AUTOMATED PULMONARY ARTERY SEGMENTATION BY VESSEL TRACKING IN LOW-DOSE COMPUTED TOMOGRAPHY IMAGES

A Thesis

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Master of Science

by

Jeremiah Anthony Wala

January 2011

ABSTRACT

Low-dose computed tomography (CT) imaging provides a method for obtaining accurate anatomical information without the full radiation exposure inherent in standard CT protocols, and is primarily used in lung cancer screening. Segmentation of the pulmonary arteries from low-dose chest CT images is a vital first step in improving computer-aided detection of frequently missed pulmonary nodules near major arteries. This thesis presents the first fully automated method for segmenting the main pulmonary arterial trees in low-dose CT images.

The correlation between the arterial and airway trees was used to develop an automated pulmonary artery seed point detector. The main basal pulmonary arteries are identified by searching for candidate vessels near known airways, using a progressive morphological opening method. The arteries are tracked into the lungs by means of a cylindrical vessel tracker that iteratively fits model cylinders to the CT image. Vessel bifurcations are detected by measuring the rate of change of vessel radii. Subsequent vessels are segmented by initiating new cylinder trackers at bifurcation points.

Quantitative analysis of both the number of arteries and veins segmented, as well as the error per vessel, was accomplished with a novel evaluation metric called the Sparse Surface (SS) metric. The SS metric was developed to capture the details of the true vessel surface while reducing the ground-truth marking burden on the human user. This metric is a unique new tool for ground truth marking and segmentation validation, with particular importance in problems with complex geometries.

The segmentation method and SS metric were applied to a dataset of seven

CT images, and achieved an overall sensitivity of 0.62 and specificity of 0.90 of all manually identified vessels. The average root mean square error between the vessel surface and the segmentation surface was 0.63 mm, or less than 1 voxel. Additionally, seed points were detected automatically for a majority (80%) of cases with labeled airways. This method is an important first step towards robust pulmonary artery segmentation and artery/vein separation in low-dose chest CT, and is the first fully automated method designed for accomplishing this task.

BIOGRAPHICAL SKETCH

Jeremiah Wala lived in Waseca, MN for 18 years before moving to Ithaca, NY in 2005 to attend Cornell University. In May 2009 he received a Bachelor of Science *magna cum laude* in Engineering Physics. He then continued at Cornell, pursuing a Master of Science degree in Applied Physics. While at Cornell, he has worked as a teaching assistant in the fields of nanotechnology, quantum and statistical physics, mathematics, and organic chemistry, and as a researcher in nanotechnology, radiobiology and medical imaging. In 2010, he began his medical studies at Harvard Medical School and the Massachusetts Institute of Technology, and is pursuing an M.D. degree.

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my advisor, Prof. Anthony Reeves, for his assistance, guidance, and vast patience as I navigated with much trial and error the research and thesis writing process. Without his support, this work would not be possible. I also want to thank my fellow researchers in the VIA group, Artit Jirapatnakul, Alberto Biancardi, Sergei Fotin, Brad Keller and Jaesung Lee, for their endless knowledge and willingness to help me sort out the bugs.

I would also like to thank my committee chair Prof. Joel Brock for his encouragement as a teaching advisor and committee member. I am very grateful to both him and the Cornell Applied Physics department for their remarkable support of my goals.

Finally, I would like to thank my parents, sister, and friends for their reassurance and encouragement throughout my time in college, graduate school, and life in general.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

The pulmonary arteries comprise the vessels that supply deoxygenated blood to the capillary beds of the right and left lungs. Each lung is supplied with a single branch of the main pulmonary artery originating at the right ventricle of the heart. A method for automatically locating these arteries and segmenting their subsequent branches is presented in this thesis. The segmentation method consists of two parts: automated seed point detection in the left and right pulmonary arteries and vessel tracking using a cylinder fitting and novel bifurcation detection method.

To quantify the performance of the segmentation algorithm, a novel evaluation metric, the Sparse Surface (SS) metric, was developed for evaluating segmentations of complex structures. This method is designed to quantify both the total number of arteries and veins segmented and the error per vessel. The sparse ground truth marking system required to implement this method constitutes an improvement over dense marking systems designed for the same purpose, significantly reducing the marking burden to the human user.

The dataset used to test the segmentation method consists of low-dose chest CT images, rather than the standard-dose CT images. The primary purpose of thoracic low-dose CT is for pulmonary nodule detection in lung cancer screening, and segmentation of the pulmonary arteries could improve computer-aided pulmonary nodule detection. However, low-dose CT images present an additional challenge to segmentation methods, as there is less information and more noise than in standard-dose CT images.

1.1 Clinical Application

Lung cancer is the leading cause of cancer death in the United States, and an estimated 1.5 million new cases are expected to be diagnosed worldwide in 2010 [9]. Early detection of lung cancer has been shown to reduce lung cancer mortality [7, 4], prompting the creation of large clinical trials of lung cancer screening using low-dose CT imaging [5]. Small malignant pulmonary nodules can be detected in low-dose CT images before they become large enough to cause symptoms, allowing for earlier treatment and improved clinical outcomes [7].

Identification of pulmonary nodules in low-dose CT is a challenging task. Computer-aided detection of lung nodules has been shown to improve the detection sensitivity of radiologists [25]. However, nodules located at vessel bifurcations, lung walls, and vessel surfaces are common sources of false negatives in automated computer detection systems. Automated segmentation of the pulmonary vessels has been shown to improve lung nodule detection sensitivity when applied to low-dose screening scans [3]. Our method is designed to segment the large arteries as they enter the lungs, a first step in improving nodule detection near the mediastinum.

After nodules are detected, they must be classified as either benign or malignant in order to make informed clinical decisions. Distinguishing between benign and malignant nodules remains a difficult task for radiologists. Computer-aided nodule classification systems have been developed to aid in this task, and rely on features such as volume, central moments, and compactness for accurate classification [19]. Nodules attached to the lung pleura and the pulmonary vessels are particularly challenging to segment and classify, as their shape and volume are difficult to measure. By defining the true blood vessel shape, segmentation of the pulmonary vasculature in low-dose CT images could provide a method of more accurately measuring and classifying nodules adjacent to vessels.

Segmentation of the pulmonary arteries also provides a first step toward obtaining highly accurate measurements of vessel radii. Patients with pulmonary arterial hypertension (PAH) have been shown to present with enlarged pulmonary arteries, as measured by manual observation of CT images [22, 17]. These enlarged vessels are at increased risk for fatal dissection and require immediate surgical intervention where possible. Automated segmentation and measurement methods have the potential to better quantify vessel enlargement in patients with PAH, providing an earlier and non-invasive means of diagnosis [14].

1.2 Computed tomography

The first clinical scans from a CT, published in 1972 by Hounsfield and Ambrose, consisted of head scans of human brains [6]. Since the inception of the single-slice scanner, CT technology has progressed to include helical CT and multi-detector row CT (MDRCT). However, because CT present a potential radiation hazard to the patient, low-dose scanning protocols are now being implemented to reduce the risk to the patient. This section will provide a brief overview of the X-ray physics and CT technology relevant to low-dose imaging.

1.2.1 X-ray physics

CT images are produced by measuring the interaction of electromagnetic radiation with matter. The intensity of radiation reaching the photodetector is dependent on the photodensity of the irradiated matter. Tissues are distinguished from one another by the rate they attenuate radiation I, given by Beer's Law:

$$I(\eta) = I_0 e^{-\mu\eta} \tag{1.1}$$

where η is the thickness of the irradiated matter and μ is the linear attenuation coefficient, which is dependent on the rate of which the material absorbs and scatters the radiation. The linear attenuation coefficient is related to the density of the material by:

$$\mu = \frac{\rho N_A}{A} \cdot \sigma_{\text{tot}} \tag{1.2}$$

where ρ is the density, N_A is Avagadro's constant, A is the atomic number, and σ_{tot} is the total photoatomic cross-section for either scattering or absorption. Thus, the attenuation coefficient is simply the number of atoms per unit volume times the interaction probability per atom.

CT scanners require high-frequency X-ray electromagnetic radiation with wavelength between 10^{-11} and 10^{-8} m (120 eV to 120 KeV) to produce high-quality images. In this regime, the photoatomic cross-section of the irradiated matter is dependent upon two major interactions: absorption via the photoelectric effect and Compton scattering.

In a photoelectric interaction, shown in Figure 1.1a, the incident radiation will eject an inner shell electron from the atom and deposit all of its energy $h\nu - \phi$ to the electron, where h is Planck's constant[1], ϕ is the energy required to free the electron, and ν is the frequency. The cross-section for the photoelectric interaction is given by Sauter:

$$\Phi_{\text{photoelectric}} = K(\nu)Z^5 \tag{1.3}$$

where Z is the atomic number and K is weighting function[2] which is dependent on the photon energy [21]. The quintic dependence of the cross-section on the atomic number provides the motivation for using contrast agents with high atomic numbers, such as barium (Z = 56) or iodine (Z= 53), which interact strongly with

[1] $h = 6.626 \times 10^{-34}$ J · s

[2] $K(\nu) = \frac{3}{2}\alpha^4 \frac{(\gamma+1)^{3/2}}{(\gamma-1)^{7/2}} \left[\frac{4}{3} + \frac{\gamma(\gamma-2)}{\gamma+1}\left(1 - \frac{1}{2\gamma\sqrt{\gamma^2-1}} \ln \frac{\gamma+\sqrt{\gamma^2-1}}{\gamma-\sqrt{\gamma^2-1}}\right)\right] \exp\left(-\pi\alpha + 2\alpha^2\left(1 - \ln \alpha\right)\right)$
where $\gamma = \frac{h\nu + m_e c^2}{m_e c^2}$ and $\alpha \approx \frac{1}{137}$ is the fine-structure constant

the radiation. However, materials with similar molecular composition (e.g. soft tissues) are difficult to distinguish in CT.

Compton scattering occurs when an incident photon scatters off an outer valence electron. Figure 1.1b shows an incident X-ray γ recoiling off a weakly bound valence electron and scattering as a lower energy photon γ'. The cross-section for Compton scattering was explored by Klein and Nishina, and is given by:

$$\sigma_{\text{Compton}} = \Phi_0 J(\nu) \tag{1.4}$$

where Φ_0 is the classical Thompson scattering constant[3], and $J(\nu)$ is a function[4] of the frequency ν [21]. Unlike the photoelectric interaction, the Compton interaction does not depend on the atomic number. The probability of a Compton interaction is dependent on the number of electrons in the material. For this reason, Compton scattering provides for contrast between tissues with different electron (mass) densities.

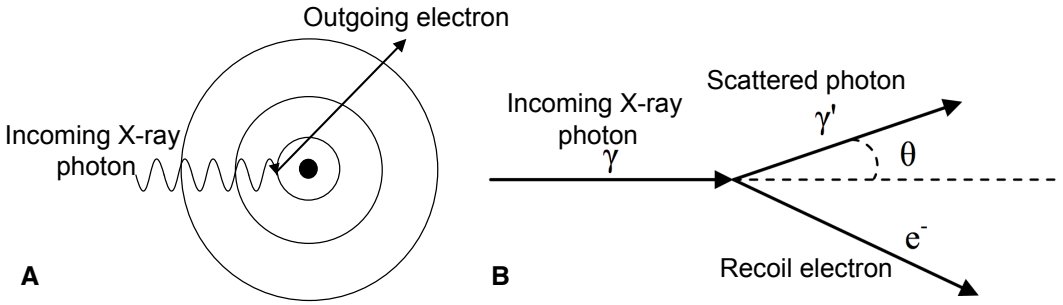

Figure 1.1: A) Photelectric interaction and B) Compton scattering of an incident X-ray. These two processes account for the attenuation of X-rays in matter, providing information to the CT scanner about the radiodensity of the irradiated tissue.

CT scans are calibrated so that materials with a given radiodensity will produce the same value across different scans. The Hounsfield scale was developed to

[3] $\Phi_0 = \frac{8\pi}{3}\left(\frac{1}{4\pi\epsilon_0}\frac{e^2}{m_e c^2}\right)^2 = 6.65 \times 10^{-29}\ \text{m}^2$

[4] $J(\nu) = 2\pi\left(\frac{1}{4\pi\epsilon_0}\frac{e^2}{m_e c^2}\right)^2\left[\left(\frac{1+\gamma}{\gamma^2}\right)\left(2\frac{1+\gamma}{1+2\gamma^2}\right) - \frac{\ln(1+2\gamma)}{\gamma} + \frac{\ln(1+2\gamma)}{2\gamma} + \frac{1+3\gamma}{(1+2\epsilon)^2}\right]$ where $\gamma = \frac{h\nu}{(m_e c^2)^2}$

describe a tissue's radiodensity relative to water and air. A value of 0 Hounsfield units (HU) is assigned to water, and -1000 HU is assigned to air. Given these two defined points, the HU value of any tissue can be found with the linear relation:

$$H_{\text{tissue}} = \frac{\rho_{\text{tissue}} - \rho_{\text{water}}}{\rho_{\text{water}}} \times 1000 \tag{1.5}$$

where ρ is the radiodensity measured by the CT scanner through the interaction processes described above. Because CT is calibrated, the soft tissue (e.g. blood) presents at around 0 HU across different scans. This provides a method to separate the class of soft tissue objects from bones ($>$ 200 HU) and lung parenchyma (-1000 HU) based on the calibrated intensity values.

1.2.2 CT scanner technology

In conventional X-ray radiography, a single immobile source irradiates a 3D object to produce a 2D projected image. CT employs a mobile X-ray source that rotates around a circular gantry to produce a series of projected images called a sinogram, which is converted via the Radon transform into the final image. Single-detector CT produces one slice per 180$^{\text{o}}$ gantry rotation. To produce multiple slice images, the patient is moved forward on the scanning table a distance equal to the slice thickness, and a new slice is acquired. Helical CT employs a continuously rotating X-ray source and constant table velocity to reduce acquisition time and minimize motion artifacts. Multi-slice CT scanners contain multiple detectors that allow for multiple slices to be acquired per gantry rotation.

CT scanners produce X-rays by accelerating electrons towards a target that interacts with the electrons to produce X-rays. Schematics of two common CT X-ray generation systems are shown in Figure 1.2 [2].

In both systems, thermal electrons created at the cathode by heating a metal filament with a current I_0 are accelerated towards the anode through a potential

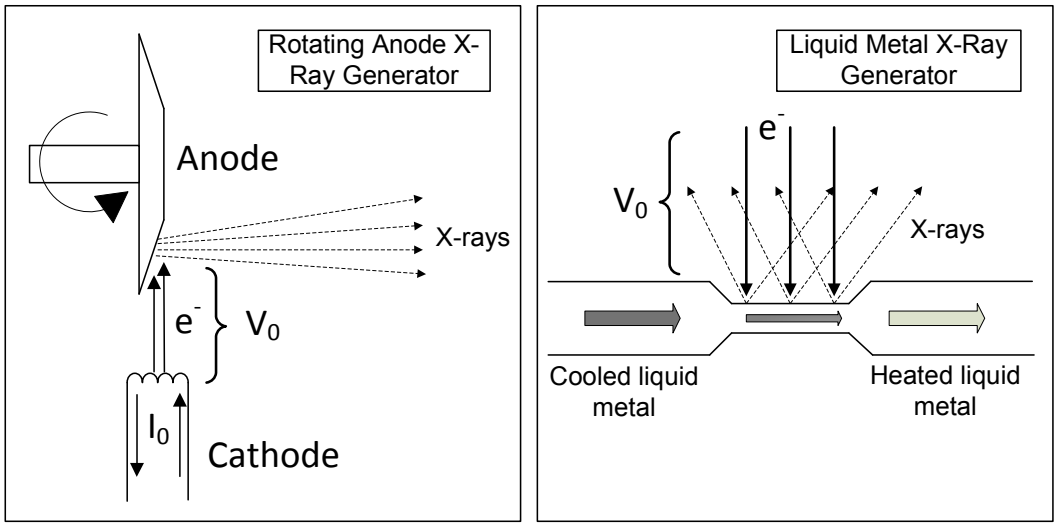

Figure 1.2: Two common systems for X-ray generation in CT scanners.

V_0, typically in the 25 kV - 150 kV range. Electrons interact with the metal anode material to produce X-rays through *bremsstrahlung* and inner shell electron ejection. Although X-rays are produced in both processes, the primary product from the electron-anode interaction is heat. Circulating liquid metal anodes were developed to counteract the problem of overheating of the anode material.

Low-dose CT scans are produced by reducing the anode current without increasing the scan time. This is achieved by decreasing the number of thermal electrons produced at the cathode. It is useful to explore how the reduction in the dose affects the quality of the signal received and thus the quality of the image. The photon statistics at the photodector can be approximated by a discrete Poisson distribution [2]. The variance is equal to the mean of the distribution, which is proportional to the total number of events n, in this case is the number of photons received at the photodetector. The signal-to-noise ratio (SNR) is:

$$\text{SNR} = \frac{\text{signal level}}{\text{noise level}} = \frac{\text{mean}}{\text{standard deviation}} = \frac{\langle n \rangle}{\sqrt{\langle n \rangle}} = \sqrt{\langle n \rangle} \tag{1.6}$$

Thus we see that the SNR is proportional to the square root of the anode current;

image quality is reduced only slowly as the radiation dose is decreased.

Naidich *et al* first explored the use of low-dose CT scanning protocols for thoracic imaging [18]. With a reduction in anode current from 120 mA to 10mA, they found that although noise was noticeably increased, organ structures within the images were still able to be robustly identified by trained radiologists. Low-dose scans provide enough detail for detecting nodules in lung cancer screening scans, where the decrease in the SNR is not significant enough to justify increasing patient exposure. The purpose of this thesis is to extend the amount of information that can be obtained in low-dose CT by automatically segmenting the pulmonary arteries.

