmentation

The fifth and final piece of work was aimed at increasing the granularity of the whole heart segmentation by identifying sub-compartments within the heart. In this work, the atlas contained up to eight labelled structures of interest with similar intensities. This posed problems for the

MAP classifier used in the previous work, and led to development of a generic framework for MAP classification based on a spatially-aware expectation maximisation (EM) algorithm.

1.2 Contributions

1.2.1 Commercial

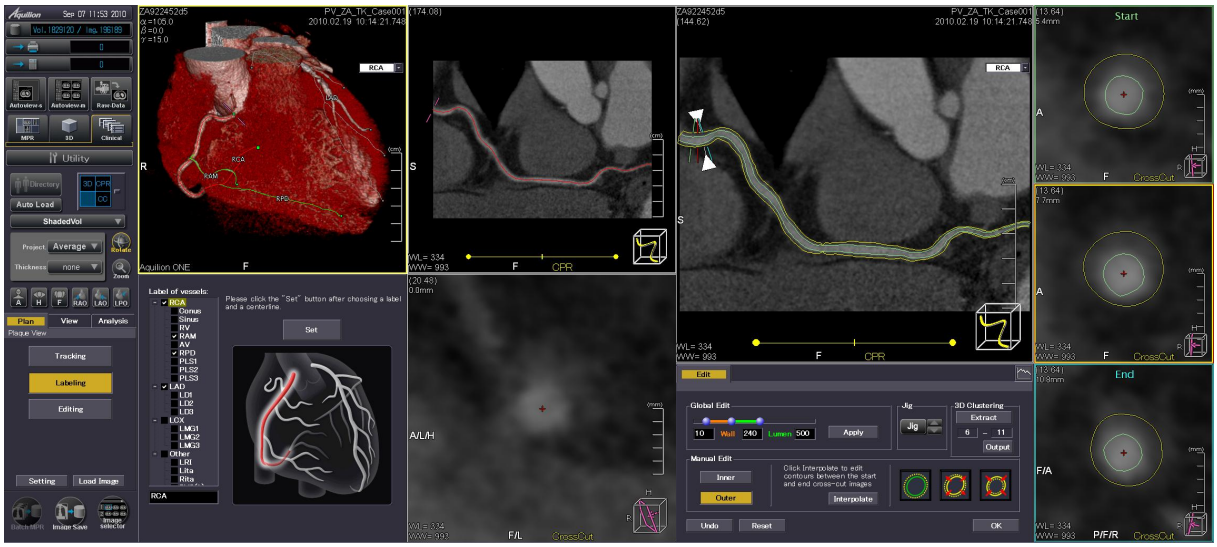

Figure 1.3: Advanced Coronary Artery Application: Toshiba's Aquilion®ONE Console
The advanced coronary artery application within Toshiba's Aquilion ONE CT scanner console performs coronary plaque assessment. The application performs automatic heart segmentation, coronary centreline tracking, coronary centreline labelling, coronary vessel segmentation and lumen analysis, allowing the clinician to validate and edit results.

- Whole Heart Segmentation and Coronary Artery Labelling: The automatic heart segmentation and coronary artery labelling algorithms are incorporated in Toshiba's Aquilion ONE console PC within an application performing advanced coronary artery analysis. Figure 1.3 shows screen-shots of this application.

- Kidney and Renal Cortex Segmentation: Toshiba's Aquilion ONE console PC contains an application for abdominal analysis. This application performs polyp-detection in the colon and perfusion analysis within the liver and kidneys. In order to carry out the latter function, it performs automatic segmentation of the kidney and renal cortex in each scan and the computed volumes are used to measure the glomerular filtration rate (GFR). Figure 1.4 shows screen-shots of the application.

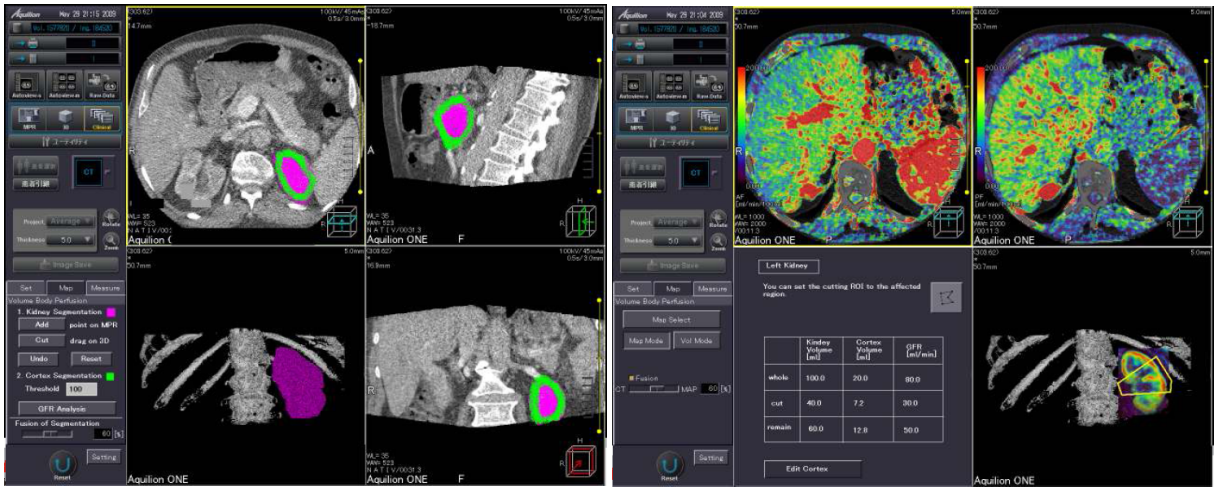

Figure 1.4: Abdomen Application: Toshiba's Aquilion®ONE Console
The abdomen application within Toshiba's Aquilion ONE CT scanner console performs colon analysis and blood-flow analysis within the abdominal organs. It measures kidney perfusion rates by carrying out automatic segmentation of the kidney and renal cortex and measuring the change in contrast-intensity over time.

1.2.2 Intellectual Property

Patents Pending

- A. Akinyemi, S. Murphy and I. Poole, *"Method and Apparatus for Classification of Coronary Artery Image Data,"*. United States Patent Application 20100082692, filing date: 24/09/2008

- A. Akinyemi, C. Plakas, J. Piper and I. Poole, *"Image Segmentation,"*. United States Patent Application 12/847372, filing date: 27/08/2010

1.2.3 Publications

- A. Akinyemi, S. Murphy, C. Roberts and I. Poole, *"Automatic Labelling of Coronary Arteries,"* in *17th European Signal Processing Conference (EUSIPCO 2009)*. Pages 1562-1566. Available: `http://www.eurasip.org/proceedings/eusipco/eusipco2009/contents/papers/1569187117.pdf`

Draft Ready: Awaiting Dataset Permissions

- A. Akinyemi, C. Plakas, J. Piper, C. Roberts and I. Poole, *"Optimal Atlas Selection Using Image Similarities in a Trained Regression Model to Predict Performance,"* in *IEEE Transactions on Medical Imaging. Drafted: September 2010*

1.3 Portfolio Organisation

The entire period of research was spread out over four years, during which individual projects were carried out, all following the same underlying theme of atlas-based segmentation. The EngD programme requires research engineers to carry out a business orientated project in addition to the themed projects, this work is presented in chapter 7. This report represents a portfolio containing a description of each individual project. The projects are categorised into technical or business projects, figure 1.5 shows layout of the chapters and their relationships.

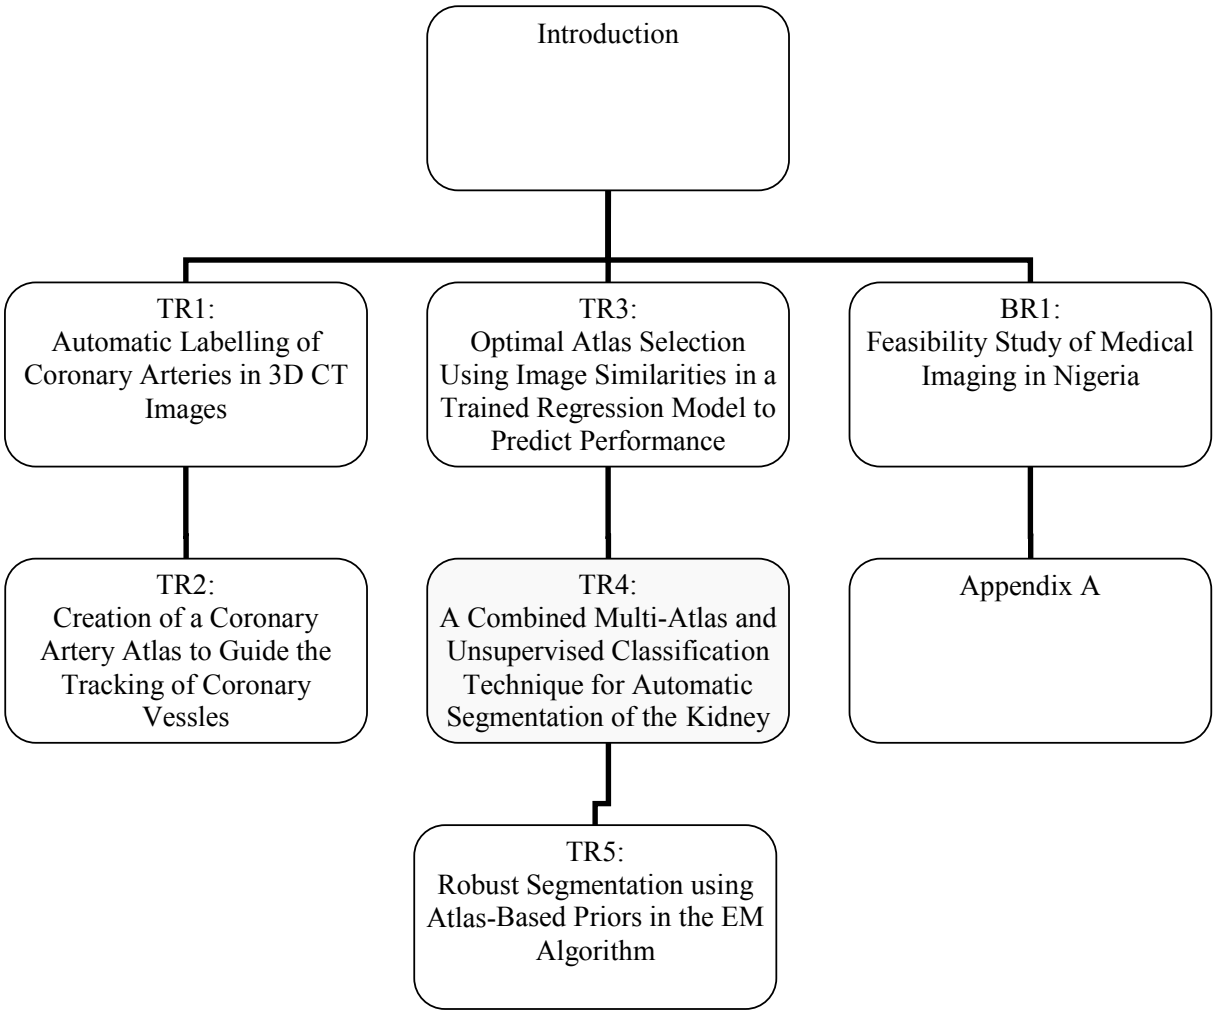

Figure 1.5: Portfolio Organisation (Reading Guide)
TR => Technical Report, BR => Business Report

References

[1] Radiological Society of North America, "Safety : Radiation Exposure in X-ray Examinations," pp. 1–4, 2009.

[2] E. Berry and A. Bulpitt, *Fundamentals of MRI: An interactive Learning Approach*. CRC Press, 2009.

[3] T. S. Curry III, M.D., E. Dowdey, James, and R. C. Murry, Jr., *ChristensenŠs Physics of Diagnostic Radiology*. Lea & Febiger, 1990.

[4] ACR, *ACR practice guideline for the performance of ultrasound evaluation of the prostate (and surrounding structures)*. American College of Radiology, 2006.

[5] ACR, *ACR practice guideline for the performance of a breast ultrasound examination*. American College of Radiology, 2007.

[6] ACR, *ACR practice guideline for the performance of obstetrical ultrasound*. American college of Radiology, 2007.

[7] ACR, *ACR practice guideline for the performance of an ultrasound examination of the abdomen and/or retroperitoneum*. American College of Radiology, 2007.

[8] J. Ruhlmann, P. Oehr, and H. Biersack, *PET in Oncology: Basics and Clinical Applications*. Springer-Verlag, 1999.

[9] P. Aljabar, R. Heckemann, A. Hammers, J. Hajnal, and D. Rueckert, "Multi-atlas based segmentation of brain images: Atlas selection and its effect on accuracy," *NeuroImage*, vol. 46, no. 3, pp. 726–738, 2009.

[10] X. Artaechevarria, A. Munoz-Barrutia, and C. Oritz-de Solorzano, "Combination Strategies in Multi-Atlas Image Segmentation: Application to Brain MR Data.," *IEEE transactions on medical imaging*, vol. 28, no. 8, pp. 1266–1277, 2009.

[11] H. Kirisli, M. Schaap, S. Klein, L. Neefjes, A. Weustink, T. van Walsum, and W. Niessen, "Fully automatic cardiac segmentation from 3d cta data: a multiatlas based approach," in *SPIE Medical Imaging 2010*, 2010.

[12] J. Lötjönen, R. Wolz, J. Koikkalainen, L. Thurfjell, G. Waldemar, H. Soininen, and D. Rueckert, "Fast and robust multi-atlas segmentation of brain magnetic resonance images," *NeuroImage*, 2009.

[13] T. Rohlfing, R. Brandt, R. Menzel, D. B. Russakoff, and C. R. Maurer, Jr., "Quo vadis, atlas-based segmentation?," pp. 435–486, Aug. 2005.

[14] T. Rohlfing, R. Brandt, C. R. Maurer Jr., and R. Menzel, "Bee brains, b-splines and computational democracy: Generating an average shape atlas," *Mathematical Methods in Biomedical Image Analysis, IEEE Workshop on*, vol. 0, p. 187, 2001.

[15] D. C. J. J. D. C. S. Aylward, "Vascular atlas formation using a vessel-to-image affine registration method," *Medical Imaging Computing and Computer-Assisted Intervention (MICCAI)*, 2003.

[16] H. Park, P. Bland, and C. Meyer, "Construction of an abdominal probabilistic atlas and its application in segmentation," *Medical Imaging, IEEE Transactions on*, vol. 22, pp. 483–492, April 2003.

[17] T. F. Cootes, G. J. Edwards, and C. J. Taylor, "Active appearance models," *Proceedings of the European Conference on Computer Vision*, vol. 2, pp. 484–498, 1998.

[18] A. Frangi, W. Niessen, D. Rueckert, and J. Schnabel, "Automatic 3D ASM construction via atlas-based landmarking and volumetric elastic registration," in *Information Processing in Medical Imaging*, vol. 2082, p. 78âĂŞ91, Springer, 2001.

[19] T. Okada, K. Yokota, M. Hori, M. Nakamoto, H. Nakamura, and Y. Sato, "Construction of hierarchical multi-organ statistical atlases and their application to multi-organ segmentation from CT images.," *Medical image computing and computer-assisted intervention : MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention*, vol. 11, pp. 502–9, January 2008.

[20] G. Vesom, N. Cahill, L. Gorelick, and J. Noble, "Characterization of Anatomical Shape Based on Random Walk Hitting Times," *MFCAâĂŹ08*, p. 117.

[21] N. Ezquerra, S. Capell, L. Klein, and P. Duijves, "Model-guided labeling of coronary structure.," in *Medical Imaging, IEEE Transactions on,*, vol. 17 of *3*, pp. 429–441, 1998.

Chapter 2

Automatic Labelling of Coronary Arteries in 3D CT Images

Abstract *Automatically assigning the correct anatomical labels to coronary arteries is an important task that would speed up work flow times of radiographers, radiologists and cardiologists, and also aid the standard assessment of coronary artery disease. However, automatic labelling faces challenges resulting from structures as complex and widely varied as coronary anatomy.*

A supervised classifier has been implemented which addresses this requirement and is capable of automatically assigning correct anatomical labels to pre-segmented coronary artery centrelines in CCTA images with 84% accuracy.

The system consists of two major phases: 1) training a multivariate Gaussian classifier with labelled anatomies to estimate mean feature vectors for each anatomical class and a covariance matrix pooled over all classes, based on a set of features; 2) generating all plausible label combinations per novel anatomy based on a set of topological and geometric rules, and returning the most likely based on the parameters generated in 1).

2.1 Introduction

2.1.1 Motivation

Heart disease is a major life-threatening disease in Europe and America at the moment; the rate of heart disease increasing in continents such as Africa and Asia every year almost at the same rate of development of these countries [1].

Cardiac CT angiography is becoming an important procedure for detecting stenoses and other pathology in coronary arteries, which are the main causes of heart attacks. In cardiac analysis, the radiologist/cardiologist is required to perform functional and structural analysis on

the coronary arteries. Advanced visualisation tools (figure 2.1) make this process more efficient by highlighting the areas of interest and even automating the necessary measurements.

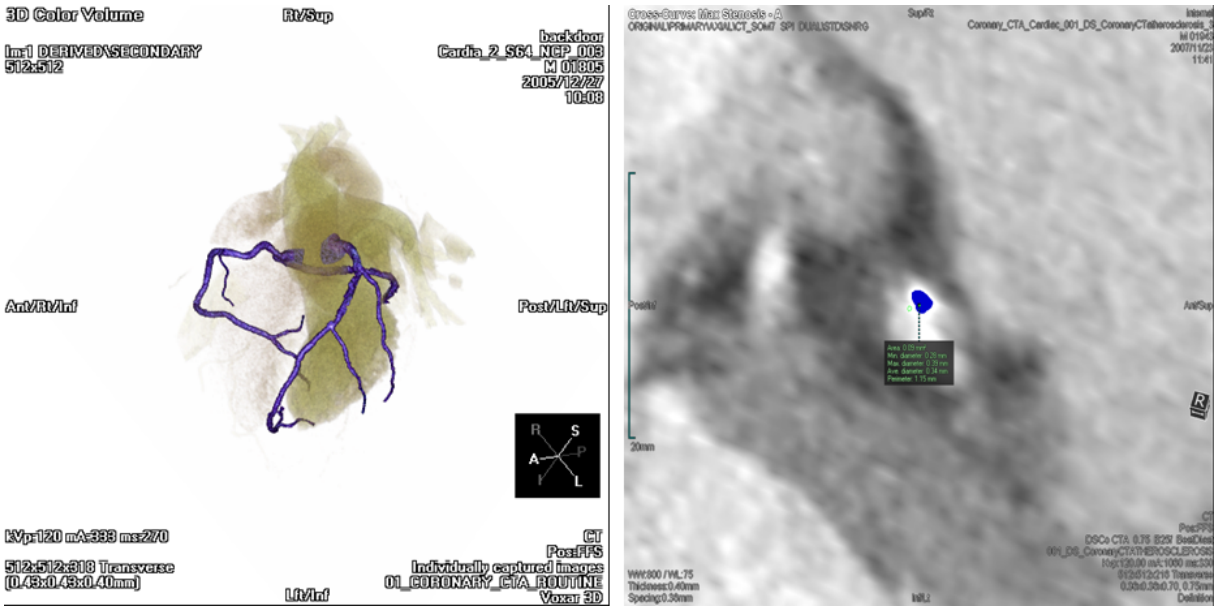

Figure 2.1: Coronary Vessel Analysis
This demonstrates the extent of clarity by which a computed tomography (CT) scan of the heart can be visualised using image analysis software incorporated in PACS. The image on the left is a volume-rendering of the CCTA image data, highlighting the segmented coronary arteries, while the image on the right demonstrates how automatic vessel analysis tools can measure stenoses in any of the arteries.

The application is intended for use by radiologists, therefore, being able to automatically locate specific arteries for analysis will assist in the reading of cardiac images, enabling automatic generation of reports containing standard locations for the measured stenoses, thus speeding-up their workflow and reducing the time to detect and diagnose coronary artery diseases. The availability standardised reports for coronary analysis can aid further data-mining applications for epidemiological studies. Human error is still very much present even in medicine, therefore an accurate automatic coronary artery labelling system can eventually reduce the risk of this type of errors in misclassification.

2.1.2 The Problem

The aim of this project is to create a system which automatically assigns the correct anatomical labels, as specified by the American Heart Association's (AHA) 17-segment model [2], to the coronary arteries. The algorithm takes as input, a pre-segmented coronary arterial tree and outputs the label for each vessel segment as shown in figure 2.2.

This report is structured as follows:

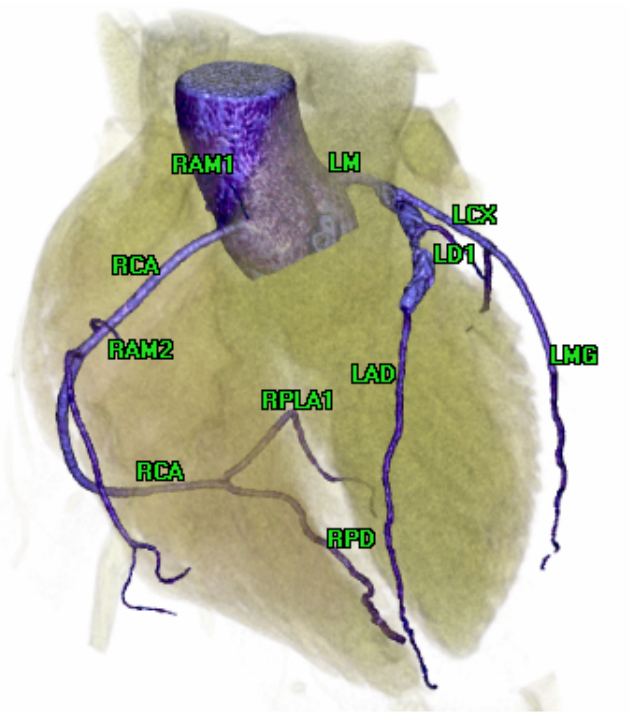

Figure 2.2: Labelled Coronary Arteries
The label assigned to each artery conforms to the AHA 17-segment model[2]. The labels are listed in section 2.2.2

- Section 2.2 introduces and discusses the algorithm in detail, describing all the components of the system used.
- Section 2.3 presents the results of the algorithm, discussing the validation method employed.
- Section 2.4 concludes and proposes further work which could be carried out to address issues relating to accuracy.

2.1.3 State of the Art

This problem has been tackled as a combined image analysis and classification task. Previous work has been carried out in attempts to solve the same problem, with some attempts aimed at different imaging modalities. In [3, 4], [5], [6], [7] and [8] the problem of coronary artery labelling in X-ray angiograms is addressed using graph-matching, while in [9, 10, 11, 12] it is tackled using rule-based approaches exploiting knowledge of the coronary anatomy.

The method introduced in [5] addresses the challenges posed by the presence of noise, artefacts, and competing structures. In this system, a model graph of a generic coronary vasculature is used as a reference, incoming datasets are then matched to this using physical features collected in [13] as extra matching criteria. In this graph, a node represents a segment of

an artery between bifurcations, and the arcs represent parent-child relationships. The method assumes that segments of a vessel between bifurcations can still possess characteristic features of the entire vessel, and hence use vessel length, lumen diameter and branch angle as similarity features in the cost function; although this may not be the case, as the individual segments do not accurately represent the entire vessel especially in terms of length and branch angles.

The method in [8] creates multiple 2-D projections of a 3-D model graph, representing a general topology of the coronary vasculature, with nodes representing bifurcations and arcs representing vessel segments. A novel arterial tree from an X-Ray angiogram is matched against these projected graphs to find the best match based on the inertia axis orientation, the inertia and the eccentricity of the nodes and arcs, after which a finer matching based on number of input/output arcs at each node is performed to assign the labels.

Rule-based approaches, such as in [10] encode coronary artery attributes and inter-segment relationships as binary constraints to eliminate label assignments.

2.1.4 Contribution

The cost functions used in the above methods do not take into account the multivariate nature of the physical features of the vessels, which is why the method proposed in this chapter uses a multivariate Gaussian classifier in order to exploit this property. Furthermore,the graph-matching approaches as used above assume a generic vasculature, but in clinical datasets, anomalies do exist that will lead to inaccuracies in such methods; hence this chapter proposes a less strict model-based approach to the problem.

The method proposed in this chapter takes into account the topological (tree) and geometric structure of the coronary vasculature and encodes such properties as binary priors. Analogous to the graphs used in [5] and [8], the input vasculature is represented as a polyline tree (2.1), which easily describes its spatial and topological properties. It exploits knowledge of the topology to generate a set of plausible labelling candidates, and then exploits knowledge of the spatial properties of each anatomical vessel class to reduce the number of candidate labelled trees. These are then scored using the physical features of each vessel in a trained classifier to find the most likely labelling.

In this chapter, the individual segments of a vessel in-between bifurcations are referred to as topological segments; the concatenation of such contiguous topological segments make up a vessel segment as shown in figure 2.4.

2.2 The Method

2.2.1 Introduction

The input to the algorithm is a CCTA dataset and adjoining automatically segmented coronary arteries, represented as two polyline trees (left and right) (figure 2.4), and the output is a mapping of each polyline to an anatomical label. A system flow diagram is shown in figure 2.3. The components of the algorithm are discussed throughout the rest of this chapter.

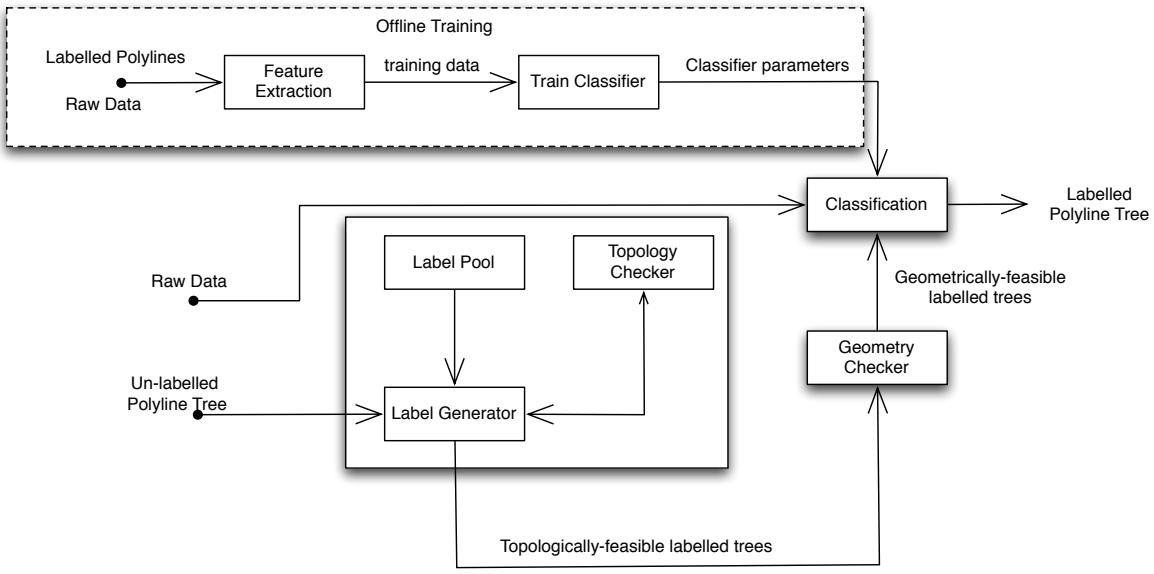

Figure 2.3: Automatic Coronary Artery Labelling
The components of the labelling algorithm and their connections within the system.

2.2.2 Label Pool

The label pool is a list of all the classes (labels) to be considered, i.e.

- LM (Left Main)
- LAD (Left Anterior Descending)
- LCX (Left Circumflex)
- LMG (Left Obtuse Marginal)
- LRI (Left Ramus-Intermedius)
- LD (Left Diagonal)
- LINSIG, RINSIG (All segments not required in specification, left and right prefixes)

- RCA (Right Coronary Artery)
- RPD (Posterior Descending Artery)*
- RPLA (Right Postero-lateral Artery)
- RAM (Right Acute Marginal Artery)

These labels are used to supply the Label Generator.
* At present this is only detected if it is present on the right coronary artery.

2.2.3 Training Data

The training data is collected by manually segmenting individual vessels from CCTA data of 42 patients, using manual vessel-tracking functionality in the VesselMetrix package from Toshiba's Voxar 3D advanced visualisation software to track their centrelines (skeletons). The vessel segments are stored as polylines:

$$polyline = \{x_0, x_1, \cdots, x_{n-1} | x_i \in \Re^3\} \tag{2.1}$$

In this context, a polyline is an ordered list of points along the centreline of a single coronary artery. The entire coronary artery vasculature follows a tree structure and therefore it is represented as a polyline-tree, see figure 2.4.

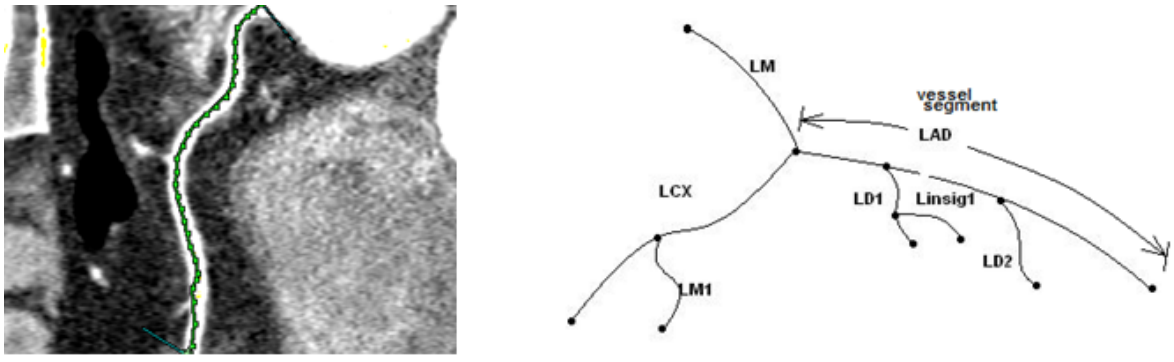

Figure 2.4: Ground-Truth format for the left coronary
The LAD for instance, is stored as a single vessel rather than as three individual segments.The image on the left shows a curved planar reconstruction of a CCTA dataset with the points making up the centreline of a coronary artery.

2.2.4 Feature Extraction

This step takes in labelled coronary vessel segments as polylines, and with the raw data in context, up to 11 features are measured and saved. Tables 2.1 and 2.2 list all the features used.

Feature Name	**Formula**	**Description**
$Length(v)$	$\sum_{i=0}^{N-1} length(t_i)$	where N is the number of topological segments, t making up the vessel v.
LogLengthProportion(v)	$Log(\frac{Length(v)}{\sum_{i=0}^{V-1} Length(v_i)})$	Where there are V vessel segments in the tree.
BranchAngle(v)(degrees)	$180 - cos^{-1}(\frac{\overrightarrow{\boldsymbol{u}}(v_i)}{\overrightarrow{\boldsymbol{u}}(v_{i-1})})$	Where $\overrightarrow{\mathbf{u}}(v_i)$ is the unit vector in the direction of vessel v_i. This is the angle made between a vessel segment and it's parent segment, belonging to another vessel class.
LogTortuosity(v)	$Log\frac{Length(v)}{v(P-1)-v(0)}$	Where P points make up a vessel segment v. This is a measure of curvature of the vessel segment; other measures exist [14].
$Diameter(t)$	$\sqrt{\frac{Volume(t)}{\pi * Length(t)}}$	Approximate diameter of a topological segment, t.
$AvgDiameter(v)$	$\frac{1}{N}\sum_{i=0}^{N-1} Diameter(t_i)$	Approximate average diameter of a vessel segment, v.
$Volume(v)$	$\sum_{i=0}^{N-1} Volume(t_i)$, where $Volume(t_i) = \lvert\{p \in \Re^3 \lvert I(p) \geq T\} \cap D_i\rvert$, where $D_i = \{p \in \Re^3 \lvert \exists k \in t_i, \lvert p-k\rvert \leq r\}$	The measured volume of the cylinder, D created by dilating the topological segment t within the data context using a spherical structuring element with radius, r greater than the maximum radius of a coronary artery, and thresholding at the mean intensity, T along the segment.
LogVolumeProportion(v)	$Log(\frac{Volume(v)}{\sum_{i=0}^{V-1} Volume(v)})$	

Table 2.1: Coronary Artery Features

Feature Name	Formula	Description
DistanceToStart(v)	$\|v(0) - v_0(0)\|$	The Euclidean distance from the root of the vessel tree (ostium) to the start of the vessel.
DistanceToEnd(v)	$\|v(P-1) - v_0(0)\|$	The Euclidean distance from the root of the vessel tree (ostium) to the end of the vessel.
DirectionWRTVesselStart(v)	$\frac{v(P-1)-v(0)}{\|v(P-1)-v(0)\|}$	X, y, and z-co-ordinates of the unit vector running from the start of the vessel to its end.
DirectionWRTOstium(v)	$\frac{v(P-1)-v_0(0)}{\|v(P-1)-v_0(0)\|}$	X, y, and z-co-ordinates of the unit vector running from the ostium to the vessel end.

Table 2.2: Coronary Artery Features Continued

These features, collected from manually labelled coronary vessel segments, are used to train the multivariate Gaussian classifier; in the application, these features when measured on novel vessel segments, form the measured feature vector for each candidate.

2.2.5 Label Generator

The label generator is a recursive function which attempts to assign to the topological segments of the coronary centreline tree all possible label combinations from the label pool. For any function of this sort, the implied computational cost is apparent, especially for trees as seen in figure 2.4, where 12^{11} different trees would be produced for scoring.
For efficiency, only topologically legal candidates (section 2.2.6) are generated. Furthermore, only geometrically plausible labelled candidates (section 2.2.7) are allowed to be scored by the classifier. It works on each level/segment of the tree as follows:

For each vessel class in the label pool

- assign the class to the current tree level
- if the assignment is topologically legal (section 2.2.6)
  - if all levels of the tree have been labelled
    - ◇ if the labelled tree is geometrically legal (section 2.2.7)

        - score configuration using classifier (section 2.2.8)
        - record best scoring labelled tree
    - else
        - Proceed to next level (recursive call)

2.2.6 Topology Checker

The label generator generates a sequence of 'legal' labelled trees. The topology checker is a rule-based system, applying certain rules pertinent to the syntax of the coronary anatomy to generate 'legal' candidates based on the topology of the un-labelled coronary tree. For the left coronary tree, the rules are as follows:

- The root segment of the Left coronary artery is the LM;
- An LM segment can not have any siblings;
- Segments labelled RCA, RPD, RAM, RPLA and RINSIG are not permitted;
- An LAD segment can only be a child of LM or an extension of itself;
- An LCX segment can only be a child of LM or an extension of itself;
- An LD segment can only be a child of LAD or an extension of itself;
- An LMG segment can only be a child of LCX or an extension of itself.
- An LRI segment can only be a child of LM or an extension of itself.
- An LRI segment can only be a child of LM.

For the right coronary tree, the rules are as follows:

- Only RCA, RAM, RPLA, RPD, and RINSIG classes are permitted;
- The root segment must be RCA;
- An RCA segment can only be a child of RCA or an extension of itself;
- An RPD segment can only be a child of RCA or an extension of itself;
- An RAM segment can only be a child of RCA or an extension of itself;
- An RPLA segment can only be a child of RCA or an extension of itself;

Rules common to both sides are as follows:

- Sibling branches can not be in the same class, except for LINSIG, RINSIG and RPLA;
- All RCA, RPD, LAD, LCX and LM segments must be contiguous;
- Any child branch of a LINSIG or RINSIG segment must be LINSIG or RINSIG, respectively.

The above rules are incorporated in a function which is accessed by the label generator and enables it to ignore 'illegal' label assignments straight-away for the sake of optimisation. For the simple 11 segment tree in figure 2.4, and taking into account the number of labels in the label pool; the number of possible label assignments is 12^{11}, but applying the above rules reduces this number to 511! These rules do not need information about the physical structure of the tree, but instead need only to consider a topological representation of the tree and the labels assigned to a single segment and its parent.

2.2.7 Geometric Checker

Vessel geometry in this case refers to the physical and spatial characteristics of the coronary vessel tree. In this algorithm, only the attributes pertaining to the centrelines of the vessels are considered here (i.e. diameter is not considered for instance), and therefore the polyline-tree representation is sufficient. The geometry checker takes in a fully-labelled polyline-tree as input and exploits certain spatial geometric rules unique to individual classes and other rules pertaining to fully-labelled trees, resulting in the elimination of more candidates. An example labelled coronary vasculature is shown in figure 2.5. The rules applied are as follows: For the LAD,

- Looking in the direction going down the vessel, all LD segments must exist on the right side of the vessel, while LINSIG segments (actually called 'Septal') must be on the left.

For the LCX,

- The vessel must traverse a direction posterior to its initial sibling (normally the LAD);
- Looking down the LCX vessel; all LMG vessels must be on the left side, and any vessels coming off the right side must be LINSIG;
- The LCX always starts within an average distance from the root of the left coronary tree.

For the LRI,

- LRI vessels can only be present when the LAD and LCX are present;
- LRI vessels must be in-between the LAD and LCX proximal segments.
- LRI vessels only occur at a trifurcation.

For the RCA,

- RCA segments can only begin within an average distance from the root of the right coronary tree.

For the RPD,

- RPD can only have RPLA or RINSIG as siblings.
- When RPLA segments are present, the RPD segment must be to the right of them, when looking down the RCA.

For the RPLA,

- RCA can never be a sibling.
- If it is an extension of an RPLA segment, then its sibling must be RINSIG.

This step in the algorithm is able to reduce the number of candidate labelled trees in the 11-segment tree (figure 2.4) from 511 to as low as 71 in some cases. The main differences between the geometry checker and the topology checker are as follows:

- Geometric checker takes in a fully-labelled polyline-tree as input, while the topology checker takes in a topological representation as input;
- Geometric checker filters based on spatial rules, while the topology checker filters based on syntactic rules.

2.2.8 Multivariate Gaussian (MVG) Classifier

In this section, *L* refers to a set of all geometrically legal labelled polyline trees, and **Z** is the set of measured feature vectors for each element of *L*.

$$L = \{l_0, l_1, ..., l_{N-1}\}$$

$$\mathbf{Z} = \{\mathbf{z}_0, \mathbf{z}_1, ..., \mathbf{z}_{M-1}\},$$

for M vessel segments in l_i.

Assuming the features measured in section 2.2.4 are normally distributed and follow a multivariate normal distribution, then a MVG classifier[15] may be used to score each l_i accordingly. Using Bayes' theorem (equation 2.2), the posterior probability, i.e. the probability of a vessel being a particular anatomical label ω_k, given its measurement vector $\mathbf{z}$ is given by:

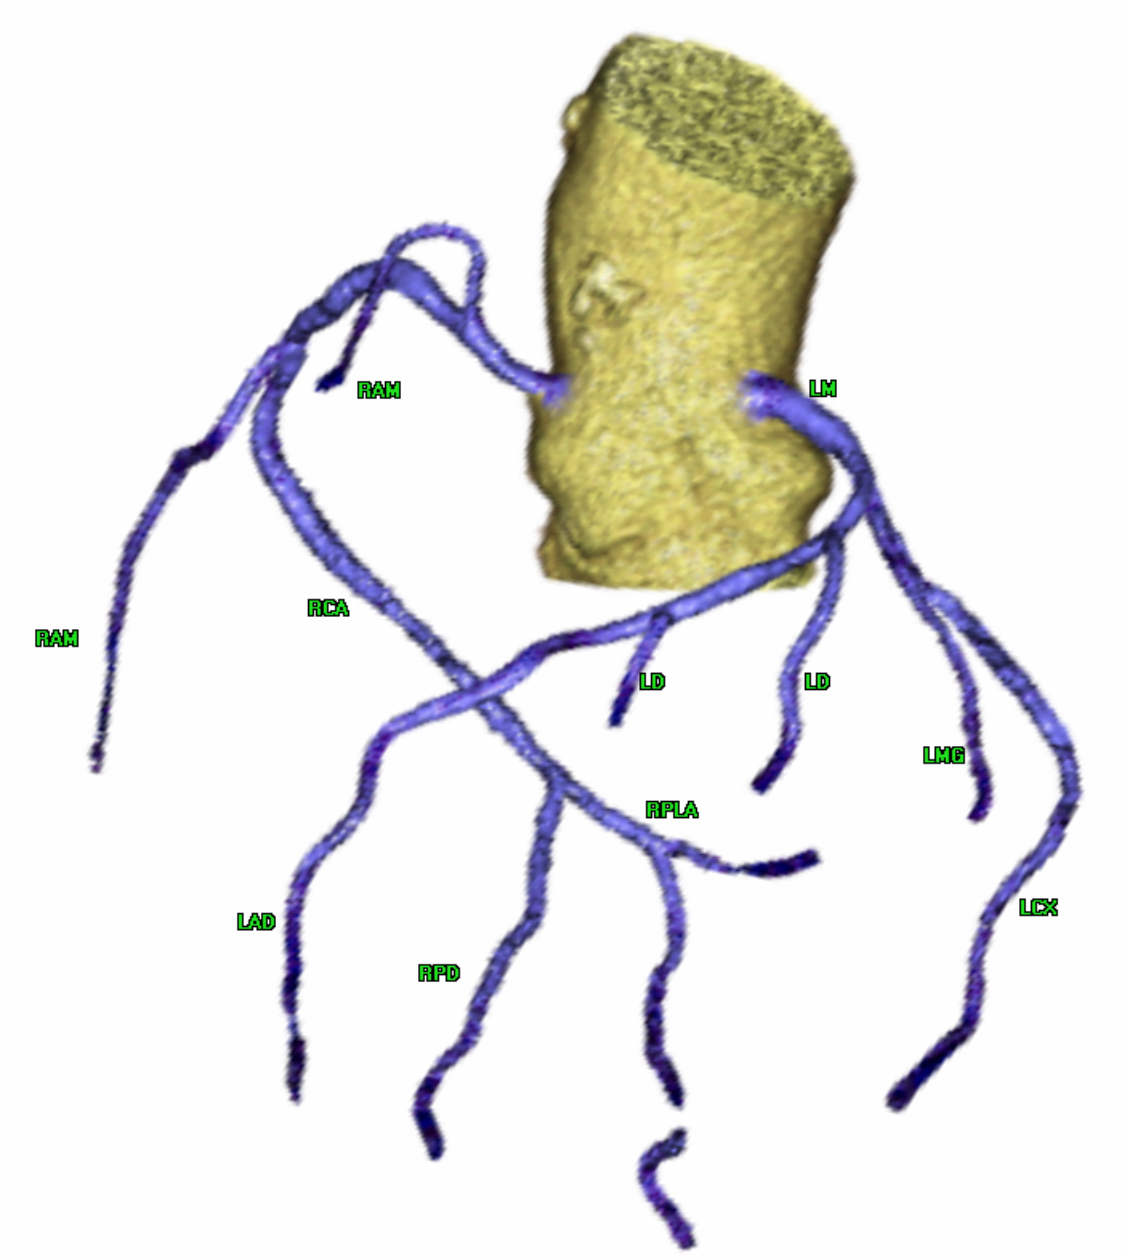

Figure 2.5: Spatial Properties of Vessel Classes
An example labelled coronary vasculature showing spatial properties of each vessel class.

$$P(\omega_k|\mathbf{z}) = \frac{P(\omega_k)P(\mathbf{z}|\omega_k)}{P(\mathbf{z})}, \tag{2.2}$$

where $P(\omega_k)$ is the prior probability of observing the label ω_k. Empirical priors are estimated based on the relative occurrences in the training data, and

$$P(\mathbf{z}|\omega_\mathbf{k}) = \frac{\mathbf{1}}{\sqrt{|\mathbf{C}|(\mathbf{2}\pi)^\mathbf{N}}}\mathbf{exp}(\frac{-(\mathbf{z}-\boldsymbol{\mu}_\mathbf{k})^\mathbf{T}\mathbf{C}^{-\mathbf{1}}(\mathbf{z}-\boldsymbol{\mu}_\mathbf{k})}{\mathbf{2}}), \tag{2.3}$$

where N is the number of features used. The unconditional probability of observing the feature vector, $\mathbf{z}$, does not depend on k by definition, and is therefore constant across all label classes. The posterior probability can then be simplified to:

$$P(\omega_k|\mathbf{z}) \propto P(\omega_k)P(\mathbf{z}|\omega_k),$$

and

$$\sum_k P(\omega_k|\mathbf{z}) = 1$$

The MVG model parameters $\boldsymbol{\mu}_k$ and $\mathbf{C}$, the mean vector for each label class and the pooled covariance matrix respectively are estimated from measurement vectors for each class computed by the feature extractor on training data (see section 2.2.3). The use of a pooled (rather than per class) covariance matrix leads to linear discriminant surfaces, while reducing the number of parameters to be estimated, thereby reducing the required number of training sets to achieve maximum accuracy based on the rule of thumb proposed in [16].

