

DIGITAL IMAGE ANALYSIS OF VERTEBRAL BODIES
FROM COMPUTED RADIOGRAPHY IMAGES

DIGITAL IMAGE ANALYSIS OF VERTEBRAL BODIES
FROM COMPUTED RADIOGRAPHY IMAGES

By

LIANNE TAN, B.ENG.

A Thesis

Submitted to the School of Graduate Studies
in Partial Fulfillment of the Requirements

for the Degree

Master of Engineering

McMaster University

Hamilton, Ontario

Canada

September 1995

© copyright 1995 by Lianne Tan

MASTER OF ENGINEERING (1995)
(Electrical and Computer Engineering)

McMASTER UNIVERSITY
Hamilton, Ontario, Canada.

TITLE : Digital image analysis of vertebral bodies
from computed radiography images

AUTHOR : Lianne Tan, B.Eng. (McMaster University)

SUPERVISORS : Dr. Zhi-quan Luo and Dr. Claude Nahmias

NUMBER OF PAGES : ix, 168

ABSTRACT

Vertebral body deformities resulting from osteoporosis are commonly observed in lateral spine radiographs. Clinically, these deformities are judged rather subjectively, both in their classification and in the assessment of the extent of deformity. In order to objectify clinical observation, morphometric measurements are a means for classifying and quantifying the presence and extent of vertebral body deformities. Drawing on the morphometric systems of Minne et al. (1988) and Smith-Bindman et al. (1991), we propose a total vertebral deformity index (TVDI) which provides a single, clinically meaningful number indicative of the extent of deformity in a vertebral column. Each of the three deformity types (wedge, biconcave, and compression) is classified and measured independently, using ratios relating measurements of anterior, mid, and posterior heights, and inferior width. Expected measurements are determined for each vertebral level, and the sum of the deviations from the measured to the expected values generates a vertebral deformity index. This proposed index addresses issues in vertebral body morphometry, including vertebral level specificity, body size differences, multiple compression effects, and multiple deformity types.

Inherent in the practical use of morphometric quantification are issues of accuracy and reproducibility of the measurements, and of the time involved in making these measurements. Digital image processing algorithms are developed to attempt automated detection and measurement of the vertebral body boundary, using Computed

Radiography images of the lateral spine. Three different methods yield varying results. In addition to the characteristics of lateral spine images (such as high intensity vertebral ridges) which are used to advantage by the algorithms, all three methods must deal with problematic anatomical characteristics such as the presence of the high intensity ribs and ilium. The first method is a series of image enhancement and thresholding steps applied to each vertebra, in order to delineate its vertebral boundary in a bi-level image. This method is fairly effective, but suffers from its dependence on thresholding. The second method uses the cross-correlation measure to detect the vertebrae, given a starting vertebra, and then uses edge-gradient tracking to trace the vertebral boundary. Results of this method are promising. There remains a weakness in the cross-correlation detector, which fails to accurately locate vertebrae that are further away from the starting vertebra. The third method is an active contouring technique called Snakes, in which the vertebral boundary is represented by a deformable spline which seeks to minimize an energy functional consisting of curvature, continuity, and image energy terms. Present results display minimal convergence onto the vertebral boundaries. Future work should refine the customization of the energy functional to produce better results. At present, all three methods require some kind of user interaction. Further development may prove fruitful in reducing user interaction to achieve a truly automated system.

“I hear but forget; I see and remember; I do and I understand.”

- based on a Native American saying

ACKNOWLEDGEMENTS

To those who made my work so interesting and exciting during these past two years :

Dr. Claude Nahmias, for his guidance and advice, and to Dr. Tom Luo, for his positive encouragement and ideas

Dr. Colin Webber, for his generous support

Scott Sloka, who devoted part of his life to this project

Nick Street, University of Wisconsin, Madison, for his support and interest

William Renwick and Eyal Reingold for their unwavering belief in me

Members of the MIND Lab, past and present : Ben Yu, for his supply of ideas, support and humour, and for sharing Mind with me even though it was on his desk; Joe Thomas, for his invaluable technical support and assistance; Bob Dony, Nek Manji, Chris Gordon, Nicholas Christoforou, Rob deKemp, and Chris Bowen

D. Bev Kenyon, programmer extraordinaire, for showing me how to program in Sunview, and for his inspiration and technical support in the home stretch

The Staff of the Department of Nuclear Medicine, for creating such a friendly and caring environment in which to work

The Staff of the Department of Radiology (MUMC), especially Michelle Cottreau and Marilyn Kereliuk, for allowing me to intrude upon their busy workdays

Cheryl Gies, to whom I am deeply grateful

WR, whose brilliant insights inspired motivation and vitality, and whose friendship was an essential ingredient of my well-being

Friends from the CRL, who helped me survive graduate schoolwork : Jun Lu, Liu Wenzhong, John Tou, Kaywan Afkhamie, Yvonne To, Lena Wu, Haiying Wang, Rupert Ho, Philippe Wu, Ban Quach, Nagula Sangary, A. Yasotharan, and others

Many unnamed friends who continuously supported me with their encouragement, thank you.

I thank Mind, Joe, Fuji, Heart, Suntan, Sunrise and Riemann, for their faithful cooperation during these past two years.

Finally, it is true that without the complete support of my parents, Phil and Grace, and my brother, Eugene, I would long ago have been lost in a dust pile.

TABLE OF CONTENTS

CHAPTER 1	INTRODUCTION	1
	References	5
CHAPTER 2	VERTEBRAL MORPHOMETRY	7
2.1	Introduction	7
2.2	Vertebral anatomy	8
2.3	Morphometric classification of vertebral deformities	14
	Smith-Bindman et al. study (1991a)	
	Hedlund and Gallagher study (1988)	
	Eastell et al. study (1991) : Graded classification	
	Minne et al. study (1988) : Spine deformity index (SDI)	
	Sauer et al. study (1991)	
	Smith-Bindman et al. study (1991b) : Index of radiographic area (IRA)	
	Our approach	
2.4	Chapter summary	32
	References	33
CHAPTER 3	X-RAY IMAGE ACQUISITION	35
3.1	Introduction	35
3.2	Physics of X-ray imaging	36
3.2.1	Image quality	40
3.3	Radiography	44
3.3.1	Conventional radiography	45
3.3.2	Digital radiography	47
3.3.2.1	Review of digital images	47
3.3.2.2	Computed Radiography	48
3.4	An experimental investigation :	
	Optimal lumbar spine X-ray image acquisition parameters	53

3.4.1	Introduction	53
3.4.2	Materials	54
3.4.3	Methods and results	55
	The method of measuring edge contrast	
	Comparing the linear attenuation coefficients of aluminum and bone	
	Determination of optimal kVp range	
	Determination of optimal mAs range	
	Effect of mAs on image statistics	
3.4.4	Discussion	63
3.4.5	Conclusions of the experiment	64
3.5	Chapter summary	65
	References	66
CHAPTER 4	DIGITAL IMAGE PROCESSING : METHODS	67
4.1	Introduction	67
4.2	Previous work	68
4.3	Image set	71
4.4	Software development environment	72
4.5	Methods	72
4.5.1	Image intensity based approach	73
4.5.2	Gradient search and cross-correlation detection	85
4.5.2.1	Image preprocessing	86
4.5.2.2	Cross-correlation vertebrae detection	87
4.5.2.3	Edge tracing and corner detection	89
4.5.3	Snakes	92
4.5.3.1	Contour initialization	95
4.5.3.2	Snakes energy minimization	96
4.6	Chapter summary	102
	References	103

CHAPTER 5	DIGITAL IMAGE PROCESSING : RESULTS	105
5.1	Introduction	105
5.2	Results of the three algorithms	106
5.2.1	Image intensity based approach (IIB)	107
5.2.2	Gradient search and cross-correlation detection (XES)	111
5.2.3	Snakes active contouring	125
5.3	Chapter summary	131
CHAPTER 6	DIGITAL IMAGE PROCESSING : DISCUSSION OF RESULTS	133
6.1	Introduction	133
6.2	Discussion of results : image characteristics	134
6.2.1	Image quality	134
6.2.2	Nature of lateral lumbar spine images	135
6.3	Discussion of results : algorithms	139
6.3.1	Performance and future directions	139
6.3.1.1	Image intensity based approach (IIB)	139
6.3.1.2	Gradient search and cross-correlation detection (XES)	141
6.3.1.3	Snakes contouring	144
6.3.2	Comparison of system and user requirements	145
6.3.3	What else do we need?	146
6.3.4	Which algorithm is best?	152
6.4	Chapter summary	152
	References	153
CHAPTER 7	CONCLUSION	155
7.1	Overview	155
7.2	Future directions	157
APPENDIX A	PROGRAM MODULES	159

CHAPTER ONE

INTRODUCTION

Osteoporosis is a condition affecting approximately 1 million Canadians today (Osteoporosis Society of Canada, 1995). It is the condition of reduced bone mass, resulting from aging, hereditary factors, dietary deficiencies, the lack of proper physical exercise, and various other factors. Since the bones of the human skeletal structure absorb the forces exerted on the human body, a loss of bone mass weakens such a structure and renders it susceptible to fracture. People with osteoporosis are at a high risk of bone fracture resulting from falls or minor injuries, or even from everyday activities such as lifting a child, which require considerable compressive and tensile strength from the skeleton. Osteoporotic fractures of the spine usually result in pain, kyphosis (hunchback effect), and loss of height, but may also be asymptomatic.

A common site for osteoporotic fractures is the vertebral body. Vertebral bodies are easily deformed by everyday stresses because the vertebrae are composed mainly of porous, trabecular bone (Fig. 1.1). In the assessment and follow-up of osteoporosis, physicians routinely request lateral X-ray images of the spine, for these can give a clear view of the skeleton. At first glance, vertebral body deformities may seem easy to detect on radiographs. However, in practice, it is a difficult task to determine from an X-ray image i) whether a vertebral body is actually deformed, and ii) to what extent it is deformed. Clinically, there exists no standard method for determining the presence and extent of vertebral deformities. Assessment is normally based on the subjective judgments of the physician or radiologist, judgments formed by comparing the shapes and relative sizes of the vertebral bodies, as seen on conventional X-ray films.

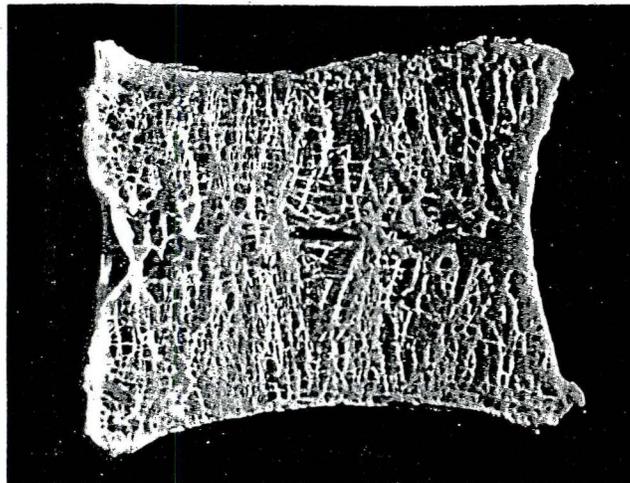


Fig. 1.1 The trabecular bone of vertebral bodies.

In order to objectify the process of assessing vertebral deformities from X-ray images, various groups have developed criteria for vertebral deformity classification. These criteria compare morphometric measurements (measurements of size and shape) of the vertebral bodies, in order to classify fractures and to quantify the presence and extent of vertebral deformities in a spine. These methods have always relied on the human observer to select appropriate points which define certain key dimensions of the vertebral body. This step is time-consuming and, more importantly, subject to inter- and intra-observer variability.

There is, therefore, a need to provide an objective means for performing measurements of vertebral body dimensions and to develop reliable deformity criteria, in order to standardize the entire process of vertebral deformity assessment from radiographs. With the ability to acquire digital X-ray images using Computed Radiography, and with the use of digital image processing techniques, our goal is to develop a computer-automated routine which performs the detection and measurement of vertebrae from digital spine radiographs. Using morphometric criteria and automated measurement, such a system would be able to provide a fast, accurate, and reproducible objective assessment of vertebral deformities. The aim is to produce, from this assessment, a single number which is indicative of the amount of vertebral deformity in a spine – a single number which is meaningful and useful for assessing and following the progression of osteoporotic vertebral fractures.

The goal of this thesis is two-fold : 1) to select morphometric criteria and combine them into a single index indicating extent of deformity, and 2) to develop software routines, based on digital image processing techniques, to automate the detection and measurement of these morphometric measures from lateral lumbar Computed Radiography spine images.

The general areas of investigation faced by this project are multidisciplinary, involving medicine, physics, and engineering. Chapter Two reviews fracture classifications and describes the selection of morphometric criteria for combining into a Total Vertebral Deformity Index. Chapter Three describes the physics of X-ray imaging, and details an experimental investigation we performed to determine the optimal X-ray imaging parameters for acquiring lateral lumbar spine Computed Radiography images. We pursued three approaches to designing an algorithm for the automatic detection and measurement of vertebrae: one, a grass-roots image intensity based approach to segmenting the vertebral outline; two, a cross-correlation vertebra detector with an edge gradient search for the vertebral body outline; and three, an active contouring algorithm called Snakes. Each method is described in detail in Chapter Four. The results of each method are presented visually in Chapter Five, and the issues plaguing each method are described and discussed in Chapter Six. Finally, Chapter Seven summarizes the goals, development, and results of the project. In addition, it includes a discussion of the clinical issues involved, and gives an overview of the broader scope of the project.

CHAPTER ONE

REFERENCES

Osteoporosis Society of Canada, Osteoporosis and Menopause Information Line, Toronto, Ontario, Canada.

Remagen W, 1989. Osteoporosis. Basle:Sandoz Ltd., p. 33.

CHAPTER TWO

VERTEBRAL MORPHOMETRY

2.1 INTRODUCTION

The vertebral bodies of the spine are primary sites for bone loss characterizing osteoporosis. The goal of quantifying vertebral body characteristics is to provide a single number which gives clinically meaningful information about the amount of vertebral deformity in a spine. Different groups have proposed criteria for classifying vertebral deformities based on morphometry, the measurement of sizes and shapes.

In this chapter, we first describe relevant vertebral anatomy. After defining the necessary terms, we review various morphometric classifications of vertebral deformities. We then develop our approach, and discuss its advantages and limitations.

2.2 VERTEBRAL ANATOMY

Spine

The human spinal column is made up of 33 vertebrae. There are 7 cervical vertebrae (labelled C1-C7), 12 thoracic (or dorsal) vertebrae (T1-T12), 5 lumbar vertebrae (L1-L5), 5 fused vertebrae composing the sacrum, and 4 fused pieces forming the coccyx (Fig. 2.1). A side view of the spine is called a lateral view. Spines imaged from the front of the body to the back are called anterior-posterior (AP) spine images (Fig. 2.2).

The vertebrae are designed expressly to support the weight of the upper body, and to protect the spinal cord (Warrick, 1976). The vertebrae of each region have different structures, depending on the type of structural load on the particular region. The two curvatures of the spine, kyphosis of the thoracic spine and lordosis of the lumbar spine, are designed to handle these loads (Fig. 2.1). In the lumbar region, the vertebrae are the largest in size, as they bear the larger forces on the skeleton.

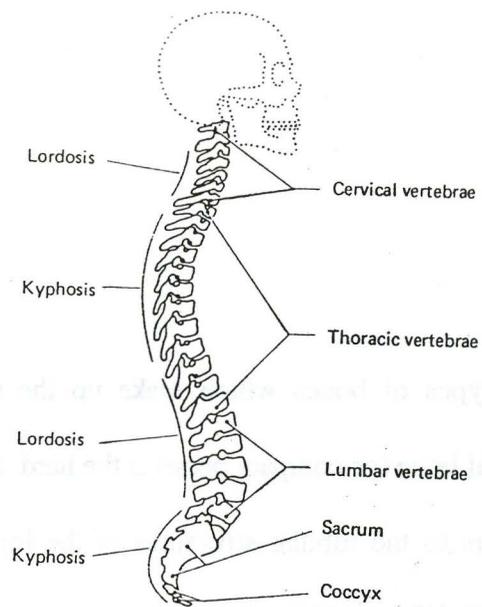


Fig. 2.1 The human spine.

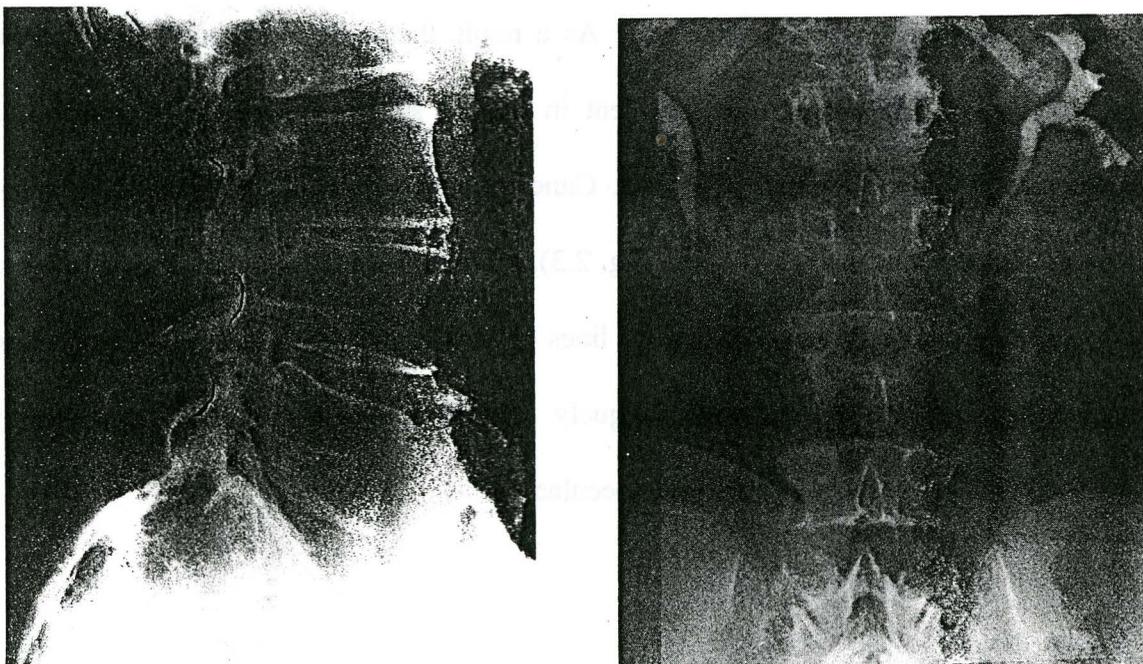


Fig. 2.2 (a) Lateral view of spine. (b) Anterior-posterior (AP) view of spine.

Bones

There are two types of bones which make up the skeleton: cortical bone and cancellous bone. Cortical bone (or compact bone) is the hard, dense bone which is familiar to most people. It refers to the tubular structures of the long bones, and makes up 80 percent of the skeleton (Griffiths, 1987). Cancellous bone (or spongy bone) is the porous bone which makes up the remaining 20 percent of skeletal bone. It is especially prominent in the vertebral bodies and in the flat bones such as the pelvis, ribs, and scapulae. Cortical bone is resorbed and formed at about 4 percent per year. In contrast, cancellous bone has a high, 20 percent per year turnover rate. As a result, the osteoporotic condition of bone mineral loss becomes much more evident in skeletal structures composed primarily of cancellous bone (eg. the vertebral bodies). Cancellous bone consists of a matrix structure of tiny bone segments called trabeculae (Fig. 2.3). Primary trabeculae are the trabeculae that are weight-bearing and aligned along the lines of mechanical stress. Secondary trabeculae are smaller, and transversely and obliquely oriented. In osteoporosis, the secondary trabeculae are lost before the primary trabeculae.

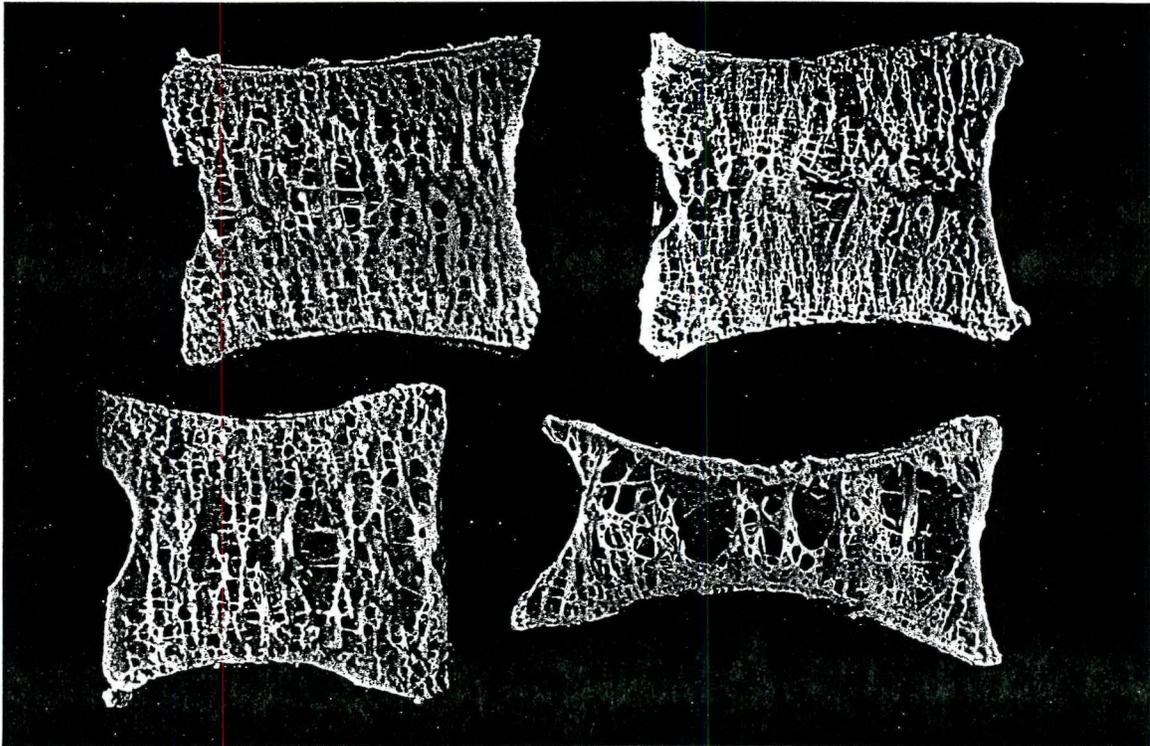


Fig. 2.3 Cross-sections of vertebral bodies showing trabecular bone structure.

Vertebrae

A vertebra consists of a vertebral body and a neural arch. The vertebral body consists of a **centrum** of cancellous bone, which is surrounded by a thin shell of cortical bone on its sides, and by end-plates on its upper (superior) and lower (inferior) surfaces. The pedicles and other processes projecting from the vertebral body constitute the neural arch. In between vertebrae are discs which allow for some movement between adjacent vertebrae (Fig. 2.4). The end-plates are composed of cartilage adjacent to the intervertebral disc, and of cortical bone adjacent to the cancellous bone of the vertebral body (Gertzbein, 1992). The intervertebral discs are made of X-ray lucent tissue (Armstrong et al., 1987), causing them to appear as empty spaces between vertebrae, on radiographs (Fig. 2.5).

The lumbar vertebrae are the largest and strongest of the spinal vertebrae (Bryan, 1970), and have large and wide vertebral bodies. Unlike the thoracic vertebrae, lumbar vertebrae are not joined to ribs. However, in lateral radiographs, the ribs from the lower thoracic vertebrae curve downwards and cause rib shadows to appear over the upper lumbar vertebrae (Fig. 2.5).

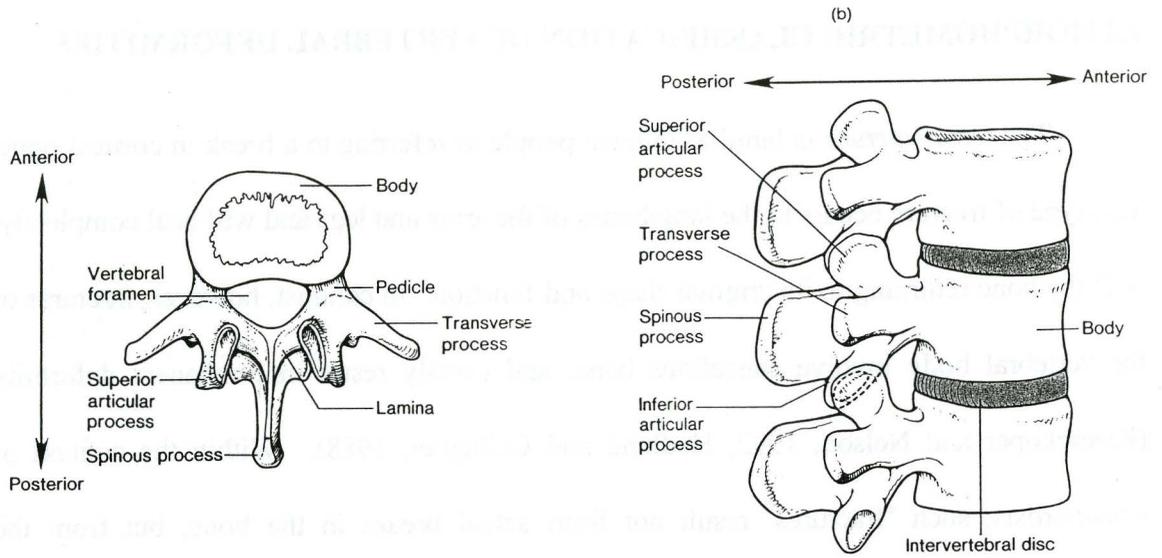


Fig 2.4 Lumbar vertebrae. (a) Top view. (b) Lateral view.