1.3 Pulmonary artery anatomy

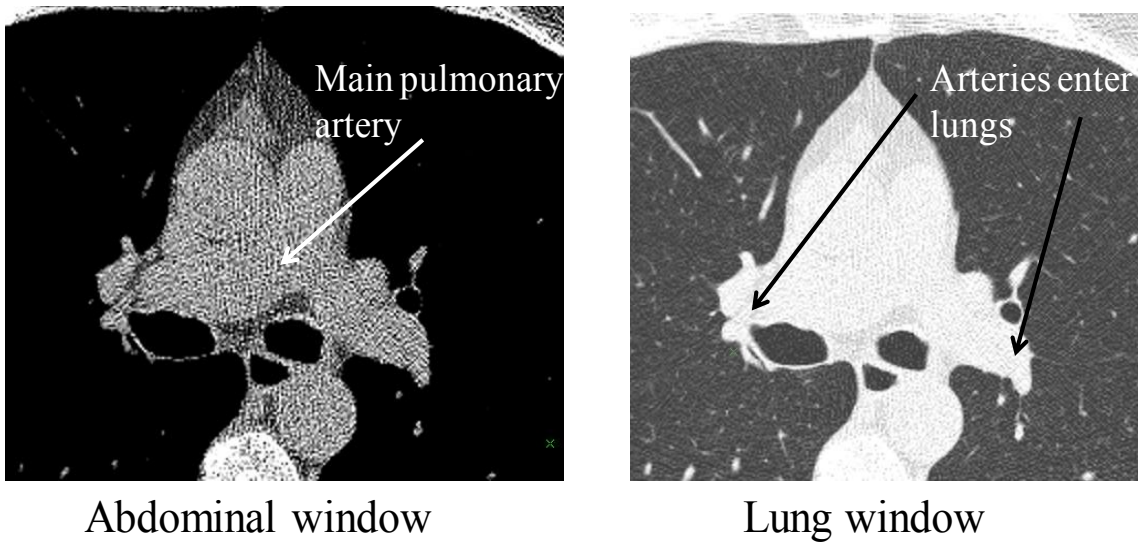

Figure 1.3: The location of the pulmonary arteries in the mediastinum and the point where they emerge into the lung parenchyma.

The pulmonary arteries begin at the right ventricle of the heart, and are part of a larger organization of organs (heart, trachea, pericardium, etc.) known as the

mediastinum, located in the center of the thoracic cavity (Figure 1.3a). Tracking the pulmonary arteries from their origin in the center of the mediastinum presents a significant challenge to automated segmentations, due to the lack of contrast between the different soft-tissue types in that region. Instead, our goal is to begin tracking of the arteries as soon as they emerge from the mediastinum into the lung parenchyma, as indicated in Figure 1.3b.

The anatomy of the airway and pulmonary arterial trees is highly complex: both the airway and arterial trees branch extensively in the lungs, with approximately seven generations appearing in the CT scan. Before attempting to develop an automated segmentation method, it is important to have an accurate picture of the pulmonary anatomy, along with its proper nomenclature. This section describes the geometry and nomenclature of the pulmonary anatomy.

Despite the individual complexity of the airway and arterial trees, a remarkable symmetry exists between the airways and the arteries. The two trees form a parallel structure called the *bronchopulmonary* tree, in which the pulmonary arteries run parallel and adjacent to the airways, forming airway/artery vessel pairs [8]. Furthermore, the pulmonary arteries bifurcated at approximately the same locations as the airways. These observations led to a parallel nomenclature for both the airways and the arteries.

The lungs are divided into five lobes, each with its own blood supply, capillary bed, and venous return. The left lung is comprised of a superior, medial, and inferior (or basal) lobe, while the right lung is divided only into a superior and inferior (basal) lobe. Each pulmonary lobe can be subdivided further into segments, where each segment is supplied by a unique bronchopulmonary vessel pair. The nomenclature for the segmental arteries and airways simply describes which segment they supply and where that segment is located within the lobe. For

example, the basal-posterior artery and airway supply the posterior segment of the basal lobe. Each segmental bronchopulmonary segment can be further divided into subsegmental, and sub-subsegmental vessel pairs.

Figure 1.4a shows an airway tree labelled with the bronchopulmonary segment nomenclature. The airway tree was generated with an automated airway segmentation by Lee and Reeves [12], along with minor manual modifications to account for segments that were visually identified in the scan but were not automatically generated. The nomenclature for the airway tree will be used to describe the pulmonary arteries running parallel and proximal to the airways. Figure 1.4b shows the symmetric nature of the airway and arterial trees. The airways were generated with the automated method by Lee and the pulmonary arteries were manually marked.

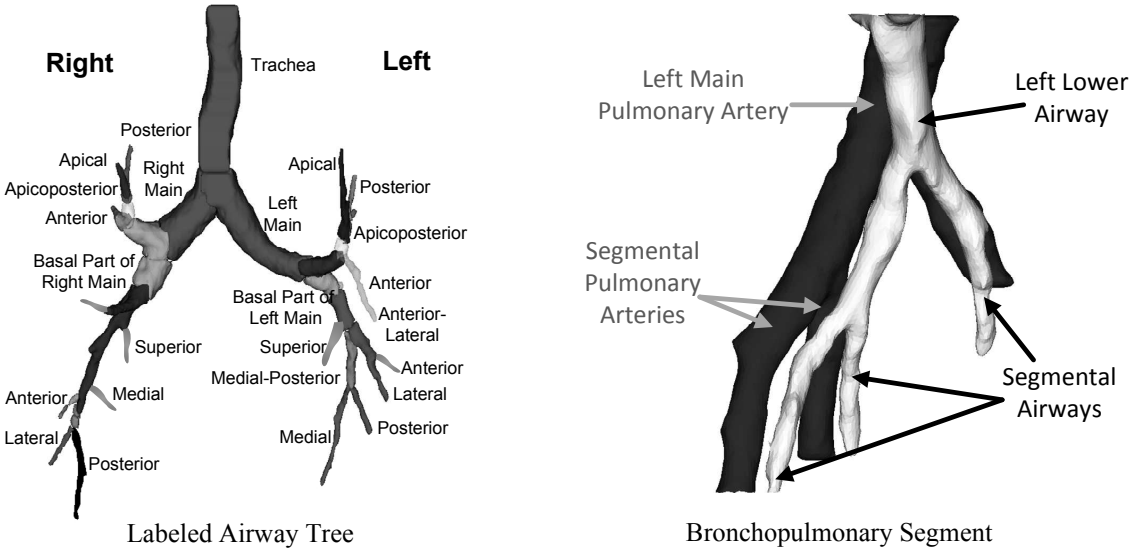

Figure 1.4: Labelled airway tree demonstrating the bronchopulmonary nomenclature used to describe the pulmonary arteries.

1.4 Issues in pulmonary artery segmentation

The pulmonary vasculature is a highly complex branching structure that presents significant challenges towards robust segmentation. Normal anatomical variability, as well as anatomic abnormality from disease, limits the amount of *a priori* information that can be used to develop a robust segmentation algorithm. Streaks from metal implants and image discontinuities from patient movement also decrease image quality and generate potential sources of error for automated methods. And as with any segmentation method in low-dose CT, image anisotropy, noise, and partial volume effects serve to complicate the task of identifying structures. This section will describe the main challenges in segmenting the pulmonary arterial tree in low-dose CT.

The main difficulty in segmenting the arterial tree in low-dose CT is the lack of contrast between the different types of dense (non-fat) soft tissue. Without physical contrast enhancement, soft tissue structures such as the blood, vascular walls, airway walls, and the heart have indistinguishable image intensity values. The current clinical protocol for CT pulmonary angiograms requires injection of 100 - 150 mL of radiocontrast [1]. This contrast allows the blood to be distinguished from airway walls and lymphatic tissue based solely on intensity information. The low-dose lung cancer screening scans used in this thesis do not have the added aid of radiocontrast.

The most frequent source of error arises from confusion with the venous vascular tree. Figure 1.5 presents a CT section of the left lower lobe containing segmental pulmonary arteries, segmental venous branches of the left inferior vein, and airway walls. Because the arteries and veins have the same intensity, separate structures appear to run together to form a complex blob. A geometrical representation is included to further indicate which tree (arterial or venous) each vessel belongs to.

The vessel were identified by manually tracing them back to their parent vessel, either the left inferior pulmonary vein or the left pulmonary artery. The boundaries between the veins, arteries and airway walls are difficult to determine, due to the lack of contrast between structures.

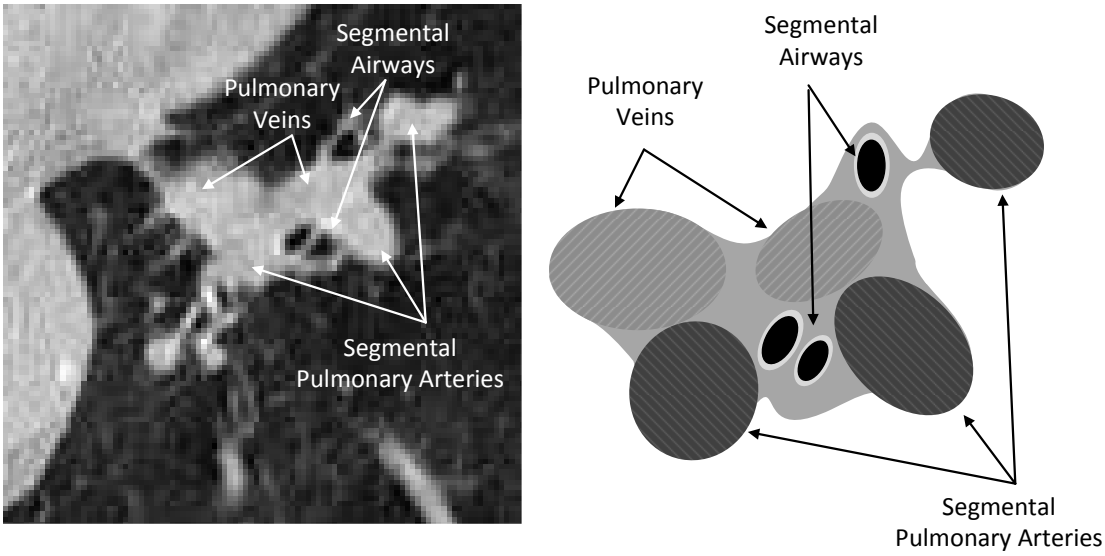

Figure 1.5: Segmentation issue caused by close proximity of arterial and venous trees

In addition, non vascular soft-tissue structures such as the airway walls and lymphatic tissue are another important source of segmentation error. Figure 1.6 shows a section of the left inferior lung windowed to soft tissue, along with a geometrical representation of the anatomy. The left medial segmental artery, the posterior branch of the left inferior vein, and the medial segmental airway wall appear as a fused isointensity object. Errors can occur in the segmentation where the airway walls bridge the venous and arterial trees, and where arteries and veins touch each other.

The morphology of the arterial tree itself further complicates the task of segmenting the vessels. Bifurcations are an especially salient source of difficulty because the conflated arterial and venous trees often present a similar geometry to

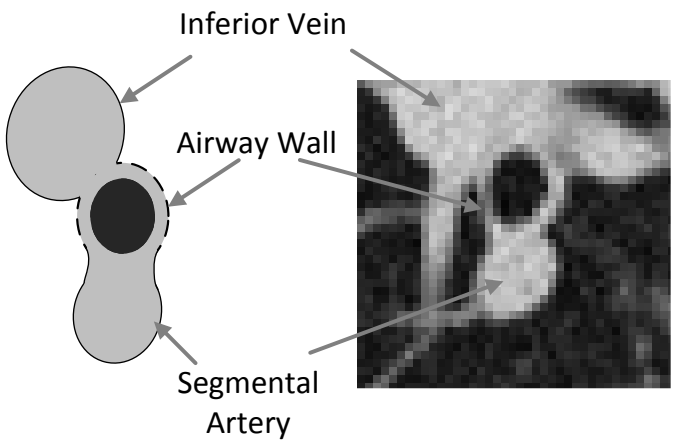

Figure 1.6: Segmentation issue caused by proximity of airway wall to arterial tree

true vessel bifurcations. Any segmentation algorithm for the pulmonary arterial tree will have two challenges at these bifurcation points: 1) finding the two child vessels and 2) avoiding segmenting extraneous structures. Figure 1.7 shows a bifurcation of the left pulmonary artery into the posterior basal and anterior basal segmental arteries. An inferior segmental vein is running posterior to the bifurcation point, and appears to be stemming from the parent artery when only the local geometry around the bifurcation point is considered.

1.5 Previous work

Although there have been several published methods for segmenting the entire pulmonary vasculature (arteries and veins), only limited work has been done on distinguishing the pulmonary arteries from the pulmonary veins in CT. Methods to segment the entire pulmonary vasculature generally rely on a combination of global image filters to identify candidate structures, followed by local operations such as vessel tracking, region growing, and morphological operations to achieve a final segmentation. This section first provides an overview of previously published

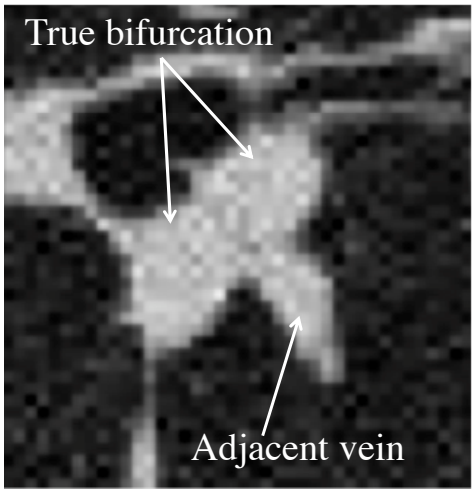

Figure 1.7: Segmentation issue caused by vessel bifurcation. In this case, an adjacent vein has the appearance of stemming from a trifurcation of the parent artery.

methods for automated pulmonary vessel segmentation, and then reviews methods specific to separating arteries from veins. Table 1.1 provides a summary of the primary methods and goals of previously published work on pulmonary vessel segmentation.

1.5.1 General pulmonary vessel segmentation methods

Several attempts to segment the pulmonary vasculature focus on detecting vessels by filtering the image with a Hessian matrix. The Hessian matrix provides an estimate of the local curvature of an image f, and is given by:

$$H_{ij}(f) = \frac{\partial^2 f}{\partial x_i \partial x_j} \tag{1.7}$$

The eigenvalues of the Hessian matrix are used to detect tubular structures and their orientations, in an effort to enhance cylindrical vessels.

Zhou (Chuan) *et al.* combined a Hessian matrix filter with a multiscale Gaussian filter to enhance vessels and their bifurcations in contrast-enhanced CT pul-

Table 1.1: Summary of previously reported methods for non-discriminantly segmenting the pulmonary vasculature and for discriminating arteries from veins.

Author	Year	Goal	Method	Data
Wood [24]	1995	PV	Region growing	1 SD (canine)
Masutani [15]	2001	PV	Connected component	12 SD w/contrast
C. Zhou [27]	2005	PV	Hessian, expectation max.	2 SD w/contrast
J. Zhou [29]	2007	PV	Hessian, bif. Detection	1 SD w/contrast
Kaftan [11]	2008	PV	Fuzzy connectivity	1 SD w/contrast
Shikata [23]	2009	PV	Hessian, vessel tracking	44 SD
Lei [13]	2001	A/V	Fuzzy connectivity (MRI)	54 (pelvic MRI)
Bemmel [26]	2003	A/V	Level-set algoritm	7 (pelvic MRI)
Saha [20]	2010	A/V	Fuzzy connectivity (CT)	2 SD
This thesis	2010	PA	Airway model, vessel tracking	10 LD

PV*: Pulm. vasculature* **A/V***: Artery/Vein separation* **PA***: Pulm. arteries*
SD*: Standard dose CT* **LD***: Low dose CT*

monary angiography (CTPA) scans [28]. The eigenvalues of the Hessian matrix were incorporated into a response function which returned values characteristic of either bifurcations or vessels, and suppressed most non-vascular structures. To overcome the challenge of detecting a wide range vessel sizes, the eigenvalues of the Hessian matrix were computed for images with different scales of Gaussian filtering; increasing the filter size enhanced larger vessels. To achieve their final segmentation, an expectation-maximization classifier based on the histogram of the lung region was applied to each voxel in a local region around the vessel candidates. Using a centerline metric (see Section 4.1) they achieved true positive rates of 93.8% and 97.0% for cases with and without lung disease.

Shikata and Sonka employed a multiscale Hessian matrix filter to enhance vessel and bifurcation structures, but combined the filter output with a vessel tracking algorithm to preserve vessel connectivity [23]. An initial segmentation was performed by locating likely vessel centerlines using the output of the Hessian matrix filter. The local distance transform was computed at each point to obtain an esti-

mate of the radius, and a threshold of -600 HU was applied within the estimated region to obtain a binary segmentation. Vessel tracking was then initiated at these seed points and proceeded in the direction of the eigenvectors of the Hessian matrix, pointed along the vessel centerlines. The vessel tracking connected centerline candidates and terminated when a bifurcation candidate, generated by the Hessian filter, was reached. The final segmentation was achieved by growing spheres at each point of the tracked centerline. Their validation criterion was similar to Zhou, except that in addition to marking vessel centerlines, false points were included outside the vessel regions. In 44 scans without contrast enhancement, Shikata identified 99% of their manually placed truth-points. However, their method did not attempt to distinguish between arteries and veins, and used high quality standard dose CT images, as compared with the low-dose images used in this thesis.

Zhou (Jinghao) and Axel *et al.* combined a multiscale Hessian matrix filter with a bifurcation detector based on the AdaBoost machine learning algorithm [29]. Similar to Shikata and Sonka, they used the eigenvalues and eigenvectors of the Hessian matrix filter to initiate vessel tracking at automatically generated seed points. To search for bifurcations, they took a cutting plane directly in front of the vessel tracker and perpendicular to the tracking direction. The cutting plane was then filtered separately by a three filters: a Gaussian filter and its first and second derivatives. The trained AdaBoost method then made a decision about whether the cutting planes represented a bifurcation or not. Their test set was limited to 50 user-selected regions of interest in the lung parenchyma of chest CT scans. They reported a 97% true-positive rate for vessel detection.

