



# VCU

Virginia Commonwealth University  
VCU Scholars Compass

---

Theses and Dissertations

Graduate School

---

2010

## Segmentation and Fracture Detection in X-ray images for Traumatic Pelvic Injury

Rebecca Smith  
*Virginia Commonwealth University*

Follow this and additional works at: <https://scholarscompass.vcu.edu/etd>



Part of the [Computer Sciences Commons](#)

© The Author

---

Downloaded from

<https://scholarscompass.vcu.edu/etd/76>

This Dissertation is brought to you for free and open access by the Graduate School at VCU Scholars Compass. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of VCU Scholars Compass. For more information, please contact [libcompass@vcu.edu](mailto:libcompass@vcu.edu).

School of Engineering  
Virginia Commonwealth University

This is to certify that the dissertation prepared by Rebecca Smith entitled SEGMENTATION AND FRACTURE DETECTION IN X-RAY IMAGES FOR TRAUMATIC PELVIC INJURY has been approved by her committee as satisfactory completion of the Dissertation requirement for the degree of Doctor of Philosophy

---

Kayvan Najarian, Ph.D., Committee Chair, Department of Computer Science

---

Krzysztof J. Cios, Ph.D., Chair of Computer Science, School of Engineering

---

Vojislav Kecman, Ph.D., Dept. of Computer Science, School of Engineering

---

Rosalyn S. Hobson, Ph.D., School of Engineering

---

Toan Huynh, M.D., Carolinas Health System

---

Kevin R. Ward, M.D., School of Medicine

---

Rosalyn S. Hobson, Associate Dean of Graduate Studies, School of Engineering

---

Russell D. Jamison, Ph.D., Dean, School of Engineering

---

F. Douglas Boudinot, Ph.D., Dean of the School of Graduate Studies

---

Date

© Rebecca Smith, 2010

All Rights Reserved

SEGMENTATION AND FRACTURE DETECTION IN X-RAY IMAGES FOR  
TRAUMATIC PELVIC INJURY

A dissertation submitted in partial fulfillment of the requirements for the degree of  
Doctor of Philosophy at Virginia Commonwealth University.

by

REBECCA SMITH

B.Eng., Imperial College (London, United Kingdom), 2003

Director: KAYVAN NAJARIAN

ASSOCIATE PROFESSOR, DEPARTMENT OF COMPUTER SCIENCE

Virginia Commonwealth University

Richmond, Virginia

April, 2010

## Acknowledgments

First, I thank my adviser, Dr. Kayvan Najarian, for his guidance and encouragement in this research project, which could not have been completed without his support. I am also grateful to all my committee members for their valuable feedback on my work. They are: Dr. Krzysztof Cios, Dr. Rosalyn Hobson and Dr. Vojislav Kecman in the VCU School of Engineering, Dr. Kevin Ward in the MCV Department of Emergency Medicine, and Dr. Toan Huynh at the Carolinas Medical Center. Special thanks go to Dr. Charles Cockrell in the MCV Department of Radiology, who has helped greatly with expert data analysis and suggestions on how to further develop this project.

To all the members of the VCU Biomedical Signal and Image Processing Lab: many thanks for making the route to a Ph.D. a lot more fun (even with all the writing). In particular, Wenan Chen and Simina Vasilache have worked with me on this project and others (and so deserve medals for their patience), while Dr. Soo-Yeon Ji has offered support and advice since the earliest stages of my Ph.D. studies.

Most of all, I thank Scott Smith - who kept things in perspective, kept trying to fix my typos, and kept the many hours of work worth it.

## Table of Contents

|  | Page |
|--|------|
| Acknowledgments . . . . .  | ii   |
| List of Tables . . . . .   | viii |
| List of Figures . . . . .  | ix   |
| Abstract . . . . .   | xiii |
| Executive Summary and Contributions . . . . .                      | xv   |
| <br>Chapter  |      |
| 1 Introduction . . . . .   | 1    |
| 1.1 Traumatic Pelvic Injury and Decision Support Systems . . . . . | 1    |
| 1.2 The Pelvic Structure . . . . .                                 | 4    |
| 1.3 Motivation Behind Study of X-Ray Images . . . . .              | 5    |
| 1.4 Significance of this Study . . . . .                           | 9    |
| 1.5 Aims of this Study . . . . .                                   | 10   |
| 2 Related Work and Approach . . . . .                              | 13   |
| 2.1 Overview . . . . .   | 13   |
| 2.2 Computer-Aided Decision Making Systems . . . . .               | 13   |
| 2.3 Medical Image Segmentation . . . . .                           | 16   |
| 2.3.1 General Approaches to Segmentation . . . . .                 | 17   |
| 2.3.2 Segmentation of X-ray Radiographs . . . . .                  | 18   |
| 2.4 Active Shape Model . . . . .                                   | 22   |

|       |   |    |
|-------|---|----|
| 2.4.1 | Shape Model . . . . .   | 23 |
| 2.4.2 | Grey-level Model . . . . .  | 25 |
| 2.5   | Approach of this study . . . . .                                      | 28 |
| 3     | Combined Spline/ASM Algorithm . . . . .                               | 32 |
| 3.1   | Overview . . . . .  | 32 |
| 3.2   | The Problems of ASM . . . . .   | 32 |
| 3.3   | Splines . . . . .   | 35 |
| 3.3.1 | Splines in Medical Image Segmentation . . . . .                       | 35 |
| 3.3.2 | Cubic Spline Interpolation . . . . .                                  | 36 |
| 3.3.3 | B-splines . . . . .   | 37 |
| 3.4   | Method . . . . .  | 40 |
| 3.4.1 | Spline interpolation in ASM deformation . . . . .                     | 40 |
| 3.4.2 | Considerations during Training . . . . .                              | 46 |
| 3.5   | Summary . . . . .   | 47 |
| 4     | Hierarchical Approach to Automatic ASM Initialization . . . . .       | 48 |
| 4.1   | Overview . . . . .  | 48 |
| 4.2   | Objective . . . . .   | 48 |
| 4.3   | Method . . . . .  | 50 |
| 4.3.1 | Detection of Femoral Shafts . . . . .                                 | 50 |
| 4.3.2 | Detection of Femoral Heads . . . . .                                  | 60 |
| 4.3.3 | Spline/ASM for Femur Detection . . . . .                              | 64 |
| 4.3.4 | Spline/ASM for Pelvic Ring Detection . . . . .                        | 65 |
| 4.3.5 | Spline/ASM for Pubis and Ischium Detection . . . . .                  | 66 |
| 4.3.6 | Hough Transform and Spline/ASM for Detection of Iliac Bones . . . . . | 66 |
| 4.4   | Summary . . . . .   | 69 |
| 5     | Fracture Detection and Displacement Measurement . . . . .             | 70 |

|       |   |     |
|-------|---|-----|
| 5.1   | Overview . . . . .                            | 70  |
| 5.2   | The Discrete Wavelet Transform . . . . .      | 70  |
| 5.3   | Ring Fracture Detection . . . . .             | 74  |
| 5.3.1 | Windowing . . . . .                           | 75  |
| 5.3.2 | Wavelet Transform . . . . .                   | 78  |
| 5.3.3 | Masking and boundary tracing . . . . .        | 80  |
| 5.4   | Symphysis pubis displacement . . . . .        | 82  |
| 5.5   | Patterns of Pelvic Injury . . . . .           | 86  |
| 5.6   | Summary . . . . .                             | 88  |
| 6     | Results of Image Processing . . . . .         | 89  |
| 6.1   | Overview . . . . .                            | 89  |
| 6.2   | Segmentation . . . . .                        | 89  |
| 6.3   | Fracture Detection . . . . .                  | 98  |
| 6.4   | Complexity and Processing Time . . . . .      | 99  |
| 6.5   | Testing on Alternative Images . . . . .       | 99  |
| 6.6   | Summary and Discussion . . . . .              | 106 |
| 7     | Image Features in Predictive Models . . . . . | 109 |
| 7.1   | Overview . . . . .                            | 109 |
| 7.2   | Decision Tree Algorithms . . . . .            | 109 |
| 7.2.1 | CART . . . . .                                | 110 |
| 7.2.2 | C4.5 . . . . .                                | 111 |
| 7.3   | Experiment 1: X-Ray Features . . . . .        | 112 |
| 7.3.1 | Experiment Overview . . . . .                 | 112 |
| 7.3.2 | Dataset . . . . .                             | 112 |
| 7.3.3 | Results . . . . .                             | 112 |
| 7.4   | Experiment 2: X-Ray and CT Features . . . . . | 115 |
| 7.4.1 | Experiment Overview . . . . .                 | 115 |



|       |                             |     |
|-------|-----------------------------|-----|
| 7.4.2 | Dataset                     | 116 |
| 7.4.3 | Results                     | 116 |
| 7.5   | Summary                     | 118 |
| 8     | Conclusions and Future Work | 120 |
| 8.1   | Conclusions                 | 120 |
| 8.2   | Future Work                 | 121 |
|       | Bibliography                | 123 |

## List of Tables

| Table Number   | Page |
|--|------|
| 6.1 Results of full segmentation testing via visual inspection. . . . .  | 98   |
| 6.2 Preliminary results of abnormality detection via visual inspection. . . . .  | 99   |
| 7.1 Classification results for Experiment 1 using leave-one-out. SVM and C4.5 testing are performed for validation purposes only, using the features selected from the dominant CART tree. . . . . | 114  |
| 7.2 Classification results for Experiment 2 using 10FCV . . . . .  | 119  |

## List of Figures

| Figure Number  | Page |
|--|------|
| 1.1 A brief overview of relevant pelvic structures. . . . .  | 5    |
| 2.1 The ASM search process calculates the fitness function $F(g_s)$ for each of a set of pixels along the profile to the model point. After multiple iterations, $F(g_s)$ should be minimal at the edge of the target structure. . . . . | 27   |
| 2.2 A schematic diagram of the overall X-ray analysis process. . . . .   | 31   |
| 3.1 Two examples of the differences in pelvic radiograph image characteristics.  | 33   |
| 3.2 Severe distortion of pubis-ischium structure seen in Standard ASM segmentation. . . . .  | 34   |
| 3.3 An example of standard spline interpolation failing on a circular structure.   | 39   |
| 3.4 A 3D plot of the monotonic, strictly increasing function of $(x, y, t)$ generated using chord-length parameterization. . . . .   | 45   |
| 3.5 A spline parameterized on chord-length deals well with curves that ‘double back’. . . . .  | 46   |
| 4.1 The effects of poor initialization on pelvic ring detection via standard ASM.  | 50   |
| 4.2 A schematic diagram of the Hierarchical Automatic ASM initialization process. . . . .  | 51   |
| 4.3 The multi-step process used in femoral shaft detection. . . . .  | 53   |
| 4.4 Images of the intermediate stages in femoral shaft detection, such as extraction of candidate shaft edges and pair-matching for candidate shafts..   | 57   |
| 4.5 The result of applying Canny edge detection to the area of the image containing the left femoral head, then applying Hough Transform to the edge-detected image. . . . .   | 63   |
| 4.6 The result of Spline/ASM detection of the right femur . . . . .  | 65   |
| 4.7 An example of automatic detection of the approximate location of the ilium, via edge detection and circular Hough Transform. . . . .   | 68   |

|      |  |     |
|------|--|-----|
| 5.1  | A single 2D DWT decomposition stage. A low-pass filter $h[n]$ and high-pass filter $g[n]$ are first applied to the rows of the image, then the columns of the output. Down-sampling by 2 occurs after each step. . . . . | 75  |
| 5.2  | A schematic diagram of the pelvic ring fracture detection process. . . . .   | 76  |
| 5.3  | Example overlapping windows around the pelvic ring, positioned according to model points placed during Spline/ASM segmentation. . . . .  | 78  |
| 5.4  | Reconstruction using the SWT approximation coefficients generates a smoothed window image. . . . .   | 79  |
| 5.5  | Reconstructions of a ring window from vertical, horizontal, and diagonal SWT detail coefficients . . . . .   | 79  |
| 5.6  | The filtering process used to create the edge window, $W_E$ . . . . .  | 80  |
| 5.7  | Example of a detected broken ring boundary, which may indicate a ring fracture. . . . .  | 81  |
| 5.8  | Calculating the width of the pubis symphysis gap using Spline/ASM landmarks. . . . .   | 83  |
| 5.9  | Calculating the vertical displacement of the left and right pubis using Spline/ASM landmarks. . . . .  | 84  |
| 6.1  | ANOVA box-plot for detection of right iliac bone. . . . .  | 91  |
| 6.2  | ANOVA box-plot for detection of right femur. . . . .   | 92  |
| 6.3  | ANOVA box-plot for detection of pelvic ring. . . . .   | 93  |
| 6.4  | An example of successful full automatic segmentation using Spline/ASM. . . . .   | 94  |
| 6.5  | A second example of successful full automatic segmentation using Spline/ASM. . . . .   | 95  |
| 6.6  | A example of only partially successful full automatic segmentation using Spline/ASM, showing incorrect detection of the left pubis bone. . . . .   | 96  |
| 6.7  | Comparing Standard ASM with Combined Spline/ASM in detection of pelvic ring and right iliac bone. . . . .  | 97  |
| 6.8  | Multiple results of fracture detection. . . . .  | 100 |
| 6.9  | An example of left cardiac ventricle detection in a MR image using the Spline/ASM algorithm. . . . .   | 102 |
| 6.10 | An example of left cardiac ventricle detection in a MR image using the standard ASM algorithm. . . . .   | 103 |

|      |  |     |
|------|--|-----|
| 6.11 | A second example of left cardiac ventricle detection in a MR image using the Spline/ASM algorithm. . . . .   | 104 |
| 6.12 | A second example of left cardiac ventricle detection in a MR image using the standard ASM algorithm. . . . .   | 105 |
| 7.1  | The dominant tree extracted by CART in injury severity prediction for Experiment 1. . . . .  | 113 |
| 7.2  | The decision tree output by the C4.5 algorithm in Experiment 2. 'PF' indicates presence of pubis fracture, and 'ICU' and 'W/H' (Ward/Home) are the two outcomes. . . . . | 118 |

## **Abstract**

### **SEGMENTATION AND FRACTURE DETECTION IN X-RAY IMAGES FOR TRAUMATIC PELVIC INJURY**

By Rebecca Smith, Ph.D.

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy at Virginia Commonwealth University.

Virginia Commonwealth University, 2010

Major Director: Kayvan Najarian  
Associate Professor, Department of Computer Science

Due to the risk of complications such as hemorrhage, severe pelvic trauma is associated with a high mortality rate. Prompt medical treatment is therefore vital. However, the complexity of the injuries can make successful diagnosis and treatment challenging. By generating predictions and recommendations based on patient data, computer-aided decision support systems have the potential to assist physicians in improving outcomes. However, no current system considers features automatically extracted from medical images. This dissertation describes a system to extract diagnostic features from pelvic X-ray images that can be used as input to the prediction process; specifically, the presence of fracture and quantitative measures of displacement.

Feature extraction requires prior identification of separate structures of interest within the pelvis. The proposed system therefore incorporates a hierarchical segmentation algorithm which is

able to automatically extract multiple structures in a single pass, using a combination of anatomical knowledge and computational techniques such as directed Hough Transform. This algorithm also applies a novel Spline/ASM segmentation method which combines cubic spline interpolation with a deformable model approach which maintains curved contours and provides local control over segmentation. In order for the proposed system to be used as a component in a computerized decision support system, segmentation is designed to be entirely automatic. Furthermore, Spline/ASM is suitable for many other segmentation applications where the objects of interest show curved contours. After successful segmentation, fracture detection is performed on the pelvic ring and pubis structures, using an algorithm based on wavelet transform, anatomical information and boundary tracing. A method is also developed to calculate quantitative measures of symphysis pubis displacement that may indicate pelvic instability and prove useful in identifying fracture patterns. Finally, X-ray features are combined with patient demographics and physiological scores for generation of predictive rules for injury severity, with promising current results. This indicates the potential diagnostic value of the extracted features, and in turn the usefulness of the proposed radiograph analysis component in a larger decision support system.

## Executive Summary and Contributions

Due to its strong association with traffic accidents, pelvic trauma affects a relatively young age range. It is also challenging to treat, and poses both a high mortality rate and the risk of life-long disability among survivors. Complications such as hemorrhage can quickly prove fatal, making early stabilization of patients vital in improving outcomes. Computer-aided decision support systems can assist surgeons and physicians in this task by generating treatment recommendations and outcome predictions, based on rapid extraction of key features from the large volume of data available for an individual patient. However, current decision-making systems for trauma are limited in scope and the variables they consider. Crucially, they do not consider the diagnostic importance of medical imaging. This dissertation describes a framework for automated processing of pelvic X-ray images which includes:

1. A novel extension to the Active Shape Model algorithm commonly used in radiograph processing, which incorporates the use of controlled spline interpolation to maintain curved contours and emphasize structure integrity. This is explained in Chapter 3.
2. A hierarchical initialization algorithm that performs automatic segmentation of multiple pelvic structures in a single pass, using the output from one step as input to the next. This is explained in Chapter 3.4
3. A method to detect fractures of the pelvic ring and horizontal fractures of the pubis and ischium using windowing and the Stationary Discrete Wavelet Transform. Details are provided in



## Chapter 5.

4. A method to automatically calculate quantitative measures of displacement based on the results of structure segmentation, which provides useful input to model generation and has potential use in identification of injury patterns. This is also explained in Chapter 5.
5. A rule-generation method that combines the extracted image features with other patient information such as demographics and physiological measures to predict injury severity. More details can be found in Chapter 7.

The proposed system is designed as a key component of a larger decision-making system for trauma, allowing rapid extraction of image features that may be missed by visual inspection or that may prove time-consuming to manually calculate, and representing them in a form suitable for use in predictive modeling. However, it has potential use as a system in its own right, particularly in remote areas where advanced resources and equipment may not be available. Furthermore, the framework of fracture detection can be applied to other body regions, and potentially adapted for industrial applications where X-ray images are widely used. Finally, the developed combined Spline/ASM algorithm has value beyond the field of medical image segmentation, as it can be applied to detection of any object with curved contours.

## CHAPTER 1 Introduction

### *1.1 Traumatic Pelvic Injury and Decision Support Systems*

Fractures of the pelvis are among the most severe injuries that can be suffered by a major trauma patient, with the most severe cases typically occurring in the context of high-energy impact injuries such as those incurred in motor vehicle collisions (MVCs) [5]. They are therefore associated with a younger demographic and polytrauma (i.e. traumatic injuries to multiple areas of the body). Various studies have estimated that 50–60% of pelvic trauma cases are due to MVCs, rising to 65–75% if car-pedestrian accidents are included [65]. High-energy trauma is becoming increasingly common; in 1986, one study reported that over 300 patients with pelvic fracture were seen in a three year period at the University of Maryland Trauma Shock Center [74], and this number has doubtlessly risen in the ensuing years.

Since pelvic fracture can lead to laceration of the surrounding organs and vascular structures, patients are at high risk of complications, which are responsible for the majority of deaths in pelvic trauma (as opposed to the original injury itself) [23]. The most common complication is hemorrhage. Patients with hemorrhagic shock have an estimated mortality rate of between 36 – 54% versus the 5.6 – 9.5% associated with general pelvic fracture [65]. In such cases, the patient may die from the hemorrhage itself, via exsanguination and shock, or due to other conditions resulting from hemorrhage such as severe infection [6]. More specifically, one survey of multiple studies found that mortality among hemodynamically unstable patients with closed

pelvic fractures stands at an average of 27%, increasing to an average of 55% for those with open pelvic fractures [27]. These deaths typically occur within the first 24 hours of injury, making early identification of at-risk patients crucial in improving outcomes [17]. Example predictors include physiological response to resuscitation, pelvic hematoma as seen on computed tomography (CT) scans, and pelvic fracture patterns presented in X-ray radiographs [23, 45]. Prompt recognition of these patterns has been shown to allow early and effective stabilization [12].

However, due to the complexity of the injuries and the many variables involved, as well as the high pace and pressure in a trauma center environment, quick and accurate diagnosis and treatment presents a challenge to physicians. This has led to considerable interest in the development of computer-assisted decision making systems - or decision support systems - which can analyze large volumes of patient information in a very short time to generate recommendations for treatment and predictions of patient outcome.

As well as potentially improving outcomes and resource allocation, these systems may also reduce the overall cost of care, which is severely impacted by inaccurate diagnoses. This was demonstrated in a study of blunt pelvic trauma, which found routine radiography to be a cost-effective method of detecting fractures that were not associated with significant clinical findings and might otherwise have remained undiagnosed [33]. Missed injuries are common in major trauma; one study conducted at a Danish level 1 trauma center found that over a three year period, 64 out of 786 patients had a total of 86 undiagnosed injuries, for an average rate of 1.3 missed injuries per patient [28]. Furthermore, injuries went undetected at multiple stages of trauma management. A decision support system capable of automatically detecting potential presence of injury and creating recommendations for treatment may be of great potential use in reducing the prevalence of these missed injuries in trauma patients.

Existing examples of computer-aided decision-making systems are described in Section 2.2, but none possess all of the following key abilities:

1. Use of *all available* patient data to generate and dynamically update predictions and recommendations at every stage of care.
2. Presentation of the reasoning behind these recommendations in an easily-understood form.
3. Incorporation of information automatically extracted from medical images.

The work in this dissertation addresses item 3, as a component used in development of a general trauma decision-support system. Prior work on this project has focused on the use of common measurements such as GCS, ISS and demographic information in prediction of patient outcome using a transparent rule generation approach [69, 31]. Although the resulting rule-bases showed high accuracy, it is vital that information from medical images be incorporated. Medical imaging is a key diagnostic tool for clinicians and surgeons and is used in a wide variety of applications including diagnosis of injury, tracking the progress of disease, planning surgical procedures, and evaluating recovery. Information that may prove of potential diagnostic and therapeutic worth in training predictive rule-bases can be extracted in quantitative form using computational methods. However, if a decision support system is to be useful in a trauma center environment, it must require only a minimal level of user input. For this reason, feature extraction must be automated as much as possible. This dissertation focuses specifically on the development of a module that analyzes and processes pelvic radiographs to extract a number of potentially useful image features, in a form that can be combined with other patient information and used as input in predictive model generation. The final system is intended to be deployed during initial triage, but will con-

tinue to be used to assess patient condition throughout every stage of care - generating updated recommendations and predictions as the patient's state changes and more data becomes available.

### **1.2 The Pelvic Structure**

The bony pelvis comprises the midline posterior sacrum and two lateral bones, each consisting of the ilium, ischium and pubis. These two halves are joined frontally by the symphysis pubis and sacroiliac ligaments to form a stable ring structure. Damage to these ligaments incurred in pelvic trauma may be associated with other severe injuries, such as those to the internal iliac artery - and even more importantly, it may indicate pelvic instability [60, 11]. For example, in cases where the anterior symphysis pubis ligaments are disrupted, diastasis of 2.5 can occur only up to 2.5cm before stability is lost, provided the posterior ligaments remain intact [66]. Furthermore, a healthy pelvis should be approximately symmetric, and asymmetry may be considered a mark of severe injury [10]. Figure 1.1 provides a labeled diagram of some pelvic structures relevant to this work. This is a very brief and general overview of pelvic anatomy but is useful in understanding the major tasks this dissertation explores.

### **1.3 Motivation Behind Study of X-Ray Images**

Plain film X-ray imaging is the oldest of the imaging modalities used in modern medicine and arguably the most prevalent, with seven out of ten US citizens having received at least one X-ray examination during their lifetime [32]. The resulting images provide a quick, low-cost means of screening and diagnosis, particularly in cases of fracture. Over the last forty years, radiography has expanded from the analysis of simple plain film images to include more sophisticated modalities such as computed tomography (CT) and positron emission tomography (PET). These offer improved image quality and definition, though at the cost of some speed and ease of collection

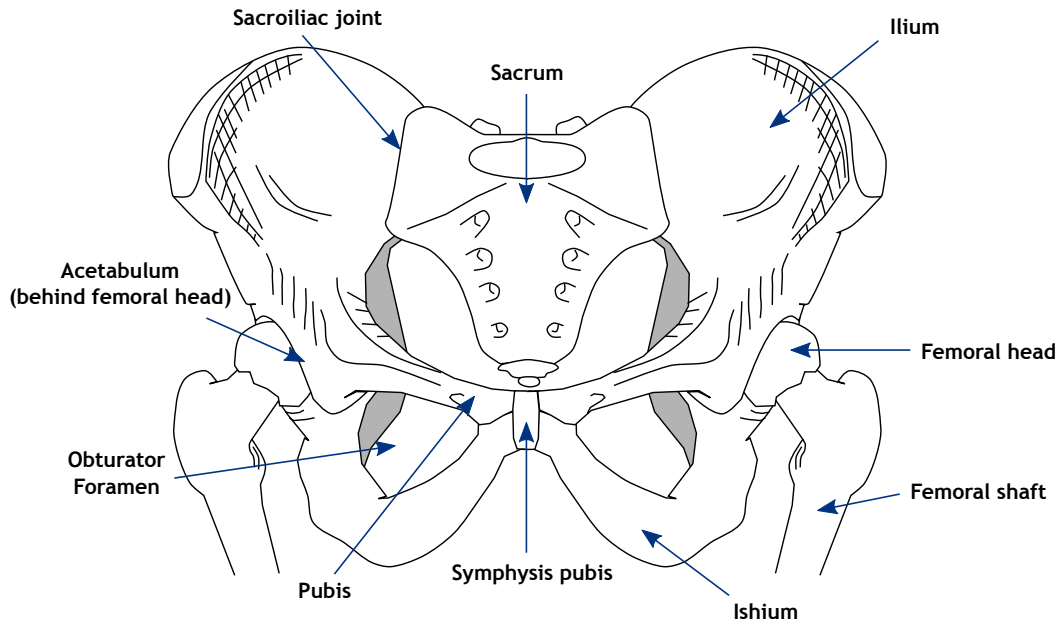


Figure 1.1: A brief overview of relevant pelvic structures.

[32].

A brief explanation of the theory behind X-ray imaging is useful in understanding some of its limitations. Medical radiographs are generated by exposing a region of the body to ionizing radiation in the form of a beam generated by an X-ray tube. The patient being examined lies between this tube (the source) and an X-ray film (the detector). Different anatomical structures absorb varying amounts of radiation from the beam, and the photons that are not absorbed pass through these structures to be sensed by the X-ray film. Visual characteristics of these structures are therefore determined by the degree to which they can block the X-rays. In the final X-ray image, bright (i.e. high intensity) regions are those exposed to fewer photons, and therefore correspond to radiopaque matter (i.e. that absorbing radiation). This includes the bone matter of the skeleton. However, since the X-ray beam effectively images anatomical structures by projecting them as shadows on the film, depth information is lost. There is no way to determine whether

one anatomical structure is in front of another, and so they overlap on the resulting image. This causes difficulties in interpretation - and, more relevantly to this dissertation, means that more complicated methods are required to process X-ray images. CT and magnetic resonance imaging (MRI) scans avoid this issue by generating a cross-section of the body as a sequence of consecutive images. X-ray images, meanwhile, present a 2-D view of a 3-D structure.