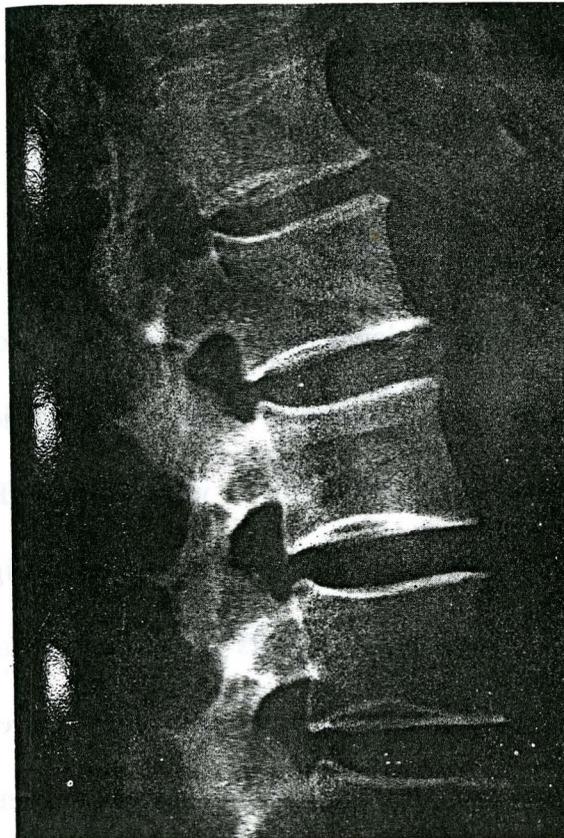


Fig. 2.5 Lateral view of the lumbar vertebrae showing rib shadows across the top two vertebrae, and the X-ray lucent intervertebral disc regions.

2.3 MORPHOMETRIC CLASSIFICATION OF VERTEBRAL DEFORMITIES

The term *fracture* is familiar to most people as referring to a break in cortical bone. This type of fracture occurs in the long bones of the arms and legs and will heal completely, with the bone returning to its original shape and function. In contrast, however, fractures of the vertebral body involve cancellous bone, and usually result in permanent deformity (Kleerekoper and Nelson, 1992; Hedlund and Gallagher, 1988). Within the context of osteoporosis, such 'fractures' result not from actual breaks in the bone, but from the resorption of bone mineral. Thus, the term *deformity* is more appropriate for describing osteoporotic vertebral body shape changes. However, since the term fracture is so commonly used in the literature to refer to vertebral deformities, both fracture and deformity will be used interchangeably in the following discussion.

There are three commonly defined types of vertebral body deformities. These are the wedge, biconcave, and compression deformities (Fig. 2.6). Wedge deformity is visible in a lateral view of the spine as a reduction of height on the anterior side of the vertebral body, so that the vertebral body appears to be wedge-shaped (Fig. 2.6a). Biconcave deformity occurs when the end-plates sag into the vertebral body (Fig. 2.6b). There is some uncertainty as to whether a vertebra that appears biconcave in a lateral radiograph is truly biconcave deformed, or whether the appearance of deformity is a result of off-centering of the X-ray beam on that vertebra (Hurxthal, 1968). Compression deformity is characterized by a relatively uniform reduction in height across the vertebral body (Fig. 2.6c).

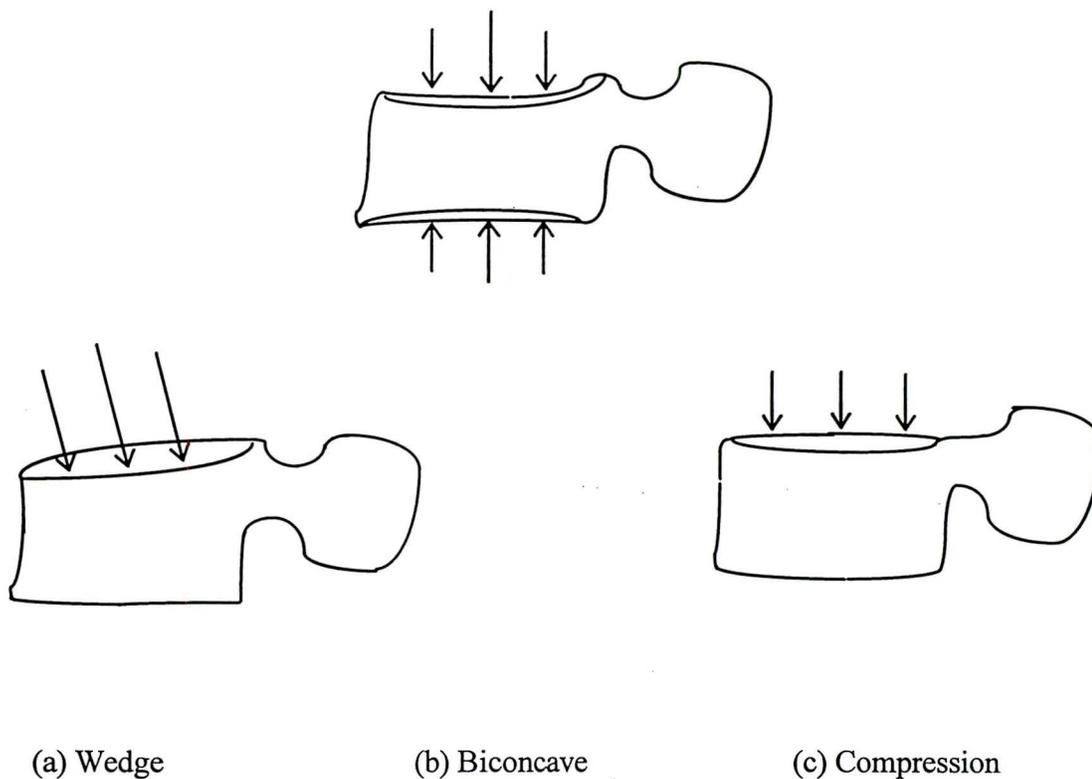


Fig. 2.6 The three commonly defined types of vertebral body deformities.

Several groups have attempted, with various levels of success, to establish reliable morphometric definitions for classifying vertebral deformities. Morphometric criteria are based on measurements of vertebral heights, widths, and areas (Fig. 2.7). There are several issues involved in generating morphometric criteria for vertebral deformity classification. A comprehensive study by Smith-Bindman et al. (1991a) will be used as a basis for introducing and comparing various commonly used morphometrics according to the issues involved. Following this discussion, several other studies and classification systems will be discussed.

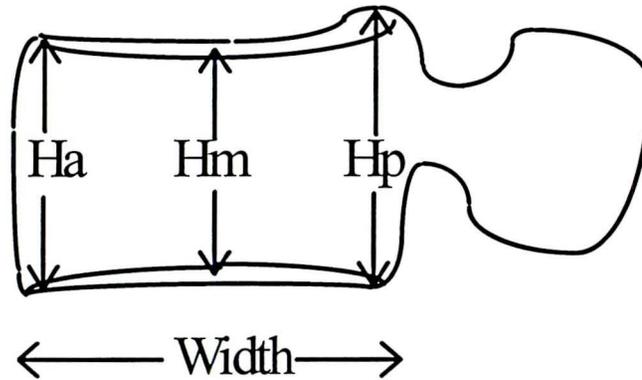


Fig. 2.7 Dimensions used in morphometric measurements of vertebral deformity.

Smith-Bindman et al. study (1991a)

Smith-Bindman et al. (1991a) conducted a comprehensive study to compare several common morphometric criteria for classifying vertebral deformities (Table II-1). Each of these criteria will be evaluated in the following discussion, according to how each accounts for issues of level-specificity, body size, multiple compressions, and multiple fracture combinations. These issues arise when using morphometric measurements to assess the severity of a vertebral deformity.

TABLE II-1. Morphometric Criteria

No.	Morphometric	Abbreviation
1	Anterior, posterior, or midvertebral height for a vertebral body	Ha, Hp, or Hm
2	Area of the polygon described by the six points placed around the margin of the vertebral body	Area
3	Heights or areas normalized by dividing these dimensions by the corresponding dimensions of T4	nHa, nHp, nHm or nArea
4	Ratios of anterior/posterior heights, mid/posterior heights	Ha/Hp, Hm/Hp normal Ha/Hp = 1.0
5	Reduction in the ratio of anterior/posterior heights compared to the normal ratio for a vertebral body	$\frac{Ha/Hp - Ha_n/Hp_n}{Ha_n/Hp_n} \times 100\%$ (level specific)
6	Differences between adjacent vertebrae in anterior, posterior, and midvertebral heights and area	$\Delta Ha, \Delta Hp, \Delta Hm, \Delta Area$
7	Ratios of anterior, posterior, and midvertebral height, wedge ratio, or area to adjacent (upper and lower) vertebrae	rHa, rHp, rHm r(Ha/Hp), rArea

Level-specificity

Level-specificity refers to the natural differences in vertebral size and shape at different levels of the vertebral column. It is important, when classifying vertebral deformities, to compare the measurements from each vertebra in the patient to

measurements from the same vertebral level in a normal population. In Criterion 4 of Table II-1, the H_a/H_p ratio for the normal population was assumed to be 1.0. Thus, this criterion does not account for natural variations in H_a/H_p at different vertebral levels. Criterion 5 compares the height ratios to the actual height ratios of the normal population, and in this way, accounts for level-specificity.

Body size

Larger people tend to have larger vertebrae (Smith-Bindman et al., 1991a). Therefore, adjustments for body size must be made. Absolute measurements such as those used for Criterion 1 are subject to higher levels of false fracture classification. Minne et al. (1988) proposed normalizing all heights to those of T4 (Criterion 3, Table II-1). While this method accounts for variations in body size in most cases, it becomes invalid if T4 is itself already deformed. By calculating ratios of heights within the same vertebra, Criterion 5 accounts for body size.

Multiple compressions

Criterion 7, Table II-1, compares ratios of measurements relative to adjacent upper or lower vertebrae. While this criterion succeeds in accounting for different body sizes, it will introduce false negative fracture classifications if multiple, successive, vertebral compressions occur along the spine. To deal with multiple compressions, some combination of criteria, or some dimension other than height must be considered.

Multiple fracture combinations

Each vertebra can experience some combination of the three deformity types. Smith-Bindman et al. also tested fracture definitions that combined criteria involving each of the three heights. If a vertebra satisfied any one of the three criteria in a particular triplet combination, then it was classified as fractured. This method was shown to be more sensitive than criteria based on single dimensions.

Criteria 2 and 6

The two remaining criteria in Table II-1, Criteria 2 and 6, differ in their limitations. Criterion 2, vertebral body area, was not found to be useful for fracture classification (Smith-Bindman et al. (1991a), Hedlund and Gallagher (1988)). Criterion 6, taking differences in dimensions, is problematic in two ways : first, it is subject to multiple compression effects because it compares the heights of adjacent vertebrae; second, it does not account for the fact that for small vertebrae, a difference in one dimension signifies a more severe deformity than the same difference in larger vertebrae. In other words, Criterion 6 does not adjust for body size effects.

Table II-2 summarizes the preceding critiques of each of the morphometric criteria in Table II-1. It is clear that the most useful of these criteria is Criterion 5, which compares the ratio of anterior to posterior heights with the same ratio at the same vertebral level in the normal population. This criterion can easily be extended to ratios of the other heights.

TABLE II-2. Critique of Table II-1 Morphometrics

No.	Critique
1	× not adjusted for body size
2	× no specific benefit
3	× body size adjusted; subject to deformed T4
4	× not level-specific
5	✓ ratio adjusts for body size;
6	× not adjusted for body size; subject to multiple compressions
7	× subject to multiple compressions

Hedlund and Gallagher study (1988)

Hedlund and Gallagher (1988) studied the usefulness of five morphometric criteria: i) wedge angle (the angle between the superior border and the vertebral width), ii) percent reduction in anterior height compared to posterior height, iii) percent difference in anterior height between adjacent vertebrae, iv) absolute anterior height, and v) vertebral body area. Of these five morphometrics, it was found that area measurements were not useful in diagnosing vertebral fractures. Moreover, it was determined that the percent difference in the anterior heights of adjacent vertebrae was the most useful of the five criteria, when used in conjunction with the absolute anterior height in cases of multiple compression. This

combination accounts for level-specificity, body size differences, and multiple compressions.

Eastell et al. study (1991) : Graded classification

Eastell et al. (1991) proposed a comprehensive classification system which includes the type, grade, and number of deformities. They made measurements of the percent change in anterior and mid heights (for wedge and biconcave deformities, respectively) relative to the posterior heights of each vertebra. For compression deformity, they measured the percent change in the posterior heights of adjacent vertebrae. As we can see, this measure does not account for multiple compressions. Subsequently, comparison with normal data provided the following classifications: normal as being within 3 standard deviations of the mean; grade 1 deformity as being between 3 and 4 standard deviations of the normal mean; and grade 2 deformity as being greater than 4 standard deviations. However, these ranges were somewhat arbitrarily chosen using thresholds of 15% and 25% reductions in vertebral height. Gallagher et al. (1988) have produced data, from a normal population, showing that the range of lower limits (mean-2SD) of percent change in anterior height relative to posterior height varies from 18.5% to 25.2%, so that a 15% reduction would in fact be within normal variation. While there is merit in defining grades of deformity, the classification system of Eastell et al. produces variable results depending on the percent of height reduction thresholds chosen.

Minne et al. study (1988) : Spine deformity index (SDI)

The spine deformity index (SDI) of Minne et al. (1988) aims at identifying the presence of vertebral deformities, as well as at quantifying the extent of deformation. First, the heights (anterior, posterior, mid) of each vertebra are measured and normalized to the respective heights of the fourth thoracic vertebra, T4. The choice of T4 was based on the observation that T4 is rarely deformed in osteoporotic spines (Minne et al.'s data set showed that 1 out of a total of 39 spines had a deformed T4). A vertebra is considered fractured if any of the three T4-normalized heights falls below the lower limit of the normal range. The deviations of the three heights from their respective lower limits are summed to give a vertebral deformity index (VDI) for each vertebra. The sum of the VDIs from T5 to L5 yields the spine deformity index (SDI) (Fig. 2.8). The index of deformity obtained by this method accounts for body size differences (by taking relative heights), for multiple compressions (by non-reliance on adjacent vertebral dimensions), and for multiple fracture combinations. As well, by comparing the values to normal data at each vertebral level, this measure is level-specific.

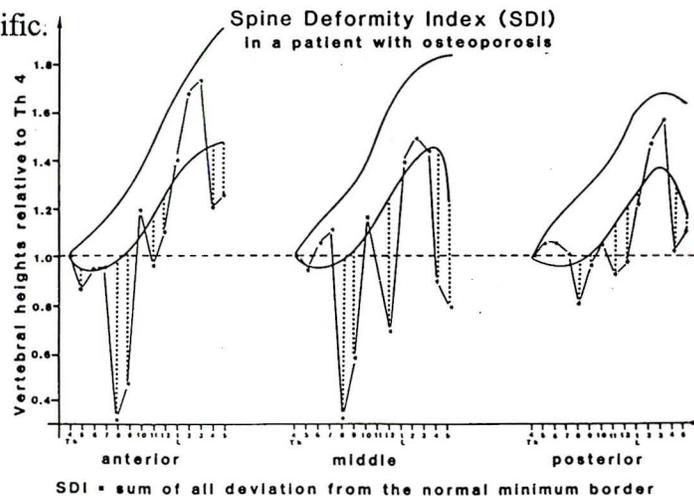


Fig. 2.8 Measured vertebral heights (dashed lines) compared to normal ranges (solid curves) of heights relative to T4. Summing the deviations (dotted lines) from the lower limit of the normal range gives the SDI for this spine.

Sauer et al. study (1991)

A study by Sauer et al. (1991) compared four classification systems for fracture identification. Three of these were based on the ratio of heights from a vertebra, or ratios of heights from adjacent vertebrae. These three systems aimed primarily at classifying vertebrae as fractured or normal, and counted the number of deformed vertebrae in a spine. The fourth system was the spine deformity index of Minne et al., which normalizes anterior, mid, and posterior heights to the respective heights of T4. Sauer et al. determined that disagreement among the four methods over whether a single vertebra is fractured or not was a result, not only of borderline normal/fractured cases, but also of weaknesses in the various fracture criteria, arising from the issues which we have discussed in the preceding paragraphs. In addition, this study showed that the number of vertebral fractures in a spine does not necessarily reflect the extent of the vertebral deformities.

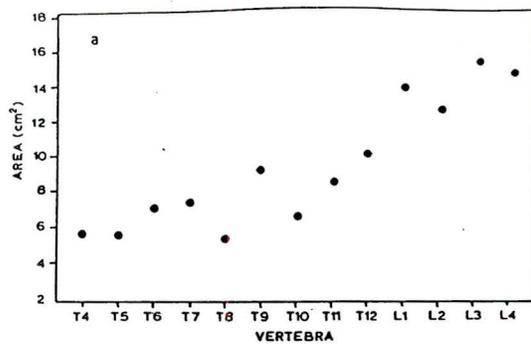
Smith-Bindman et al. study (1991b) : Index of radiographic area (IRA)

Along similar lines as the Minne et al. SDI index, Smith-Bindman et al. (1991b) developed an Index of Radiographic Area (IRA). Like the SDI, the IRA measures anterior, mid, and posterior heights and normalizes them to T4 heights. Additionally, the IRA calculates the area of the polygon defined by these points. These vertebral areas are plotted against their corresponding vertebral levels. Then, using any criterion or combination of criteria (including those we have already discussed), each vertebra is classified as normal or abnormal. (It was found that the final IRA values were quite independent of the choice of

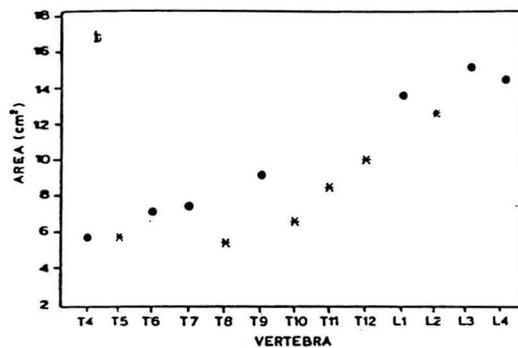
this classification criterion, and thus, any reasonable criterion can be used.) On the plot, the areas of the 'normal' vertebrae are fitted by a fourth order polynomial, which interpolates expected normal areas for the vertebrae that are classified as abnormal. The expected normal areas and the observed areas are then used in the calculation of the IRA score. Either an absolute IRA score or a percentage IRA score can be calculated, where the absolute IRA score is the sum of the relative reductions in area of the abnormal vertebrae, and the percentage IRA score is the mean of the ratios of observed over expected areas (Fig. 2-9 (a)-(e)). Whereas the SDI uses the *minimum* values of the normal population as reference normal values, the IRA uses expected areas calculated specifically for each person, according to the fourth order polynomial model. In this way, there is minimal reliance on the absolute range of normal data, and minimal effect from natural individual variations in vertebral dimensions. Because the IRA measure depends on polynomial fitting of the undeformed vertebral areas, enough of these must be present to produce an appropriate curve.

Our approach

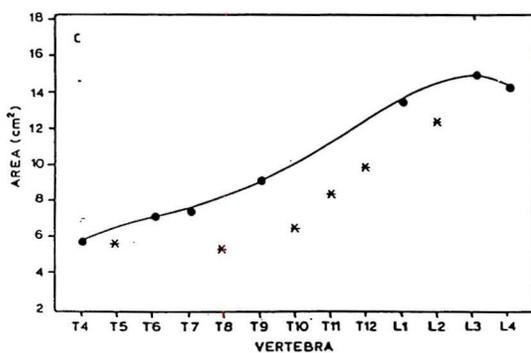
Traditionally, there have been two ways of assessing vertebral deformities using morphometry. One approach aims primarily for the classification of vertebra as normal or fractured, and then counts the number of fractured vertebrae as an indication of the amount of vertebral deformity in a patient. The other approach to quantifying vertebral deformities, as we have seen, is by calculating an index which measures in some way the extent of the



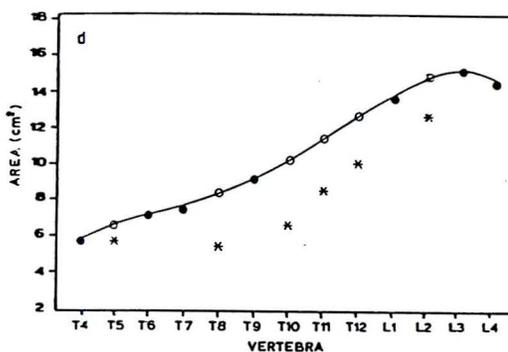
(a)



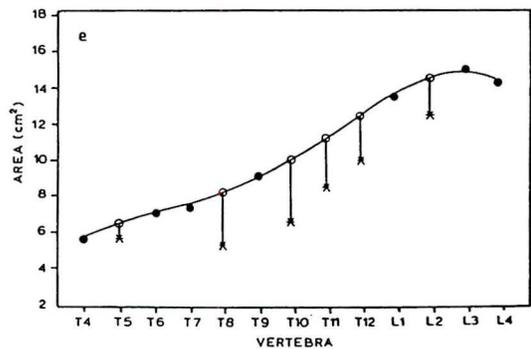
(b)



(c)



(d)



(e)

Fig. 2.9 Steps in computing the IRA :
 (a) Measured vertebral areas for a spine with several deformed vertebrae. (b) Classification of normal (●) and abnormal (*) vertebrae. (c) 4th order polynomial fit to the normal vertebrae. (d) ○ = normal expected areas for all deformed vertebrae. (e) Deviations of abnormal from normal areas.

vertebral deformities in a spine. Both approaches are related, and in fact, the second approach normally includes the first approach as a preliminary step. As Sauer et al. have shown, a given number of vertebrae classified as fractured could, in fact, represent vastly differing extents of vertebral deformation. Knowledge of the number of fractures is useful for epidemiological studies and as eligibility criteria for clinical trials (Eastell et al.). However, the clinical purpose of our work suggests that the number of fractured vertebrae in a spine is insufficient information for assessing the spinal osteoporotic condition. A patient may reveal the same number of fractures in a followup assessment six months later, but in fact, may be experiencing more severe pain or kyphosis as a result of more pronounced deformation of already deformed vertebrae. Therefore, the purpose of our work, to provide a single number useful for assessing vertebral deformities in osteoporotics, requires that the second more comprehensive index approach be used. A combination and modification of the SDI and IRA methods may prove most useful. Such a system will be outlined below.

First of all, it would be useful to classify each deformity type separately. The following morphometrics may be defined : for wedge deformity, the ratio of anterior to posterior height; for biconcave deformity, the ratio of mid to posterior height; for compression deformity, the ratio of posterior height to inferior width (Fig. 2.10). We recognize the first two morphometrics as body size independent ratios. We introduce the third morphometric for describing compression deformity which, unlike wedging or biconcavity, exhibits a shrinking of the normally stable posterior height. There is no

evidence that vertebral width experiences significant variation in vertebral deformation (Smith-Bindman et al., 1991b). Consequently, all three morphometrics are body-size independent and unaffected by multiple compression effects.

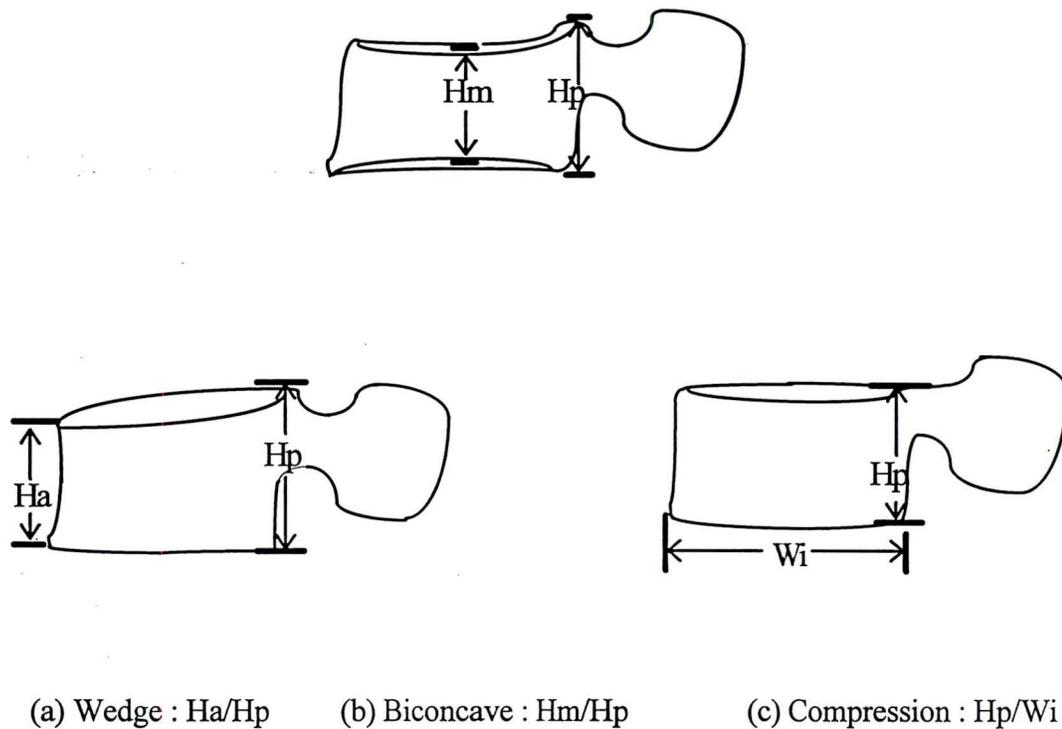


Fig. 2.10 Morphometric measurements for proposed approach.

A combination of the SDI and IRA procedures can then be followed. Following the SDI procedure, each morphometric measurement is plotted separately, against vertebral level, to produce three graphs, one for each deformity type. Then, each individual measurement is determined to be normal or abnormal, using any suitable criterion (such as a lower limit

threshold, for example). On each of the three graphs, polynomial fitting of the normal measurements is then done to estimate expected values for the abnormal vertebrae, as in the IRA method. The order of the polynomial can be empirically chosen, depending on the number of vertebrae measured. The deviations of the observed from the expected values are then used in the calculation of the index.

for each vertebral level, i :

$$VDI_i = d_{wedge} + d_{biconcave} + d_{compression} \Leftrightarrow \text{Vertebral deformity index (Eq. 2.1)}$$

for the spine :

$$TVDI = \sum_i VDI_i \Leftrightarrow \text{Total VDI (Eq. 2.2)}$$

The TVDI provides the single number indicative of the amount of vertebral deformation in a spine. This index can be used for quick and easy comparisons of the progression of deformities, in followup assessments of the osteoporotic condition in the spine. Calculating each VDI individually also provides additional information to the clinician as to the location and severity of the deformities – information which would be useful for making treatment decisions.