Although most automated approaches towards segmenting the pulmonary vasculature have focused on identifying vessel candidates with the Hessian matrix, several other approaches have been attempted. Kaftan and Ach applied a soft-

tissue threshold to the lungs to generate candidate seed points [10]. The fuzzy connectivity of neighboring pixels to the seed points was computed and thresholded to obtain the final segmentation with a sensitivity of 89% at a specificity of 98%. Masutani and Doi implemented hysteresis thresholding and connected component analysis to generate seed points for the major pulmonary vessels [15]. These seed points were used to initiate region growing within a segmented lung region. They employed their segmentation to limit the search space for a pulmonary embolism detector, which achieved 85% sensitivity with 2.6 false positives per case over 19 cases [16]. Wood and Hoffman performed manually seeded region growing in both the pulmonary vessels and the lung airways. The binary segmentation was thinned to search for bifurcations [24].

1.5.2 Pulmonary artery segmentation methods

Separating arteries from veins is a challenging task. Lei and Udupa used a multiseeded fuzzy connectivity algorithm to separate a conflated arterial and venous tree in contrast-enhanced magnetic resonance angiography [13]. After an initial fuzzy segmentation of the entire vasculature, the connectedness of each point in the segmentation to either the arterial or venous seed points was computed. The separation was only validated visually.

Bemmel *et al.* implemented a level-set algorithm to separate arteries and veins in magnetic-resonance images [26]. A Hessian matrix filter was used to detect the centerlines of vessel candidates. A function representing the boundary of the vessel was fit around each centerline using the level-set algorithm to obtain the final segmentation. Using a volume-based similarity metric with manually created ground-truth regions, they achieved a similarity of 0.94.

Although these previous methods were implemented in MRI scans, they are

potentially generalizable to other imaging modalities such as CT. However, Saha and Hoffman presented an approach specifically in pulmonary CT images [20]. Their method is based on multiseeded fuzzy connectedness, employed by Lei and Udupa for the same task in magnetic resonance angiography (MRA) images. The pulmonary vessels are manually seeded as either arteries or veins at their centerlines to obtain a fuzzy connectivity map. The map is analyzed to obtain an approximation for the size of the vessels, and a local morphological erosion of this size is applied to separate the fused vessels. The separated, eroded vessels are dilated to their original size, but each voxel in the dilated region is assigned to either the vein or artery based on their proximity to the newly identified vessel center. Validation was performed on two CT scans without contrast enhancement. Several thousand points on the centerlines of the arteries and veins were manually marked to establish the sensitivity and false positive rate of the pulmonary artery segmentation. The quality of the segmentation increased with the number of seed points, and a sensitivity of 99% and false positive rate of 0.5% were reported with 66 seed points.

The methods aimed at segmenting only the arterial trees require anatomical knowledge to properly seed the segmentation algorithms. This is markedly different from methods to segment the entire pulmonary vasculature, which require only morphometry and intensity clues inherent in the vessels themselves. Separating the arteries required a trained user familiar with anatomy to use other clues (parent vessel identity, proximity to known anatomical structure, etc.) to inform the method about how to classify the segmented vessels. Thus, although general pulmonary vessel segmentation methods can provide a significant first step towards artery segmentation, they are fundamentally unable to distinguish and classify conflated isointensity trees.

1.6 Outline

The primary goal of this work is to present an automated method for identifying the parent pulmonary arteries as they enter the lungs, and automatically track them into the lung parenchyma. Chapter 2 describes the method for automatically seeding the pulmonary arteries and the experiment designed to evaluate the method. The results and discussion of the seed point method are presented at the end of the chapter. Chapter 3 describes the method for iteratively tracking the arteries along their lengths and detecting their bifurcations. In Chapter 4, the Sparse Surface Metric is first defined in general terms, and then is applied to the problem of artery segmentation. Furthermore, an overview of previously reported validation metrics is presented. Chapter 5 describes the experiment used to train and test the vessel tracking and bifurcation detection method and presents the results. Finally, a concluding discussion and suggestions for future work are provided in chapter 6.

CHAPTER 2

SEED POINT DETECTION

This chapter describes our work in automated pulmonary artery seed point detection for use with vessel tracking. The first section describes the method for automatically finding the location and orientation of the pulmonary arteries as they enter the lungs. The details of the detection experiment, including the visual training and parameter optimization, are then presented. Finally, the results of the experiment are provided along with a discussions of our findings.

2.1 Method

This section will describe a method of automatically generating seed points and vessel orientations for the basal part of the left and right main pulmonary arteries, using as an axiom the results of a previously reported method of segmenting and labeling the airways by Lee and Reeves [12].

2.1.1 Anatomical model

As described in Section 1.3, the pulmonary arteries run parallel and adjacent to the airways. This anatomical correlation provides strong evidence for the location of the pulmonary arteries, and is especially useful in CT images without physical contrast enhancement where isointensity arteries and veins are only distinguishable by their anatomical locations. In order to utilize this bronchopulmonary model for locating the pulmonary arteries, the airway tree must be robustly segmented. Furthermore, each segmented branch must be uniquely labelled, with each airway receiving the same label across every case. This ensures that the corresponding pulmonary arteries are also correctly identified across different cases.

Lee and Reeves provide a fully-automated method for accomplishing this task

in low-dose chest CT [12]. Airway segmentation is performed by 3D region growing from a seed point in the trachea, generated automatically by searching for a large, low-intensity component 50 mm below the most superior axial frame. The airway region growing is performed within an advancing cylindrical region of interest. In the case of a leak outside the ROI, the homogeneity parameter for growing is narrowed, and a more conservative region growing is performed. After the segmentation is completed, the airway tree is skeletonized and pruned with a 3D thinning algorithm. Each segment of the skeleton is then uniquely labelled according to its relationship to the trachea and the other branches. A diagram of the airway labels, along with the relevant local vascular anatomy, is shown in Figure 2.1.

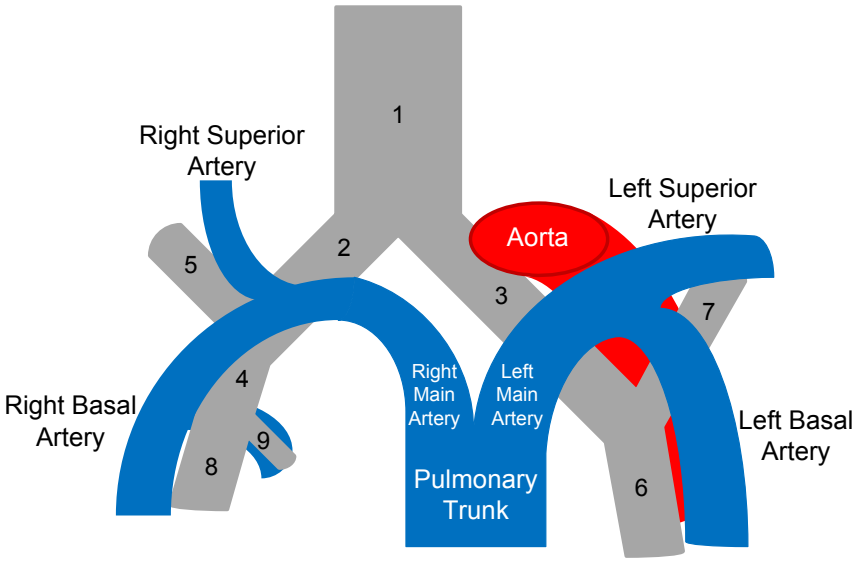

Figure 2.1: Illustration of the relationships of the labelled airways to the pulmonary arteries

The strongest anatomical correlation between the airways and the pulmonary arteries occurs along the left and right basal airway segments, labelled as segments 6 and 8 respectively. With very high probability, the left and right basal arteries will be located parallel and lateral to their corresponding airway segments. Fig-

ure 2.2 shows this relationship between the airways and arteries in both the left and right lungs. The airway segmentation by Lee is superimposed on the image as indicated.

In addition to their strong anatomical correlation with the airways, the basal arteries also exhibit high contrast against the lung parenchyma. Although the basal arteries are located near the soft tissue mediastinum (heart, great vessels, and surrounding tissue), they penetrate deep enough into the lung parenchyma that they can be distinguished from the isointensity great vessels of the mediastinum. Thus, the basal arteries were also chosen as candidates for automated seed-point generation because their location outside the mediastinum is most likely to produce a lower false positive rate; other non-artery vessels are not likely to be confused with the pulmonary arteries. Figure 2.2 shows that the basal arteries are separated from the heart and aorta by lung parenchyma. Notice, however, that in the left lower lobe, the inferior pulmonary vein runs near the basal artery, creating a potential hazard to automated seed-point detection and pulmonary vessel tracking.

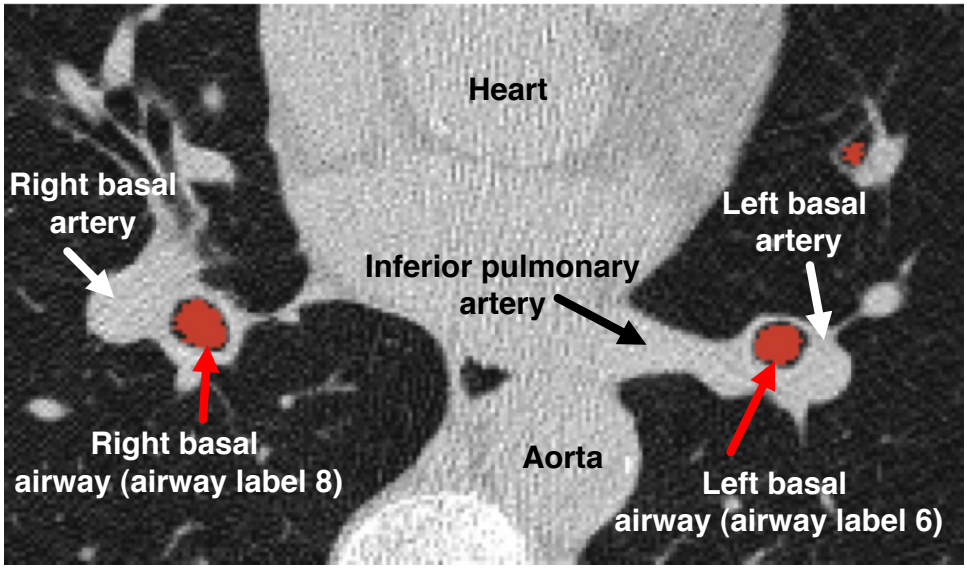

Figure 2.2: Axial slice showing the location of the left and right basal arteries relative to the airway segmentation and the lung parenchyma.

2.1.2 Region of interest construction

Given a labelled airway, the next step is to limit the search space for the pulmonary artery. This was accomplished by constructing a region of interest (ROI) around the basal airways. Lee provides a method for creating a parametrized cylindrical ROI around the airway. The segmented airway is fit with a cubic spline $\vec{S}(n)$ in three dimensional space:

$$\vec{S}(n) = s_x(n)\hat{x} + s_y(n)\hat{y} + s_z(n)\hat{z} \tag{2.1}$$

$$s_i = a_i n^3 + b_i n^2 + c_i n + d_i \tag{2.2}$$

where the parameter n defined on the boundary [0,1]. Reconstructed views are created orthogonal to the airway by taking cutting planes at constant intervals of n on [0,1]. Each cutting plane becomes a single frame in a reconstructed space representing the ROI. Each frame of the ROI reconstruction is associated with a unique coordinate system $x_i^{(n)}$ that facilitates the mapping between the CT space and the ROI. The airway reconstruction process is shown in Figure 2.3. The image on the left is a diagram in CT space of the local anatomy around the airway. The image on the right is a single frame from the reconstruction of the left basal airway.

Seed points generated in the ROI must be transformed back to the original CT image space to initiate the vessel tracking algorithm. Let $\hat{n}_i'$ be the ROI basis vectors of the slice taken at the S(n) location of the spline, and $\hat{e}_i$ be the basis vectors of the original CT image space. The transformation from a coordinate in the ROI to the CT image space is given by the transformation matrix:

$$T_n = \begin{pmatrix} \hat{n}_1' \cdot \hat{e}_1 & \hat{n}_1' \cdot \hat{e}_2 & \hat{n}_1' \cdot \hat{e}_3 \\ \hat{n}_2' \cdot \hat{e}_1 & \hat{n}_2' \cdot \hat{e}_2 & \hat{n}_2' \cdot \hat{e}_3 \\ \hat{n}_3' \cdot \hat{e}_1 & \hat{n}_3' \cdot \hat{e}_2 & \hat{n}_3' \cdot \hat{e}_3 \end{pmatrix} \tag{2.3}$$

The vector $\hat{n}_3'$ is the tangent vector to the cubic spline at $S(n)$, and $\hat{n}_1'$ and $\hat{n}_2'$ are basis vectors constructed to be orthogonal to $\hat{n}_3'$. Thus, the T matrix for

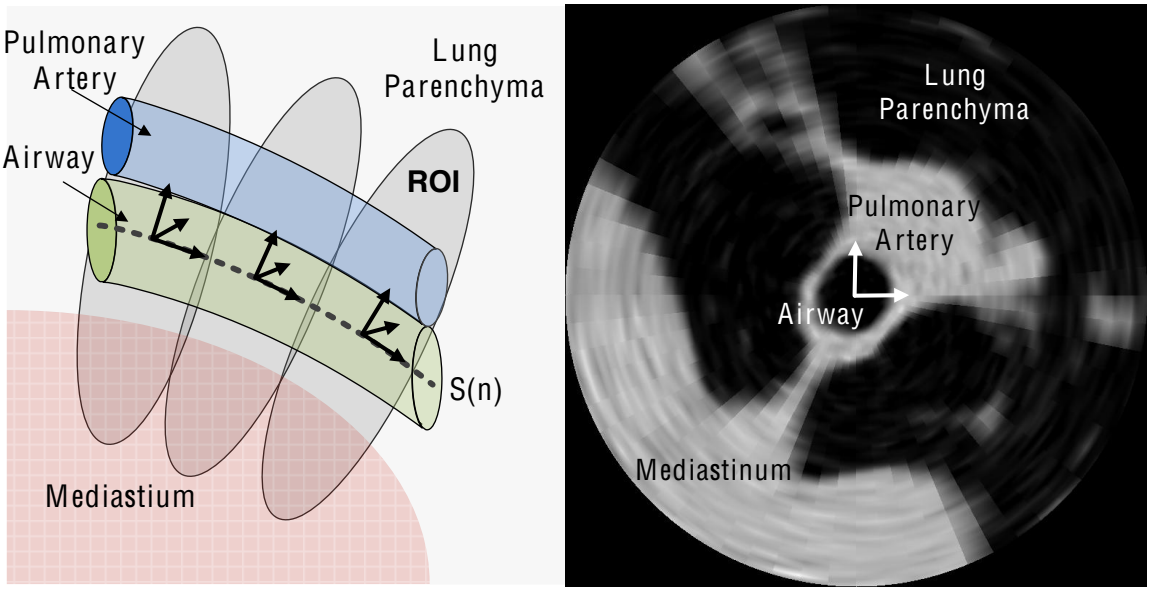

Figure 2.3: Diagram of the airway ROI construction process and a single frame from the output.

transformations back to image space can be computed given the spline parameters a_i, b_i, c_i, d_i for the three coordinates. The final transformation from the primed ROI space to the unprimed image space is given by:

$$x_i = T_{ij} x_j^{(n)}. \tag{2.4}$$

2.1.3 Artery identification

The algorithm for identifying the pulmonary arteries within the ROI around the labelled airway is based on the following observations:

1. There is very limited contrast between soft tissue structures, and structures are best identified by their morphology and anatomical location.
2. The main basal pulmonary arteries extend from the mediastinum into the lungs, where they are bordered by lung parenchyma.

3. The main basal pulmonary arteries are parallel and adjacent to the main basal airways.

The ROI around the airways is composed of either isointensity soft tissue (airway walls, blood vessels, lymphatic tissue) with no edge information, or lung parenchyma. Based on this observation it becomes appropriate to work in binary space (soft tissue or lung parenchyma) using only geometric and anatomic information for finding the arteries. Figure 2.4 shows a histogram of the ROI around the airway. The peak at 0 HU corresponds to soft tissue, and the peak at -900 HU is the lung parenchyma. A binary representation of the ROI is created by thresholding at a value of -400 HU, midway between the lung parenchyma and soft tissue histogram peaks.

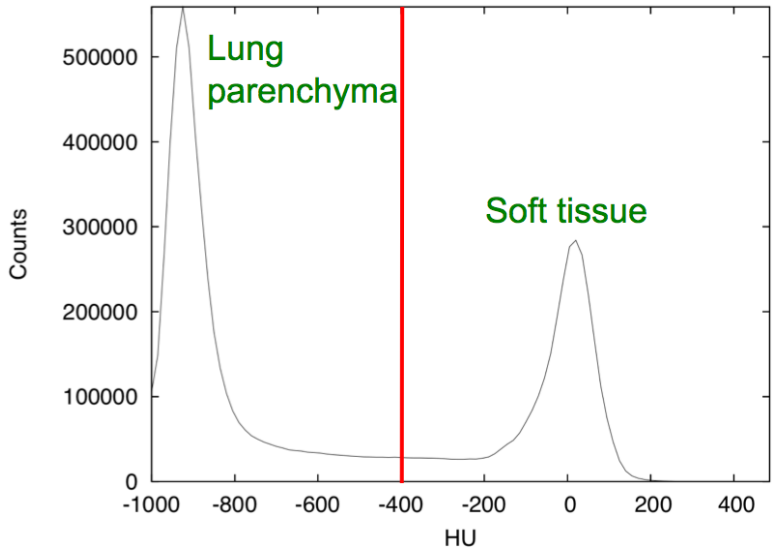

Figure 2.4: Histogram of the ROI around the left basal airway. A threshold of -400 HU was applied to separate the peaks.

The second observation indicates that because the arteries are bordered primarily by lung parenchyma, it is possible to separate the pulmonary arteries from the mediastinum by looking for structures within the ROI surrounded by lung

parenchyma. The third observation provides a key to identifying the separated soft tissue structures. The pulmonary artery will be, with high probability, the structure nearest to the airway that runs parallel with it.