A merit score can then be assigned to a labelled polyline tree containing V vessel segments using equation 2.2, by summing the posterior probability over all segments:

$$\textit{Merit} = \sum_{v=0}^{V-1} P(\omega_k|\mathbf{z}_v) \tag{2.4}$$

The quantity computed in equation 2.4 is biased towards tree configurations with more vessel segments, this can be normalised by instead summing the posterior probability per unit length of each vessel. Equation 2.4 becomes:

$$P(l_i|\mathbf{Z}_i) = \sum_{v=0}^{V-1} P(\omega_k|\mathbf{z}_v) * 10^{LogLengthProportion_v} \tag{2.5}$$

The labelling configuration with the highest merit is chosen:

$$l^* = \underset{l_i}{\operatorname{argmax}}\, P(l_i|\mathbf{Z}_i)$$

2.2.9 Summary

The method described above takes advantage of syntactic approaches to pattern recognition; while exploiting the spatial and physical features of the various classes by using a multivariate Gaussian classifier to score the labelled candidates. The steps comprising the entire algorithm are shown in figure 2.1. In the training phase of the algorithm, the feature extractor measures a list of features of manually-labelled trees; the classifier uses these features to generate mean vectors and covariance matrices for each class. The label generator, along with the topology checker and geometry checker, generates a set of plausible labelled trees; the classifier uses its generated parameters to compute the merit for each labelled candidate and returns the highest scoring.

2.3 Results and Discussion

2.3.1 Introduction

The labelling algorithm described in Chapter 2.2 was tested on CCTA images from several modality vendors. The MVG parameters were estimated using vessel feature data from 42 datasets. This chapter describes the validation metrics used and the results obtained. Several issues were encountered due to anatomical anomalies present in few of the test data; this chapter discusses these issues and describes the impact on labelling performance.

2.3.2 Validation Metrics

Leave-one-out (LOO) validation was carried out on the 42 training datasets by excluding the test dataset during classifier training. The performance of the labelling algorithm was assessed on a topological segment basis. The generated label of each segment was compared against its corresponding ground-truth label and assigned a binary value based on the agreement between the two labels.

The results are presented in a confusion matrix (figures 2.6 and 2.7), therefore the accuracy can be defined as:

$$Accuracy = \frac{\sum_{i=0}^{N-1} \mathbf{C}_{ii}}{\sum_{i=0}^{N-1} \sum_{j=0}^{N-1} \mathbf{C}_{ij}},$$

for N columns in the square matrix $\mathbf{C}$.

Subjective validation was also carried out by a radiologist using previously unseen data.

2.3.3 Ground-Truth

For validation, two sets of ground-truth were used, derived from the 42 datasets.

The first set comprised of manually segmented coronary artery centrelines labelled by medically qualified personnel (as described in section 2.2.3). The second set comprised of output from an in-house automatic coronary centreline tracking algorithm with labels transferred automatically from the first set.

2.3.4 Results

The confusion matrix in figure 2.6 shows the overall accuracy of the algorithm using manually-segmented polyline trees as input and validating against their manually-labelled counterparts.

True Class	Assigned Class													
	LM	LAD	LCX	LRI	LD	LMG	Linsig	RCA	RAM	Rinsig	RPD	RPLA	Total	Accuracy (%)
LM	42	0	0	0	0	0	0	0	0	0	0	0	42	100
LAD	0	111	0	0	1	0	2	0	0	0	0	0	114	97.37
LCX	0	0	55	0	0	1	4	0	0	0	0	0	60	91.67
LRI	0	0	0	5	3	0	0	0	0	0	0	0	8	62.5
LD	0	1	4	6	60	0	4	0	0	0	0	0	75	80
LMG	0	0	0	1	0	15	2	0	0	0	0	0	19	78.95
Linsig	0	2	3	1	7	3	28	0	0	0	0	0	44	63.64
RCA	0	0	0	0	0	0	0	90	0	0	0	1	91	98.9
RAM	0	0	0	0	0	0	0	0	25	12	0	0	37	67.57
Rinsig	0	0	0	0	0	0	0	1	7	34	1	14	57	59.65
RPD	0	0	0	0	0	0	0	0	0	3	13	4	20	65
RPLA	0	0	0	0	0	0	0	0	0	2	3	32	37	86.49
Overall	accuracy: 84.44%													

Figure 2.6: Confusion Matrix: Manually Segmented Input
Confusion matrix showing overall accuracy for 42 datasets. The accuracy is measured per topological segment, i.e. each vessel is split up at bifurcations. The green coloured cells show the accuracy for the major vessels of interest.

It is now clear from this figure, that 111 of 114 LAD segments were correctly classified; hence the accuracy for the LAD segments is 97.3%. It should be noted however, that the accuracy presented here shows the objectively calculated performance of the algorithm. For clinical use, certain misclassifications are tolerated. For instance, misclassifying LINSIG as LD (see (Linsig, LD)) in 7 of the 44 LINSIG segments is permissible because distal LD segments may have been classified as insignificant while the ground-truth was collected; same goes for the LMG segments. The inverse is also permissible for distal LD segments; thus the need for subjective validation. It can also be seen that the RCA segments have been correctly classified in 90 of the 91 cases, with a terminal segment being misclassified as the RPLA.

Figure 2.7 shows the results of the algorithm run using automatically segmented vessels as input. In this case, a label-transfer algorithm is used to automatically transfer labels from the manually-labelled/manually segmented vessels to the automatically segmented ones using the overlapping volume as criteria. This results in more LINSIG and RINSIG segments, as non-overlapping segments were labelled as such. The pink cells show a popular source of error; the LRI segments being misclassified as LD or LMG. This is because the topological rule which states that the LRI has to occur at a trifurcation had to be relaxed for this experiment configuration because the tree-tracking algorithm used is unable to detect trifurcations. Without

True Class	Assigned Class													
	LM	LAD	LCX	LRI	LD	LMG	Linsig	RCA	RAM	Rinsig	RPD	RPLA	Total	Accuracy (%)
LM	41	2	0	0	0	0	1	0	0	0	0	0	44	93.18
LAD	0	88	0	0	1	0	7	0	0	0	0	0	96	91.67
LCX	0	0	32	0	8	2	9	0	0	0	0	0	51	62.75
LRI	0	0	1	2	1	2	1	0	0	0	0	0	7	28.57
LD	0	3	0	1	40	0	2	0	0	0	0	0	46	86.96
LMG	0	0	2	0	2	4	0	0	0	0	0	0	8	50
Linsig	0	6	11	2	7	5	40	0	0	0	0	0	71	56.34
RCA	0	0	0	0	0	0	0	86	0	10	1	2	99	86.87
RAM	0	0	0	0	0	0	0	1	7	11	0	0	19	36.84
Rinsig	0	0	0	0	0	0	0	7	5	92	1	8	113	81.4
RPD	0	0	0	0	0	0	0	0	0	8	5	2	15	33.33
RPLA	0	0	0	0	0	0	0	0	0	7	1	18	26	69.23
Overall	accuracy: 76.47%													

Figure 2.7: Confusion Matrix: Automatically Segmented Input
Confusion matrix showing overall accuracy for the algorithm using automatically segmented vessels as input. The labels were transferred automatically from the manual segmentations to the automatic segmentations and are prone to error.

this constraint, the LRI and the LD are easily confused. The results are summarised in table 2.3.

Vessel Class	Accuracy Mode 1(%)	Accuracy Mode 2(%)
LM	100.0	93.2
LAD	97.4	91.7
LCX	91.7	62.7
LRI	62.5	28.6
LD	80	87.0
LMG	78.9	50.0
LINSIG	63.6	53.3
RCA	98.9	86.7
RAM	67.56	36.8
RPD	65.0	33.3
RPLA	86.5	69.2
RINSIG	59.7	81.4
Overall	84.4	76.5

Table 2.3: Results for both modes of validation
A summary of the per-class accuracy computed from the confusion matrix.

Subjective validation of the algorithm was carried out by a clinically trained person, whose requirement was that the algorithm met functional requirements for clinical use. This lowers the cost of the following misclassifications:

- Distal LAD segments labelled as LINSIG
- LMG segments labelled as LINSIG
- Distal LD segments labelled as LINSIG
- Certain LINSIG segments labelled as LD and LMG

- Certain RINSIG segments labelled as RAM, RPLA, RPD
- RAM segments labelled as RINSIG

2.3.5 Dealing with Anatomical Variation

A particularly interesting case supporting the variability of coronary anatomy has been presented to the labelling algorithm. A congenital anomaly exists in $\approx 1\%$ of human beings [17, 18] whereby the circumflex artery does not emerge from the left main coronary artery (LM), but instead emerges near the root of the right coronary artery and follows a retro-aortic course to supply the postero-lateral left ventricle(LV) (which is always supplied by the circumflex artery). This is demonstrated in figure 2.8.

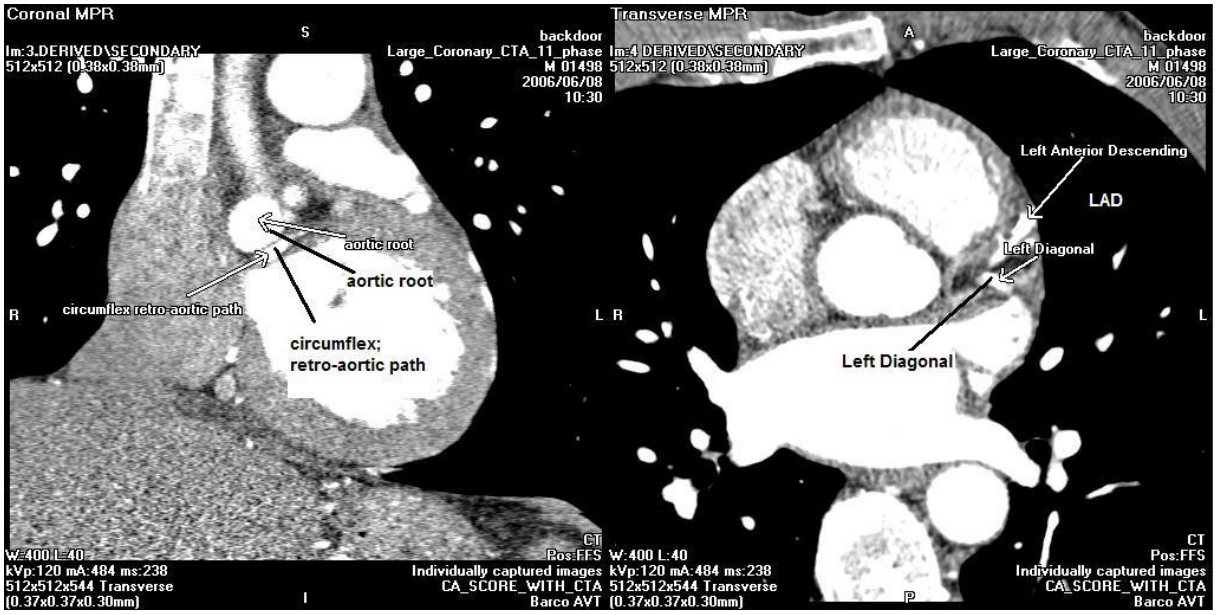

Figure 2.8: Retro-Aortic Circumflex Artery
MPR views showing anomalous case of a retro-aortic circumflex artery. On the left artery (right image), the first artery to come off the LM is the LD instead of the LCX in this case.

In this case, features such as distance from the root, average diameter and branch angle are quite similar for the LD and the LCX. The test data contains two incidents of this retro-aortic circumflex artery, but in both cases the algorithm incorrectly labels the first diagonal (LD) artery emerging from the left main (LM) artery as the left circumflex (LCX) artery.

2.3.6 Summary

The automatic labelling algorithm performs at an average accuracy of 84% when labelling vessels segmented manually; whilst performing at 76.47% when labelling vessels segmented

automatically with the labels transferred from the ground-truth segmentations, but this mode is subject to errors existent in the tree-tracker. In the case of subtle anomalous anatomies, for example, in the presence of a retro-aortic circumflex artery, it is still incapable of accurate classification. It will require heuristic post-processing and more training examples to tackle that particular issue.

2.4 Conclusion and Further Work

2.4.1 Conclusion

A method has been developed for labelling segmented coronary vessels in CCTA images. The approach can be seen as a two-step process: 1) knowledge-based assignment, whereby all plausible labelled trees are generated and 2) statistical classification, whereby the most likely labelling is chosen based on closeness to parameters modelled by a multivariate Gaussian classifier. Due to the paucity of training data compared to the relatively large feature set, a pooled covariance matrix was used as the model to reduce the number of estimated parameters in the classifier.

The method is tested on 42 CCTA datasets of varying coronary anatomy using two different sources as input: 1) manually-segmented vessels 2) automatically-segmented vessels, and yields an overall level of agreement of 84% in case 1 between the labels assigned by the human expert and those assigned by the algorithm. In case 2, the system yields an overall level of agreement of 76% between the labels automatically transferred to the tracked vessels and those assigned by the algorithm.

Further subjective validation is carried out by a radiologist, applying clinical requirements to the assessment, and in this case the algorithm is deemed useful due to high accuracy for proximal segments of the major arteries, provided the user-interface is set-up in such a way that the user is able to validate and change the automatic labelling.

2.4.2 Further Work

Further tasks to be carried out on this algorithm to improve its accuracy are as follows:

- As a supervised classifier relying on multi-variate properties of the features used, the system requires more training data to reduce the possibility of over-fitting and to allow for more accurate covariance estimates.
- Add more vessel classes to reduce variance of LINSIG and RINSIG classes:
  - At present the distributions of measured features of these classes are falsely wide, because various anatomical classes are labelled as such in the training data.

- Implementation of 'Reject' class for label assignments with likelihoods not 'close' to any of the defined classes..
- Train the classifier separately for the left and right sides:
  - This would reduce the number of classes required by the classifier, hence theoretically increasing its accuracy.
- Addition of positioning features to improve the accuracy of the classifier:
  - This could potentially reduce the misclassification of LD and LCX especially in the anomalous cases as vessel path is a strong discriminator for these two vessels.
- Normalisation of features to a common space:
  - The polyline trees (both training and novel) should be warped to the space of a reference dataset, before measuring feature vectors.

References

[1] *The Global Burden of Disease*. World Health Organisation, 2008.

[2] U. Hoffmann, A. Pena, R. Cury, S. Abbara, and M, "Cardiac ct in emergency department patients with acute chest pain," *Radiographics*, pp. 963–979, 2006.

[3] a. C. Dumay, J. J. Gerbrands, and J. H. Reiber, "Automated extraction, labelling and analysis of the coronary vasculature from arteriograms.," *International journal of cardiac imaging*, vol. 10, pp. 205–15, September 1994.

[4] A. C. M. Dumay, R. J. V. D. Geest, J. J. Gerbrands, E. Jansen, and J. H. C. Reiber, "Consistent inexact graph matching applied to labelling coronary segments in arteriograms," *Radiology*, 1992.

[5] N. Ezquerra, S. Capell, L. Klein, and P. Duijves, "Model-guided labeling of coronary structure.," in *Medical Imaging, IEEE Transactions on,*, vol. 17 of *3*, pp. 429–441, 1998.

[6] K. Haris, N. Efstratiadis, S.N.and Maglaveras, C. Pappas, J. Gourassas, and G. Louridas, "Model-based morphological segmentation and labelling of coronary angiograms.," in *IEEE Transactions on medical imaging*, vol. 18 of *10*, pp. 1003 – 1015, October 1999.

[7] K. Haris, S. Efstratiadis, N. Maglaveras, C. Pappas, J. Gourassas, and G. Louridas, "Automatic labelling of the coronary tree using a three dimensional reference prior model," in *Medical Imaging, IEEE Transactions on*, vol. 18, pp. 1003 – 1015, 1999.

[8] C. C. I. M. G. Finet, "Automatic labelling of the coronary tree using a three dimensional reference prior model.," *Computers in Cardiology*, vol. 5, 1998.

[9] M. Garreau, J. L. Coatrieux, R. Collorec, and C. Chardenon, "A knowledge-based approach for 3-d reconstruction and labeling of vascular networks from biplane angiographic projections.," *IEEE transactions on medical imaging*, vol. 10, pp. 122–31, January 1991.

[10] C. Smets, F. Van de Werf, P. Suetens, and a. Oosterlinck, "An expert system for the labeling and 3d reconstruction of the coronary arteries from two projections.," *International journal of cardiac imaging*, vol. 5, pp. 145–54, January 1990.

[11] S. A. Stansfield, "A Rule-Based Expert System for Automatic Segmentation of Coronary Vessels From Digital Subtracted Angiograms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, pp. 188–199, 1986.

[12] S. Tsuji and H. Nakano, "Knowledge-based identification of artery branches in cineangiograms," in *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, vol. 2, p. 710Ű715, 1981.

[13] J. D. Jr, B. Brown, E. Bolson, and H. Dodge, "Lumen diameter of normal human coronary arteries: Influence of age, sex, anatomic variation, and left ventricular hypertrophy or dilation," in *Circulation, American Heart Association*, vol. 86, pp. 232–246, 1992.

[14] E. C.U., G. R., and N. D.L., "A comparison of quantitative measures of arterial tortuosity using sine wave simulations and 3d wire models.," *Physics in Medicine and Biology*, vol. 45, pp. 2593–2599(7), 2000.

[15] F. van der Heijden; R. P. W. Duin; D. De Ridder; D. M. J. Tax, *Classification, Parameter Estimation and State Estimation: An Engineering Approach using MATLAB*. Wiley, 2005.

[16] J. Piper, I. Poole, and a. Carothers, "Stein's paradox and improved quadratic discrimination of real and simulated data by covariance weighting," *Proceedings of the 12th IAPR International Conference on Pattern Recognition (Cat. No.94CH3440-5)*, pp. 529–532, 1994.

[17] N. Kobayashi, T. Takebayashi, S. Saito, T. Terashita, T. Shimokawa, and S. Matsuda, "Coincidence of a right retroaortic circumflex coronary artery and a right-sided aortic arch with a retroesophageal course of the left subclavian artery (arteria lusoria)," *Clinical Anatomy*, vol. 19, pp. 354–357, August 2005.

[18] P. Koşar, E. Ergun, C. Oztürk, and U. Koşar, "Anatomic variations and anomalies of the coronary arteries: 64-slice CT angiographic appearance.," *Diagnostic and interventional radiology (Ankara, Turkey)*, vol. 15, pp. 275–83, December 2009.

Chapter 3

Creation of a Coronary Artery Atlas to Guide the Tracking of Coronary Vessels

Abstract *Coronary anatomy varies from person to person, although general patterns of this anatomy co-exist alongside anomalous patterns among individuals. This paper proposes a method to represent these anatomical patterns using a coronary artery centreline atlas, and demonstrates how this atlas can be applied to an existing coronary centreline tracking algorithm to reduce the occurrence of non-arterial vessels from 8.4% to 5.0%, thus improving the specificity of the algorithm.*

The atlas is created by registering a set of CCTA datasets, each with manually-segmented coronary artery centrelines to a reference CCTA dataset, with a manually segmented heart, and warping these centrelines onto the reference heart. Statistical features relating to density and direction are then extracted from this atlas and used through a coronary vessel classifier to assess the probability of a tracked centreline being that of a true coronary artery centreline. The unlikely candidates are pruned leaving only true coronary arteries.

3.1 Introduction

Close to 17 million people worldwide die annually from heart disease; it is essential to have the ability to detect and diagnose the symptoms of these diseases as early and as accurately as possible. CT angiography is becoming an increasingly important diagnostic tool for this purpose.

Automatic identification of coronary arteries from medical images provides a quick and efficient way for cardiologists to assess the risk of coronary artery disease in patients; although these segmentation algorithms are subject to error and may produce additional incorrect vessels, which could end up being misleading.

3.1.1 Motivation

Automatically delineating the coronary arteries and presenting them to the radiologist, for example by curved MPR, allows them to make measurements of stenoses within the arteries and assess the risk of CAD (figure3.1) quicker than it would take if they had to manually segment the arteries [1, 2, 3]. Accurate segmentation of the coronary arteries, therefore can reduce the time taken to assess the risk of CAD.

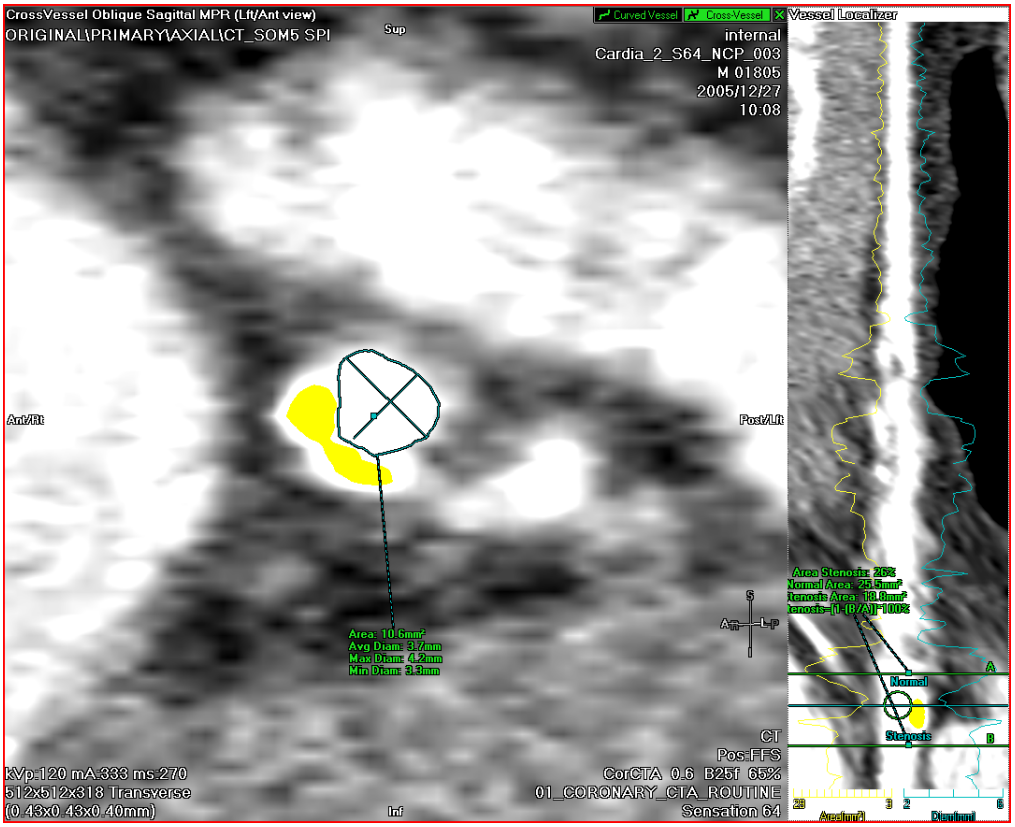

Figure 3.1: Automatic Coronary Analysis
The image on the left shows a cross-sectional view of a stenosed section of a segmented coronary artery from a CCTA dataset. In the straightened vessel view on the right, the artery is automatically analysed for stenosis.

A coronary centreline tracking algorithm is required in order to provide this functionality in automatic cardiac analysis software. Such algorithms, are qualified by the level of accuracy with which they segment the coronary arteries. This accuracy is defined in terms of sensitivity and specificity. Sensitivity, in this case is defined as the proportion of coronary arteries that are correctly identified, while specificity refers to the proportion of non-arterial vessels that are correctly ignored (Equation 3.2). Figure 3.2 shows two cases where the specificity of the applied centreline tracking algorithm is reduced.

$$sensitivity = \frac{\text{Volume of correctly identified arterial vessels}}{\text{Total volume of segmented vessels}} \tag{3.1}$$

$$specificity = \frac{\text{Volume of correctly identified non-arterial vessels}}{\text{Total volume of non-arterial vessels}} \tag{3.2}$$

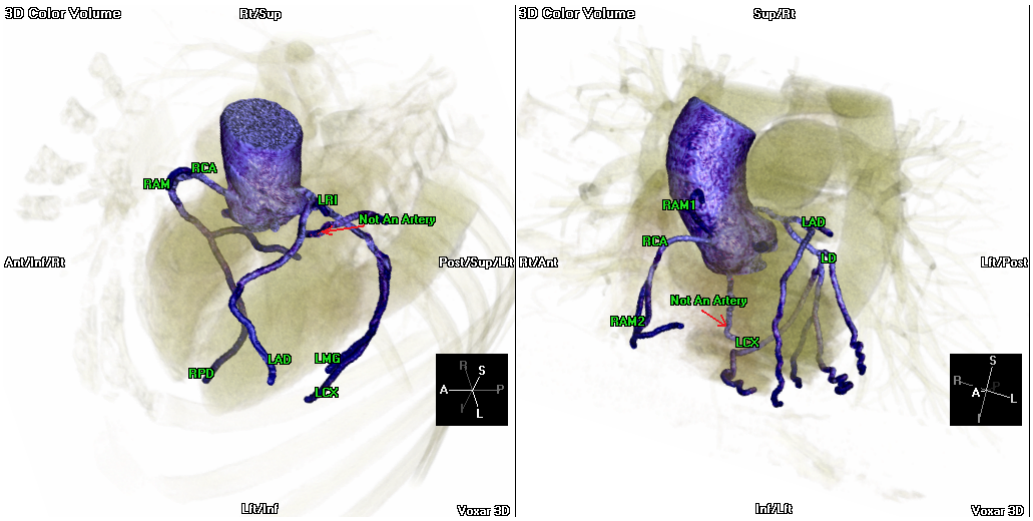

Figure 3.2: Specificity Errors in Coronary Artery Segmentation
The Volume Rendered (VR) image on the left shows an error whereby a coronary vein is identified as a branch of the left anterior descending artery. In the image on the right, a section of the right ventricle has been identified as part of the right coronary tree.

The aim of this work is to demonstrate that knowledge of coronary anatomy applied through an atlas can improve the specificity of automatic coronary artery centreline tracking algorithms by identifying non-arterial vessels in the tracked output. This report discusses a method for the creation of a coronary artery centreline atlas and its application to a system for tracking coronary artery centrelines from cardiac angiograms. The chapters are arranged as follows:

- Section 3.1.2 reveals the state of the art in vessel atlas creation.
- Section 3.2 presents the proposed method, module by module.
- Section 3.3 presents results, describing the validation method adopted.
- Section 3.4 concludes by discussing further work proposed.

3.1.2 State of The Art

Segmentation of coronary artery centrelines from medical imaging modalities falls under the general problem of vessel segmentation. This problem has been adequately categorised by

[4] into pattern recognition, model-based, tracking-based, artificial intelligence-based, neural network-based, and miscellaneous tube-like object detection approaches.

The problem of creating a vessel atlas to represent anatomical *a priori* knowledge for application to vessel segmentation algorithms has been addressed in [5], [6], [7], and [8] as a process requiring registration of mutliple datasets to a reference dataset. In [5] and [6], a method whereby a vascular image volume is picked as a template and other segmented vascular images are mapped onto it using a vascular-model to image affine registration method is proposed. Here, the model is represented as a set of points in the centreline of the vessels with their radii and medialness [9] and it is used to detect anomalous and pathogenic blood vessels in liver and brain. The use of affine transformation as the only form of registration for elongated structures such as vessels, which vary widely in tortuosity, can lead to sub-optimal mappings. In [7] and [8], an atlas of the brain vasculature, which contains information about vessel size, orientation of individual vessels, and the probability of finding vascular structures at a given position is created from Magnetic Resonance Angiograms (MRA). The atlas is created using rigid and non-rigid registration on automatically segmented vascular image datasets, using one dataset as a reference whilst registering the others to it. The method used to construct the atlas relies on automatically segmented vessels which in turn exposes the atlas to errors from the segmentation such as false branches, which should be detected by the vascular atlas in the first place.

3.1.3 Contribution

This report is not concerned with vessel tracking algorithms per se, rather it concerns the construction and application of a vessel atlas to improve the specificity of such algorithms.

The method proposed in this paper builds on [7], but instead uses manually delineated coronary artery centrelines provided by medical trained personnel; this ensures that the entire centreline database contains only coronary artery centrelines, providing a suitable basis for building a model. The method selects a reference CCTA dataset and warps the remaining centrelines onto it following registration of the raw volume-pairs to consider the entire context. The probability atlas described in [7] tends to produce sparse data from limited training sets, therefore this paper proposes a probability density atlas created by a union of the transformed centreline-trees with smoothing performed using an adaptive anisotropic kernel density estimation method in order to maintain the shape of the vasculature. The orientation atlas provided in [7] is replaced with the direction of centreline points in each vessel represented as unit vectors to provide more useful information.

3.2 Method

3.2.1 Introduction

This chapter describes the implementation of atlas-guided tracking of coronary artery centrelines. A flow diagram of the overall method is provided in figure 3.3, and the rest of the chapter discusses in greater detail the components of this diagram.

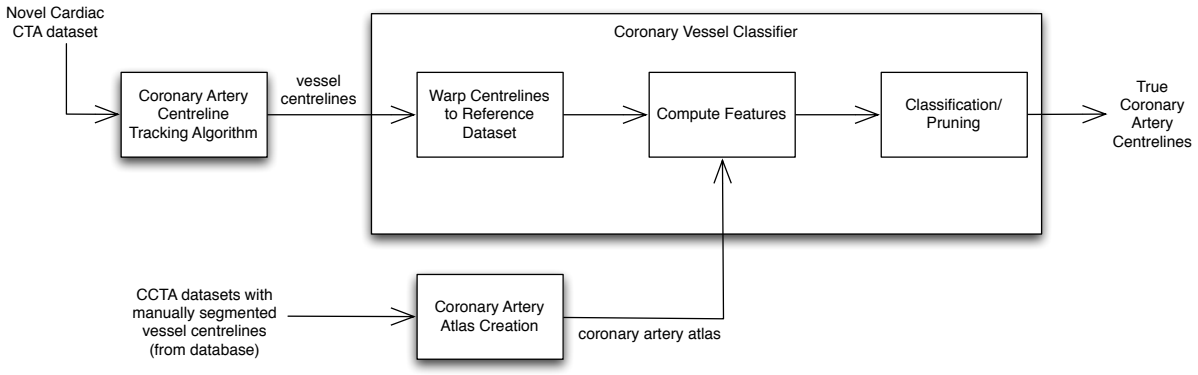

Figure 3.3: Atlas-Guided Coronary Artery Centreline Tracking
The input is a CCTA image, and the final output are the artery centrelines having fewer non arterial vessels.

3.2.2 Training Data

The training data consists of CCTA image datasets collected from 42 patients. In each of these datasets, the coronary artery centrelines are manually segmented by medically-trained personnel using commercially available advanced visualisation software. Each vessel is stored as a polyline (see equation 3.3), with the entire arterial tree represented by a polyline-tree.

$$polyline = \{x_0, x_1, \cdots, x_{n-1} | x_i \in \Re^3\} \tag{3.3}$$

In this context, a polyline is an ordered list of points along the centreline of a single coronary artery. The entire coronary artery vasculature is represented as a polyline-tree, see figure 3.4.

3.2.3 Coronary Artery Atlas Creation

The intent of the coronary artery atlas is to provide statistical information about the arterial vasculatures present in the training data, and in the presence of a large amount of training data, information about the vasculature of a specific population. Its construction takes place in two phases; the first phase involves warping all the training data onto a reference space, while the

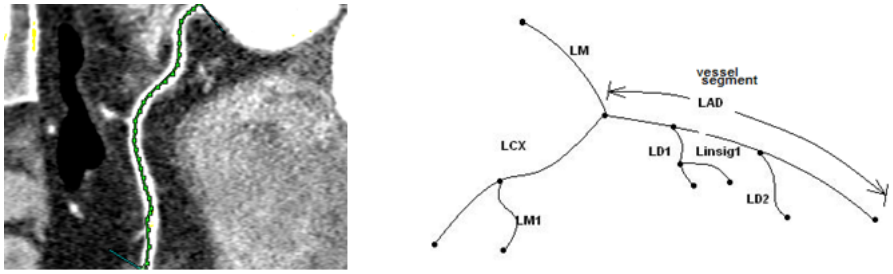

Figure 3.4: Training Data
The image on the left shows a curved planar reconstruction of a CCTA dataset with the points making the centreline of a coronary artery. The sketch on the right is an example of an entire artery vasculature represented as a polyline-tree.

second phase computes quantitative information about the atlas, which can be used to classify novel vessels.

Representing the Training Data in a Reference Space

From a database, D, of CCTA datasets containing manually segmented coronary arteries, a dataset is arbitrarily selected as a reference d_{ref}.

$$D = \{\{d_0, \{p_0\}\}, \cdots, \{d_{N-1}, \{p_{N-1}\}\}\},$$

where d_i is volume data, $d_i \in \Re^3 \rightarrow \Re$, $\{p_i\}$ is a polyline tree, and $\{p_i\} \subset d_i$.

Each of the remaining $N-1$ datasets are registered with d_{ref}, using a multi-scale affine registration algorithm τ that produces the best matching between the two volumes based on 3-D translation, anisotropic scaling and rotation.

$$\tau : (\Re^3 \rightarrow \Re) \times (\Re^3 \rightarrow \Re) \rightarrow (\Re^3 \rightarrow \Re^3)$$

$$\tau(d_i, d_{ref}) = \mathbf{T}_i$$

In the above definition, $\mathbf{T}$ is an affine transformation matrix representing the best linear transformation required to map d_i onto d_{ref}.

$$\mathbf{T} = [t_{i,j}]_{3\times 4}$$

$$\forall x \in d_i(x), \quad d_i(x)(\mathbf{T}x) \approx d_{ref}(x)$$

The affinely transformed $N-1$ datasets are non-linearly registered with d_{ref} using multi-scale non-rigid registration, which aims to find the warp field which maximises mutual information within fluid and elastic constraints[10], to take account for changes in shape across the different datasets. This produces a warp-field W_i mapping the dataset d_i to d_{ref}, with the initial warp-field

derived from $d_i(x)(\mathbf{T}x)$.

$$\forall x \in d_i(x), \quad d_i(x)W_i(x) \approx d_{ref}(x)$$

The segmented polyline trees, $\{p_i\}$ are warped onto the space of d_{ref} using W_i.

The above procedure (figure 3.5) produces N-1 coronary trees mapped into the space of one heart, and this is referred to as the coronary artery atlas (figure 3.6).

$$A = \bigcup_{i=1}^{N-1} \{p_i\}(W_i) \tag{3.4}$$

Extracting Statistical Features from the Atlas

The images in figure 3.6 show the coronary artery atlas, which can be perceived as a source of anatomical information about the coronary arteries present in the training set. In order to exploit this information, statistically descriptive features need to be defined from the atlas. In [7], the density, orientation, and vessel thickness are used to describe vessels in the brain; in our method, the probability density and direction are used to describe the coronary arteries.

Probability Density The probability density of the atlas can be visualised as a 3D image, with values being the probability of finding a coronary artery centreline at any point within it. In order to compute a probability density function (PDF) from the coronary atlas, each centreline point is regarded as a random variable, and kernel density estimation is used to approximate the distribution of this random variable. The conventional Parzen estimator proposes that the value of the p.d.f at a point x is the average of all kernels centred at each training sample:

$$\hat{f}(x) = \frac{1}{N} \sum_{i=1}^{N} K_h(x, x_i), \quad x \in \Re^3, \tag{3.5}$$

where $K_h(x, x_i)$ is an isotropic, usually Gaussian kernel [11], with fixed width h.

$$K_h(x, x') = (\frac{1}{2\pi h^2})^{\frac{1}{2}} exp(-\frac{\|x - x'\|^2}{2h^2}) \tag{3.6}$$

The Parzen method using an isotropic kernel implies that the distribution will be spread equally in all 3 dimensions (x, y, z) at every training point; this is not desirable in the case of coronary anatomy, as the coronary vasculature is known to have a curved surface, therefore isotropic kernel density estimation will lead to inaccurate probabilities in directions normal to this surface. In order to avoid this situation, the Parzen estimator is used with an adaptive

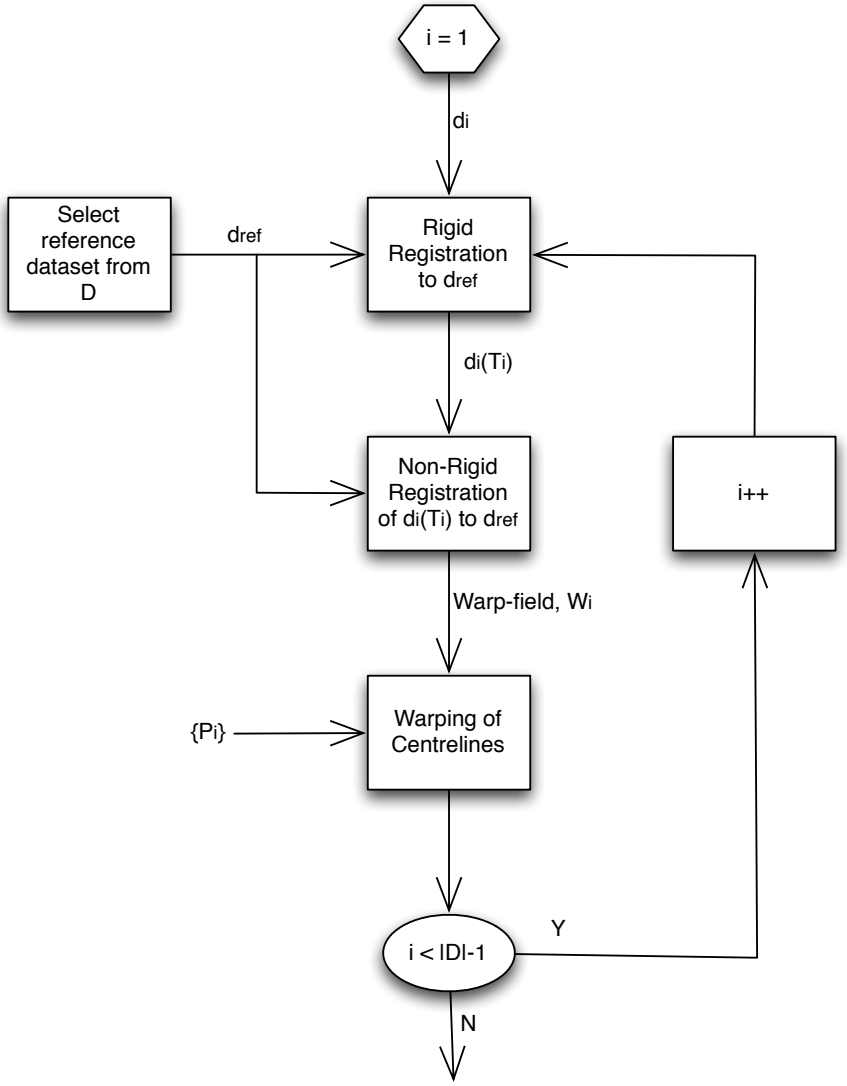

Figure 3.5: Artery Atlas Creation

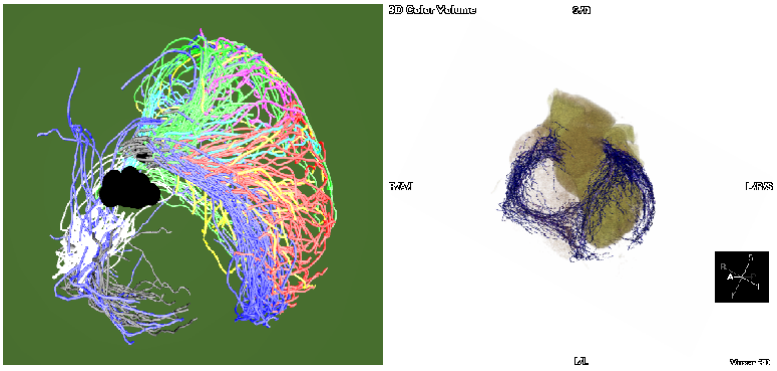

Figure 3.6: The Coronary Artery Atlas
The figure on the left shows N coronary trees registered to the space of a heart. The black dots indicate the centre of gravity of the combined (left and right) vasculature for each pair of coronary trees, while the different colours indicate the anatomical label of each vessel. The image on the right shows a volume-rendering of the registered centrelines overlayed on the contrast-enhanced blood in the reference heart.

anisotropic kernel as proposed in [11].

$$\hat{f}(x) = \frac{1}{N}\sum_{i=1}^{N} K_i(x, x_i) \tag{3.7}$$

In equation 3.7, $K_i(x, x_i)$ is the adaptive anisotropic kernel centred at a training sample point, x_i and is computed from the distribution, Σ_i at the point x_i.

$$K_i(x, x_i) = \frac{1}{|2\pi\Sigma_i|^{\frac{1}{2}}} exp(-\frac{1}{2}(x - x_i)^T \Sigma_i^{-1}(x - x_i)) \tag{3.8}$$

$$\Sigma_i = \alpha\mathbf{1} + \sum_{j=1}^{N} K_h(x_i, x_j)(x_i - x_j)(x_i - x_j)^T, \tag{3.9}$$

where α is a regularisation parameter that controls the magnitude of the probability in directions orthogonal to the principal direction of the distribution (see figure 3.7). It is not necessary to compute the distribution Σ_i exactly as described above, which considers the influence of all the training points, as this is not only computationally intensive, but may fail to capture regional shape characteristics, therefore we use a local neighbourhood approach.

S_i is a spherical neighbourhood of fixed radius, r centred at sample point x_i:

$S_i = \{x||x - x_i| \leq r, x \in \textit{Atlas}\}$.

The distribution is now calculated within S_i as:

$$\Sigma_i = \alpha\mathbf{1} + \sum_{x \in S_i} K_h(x_i, x)(x_i - x)(x_i - x)^T, \tag{3.10}$$

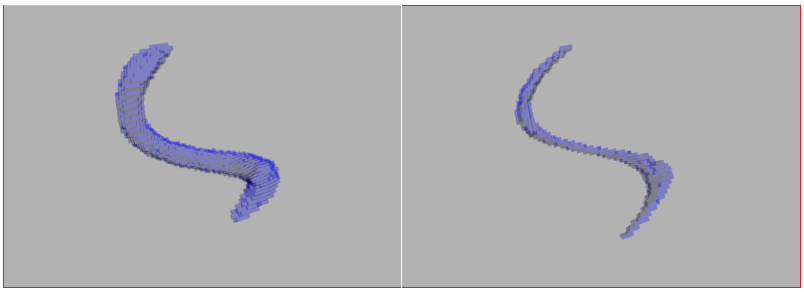

Figure 3.7: Adaptive Anisotropic Kernel Density Estimation
The example images are generated by applying the adaptive anisotropic kernel to two helical lines parallel to each other along their entire lengths. The image on the left shows α set at 1.0, while the image on the right shows α set at 0.02.

Figure 3.8 shows the created density atlas.

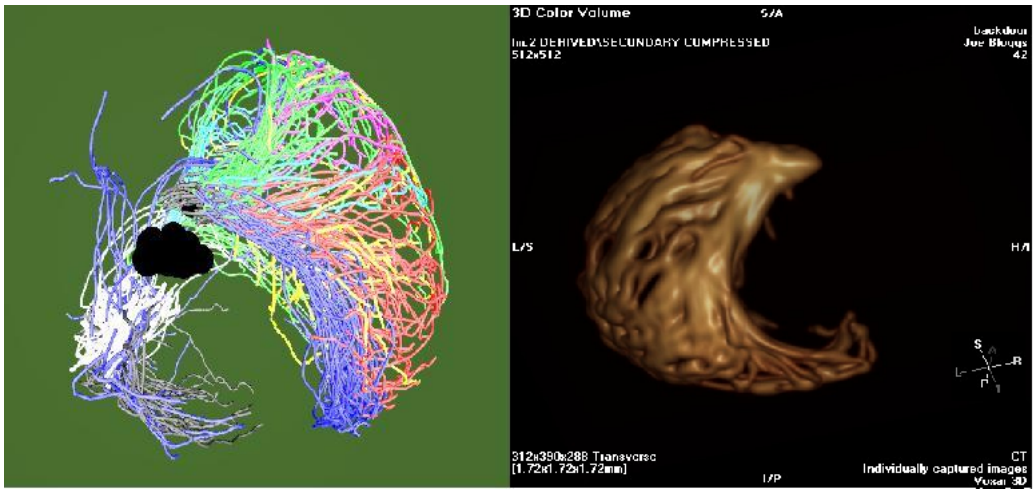

Figure 3.8: Probability Density and Directional Features of the Atlas
The image on the right represents the probability of a coronary artery being at any position within the atlas space, while the one on the left shows the direction of each artery at each point within the space.

Direction The directional information can be visualised as a 3D vector image representing the directions of each point within each coronary artery centreline contained in the atlas. The direction is the unit vector at point x_i in the direction of point x_{i+1} in the centreline of an artery, p_i. The direction at the last point is computed by extrapolation.

$$\hat{\mathbf{x}}_i = \frac{x_{i+1} - x_i}{|x_{i+1} - x_i|} \tag{3.11}$$

These two feature images contain the relevant information that can be used to identify a coronary vessel as an artery (see figure 3.8).

3.2.4 Coronary Vessel Classifier

The computed features discussed in the previous section are used to classify candidate centrelines representing vessel segments. In order to do this, the output polyline tree from the coronary artery centreline tracking algorithm is warped onto d_{ref}, using the same affine and non-rigid registration strategy described in section 3.2.3.
Two quantities are computed in this case:

- Density-based probability of a vessel-centreline (PDV).
- Congruence of the direction of a vessel centreline and the directional atlas (CDV).