This method of vertebral deformity quantification is based on a combination of the SDI and IRA scoring procedures, with the exception of the morphometric measure for compression deformity. Aside from the inherent similarities in all three methods, in terms

of advantages and limitations, there are several significant differences which were designed to make our method applicable to our circumstances. The most significant difference is in the normalization of our morphometric measurements. Two related possible situations prompted the decision not to normalize heights to T4 heights, as was done by both SDI and IRA groups. The first situation, a deformed T4 (acknowledged by both SDI and IRA groups), and the second, the possible unavailability of thoracic spine radiographs, render normalization to T4 impractical. As well, by normalizing to T4, the accuracy of the ratios are dependent on the accuracy of the points chosen to define T4's vertebral dimensions. We, therefore, normalize each of the three measurements to each vertebra's posterior height. Doing so does not affect assessment of any of the three deformities : the posterior side of the vertebra is supported by the neural arch and its attached processes, thereby ensuring that a wedge deformity is always a reduction in anterior height over posterior height, and not vice versa; biconcave deformity affects only the midline of the vertebral body; and compression deformity measurements are stabilized by the inferior width measurement. Another difference in the methods arises from the possibility that one or more vertebrae may not be clearly delineated, or a vertebral corner may be clipped from the image, so that one or more of the heights cannot be measured. Thus, the total number of vertebrae is no longer regular across patients, and TVDI scores cannot be compared with each other. We therefore normalize the TVDI by the number of measured vertebrae.

As already stated, one possible problem resulting from polynomial modelling of the measurement points is that enough 'normal' vertebrae must be identified for a reasonable polynomial fit to be achieved. This problem was not mentioned in either the SDI or the IRA reports. If, for example, all of the vertebrae visible on a particular radiograph were deformed, there would be no normal measurements left for the polynomial fitting. In such a case, the method could conceivably revert to the SDI method of using the lower limit of the range of measurements from the normal population.

One benefit of morphometric quantification of vertebral deformities is that, because only a certain number of points need to be defined on a vertebra, the procedure need not be done by radiologists, but can be done by well-trained technologists (Smith-Bindman et al., 1991b). On the other hand, the task of defining a set of points on each vertebra in a radiograph is significantly time consuming, requiring around 15 to 30 minutes per spine (Smith-Bindman et al., 1991b). Furthermore, regardless of who the observer is, it is sometimes quite difficult to decide where to place a point. For example, the midpoint on the superior endplate, used for measuring mid height, may be ambiguous, because there are sometimes two shadows of the vertebral cortex in the superior and the inferior horizontal planes (Fig. 2.11). For problematic situations such as this, detailed protocols for locating the most representative points must be drafted ahead of time and consistently followed.

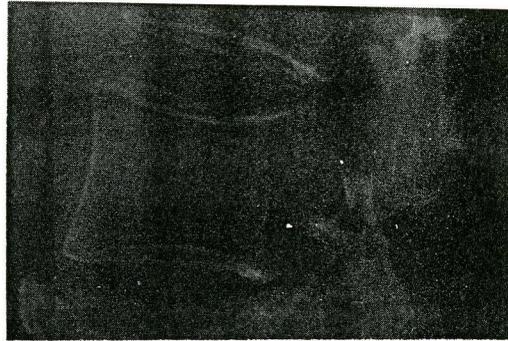


Fig. 2.11 Two cortical outlines are apparent on both the superior and inferior boundaries of this vertebral body.

In spite of the preceding concerns, the method developed here for quantifying vertebral deformities is deemed to be potentially useful. It takes the more comprehensive view that it is more meaningful and useful to quantify the *extent* of deformation rather than just to classify fractured and normal vertebrae. By providing a deformity index for each vertebral level, the location and severity of the deformities are identified. By calculating a total deformity index for a group of vertebrae, a single meaningful number is generated.

2.4 CHAPTER SUMMARY

In this chapter, we defined terms from skeletal and vertebral anatomy, and discussed the meaning of the term *fracture*. Preference for the term *deformity* was explained in the context of vertebral bone loss leading to vertebral body shape changes. Having defined the relevant anatomical and clinical terms, we surveyed several systems for the morphometric assessment of vertebral deformities, and discussed issues of level-specificity, body size adjustment, multiple compression fractures, and multiple fracture combinations. We compared systems using a binary (fractured/normal) classification, with systems using a continuous numerical index which provides information on the extent of the vertebral deformities. In addition to various other systems, we discussed the spine deformity index (SDI) of Minne et al., and the index of radiographic area (IRA) of Smith-Bindman et al., and from these proposed an index which is more suitable to our available data. The calculation of this index is based on both the SDI and IRA procedures, but does not require normalization to T4 dimensions, and does not assume a constant total number of vertebrae. The level-specific indices and the total spine index provide meaningful and useful measures for assessing the presence and progression of vertebral deformities in a spine. Inherent in morphometrical quantification are issues of reproducibility of the choice of points, of the time involved in marking these points, and of measurement accuracy. It is with these concerns in mind that we aim to develop an efficient computer-automated procedure which will make consistent decisions and precise measurements for vertebral deformity assessment.

CHAPTER TWO

REFERENCES

- Armstrong P, Wastie ML, Brenbridge A, Norman AG, 1987. Diagnostic imaging. 2nd ed. Oxford: Blackwell Scientific Publications.
- Bryan GJ, 1970. Diagnostic radiography: a concise practical manual. Edinburgh: E&S Livingstone.
- Eastell R, Cedel SL, Wahner HW, Riggs BL, Melton LJ III, 1991. Classification of vertebral fractures. *J Bone Miner Res* 6:207-215.
- Gallagher JC, Hedlund LR, Stoner S, Meeger C, 1988. Vertebral morphometry: normative data. *Bone Miner* 4:189-196.
- Gertzbein SD, 1992. Fractures of the thoracic and lumbar spine. Baltimore, MD: Williams & Wilkins.
- Griffiths HJ, 1987. Basic bone radiology. 2nd ed. Norwalk, CT: Appleton & Lange.
- Hedlund LR, Gallagher JC, 1988. Vertebral morphometry in diagnosis of spinal fractures. *Bone Miner* 5:59-67
- Hurxthal LM, 1968. Measurement of anterior vertebral compressions and biconcave vertebrae. *Am J Roentgenol* 103:635-644.
- Kleerekoper M, Nelson DA, 1992. Vertebral fracture or vertebral deformity? *Calcif Tissue Int* 50:5-6.
- Minne HW, Leidig G, Wüster C, Siromachkostov L, Baldauf G, Bickel R, Sauer P, Lojen M, Ziegler R, 1988. A newly developed spine deformity index (SDI) to quantitate vertebral crush fractures in patients with osteoporosis. *Bone Miner* 3:335-349.
- Remagen W, 1989. Osteoporosis. Basle:Sandoz Ltd., p. 4
- Sauer P, Leidig G, Minne HW, Duckeck G, Schwarz W, Siromachkostov L, Ziegler R, 1991. Spine deformity index (SDI) versus other objective procedures of vertebral fracture identification in patients with osteoporosis: a comparative study. *J Bone Miner Res* 6:227-238.
- Smith-Bindman R, Cummings SR, Steiger P, Genant HK, 1991. A comparison of morphometric definitions of vertebral fracture. *J Bone Miner Res* 6:25-34.

Smith-Bindman R, Steiger P, Cummings SR, Genant HK, 1991. The index of radiographic area (IRA) : a new approach to estimating the severity of vertebral deformity. *Bone Miner* 15:137-150.

Van de Graaff KM, 1992. *Human Anatomy*. 3rd ed. Dubuque, IA:Wm. C. Brown Publishers, p. 169.

Warrick CK, 1976. *Anatomy and physiology for radiographers*. 5th ed. Edward Arnold Ltd.

CHAPTER THREE

X-RAY IMAGE ACQUISITION

3.1 INTRODUCTION

Medical images are generally acquired for diagnostic purposes. Therefore, image quality as determined by clinicians may differ from the image quality that may be desirable for digital image processing and analysis. It is important when doing any signal processing, to first obtain as strong a signal as possible in the data. The signal in our study is the image intensity of the vertebral bone. While maximizing this signal, we also want to minimize all other signals, including those from the soft tissues and other foreign objects in the body. It is preferable to maximize image quality at the time of acquisition, in order to obtain images of uniformly similar quality, rather than to develop digital processing methods robust

enough to handle a wide range of image quality. In X-ray imaging, the intensity levels in the image depend upon the amount of attenuation of X-ray energy by the different body tissues. The physical parameters controllable by the technologist are the distribution of X-ray energies in the X-ray beam, and the intensity of the X-ray beam, which are determined by kVp (peak kilovoltage) and mAs (milliamperere-seconds), respectively. In this chapter, we first discuss the physics of X-ray imaging, as well as various factors influencing the quality of X-ray images. Then, we briefly describe conventional radiography technology, review the concept of the digital image, and outline the principles and advantages of Computed Radiography (CR) technology. Finally, we discuss in detail an experimental investigation in which we attempted to determine the optimal X-ray imaging parameters, kVp and mAs, for acquiring lateral lumbar spine images.

3.2 PHYSICS OF X-RAY IMAGING

X-rays are generated when fast-moving electrons collide with an object. This is achieved by applying an electric potential (usually 50-150kV) across the two elements (cathode and anode) in a vacuum tube. A filament heated to incandescence is the source of electrons. The electric potential attracts the free electrons from the filament to the anode, causing the electrons to accelerate through the tube towards the anode. There, a high heat dissipating material (tungsten) is attached to the anode. This block of tungsten becomes the site of the high-speed collisions which release X-ray energy (Armstrong et al., 1987). X-rays are simply photons of energy in the X-ray range of the electromagnetic spectrum.

The quantity of X-rays produced by the X-ray tube is controlled by the mAs parameter:

$$\text{mAs} = (\text{mA}) * (\text{s}) \quad (\text{Eqn 3.1})$$

Number of electrons = tube current * duration of exposure
produced in the tube

The peak kilovoltage (kVp) applied across an X-ray tube (Fig. 3.1) controls the distribution of energies in the X-ray beam (Fig. 3.2). The higher the kVp, the higher the electric potential across the tube, and the greater the kinetic energy gained by the electrons. A small proportion (less than 1%) of this kinetic energy is transformed by the high-speed collision into X-ray energetic photons. The rest of the energy is released as heat into the high heat absorbing tungsten block. Thus, a high kVp produces more higher energy X-rays in the X-ray beam.

The fate of X-ray photons depends on how they interact with the matter that is being irradiated. X-rays can be either a) absorbed by the matter (photoelectric interaction), b) scattered or deflected by the matter (Compton interaction), or c) penetrate right through the matter (no interaction). When X-ray photons are absorbed or scattered by matter, the intensity of the X-ray beam decreases, ie. the beam is attenuated. The relative rate of X-ray beam attenuation is expressed by a linear attenuation coefficient, μ . This value, in mm^{-1} , represents the probability that a photon will interact with the matter as it travels through, and depends on three main factors : a) the energy of the photon, b) the density of the matter,

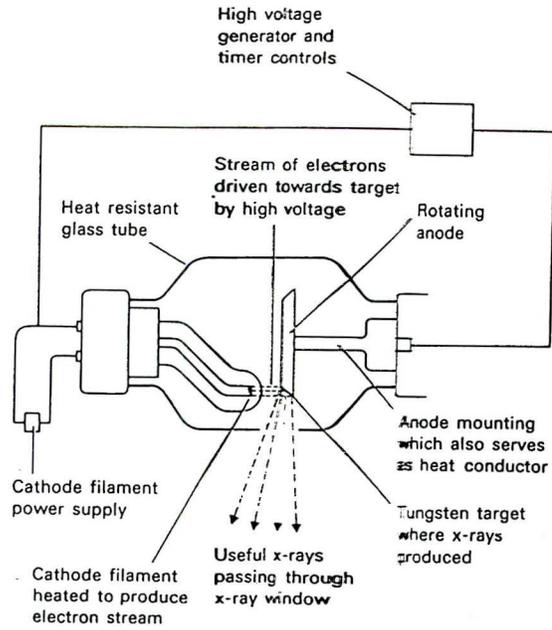


Fig. 3.1 X-ray tube.

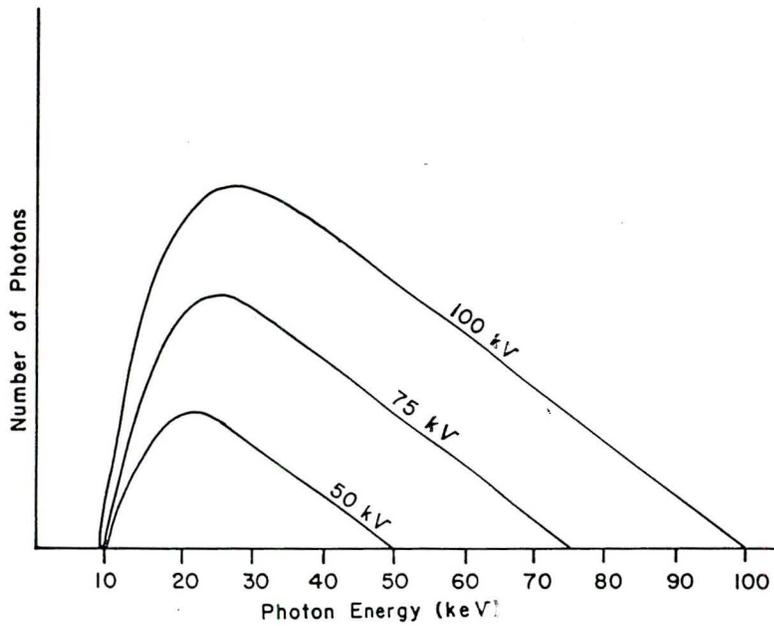


Figure 6. Comparison of photon energy spectra produced at different kVp values.

Fig. 3.2 Comparison of photon energy spectra produced at different kVp values.

and c) the atomic number of the matter. It is important to remember that linear attenuation coefficient values are not constant values, but must be taken at the particular photon energies being considered. While low energy (soft) X-rays are absorbed by the object being imaged, high energy (hard) X-rays are able to penetrate the object and expose the film or image plate behind the object. There is a directly proportional relation between object density and linear attenuation coefficient, because the concentration of electrons in a material depends primarily on its physical density (Sprawls, 1977). The atomic number, Z , is the number of protons in an atomic nucleus, and is equivalent to the number of electrons in a neutral atom. It gives a relative indication of the amount of electron interaction that will cause X-ray attenuation by that element. When a material is composed of more than one element, an effective atomic number is used. The four basic types of tissue in the human body each absorb X-ray photons to different degrees. These tissue types are : gas, fat, soft tissues other than fat, and calcified structures (such as bone). X-ray photons interact minimally with gas or air atoms and molecules, and therefore most of the photons pass through gas or air to cause the darkest image shadow on the radiographic image. In contrast, bone, composed primarily of calcium (which has a relatively high atomic number of 20) has high X-ray attenuation, and appears white on radiographs. Soft tissues, such as muscle, blood, etc., but excluding fat, have similar attenuation coefficients to each other, and are classed together. They appear in the same range of grey level intensities on the image. It is useful to note that because soft tissue (excluding fat) is approximately 75% water, it has similar attenuation characteristics as water, and is therefore usually simulated by water, in a body phantom (Hendee and Ritenour, 1992). Fat absorbs slightly less X-ray

energy than the other soft tissues, and so appears slightly darker in the image (Armstrong et al., 1987). It is precisely because of differential X-ray attenuation abilities that an X-ray beam passing through a region of the body is able to image different tissue structures.

Table III-1 Effective atomic number and physical density of the main tissue types

Material	Eff. atomic no.	Density (g/cm ³)
air	7.6	1.205x10 ⁻³
water	7.4	0.9982
soft tissue (muscle, blood, etc.)	7.4	1.04
fat	5.9-6.3	0.92
bone	11.6-13.8	1.85

[Source : Attix FH, 1986. Introduction to Radiological Physics and Radiation Dosimetry, John Wiley & Sons, Inc., p.531]

3.2.1 IMAGE QUALITY

In general, the quality of an X-ray image refers to the visibility of structures in the image, especially for diagnostic purposes. Several physical factors affect the signal-to-noise ratio in the image, resulting in improved or degraded image quality.

mAs. As the tube current, mA, is varied, the quantity of X-rays produced varies proportionally, as given by Equation 3.1. Changing mAs increases or decreases beam intensity, resulting in better or worse image statistics, respectively, but not affecting image contrast (Fig. 3.3).

kVp. Changing *kVp* has two major effects on the resulting image:

- i. increasing *kVp* produces higher energy X-rays which are more penetrating, hence decreasing the relative intensities among tissue types
- ii. increasing *kVp* causes beam intensity to increase

Film graininess is reduced by raising the *kVp* so that a higher proportion of penetrating X-rays are generated in the X-ray beam energy spectrum. More penetrating X-rays means that more signals are captured on the film or image plate, and therefore less graininess results from not having enough signals. Thus, the first two effects combine to produce smoother, less noisy images. The first effect results in decreasing the contrast between the different tissue types in the image (Fig. 3.4). The second effect is actually a fault of X-ray tubes, and not a theoretically predictable phenomenon.

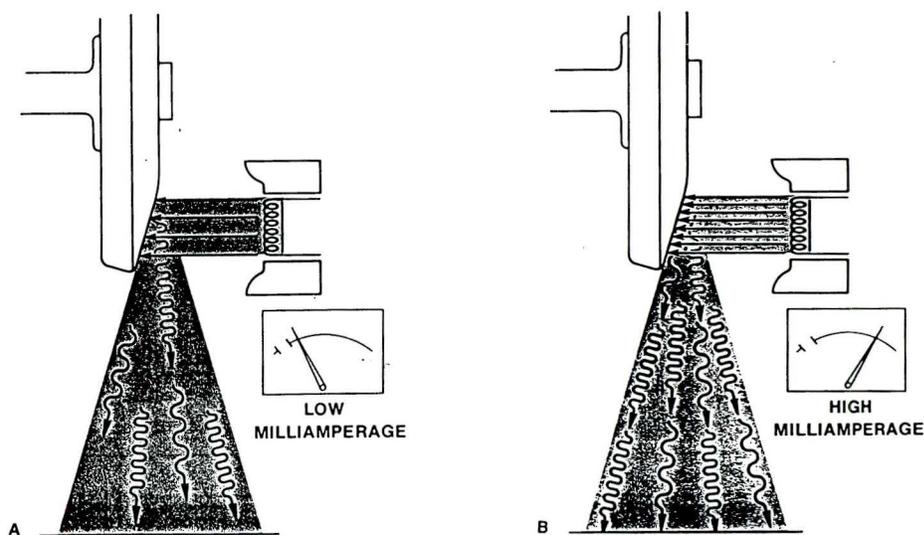


Fig. 3.3 Effect of mAs on the X-ray beam energies.

As the mA is decreased (A) or increased (B), the number of electrons made available for use in the generation of the X-ray beam decreases or increases and the intensity of the X-ray beam changes accordingly.

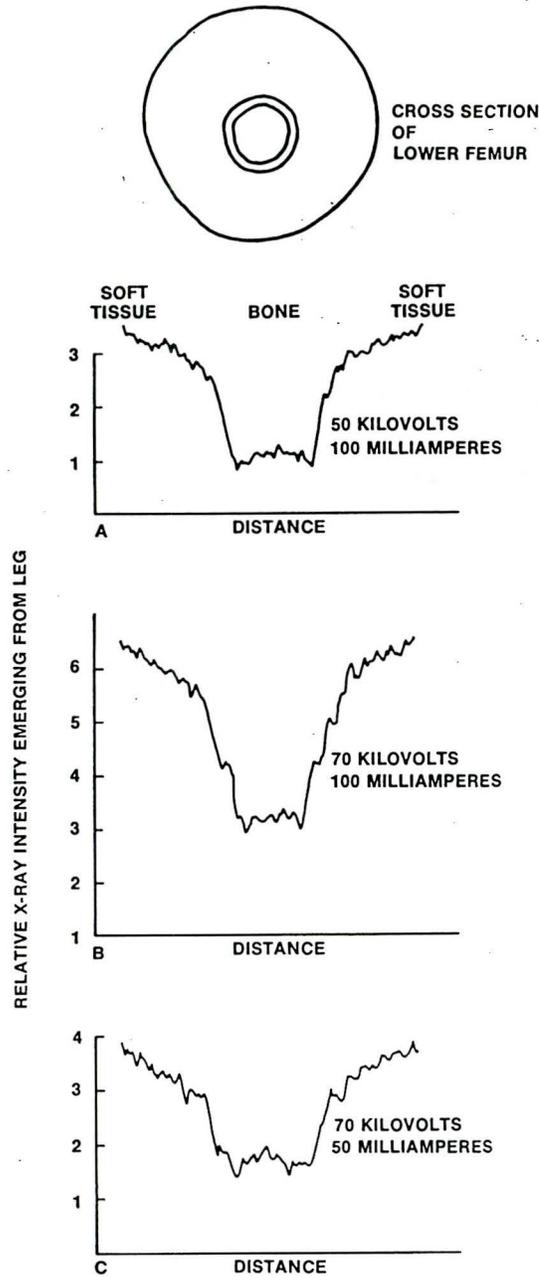


Fig. 3.4 Effect of kV and mA on image contrast.

Three graphs illustrating the different relative X-ray intensities emerging from the cross-section of a leg, resulting from changes in kilovoltage and milliamperage. Comparing the intensity distribution for a 50 kV beam (a) with that of a 70 kV beam (b) shows that as kV is increased, the beam intensity and penetrating power increase and subject contrast decreases (from 3 to 2 in this example). Comparison of the patterns resulting from reducing milliamperage from 100 mA (b) to 50 mA (c) shows that X-ray intensity has been halved but that the corresponding intensities bear the same ratio to each other — that is, contrast is not altered by changes in milliamperage.

Scattered radiation. Scattered radiation is radiation that is dispersed in oblique directions by the object being imaged. Besides being undesirably absorbed by body tissues that are not meant to be imaged, scattered radiation creates unwanted noise in the image. Grids are used to filter out scattered radiation. Grids are made of lead strips aligned in the desired directions (usually parallel to the X-ray beam). A grid is placed between the object and the film or image plate, thereby allowing only the components of the beam in the desired directions to pass through to the film (Fig. 3.5).

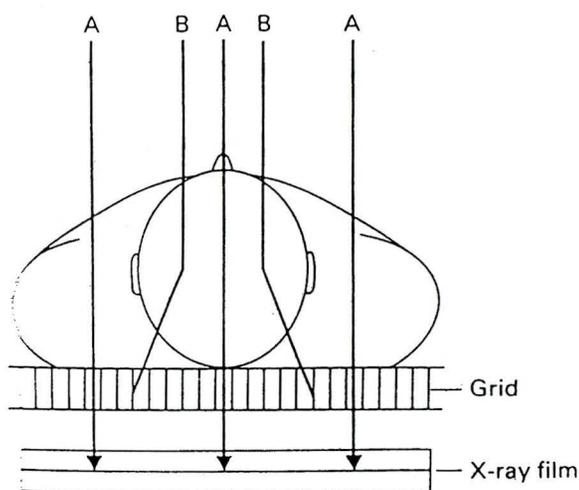


Fig. 3.5 Principle of the grid.

The beams labelled A pass through both the patient's body and the grid to expose the film. The two beams labelled B are scattered within the body : they are absorbed by the slats of the grid and therefore do not reach the film.

Clinical considerations. The acquisition parameters chosen for an X-ray examination depend on the region of the body being examined, the structures to be imaged, and the physical condition of the patient (Bryan, 1970). Patient size (thickness) may require that a higher kVp be used in order to produce sufficient tissue contrast in the image. A patient suffering from osteoporosis will have decreased bone density, and therefore may require a reduced kVp for an image of greater contrast, although this may lead to greater absorbed radiation. The ability of x-radiation to alter biological cells is a hazard of X-ray imaging, and therefore any determination of clinical X-ray acquisition parameters must consider the dose to the patient. It is imperative to generate as low as possible a radiation dose to the patient, while generating enough X-ray intensity for an image to be captured on the film or imaging plate. This intensity balance is controlled by the mAs parameter. When setting the kVp, a compromise must be found between a kVp low enough to yield sufficient image contrast, yet high enough to reduce absorbed patient dose and film graininess.

3.3 RADIOGRAPHY

Radiography is the process of capturing, storing, and displaying an X-ray image. Two technologies are currently in use: conventional radiography, and the newer computed radiography.

3.3.1 CONVENTIONAL RADIOGRAPHY

Conventional radiography, the technology of the past century, relies on the principles of photography for capturing X-ray exposures onto film. The film is inserted into a film cassette, which is then placed on the side of the patient away from the X-ray tube. Silver halide crystals in emulsion form the light-sensitive component of the film. X-ray energy being absorbed by the film results in the release of electrons. These electrons are trapped in the crystal lattice, and attract and neutralize the mobile silver ions (Ag^+) in the lattice, resulting in deposits of metallic silver in the emulsion. These deposits across the film form the latent image. When the film is developed in chemical solutions, the silver is fixed. The degree of blackening in a region of the film is relative to the amount of metallic silver deposited in that region, which in turn is dependent on the number of X-ray photons absorbed in that region. Hence, an image is captured on film through this photographic process. Intensifying screens are also routinely used so that the radiation exposure required to generate an image may be reduced (Hendee and Ritenour, 1992).

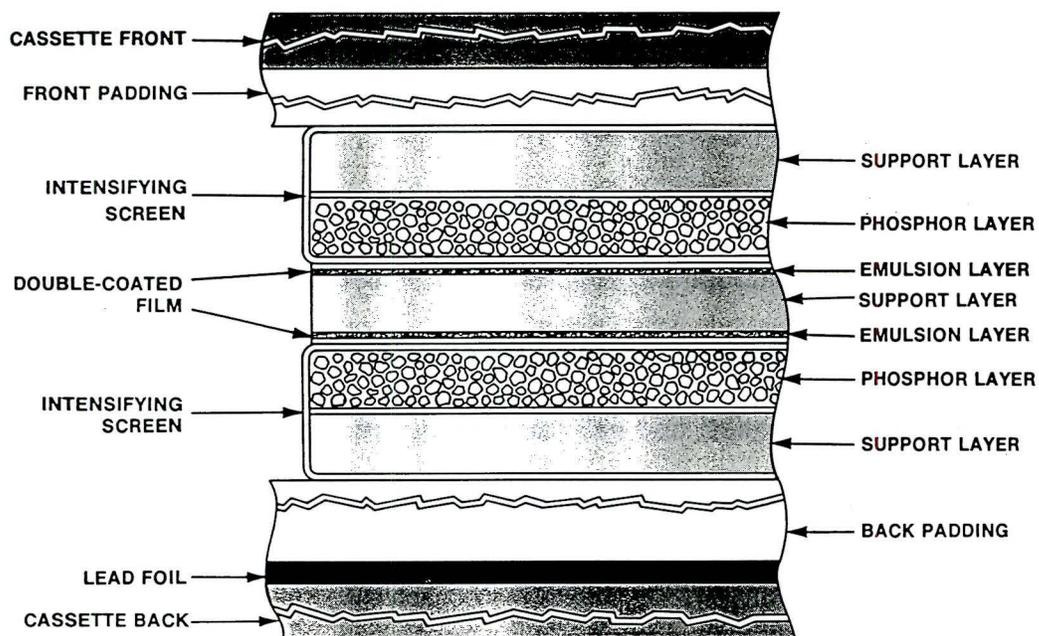


Fig. 3.6 Medical radiographic film typically used in a film-screen combination.