It may therefore seem wiser to rely solely on CT imaging in diagnosing trauma; it can be used to generate 3D models, offers a higher level of detail and provides clear delineation of bone from soft tissue. For these reasons, the total number of CT scans performed annually in the United States has risen over the past thirty years from approximately 3 million to almost 70 million [1]. However, there are caveats. The additional exposure to ionizing radiation associated with CT may pose an additional cancer risk to patients, particularly those who receive multiple scans (as can happen in follow-up care for trauma). For example, a chest CT scan typically delivers over 100 times the radiation of a conventional frontal and lateral chest radiograph [42]. While the FDA estimates that a CT scan is associated with an increased cancer risk of approximately 1 in 2000, other studies have found the risk to be as high as 1 in 80 among young women who undergo multiphase abdomen and pelvis CT imaging [62]. Though further analysis and studies are required, current results suggest that limiting the number of CT examinations may be wise until the risks are fully understood.

There are also issues specific to the use of CT in trauma. CT scans take longer to perform than conventional radiographs, and while they pose less disruption to the patient than a MRI scan, they may still be unsuitable for use in unstable patients with severe traumatic injuries. Conflicting opinions exist as to the utility of anteroposterior (AP) radiographs in pelvic injury, with proponents of CT imaging pointing to its ability to detect bone fragments and hemorrhage; as an example, a

2006 study found that the sensitivity of routine pelvic radiographs was 78% of that of CT scans in identifying fracture across multiple bone structures. However, the authors also stated that patients who are hemodynamically unstable or critically injured should undergo plain-film X-ray imaging to aid in prompt stabilization [47]. Another study conducted in 2007 evaluated the success rate of radiographs in identifying fractures to specific pelvic structures. Results showed that while AP radiographs were not useful in detecting fractures of the sacrum and iliac wing, they showed good sensitivity in identifying anterior injuries such as those to the pubic rami [38]. Resnik et al. found that AP radiographs are an efficient way to quickly determine the need for external fixation in cases of acute pelvic trauma, and can identify most clinically significant fractures and dislocations [54]. Finally, Niwa et al. recommend the use of plain film radiography in predicting hemorrhage sites in hemodynamically unstable patients; in keeping with the findings in [38], accuracy was highest in patients with anterior fracture [46]. Though these studies differ in their overall findings, the general agreement is that in cases of traumatic pelvic injury, particularly in the larger setting of polytrauma, pelvic radiography is a useful initial diagnostic step. This group covers the most severely injured patients, some of whom are at risk of hemorrhage due to pelvic instability and/or require stabilization prior to CT imaging. In these situations, X-ray images allow prompt visual evaluation while minimizing disruption to the patient, and may indicate the need for angiography if hemorrhage is suspected. The Advanced Trauma Life Support Protocol used by many trauma centers therefore includes an AP view pelvic radiograph as an adjunct to primary survey of polytrauma [48]. For these reasons, a module for the analysis of X-ray images should be considered a fundamental component of a computer-aided decision support system for traumatic injury.



#### **1.4 Significance of this Study**

The significance of this study has both clinical and computational aspects. In a clinical setting, the major contribution lies in the development of an x-ray image analysis component intended for use in a larger computerized decision support system for trauma. If potentially useful diagnostic features can be extracted from pelvic radiographs, they can then be combined with patient variables from other sources and used as input to predictive models. The resulting computer-assisted decision support system can thus evaluate a patient's state at every stage of care and provide dynamically generated predictions and recommendations to physicians. Though similar systems and predictive models have been designed which incorporate image features into their generated decisions, they rely on human visual inspection - such as requiring a radiologist to interpret the image and then answer a series of simple yes/no questions about their findings (CITE). A decision support system with the ability to automatically extract key features directly from medical images would truly be novel in the field.

The work in this dissertation also offers a useful extension to the existing array of deformable model segmentation techniques. The use of B-spline interpolation to control a model's deformation and maintain curvature is a novel technique that can be applied to other imaging modalities and to applications outside the medical field. Prior use of splines has been confined to active contour models with no statistical component. Furthermore, only one other study has so far attempted to address the shortcomings of ASM in preserving smooth curves.

Finally, the techniques developed for automated fracture detection in radiographs offer clinical and non-clinical innovations. This is a problem that has received relatively little attention in the past, and those studies that exist have focused on simple anatomical structures. Automation also

remains challenging and is frequently dependent on pre-calibrated images (as described in Chapter 2). Though the solution described in this dissertation is specific to pelvic radiographs, the concept of hierarchical segmentation can be generalized to other areas. More significantly, the techniques developed for fracture detection have potential value in industrial applications.

### **1.5 Aims of this Study**

The ultimate aim of this study is to develop an image analysis component for a computer-assisted trauma DSS which can extract potentially useful diagnostic features from plain-film pelvic radiographs. These will later be combined with demographic information, physiological measurements and features extracted from images created via other modalities to form a dataset characterizing a patient's condition. By applying these techniques to a large set of prior cases, a dataset can be created to train multiple predictive models. When given images and signals collected from a new patient, the system will use the same methods to extract relevant features which will be fed as input to the trained models, generating treatment recommendations and outcome predictions.

In developing the image analysis module, this work focuses on the following tasks:

1. Segmentation and detection of key pelvic structures
2. Calculation of quantitative displacement measures
3. Detection of fracture in the pelvic ring and pubis
4. Identification of potential injury patterns

This comprises a set of potentially useful features that are likely to be visible in pelvic radiographs and can be represented in a form suitable for machine learning (e.g. as categorical or

numerical variables). Although other features such as the presence of sacral fracture are important in diagnosing pelvic trauma, AP radiographs may not provide sufficient detail [38]. This information is instead more likely to be obtained using a separate CT image analysis component.

It is vital to remember that the methods developed for these tasks must be *fully automated* if they are to be of use in a computerized DSS for trauma. Furthermore, since different structures will be analyzed for different visual characteristics, segmentation must be performed before any other analysis can take place. This creates the need for an automated segmentation approach able to handle pathological cases and the complications inherent in x-ray imaging, which will be described in Chapter 4.

The remaining chapters in this dissertation are organized as follows. Chapter 2 provides a survey of related work in the field and a detailed explanation of the basic Active Shape Model (ASM) segmentation algorithm which this work extends. The nature of this extension is described in Chapter 3, which explores a novel combination of cubic B-splines and ASM which controls model deformation in order to maintain curvature of detected structures. Chapter 4 presents a hierarchical algorithm which automates the segmentation process in the presence of multiple image artifacts and uncertain position of structures. Automatic fracture detection of the pubis and pelvic ring is explained in Chapter 5, along with a method used to calculate quantitative measurements that may be of use in identifying injury patterns. Results are presented in Chapter 6. Although it is not the primary focus of this dissertation, creating predictive rule-bases is the ultimate goal of the overall system; Chapter 7 therefore describes two recent experiments which combined the extracted X-ray features with other patient information to create predictive models of injury severity. The final chapter presents the conclusions of this study, along with suggestions for future enhancements.

## CHAPTER 2 Related Work and Approach

### **2.1 Overview**

This chapter provides an overview of relevant previous work in medical image analysis, focusing specifically on segmentation and fracture detection. Though X-ray analysis is the primary concern, studies involving the CT and MRI modalities are also examined. Computer-aided decision-support systems are briefly discussed, and the chapter concludes with a detailed description of the standard Active Shape Model (ASM) algorithm, which this dissertation adapts and incorporates into the X-ray segmentation approach in order to overcome certain key limitations.

### **2.2 Computer-Aided Decision Making Systems**

As stated in Section 1.1, computer-aided decision-making systems - or decision support systems (DSS) - have attracted a great deal of interest as a potential means to significantly improve patient diagnosis, treatment, and outcomes. This is of particular relevance in the field of trauma care, which requires rapid decisions based on a complex array of patient information consisting of many variables. Computerized systems provide a means to rapidly process large volumes of such data and offer advice to physicians based on the knowledge extracted from prior cases.

The theory of medical decision support systems involves the generation of predictive models that take patient data (or subsets of data) as input and, based on a large database of prior cases, evaluates the likelihood of specific outcomes. Individual models are typically specialized toward

specific tasks, such as predicting whether a patient will survive or whether they will benefit from a certain drug. The combination of multiple models, each able to answer some relevant question, forms a decision support system. Relevance is determined by the overall goal of the system and the field in which it is to be applied. Currently, no DSS has been designed for pelvic trauma in particular. However, several studies have developed models intended to predict the survival of patients with traumatic brain injuries (TBI), including two basic examples proposed by Signorini [59] and Rovlias [55]. Although both are efficient and simple to use, they incorporate only a limited range of patient variables, severely restricting their ability to output accurate and reliable predictions. Despite this, Rovlias still demonstrated the potential usefulness of the CART decision tree algorithm in creating transparent recommendations and allowing visualization of its predictions.

Several recent studies have explored the use of logistic regression and feed-forward neural networks, again in the field of TBI [36, 29, 56]. The results suggest that neural networks may outperform logistic regression; however, both methods remain questionable in their statistical reliability. Furthermore, neural networks adopt a 'black-box' approach that obscures the knowledge extracted from the input data. Logistic regression models present similar difficulties due to their focus on probabilities and numerical output. This lack of transparency may have a negative impact on physician trust in the final DSS, as there is no simple method by which a human user can view the reasoning behind its predictions. As stated above, the grammatical rules created by decision tree algorithms such as CART offer transparency by explaining which patient variables influenced a given prediction or recommendation. Due to this transparency and their level of accuracy, decision tree methods are popular in medical applications for predictive tasks. For example, a 2007 study examined the performance of multiple machine learning methods in predicting the outcome of TBI, and found that logistic regression and decision tree analysis offered the best re-

sults [50]. A similar study on TBI outcome prediction found that decision tree analysis confirmed some results of logistic regression and challenged others, and concluding that applying decision trees to observational data can potentially uncover additional diagnostic knowledge [2]. Chapter 7 provides more detail on both the CART and C4.5 decision tree algorithms, and how they can be employed in decision-support tasks. The use of CART has been emphasized in previous studies, along with other methods such as multivariate adaptive regression splines (MARS) [57, 34]. Both approaches provide informative models and easy visualization, rendering them highly suited to medical application.

As technology improves and more sources of patient data become available, the limitations of systems that use only a small subset of information will become increasingly clear. Currently, there is no widely-used DSS capable of integrating and processing all relevant patient information to generate accurate decisions - and more specifically, there is no system which incorporates clinically relevant features automatically extracted from medical images. Few systems acknowledge the importance of medical imaging in the diagnostic process at all, and those that do require manual inspection of images and input of observations into the predictive model [64, 24]. Omitting medical images fundamentally restricts the performance of any DSS - and if the system is to be automated, features must be extracted via computational methods rather than by sight. Section 1.3 described the numerous imaging modalities available to physicians, each of which requires different techniques for extraction of key information in quantitative form. However, a common theme in processing these images is the need to detect structures of diagnostic and therapeutic interest. The following section therefore presents a brief overview of medical image segmentation, with particular focus on the techniques previously applied to radiographs.

### **2.3 Medical Image Segmentation**

Medical image segmentation has many practical applications, whether as part of a decision support system or as some other diagnostic or quantitative analysis tool. It has consequently attracted a great deal of attention resulting in a vast literature of published studies and computational methods. A comprehensive review is beyond the scope of this dissertation; however, a brief overview of general techniques will provide useful context for the specific area of X-ray image segmentation.

#### *2.3.1 General Approaches to Segmentation*

In their most general sense, medical image segmentation methods can be broken down into three categories: automated, semi-automated, and manual. Manual segmentation covers the traditional techniques used by human experts, which are typically accurate but also time-consuming. Semi-automated segmentation requires a degree of user input, such as in the initialization process or evaluation of results; an example would be specifying parameters to a pixel clustering algorithm. Automatic segmentation is the most difficult of the three to achieve, as it must be performed without any human input. However, it is arguably also the most useful; not only does it provide the only practical means of processing large numbers of images, but it also has the potential to be much faster than segmentation involving human input, provided that the algorithms used are suitably efficient [58]. While the first benefit is not strictly necessary in the application described in this dissertation, the second is crucial in a busy trauma center environment.

In a recent review of methods for medical image segmentation, Withey divides them into three generations where each adds an additional level of algorithmic complexity [71]. The first generation groups together low-level techniques based primarily on pixel intensity values, such as binary thresholding using Otsu's selection method [49], edge tracing, and region growing based

on homogeneity of adjacent pixels. These approaches make no use of prior information and are of limited use in segmenting structures which show intensity inhomogeneity and/or have intensity values to other unrelated objects. Methods of the second generation incorporate statistical analysis and elements of uncertainty - for example, mixture models which model each pixel in an image as belonging to a known class of a predefined set, and neural networks for image pattern recognition. This group also includes deformable models such as active contours or snakes, which expand and contract in order to converge toward specific features in an image [40]. The third generation expands on this approach through the design of deformable models which incorporate a prior information; two of the best known are Active Shape Model (ASM) and Active Appearance Model (AAM) [16]. Both are supervised learning methods, and require a set of labeled training images in which some structure of interest is represented by a series of landmarks manually placed along its boundaries. Shape and intensity information is extracted across the training set to form a statistical model representing the desired structure. As the algorithm searches for the same structure in a new, unseen image, this model is used to evaluate the plausibility of each possible deformation. This use of higher-level knowledge addresses a key limitation of the second generation deformable models, which are more vulnerable to variations between images and the effects of noise, and has led to the widespread use of ASM, AAM and many proposed variants in medical image segmentation [22, 41, 76]. The majority of recently developed methods for segmenting radiographs are based on these algorithms, as will be discussed in the following section.

### *2.3.2 Segmentation of X-ray Radiographs*

Compared to other imaging modalities, segmentation of radiographs has received relatively little attention. This may be due to the additional complexities and significant variability in quality inherent to X-ray images. For instance, statistical fluctuations in the photon density of the X-ray



beam itself generate 'quantum noise', resulting in contrast and intensity variations that reduce the quality of the resulting image [4]. Furthermore, different tissues may have similar absorption rates, leading to blurred edges and no clear distinction between organs, skin and bone. Other challenges relate to the patient; for example, bone density varies between individuals, so radiographs of the same region of bone from two different patients may show very different intensity characteristics. A patient with osteoporosis will have reduced bone density, and bone matter in the affected area will therefore appear darker than normal in X-ray images. Similarly, though the general structure remains the same, precise shape of bones may differ from one patient to the next. This is complicated further by pose variability. In cases of traumatic injury, it may not be feasible to move the patient into the ideal position when taking X-ray images, particularly if he/she is in an unstable condition or has sustained multiple trauma. The location and orientation of specific structures may therefore vary from one X-ray image to the next, even when collected from the same patient. An issue of particular relevance in pelvic radiographs is overlap between bones, which causes blurred boundaries and difficulty delineating bone matter from soft tissue and organs. Pelvic X-rays are prone to another specific complication: the presence of gas inside the colon, which causes dark shadows to appear over the iliac fossa.

One early attempt at radiograph segmentation by Manos [39] employs region-growing and merging according to size, similarity, and connectivity, then assigns labels corresponding to each region's intensity characteristics. The effectiveness of this approach is limited as it considers neither spatial information nor existing knowledge of anatomy. Another group of early studies addresses the problem of identifying lung regions in chest radiographs, applying methods such as Markov random field models [70], rule-based heuristics [51], and classifiers based on local features [21] with varying levels of success. In recent years, however, statistical deformable models

have become the dominant approach in both automated and semi-automated methods. Example applications include detecting vertebrae fracture [43] and isolating lung fields in chest radiographs [68, 72]. Chen et al. apply ASM to automatic segmentation of the patella from lateral radiographs of the knee joint [13], initializing the model with a genetic algorithm that determines the correct starting position for the model. However, the patella is a relatively simple structure with clearly defined edges in the X-ray image, and it is unclear how well the method would perform in the presence of multiple overlapping structures and poorly defined bone boundaries. Unfortunately, this caveat applies to many of the studies that use ASM or AAM to detect objects of interest in pelvic X-ray images. For example, the automated segmentation method proposed by Boukala [7] treats the pelvis as a single structure and cannot distinguish between individual bones; furthermore, the shape context descriptors it uses to initialize the model are vulnerable to spurious edges in the input images. Ding [19] explores the use of an atlas-based approach in automatically segmenting the femurs from AP pelvic radiographs. The atlas defines the left and right femurs and a third shape that combines the bony pelvic structure, and the algorithm searches for correspondence between these shapes and the edges in the X-ray image. Results are promising, but the success of the method is dependent on the edges extracted from the X-ray, and it may prove difficult to determine parameters for the edge extraction operator which can produce similar results across a wide range of images. Furthermore, the femur is a relatively simple structure. The atlas approach is less suitable for segmentation of pelvic bones since it relies on integrity of the desired structure; an unreasonable assumption in pathological cases. The method of Chen et al. [15] also segments femurs from pelvic X-ray images, this time using an active contour model (or snake model) with a curvature constraint. Automatic initialization is achieved based on the position of the femoral shafts, using a technique similar to that described later in this dissertation, and is reasonably accu-

rate. However, the vulnerability of the snake model to false edges occasionally leads to incorrect structure detection; an effect which is likely to be far more pronounced in segmenting pelvic bones due to the complexity of the structure. Two other studies explore the use of texture analysis in fracture detection - but as with several of the other studies presented here, focus only on the femurs [73, 37]. Finally, Dong and Zheng propose an approach which automatically segments the femurs from 2-D pelvic radiographs using 3-D statistical models [20]. Since dense 3-D meshes make automation computationally expensive, the algorithm represents the 3-D structure using a simple component model; however, their method requires pre-calibrated X-rays and performs poorly in the presence of occluded edges, which could impact performance in pelvic structure segmentation.

Although these studies represent only a portion of the literature concerning ASM/AAM and radiograph segmentation, they follow a general trend: namely, the focus on detecting relatively simple structures with well-defined edges. Images of more complex regions with multiple overlapping bones present a greater technical challenge, which may explain the decision to treat the bony pelvis as a single shape in two of the studies mentioned in the previous paragraph. This shortcoming is addressed by the method presented in this dissertation, which is able to automatically extract the individual structures that comprise the pelvis using a novel adaptation of the ASM model. The following section therefore provides a brief overview of the standard ASM algorithm, as background to the new approach.

#### **2.4 Active Shape Model**

As stated in 2.3.2, multiple studies have applied deformable models to radiograph segmentation. In their most general sense, model-based methods locate structures in images using a prior model of the form this structure is likely to take, based on its characteristics in a set of training images.

When given a new, unseen image, the algorithm attempts to find the best match to this model. The exact form of this model depends on the algorithm used; Active Shape Model (ASM) takes a statistical approach that uses a training set of labeled images to determine plausible variations of the desired structure. ASM requires the construction of two separate models representing the structure of interest - one concerning its shape, and the other its appearance. Both of these will now be described in more detail.

#### 2.4.1 Shape Model

The first step in constructing the shape model is to determine a set of landmark points which describe the structure of interest in every training image. These are usually placed with the help of experts, effectively incorporating domain knowledge into the model. The major landmarks are located at edge junctions and points of high curvature, with intermediate landmarks placed at equal intervals between them to provide more precise definitions of the structure's contours. All landmarks are placed into a vector recording their connectivity. A set of  $n$  landmarks defining a given object can be represented as a  $2n$  element vector  $\mathbf{x}$  given by

$$\mathbf{x} = [x_1, \dots, x_n, y_1, \dots, y_n]^T \quad (2.1)$$

where  $(x_i, y_i)$  provides the coordinates of landmark  $\mathbf{x}$ . One such vector is generated for each training image, and so a set of  $m$  training images  $I_i$  will result in  $m$  vectors  $\mathbf{x}_j$  ( $j = 1, \dots, m$ ). At this stage of training, all landmark positions are still relative to their corresponding image. To make the shape model independent of scale, position and orientation, a transformation must be found that aligns the landmark vectors (or shapes) into a common coordinate frame. The best transformation is one that minimizes the sum of distances  $D$  between each shape and the mean,

calculated as

$$D = \sum |\mathbf{x}_i - \bar{\mathbf{x}}|^2 \quad (2.2)$$

First, each shape  $\mathbf{x}_i$  is translated so its center of gravity lies at the origin  $(0, 0)$ . One shape  $\mathbf{x}_i$  is then chosen as the starting estimate of the mean,  $\bar{\mathbf{x}}_0$ , with which all other shapes  $\mathbf{x}_j$  ( $i \neq j$ ) are then aligned. A mean of the aligned shapes is estimated as  $\bar{\mathbf{x}}_1$ , aligned with  $\bar{\mathbf{x}}_0$ , and scaled so  $|\bar{\mathbf{x}}_1| = 1$  (to constrain scale and orientation). This process is repeated until the estimated mean shows no significant change between iterations. Alignment of an individual shape to the mean is performed by rotating and scaling the shape in the tangent space to the mean. Mathematically, aligning  $\mathbf{x}_i$  with  $\bar{\mathbf{x}}$  requires finding values for the scale parameter  $s$  and rotation angle  $\theta$  which fulfill

$$\min |T_{s,\theta}(\mathbf{x}_i) - \bar{\mathbf{x}}|^2 \quad (2.3)$$

where  $T_{s,\theta}$  is a shape transformation operator defined by  $s$  and  $\theta$ . In other words, values which minimize the sum of square distances between points on the estimated mean and the transformed shape. Once the training shapes are aligned, they form a point cloud in a  $2n$ -dimensional space which can be modeled as a probability distribution, enabling evaluation of the plausibility of new shapes. To simplify the model, Principle Component Analysis (PCA) is used to reduce the dimensionality of the dataset, after which any shape  $\mathbf{x}_i$  can be approximated by

$$\mathbf{x}_i \approx \bar{\mathbf{x}} + \mathbf{P}\mathbf{b} \quad (2.4)$$

where  $\mathbf{P}$  contains the eigenvectors of the PCA covariance matrix, and so defines the coordinate frame aligned with the point cloud. A shape in this frame is defined by  $\mathbf{b}$ ; altering its elements therefore alters the shape. The degree of these alterations is controlled by applying limits to  $\lambda_i$ , the variance of the  $i^{\text{th}}$  element of  $\mathbf{b}$  across the entire training set of images.

Since  $\mathbf{b}$  defines a shape in model space, a Euclidean transformation is required to map  $\mathbf{b}$  to image space. The mapping defines the positions of the shape landmarks as

$$\mathbf{X} = T_{X_t, Y_t, s, \theta}(\bar{\mathbf{x}} + \mathbf{P}\mathbf{b}) \quad (2.5)$$

where  $s$  and  $\theta$  are the scale and rotation parameters and  $(X_t, Y_t)$  defines the translation parameter of the transformation operator  $T_{X_t, Y_t, s, \theta}$ . Together, these form the set of pose parameters  $\mathbf{Q}$ . The ‘best’ parameters  $\mathbf{Q}$  that fit a model instance to a new set of image points  $\mathbf{Y}$  (representing the shape in a new image  $I$ ) is obtained by minimizing the sum of squared distance between the corresponding model and image points. In other words, finding  $\mathbf{Q}$  to fulfill

$$\min |\mathbf{Y} - T_{\mathbf{Q}}(\bar{\mathbf{x}} + \mathbf{P}\mathbf{b})|^2 \quad (2.6)$$

Full details of the algorithm can be found in [16], but a brief description will be provided here. ASM adopts an iterative approach which initializes the model parameters  $\mathbf{b}$  to those for the mean shape  $\bar{\mathbf{x}}$ . Equation 2.4 is used to generate the positions of the model points,  $\mathbf{x}$ , which are then aligned with the points  $\mathbf{Y}$  found in the image. The image points are then projected into the model space, so the parameters in  $\mathbf{b}$  can be updated to match them. The process is repeated until there is no significant change in  $\mathbf{b}$  between iterations.

#### 2.4.2 Grey-level Model

Using the approach just described, a set of parameters  $\mathbf{b}$  can be used to generate an instance of the shape model in an image  $I$ . Since the goal is to locate a target structure in  $I$ , a method must be found to measure how well the shape generated by the parameters  $\mathbf{b}$  matches this target structure. In other words, a fitness function  $F(\mathbf{b})$  is required; the set of parameters  $\mathbf{b}_i$  which minimizes  $F$  corresponds to the shape in  $I$  that best matches the target structure. The simplest approach would

be to assume that the shape models are constructed around strong edges in the training images. In this case,  $F$  would be based purely on the distance between each model point and the nearest strong edge in  $I$ . However, this does not allow for weaker edges defining internal boundaries or some other aspect of the desired structure. Instead, ASM builds a statistical grey-level model based on pixel intensity values around each landmark point in the training set.

When the grey-level model is built, each landmark point  $\mathbf{l}_i = (x_i, y_i)$  is considered in turn. Assume the process begins with a training image  $I_m$ . ASM focuses on the intensity values along the profile normal to the model boundary at  $\mathbf{l}_i$ ; specifically,  $k$  pixels are sampled along this profile either side of  $\mathbf{l}_i$ . The derivatives of this set of  $2k + 1$  samples are stored in a vector  $\mathbf{g}_i$  (sampling the derivatives instead of the raw intensity values limits the impact of global intensity changes between images). This process is repeated for all  $n$  images in the training set. Since it is assumed that the resulting  $n$  normalized vectors  $\mathbf{g}_i$  ( $i = 1, \dots, n$ ) follow a multivariate Gaussian distribution, the region around the landmark  $\mathbf{l}_i$  can be characterized by their mean  $\bar{\mathbf{g}}$  and covariance  $S_g$ . This generates an appearance model for each individual landmark in the shape model. The fitness function  $F$  can thus be formulated as

$$F(g_s) = (\mathbf{g}_s - \bar{\mathbf{g}})^T S_g^{-1} (\mathbf{g}_s - \bar{\mathbf{g}}) \quad (2.7)$$

where  $g_s$  is the sample being tested. It can be seen that  $F$  represents the Mahalanobis distance of the sample from the appearance model mean.