Density-based probability of a vessel-centreline

This is computed from the PDF, (equation 3.7), using a neighbourhood strategy.

$$p(x_i \in R_i) \approx \sum_{x \in R_i} \hat{f}(x), \tag{3.12}$$

where R_i is a spherical neighbourhood of fixed radius, r centred at x_i, such that

$$R_i = \{x || x - x_i| \leq r, x \in \mathbb{Z}^3\} \tag{3.13}$$

The probability of a candidate polyline being a coronary artery segment, i.e. belonging to the atlas can be described as:

$$PDV = \frac{1}{N} \sum_{i=0}^{N-1} p(x_i \in R_i), \tag{3.14}$$

for N points in the polyline.

Directional Congruence

This is computed from the directional information (equation 3.11), using a similar neighbourhood strategy.

$$CDV = \frac{1}{N} \sum_{i=0}^{N-1} \sum_{x \in R_i} \hat{\mathbf{x}}_i . \hat{\mathbf{x}}, \tag{3.15}$$

where $R_i = \{x || x - x_i| \leq r, x \in \textit{Atlas}\}$.

This gives a measure of agreement with the neighbouring atlas points based on direction. A $CDV < 0$ implies that the candidate polyline is anti-parallel to the arterial centrelines in a local neighbourhood within the atlas. Such polylines would naturally be classified as non-arterial

vessels; therefore the dot-product is clipped at 0, in order to make the CDV lie between 0 and 1, representing a pseudo-probability.

Classification

The probability of a vessel being a true coronary artery can be obtained in a number of ways based on each of the measured features discussed in the previous section. The simplest approach is to use the computed PDV and CDV values directly, as they scale between 0 and 1. PDV represents the probability of a candidate vessel being part of a coronary artery, given its spatial co-ordinates within the atlas, while CDV is treated as the probability of a candidate vessel being part of a coronary artery given its average direction within the atlas.

In order to detect the non-arterial vessel-segments produced by the vessel tracking algorithm, thresholding based on the chosen quantity is used. The output of the tracking algorithm is warped onto the reference space prior to classification. Pruning is performed recursively by removing child nodes of the polyline tree with the chosen quantity (*PDV* or *CDV*) less than a trained threshold K.

The output is a polyline tree containing all centrelines that have a CDV or PDV above K, indicating that they are true coronary artery centrelines based on the atlas.

3.2.5 Summary

This chapter has described the method used to detect and remove non-arterial vessels from automatically tracked coronary vessels using a statistical coronary artery atlas. The process of atlas creation uses affine and non-rigid registration to warp a set of manually segmented coronary artery centreline-trees to a reference; the candidate vessels are also warped to this reference and classified as being arterial or non-arterial coronary vessels.

3.3 Results and Discussion

3.3.1 Introduction

This chapter discusses the results obtained from the method described in chapter 3.2. Various approaches to coronary vessel classification are proposed in section 3.3.4, although none have been validated at this point, and so only the results from the applied methods are discussed.

3.3.2 Validation Method

20 CCTA datasets with manually segmented coronary artery centrelines are selected to be used as validation data. The ground-truth obtained from these datasets are collected by medically-trained

personnel in a manner such that all coronary arteries (including small arteries) present in the dataset, are included. This allows for accurate assessment of the proposed method in dealing with non-arterial branches.
The metrics used for validation of the coronary artery tracking algorithm are:

- False Positive Ratio (FPR): This is the volume ratio of false vessels to ground-truth vessels. False vessels in this case refer to vessels that are included in the generated output, but absent in the ground-truth.

$$FPR = \frac{|\oplus_d \{T'_{GEN}\} \setminus \oplus_d \{T'_{GT}\}|}{\oplus_d \{T'_{GT}\}},$$

  where T' is a connected binary image representing points in the polyline tree, and d is the size of a spherical structuring element used for dilation. False positive ratio is inversely proportional to the specificity of the overall system.

- False Negative Ratio (FNR): This is the volume ratio of missing vessels to ground-truth vessels. Missing vessels in this case are the vessels which are not present in the generated output, but present in ground-truth. this measure is inversely proportional to the sensitivity of the overall system.

$$FNR = \frac{|\oplus_d \{T'_{GT}\} \setminus \oplus_d \{T'_{GEN}\}|}{|\oplus_d \{T'_{GT}\}|}$$

For the purpose of this work, false vessels, in the context of FPR, are treated as non-arterial vessels, therefore the above metrics are still meaningful.

3.3.3 Results

The aim of any classification algorithm employed here is to reduce FPR as much as possible, leaving FNR unchanged.

Results using Measured Quantities Directly

The CDV and PDV (equations 3.15 and 3.14, respectively) computed in chapter 3.2 can be used directly as probabilities of a candidate vessel centreline belonging to a coronary artery as they both represent relationships between candidate centrelines and the coronary artery centrelines via the atlas. In this approach, a threshold is selected based on Leave One Out (LOO) experiments on the 20 datasets selected for validation. Receiver Operating Characteristic Curves based on FPR and FNR are shown in figure 3.9.

It can be seen from figure 3.9 that applying CDV directly is more radical at detecting false vessels, whilst still maintaining the true vessels. From figure 3.9 , applying a threshold of 0.1

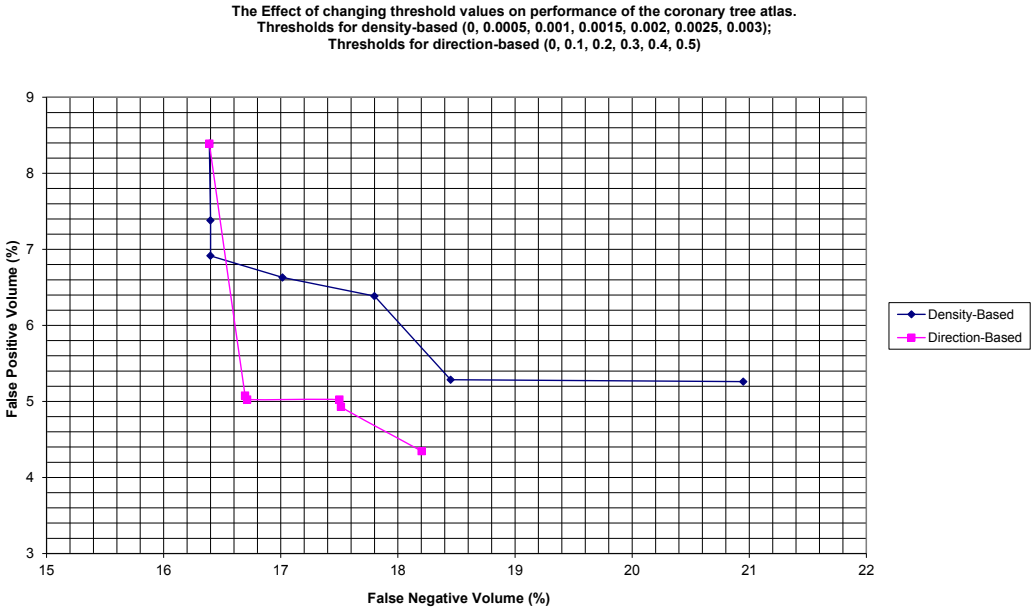

Figure 3.9: Receiver Operating Characteristic using CDV and PDV directly.
The thresholds are increased as the curves move from left to right; the left-uppermost point indicating the results when no pruning is carried out.

reduces reduces FPR from 8.39% to 5.02%, while increasing FNR from 16.5% to 16.67%. Thresholding based on the PDV however, produces a reduction in FPR from 8.39% to 6.92%, while increasing FNR from 16.39% to 16.4%.

3.3.4 Discussion

The classification method employed in section 3.2.4 could be regarded as a trivial approach out of many possible classification techniques using the data available.

The histograms in figures 3.10 and 3.11 indicate that the training data is almost normally distributed based on CDV and PDV. Assuming a normal distribution in both cases, the probability of a candidate coronary vessel centreline being part of a coronary artery can be computed given its CDV or PDV, z, and its normalised distance to the mean, μ, of any of these quantities measured on training data.

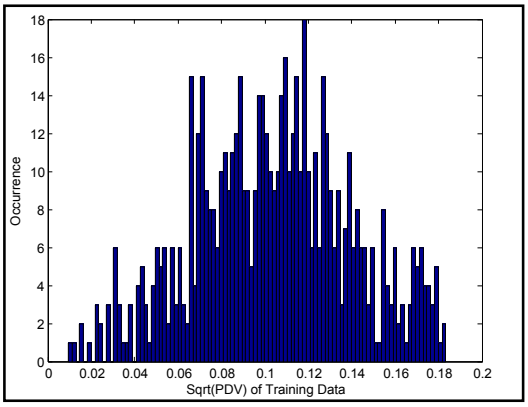

Figure 3.10: PDV Distribution
Distribution of Probability Density of a coronary vessel. This is computed on the vessel centrelines of the coronary vessel atlas.

$$P(z|\omega_k) = \frac{1}{\sigma\sqrt{2\pi}} exp(-\frac{(z-\mu_k)^2}{2\sigma_k^2}) \tag{3.16}$$

In the above equation, the quantity $P(z|\omega_k)$, also known as likelihood [12] refers to the probability of observing the measured feature z, given a vessel of class ω_k (coronary artery, in this case). The quantity required here is the posterior probability; i.e. the probability of a vessel being a true coronary artery ω_k, given a measured quantity z, and this is computed using Bayes' theorem:

$$P(\omega_k|z) = \frac{P(\omega_k)P(z|\omega_k)}{P(z)} \tag{3.17}$$

The prior probability of observing an arterial vessel, $P(\omega_k)$ is rather difficult to estimate accurately, even if a record of non-arterial vessels generated by the tracker is available. Hence,

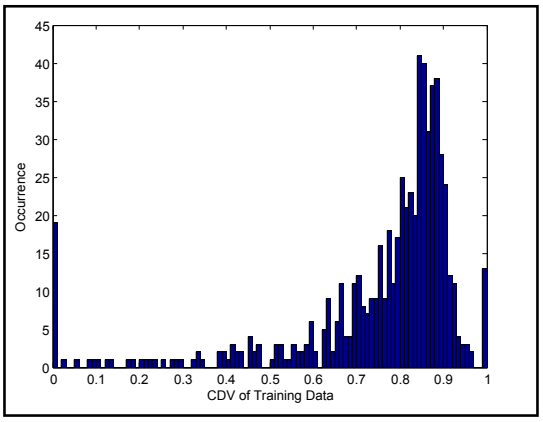

Figure 3.11: CDV Distribution
Distribution of Directional Congruency of the points within each vessel. This figure is computed for the vessel centrelines in the coronary vessel atlas.

if equal priors are assumed for both arterial and non-arterial vessels, then cancelling the prior probability quantity $P(\omega_k)$ and the normalizing constant $P(z)$, and leaving $P(z|\omega_k)$ indicates that the likelihood in equation 3.16 can be used to correctly classify the coronary vessels.

The two quantities (CDV and PDV) can also be combined exploiting any multi-variate behaviour, revising equation 3.16 as follows:

$$P(\mathbf{z}|\omega_k) = \frac{1}{\sqrt{|\mathbf{C}_k|2\pi}} exp(\frac{-(\mathbf{z}-\boldsymbol{\mu}_k)^T \mathbf{C}_k^{-1}(\mathbf{z}-\boldsymbol{\mu}_k)}{2}) \tag{3.18}$$

In the above equation, $\mathbf{z}$ now represents the measured feature vector, $\boldsymbol{\mu}_k$ represents the mean vector for the class ω_k (coronary artery) and $\mathbf{C_k}$ is the covariance matrix. If μ and σ were both estimated for non-arterial vessels, then the classification can be expressed as a maximum-likelihood problem, whereby the class assigned to a candidate vessel is described as such:

$$\arg\max_{\omega_k} P(\mathbf{z}|\omega_k) \tag{3.19}$$

In the absense of these parameters, iterative approaches to segmentation, such as the EM algorithm [13] can be used to perform the classification based on the measured CDV and PDV alone. The EM algorithm iterates over two steps until convergence. The E step computes the expected log-likelihood based on the current estimates of the mean and standard deviation for arterial and non-arterial vessels, while the M step computes the values for these parameters, which maximise the log-likelihood computed in the E step.

3.4 Conclusion and Further Work

3.4.1 Conclusion

The work presented in this paper demonstrates that an atlas, representing knowledge of the coronary anatomy can be constructed using a registration strategy. Sections 3.2.4 and 3.3.3 demonstrate one possible means by which such an atlas can be applied to the output of a coronary vessel tracking algorithm and improve its specificity by identifying and removing non-arterial vessels.

3.4.2 Further Work

The main focus of this work was to create a coronary artery atlas capable of providing information about coronary anatomy of the represented population. In this regard, 42 datasets is a starting point, although much more data is needed in order to arrive at the intended goal.

Coronary anatomy is divided into 3 main circulation patterns, namely, left dominant, right dominant and balanced circulation anatomies. Approximately 70% of the general population are right-dominant, 20% are balanced, while 10% are left-dominant. As an extension to this work, it would be beneficial to separate the atlas into three, one representing each anatomical category.

The results presented in figure 3.9 were obtained by applying various thresholds to the measured CDV and PDV; this exposes the algorithm to over-fitting errors, which can be avoided by a larger validation set and a more sophisticated classification algorithm like the EM algorithm as suggested in section 3.3.4.

The direction and density features express useful information about the coronary vessel centrelines, but the classification procedure used in this chapter has not exploited both features simultaneously. Future work would also aim to combine these in a suitable framework, for example, using the multivariate Gaussian approach discussed in section 3.3.4.

References

[1] N. E. Manghat, G. J. Morgan-Hughes, A. Marshall, and C. Roobottom, "Multi-detector row computed tomography: imaging the coronary arteries.," *Clinical radiology*, vol. 60, pp. 939–52, September 2005.

[2] P. M. A. van Ooijen, K. Y. Ho, J. Dorgelo, and M. Oudkerk, "Coronary Artery Imaging with Multidetector CT: Visualization Issues," *Radiographics*, vol. 23, November 2003.

[3] P. M. A. van Ooijen, B. M., and K. P., "System for Segmentation and Selective Visualization of the Coronary Artery Tree for Evaluation of Stenosis, Soft Plaque and Calcification in Cardiac CTA," *Imaging Decisions MRI*, vol. 8, pp. 25–30, September 2004.

[4] C. Kirbas and F. Quek, "Vessel extraction techniques and algorithms: a survey," in *Bioinformatics and Bioengineering, 2003. Proceedings. Third IEEE Symposium on*, pp. 238–245, March 2003.

[5] D. C. J. J. D. C. S. Aylward, "Vascular atlas formation using a vessel-to-image affine registration method," *Medical Imaging Computing and Computer-Assisted Intervention (MICCAI)*, 2003.

[6] D. Cool, D. Chillet, J. Kim, J. Guyon, M. Foskey, and S. Aylward, "Tissue-based affine registration of brain images to form a vascular density atlas," *Lecture Notes in Computer Science*, p. 9Ű15, 2003.

[7] N. Passat, C. Ronse, J. Baruthio, J.-P. Armspach, and C. Maillot, "Cerebral vascular atlas generation for anatomical knowledge modeling and segmentation purpose," *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, vol. 2, pp. 331–337, 2005.

[8] N. Passat, C. Ronse, J. Baruthio, J.-P. Armspach, and C. Maillot, "Magnetic resonance angiography: from anatomical knowledge modeling to vessel segmentation.," *Medical image analysis*, vol. 10, pp. 259–74, 2006.

[9] S. Aylward and E. Bullit, "Initialization, noise, singularities, and scale in height ridge traversal for tubular object centerline extraction," *IEEE Transactions on Medical Imaging*, vol. 21, no. 2, 2002.

[10] W. R. Crum, D. L. G. Hill, and D. J. Hawkes, "Information theoretic similarity measures in non-rigid registration," in *IPMI*, pp. 378–387, 2003.

[11] T. B. B. R. D. C. H.-P. Seidel, "Nonparametric Density Estimation with Adaptive, Anisotropic Kernels for Human Nonparametric Density Estimation with adaptive, Anisotropic Kernels for Human Motion Tracking," *Lecture Notes in Computer Science*, vol. 4814/2007, pp. 152–165, 2007.

[12] F. van der Heijden; Robert Duin; Dick de Ridder; David M. J. Tax, *Classification, Parameter Estimation and State Estimation: An Engineering Approach Using MATLAB*. John Wiley and Sons, 2004.

[13] G. J. Mclachlan and T. Krishnan, *The EM Algorithm and Extensions (Wiley Series in Probability and Statistics)*. Wiley-Interscience, October 2007.

Chapter 4

Optimal Atlas Selection Using Image Similarities in a Trained Regression Model to Predict Performance.

Abstract *An atlas in the context of atlas-based segmentation refers to a pre-selected image with labelled anatomical regions of interest. Atlas-based segmentation is the propagation of these labels to a novel image, usually after both images have been registered. In such applications, the choice of image to use as the atlas is very important. The overall goal of an atlas is to be as typical as possible to an anatomical category, but in practice there exists variability in human anatomy. One solution to maintain consistent segmentation accuracies is to have multiple atlases, representing several categories of anatomical variation, with a system for selecting the most appropriate atlas at the time of segmentation. This paper describes a method for selecting an atlas using a linear regression model to predict the segmentation accuracy based on image similarity metrics. It goes further to present an offline method for automatically selecting a set of atlases, representative of the training set to be used during segmentation; all of this illustrated by segmentation of the heart in 3D CT images.*

4.1 Introduction

4.1.1 Clinical Motivation

The problem of selecting the best atlas for segmentation of a particular novel image dataset out of a database of potential atlases arises because of the variability that exists within human anatomy. This variation in structure of an organ across different patients may be due to pathology, for instance patients suffering from heart failure usually have enlarged hearts relative to the thoracic cavity. It may also be due to subject's lifestyle, e.g. athletes have relatively large hearts, size and

demographics.

Anatomical variation in the thoracic cavity is a good example of general differences that exist among patients. Cardio-thoracic ratio (CTR) is a measure of the diameter of the heart with respect to the diameter of the thoracic cavity defined for chest X-Ray images(See figure 4.1); this measure has been found to be a good indicator of heart pathology [1]. It has also been reported in [1], [2] and [3], that this ratio not only varies based on pathology, but also varies with age, body size and ethnicity. Figure 4.2 demonstrates how this variability in CTR manifests itself in cross-sectional imaging modalities; in this case, chest CT. This poses a problem to atlas-based segmentation approaches that rely on a single atlas (figure 4.3), and so a segmentation approach which takes these differences into account is required.

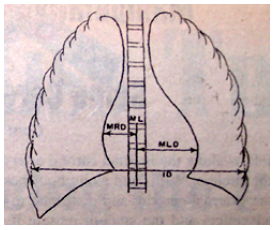

Figure 4.1: Cardiac Thoracic Ratio [4]
MRD = maximum transverse diameter of the right side of the heart, which is a line drawn from the midline of the spine to the most distant point of the right cardiac margin. ML= mid-line of the spine. MLD = maximum transverse diameter of the left side of the heart, which is a line drawn from the mid-line of the spine to the most distant point of the left cardiac margin. ID= greatest internal diameter of the thorax. TD = MRD + MLD. Cardiothoracic ratio = TD/ID

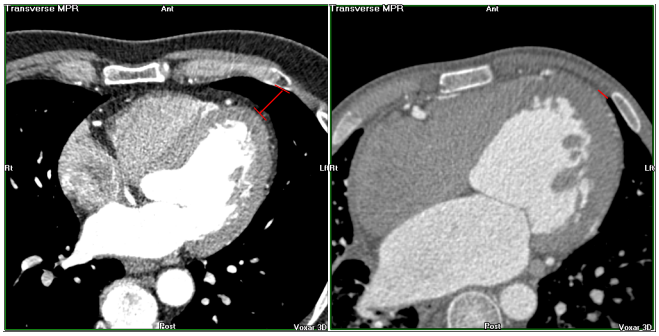

Figure 4.2: Heart-Shape Variability
The images show axial projections of CT angiograms taken from two patients. In the image on the left, there is a relatively large cavity in between the pericardium and the rib-wall as compared to the image on the right.

4.1.2 The Problem

A system is required which is capable of selecting the most similar atlas image to a novel image from a set at the time of segmentation. A system is also required for selecting an optimal set of

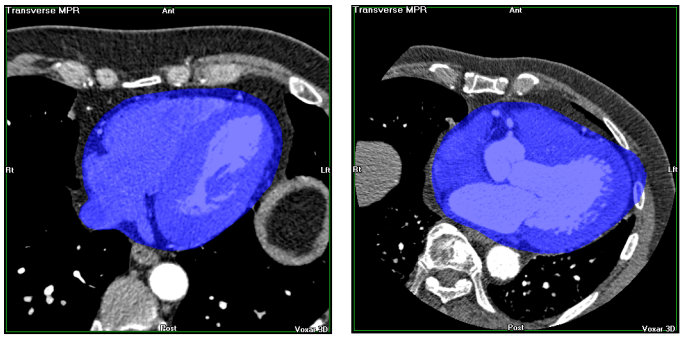

Figure 4.3: Single-Atlas Errors due to Variation in CTR
In this example, the image on the left is used as an atlas to segment the heart in the image on the right, which has a higher CTR than the atlas image. The result is a gross over-segmentation into the ribs.

atlases capable of representing a wide range of anatomical variations.

4.1.3 State of The Art

The evaluation of selection strategies in [5] demonstrates that selecting the atlas which produces the highest normalised mutual information after non-rigid registration to the novel dataset, results in higher average segmentation accuracy on the experiment set than the highest average accuracy obtained by using any atlas in the set on its own. Subsequent approaches in [6] extend this method by computing the similarity metric in a region of interest; this is particularly useful in situations where certain sub-anatomies are likely to vary across subjects, such as the right anterior cingulate cortex and right amygdala in the brain.

Other similarity metrics can be used. In [7], the magnitude of the deformations from the non-rigid registration warp field, which transforms each atlas dataset to the novel dataset is used to select an atlas, with the assumption that the most similar dataset will require the least amount of deformation.

An alternative approach to multi-atlas based segmentation is decision-fusion from multiple classifiers as reported in [8], [9], [10], [5] and [11], where each atlas is treated as a separate classifier and each pixel in the novel image assigned a class label following registration. The problem is to find a combination of the individual segmentations that produces the most accurate overall segmentation. In [10], the vote rule, whereby the final class assigned to each pixel is the modal class from all the individual segmentations, is compared with the sum rule, whereby the final class assigned to each pixel is the class that maximises the probability of the pixel belonging to it, averaged over all the atlases. The method of combining atlases using either the vote or sum rule is reported in [5] to produce better segmentation results than by simply selecting the most similar atlas.

A problem with vote counting strategies arises in cases where multiple anatomical categories

exist, and the category represented by the novel dataset has minority representation amongst the group of atlases used. In [8], this issue is addressed by limiting the set of atlases to those that are most similar to the novel dataset using image similarity and meta-information, such as patient's age, gender and clinical status. In [12], selection of the set is addressed as a clustering issue, and the mean-shift algorithm is used to categorise a large atlas database. This implies that a novel dataset will be best segmented using atlases within its category.

4.1.4 Contribution

The method proposed in this paper draws upon the findings reported in [8], [13] and [5], stating that an atlas selected on the basis of image similarity produces better results than is possible with a single atlas used across all novel datasets. This paper extends this idea by acknowledging that image similarity measures such as mutual information and cross-correlation can be used to predict the segmentation accuracy. However, it is possible that the mutual information between one atlas and the novel dataset may be less than that computed using a second atlas, but the first atlas nevertheless generates a more accurate segmentation. What is needed is a measure that is comparable across the entire set of possible segmentations. Here, instead of using image similarity measures directly to compare segmentations, they are used to predict the accuracy of the segmentation. This predicted accuracy is then directly comparable across the set of segmentations. The system is based on a linear regression predictor model, developed from training data.

By applying this system to whole heart segmentation in CT and using the correlation coefficient within a region around the boundary of the heart as a predictor, it was observed that a set of four atlases, covering different categories available in the training set, was able to produce a higher segmentation accuracy than using any single atlas.

4.2 The Method

4.2.1 Introduction

This section describes all aspects involved in the selection of an optimal atlas during multi-atlas based segmentation. A description of multi-atlas based segmentation (figure 4.4), which uses the optimal atlas selection is provided first. The remaining sections describe the process of selecting an optimal set of atlases from a training set.

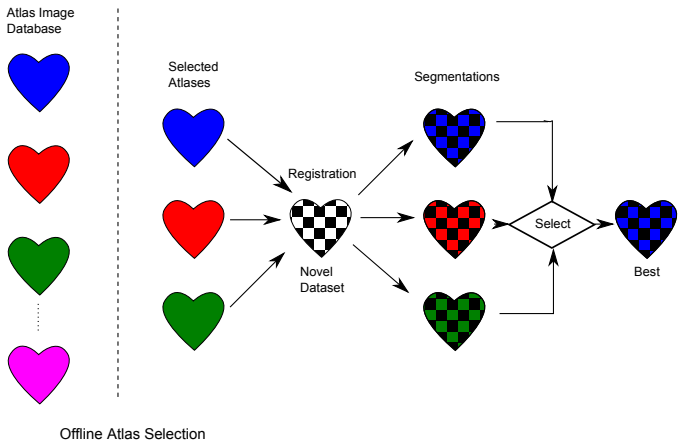

Figure 4.4: Multi-Atlas Based Segmentation
The atlases in use in the application are selected offline from a database of image atlases. Each of the chosen atlases produces a segmentation when aligned with the novel image, therefore a mechanism is required, which is capable of predicting the best segmentation.

4.2.2 Multi-Atlas Based Segmentation

Assuming a set $A = \{a_0, ..., a_{M-1}\}$ of M atlases has been provided, segmentation of the required anatomical region in a novel image T begins by registering each $a_i (i = 0, ..., M-1)$ with the novel image in order to produce a set of warp-fields $W = \{W_0, ..., W_{M-1}\}$, where W_i maps A_i onto T.

The registration process is a multi-scale approach, performing 2 passes of rigid registration in order to find the transformation that maximises the mutual information between the pair of images. The transformation parameters used here are 3D translation, anisotropic scaling, and 3-dimensional rotation (roll, pitch and yaw). Once this initial alignment is achieved, 2 passes of non-rigid registration are carried out in order to find a deformation field which maximises mutual information, within fluid and elastic constraints[14].

Using the warp-field W_i, the cross-correlation between $W_i(a_i)$ and T can be measured. The cross-correlation between any $W_i(a_i)$ and T is not used directly to determine the segmentation accuracy, because this measure by itself is not directly comparable across all possible pairs of $W_i(a_i)$ and T. In order to get a more standard measure for selecting the best out of M possible segmentations, a linear regression model specific to each atlas is used. The function is of the form $y = mx + c$, where x is the value of the cross-correlation and y is then the predicted Jaccard overlap (equation 4.1); m and c are stored as part of the atlas.

$$Jaccard(A, B) = \frac{Vol.(A \cap B)}{Vol.(A \cup B)}, \tag{4.1}$$

where A and B are the generated and ground-truth segmentations, respectively.

The labelled anatomical region from the a_i that predicts the highest segmentation accuracy,

is propagated onto T using W_i, resulting in a segmentation.

$$a^* = \underset{a_i}{\operatorname{argmax}}\{m_{a_i}\mathit{corr}(W_i(a_i), T) + c_{a_i}\}$$

Section 4.2.4 describes how these regression functions are obtained, while section 4.2.7 describes a method for selecting an appropriate atlas set A based on available training data.

4.2.3 Materials

Cardiac CT angiography (CCTA) datasets obtained from hospitals in Japan, U.S.A, and the U.K. were used in this experiment. The patients therefore varied not only in ethnicity, but also in age, sex and pathology, although demographic information was not present in the DICOM meta-data.

Each scan was captured as part of a cardiac perfusion protocol, from the 65% phase of the cardiac cycle. Different scanners were used depending on the source of data, but all images were captured following administration of a contrast agent. The resultant images were 3-D volumes containing at least 300 slices with 512x512 pixels per slice, and a slice thickness of 0.5mm with typical intra-slice pixel spacing of 0.4mm. Each image therefore contained the entire heart and thoracic cavity.

Whole heart segmentations, represented as masks containing the entire heart as one fully-connected domain, were collected by a clinically qualified individual on 21 of these datasets. In each dataset, the distance along the long axis between the apex of the heart and the rib-wall was measured (figure 4.5) as well as the CTR. The CTR was measured by dividing the longest diameter of the heart by the longest diameter of the chest cavity. It was observed that in datasets 14-21, the distance between the apex and the rib-wall was close to zero, while the other datasets generally demonstrated greater distances. The CTR on the other hand was less discriminating, although the highest values for this measure were observed on this set. Demographics are unknown for these datasets, however datasets 14-21 were received from hospitals in Japan while the rest were received from hospitals in U.S.A and the U.K.

Atlases were created from each image and mask combination by constraining the image to a region of interest containing the heart mask morphologically dilated by 6mm to contain a degree of context.

4.2.4 Training Parameters

Each dataset in turn is used as an atlas to segment the heart in all the other datasets using the registration method outlined in section 4.2.2. The final mutual information between the atlas dataset and the novel dataset, along with the cross-correlation between the two aligned volumes, the cross-correlation between the two volumes measured on a region of interest about the apex

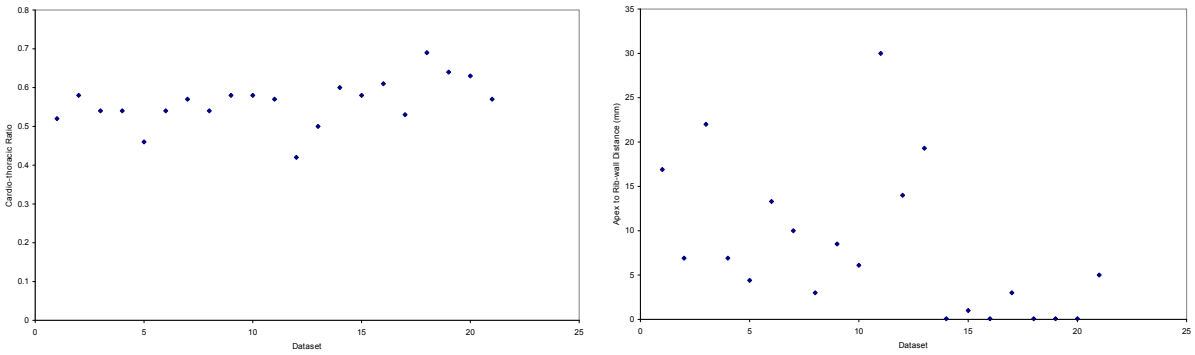

Figure 4.5: Cardiac-Thoracic Ratio and Apex to Rib-Wall Distance for Patient Data
Datasets 14-21 were received from hospitals in Japan. The distance between the heart's apex and the rib-wall is generally much lower in this set than on the remaining data. CTR is not a clear discriminant between the two sets of data, although the highest values are observed within datasets 14-21.

of the heart, and the final segmentation accuracy measured using the Jaccard overlap (4.1) are collected from each segmentation.

These measurements make up the training parameters and they are presented in seperate cross-segmentation matrices, **J** and **CC**(figure 4.6), where $\mathbf{J}(i,j)$ is the Jaccard overlap produced when atlas a_j is used to segment the heart in dataset t_i, and $\mathbf{CC}(i,j)$ is the cross-correlation measured at the optimal alignment between atlas dataset a_j and novel dataset t_i. The diagonal elements represent parameters collected from self-segmentation.

	1798_00	1803_02	1805_00	1831_00	1686_00	1688_001
01798_006	0.966	0.887	0.915	0.86	0.914	0.903
01803_025	0.864	0.951	0.875	0.887	0.833	0.844
01805_003	0.901	0.889	0.959	0.925	0.859	0.867
01831_004	0.863	0.905	0.929	0.954	0.885	0.885
01686_001	0.898	0.846	0.899	0.913	0.956	0.878
01688_001	0.934	0.9	0.937	0.933	0.922	0.961

Jaccard Overlap Matrix

	1798_00	1803_02	1805_00	1831_00	1686_00	1688_001
01798_006	0.928	0.855	0.911	0.794	0.878	0.911
01803_025	0.822	0.92	0.901	0.85	0.846	0.874
01805_003	0.838	0.795	0.927	0.804	0.842	0.854
01831_004	0.882	0.829	0.907	0.934	0.879	0.869
01686_001	0.816	0.754	0.855	0.828	0.918	0.817
01688_001	0.863	0.786	0.854	0.8	0.821	0.917

Cross-Correlation Matrix

Figure 4.6: Training Parameter Matrices
The datasets names are displayed as row and column headers in the matrix. The columns represent the atlas datasets, while the rows represent the novel datasets; therefore, the diagonal elements are self-segmentation parameters. The blue and red coloured cells show cases where CC is higher between one atlas and a novel dataset, but the resulting Jaccard overlap is lower.

4.2.5 Regression Model

The parameter model used in this system is a univariate least-squares regression function measured on each similarity measure for each atlas in the form $\hat{y} = \hat{\beta}_1 x + \hat{\beta}_0$, correlating mutual information or correlation coefficient with segmentation accuracy (figure 4.7). Equation 4.2 therefore predicts the segmentation accuracy obtained if using atlas dataset a_j to segment the heart in dataset t_i based on the measured correlation coefficient between the two at the

region-of-interest.

$$J^*_{a_j}(t) = \mathbf{CC}(i,j) * \hat{\beta}_{1_{a_j}} + \hat{\beta}_{0_{a_j}} \tag{4.2}$$

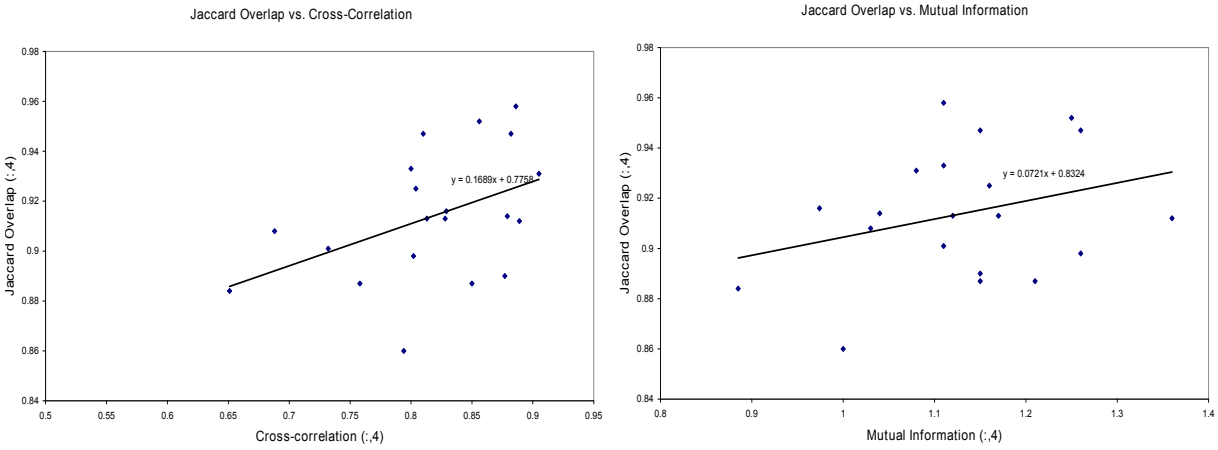

Figure 4.7: Segmentation Overlap vs. Similarity Metric
This regression model is built up for each atlas based on values collected from the relevant column in the cross-segmentation matrix. Both images show regression functions for the same atlas using different similarity metrics as predictors; the slope is greater using CC, assuming similar residuals implies that CC is a better predictor. It is arguable that a linear regression model is the best line fitting technique for the observed data, but this model was chosen purely for its simplicity and because of the small size of the training set.

4.2.6 Atlas Statistics

Using the per atlas regression functions, J^*_a and the Jaccard overlap matrix $\mathbf{J}$, this stage uses the regression functions to predict the segmentation overlap and compares the predicted value $J^*_{a_j}(t_i)$, to the actual value collected during training $\mathbf{J}(i,j)$. The following information is obtained (figure 4.8):

- Correlation:

  $$\rho(J^*_{a_j}(t_i), \mathbf{J}(i,j))$$

  A value close to $+1$ indicates strong prediction power.

- Mean absolute error:

  $$\frac{1}{N}\sum_{i=0}^{N-1} |J^*_{a_j}(t_i) - \mathbf{J}(i,j)|$$

  This is the average of the discrepancy between the predicted overlap and the actual overlap produced by the particular atlas over all the datasets. The best single atlas is generally the atlas with the lowest mean error.

- Standard deviation of the error:

$$\sigma(J^*_{a_j}(t_i) - \mathbf{J}(i,j))$$

  This is a measure of the distribution of the discrepancy between the predicted overlap and the actual overlap produced by the particular atlas over all the datasets. A low value of this statistic indicates a stable atlas.

- Max absolute error:

$$\max_i |J^*_{a_j}(t_i) - \mathbf{J}(i,j)|$$

  The maximum discrepancy between the predicted overlap and the actual overlap produced by the particular atlas over all the datasets.

- Best-segmentation atlases:

$$\operatorname*{argmax}_{a_j} \mathbf{J}(i,j)$$

  For each novel dataset, the atlases which performed in the top three in terms of segmentation accuracy are highlighted. This is intended to provide information about which datasets are better suited for certain categories (figure 4.9).

4.2.7 Offline Multi-Atlas Selection

The aim of this step is to select the set of M atlases from a set of N training datasets that produces the highest overall average segmentation accuracy, while using the regression function to decide on which atlas to use on each novel dataset. The end goal is to find out whether this combination represents all categories present in the training population, and whether it is capable of producing a more accurate segmentation than using only one atlas.

In this case, M is increased sequentially from 1, leading to $\binom{N}{M}$ combinations of atlases in each iteration. The procedure is as follows:
For each combination A *where* $A = \{a_0, ..., a_{M-1}\}$,

- For each t_i of the remaining N-M training datasets,
  - Compute the regression parameters $\hat{\beta}_{1SIM_{a_m}}$ and $\hat{\beta}_{0SIM_{a_m}}$ for each atlas a_m in A (see equation 4.2), excluding t_i from the observations.
  - Predict the overlap for dataset t_i:

$$J^*_{a_m}(t_i) = \hat{\beta}_{1SIM_{a_m}} \mathbf{SIM}(\textit{index}(t_i), \textit{index}(a_m)) + \hat{\beta}_{0SIM_{a_m}}$$

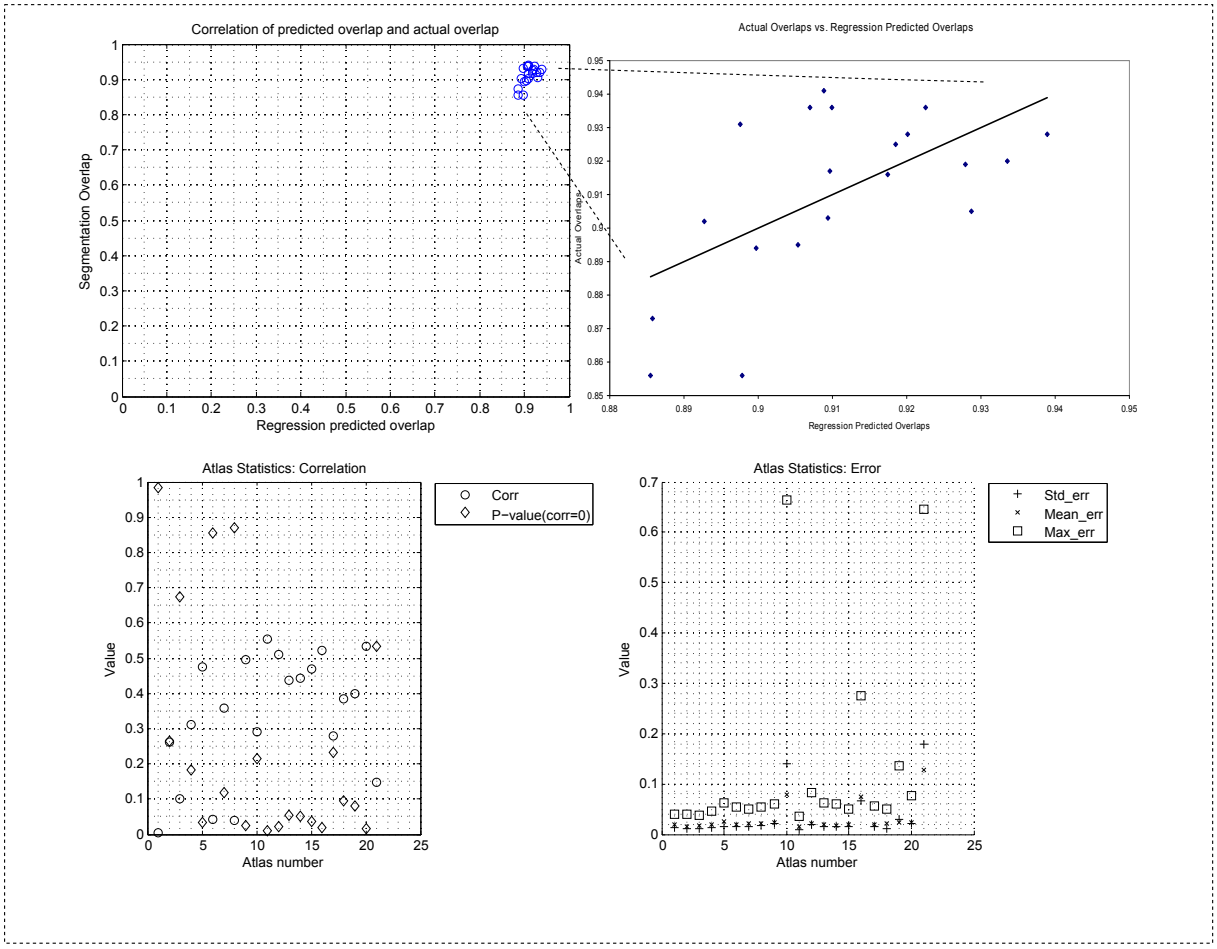

Figure 4.8: Atlas Statistics: Correlation and Error

Atlas number 19 appears to be a good predictor because it has both high correlation and low mean absolute error between its predicted overlaps and actual overlaps produced.

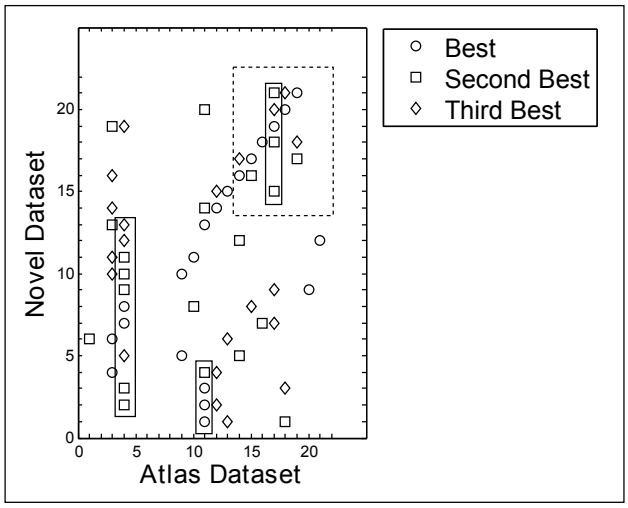

Figure 4.9: Best Segmentation Matrix
Datasets 9, 12, and 16-21 all belong to the group with high CTRs. The cluster in the top-right corner indicates that atlases from this category, particularly atlas 19 are the best for segmenting hearts within this group. Atlas number 4 generates the best segmentation accuracy over datasets [2-14].

SIM is the cross-segmentation matrix containing the image similarity measure used to predict overlap.

– Select a_m with the highest predicted overlap:

$$a_m^* = \underset{a_m}{\operatorname{argmax}}\, J_{a_m}^*(t_i)$$

– Collect the actual overlap produced by a_m^* on t_i from the overlap matrix: $\mathbf{J}(\mathit{index}(t_i), \mathit{index}(a_m^*))$

- Compute the average segmentation overlap:

$$J_A = \frac{1}{N-M} \sum_{i=0}^{N-M-1} \mathbf{J}(\mathit{index}(t_i), \mathit{index}(a_m^*))$$

The combination with the highest overlap is chosen.

$$A^* = \underset{A}{\operatorname{argmax}}\, J_A \tag{4.3}$$

4.3 Results and Discussion

The entire system described in section 4.2 was implemented in Matlab, and all the required information made available offline in the cross-segmentation matrices (figure 4.6).

The best set of 2 atlases consists of atlas numbers 4 and 19; these datasets are significantly different in terms of heart shape and CTR (see figure 4.10), representing the 2 categories identified. From figure 4.9, it is evident that atlas number 4 is more suitable for segmenting datasets 2 to 14, while atlas number 19 is more suitable for segmenting datasets 16 to 21. These datasets have already been identified as separate groups based on distance of the apex to the rib-wall in each case. When M is increased to 3, atlas 13 joins the set; and from figure 4.9 it can be seen that this dataset is suitable for segmenting datasets 1 to 4.