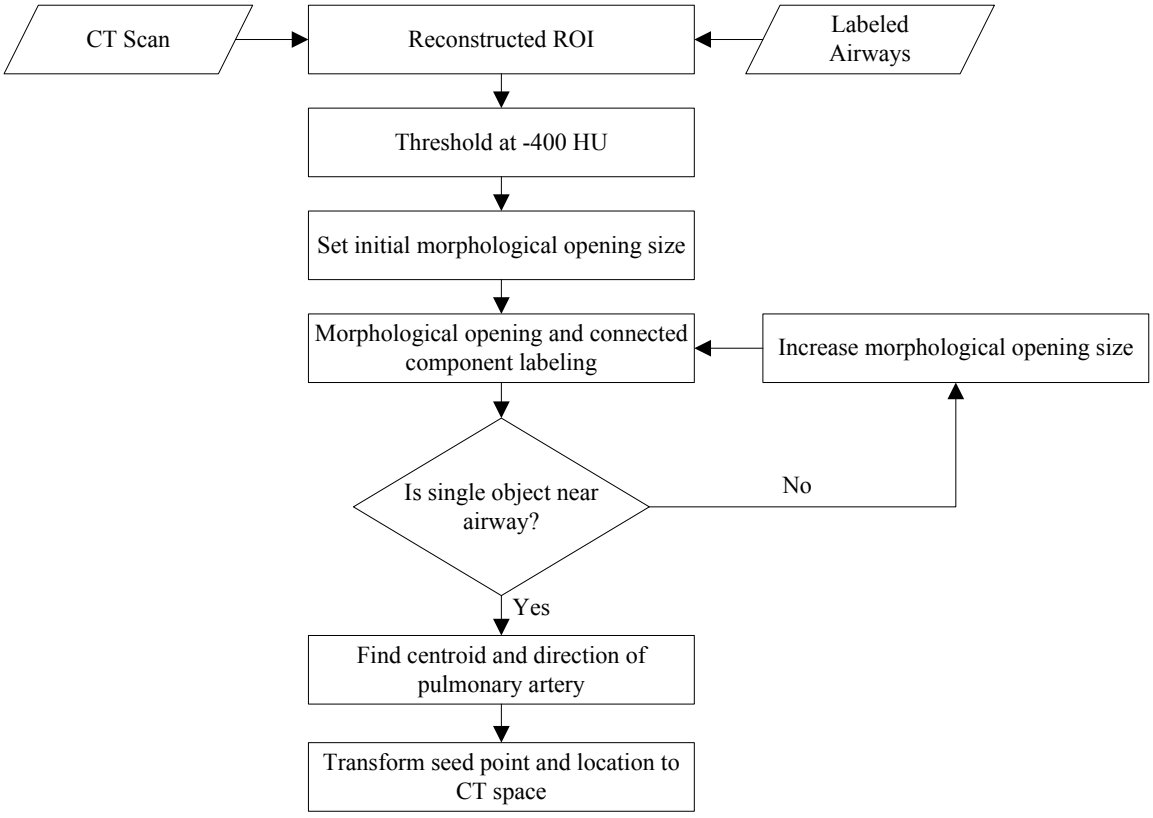

Figure 2.5: Process diagram for the identification of the main basal pulmonary arteries within the ROI

The process for locating the pulmonary arteries within the reconstructed ROI is outlined in Figure 2.5. The soft tissue structures are separated using a 3D morphological opening operation with an ellipsoidal kernel $S^{(0)}$ set to a small initial size. The purpose of the morphological opening is to remove the connections between the mediastinum and the pulmonary arteries, effectively separating the two larger objects (Figure 2.6a). The resulting objects are labelled with a 3D connected component algorithm, a method for assigning each separate object a

unique label. The connected component algorithm used assumes a 26-connected foreground and a 6-connected background.

Each unique object is then checked to see if it is a possible candidate for the pulmonary artery. Objects which extend past 4 cm from the center of the airway are likely to contain mediastinum and are eliminated. Of the objects remaining after the elimination, the component that is closest to the airway center is identified as the pulmonary artery, as predicted by the bronchopulmonary model. If no components are found, it is likely that the morphological operation opening was too small to separate the pulmonary artery from the mediastinum – both structures were eliminated together as a single object extending too far from the airway. This situation is diagrammed in Figure 2.6a. In this case, the size of the kernel is increased ($S^{(1)}$) and the process is repeated until the pulmonary artery is separated and identified, as shown in Figure 2.6b. The algorithm for isolating the artery from the mediastinum is shown in Algorithm 1.

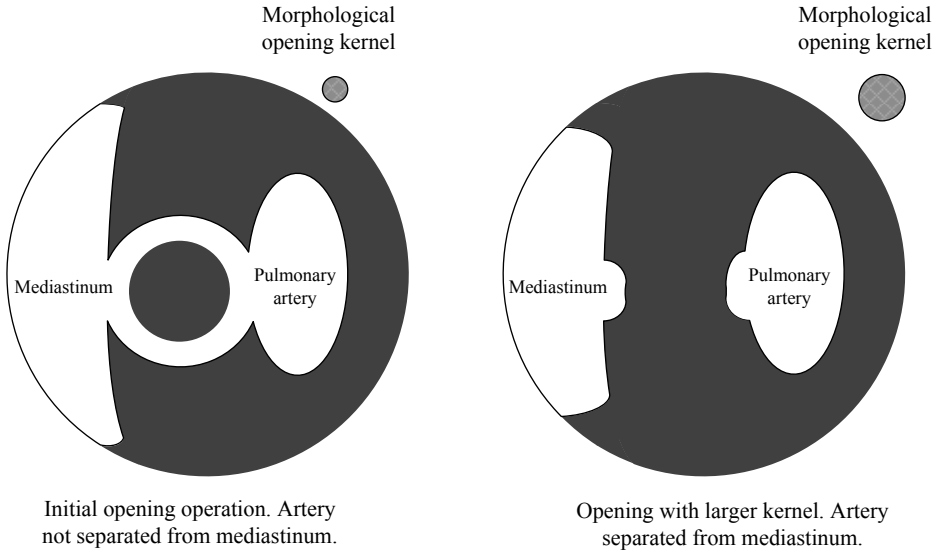

Figure 2.6: Diagram of the progressive opening operation for separating the pulmonary artery from adjacent structures within the ROI

Algorithm 1 Separate and isolate pulmonary artery from mediastinum

$S \leftarrow$ **ellipse**$(s_x^{(0)}, s_y^{(0)}, s_z^{(0)})$
while $s_x < s_{\text{limit}}$ **do**
  $B \leftarrow A \circ S$ {Morphological Opening}
  $C \leftarrow$ **label** B {Connected Component Label}
  for all Voxels $c_i \in C$ **do**
    if $\left(c_i^{(x)}\right)^2 + \left(c_i^{(y)}\right)^2 < r_{\text{lim}}^2$ **then**
      for all $c_j^{(L)}$ **equals** $c_i^{(L)}$ **do**
        remove c_j
      end for
    end if
  end for
  $r_{\text{min}}^2 = \min \left| \left(c_i^{(x)}\right)^2 + \left(c_i^{(y)}\right)^2 \right|$ {Find label with min distance to airway}
  for all Voxels $c_i \in C$ **do**
    if $c_i^{(L)} \neq c_{\text{min}}^{(L)}$ **then**
      remove c_i
    end if
  end for
  if C **equals** $\emptyset$ **then**
    $s_x \leftarrow s_x + 1$
    $s_y \leftarrow s_y + 1$
  else
    return C
    break
  end if
end while

The location of the separated pulmonary artery is defined as the location of the centroid of the candidate object at the most proximal (close to heart) frame in the ROI. This ensures that vessel tracking begins as close to the start of the vessel as possible, increasing the extent of the segmentation. The centroid coordinate $(\overline{x}, \overline{y})$ of a binary object with N pixels is defined as:

$$(\overline{x}, \overline{y}) = \frac{1}{N} \sum_{i}^{N} (x_i, y_i) \tag{2.5}$$

where the centroid is computed in the ROI coordinate system.

Similarly, a second seed point is identified by computing the centroid of the artery in the most distal (far from heart) frame it appears in the ROI. The initial orientation of the vessel is defined as the vector running from the proximal seed point to the distal seed point. Finding a seed point and approximate direction for the basal pulmonary arteries allows the vessel tracking algorithm to perform the segmentation of the identified pulmonary artery. As discussed in Chapter 3, vessel segmentation is performed by automatically optimizing a cylinder model to the vessel; the initial direction need only be an approximation.

2.2 Training and Experiment Design

Both training and experimental testing of the seed point method were done using visual validation. The first five cases in our low-dose CT dataset were selected for optimizing two parameters: kernel height s_z and ROI radius d_{ROI}. For each parameter set, the automatically generated seed points were used to initiate vessel tracking. The vessel tracking result was inspected to ensure that it 1) began in the correct vessel and 2) started cylinder tracking along that vessel into the lung parenchyma. The optimal value of the s_z was found to be 1 cm, which was sufficient to remove most small structures from the ROI without eliminating the vessel. The optimal value of d_{ROI} was 4 cm, which was large enough to capture the arteries with the ROI.

The process diagram for the seed point identification experiment is presented in Figure 2.7. The full, undocumented low-dose CT dataset was used to evaluate the detection method on a larger scale (50 cases). Visual validation of the tracking initiation was used to confirm the success of the detection method for each seed point pair. Additionally, visual inspection of the labeled airway tree was performed to evaluate the detection method success rate relative to the number of correctly

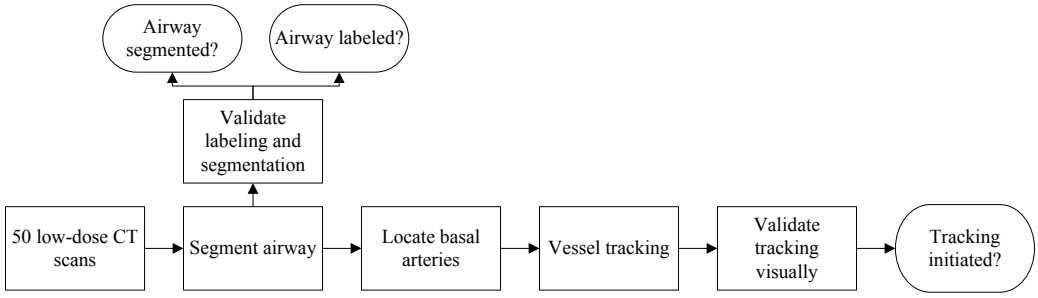

Figure 2.7: Diagram of the seed point identification experiment

generated ROIs.

2.3 Results

The results of 50 case seed point identification experiment are given in Table 2.1. The overall success rate for initiating automated tracking given a correctly generated ROI is 80%. The method was significantly more successful for the left basal artery (88%) than the right (74%). Without correcting for mislabeled airways, the success rate decreases to 72% in the left lung and 50% in the right, with an overall result of 60/100 arteries seeded.

Table 2.1: Qualitative results for the basal artery identification method, tested on 50 cases.

Dataset	Airway Segmented	Airway Labeled	Tracker Initiated
50 cases (Left only)	47	41	36 (88%)
50 cases (Right only)	47	34	25 (74%)
50 cases (Total)	94	75	60 (80%)

2.4 Discussion

An example of a successful seed point detection is shown in Figure 2.8. The images were generated by taking a single frame from the left and right basal artery ROIs, with the detected artery overlaid. In this case, both the left and right basal pulmonary arteries were correctly separated from the mediastinum, allowing for automated seed point generation. The basal arteries exhibited a clear anatomical separation from the mediastinum as they entered the lungs, which allowed for their identification as in a majority of the cases.

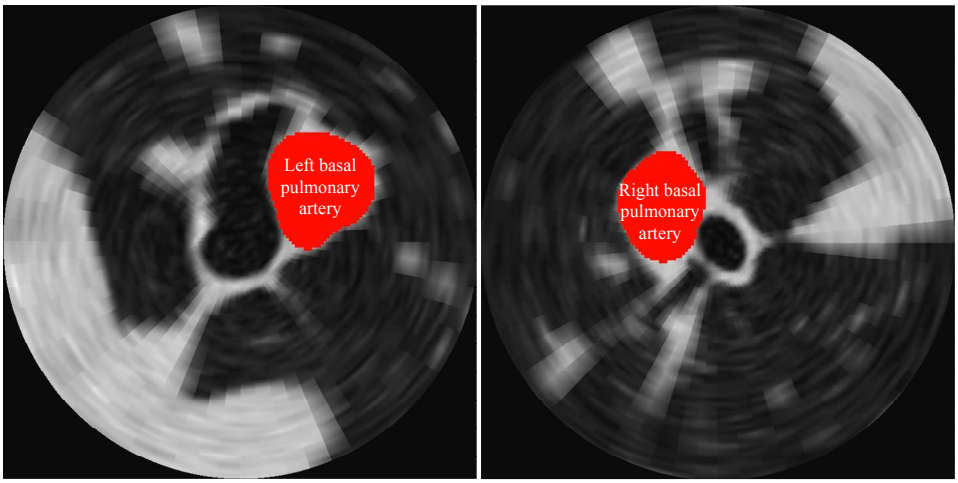

Figure 2.8: Successful seed point detection. The result of the morphological opening is overlaid on the images.

The seed point detector was not successful in all cases, and an example of a single frame from a failed case is shown in Figure 2.9. The artery never sufficiently separated from the mediastinum, and the morphological opening was unable isolate the artery. After elimination of ROI border objects, no components remained and no seed point was able to be identified.

The seed point detection method was successful in initiating tracking in a majority of cases, and represents the first reported method for automatically seeding

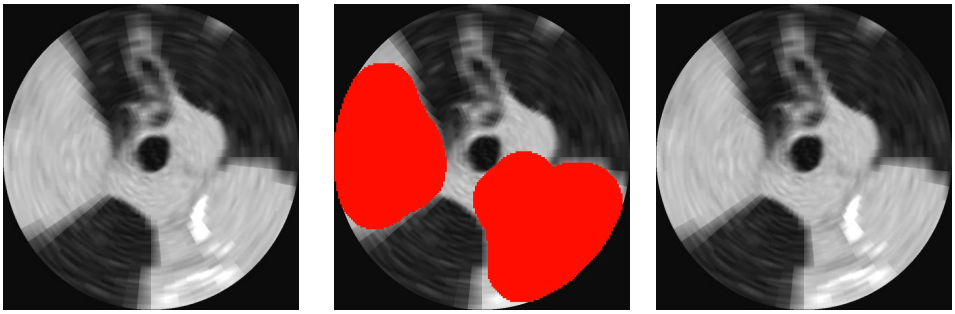

Figure 2.9: Failed seed point detection. The morphological opening does not separate the artery from the mediastinum. Connected component analysis eliminates the morphological opening, as they are not viable artery candidates. The final result is that no seed point is detected.

the pulmonary arteries. The method for locating the pulmonary artery is highly dependent on the airway labeling method by Lee. Although successful in a majority of cases, further refinement is needed in both the airway labeling and artery extraction methods.

The primary difficulty encountered by the artery extraction method occurs when the vessels are not sufficiently physically separated from the mediastinum within the ROI. The primary cause of this is that the basal airway ends before the pulmonary arteries emerge into the lung parenchyma where they can be identified. The decreased performance in the right basal artery (74%) compared with the left (88%), is most likely due to the shorter length of the right basal airway, illustrated in Figure 1.4 in Section 1.3. The artery identification method failed if the basal airway did not extend into the lung parenchyma, as was more often the case with the shorter right basal airway. One possible solution would be to extrapolate the ROI further into the lungs until the airway and artery reliably separated from the mediastinum.

Another important limitation of this method is that the parent arteries to the

left and right superior lobes are not surrounded by lung parenchyma, and are not identifiable by the method presented in this thesis. A possible extension of this work would be do develop a method for automated seed-point generation in the smaller parent branches of the upper lobes.

CHAPTER 3

METHOD FOR TRACKING THE ARTERIAL TREE

This chapter describes a method for tracking the pulmonary arteries along their lengths and into the lungs. There are two main components to this method: vessel tracking and bifurcation detection. Vessel tracking requires an initial seed point and direction, and thus builds on the method for locating the pulmonary arteries. The bifurcation detector builds on the vessel tracking method by finding the location and orientation of child vessels to be tracked.

3.1 Preliminary steps

Before applying the vessel tracking algorithm to the CT images, three preliminary steps were performed to reduce noise and decrease the signal from unwanted structures. The first step was to apply a 3×3 median filter to the image to reduce the noise inherent in low-dose scans. The CT scans were then thresholded at a value of -400 HU, to separate the soft tissue (0 HU) from the lung parenchyma (-900 HU) within the lungs. It is important to note that structures such as bone (400 HU) and fat (50 HU) were not eliminated, and took on the same values as the soft tissue after thresholding. However, because these structures do not appear in the lungs where the tracking occurs, they are highly unlikely to affect the results of the segmentation.