Searching for the target structure in a new image is an iterative process that fits the shape model  $\mathbf{X}$  to the image based on this set of grey-level models. Again, each landmark  $\mathbf{l}_i$  is considered in turn. The goal is to find the ‘best’ new position for  $\mathbf{l}_i$  by calculating  $F(g_s)$  for each of a set of points along its normal profile. The point with the lowest value of  $F$  - i.e. the lowest Mahalanobis

distance from the model mean - is chosen as the new position for  $\mathbf{l}_i$ . The Mahalanobis distance is preferred as it incorporates correlations present within the data and is scale-invariant. The search process is repeated separately for each landmark point. Figure 2.1, inspired by that presented in [16], provides a visual explanation.

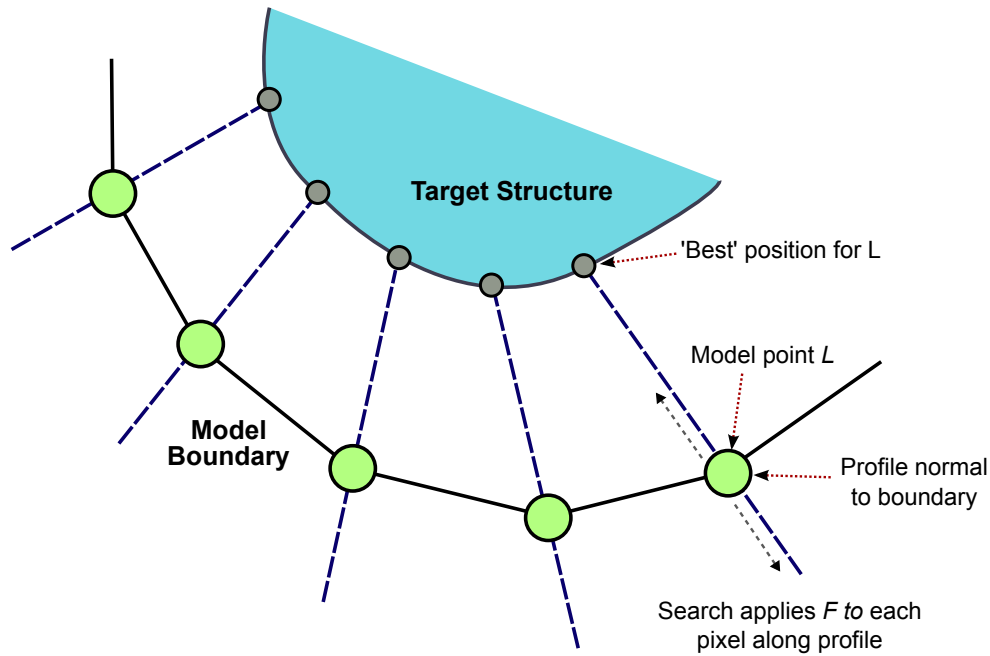


Figure 2.1: The ASM search process calculates the fitness function  $F(g_s)$  for each of a set of pixels along the profile to the model point. After multiple iterations,  $F(g_s)$  should be minimal at the edge of the target structure.

Finally, the pose parameters  $\mathbf{Q}$  and shape parameters  $\mathbf{b}$  are updated to fit the shape model to the new landmark positions. This process constitutes a single search iteration, and is repeated until there is no longer any significant change in  $\mathbf{b}$  (which corresponds to no significant change in the positions of the landmarks). ASM is often used in a multi-resolution form, where each image is decomposed across several levels, creating a series of decreasing resolution images, each of which is a smoothed and subsampled version of the last. These correspond to levels of a Gaussian pyramid. The search process begins at the highest level of the pyramid, i.e. with the lowest



resolution image. When deformation converges at some level  $m$ , the algorithm progresses to level  $m - 1$ , transferring the intermediate search results into a higher resolution image. In this way, the location of the identified shape is refined over progressively higher resolutions; the original image is considered only at the final stage of the process when no major changes are expected in the model point positions. Multi-resolution analysis offers reduced processing times, as the starting search space is greatly reduced. It is therefore of use in time-sensitive applications such as analyzing X-ray images for trauma.

### **2.5 Approach of this study**

The focus of this dissertation is on extracting potentially useful diagnostic features from pelvic radiographs, such as the presence of fracture and quantitative measures of displacement. If these features are to be used in an effective trauma decision support system, they must be extracted quickly and without any user input. This first requires that the key structures of the pelvis be automatically segmented from the image for separate analysis. Numerous issues complicate this in a clinical environment; the novelty of the approach described in this dissertation is its ability to deal with variations in both the relative location of the x-ray machine and the horizontal and vertical position of the patient on the table. Unlike other segmentation algorithms, the novel method in this study also performs successfully in the presence of fracture and other pathologies, as is necessary for clinical use.

Segmentation is performed using a hierarchical approach that begins with identification of the femoral shafts and uses the information obtained at each successive stage as input to the next. This process identifies and extracts the left and right iliac bones, the outline of the pelvic ring, the left and right femurs, and the left and right pubis-ischium structures. Quantitative measures of

displacement can then be calculated, such as increased width of the symphysis pubis, and structures can also be analyzed for the presence of fracture. This dissertation focuses on detecting fractures of the pubis and pelvic ring, as these may prove useful in identifying established pelvic injury patterns. Figure 2.2 outlines the entire process of X-ray analysis, beginning with an antero-posterior (AP) X-ray of the pelvis as input. The analysis identifies the presence and location of fracture in the pelvic ring, along with horizontal fractures of the ischium and pubis, and calculates quantitative measures of pubis symphysis displacement. These features are combined with other patient data (such as demographics and physiological scores) and used as input to predictive rule generation. Although other image information is needed in order to characterize the full range of patterns, these features may also be potentially useful in identifying basic injury patterns.

The novelties of the approach are as follows:

- The ability to distinguish between the individual components of the pelvic structure, a necessary first step in detecting abnormalities
- A hierarchical initialization process which avoids the need for manual interaction
- Extraction of potentially useful diagnostic features such as fracture and displacement measures which can be used in the generation of predictive models for patient diagnosis and treatment outcomes

These three capabilities fulfill the essential needs of an image analysis module in a computerized decision support system. The following chapters will explain each component in more depth, beginning with the combined Spline/ASM algorithm which introduces an additional constraint on deformation to address the issue of blurred edges and overlapping bones.

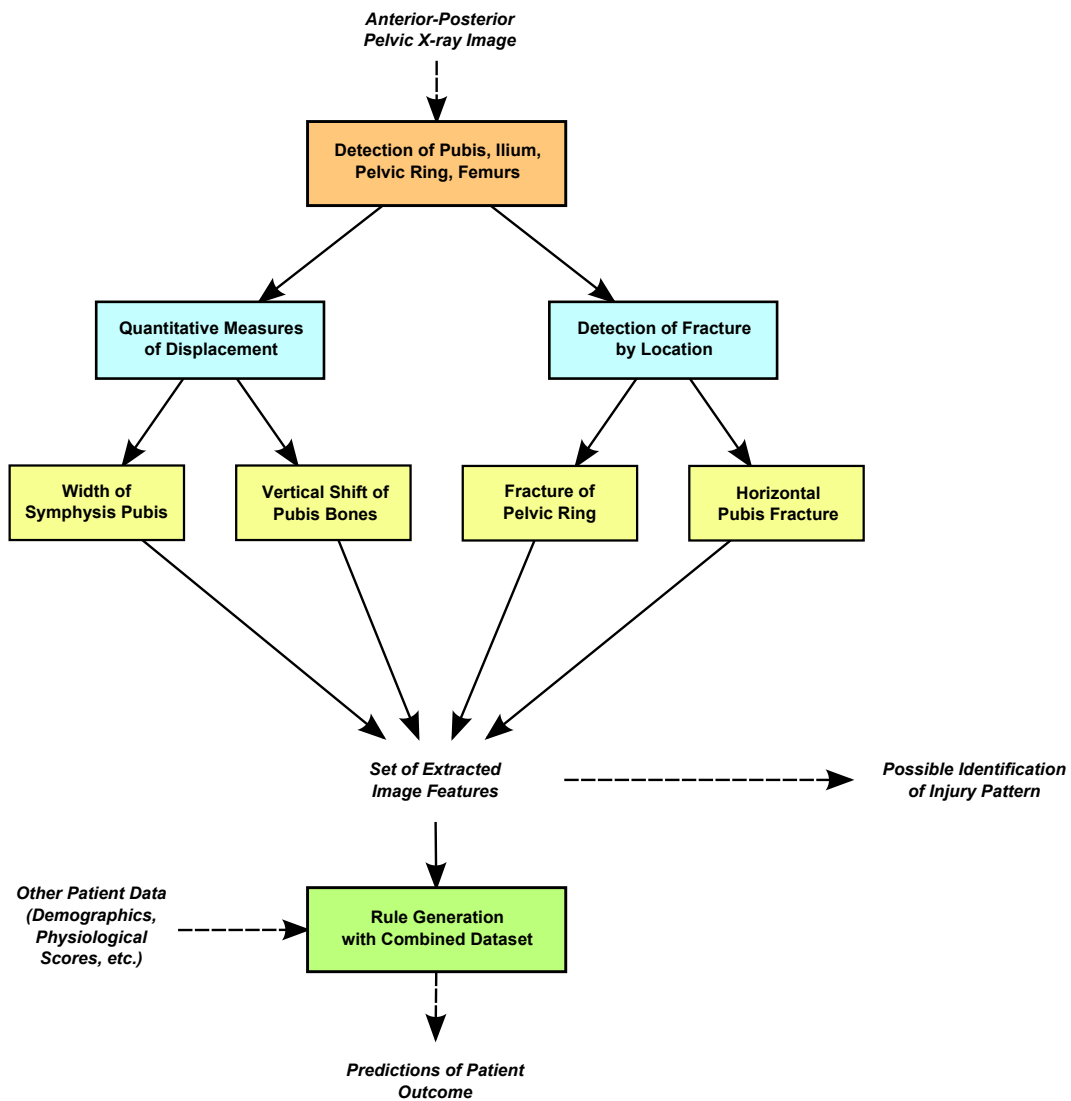


Figure 2.2: A schematic diagram of the overall X-ray analysis process.

## CHAPTER 3 Combined Spline/ASM Algorithm

### 3.1 Overview

This chapter presents a novel extension of the basic ASM algorithm outlined in 2.4 which provides a means to maintain structure curvature in the presence of false edges through the use of spline interpolation. The result is a combined Spline/ASM algorithm that preserves curves along the edges of objects during the deformation process, which is of particular use in segmentation of pelvic structures such as the left and right iliac bones, but is also applicable to other curved structures. Both cubic splines and B-splines are considered. Section 3.2 presents the basic concept behind the Spline/ASM algorithm, while Section 3.3 provides brief background detail on splines, spline interpolation, and their use in image segmentation. A detailed description of the algorithm is presented in Section 3.4, including more details on spline formation, and Section 3.5 summarizes the algorithm's approach and impact.

### 3.2 The Problems of ASM

The use of ASM in segmenting specific pelvic structures has shown promising results [61]. However, one of the major challenges in segmenting plain-film pelvic radiographs is the wide variation in image characteristics. In particular, multiple factors affect intensity contrast and sharpness, with some images appearing blurred and indistinct, and others highlighting variations in both bone texture and soft tissue. Figures 1(a) and 1(b) provide evidence of the extent of these variations.

Furthermore, the overlap of bones within the pelvis can create false edges in the X-ray image to which standard ASM is particularly vulnerable. Together, these two factors often cause




Figure 3.1: Two examples of the differences in pelvic radiograph image characteristics.

incorrect convergence and severe distortion of the detected shape. Since the problem lies in over-deformation of the starting shape model, the most obvious solution is to strongly limit the degree of deformation by reducing the  $\sigma$  parameter. Unfortunately, this is only an option when the shape in the image under consideration is highly similar to those used to train the shape model. This assumption cannot be made for pelvic radiographs, due to the differences in patient and machine position, individual variations in patient anatomy, and variability in pathological cases.

The example segmentation result presented in Fig. 3.2 hints at the underlying problem. While the landmarks along the upper edge of the superior pubic ramus are correctly placed, those outlining the obturator foramen show incorrect convergence, and the underside of the ischium shows extreme deformation. However, in the latter case, the three landmarks at either end of the corresponding edge are placed with reasonable accuracy. The issue seems to be that during deformation, standard ASM does not put sufficient emphasis on the relationship between neighboring landmarks. As explained in Section 2.4, the deformation step of the ASM algorithm considers each landmark point separately and attempts to find its ‘best’ new position. This search depends

largely on the grey-level intensity model, with the influence of the shape model being limited to the overall structure of the shape. Individual landmarks are consequently allowed to move far from their neighbors, and the deformation of separate edges within the shape (such as the ischium underside in Fig 3.2) is not sufficiently controlled.




Figure 3.2: Severe distortion of pubis-ischium structure seen in Standard ASM segmentation.

In practice, many different techniques could be used to implement these restrictions. The choice depends on the general shape of the structure the user wishes to extract from the image. In the case of pelvic X-ray segmentation, the majority of the structures of interest have curved edges; for example, the pubis-ischium structure shown in Fig. 3.2, and the pelvic ring and iliac bones shown in Fig. 1.1. Fracture does not typically disturb the general shape of each structure; exceptions are severe cases of pelvic disruption (in which case the severity of injury will be obvious), and fracture to the left and right pubis (which are therefore subjected to weaker spline constraints in this study).

Assume a subset of consecutive landmarks  $S$  is used to represent one of these curved edges.  $S$  becomes analogous to a sequence of input data points which can be approximated by an interpolating spline - providing a mathematical method to express the connections between one landmark

and the next. Though multiple forms of spline exist, in a general sense they can all be considered as mathematical models that associate a continuous representation of a curve or surface with a discrete set of points in a given space. They are constructed from piecewise polynomials fitted over the intervals between the specified control points, where the coefficients of each polynomial dictate the overall shape of the spline. The aim of spline interpolation is to find the polynomials that best fit the control points and thus interpolate some existing function  $f$ . This application focuses on the use of cubic splines, which are described in the following section.

### 3.3 Splines

#### 3.3.1 Splines in Medical Image Segmentation

One of the most popular uses of B-splines in segmentation are as B-spline snakes, or active contours [9]; they require fewer parameters and also offer local control, as described in Section 3.3.3. They are also used in image interpolation, where their high precision makes them particularly suited to analysis of medical images [35]. In one particular application, Stammberger et al [63] explored their use in measuring cartilage. Another example is the method of Arikidis et al. [3] which uses B-splines to detect microcalcifications in mammograms, images which present similar segmentation issues to conventional radiographs.

#### 3.3.2 Cubic Spline Interpolation

A cubic spline is formed from a series of piecewise cubic polynomials passing through a set of control points, where adjacent points are joined by a single cubic polynomial. Cubic splines are popular in interpolation tasks, as they can be used to fit a smooth curve to a series of data points. Consider a set of data points  $\mathbf{D}_k = x_k, y_k$  (where  $0 \leq k \leq n$ ), for some unknown function  $y = f(x)$ . This creates  $n$  intervals between points. Cubic spline interpolation creates a continuous curve  $\mathbf{S}(x)$

passing through each point can be formed from  $n$  piecewise cubic polynomials  $S_k(x)$  connecting successive data points, as follows:

$$\mathbf{S}(x) = \begin{cases} S_0(x) & x \in [x_0, x_1] \\ S_1(x) & x \in [x_1, x_2] \\ \dots & \\ S_{n-1}(x) & x \in [x_{k-1}, x_k] \end{cases} \quad (3.1)$$

where  $S_k(x)$  is defined as:

$$S_k(x) = a_k(x - x_i)^3 + b_k(x - x_i)^2 + c_k(x - x_i) + d_i \quad (3.2)$$

Note that several conditions must be imposed on  $\mathbf{S}(x)$  for it to be considered a cubic interpolating spline. First, it must pass through all the given data points. To ensure smoothness, it is also required that each  $S_k(x)$ , its derivative, and its second derivative all be continuous at the two control points defining its interval. In other words, each spline segment must be smoothly joined to the next, creating the appearance of a seamless curve.

### 3.3.3 B-splines

In data interpolation applications, cubic splines can be represented in B-spline form. B-splines (or basis splines) are generalizations of Bezier curves, and provide a spline representation that depends only upon a set of control points and the degree of the piecewise polynomial segments between them [67]. The control points define the shape that the spline should follow and moving these points therefore alters the shape of the curve. Each control point is associated with a B-spline basis function, which is non-zero only over a number of adjacent intervals. This allows local control, one of the main advantages of B-splines over Bezier curves; in other words, moving a control point only alters the local region of the curve. The shapes of the basis functions are



defined by a knot vector - or more precisely, the spacing between individual knots. Moving the knots is therefore another way to change the overall shape of the spline.

Given a set of  $(n + 1)$  control points  $P_0, \dots, P_n$  and a vector of knots  $\mathbf{u} = u_0, \dots, u_m$ , the cubic B-spline  $\mathbf{S}(\mathbf{u})$  is a parametric curve defined as

$$\mathbf{S}(\mathbf{u}) = \sum_{i=0}^{k-1} P_i b_{i,3}(\mathbf{u}) \quad (3.3)$$

where  $b_{i,n}$  are the B-spline basis functions,  $i$  is the control point number,  $\mathbf{S}(\mathbf{u})$  consists of  $(k-2)$  spline segments between these points, and  $m = n + 4$ . The point on the spline corresponding to a specific knot is called a knot point, and these divide the spline into segments (each of which is a cubic Bezier curve).

Consider a set of  $(n + 1)$  input data points  $\mathbf{D}_0, \mathbf{D}_1, \dots, \mathbf{D}_n$ , where  $D_k = (x_k, y_k)$ . Cubic interpolation of this data requires finding a B-spline curve  $\mathbf{S}(\mathbf{u})$  with a parameter vector  $\mathbf{t} = t_0, \dots, t_n$ , such that  $\mathbf{D}_k = \mathbf{S}(t_k)$  for all  $0 \leq k \leq n$ . These parameters are used to compute the knot vector and so their choice strongly influences the shape of the curve. One selection method is uniform spacing. Assuming the domain of the curve is  $[a, b]$ , the parameters are defined as follows:

$$t_i = \begin{cases} a & i = 0 \\ a + i \left( \frac{b-a}{n} \right) & 0 < i < n \\ b & i = n \end{cases} \quad (3.4)$$

It can be seen that the parameters (and therefore the computed knots) are equally spaced. Uniform spacing is the simplest approach and typically offers the tightest interpolation. It also mimics the spacing of the points on the trained shape model, which are roughly the same distance apart.

However, there is an intrinsic issue with the use of splines in ASM. All the model shapes

of pelvic structures constructed during ASM training contain non-distinct data points - i.e. there may be two different values of  $y$  for the same value of  $x$ . For example, the pelvic ring forms a roughly elliptical structure. Standard interpolation between the elements of two vectors  $x$  and  $y$  (containing the  $x$ -coordinates and  $y$ -coordinates of the landmark points respectively) expects a one-to-one correspondence between  $x$  and  $y$  values. Figure 3.3 demonstrates this situation; the spline  $Sp1$  covers a subsequence of the pelvic ring where there is no ‘doubling back’ and therefore no issue with standard interpolation. However,  $Sp2$  covers a larger subsequence, where two separate points on the defined curve can share the same  $x$ -coordinate but have different  $y$ -coordinates. The interpolating spline is required to pass through every given data-point, which results in an oscillating function that bears no resemblance to the true shape. Furthermore, consider two points on a line close to the vertical. Considering this in terms of image coordinates, a small range of  $x$  values will correspond to as many  $y$  values as there are pixels in the line, which causes similar problems in interpolation. This is an issue that is likely to affect any object represented by a model with a closed contour, and will be addressed in the following section.




Figure 3.3: An example of standard spline interpolation failing on a circular structure.

### 3.4 Method

#### 3.4.1 Spline interpolation in ASM deformation

The idea behind the combined Spline/ASM algorithm is to strengthen the connection between successive model points, where these points form a smooth curved contour. The curve can be relatively weak; the aim is to create a smooth approximation of the contour in spline form. This approach is not suitable for objects with hard corners and non-smooth edges, but works well for the curved structures of the pelvis, particularly the pelvic ring.

Spline/ASM maintains curves by incorporating splines into the ASM quality of fit function described in Eqn. 2.7 in Section 2.4. When considering a model point  $\mathbf{I}_i = (x_i, y_i)$ , where  $i = 1, \dots, k$  and  $k$  is the total number of points that form the model, an cubic interpolating spline is constructed using  $m$  model points on each side of  $\mathbf{I}_i$  (excluding  $\mathbf{I}_i$  itself). In other words, a sequence

of points  $\mathbf{p}$  is used for interpolation, where:

$$\mathbf{p} = l_{i-m}, \dots, l_{i-2}, l_{i-1}, l_{i+1}, l_{i+2}, \dots, l_{i+m} \quad (3.5)$$

The exact value of  $m$  can be altered by a system user, but is preset to a value appropriate relative to  $k$ . A spline function  $\mathbf{S}(\mathbf{u})$  is constructed according to the constraints specified in Section 3.3, with the points in  $\mathbf{p}$  providing the data  $D_k$  to be interpolated.  $\mathbf{S}(\mathbf{u})$  should form a smooth curve that passes through each of the points in  $\mathbf{p}$ . Using  $x_i$ , the x-coordinate of point  $\mathbf{l}_i$ ,  $\mathbf{S}(\mathbf{u})$  predicts a corresponding value of  $y$  along this curve, returning a predicted point location  $\mathbf{l}_p = (x_p, y_p)$ .

When deciding on the best new location for  $\mathbf{l}_i$  during an iteration of model deformation, the difference in position between  $\mathbf{l}_i$  and  $\mathbf{l}_p$  is included in the quality of fit function. As in the standard ASM algorithm,  $a$  points  $q_j$  (where  $j = 1, \dots, 2a$ ) are sampled along the profile either side of  $\mathbf{l}_i$  and the quality of fit is calculated for each. At each profile point  $\mathbf{q}_j = (x_j, y_j)$ , an error measure  $e$  is obtained between  $\mathbf{q}_j$  and  $\mathbf{l}_p$  as follows:

$$e = |x_j - x_p| + |y_j - y_p| \quad (3.6)$$

The quality of fit function then becomes:

$$f(g_s) = (g_s - \bar{g})^T S_g^{-1} (g_s - \bar{g}) + \alpha * e \quad (3.7)$$

where the scaling measure  $\alpha$  is chosen via experimental results; since the sample derivatives used to construct the grey-level model tend to be lower in value than possible pixel distances,  $\alpha$  is typically very small. For the application described in this dissertation, the optimal value of  $\alpha$  was found to be  $10^{-5}$ . Note that although the standard ASM fitness function is calculated using the Mahalanobis distance measure, this relates to the fit with the statistical shape and grey-level

models. The spline penalty term is purely spatial in  $x - y$  pixel coordinates, and an error measure based on Euclidean distance is therefore acceptable. Using this approach,  $l_i$  is penalized for strong movements away from the curve specified by its neighboring points. However, smaller degrees of movement are allowed, since the interpolating spline  $\mathbf{S}(\mathbf{u})$  for a given  $\mathbf{l}_i$  is recalculated during each iteration of deformation using the updated model point positions. This provides the local control that is lacking in standard ASM, which only allows control over deformation of the entire shape model rather than individual segments. Additionally, splines are calculated using the new locations of neighboring points calculated during the current iteration. When Spline/ASM is searching for the best new position for point  $\mathbf{l}_i$ , the best positions for points  $(\mathbf{l}_{i-m}, \dots, \mathbf{l}_{i-2}, \mathbf{l}_{i-1})$  have already been determined. These new positions are used in calculating  $\mathbf{S}(\mathbf{u})$  rather than the positions at the previous iteration of deformation.

Spline/ASM introduces several other constraints. First is the concept of *control points*, which are specified during training and incorporated into the resulting ASM model. A control point is a model point that marks a hard angle in the shape. For example, consider the two vertical edges of the left and right pubis; the transition between these and the smooth curve of the upper ramus would be labeled as a control point. If a control point  $\mathbf{c}_k$  is encountered when constructing the sequence  $\mathbf{p}$  of neighboring points used for interpolation, the sequence ends there. Furthermore, the interpolation error measure is not applied when considering the best new position for  $\mathbf{c}_k$ . This performs the opposite task of spline interpolation in maintaining any hard angles of the shape. Control points also include the first and last points in non-closed shapes.

A second constraint is used to deal with model contours that are approximately vertical, such as the femoral shafts. When building  $\mathbf{p}$  for a given model point  $\mathbf{l}_i$ , the gradient between neighboring points is calculated. If this exceeds a predefined threshold,  $\mathbf{l}_i$  is considered a *vertical point* and

the interpolation measure is not used. This handles one of the situations where non-distinct data points occur. However, this will not detect whether  $\mathbf{p}$  ‘doubles back’ - i.e. some members of  $\mathbf{p}$  are close in  $x$ -position but distant in  $y$ -position. Three methods are available to deal with this problem. The first involves specifying a set of *pivot points* during model training; these behave similarly to control points, except that they are subject to the interpolation error penalty. The user can place these at locations where a curve interpolating a given sequence  $s$  is likely to double back on itself, and they will terminate sequence construction in the same way as control points. The second choice involves simply not interpolating a point  $\mathbf{l}_i$  if its surrounding sequence  $\mathbf{p}$  contains non-distinct data points. This is a simple and quick approach that requires no alteration in initial training. However, it limits curvature preservation.

A third method employs chord length parameterization instead of uniform spacing. A chord in plane geometry is defined as the line segment joining two points on the same curve. If the interpolating spline curve closely follows the polygon created by the data points, chord-length parameterization approximates arc-length parameterization between two adjacent data points. If the domain is then subdivided according to the distribution of chord lengths, the results approximates arc-length parameterization for the entire spline.