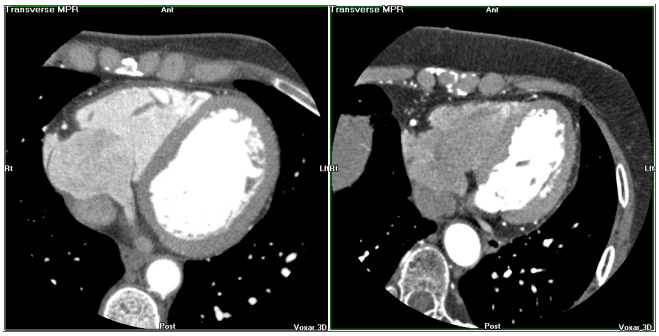

Figure 4.10: Best Combination of 2 Atlases for Heart Segmentation
left, right: 4, 19. The size of the hearts with respect to the thoracic cavity in each of these datasets is significantly different, and is easily seen from the distance between the apex of the heart and the rib-wall.

The multi-atlas segmentation simulator (section 4.2.7) produced a set of 4 atlases which generated an overall average segmentation accuracy of 92.8% compared to the best single atlas, which generated an average accuracy of 91.5%. This performance was achieved by using the cross-correlation at a region containing the heart apex and the rib-wall (figure 4.11).

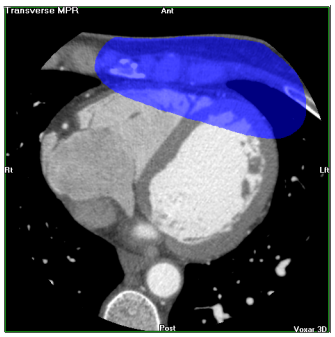

Figure 4.11: ROI over which image similarity is used to predict segmentation performance.
The image similarity metric is computed between the atlas and the novel datasets over the region specified in blue.

Figure 4.12 shows the average segmentation overlap generated by the simulation for M ranging from 1 to 5 inclusive. The uppermost line shows the best possible segmentation overlaps that can be generated if the most suitable atlas was selected from the set of M in every case. This

chart suggests that the Region of Interest (ROI) based measures, particularly cross-correlation are better predictors than the same similarity metrics measured on the entire volumes.

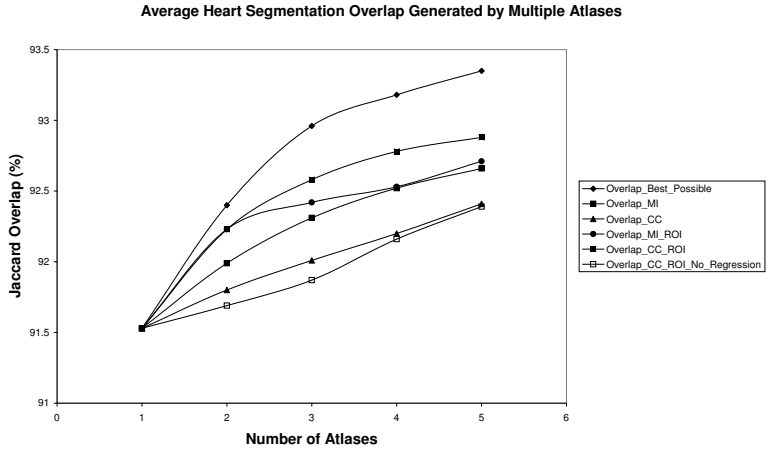

Figure 4.12: Average Heart Segmentation Overlap Generated by Multiple Atlases
The topmost line shows the average heart segmentation accuracy generated if the atlases were selected based on the known actual overlaps, rather than predicted overlaps. The chart suggests that ROI computed metrics are better predictors than the same metrics computed on the entire image volume.

4.4 Conclusion and Further Work

This chapter presents a method for optimal atlas selection from a set of candidate atlas images. The underlying method is applied at two stages in multi-atlas based segmentation. In the first stage, it is used offline to select an optimal multi-atlas set from a database of atlas images, while in the second stage it is used during application to select an optimal atlas from within the set.

Each atlas contains an image, a mask containing the labelled heart and a pair of linear regression parameters. The parameters are obtained from a training set of images by co-registering each pair of images and recording the resulting similarity measures and Jaccard overlaps of the generated segmentation. The optimal atlas from a set is selected following registration by measuring the cross-correlation between each atlas image and the novel image and applying the respective linear regression slopes and intercepts to predict the Jaccard overlap. The atlas image that predicts the highest segmentation accuracy is selected.

The results described in section 4.3 demonstrate that a multi-atlas approach performs better on average than a single atlas method, given the same registration method. The multi-atlas set

selected based on the data demonstrates that this method is capable of identifying categories of anatomic variation present in the training set.

Further work is required however, for instance the current system uses cross-correlation, measured over a specific region, as the predictor for segmentation overlap. It is sufficient to solve the problem posed at the start of this paper, but the particular ROI used is ad-hoc and therefore may not be the best predictor for selecting atlases. This can be seen in figure 4.12 by the increasing divergence between the *Overlap_CC_ROI* and *Overlap_Best_Possible* curves; a better predictor is required to reduce this distance. A shell-region around the heart (figure 4.13) may be worth trying as a potentially more general ROI.

The current method uses a linear least-squares regression model to fit the data, however figure 4.7 indicates that this may not be the best model for the data. Prediction capabilities of other parametric functions need to be investigated, but this requires more training data in order to avoid over-fitting errors.

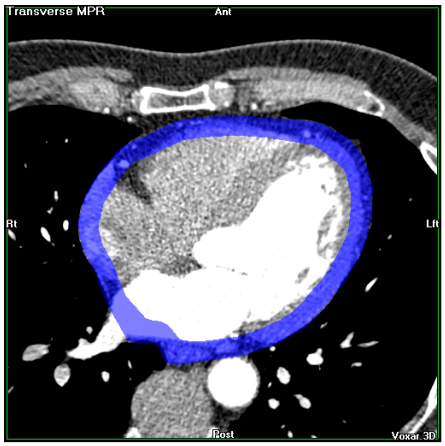

Figure 4.13: Shell Region Around the Heart
The image similarity metrics discussed may be measured on the region in blue to gain more general information about the quality of the final segmentation.

Time considerations also exist: if the system is to be applied in practice, it would require multiple registrations. This will most certainly increase the time taken to segment a single case, although it could be made more efficient in a number of ways, e.g. measuring the similarity after rigid registration, and performing the registrations in parallel.

References

[1] Nickol and Wade, "Radiographic heart size and cardiothoracic ratio in three ethnic groups: a basis for a simple screening test for cardiac enlargement in men," *Br J Radiol*, vol. 55, pp. 399–403, 1982.

[2] M. Mikado, U. Keiko, B. Akira, I. Atsushi, M. Tooru, F. Takehiko, K. Shozo, and K. Terukazu, "Relationship between cardio-thoracic ratio and body height and weight, blood pressure, and pulse rate in university students," *Journal of Health Science*, vol. 24, pp. 17–22, 2002.

[3] A. R. P. Walker D.Sc, B. D. Richardson M.Sc., M. Wadvalla, and B. F. Walker, "Cardio-thoracic ratio in negroes in southern africa," *Postgraduate Medical Journal*, vol. 48, pp. 584–589, October 1972.

[4] D. Pate, "Chest x-ray," *Dynamic Chiropractic*, vol. 12, pp. 1–3, 1994.

[5] T. Rohlfing, R. Brandt, R. Menzel, D. B. Russakoff, and C. R. Maurer Jr., "Quo vadis, atlas-based segmentation?," in *The Handbook of Medical Image Analysis – Volume III: Registration Models* (J. Suri, D. L. Wilson, and S. Laxminarayan, eds.), ch. 11, pp. 435–486, New York, NY: Kluwer Academic / Plenum Publishers, Aug 2005.

[6] M. Wu, C. Rosano, P. Lopez-Garcia, C. S. Carter, and H. J. Aizenstein, "Optimum template selection for atlas-based segmentation," *Neuroimage*, vol. 34, pp. 1612–1618, 2006.

[7] O. Commowick and G. Malandain, "Efficient selection of the most similar image in a database for critical structures segmentation," in *Proc. Medical Image Computing and Computer Assisted Intervention (MICCAI'07), Part II* (N. Ayache, S. Ourselin, and A. Maeder, eds.), no. 4792, pp. 203–210, October 2007.

[8] P. Aljabar, R. Heckemann, A. Hammers, J. Hajnal, and D. Rueckert, "Multi-atlas based segmentation of brain images: Atlas selection and its effect on accuracy," *NeuroImage*, vol. 46, no. 3, pp. 726–738, 2009.

[9] H. Kirisli, M. Schaap, S. Klein, L. Neefjes, A. Weustink, T. van Walsum, and W. Niessen, "Fully automatic cardiac segmentation from 3d cta data: a multiatlas based approach," in *SPIE Medical Imaging 2010*, 2010.

[10] J. Kittler and F. Alkoot, "Sum versus vote fusion in multiple classifier systems," *PAMI*, vol. 25, pp. 110–115, January 2003.

[11] S. K. Warfield, K. H. Zou, and W. M. Wells, "Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation," *Medical Imaging, IEEE Transactions on*, vol. 23, pp. 903–921, 2004.

[12] D. Blezek and J. Miller, "Atlas stratification," *Med Image Anal*, 10 2007.

[13] J. Lötjönen, R. Wolz, J. Koikkalainen, L. Thurfjell, G. Waldemar, H. Soininen, and D. Rueckert, "Fast and robust multi-atlas segmentation of brain magnetic resonance images," *NeuroImage*, 2009.

[14] W. R. Crum, D. L. G. Hill, and D. J. Hawkes, "Information theoretic similarity measures in non-rigid registration," in *IPMI*, pp. 378–387, 2003.

Chapter 5

A Combined Multi-Atlas and Unsupervised Classification Technique for Automatic Segmentation of the Kidney and Renal Cortex in Low-Dose CT Data.

Abstract *Automatic segmentation of the kidney and renal cortex in low-dose CT scans can serve as a useful aid to nephrologists during kidney transplant planning and can also be a prerequisite to automating renal perfusion measurements, in order to quantify kidney function.*

There are various published methods for successfully segmenting the kidney and other abdominal organs from MRI, and non-contrast enhanced CT images. This paper describes a method for segmenting the kidneys and renal cortices from low-dose CT images, which tend to have a low signal to noise ratio.

The method proposed in this paper makes use of multi-atlas based segmentation, selecting the most suitable atlas to the subject from a set during run-time. A refinement step estimates tissue parameters for the renal cortex and medulla using the expectation maximisation algorithm and combined with the spatial probabilities produced by the registered atlas in a maximum a-posteriori framework, each pixel is classified. The classification framework also takes into account the hierarchical relationship between the kidney and renal cortex and exploits this property to increase segmentation accuracy.

The proposed method runs in 10 seconds and has been tested on 22 datasets, demonstrating an average voxel overlap between the automatic segmentations and ground-truth segmentations of 85% and 72% for the kidneys and renal cortices respectively.

5.1 Introduction

5.1.1 Clinical Motivation

Organ segmentation from medical images is an important first step for further medical analyses, usually a pre-processing step for automatic measurements. For clinicians interested in VR views, providing automatic segmentations of the organs can provide a visualisation aid, allowing un-distracted views of the organs, saving the time required to perform manual delineation.

In the year 2008-2009, nearly 3,000 kidney transplants were performed in the UK and 12,000 in the United States[1, 2]. Transplant planning aims at matching the recipient to a donor and carrying out steps to ensure successful surgery. Medical imaging can be used to non-invasively assess the kidney volumes, the number of renal arteries in both kidneys, the extent of damage of kidneys from living donors, amongst other important issues [3]. An accurate automatic segmentation tool can assist radiologists with this task by increasing visibility of the kidneys and aiding automatic measurements.

Renal perfusion is a common procedure in nuclear medicine, used to detect abnormalities in kidney function, which are usually indicators of renal artery stenosis, tumours, or acute rejection as a result of a failed transplant. This procedure is becoming increasingly popular in MR using Dynamic Contrast Enhanced (DCE-MR) imaging, where the glomerular filtration rate (GFR) is well correlated with the injected Gadolinium-diethylenetriamine pentaacetic acid (Gd-DTPA) clearance rate [4, 5].

The prevalence of CT scanners, due to lower cost compared to MR scanners is motivation for implementing renal perfusion procedures in CT. It is now possible due to availability of wide detector scanners with increased spatial resoultion, such as Toshiba's Aquilion ONE. Whole organs can be scanned in one gantry rotation (0.3 seconds), thereby limiting the dangers of over-exposure of patients to ionising radiation.

Carrying out this procedure in CT requires a series of scans, following injection of contrast agent into the patient. In order to capture multiple phases without posing more danger to the patient from ionizing radiation, the scans are taken with just enough X-Ray dose to guarantee a diagnostic image. This results in an effective total radiation dose of $\approx 10 - 15$mSv per exam [6](compared to ≈ 0.7mSv for a single pelvic X-Ray scan) and affects the resultant image by increasing its level of noise as shown in figure 5.1.

The ability to quantify kidney perfusion in CT is equivalent to measuring the GFR. This flow can be measured fairly accurately in CT angiography by the contrast media clearance rate, and relies on accurate measurements of kidney and renal cortex volumes [7, 8], which can easily be provided given accurate segmentations of both anatomies.

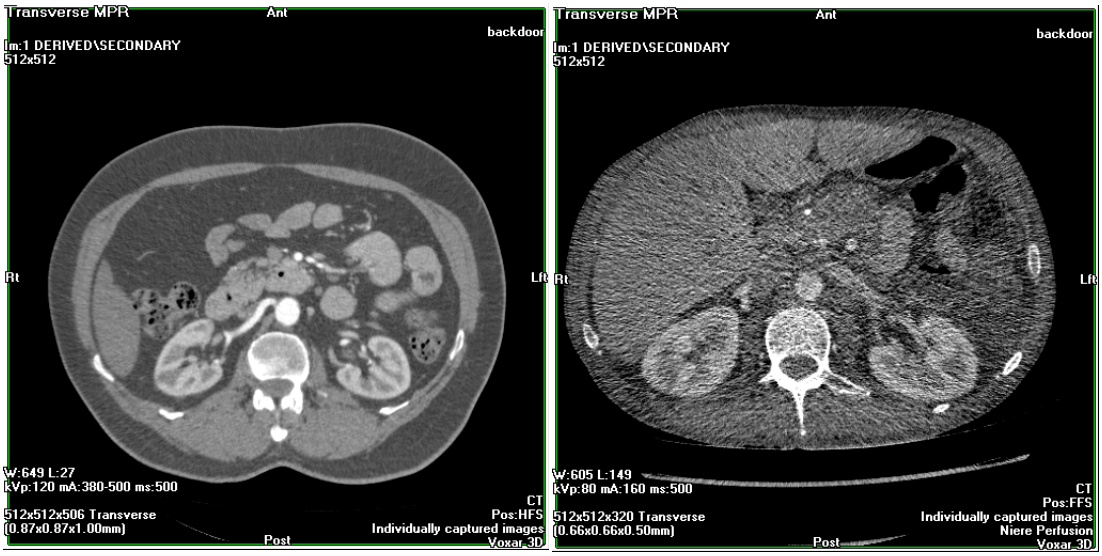

Figure 5.1: Normal vs. Low Dose CT
The image on the left was captured from a scan taken at 380mA, while the one on the right was from a renal perfusion series, taken at 100mA.

5.1.2 The Problem

Low-dose CT poses a problem for automatic segmentation algorithms due to the images having lower signal to noise ratio than regular-dose CT data. The kidneys and other abdominal organs are structurally and spatially variable across patients; this also poses a challenge to atlas-based segmentation techniques.

Minimising the radiation dose while scanning a patient is advantageous however, therefore developing an algorithm which can successfully segment anatomy of interest under these conditions is necessary.

A method is required, which deals with both the noise generated by low-dose CT and the anatomical variation that exists within abdominal organs across individuals, to produce accurate segmentations of the kidneys and renal cortices (figure 5.3).

5.1.3 State of the Art

The problem is essentially a segmentation task, hence there are a lot of published segmentation methods targeted at other organs and modalities, although several attempts at kidney segmentation have been made; for instance [9, 10, 11] perform kidney segmentation on Dynamic Contrast Enhanced Magnetic Resonance Images (DCE-MRI). Kidney segmentation was carried out as part of multiple abdominal organ segmentation on non-contrast enhanced CT images in [12], and on CT images in [13, 14].

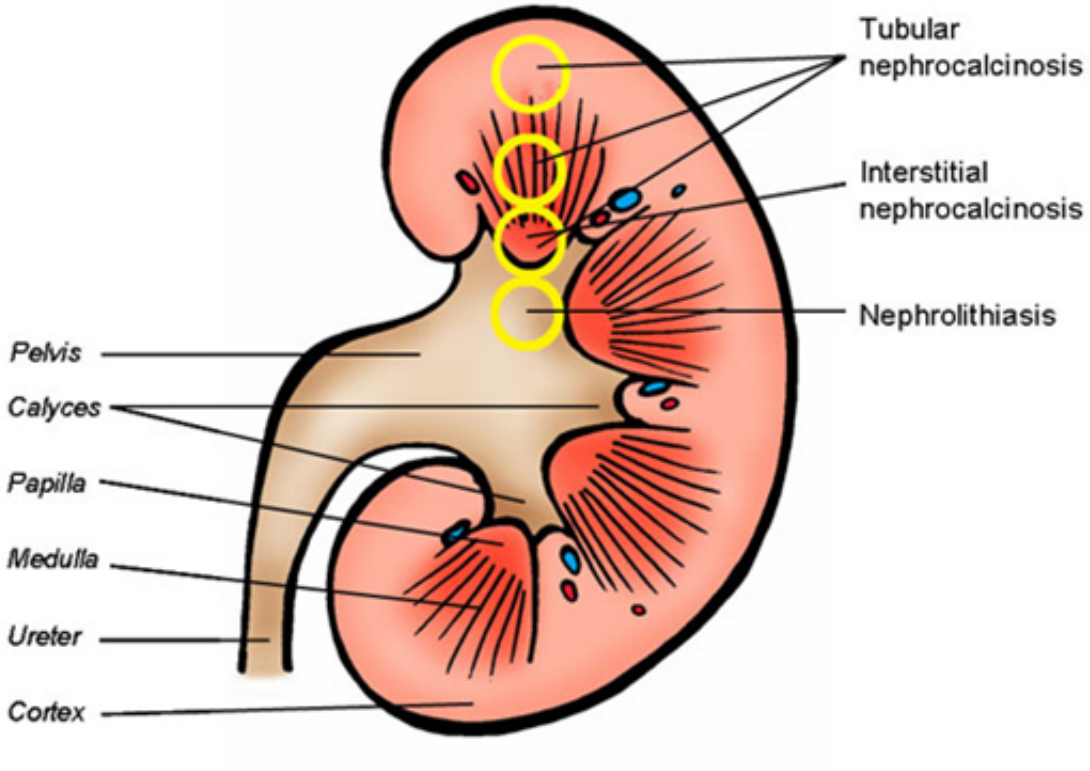

Figure 5.2: Kidney Anatomy
Kidney anatomy highlighting the medulla, renal cortex and pelvis. Image courtesy of glycoforum (www.glycoforum.gr.jp).

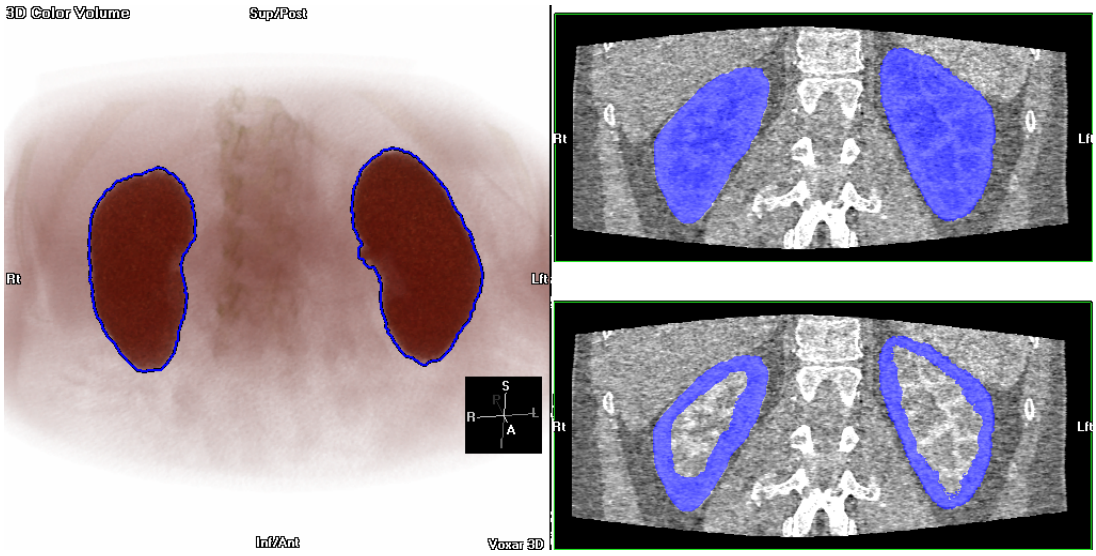

Figure 5.3: Kidney and Cortex Segmentation

Atlas-Based Methods

Atlas based segmentation is the process of mapping a set of labelled anatomical regions from a representative image, termed the atlas, onto a novel image following an initial alignment process, usually involving volume registration. Atlas-based methods have been used to segment various anatomical regions, and are either used alone or to initialise other voxel-based classification methods. These methods are increasingly popular because they provide segmentation algorithms with prior anatomical knowledge about the anatomy of interest in a simpler way than popular segmentation algorithms such as level-set methods.

Average Atlas In [15], an average bee brain atlas is created by registering a set of images to an iteratively created reference image. It demonstrates that the average atlas approach accounts for variation across subjects and so produces a higher segmentation accuracy than the single atlas approach. In [12], a probabilistic abdominal multi-organ atlas is created by registering multiple images to a reference image and propagating their labels onto a common space. This produces a vector-valued atlas mask, providing a probability for each registered novel pixel belonging to each of the abdominal organs given its spatial location.

Optimal Atlas Selection The comparison of various atlas-based methods in [16] demonstrates that a multi-atlas approach, whereby the most suitable atlas is selected following non-rigid registration of candidate atlases with the novel images, is capable of producing higher segmentation accuracy than methods that rely on a single or averaged atlas. In [17], this multi-atlas approach is used, and for each structure in the brain, the candidate atlas image with the highest normalised mutual information is selected.

Atlas Combination Multi-atlas approaches that combine the segmentations produced by each individual atlas to produce a final segmentation have also been reported to produce accurate segmentations [18, 19, 20, 21, 22, 23, 24]. In [18], the set of atlases to use for segmentation is limited to the set of images that are most similar to the novel image, with similarity based on image and meta-data (e.g. gender and age) similarity. In [20], the individual heart segmentations are combined using a majority vote rule, whereby the label assigned to a pixel is the modal label from all candidate labellings. Shape-based averaging of candidate segmentations, using signed Euclidean distance transforms [25, 26], or the Poisson transform [27] have also been demonstrated to produce smoother and more accurate segmentations than voting-based combinations. Assuming each candidate atlas performs differently depending on the novel image, then these segmentation performances may be iteratively estimated and used to weight each candidate segmentation as shown in [23, 24]. A similar approach by [19] demonstrates that local combination strategies outperform global methods, particularly in brain tissue segmentation. In

[21, 22], the labels from each segmentation are combined to provide a prior probability mask which is then used in voxel classification.

Combined Atlas-Based and Classification Methods

The atlas provides prior anatomical information about the organs of interest, even more so in the average and multi-atlas cases, where the warped atlas represents a spatial probability distribution of the anatomy of interest. Methods exploit this extra information by using the output after registration of the atlas to initialise supervised or un-supervised classification algorithms.

Bayesian Methods In [12], the probabilistic atlas provides the prior probabilities required for MAP classification, while the tissue means and variances are obtained from voxels with high probabilities of belonging to the respective tissues. In [28, 21] however, the tissue model parameters are iteratively modified using the Expectation Maximisation (EM) algorithm. Both methods employ a regularisation stage to ensure smoothness and connectivity of the resulting segmentations.

Markov Random Fields Markov Random Fields (MRF) are a popular choice for regularisation as they enforce smoothness by penalising dissimilar adjacent labels. In [12], the MRF is modelled by a Gibbs distribution and multiplied by the posterior probability function to form an energy functional which is then iteratively maximised.

Graph-cuts have been reported in [29] to be useful for minimising energy functionals, and so in [30, 21], a graph-cuts algorithm is used to minimise a MRF energy functional following non-rigid registration using a multi-atlas. The energy functional[31] consists of a data-term equivalent to the negative logarithm of the posterior probability, and a smoothness term taking the neighbourhood of each voxel into account. The data-term effectively defines the voxel-terminal node edge (t-link) weights, while the smoothness term defines the voxel-voxel node edge (n-link) weights of the graph.

In [10], automatic kidney segmentation in DCE-MRI is performed using a graph-cuts approach, following registration, to minimise an energy functional which takes into account shape in addition to intensity and neighbourhood information. Training shapes are aligned using registration, and the variability modelled using a distance probabilistic model estimated by a Poisson distribution. The intensity term is modelled by a Gaussian distribution, with parameters estimated using the EM algorithm, and the neighbourhood term is modelled by a MRF.

Non-Bayesian Methods Non-Bayesian approaches have also been used in conjunction with atlas-based methods. For instance in [14], fuzzy connectedness segmentation is initialised by the atlas-based registration for abdominal organs. In [32], the EM algorithm is extended with

level-set evolution to allow for tumours which are not necessarily present in the atlas. Kidney segmentation in DCE-MR images is performed in [9, 11] using an atlas-based approach followed by level-set evolution. The atlas in [11] is averaged from a set of images and represented as a binary image containing the kidney and a signed distance map density function. The energy functional to be minimised penalises deviations from intensity and distance map distributions, both estimated using a modified EM algorithm.

Features other than raw intensity can be used in the classification phase of segmentation; gradient and scale-space Gaussian derivatives are used with a k-Nearest Neighbour (kNN) classifier in [22] to segment the liver from CT data.

Active-Shape Models Active shape models (ASM) have also been used in [25, 33, 27] following atlas-based procedures. In [25], the training shapes are initially aligned using atlas registration, but the final segmentation is carried out using ASM. A combined segmentation algorithm, which combines atlas-based segmentation and ASM is proposed in [33]. The method takes into account the hierarchical relationship between the liver, vena-cava and gall bladder when creating the atlas and multi-organ ASM, also using the ASM to refine the atlas-based segmentations in an iterative framework. A similar approach was proposed in [13]; in this method, the point-correspondence required to create the ASM is achieved using non-rigid registration to a reference image, while the actual segmentation involves fitting an associated gray-level appearance model to the novel image with regularisation by the ASM.

5.1.4 Contribution

The above segmentation methods have all been shown to be successful on various areas of anatomy and modalities. Our method focuses atlas-based segmentation combined with voxel-based classification on kidney and renal cortex segmentation, specifically in low-dose contrast-enhanced CT images.

The average atlas approach, when applied to organs with a lot of relative motion, such as the kidney, produces a blurred intensity atlas image, which loses vital intra-organ and boundary information. Our method makes use of a multi-atlas, adopting an optimal atlas selection strategy which uses cross-correlation between each atlas and the novel image to predict the segmentation performance via an atlas-specific linear regression model [34].

We then use the atlas-based segmentation to provide spatial priors for a MAP classification phase, obtaining Gaussian model parameters for the kidney and renal cortex voxel intensities using the EM algorithm.

The method described in [33] encodes inter-organ relationships in a probabilistic atlas and statistical shape model; our method also takes into account this constraint and uses a multi-organ atlas, consisting of the kidney and renal cortex. The renal cortex is wholly contained in the entire

kidney; this relationship is enforced using a diffeomorphic registration algorithm and also during an extended MAP classification phase by using the resulting renal cortex segmentation to refine the whole kidney segmentation. We are able to demonstrate that gross kidney segmentation is improved by applying inter-organ constraints in this manner.

Methods which have been successful in segmenting the kidney, such as [12, 13], require a degree of user-interaction. We present a fully-automatic method, with an initial coarse rigid registration phase to transform the atlas to the centre of the kidney in the novel image. Furthermore, our implementation has been optimised for multi-core CPUs and takes only 10 seconds to segment a pair of kidney and renal cortex on a DELL Precision T5500 machine, thus increasing its usability.

5.2 Method

5.2.1 Introduction

This chapter describes the method in detail, breaking the system up into components as shown in figure 5.4.

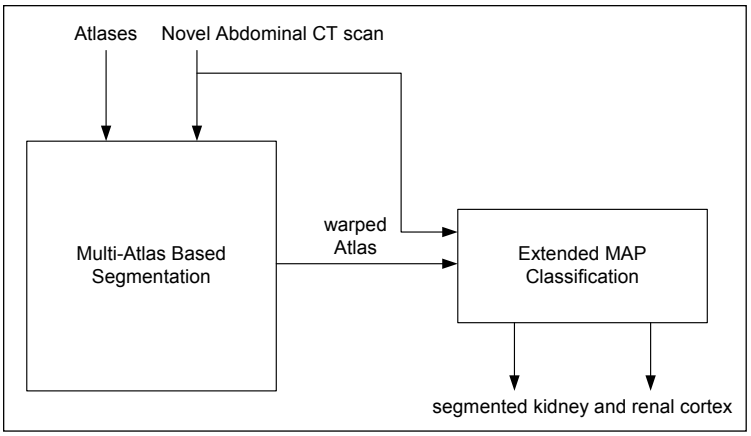

Figure 5.4: Method for Automatic Kidney and Renal Cortex Segmentation

5.2.2 Materials

Abdominal contrast-enhanced 4D CT studies from 24 subjects are used in this work, each study containing 20 phases. The scans were taken using Toshiba's Aquilion ONE 320-slice scanner, which covers a vertical range of 16cm, hence both kidneys in one gantry rotation.

The patients were scanned in the Feet-First-Supine (FFS) position following administration of an Iodine-based contrast agent, with an average X-ray beam intensity of $60mA$. Different

acquisition protocols were used however; 14 studies were acquired using a breath-hold protocol, while the remaining 10 were acquired using a shallow-breathing protocol.

Manual segmentations of each kidney and its renal cortex were collected on the early arterial phase for all 24 studies by clinically trained persons using the MPR manual shape selection tool in Toshiba's Voxar 3D advanced visualisation software, therefore four separate segmentations exist for each dataset. These are used as ground-truth to quantify the segmentation accuracy of the algorithm described below and as masks constituting the atlas.

5.2.3 Pre-processing

The algorithm is designed to segment the left kidney and its renal cortex independent of the right and vice-versa, therefore a few preparatory steps can be performed. Assuming the algorithm is to be run mainly on the type of data described in section 5.2.2, i.e. abdominal CT scans from the Aquilion ONE scanner, then it becomes straightforward to consistently split the images into two halves, each containing one kidney. This increases the speed and robustness of registration by limiting the search space.

The dynamic range of the images is also reduced prior to registration, and our experiments have demonstrated that clamping the intensities of each image to the range [0, 300]*HU* improves the segmentation accuracy. This reduction in dynamic range has the effect of increasing the saliency of the organs particularly in this sort of noisy data.

5.2.4 Multi-Atlas Based Segmentation

An atlas in the context of atlas-based segmentation refers to a pre-selected image with labelled anatomical regions of interest. Multi atlas-based segmentation requires a set of atlases. Each atlas in this case consists of a downscaled 3D intensity image constrained to the region containing either the left or right kidney with a little extra context, controlled by a dilation parameter d (set to 2mm) applied to the kidney mask. The intensity images are iteratively smoothed by a Gaussian filter and down-sampled to reduce noise and eliminate aliasing, while allowing for quicker registration. Two masks are held in each atlas, one representing the entire kidney, and the other representing the renal cortex. Figure 5.6 shows a typical right kidney atlas.

Assuming a set $A = \{A_0, ..., A_{M-1}\}$ of M atlases has been provided, segmentation of the required anatomical region in a novel image T begins by registering each $A_i (i = 0, ..., M-1)$ with T in order to produce a set of warp-fields $W = \{W_0, ..., W_{M-1}\}$, where W_i maps A_i onto T. The registration process used is explained in section 5.2.4.

Using the warp-field W_i, the cross-correlation between $W_i(A_i)$ and T can be measured. The image similarity between any $W_i(A_i)$ and T is not used directly to determine the segmentation accuracy, because this measure by itself is not directly comparable across all possible pairs of

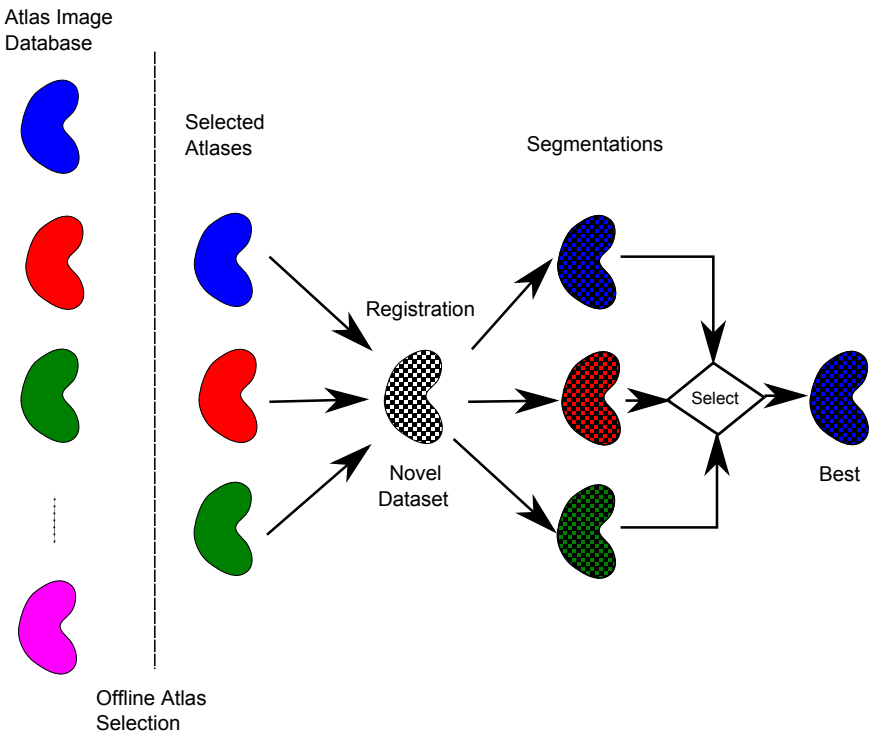

Figure 5.5: Multi Atlas-Based Kidney Segmentation

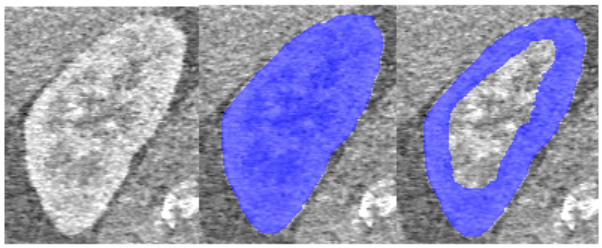

Figure 5.6: Right Kidney Atlas
Left to right: Constrained CT image, renal cortex mask, entire right kidney mask in sagittal MPR.

$W_i(A_i)$ and T. In order to get a more standard measure for selecting the best out of M possible segmentations, a linear regression model specific to each atlas is used. The function is of the form $y = mx + c$, where x is the value of the cross-correlation and y is then the predicted Jaccard overlap (equation 5.1); m and c are stored as part of the atlas.

The labelled anatomical region from the A_i that predicts the highest segmentation accuracy, is propagated onto T using W_i, resulting in a segmentation.

The set A is selected offline from all training data using the multi-atlas selection procedure described in [34].

$$Jaccard(A, B) = \frac{Vol.(A \cap B)}{Vol.(A \cup B)} \tag{5.1}$$

This value ranges between 0 and 1, with 1 representing a perfect segmentation.

Registration

The aim of registration is to find a mapping W between a reference R and a target image T, that minimises a pre-defined cost function. The registration process used here is a multi-scale approach, performing 2 passes of rigid registration, followed by 2 passes of non-rigid registration.

The rigid registration finds the best 3D translation, anisotropic scaling and roll, pitch and yaw components to transform the atlas image R onto the novel image T. The similarity metric chosen to drive the registration is normalised mutual information (NMI). Assuming image intensities are normally distributed, the joint entropy can be estimated from the determinant of the covariance matrix between the two overlapping image pairs [35], and a variance-based approach can be used to compute NMI.

$$\text{NMI}(R, T) = \frac{h(R) + h(T)}{h(R, T)},$$

where

$$\begin{aligned} h(R) &= \frac{1}{2} ln(2\pi e \sigma_R^2), \\ h(T) &= \frac{1}{2} ln(2\pi e \sigma_T^2), \\ h(R, T) &= \frac{1}{2} ln\{(2\pi e)^2 |\Sigma|\}, \end{aligned}$$

and

$$\Sigma = \begin{pmatrix} \sigma_R^2 & \sigma_{R,T}^2 \\ \sigma_{R,T}^2 & \sigma_T^2 \end{pmatrix}$$

Powell's optimisation method [36] is used to find the affine transformation matrix **A** that

maximises NMI between the transformed R and T.

$$\mathbf{A}^* = \underset{\mathbf{A}}{\text{argmax}}\{\text{NMI}(\mathbf{A}.R, T)\}$$

A rapid initial search phase is performed in order to align the atlas and novel images such that the centre of the atlas image, which is also the centre of the kidney, is aligned with the centre of the kidney in the novel image. It works by initialising the Powell search algorithm at evenly spaced points along the anterior-posterior direction in the novel volume (Figure 5.7), and returns the best transform resulting from all starting points.

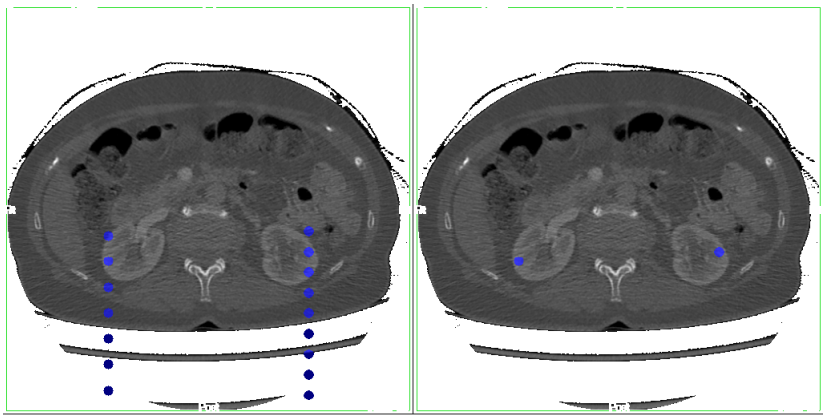

Figure 5.7: Rapid Initial Search
Powell's optimisation algorithm is initialised at each of the blue points in the image on the left. The best resulting transforms are shown on the image on the right. These points are used as the starting points for the full rigid registration.

This phase is designed to be fast rather than accurate, and so a number of compromises are made. The full rigid registration operates on images downscaled by a factor of 8 and then by a factor of 4, while the rapid initial search phase operates on images downscaled by a factor of 8 only. In the full rigid registration, 5000 voxel samples are taken from the overlapping region between R and T and used to compute NMI, but in the rapid initial search case, only 1000 voxel samples are taken. The affine components in the full rigid registration include anisotropic scaling, while the rapid initial search uses isotropic scaling, reducing the number of search dimensions by 2.

The non-rigid phase aims to find the warp-field W that maximises mutual information MI between T and the transformed reference R'. Non-rigid registration allows non-linear deformation of the image, making it an ill-posed problem, therefore some form of regularisation is required. Our method finds the W that maximises MI between R' and T within fluid and elastic constraints.

The warp-field is initialised from the final affine transformation in the rigid phase and evolved

until convergence is achieved, with the warp-field at iteration i given by:

$$W_i = E \otimes (W_{i-1} + kV \otimes F_{i-1}),$$

where F_i is the current force-field, i.e. the derivative of MI at each point in R computed from a joint intensity histogram using the method described in [37]. V is the viscous fluid constraint, modelled by a Gaussian filter, and E is the elastic constraint also modelled by a Gaussian filter.

5.2.5 Maximum a-Posteriori Classification

As indicated in figure 5.4, following the multi-atlas based segmentation, the atlas masks (whole kidney and cortex) are now available in the spatial domain of the novel image. The goal of this step is further refinement of these segmentations. An outline of the process is shown in figure 5.8.

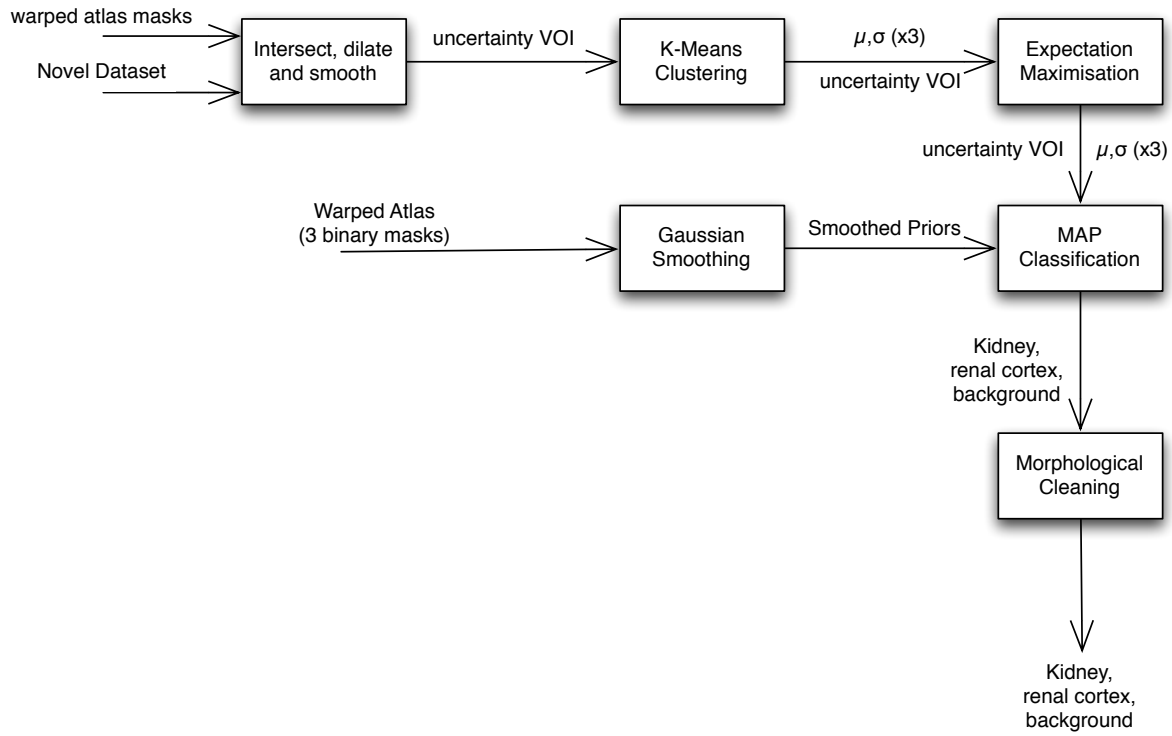

Figure 5.8: Voxel-Based Classification

Notations The label or segmentation field is denoted by $\mathbf{L} = \{l_0, ..., l_{N-1}\}$, the observed novel data by $\mathbf{I} = \{i_0, ..., i_{N-1}\}$, and the warped atlas by $\mathbf{A} = \{a_0, ..., a_{N-1}\}$, where N is the total number of voxels. Elements of $\mathbf{L}$, $\mathbf{I}$, and $\mathbf{A}$ are arranged by spatial position within a 3-D grid, and denoted by $x \in \Re^3$. The sample space of $\mathbf{L}$ is denoted by $\Omega_l = \{l : l_x \in \{1, 2, 3\}, \forall x\}$. Labels 1, 2, and 3 are cortex, medulla, and background, respectively.

The warped atlas masks should provide spatial probabilities for the relevant tissues, indicating $P(l_x = k|x)$, for $k = 1, 2, 3$. The atlas contains renal cortex and whole kidney masks however, and so the medulla mask can be obtained by a simple set difference.

$$\text{medulla} = \text{wholekidney} \setminus \text{cortex}$$

The warped atlas **A** can now be interpreted as vector-valued, with $a_x = \{a_{x1}, a_{x2}, a_{x3}\}$ and the information as

$$P(l_x = k) = a_{xk},$$

$$a_{x3} = 1 - \sum_{k=1}^{k=2} a_{xk}$$

The selection of the optimal atlas in the atlas-based segmentation phase implies binary spatial priors, i.e. $P(l_x = k) \in \{0, 1\}$. To account for segmentation errors in the registration stage and provide a range of probabilities between 0 and 1, **L** is extended spatially by morphological dilation of **A** by a spherical structuring element of radius $r$mm (set to 2mm) and **A** is smoothed accordingly by a Gaussian filter with a standard deviation of r.

Estimating Tissue Model Parameters The elements of **I** are smoothed with a 1.5 pixel s.d. Gaussian kernel and k-means clustering is applied to find 3 cluster means corresponding to the named tissue types. These tissue mean intensities are used to initialise a standard EM algorithm, which in this case assumes normally distributed intensities and models the Gaussian class-specific intensity distributions for 3 tissue types as shown in figure 5.9. This assumption is based on the renal cortex and medulla having almost homogeneous intensities in contrast-enhanced CT.