The third preliminary step was the partial removal of the airway walls, given in Algorithm 2. As described in Section 1.4, the airways run parallel to the vessels, and the isointensity airway walls have the potential to be confused with the pulmonary arteries. To remove the airway walls, the binary airway segmentation A provided by Lee was dilated with a $5 \times 5 \times 3$ ellipsoidal structuring element in the anisotropic image space. This corresponds to an approximately 2 mm expansion

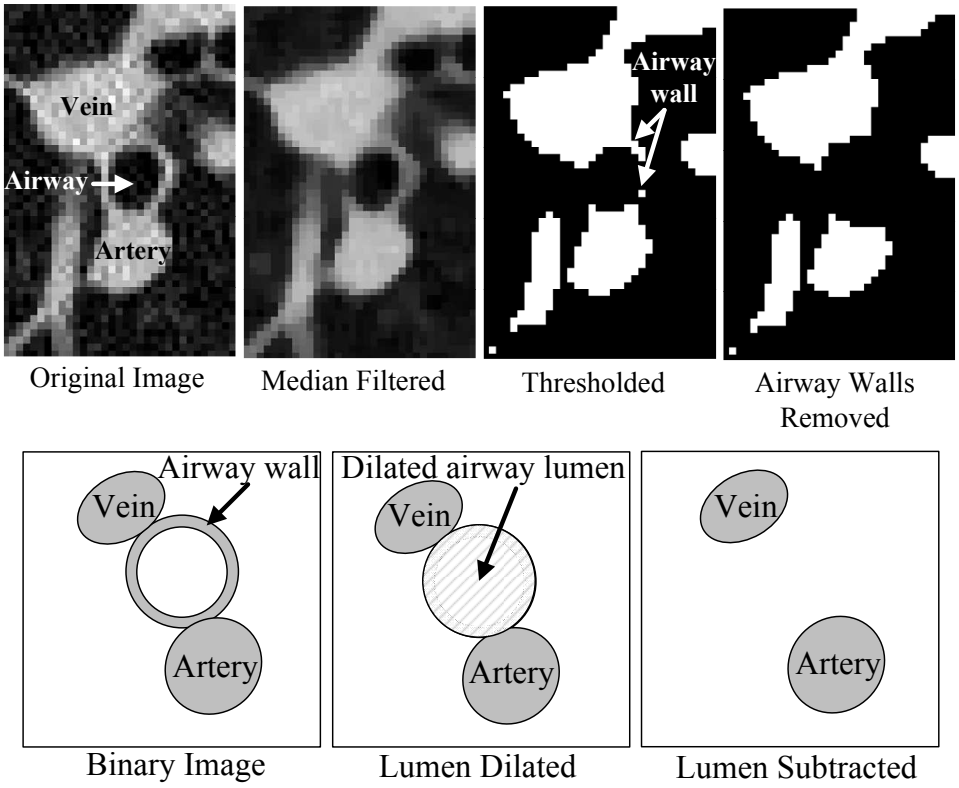

Figure 3.1: Pre-processing of CT images reduces noise, eliminates airway walls, and converts to binary image.

of the airways in isotropic space. The airway walls were removed by subtracting the dilated airway from the thresholded CT image. Figure 3.1 illustrates the application of this operation to a 3rd generation bronchopulmonary segment.

3.2 Cylinder matching metric

A single vessel can be modeled as a series of discrete cylinders as shown in Figure 3.2. This model is based on the observation that vessels are approximately cylindrical in shape and curve only slowly along their lengths. Vessel tracking is the task of recreating the vessel by progressively fitting model cylinders in a binary 3D discrete CT image space W with voxels $w_1, w_2, ...w_i$.

Algorithm 2 Threshold image and remove airway walls

get W {Binary CT image}
get A {Airway segmentation}
for all $w_i \in W$ **do**
  if $w_i > w_{\text{thresh}}$ **then**
    $w_i \leftarrow 1$
  else
    $w_i \leftarrow 0$
  end if
end for
$S \leftarrow (5, 5, 3)$ {Structuring element}
$B \leftarrow A \oplus S$ {Dilate segmentation}
return $W \wedge \overline{B}$

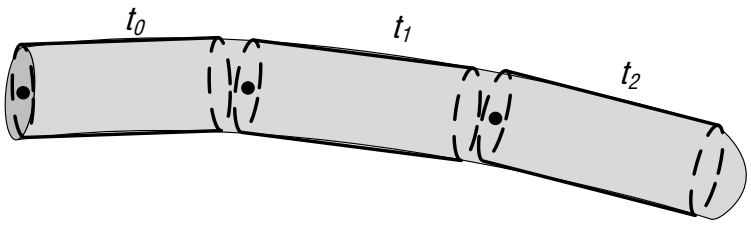

Figure 3.2: Modeling of a blood vessel with discrete cylinders

To determine the level of similarity between a model cylinder and the actual vessel, a similarity metric must be defined. The metric must return a maximum score when the model cylinder fits directly over the vessel, and minimized when there is no overlap. Additionally, because we are searching for a volume representation of the vessel, it is important to define a similarity metric that gives a greater score to model cylinders that encompass a larger volume of the actual vessel.

To construct a similarity metric, define the set C of all voxels in a model cylinder with a given location and orientation, and the set W of voxels in the thresholded image. The similarity score between the cylinder and the image is computed by Algorithm 3. A penalty of -5 was selected to prevent the cylinder tracker from jumping between adjacent soft tissue structures; model cylinders that span areas of lung parenchyma are highly penalized. The *ad hoc* value of -5 inserted enough

gain into the system to prevent leaks across parenchyma without overly penalizing well-fitting cylinders. The similarity metric is *not* normalized to the size of the cylinder, giving model cylinders that capture more of the vessel a higher similarity score, increasing the extent of the vessel segmentation.

Algorithm 3 Match cylinder to CT image

```
get W {Binary CT image}
get C {Model cylinder}
M ← 0{Cylinder match score}
for all w_i ∈ W do
  if w_i ∈ C then
    if w_i^(L) > 0 then
      M ← M + 1
    else
      M ← M − 5
    end if
  end if
end for
return M
```

3.3 Iterative vessel tracking

The iterative vessel tracking process is outlined in Figure 3.3. The algorithm begins with an initial seed point x_0 and normalized direction $\hat{d}'_0$. The initial seed point is fixed as the start point of the first model cylinder. From this initial point, model cylinders of varying orientations and radii are evaluated for their similarity to the vessel. The initial radius r_0 is optimized on the interval [2 mm, 12 mm], which is sufficient to cover both abnormally large and small vessels. The search space for finding the optimal orientation of the model cylinder is limited to the solid angle $\Omega = 2\pi$, parametrized on the azimuthal angle α and the polar angle β. The α parameter was discretely sampled between 0 and $\pi/2$ radians at a sampling step size $\triangle\alpha$ of $\pi/32$. Similarly, the β parameter was sampled between 0 and 2π at step

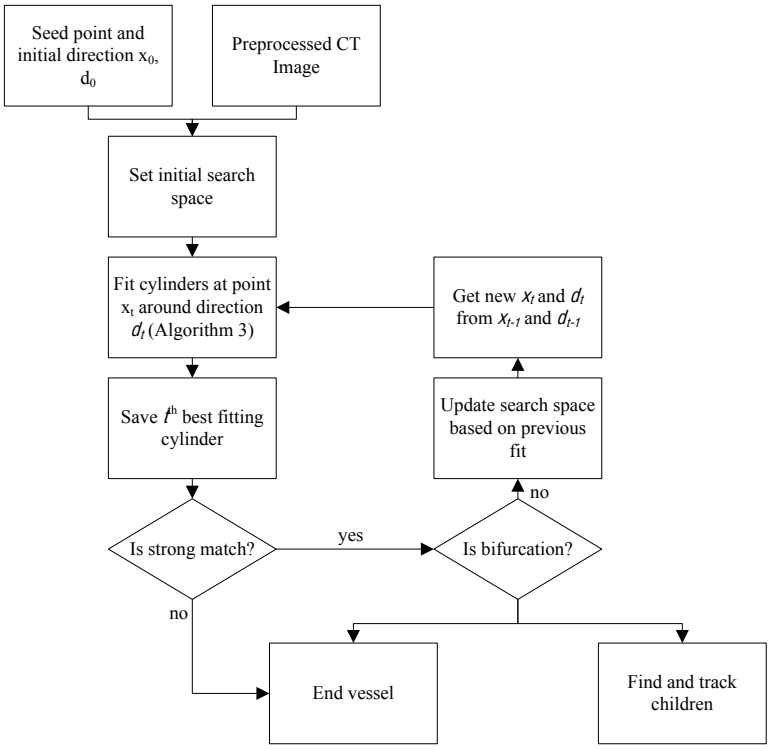

Figure 3.3: Iterative vessel tracking process diagram.

sizes $\triangle\beta$ of $\pi/32$. The step size of $\pi/32$ radians was selected to keep the run-time to roughly 3 minutes per lung.

Finally, in order to limit the size of the model fitting parameter space, the height of each model cylinder is set to a fixed value of h_0. The model cylinder which returns the greatest similarity score is defined as the initial cylinder, with starting location x_0, radius r_0, height h_0, and direction $\hat{d}_0$. The algorithm used to implement this process is given in Algorithm 4.

After the initial model cylinder is fit, the process repeats for a new initial position x_1, shown in Figure 3.4. For the t^{th} iteration of cylinder fitting, the initial point x_t is defined as:

$$x_t = (\Delta_{\text{step}})\hat{d}_{t-1} + x_{t-1} \tag{3.1}$$

where $\hat{d}_{t-1}$ is the normalized direction of the previous model cylinder, and Δ_{step} is

Algorithm 4 Fit t cylinder to vessel

Require: $t > 0$
  get W {CT image}
  $r_{\min} \leftarrow 0.7 r_{t-1}$
  $r_{\max} \leftarrow 1.3 r_{t-1}$
  $M_t \leftarrow -1000$ {Initial match score for t cylinder}
  generate $\hat{o}_1$ **such that** $\hat{o}_1 \perp \hat{d}_t$
  generate $\hat{o}_2$ **such that** $\hat{o}_2 \perp \hat{d}_t, \hat{o}_1$
  for $r = r_{\min}$ **to** $r_{\max}$ **by** $\triangle r$ **do**
    for $\alpha = 0$ **to** $\pi/4$ **by** $\triangle \alpha$ **do**
      for $\beta = 0$ **to** 2π **by** $\triangle \beta$ **do**
        $\hat{a}_1 \leftarrow \hat{o}_1 \cos\beta + \hat{o}_2 \sin\beta$
        $\hat{a}_2 \leftarrow \hat{d}_t \cos\alpha + \hat{a}_1 \sin\alpha$
        $x'_t \leftarrow x_{t-1} + \hat{a}_2 h_0$
        $C \leftarrow$ cylinder of radius r between x_{t-1} and x'_t
        $M \leftarrow$ **match** (C, W)
        if $M > M_t$ **then**
          $M_t \leftarrow M$
          $x_t \rightarrow x'_t$
          $r_t \leftarrow r$
          $\hat{d}_t \leftarrow \hat{a}_2$
        end if
      end for
    end for
  end for
  return $x_t, r_t, M_t, \hat{d}_t$

the distance the next cylinder is progressed from the previous one. The orientation search space remains bounded by $\Omega = 2\pi$. Additionally, because the radius is expected to vary only slowly between model cylinders, the radius search space is limited as:

$$r_t \in [0.7 r_{t-1}, 1.3 r_{t-1}] \tag{3.2}$$

The full set of vessel tracking parameters is given in Table 3.1. The cylinder radius r and orientation α and β are optimized for each cylinder at run time. The incremental step size δ_{step} and cylinder height h_0 are fixed at values optimized durning training (Section 5.2).

Iterative vessel tracking terminates when the tracker either fails to find a strong

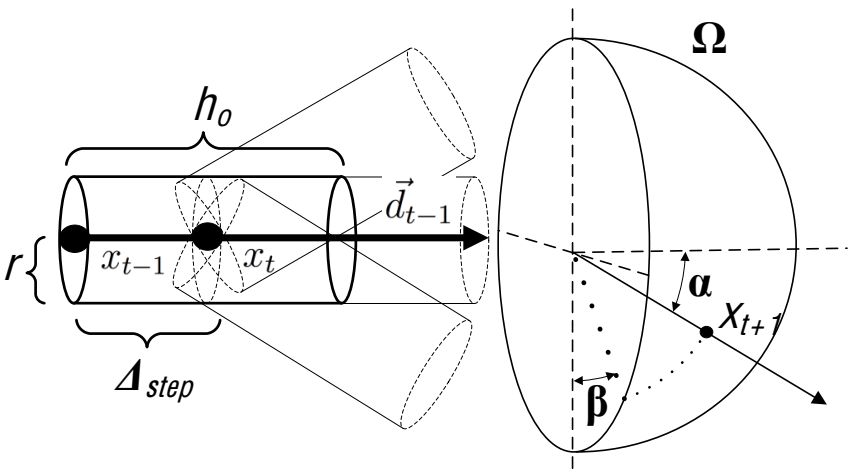

Figure 3.4: Iterative fitting of the model cylinders

Table 3.1: Static and run-time optimized parameters of the vessel tracking method

Parameter	Description	Search space or fixed value	Step size
h_0	Cylinder height	15 mm	-
Δ_{step}	Incremental step	$0.20h_0$	-
r	Cylinder radius	$[0.7, 1.3]r_{t-1}$	$\triangle r = 0.25$ mm
α	Azimuthal angle	$[0, \pi/2]$	$\triangle\alpha = \pi/32$
β	Polar angle	$[0, 2\pi]$	$\triangle\beta = \pi/32$

match or a bifurcation is detected. A strong match is defined as having at least 50% of the voxels in the model cylinder be soft-tissue. This criterion fails to be met when either the vessel becomes too small to give a strong signal, as in the case of higher generational subsegmental arteries, or the tracker has fallen off the vessel. The condition for termination by vessel bifurcation relies on the method outlined in the following section.

3.4 Bifurcation detection

Child vessels were segmented by implementing a bifurcation detector. The detector was designed to identify where on the parent branch to look for potential child

branches, and to initiate tracking of the child vessels at the proper location and direction. Detection is primarily complicated by the proximity of arteries to veins. Thus, it is important to develop a bifurcation detector with a low false-positive rate to maintain the separation of the venous and arterial trees.

Automated bifurcation detection is performed by comparing the actual vessel geometry to a bifurcation model. Child vessel trackers are only initiated all the parameters of the model are met. This serves to prevent erroneous segmentations at points that are unlikely to contain true vessel branching. The bifurcation model is constructed from the following observations:

1. The radii of the child vessels are measurably smaller than the radii of their parent vessel.
2. The angle between child vessels is no less than 30° and no more than 90°.
3. The angle between the child vessels and the parent is greater than

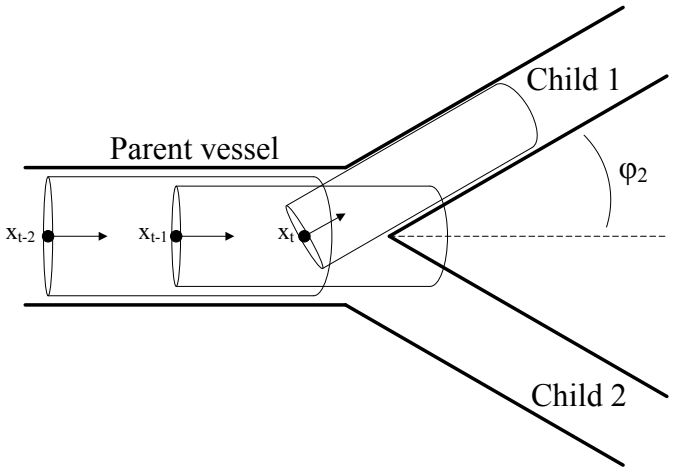

Figure 3.5: Bifurcation detection using model cylinder radius change

The first observation indicates the bifurcation detector to search for child vessels when the parent vessel tracker undergoes a large decrease in radius. As indicated in

Figure 3.5, when the parent vessel tracker encounters a bifurcation, the next model cylinder will fit onto one of the child vessels with a smaller radius. Because the cylinder matching metric highly penalizes oversegmentation, a cylinder straddling a parent and child vessel will take on the smaller radius of the child vessel. Using this model, a bifurcation candidate is generated when the ratio of the current radius to the radius at $\frac{1}{2}h_0$ upstream, $r_0/r_{\frac{h_0}{2}}$, falls below radius change threshold δ_{radius}, a bifurcation candidate is detected at point x_t. This calculation is presented in Algorithm 5.

Algorithm 5 Bifurcation detection

$t_{\text{check}} \leftarrow h_0/(2f_{\text{step}})$
Require: $t \geq t_{\text{check}}$
  if $r_t/r_{t_{\text{check}}} \leq \delta_{\text{radius}}$ **then**
    $t \leftarrow t_{\text{lim}} + 1\{\text{Stops parent vessel}\}$
    return $1\{\text{Bifurcation detected}\}$
  else
    return 0
  end if

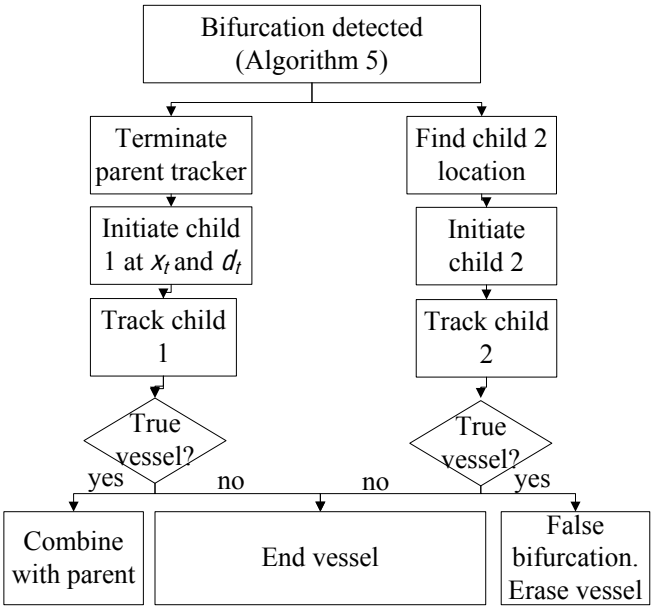

Figure 3.6: Process diagram for locating child vessels at bifurcation point.