Let the provided data points around the ASM model be  $\mathbf{D}_0, \mathbf{D}_1, \dots, \mathbf{D}_n$ . The length between  $\mathbf{D}_{i-1}$  and  $\mathbf{D}_i$  is  $|\mathbf{D}_i - \mathbf{D}_{i-1}|$ . The length of the entire polygon,  $L$ , is given by:

$$L = \sum_{i=1}^n |\mathbf{D}_i - \mathbf{D}_{i-1}| \quad (3.8)$$

Consider the chord length from  $\mathbf{D}_0$  to data point  $\mathbf{D}_k$ , and let  $L_k$  be the ratio of this length over  $L$ , given by

$$L_k = \frac{\sum_{i=1}^k |\mathbf{D}_i - \mathbf{D}_{i-1}|}{L} \quad (3.9)$$

If the domain of the curve is  $[0, 1]$ , then parameter  $t_k$  should be located at the value of  $L_k$ . The parameters used to compute the knot vector are therefore defined as follows:

$$t_i = \begin{cases} 0 & i = 0 \\ \frac{1}{L} \left( \sum_{j=1}^i |\mathbf{D}_j - \mathbf{D}_{j-1}| \right) & 0 < i < n \\ 1 & i = n \end{cases} \quad (3.10)$$

As before, during ASM search the spline curve is constructed using  $m$  points either side of  $\mathbf{l}_i$ , the control point under consideration, and the polynomial spanning the interval  $[\mathbf{l}_{i-1}, \mathbf{l}_{i+1}]$  is extracted. Interpolation is then performed for multiple values in  $[t_{i-1}, t_{i+1}]$  (typically around 100) to generate a vector  $\mathbf{t}_s$  of  $t$ -values spanning the interval. This allows the  $x$  and  $y$  coordinates along the interval  $[\mathbf{l}_{i-1}, \mathbf{l}_{i+1}]$  to be considered as separate functions of  $t$ :

$$y = f(t) \quad (3.11)$$

$$x = g(t) \quad (3.12)$$

Cubic spline interpolation is then performed separately for each function, creating two vectors  $\mathbf{x}_s$  and  $\mathbf{y}_s$  along the polynomial spanning  $[\mathbf{l}_{i-1}, \mathbf{l}_{i+1}]$ . The  $y$ -coordinate of the predicted point  $l_p$  can then be found via interpolation using  $\mathbf{x}_s$ ,  $\mathbf{y}_s$  and  $x_i$  (the  $x$  coordinate of  $\mathbf{l}_i$ ). Incorporating the third parameter  $t$  avoids the problem of oscillation affecting the uniform spacing scheme (which treats  $y$  as a function of  $x$ ) by moving from 2-D to 3-D space, providing monotonic parameterization.

A visual representation is provided in Fig. 3.4. Here, the stars in the plane  $t = 0$  represent the  $(x, y)$  coordinates of the pelvic ring ASM model. As a closed shape, this will be affected by the ‘doubling back’ problem; it will also fail completely in the situation where a single  $x$ -value corresponds to two  $y$ -values. The hollow circles represent the points plotted using the values  $x_j$ ,  $y_j$  and  $t_j$ , where  $x_j$  and  $y_j$  are obtained separately using  $t_j$  as described in Eqn. 3.11.

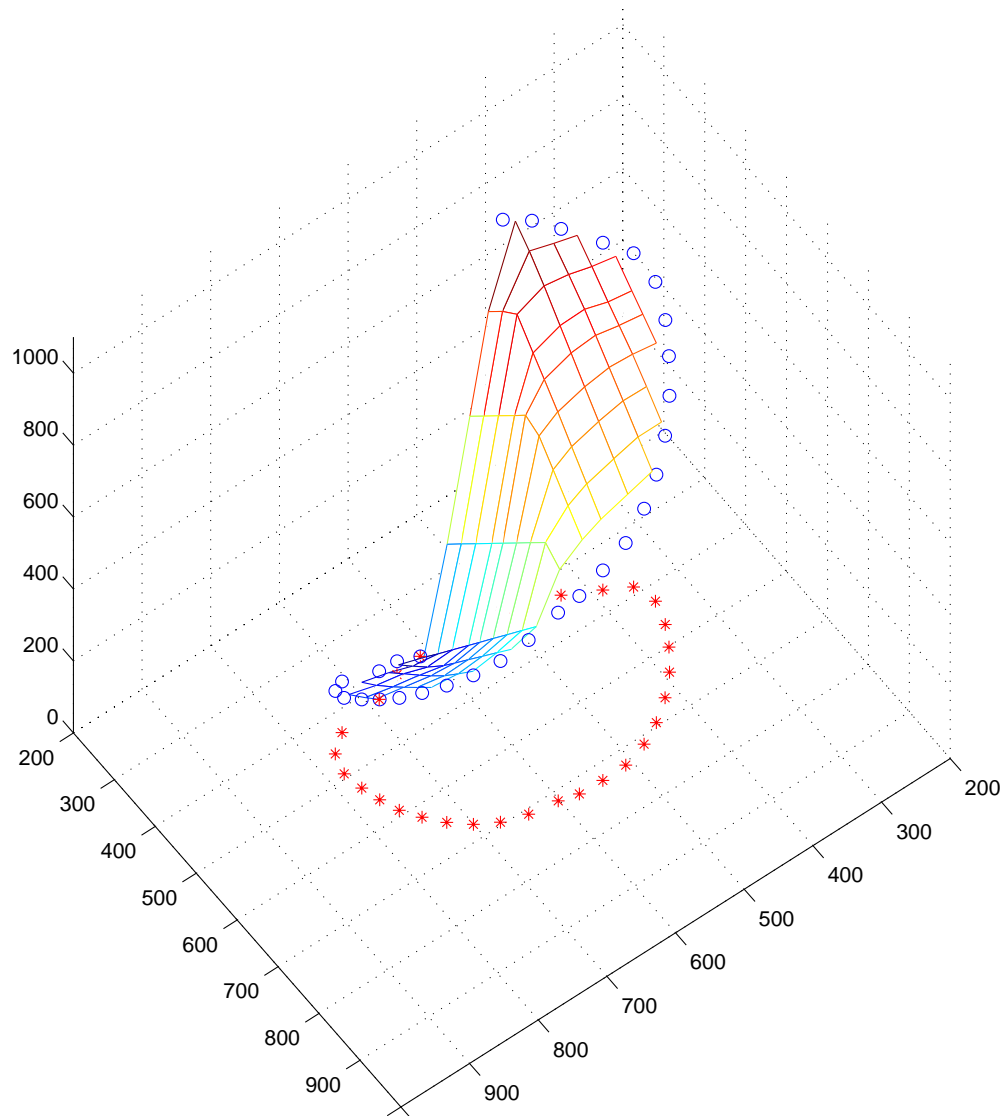


Figure 3.4: A 3D plot of the monotonic, strictly increasing function of  $(x, y, t)$  generated using chord-length parameterization.

Figure 3.5 shows successful interpolation of non-distinct data points around the pelvic ring using chord length parametrization, in contrast to the oscillating function in Fig. 3.3. The results presented in this dissertation were obtained using the uniform spacing scheme, as the chord-length approach requires further development. However, it is expected to provide improved performance in the future.



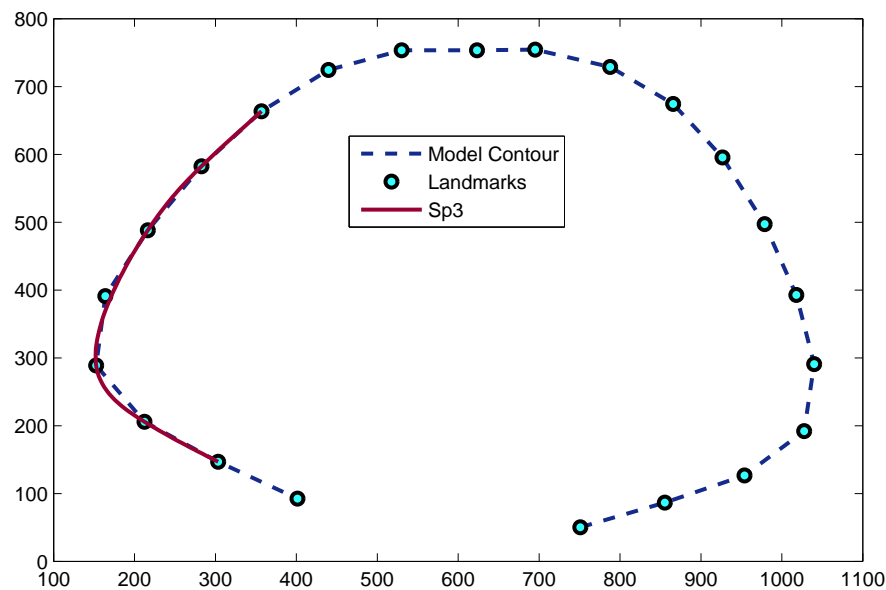


Figure 3.5: A spline parameterized on chord-length deals well with curves that ‘double back’.

### 3.4.2 Considerations during Training

The proposed system is intended to be trained prior to deployment; it then applies the learned models to new images. One of the important considerations during training is determining how many landmark points should be used to define each structure. The ASM framework has no standard method for this, since it depends on the application, the segmentation precision required, the permitted execution time, and multiple other factors. In this application, speed is the most crucial aspect. Though an increased number of landmark points typically provides smoother detected shapes, this comes at a cost of increased search time. Since segmentation in this application is intended as a way to identify general regions for analysis, rather than being an end goal in itself, smoothness is a lower priority than speed. The number of landmark points is therefore kept relatively low for all structures of interest.

Spline/ASM requires additional decisions: specifically, the number of points to be used in

spline construction. If uniform knot spacing is used, the developer training the system must also decide on appropriate positions for control and pivot points (another reason why chord-length parameterization must be further explored). The choice is expected to be task-specific. In this application, all parameters were determined via experimentation. Though this is a time-consuming process, it need only be performed prior to system deployment and requires no input from end users in the trauma center environment.

### **3.5 Summary**

This chapter presents a novel extension of the basic ASM algorithm to incorporate the use of splines to maintain curvature of the desired structures. The combined Spline/ASM algorithm preserves curves along object contours during the deformation process by penalizing strong deviations and thus emphasizing the relationship between neighboring model points. This dissertation focuses on its use in detecting pelvic structures, but it is of potential use in accurate segmentation of any strongly curved object.

## CHAPTER 4 Hierarchical Approach to Automatic ASM Initialization

### 4.1 Overview

This chapter introduces a novel hierarchical method that automatically initializes the Spline/ASM algorithm described in Chapter 3 for all the pelvic structures of interest. Using this method not only removes the need for manual placement of the structure models, but also enables full segmentation to be performed in a single pass. The approach is based on a combination of directed Hough Transform (DHT), Spline/ASM, edge detection, gradient filtering and statistical distributions of structure dimensions. An overview is provided in Section 4.2, followed by a more detailed description in Section 4.3 and a summary in Section 4.4.

### 4.2 Objective

Numerous issues complicate the automatic segmentation of plain-film pelvic X-rays in a clinical environment. These include:

- Uncertain horizontal and vertical position of the patient on the table
- Varying location of the x-ray machine relative to the patient
- The need to avoid disruption to the patient in emergency diagnosis situations

Chapter 3 introduced a novel adaptation and extension of standard ASM that incorporates spline interpolation to maintain the curvature of structures of interest. This addresses the problem of convergence to false edges. However, standard ASM suffers from a second disadvantage: its

high sensitivity to initialization. In order to detect a target structure, the corresponding mean shape model must be assigned a starting position somewhere in the image. The algorithm then attempts to fit the shape model to the structure over a number of iterations. If the shape model is incorrectly placed - for example, only partially overlapping the structure of interest - then ASM will instead converge to the nearest edges matching the corresponding grey-level model. Detection is therefore likely to fail. This is particularly relevant when dealing with pelvic X-ray images, which typically have multiple false edges due to overlapping bones. A result of inaccurate ASM initialization can be seen in Fig. 4.1. To human eyes, the starting position of the shape model does not appear significantly different to the location of the target structure; however, ASM's high sensitivity to false edges results in poor detection, particularly at the top of the pelvic ring.

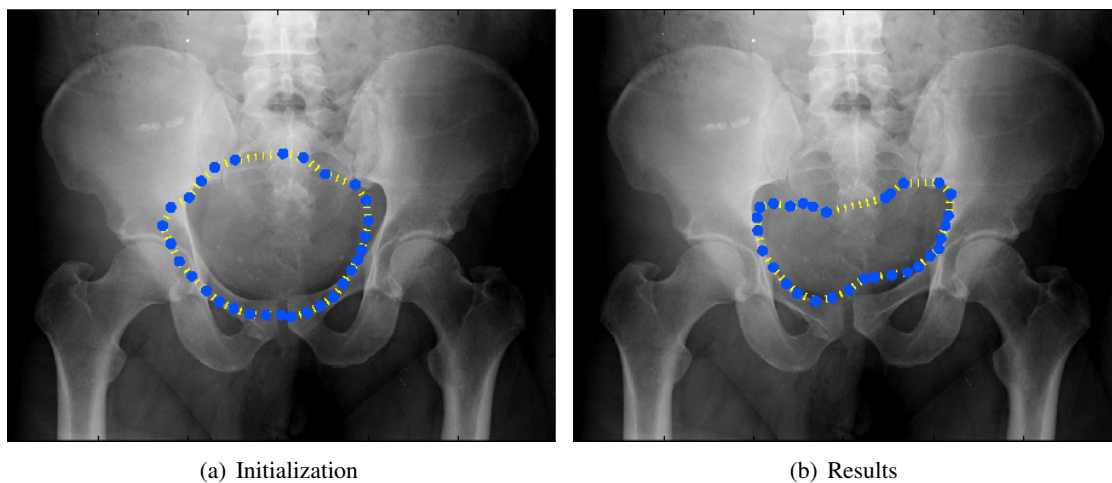


Figure 4.1: The effects of poor initialization on pelvic ring detection via standard ASM.

While the Spline/ASM algorithm prevents gross distortion of the model shape, it cannot compensate for an inaccurate starting location. Deformable model algorithms are based on pixel intensity values and target structure shape; they have no knowledge of the structure of the chosen image. The hierarchical initialization method compensates for this by using anatomical knowl-

edge and prior segmentation results to sequentially place the mean shape models for each individual structure of interest, starting with the femoral shafts. The femurs are solid, inflexible bones with high density, and are therefore clearly visible in the majority of pelvic X-ray images. The knowledge of each shaft's width and position is then used in locating the femoral heads. Once both the head and shaft positions are known, the Spline/ASM femur shape model can be correctly placed. Spline/ASM detection of the pelvic ring is then performed using multiple reference points, followed by detection of the left and right pubis-ischium structures. The location of the pelvic ring is then used to estimate the approximate positions of the left and right ilium, which are used in turn to place the corresponding Spline/ASM shape models. The schematic diagram in Fig. 4.2 outlines the initialization process, including how the result of each step is used to initialize the next.

### **4.3 Method**

#### *4.3.1 Detection of Femoral Shafts*

The femurs are among the most distinct structures in a pelvic x-ray, and typically suffer little deformity except in cases of severe femoral fracture. The initialization method therefore uses the left and right femoral shafts as reference points to determine the patient's horizontal displacement from the center, allowing location of the femoral heads and accurate placing of the femur shape models for Spline/ASM. Since accurate shaft detection is crucial to successful segmentation, this first stage is the most intricate of the initialization process. Figure 4.3 outlines the individual steps taken to accurately identify the femoral shafts, which will now be explained in more detail.

##### *4.3.1.1. Constructing the vertical emphasis edge map*

Since the femurs are typically located at the base of the X-ray image, the lowest 15% of rows are extracted as the region of interest  $I$  (as shown in Fig. 4(a)). Following analysis of the training

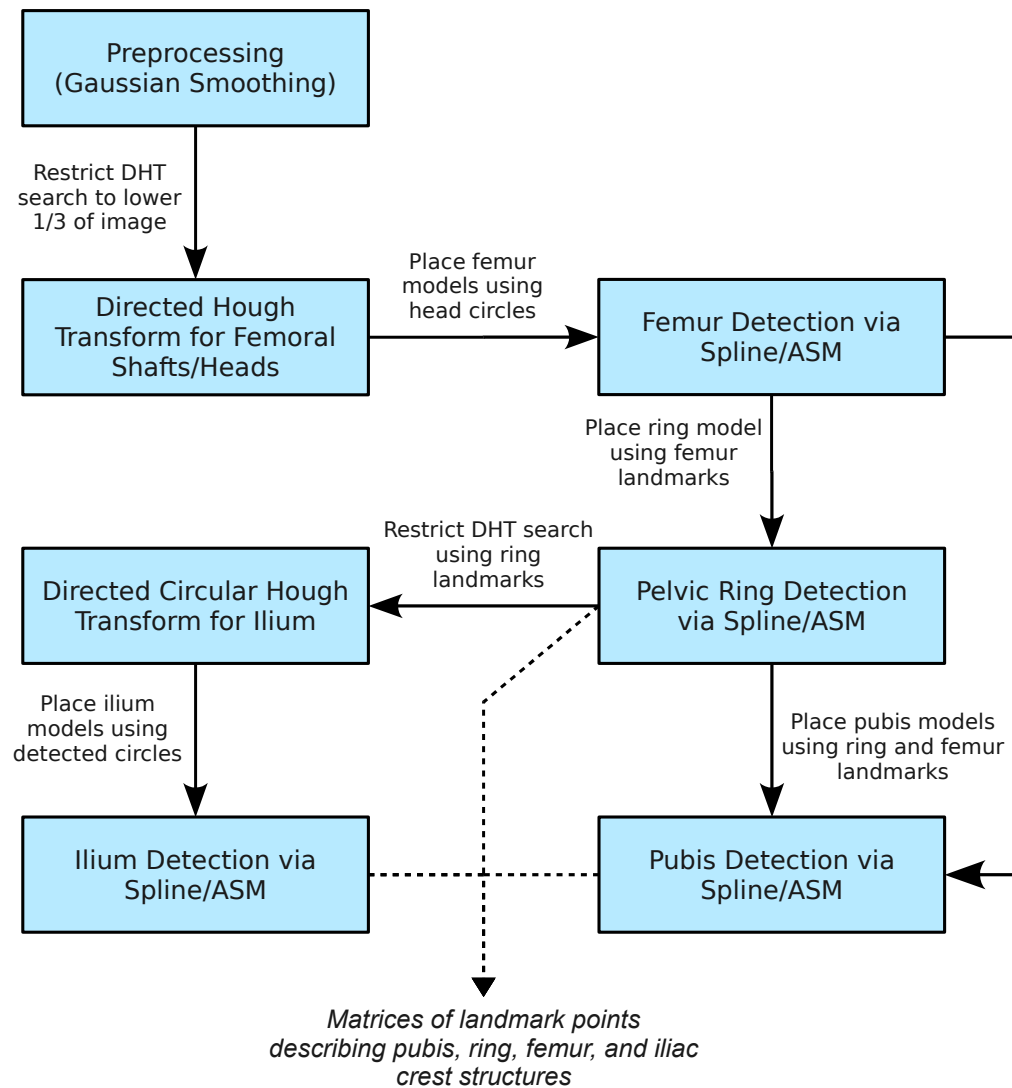


Figure 4.2: A schematic diagram of the Hierarchical Automatic ASM initialization process.

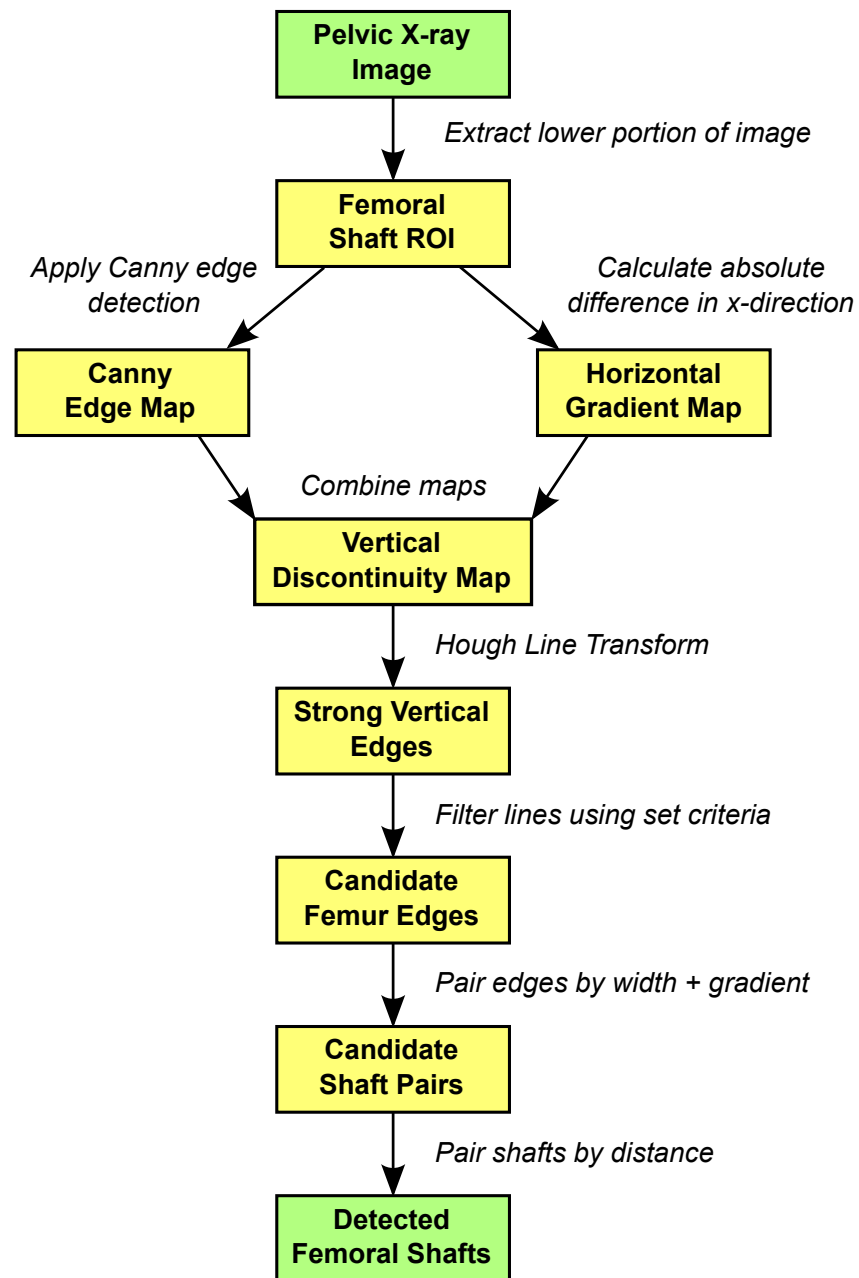


Figure 4.3: The multi-step process used in femoral shaft detection.

set, this was determined as the region most likely to contain a recognizable portion of the femoral shafts. Two separate operators are then applied in succession. The first, Canny edge detection, identifies the strongest edges in  $I$  by searching for local maxima of its gradient, calculated using the derivative of a Gaussian filter. Let  $I = (x, y)$  be the input image (the region of interest), treated as a function of  $x$  and  $y$ , and let  $G(x, y) = e^{-\frac{x^2+y^2}{2\sigma^2}}$  be a circular 2-D Gaussian function. A smoothed image  $I_s$  is formed as:

$$I_s(x, y) = G(x, y) * I(x, y) \quad (4.1)$$

where  $*$  is the discrete convolution operator, given by:

$$f(x, y) = h(x, y) * g(x, y) = \sum_{i=-n}^n \sum_{j=-m}^m h(i, j)g(x - i, y - j) \quad (4.2)$$

In the Gaussian smoothing case, input image  $I(x, y)$  is convolved (filtered) with the mask  $G(x, y)$ , and  $n$  and  $m$  define the domain of the kernel - i.e. the size of the pixel neighborhood involved in smoothing.

This removes noise present in the raw image which might interfere with accurate edge detection. The gradient magnitude  $M$  and gradient angle  $\theta$  of  $I_s$  are then calculated for every position  $(x, y)$  in the image, as follows:

$$M(x, y) = \sqrt{(G_x^2 + G_y^2)} \quad (4.3)$$

$$\theta(x, y) = \arctan \left[ \frac{G_y}{G_x} \right] \quad (4.4)$$

where  $G_x = \frac{\delta I_s}{\delta x}$  and  $G_y = \frac{\delta I_s}{\delta y}$ .  $M$  contains large ridges around local maxima in  $I$  which represent potential edges, which are thinned via non-maxima suppression in the horizontal, vertical



and two diagonal directions. A  $3 \times 3$  grid (the pixel neighborhood mentioned above) is then passed over every position in  $M$ , where a grid center  $M(x, y)$  is suppressed if its value is less than either of its neighbors in the direction specified by the rounded value of  $\theta(x, y)$ . This grid size was chosen following a series of experiments using a range of dimensions. The suppression process thins down wide edges to those positions with the highest gradient magnitudes, and the resulting thinned edge map  $I_N$  undergoes hysteresis thresholding and connectivity analysis. These techniques detect ‘strong’ and ‘weak’ edge pixels using high and low magnitude thresholds  $T_H$  and  $T_L$ . All ‘strong’ pixels are assumed to define edges, with only those ‘weak’ pixels connected to these edges being retained. In this application, the ratio of  $T_H$  to  $T_L$  is set lower than the standard recommendation made by Canny, and a relatively high standard deviation  $\sigma$  is used to construct  $G(x, y)$ ; this deals with the noise present in X-ray images and detects only the strongest edges likely to represent femur contours. Again,  $T_H$  and  $T_L$  were determined by a series of experiments using different threshold values and ratios; the final parameters are set as  $T_H = 0.3$  and  $T_L = 0.2$ , with  $\sigma = 4$  used to define the Gaussian kernel  $G(x, y)$ .

However, as can be seen in Fig. 4(b), superfluous edges are still returned that may impact shaft detection. Although these could be removed by further reducing the threshold ratio, this would likely result in a loss of generalization. As stated in 2, X-ray radiographs vary greatly in contrast and sharpness, and if the thresholds  $T_H$  and  $T_L$  are too specific then they are likely to give poor performance on other images. Therefore, a refined edge map is generated by combining the results of Canny edge detection with the horizontal gradient magnitude of  $I$ . The numerical gradient of  $I$  is defined as:

$$\nabla I = \frac{\delta I}{\delta x} \hat{\mathbf{i}} + \frac{\delta I}{\delta y} \hat{\mathbf{j}} \quad (4.5)$$

where the  $x$  component represents the differences between consecutive pixels across all columns

of  $I$ . Strong positive or negative values indicate gradient discontinuities, as shown in 4(c). The absolute value of  $\frac{\delta I}{\delta x}$  corresponds to the horizontal gradient magnitude, and is combined with the Canny edge map  $E_c$  to generate a new vertical discontinuity map  $E_M$ .

$$E_M = \frac{\delta I}{\delta x} \times E_c \quad (4.6)$$

This highlights those edges with highest gradient magnitude in the horizontal direction. Figure 4(d) presents an example discontinuity map; the shaft edges are preserved, while the artifact edges caused by non-bone matter are reduced in strength. When  $E_M$  is thresholded to a binary map, the shafts form clear straight lines which can be identified via Directed Hough transform.