3.3.2 DIGITAL RADIOGRAPHY

In digital radiography, either an analog image is frame-grabbed (as in fluoroscopy), or the latent image is converted to digital form by laser-scanning (as in computed radiography). The advantage of digital radiography is the separation of the detector from the display (Brody, 1984). Conventional film radiography uses the same sheet of film for detecting the image and displaying it, and is thus constrained to the use of film having properties that are a compromise for both X-ray photon detection and image display. In contrast, when detector and display are separated, each may be optimized to its own process, and the detected image can be further manipulated before the final display. In this section, we focus on the Computed Radiography technology used to acquire the images for our work. But first, the concept of the digital image will be reviewed and its associated terms defined.

3.3.2.1 REVIEW OF DIGITAL IMAGES

A digital image is composed of small rectangular blocks called pixels (from '*picture elements*'). Each pixel carries a number which, when mapped to a colour table, identifies a colour or a specific shade of a colour. The range of colours used in a particular image is called the colourmap of the image. A typical colourmap, used especially for X-ray images, is the k-level grey-scale, which represents degrees of blackness and whiteness by k gradually varying shades of grey. A digital image is represented by a rectangular matrix,

whose dimensions, $m \times n$, indicate the image size in pixels. M is the number of pixels in one column of the image (that is, the vertical dimension), and n is the number of pixels in one row (that is, the horizontal dimension). Each matrix element is then stored with the colour number (pixel value) of the pixel at that matrix location. Digital images can then be easily manipulated, transformed, or altered, by changing the pixel values stored in the matrix. This manipulation of pixel values by mathematical transformations is the art of digital image processing.

In general, *pixel value* refers to a number carried by a pixel, and does not indicate to what colourmap that number belongs. The term *grey value*, or *grey level*, signifies that the colour number is in reference to a grey-scale colourmap, where the range of grey levels available, k , is usually known a priori. In this thesis, pixel value, grey value, and grey level will be used interchangeably. We will discuss our images only with reference to the 1024 grey-scale, which is specific to our digital X-ray image representation.

3.3.2.2 COMPUTED RADIOGRAPHY

Computed radiography (CR) is a relatively new technology employing a re-usable photostimulable storage phosphor (PSP) imaging plate (IP) and a laser reader. The imaging plate, which replaces the film-screen cassette of conventional radiography, consists of a photostimulable phosphor, BaFX (barium fluoro-halide), applied onto a polyester film base. The inorganic crystals can store the energy dissipated in the film by the X-rays. When the

imaging plate is stimulated with a He-Ne laser beam, the crystals emit light proportional to the stored X-ray exposure. The optical signals are converted into electrical signals and digitized with an analog-to-digital converter. At this point, this raw digital image may be captured as a data file by an external computer linked to the CR system. Within the CR system (Fig. 3.7), the raw image is then passed through image processing functions specific to the particular X-ray examination. The modified digital information is then converted back to analog signals, and then to optical signals, which modulate a laser beam which exposes the film (Fig. 3.8). The after-image remaining on the IP is then erased by light energy, and the IP becomes re-usable (Fig. 3.9).

First introduced by Fuji Medical Systems in 1983, the present widespread use of CR is a result of several main benefits. Of significance are the primary advantages of the PSP imaging plate :

- i) the imaging plate has a 10000:1 dynamic range, compared to 100:1 for conventional film-screen detectors (Chotas and Ravin, 1992), so that a wider range of exposures can be discriminated
- ii) the imaging plate has a linear response, that is, the digital pixel values are linearly related to the X-ray exposure intensity, to provide unbiased storage of image information
- iii) the high sensitivity of the PSPs to X-ray energy can result in reduced radiation dose to the patient.

In addition, the CR system provides built-in image processing functions for contrast and density optimization tailored to various X-ray examinations, including chest, abdomen, bone, spine, contrast medium imaging, and tomography, for the optimal display of the image on film.

The images acquired for our work were the raw digital images captured at point B in Figure 3.8, using the Fuji AC-1 Computed Radiography system (Fig. 3.7).

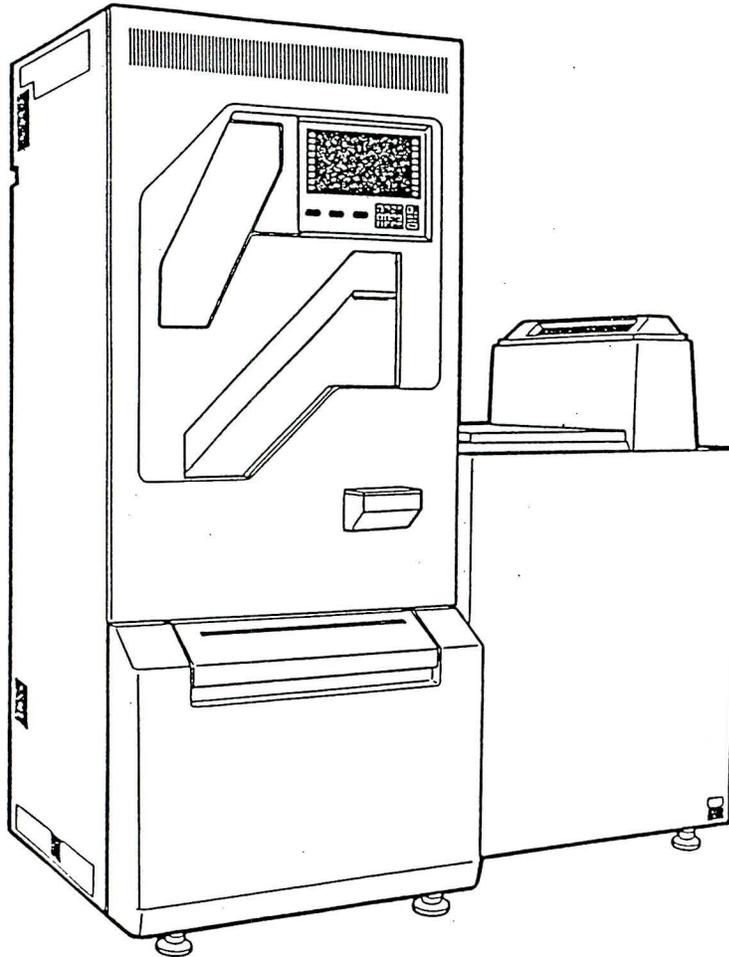


Fig. 3.7 Fuji Computed Radiography System FCR AC-1.

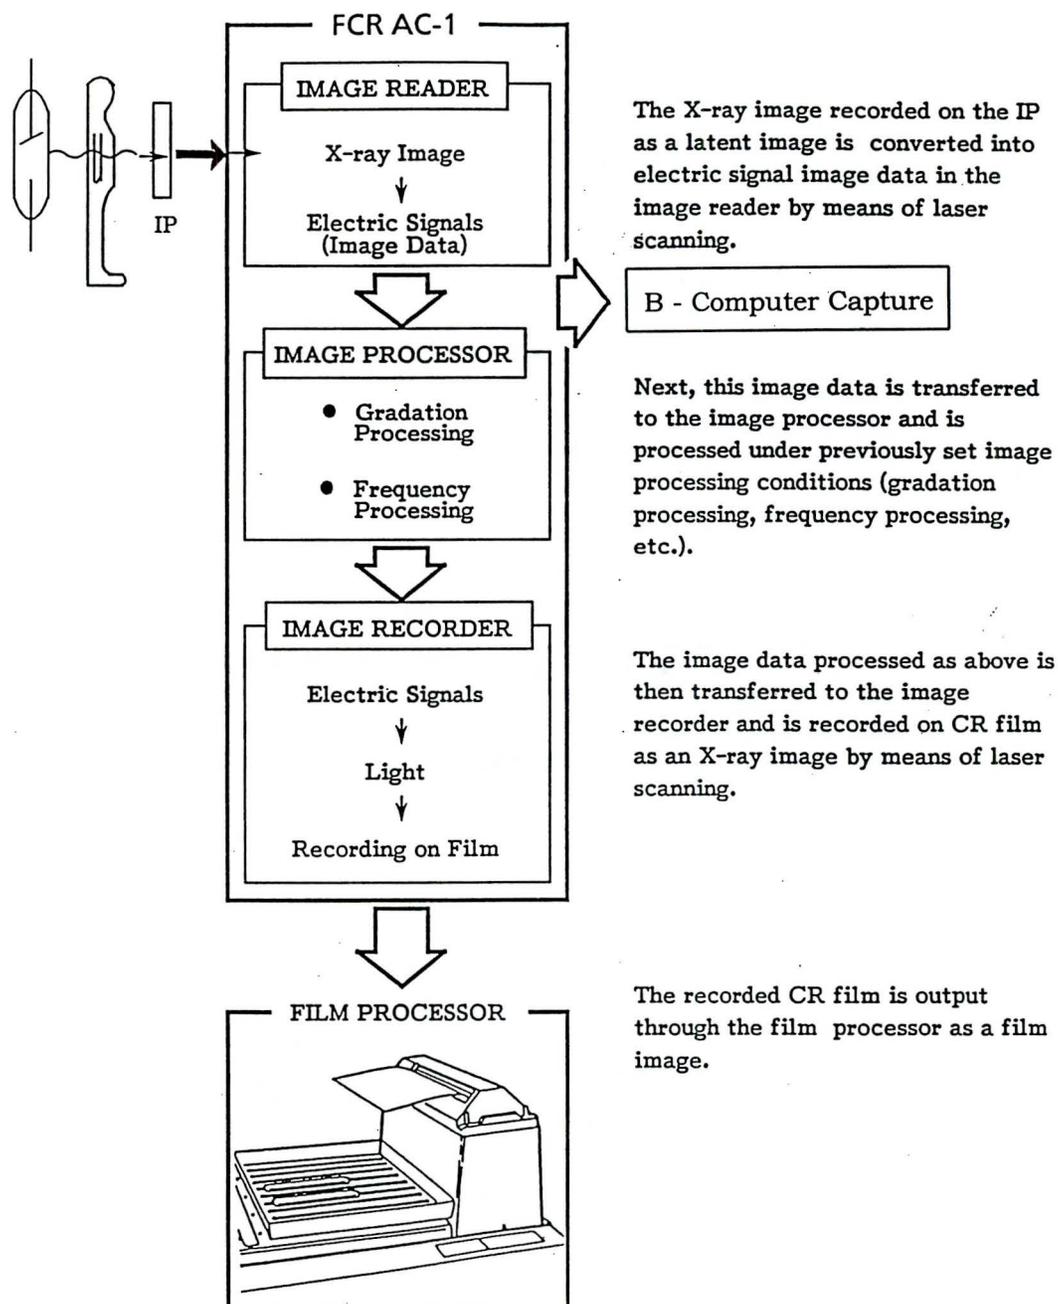


Fig. 3.8 FCR AC-1 image acquisition and processing flow.

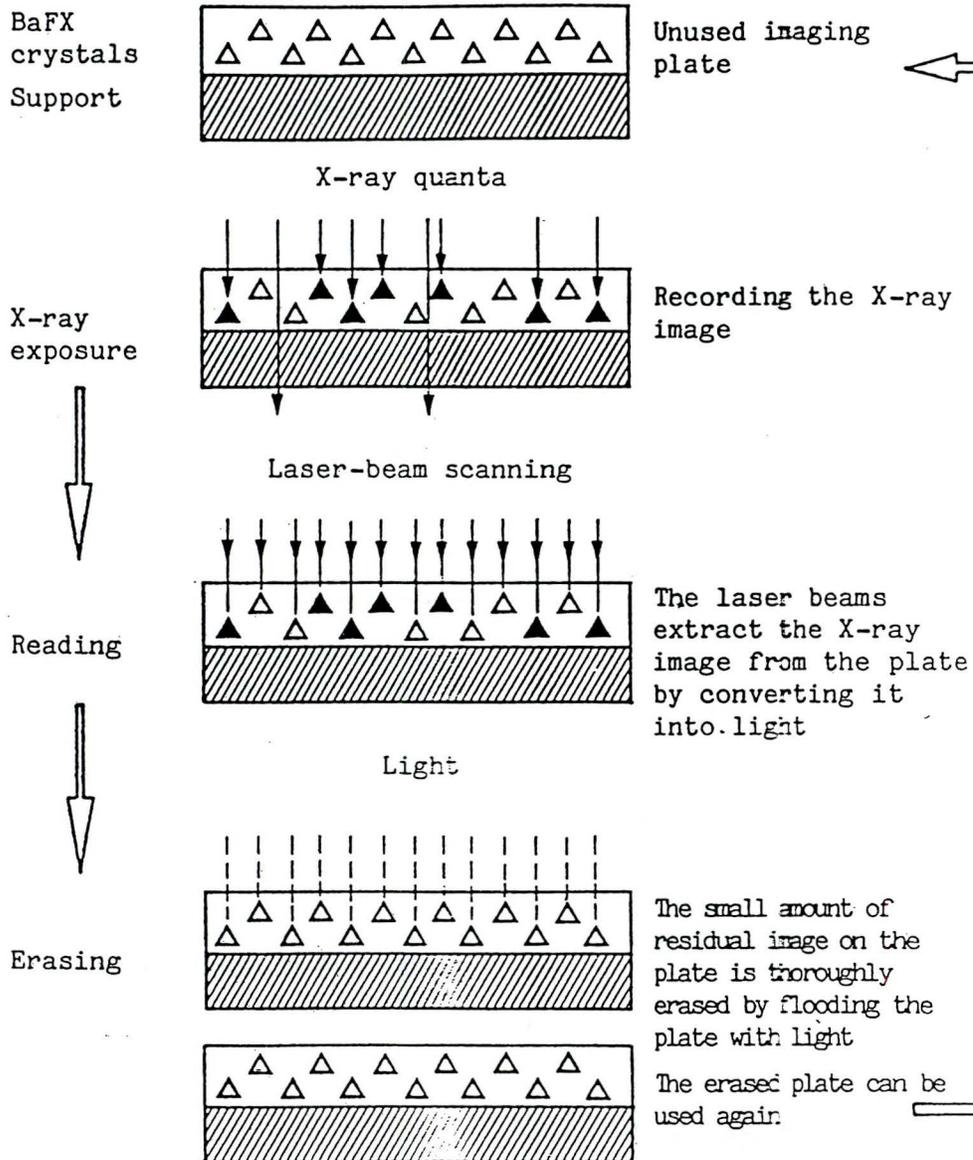


Fig. 3.9 Principles involved in the image recording, reading and erasing on the photostimulable storage phosphor imaging plate developed by Fuji.

3.4 AN EXPERIMENTAL INVESTIGATION :

OPTIMAL LUMBAR SPINE X-RAY IMAGE ACQUISITION PARAMETERS

3.4.1 INTRODUCTION

The complexity and ultimate success of any automated image analysis system depends on the quality of the images it is expected to process. For our goal of automating lumbar vertebrae detection and measurement, it is desirable to acquire the best possible image in terms of image contrast. Image contrast, in this context, is defined to be the relative utilization of the set of all possible grey levels that may be represented by the image format. If the entire range of pixel values allowable in that format is DR (Dynamic Range), and the range of pixel values present in a particular image is R, then an image with low contrast has a small R/DR ratio, much smaller than 1. That is, the image values comprise only a small subset of the set of all possible pixel values. On the other hand, an image of high contrast has a very wide range, so that the image contains pixel values filling most of the dynamic range, ie. $R/DR \rightarrow 1$. As we have discussed earlier, the adjustable acquisition parameters which affect image quality are X-ray tube potential (kVp), and the quantity of photons generated (mAs). An important effect of increasing kVp is the reduction in image contrast between different tissue types. The amount of photons produced by the X-ray tube, in units of mAs, determines the intensity of the X-ray beam. Image contrast should not be affected by mAs; however the graininess, or relative signal-to-noise content in the image will affect image quality and the success of digital image processing algorithms.

An experimental investigation was carried out to determine the optimum kVp and mAs ranges for maximum image contrast at the bone-tissue edge in X-ray images of the lateral lumbar spine. The four parts of the experimental investigation were :

- 1) to compare the linear attenuation coefficients of aluminum and bone
- 2) to determine the optimum kVp range for maximum bone-tissue contrast
- 3) to determine the optimum mAs range for maximum bone-tissue contrast
- 4) to evaluate the effects of mAs on the uniformity of image intensity.

3.4.2 MATERIALS

To simulate the body experimentally, a circular plastic phantom filled with water (which simulates tissue) was used. The phantom was 16 cm thick with a 30 cm diameter. An aluminum step wedge (an aluminum block with 11 steps ranging from 5 to 34 mm in height) was immersed in the water in the phantom shell to simulate the spinal column. The general X-ray machine in Room 8 of the Department of Radiology at MUMC was the X-ray source. The CR image plates were processed by the Fuji AC-1 Computed Radiography unit (Fig. 3.7), and the digital information stored as an image file by the Sun SPARCstation IPX attached to the CR unit. All analyses were done using a SPARCstation IPX running MUMC Display and Matlab™.

3.4.3 METHODS AND RESULTS

Eighteen images of the aluminum step wedge immersed in water were acquired at different combinations of mAs and kVp (Table III-2). The kVp settings were limited to the discrete values available on the control dials of the Siemens Tridoros-Optimatic 800 X-ray machine (Table III-3).

Image analysis was performed in MUMC Display, using the built-in histogram and image smoothing routines. Further calculations and graphical output were done using MATLAB. The aluminum step wedge had 11 steps, ranging in height from 5 to 34 mm (Fig. 3.10). In our analyses, we focused on steps 1 (3.4 cm), 5 (2.3 cm) and 8 (1.4 cm), as they were generally easier to identify in the image, and they provided a range of thicknesses for comparison of the results. An *edge* was indicated by the boundary between aluminum and water. This edge simulates a boundary between vertebral bone and the soft tissue surrounding it.

The method of measuring edge contrast

The histogram of pixel values along a straight line lying perpendicularly across an edge was plotted using the Plot A Slice function in MUMC Display. The *edge gradient* was measured as the absolute difference between the pixel values of two points, one on each side of the edge. Generally, these points were within two to four pixels on either side of the

edge. These points were chosen subjectively using a high pixel value point on the aluminum side of the edge, and usually the lowest pixel value point closest to the edge on the water side. Before the edge gradient was measured, the histogram was smoothed using either a 5-point or a 7-point average filter, depending on whether the histogram was judged very noisy or already fairly smooth. This smoothing was done to yield more precise and reproducible choices of the points to be measured on either side of the edge. In most cases, at least two measurements were done for each edge; sometimes, up to five measurements were done where reproducibility was uncertain.

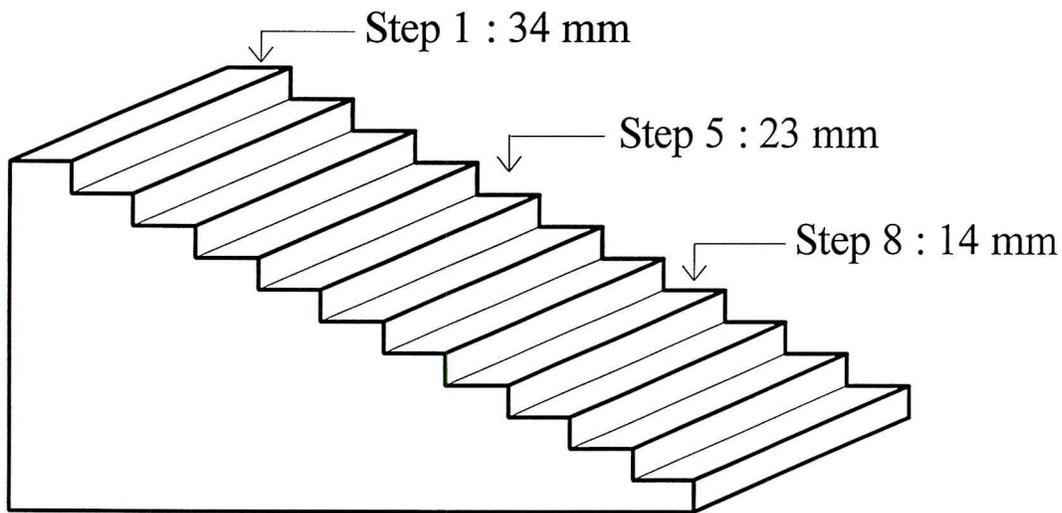


Fig. 3.10 Aluminum step wedge

Table III-2. Parameters of the acquired images

100 mAs	63	70	73	77	81	85	96	117 kVp				
77 kVp	10	20	25	32	40	50	64	80	100	120	160	mAs

Table III-3. General X-Ray Room No. 8 parameter settings

kVp	mAs																	
60	8	10	12	16	20	25	32	40	50	64	80	100	120	160	200	250	320	400
63	8	10	12	16	20	25	32	40	50	64	80	100	120	160	200	250	320	400
66	8	10	12	16	20	25	32	40	50	64	80	100	120	160	200	250	320	
70	8	10	12	16	20	25	32	40	50	64	80	100	120	160	200	250	320	
73	8	10	12	16	20	25	32	40	50	64	80	100	120	160	200	250	320	
77	8	10	12	16	20	25	32	40	50	64	80	100	120	160	200	250	320	
81	8	10	12	16	20	25	32	40	50	64	80	100	120	160	200	250		
85	8	10	12	16	20	25	32	40	50	64	80	100	120	160	200	250		
90	8	10	12	16	20	25	32	40	50	64	80	100	120	160	200	250		
96	8	10	12	16	20	25	32	40	50	64	80	100	120	160	200	250		
102	8	10	12	16	20	25	32	40	50	64	80	100	120	160	200			
109	8	10	12	16	20	25	32	40	50	64	80	100	120	160	200			
117	8	10	12	16	20	25	32	40	50	64	80	100	120	160	200			
125	8	10	12	16	20	25	32	40	50	64	80	100	120	160				

Comparing the linear attenuation coefficients of aluminum and bone

It was necessary to first verify that aluminum is a good substitute for bone in X-ray experiments. Since, in clinical practice, the X-ray source tube energy used for spine examinations is no more than 150 kVp, the linear attenuation coefficients of bone and aluminum were compared in the 10-150 keV photon energy range (Fig. 3.11). The numerical values are shown in Table III-4. The linear attenuation coefficients were

calculated by multiplying the mass attenuation coefficients in this energy range (obtained from Hubbell, 1982) by the physical density of the material. The density of bone was taken to be from 1.65-1.85 g/cm² (Hendee and Ritenour, 1992), and the mass attenuation coefficients used were the ICRP standards for cortical bone (Hubbell, 1982). Although we are studying vertebrae, which is composed mostly of trabecular bone, we measure image contrast at the vertebral body boundary, which is a thin layer of cortical bone. Figure 3.11 shows the consistent similarity in behaviour of the linear attenuation coefficients of both materials, in the 10-150 keV energy range. Table III-4 details the consistent proportionality of 1.32±6% between the linear attenuation coefficient for aluminum and that for bone. These results indicate that in the 10-150 keV range, aluminum is consistently representative of cortical bone in its X-ray attenuation properties.

Determination of Optimal kVp Range

To begin searching for the optimal kVp range, images were acquired at 63, 70, 81, 96 and 117 kVp at a fixed mAs of 100. 100 mAs was chosen to ensure enough X-ray photons for an image to be produced. The edge between aluminum (bone) and water (tissue) at Step 5 of the aluminum wedge was analysed in these 5 images. Between 70 and 81 kVp, a sharp drop in the contrast at the edge was apparent, so another 2 images were acquired using the remaining available kVp settings within this range, namely 73 and 77 kVp. As well, an additional image at 85 kVp was acquired for verifying the low edge contrast seen at 81 kVp. It was apparent that a distinct peak edge gradient existed, at 77

kVp, for this step. The edges at two other steps, Step 1 (3.45 cm) and Step 8 (1.4 cm), were investigated as well, by measuring the edge gradients on these same eight images. It is clear, in Figure 3.12, that high edge gradients are produced in the 63-85 kVp range, but the contrast drops off after 85 kVp. The three curves show this trend for the three steps measured. However, there is also significant variation in the edge gradient within the 63-85 kVp range, with the distinct peak evident for Step 5 no longer evident for Steps 1 and 8. Repeated measurements in this range show that the variation does not result in large part from measurement error. Thus, other factors must be involved in the non-linear behaviour of the edge contrast at the various step thicknesses. A misaligned X-ray beam centre might cause such non-linearity in the image. However, since the high edge gradient values are concentrated in this range, we can consider the 70-85 kVp range as the optimal for producing the highest edge contrast. Once this optimal range was localized, 77 kVp was then chosen as the fixed point from which to search for the optimal mAs range.

Determination of optimal mAs range

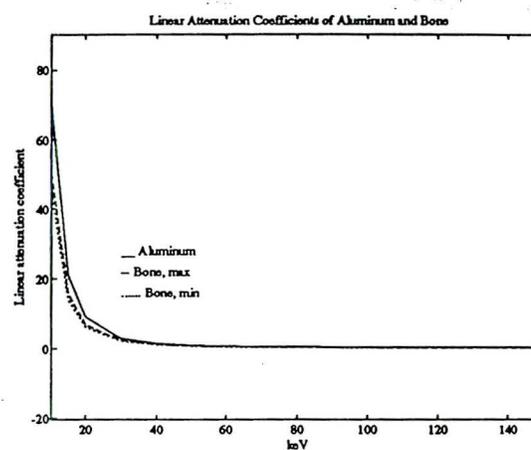
The edge gradient of Step 5, in each image acquired at 77 kVp, was measured and related to the mAs which produced each image (Table III-6). Below 25 mAs, no visible image was produced, indicating that there were not enough photons penetrating the phantom to generate an image. From 25 to 160 mAs, except for the surprisingly high value at 25 mAs, the edge gradient values fluctuated within 5.1% of their mean value, with no noticeable trend with respect to mAs (Fig. 3.13).

Effect of mAs on image statistics

It was important to examine the effect of mAs on the image statistics as these will influence the performance of digital image processing techniques. An area of water at the top of the phantom was selected as a uniform region of interest for measuring the image statistics using the ROI Statistics routine of MUMC Display, which computes the mean and standard deviation of the pixel values in a defined region. Figure 3.14 shows the decrease in standard deviation of the pixel values as mAs increases. Curve fitting of these data from the nine images of different mAs indicates an inflection point between the image statistics of the 50 mAs and the 64 mAs images. Below 50 mAs, the spread of pixel values decreases quickly, indicating a significant gain in pixel value uniformity. Above about 57 mAs, no significant gain in the uniformity of image intensity is observed.