The process for finding the child vessels at bifurcation candidate points is outlined in Figure 3.6. When the parent tracker detects a bifurcation at x_t, a new child vessel (child 1) is initiated at x_t and in the direction $\vec{d_t}$. The next step is to search for the second child vessel, and begin a vessel tracker along its length. Observations 2 and 3 of the bifurcation model indicate where to look for bifurcation candidates. Let the orientation of the parent vessel at the bifurcation be $\hat{d}_p$, and the orientation of children 1 and 2 at the bifurcation be $\hat{d}_1$ and $\hat{d}_2$ respectively. Child vessel 2 will satisfy the bifurcation model if:

$$\cos 30^{\circ} < \left(\hat{d}_1 \cdot \hat{d}_2\right) < \cos 90^{\circ} \qquad \text{and} \qquad \left(\hat{d}_2 \cdot \hat{d}_p\right) < \cos 90^{\circ} \tag{3.3}$$

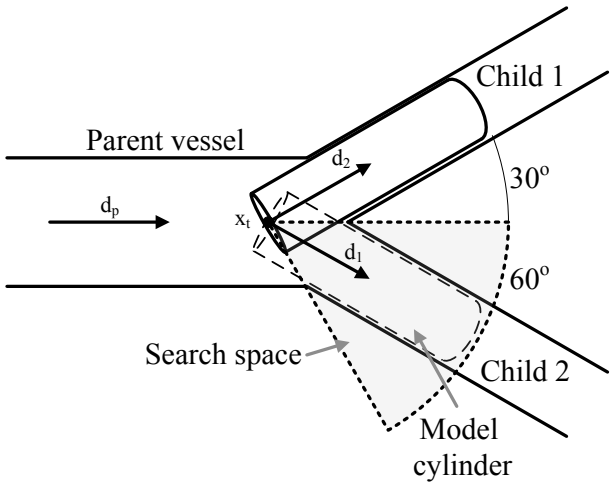

Figure 3.7: Search space for locating the second child vessel at a bifurcation point.

The search space that satisfies these conditions is illustrated in Figure 3.7. Algorithm 6 explicitly states how the search space is sampled for the sister vessel, using a rotation R^1 about direction $\hat{o}$ by angle α. The orientation of the best fitting model cylinder within this search space becomes the initial orientation of a new vessel tracker for child 1, with initial point x_t.

$$^1R = \begin{pmatrix} o_x^2 + (1-o_x^2)\cos\alpha & o_x o_y(1-\cos\alpha) - o_z\sin\alpha & o_x o_z(1-\cos\alpha) + o_y\sin\alpha \\ o_x o_y(1-\cos\alpha) + o_z\sin\alpha & o_y^2 + (1-o_y^2)\cos\alpha & o_y o_z(1-\cos\alpha) - o_x\sin\alpha \\ o_x o_z(1-\cos\alpha) - o_y\sin\alpha & o_y o_z(1-\cos\alpha) + o_x\sin\alpha & o_z^2 + (1-o_z^2)\cos\alpha \end{pmatrix}$$

Algorithm 6 Find sister vessel orientation bifurcation point

get W {Preprocessed CT}
get $\hat{d}_1$ {First child direction}
get $\hat{d}_p$ {Parent direction}
get x_0 {Bifurcation point}
$\hat{o}_1 \leftarrow \hat{d}_1 \times \hat{d}_p / \left|\hat{d}_1 \times \hat{d}_p\right|$
for $\theta = 0$ **to** 2π **by** $\triangle\theta$ **do**
  for $\phi = \phi_{\min}$ **to** $\phi_{\max}$ **by** $\triangle\phi$ **do**
    $\hat{o}_2 \leftarrow R(\theta, \hat{d}_1)\hat{o}_1$
    $\hat{d}_2 \leftarrow R(\phi, \hat{o}_2)\hat{o}_2$
    if $\hat{d}_2 \cdot \hat{d}_p > \cos\phi_{\text{parent}}$ **then**
      $x_1 \leftarrow x_0 + h_0\hat{d}_2$
      $C \leftarrow$ cylinder of radius $r_{\text{child 1}}$ between x_0 and x_1
      if match $(C, W) > M_{\max}$ **then**
        $M_{\max} \leftarrow$ **match** (C, W)
        $\hat{d}_{\text{sis}} \leftarrow \hat{d}_2$
      end if
    end if
  end for
end for
return $\hat{d}_{\text{sis}}$

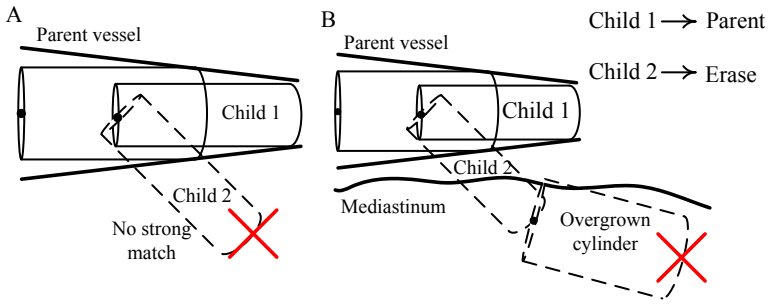

Figure 3.8: False positive bifurcation detection and error correcting. A) The child vessel does not find any strong matches. B) The child vessel leaks onto the mediastinum.

In the case of a false positive for bifurcation detection, child 2 may terminate after only a few iterations. This situation is presented in Figure 3.8a. To correct the error generated by the false bifurcation detection, the abruptly terminated child 2 vessel is erased, and child 1 continues to track, but as a continuation of the parent

vessel. Additionally, a leak detector was implemented to remove false positive child vessel returning large similarity scores when they leak onto the large mediastinum. A leak is detected when a model cylinder is fit with a radius over 150% of the radius of the model cylinder beginning at one full cylinder length upstream. A diagram of the case of a leaking child vessel is illustrated in Figure 3.8b.

CHAPTER 4

SPARSE SURFACE EVALUATION METRIC

The pulmonary arteries are a highly complex branching structure, and quantifying the quality of a vessel segmentation is a difficult task. Modern CT scanners produce images with a total number of voxels on the order of 10^8, with roughly 10^5 voxels belonging to the pulmonary arteries in a typical thoracic scan. Capturing both the full geometry and finer voxel-level detail of the vessel tree for validation requires extensive manual markings. This places a practical limit on the number of cases that can be used for validation, indicating a tradeoff between the size of the dataset and the level of detail of the analysis.

This chapter will first review the major evaluation metrics that have been previously used for the problem of pulmonary vessel segmentation. A novel validation metric called the Sparse Surface Metric will then be described. The SS method was designed to provide quantitative evaluation of both the total number of vessels segmented and the accuracy per vessel, while reducing the number of manual markings required. The metric will be defined here in general terms, and is capable of being extended to other segmentation problems with complex geometries.

4.1 Previously reported evaluation metrics

Table 4.1 provides a summary of the relevant previous work done in pulmonary vessel segmentation validation. Full volume based comparisons, where the ground truth is marked to produce a complete volume representation of the desired structure, are able to quantify both the detail and extent of a segmentation; however, they require the greatest number of manually marked points.

Saha et al [20] performed the most extensive volume-based ground-truth marking reported for pulmonary vessel segmentation – an expert classified 8000-10000

Table 4.1: Summary of previously reported metrics for evaluating pulmonary vessel segmentations.

Author	Year	Segmentation	Method	# of cases
Saha [20]	2010	Pulm. A/V separation	Full volume	2
Bemmel [11]	2008	Abdominal A/V separation	ROI volume	7
Kaftan [11]	2008	Pulm. vasc.	ROI volume	5
C. Zhou [27]	2005	Pulm. vasc.	Centerline	2
Shikata [23]	2009	Pulm. vasc.	Centerline	44
Masutani [15]	2001	Pulm. vasc.	Visual	12
This thesis	2010	Pulm. arteries	Sparse surface	10

voxels in each lung, per CT scan, as either belonging to the arterial or venous trees. Quantitative validation was performed by computing the sensitivity of artery identification:

$$\text{Sensitivity} = \frac{|G_{\text{artery}} \cap S|}{|G_{\text{artery}}|} \tag{4.1}$$

weighed against false positives and misses (false negatives):

$$\text{False} = \frac{|G_{\text{vein}} \cap S_{\text{artery}}|}{G_{\text{artery}}} \qquad \text{Miss} = \frac{|G_{\text{artery}} - S_{\text{artery}} \cap S_{\text{vein}}|}{G_{\text{artery}}} \tag{4.2}$$

where G is the set of ground truth points manually marked to capture the volume of the arteries, and S is the set of voxels belonging to the segmented volume. The strength of their validation system was that it contained a method of quantifying both the extent and accuracy of the artery segmentation. However, ground truth markings were only performed for two CT cases, suggesting that this method is not suitable for application to larger datasets.

Similar volumetric manual marking were implemented by Kaftan et al [11] to compute the sensitivity and specificity of their pulmonary vessel segmentation algorithm. To reduce the number of manually placed ground truth points, they limited their markings to 6 ROIs placed at different points within the lungs. Similarly, Bemmel et al [26] calculated a true positive fraction and false positive fraction for seven cases by marking ground-truth volumes within 6 manually selected ROIs per

case. Although limiting the extent of the ground truth markings to selected ROI allowed for a larger dataset, the markings for each scan do not capture the full geometry of the vessels – the choice of where to place the ROIs has the potential to significantly alter the results.

Zhou (Chuan) et al [27] reduced their number of ground-truth points by marking only vessel centerlines rather than defining vessel volumes. In their study, a total of 9421 manually marked points were placed at the centerlines of arteries and veins across two CT scans. This allowed them to compute the sensitivity and specificity of their algorithm up to the 7th generation of vessels, while requiring significantly fewer points than volumetric methods. Shikata et al [23] also implemented centerline marking by identifying over 1000 points at vessel centerlines for each scan. To estimate the false positive rate, they inserted 150 randomly placed points per scan outside of the pulmonary vessels.

Although centerline marking is valuable metric for quantifying the extent of a segmentation, it cannot provide information on the quality of the segmentation at points that lie off of the centerlines. Erroneous segmentation of lymphatic tissue, the airways, or other neighboring structures, are not penalized by a centerline marking system.

4.2 Sparse surface metric

We propose a novel surface-based method called the sparse surface (SS) metric for validating the results of our pulmonary artery segmentation method. The goals in developing a new validation criteria to be applied to pulmonary artery segmentation were three-fold:

1. To develop a metric and marking system to quantify the number of vessels segmented (i.e. extent of segmentation).

2. To use the same markings to also evaluate the error per vessel (i.e. the accuracy of segmentation).

3. To minimize the number of manual ground truth markings required of an expert human user.

Dense volumetric markings accomplish the first two goals, but require a significant number of manual points. ROI markings are able to find the error per vessel, but are not able to evaluate the overall extent of the segmentation. Conversely, centerline markings capture the extent, but not the accuracy, of the segmentation.

The SS method accomplishes the three listed goals by defining a mapping from ground truth points marked at the surface of vessels to the surface of the segmentation. By mapping from the ground truth to the segmentation, ground truth can be marked only sparsely – there is no penalty incurred by segmentation surface points that are not close to ground truth points. The full geometry of the vascular tree can be captured by marking a limited number of surface points for each vessel, allowing for the calculation of the extent of the segmentation. Similarly, the The following section formally defines the SGT metric.

4.3 Formal definition

Let V be the set of voxels of the segmented volume. The boundary of V is the set ∂V of voxels that are 6-connected to a background voxel. The set S of surface points of a segmented volume V is then defined as:

$$S = \{x | x \in \partial V\} = \{s_1, s_2, ..., s_k\} \tag{4.3}$$

Let G be the set of ground-truth points, where the members g_i were individually identified by manual markings. The set G contains n subsets O_n representing

the ground-truth markings for n objects. The SGT metric is constructed from a mapping f of O_n to M_n such that for every element o_i of a set O_n:

$$f(o_i) = m_i = a_i \min_{j=1,\ldots,k} \{\|(o_i, s_j)\|_E\} \tag{4.4}$$

where the parameter a_i describes whether the point is interior or exterior to the segmentation:

$$a_i = \begin{cases} -1 & \text{if } m_i \in V \quad \text{(interior)} \\ \\ 1 & \text{otherwise} \quad \text{(exterior)} \end{cases} \tag{4.5}$$

The isotropic distance metric $\|(a, b)\|_E$ is the positive semi-definite 3D Euclidean distance between two points:

$$\|(a, b)\|_E = \sqrt{\left[\frac{(a_x - b_x)}{x_r}\right]^2 + \left[\frac{(a_y - b_y)}{y_r}\right]^2 + \left[\frac{(a_z - b_z)}{z_r}\right]^2} \tag{4.6}$$

where (x_r, y_r, z_r) are the lengths the vectors defining a unit cell in anisotropic space (i.e. the length, width and height of a voxel).

The set M_n now holds the values of the Euclidean distance from each ground truth point for object n to the nearest surface point of the entire segmentation, with positive values corresponding to ground truth points that lie outside the segmentation and negative values for interior points. An illustration of the SGT metric being applied to an example segmentation is provided in Figure 4.1.

The distribution of distance elements in each of the sets M_n is expected to take on a normal unimodal distribution. The ideal algorithm will return a narrow distribution centered around zero. For an over-segmentation, the distribution will be shifted towards negative distances, and for an under-segmentation, will be shifted towards positive distances.

Quantitative information about the segmentation can be obtained by computing statistics on the sets M_n. The average signed distance (ASD) $\overline{M_n}$ of the set

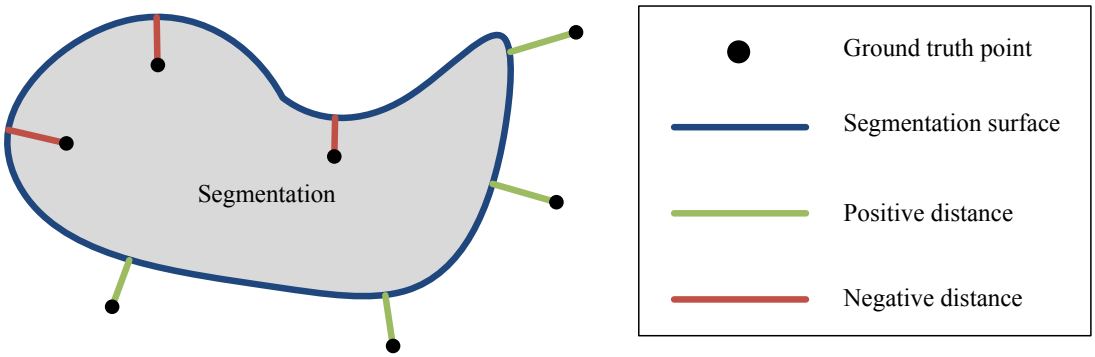

Figure 4.1: Diagram illustrating the SS metric. The illustrated distances comprise the set M.

M_n of order N can be computed as:

$$\overline{M_n} = \sum_{i}^{N} \frac{m_i}{N} \tag{4.7}$$

The average signed distance is a measure of the overall bias of the segmentation – an ASD near zero results from data whose collective distance between the segmentation surface and points outside the segmentation equals the collective distance for points inside the segmentation. To evaluate the magnitude of the average error per truth point, the root mean square distance (RMSD) can be calculated for each object distance error set M_n.

$$M_{\text{nRMS}} = \sqrt{\sum_{i}^{N} \frac{m_i^2}{N}}. \tag{4.8}$$

Additionally, the maximum amount of over- or under-segmentation can be found by calculating the maximum negative distance (over-) and maximum positive distance (over-). These measures classify which error is the most severe. Table 4.2 summarizes the statistics used in the SS method.

Given a set of surface markings for each object, the distance error statistics obtained using the SS metric as defined above can be used to achieve our goals of evaluating the number of objects segmented, and the error per object. An object

Table 4.2: Statistical measurements used in the SS evaluation metric

Statistic	Measure of
Average signed distance (ASD)	Accuracy/bias of segmentation
Root mean square distance (RMSD)	Magnitude of distance error
Max negative distance	Point of greatest over-segmentation
Max positive distance	Point of greatest under-segmentation

may be classified as having been segmented if the RMS truth-to-surface error is small, signifying that the segmentation did not stray far from the object. The bias, magnitude, and extreme values of the error can then be calculated for each segmented object.

4.4 Comparison with dense markings

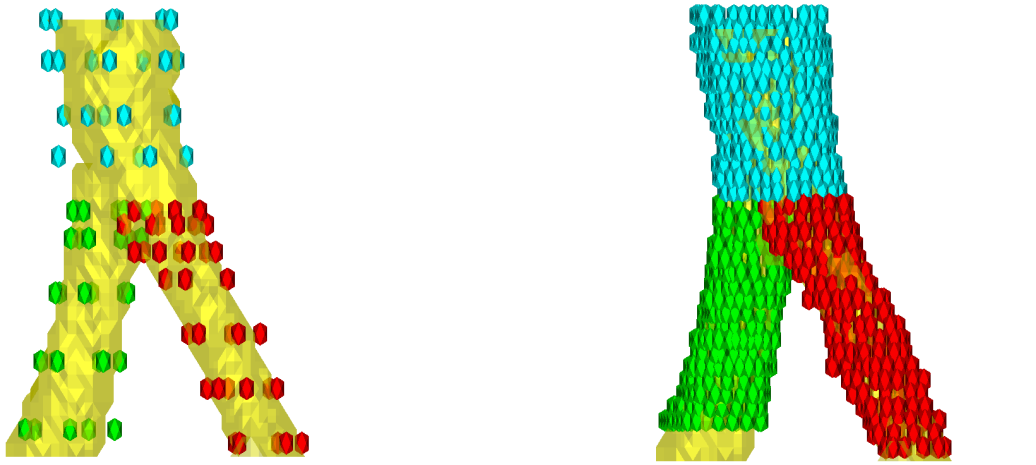

Figure 4.2: Manual SS markings overlaid on the left main pulmonary artery and its two child branches

To evaluate the new SS method, case 1 was documented with 4000 dense marking and 300 sparse markings. Figure 4.2a shows the sparse markings for three pulmonary arteries, shown here overlaid on the surface of a segmentation. Fig-

ure 4.2b shows surface markings for a dense, volumetric ground-truth on the same vessels. The ratio in the number of points between the dense markings and the sparse markings is 10:1.