#### 4.3.1.2. Locating shaft edge candidates

The Standard Hough Transform is a feature extraction method best suited to detecting regular geometric shapes with clear parametric forms. The straight lines of the femoral shafts can be described by the following parametric equation:

$$\rho = x \cos \theta + y \sin \theta \quad (4.7)$$

where  $\rho$  is the distance from the origin to the line along a vector perpendicular to the line and  $\theta$  is the angle between this vector and the x-axis. The lines that pass through a point  $(x, y)$  in the image space correspond to a sinusoidal in polar Hough space (the  $(\rho, \theta)$  plane). If the sinusoidals for two points  $(x_0, y_0)$  and  $(x_1, y_1)$  overlap, the location where they intersect correspond to lines in the image space that pass through both points. Therefore, straight lines in the image can be detected as a series of overlapping sinusoidals which intersect at the same values of  $\rho$  and  $\theta$  (the parameters for the line). Figure 4(e) shows the lines detected via Hough Transform in the example femur image. The second stage of Directed Hough Transform filters these lines according to a set

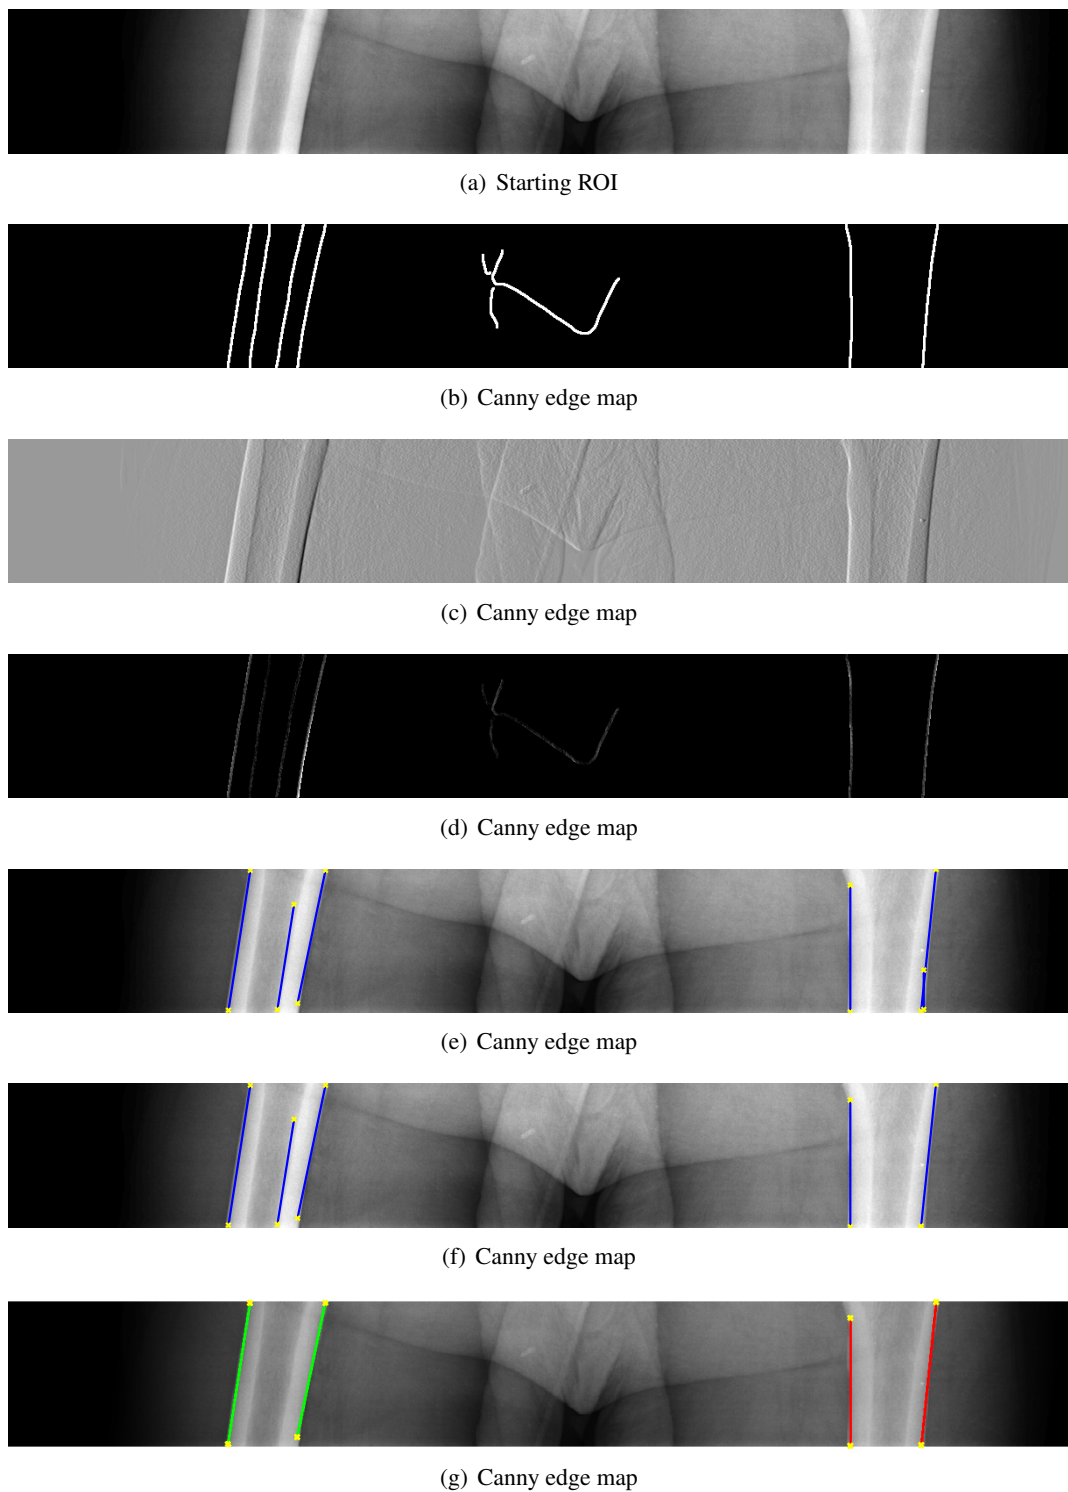


Figure 4.4: Images of the intermediate stages in femoral shaft detection, such as extraction of candidate shaft edges and pair-matching for candidate shafts..

of predefined criteria:

1. The absolute value of a line's  $\theta$  parameter must not exceed  $60^\circ$ .
2. Lines with outlying  $\theta$  values are removed.
3. If multiple lines overlap, only the longest is retained.

Condition 2 is implemented by calculating the absolute mean  $\theta$  value across all detected lines and removing those which deviate too strongly. This is used in place of a more restrictive threshold in Condition 1, as it offers more flexibility in cases where the images are rotated and the femurs appear at an angle. Despite the three conditions, the algorithm should still prove general enough to correctly detect the majority of femoral shaft edge candidates; the exception would be situations where one of the femurs is at a drastic angle, which would contravene Condition 1. This filtering step returns a set of candidate shaft edges.

#### 4.3.1.3. Pairing edges into candidate shafts

The set of candidate edges undergoes filtering according to two criteria. First, they must have the potential to represent a shaft. This is decided according to a shaft width distribution calculated across the training set, following the method developed by Chen *et al.* [15]. Femoral shaft width varies from patient to patient due to differences in build; however, the probability of two lines forming a shaft contour can be estimated based on the distance between them across a set of training images, as modeled by a Gaussian  $G_w$ . The probability  $p_i$  that the pair of lines  $i$  in the test image form a shaft contour based on width alone is given by:

$$p_i = G_w(w_i | \mu_w, \sigma_w) \quad (4.8)$$

where  $w_i$  is the distance between the two lines in the pair, i.e. the expected shaft width. Chen proposes incorporating pixel intensity information into the distribution, since it is intuitive that the intensity gradient of the leftmost shaft edge should change from dark to bright and vice versa for the rightmost shaft edge. The mean magnitude of these gradients should also be large if the line represents a true bone contour.

However, the approach proposed in this dissertation instead filters out lines with weak horizontal gradient magnitudes during initial edge map generation. The candidate shaft edge set shown in Fig. 4(f) is characteristic of most images processed using the initialization algorithm. Furthermore, the method described in [15] calculates the mean intensity gradient magnitudes only of the points along the detected line. If the bone contour is blurred or the line does not perfectly correspond to the image contour, this may cause genuine shaft edge pairs to be incorrectly rejected.

The approach in this dissertation instead evaluates intensity gradient changes via a ‘beam’ approach, which examines a specified range of columns either side of the detected line and calculates the average pixel intensity value for each. These are averaged again to give mean intensity values  $m_l$  and  $m_r$  for the left and right sides of each candidate shaft edge. Edges with the highest intensity differences, calculated as  $|(m_l - m_r)|$ , are favored as these represent the strongest change in intensity values. This can be used to decide between two candidate edge pairs that exist in close proximity, as in Fig. 4(f) (where the edge marking the transition between high density and low density bone will be rejected due to the lower intensity difference). The direction of the change is also taken into account; a line with a sharp increase in intensity from left to right can never be the rightmost edge in a shaft pair.

#### 4.3.1.4. Pairing shafts into candidate femur pairs

After filtering, the candidate edge pairs are ranked by their probability  $p_w$  (calculated using the width distribution as described in Eqn. 4.8), and the top four pairs are retained as shafts. Based on their location, these shafts are then used to build candidate shaft pairs representing the left and right femurs. In practice, there are typically only two shafts left at this point and candidate shaft pair filtering is unnecessary.

However, to ensure that the femurs are correctly identified in a situation where more than two shafts are detected, a probability distribution is constructed based on the distance between the left and right shafts in each shaft pair. Similar to the shaft width distribution defined in Eqn. 4.8, the shaft distance distribution is modeled as a Gaussian. The probability  $p_j$  that a pair of detected shafts  $j$  accurately match the shafts in the actual image is given by:

$$p_j \propto G_d(d_j | \mu_d, \sigma_d) \quad (4.9)$$

where  $d_j$  is the distance between the inner contours of the left and right femoral shafts. The shaft pair with the highest value of  $p_j$  is taken to represent the final detected shafts. The results presented in Chapter 6 include exactly how many of the test cases required this final filtering step - but generally speaking, the earlier stages of the shaft detection process are successful in filtering out all but the two actual femur shafts.

#### 4.3.2 Detection of Femoral Heads

Once the positions of the femoral shafts are known, the patient's approximate horizontal displacement is calculated as the horizontal distance between the center point of the shafts and the center of the image. This information is later used in accurate initialization of the pelvic ring Spline/ASM

model, as it deals successfully with situations where the femoral head may have been pushed inside the ring due to impact injury. However, the patient's vertical position within the X-ray image can also vary widely; in some images the shafts will cover up to half of the rows, while in others they may only be visible in the lowest few. This is one reason for the restricted region of interest used for shaft detection (as described in Section 4.3.1). Additional information is therefore required to identify the patient's approximate vertical position, which is obtained via identification of the femoral heads. Vertically, these are positioned at the base of the pelvic ring and maintain their structure even in severe injury - making detection relatively simple.

The femoral heads are roughly circular, meaning their approximate shape can be defined via a parametric equation. Consequently, directed Hough Transform is again a suitable method for detection. A circle with radius  $R$  and center  $(a, b)$  can be described with the Cartesian equation:

$$(x - a)^2 + (y - b)^2 = r^2 \quad (4.10)$$

which can then be written in the following parametric forms:

$$x = a + R \cos(\theta) \quad (4.11)$$

$$y = b + R \sin(\theta) \quad (4.12)$$

where  $\theta$  represents the angle that the vector from  $(x, y)$  to the origin makes with the  $x$ -axis.

For any single pixel  $(x_i, y_i)$ , these equations are variable in three dimensions:  $a$ ,  $b$ , and  $R$ . Each edge pixel in the original image therefore corresponds to a cone in the 3-D Hough parameter space (as can be seen by rearranging Eqn. 4.10). Similar to the Hough line transform, edges are defined by the overlap of their representations in parameter space; pixels belong to the same circle in image space if their cones share a common intersection point.

However, Hough Circle Transform is computationally expensive, due to the higher dimension of the parameter space. To ensure that circle detection can be performed with sufficient speed in the chosen application, several restrictions are used. The most intuitive is the restriction on the  $(a, b)$  search space to a limited window around the femoral shafts. This varies depending on the image, but its width  $w_w$  and height  $w_h$  can be defined as:

$$w_w = 3w_s; w_h = 2I_h; \quad (4.13)$$

where  $w_s$  is the pixel width of the detected femoral shaft at the base of the image and  $I_h$  is the total pixel height of the image. These constraints on window size were determined via experimentation and are relatively lax, but allow for considerable variation in patient position. The computational cost of the Hough transform can be further reduced by restricting the circle radii search range. Though multiple searches are conducted across this range,  $R$  is fixed during each individual search. Only  $a$  and  $b$  must be found, and so each search occurs in 2-D parameter space. The radii search range  $S_r$  covers  $(0.5w_s - 1.5w_s)$ , again defined in terms of pixels. Searches in images with wider shafts will consequently involve a wider range and therefore require more time; however, the constraints still reduced the search time to an average of 2.5 seconds across the test set, with the longest search taking 3.3 seconds (both rounded to 1 d.p.). The search is performed on an edge map generated using the Canny detection operator, as described in subsection 4.3.1. In this situation, the high and low thresholds used in non-maxima suppression are less restrictive, since the femoral heads are not as well-defined as the shafts. An example result of femoral head detection is presented in Fig. 4.5, where it can be seen that the approximate position of the femoral head is correctly identified.






Figure 4.5: The result of applying Canny edge detection to the area of the image containing the left femoral head, then applying Hough Transform to the edge-detected image.

### 4.3.3 *Spline/ASM for Femur Detection*

The approximate positions of the femoral heads are then used to place the Spline/ASM shape models for both the left and right femurs. In turn, the located femurs are used to initialize placement of the models for the pelvic ring and the left and right pubis-ischium structures. A simple method of placing the femur shape model would be to use the circumference of the circle approximately covering the femoral head. However, since the Hough Transform is prone to overestimating the radius of the head due to false edges in the search window, this circumference may be inaccurate. Instead, the model is placed based on the circle center; this requires extra calculations, but offers a reasonably stable reference point. The change in overall shape is limited to  $4\sigma$ , allowing the model to deform to fit femurs at various angles without losing its base structure. Image subsampling using the Gaussian pyramid caused distortion in detection results, and was therefore not used.

Experiments were performed using two different landmark placement schemes for each structure. In both schemes, landmark points were placed at roughly equal intervals around the structure edge; however, the second used only half the points of the first. This was repeated for all structures, but for reasons outlined in Chapter 6, the increased number of landmarks did not offer significant benefits. Figure 4.6 presents an example of femur detection using ASM with the reduced landmark scheme.

### 4.3.4 *Spline/ASM for Pelvic Ring Detection*

Identifying the femurs enables correct initialization of pelvic ring and pubis-ischium models. Across all training images, it was found that the distance between each femoral head and the pelvic ring (proportional to the image width) fell within a set range. Furthermore, the boundary of

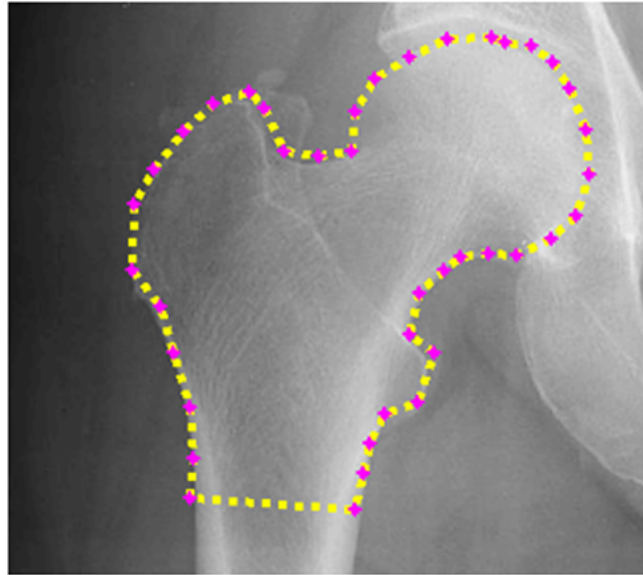


Figure 4.6: The result of Spline/ASM detection of the right femur

the pelvic ring is marked by a transition from bright bone matter to darker soft tissue. Therefore, a horizontal ‘beam’ search is conducted for sudden changes in pixel intensity, beginning at the uppermost landmark on each femur model (located near the top of the femoral head). This search covers  $I_w/10$  columns and  $H_r$  rows, where  $I_w$  is the pixel width of the image and  $H_r$  is the pixel radius of the detected femoral head. The Spline/ASM ring model is then placed by matching a set number of landmarks along the points with the steepest gradient changes. This approach deals reasonably well with situations where the femoral head is pushed inside the pelvic ring. Note that prior detection of the femurs is vital for locating the pelvic ring; automatic segmentation cannot begin with the ring since its horizontal and vertical position in the x-ray image is unknown. Deformation of the ring shape model is again limited to  $4\sigma$ , but the search range along the profile to each point is reduced to prevent the model aligning to the upper edge of the obturator foramen (the gap below the superior pubic ramus underneath the pelvic ring). Two pyramid levels are used to subsample the image, since the simple shape of the ring remains intact at the lower image

resolution.

#### 4.3.5 *Spline/ASM for Pubis and Ischium Detection*

During segmentation, the ischium and pubis on each side are treated as a single structure. Unlike the other structures, in this case the Spline/ASM model consists of two separate contours; one marking the outer edge of the bones and the other the obturator foramen. Since the pubis-ischium structure shares a common edge with the pelvic ring, the landmarks detected in the previous step are used in placement of the shape model. Furthermore, deformation from the points along the upper edge is penalized more severely. The center of the femoral head is also treated as a reference point.

#### 4.3.6 *Hough Transform and Spline/ASM for Detection of Iliac Bones*

The final structures to be detected are the left and right iliac bones. In a standard AP view pelvic X-ray, these appear approximately circular, and therefore Canny edge detection and directed Hough transform can again be used to identify their approximate location. The search windows are defined as the top left and top right quadrant of the image, with detection and segmentation being performed separately for each of the two sides. The range of radii to be searched is harder to define. Following examination of the dataset, it is set as  $(0.5h_p - 1.5h_p)$ , where  $h_p$  is the distance between the lowest and highest landmark points on the detected pelvic ring. Since the iliac bones are much larger than the femoral heads, the computational cost of the search is much higher - however, this is offset by testing radii at 2-pixel intervals (i.e.  $r = \dots, 60, 62, 64, \dots$ ), effectively halving the range.

As with femur detection, the iliac shape models are placed according to the center of the located circle to avoid the impact of radius over-estimation. Based on the results of a series of

experiments, shape deformation is limited to  $4\sigma$ , since although the edges of the iliac bones are fairly distinct the model is still vulnerable to image noise and the false edges caused by the iliac crests. An example result is presented in Fig. 4.7.


Note that if the patient is rotated on the table, one bone may appear smaller in size or hooked rather than circular. This may cause errors in detection. However, at least one bone maintained a circular appearance across all images in the dataset; the range of patient pose angles present within AP pelvic X-rays does not seem sufficient to cause simultaneous distortion of both sides. Since the pelvis should maintain bilateral symmetry [10], it is intuitive that if the approximate location of one bone can be identified, this can be used to estimate the position for the opposite side. However, since severe trauma can result in pelvic asymmetry [30], absolute symmetry cannot be assumed when placing the model. Instead, a window is identified that is likely to contain the non-detected iliac bone and an edge map is created to place the Spline/ASM model. This is the final stage of the hierarchical pelvic segmentation process.

#### **4.4 Summary**

This chapter presents a novel hierarchical method for automatic segmentation of key pelvic structures. Specifically, it combined anatomical knowledge and multiple computational techniques to deal with the high sensitivity of deformable models to their initialization. Starting with detection of the femoral shafts, the information obtained at each step of the segmentation process is then used to begin the next. When applied with the combined Spline/ASM algorithm described in Chapter 3, this method offers highly accurate detection of key pelvic structures, as supported by the results presented in Chapter 6.



(a) Result of edge detection



(b) Result of circular Hough transform

Figure 4.7: An example of automatic detection of the approximate location of the ilium, via edge detection and circular Hough Transform.

## CHAPTER 5 Fracture Detection and Displacement Measurement

### 5.1 Overview

This chapter explains the wavelet-based method developed for fracture detection in the pelvic ring and pubis/ischium structure, to be applied after the segmentation process described in Chapter 4. It also explores the importance of injury patterns in successful treatment of pelvic trauma, and the diagnostic signs that can be extracted from AP view radiographs of the pelvis. Section 5.2 presents an overview of the Discrete Wavelet Transform (DWT), which forms part of the fracture detection algorithm described in Section 5.3. A method of automatically calculating quantitative measures of pelvic displacement is explained in Section 5.4, while Section 5.5 describes key patterns of pelvic injury and how the features extracted can be used in determining the basic form. A summary of the chapter is presented in Section 5.6.

### 5.2 The Discrete Wavelet Transform

Understanding the developed method of fracture detection first requires an understanding of the wavelet transform (WT), which transforms an input signal into a time-frequency representation. This may be more compact or assist in signal analysis. The one-dimensional continuous wavelet transform (CWT) of an input signal  $x(t)$  is obtained by convolving it with a set of wavelet functions. These wavelets are shifted and scaled versions of some mother wavelet  $\psi(t)$  and are therefore known as daughter wavelets. Calculating the CWT involves calculating the correlations between these daughter wavelets and local portions of  $x(t)$ . Choosing the best mother wavelet for a given application is an open problem, but a useful heuristic is to pick one similar to the input signal.

The normalized daughter wavelet for given scale and shift parameters  $a$  and  $b$  can be written as

$$\psi_{a,b}(t) = \frac{1}{\sqrt{|a|}} \psi\left(\frac{t-b}{a}\right) \quad (5.1)$$

The WT is then expressed as

$$T(a, b) = \int_{-\infty}^{\infty} x(t) \psi_{*a,b}(t) dt \quad (5.2)$$

where  $*$  represents the complex conjugate operator. Equation 5.2 is equivalent to the convolution of the wavelet function and the signal  $x(t)$ . In the CWT, the mother wavelet is continuously shifted and scaled over  $x(t)$ , which is not feasible in a digital setting and also carries a high level of redundancy. The Discrete Wavelet Transform (DWT) circumvents this issue by accepting continuous signals as input but only applies discrete shift and scale values. If a suitable set of values is used to sample  $x(t)$ , then the inverse DWT can be used to reconstruct the original continuous signal in full. The DWT provides a means of multi-resolution analysis (MRA), applying a scaling function  $\phi(t)$  to create a series of approximations of the input signal, each differing in factor 2 from its nearest neighbors. The wavelets encode the difference in information (i.e. the detail) between neighboring approximations.

Let  $m$  and  $n$  be integer values controlling scale and shift respectively. The scaling function is given by

$$\phi_{m,n}(t) = 2^{-\frac{m}{2}} \phi(2^{-m}t - n) \quad (5.3)$$

where

$$\phi(t) = \sum_k c_k \phi(2t - k) \quad (5.4)$$

Equation 5.4 is the scaling equation, where  $\phi(2t - k)$  is a contracted and shifted version of  $\phi$



and  $c_k$  are the scaling coefficients. The scaling function  $\phi_{m,n}$  is associated with smoothing of the input, and so is convolved with  $x(t)$  to create the approximation coefficients  $S_{m,n}$ :

$$S_{m,n} = \int_{-\infty}^{\infty} \psi_{m,n}(t) dt \quad (5.5)$$

The set of values of  $S_{m,n}$  for a specific scale  $m$  form the discrete approximation of  $x(t)$  at that scale (or level).

A discretized version of the wavelet function for scale and shift values  $m$  and  $n$  has the form

$$\Psi_{m,n}(t) = \frac{1}{\sqrt{|a_0^m|}} \Psi\left(\frac{t - nb_0 a_0^m}{a_0^m}\right) \quad (5.6)$$

$$= a_0^{-\frac{m}{2}} \Psi(a_0^{-m} t - n) \quad (5.7)$$

where  $a_0$  and  $b_0$  are fixed scale and shift step parameters used in discretizing the parameters  $a$  and  $b$  used in Eqn. 5.1. The discrete wavelet transform of  $x(t)$  can now be written as

$$T_{m,n} = \int_{-\infty}^{\infty} \psi_{m,n}(t) dt \quad (5.8)$$

In this case,  $T_{m,n}$  form the set of detail coefficients;  $x(t)$  can be expressed by adding the approximation of  $x(t)$  at some arbitrary level  $m_0$  to the sum of the detail coefficients created at levels  $m_0$  down to  $-\infty$ . If the signal detail at level  $m$  is written as

$$d_m(t) = \sum_{-\infty}^{\infty} T_{m,n} \Psi_{m,n}(t) \quad (5.9)$$

then  $x(t)$  is given by

$$x(t) = x_{m_0}(t) + \sum_{m=-\infty}^{m_0} d_m(t) \quad (5.10)$$

This is important in understanding the signal decomposition process used in MRA via DWT.

Assume we have an input signal  $x(t)$  which can be written as  $S_{0,n}$  (i.e. the approximation of itself

at level 0). We compute  $S_{1,n}$  and  $T_{1,n}$  as

$$S_{1,n} = \frac{1}{\sqrt{2}} \sum_k c_k S_{0,2n+k} \quad (5.11)$$

$$T_{1,n} = \frac{1}{\sqrt{2}} \sum_k b_k S_{0,2n+k} \quad (5.12)$$

where  $c_k$  and  $b_k$  are the scaling coefficients used for the scaling and wavelet functions respectively. In the same way,  $S_{2,n}$  and  $T_{2,n}$  can then be calculated using  $S_{1,n}$ . At each level  $m$ , the approximation created at level  $(m - 1)$  is decomposed using Eqns. 5.11.

The one-dimensional DWT can be extended to 2-D DWT for application to 2-D signals such as images (where the signal is formed by the image's array of pixels). This requires 2-D scaling and wavelet functions which take their parameters from the pixel positions  $(x, y)$ . Each level of decomposition  $m$  creates a decomposition array which can be divided into four sub-matrices:  $\mathbf{S}_m$  (containing the approximation coefficients) and  $\mathbf{T}_m^h$ ,  $\mathbf{T}_m^v$  and  $\mathbf{T}_m^d$  (containing the horizontal, vertical and diagonal detail coefficients respectively). The detail coefficients can be used to reconstruct discrete detail arrays at the scale of the input image, defined as  $D_m^h$ ,  $D_m^v$  and  $D_m^d$ . The sum of these arrays forms a combined detail array. Remember that each level of decomposition in 1D-DWT involved down-sampling the input by factor 2; the same is repeated in 2D-DWT in both dimensions, hence the creation of four sub-matrices. In other words, if the input image is of size  $(i \times j)$ , the approximation created at the first level of decomposition is of size  $(\frac{i}{2} \times \frac{j}{2})$ .