Table III-4. Linear attenuation coefficients

keV	Linear attenuation coefficients			
	Bone		Al	Al/Bone*
	min	max*		
10	44.39	49.77	69.56	1.40
15	14.06	15.77	21.11	1.34
20	6.24	7.00	9.14	1.31
30	2.09	2.34	3.00	1.28
40	1.06	1.18	1.52	1.28
50	0.68	0.76	0.98	1.29
60	0.51	0.57	0.74	1.30
80	0.37	0.41	0.54	1.32
100	0.31	0.34	0.46	1.33
150	0.25	0.28	0.37	1.35

Fig. 3-11. Linear attenuation coefficients of aluminum and bone**Table III-5.** Edge gradients at 100 mAs

kVp	Edge Gradients (grey levels)		
	Step 1	Step 5	Step 8
63	165.67	146.00	92.00
70	154.00	151.50	99.00
73	162.33	158.25	98.00
77	160.00	162.50	85.00
81	157.00	140.50	86.50
85	166.00	132.00	104.00
96	152.00	127.75	81.00
117	141.33	113.50	67.50

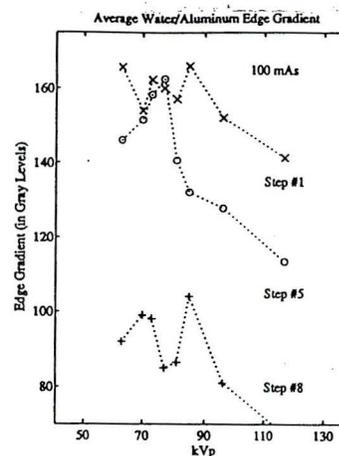
Fig. 3.12. Edge gradients vs. kVp

Table III-6 Edge gradient vs. mAs

mAs	Step 5	% from mean (mean = 150.4)
25	164	9.0
32	143	4.9
40	143	4.9
50	148	1.6
64	147	2.3
80	158	5.1
100	143	4.9
120	145	3.6
160	148	1.6

Fig. 3.13 Edge gradients vs mAs

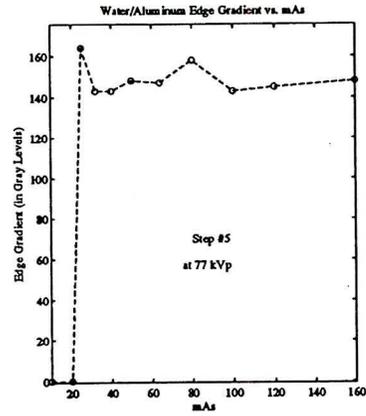
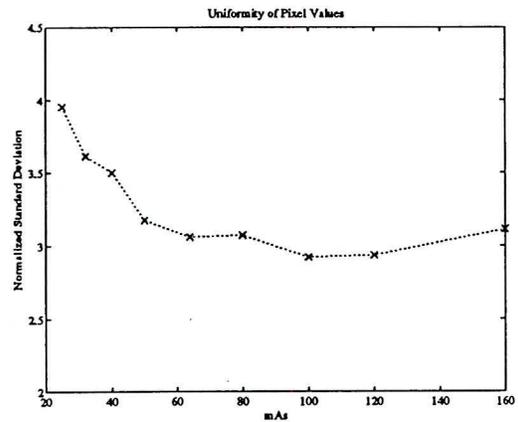


Table III-7 Standard deviations of pixel values

mAs	Std Dev	Mean	% (std/mean)
25	27	683	3.95
32	25	692	3.61
40	24	686	3.50
50	22	693	3.17
64	21	686	3.06
80	21	683	3.07
100	20	684	2.92
120	20	681	2.94
160	21	675	3.11

Fig. 3.14. Uniformity of image intensity



3.4.4 DISCUSSION

The curves of Figure 3.12 interpolate linearly between data points, and may not represent the true behaviour of the edge gradient in the kVp range above 85 kVp, where images were acquired at coarse intervals of kVp. In contrast, the fine intervals examined between 70-85 kVp caused the trough to be seen (especially pronounced for Step 8). It is not known what this trough may have resulted from. The limitations on the data points acquired were primarily the limitation of the X-ray machine in providing only discrete parameter settings. The cost of film motivated the coarse searching process preceding the finer search for the optimal kVp.

In controlling the intensity of the X-ray beam, mAs does not affect contrast of tissue types, as the intensities through the body vary linearly with the intensity of the beam (Kodak, 1980). However, given that enough photons were generated to produce an image, it was important to examine the effect of mAs on the image statistics, which in turn affect the performance of digital image processing techniques. The smoothing performed on the images, before edge gradient measurement, helped to reduce the noise in the pixel value histograms at each edge, and thus reduced measurement error. Nonetheless the noise in the raw image is still largely dependent on the X-ray beam intensity, hence mAs. Large standard deviations in the pixel values of a uniform region of an image indicate low signal-to-noise content, which may cause image processing algorithms to falsely identify features in noisy images.

3.4.5 CONCLUSIONS OF THE EXPERIMENT

For maximum contrast between bone and tissue in a lateral lumbar spine CR image, the optimal kVp range is roughly 70-85 kVp. Although mAs does not affect tissue contrast in the image, given that at least 25 mAs are needed to produce a visible image, a minimum of about 60 mAs is recommended for improved image signal-to-noise characteristics. It was determined that, in the 10-150 keV photon energy range which is associated with an applied X-ray tube potential range of 60-150 kVp, aluminum has approximately 1.3 times the linear attenuation coefficient of bone, with small variations in this factor depending on the range of photon energy in the X-ray beam. Thus it was verified that aluminum is a good substitute for bone in X-ray attenuation investigations.

It is essential, in determining optimal parameter settings for image processing purposes, that the radiation dose to the patient be maintained as low as possible. Thus, the optimal kVp value would need to be as high as possible to reduce absorbed patient dosage and film graininess, yet low enough to yield sufficient image contrast. Optimal mAs values would need to be as low as possible to keep dose to the patient low, yet high enough to produce good image statistics. In considering both clinical diagnostic and digital analysis requirements, it is suggested that 85kVp and 55-60 mAs would be desirable and feasible parameter settings for the clinical acquisition of lateral lumbar spine CR images. In clinic practice, however, these recommended values would be subject to wide variation, depending on factors of dosage and image quality required. Presently, values of 96 kVp and 50 mAs are used for the average-sized patient; however, these values may range from

75-100 kVp, and from 50-100 mAs, depending on the body size of the patient and the condition of the spinal bone being imaged (Kereliuk, 1994). Therefore the recommended values suggested in these conclusions are indicative only of optimal and feasible CR imaging parameter values for the development of computer-aided diagnosis systems for vertebral deformities and other spinal disorders.

3.5 CHAPTER SUMMARY

The ability of X-rays to penetrate different tissues to different degrees, and to cause certain materials to fluoresce and expose photographic film, are the primary advantages of X-rays for medical imaging. X-ray imaging is especially suited for viewing skeletal structures, since bone is the highest X-ray-attenuating medium occurring naturally in the human body. In this chapter, we studied the ways in which the parameters of the X-ray tube can be varied to yield different image intensity levels and contrasts between bone and soft tissue in a radiograph. The main benefit of improving image quality at the time of image acquisition is in the computation complexity that is avoided by having software dedicated to performing its specific task (detection, measurement, etc.) rather than to the overhead of image quality improvement. It is more practical to obtain images of uniformly similar image quality at the time of acquisition than to create digital processing methods robust enough to handle images with a wide variation of image quality. At MUMC, Computed Radiography technology is in use alongside traditional film-screen radiography. It is with Computed Radiography that we are able to pursue digital image processing techniques for automating the measurement of vertebral deformities from lateral lumbar CR spine images.

CHAPTER 3

REFERENCES

Armstrong P, Wastie ML, Brenbridge A, Norman AG, 1987. Diagnostic imaging. 2nd ed. Oxford: Blackwell Scientific Publications.

Brody WR, 1984. Digital Radiography. New York: Raven Press.

Bryan GJ, 1970. Diagnostic radiography: a concise practical manual. Edinburgh: E&S Livingstone.

Chotas HG, Ravin CE, 1992. Digital radiography with photostimulable storage phosphors: control of detector latitude in chest imaging. *Investig Radiol* 27:822-828.

Eastman Kodak Co., 1980. The fundamentals of radiography, 12th ed., Rochester: Eastman Kodak Co.

Economic Commission for Europe, Geneva, 1987. Digital imaging in health care. New York: United Nations Publication ECE/ENG.AUT/25.

Fuji Photo Film Co., Ltd., 1989. Fuji computed radiography system FCR AC-1 operation manual. Tokyo: Fuji Photo Film Co., Ltd.

Hendee WR, Ritenou R, 1992. Medical Imaging Physics, 3rd ed., St. Louis: Mosby - Year Book, Inc.

Hubbell JH., 1982. Photon mass attenuation and energy-absorption coefficients from 1 keV to 20 MeV. *Int J Appl Radiat Isot* 33:1269-1290.

Kereliuk, M, RTR, Senior Technologist, Department of Radiology, McMaster University Medical Centre, Hamilton, Ontario, Canada.

Long BW, 1990. Image enhancement using computed radiography. *Radiography* 61:276-280.

Schulz RJ, 1977. Diagnostic X-ray physics. New York: GAF Corp.

Sprawls P, 1977. The physical principles of diagnostic radiology. Rockville: Aspen Publishers, Inc.

Tateno Y, Iinuma T, Takano M (Eds.), 1987. Computed Radiography. Tokyo: Springer-Verlag.

Tucker DM, Souto M, Barnes GT, 1993. Scatter in computed radiography. *Radiology* 188:271-274.

CHAPTER FOUR

DIGITAL IMAGE PROCESSING : METHODS

4.1 INTRODUCTION

Digital image processing comprises methods for manipulating the pixel values of digital images in order to enhance selected image features, or to extract information from the image which may be useful for comparing images for some purpose. Although a body of algorithms is available for performing general tasks, such as thresholding or noise reduction through smoothing, the development of digital image processing techniques and algorithms is often an application-specific process. Creating an algorithm to solve a specific problem requires combining known techniques and creating new ones, in a process which requires a great deal of creativity and insight into the nature of the image. This is particularly true when dealing with real-life images, in which complex objects or scenes are represented with many shades of grey or colour.

An approach to developing an image processing solution for the automatic detection and identification of a certain object type involves two components : image pre-processing, and search engine development. The object of image pre-processing is to enhance desired image features as much as possible. Pre-processing usually involves noise suppression and contrast enhancement. A search engine is the measure used by the automatic detector. It is often some sort of edge detector which tracks the edges of the desired object.

In this chapter, past work on the problem of automating the detection and measurement of vertebral dimensions is discussed. Then, the specifications of the image set are given, and the software development environment and specifications explained. The three methods that we pursued for automating vertebral body detection and measurement are then described in detail. One is an image intensity based approach, the second is edge gradient and cross-correlation based, and the third, a fast contouring technique based on the minimization of an energy functional.

4.2 PREVIOUS WORK

For our particular problem of automatically detecting and measuring vertebral bodies, no previous published work is available, unlike the plethora of papers resulting from research in automating the detection of various lung disease processes from chest radiographs, and the development of other automated techniques specific to heart disease

and mammography (Doi et al., 1992). Only a private report by J.S. Sloka (1993) documents attempts at automatic vertebral body detection and measurement. Sloka describes his use of masks that simulate vertebral body corners, and the convolution of these masks throughout the image, in an attempt to detect the vertebral body corners. A convolution mask will produce a maximum value where the data set most corresponds to the mask model. Of the various masks that were tried, none produced useful results. It was determined that many different masks for each of the four vertebral corners would be required to deal with the variability in the rotation and location of the spinal vertebrae.

In the same work by Sloka, a global segmentation approach was also attempted. It was first observed that lumbar spine images possess an illumination gradient. This is a direct result of variations in the body thickness and bone density of the patient from the thoracic to the lumbar regions, and from the lumbar to the sacral regions of the body. The gradient is most apparent in the sacral region, which has a very high intensity as a result of the high X-ray attenuation of the larger sacral bones (hip, pelvis). This brightness in the sacral region often buries the fifth lumbar vertebra, and makes it especially difficult to automatically extract this vertebra from the sacral skeleton. As a result, global thresholds to separate the bone from the soft tissue in the image were not successful, even after the illumination gradient was removed by subtracting a large (49 x 49) mean-filtered image from the original.

and mammography (Doi et al., 1992). Only a private report by J.S. Sloka (1993) documents attempts at automatic vertebral body detection and measurement. Sloka describes his use of masks that simulate vertebral body corners, and the convolution of these masks throughout the image, in an attempt to detect the vertebral body corners. A convolution mask will produce a maximum value where the data set most corresponds to the mask model. Of the various masks that were tried, none produced useful results. It was determined that many different masks for each of the four vertebral corners would be required to deal with the variability in the rotation and location of the spinal vertebrae.

In the same work by Sloka, a global segmentation approach was also attempted. It was first observed that lumbar spine images possess an illumination gradient. This is a direct result of variations in the body thickness and bone density of the patient from the thoracic to the lumbar regions, and from the lumbar to the sacral regions of the body. The gradient is most apparent in the sacral region, which has a very high intensity as a result of the high X-ray attenuation of the larger sacral bones (hip, pelvis). This brightness in the sacral region often buries the fifth lumbar vertebra, and makes it especially difficult to automatically extract this vertebra from the sacral skeleton. As a result, global thresholds to separate the bone from the soft tissue in the image were not successful, even after the illumination gradient was removed by subtracting a large (49 x 49) mean-filtered image from the original.

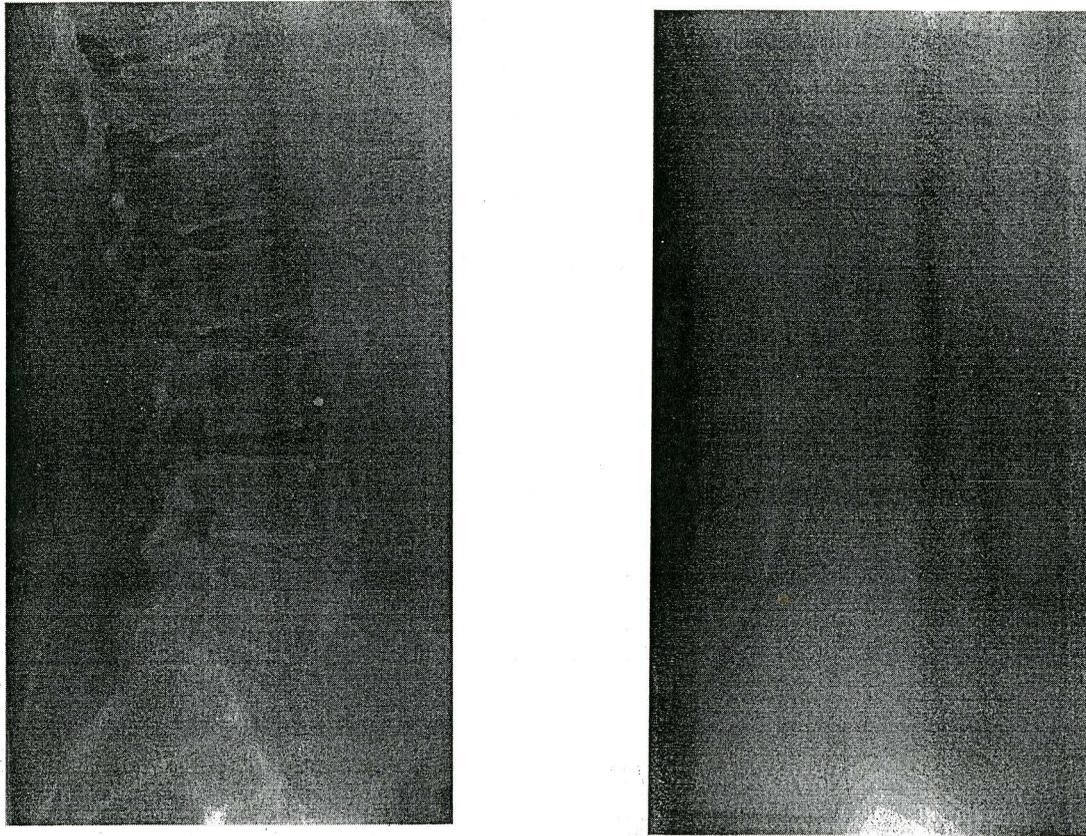


Fig. 4.1. (a) Original lateral lumbar spine image. (b) Illumination gradient in image.

4.3 IMAGE SET

Seven lateral lumbar spine images from six different patients were acquired digitally onto a Sun SPARCstation from a Fuji AC-1 Computed Radiography system (Fig. 3.8). The digital images, referred to here as FUJI images, are of size 2048 x 2048 pixels, with 2 bytes per pixel, thus requiring over 8.3 MBytes of storage space for each image. This image size is the largest of the medical images from all of the imaging modalities used at McMaster University Medical Centre (MUMC), which include 512 x 512 CT images, 256 x 256 MRI images, and 128 x 128 PET images. The FUJI spine images have a maximum range of 1024 grey levels (grey levels 0 to 1023).

Although none of the images were clinically requisitioned for assessing osteoporosis, they provided a starting base for algorithm development. Although the images would have been acquired with typical X-ray parameter ranges of 75-100 kVp and 50-100 mAs, they represented a wide range of image quality. Two of the images were of poor quality, being so underexposed that the intensity level of the tissues was near that of the bone; four of the images were of average quality, with reasonable contrast between bone and tissue in most areas; and one excellent image, image RH, displayed the anatomical structures with excellent contrast gradations. This image was used for initial testing of the algorithms.

4.4 SOFTWARE DEVELOPMENT ENVIRONMENT

The algorithms were developed in C with the Unix cc compiler, and tested on a Sparc IPX workstation running SunOS version 4.1.3 and the Sunview graphical user interface. An in-house image viewer, MUMC Display¹, provided the customized graphical user interface in which the algorithms were integrated and tested. The program modules developed in this work are listed and described in Appendix A. A beneficial feature of MUMC Display is its integrated database of all of the images available on the medical imaging network at MUMC. In this environment, the FUJI images are displayed on a 1152 x 900 pixel screen with a range of 125 grey levels (ie. 7 bits, with the remaining 3 levels reserved for Sunview screen colours).

4.5 METHODS

Three very different methods were studied : i) a grass-roots image intensity based approach; ii) gradient search and cross-correlation detection; and iii) a dynamic contouring algorithm called Snakes. Each method will be discussed in detail in the following sections. The subroutines developed for all three methods are listed and described in Appendix A.

¹ developed by D.B. Kenyon, Department of Nuclear Medicine, Chedoke-McMaster Hospitals

4.5.1 IMAGE INTENSITY BASED APPROACH

Previous observations of the nature of the lateral spine images revealed that global thresholding as a segmentation approach is not successful because of the illumination gradient apparent in lumbar spine images (see section 4.2). As a result, we developed a local approach to image pre-processing. Each vertebra is processed separately, in a 400 x 400 rectangular image clip which is determined by creating a 400 x 400 rectangle around a user-chosen vertebral centre. This image clip is processed with a series of operations designed to enhance image quality and to maximize the absolute value of the image intensity at the vertebral body boundaries, while minimizing the image intensity everywhere else. The resulting image clip is ideally a bi-level image, containing only an outline of the vertebral body in white (the maximum grey level). The goal is to use image intensity as a search engine, based on the observation that the boundaries of the vertebral body are usually brighter than both the surrounding tissue and the interior (centrum) of the vertebral body. As well, the superior and inferior boundaries of the vertebral body are normally brighter than the lateral edges. These observations arise from the structure and composition of the vertebrae. We recall that each vertebra has two endplates, a superior and an inferior (see section 2.2). The endplates are regions composed of dense, cortical bone, which yields a higher image intensity than the sparse, trabecular bone comprising the centrum of the vertebral body. A search routine would track the maximum image intensity pixels in the pre-processed image clip to extract the coordinates of the vertebral body outline, and then determine the salient features of corners and midpoints by simple calculations.

Algorithm Pseudocode

- 1 : Image clipping
- 2 : Median filtering
- 3 : Linear histogram stretching
- 4 : Piecewise linear contrast enhancement
- 5 : Edge filtering
- 6 : Linear stretching
- 7 : Thresholding
- 8 : Mode filtering
- 9 : Search engine

Each of these steps in the algorithm is described in the following sections. The corresponding images in Figure 4.2 will be referenced.

Image clipping

The user selects a point in the centre of the desired vertebra. This point does not have to be precisely the centre, but as good an approximation to it as possible. A 400 x 400 rectangular image clip around this point is copied for processing. This 400 x 400 size was selected because it fully covers a single vertebra (even if the vertebra is oriented at an angle) in all of the spines in our image set (Fig. 4.2(b)).

Median Filtering

Median filtering attempts to remove noise in the image while minimizing any resulting loss of high frequency edge signals. At each pixel in the image clip, the median grey level of a 5 x 5 rectangular neighbourhood of pixels around that pixel is calculated and placed in another image matrix of the same size as the image clip. In this way, a new image clip is created, containing the median grey levels of 5 x 5 neighbourhoods successively passed through the original image clip (Fig. 4.2 (b)). It is usually necessary to sort a list of pixel values in order to find the median value of the list. To eliminate the computationally costly sorting step, we first calculate the average grey level in the neighbourhood, and then choose as the median the pixel value which is closest to the average (Yu, 1994). In some cases, this value does equal the average value, but the benefit of median filtering is that we are always using a pixel value that is already present in that neighbourhood.

Linear histogram stretching

After noisy pixels have been 'smoothed' by median filtering, we then maximize the range of pixel values in the image to fit the full range of pixel values available (in this case, 1024 levels) (Fig. 4.2 (c)). Even though the FUJI images are initially already linearly stretched by the FUJI processor, it is always desirable to repeat this step after performing operations which alter the pixel values in the image. First, we find the

maximum and minimum grey values in the image clip, MAX and MIN, respectively. The image pixel value range is then (MAX-MIN). Then, we find the ratio of the full available range to the image range : ie., $1024/(MAX-MIN)$. A shift of MIN levels is performed on each pixel value before being rescaled linearly by the factor $1024/(MAX-MIN)$. The linear transformation is :

$$IMAGE_{out} [i][j] = (IMAGE_{in} [i][j] - MIN) * \frac{1024}{(MAX - MIN)} \quad (\text{Eqn. 4.1})$$

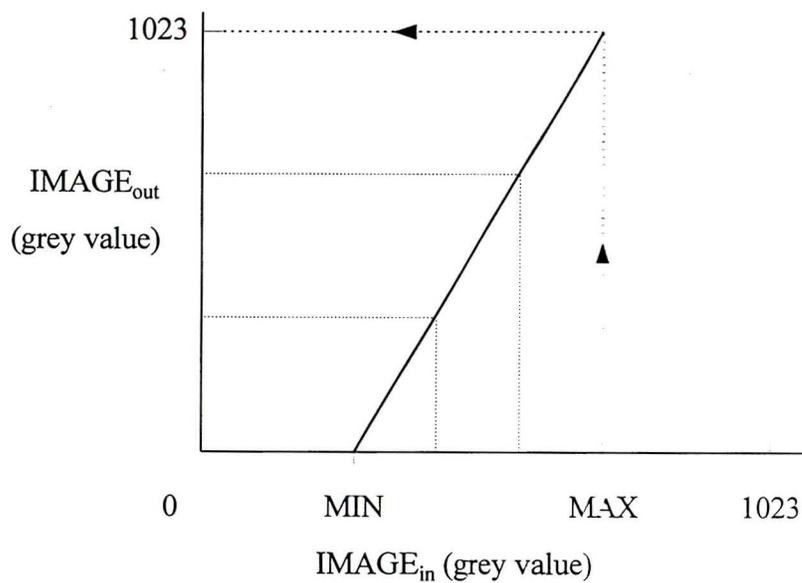


Fig. 4.3 Linear histogram stretching transformation.

Piecewise linear contrast enhancement

In order to improve the contrast in the images, so that the soft tissue structures may be more clearly delineated from the bony structures of interest, piecewise linear contrast enhancement is performed to amplify the high intensity pixel values (mostly bone) and to reduce the low intensity values (soft tissue and background) (Fig. 4.2 (d)). Piecewise linear contrast enhancement is in effect a combination of thresholding and linear stretching. The Fuji AC-1 system provides a non-linear contrast enhancement curve, Type O, that has been optimized for bone radiographs (Fuji Photo Film Co., 1989). This curve maps grey values to photographic density values to achieve maximum contrast between bone and tissue on the image that is printed onto film.

The histograms of our image set (Fig. 4.4) and the Fuji Type O curve (Fig. 4.5) were empirically studied to determine appropriate breakpoints for piecewise contrast enhancement. Breakpoints are the pixel values at which thresholding or a change in the linear stretch factor is desired. Two breakpoints were chosen, one at pixel value 400 and another at pixel value 700 (Fig. 4.6). Below a grey level of 400, most of the pixels belong to the exposed background and some belong to the soft tissues. Above a grey level of 700, most of the pixels belong to bony structures or artificial objects such as metal implants. Between these two grey levels is the range for most of the soft tissue structures. Some of the thinner trabecular bone structures are also within this range.

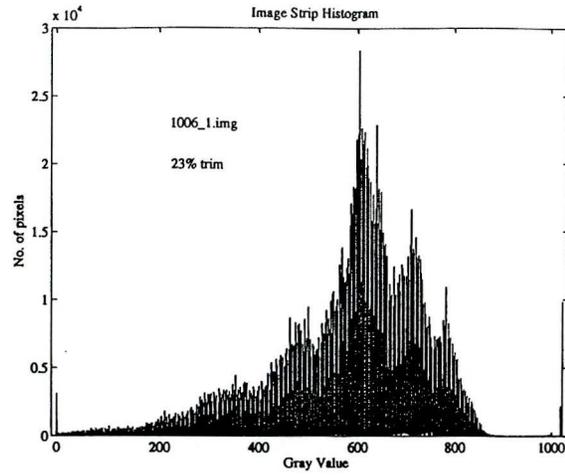


Fig. 4.4 Pixel value histogram of a typical lateral spine image from the image set.