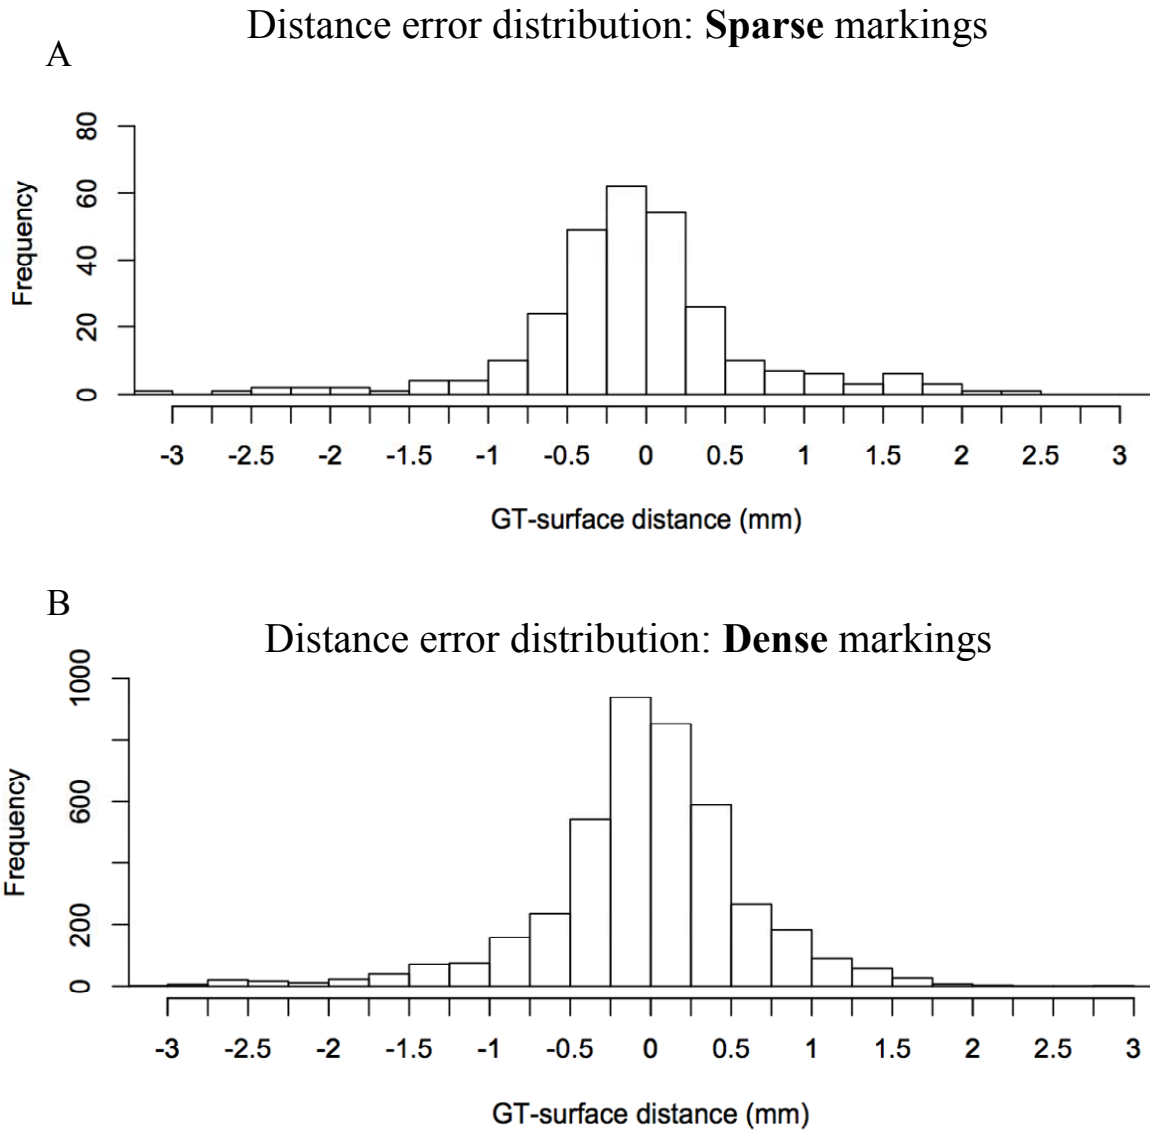

Figure 4.3: Histograms of the truth-to-surface error applied to case 1 for a) sparse markings and b) dense markings.

The distributions of truth-to-surface distance errors for both sparse and dense markings are presented in Figure 4.3. Both histograms exhibited a normal distribution, and an F-test indicated the distributions had no statistical difference in

variation. This allowed an unpaired t-test to be performed to test the hypothesis that both samples (dense and sparse) are drawn from the same population (the true surface). A small difference in means between the two samples populations indicates they are likely the same population.

The 95% confidence interval of mean differences between the population containing the sparse markings and the population containing the dense markings was -0.06 mm to 0.12 mm. The hypothesis that both samples represent the same population (i.e. the true vessel surface) is well supported. Thus we find that the by applying the SS method to sparse markings, we can decrease the manual marking burden by at least one order of magnitude and still robustly evaluate the segmentation extent and accuracy.

CHAPTER 5

VESSEL TRACKING OPTIMIZATION

This chapter discusses the training and optimization of the methods used in this thesis, as well as presents the details of the experiments used to evaluate the segmentations.

5.1 Dataset

The dataset used for this experiment consisted of 50 low-dose thoracic CT scans. From these 50 cases, 10 cases with accurate airway segmentations and labels were selected for ground truth marking. The scans were generated with an 8 slice helical CT scanner operating under a low-dose protocol with beam current of 40 mA. The full set of scan and image parameters for the dataset are presented in Table 5.1. The scans were tagged with standard Digital Imaging and Communications in Medicine (DICOM) headers, with the patient-identifying information removed. The image resolution information, used in our method to implement isotropic methods in an anisotropic image, were automatically retrieved from the DICOM header.

Table 5.1: CT scanning parameters for the 50 cases in the low-dose data set.

Parameter	Value
CT scanner model	GE LightSpeed Ultra
Accelerating voltage	120kV
Anode current	40 mA
Gantry aperture	70 cm
Number of slices	8
Gantry revolution time	500 ms
Axial image size	512×512
Image format	16 bit
Slice thickness	1.25 mm
In-plane resolution	0.55 - 0.82 mm
Patient position	Feet first supine

5.2 Performance evaluation with the SS metric

Ground truth points were marked on axial slices of CT scans windowed to soft tissue. An approximately homogenous mat of ground truth points was created by marking vessels at roughly 4 mm intervals around their circumference for roughly every third frame, corresponding to about a 4 mm interval in the axial direction. Each vessel was marked separately and given a unique label. Only left and right basal pulmonary arteries child vessels with radii greater than 1 mm were marked. Additionally, veins with radii > 1 mm in the conflated inferior pulmonary venous tree were individually marked.

Each marked vessel was classified as either having been segmented or missed by measuring the RMS error per vessel. A vessel was classified as having been segmented if the RMS the distance was below 2.0 mm, which corresponds to an average error of roughly 3-4 voxels. The threshold of 2.0 mm is an *ad hoc* value that represents the upper error bound for vessels visually confirmed to be segmented. Using this method, each vessel was classified into one of four categories, given Table 5.2, for evaluating the performance of the algorithm.

Table 5.2: Vessel classification for evaluating algorithm performance

Category	Type of vessel	Criterion
True Positive (TP)	Artery	RMS error less than 2.0 mm
False Negative (FN)	Artery	RMS error greater than 2.0 mm
True Negative (TN)	Vein	RMS error greater than 2.0 mm
False Positive (FP)	Vein	RMS error less than 2.0 mm

The primary goal of the method was to maximize the number of arteries segmented while minimizing the number of veins segmented. To quantify the performance of the method using these criteria, the overall sensitivity and specificity of

the algorithm were calculated:.

$$\text{Sensitivity} = \frac{\text{TP}}{\text{TP} + \text{FN}} \qquad \text{Specificity} = \frac{\text{TN}}{\text{TN} + \text{FP}} \tag{5.1}$$

The sensitivity is a measure of the ability of the method to identify the arteries, while the specificity is a measure of the ability of the method to separate the arteries from the veins.

The secondary goal was to minimize the error between the segmented vessel shape and the true vessel shape for each identified artery. Quantification of this error was also performed using the SS metric. The bias, magnitude, and extreme values of the truth-to-surface error were calculated using the error statistics given in Section 4.3.

5.3 Training and Optimization

Three CT scans were selected for algorithm training and parameter optimization. Three parameters were selected for optimization: cylinder height h_0, bifurcation detection threshold δ_{radius}, and the incremental step size Δ_{step}. Table 5.3 gives the values of parameter space that were explored for optimization.

The tracking and bifurcation detection algorithms were trained to maximize the number of arteries and minimize the number of veins segmented (i.e. maximize segmentation extent). To find the optimal set of parameters, a training score was constructed which gives the fraction of vessels correctly identified by the algorithm:

$$\text{Training Score} = \frac{\text{True Positives} + \text{True Negatives}}{\text{Total number of vessels marked}} \tag{5.2}$$

The optimal set of parameters which maximize this training score is given in Table 5.3.

Additionally, the sensitivity and specificity of the algorithm for each parameter set were computed. Figure 5.1 shows the effect on segmentation and training

performance when each of the three trained parameters is varied from the optimal value.

Table 5.3: Parameter sets used in training on 3 documented CT images.

Parameter	Range	Step size	Optimal value
h_0	10 mm - 30 mm	5 mm	15 mm
δ_{radius}	0.75 - 0.95	0.05	0.90
Δ_{step}	10% - 30% h_0	10%	20%h_0

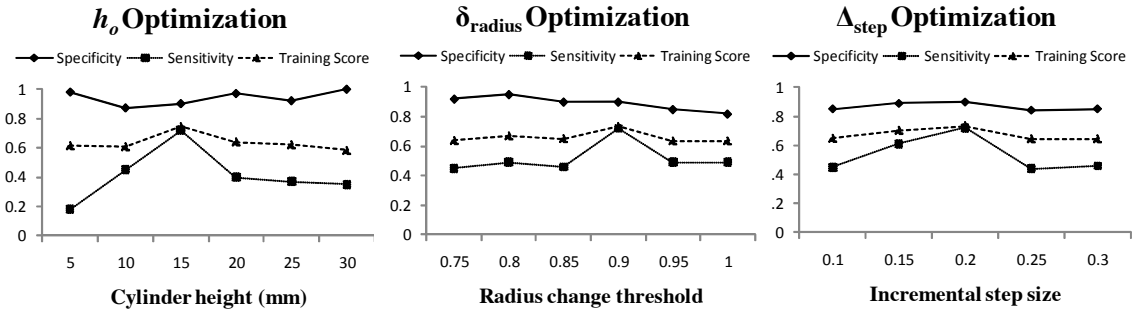

Figure 5.1: Training results for the cylinder height, bifurcation detection, and incremental step size parameters. Each plot shows the effect of varying one parameter while retaining the other two at their optimal values.

The parameter that had the largest effect on the outcome was the cylinder height h_0. Visual inspection of the training tracking results indicate that cylinders which are too short are sensitive to noisy regions and are unable to extract the vessel structure. Conversely, cylinders that were not sensitive enough to detect most bifurcations, reducing the overall sensitivity. Similarly, incremental step sizes above the optimal value of $0.20h_0$ were too coarse detect many bifurcations, decreasing the sensitivity.

The radius change threshold parameter δ_{radius} was optimized to a value of 0.90, which gave a marked improvement in algorithm performance over a much weaker value of 1.0. This indicates that limiting the search for child vessels to geometries

where the radius changes abruptly improves overall vessel identification performance.

5.4 Experiment design

The process diagram for the experiment is given in Figure 5.2. The primary purpose of the experiment was to evaluate the ability of the algorithm to segment the arterial tree while avoid venous segmentations. However, the SS metric also allows for analysis of truth-to-surface distance per vessel to evaluate the accuracy of the segmentation for identified vessels.

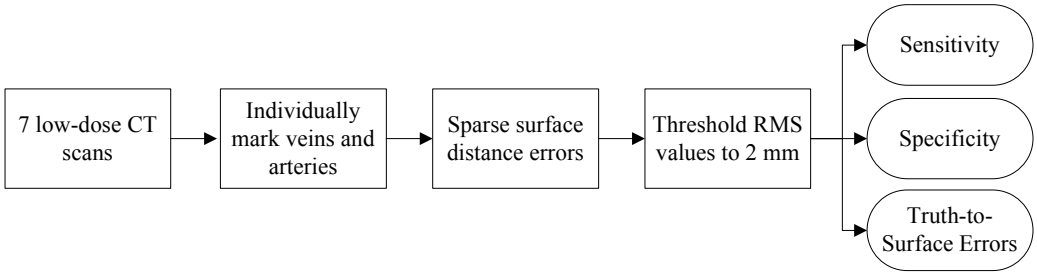

Figure 5.2: Process diagrams vessel tracking experiment

Sparse surface documentation of the veinous and arterial trees was performed on 7 cases in addition to the 3 training cases. Seed points in the left and right lungs were successfully automatically generated for 19/20 of the arterial trees across the ten cases. The SS metric was applied as outlined in Section 4.3 to compute the sensitivity and specificity of the algorithm. Additionally, the truth-to-surface distance errors were calculated for each vessel to obtain the overall distance error distribution for segmented vessels. Additionally, the maximum extent of over- and under-segmentation, the RMS error, and the overall bias (ASD) were calculated for the 19 segmented basal parent arteries.

CHAPTER 6

VESSEL TRACKING RESULTS AND DISCUSSION

This chapter presents the results of the pulmonary artery segmentation experiment (Section 5.4) performed on the ten documented cases. The ability of the method to identify arteries and avoid veins is evaluated first, followed by analysis of the truth-to-surface errors for the parent vessels. Additionally, this chapter discusses the main sources of error and provides a discussion of the tracking method.

6.1 Results

The sensitivity and specificity of the tracking method for the three training cases and seven test cases is shown in Figure 6.1. The average sensitivity across the seven test cases was 0.64, while the average specificity was markedly higher at 0.90. The average sensitivity across the three training cases was 0.60, with a specificity of 0.90.

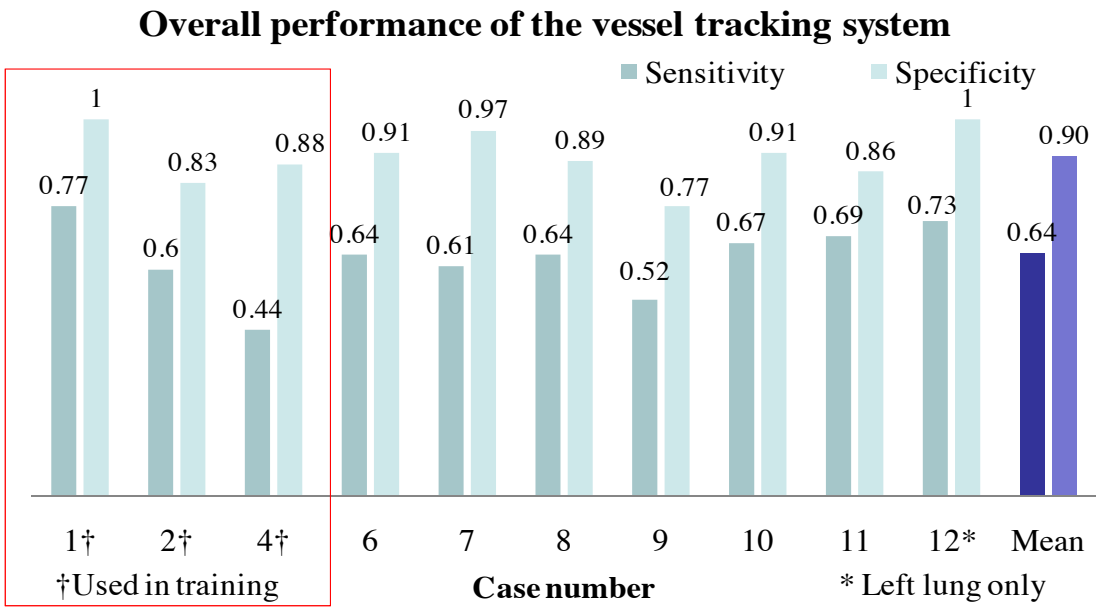

Figure 6.1: Sensitivity and specificity for the ten documented cases.

In addition to evaluating the specificity and sensitivity, the sparse surface metric provides for analysis of the segmentation error per vessel. As described in Chapter 4, four different statistics were calculated on the truth-to-surface errors for the segmented vessels. Of particular interest are the parent basal arteries where the seed points are detected. The plot in Figure 6.2 presents the average signed distance, root mean square distance, and distances of greatest over- and under-segmentation for the basal arteries across the ten documented cases.

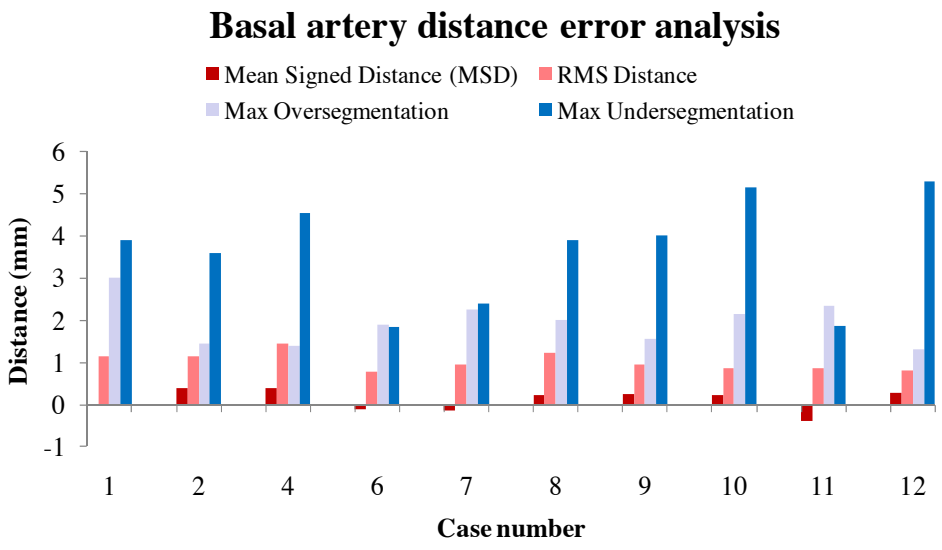

Figure 6.2: Results of the artery identification method for the LBPA and RBPA

The overall bias in the segmentation across the ten cases is 0.11 mm (approximately 1/4 of a voxel). This indicates that neither isotropic dilation nor erosion of the segmentation would improve results, as the bias is significantly less than one voxel. The mean RMS error per vessel is 1.01 mm (approximately 1 voxel). The SS metric also indicates that the greatest source of error arises from under-segmentation.