Mallat defined a method of calculating DWT decompositions via a succession of quadrature mirror filters. Figure 5.1 shows how this approach can be applied to images. Decomposition first treats the rows of the image as a 1-D signal, then the columns;  $g[n]$  and  $h[n]$  are the quadrature mirror filters, where  $h[n]$  is a low-pass filter and  $g[n]$  is the mirror filter of  $h[n]$  (and therefore a high-pass filter). The process creates an approximation component and three detail components,

as described in the previous paragraph, which can be reconstructed into arrays at the scale of the input image.

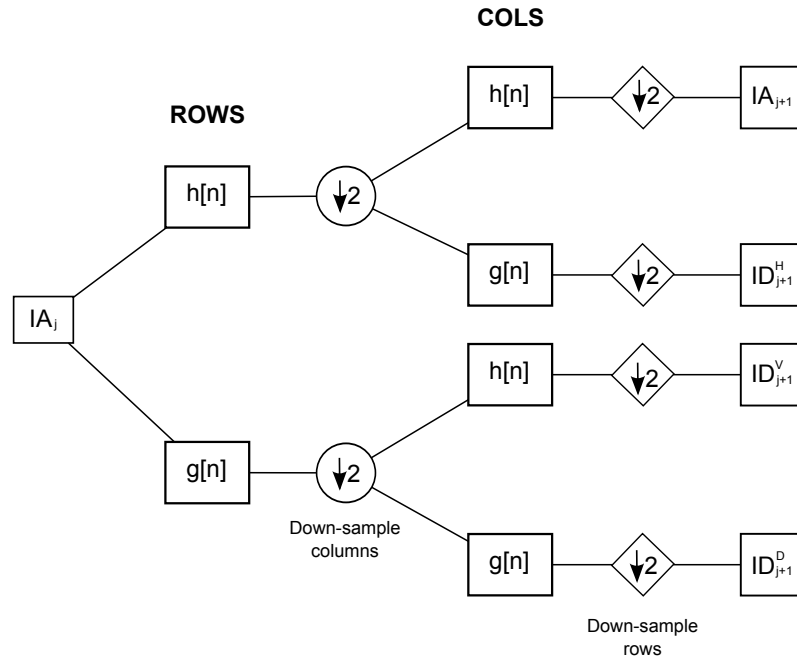


Figure 5.1: A single 2D DWT decomposition stage. A low-pass filter  $h[n]$  and high-pass filter  $g[n]$  are first applied to the rows of the image, then the columns of the output. Down-sampling by 2 occurs after each step.

### 5.3 Ring Fracture Detection

This dissertation presents an automated method for detecting fractures around the pelvic ring, which creates a series of overlapping windows based on the boundaries extracted during initial Spline/ASM segmentation and tests each one for contour discontinuities which may indicate fracture. The method is also applicable to identifying horizontal fractures of the pubic rami, ischium and obturator foramen. In particular, the left and right upper pubis form part of the stable ring structure, and so fractures in that area are detected during analysis of the pelvic ring structure.

Figure 5.2 presents an overview of the fracture detection process.

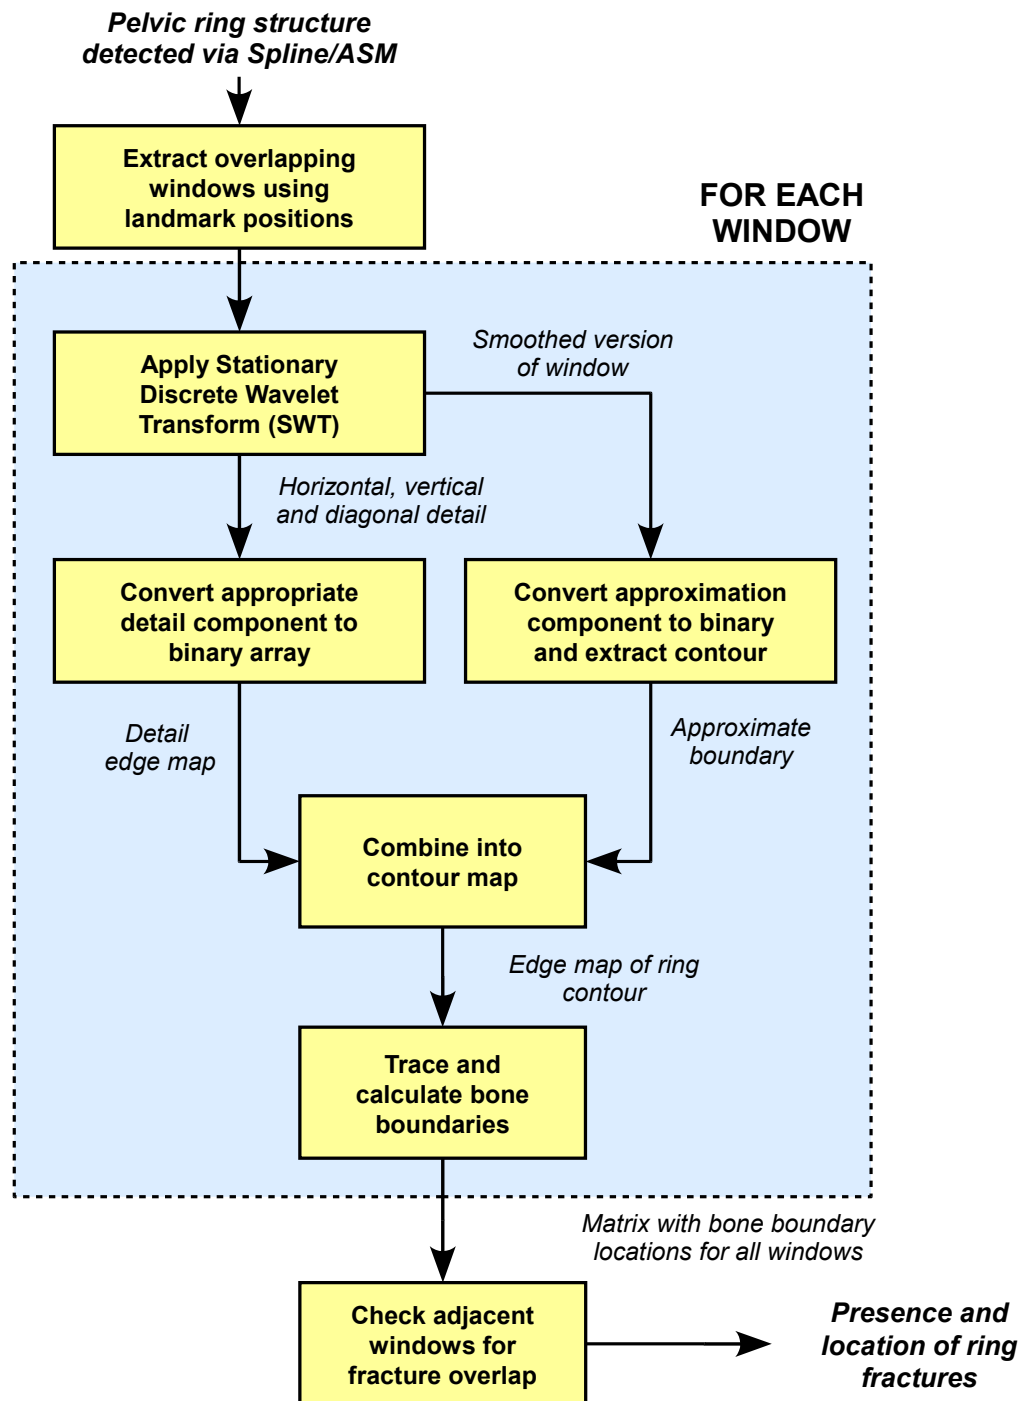


Figure 5.2: A schematic diagram of the pelvic ring fracture detection process.

### 5.3.1 Windowing

The appearance of ring fractures in an AP pelvic radiograph depends on their severity. Unlike major fractures, minor fractures may not severely distort the edge of the ring; instead they may appear as dual edges, or a single edge segment that is slightly blurred compared to its neighbors. While the Spline/ASM model created during the hierarchical segmentation process (described in Chapter 4) is capable of detecting large distortions in the ring structure, it is not sufficiently sensitive to subtler changes. The first step in fracture detection is therefore to refine the boundary of the pelvic ring and obtain a more detailed view of any discontinuities. This is performed via a series of wavelet transforms for noise reduction and edge detection; however, applying these transforms to the entire ring is both impractical and unlikely to achieve good results due to local intensity variations. Instead, the ring boundary is divided into a series of overlapping windows whose size and location are determined by the coarse boundary detected using the Spline/ASM model. Each window covers a subsequence of consecutive model landmarks; the length of this subsequence is decided during initial Spline/ASM training, as the optimal value depends on the number of landmarks used to describe the models. Although using non-overlapping windows would reduce the total number of windows to be processed and hence lower the computational cost, fractures crossing the boundaries between windows may not be detected. As shown in Fig. 5.3, each window has an overlap region at either end set to a specific number of landmarks (again decided during training). Since individual windows are defined in terms of landmark sequences, their pixel dimensions in the unseen image under analysis are dependent on the results of Spline/ASM model fitting. Sizes are therefore not equal between one window and the next, but this does not appear to negatively impact performance, and all sections of the ring edge are still processed. After windowing has been performed, three steps - wavelet transform, masking and boundary tracing - are

performed on each individual window. The windowing method is also applicable to the contours of the obturator ring and the lower edge of the ischium, since these structures are also outlined by landmark points after automated Spline/ASM segmentation.




Figure 5.3: Example overlapping windows around the pelvic ring, positioned according to model points placed during Spline/ASM segmentation.

### 5.3.2 Wavelet Transform

The input window  $W$  is first decomposed using the 2-D Stationary Discrete Wavelet Transform (SWT). SWT is a redundant, translation-invariant version of DWT which does not decimate the coefficients at every level of decomposition, and where the filters at level  $i$  are up-sampled versions of those at level  $(i-1)$  (where  $i > 0$ ). As with the 2-D DWT, decomposition outputs approximation and horizontal, vertical and diagonal detail coefficients. In this application, three levels of decomposition are applied to window  $W$  using the *db1* (or Haar) wavelet. First, the level 3 approximation coefficient is used to reconstruct a smoothed version  $W_s$  of the input window  $W$ , as shown in Fig.

5.4




Figure 5.4: Reconstruction using the SWT approximation coefficients generates a smoothed window image.

The level 3 detail coefficients,  $\mathbf{T}_3^h$ ,  $\mathbf{T}_3^v$  and  $\mathbf{T}_3^d$ , are then extracted and used to reconstruct detail arrays  $D_h$ ,  $D_v$  and  $D_d$ . Figure 5.5 presents example images.




Figure 5.5: Reconstructions of a ring window from vertical, horizontal, and diagonal SWT detail coefficients

The reconstructed detail arrays highlight high frequency variations in the image - including the edge of the ring. The coefficient set that best describes the edge depends on the angle of the edge, and therefore on the location of the window around the ring. It can be seen in Fig. 5(b) that

the vertical detail coefficients most clearly define the edge in this case. In practice the choice is handled automatically during the initial windowing process, as each window is labeled according to the set of coefficients most likely to be suitable for ring contour detection. The labeling scheme is defined in advance based on prior examples to avoid the need for manual interaction by the user.

### 5.3.3 Masking and boundary tracing

The next step creates a binary version  $D_B$  of the chosen detail array (here, the vertical reconstruction shown in 5(b)). This will contain the pelvic ring contour as a white line but will also include other extraneous edges. To filter these out, a mask  $A_B$  is formed by converting the smoothed window  $W_S$  to a binary image using Otsu's threshold [49], then extracting the contour (which approximates that of the ring). By combining  $A_B$  and  $D_B$ , the unwanted edges are removed to create an edge window  $W_E = A_B \times D_B$ . Figure 5.6 shows each of these windows.

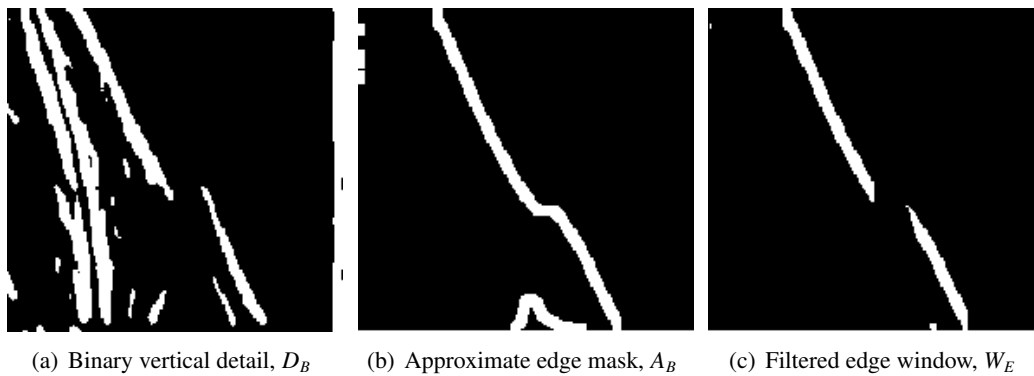


Figure 5.6: The filtering process used to create the edge window,  $W_E$

The final task is to trace the edges in  $W_E$  and detect any discontinuities. First, morphological opening is applied to all objects in the image with area below a specific threshold. This was pre-defined in testing as 0.5% of the window area, but the most suitable value must be decided during training based on the number of landmarks used to generate each window. Thresholding removes



small artifacts that may interfere with edge tracing. Finally, the remaining edges are traced using the 8-neighborhood of each pixel and returned as a matrix of pixel positions. The traced edges represent the bone contours of the pelvic ring; if there is no fracture, the window will therefore contain a single uninterrupted boundary. Otherwise, there will be multiple boundaries as shown in Figure 5.7, depending on the types and number of fracture. Currently, fracture presence is a binary value (i.e. present or not present). Future work will examine a finer level of fracture classification based on location and severity, which may provide more useful information in rule generation.

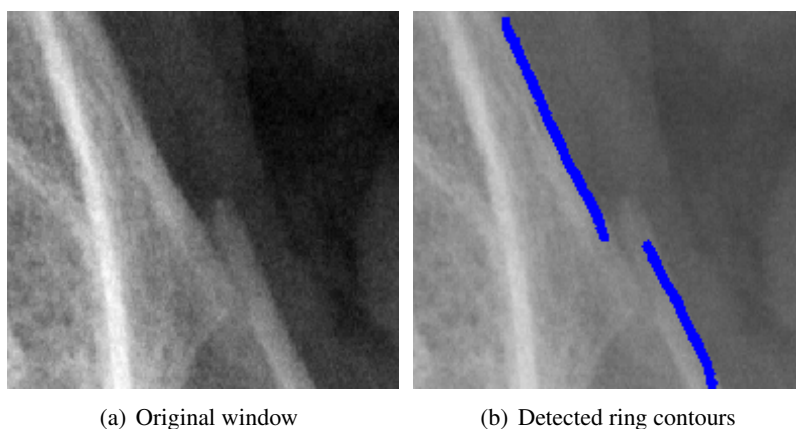


Figure 5.7: Example of a detected broken ring boundary, which may indicate a ring fracture.

#### 5.4 *Symphysis pubis displacement*

Discussion with radiologists and a review of the literature has indicated that the horizontal and vertical displacement of the left and right pubis may be important quantitative measurements in determining pelvic injury severity. At the time these measurements are calculated, it is assumed that the left and right pubis have been detected using the Spline/ASM algorithm as outlined in Chapter 3. This returns a detected object as a sequence of consecutive landmarks around its boundary, in  $(x, y)$  coordinate form. Two displacement measures between the left and right pubis can therefore

be calculated based on the positions of these landmarks; specifically, those on the edges bridged by the symphysis pubis.

The ‘Horizontal Gap’ feature is defined as the horizontal distance between the left and right pubis. Figure 5.8 describes visually how this is determined. The number of landmarks on the edge of the gap will vary depending on how the Spline/ASM model was trained prior to the system being deployed. However, it will always be equal for both the left and right sides.

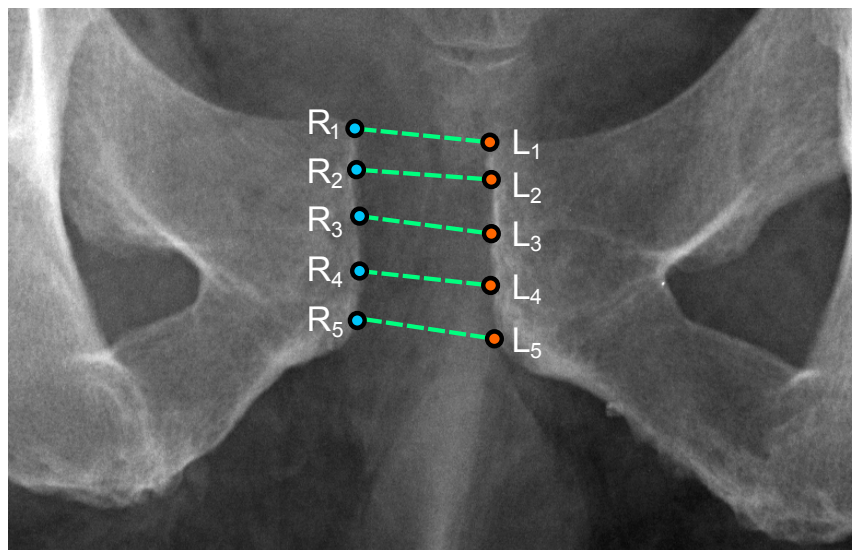


Figure 5.8: Calculating the width of the pubis symphysis gap using Spline/ASM landmarks.

When measuring displacement, the algorithm considers the landmarks as pairs  $(L_i, R_i)$ , and calculates the difference in x-coordinates between the pair members. The differences are summed and averaged over the number of landmarks  $n$  (in Fig. 5.8,  $n = 5$ ). If  $L_i = [x_k, y_k]$  and  $R_i = [x_j, y_j]$  form a landmark pair where  $i = 1, \dots, (n)$ , let  $L_i[x]$  and  $R_i[x]$  represent the  $x$  coordinates of  $L_i$  and  $R_i$  respectively. The horizontal symphysis gap  $g_H$  is calculated as:

$$g_H = \frac{1}{n} \sum_{i=1}^n |L_i[x] - R_i[x]| \quad (5.13)$$

Figure 5.9 shows how vertical pubis displacement is determined; to save space, landmarks 2-4 on each side are not labeled, but are still used in calculation. Where  $L_i$  and  $R_i$  are defined as above and  $L_i[y]$  and  $R_i[y]$  represent their  $x$  coordinates, vertical displacement  $g_V$  is calculated as:

$$g_V = \frac{1}{n} \sum_{i=1}^n |L_i[y] - R_i[y]| \quad (5.14)$$




Figure 5.9: Calculating the vertical displacement of the left and right pubis using Spline/ASM landmarks.

Since the sides of the symphysis gap appear as vertical edges in the image, the model deformation process can cause variations in vertical landmark spacing. This is because the vertical distance between consecutive landmarks can increase and decrease without a significant overall effect on segmentation, whereas horizontal landmark movement converges to the edge. The vertical differences in landmark position (i.e. differences in  $y$  coordinates) can be safely omitted when calculating the horizontal gap (as in Eqn. 5.13) but are crucial to calculating the vertical displace-

ment. However, the use of averaging over all landmarks on the edge (as in Eqn. 5.14) should compensate for these variations. The alternative - using only landmark pairs  $(L_1, R_1)$  and  $(L_n, R_n)$  - is vulnerable to inaccuracies in the model deformation process.

Equations 5.13 and 5.14 return measurements in terms of pixels. Since the area of the radiograph occupied by the pelvic structure can vary from one image to the next, these values are normalized by the average length of the diagonals crossing the pelvic structure (each defined by two reference points on the femoral head and the opposite iliac bone). This approach was used in the experiments described in Chapter 7, as a method of measuring distances in centimeters was unavailable.

However, this situation has changed as more DICOM (Digital Imaging and Communications in Medicine) format images have become available. The DICOM standard specifies key aspects of medical imaging, such as handling of information and transmission protocols. It also defines a file format to be used in digital image storage which incorporates not only the image data itself but also related text and numerical information. A DICOM file is formed as a single data object consisting of multiple attributes - such as the patient ID, the data the scan was collected, and the imaging modality - where the final attribute contains the image pixel data. The *Pixel Spacing* attribute contains the physical distance in the patient corresponding to the distance between the centers of two immediately neighboring pixels in the image. The radiograph processing component currently contains a function which extracts the Pixel Spacing attribute and uses it to convert the values returned in Eqns. 5.13 and 5.14 to physical distances in the patients, expressed in millimeters. Not only is this form more easily understood in a medical context, but it also prevents small inaccuracies in segmentation of the iliac bones and the femurs from affecting calculation of normalized pixel distances. Future experiments will therefore use the millimeter form in rule

generation. The DICOM format offers other benefits that may prove useful in future development of the work in this dissertation. For example, images are stored at 12-bit depth rather than the 8-bit depth used in common file formats such as PNG and JPEG. This can potentially provide better definition of detail; however, it may not have much impact in low-resolution X-ray images. Testing is needed to determine whether using the 12-bit form will be beneficial. For ease and speed of processing, the work presented in this dissertation down-samples the images to 8-bit depth. More information on the DICOM standard can be found in [44].

### **5.5 Patterns of Pelvic Injury**

As stated in the introductory chapter to this dissertation, identification of pelvic fracture patterns is vital in early stabilization [12, 23]. Numerous classifications of these patterns have been proposed [10]. One of the mostly widely used is based on the mechanism of injury, as this assists physicians and surgeons in applying the most appropriate external fixation force [75]. This scheme separates pelvic trauma by mechanism into three main types: anterior compression, lateral compression, and vertical shear.

Anterior compression (AP) injuries involve a force applied to the pubis or posterior pelvis, and are common in motor vehicle collisions (MVCs). Since it typically involves external rotation of the ilium, the AP pattern is also referred to as ‘open book’. Diagnostic signs include vertical pubic rami fractures and diastasis of the symphysis pubis and sacroiliac joints. Three AP sub-types have been defined, according to injury severity; the most severe involves disruption of all the sacroiliac ligaments and separation of the iliac wing from the sacrum. Lateral compression (LC) injuries are the most common, caused by a lateral blow to either side of the pelvis resulting in internal iliac rotation. Again, they are classified into three sub-types, but general diagnostic

features include: sacral fracture, horizontal or coronal pubic rami fracture, iliac wing fracture, diastasis of the sacroiliac joints, and overlapping fractures of the pelvic ring. The final group are vertical shear (VS) injuries, such as those due to falls from a height. Signs include symphysis pubis and sacroiliac joint diastasis, iliac and sacral fractures, and vertical fractures of the pubic rami. There are more complex injury patterns, as well as other classification systems, but these may require more images for full analysis (e.g. CT scans and/or inlet and outlet X-ray views).

The diagnostic signs in this classification scheme include several that are visible on AP radiographs. Detection of pelvic ring and pubis fracture has been described in Chapter 5. Fracture of the ilium and sacrum are more challenging to detect on AP pelvic X-rays and are left as future work. However, all three injury types involve some degree of ligament damage; more specifically, pelvic instability can be defined by two types of displacement, rotational and vertical [65]. In an AP radiograph of the pelvis, this can be evaluated by comparing the positions of the left and right bones of the pubis as well as identifying potential abnormalities around the sacroiliac joints. An axial CT scan offers a clearer view of displacement, particularly concerning the sacroiliac ligaments; however, in situations where a scan cannot be obtained, features extracted from X-ray images may provide useful diagnostic information and should therefore be included as input to the decision-making process.

The criteria used to define these various patterns suggests that the calculated quantitative displacement measures may prove useful in pattern identification. As stated above, diastasis of the symphysis pubis may indicate severe injury, and is quantified by the 'Horizontal Gap' measure. Vertical displacement can potentially indicate a vertical shear injury, and is quantified by the vertical displacement measure. Additional work not presented in this dissertation has focused on detecting diastasis of the sacral joints; as stated, this is associated with multiple patterns of injury

and pelvic instability. The method is not yet robust enough for evaluation, but shows promise - particularly if it can be verified via analysis of corresponding CT scans.

### **5.6 Summary**

This chapter describes the extraction of potentially useful diagnostic features after segmentation of key pelvic structures. Specifically, it presents an algorithm developed to detect fracture of the pelvic ring and horizontal fracture of the pubis that can be adapted to other regions of the body, and potentially other X-ray analysis tasks involving discontinuity detection (such as in industrial applications). It also presents a method for automatic calculation of displacement measures based on initial segmentation results, and suggests how these features may be useful in a future system component that identifies injury patterns.

## CHAPTER 6 Results of Image Processing

### 6.1 Overview

This chapter presents the results of multiple tasks included in the developed X-ray analysis module. Segmentation is covered in Section 6.2 (including ANOVA evaluation to test the effectiveness of manual initialization), and fracture detection of the pelvic ring in Section 6.3 fracture detection of the pelvic ring. Example images are provided of all results. Section 6.5 demonstrates how the Spline/ASM algorithm can be used in other medical imaging applications, and the chapter ends with a discussion of the combined results. Note that the results of predictive model generation are addressed separately in Chapter 7.

### 6.2 Segmentation

Since the purpose of the hierarchical automatic segmentation algorithm is to automatically initialize detection of key pelvic structures, performance is compared versus manually initialized segmentation. A second set of examples was generated by a user manually placing the model template within the desired structure in the x-ray image. Both models were trained using Spline/ASM. Across a set of 20 images taken from male pelvic trauma patients, three key structures - the right ilium, right femur and pelvic ring - were manually labeled and taken as reference shapes. For each image, the difference in area between this reference shape and the shape detected via the initialization algorithm was calculated and normalized by the area of the reference shape. In other words, where  $A$  is the reference shape (i.e. a set of pixels) and  $B$  is the shape detected automatically, the



normalized error measure  $e$  is calculated as:

$$e = \frac{(A - B) \cup (B - A)}{\text{area}(A)} \quad (6.1)$$

The same error measure was calculated for the reference shape and the shape detected via manually-initialized ASM. ANOVA was then performed for all three structures to determine whether there is a significant difference between results obtained via manual initialization and results obtained using the hierarchical automatic initialization algorithm. When comparing results for pelvic ring detection, ANOVA calculates a p-value of 0.8431. For iliac crest detection,  $p = 0.0776$ , and for femur detection  $p = 0.6078$ . These p-values indicate no significant difference between results for manual initialization and the automatic method. It should be noted that in all three cases, the sum of normalized error measure for the automated method is less than the sum for manual initialization. This is particularly noticeable in detection of the right ilium, explaining the lower p-value for that ANOVA test. Figures 6.1, 6.2 and 6.3 present ANOVA box-plots for right iliac bone detection, right femur detection and pelvic ring detection respectively. It can be seen that the automatic initialization offers slightly more consistent results with lower variance.