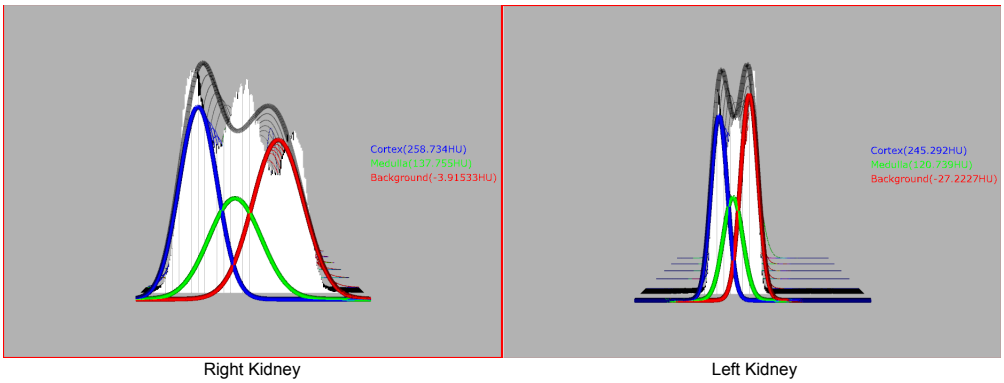

Figure 5.9: Gaussian-modelled Intensity Distributions
The underlying intensity histogram is shown above in white, while the coloured curves refer to the Gaussian-modelled intensity distributions of each of the 3 tissues types. Left to right: Background, Medulla, Renal Cortex. Multiple graphs in the background are for each iteration of the EM algorithm.

Bayesian Framework Assuming the observed random variable I, i.e. $i_0, i_1, ..., i_{N-1}$ are conditionally independent given the true segmentation L and that the probability distribution of each i_i given L depends only on l_i, then Bayes' theorem can be applied to find the probability of an element of **L** truly belonging to a tissue type k, given the observed data **I**.

$$P(l_x = k|i_x) = \frac{P(i_x|l_x = k)P(l_x = k)}{P(i_x)}, \tag{5.2}$$

where $P(i_x|l_x = k)$ is provided by the EM algorithm, using its estimated means, μ_k and variances, σ_k^2 of the assumed Gaussian intensity distributions for each tissue class.

$$P(i_x|l_x = k) = \frac{1}{\sqrt{2\pi\sigma_k^2}} exp(\frac{-(i_x - \mu_k)^2}{2\sigma_k^2}) \tag{5.3}$$

The unconditional probability of observing the data, $P(i_x)$ is independent of k by definition, thus the term becomes absorbed into the constant of proportionality, and equation 5.2 becomes:

$$P(l_x = k|i_x) \propto P(i_x|l_x = k)P(l_x = k), \tag{5.4}$$

and

$$\sum_k P(l_x = k|i_x) = 1$$

Voxel Classification Equation 5.4 can be applied to every voxel within the dilated warped atlas domain to compute the posterior probability of its true label being either cortex, medulla, or background given its value and position within the image. In order to enforce the hierarchical relationships and hence segment the entire kidney and the renal cortex, a new segmentation field **L'** is required, with sample space $\Omega_{l'} = \{l' : l'_x \in \{1, 2, 3\}, \forall x\}$. Labels 1, 2, and 3 in this case are cortex, kidney, and background, respectively. Labels are assigned to **L'** as follows:

- **Kidney:**

$$l'_x = 2 \textbf{ if } (P(l_x = 1|i_x) + P(l_x = 2|i_x) + \alpha \geq 0.5)$$

- **Cortex $\subset$ Kidney:**

$$l'_x = 1 \textbf{ if } ((l'_x = 2) \textbf{ and } P(l_x = 1|i_x) + \beta > P(l_x = 2|i_x)),$$

where α and β are user-specified to control the inner and outer extent of the cortex and kidney respectively.

Post-processing The steps taken during the voxel classification phase, described above, produce segmentations of the entire kidney and of the renal cortex. These segmentations may

contain spikes and holes and need to be cleaned.

Applying a morphological opening, followed by a fill to the whole kidney segmentation, reduces spikes and fills any holes that may have arisen as a result of the voxel classification. The renal cortex is normally a rather thin domain, and so a morphological closing followed by a fill is applied to the segmentation instead. Connected component analysis is then carried out on both segmentations to obtain the largest connected component, and the cortex is further refined to remove spikes by

$$\text{cortex} = \text{cortex} \cap \text{kidney}$$

This ensures that the cortex is wholly contained within the kidney as naturally dictated by the anatomy.

5.3 Results and Discussion

5.3.1 Introduction

The method described in the previous section was implemented in C++ and deployed on a Dell Precision T5500 windows machine containing a 2GHz quad-core Intel E5504 Xeon CPU and 6GB of RAM.

This section describes the validation method employed and the results obtained.

5.3.2 Segmentation of the Kidney and Renal Cortex from CT Data

Both the left and right kidneys with renal cortices were segmented in 22 cases. The segmentations were performed on the study phase that captured the early arterial phase of contrast perfusion. This phase was chosen because it sufficiently highlights the renal cortex relative to the medulla.

The best set of 2 atlas images from 24 images was selected using the method described in [34], while the remaining 22 images made up the validation set. The accuracy was quantified by measuring the overlap between the generated segmentation and the manual segmentation, using the Jaccard overlap (equation 5.1). The sensitivity and specificity were also quantified using the true positive fraction (TPF) and positive predictive value (PPV) respectively:

$$\text{TPF} = \frac{|A \cap B|}{|B|} \tag{5.5}$$

$$\text{PPV} = \frac{|A \cap B|}{|A|}, \tag{5.6}$$

where A is the generated segmentation and B is the ground-truth segmentation.

Subjective evaluation was also carried out by a clinical associate in order to assess its usefulness in a hospital setting. This evaluation was performed on 12 unseen cases, where no

ground-truth segmentation had previously been collected. An opinion score, ranging from 0 to 10 (10 being perfect) was assigned to each segmentation (left kidney, right kidney, left renal cortex, right renal cortex). The mean opinion scores (MOS) are shown in figure 5.10.

Figure 5.10 shows the objective test results averaged across 22 datasets. Ground-truth kidney segmentations are available for all 22 cases, but renal cortex ground-truth is only available for 8 of these cases. The cortex accuracy is lower than that of the entire kidney in both left and right cases, and the definition of the ground-truth accounts somewhat for this. The cortex ground-truth segmentation does not include any of the cortical columns, but the segmentation algorithm by design includes an extent of the cortical columns, controllable by a parameter α as described in section 5.2.5.

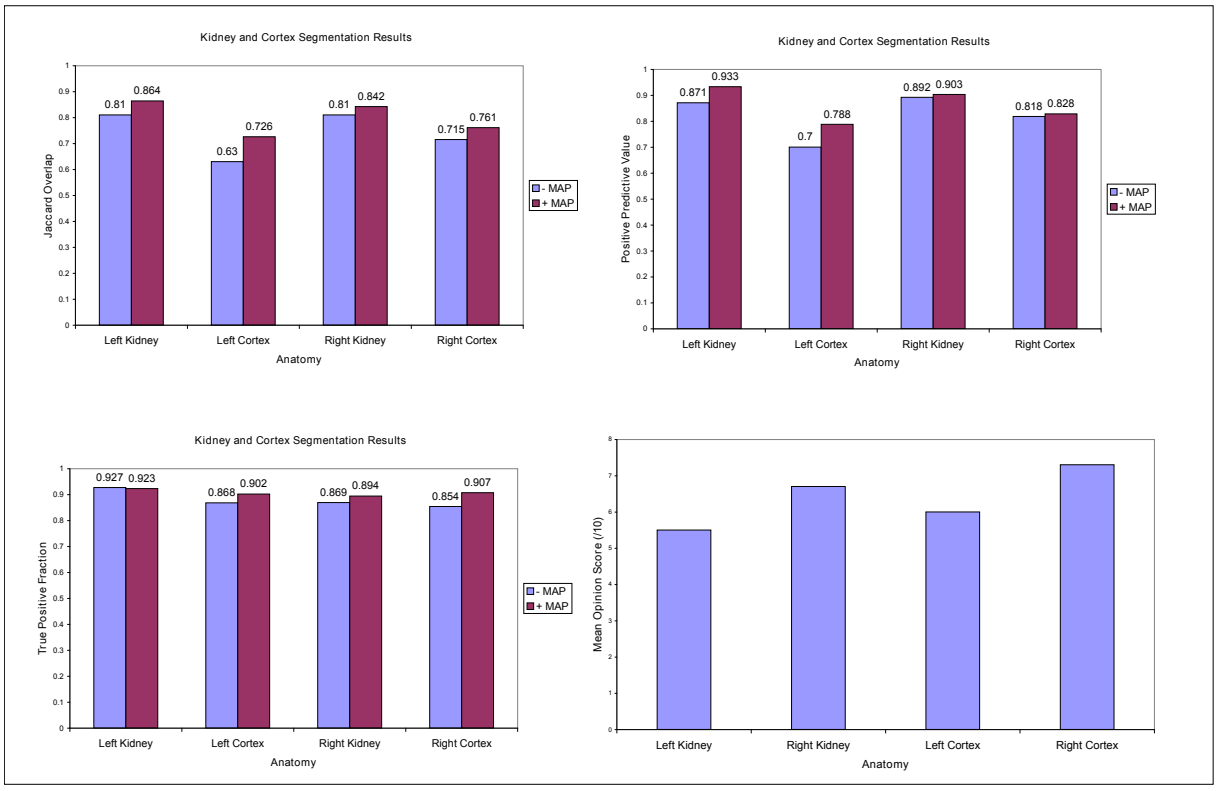

Figure 5.10: Kidney And Cortex Segmentation Results, $\alpha = 0$
Run-time for cortex segmentation is not available, as cortex segmentation is an integral part of the entire kidney segmentation. The results are shown for the inclusion and exclusion of the voxel classification phase.

5.3.3 Effect of Atlas Selection

The offline atlas selection method proposed in [34] is used to select an optimal set of atlases, $A = \{a_0, ..., a_{M-1}\}$ based on the training data D available.

$$A_{opt} = \underset{A \in D^M}{\operatorname{argmax}} \frac{1}{N-M} \sum_{i=1, d_i \notin A}^{n} J'(d_i, A), \tag{5.7}$$

where $N = |D|$ is the number of training images available,

$$J'(d_i, A) = J(d_i, \underset{a \in A}{\operatorname{argmax}}\, J^*(d_i, a)), \text{and} \tag{5.8}$$

$$J^*(d_i, a) = m_a.\mathrm{CC}(d_i, a) + c_a \tag{5.9}$$

$J(x, y)$ in this case is the resulting Jaccard overlap between the segmented kidney in image x and its ground-truth, while using atlas image y to perform the segmentation. m_a and c_a are the linear regression parameters stored as part of the atlas. $CC(x, y)$ is the cross-correlation between the atlas image y and novel image x after transformation to the same spatial frame. Figure 5.10 shows segmentation performance while using the optimal set of 2 atlases in the multi-atlas framework.

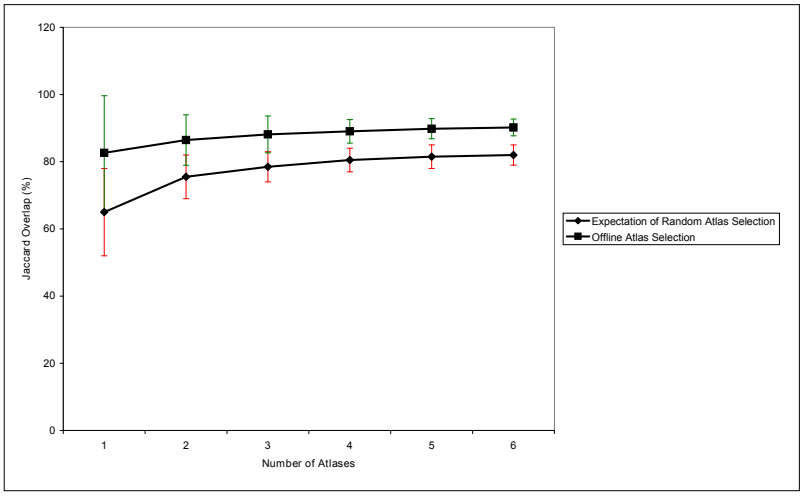

Figure 5.11: Effect of the Number of Atlases on Kidney Segmentation
The upper line shows the Jaccard overlap averaged over all validation data when selecting the optimal atlases for each set size (equation 5.7), while the lower line shows the mean of the Jaccard overlap averaged over all training data for all possible random selections of atlases (equation 5.10).

Figure 5.11 demonstrates that the segmentation accuracy of the entire kidney increases as the number of atlases in the multi-atlas set increases. The figures also demonstrate the added benefit of using the offline atlas selection to select the set of atlases (equation 5.7), as indicated by the

upper lines. The lower line refers to the expectation E of the overlaps averaged over all training datasets for a multi-atlas set selected at random:

$$E = \frac{1}{\binom{|D|}{|M|}} \sum_{A \in D^M} [\frac{1}{N-M} \sum_{i=1, d_i \notin A}^{N} J'(d_i, A)] \tag{5.10}$$

What the figures do not show, is the impact on run-time caused by adding an extra atlas to the set. In this implementation, the atlas-novel image registrations are performed sequentially; this sets a trade-off between potential accuracy and practicality of the application. In our case:

$$\textit{Runtime} \approx 7 + 2M$$

The choice of 2 atlases arose as a compromise for the execution time of the segmentation, as ≈ 10 seconds has been deemed acceptable by our clinical acceptance team.

5.4 Conclusion and Further Work

5.4.1 Conclusion

This paper has proposed an atlas-based method of segmenting the kidneys and renal cortices from low-dose contrast enhanced abdominal CT scans with average accuracies of 85.3% and 74.3% respectively.

The results in figure 5.10 show an average increase in accuracy of 4.3% for the kidneys and 7.3% for the renal cortices gained by using the unsupervised voxel-based classification technique following registration. Figure 5.11 points out clear advantages to using our optimal atlas selection approach (equation 5.7); both at the offline stages, where it is used to select a set of atlases to be used for registration, and during the actual application, where it is used to select the best segmentation from a list of candidates.

The proposed method is fully automatic, requiring no user interaction; hence it has potential to be deployed on scanner workstations as part of pre-processing software for kidney analysis. Furthermore, its current implementation takes only 10 seconds to segment each kidney/renal cortex pair, making it an acceptable alternative to manual segmentation or at the very least, a starting-off point which can be adjusted by the radiologist as required.

Limitations of this method lie mainly in its overall accuracy. As shown by the mean opinion scores in figure 5.10, the left kidney and left cortex scored 5.5 and 6.0 out of 10 respectively. Higher scores on unseen data is desirable for an application intended for clinical setting. In order to achieve this using the current approach, the registration phase has to be improved upon.

5.4.2 Further Work

One possible avenue for improvement is in the rapid initial search phase of the registration, whereby the atlas is placed at the centre of the kidney as a starting point for full multi-scale rigid registration. The current method searches from positions along the anterior-posterior line in the middle of each half-image, as shown in figure 5.7. This can be extended to include search points evenly spaced within a cuboid centred at the middle of the half-image. Another approach involves making use of the available training data and sampling search points within the distribution of points marking the centre of the kidney in each training image.

The atlas-supplied spatial priors are representative of the more similar atlas image to the novel image at the moment. For this reason, we apply a Gaussian smoothing to the domain represented by the atlas and obtain fuzzy prior probabilities in the process. This smoothing step can be avoided if the atlas mask was indeed a probability mask made up from several organ masks spatially normalised to the atlas image. It has the added advantage of being more representative of anatomical variations than our current mask.

The voxel-based classification step applies k-means clustering to estimate the initial means, after which the expectation maximisation algorithm is employed to discover the optimal parameters of each distribution by optimising the likelihood (equation 5.3). A more elegant and possibly accurate solution to this problem would be to eliminate the dependency on k-means clustering by obtaining the initial means of each tissue type from the warped atlas masks and feed the EM algorithm with prior probabilities of observing each tissue type at a particular voxel position, thereby optimising the expectation of the posterior probability instead[38].

The algorithm currently takes 10 seconds when using 2 atlases and rises by $\approx 2secs$ for every extra atlas added to the multi-atlas set. However, there is no reason why the registrations are not performed in parallel; performing each atlas-novel image registration concurrently will reduce the execution-time costs associated with extra atlases.

References

[1] Health Resources and Services Administration, “OPTN/SRTR Annual Report,” 2008.

[2] NHS Blood And Transplant, “Transplant activity in the UK 2008-2009,” 2009.

[3] A. Nadalo, Lennard, M. Dickerman, Richard, M. Slonim, Suzanne, and R. Brinker, Karl, “Kidney transplantation, surgical complications: Imaging,” 2010.

[4] V. S. Lee, H. Rusinek, M. E. Noz, P. Lee, M. Raghavan, and E. L. Kramer, “Dynamic three-dimensional MR renography for the measurement of single kidney function: initial experience.,” *Radiology*, vol. 227, pp. 289–94, April 2003.

[5] R. K. Sharma, R. K. Gupta, H. Poptani, C. M. Pandey, R. B. Gujral, and M. Bhandari, "The magnetic resonance renogram in renal transplant evaluation using dynamic contrast-enhanced mr imaging.," *Transplantation*, vol. 59, no. 10, pp. 1405–9, 1995.

[6] Radiological Society of North America, "Safety : Radiation Exposure in X-ray Examinations," pp. 1–4, 2009.

[7] N. Hackstein, H. Cengiz, and W. S. Rau, "Contrast media clearance in a single kidney measured on multiphasic helical CT: results in 50 patients without acute renal disorder.," *AJR. American journal of roentgenology*, vol. 178, pp. 111–8, January 2002.

[8] F. G. Sommer, "Science to Practice : Can Single-Kidney Glomerular Filtration Rate Be Determined with Contrast-enhanced CT ?," *Radiology*, vol. 242, no. 2, pp. 325–326, 2007.

[9] H. Abdelmunim, A. Farag, W. Miller, and M. AboelGhar, "A kidney segmentation approach from DCE-MRI using level sets," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, 2008. CVPR Workshops 2008*, no. iii, p. 1âĂŞ6, 2008.

[10] A. M. Ali, A. A. Farag, and A. S. Ell-Baz, "Graph cuts framework for kidney segmentation with prior shape constraints.," *Medical image computing and computer-assisted intervention : MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention*, vol. 10, pp. 384–92, January 2007.

[11] S. E. Yuksel, A. El-baz, and A. A. Farag, "A Kidney Segmentation Framework for Dynamic Contrast Enhanced Magnetic Resonance Imaging," *Journal of Vibration and Control*, vol. 13, pp. 1505–1519, 2007.

[12] H. Park, P. Bland, and C. Meyer, "Construction of an abdominal probabilistic atlas and its application in segmentation," *IEEE Trans. Med. Imaging*, vol. 22, no. 4, pp. 483–492, 2003.

[13] M. Spiegel, D. A. Hahn, V. Daum, J. Wasza, and J. Hornegger, "Segmentation of kidneys using a new active shape model generation technique based on non-rigid image registration.," *Computerized medical imaging and graphics : the official journal of the Computerized Medical Imaging Society*, vol. 33, no. 1, pp. 29–39, 2009.

[14] Y. Zhou and J. Bai, "Multiple Abdominal Organ Segmentation: An Atlas-Based Fuzzy Connectedness Approach," *IEEE Transactions on Information Technology in Biomedicine*, vol. 11, no. 3, pp. 348–352, 2007.

[15] T. Rohlfing, R. Brandt, C. R. Maurer Jr., and R. Menzel, "Bee brains, b-splines and computational democracy: Generating an average shape atlas," *Mathematical Methods in Biomedical Image Analysis, IEEE Workshop on*, vol. 0, p. 187, 2001.

[16] T. Rohlfing, R. Brandt, R. Menzel, D. B. Russakoff, and C. R. Maurer Jr., "Quo vadis, atlas-based segmentation?," in *The Handbook of Medical Image Analysis – Volume III: Registration Models* (J. Suri, D. L. Wilson, and S. Laxminarayan, eds.), ch. 11, pp. 435–486, New York, NY: Kluwer Academic / Plenum Publishers, Aug 2005.

[17] M. Wu, C. Rosano, P. Lopez-Garcia, C. Carter, and H. Aizenstein, "Optimum template selection for atlas-based segmentation," *NeuroImage*, vol. 34, no. 4, p. 1612âĂŞ1618, 2007.

[18] P. Aljabar, R. Heckemann, A. Hammers, J. Hajnal, and D. Rueckert, "Multi-atlas based segmentation of brain images: Atlas selection and its effect on accuracy," *NeuroImage*, vol. 46, no. 3, pp. 726–738, 2009.

[19] X. Artaechevarria, A. Munoz-Barrutia, and C. Oritz-de Solorzano, "Combination Strategies in Multi-Atlas Image Segmentation: Application to Brain MR Data.," *IEEE transactions on medical imaging*, vol. 28, no. 8, pp. 1266–1277, 2009.

[20] H. Kirisli, M. Schaap, S. Klein, L. Neefjes, A. Weustink, T. van Walsum, and W. Niessen, "Fully automatic cardiac segmentation from 3d cta data: a multiatlas based approach," in *SPIE Medical Imaging 2010*, 2010.

[21] J. Lötjönen, R. Wolz, J. Koikkalainen, L. Thurfjell, G. Waldemar, H. Soininen, and D. Rueckert, "Fast and robust multi-atlas segmentation of brain magnetic resonance images," *NeuroImage*, 2009.

[22] E. van Rikxoort, Y. Arzhaeva, and B. van Ginneken, "Automatic segmentation of the liver in computed tomography scans with voxel classification and atlas matching," in *3D Segmentation In The Clinic: A Grand Challenge* (T. Heimann, M. Styner, and B. van Ginneken, eds.), pp. 101–108, 2007.

[23] T. Rohlfing, D. B. Russakoff, and C. R. Maurer Jr., "Performance-based classifier combination in atlas-based image segmentation using expectation-maximization parameter estimation," *IEEE Trans. Med. Imag*, vol. 23, pp. 983–994, 2004.

[24] S. K. Warfield, K. H. Zou, and W. M. Wells, "Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation," *Medical Imaging, IEEE Transactions on*, vol. 23, pp. 903–921, 2004.

[25] A. Frangi, W. Niessen, D. Rueckert, and J. Schnabel, "Automatic 3D ASM construction via atlas-based landmarking and volumetric elastic registration," in *Information Processing in Medical Imaging*, vol. 2082, p. 78âĂŞ91, Springer, 2001.

[26] T. Rohlfing and C. R. Maurer, "Shape-based averaging.," *IEEE transactions on image processing : a publication of the IEEE Signal Processing Society*, vol. 16, pp. 153–61, January 2007.

[27] G. Vesom, N. Cahill, L. Gorelick, and J. Noble, "Characterization of Anatomical Shape Based on Random Walk Hitting Times," *MFCAâĂŹ08*, p. 117.

[28] M. Lorenzo-Valdes, D. Rueckert, R. Mohiaddin, and G. Sanchez-Ortiz, "Segmentation of cardiac MR images using the EM algorithm with a 4D probabilistic atlas and a global connectivity filter," *Proceedings of the 25'h Annual Intemational Conference of the IEEE EMBS*, no. Lv, pp. 626–629, 2003.

[29] V. Kolmogorov and R. Zabih, "What energy functions can be minimized via graph cuts?," *IEEE transactions on pattern analysis and machine intelligence*, vol. 26, pp. 147–59, February 2004.

[30] F. V. D. Lijn, T. D. Heijer, M. M. Breteler, and W. Niessen, "Hippocampus segmentation in MR images using atlas registration, voxel classification, and graph cuts," *NeuroImage*, vol. 43, pp. 708–720, 2008.

[31] R. Wolz, P. Aljabar, D. Rueckert, R. A. Heckemann, and A. Hammers, "Segmentation of subcortical structures in brain MRI using graph-cuts and subject-specific a-priori information," *IEEE Internaional Symposium on Biomedical Imaging - From Nano to Macro*, pp. 470–473, 2009.

[32] G. Gerig, N. Moon, S. Ho, and E. Bullitt, "Model-Based Segmentation of Brain Tissue and Tumour," in *Annu. Int. Conf. IEEE Eng. Med. Biol.*, 2002.

[33] T. Okada, K. Yokota, M. Hori, M. Nakamoto, H. Nakamura, and Y. Sato, "Construction of hierarchical multi-organ statistical atlases and their application to multi-organ segmentation from CT images.," *Medical image computing and computer-assisted intervention : MICCAI ... International Conference on Medical Image Computing and Computer-Assisted Intervention*, vol. 11, pp. 502–9, January 2008.

[34] A. Akinyemi, I. Poole, J. Piper, C. Plakas, and C. Roberts, "Optimal Atlas Selection Using Image Similarities in a Trained Regression Model to Predict Performance.," tech. rep., Toshiba Medical Visualisation Systems, Edinburgh, 2009.

[35] N. Ahmed and D. Gokhale, "Entropy expressions and their estimators for multivariate distributions," *IEEE Transactions on Information Theory*, vol. 35, no. 3, p. 688âĂŞ692, 1989.

[36] J. Nocedal and S. J. Wright, *Numerical Optimization*. Springer, August 2000.

[37] W. R. Crum, D. L. G. Hill, and D. J. Hawkes, "Information theoretic similarity measures in non-rigid registration," in *IPMI*, pp. 378–387, 2003.

[38] M. Murgasova, D. Rueckert, D. Edwards, and J. Hajnal, "Robust segmentation of brain structures in MRI," *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pp. 17–20, June 2009.

Chapter 6

Robust Segmentation using Atlas-Based Priors in the EM Algorithm.

Abstract *This chapter presents a segmentation method combining registration and intensity based classification within an atlas framework. The registration makes use of a multi-atlas, with the most similar atlas to the subject data selected after registration to provide an initial segmentation. The Expectation Maximisation algorithm is used to perform intensity-based classification in order to fine-tune the detail of the initial segmentation. This modified EM algorithm is equipped with spatial priors derived from the initial segmentation to give it anatomical information necessary for segmenting multiple structures in the subject data.*

Evaluation of this combined approach to segmentation demonstrates that the method is indeed robust, as results from applying the method to kidney and renal cortex segmentation from CT data with low SNR imply. The method was also applied to the segmentation of seven sub-structures within the heart on CCTA data to demonstrate its effectiveness at segmenting multiple structures displaying similar intensities. The main advantage of this method lies in its versatility; the algorithm was able to tackle both presented problems with only a change of atlas required.

The results presented in this paper demonstrate that using the atlas-based priors in the EM algorithm not only improves segmentation accuracy over methods using registration only or the conventional EM algorithm, but it also provides an elegant and general framework for performing segmentation.

6.1 Introduction

6.1.1 Motivation

The task of segmenting anatomical structures from medical images is a challenging one, with a few categories of methods designed to tackle particular anatomical structures from specified imaging modalities. Atlas-based registration methods are popular because atlases provide specific anatomical knowledge to the segmentation process in a straightforward manner, making them widely applicable to many large structures within the human anatomy.

Intensity-based classification methods are used to fine-tune the segmentations provided by the atlas-based registration. A robust method to achieve this goal is the Expectation Maximisation (EM) algorithm [1, 2]. Being an unsupervised classifier, it relies purely on the subject data and is proven to converge on a segmentation that maximises the likelihood of observing the subject data given a set of tissue models [3]. Furthermore, it requires no training, making it unbiased to a particular population demographic.

The limitations of the EM classification arise when the task involves segmentation of a large number of anatomical structures of varying sizes. The challenge here is that, with a large number of structures to segment, the chances of each structure exhibiting similar intensity distributions is increased. It becomes impossible for any automatic segmentation algorithm to correctly identify all of these structures without any prior anatomical information.

6.1.2 Aim

The aim of this paper is to present a method for supplying the EM algorithm with prior anatomical information through the use of an atlas. It demonstrates the versatility of a combined segmentation algorithm that employs this tactic by segmenting the kidneys and renal cortices from low-dose abdominal CT data, and a set of sub-structures within the heart from cardiac CT angiography data.

6.1.3 State of the Art

This paper is based on work done in [4]. In [4], the EM algorithm is fed with spatial priors provided by a probabilistic atlas made up by registering a set of candidate atlases onto a template atlas. In the expectation phase of the EM algorithm, the tissue likelihood at each voxel is weighted by the prior probability of observing the tissue at that point, obtained from the probabilistic atlas. This extended algorithm was used to segment up to eleven sub-structures in the brain from MRI data, and was reported to be more accurate than both purely registration-based segmentation and EM-based segmentation without anatomical priors.

Our method applies this principle to both noisy abdominal CT and cardiac CT angiography data. In our paper, we achieve fast robust atlas-based registration by using a multi-atlas followed by optimal atlas selection based on image similarity [5] rather than combining candidate segmentations as performed in [4, 6]. The optimal atlas selection ensures that the anatomical priors provided to the EM algorithm are biased towards the subject data.

6.2 Method

6.2.1 Applying Prior Information to EM

The aim of EM-based segmentation is to assign to each voxel in the observed data, a tissue label that maximises the probability of observing its intensity given estimated tissue distribution models, measured by the total log-likelihood.

$$\ln p(\mathbf{X}|\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\sigma}) = \sum_{i=0}^{N-1} \ln \sum_{k=0}^{K-1} p_k p(x_i|\mu_k, \sigma_k) \tag{6.1}$$

It performs this in an iterative manner, estimating the tissue model parameters and computing the resultant overall log-likelihood at each step.

The process is split into two phases, and expectation (E-) step and a maximisation (M-) step. Assuming intensity homogeneity within the respective tissues, which is mostly the case in human anatomy, then the tissue intensity distributions can be modelled as Gaussian distributions with parameters σ and μ (standard deviation and mean, respectively).

The probability of a voxel, v_i in image $\mathbf{X}$, belonging to a tissue k as described by (σ_k, μ_k) at iteration n is $p(k|x_i)^{(n)}$. Therefore, given initial estimates of $\sigma_k^{(0)}, \mu_k^{(0)}$ for all K tissues, the E-step at iteration n computes the posterior probability $p(k|x_i)$ for each pixel given $\sigma_k^{(n)}, \mu_k^{(n)}$ as follows:

E-Step:

$$p(k|x_i)^{(n)} = \frac{p(x_i|\sigma_k^{(n)}, \mu_k^{(n)}) p_k^{(n)}}{\sum_{k=0}^{K-1} p(x_i|\sigma_k^{(n)}, \mu_k^{(n)}) p_k^{(n)}} \tag{6.2}$$

The M-step computes the parameters $\sigma_k^{(n+1)}$, $p_k^{(n+1)}$ and $\mu_k^{(n+1)}$ that maximise equation 6.1, setting the first derivative of $\ln p(\mathbf{X}|\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\sigma})$ with respect to each of these terms to zero.

M-Step:

$$\mu_k^{(n+1)} = \frac{\sum_{i=0}^{N-1} p^{(n)}(k|x_i) x_i}{\sum_{i=0}^{N-1} p^{(n)}(k|x_i)} \tag{6.3}$$

$$\sigma_k^{(n+1)} = \sqrt{\frac{\sum_{i=0}^{N-1} p^{(n)}(k|x_i)(x_i - \mu_k^{(n+1)})^2}{\sum_{i=0}^{N-1} p^{(n)}(k|x_i)}} \tag{6.4}$$

$$p_k^{(n+1)} = \frac{1}{N} \sum_{i=0}^{N-1} p^{(n)}(k|x_i) \tag{6.5}$$

The E- and M-step are then iterated until convergence, i.e.

$$\frac{\ln p^{(n+1)}(\mathbf{X}|\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\sigma}) - \ln p^{(n)}(\mathbf{X}|\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\sigma})}{\ln p^{(n)}(\mathbf{X}|\mathbf{p}, \boldsymbol{\mu}, \boldsymbol{\sigma})} > T,$$

where T is a user-defined threshold.

In equation 6.1, the term p_k is the prior probability of observing the tissue type k or the mixing probability. The conventional EM algorithm defines it as equivalent to the proportion of the entire data assigned to k (equation 6.5). This is rephrased for our purposes as the prior probability of observing the tissue type k at voxel v_i. This term can be appropriately substituted for the output of atlas-based segmentation, p_{ik}^{atlas}, following propagation of the atlas masks onto the subject data and can be held as constant over all iterations, biasing the classification of the data towards anatomical structures demarcated by the atlas.

6.2.2 Combination with Multi-Atlas Based Segmentation

An atlas, in the context of atlas-based segmentation, is an image (or image volume in the 3D case) with an associated set of manually delineated structures of interest. A multi-atlas is a set $A = \{a_0, ..., a_{M-1}\}$ of such atlases.

Figure 6.1 depicts the multi-atlas based segmentation approach as described in detail in [7].

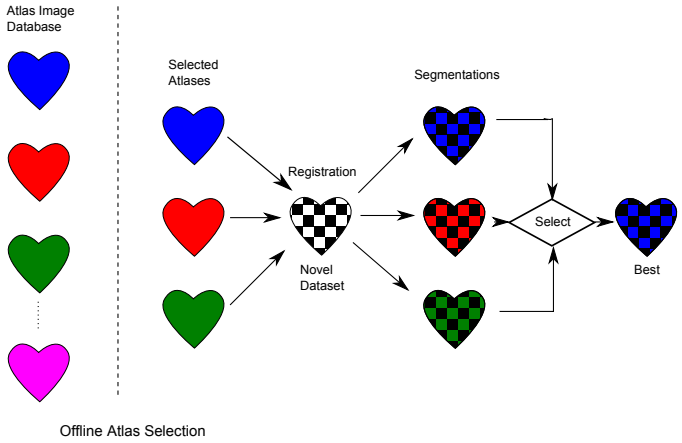

Figure 6.1: Multi-Atlas Based Segmentation

The atlases in use in the application are selected offline from a database of image atlases. Each of the chosen atlases produces a segmentation when aligned with the novel image, therefore a mechanism is required, which is capable of selecting the best segmentation.

The best segmentation is obtained from the registered atlas image, a^* which predicts the highest Jaccard overlap using an image similarity metric in a linear regression framework.

$$a^* = \underset{a}{\text{argmax}}\{m_a.\text{SIM}(X, a) + c_a\},$$

where X is the novel image; m_a and c_a are atlas specific linear regression slope and intercept respectively, obtained through an offline training process [5].

The end result of this step is a set of binary mask overlays on the novel image for each anatomical structure of interest. To account for inaccuracies in the registration, these binary masks are Gaussian smoothed over an expanded area to obtain soft probabilities p_{ik}^{atlas}, where $0 \leq p_{ik} \leq 1$ and $\sum_{k=0}^{K-1} p_{ik} = 1$.

p_{ik}^{atlas} is the atlas-based probability of observing tissue type k at spatial position i within X. This softening of the segmentation allows the EM classifier to refine the detail of the segmented structures, particularly close to the edge voxels. The initial Gaussian parameters $\mu_k^{(0)}$ and $\sigma_k^{(0)}$ are obtained from any voxel in X with $p_{ik}^{atlas} \geq 0.8$.

6.3 Experiments

6.3.1 Introduction

The combined multi-atlas segmentation and EM classification algorithm discussed in chapter 6.2 was used to segment anatomical structures from CT data in two different scenarios. This chapter describes both segmentation tasks and the results are presented in chapter 6.4.

6.3.2 Kidney and Renal Cortex Segmentation

The aim of this task is to segment the kidneys and their renal cortices from low-dose abdominal perfusion CT scans (figure 6.2).

Abdominal contrast-enhanced 4-D CT studies from 24 subjects were used in this work, each study containing 20 phases. The scans were taken using Toshiba's Aquilion ONE 320-slice scanner, which covers a vertical range of 16cm, hence both kidneys in one gantry rotation.

The patients were scanned in the FFS position following administration of an Iodine-based contrast agent, with an average X-ray beam intensity of $60mA$. Different acquisition protocols were used however; 14 studies were acquired using a breath-hold protocol, while the remaining 10 were acquired using a shallow-breathing protocol.

Manual segmentations of each kidney and its renal cortex were collected on the early arterial phase for all 24 studies by clinically trained persons using the MPR manual shape selection tool in Toshiba's Voxar 3D advanced visualisation software, implying four separate segmentations per dataset. These were used as ground-truth to quantify the segmentation accuracy of the algorithm and as masks constituting the atlas. The multi-atlas used in this experiment was made up of two

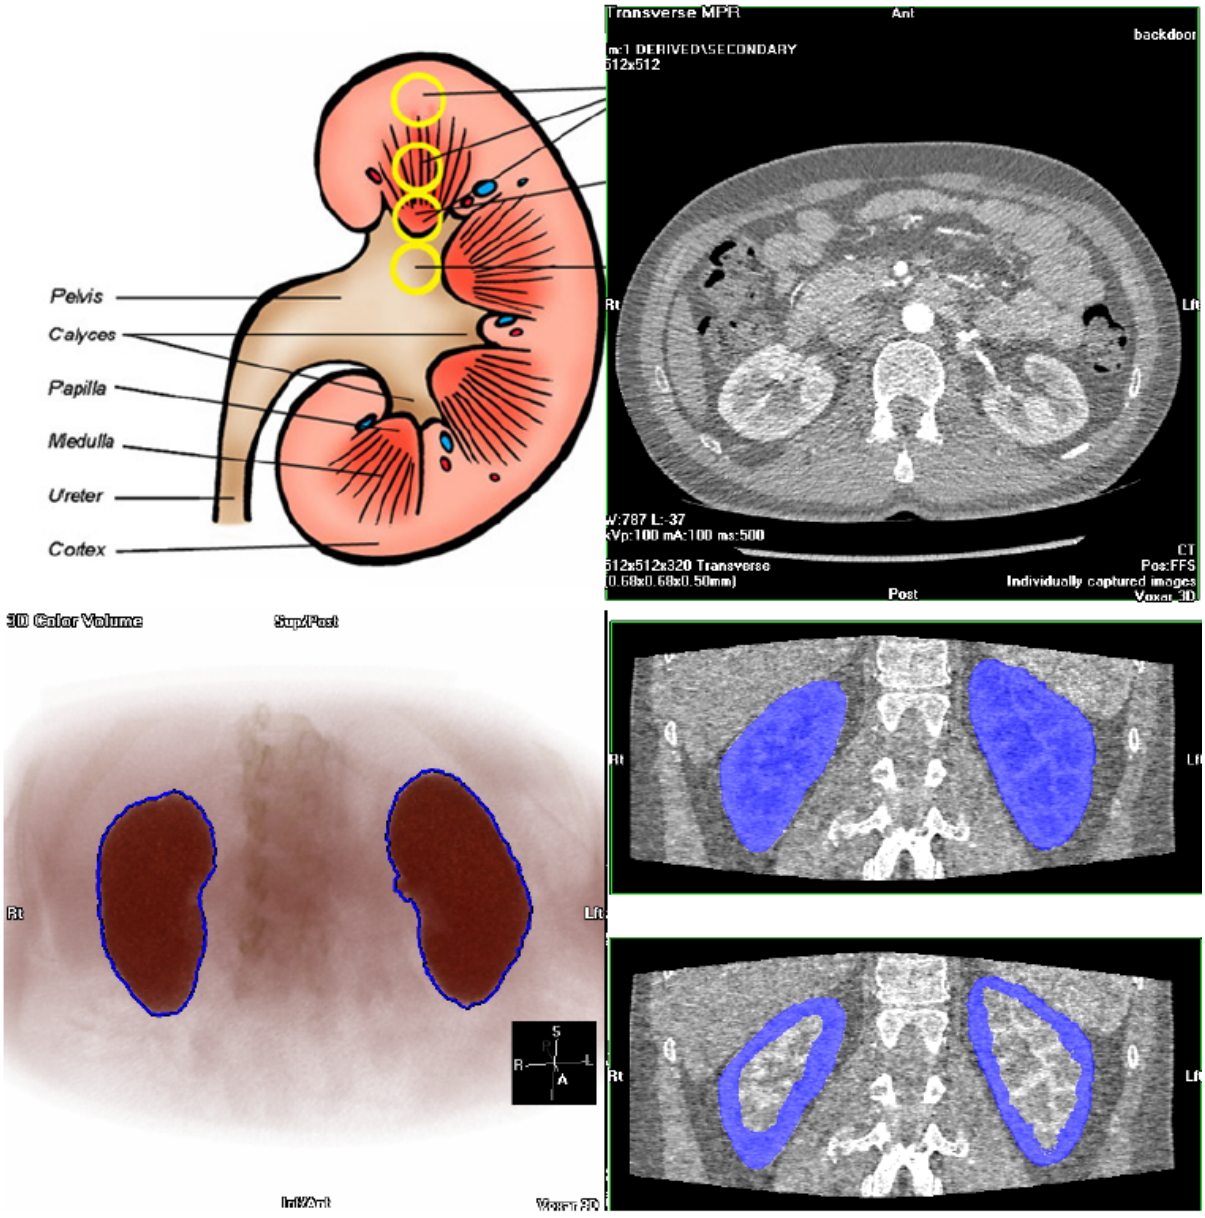

Figure 6.2: Kidney and Cortex Segmentation
Kidney anatomy image courtesy of glycoforum (www.glycoforum.gr.jp).

of these cases.

For the EM classification phase, three tissue classes were to be identified: renal cortex, renal medulla and background. The renal medulla was obtained from the set difference between the atlas-segmented whole kidney and renal cortex, based on the assumption that the cortex is wholly contained within the kidney. The background was derived by subtracting the whole kidney domain from the bounding box containing the whole kidney, dilated by $4mm$.

6.3.3 Multi-Compartment Heart Segmentation

The aim of this task is to segment the whole heart and the following sub-structures from 3D CT cardiac perfusion scans (figure 6.3):

- Aortic Root
- Left Atrium
- Right Atrium
- Left Ventricle Endocardium
- Left Ventricle Myocardium
- Right Ventricle Endocardium
- Right Ventricle Myocardium

The multi-atlas in this case contains only one image due to the amount of work required to manually segment the above regions. The atlas image in use is a scan captured from the 60% phase of the cardiac cycle of a male patient, following administration of iopramol (Bayer Schering Pharma, UK) contrast agent. The phase was selected which contained enough contrast within the right ventricle and atrium. The scans were taken with X-Ray beam current of 360mA using Toshiba's Aquilion helical CT scanner, with 396 slices captured at 0.5mm per slice resulting in a 512x512x396 volume.

The whole heart and the above structures were manually segmented in the atlas image by a clinically trained individual using the MPR manual shape selection tool in Toshiba's Voxar 3D advanced visualisation software. The left and right ventricles were segmented instead of the myocardia as they could be derived from the set difference between the respective ventricles and endocardia following registration.

Each mask of the segmented structures after registration was smoothed using a Gaussian filter of kernel size $\sigma = 3mm$ to account for the inaccuracies in the registration and to allow more weight to the intensity-based classification. The background domain presented to the EM

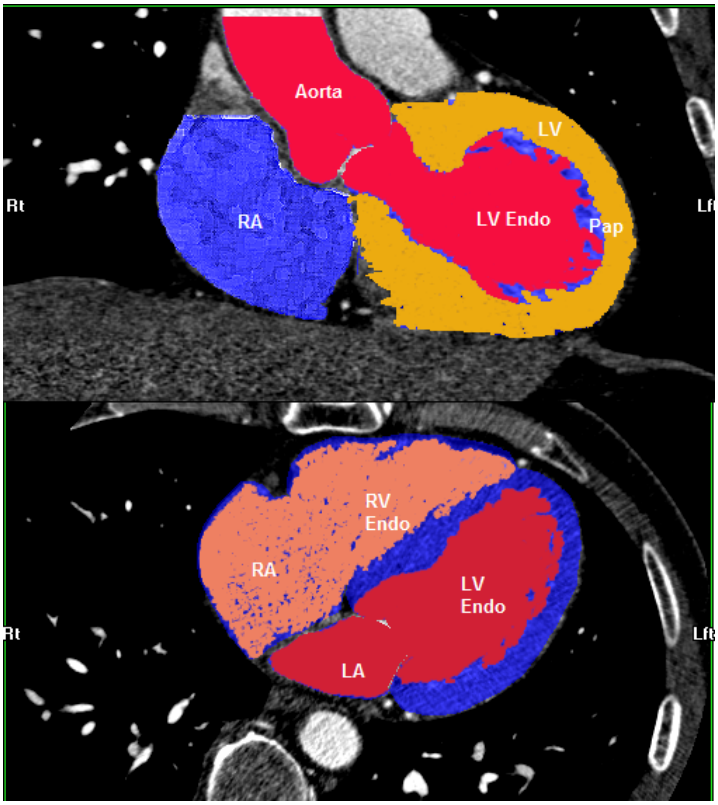

Figure 6.3: Cardiac Sub-Structures

classification phase is derived by subtracting the above listed structures of interest from the whole heart domain.

Thirty-five similar cardiac perfusion datasets were available as test data, although no ground-truth was collected on these images. Subjective validation was therefore performed by clinically trained personnel on 5 of these images selected at random.