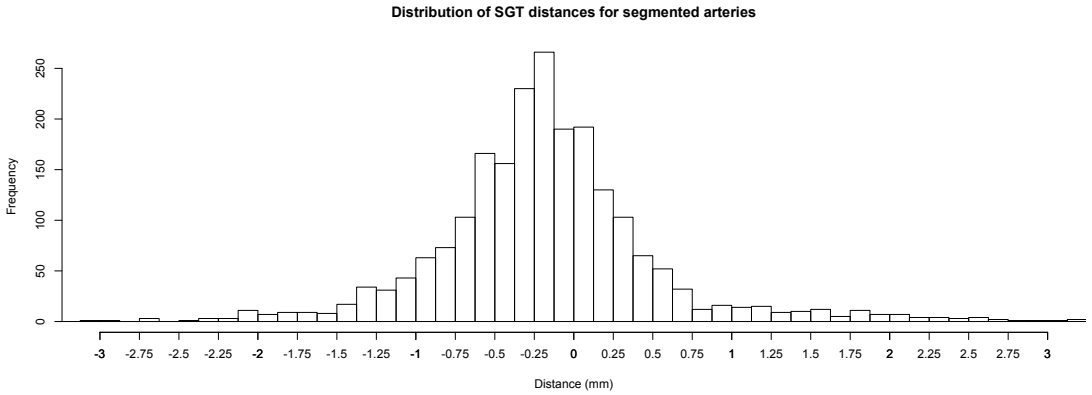

Figure 6.3: Distribution of truth-to-surface distance errors for correctly identified vessels using sparse markings across 10 cases.

The overall distribution of distance errors for segmented vessels is given in Figure 6.3. The average signed distance error was -0.15 mm, and the average RMS error was 0.63 mm.

6.2 Discussion

Visualizations of two arterial tree segmentations are given in Figure 6.4 along with the manual sparse truth markings for the arteries. The first segmentation shown provides an example of a case from the right lung with a high false positive and low true positive rate. In this case, a large child artery was missed, causing subsequent child branches to also be missed. Additionally, there is leaking onto the inferior pulmonary vein, and the separation of the arterial and venous trees is not fully maintained. The second segmentation represents a highly successful case from the left lung. The arterial tree is fully detected while maintaining full separation of the arterial and venous trees.

Figure 6.5 shows an example of the most significant source of error: leaking onto the inferior pulmonary vein. In both the right and left lungs, the large inferior

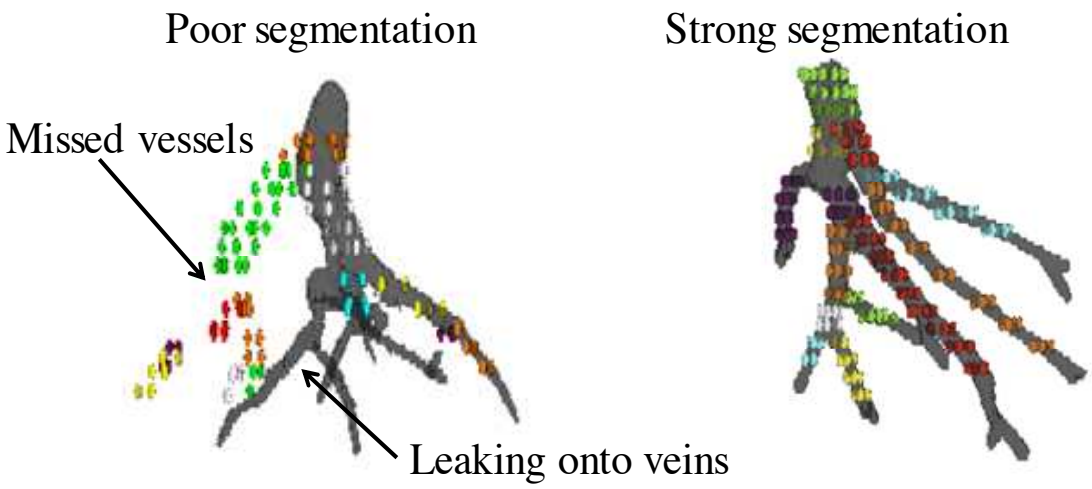

Figure 6.4: A poor segmentation in the right lung (case 9) and a strong segmentation in the left lung (case 10). Sparse surface truth markings are overlaid on the segmentation.

pulmonary veins runs directly posterior to the main arterial bifurcation. Errors at the inferior pulmonary vein are particularly, as they redirect the tracker onto a large parent vein, resulting in the segmentation of the child veins as well. The three frames in Figure 6.5 show the development of an inferior pulmonary vein leak as the tracker progresses into the lung parenchyma. In the first frame, the true bifurcation is missed, and the second child vessel is initiated on the inferior vein, as shown in the second frame. The result is a segmentation of child veins, as shown in the third frame.

The segmentation was performed in binary space, which reduced both the signal from the vessels and the noise inherent to low dose scans. The arteries were separated from the veins through vessel morphology and relative anatomical location alone. This approach is best suited for scans without contrast enhancement, where vessel edge information is lacking.

The vessel tracking algorithm was able to successfully segment vessels of greatly different radii. By fitting the size of the model cylinders to the vessel tree, the

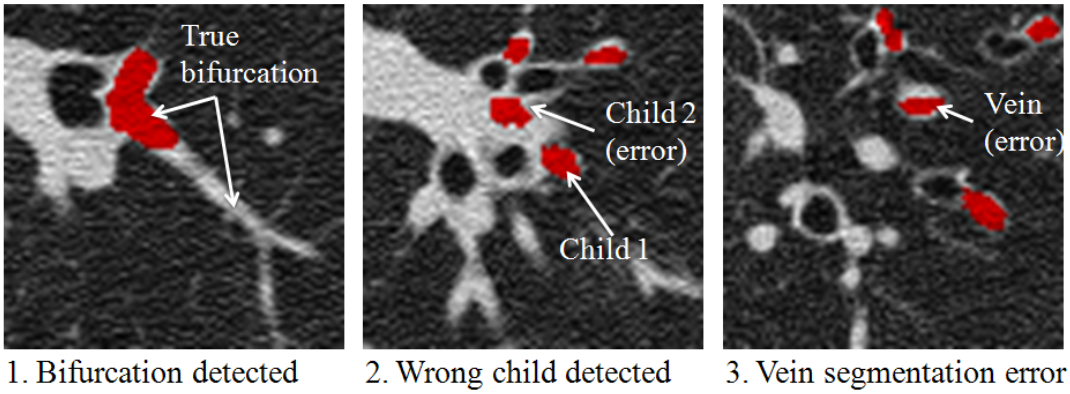

Figure 6.5: Three frames from a single CT image showing the progression of a leak onto the inferior pulmonary vein.

tracker was able to segment both the large parent arteries of the inferior lobes and the small subsegmental arteries deep in the lung parenchyma. However, as the vessel radii approached the smallest possible cylinder, the tracker was unable to detect bifurcations. The method was strongest in the larger vessels where Murray's law of vessel bifurcation is more closely followed. As vessel sizes decrease, blood flow becomes more non-Newtonian, and child vessel sizes become larger relative to the parent. Additionally, the smaller vessels deep in the lung parenchyma were more prone to obfuscation by noise because of their small size.

In most cases, the tracker was able to separate the arteries from the veins by searching for cylindrical objects in a limited size and orientation search space defined by the bifurcation model. While this approach was able to distinguish nearly identical vessels based only on their relationship to other arteries, it is unable to locate vessels that fall outside of the search space. Thus, arteries were missed if they originated at apparent trifurcations, or had unusually small radii relative to their sister branches. Relaxing the bifurcation model results in few missed segments, but more false positives (i.e. segmented veins).

The main source of error in vessel segmentation occurred at the left and right

inferior pulmonary veins. In cases where the arteries were highly conflated with the veins and airway walls, due either to high noise or unusual anatomy, the vessel tracker leaked onto the pulmonary veins and tracked them into the parenchyma. One possible method for reducing the false positive rate would be to initiate a vessel tracker at the main inferior pulmonary arteries, and exclude from the arterial segmentation vessels that were highly likely to belong to the venous tree.

CHAPTER 7

CONCLUSION

Pulmonary artery segmentation has been suggested as a method to improve diagnosis of both pulmonary embolism and pulmonary arterial hypertension. Although much work has focused on segmenting the arteries and veins as one structure, very little previous work has been performed to distinguish the two trees. In this work, a method was developed for automatically segmenting the pulmonary arteries in low-dose CT, and was applied to the basal pulmonary arteries and their child vessels.

7.1 Contributions

The method in this thesis is the first to automatically separate the arteries from the veins in thoracic CT. Automated segmentations eliminate inter-user variability, and provide the possibility of obtaining novel and quantitative anatomical information. Additionally, the presented method is the first to use an airway segmentation to eliminate airway walls prior to vascular segmentation.

A cylinder tracker was used to perform the entire segmentation, which allows vessels to be tracked individually to separate them from fused structures. Vessels were tracked based on anatomical and morphological information alone, as edge information between soft tissue structures is highly suppressed in low-dose CT. Child vessels were identified by comparing the local vessel morphology to a bifurcation model. Child vessels were searched for only at points where the parent vessel underwent a large change in radius, suggesting a bifurcation point. Separation of the arteries and veins was frequently maintained by excluding child vessels at highly unlikely angles from the parent and other child vessel.

A novel method for automatically locating large parent arteries was developed

and tested on the basal arteries. An automated airway segmentation was used to provide a reference point for locating the basal arteries, and a region of interest was constructed around the basal airways to limit the search space for the arteries. A method for extracting the location and orientation of the basal artery was used to automatically seed the vessel tracker for segmenting the left and right lower lobes.

Validation of complex structures is a challenging task, and a novel surface-based ground truth marking system and validation metric were defined for use with complex structures. Sparse surface markings capture significant vessel detail while require fewer markings that traditional dense surface markings. This allows for validation on a greater number of cases given the same amount of effort. The sparse surface markings were shown to be an accurate sampling of the true surface, justifying the use of spares markings over dense markings.

The artery identification method was validated on 50 cases (100 basal arteries) with a success rate of 60% overall, and 80% when ROIs were provided from the airway segmentation. The tracking algorithm achieved an overall sensitivity rate of 0.64 and specificity rate of 0.90 when applied to 10 cases with correctly labeled airways. Although the success rate is lower than the previously reported method by Saha, the presented method has the advantage that it is completely automated, and works in more noisy low-dose CT scans.

7.2 Future Work

Automated detection of parent vessels in the left and right upper lobes is the next appropriate step to take in furthering this work. In this work, only the largest arteries to enter the lung parenchyma were identified. Such a method could employ a similar ROI construction or use another algorithm to locate the

upper lobe arteries.

The sensitivity of the algorithm could be improved by developing a trifurcation detector, or developing a method for identifying small subsegmental arteries missed by the tracker. Additionally, the method could be paired with a global Hessian matrix filter to amplify the signal from vessels and reduce tracking error.

Finally, the vessel tracker presented in this thesis could also be tested on the pulmonary veins in low-dose CT. Additionally, the tracker could be used to segment branching biological networks encountered in alternative imaging modalities, including standard-dose CT and magnetic resonance imaging.

BIBLIOGRAPHY

[1] David Anderson, Susan Kahn, and et al. Marc Rodger. Computed tomographic pulmonary angiography vs ventilation-perfusion lung scanning in patients with suspected pulmonary embolism: a randomized controlled trial. *JAMA*, 298:2743–2753, 1984.

[2] Thorsten Buzug. *Computed Tomography: From photon statistics to modern cone-beam CT*. Springer, 2008.

[3] P Croisille, M Souto, M Cova, S Wood, Y Afework, J E Kuhlman, and E A Zerhouni. Pulmonary nodules: improved detection with vascular segmentation and extraction with spiral CT. *Radiology*, 197:397–401, 1995.

[4] B Flehinger, M Kimmel, and M Melamed. The effect of surgical treatment on survival from early lung cancer: implication for screening. *Chest*, 101:1013–1018, 1992.

[5] C. Henschke and D. Yankelevitz. CT screening for lung cancer: update 2007. *Oncologist*, 13:65–78, 2008.

[6] Godfrey Hounsfield. Computerized transverse axial scanning (tomography): Part 1. description of system. *British Journal of Radiology*, 46:1016–1022, 1973.

[7] The International Early Lung Cancer Action Program Investigators. Survival of patients with stage I lung cancer detected on CT screening. *N Engl J Med*, 355:1763–1771, 2006.

[8] Chevalier Jackson and John Huber. Correlated applied anatomy of the bronchial tree and lungs with a system of nomenclature. *Chest*, 9:319–326, 1943.

[9] A. Jemal, R. Siegel, J. Xu, and E. Ward. Cancer Statistics, 2010. *Cancer J Clin*, 60:277–300, 2010.

[10] Jens Kaftan and Atilla Kiraly an Annemarie Bakai. Fuzzy pumonary vessel segmentation in contrast enhanced CT data. *Medical Imaging 2008: Image Proceedings*, 6914:69151Q–1–12, 2008.

[11] Jens N. Kaftan, Annemarie Bakai, Marco Das, and Til Aach. Locally adaptive fuzzy pulmonary vessel segmentation in contrast enhanced CT. *Proceedings*

5th international symposium on biomedical imaging: from nano to macro, pages 101–104, 2008.

[12] Jaesun Lee and Anthony Reeves. Segmentation of the airway tree from chest CT using local volume of interest. *Second International Workshop of Pulmonary Image Analysis*, pages 353–364, 2009.

[13] Tianhu Lei, Jayaram Udupa, Punam Saha, and Dewey Odhner. Artery-vein separation via MRA - An image processing approach. *IEEE Transaction on Medical Imaging*, 20:689–703, 2001.

[14] MG Linguraru, JA Pura, RL Van Uitert, N Mukherjee, RM Summers, C Minniti, MT Gladwin, G Kato, RF Machado, and RJ Wood. Segmentation and quantification of pulmonary artery for noninvasive CT assessment of sickle cell secondary pulmonary hypertension. *Medical Physics*, 37:1522–1532, 2010.

[15] Yoshitaka Masutani, Heber MacMahon, and Kunio Doi. Automated segementation and visualization of the pulmonary vascular tree in spiral CT angiography: an anatomy-oriented approach based on three-dimensional image analysis. *Journal of Computer Assisted Tomography*, 25:587–597, 2001.

[16] Yoshitaka Masutani, Hever MacMahon, and Kunio Doi. Computerized detection of pulmonary embolism in spiral CT angiography based on volumetric image analysis. *IEEE Transaction on Medical Imaging*, 21:1517–1523, 2002.

[17] D Montani, L.C. Price, B. Girerd, T. Chinet, P. Lacombe, G. Simonneau, and M. Humbert. Fatal rupture of pulmonary arteriovenous malformation in hereditary haemorrhagic telangiectasis and severe PAH. *European Respiratory Review*, 18:42–46, 2009.

[18] David Naidich, Christopher Marshall, Christopher Gribbin, Ronal Arams, and Dorothy McCauley. Low-Dose CT of the lungs: preliminary observations. *Radiology*, 175:729–731, 1990.

[19] A. Reeves and W. Kostis. Computer-aided diagnosis of small pulmonary nodules. *Semin Ultrasound CT MR*, 21:116–128, 2000.

[20] Punam Saha, Zhiyun Gao, Sara Alford, and Milan Sonka. Topomorphological separation of fused isointensity objects via multiscale opening: separating arteries and veins in 3-D pulmonary CT. *IEEE Transaction on Medical Imaging*, 29:840–851, 2010.

[21] E Segre. *Experimental Nuclear Physics*, volume 1. Thieme, 1953.

[22] Ömer Senbaklavaci, Yukihiro Kaneko, Anna Bartunek, Christof Bruner, Erol Kurkciyan, Patrick Wunderbaldinger, Walter Klepetko, Ernst Wolner, and Werner Mohl. Rupture and dissection in pulmonary artery aneurysms: Incidence, cause and treatment - Review and case report. *Journal of Thoracic Cardiovascular Sugery*, 121:1006–1008, 2001.

[23] HIdenori Shikata, Geoffrey McLennan, Eric Hoffman, and Milan Sonka. Segmentation of pulmonary vascular trees from thoracic 3D CT images. *International Journal of Biomedical Imaging*, 2009, 2009.

[24] Elias Zerhouni Susan Wood, John Hoford, Eric Hoffman, and Wayne Mitzner. Measurement of three-dimension lung tree structures by using computed tomography. *Journal of Applied Physiology*, 79:1687–1697, 1995.

[25] S. Teague, G. Trilikis, and E. Dharaiya. Lung nodule computer-aided detection as a second reader: influence on radiology residents. *Journal of Computer Assisted Tomography*, 34:35–39, 2010.

[26] Cornelis van Bemmel, Luuk Spreeuwers, Max Viergever, and Wiro Niessen. Level-set-based artery-vein separation in blood pool agent CE-MR angiograms. *IEEE Transactions on Medical Imaging*, 22:1224–1234, 2003.

[27] Chuan Zhou, Heang-Ping Chan, Smita Patel, Philip Cascade, Berkman Sahiner, Lubomir Hadjiski, and Ella Kazerooni. Preliminary investigation of computer-aided detection of pulmonary embolism in three-dimensional computed tomography pulmonary angiography images. *Academic Radiology*, 12:782–792, 2005.

[28] Chuan Zhou, Heang-Ping Chan, Berkman Sahiner, Lubomire Hadjiski, and Aamer Chughtai. Automatic multiscale enhancement and segmentation of pulmonary vessels in CT pulmonary angiography images for CAD applications. *Medical Physics*, 34:4567–4577, 2007.

[29] Jinghao Zhou, Sukmoon Chang, Dimitris Metaxas, and Leon Axel. Vascular structure segmentation and bifurcation detection. *IEEE International Symposium on Biomedical Imaging*, pages 872–875, 2007.