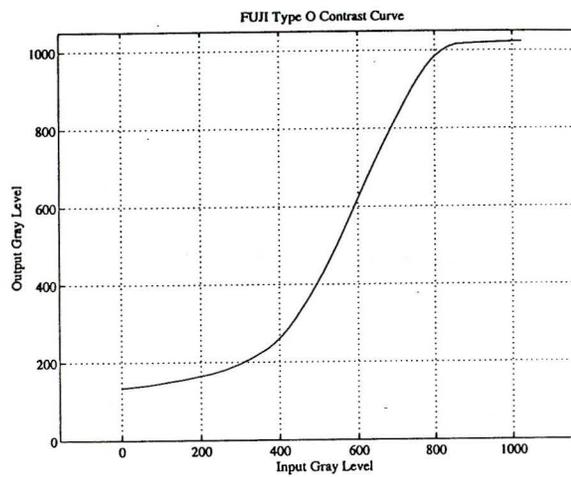


Fig. 4.5 Non-linear contrast enhancement transformation (Fuji Type O curve).

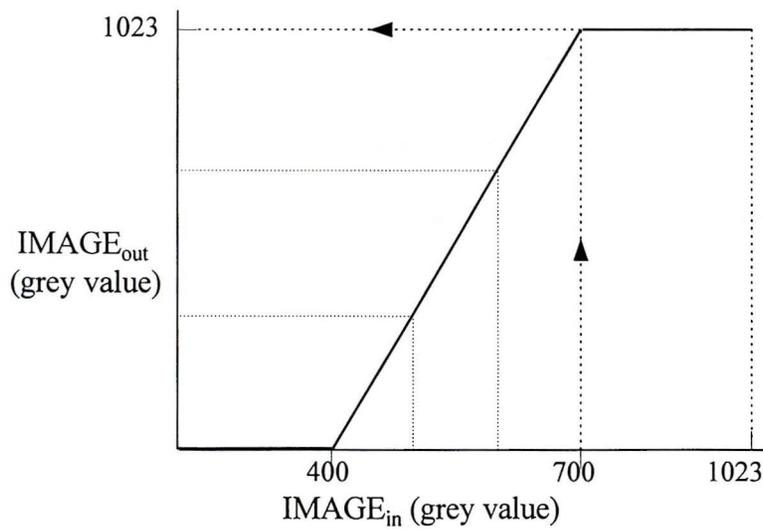


Fig. 4.6 Piecewise linear contrast enhancement transformation.

Edge filtering

Once the contrast has been improved, the image values are better suited to edge filtering. Here, an edge mask of size 3 x 3 is convolved throughout the image clip. This mask is a sum combination of horizontal and vertical Sobel edge masks (Fig. 4.7). The net result is a mask which detects horizontal, vertical, and diagonal edges by amplifying the pixels which are aligned with the mask's edge directions (horizontal, vertical, diagonal). The mask is applied at each pixel value in the image clip (except those at the borders of the image) (Fig. 4.2 (e)).

-1	-2	-1
0	0	0
1	2	1

Horizontal

-1	0	1
-2	0	2
-1	0	1

Vertical

Fig. 4.7 Sobel edge masks.

Linear stretching

Linear histogram stretching as described earlier is applied again to maximize the range of pixel values generated by the edge filter, whose small gradient values tend to produce a darkened image (Fig. 4.2 (f)).

Thresholding

In order to obtain a bi-level image, the image clip is thresholded with a single threshold value, T , determined empirically to be two times the average pixel value in the image clip (Fig. 4.2 (g), Eqn. 4.2). The average value is calculated using all pixel values in the image clip, except those that are equal to 0 or 1023, because these two extreme values dominate the histogram after the image has been edge filtered and linearly stretched. Any pixel value greater than or equal to T is set to 1023 (white), and any pixel value smaller than T is set to 0 (black) (Eqn. 4.3).

$$T = 2.0 * \text{Average pixel value in image clip, IMAGE}[i][j] \neq 0 \text{ or } 1023 \quad (\text{Eqn. 4.2})$$

$$\text{IMAGE}_{\text{out}}[i][j] = \begin{cases} 0, & \text{if IMAGE}_{\text{in}}[i][j] \geq T; \\ 1023, & \text{if IMAGE}_{\text{in}}[i][j] < T \end{cases} \quad (\text{Eqn. 4.3})$$

Mode filtering

Once a bi-level image has been formed by thresholding, a mode filter is applied to remove the isolated pixels which appear mostly within the boundaries of the vertebral body (Fig. 4.2 (h)). Naturally, after thresholding to a bi-level image, the mode will be at either 0 or 1023. The mode filter removes isolated pixels by replacing a pixel value with the most common pixel value in the pixel's 5 x 5 neighbourhood.

Development of search engine

A search engine must subsequently be developed to detect the locations of the corners of the vertebral body. Future work must, however, first concentrate on the preceding steps for obtaining a bi-level image. A possible tracking algorithm is proposed below. It is based on the assumption that a bi-level image has been achieved, yielding a fairly clear white vertebral body outline on a black background.

Given the starting centre point chosen by the user, the search algorithm travels out in all four directions to locate white edge pixels. If for each direction no edge pixel is found within the boundary of the image clip, then a new starting centre point is determined by horizontal or vertical shifting. Using each detected edge point, a rectangular 3x3 neighbourhood is examined for any connected white points.

Corners

Effectively, we are starting the search in two opposite directions from each of the four starting edge points. Intersections of search tracks are stop criteria for each other, and would identify initial corner points. A small region around each initial corner point is selected. A trace in all directions for detected edge points is performed. Where a significant angle exists between a set of three contiguous points, a potential corner point is

identified. Out of these potential corner points, the one identified with the largest angle is selected as the sole corner point.

Gaps

The closest white point in an edge pixel's neighbourhood is taken as the next point on the edge. In order to accommodate gaps in the edge, if there is no white point in the neighbourhood, the neighbourhood is radially enlarged and examined again, until a white point is found. A linear interpolation between a newly found white point and the previous point would fill in gaps linearly, as long as the gap is smaller than say a five pixel distance. If it is greater, then the search skips this gap temporarily, to resume edge tracking starting from the newly found point (on the other side of the gap). Assuming at least two connected points are found on this side, the slopes of the tracks on either side of the gap are computed, and straight lines drawn into the gap from both sides, until the lines meet. In this way, the gap is filled based on the tendencies of the edge segments on either side of the gap. If gaps exist at the corners, this method can also create piecewise linear corners instead of cutting through them.

Measurement points

Once the corner points are identified, the user may confirm each of them, or manually correct incorrect points. Straight line segments are then automatically drawn on

screen, connecting the four corner points. The midpoints of these line segments are then computed. To locate the midpoint on the actual vertebral edge, the slope of the straight line segment is calculated, and the slope of the perpendicular to this line is computed. Using the initial midpoint as a point on this perpendicular, the edge point nearest in euclidean distance is found. Its coordinates tell us which direction (in or out) to search from the initial midpoint to the real midpoint. A loop of coordinates of each point on the perpendicular, in this direction, are matched to existing stored edge points. When there is a match, the point is identified as the true midpoint (ie. on the actual edge of the vertebra). By retracing the already detected vertebral body edge from midpoints to corner points, the pixel distances of the various segments of the vertebral body boundary are determined. These distances do not need to be converted to true distances if ratios of measurements are the morphometric criteria used for classification of deformity.

This search engine runs on the assumption that the image has been well-processed, and a very good and accurate bi-level image delineating the true vertebral outline has been obtained.

4.5.2 GRADIENT SEARCH AND CROSS-CORRELATION DETECTION

This method, developed by J.S. Sloka (1993), aims for full automation of the vertebral detection and measurement process. The algorithm is in two parts : i) vertebral body detection using a cross-correlation measure, and ii) vertebral corner detection using

edge gradients as the search engine. The algorithm is performed on the entire lateral spine image. The user first identifies the top vertebra in the image by placing a box around it. This box is then used by the cross-correlation detector to locate the next lowest vertebra in the image, which in turn is used to locate the third vertebra in the image, and so on. (Currently, the number of vertebrae to be located in the image is set into the program, but this can easily be changed to provide user input for this number, or to have the program include a stopping criterion when the bottom of the image is reached.) Then, at each vertebra, an edge gradient search is performed to track the edges of the vertebral body. Angle criteria are used to identify corners.

METHOD

- 1 : Image preprocessing
- 2 : Cross-correlation vertebrae detection
- 3 : Edge tracing and corner detection

Each of these parts is described in detail in the following sections.

4.5.2.1. *Image preprocessing*

In order to eliminate high frequency noise, the image is passed through a 5 x 5 median filter, where the centre pixel of a 5 x 5 neighbourhood is assigned the median pixel value of that neighbourhood.

4.5.2.2. Cross-correlation vertebrae detection

1 - The user begins by selecting two points defining opposite corners of a rectangle enclosing the top vertebra in the image. The image within this rectangle is stored in a matrix, V_1 , of size r (vertical dimension) by c (horizontal dimension).

2 - A second rectangle, V_2 , of the same size is then automatically located at 10% of V_{dy} pixels below the first rectangle.

3 - As the second rectangle is moved within a neighbourhood of ± 3 pixels in both horizontal and vertical directions, cross-correlation coefficients between the second rectangle and the first are calculated.

$$\text{Cross-correlation coefficient} = \frac{\text{mean}(V_1 * V_2) - \text{mean}(V_1) * \text{mean}(V_2)}{\text{variance}(V_1) * \text{variance}(V_2)}$$

$$\text{where mean}(V_k) = \frac{\sum_{j=1}^c \sum_{i=1}^r V_k(i,j)}{r * c}$$

$$\text{mean}(V_1 * V_2) = \frac{\sum_{j=1}^c \sum_{i=1}^r [V_1(i,j) * V_2(i,j)]}{r * c}$$

$$\text{variance}(V_k) = \frac{\sum_{j=1}^c \sum_{i=1}^r [V_k(i,j) * V_k(i,j)]}{r * c} - [\text{mean}(V_k)]^2$$

4 - The coordinates of the location of highest cross-correlation between V_1 and V_2 are chosen as the coordinates of the next new starting point of V_2 .

5 - Steps 3 to 4 are repeated until a predetermined number of iterations has been reached, or until no new starting point for V_2 is found, whichever situation occurs first.

CROSS-CORRELATION DETECTOR ALGORITHM PSEUDOCODE

```

/* DEFINE RECTANGLE AROUND FIRST VERTEBRA */
input two points (X11,Y11) and (X21,Y21)
c = X21 - X11;
r = Y21 - Y11;

V1[ ][ ] = Image of size r by c starting at (X11,Y11)

FOR (I = 2; I <= NUMBER OF VERTEBRAE ; I = I + 1)
{
    /* LOCATE NEXT RECTANGLE BELOW */

    X1I = X1I-1; Y1I = Y1I-1 - 1.1*r;
    X2I = X2I-1; Y2I = Y2I-1 - 1.1*c;

    /* DO CROSS-CORRELATION SEARCH */

    MAXCORR = 0;
    ITERATIONS = 0;
    DO
    {
        FOR (x = -3; x < 4; x = x + 1)
        { FOR (y = -3; y < 4; y = y + 1)
            { VITEMP[ ][ ] = Image of size r by c, starting at (X1I+x,Y1I+y)
                /* Multiply corresponding elements of the two matrices */
                VIPROD[ ][ ] = Multiply(VI-1[ ][ ], VITEMP[ ][ ])

                /* Calculate cross-correlation */
                XCORR = Crosscorr(VI-1, VIPROD)

                /* Store location if XCORR is larger than previous XCORR */
                if (XCORR > MAXCORR)
                { dxloc = x;
                    dyloc = y;
                    MAXCORR = XCORR;
                }
            }
        }
    }

    /* ASSIGN NEW STARTING POINT TO THAT OF MAXCORR */
    IF ( (dxloc == 0) OR (dyloc == 0) )
        ITERATIONS = STOPNUM;
    ELSE
    { X1I = X1I + dxloc;
        Y1I = Y1I + dyloc;
    }
} UNTIL ( iterations = stopnum ) OR ( no location of higher XCORR is found )
}

```

4.5.2.3. Edge tracing and corner detection

Secondary image pre-processing

The vertebral box identified by the cross-correlation detector is passed through secondary image pre-processing consisting of :

- i) linear histogram stretching
- ii) horizontal edge filtering
- iii) vertical edge filtering

The linear stretching is done to locally enhance the contrast in the identified box so that the subsequent edge filtering can be more effective. The edge enhancing masks which are convolved at each pixel across the image are :



Fig. 4.8 Edge enhancing masks.

Edge tracking

The algorithm uses a similar search tactic for each vertebral corner. First, a box 1/3 smaller than the vertebral box is computed within that vertebral box. The corners of the smaller box are used as starting points for the vertebral corner search. From each starting point, the search moves along a vertical line either up or down (towards the nearest horizontally-oriented edge), locating the pixel with the highest edge gradient (sum of the magnitudes of the horizontal and vertical edge gradients). This pixel is presumed to be on the horizontally-oriented vertebral edge. The search tracks this edge and moves from here towards the desired corner by comparing the edge gradients of the pixels in a directionally determined search neighbourhood corresponding to the direction of the desired corner (Fig. 4.9).

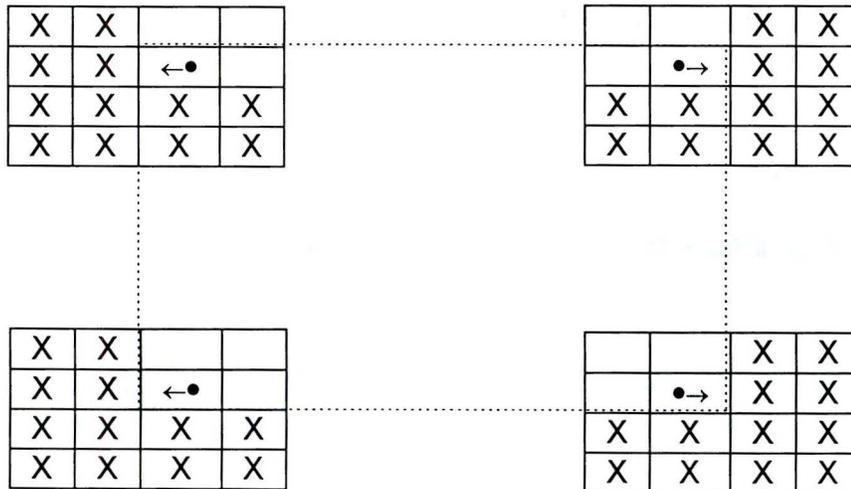


Fig. 4.9 Search neighbourhoods for the four corners.

Corner detection

At each new vertebral edge pixel, the edge gradient angle, as well as the edge gradient magnitude, is calculated. The edge gradient magnitude at each pixel is the sum of the values from the convolution of the directional edge enhancing filters (Fig. 4.8) with the neighbourhood around that particular pixel. The edge gradient angle is given by :

$$\theta (x,y) = \tan^{-1}(G_y/G_x) \quad (\text{Eqn. 4.4})$$

where θ is relative to the x-axis,

G_y is the vertical edge gradient,

G_x is the horizontal edge gradient

This edge gradient angle is perpendicular to the orientation of the edge at (x,y) (Gonzalez and Woods, 1992). For convenience, we will discuss the rest of the corner detection with reference to the actual edge angle, keeping in mind that it is easily calculated by taking the perpendicular orientation to the edge gradient vector.

The edge angle at each new edge pixel is tested against an angle threshold of 35° to the horizontal ($\pm 35^\circ$ relative to the positive or negative x-axis, depending on which of the four corners is being detected). Assuming a rectangularly shaped vertebral body with sides parallel to the Cartesian axes, we assume that a corner has been reached if the edge angle is greater than the angle threshold. However, the vertebrae are most often rotated at some orientation, based on the natural curvature of the spine. To correct for this rotation,

indicates smoothness in the contour and a high value results from sharp corners. E_{image} is the energy of a desired image feature. Whereas E_{cont} and E_{curv} are geometric properties inherent in the desired contour model, E_{image} is actually the image-dependent force driving the snake's slithering towards actual objects in the image. The values of the energy term weights, α , β , γ , influence the snake towards one or another of the energy features. These energy terms and weights will be discussed in detail in Section 4.5.3.2.

Ideally, the snake points should be updated to new locations at which the total snake energy is minimized. This means that the points are updated only at the completion of an iteration down the snake, that is, an integration over the arclength of the snake. However, to reduce computation time, we chose to update the snake points according to a Greedy algorithm (Williams and Shah, 1990), which immediately accepts new locations if they satisfy local minimization. That is, a small neighbourhood, say of 9×9 pixels around a snake point, is examined for the pixel with a minimum sum of the weighted energy terms. The current snake point is then updated to this new locally optimal location. Then, a neighbourhood around the next current snake point is examined, and the point is updated right away if a new locally optimal point is found. This process is repeated at each point in the snake until all points satisfy local minimization of the weighted energy sum.

Why Snakes? In contrast to other segmentation methods, Snakes has the ability to fill in gaps in a contour, with the use of the continuity and curvature terms. This is especially advantageous in complex images such as radiographs, where there may be areas

of very low contrast in the image. Furthermore, typical segmentation techniques such as region-growing and thresholding tend to leave gaps in the image (Street, 1994).

METHOD

1 : Contour initialization

2 : Snake energy minimization

Each of these steps is now discussed.

4.5.3.1 *Contour initialization*

The user first initializes a discrete contour by dragging the mouse around the desired contour. The points on this snake are stored as nodes in a dynamically allocated linked list data structure (Horowitz et al., 1993) (Fig. 4.10). Each node contains the following four fields: x coordinate, y coordinate, grey level, and a pointer to the next node in the list.

Fig. 4.10 Linked list data structure for storing and updating snake (in C)

```
typedef struct snake_node
{ int x;
  int y;
  unsigned short grey_level;
  struct snake_node *nextnode;
} SNAKE_NODE;
```

When the user has completed the contour definition, the node list is 'closed', ie. the last node collected is linked to the first node in the list. Next, the node list is examined for any redundant, repeated nodes, which are removed.

4.5.3.2 Snakes energy minimization

- a. Snake energy terms**
- b. Other parameters**
- c. Snakes algorithm pseudocode**

Snake energy terms

Econt : Continuity of discretized contour

The Econt term aims for continuity of the contour, ie. spline continuity. In our discretized contour, Econt regulates even spacings between adjacent snake points. This is important because snake points tend to bunch up around the local minima of the higher-weighted image energy term (when adjacent neighbourhoods overlap, adjacent snake points may move to the same minimum).

The Euclidean distance between a snake point and each of the points in its neighbourhood is calculated. Econt is then the value of the difference between this adjacent distance and the average adjacent spacing of the snake. Minimizing Econt produces even spacings between adjacent snake points.

The average distance between adjacent snake points is :

$$\underline{d} = \frac{\sum_{i=1}^{N-1} |P_{i+1} - P_i|}{N - 1} ,$$

where $P(x,y)$ is a snake point,

N is the total number of snake points,

$|P_{i+1} - P_i|$ is the Euclidean distance between P_{i+1} and P_i ,
that is, $\sqrt{(x_{P_{i+1}} - x_{P_i})^2 + (y_{P_{i+1}} - y_{P_i})^2}$

At each point in the neighbourhood,

$$E_{\text{cont}_n} = \underline{d} - |P_s - P_n|$$

where $|P_s - P_n|$ is the Euclidean distance from the current snake point, P_s ,

to the candidate point, P_n , in the neighbourhood

E_{curv} : Curvature of contour

This term measures the smoothness in the curvature of the contour. The Euclidean distances of the radii from the snake's centre of mass to each of the two snake points adjacent to the current snake point candidate are calculated. The average of these two radii is used to minimize the radius to the new snake point.

Let $P_c(x_c, y_c)$ represent the centre of mass of the snake, where

$$x_c = \frac{\sum_i^N x_i}{N} \quad \text{and}$$

$$y_c = \frac{\sum_i^N y_i}{N}$$

where N is the total number of points in the snake
 x_i is the x coordinate of each snake point
 y_i is the y coordinate of each snake point

Then,

$$r_{av} = \frac{|P_{s-1} - P_c| + |P_{s+1} - P_c|}{2}$$

where P_s is current snake point
 P_c is the centre of mass of the snake
 P_{s-1} and P_{s+1} are the two snake points adjacent to the current snake point

Finally,

$$E_{curv} = |r_{av} - |P_n - P_c||$$

where P_n is a candidate snake point in the neighbourhood.

Eimage : Image features

Image features such as edges are used to drive the snake towards the desired objects in the image. In our application, we use the vertebral body boundary as the desired contour. A 3 x 3 pixel Sobel edge mask (Gonzalez and Woods, 1992) is convolved across the

neighbourhood of each snake point. E_{image} is the value of the edge gradient at each neighbourhood point.

-2	-2	0
-2	0	2
0	2	2

Fig. 4.11 Sobel edge mask

Energy term weights :

α weights E_{cont}
 β weights E_{curv}
 γ weights E_{image}

$$\text{for } E_T = (1/N) \sum_{i=1}^N \{ [\alpha * E_{\text{cont}_i}] + [\beta * E_{\text{curv}_i}] + [\gamma * E_{\text{image}_i}] \}.$$

These weights are empirically determined. For example, $\alpha=1$, $\beta=1$, $\gamma=-1.5$. We found the α for this set of weights too strongly restricting the snake behaviour to the even spacing continuity constraint. The set of weights $(\alpha, \beta, \gamma) = (0.7, 0.7, -1.5)$ proved better for the snake's convergence, but this α was slightly too weak, causing the snake to spread out unevenly. A reasonable set of weights was $(1.1, 1.0, -1.8)$. Generally, the magnitude of γ is set greater than α and β , so that the snake moves towards features of the image, and 'interpolates' its shape based on the continuity and curvature conditions only where image features are weak (eg. no apparent edge).

Other parameters

Other parameters which needed to be empirically determined were :

- i) $M \times M$ neighbourhood size to search around each snake point
- ii) `Max_iters`, the maximum allowable number of iterations around the snake
- iii) `Moved_limit`, the threshold total number of points moved in an iteration for a subsequent iteration to be initiated

M was arbitrarily set to 9 for our snake algorithm. This neighbourhood size is large enough to provide a good selection of candidate points, however, its size also greatly affects the runtime of the Snakes algorithm. `Max_iters`, when small (between 20 and 25) did not seem enough for the snake to converge. We used a value of 50 for `Max_iters`. If much larger, the snake actually spreads out further and further away from the initial contour, probably because of energy weighting on the E_{cont} and E_{curv} . `Moved_limit` determines the number of points moved in an iteration of the algorithm for the snake to be considered significantly different. It would depend on the number of snake points in a typical snake. We set `Moved_limit` to 2 for initial testing.

Snakes algorithm

SNAKES ALGORITHM PSEUDOCODE

Set α, β, γ (energy weights)

Set M (size of search neighbourhood)

Set Max_iters (maximum number of iterations)

Set Moved_limit (threshold number of snake points moved in an iteration before another iteration is performed)

Head and tail snake nodes are always retained by two pointers

Snake_length (number of nodes in the snake) is always maintained

1. Contour initialization (Section 4.5.3.1)
2. For each iteration down the snake:
 - i. Find average spacing between adjacent points
 - ii. Find centre (of mass) of snake contour
 - iii. For each point on the contour:
 - a. Calculate the E_{cont} , E_{curv} , and E_{image} at each point in the $M \times M$ neighbourhood of the contour point.
 - b. Weight the energy terms and sum to form total energy (Eqn. 4.5b)
 - c. Find the neighbour location with the minimum total energy
 - d. Move current snake point to this new location ("Greedy algorithm")
 - iv. Reinitialize the head and tail pointers
3. Stop iterating if
 - A. fewer than Moved_limit points were moved in previous iteration
 - and*
 - B. Max_iters , maximum preset number of iterations, is reached
4. Display "converged" snake

4.6 CHAPTER SUMMARY

In this chapter, we discussed digital image processing approaches to developing a solution for automatic vertebral detection and measurement. We first discussed previous work, and then outlined the specifications of the image set and the software development environment. Following that, we described in detail three algorithms for the detection and measurement of vertebral body dimensions. The first was an approach based on pre-processing image clips of individual vertebrae until vertebral body outlines are clearly delineated in bi-level images. The proposed search engine then tracks the edges and identifies corners and midpoints using the high image intensity vertebral body outline created by the pre-processing procedure. The second method used an edge gradient search to detect vertebral body edges, and identified corner points based on angle criteria. As a step towards full automation, this method also includes a cross-correlation detector to identify successive vertebrae, once a starting vertebra is identified by the user. The third method, Snakes, was based on the minimization of an energy functional, with energy terms representing spline continuity, contour curvature, and specific image features.

In the next chapter, Chapter Five, the results of applying these algorithms to the image set are presented. Then, in Chapter Six, we discuss these results, compare the performance of each algorithm, and compare their strengths and weaknesses, similarities and differences.

CHAPTER FOUR

REFERENCES

- Doi K, Giger ML, MacMahon H, Hoffmann KR, Nishikawa RM, Schmidt RA, Chua KG, Katsuragawa S, Nakamori N, Sanada S, Yoshimura H, Metz CE, Montner SM, Matsumoto T, Chen X, Vyborny CJ, 1992. Computer-aided diagnosis: development of automated schemes for quantitative analysis of radiographic images. *Seminars in Ultrasound, CT, and MRI* 13:2:140-152.
- Franks LE, 1981. *Signal space theory*. Stroudsburg, PA: Dowden & Culver, Inc.
- Fuji Photo Film Co., Ltd., 1989. *Fuji computed radiography system FCR AC-1 operation manual*. Tokyo: Fuji Photo Film Co., Ltd.
- Gonzalez RC, Woods RE, 1992. *Digital image processing*. Reading, MA: Addison-Wesley Publishing Co. Inc.
- Green WB, 1983. *Digital image processing*. New York: Van Nostrand Reinhold Co. Inc.
- Horowitz E, Sahni S, Anderson-Freed S, 1993. *Fundamentals of data structures in C*. New York: WH Freeman & Co.
- Kahaner D, Moler C, Nash S, 1989. *Numerical methods and software*. Englewood Cliffs, NJ: Prentice-Hall.
- Kass M, Witkin A, Terzopoulos D, 1988. Snakes: active contour models. *Int J Computer Vision* 1:321-331.
- Lastman GJ, Sinha NK, 1989. *Microcomputer-based numerical methods for science and engineering*. New York: Saunders College Publishing.
- Parker JR, 1994. *Practical computer vision using C*. New York: John Wiley & Sons Inc.
- Sloka JS, 1993. *The detection of spinal vertebra characteristics*. Private report.
- Street WN, 1994. *Cancer diagnosis and prognosis via linear-programming-based machine learning*. University of Wisconsin-Madison PhD dissertation.
- Williams DJ, Shah M, 1990. A fast algorithm for active contours. *Proc Third Int Conf Computer Vision (Osaka, Japan)*, 592-595.
- Yu SK, 1994. *Modified median filtering*. Personal communication.