6.4 Results and Discussion

6.4.1 Kidney and Renal Cortex Segmentation

Figure 6.4 compares the discussed method to atlas-based segmentation with and without EM-based classification. The validation data consists of 22 abdominal CT datasets as described in 6.3.2, of which left and right kidney ground-truth are available for all 22, and renal cortices available for 8. The performance metrics used are the Jaccard overlap (equation 6.6), the true-positive fraction (as a measure of under-segmentation) and the positive-predictive value (as a measure of over-segmentation).

$$Jaccard(A, B) = \frac{Vol.(A \cap B)}{Vol.(A \cup B)} \tag{6.6}$$

The segmentation results for the whole kidney do not report any significant differences, and while there appears to be an increase in the Jaccard overlap for the renal cortex segmentation, these are not statistically significant (using a student t-test with a significance level $\alpha = 0.05$).

The clearest benefit of using the atlas-supplied priors in the EM algorithm is indicated in the bottom right-hand figure in figure 6.4. The parameters for the tissue intensity distributions are measured on the respective ground-truth segmentations and compared against the estimated parameters. It demonstrates that the use of spatial priors in the EM algorithm leads to better estimates of the parameters, particularly of the renal medulla. This improvement in the estimation of the renal medulla leads to greater confidence in segmenting the inner boundary of the renal cortex (see kidney anatomy in figure 6.2), and hence an improvement in the renal cortex segmentation accuracy.

Figure 6.5 shows the estimated distributions of the kidney and renal cortex obtained by including the spatial priors in the EM algorithm.

6.4.2 Multi-Compartment Heart Segmentation

The advantage of this method is displayed in this particular application, whereby different tissue types share similar intensities (aortic root and left ventricular blood-pool, for instance). Figure 6.8 shows the estimated Gaussian components for each tissue class in question. The top row was obtained using the discussed method, while the bottom row was obtained by a 'blind' run of the EM algorithm without any spatial priors following registration. Without any anatomical constraints, the distributions are then assigned almost randomly. This can be avoided with some set of constraints on the expected ordering of intensities, but the method becomes more ad-hoc and has to be re-implemented for each application case.

The estimated background distribution naturally has a much larger variance than the others. This is because it does not follow the same assumption of intensity homogeneity as the other structures of interest. The definition of the background domain in section 6.3.3, implies that it contains high intensity structures such as the pulmonary veins and coronary arteries as well as low intensity features like fat. Highlighting more structures within the atlas would reduce background to fat and other low-intensity tissues, thereby increasing the segmentation accuracy on the whole. The segmentation results are shown in figures 6.6 and 6.7.

Subjective validation was carried out by a clinically trained individual on 5 datasets selected at random; the results are shown in figure 6.9. It is clear from figures 6.9 and 6.7 that the segmentation of the right myocardium is suboptimal. This points to a weakness in the way the atlas-based priors were presented to the EM classification phase. As described in section 6.3.3, the myocardium segmentation is the domain difference between the segmented ventricle and endocardium. This is sufficient for the left side of the heart, as the left ventricle myocardium is a much thicker structure than the right and highly discernible from the blood pool (as seen in figure

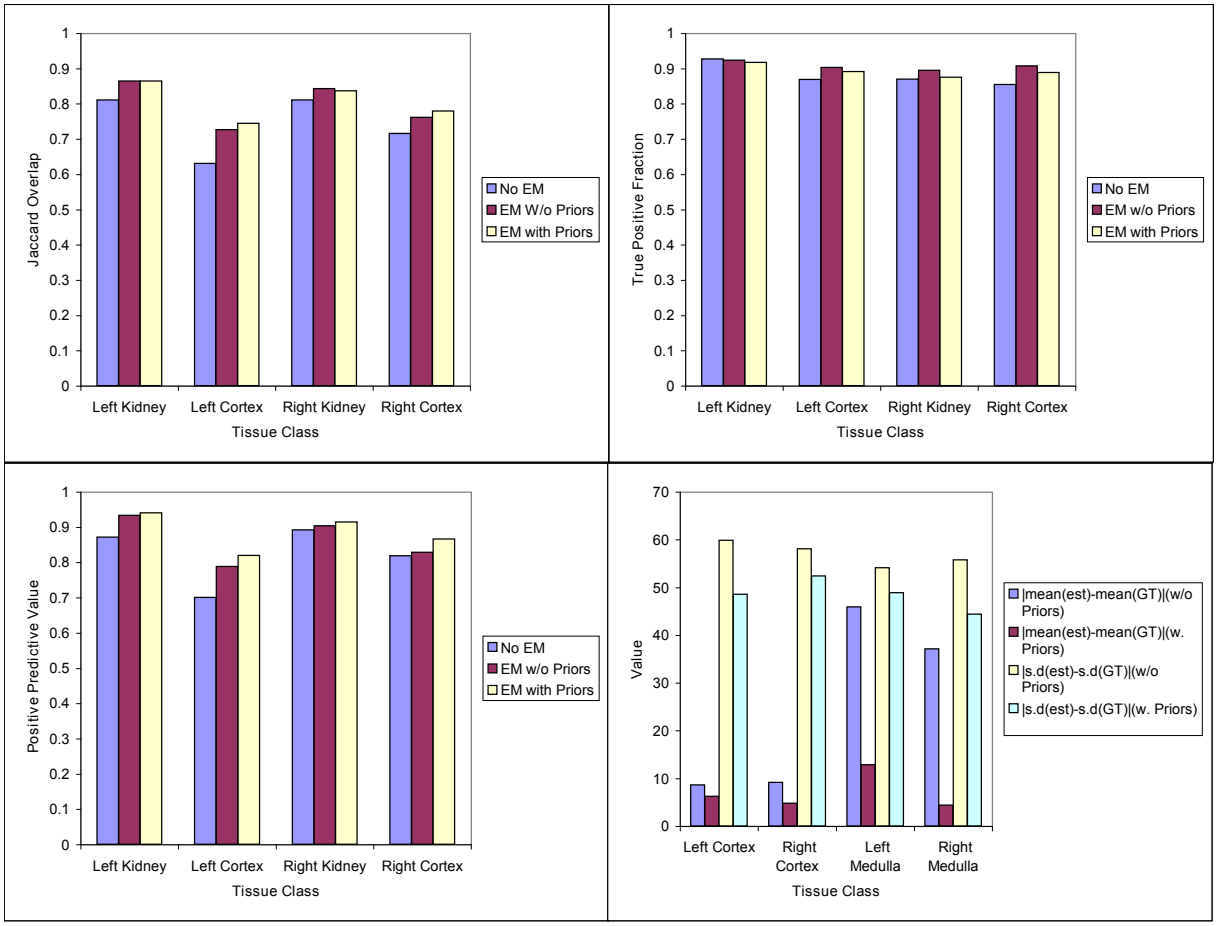

Figure 6.4: Kidney and Renal Cortex Segmentation Results

The results demonstrate a clear benefit of using EM-based classification; based on a validation set of 22 datasets for the kidney segmentation, the difference in segmentation accuracy when using priors is not statistically significant (based on a students t-test with $\alpha = 0.05$). There is an improvement in the absolute error between the estimated tissue parameters and the expected parameters gained when using spatial priors in the EM algorithm. The improvement in the estimation of the renal medulla is statistically significant and has led to increased segmentation accuracy of the renal cortices.

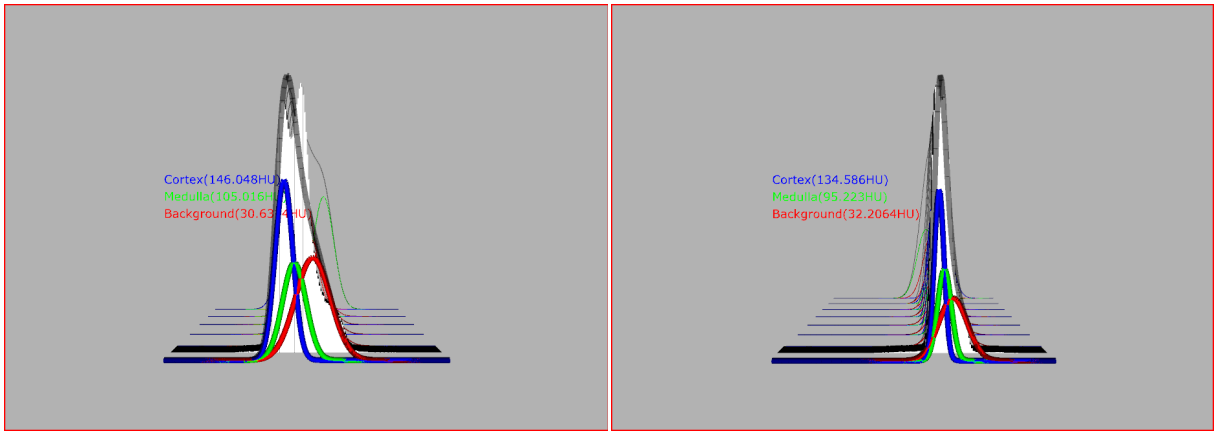

Right Kidney Left Kidney

Figure 6.5: Estimated Abdominal Tissue Distributions

The heights of each component corresponds to its proportional composition within the data, estimated using maximum posterior probabilities. The white plot is the underlying histogram of the data, and the grey line is the sum of all individual components. Multiple plots in the background show the estimated distributions at each iteration.

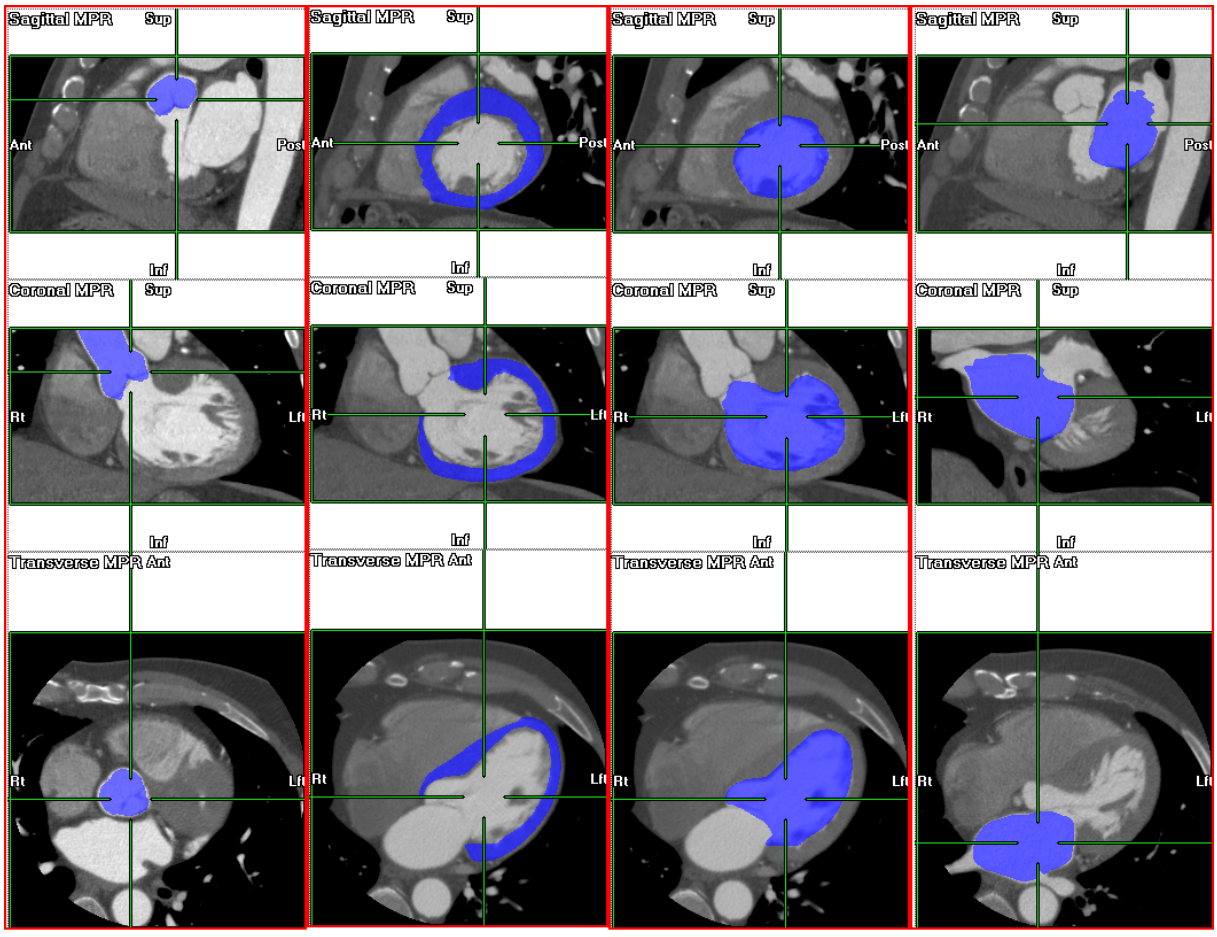

Aortic Root Left Ventricle Myocardium Left Ventricle Endocardium Left Atrium

Figure 6.6: Segmented Sub-Structures of the Heart (Left View)

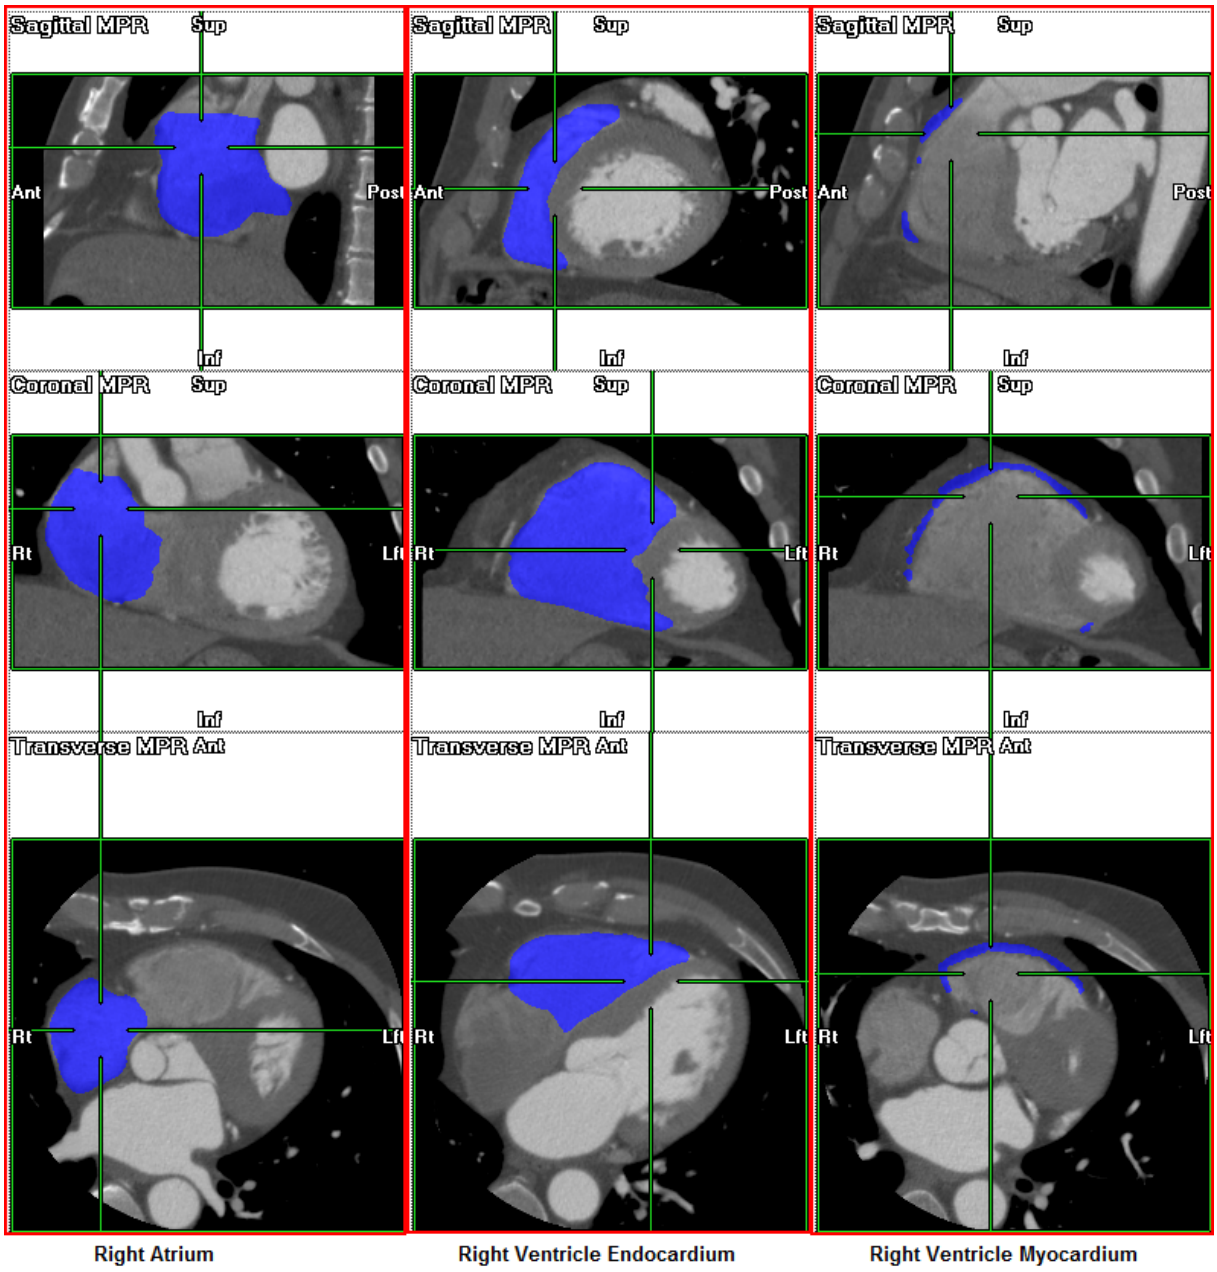

Figure 6.7: Segmented Sub-Structures of the Heart (Right View)

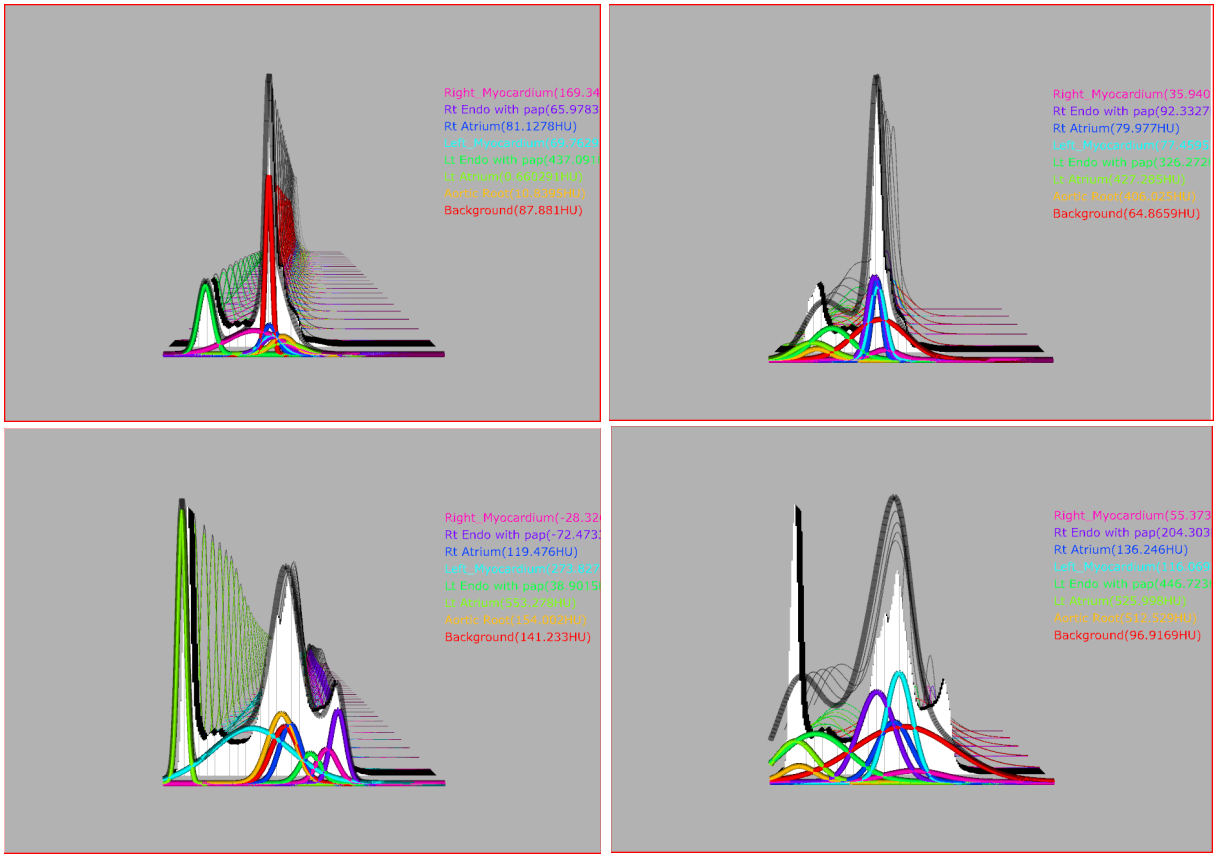

Figure 6.8: Estimated Heart Tissue Distributions

The left column shows the estimated distributions obtained using the discussed method, while the right column was obtained by running the EM algorithm without priors. The heights of each component corresponds to its proportional composition within the data, estimated using maximum posterior probabilities. The white plot is the underlying histogram of the data, and the grey line is the sum of all individual components. Multiple plots in the background show the estimated distributions at each iteration. The mix-up of distributions in the bottom image can be avoided by ad-hoc methods, but applying spatial priors achieves this in a much simpler manner.

6.3). On the right side however, this section is much thinner and except in exceptional cases, the contrast difference between it and the blood-pool is obscure. This leads to a disconnected domain for the right myocardium post-registration, implying misleading priors, and this affects the segmentation. Constraints could be added to enforce connectedness and further anatomical knowledge about the expected thickness of the right myocardium, but that is beyond the scope of this work.

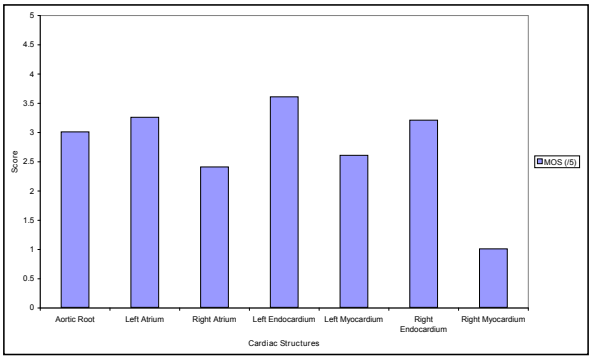

Figure 6.9: Heart Sub-Structures: Subjective Validation Results
Subjective evaluation carried out by a clinically-trained individual on 5 datasets selected at random. The right myocardium is the poorest segmented structure.

6.5 Conclusion

This paper presents a robust method for segmentation, combining atlas-based registration and voxel-based classification using the EM algorithm. The method allows spatial information, provided by the atlas to be incorporated in the EM algorithm. This extra information allows the EM algorithm to distinguish between different tissues with similar intensity distributions. The method employs multi-atlas based segmentation, registering multiple atlases to the subject image and selecting the segmentation from the atlas image that is most similar to the subject image.

The paper demonstrates the versatility of this method by applying it to two different segmentation applications. The first application, kidney and renal cortex segmentation from low-dose abdominal CT data, tests the robustness of the algorithm by applying it to noisy data. The second application, cardiac sub-structure segmentation, applies the algorithm to segment multiple structures within the heart without encoding any extra prior information except that provided by the atlas.

Both experiments demonstrate that there is an advantage to be gained by incorporating spatial priors in the EM algorithm. Tests carried out on kidney segmentation indicate that the estimated distribution models are more accurate when this information is applied, and this led to better segmentation of the renal cortex. Subjective validation was performed in the case of the cardiac

sub-structures. In this scenario, the various structures of interest share similar intensity values; this caused problems for the naive EM implementation, but upon addition of spatial priors, the correct models were estimated.

References

[1] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics*. Springer, 1st ed. 2006. corr. 2nd printing ed., October 2007.

[2] C. Tomasi, "Estimating Gaussian Mixture Densities with EM: A Tutorial," *Duke University*, pp. 1–8, 2005.

[3] N. A. Thacker, M. Pokri, and A. J. Lacey, "Model Selection and Convergence of the EM Algorithm," *Biomedical Engineering*, 2004.

[4] M. Murgasova, D. Rueckert, D. Edwards, and J. Hajnal, "Robust segmentation of brain structures in MRI," *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pp. 17–20, June 2009.

[5] A. Akinyemi, I. Poole, J. Piper, C. Plakas, and C. Roberts, "Optimal Atlas Selection Using Image Similarities in a Trained Regression Model to Predict Performance.," tech. rep., Toshiba Medical Visualisation Systems, Edinburgh, 2009.

[6] H. Park, P. Bland, and C. Meyer, "Construction of an abdominal probabilistic atlas and its application in segmentation," *IEEE Trans. Med. Imaging*, vol. 22, no. 4, pp. 483–492, 2003.

[7] A. Akinyemi, "A Combined Multi-Atlas and Unsupervised Classification Technique for Automatic Segmentation of the Kidney and Renal Cortex in Low-Dose CT Data .," tech. rep., Toshiba Medical Visualisation Systems, 2010.

Chapter 7

Feasibility Study of Medical Imaging in Nigeria

Abstract *Medical imaging is defined as the techniques and processes used to create images of the human body (or parts thereof) for clinical purposes (i.e. medical procedures seeking to reveal, diagnose or examine disease) or medical science (including the study of normal anatomy and physiology).*

Medical imaging has proven to be a useful tool in diagnosing and treating some of the world's deadliest diseases, but its presence in Africa is much lower than its prevalence in Europe, America and Asia. This report examines this issue, indicating that medical imaging technology can be used to diagnose and control some the most burdensome diseases facing Nigeria.

Adoption of medical imaging equipment requires a large amount of capital, and countries such as the United Kingdom are able to support the industry by providing reimbursements to institutions involved in radiology. This report examines the business of radiology as it is practiced in Nigeria. The Nigerian market is developing, with opportunities in the near future for any company, local or international, with the necessary investment capital. The report outlines some of the key measures proposed by the Nigerian government to enhance the growth of medical imaging and the development of the healthcare industry; some of these measures include partnerships with private corporations and grant-funding for certain projects.

7.1 Acknowledgment

Thanks to the entire image analysis team for allowing me the time to focus on creating this report. Many thanks to Dr. George Burns for supervising the entire process and letting me pursue this particular project.

I have to thank Brady Anderson, Tristan Lawton and Aidan Wellington for their invaluable guidance throughout the entire project. Thanks to the people in Nigeria who supplied me with

the information I needed to write this paper: Dr. Kolapo Akinyemi, Dr. Ritgak Dimka, Dr. Ihimekpen and Dr. Kunle Irojah.

7.2 Introduction

The aim of this report is to present Nigeria as a possible market for medical imaging technology. This involves taking a look at the burden of disease in the country and evaluating the efficacy of the established imaging modalities in diagnosing and monitoring these diseases. The report discusses the macroeconomic factors of the Nigerian market, in order to provide insight on the general business climate of the country. It then hones in on the healthcare system and the medical imaging industry in particular, to present some of the key metrics such as healthcare funding, which can stimulate further research required to pursue an actual opportunity.

The structure of the report is as follows:

- Section two discusses the prominent medical imaging modalities, outlining their principles of operation, application areas and limitations.
- Section three presents an epidemiology study of Nigeria, matching the most burdensome diseases to the modalities discussed.
- Section four discusses the different business models for radiology, as practiced in developed markets.
- Section five presents a market analysis of Nigeria, touching on macroeconomic factors and focusing on the healthcare industry.
- Section six summarises the whole report, outlining suggestions for further work.

7.3 Medical Imaging Technology

7.3.1 Introduction

The aim of this section is to expose the various medical imaging modalities available for diagnostics. The major modalities and some historical/less prevalent technologies are discussed, outlining the underlying principles related to each method and their most efficient application areas. The section summary provides a table summarising the documented information.

7.3.2 Imaging Modalities

The various diagnostic medical imaging modalities available are as follows:

Electron Microscopy

This modality uses electrons as the source of illumination and provides up to x2M magnification. It is used mainly in anatomic pathology to identify organelles within their cells; it is also used in diagnosis of kidney disease.

Projection Radiography (X-rays)

For a while, this has been the most prevalent imaging modality due to price and ease of use. It works by transmitting x-rays through the body, understanding that different parts of the anatomy have different absorption properties, with bones being the highest. This makes projection radiography most suitable for determining the type or extent of fractures, dislocations, and arthritis. It is also used (much less so now) to visualise both benign and malignant tumours, and finds some use in chest pathology.

The risks associated with frequent use of this method come from repeated exposure to ionising radiation from the x-rays.

Fluoroscopy

This modality provides real-time images of internal structures by employing a constant input of x-rays at a relatively low dose rate, using contrast agents such as barium and iodine. As a result of this, it is mainly used for image-guided procedures.

Computed Tomography

The first CT scanner was developed in 1972 by Sir Godfrey Hounsfield. Since then, it has become one of the most essential medical imaging techniques through its applicability to a range of clinical situations.

Principles of Operation The CT scanner in figure 7.1 consists of a patient table; a rotating x-ray source and rotating detectors enclosed in a gantry. The emitted x-ray beams are attenuated by the internal organs and the level of attenuation measured by the ring of detectors. Modifications of this basic underlying technology have led to the development of helical CT, in which the patient table is moved through the gantry as the x-ray tubes rotate, thus producing a volume of contiguous slice data. Multi-slice/multi-detector CT scanners are now quite common; these utilise the same principles of the helical scanner, but contain multiple rows of detector rings. This allows the scanner to capture multiple slices per rotation; thus increasing the anatomical area imaged in a fixed time. Toshiba's Aquilion ONE 320-slice CT scanner can cover a 16cm vertical range, capable of imaging an entire organ, in one gantry rotation.

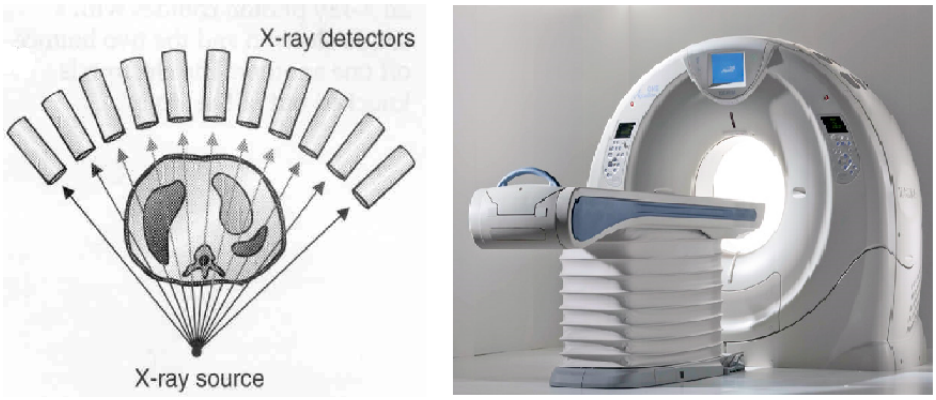

Figure 7.1: Computed Tomography Scanner
The image on the right shows a CT scanner [1]. The image on the left is a diagrammatic representation of the components in the gantry; the CT scanner is basically composed of x-ray sources and x-ray detectors [2].

Each pixel in the reconstructed image is a measure of the average attenuation of the x-rays at that particular point in the anatomy, as measured by the detectors; this is referred to as the Hounsfield Unit (HU). Hounsfield units are relative to the attenuation rate for water; therefore water has a value of 0 HU. Examples of HU values for various tissue-types are shown in figure 7.2.

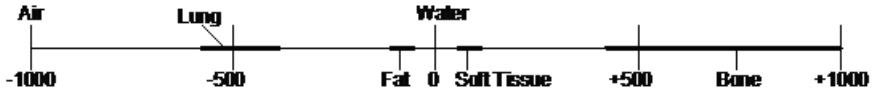

Figure 7.2: The HU Range of Values
Air has a value of -1000 Hounsfield Units [3].

The entire range of HU values is not visible at a particular instant by the human eye, and so it is possible to map a specific range of HU values to the range of the display equipment. This is known as windowing, and the window-level is the HU value at the centre of this window. For example, in a CT examination of the chest, a window level of +40 and a window-width of 350 are chosen to image the soft-tissue.

Visualisation software packages have made it possible to reconstruct the CT images in various ways suited to the particular clinical application. Multi-planar reformatting (MPR) and 3-D volumetric data are the two basic forms of visualisation, allowing users to view the anatomy as a series of contiguous slices for internal tissue analysis, or for surface and texture analysis, respectively.

In order to visualise blood vessels clearly, contrast agents which increase the attenuation rate of the blood are applied to the patient, either orally or intravenously. This is particularly useful

for vessel analysis. The procedure is termed CT Angiography, and the most commonly used of these contrast agents are iodine-based solutions.

Applications of CT CT imaging is readily available in most hospitals due to its excellent image resolution, which makes it applicable to a wide range of clinical procedures. Advances in visualisation technology have made it more convenient to use; for instance, orthogonal and oblique planes can be viewed with ease. Computer-aided detection tools have been developed for CT to speed up reading and analysis of the images. Some applications of CT are in:

- Head and Neck Imaging: This is typically performed using intravenous contrast, and is used to evaluate patients with neck masses, and head and neck cancer.
- Cardiac Imaging: CT angiography is commonly used with 64 slice CT scanners to detect plaque in coronary arteries. This makes it suitable for general vascular imaging.
- Pulmonary Imaging: High resolution CT images can be used to evaluate lung disease and lung cancer.
- Functional Imaging: CT images are superimposed on PET images to add structural information to the functional information provided by PET.
- Abdominal and Pelvic Imaging: CT is the most widely used modality for imaging the liver, kidneys, bowels and pelvis; although ultrasound is preferred for imaging the female reproductive system. Most of these scans are performed with intravenous contrast agents[4].

Limitations of CT

- Radiation: The use of short wavelength x-rays cause ionisation in body tissues, damaging cells over constant exposure. A single abdominal CT scan delivers up to $10mSv$ to the patient, compared with $\approx 3mSv$ average yearly background radiation dose in the United States [5].
- Brain imaging: MR has proved to be a more sensitive imaging technology for evaluating early ischemic strokes and grey/white matter atrophy.
- A multi-slice CT scanner with high enough resolution to perform most of the applications will cost at least US$700k, making it more expensive than safer and more portable technologies such as ultrasound.
- 5% of patients experience mild allergic reactions to contrast agents [4]. One potentially severe complication in the application of intravenous contrast is acute renal failure.

Magnetic Resonance Imaging

Figure 7.3 is a cut-out view of a typical MRI scanner showing its components.

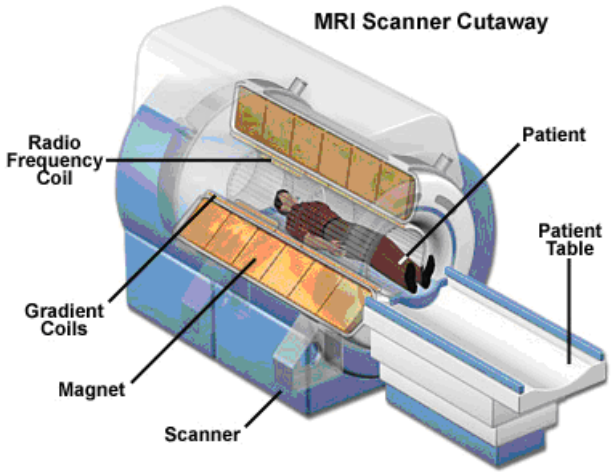

Figure 7.3: Components of an MRI scanner
The coils are used to control temporal and spatial localisation [6].

Principles of Operation In MRI, a magnetic field is applied to align the hydrogen nuclei in the body. Hydrogen is chosen because it is abundant in the human body in form of water and fat. The hydrogen nucleus possesses the 'spin' property which causes it to precess at a known frequency (termed the Larmor frequency) when an external magnetic field is applied, and all the nuclei produce a net magnetisation vector pointing in the direction of this external magnetic field [7].

In order to produce an image signal, an RF pulse is applied at the Larmor frequency, so the precessing nuclei absorb the energy from this pulse until no component of the net magnetisation vector exists parallel to the magnetic field; this is known as saturation. Once this pulse is switched off, the relaxation process emits a wave that is picked up by an RF receiver. Relaxation is divided into longitudinal (or T1) and transverse (or T2) relaxation, and the time taken for each of these are unique to the type of tissue. This produces a signal for the entire body; this signal is localised to a particular slice and 2-D co-ordinates by using three gradient magnetic fields.

The gradient coils apply three gradient fields in orthogonal directions that interact with the main magnetic field:

- The slice-selection field, to localise the signal to a 2-D plane;
- Frequency encoding field, to localise the signal to a column in the 2-D plane;
- Phase-encoding field, to localise the signal to a row within a column.

The RF receiver detects a signal at a time during T2-relaxation; and various sequences for this image acquisition stage exist. In the spin-echo (SE) sequence, for instance, the RF pulse is applied a second-time at time TR during the T1 relaxation time, and the T2 relaxation signal measured after a set time TE. The choice of TR and TE controls the contrast of the captured image, and is referred to as T1- T2-weighting [7]. Controlling TR and TE and varying magnetic field strength, makes MRI an art-form in itself, capable of imaging all kinds of anatomy for both structural and functional imaging.

Applications of MR The main advantage of MRI over x-rays and CT imaging is that the RF pulses applied to the body are not of high enough frequency to cause ionisation, and hence the risks of tissue damage associated with the other two methods are not applicable.

MRI is used for both structural and functional imaging. It is used in functional brain imaging, muscle imaging, liver, kidney, and lung imaging [8]. Angiographies are also performed in MRI making it useful for cardiac analysis and perfusion studies. The structural information obtained from MRI can also be overlaid with functional information from PET (Positron Emission Tomography) images to provide more useful information, especially in oncology.

MRI is useful for all the above forms of imaging because the quality of the image produced, i.e. contrast levels for anatomy of interest, image resolution, etc, are all controlled by physical parameters on the machine. It is therefore an active area of research in imaging at the moment.

Limitations of MRI The major downside to this technology is its cost; MRI scanners are by far the most expensive of the structural medical imaging scanners, with an average price of US$1.5M for the basic 1.5 Tesla scanner. For the patient, switching the gradient fields on and off constantly can create a very noisy experience. In cases of trauma, it is very difficult to prepare the patient for a scan, with preparation along the lines of checking for replaced hips, pace-makers and other items that will disrupt the magnetic field.

Medical Ultrasonography (Ultrasound)

Principles of Operation The underlying principle of ultrasound is the piezo-electric effect, which transforms an electrical voltage to high frequency sound waves and vice-versa.

An ultrasound scanner consists of a transducer, normally on a hand-held probe. The sound frequencies produced by this transducer are transmitted into the body; the same transducer (or a separate one on the same probe) is used to detect the reflected sound, or echo. The time taken for the echo to be received provides information about the depth of the reflecting anatomy (see figure 7.4).

Tissues have unique acoustic impedance, and the relative acoustic impedance between adjacent tissues determine the proportion of the signal that will be reflected. For example, air

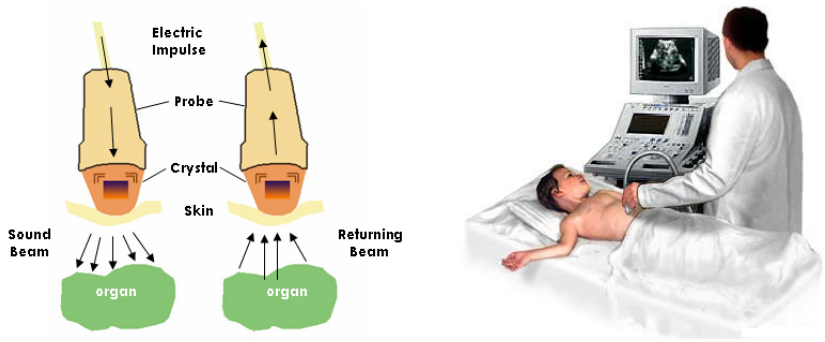

Figure 7.4: The principles of ultrasound
The image on the left shows transmitting and receiving ultrasound transducers at the ends of hand-held probes. The image on the right shows the scanning procedure being carried out; the image is obtained in real-time [9].

has an acoustic impedance of 0.0004 Rayls, and muscle is about 1.70 Rayls [10]; this implies that the muscle-air interface will reflect practically the entire signal. This is why transducers are tightly coupled to the patient with water-based gel to allow the signal to travel into the body.

Frequencies used in medical ultrasound are in the range of 1 - 10 MHz [10]. High frequency ultrasound has the advantage of producing better resolution due to the magnitude of the reflected signal, but more of the signal energy is absorbed by tissues, reducing its penetration potential. On the other hand, low frequencies produce relatively coarse resolutions, but are capable of imaging deeper tissues within the body.

The application of the Doppler effect to ultrasound makes it possible to measure blood-flow within arteries. The Doppler effect is essentially what makes the sirens of fire-engines appear to change in pitch as they get closer to the listener [10]. Doppler ultrasound is applied as continuous-Doppler or pulsed-Doppler. The basic idea is that a known frequency is applied, targeted at the artery of interest, and the received frequency is measured, with the difference in frequency relating to the velocity of the red blood cells in the blood within the artery.

Visualisation in ultrasound has evolved from seeing frequency spikes to visualising the entire scan in 3-D in real-time; this is one of the main advantages of ultrasound.

Applications of ultrasound Ultrasound is the cheapest of the cross-sectional imaging modalities around at the moment, making it very efficient as a technology for initial evaluation. Ultrasound scanners are also portable, allowing it to be used in cases where the patient cannot be prepared for a CT or MR scan. Unlike CT, the signals emitted into the body in ultrasound imaging are non-ionising, and so there are no known risks attached to this modality.
Some of the applications of ultrasound are as follows:

- Obstetrics: By far the most popular use of ultrasound today. It is used for a number of

purposes such as; imaging the foetus, in order to rule out abnormalities; evaluating the female pelvic organs during pregnancy; and general monitoring of the pregnancy.

- Vascular ultrasound: This combines Doppler ultrasound with real-time imaging to assess blood flow in the arteries. It is used to detect abdominal aneurysms, arterial stenoses, carotid occlusive disease [4], and renal vascular diseases. High frequency continuous-Doppler is applied to superficial arteries like the carotids and arteries in the limbs, whilst pulsed-Doppler is applied to the heart. Vascular ultrasound is used in echo-cardiography to assess coronary and left-ventricle functionality.

- Testicular ultrasound: Ultrasound is the most favourable technology for testicular imaging. High frequency ultrasound is used to assess testicular abnormalities.

- Abdominal ultrasound: Ultrasound is popular in imaging the liver, gall bladder, spleen, pancreas, and kidney. Although it is being replaced by CT and MRI now, it is still a popular tool for initial assessments.

Limitations of ultrasound

- Relatively poor resolution compared to other cross-sectional imaging modalities like CT and MRI.

- It is highly operator-dependant; the quality of the image depends not only on the quality of the transducer used, but also the technologist, and the interpreting physician.

Nuclear Imaging

Principles of Operation Nuclear imaging has been around since before CT or MRI; a number of people from as far back as 1903 [11] have contributed to its development.

The underlying assumption in nuclear imaging is that tissues in the body exhibit characteristic absorption rates for substances (primarily nutrients). In order to exploit this assumption, the specific nutrients, in the form of pharmaceuticals, are attached to radionuclides to make up radiotracers, or radiopharmaceuticals. These radiotracers accumulate in the organ of interest and start to decay, emitting gamma rays. The emitted gamma rays are detected to form an image displaying the relative absorption rates of the radiotracer in the body. A common radionuclide used in nuclear imaging is Technetium-99m [12]; when this is attached to a pharmaceutical like Sestambi, it accumulates in the myocardium, enabling myocardial imaging.

The radionuclides are produced in cyclotrons, preferably close to the imaging lab. This is an important consideration in the planning of nuclear imaging centres, because of the decays associated with radioactive substances.

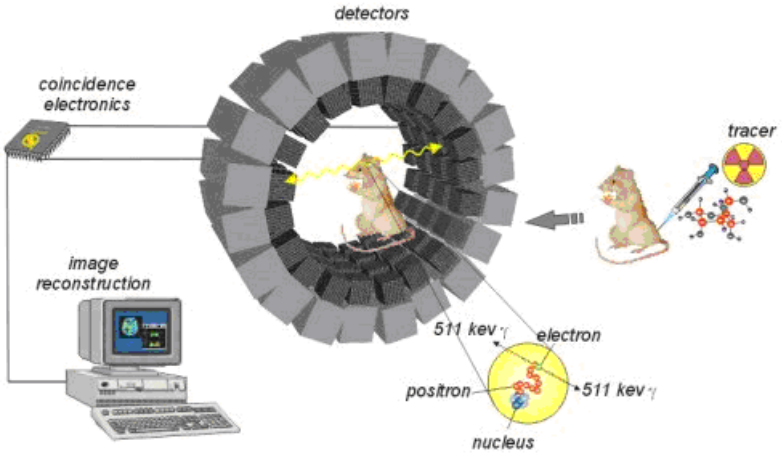

Figure 7.5: Nuclear imaging procedure
PET is shown here, whereby the tracer is injected into the patient, and the photons are used to reconstruct the image [13].