Via visual inspection, detection of each structure was classified across the fifteen test images into three categories: Good, Acceptable, and Unacceptable. These categories were determined via consultation with a trauma physician, and are subjective. However, the algorithm is trained on a set of human-labeled images, and since the aim is to mimic manual segmentation, visual inspection can arguably be considered a gold standard. Results of visual inspection therefore provide useful feedback on the current performance of the algorithm. Results are presented in Table 6.1. The crests, femur and pelvic ring are almost always detected to at least an acceptable standard and segmentation is typically rated as good. The exceptions are the left and right pubis




Figure 6.1: ANOVA box-plot for detection of right iliac bone.

bones; this may be due to the greater degree of deformation that can occur in severe pubis fracture, particularly those that are vertical. It is expected that this can be resolved by training on a larger set of images which better captures plausible shape variations.

Table 6.1: Results of full segmentation testing via visual inspection.

| Structure        | Good | Acceptable | Unacceptable |
|------------------|------|------------|--------------|
| Left Iliac Crest | 11   | 4          | 0            |
| Right Iliac      | 12   | 2          | 1            |
| Pelvic Ring      | 10   | 5          | 0            |
| Left Pubis       | 9    | 1          | 5            |
| Right Pubis      | 10   | 1          | 4            |
| Left Femur       | 12   | 3          | 0            |
| Right Femur      | 11   | 3          | 1            |

Figures 6.4 and 6.5 present two sample images processed using the automatically initialized




Figure 6.2: ANOVA box-plot for detection of right femur.

Spline/ASM segmentation. It can be seen by visual inspection that all structures of interest are successfully detected. Figures 6.6 shows a more problematic case where the left pubis is not correctly segmented. In situations such as this, symmetry checks can detect failed segmentation.

To illustrate individual performance of Spline/ASM, Fig.6.7 presents two sample results comparing the performance of standard ASM with Spline/ASM in segmenting individual structures. The pelvic ring and right ilium are both structures with strongly curved contours and so are likely to benefit from Spline/ASM. The difference is most pronounced in iliac bone segmentation, as shown in Fig. 7(c) and Fig. 7(d). To compare the time required for Spline/ASM segmentation versus standard ASM, performance across 12 images was compared in pelvic ring segmentation, with each image being tested twice to account for variations due to other processes running on the computer.




Figure 6.3: ANOVA box-plot for detection of pelvic ring.

The significance of the number of landmarks used in defining shape models was also explored. The right femur, right pubis, right iliac bone and pelvic ring were all assigned two separate landmark schemes, with the second using approximately double the landmarks of the first. Each additional landmark was placed at equal intervals between pairs of original landmarks. It was found that the average processing time increased to between 150-200% of the reduced landmark scheme, with no significant improvement in segmentation results. Although contours were more precisely defined, the goal of segmentation in this application is to identify structures for further analysis, such as in fracture detection. Since the final decision-making system is intended for use in trauma care, increased processing time is unacceptable if there is no significant benefit to analysis.




Figure 6.4: An example of successful full automatic segmentation using Spline/ASM.




Figure 6.5: A second example of successful full automatic segmentation using Spline/ASM.




Figure 6.6: A example of only partially successful full automatic segmentation using Spline/ASM, showing incorrect detection of the left pubis bone.




Figure 6.7: Comparing Standard ASM with Combined Spline/ASM in detection of pelvic ring and right iliac bone.



### 6.3 Fracture Detection

Statistical evaluation of fracture detection is difficult at this stage. At present the algorithm detects abnormalities that may indicate the presence of fracture, but further input from experts is required to determine which of these are genuinely fractures. For example, it is necessary to distinguish between severe ring fractures and those that may be caused by overlapping structures or X-ray artifacts; for now, such cases are considered only as potential abnormalities. Since there are multiple types of fracture, full evaluation also requires a larger dataset containing more examples of each type. However, preliminary results of pelvic ring fracture detection are presented in Table 6.2. Twenty-four extracted windows were considered, equally divided into three groups (Suspected Minor Fracture, Suspected Larger Abnormality, and No Fracture) . As stated, the classifications must be confirmed before further development and future full evaluation can occur. Three visual examples of ring fracture detection can be seen in Fig. 6.8: the first involving small ring fracture, the second a larger potential abnormality, and the third no fracture at all. Future testing of this component should explore sensitivity and specificity measures; the former is more crucial than the latter in order to avoid missed injuries.

Table 6.2: Preliminary results of abnormality detection via visual inspection.

| Type                         | Correctly Detected | Not Correctly Detected |
|------------------------------|--------------------|------------------------|
| Suspected Minor Fracture     | 6                  | 2                      |
| Suspected Larger Abnormality | 7                  | 1                      |
| No Fracture                  | 6                  | 2                      |

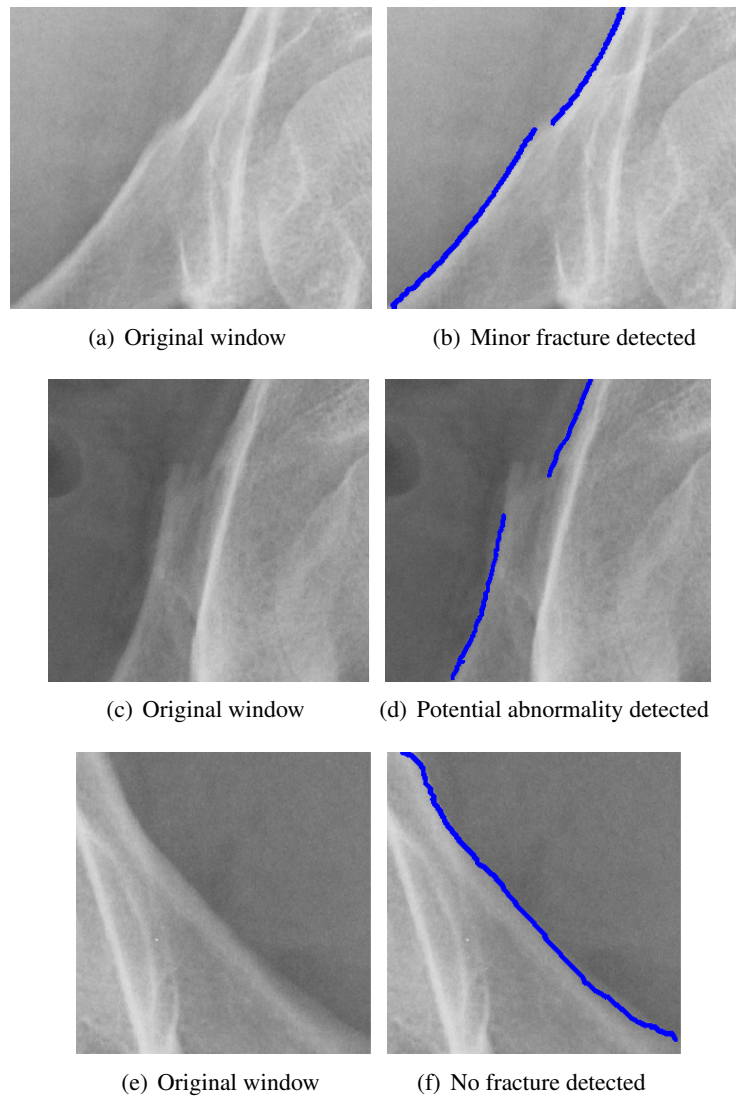


Figure 6.8: Multiple results of fracture detection.

#### **6.4 Complexity and Processing Time**

Testing of Spline/ASM segmentation (as described in Section 6.2) was performed on an Intel quad-core machine with 4GB RAM. Average run times were very similar: 7.15 seconds for Spline/ASM versus 7.85 for standard ASM (rounded to 2 d.p.). The slightly lower average for Spline/ASM is likely due to two images with several prominent false edges around the ring. In these cases, Spline/ASM maintained ring curvature and quickly converged to the correct contour, whereas standard ASM attempted to move toward the false edges.

#### **6.5 Testing on Alternative Images**

This section presents the results of applying the Spline/ASM algorithm to cardiac MR images in order to segment the left ventricle, which shows strong curvature. Performance is compared with that of standard ASM. This experiment gives a general indication of the algorithm's potential use in segmentation of other medical images. The dataset consisted of 14 cardiac MR images, focused on the left ventricle, and was published by the Informatics and Mathematical Modeling department at the Technical University of Denmark in 2002. Each MR image was first resized to 512x512 pixels, after which five of the images were each labeled with 35 landmark points and used for training, with nine reserved for testing. Though this is insufficient for full statistical evaluation, it gives a general impression of the performance of the Spline/ASM algorithm on a different type of image. Both the basic ASM and combined spline/ASM algorithm were tested, and the results were compared by sight.

The two algorithms offered comparable and accurate results in five of the test cases, which indicates that the combined algorithm is as robust as the standard algorithm. In two of the cases, results for both algorithms were inaccurate. In the two remaining cases, the performance of the

combined algorithm exceeded that of the original. The results for the first case are presented in Figure 6.9 (segmentation using combined Spline/ASM algorithm) and Figure 6.10 (segmentation using basic algorithm). It can be seen by sight that, although both algorithms do capture the boundary, the combined/spline ASM algorithm offers more accurate landmark placement. The results for the second case are presented in Figure 6.11 (segmentation using combined spline/ASM algorithm) and Figure 6.12 (segmentation using basic algorithm). In this case the difference between the two is more obvious. Though there are slight inaccuracies in the boundary determined by the combined algorithm (most noticeably to the left side of the ventricle), the results obtained using the basic ASM algorithm are noticeably worse.




Figure 6.9: An example of left cardiac ventricle detection in a MR image using the Spline/ASM algorithm.




Figure 6.10: An example of left cardiac ventricle detection in a MR image using the standard ASM algorithm.




Figure 6.11: A second example of left cardiac ventricle detection in a MR image using the Spline/ASM algorithm.




Figure 6.12: A second example of left cardiac ventricle detection in a MR image using the standard ASM algorithm.

## 6.6 Summary and Discussion

The results presented in this chapter show promising performance given the data available, indicating that the basic framework of the system is sound. At this stage further evaluation requires not only larger datasets, but also extended involvement of experts who can provide information on dataset images and verify the accuracy of the developed algorithms. Currently, assessment is based upon a number of images analyzed by a participating radiologist, along with information provided by him on particular characteristics of fracture and reference images from medical literature. This has been very useful in verifying the concepts behind the developed algorithms and demonstrating the tasks that can be performed. The next step in development is to pursue cooperation from experts on a larger scale. This will enable regular comprehensive statistical assessment of performance as the algorithms are refined and additional processing components are developed.

Current evaluation of the overall automatic segmentation algorithm demonstrates performance at least as accurate as manual model placement, and a higher level of consistency between images. This suggests that the Spline/ASM algorithm is performing well in detection of individual structures, and that the hierarchical initialization algorithm is successfully locating the correct position for the models at each step. It copes well with labels, text and other artifacts present on a number of the X-ray images, which presented difficulties at the time of my dissertation proposal. The proposal also considered the possibility that increasing the number of landmarks might improve performance; in practice, this offered little benefit versus the increase in processing time. It is important to remember that segmentation is not the final goal in this application, but a first step required before individual structures can be separately analyzed; very high precision is therefore not as crucial as speed, given the intended use in treatment of trauma. Errors may still occur in



some situations, such as when the femurs are severely fractured. However, it is very likely that these injuries would be immediately obvious, rather than something observed by the X-ray analysis component. Time complexity for the segmentation algorithm is acceptable; when tested on an Intel quad-core machine with 4GB RAM, running time for a single image was approximately 40 seconds. This has been reduced from the 1 minute run-time stated in the dissertation proposal by imposing additional constraints on Hough circle detection for the left and right ilium (the most time-consuming step). However, Hough circle detection is still  $O(n^3)$  in complexity, and is expected to remain the most time-consuming step of the algorithm. Hough line detection is an  $O(n^2)$  operation, while Spline/ASM is of linear complexity (as confirmed by the approximately linear increase in execution time that results from an increase in the number of landmark points).

As stated, full evaluation of fracture detection performance is not feasible at this stage. However, preliminary results are promising, and - if the criteria used are correct - show good accuracy. Potential abnormalities are consistently detected, though these need to be verified as fracture or non-fracture before they can be firmly classified. The fracture detection algorithm is flexible and uses a simple set of parameters, suggesting that refining its behavior according to expert opinion may not be too difficult. Computational complexity is also acceptable; while each transform step of a matrix-based DWT algorithm is an  $O(n^2)$  operation, it is typically implemented using a linear algebra approach which reduces this to  $O(n)$ .

The next chapter of this dissertation examines the predictive model aspect of this work - or more specifically, how the extracted image features can be used in generating outcome predictions for pelvic trauma patients.

## CHAPTER 7 Image Features in Predictive Models

### 7.1 Overview

This chapter describes two previous experiments incorporating medical image features into the generation of predictive models for pelvic injury severity. Experiment 1 was performed and published jointly with Wenan Chen [14] and focuses on the combination of X-ray features with demographic information and physiological scores. Experiment 2 was performed jointly with Simina Vasilache and has been submitted for publication; it extends the approach of Experiment 1 by also including image features present in CT scans. As stated in Chapter 1, the work in this dissertation is intended to form part of a more general system able to generate grammatical rules for outcome prediction in cases of traumatic injury. The experiments therefore use the CART and C4.5 decision tree algorithms, which are briefly described in the following section.

### 7.2 Decision Tree Algorithms

Decision tree algorithms generate predictive models of some target variable based on a specific set of input variables. These models are formed as trees, where branches express conjunctions of input variables leading to the final classifications (values of the target variable) at the leaves. Each of the nodes between the root and leaves corresponds to an input variable which is split on at that node, and the split creates branches to child nodes or leaves according to specific values (or ranges of values) of the node variable. Tracing the branches from the root down to a specific leaf thus generates a predictive IF-THEN rule for that particular leaf, where the THEN-result is some value of the target variable and the IF-condition is combination of input variables and their split values

at each of the traversed nodes. The structure of a generated decision tree (and consequently the extracted rules) depends on the criteria used in choosing the best variable to split on at each node, which vary between algorithms. CART (Classification And Regression Tree) and C4.5 are two widely used decision-tree methods and will be explained in more detail.

### 7.2.1 CART

The CART algorithm constructs binary decision trees (where each node is split into two child nodes) using concepts from information theory [8]. Gini impurity is used as the splitting criterion when growing the tree, with the ‘best’ variable to split on being chosen as the one that most reduces impurity (heterogeneity of the dataset). The Gini impurity of a set of data measures how often a randomly chosen member of the set would be incorrectly classified if its classification was randomly chosen according to the distribution of classes in the subset. Assume there are  $N$  distinct classes for the target variable  $Y$ , and let  $f_i$  represent the proportion of items in the dataset of class  $k$ . The Gini impurity index is given by

$$I_G = 1 - \sum_{i=1}^m f_i^2 \quad (7.1)$$

Even though CART is an older algorithm, its handling of missing data and categorical variables make it a popular choice in machine learning and data mining, and it is still used in a variety of clinical research applications [25, 55].

### 7.2.2 C4.5

Similarly to CART, C4.5 generates binary decision trees using information theory concepts [53]. In this case, the splitting criterion is information gain, or the change in information entropy from a prior state to a new state incorporating additional information - such as that obtained from a new

data variable. When constructing a C4.5 tree, variables with high information gain measures are preferred as they are likely to be most relevant in rapidly reaching a value for the target variable. The information gain of an input variable  $X$  with respect to a target variable  $Y$  represents the reduction in uncertainty about the value of  $Y$  caused by knowing the value of  $X$ . This uncertainty is measured using entropy. Information gain  $IG(Y, X)$  is given by

$$IG(Y, X) = H(Y) - H(Y|X) \quad (7.2)$$

where  $H(Y)$  is the entropy of  $Y$  and represents the uncertainty about  $Y$ 's value, while  $H(Y|X)$  is the conditional entropy of  $Y$  given  $X$  and represents the uncertainty about  $Y$ 's value when  $X$  is known. Information gain can also be expressed in terms of joint entropy  $H(X, Y)$ . C4.5 is an extension of the ID3 algorithm, designed to be more computationally efficient and avoid issues with over-fitting of data [52].

### **7.3 Experiment 1: X-Ray Features**

#### *7.3.1 Experiment Overview*

Experiment 1 combines X-ray image features with demographic information and physiological scores to generate rules for injury severity prediction, based on a patient's predicted length of stay in the Intensive Care Unit (ICU). In accordance with a previous study, 'severe' corresponds to  $\geq 3$  days in ICU, and 'non-severe' to  $\leq 2$  days. The rules are generated using the CART algorithm described in Section 7.2.

#### *7.3.2 Dataset*

The dataset used in this experiment was provided by Carolinas Healthcare System and consists of 33 male pelvic trauma patients. All patients were admitted to ICU and survived treatment. Admis-

sion x-ray images were available for all cases. The dataset features included: Patient Age, Patient Sex, Injury Severity Score (ISS), Mechanism of Injury (e.g. Fall, MVC), Transport Mode (Ambulance or Helicopter), PH Airway, PH Fluids. All of these were stated to be potentially important in predicting patient outcome during discussion with experts. These features were combined with the following features extracted from pelvic radiographs: Horizontal Gap, Vertical Displacement, Symmetric Displacement, Pelvic Ring Fracture and Pubis-Ischium Fracture.

### 7.3.3 Results

Due to the small dataset size, CART with leave-one-out selection was used, generating 32 trees in each run. Since a stable pattern is required to represent the dataset, the most frequently generated tree was selected - referred to from now on as the dominant tree. Across all 33 runs, 11 distinct trees were generated; the dominant tree is shown in Figure 7.1 and appeared 20 times. This tree was then applied to the whole dataset, and the resulting sensitivity, specificity and accuracy measures were recorded. Note that the dominant tree has only two levels; this is due to the small dataset used, as including many features when few samples are available is likely to result in over-specific rules.

C4.5 was then used to verify the extracted rules by generating a tree using the ISS and Horizontal Gap features that appear in the dominant CART tree. The resulting C4.5 tree has the same structure and the split points are close to those specified by CART. C4.5 and CART generated very similar rules from the pelvic injury dataset, indicating that the patterns extracted from the data are genuine and not dependent on the machine learning algorithm used. To check whether CART offers comparable accuracy to other machine learning techniques in this application, the Support Vector Machine (SVM) method was also used to classify the dataset using the selected ISS and Horizontal displacement features with a linear kernel. Again, leave-one-out was used to compen-

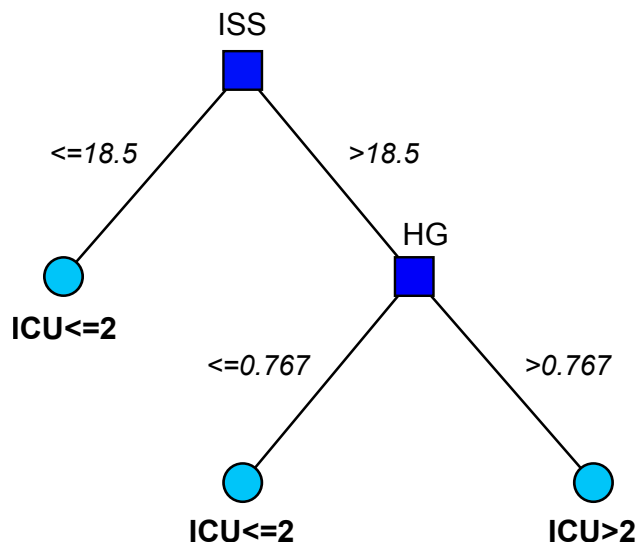


Figure 7.1: The dominant tree extracted by CART in injury severity prediction for Experiment 1.

sate for the small dataset size. Results for the rules generated by the selected CART decision tree, rules generated by C4.5 tree, and the SVM method trained on the dataset are presented in table 7.1. Note that both C4.5 and SVM verification were performed after CART-based feature selection via extraction of the dominant tree.

Table 7.1: Classification results for Experiment 1 using leave-one-out. SVM and C4.5 testing are performed for validation purposes only, using the features selected from the dominant CART tree.

| Algorithms          | CART  | C4.5  | SVM   |
|---------------------|-------|-------|-------|
| Average Sensitivity | 100%  | 100%  | 84.2% |
| Average Specificity | 69.2% | 61.5% | 69.2% |
| Average Accuracy    | 87.5% | 84.4% | 78.1% |

The results indicate that the dominant decision tree extracted by the CART algorithm produces rules with good sensitivity, specificity and accuracy in predicting a patient's expected length of ICU stay. Specificity is lower, but this is likely due to the small dataset available at the time of

the experiment. The results for SVM are poorer than would be expected, even after repeating the experiment with Gaussian kernels (which gave poorer performance than linear kernel). However, it must be remembered that SVM classification was applied only to the features extracted from the dominant CART tree, as previously described. The aim was to use both SVM and C4.5 to verify the results provided via CART - i.e. to confirm that the extracted pattern is not algorithm-specific - rather than testing their individual performance in the classification task. Applying SVM with various Gaussian kernels directly to the dataset typically provided better average performance than CART. However, SVM is not considered a suitable method for outcome prediction in this task due to its lack of transparency. It is also likely that CART results will improve as more data becomes available, allowing a larger and potentially more accurate rule-base to be generated. This will also enable leave-one-out to be replaced with 10-fold cross-validation during assessment.

Most significantly in terms of the study in this dissertation, the dominant CART tree selected one of the features extracted from the X-ray image as the second most important variable in predicting patient's ICU length of stay. As stated in 5, a change in the horizontal distance between the left and right pubis bones is a significant component in patterns of pelvic injury, as it can indicate pelvic instability when combined with other features. This provides a compelling argument for the incorporation of image features into the decision-making process for diagnosis and treatment.

#### **7.4 Experiment 2: X-Ray and CT Features**

##### *7.4.1 Experiment Overview*

Experiment 2 combines both X-ray and CT image features with demographic information and physiological scores, again to generate rules for injury severity prediction. In this case, severity is based on whether a patient is sent to ICU. The rules are generated using the C4.5 algorithm described in Section 7.2.

#### 7.4.2 Dataset

The dataset was again provided by Carolinas Healthcare System. More cases were available, comprising a total of 45 pelvic trauma patients of both genders who all survived treatment. Of these patients, 20 were sent to ICU after initial diagnosis in the E.R., and the remaining 25 were either sent home or to a hospital ward. Both X-ray images and CT scans taken on the day of admission are available for each patient. The demographic and physiological features included are: Patient Age, Patient Sex, Injury Severity Score (ISS), Mechanism of Injury (e.g. Fall, MVC) and Glasgow Coma Score (GCS). These were combined with the following extracted image features: presence of pubis fracture, presence of fracture in sacrum, pubic symphysis width, vertical displacement of pubis, pelvic ring symmetry displacement, and potential presence of hemorrhage.

#### 7.4.3 Results

Since the combined dataset contained a relatively wide range of features for the number of samples available, feature selection was performed to identify the most useful attributes and remove those that are irrelevant to the learning task. Feature selection typically results in a more robust predictive model and helps avoid over-fitting. This experiment used a correlation-based method which selects a subset of features based on both their individual predictive ability and the degree of redundancy between them [26]. Results indicated that multiple features extracted from the X-ray and CT images may prove useful in prediction of injury severity. Ranking of features via correlation-based subset selection selected ISS, GCS and presence of pubis fracture as the three features statistically significant for the prediction task. It should be noted that fracture of the sacrum and potential presence of hemorrhage also ranked in the top six. These results correspond to knowledge of pelvic injury patterns; as stated in Section 7.3, fracture of the pubis can not only cause damage to



the surrounding organs and tissues but may also be associated with pelvic instability. The same is true of damage to the sacrum - and, as stated in Chapter 1, hemorrhage has a significant impact on patient outcome. Their absence from the filtered feature set may be due to their infrequency in the dataset. As more data become available, it is likely that there will be more examples of these conditions, and they are expected to become more statistically significant in predicting patient outcome.

The three selected features - ISS, GCS, and presence of pubis fracture - were used as input to the C4.5 algorithm, generating the tree shown in Figure 7.2. To verify the accuracy of the extracted rules, the same input features were used as input to the CART algorithm, generating a second tree. This is identical in structure to the C4.5 tree, though ISS and GCS are split on 26.5 and 14.5 respectively. The close correspondence in structure indicates that the patterns extracted from the data are genuine. The results of both methods were assessed using ten-fold cross-validation, and the accuracy, specificity and sensitivity values are presented in Table 7.2. It can be seen that C4.5 offers the best performance, with good sensitivity and overall accuracy. Specificity is reduced, but this is likely due to the relatively small dataset and the imbalance between the two classes.

Table 7.2: Classification results for Experiment 2 using 10FCV

|                     | C4.5  | CART  |
|---------------------|-------|-------|
| Average Sensitivity | 84.0% | 84.0% |
| Average Specificity | 70.0% | 60.0% |
| Average Accuracy    | 77.8% | 73.3% |

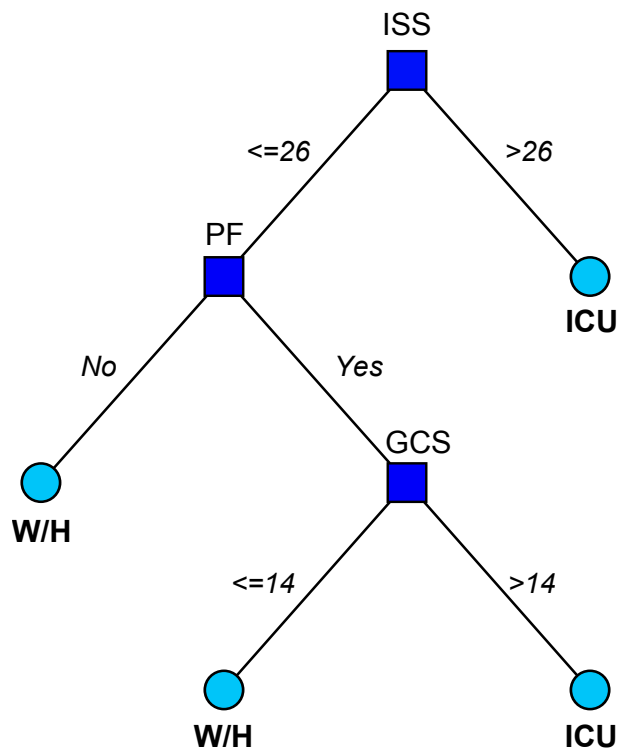


Figure 7.2: The decision tree output by the C4.5 algorithm in Experiment 2. 'PF' indicates presence of pubis fracture, and 'ICU' and 'W/H' (Ward/Home) are the two outcomes.