CHAPTER FIVE

RESULTS

5.1 INTRODUCTION

In the previous chapter, we described in detail the three approaches and algorithms for detecting vertebral body boundaries and corners. These methods are quite different from each other, and their results differ accordingly. In this chapter, we present the results of these algorithms, and point out salient observations in the output images.

5.2 RESULTS OF THE THREE ALGORITHMS

The image gallery of the following pages summarizes the results of the three algorithms we pursued for automating the detection of vertebral body boundaries from lateral spine images. The image intensity based algorithm will be abbreviated as IIB; the gradient search and cross-correlation algorithm as XES; and the Snakes contouring algorithm will be referred to as Snakes. The following results are grouped according to the original source images, which are identified as RH, TS, WJ, RG, BC1, BC2, and AG. This set of images represented a wide range of image quality. Some of the images were of poor quality, being so overexposed that the intensity level of the tissues was near that of the bone; some of the images were of average quality, with reasonable radiographic contrast between bone and tissue in most areas; and one image displayed the anatomical structures with excellent contrast. Figures 5.1 to 5.3 are representative results from the image intensity based algorithm (IIB) for isolating the vertebral body boundary. The results of processing most of the spine images in the image test set are given in Figures 5.4 to 5.16. Two example results of applying the Snakes contouring algorithm are given in Figures 5.17 and 5.18.

5.2.1 IMAGE INTENSITY BASED APPROACH (IIB)

Figures 5.1 to 5.3 are representative results from the image intensity based algorithm for isolating the vertebral body boundary.

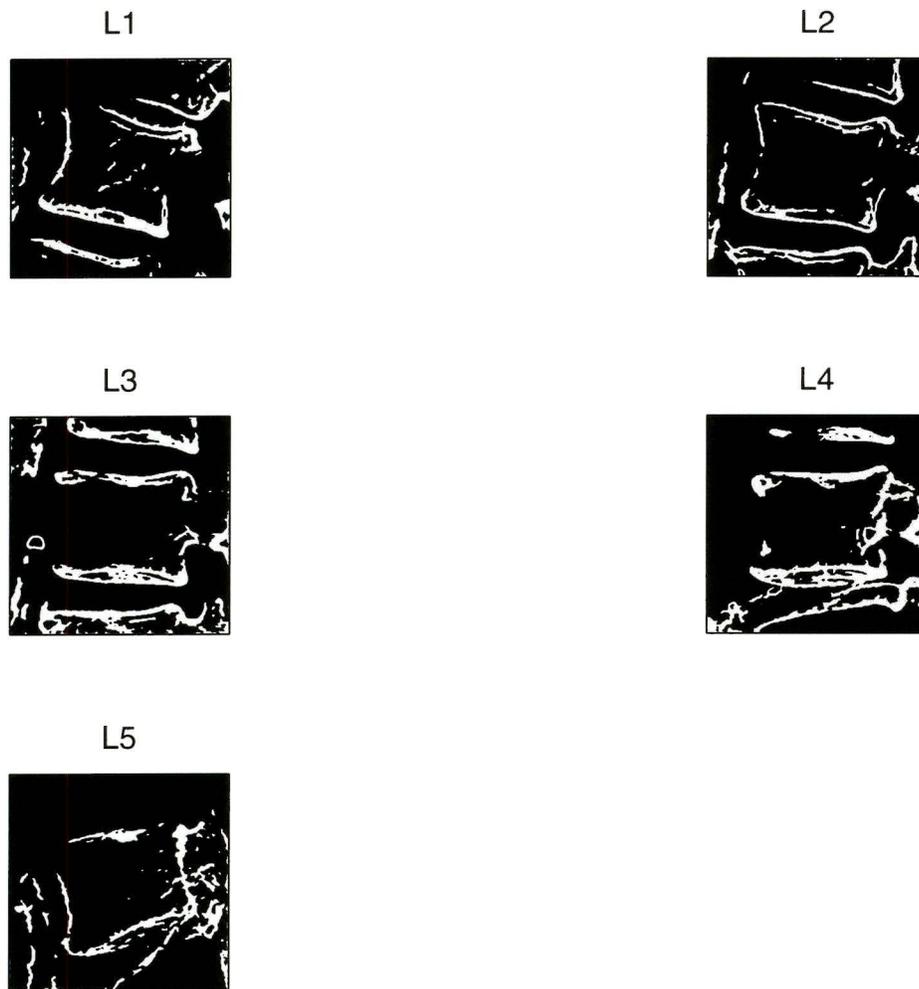


Fig. 5.1 [RH] : The vertebral body boundaries of the five lumbar vertebrae of this spine are successfully isolated. As a result of higher intensity endplate signals, the horizontal edges of the boundaries are clearly stronger and more connected than the vertical edges. Large gaps at two corners (on L1 and on L5) are apparent. Double ridges have been detected, though not always well connected. For the most part, the signal from the centrum, or interior of the vertebral body, has been effectively removed.

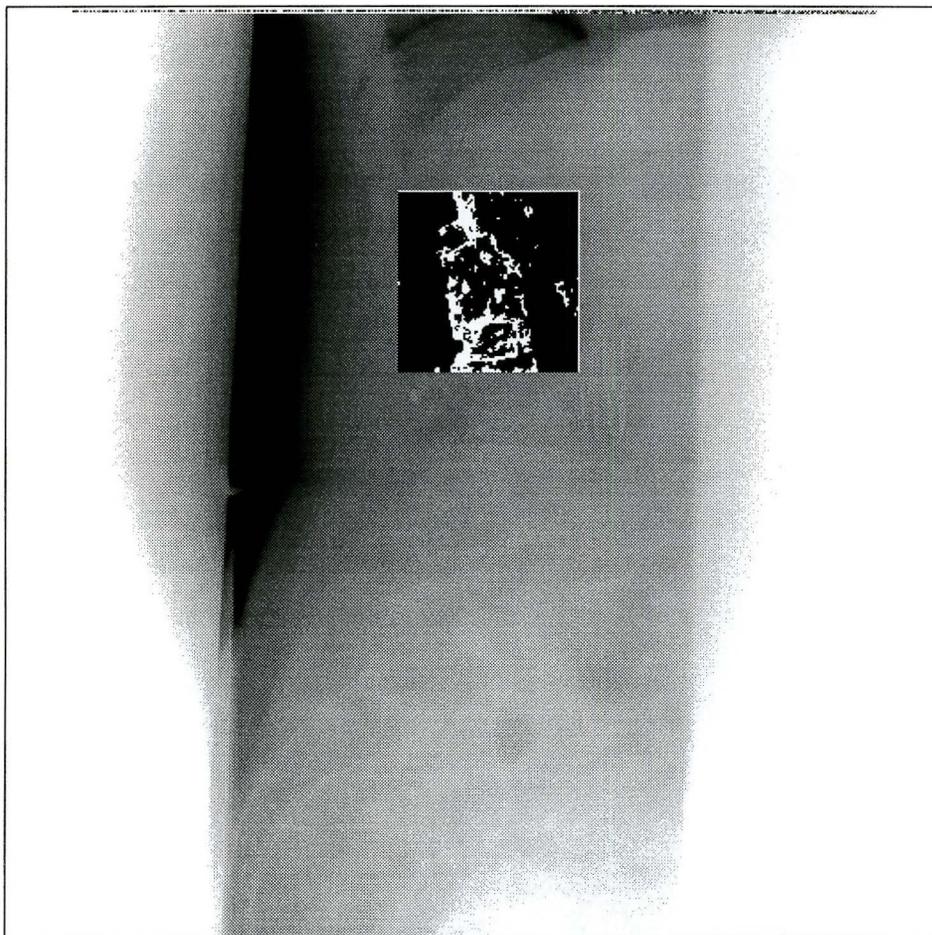


Fig. 5.2 [TS] : An unsuccessful isolation of the vertebral body boundary : rib shadows crossing diagonally through the vertebral bodies have stronger edge gradient signals than the actual vertebral body edge. Note also that in this image, the dark lung shadow overcasts the top vertebral bodies.

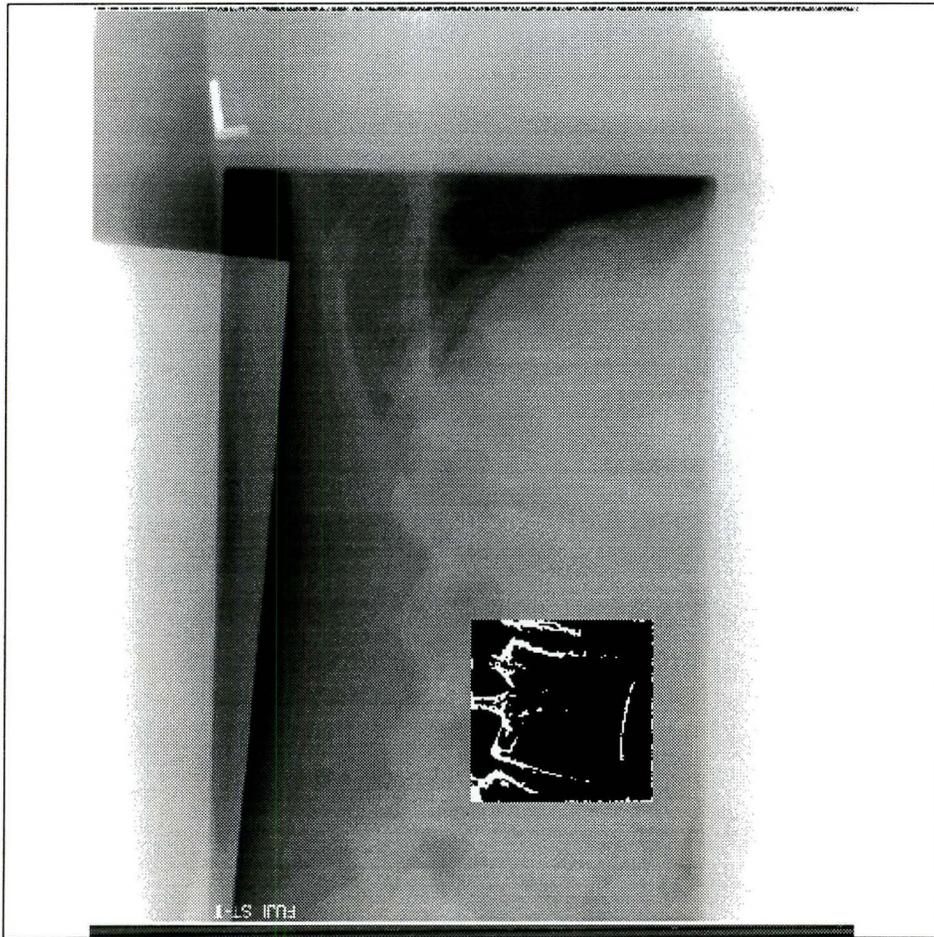


Fig. 5.3 [WJ] : The uniformity in image intensity of this image allows IIB to work fairly successfully. The threshold value may be too high, causing gaps near some of the corners.

5.2.2 GRADIENT SEARCH AND CROSS-CORRELATION DETECTION (XES)

The results of processing some of the test images with the XES algorithm are given in Figures 5.4 to 5.16.

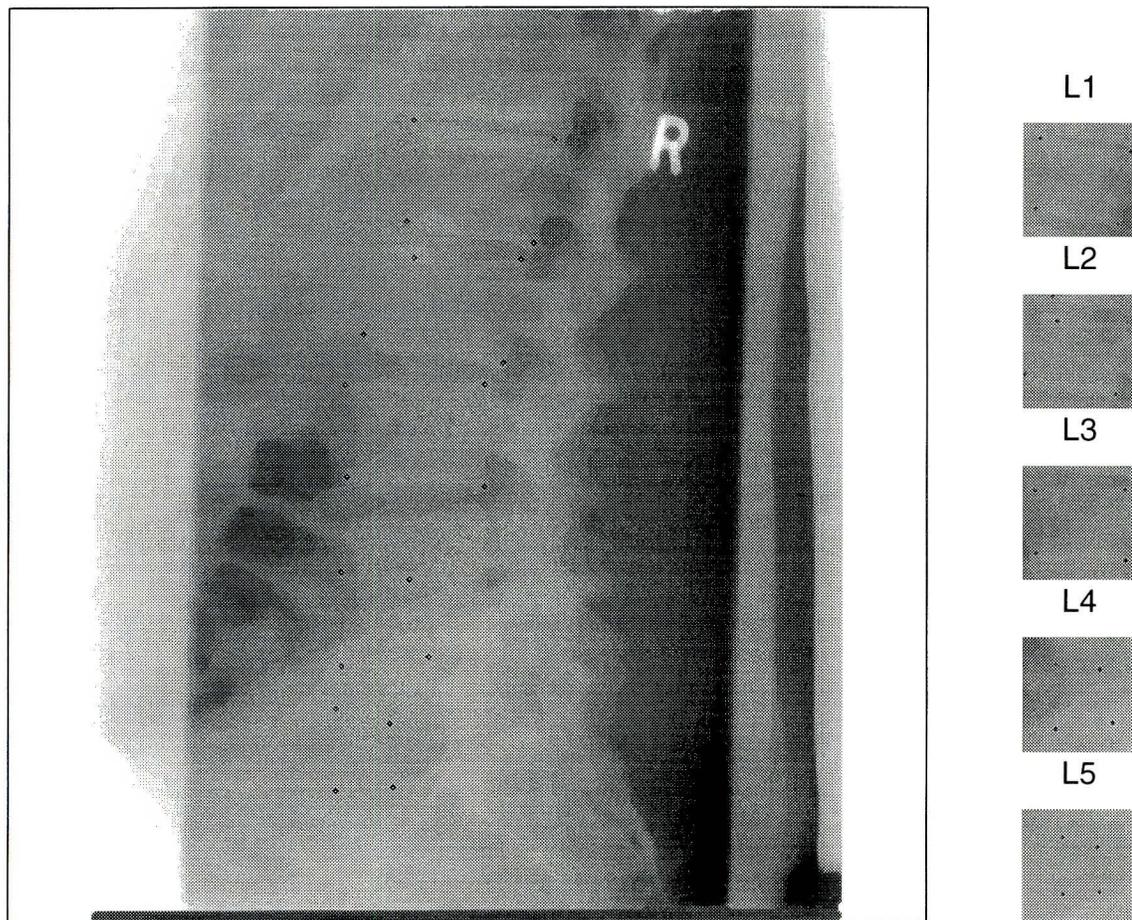


Fig. 5.4 [RG] : The acquired image has average contrast between tissue structures, and displays the vertebrae well. The first three vertebrae are well detected by the algorithm. L4 and L5 are hindered by the iliac crest and the ilium, and are poorly localized by the cross-correlation vertebra detector. It appears that the search has moved too far down and missed L4.

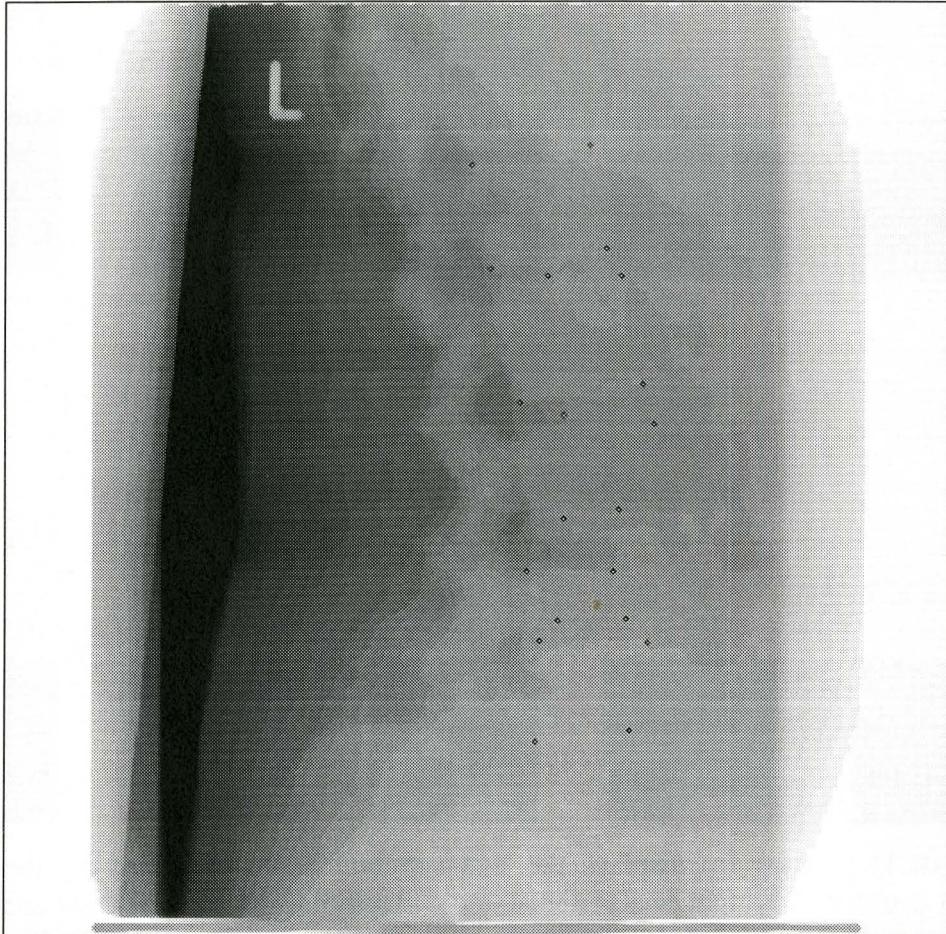


Fig. 5.5 [BC1] : This image displays clear lumbar vertebrae, and the overall contrast is quite good. The two small white points on the first vertebra indicate the starting points. The overall corner detection fails for one corner of the second vertebra, and for all subsequent vertebrae. The next figure shows the individual vertebrae.

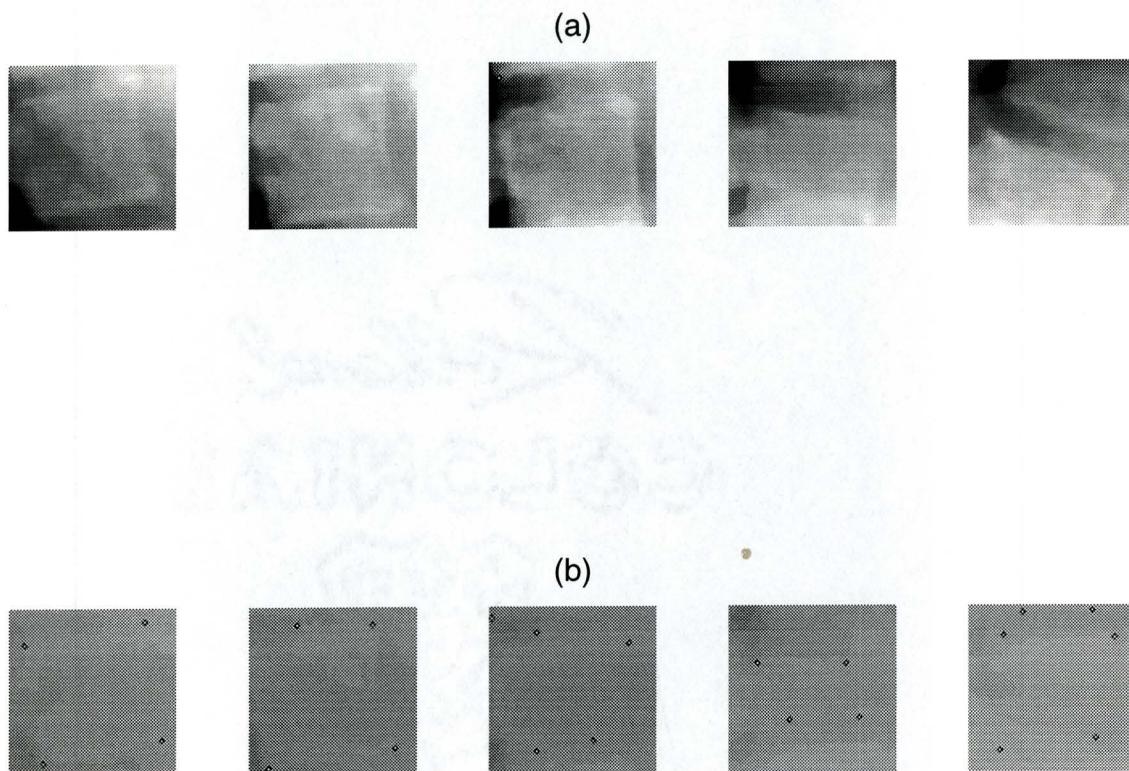


Fig. 5.6 [BC1] : Row (a) displays the five vertebral boxes identified by the cross-correlation detector. The first two vertebrae are correctly localized. The third and fourth vertebrae are cut off at the inferior ridge. The fifth vertebra is buried in the high intensity region of the ilium. The fifth vertebral box incorrectly encloses the upper and lower halves of adjacent vertebrae. Perhaps the number of iterations of the cross-correlation search is constrained too soon, and more iterations would allow the box to move further down. Row (b) displays the detected vertebral body corners, given the vertebral boxes identified in row (a). Corner detection on the first vertebra is good. On the second vertebra, the gradient search is misled by the strong rib shadow evident in the corresponding image in row (a). The iliac crest is incorrectly identified as the lower ridge of the fourth vertebra. Corner detection on the fifth vertebra is incorrect because of an incorrectly localized vertebral box. L5 is buried by the high intensity iliac signal, and is rotated at a greater angle than the previous vertebrae.

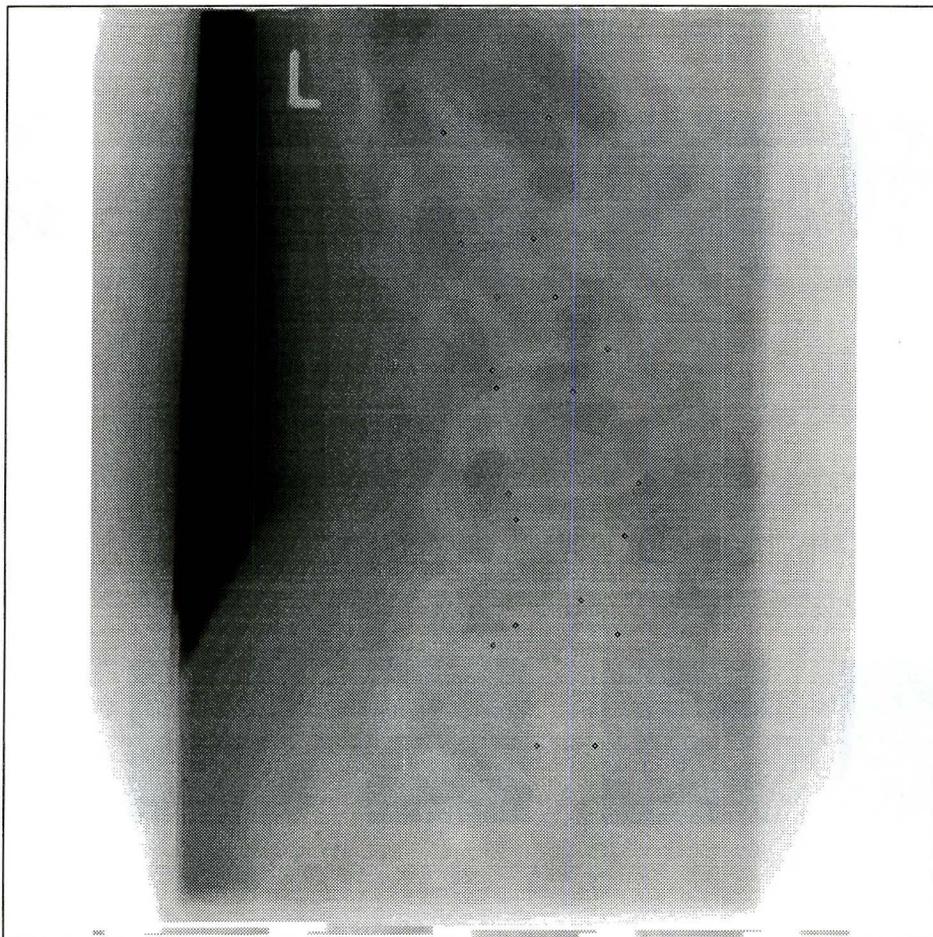


Fig. 5.7 [BC2] : This image has very good contrast between the tissue structures, and displays the lumbar vertebrae clearly. Present in the image are many gas bubbles causing dark blotches along the spine. The output of XES shows many incorrectly identified corner points. The next figure displays the individual vertebrae.

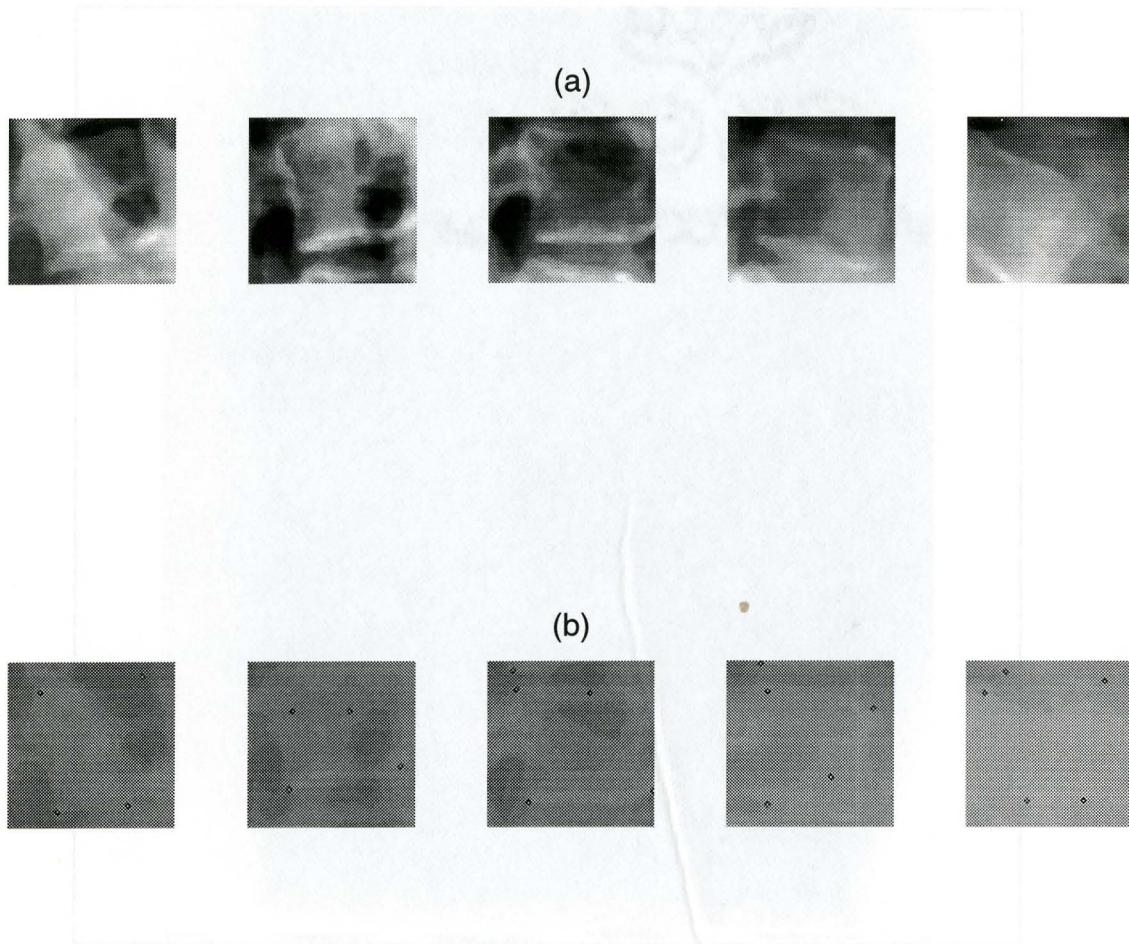


Fig. 5.8 [BC2] : Row (a) displays the five vertebral boxes identified by the cross-correlation detector. The vertebrae have been localized better than in image BC1, probably because of the better contrast in BC2. However, the detected corners, as shown in row (b), are worse than in BC1, because of the gradations of the bowel gas blotches over and across the vertebral bodies.