In order to form an image, gamma cameras are placed around the area of interest to detect the emitted radiation. These cameras contain fluorescent crystals, which change the gamma rays to light photons. The light photons react with a photo-electrode to generate electrons. Therefore, an image is the geographical location of the radiopharmaceutical within the body [14]. Figure 7.5 depicts the nuclear imaging procedure.

Applications of Nuclear Imaging Nuclear imaging is used as a diagnostic imaging technique for certain conditions that otherwise would require exploratory surgery. The main advantage of nuclear imaging is that it provides unique information that cannot be obtained by other imaging techniques. Some of its application areas include:

- Detection of brain abnormalities
- Measuring thyroid function
- Analysing kidney function
- Evaluating bones
- Detecting the presence/spread of cancer [12]

The most popular branch of nuclear medicine in use in oncology imaging is Positron Emission Tomography (PET).

Positron Emission Tomography PET uses Positron emitters, such as Fluorine, attached to pharmaceuticals, in order to detect the uptake sites within the body. In PET, the positrons emitted combine with the electrons in the surrounding tissue briefly; and then a process called annihilation occurs, releasing two photons in opposite directions. These photons are measured to reconstruct the image. The most common use of PET is in oncology, and so the most commonly used pharmaceutical is fluorodeoxyglucose (FDG). FDG is a glucose analogue, and the theory here is that cancerous cells are highly energetic, requiring a lot of glucose for metastases. FDG is labelled with fluorine-18 to produce 18F-FDG. This radionuclide has a half-life of 110 minutes [15], therefore it needs to be produced on-site to be clinically useful. PET is used in oncology for the following tasks:

- Distinguishing between benign and malignant tumours
- Establishing the site of the cancer
- Measuring response to therapy
- Identifying the primary site for a tumour for biopsy.
- Cancers currently investigated by PET: Head & neck tumours, thyroid carcinomas, pulmonary nodules, lung cancer, breast cancer, pancreatic cancer, colorectal cancer, ovarian cancer, testicular tumours, Hodgkin's disease, and brain tumours [15].

Limitations of PET/Nuclear Imaging Some of the limitations associated with PET/Nuclear imaging are as follows:

- The radiotracers could take anywhere between a few seconds and a few days to accumulate at the tissue of interest; this could lead to a time-consuming process.
- Although the radiation dose associated with the radionuclides is quite low (10mSv for Fluorine-18[15]), the risk of dying from radiation-induced disease is estimated at 5-6 per 10,000 [15].
- The procedure is not normally performed during pregnancy, to prevent complications.
- PET images provide functional information only, and so need to be combined with CT or MR images to present structural information to the reader.
- A PET scanner costs about US$2M [16], and the cyclotron an additional US$2.4M; making PET a very expensive imaging technology to consider.

7.3.3 Summary

The available diagnostic medical imaging modalities have been described in this section; table 7.1 summarises their relative features.

Imaging Modality	**General Principles**	**Application Areas**	**Associated Shortcomings/Risks**
Electron Microscopy	Electrons used to provide illumination	Virology, anatomic pathology, 3D tissue imaging	Expensive to maintain, cannot image internal organs
Fluoroscopy	X-ray source and fluorescent screen	Image-guided procedures (some surgery)	Exposure to ionising radiation through x-rays
Positron Emission Tomography	Detect photons emitted by positron-emitting radionuclides	Cancer detection, and monitoring response to therapy	High cost of cyclotron to create radionuclides
Projectional Radiography	X-rays attenuated differently by parts of anatomy	Detect bony anomalies, soft-tissue anomalies (not very efficient)	Ionising radiation through x-rays, limited to bones
Computed Tomography	Rotating x-rays to form image slice.	Chest, pulmonary, cardiac, abdominal	Ionising radiation through x-rays
Magnetic Resonance Imaging	Uses magnets to polarise hydrogen nuclei in tissues	Neurological, musculoskeletal, cardiovascular and oncology imaging	The magnet is very expensive; equipment is large.
Ultrasound	High frequency sound waves are reflected by anatomy	Foetal imaging in pregnant women, abdominal organs, breast, muscles	Limited depth penetration, operator-dependent quality

Table 7.1: Medical Imaging Modalities
Summarising general principles of operation, major application areas, shortcomings/associated risks.

7.4 Epidemiology

7.4.1 Introduction

The overall aim of this section is to express understanding of the burden of disease in Nigeria, and to provide evidence of how medical imaging technology can address the issue, potentially improving the general state of health in the country.

Global epidemiology surveys are carried out periodically by a host of organisations, and these provide documented data indicating the relative burden of a specific disease or group of diseases amongst the different countries involved. From this sort of data, it is possible to deduce patterns that account for the difference in disease burdens across countries, and ascertain the impact of medical imaging technology on the state of health in a community.

This section begins by considering the type of measurements that need to be performed on the population in order to understand the real burden of diseases. It goes on to discuss the actual data collected based on the relevant measurements, and any interesting observations. Based on the data collected, the following section attempts to match up the key diseases to diagnostic imaging modalities, providing information on the benefits of medical imaging technology in controlling the diseases that are affecting Nigeria.

7.4.2 Measures of Disease Burden within a Population

Diseases have different characteristics based on the way they affect human beings; some diseases are life-threatening, whilst others although not being fatal, cause long-periods of inactivity or even permanent disability.

In order to assess the burden of disease within a population, it is necessary to consider not only the mortality rate of the particular disease, but also the equivalent years lost due to affliction caused by the disease, the morbidity rate, and the financial costs associated with the disease.

The World Health Organisation (WHO), in the Global Burden of Disease book [17], uses two metrics capable of capturing the pertinent information. The Disability Adjusted Life Year (DALY) sums the Years of Life Lost due to premature mortality (YLL) and the Years Lost due to Disability (YLD) for incident cases of the health condition.

By the definition provided in [17],

$$YLL = N \times L*$$

* N is the number of deaths and L is the standard life expectancy at age of death in years. This measure gains credibility because in its computation, the life-expectancy of the different age-groups within the population have been considered, therefore Years of Life Lost accurately

represents the mortality associated with a given health condition.

$$YLD = I \times DW \times L*$$

* I is the number of incident cases, DW is the disability weight, and L is the average duration of the case until remission or death in years.

The disability weight is a number between 0 and 1 inclusive, indicating the level of severity of a particular health condition; 0 indicates perfect health and 1 indicates death.

It is worth noting that none of these measures take into account social preferences, and are therefore recorded as per observation. It has also been argued in [18] that DALY does not account for broader aspects of burden, for example, deterioration in quality of life, and the emotional and physical impacts on families. Research is still being carried out to determine more accurate measures of the burden of disease, but the DALY is so far the most acceptable measure and is used in the rest of this section.

7.4.3 The Burden of Disease in Nigeria

The World Health Organisation (WHO) carries out a periodical epidemiological survey, funded by the Bill and Melinda Gates Foundation. The survey covers each of the United Nations (UN) member states, and the data is recorded in the book titled: "The Global Burden of Disease". The survey creates 20 categories for health conditions, with categories for intentional and unintentional injuries; the measures used here are Deaths, Death rates per 100,000, DALY, and DALY rates per 100,000.

Based on the GBD 2002 reports, Nigeria has the 7th highest morbidity rate in West Africa. Communicable, maternal, perinatal and nutritional conditions account for the highest mortality/morbidity rates, of which infectious and parasitic diseases are the highest sub-group, with HIV/AIDS, malaria and measles causing the greatest burden in this category. Under communicable diseases, lower respiratory infections hold the second highest DALY rate, with maternal and perinatal conditions, such as low-birth weight and birth trauma closely following.

Non-communicable diseases account for about a fifth of the total DALY rate; of these, neuropsychiatric conditions are the highest, followed by cardiovascular diseases, sense-organ diseases (particularly cataracts), respiratory diseases and cancers. Of all cancers, breast cancer ranks the highest in terms of morbidity, followed by liver cancer and then prostate cancer; with colon cancer, stomach cancer and leukaemia closely following. On the other-hand, in terms of mortality rates, prostate cancer ranks the highest, followed by breast cancer, and then liver cancer. Under this same measure, cardiovascular diseases are the deadliest non-communicable disease-group in Nigeria. These figures are represented in figures 7.6 and 7.7.

The full tables containing the figures are available in figures A.1 and A.2 in appendix A.

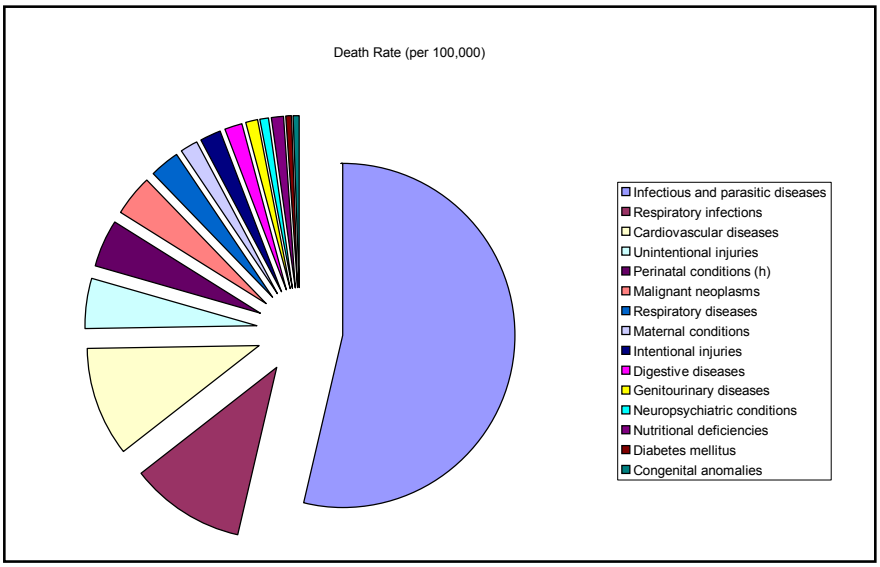

Figure 7.6: Death Rates per 100,000 people in Nigeria
The ranking still remains the same, except for Neuropsychiatric conditions not being as high up in death-rate ranking as it is in the DALY-rate ranking. Cardiovascular diseases are the third deadliest disease group in Nigeria.

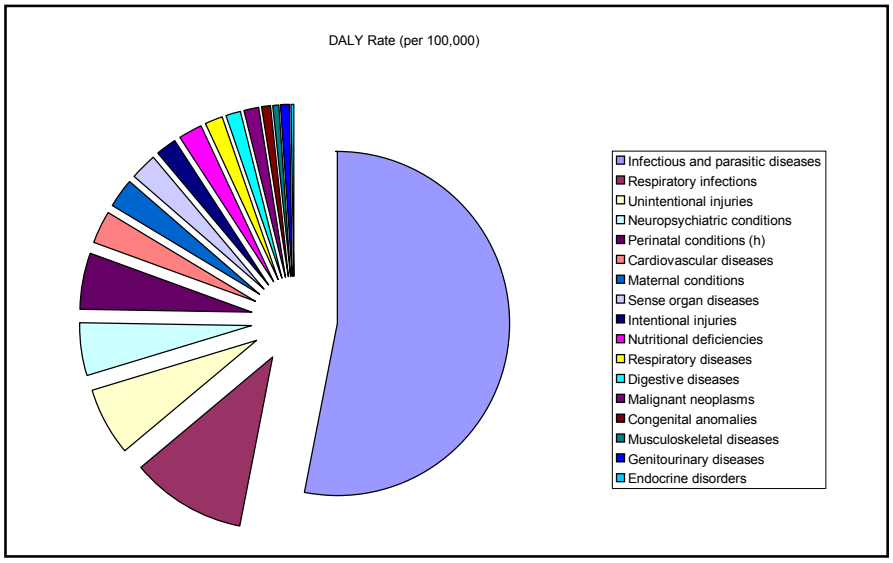

Figure 7.7: DALY Rates per 100,000 people in population for Nigeria
The ranking still remains the same, except for Neuropsychiatric conditions not being as high up in death-rate ranking as it is in the DALY-rate ranking. Cardiovascular diseases are the third deadliest disease group in Nigeria.

7.4.4 Medical Imaging and the Burden of Disease in Nigeria

Tables A.1 and A.2 in appendix A provide information matching the top 6 disease categories (based on both DALY and Death rates) to the imaging modalities discussed in section 7.3. Table 7.8 summarises this table, showing the applicability of the major imaging modalities to these disease categories.

	Inf	RI	MC	MN	NC	CV	UI	PC
Ultrasound								
MR								
CT								
PET								
X-Rays								
DALY Rate	29673	6034	1589	748	2858	1750	3631	2830
Death Rate	884	182	31	65	16	167	81	74

Figure 7.8: Applicability of medical imaging modalities to the major disease categories in Nigeria

Red implies inapplicability, while green implies applicability. Inf: Infectious and parasitic Diseases; RI: Respiratory Infections; MC: Maternal Conditions; MN: Malignant Neoplasms; NC: Neuropsychiatric Conditions; CV: Cardiovascular Diseases; UI: Unintentional Injuries; PC: Perinatal Conditions.

It is worth noting from figures 7.6 and 7.7, that infectious and parasitic diseases carry the greatest burden. Within this category are HIV AIDS, malaria, measles, diarrhoea and tuberculosis. At present the detection rate for TB is 27% compared to the global goal of 80%, placing Nigeria in fourth place among the 22 high-TB-burden countries worldwide. Malaria accounts for 60% of all outpatient visits and 30% of childhood deaths, costing the country an estimated US$88M annually.

The data provided in the appendix points out that medical imaging modalities do not offer much in the way of diagnosing or monitoring AIDS, malaria, diarrhoea and measles, which hold the highest burden under this category, and in fact if excluded, will bring this category on par with respiratory infections. Diagnosis and monitoring of HIV, malaria, diarrhoea and measles are beyond the scope of this report, and therefore are not discussed here, although the two latter diseases are negatively correlated with the levels of hygiene and nutrition in the country.

Figure 7.8 and table 7.2 show the applicability of the imaging modalities discussed in section 7.3 to the most deadly conditions in Nigeria. These figures show that x-ray is the only modality useful to target tuberculosis. For all the other disease groups, CT and MR seem to be equally applicable except in the case of obstetrics, in which ultrasound is globally preferred.

In order to target all the diseases listed in table 7.2, the following equipment is at least required:

- Tuberculosis:

Disease Category	**Specific Disease**	**Diagnostic Modality**
Infectious diseases	Tuberculosis	X-Ray and CT
Respiratory infections	Pneumonia	All modalities except PET
Cardiovascular diseases	Ischaemic heart disease	All modalities except X-ray
Unintentional injuries	Road traffic accident (trauma)	All modalities except PET
Perinatal conditions	Low birth weight and birth trauma (obstetrics)	Ultrasound and MRI
Malignant Neoplasms	Breast cancer, Prostate cancer	Mammography, All modalities

Table 7.2: The top 6 deadly diseases in Nigeria (2006).
The ranking is based on data collected by WHO; more diseases are prevalent under these sub-categories, but the table only shows those which are easily diagnosed by the major imaging modalities.

  - X-Ray: Fixed digital(computed radiography or direct radiography) or plain-film x-ray scanner.

- Ischaemic heart disease:

  - CT: Requires 64+ -slice multi-detector helical CT with visualisation workstation[19].
  - Ultrasound: Echocardiography and Doppler capabilities for arterial imaging.

- Trauma:

  - CT: Requires multi-detector CT with at least 16-slice resolution[20][21] and visualisation equipment (PACS), as angiography is recommended.
  - Ultrasound: Ultrasound is suitable for unstable patients with abdominal trauma[20]. It requires real-time scanners with Doppler capabilities, using sector or linear transducers with mean frequencies between 2 and 5 MHz (for adults)[22].
  - X-Ray: Requires portable x-ray scanner, which could be digital or plain-film.

- Obstetrics:

  - Ultrasound: Requires real-time sonography, using 3-5MHz abdominal transducer or 5-10MHz vaginal transducer[23].

- Breast cancer:

  - X-Ray: Mammography could be plain-film, digital with viewing resolution of at least 5MP, or tele-mammography[24].
  - Ultrasound: Follow-up findings from mammography requires real-time linear array scanner operating at 10MHz or higher[25].

- Prostate cancer:

  - Ultrasound: Used as follow-up to prostate screening requires a real-time transrectal transducer with a frequency of at least 6MHz[26].

The above list ignores MRI and PET modalities although they are capable of detecting some of the diseases listed above. This is just because of the high cost of these two modalities, and so proving that they are not explicitly required to reduce the burden of disease in Nigeria.

7.4.5 Summary

This section begins by describing one of many possible methods to measure the burden of diseases within a population. Epidemiological data provided by WHO is used to describe the state of health in Nigeria. It can be seen that the most burdensome diseases are on-going areas of research, which may or may not be involved with the field of medical imaging. The last section points out the applicability of the medical imaging modalities discussed in section 7.3 to the diagnosis and monitoring of the major disease categories; more data is provided in appendix A.

7.5 The Business of Radiology

7.5.1 Introduction

Following discussions about the technology available for radiology and the diseases that could be addressed by these pieces of technology, this section provides an insight into the way radiology practices are run as businesses. The section outlines the various business categories as practiced in the UK and the USA, as both countries have suitably mature healthcare industies.

7.5.2 Business Models for Radiology

Imaging centres are very expensive to equip and run; for this reason, medical institutions need to meticulously consider the need for this service.

The business of radiology is driven by the demands of the customer who are in this case, the referring physicians. They are therefore structured based on a few driving factors, such as:

- Capacity management: Due to the high capital cost of imaging equipment, health institutions need to ensure maximum utilisation in order to balance out this fixed cost.
- Convenience: Referring physicians and patients need to have convenient access to the radiology service.
- Turn-around time: Referring physicians measure quality by a number of factors including fast delivery of results. Other factors such as the diagnostic quality of images, succinctness of reports, and the boasted level of technology are very important driving factors.
- Breadth of Service: It may be an advantage to the referring physicians to have a one-stop-shop for all imaging modalities. With this set-up, scans can be carried out on different modalities to increase the quality of diagnosis.

Therefore, in all healthcare systems, ranging from private to public systems, there are four classes of radiology solutions that exist to address these issues.

Tele-radiology

Tele-radiology involves the distribution of radiological patient images of various modalities from one location, usually the scanning centre, to another for reading. This service comes about due to shortage of trained or specialised radiologists, e.g. MRI, neuro, paediatric, etc, in the scanning centres. It provides operating leverage to institutions concerned. Tele-radiology practices make use of PACS as an integral component of their business, as it allows the scans to be transported from/to the referring physicians with ease.

These practices are mostly called upon at times when health institutions have a back-log of scans, they are efficient at reducing turn-around-time for the referring institutions in this case.

Tele-radiology practices can be further divided into categories, based on the level of service provided.

- Stand-Alone Tele-Radiology[27]: These are radiology organisations that own their own distributed imaging centres with a centralised reading centre. These organisations tend to focus on a radiology sub-specialty, and provide primary diagnosis. Examples of companies that run this structure in the United States are “Proscan imaging”, “MedTel International”, “Centre of Diagnostic Imaging”, and “American Radiology Services”.
- On-Call Coverage (a.k.a. Night-Hawk): These are organisations that provide extended coverage for radiology departments. They are essentially in place to tackle variable capacity in customers’ institutions. It may be quite expensive to have a full-time radiologist on the staff during off-peak hours (over-night, for instance) in a hospital, and so such “Night-hawks” come in handy to perform reading during these hours. These could be practiced by

individual radiologists or by organisations with staff of accredited radiologists. Common services provided by such structures are preliminary reads, but sub-specialty and final interpretation services are also common, depending on the service supplier. Examples of organisations practising on-call radiology are "NightHawk Radiology Services" and "Virtual Radiologic Consultants".

This form of tele-radiology lends itself naturally to international radiology, whereby practitioners can take advantage of the differences in time-zones.

Tele-radiology organisations make their revenue by charging the referring institutions directly.

Outpatient Imaging Centres

This scenario usually involves a group of radiologists coming together to start a business providing a complete package, e.g. scanning, reading, reporting and sometimes diagnosis, as a service to hospitals and other medical institutions. This is usually a dedicated business, and usually provides no other service. In order to achieve their independence, outpatient imaging centres need to have their own imaging equipment, i.e. modalities and PACS, to support their intended capacity.

These organisations could be either small businesses or large international chains, and usually provide their services to multiple hospitals, thus alleviating the problem of under-utilisation of medical imaging equipment. This saves their clients the need to hire full-time technologists and radiologists, while still maintaining good quality of diagnostic imaging. Outpatient imaging centres need to cater to all the needs of their clients. To do this, they normally have to boast the latest in imaging technology, support multiple modalities, as well as target the most burdensome diseases affecting their areas of operation. In some cases, these companies also provide laboratory and clinical services as part of their product portfolio. Examples of outpatient imaging companies are "Center for Diagnostic Imaging" and "Austin Radiological Associates" in the United States; "Union Diagnostic and Clinical Services", and "Image Diagnostics" in Nigeria.

Out-patient imaging centres make their revenue by charging the clients directly. Clients in this case may be individuals, insurance companies, or other health institutions. In the case of the United States, these organisations charge Medicare or patients' insurance companies for each scan carried out. In the UK, the same procedure applies with the NHS for authorised practices.

Hospital-Adjoined Imaging Centres

The imaging centre may be part of the hospital in terms of administration, but exist in a separate location and act as an independent entity, providing services not only to the parent hospital, but to other medical institutions.

This set-up is typical in situations where the hospital is able to procure the huge funds required to set up one of these centres, but can not guarantee the required capacity to utilise the equipment optimally. It is therefore most convenient for the parent hospital, as these centres are in close proximity.

In this case, the revenue model is split into two parts. The clients can be charged directly, in which case, the revenue model is similar to that of an out-patient imaging centre. The parent hospital may be charged outright for scans carried out at the centre, but a budget may already be allocated to the centre from the overall running costs of the parent hospital.

Fully Incorporated Imaging Centres

This is typical in general and teaching hospitals, whereby the radiology department exists for this purpose. In this set-up, the revenue model is tied directly to the revenue model of the entire hospital.

7.5.3 Summary

This section has split the practice of radiology into four major classes: tele-radiology, out-patient radiology, hospital-adjoined radiology, and fully-incorporated radiology practices. These classes differ mainly by the way they provide their services, and the extent of radiology work-flow provided. The classes are all driven by the needs of the referring physician, who is the primary customer of the radiology business. Each class has a different revenue model, although these models tend to cross over.

7.6 Market Analysis

7.6.1 Introduction

This section introduces Nigeria, providing information on key factors that describe the country; this description is based on a PEST (Political, Economical, Social, and Technological) analysis. The section goes on to present an overview of the healthcare industry, highlighting key market metrics, in order to provide information on the current state of the market.

Nigeria, or the Federal Republic of Nigeria, is a federal constitutional republic comprising thirty-six states and one Federal Capital Territory, Abuja. The country is located in West Africa, and shares borders with Niger in the north, Benin Republic in the west, Cameroon in the east, and Chad in the north-east. Its coast lies in the Gulf of Guinea, which is part of the Atlantic Ocean (see figure 7.9). Nigeria is the most populous country in Africa with a population of 144,720,000; it is followed by Ethiopia, with a population of 79,221,000.

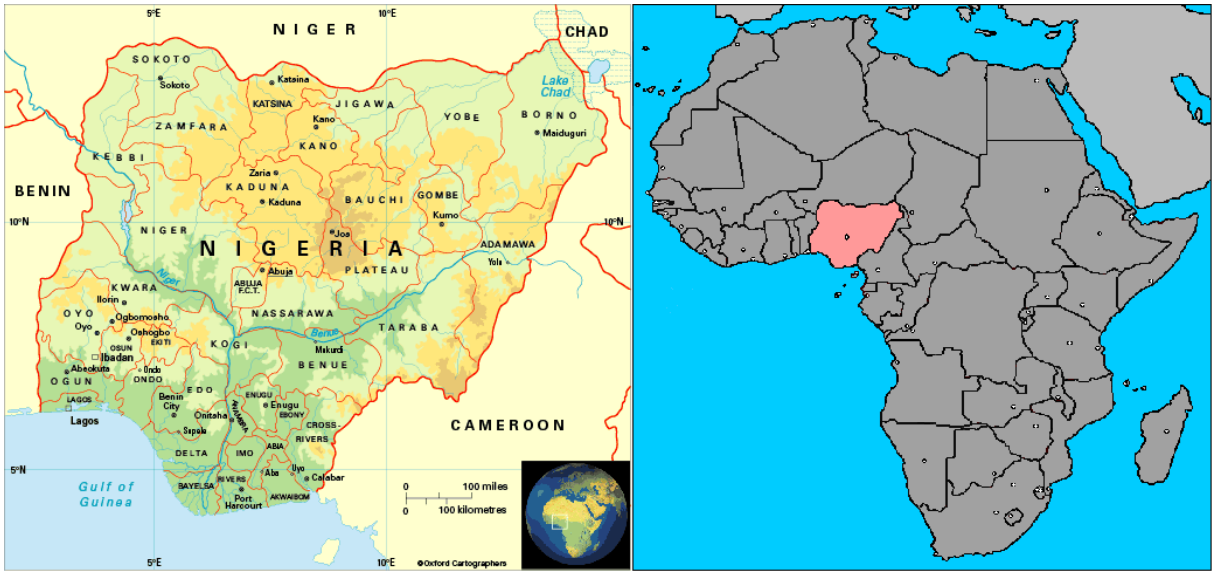

Figure 7.9: Map of Nigeria. [28][29]

Nigeria has the 9th largest income disparity in Africa, with a GINI index of 50.6; ranking lower than South Africa, but much higher than Ghana, which has a GINI index of 30.0. The GDP of Nigeria as of 2008 is US$214Bn, with the per-capita GDP being US$1,450. This ranks Nigeria just below South Africa as the African country with the second largest economic output, although inequitable income distribution results in a rank of 16th highest based on per-capita GDP. Despite this ranking, the World Bank ranks Nigeria as 118/181 in terms of the ease of doing business. The standard currency of Nigeria is the Naira, symbol (₦), and as of April 2009 exchanges for 145 US Dollars.

Nigeria received its independence from the United Kingdom in 1960, it is a member of the Commonwealth of Nations; although there was a ban in place from 1995 to 1999, Nigeria has since been a majority-rule democracy. The current ruling political party is the Peoples' Democratic Party (PDP); the current president is Umaru Yar'Adua, and the vice president is Goodluck Jonathan.

Table 7.3 shows an overview of Nigeria's demographic and healthcare data.

7.6.2 PEST Analysis

The PEST analysis allows for thorough investigation of macro-environmental factors that indicate the issues faced when conducting business in Nigeria.

- Political factors: These indicate the impact of the Nigerian government on the healthcare industry.
- Economic factors: These indicate the economic status of Nigeria, with particular focus on

Total Population (million)	144.7
GDP ($ Billion)	214
Total Fertility rate (births/woman)	5.5
Life Expectancy at Birth Males	48
Life Expectancy at Birth Females	49
Under 5 mortality rate (per 1000 live births)	97.1
Deaths due to HIV/AIDS	311
HIV Prevalence among adults aged 15+ years (%)	5.1
Physicians (% of Population)	0.023842
Nurses (% of Population)	0.145197
Total Expenditure on Health ($ Billion)	3.2
Total Expenditure on Health (% of GDP)	6.941432
State Expenditure on Health (% of Total Expenditure on Health)	25

Table 7.3: Nigeria: Demographics and Healthcare Data.

healthcare spending and other key economic indicators.

- Social factors: These provide information on the culture and habits of Nigerians, as well as demographics; focus is placed on Health indicators.

- Technology factors: These indicate the readiness of the Healthcare industry to incorporate new technology.

Political Factors

Nigeria went under military rule from 1966 up until 1999; during this time, there were high levels of unrest in the nation, with corruption deemed at an all time high towards the end of the regime. The current president, Umaru Yar'Adua came into power in 2007 and since then radical development plans have been proposed in form of a 7-point plan focusing on energy, security, wealth creation, education, land reform, mass transit, and the Niger delta conflict.

Nigeria is also witnessing the re-emergence of a middle-class[30], which indicates more employment opportunities for individuals, creating more levels of income. Poverty and illiteracy levels are still high in rural areas of the country, where access to communications and other infrastructure is lacking.

In 2005, Nigeria received a US$18 billion write-off, constituting 60% of its international debt by the Paris Club. The savings from that are accounted for in the budget as being utilised for achieving millennium development goals (MDGs) in the country and as such, an extra US$750 million is allocated yearly to health, energy, education and agriculture projects as capital expenditure.

The 2009 proposed budget estimates US$1.37 billion will be spent on public service reform and the state pension service. Spending on government departments, ministries and agencies

is totalled at US$15 billion. 40% of this sum is allocated to capital expenses, while the rest is allocated to recurrent expenses such as personnel costs. Table 7.4 shows the sectors that obtained the largest allocation of capital expenditure.

Department	Amount N(Billion)	Amount US$(Billion)
Transportation	114.2	0.787
Energy	110.3	0.76
Agriculture and Water	80.9	0.56
Health	34.5	0.24
Federal Capital Territory	30.6	0.21
Military	27.7	0.19
Interior	25.6	0.17
Education	23.8	0.16

Table 7.4: Capital Expenditure Rankings (Nigeria 2009 budget).
Capital Expenditure on health is ranked 4th below agriculture, energy and transportation, which are also industries in states of emergency.

The figures in Table 7.5 give the political indicators for Nigeria. The figures on public expenditure were obtained from the 2009 budget. These figures are low compared to benchmark African nations like South Africa, who spend 3.2% of their GDP on health and 5.2% on education.

Human Development Index	154/179
Public Health Expenditure (% GDP)	0.5
Public Education Expenditure (% GDP)	0.7
Military Expenditure (% GDP)	0.7
Debt Service (% GDP)	2.0

Table 7.5: Political Indicators
The Human Development Index is collected from the United Nations Development Program reports on HDI. The public spending figures are derived from the Nigerian national budget archive[31][32].

The Human Development Index (HDI) measures the average progress of a country in human development. It takes into account the life expectancy at birth, the adult literacy rate, and the per-capita GDP (measured by purchasing power parity) among other factors.

Economic Factors

Since the oil boom in the late 1960's, Nigeria has grown to be almost solely reliant on oil revenues to feed the country's budget. As oil prices start to drop, from a high of over $100 to a current price of $48, Nigeria has seen its revenue deplete drastically. This has led to moves in the budget to raise revenue by issuing government bonds, and a move towards restoring the once-great agricultural output of the nation. Agricultural produce such as cocoa, ground-nut and

cotton were once major exports before the oil boom; it is now the case that the finished products of these raw materials are being imported.

The nation is faced with a crippling power production problem. The National grid is operated by the Power Holding Company of Nigeria (PHCN), although smaller supply companies exist since the deregulation of the power sector in 2005. The main source of electricity to the grid is hydro-electric power supplied from Kainji dam and Kurra Falls. The bulk of power plants and transmission facilities were built in the 1950s and 1960s. Little investment and maintenance has left the infrastructure in a state of emergency, leaving the plants operating at as low as one-third of their installed capacity. The government plans to double the power output to 6000 Mega-watts by the end of 2009, although other sources of energy such as bio-fuels, solar and wind are being invested in through public-private-partnerships (PPPs).

The Nigerian banking industry witnessed a boom during the last three years, after a recapitalisation legislation passed during the term of former president Olusegun Obasanjo, which forced consolidation of smaller banks forming more reliable banking institutions. This led to abnormal increases in the Nigerian stock exchange index during that period, although the global recession is affecting the industry now because international investors are withdrawing money from the country.

Table 7.6 provides an overview of the economic indicators for Nigeria. The GDP figures rank Nigeria as Africa's second largest economy next to South Africa.

Gross Domestic Product (GDP) (International $Billion)	315
GDP per capita annual growth rate (%)	0.8
Average annual change in CPI	11.24
GDP per Head (US$) - poorest 20% (% of GDP)	4.4
GDP per Head (US$) - richest 20% (% of GDP)	55.7
GINI index	50.6

Table 7.6: Economic Indicators
The figures are normalised by purchasing power parity (PPP). An international dollar has the same purchasing power in a given country as a US dollar in the USA.

Social Factors

Nigeria has been a source of international concern in recent years; this is due to an inefficient social system within the country, despite its huge oil revenues. The GDP has been growing at a steady rate since the oil boom in the 1970's, but the social sector, particularly the health sector has witnessed a decrease in the proportion of the government's budget allocated to it (as seen in figure 7.10).

The key health indicators in table 7.8 show that private expenditure on health is over twice the government's expenditure on health, and that 90% of this expense is paid for outwith any

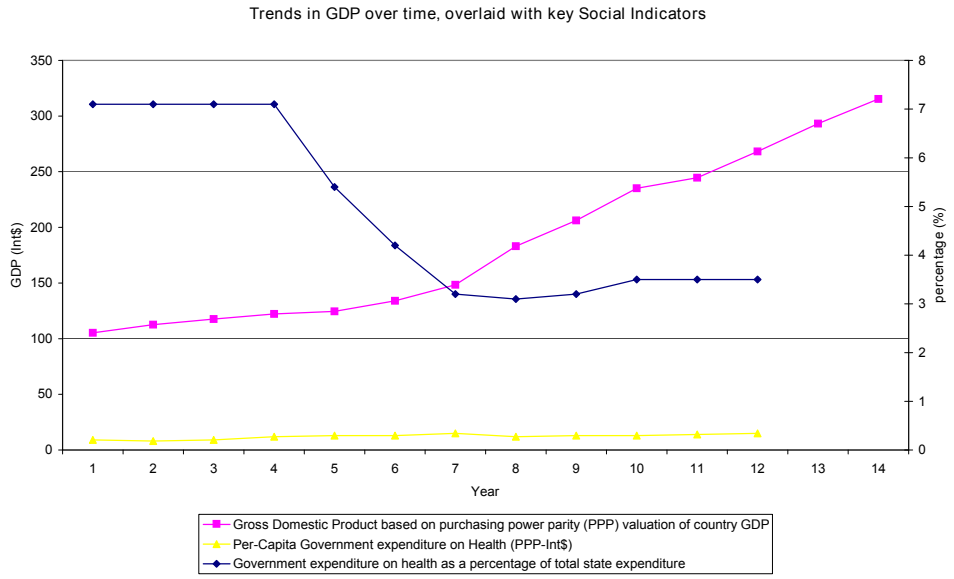

Figure 7.10: Trends in GDP growth over-time.
Government expenditure on health as a percentage of total government expenditure has been reducing, although the GDP has been increasing.

health plans. This naturally implies great disparities in primary health as available to different income levels of the society.

Per capita Total Expenditure on Health at average exchange rate (US$)	32
Per capita State Expenditure on Health at average exchange rate (US$)	10
Hospital Beds (% population)	0.05
Population annual growth rate (%)	2.4
Total Fertility Rate (per woman)	5.5
Adult literacy rate (%)	69

Table 7.7: Key Social Indicators
Data obtained from the World Health Organisation.

The fertility rate in Nigeria is high relative to other developing countries, for instance in South Africa, the fertility rate per woman is 2.7, compared to 5.5 in Nigeria. Moreover, the per-capita state expenditure on health in South Africa is US$157 compared to US$10 in Nigeria as shown by the social indicators in table 7.7. The fertility rate is geographically variable within Nigeria, with the North having a higher fertility rate per woman than the south. This variability is also present between rural and urban areas.

In terms of education, Nigeria, despite having the second largest economic output in terms of GDP, is ranked 18th in Africa, based on the adult literacy rate. Senegal is ranked first in this category.

At the current growth rate of the population, it is clear that the social system is in need of drastic reform if the welfare of the Nigerian populace is to improve.

Total Expenditure on Health (% GDP)	4.1
Government Expenditure on Health (% Total Expenditure on Health)	30.5
Private Expenditure on Health (% Total Expenditure on Health)	69.9
General Government Expenditure on Health (% Total State Expenditure)	3.5
External Resources for health (% Total Health Expenditure)	5.6
Social Security Expenditure on Health (% Total State Health Expenditure)	0
Out-of-Pocket Expenditure (% Private Expenditure on Health)	90.4
Private Pre-paid Plans (% Private Expenditure on Health)	6.7

Table 7.8: Key Health Indicators
Data obtained from the World Health Organisation, based on 2005 records.

Technological Factors

In recent years, there has been a sharp growth in the use of mobile phone technology in Nigeria (figure 7.11). This has evolved from simple phone-call only usage, to the adoption of data intensive tasks such as, 3G web-browsing and Blackberry/general push-email technology. This has seen the communications sector strengthen, with emergence of local mobile telecommunications companies. Initially, the only companies involved in the market in Nigeria were South Africa's MTN and ECONET, and Nigeria's fixed line operator NITEL. Since then, there are now six locally operating mobile providers in the country, with Globacom Limited and MTel being the only two indigenous companies in the sector.

To promote equitability in the use of mobile telephones, other companies offering cheaper mobile telecommunications technology have sprung up. Multilinks Limited is a company offering SIM-free CDMA mobile phones for a fixed price of US$10; this has led to a nationwide adoption of mobile telephony.

The internet situation is increasing at a much slower rate, one reason for this being the relatively low internet services available for residential purposes. Mobile telecommunications companies have exploited the inefficiency and low prevalence of fixed-line telephony by providing 3G and GPRS internet packages for laptop and desktop users, although the rates charged monthly for this service are much higher than broadband internet rates in Europe and America; generally about US$68 per month for unlimited access. Fixed-line dial-up services are still being provided at a much cheaper rate, although the bandwidth offered is below 28kb per second. Broadband internet is becoming increasingly popular. This service uses satellite technology, VSAT, mainly because the telephone lines cannot support the high-speeds which are characteristic of broadband internet. This is the prevalent form of internet amongst business and other corporate bodies. In all, there are over twenty companies providing internet in Nigeria; but the

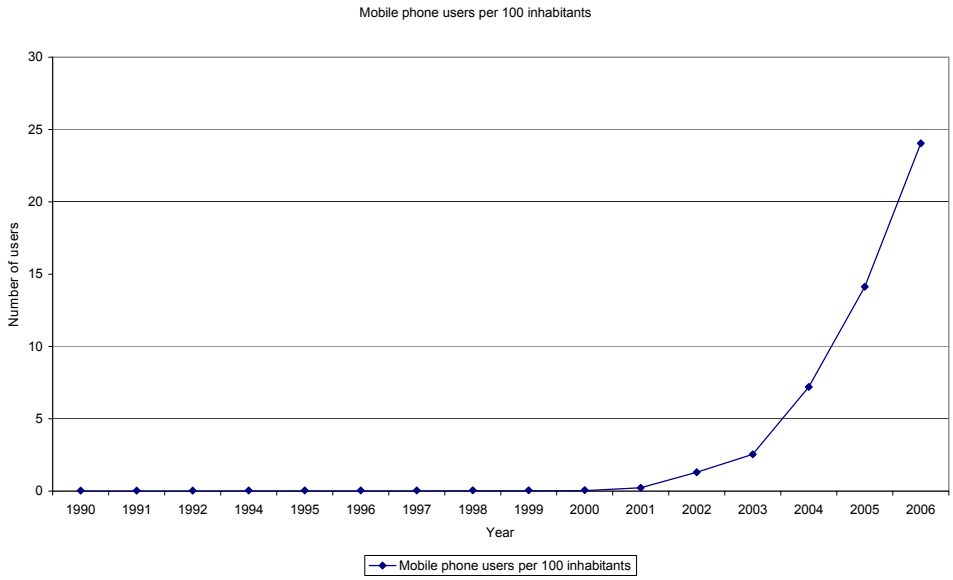

Figure 7.11: The growth in users of mobile phones in Nigeria.
The sharp rise from 2004 was as a result of huge drops in the prices of mobile phones and SIM cards. The relatively poor quality of fixed-line telephony is also an aid to this rise.

low rate of adoption is predicated by poverty, implying an inequitable distribution of computers across the nation. The Nigerian government has been involved in some of the schemes put in place by international organisations to provide personal computers to children at an early age. The government has invested in schemes such as One Laptop per Child (OLPC) and Intel's Classmate programme, with an aim to tackle computer illiteracy in Nigerians from a young age.

The banking sector, following their recent boom, have advanced sufficiently on the technology front; allowing for much easier transactions with international banks. This change has resulted in the emergence of automatic teller machines nationwide, provision of 'plastic' as a new means of payment, and internet/mobile-telephone banking.

Network Readiness Index (rank)	3.45 (rank = 90/133)
Main Telephone lines per 100 inhabitants	0.7
Mobile phone subscribers per 100	12.9%
Cost of a 3-minute mobile phone call (US$)	0.22
Personal Computers per 1000	6.02%
Internet users per 1000	6%

Table 7.9: Technology Indicators

The network readiness index measures the propensity of a country to exploit the opportunities offered by information and communications technology. It is a term developed by the World Economic Forum, as a means to rank countries in terms of technological advancement. This and

other technological indicators are presented in table 7.9.

7.6.3 Healthcare Industry Overview

The state of healthcare in Nigeria is beginning to recover from an all-time low, when in 2000 it was ranked 187^{th} out of the 191 member states of the WHO. The healthcare system, as in most countries, is primarily the responsibility of the government; although private organisations still play a major role in the system. Nigeria maintains three tiers of healthcare:

- Primary healthcare: This provides essential healthcare as a community-wide service, and is mainly supplied by the local government and private health centres.
- Secondary healthcare: Handles referrals from the primary healthcare centres, normally dealing with more specialised cases and emergency cases. This tier is provided by the state government and private health institutions.
- Tertiary healthcare: Handles highly specialised treatments. This tier is provided by specialist private hospitals and federal government teaching hospitals.

The Federal Ministry of Health is the arm of government which deals with affairs of health in the country. The minister for health has pointed out that the primary healthcare sector currently caters for 20% of its potential load due to a number of factors including loss of confidence by its customers, inefficient facilities, and a lack of health management information systems. The ministry has since put in place a National health investment plan aimed at improving health infrastructure and providing suitable medical equipment to state-owned health institutions. At a meeting of African leaders in 2001, all governments agreed a target of 15% of total government allocations to the healthcare system; in order to achieve this, a number of instruments have been set up. The National Health Insurance Scheme (NHIS) kicked-off properly in 2005, and since then has achieved 60% coverage for government employees, which accounts for 30% of Nigerians. A conditional grant scheme has also been put in place, funding private companies willing to pursue healthcare development in rural areas.

The government, in a move towards wealth creation, has also endorsed public-private partnerships in the healthcare sector; they are mostly targeted at capacity-building schemes that aim to provide measurable output.

The burden of disease as seen in Nigeria is predicated on the degree of poverty, therefore special policies need to be put in place to target the population living in rural areas with poor communication and transportation infrastructure. This has spawned several pro-poor schemes giving free ante-natal, newborn and child healthcare services. This, in particular is aimed at achieving the millennium development goal, which states a mortality rate of 74 per 100,000

live births by 2015(this figure is currently 800 per 100,000 live births, one of the highest in the world).

Public Healthcare

The public healthcare system constitutes the health institutions and health management organisations (HMOs) funded solely by the government. In Nigeria, the three tiers of government are:

- Local Government: Catering for the sub-state constituencies. These are usually small communities, and each state could have about 20 such local government areas.
- State Government: On the state level; these are headed by the governor, and each state can function independent of the Federal Government to an extent. Nigeria has 36 states.
- Federal Government: The top tier of government, which oversees the affairs of the entire country.

As of 2008, there were over 9000 public health institutions in Nigeria with a majority run by the local governments (figure 7.12).

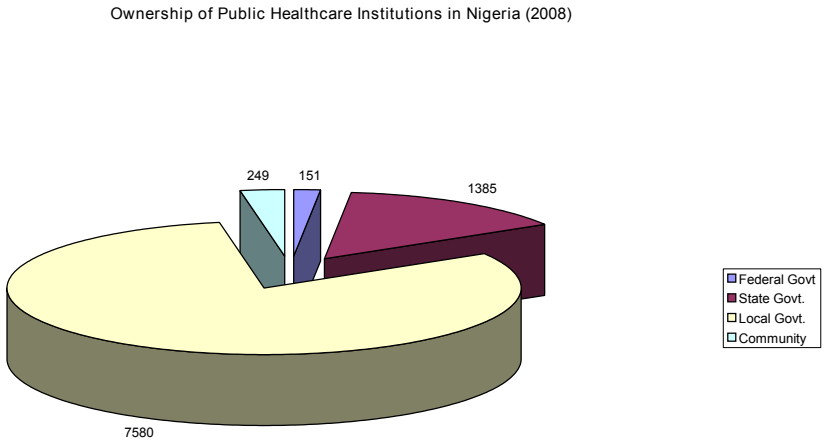

Figure 7.12: Ownership distribution of public health institutions in Nigeria (2008)
Local governments run the majority of health institutions. Source: The National Bureau of Statistics.

Public health bodies are fully funded by the government as part of the yearly budget. This funding accounts for personnel, equipment, and all other logistical processes. There are 48 teaching hospitals nationwide, with 15 major teaching hospitals on the state-level. These hospitals were targeted as vital to the state of health in the country, and the government pledged

a revitalisation package which saw amongst other things, the acquisition of MRI scanners in 6 of these teaching hospitals in 2005.