### 7.5 Summary

The two experiments presented in this chapter offer promising results, and indicate the potential diagnostic value of X-ray image features in a decision-making system for pelvic trauma. In particular, Experiment 1 found the width of the symphysis pubis to be significant, which is supported by anatomical knowledge of pelvic instability. Meanwhile, the pubis fracture feature selected in Experiment 2 may indicate instability; these fractures can also result in damage to surrounding soft tissue. Feature extraction for both experiments was impacted by two factors: an insufficient number of examples of certain injuries, and the inclusion of both genders in the dataset. Since Spline/ASM models should be trained separately on each gender to account for differences in

anatomy, this further reduces the number of training examples. It should also be noted that the dataset used in the first experiment is smaller than the second, as more cases became available in the period between them. As even more data is obtained in the future it will become important to repeat these experiments on a regular basis in order to obtain more reliable evaluation of results and refine the methods used in feature extraction. Furthermore, a larger number of available cases will enable more features to be considered in the rule generation process, which is likely to result in a larger and more accurate rule-base.

## CHAPTER 8 Conclusions and Future Work

### 8.1 Conclusions

This dissertation provides a framework for an X-ray analysis component of a computer-aided decision-support system, comprising automatic segmentation and extraction of potentially useful diagnostic features. First, a hierarchical initialization algorithm is applied that automatically detects each structure of interest in the radiograph in a single pass. Approximate structure locations are identified using directed Hough Transform, with segmentation being performed using Spline/ASM, an extension of the standard Active Shape Model algorithm which maintains curved contours through the use of cubic B-spline interpolation. Based on the detected position of the pubis bones, quantitative measures of symphysis pubis displacement are automatically calculated. Detection of potential fractures is then performed by creating a series of overlapping windows using Spline/ASM model points, and analyzing each one for contour discontinuities. The proposed method offers multiple benefits over existing approaches to pelvic radiograph segmentation and bone fracture detection, including the ability to detect multiple separate structures rather than a single object representing the pelvis. It also performs full segmentation of all structure in a single pass, with no need for any interaction or additional input by the user other than the patient's gender, and addresses the problem of false edges by introducing additional deformation constraints. Furthermore, the hierarchical approach to segmentation means that feature extraction algorithms can be independently designed and improved for each individual structure. Results so far are promising, with the automated segmentation approach offering consistent performance on the data

available. Detection of potential fractures is also in a good preliminary state. This indicates that seeking expert verification in full statistical evaluation is now appropriate, in order to verify characteristics of true fractures and further develop the system to be robust for practical application. Two experiments have been performed exploring the prediction of injury severity, using datasets which combine extracted image features with demographic information and physiological scores. In both experiments, features extracted from X-ray images are shown to be useful in accurately predicting outcomes.

## **8.2 Future Work**

In order to make it suitable for use in polytrauma, many extensions are being considered for the trauma decision support system as a whole. These include incorporation of existing modules for evaluating traumatic brain injuries, and development of techniques for analyzing injuries to other anatomical regions such as the chest. However, I will focus here on possible extensions to the X-ray image processing component.

The first goal is to obtain expert assistance from multiple radiologists in labeling the full dataset. This is likely to prove difficult logistically due to the expected volume of data and the time demands involved. However, more rigorous testing demands a gold standard, which in this application is manual evaluation from multiple experts (to compensate for individual variation in opinions). Since the potential of the system has been demonstrated through experimental testing, this may encourage future collaborations. Related to this is the need to obtain more examples concerning sacral ligament damage and established injury patterns. This will allow the injury pattern detection component to be developed in more detail and subjected to full evaluation. It may also allow the incorporation of clinical signs into injury pattern detection. Although this removes

full system automation since these signs are based on physical examination of the patient, when combined with the quantitative measurements provided by the imaging component and a large set of prior cases, they may prove valuable in verifying suspected patterns based on prior cases.

Concerning the technical aspects, an important future task is to improve the robustness of the automatic initialization process. Performance is currently very good even in the presence of artifacts such as labels, text and measurement marks. However, several of the images in the dataset collected using from portable devices showed extensive patient rotation or non-alignment of the femurs. Although the capacity to deal with basic rotation has been included in the current system, further work is needed to handle the cases where the femurs are non-parallel (due to positioning of the patient's limbs). Likewise, though the system can determine loss of structure symmetry, quantitative measurements of these changes would be more useful in generating predictive models. Finally, the proposed system can be adapted to other X-ray viewing angles, though the potential clinical utility of doing so must be established first.

## Bibliography

## Bibliography

- [1] E. S. Amis, P. F. Butler, K. E. Applegate, S. B. Birnbaum, L. F. Brateman, J. M. Hevezi, F. A. Mettler, R. L. Morin, M. J. Pentecost, G. G. Smith, K. J. Strauss, and R. K. Zeman. American College of Radiology white paper on radiation dose in medicine. *J Am Coll Radiol*, 4:272–284, May 2007.
- [2] P. J. Andrews, D. H. Sleeman, P. F. Statham, A. McQuatt, V. Corruble, P. A. Jones, T. P. Howells, and C. S. Macmillan. Predicting recovery in patients suffering from traumatic brain injury by using admission variables and physiological data: a comparison between decision tree analysis and logistic regression. *J. Neurosurg.*, 97:326–336, Aug 2002.
- [3] N. S. Arikidis, S. Skiadopoulos, A. Karahaliou, E. Likaki, G. Panayiotakis, and L. Costari-dou. B-spline active rays segmentation of microcalcifications in mammography. *Medical Phys.*, 35(11):5161–5171, 2008.
- [4] J.L. Ball and T. Price. *Chesneys' Radiographic Imaging*. Wiley-Blackwell, 6th edition, 1995.
- [5] A. M. Basta, C. C. Blackmore, and H. Wessells. Predicting urethral injury from pelvic fracture patterns in male patients with blunt trauma. *J. Urol.*, 177:571–575, Feb 2007.
- [6] Y. Ben-Menachem, D.M. Coldwell, J.W. Young, and A.R. Burgess. Hemorrhage associated with pelvic fractures: causes, diagnosis, and emergent management. *American J. Roentgenology*, 157(5):1005–1014, November 1991.
- [7] N. Boukala, E. Favier, B. Laget, and P. Radeva. Active shape model based segmentation of bone structures in hip radiographs. *Industrial Technology, 2004. IEEE ICIT '04. 2004 IEEE International Conference on*, 3:1682–1687 Vol. 3, Dec. 2004.
- [8] L. Breiman. *Classification and Regression Trees*. Chapman & Hall, Boca Raton, 1993.
- [9] P. Brigger, J. Hoeg, and M. Unser. B-spline snakes: a flexible tool for parametric contour detection. *Image Processing, IEEE Transactions on*, 9(9):1484 –1496, sep 2000.
- [10] B.D. Browner, J.B. Jupiter, A.M. Levine, and P.G. Trafton. *Skeletal Trauma: Basic Science, Management, and Reconstruction*. Saunders Book Company, 3rd edition, 2003.
- [11] R. W. Bucholz. The pathological anatomy of Malgaigne fracture-dislocations of the pelvis. *J Bone Joint Surg Am*, 63:400–404, Mar 1981.
- [12] A. R. Burgess, B. J. Eastridge, J. W. Young, T. S. Ellison, P. S. Ellison, A. Poka, G. H. Bathon, and R. J. Brumback. Pelvic ring disruptions: effective classification system and treatment protocols. *J Trauma*, 30:848–856, Jul 1990.



- [13] H. C. Chen, C. H. Wu, C. J. Lin, Y. H. Liu, and Y. N. Sun. Automated segmentation for patella from lateral knee X-ray images. *Conf Proc IEEE Eng Med Biol Soc*, 2009:3553–3556, 2009.
- [14] W. Chen, R. Smith, S. Vasilache, K. Najarian, K. Ward, C. Cockrell, and J. Ha. Traumatic pelvic injury outcome prediction by extracting features from relevant medical records and x-ray images. *Bioinformatics and Biomedicine, IEEE International Conference on*, pages 291–294, 2009.
- [15] Y. Chen, X. Ee, W. K. Leow, and T. S. Howe. Automatic extraction of femur contours from hip x-ray images. *Computer Vision for Biomedical Image Applications*, 3765:200–209, 2005.
- [16] T.F. Cootes and C.J. Taylor. Statistical models of appearance for medical image analysis and computer vision. In *Proc. SPIE Medical Imaging*, volume 42, 2001.
- [17] D. Demetriades, M. Karaiskakis, K. Toutouzas, K. Alo, G. Velmahos, and L. Chan. Pelvic fractures: epidemiology and predictors of associated abdominal injuries and outcomes. *J. Am. Coll. Surg.*, 195:1–10, Jul 2002.
- [18] P. Dierckx. *Curve and Surface Fitting with Splines*. Oxford University Press, 1993.
- [19] F. Ding, W.K. Leow, and T.S. Howe. Automatic segmentation of femur bones in anterior-posterior pelvis x-ray images. In *Computer Analysis of Images and Patterns, 12th Int. Conf. (CAIP 2007)*, pages 205–212, 2007.
- [20] X. Dong and G. Zheng. Automatic extraction of femur contours from calibrated x-ray images: A bayesian inference approach. In *Biomedical Imaging: From Nano to Macro, 2008. ISBI 2008. 5th IEEE International Symposium on*, pages 57–60, may 2008.
- [21] J. Duryea and J.M. Boone. A fully automated algorithm for the segmentation of lung fields on digital chest radiographic images. *Medical Phys.*, 22(2):183–191, February 1995.
- [22] N. Duta and M. Sonka. Segmentation and interpretation of MR brain images: an improved active shape model. *IEEE Trans Med Imaging*, 17:1049–1062, Dec 1998.
- [23] B. J. Eastridge, A. M. Starr, J. P. O’Keefe, and E. Grant. The importance of fracture pattern in guiding therapeutic decision-making in patients with hemorrhagic shock and pelvic ring disruptions. *J. Trauma*, 53(3):446–450, September 2002.
- [24] R. J. Forsyth, R. C. Parslow, R. C. Tasker, C. A. Hawley, K. P. Morris, K. Morris, R. Appleton, M. Crouchman, R. Forsyth, C. Hawley, M. Marsh, P. May, P. McKinney, J. Middleton, R. Parslow, J. Punt, T. Ralph, R. Tasker, R. Sarginson, A. Ferguson, E. Keane, K. Morris, J. Fraser, J. Alexander, S. Ferguson, M. Kenny, D. Lutman, A. Durward, H. Klonin, A. Shefler, C. Killick, D. Prior, L. Edwards, Y. Egberongbe, T. Chater, M. Darowski, P. Barry, P. Khandelwal, A. Robinson, R. Forsyth, B. Taylor, A. Maunganidze, J. Purday, M. Lo, D. Simpson, P. Cullen, P. Withington, D. Stewart, M. Samuels, P. Tomlin, J. Bellin, T. Ralph,

- C. Boyles, D. Snaddon, A. Wagstaff, S. Skellett, M. Gajraj, M. Christie, and E. Wright. Prediction of raised intracranial pressure complicating severe traumatic brain injury in children: implications for trial design. *Pediatr Crit Care Med*, 9:8–14, Jan 2008.
- [25] M. Garzotto, T. M. Beer, R. G. Hudson, L. Peters, Y. C. Hsieh, E. Barrera, T. Klein, and M. Mori. Improved detection of prostate cancer using classification and regression tree analysis. *J. Clin. Oncol.*, 23:4322–4329, Jul 2005.
- [26] M. A. Hall. *Correlation-based Feature Subset Selection for Machine Learning*. PhD thesis, University of Waikato, Hamilton, New Zealand, 1998.
- [27] M. J. Heetveld, I. Harris, G. Schlaphoff, and M. Sugrue. Guidelines for the management of haemodynamically unstable pelvic fracture patients. *ANZ J Surg*, 74:520–529, Jul 2004.
- [28] S. Houshian, M. S. Larsen, and C. Holm. Missed injuries in a level I trauma center. *J Trauma*, 52:715–719, Apr 2002.
- [29] M. H. Hsu, Y. C. Li, W. T. Chiu, and J. C. Yen. Outcome prediction after moderate and severe head injury using an artificial neural network. *Stud Health Technol Inform*, 116:241–245, 2005.
- [30] B. Isler and R. Ganz. Classification of pelvic ring injuries. *Injury*, 27(Supplement 1):3–12, 1996.
- [31] S.Y. Ji, T. Huynh, and K. Najarian. A comparative medical informatics approach to traumatic pelvic injuries. In *Bioinformatics and Biomedicine, 2007. BIBM 2007. IEEE International Conference on*, pages 188–193, nov. 2007.
- [32] S. Kane. *Introduction to Physics in Modern Medicine*. CRC Press, 2002.
- [33] P. P. Kaneriya, M. E. Schweitzer, C. Spettell, M. J. Cohen, and D. Karasick. The cost-effectiveness of routine pelvic radiography in the evaluation of blunt trauma patients. *Skeletal Radiol.*, 28:271–273, May 1999.
- [34] P.M. Kuhnert, K. Do, and R. McClure. Combining non-parametric models with logistic regression: an application to motor vehicle injury data. *Computational Statistics and Data Analysis*, 34(3):371–386, 2000.
- [35] T.M. Lehmann, C. Gonner, and K. Spitzer. Addendum: B-spline interpolation in medical image processing. *Medical Imaging, IEEE Transactions on*, 20(7):660–665, July 2001.
- [36] Y. C. Li, L. Liu, W. T. Chiu, and W. S. Jian. Neural network modeling for surgical decisions on traumatic brain injury patients. *Int J Med Inform*, 57:1–9, Jan 2000.

- [37] V.L.F Lum, W.K. Leow, Y. Chen, T.S. Howe, and M.A. Png. Combining classifiers for bone fracture detection in x-ray images. In *IEEE Int Conf. on Image Processing (ICIP 2005)*, pages 1149–1152, 2005.
- [38] K. Lunsjo, A. M. Tadros, A. Hauggaard, R. Blomgren, J. Kopke, and F. M. Abu-Zidan. Acute plain anteroposterior radiograph of the pelvis is not useful in detecting fractures of iliac wing and os sacrum: a prospective study of 73 patients using CT as gold standard. *Australas Radiol*, 51:147–149, Apr 2007.
- [39] G.K. Manos, A.Y. Cairns, I.W. Rickets, and D. Sinclair. Segmenting radiographs of the hand and wrist. *Computer Methods and Programs in Biomedicine*, 43(3-4):227–237, June 1993.
- [40] T. McInerney and D. Terzopoulos. Deformable models in medical image analysis: a survey. *Med Image Anal*, 1:91–108, Jun 1996.
- [41] R. Medina, A. Bravo, P. Windyga, J. Toro, P. Yan, and G. Onik. A 2-D Active Appearance Model For Prostate Segmentation in Ultrasound Images. *Conf Proc IEEE Eng Med Biol Soc*, 4:3363–3366, 2005.
- [42] F. A. Mettler, W. Huda, T. T. Yoshizumi, and M. Mahesh. Effective doses in radiology and diagnostic nuclear medicine: a catalog. *Radiology*, 248:254–263, Jul 2008.
- [43] Roberts M.G., T.F. Cootes, and J.E. Adams. Vertebral morphometry: semiautomatic determination of detailed shape from dual-energy x-ray absorptiometry images using active appearance models. *Invest Radiol.*, 41(12):849–859, December 2006.
- [44] P. Mildenerger, M. Eichelberg, and E. Martin. Introduction to the DICOM standard. *Eur Radiol*, 12:920–927, Apr 2002.
- [45] P. R. Miller, P. S. Moore, E. Mansell, J. W. Meredith, and M. C. Chang. External fixation or arteriogram in bleeding pelvic fracture: initial therapy guided by markers of arterial hemorrhage. *J Trauma*, 54:437–443, Mar 2003.
- [46] T. Niwa, S. Takebayashi, H. Igari, N. Morimura, K. Uchida, M. Sugiyama, and S. Matsubara. The value of plain radiographs in the prediction of outcome in pelvic fractures treated with embolisation therapy. *Br J Radiol*, 73:945–950, Sep 2000.
- [47] A. K. Obaid, A. Barleben, D. Porral, S. Lush, and M. Cinat. Utility of plain film pelvic radiographs in blunt trauma patients in the emergency department. *Am Surg*, 72:951–954, Oct 2006.
- [48] American College of Surgeons Committee on Trauma. *Advanced Trauma Life Support: ATLS Instructor Manual*. American College of Surgeons, 2004.
- [49] N. Otsu. A threshold selection method from gray level histograms. *IEEE Trans. Systems, Man and Cybernetics*, 9:62–66, March 1979.

- [50] B. C. Pang, V. Kuralmani, R. Joshi, Y. Hongli, K. K. Lee, B. T. Ang, J. Li, T. Y. Leong, and I. Ng. Hybrid outcome prediction model for severe traumatic brain injury. *J. Neurotrauma*, 24:136–146, Jan 2007.
- [51] E. Pietka. Lung segmentation in digital radiographs. *Journal of Digital Imaging*, 7(2):79–84, May 1994.
- [52] J. R. Quinlan. Induction of decision trees. *Mach. Learn.*, 1(1):81–106, 1986.
- [53] J.R. Quinlan. *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Francisco, 1993.
- [54] C. S. Resnik, D. J. Stackhouse, K. Shanmuganathan, and J. W. Young. Diagnosis of pelvic fractures in patients with acute pelvic trauma: efficacy of plain radiographs. *AJR Am J Roentgenol*, 158:109–112, Jan 1992.
- [55] A. Rovlias and S. Kotsou. Classification and regression tree for prediction of outcome after severe head injury using simple clinical and laboratory variables. *J. Neurotrauma*, 21:886–893, Jul 2004.
- [56] A. I. Rughani, T. M. Dumont, Z. Lu, J. Bongard, M. A. Horgan, P. L. Penar, and B. I. Tranmer. Use of an artificial neural network to predict head injury outcome. *J Neurosurg*, Dec 2009.
- [57] M. E. Segal, P. H. Goodman, R. Goldstein, W. Hauck, J. Whyte, J. W. Graham, M. Polansky, and F. M. Hammond. The accuracy of artificial neural networks in predicting long-term outcome after traumatic brain injury. *J Head Trauma Rehabil*, 21:298–314, 2006.
- [58] N. Sharma and L. M. Aggarwal. Automated medical image segmentation techniques. *J Med Phys*, 35:3–14, Jan 2010.
- [59] D. F. Signorini, P. J. Andrews, P. A. Jones, J. M. Wardlaw, and J. D. Miller. Predicting survival using simple clinical variables: a case study in traumatic brain injury. *J. Neurol. Neurosurg. Psychiatr.*, 66:20–25, Jan 1999.
- [60] K. Smith, Y. Ben-Menachem, J. H. Duke, and G. L. Hill. The superior gluteal: an artery at risk in blunt pelvic trauma. *J Trauma*, 16:273–279, Apr 1976.
- [61] R. Smith and K. Najarian. Automated segmentation of pelvic bone structure in x-ray radiographs using active shape models and directed hough transform. In *BIBMW '08: Proceedings of the IEEE International Conference on Bioinformatics and Biomedicine Workshops*, pages 56–63, 2008.

- [62] R. Smith-Bindman, J. Lipson, R. Marcus, K. P. Kim, M. Mahesh, R. Gould, A. Berrington de Gonzalez, and D. L. Miglioretti. Radiation dose associated with common computed tomography examinations and the associated lifetime attributable risk of cancer. *Arch. Intern. Med.*, 169:2078–2086, Dec 2009.
- [63] T. Stammberger, F. Eckstein, M. Michaelis, K. H. Englmeier, and M. Reiser. Interobserver reproducibility of quantitative cartilage measurements : Comparison of b-spline snakes and manual segmentation. *Mag. Res. Imaging.*, 17(7):1033–1042, 1999.
- [64] O. Tasaki, T. Shiozaki, T. Hamasaki, K. Kajino, H. Nakae, H. Tanaka, T. Shimazu, and H. Sugimoto. Prognostic indicators and outcome prediction model for severe traumatic brain injury. *J Trauma*, 66:304–308, Feb 2009.
- [65] N. H. Theumann, J. P. Verdon, E. Mouhsine, A. Denys, P. Schnyder, and F. Portier. Traumatic injuries: imaging of pelvic fractures. *Eur Radiol*, 12:1312–1330, Jun 2002.
- [66] M. Tile. *Fractures of the Pelvis and Acetabulum*. Lippincott Williams & Wilkins, 3rd edition, 2003.
- [67] M. Unser, A. Aldroubi, and M. Eden. B-spline signal processing. i. theory. *Signal Processing, IEEE Transactions on*, 41(2):821 –833, feb 1993.
- [68] B. van Ginneken, M. B. Stegmann, and M. Loog. Segmentation of anatomical structures in chest radiographs using supervised methods: a comparative study on a public database. *Med. Image Anal.*, 10:19–40, Feb 2006.
- [69] S. Vasilache, R. Smith, S.Y. Ji, K. Najarian, and T. Huynh. Outcome prediction in traumatic pelvic injuries using maximum similarity and quality measures. In *Information Reuse and Integration, 2008. IRI 2008. IEEE International Conference on*, pages 82 –85, July 2008.
- [70] N. Vittitoe. Identification of lung regions in chest radiographs using markov random field modeling. *Medical Physics*, 25(6):976–985, June 1998.
- [71] D.J. Withey and Z.J. Koles. Medical image segmentation: Methods and software. In *Noninvasive Functional Source Imaging of the Brain and Heart and the International Conference on Functional Biomedical Imaging, 2007. NFSI-ICFBI 2007. Joint Meeting of the 6th International Symposium on*, pages 140 –143, oct. 2007.
- [72] T. Xu, M. Mandal, R. Long, and A. Basu. Gradient vector flow based active shape model for lung field segmentation in chest radiographs. *Conf Proc IEEE Eng Med Biol Soc*, 2009:3561–3564, 2009.
- [73] D.W. Yap, Y. Chen, W.K. Leow, T.S. Howe, and M.A. Png. Detecting femur fractures by texture analysis of trabeculae. In *Int. Conf. on Pattern Recognition (ICPR 2004)*, pages 730–733, 2004.

- [74] J. W. Young, A. R. Burgess, R. J. Brumback, and A. Poka. Pelvic fractures: value of plain radiography in early assessment and management. *Radiology*, 160:445–451, Aug 1986.
- [75] J. W. Young and C. S. Resnik. Fracture of the pelvis: current concepts of classification. *AJR Am J Roentgenol*, 155:1169–1175, Dec 1990.
- [76] S. K. Zhou, J. Shao, B. Georgescu, and D. Comaniciu. Pairwise active appearance model and its application to echocardiography tracking. *Med Image Comput Comput Assist Interv*, 9:736–743, 2006.

## VITA

Rebecca Smith was born on July 21, 1982, in London, England, and holds United Kingdom citizenship. She received her Bachelor of Engineering in Computing from Imperial College London in 2003. Between January and December 2007 she was employed as a Graduate Teaching Assistant at Virginia Commonwealth University, and then as a Graduate Research Assistant between January 2008 and May 2010.

### **List of Relevant Publications:**

- R. Smith and K. Najarian, *A Hierarchical Method Based on Active Shape Models and Directed Hough Transform for Segmentation of Noisy Biomedical Images; Application in Segmentation of Pelvic X-ray Images*, *BMC Medical Informatics and Decision Making*, 9(Suppl 1):S2, Nov. 2009.
- W. Chen, R. Smith, S.Y. Ji, K. Ward, and K. Najarian, *Automated Ventricular Systems Segmentation in Brain CT Images by Combining Low-level Segmentation and High level Template Matching*, *BMC Medical Informatics and Decision Making*, 9(Suppl 1):S4, Nov. 2009.
- S. Y. Ji, R. Smith, T. Huynh and K. Najarian, *A Comparative Analysis of Multi-Level Computer-Assisted Decision Making Systems For Traumatic Brain Injuries*, *BMC Medical Informatics and Decision Making* 9:2, 2009.
- W. Chen, R. Smith, N. Nabizadeh, K. Ward, C. Cockrell, J. Ha, and K. Najarian, *CT Brain Image*

- Texture Feature Extraction for ICP Prediction*, to appear at the International Conference on Image and Signal Processing (ICISP 2010), July 2010.
- R. Smith, K. Ward, C. Cockrell, J. Ha and K. Najarian, *Detection of Fracture and Quantitative Assessment of Displacement Measures in Pelvic X-ray Images*, IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2010), March 2010.
- W. Chen, R. Smith, S. Vasilache, K. Najarian, K. Ward, C. Cockrell and J. Ha, *Traumatic Pelvic Injury Outcome Prediction by Extracting Features from Relevant Medical Records and X-ray Images*, IEEE Conference of Bioinformatics and Biomedicine (IEEE BHI 09), Nov 2009.
- R. Smith and K. Najarian, *Splines and Active Shape Model for Segmentation of Pelvic X-Ray Images*, IEEE International Conference on Complex Medical Engineering (ICME), Tempe, AZ, April 9-11, 2009.
- W. Chen, R. Smith and K. Najarian, *Automated Segmentation of Lateral Ventricles in Brain CT image*, IEEE Conference of Bioinformatics and Biomedicine (IEEE BHI 2008), Biomedical and Health Informatics Workshop, Nov 2008.
- R. Smith and K. Najarian, *Automated Segmentation of Pelvic Bone Structure in X-Ray Radiographs Using Active Shape Models and Directed Hough Transform*, IEEE Conference of Bioinformatics and Biomedicine (IEEE BHI 2008), Biomedical and Health Informatics Workshop, Nov 2008.
- S. Vasilache, R. Smith, S.Y. Ji and K. Najarian, *Outcome Prediction in Traumatic Pelvic In-*



*juries Using Maximum Similarity and Quality Measures*, IEEE International Conference on Information Reuse and Integration (IEEE IRI 2008), July 2008.

S.Y. Ji, K. Najarian, R. Smith, and T.Huynh, *Computer-Aided Traumatic Pelvic Injury Decision Making*, Southern Association for Information Systems (SAIS) March 2008.