Private Healthcare

The private healthcare system in Nigeria consists of all health institutions and HMOs which are neither funded nor controlled by the government. These exist as independent entities, and unlike South Africa, there are no "super" institutions that have majority shares of the market. As of 2008, there were over 7000 of these institutions in operation in the country, among these are several specialist institutions (figure 7.13).

A class of private healthcare institutions commonly known as non-government organisations (NGO) also exists. These institutions do not necessarily fit into the three tiers of healthcare as defined, but mostly engage in research activities tailored to the general health of the population. Activities such as HIV/TB research and monitoring, and a lot of these organisations are internationally funded.

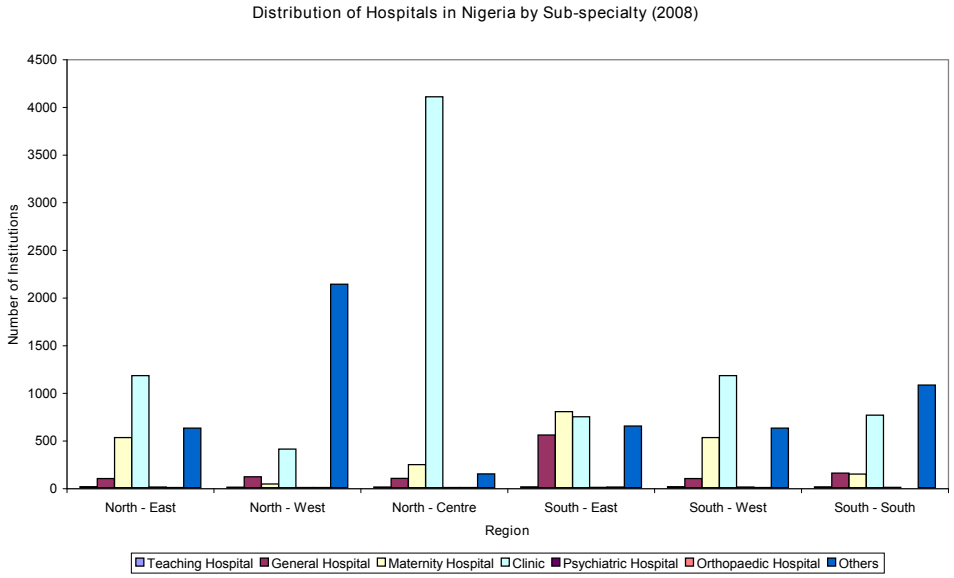

Figure 7.13: Distribution of Hospitals in Nigeria by Sub-specialty (2008)
Clinics are still the most prevalent form of private health centres, although maternity hospitals are the second most popular specialtist institutions.

Private health institutions in Nigeria are financially responsible for their personnel and equipment, but they gain revenue through reimbursements by the NHIS and other private health insurance schemes, as well as out-of-pocket fees paid by non-registered patients. In order to be registered as a health provider in Nigeria, private institutions need to obtain a license to practice from the state's health board; this licence is renewed yearly to maintain standards and up-to-date

records.

Figures 7.14 and 7.15 show the ratio of publicly owned to privately owned healthcare institutions around the country.

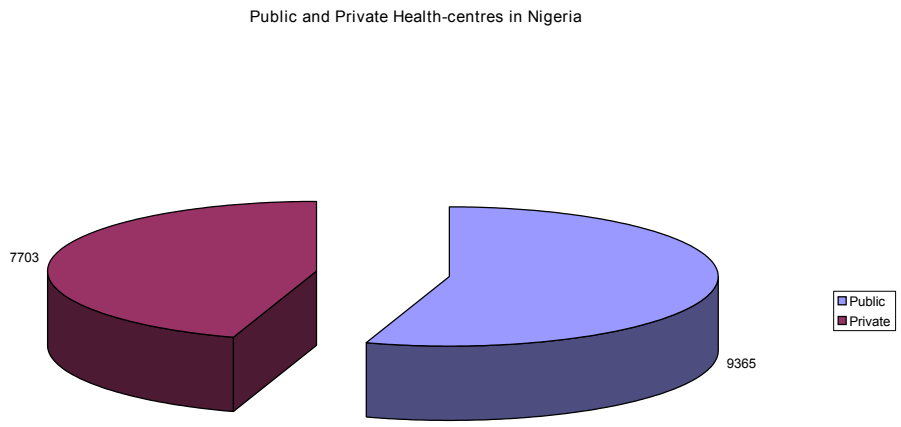

Figure 7.14: Distribution of Ownership of health establishments in Nigeria. Source: The National Bureau of Statistics.

National Health Insurance Scheme

The NHIS is a body set up by the Federal Government of Nigeria to provide healthcare to the general population in a cost-effective and efficient way, providing room for private healthcare providers to earn revenue and compete in a healthy manner.

The NHIS act was passed in 1999, but it officially kicked off in 2005. The first phase of operation was to cater for core civil servants, i.e. staff of public corporations, police, military personnel and uniformed paramilitary personnel. To date, it has achieved 60% coverage of this population sector.

The NHIS has since created several packages tailored to the various sectors of the economy. For the organised private sector, it is a mandatory requirement that all participants contribute 15% of their basic salaries to the scheme; 10% paid by the employer and the rest paid by the employee. This covers registrants for out-patient care, diagnostic tests, drugs, maternity care, immunisation, dental, and eye examinations. For the rural areas, the scheme works on a community basis, whereby a group of at least 500 individuals come together and contribute a monthly flat-fee according to healthcare requirements.

Health Management Organisations (HMOs) function as brokers in this system. Private and public health institutions register with these bodies as healthcare providers, and a list is made available to every registered patient member of the NHIS to choose their provider. This allows healthy competition amongst providers, as reputation is the driving force behind client choices.

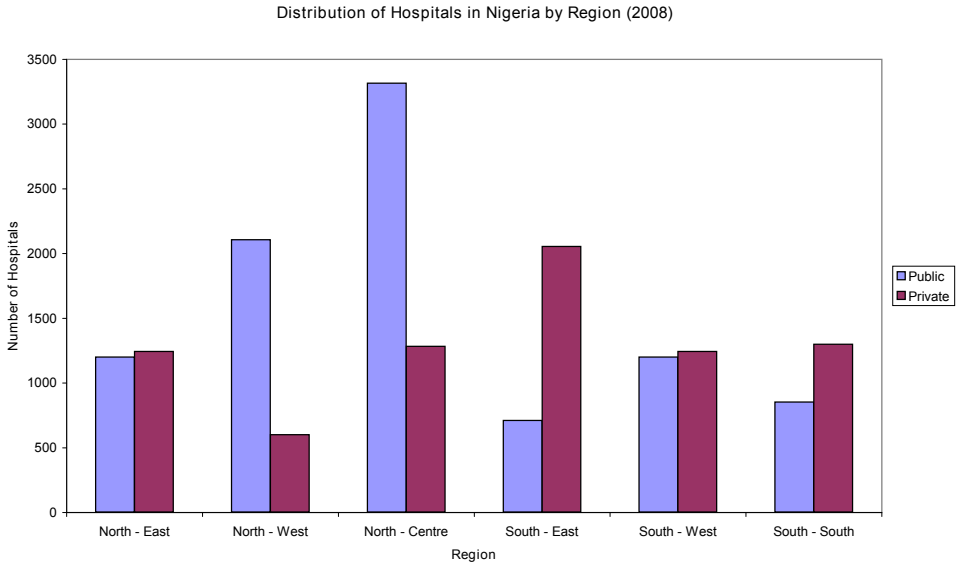

Figure 7.15: Distribution of hospitals in Nigeria by ownership
The north-centre region of Nigeria contains more public hospitals than any other region of the country. This is also the region where the capital, Abuja is located (See map in figure 7.9). The south-south region is the oil producing region of Nigeria. Source: The National Bureau of Statistics.

The providers receive remunerations per head from the NHIS via the HMO; other services such as prescriptions and diagnostic tests are classed as fee-for-service, and reimbursements are paid for such services.

7.6.4 Medical Imaging Market

The X-Ray and CT imaging activities in Nigeria are regulated by Nigeria Nuclear Regulatory Agency (NNRA). They require a registration fee of about US$450 a year for every X-Ray and CT scanner in operation in the country. Importation of all medical equipment is regulated by the National Agency for Food and Drug Administration and Control (NAFDAC). This is equivalent to the Food and Drug Administration (FDA) in the United States.

Medical Imaging activities are spread across both private and public sectors. In the public sector, it mainly takes the form of radiology departments in the general and teaching hospitals. In the private sector, several hospitals and clinics run radiology departments, and operate through referrals from public and other private institutions. Outpatient imaging centres do exist however, and these operate strictly by referrals, offering general laboratory services as part of their product portfolio.

Revenue models for imaging centres in Nigeria

Of all the radiology models presented in section 7.5, the most prevalent forms of the business in Nigeria are fully incorporated imaging centres, hospital-adjoined imaging centres, and outpatient imaging centres. The sources of revenue for these practices are as follows:

- **Fully incorporated imaging centres**: An example of this model is seen in the National hospital, Abuja and other teaching hospitals across the country. This is a government-owned hospital, and so all its funding is provided by the state. This allows it to offer services at a relatively lower price than private institutions. These institutions are also on the NHIS list of registered health providers, and so receive fee-for-service reimbursements for all diagnostic imaging performed. Services provided to non-NHIS patients warrant out-of-pocket payments from these patients. Although normally these institutions would only provide radiology services to their own patients, there are a lot of primary-care clinics without these facilities, and so they frequently receive referrals for diagnostic imaging.

- **Hospital-adjoined imaging centres**: This model is prevalent in the private sector, mainly among the bigger hospitals. Examples of this model can be seen in Kolstaf medical centre and Zankli medical centre, both located in Abuja. These are primary-care providers that also provide radiology and laboratory services. As accredited NHIS health providers, they receive fee-for-service reimbursement for each scan performed. Due to the premature stage of the NHIS, a bulk of revenue comes in from private health insurance companies and out-of-patient payments by uninsured patients. In order to maximise capacity, such institutions tend to actively seek referrals from other hospitals, as well as the national hospitals (in times of equipment failure, etc).

- **Outpatient imaging centres**: An increasingly popular model in large cities in Nigeria. Examples of such institutions are Union Diagnostic and Clinical Services in Lagos, and Image Diagnostics in Port-Harcourt. These institutions operate a strictly referral-based service, and their niche allows them to provide a wide range of imaging modalities. Their revenue comes mainly in form of reimbursements from the referring institutions, because they are able to make special arrangements for patient referrals. Outpatient imaging centres located in cities with foreign companies and thus high expatriate population, normally make arrangements with such companies to act as health providers for their employees, hence receiving reimbursements from these firms.

7.6.5 Summary

Nigeria is categorised on most standards as a developing economy; with most of its income generated from oil revenues, the government is actively pursuing diversification in the wake

of falling oil prices. The government of Nigeria is now ten years into democracy, and already improvements and stability are visible. The Federal ministry of health, which is the arm of government responsible for healthcare, is aware of the developmental problems with the system and is implementing policies to improve healthcare infrastructure within the country. Poverty is the underlying problem crippling the market at the moment; preventing healthcare from being efficient in the rural areas. There are now growing numbers of health insurance bodies, including the NHIS, which is run by the government. These bodies are not only providing guaranteed healthcare to participants across all income levels, but also provide health providers with a source of income and an incentive to keep developing their services. The medical imaging market is still growing in Nigeria, with all of the equipment being imported from outside Africa.

7.6.6 Data Sources

World Health Organization (WHO), United Nations (UN), International Monetary Fund (IMF), World Bank, African Development Bank, Nigerian National Bureau of Statistics, Central Intelligence Agency (CIA) factbook.

7.7 Summary

Medical imaging is defined as the techniques and processes used to create images of the human body, and it is used in modern medicine as a means of diagnosing, monitoring and treating diseases. The medical imaging modalities of interest in this study are x-ray, CT, ultrasound, MRI and PET. Picture Archival and Communication Systems (PACS) are used to display and share image data of patients; these were not discussed in this report due to time constraints.

Epidemiology studies on Nigeria as provided by WHO, show that the seven most burdensome disease categories are infectious diseases, respiratory infections, unintentional injuries, neuropsychiatric conditions, perinatal conditions and cardiovascular diseases. Studies show that the imaging modalities presented here are efficient at diagnosing specific conditions from each of these categories. The report shows that if all these conditions are to be targetted in a cost-effective manner, a medical imaging centre would require x-ray, CT, and ultrasound. Providing the scanners alone for these modalities would cost about US$2 million making a capital-intensive enterprise, especially once building and running costs are factored in.

Medical imaging centres can be run as any of four different business models namely; tele-radiology, outpatient, fully integrated, and hospital-adjoined. The most popular forms of the business in Nigeria are the three latter models; the lack of tele-radiology practices may be due to a shortage of modality equipment in the first place, or a general inefficiency of the underlying communications infrastructure.

The population of Nigeria is about 140 million, spread across thirty-six states. There is a large income disparity between people living in the urban areas and those living in the rural areas, this goes as far as to affect the quality of health available in both settings. The government of Nigeria has been under democratic rule for ten years now, and growth in the economy is visible. This has resulted in more funds being allocated to improving public healthcare, and making it affordable to the low-income population. A national health insurance scheme is now in place, along with other private health insurance schemes. These guarantee income for private medical practitioners, with full reimbursements paid for radiology services performed.

As a result of the stability in the political environment and the improving economy, international organisations like the International Finance Corporation (IFC), the World Health Organization (WHO) and Africa Development Bank (AFDB) have provided finance for commercial healthcare projects in Nigeria. The government is also engaging in PPPs to enhance healthcare, especially in rural areas.

Most of the data used in this report was provided by international organisations. The epidemiology data was provided by the World Health Organization by estimating national data from data collected from few parts of the country. This data was used regardless, because there is no such data available publicly from local sources, making it difficult to conduct such a study from outside the country. This points to a desperate need for healthcare information records to be held at a central location and made available for such research if the state of health is to be improved.

In order to take this report closer to being a complete feasibility study from which implementation ideas can be derived, more information is required. Information concerning the modality distribution across the country is not presented here due to difficulty in obtaining the data. Other information regarding the average prices per scan have also not been presented, as this varies across the country. Certainly, for this particular case, i.e. Nigeria, local data gathering has to be performed.

References

[1] T. M. S. Europe, "Aquilion-one." http://www.toshiba-medical.eu/en/Our-Product-Range/CT/Systems/Aquilion-ONE/ (accessed: 06/04/09).

[2] Hong-Kong University, "Revealing the magic in everyday life." http://www.physics.hku.hk/~phys0607/lectures/chap06.html (accessed: 06/04/09).

[3] "Chapter 1: Introduction to CT physics." http://docs.ksu.edu.sa/PDF/Articles27/Article270699.pdf (accessed 06/04/09).

[4] College of Medicine: University of Illinois, "Basic concepts in diagnostic imaging: Modalities in radiology." https://www-s.med.uiuc.edu/m2/icm/BasicconceptsindiagnosATversion1.pdf (accessed: 06/04/09), April 2009.

[5] Radiological Society of North America, "Safety : Radiation Exposure in X-ray Examinations," pp. 1–4, 2009.

[6] Florida State University, "Magnet lab: National high magnetic field laboratory." http://www.magnet.fsu.edu/education/tutorials/magnetacademy/mri/ (accessed: 05/04/09).

[7] E. Berry and A. Bulpitt, *Fundamentals of MRI: An interactive Learning Approach*. CRC Press, 2009.

[8] C. Westbrook, *Handbook of MRI Technique*. Blackwell Science Ltd, 1994.

[9] L. Shenzen Shenchao Transducer Co., "Human resources." http://www.probe-china.com/en/r-03.htm (accessed 08/04/09).

[10] T. S. Curry III, M.D., E. Dowdey, James, and R. C. Murry, Jr., *ChristensenŠs Physics of Diagnostic Radiology*. Lea & Febiger, 1990.

[11] Society of Nuclear Medicine: Resource Centre. http://interactive.snm.org/index.cfm?PageID=1107&RPID=924 (accessed: 15/04/09).

[12] Radiology, "General nuclear medicine." http://www.radiologyinfo.org/en/info.cfm?pg=gennuclear (accessed: 15/04/09).

[13] Themesotheliomalibrary.com, "Authoritative library on mesothelioma lung cancer." http://www.themesotheliomalibrary.com/diagnosis-techniques.html (accessed: 15/04/09).

[14] T. S. Stashak, *AdamsŠ Lameness in Horses*. Williams and Wilkins, 5.

[15] J. Ruhlmann, P. Oehr, and H. Biersack, *PET in Oncology: Basics and Clinical Applications*. Springer-Verlag, 1999.

[16] Intercollegiate Standing Committee on Nuclear Medicine, "Positron emission tomography: A strategy for provision in the uk," 2003.

[17] World Health Organisation, *The Global Burden of Disease*. World Health Organisation, 2008.

[18] S. B. Thacker, "Measuring the public's health," *Public Health Reports*, vol. 121, January-February 2006.

[19] ACR, *ACR practice guideline for the performance and interpretation of cardiac computed tomography (CT)*. American College of Radiology, 2006.

[20] ACR, *ACR Appropriateness Criteria: Suspected Abdominal Trauma*. American College of Radiology, 2005.

[21] ACR, *ACR Appropriateness Criteria: Suspected Spine Trauma*. American College of Radiology, 2007.

[22] ACR, *ACR practice guideline for the performance of an ultrasound examination of the abdomen and/or retroperitoneum*. American College of Radiology, 2007.

[23] ACR, *ACR practice guideline for the performance of obstetrical ultrasound*. American college of Radiology, 2007.

[24] ACR, *ACR practice guideline for the performance of screening and diagnostic mammography*. American College of Radiology, 2008.

[25] ACR, *ACR practice guideline for the performance of a breast ultrasound examination*. American College of Radiology, 2007.

[26] ACR, *ACR practice guideline for the performance of ultrasound evaluation of the prostate (and surrounding structures)*. American College of Radiology, 2006.

[27] S. K. Mun, PhD, "Teleradiology and Emerging Business Models," *Journal of Telemedicine and Telecare*, vol. 11, pp. 271–275, 2005.

[28] "Luventicus maps." `http://www.luventicus.org/maps/africa/nigeria.html` (Accessed: 23/04/09).

[29] The Commonwealth, "The commonwealth secretariat.." `http://www.thecommonwealth.org/YearbookHomeInternal/138917/` (Accessed: 23/04/09).

[30] National Bureau of Statistics, Nigeria, "The middle class in nigeria: Analysis of profile, determinants and characteristics (1980-2007)," *Economic Reform & Governance Project*, July 2007.

[31] "Federal Government of Nigeria 2009 Budget Preparation and Submission Call Circular.," August 2008.

[32] O. Obasanjo, “2007 budget speech,” *Joint Session of the National Assembly Abuja*, October 2006.

Chapter 8

Conclusions and Further Work

The technical reports in this thesis describe implemented solutions to all of the projects tasked with over the course of the EngD programme. Each chapter presents the results obtained from both the objective validation, using known experimental techniques and subjective evaluation by clinical personnel. The layout of this thesis shows the evolution of atlas-based techniques to match the problem at hand. The proposed methods show promising results with clear room for improvement. Suggested avenues for improvement are outlined in the individual chapters. This thesis demonstrates that atlas-based techniques are suitable for solving a range of clinical segmentation problems.

This chapter presents the conclusions from each chapter, summarising the methods to show a progression of the methods applied. The further work suggested in the chapters of this thesis share common themes, and these are discussed at the end of the chapter.

8.1 Conclusions

Automatic Labelling of Coronary Arteries in 3D CT Images Chapter 2 describes a method for labelling segmented coronary vessels in CCTA images. The method is a two-step process: 1) knowledge-based assignment, whereby all plausible labelled trees are generated and 2) statistical classification, whereby the most likely labelling is chosen based on closeness to parameters modelled by a multivariate Gaussian classifier.

This was the first piece of work undertaken within the EngD programme, and as such the concept of an atlas is less explicit. The atlas in this case is a syntactic model, representing the topological and geometric relationships between arterial segments (see figure 8.1).

The method was validated against ground-truth on a set of CCTA datasets of varying coronary anatomy. Further subjective validation was carried out by a radiologist, based on clinical requirements. The algorithm was deemed useful due to high accuracy for proximal segments of the major arteries, providing the user-interface allowed the user to validate and change the

automatic labelling in bad cases.

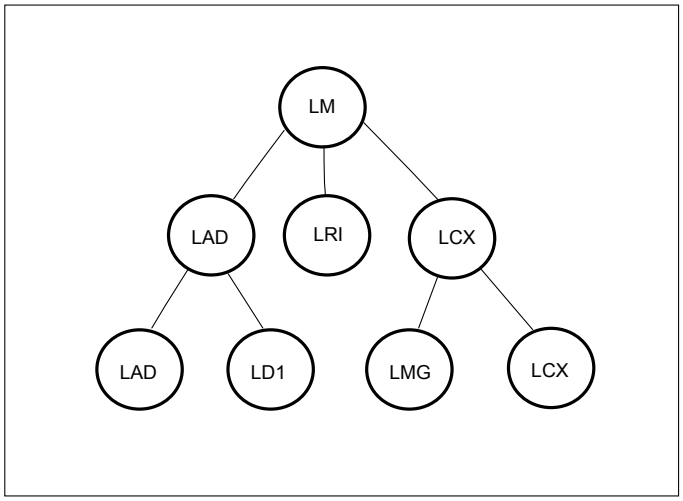

Figure 8.1: Syntactic Atlas Used for Assigning Labels
Only combinations of labelled polyline trees conforming to this topology were generated and assigned to the classifier for scoring.

Creation of a Coronary Artery Atlas to Guide the Tracking of Coronary Vessels The work presented in chapter 3 demonstrates that an atlas, representing knowledge of the coronary anatomy can be constructed using a registration strategy. The chapter presents one possible means by which such an atlas can be applied to the output of a coronary vessel tracking algorithm and improve its specificity by identifying and removing non-arterial vessels.

This was the first piece of work in the thesis to address anatomical variability within a group using a structural atlas. A reference CT image was selected from a set, while manually segmented arterial centrelines from the remaining images within the set were transformed to this reference space following registration. An adaptive anisotropic smoothing step was included to mitigate the impact of over-fitting caused by using limited number of training datasets (see figure 8.2).

Optimal Atlas Selection Using Image Similarities in a Trained Regression Model to Predict Performance Chapter 4 presents a method for optimal atlas selection from a set of candidate atlas images. The aim of a multi-atlas is to target variations in anatomy by building a set of atlases representative of the variability in the training set. Atlas selection arises in two contexts; firstly there is the offline selection of the datasets to include in the multi-atlas set; and secondly during application to a novel image, the best fitting atlas from the set is selected.

Each atlas contains an image, a mask containing the labelled heart and a pair of linear regression parameters (see figure 8.3). The latter parameters are obtained from a training set of images by co-registering each pair of images and recording the resulting similarity measures and Jaccard overlaps of the generated segmentation. The regression parameters represent a linear

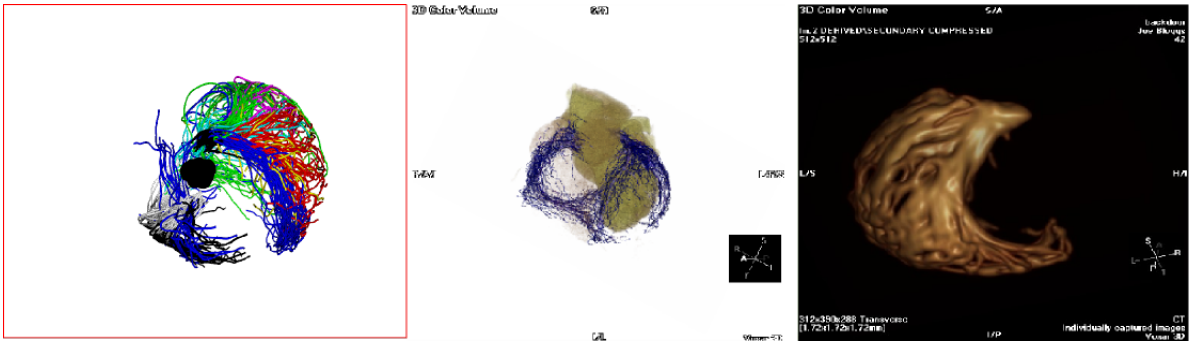

Figure 8.2: Addressing Anatomical Variability Using an Atlas of the Coronary Vasculature
The atlas was constructed by registering manually segmented coronary trees from 42 datasets onto a reference dataset.

relationship between image cross correlation (an internal measure not requiring a mask) and the Jaccard overlap of true and computed masks, thus the latter can be predicted from the former.

The optimal atlas from a set is selected following registration by measuring the cross-correlation between each atlas image and the novel image and applying the respective linear regression slopes and intercepts to predict the Jaccard overlap. The atlas image that predicts the highest segmentation accuracy is selected.

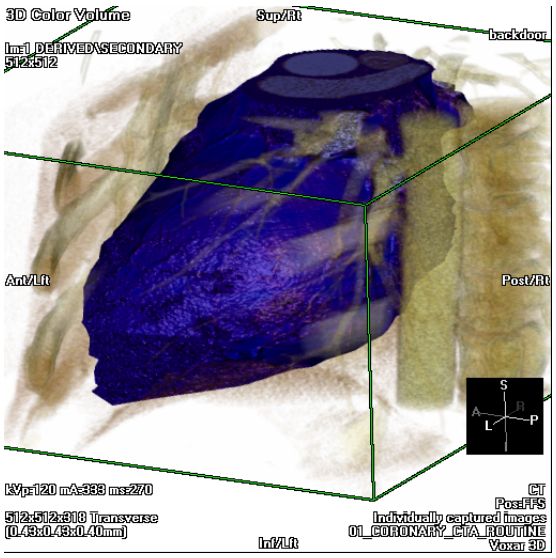

Figure 8.3: Modality Specific Atlas
The atlas is a CT volume with a manually labelled mask of the whole heart.

The results demonstrate that a multi-atlas approach performs better on average than a single atlas method, given the same registration method. This was tested on a set of cardiac CT images with visible variations in the cardiac anatomy, and the multi-atlas set selected contained images from each of the pre-identified categories. The next chapter applies this method to segmenting the kidney and renal cortex from noisy CT data with positive results.

A Combined Multi-Atlas and Unsupervised Classification Technique for Automatic Segmentation of the Kidney and Renal Cortex in Low-Dose CT Data Chapter 5 describes an atlas-based method of segmenting the kidneys and renal cortices from low-dose contrast enhanced abdominal CT scans using a multi-atlas with the optimal atlas selection procedure described in chapter 4.

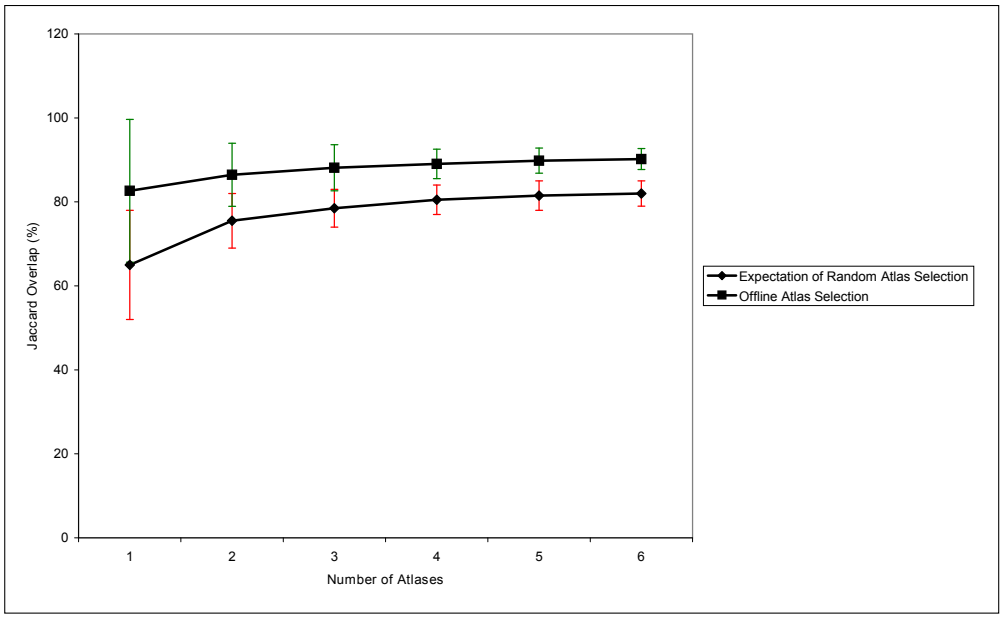

Figure 8.4: Benefits of Optimal Multi-Atlas Selection Strategy
The upper line shows the Jaccard overlap averaged over all validation data when selecting the optimal atlases for each set size (equation 5.7), while the lower line shows the mean of the Jaccard overlap averaged over all training data for all possible random selections of atlases (equation 5.10).

Figure 8.4 demonstrates the advantage of using the optimal atlas selection approach; both at the offline stages, where it is used to select a set of atlases to be used for registration, and during the actual application, where it is used to select the best segmentation from a list of candidates. This was the first piece of work in the thesis that incorporated multiple structures in the atlas, the whole kidney and the renal cortex in this case (see figure 8.5).

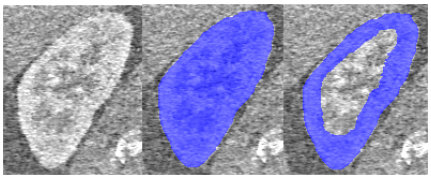

Figure 8.5: Multi-Object Atlas
The atlas is still a CT volume (left), but in this case contains 2 labelled structures, the whole kidney (middle) and the renal cortex (right).

This chapter discusses an unsupervised voxel-based classification as a post-processing step,

using the expectation-maximisation (EM) algorithm to estimate the intensity distributions of each structure. The results show an increase in accuracy for both structures gained due to the post-processing.

Robust Segmentation using Atlas-Based Priors in the EM Algorithm Chapter 5 exposed an important limitation to the post-processing step. The method relied on prior knowledge of the relative intensities of the renal cortex and the medulla, therefore the intensity distributions estimated by the EM algorithm were assigned to the respective structures following a specific order. Chapter 6 describes a robust method for segmentation, which incorporates spatial information from the atlas in the EM algorithm. The extra information allows the EM algorithm to distinguish between different tissues with similar intensity distributions.

The versatility of this method is demonstrated by applying it to two different segmentation applications. The first, kidney and renal cortex segmentation from low-dose abdominal CT data, demonstrated the robustness of the algorithm by applying it to noisy data. The second, cardiac sub-structure segmentation, uses the method to segment multiple structures within the heart with no explicit encoding of anatomical prior information other than that provided by the atlas.

Both experiments demonstrate that there are advantages to be gained by incorporating spatial priors in the EM algorithm. Experiments carried out on kidney segmentation produced more accurate estimates of the distribution models, and this led to better segmentation of the renal cortex. Due to difficulty in obtaining a suitable quantity of ground-truth for cardiac sub-structures, only subjective validation was performed in the second case. In this scenario, the various structures of interest shared similar intensity values, this led to random assignment of classes to distributions with the naive EM implementation. Upon inclusion of spatial priors, the correct assignment was made.

8.2 Further Work

Expansion of training and validation sets Each algorithm described in this thesis relies on a training phase to either select an atlas (as in chapters 4, 5 and 6), or to train a statistical classifier (as in chapters 2 and 3). In chapter 4 for instance, selection of an optimal multi-atlas set assumes that all anatomical variations are present in the training set. Chapter 2 estimates per-class covariance matrices of a set of features, the accuracy of this estimate is directly proportional to the number of training data available.

Chapter 3 aims to create a statistical domain capable of providing information about coronary anatomy of a population subset, but coronary anatomy is divided into three main circulation patterns, namely left dominant, right dominant and balanced circulation anatomies. It is therefore desirable to have an atlas for each circulation pattern, and this requires a lot more training data

for each category to reduce over-fitting.

The datasets used in validation were drawn from a small number of institutions, biasing the results towards several factors such as the scanning protocols used, ethnicity of the patients and the predominant pathology at these institutions. The largest validation set in this thesis contained 42 datasets, others contained as few as 20 datasets. This paucity of validation data led to high variance in the accuracy figures reported. Further work will be to evaluate these methods on a wider variety of real world data.

Enrichment of class set Every chapter in this thesis attempts to solve a segmentation problem, whereby voxels in the image are assigned a label from a pool of possible classes. The pool is chosen as the set of classes relevant to the end-user, as a result several possible classes with varying properties are bundled into one class. Examples are the LINSIG and RINSIG classes in chapter 2, 'non coronary artery' vessels in chapter 3, and 'background' in chapters 5 and 6. This creates feature distributions with falsely large variances, diminishing the overall quality of the segmentation.

This effect can be reduced by identifying more classes in the training data; for example, the background class in the heart segmentation in chapter 6 can be split further into auricles, pulmonary veins, superior vena-cava and pulmonary artery. The aim is to arrive at classes with homogeneous properties, implying tighter feature distributions and thence more accurate classification.

Improved Alignment Atlas-based segmentation fundamentally relies on the automatic alignment of a novel image to the atlas. In this thesis, alignment is achieved using registration.

Chapters 5 and 6 discuss a method for refinement of the registration-based segmentation by combining intensity information with the spatial priors provided by the registration. The warped atlas masks, represented as binary images are smoothed by an isotropic Gaussian kernel to provide a soft segmentation which aims to address any inaccuracies in the registration. This assumes the error is evenly distributed across the structure; but this is not always true as demonstrated by the over-segmentation errors at the apex-rib wall boundary of the heart in chapter 4, and hence the use of a multi-atlas. This approach can be extended by pre-aligning several manually-segmented masks to the atlas image and storing a probabilistic mask as part of each atlas. This mask will provide the soft segmentation needed in a form possibly capable of addressing anatomical variability at a more granular level especially in sub-structure segmentation.

Several publications reviewed in this thesis attempt multi-compartmental segmentation by using local affine registration, whereby each sub-structure is registered independently, following a global affine registration step. These publications have already demonstrated that this approach produces a more accurate final alignment. Further work will be to incorporate a local rigid registration phase in the multi-compartment heart segmentation.

In chapter 5, the rapid initial search phase of the registration is a critical step, which aims at placing the atlas at the centre of the kidney in the novel image as a starting point for full multi-scale registration. The current implementation initialises a constrained powell search at evenly spaced positions along the anterior-posterior axis at the centre of each half-volume. Initialising these searches at points placed evenly within a 3-D grid about the centre of each half-volume will be more robust to variations in the scan field of view.

Majority of the execution time of each algorithm discussed in the thesis is spent during registration, thus any speed-up in registration will improve the usability of the algorithms, particularly multi-atlas based approaches. The multi-atlas based approach described in chapters 4 and 5 perform sequential registrations of each atlas to the novel image. These registrations are completely independent and should be performed in parallel, taking advantage of multi-core processors to provide a speed-up.

Investigation of more powerful features for classification Classification of voxel data relies on measurable features that best describe the properties of each label, therefore feature selection is very important.

Chapter 2 uses 14 features to assign each arterial centreline to its anatomical label. These features were chosen by observing the data and discerning the physical features that best described each vessel type, but there were still several predominant misclassifications appearing during the evaluation of the algorithm. One of these was the classification of an LD branch as an LCX artery, specifically in cases where the left circumflex artery was absent. This source of error may be diminished by using stronger positional features. The path followed by the vessel with respect to the heart chambers, is an example. Such a feature will clearly distinguish between these two vessel types in the absence of an LCX artery.

In chapters 5 and 6, the voxel intensity was the only feature used to classify the data in the post-processing step. These features seemed to work well, and were also easily available from the raw data. It would be worthwhile investigating texture-based features in multi-dimensional analysis. In chapter 3, the direction and density features express useful information about the coronary vessel centrelines, but the classification procedure used in this chapter did not exploit both features simultaneously. Further work would be to combine these in a suitable framework, for example, using a multivariate Gaussian approach.

8.3 Future Prospects

Each chapter of this thesis presents a list of further work that can be carried out to improve upon the piece of work concerned, however each project presents atlas-based methods as applied to a specific segmentation problem.

All the work carried out in this thesis was performed on CT data, due to product requirements at the time. A more complete solution will be geared towards a modality-agnostic atlas-based segmentation product. Modality-agnostic implying that any created atlas can work in any modality, and more desirable is an atlas image expunged of any modality-specific information.

Appropriate productisation of the entire atlas-based package is desirable. The culmination of the years of research in atlas-based segmentation hints at a generic framework for atlas-based segmentation, as indicated by its application to segmentation of multiple sub-anatomical structures within the heart. Being able to build a single or multi-atlas of any set of anatomical structures and use the created atlas to carry out segmentations in novel images without programmatic intervention will be a valuable asset to clinicians and particularly clinical researchers.

Appendix A

Burden Of Disease in Nigeria

Death Rate Rank	Disease Group and disease name	Death Rate (per 100,000)
1	Infectious and parasitic diseases	883.62
	HIV/AIDS	257.56
	Malaria	180.96
	Measles	91.52
	Tuberculosis	62.11
2	Respiratory infections	181.76
	Lower respiratory infections	181.26
3	Cardiovascular diseases	166.66
	Cerebrovascular disease	57.84
	Ischaemic heart disease	53.57
4	Unintentional injuries	80.60
5	Perinatal conditions (h)	73.94
	Low birth weight	31.88
6	Malignant neoplasms	65.10
	Prostate cancer	11.36
	Breast cancer	8.78
7	Respiratory diseases	42.46
	Chronic obstructive pulmonary disease	17.79
8	Maternal conditions	30.68
9	Intentional injuries	30.64
10	Digestive diseases	25.90
	Cirrhosis of the liver	9.04
11	Genitourinary diseases	19.66
	Nephritis and nephrosis	18.61
12	Neuropsychiatric conditions	15.76
13	Nutritional deficiencies	15.11
14	Diabetes mellitus	11.75
15	Congenital anomalies	7.91
16	Endocrine disorders	4.16
17	Musculoskeletal diseases	1.98
18	Other neoplasms	1.43
19	Skin diseases	0.02
20	Oral conditions	0.00
21	Sense organ diseases	0.00

Figure A.1: Death-Rate Ranking in Nigeria

DALY Ranking	Disease Group	DALY Rate (per 100,000)
1	Infectious and parasitic diseases	29672.88
	HIV/AIDS	7614.27
	Malaria	6505.25
	Diarrhoeal diseases	3623.95
	Measles	3206.64
	Tuberculosis	1626.00
2	Respiratory infections	6033.58
	Lower respiratory infections	5980.13
3	Unintentional injuries	3631.28
4	Neuropsychiatric conditions	2858.23
	Unipolar depressive disorders	635.15
5	Perinatal conditions (h)	2830.39
	Birth asphyxia and birth trauma	1284.43
6	Cardiovascular diseases	1750.34
	Cerebrovascular disease	589.50
7	Maternal conditions	1589.15
8	Sense organ diseases	1349.44
	Cataracts	782.10
9	Intentional injuries	1183.79
10	Nutritional deficiencies	1124.78
	Protein-energy malnutrition	799.77
11	Respiratory diseases	921.01
12	Digestive diseases	828.45
	Cirrhosis of the liver	160.00
13	Malignant neoplasms	747.99
	Breast cancer	109.67
	Liver cancer	72.85
14	Congenital anomalies	521.29
15	Musculoskeletal diseases	354.91
16	Genitourinary diseases	352.70
17	Endocrine disorders	207.93
18	Diabetes mellitus	180.80
19	Skin diseases	93.34
20	Oral conditions	86.67
21	Other neoplasms	26.37

Figure A.2: DALY-Rate Ranking in Nigeria

	Infectious and parasitic diseases	**Respiratory infections**	**Maternal conditions**
Ultrasound		Lung ultrasound for early detection of pneumonia[1], monitoring lung disease[2]	Obstetric ultrasound for pregnancy monitoring and fetus imaging
MRI		Pulmonary imaging for lung diseases[3]	Obstetric MRI[4] provides high resolution for pelvic and foetal imaging[5].
CT	Chest CT less popular than x-rays for tuberculosis.	Pulmonary CT for diagnosis of pneumonia[6].	Not popular due to ionizing radiation.
PET			
X-Ray	Chest x-ray for tuberculosis	Chest x-ray for pneumonia	

Table A.1: Medical Imaging Modalities and the Burden of Disease in Nigeria

References

[1] D. Lichtenstein and O. Peyrouset, "Is lung ultrasound superior to CT? The example of a CT occult necrotizing pneumonia," *Intensive Care Medicine*, vol. 32, pp. 334–335, 2006.

[2] D. A. Lichtenstein, "Lung ultrasound in the intensive care unit," *Recent research developments in respiratory and critical care medicines*, vol. 1, pp. 83 – 93, 2001.

[3] N. L. Muller, "Value of mr imaging in the evaluation of chronic infiltrative lung diseases: Comparison with ct," *American Journal of Radiology*, pp. 1205–1209, 1992.

[4] D. Levine, "Obstetric mr imaging," *Radiology*, vol. 211, pp. 607–17, 1999.

[5] S. McCarthy, "Magnetic resonance imaging in obstetrics and gynaecology," *Magnetic Resonance Imaging*, vol. 4, pp. 59–66, 1986.

[6] H. Syrjala, "High resolution computed tomography for the diagnosis of community-acquired pneumonia," *Clinical Infectious Diseases*, vol. 27, pp. 358–63, 1998.

[7] M. R. Williamson, *Essentials of Ultrasound*. Saunders Company, 1996.

	Neuropsychiatric conditions	**Cardiovascular diseases**	**Unintentional Injuries**
Ultrasound	Detection of germinal matrix haemorrhages in neonatal brains[7]	Echocardiography for diagnosis and prognosis [8], myocardial contrast echocardiography, duplex ultrasound for vascular imaging[7]	Trauma ultrasound for detecting haemorrhages in the abdomen, pericardium and pelvis[9].
MRI	Brain imaging studies are now increasingly carried out with functional MRI[10]. Spinal cord evaluation.	MRI angiography to assess blood flow in coronary arteries. Assess progress of heart disease, and measure heart function. Head/Neck MRI to analyse carotid arteries.	MRI is excellent for soft-tissue imaging, and is used extensively in musculoskeletal injuries
CT	Brain imaging studies are less sensitive in CT than MR [10].	Brain CT for early ischemic strokes; CT angiography for stenosis assessment in coronaries. Chest CT to evaluate heart function. Head and neck CT[10].	Multi-detector CT popular for spine imaging and musculoskeletal imaging[11]. CT for detecting brain haemorrhage in head trauma[12].
PET	Functional brain imaging carried out with PET and PET/MR to detect brain abnormalities.	Functional imaging of the heart with structural info provided by CT.	
X-Ray			Detecting fractures and muscle injuries.

Table A.2: Medical Imaging Modalities and the Burden of Disease in Nigeria continued

	Malignant neoplasms	**Perinatal Conditions**
Ultrasound	Breast cancer detection/monitoring, low frequency US for cysts and tumours	Obstetric ultrasound. Detection of germinal matrix haemorrhages in neonatal brains [7]
MRI	MRI combined with PET for oncology imaging. MRI angiography to visualise tumours. MRI becoming increasingly popular for mammography[13].	Obstetric MRI[4] provides high resolution for foetal imaging.
CT	CT combined with PET to provide structural information in oncology imaging. CT mammography for detecting smaller tumours[14]. Visualisation of lung, liver and lymph node lesions[15].	Not popular due to ionizing radiation.
PET	The most popular modality for oncology imaging[16].	
X-Ray	Screening and diagnostic breast mammography	

Table A.3: Medical Imaging Modalities and the Burden of Disease in Nigeria

[8] W. F. Armstrong, "Echocardiography and coronary artery disease: Current and future applications," 1987.

[9] A. Ng, "Focused assessment with sonograph for trauma (fast)." `http://www.trauma.org/archive/radiology/FASTintro.html` (accessed: 06/04/09), April 2009.

[10] College of Medicine: University of Illinois, "Basic concepts in diagnostic imaging: Modalities in radiology." `https://www-s.med.uiuc.edu/m2/icm/BasicconceptsindiagnosATversion1.pdf` (accessed: 06/04/09), April 2009.

[11] K. Ohash, "Musculoskeletal CT: Recent Advances and Current Clinical Applications," *Radiologic choices of North America*, vol. 47, pp. 387–409, 2007.

[12] E. A. Elgamal, "Evaluation of the Canadian CT Head Rule for Minor Head Trauma in a Tertiary Referral Institution," *European Journal of Trauma*.

[13] C. Lehman, "Role of MRI in screening women at high risk for breast cancer," *Journal of Magnetic Resonance Imaging*, vol. 24, pp. 964–70, 2006.

[14] C. H. Chang, "Computed Tomographic Mammography using a conventional body scanner," *American Journal of Roentgenology*, vol. 138, pp. 553–558, 1982.

[15] Siemens, "Syngo ct oncology." `https://www.medical.siemens.com/`(accessed 04/04/2009).

[16] J. Ruhlmann, P. Oehr, and H. Biersack, *PET in Oncology: Basics and Clinical Applications*. Springer-Verlag, 1999.