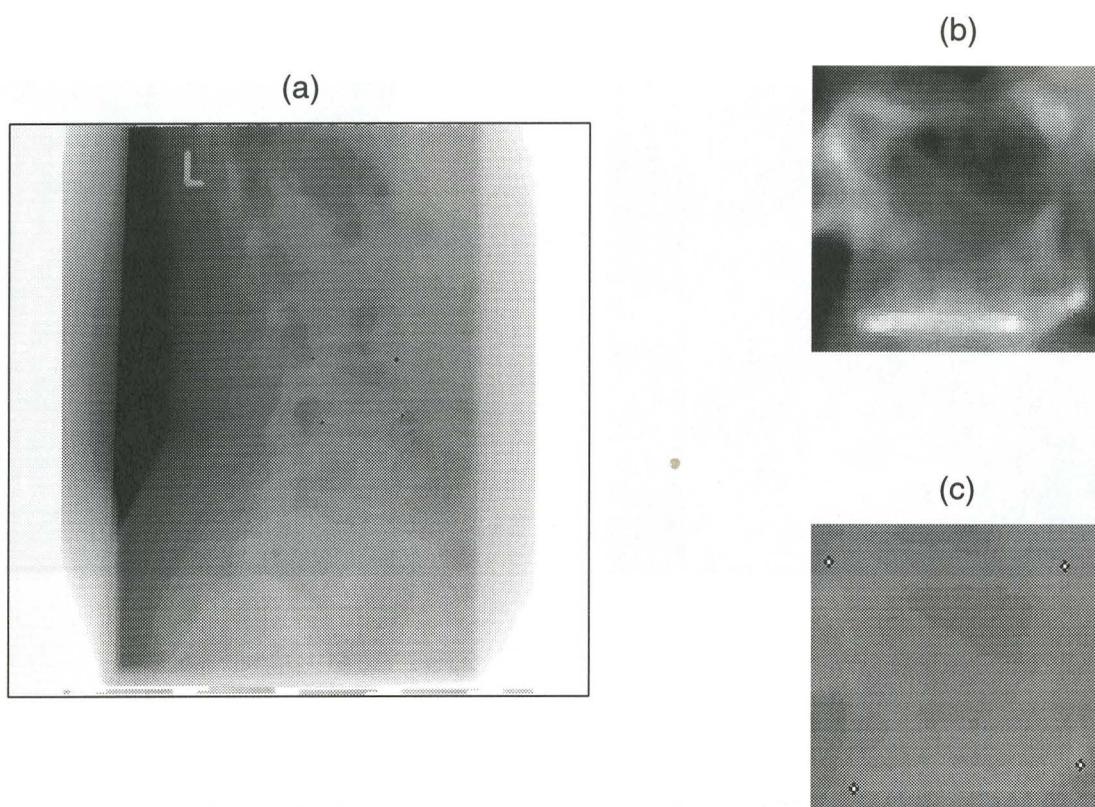
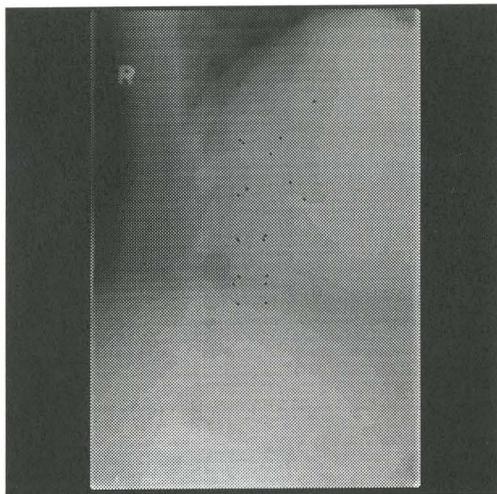
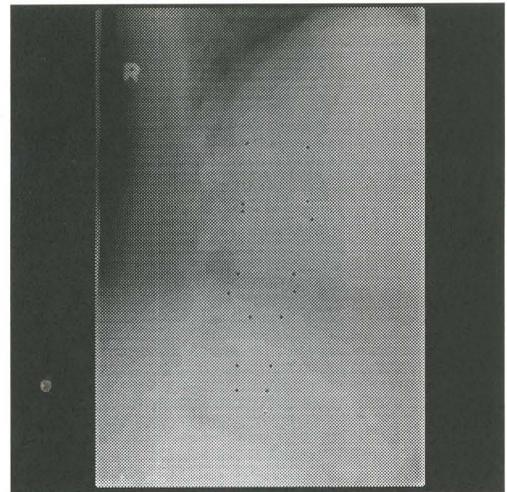


Fig. 5.9 [BC2] : This set of figures shows the result of processing L3 with only the gradient search corner detector. Fig. 5.9(a) displays the view of the spine with the corners identified on L3. Fig. 5.9(b) is the contrast enhanced vertebral box selected by the user, and Fig. 5.9(c) displays the corners detected on this vertebra. The result is very good, and it demonstrates the dependence of the success of the gradient corner search on a well-localized and well-centred vertebral body box.



(a)



(b)

Fig. 5.10 [AG] : AG is an image with well-formed vertebrae but poor contrast between the bone and the soft tissue. In Fig. 5.10(a), the starting box was chosen around the top vertebra. Pronounced rib shadows over L1-L3 are leading the search algorithm astray. Only a few corners are correctly identified. In Fig. 5.10(b), L2 was chosen to be the starting box. The corners of L2 and L3 are now quite well-detected, and both the vertebrae and the corners of L4 and L5 are missed completely. L4 is seen to be overshadowed by the intensity variations at the waistline (see next figure).

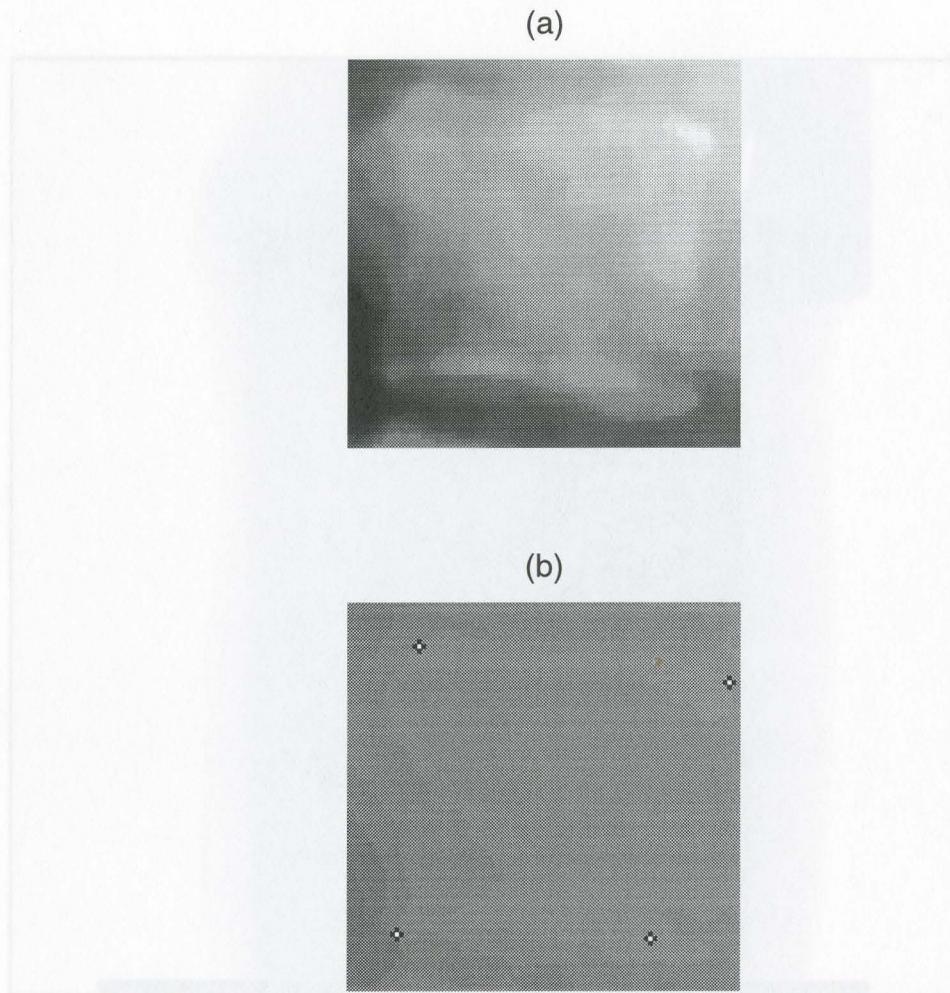


Fig. 5.11 [AG] : The result of selecting L4 and performing only the gradient search corner detection stage of the algorithm. Fig. 5.11(b) displays the corners detected on the vertebra as shown in Fig. 5.11(a). Good detection is achieved, therefore demonstrating that the misdetections of L4 in Fig. 5.10(a) and (b) result from incorrect vertebra localization by the cross-correlation detector. Again, this error may be caused by the intensity variations in the hip region of the spine.

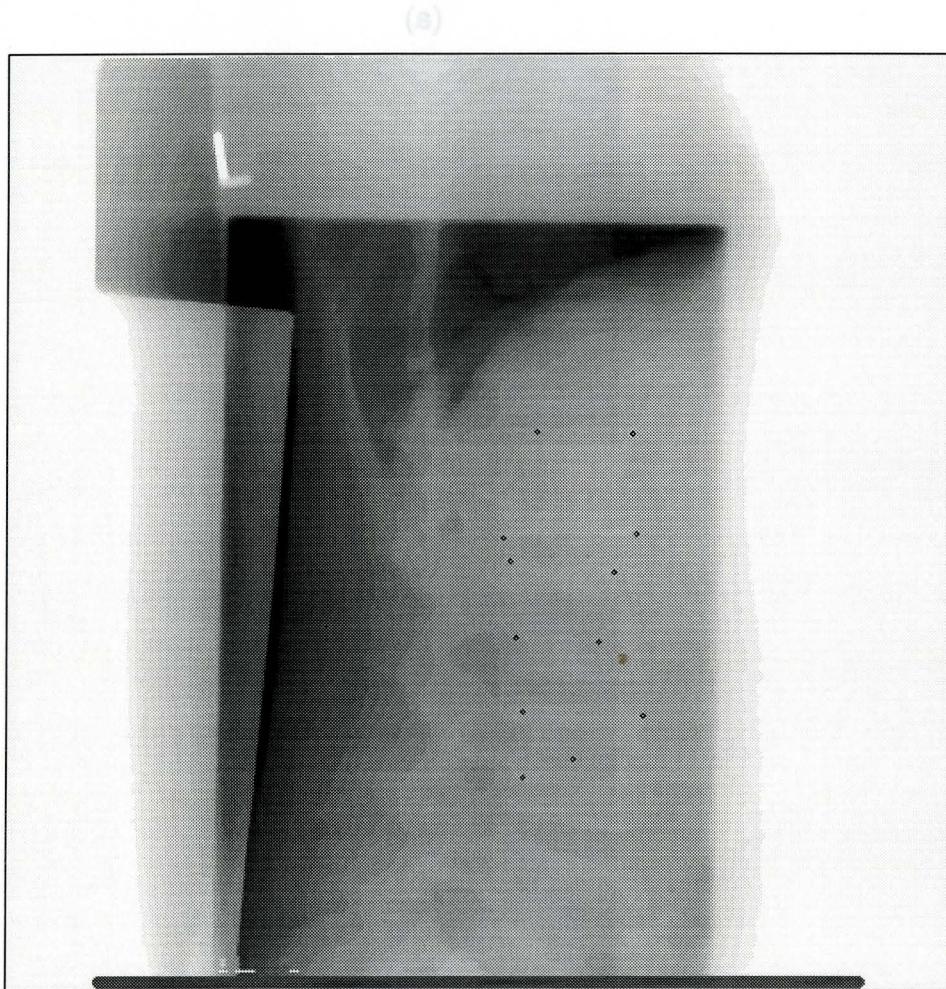


Fig. 5.12 [WJ] : WJ is an image of poor contrast, but large, well-formed vertebrae which are fairly parallel and horizontally oriented. This image is the result of XES on image WJ.

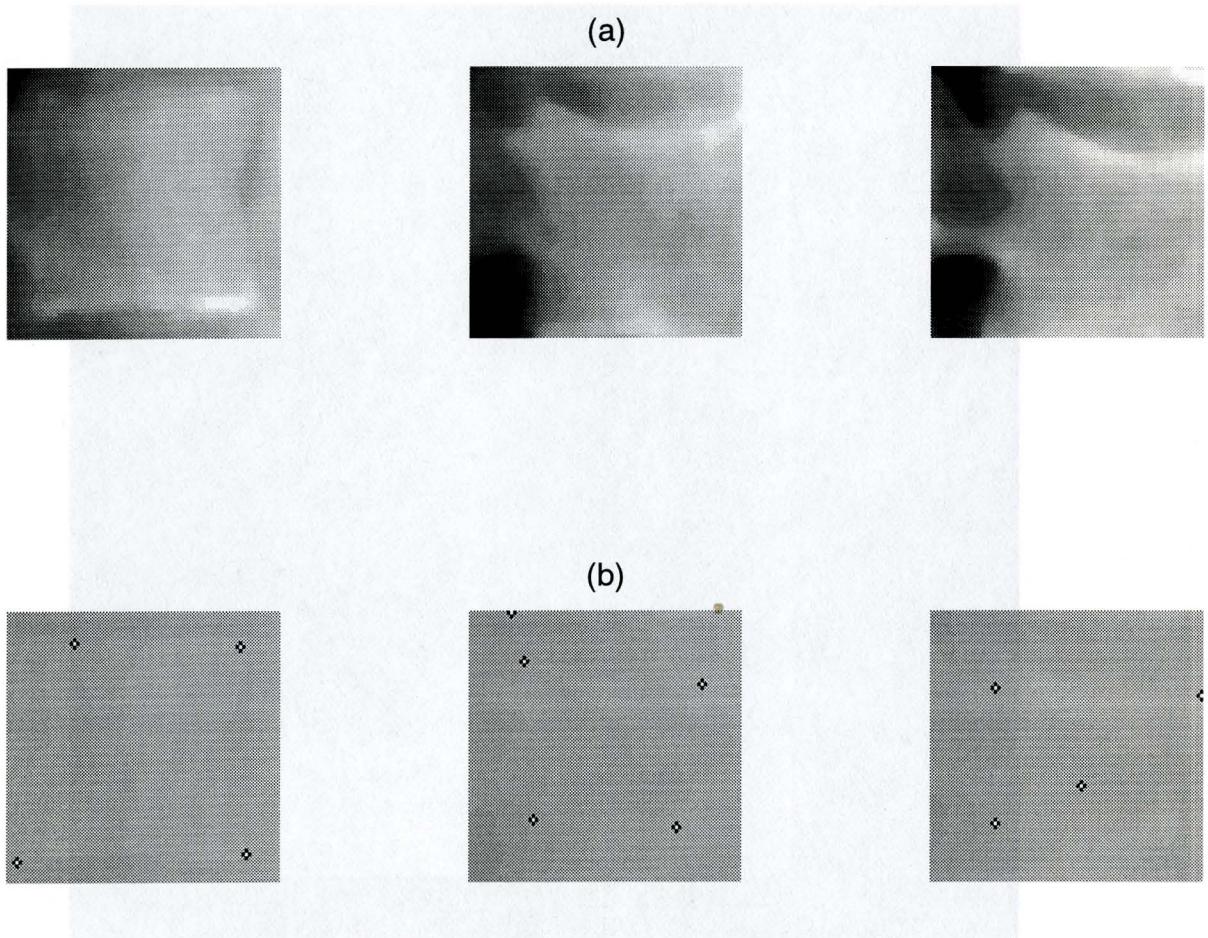


Fig. 5.13 [WJ] : This series of images shows that the cross-correlation detection does not perform as well as might be expected from an image with good alignment and similar orientation of the vertebrae. The vertebral boxes identified as in row (a) do not enclose the vertebral bodies, other than for the starting vertebra. Perhaps the initial box size is not large enough to enclose subsequent vertebrae, which require larger boxes because of their rotated orientation.



Fig. 5.14 [RH] : RH is an example of a fine radiograph, displaying excellent contrast levels. This image shows that the XES algorithm as a whole has some measure of success in detecting and identifying the corners on all vertebrae, given starting points for the first vertebra (the small white points). The next figure shows individual vertebrae.

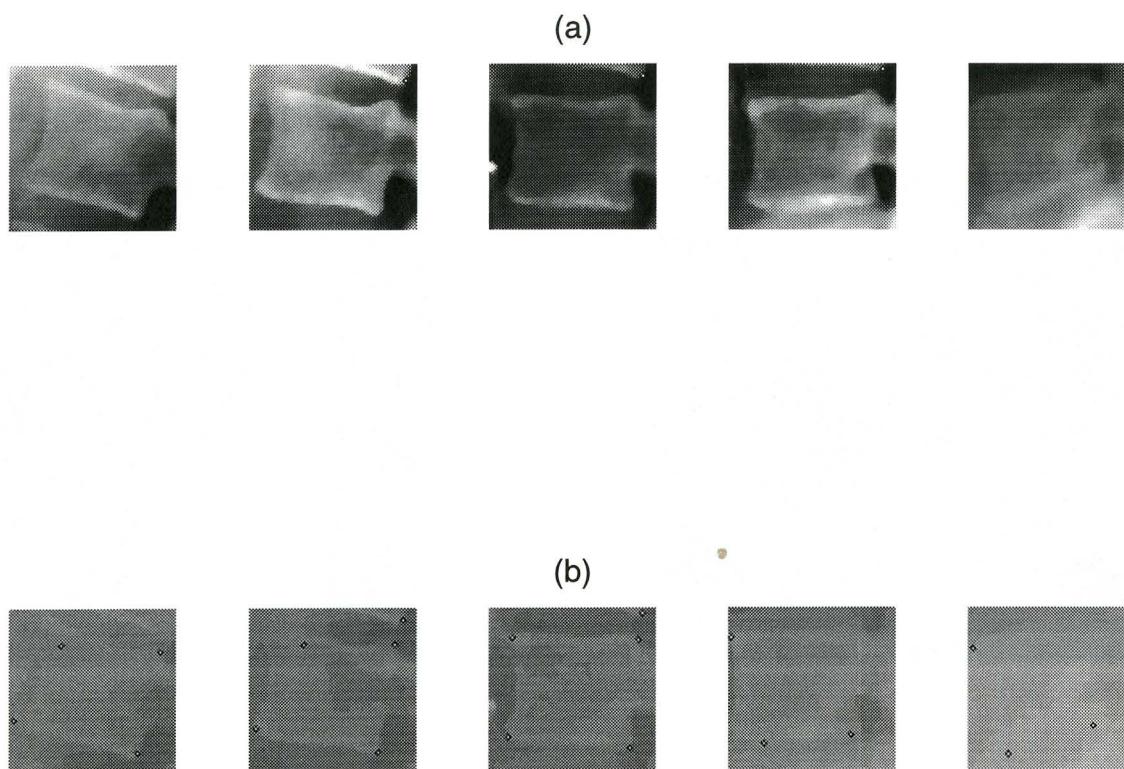


Fig. 5.15 [RH] : Here, the individual vertebrae of RH are shown. Row (a) shows that the cross-correlation detection has been very successful, enclosing each vertebra fully and centrally. Row (b) and Figure 5.14 show that, for L1, a clear diagonal rib shadow across the vertebral body did not affect the corner detection. In L2, variations in intensity caused misdetection of the top left corner. The corners of L3 were excellently detected. In L4 and L5, the search for the top right corner has been led astray by strong edge signals from adjoining pedicles.



(a)



(b)

Fig. 5.16 [RH] : This figure compares the results of XES on RH, starting at L1 in Fig. 5.16(a) and at L2 in Fig. 5.16(b). There are no apparent differences in the the corners identified on L2 through L5 in both images. Unlike image AG, this image is not sensitive to the starting vertebra.

5.2.3 SNAKES ACTIVE CONTOURING

Two example results of applying the Snakes contouring algorithm are given in Figures 5.17 and 5.18.



Fig. 5.16 [RG] : The results of applying Snakes contouring to vertebrae of RG. The red contour is the user-initialized contour, and the yellow dots constitute the result of the Snakes algorithm. The (yellow) snake points are scattered and little or no convergence onto the vertebral body boundaries is seen.



Fig. 5.16 [RG] : The results of applying Snakes contouring to vertebrae of RG. The red contour is the user-initialized contour, and the yellow dots constitute the result of the Snakes algorithm. The (yellow) snake points are scattered and little or no convergence onto the vertebral body boundaries is seen.

5.3 CHAPTER SUMMARY

In this chapter, we reviewed the results of processing spines and vertebrae with each of the three algorithms described in Chapter Four. The observations described here will be discussed in the following chapter.

CHAPTER SIX

DIGITAL IMAGE PROCESSING : DISCUSSION OF RESULTS

6.1 INTRODUCTION

In this chapter, we evaluate the performance of the three algorithms for vertebral body detection. The results from Chapter Five will first be summarized and discussed according to common observations based on image quality and the nature of lateral spine images. Then, we compare the strengths and weaknesses of the three algorithms at their present stage, and where possible, suggest improvements.

6.2 DISCUSSION OF RESULTS : IMAGE CHARACTERISTICS

From Chapter Five, we see that the performance of the algorithms can be evaluated according to two main groups of observations : those relating to image quality, and those relating to the nature of spine radiographs. While the image intensity based method aims to separate the boundary signal from all other signals from the vertebra, the cross-correlation method localizes a vertebra by its region similarity to the previous detected vertebra. Consequently, both methods are very sensitive to image quality and the presence of unwanted signals. In this discussion, the image intensity based approach is abbreviated as IIB, and the edge gradient and cross-correlation detection method as XES. Snakes contouring is referred to as Snakes.

6.2.1 Image quality

Our images had three general levels of contrast : good, average, and poor (see Figures 5.14, 5.4, and 5.10 respectively). Because of these variations across images, an important part of the detection algorithms was image enhancement. We noted significantly better boundary and corner detection by the IIB and XES algorithms when the image had good contrast. When there is good contrast, the differences in grey level of the tissue structures are greater, and a better classification of the structures according to image intensity is achieved.

6.2.2 Nature of lateral lumbar spine images

We summarize here the common features of a lateral lumbar spine image which led to most of the problems encountered by the detection algorithms. These are 'facts' of lateral lumbar spine images, and will require special consideration in the further development of the algorithms.

Ribs

Rib shadows are visible in most spine images. The ribs extend down from the lower thoracic vertebrae and cross diagonally across the upper lumbar vertebral bodies. As we saw in Figure 5.2, the strong edge signals from the ribs can lead the algorithms to terminate prematurely at rib edges instead of at vertebral body edges. As well, there lies a potential challenge in classifying rib signals apart from the edges of wedge-shaped vertebrae. Although rib shadows generally fall at a sharper angle than the edges of normal vertebrae, it is possible that they may lie in the same orientation as the sloped edges of wedge-shaped vertebrae.

Ilium

In general, the ilium (hip bone), and especially the iliac crest (the curved boundary of the ilium), overshadow the fourth and especially the fifth lumbar vertebrae. We noted in the results that the performance of XES was often favourable for L1-L3, but dropped significantly for L4-5, both in the vertebra localization and in the corner detection (Fig. 5.4). An additional difficulty in this region of the spine is the strong edge signal from the iliac crest curving through a vertebral body (Fig. 5.6). This curve may even be of the same orientation as a vertebral boundary, thus making it difficult to place angle constraints on the detection algorithm in this region.

Rotation of vertebrae

The axes of alignment of the vertebrae follow the natural curvatures of the spine in kyphosis and lordosis. As well, depending on the amount of flexion by the patient, there is a wide variation in the angle of rotation of each vertebra (Fig. 5.8(a)). Such variation from one vertebra to the next may challenge the cross-correlation vertebral detector.

Bowel gas

Bowel gas bubbles may be present in an image. These dark blotches may interfere with the visibility of vertebrae, and can also cause the algorithms to detect a false edge and terminate. Bowel gas shadows caused major detection errors with image BC2 (Fig. 5.7).

Lung shadow

In some of the images the bottom of the dark lateral lung shadow covered parts of or whole vertebral bodies (Fig. 5.12). This lung region obscures the top vertebrae, and may cause misdetection of vertebral boundaries and corners.

Double ridges

An advantage of the IIB method is that, because it highlights the strongest vertebral body edge signals, it also detects any double ridges that appear (see L1 and L3, Fig. 5.1). This may be very useful for identifying accurate midheight locations. Snakes contouring can produce only a single line contour.

Pedicles

Strong edge signals from the adjoining pedicles may be incorrectly determined by the algorithms to be vertebral body boundary signals.

Artifacts

In general, artifacts are any objects in the image which do not belong in the normal context, and which cause unwanted signals. With BC2, we noted a large process of a very high intensity descending through the middle of L5 (Fig. 5.7). As a result, the XES search incorrectly terminated at one edge of this signal. In TS (Fig. 5.2) and RH (Fig. 5.14), arterial nets cause high intensity signals near to but not touching the vertebral bodies. Artifacts would not generally be significant to a method such as Snakes contouring, but would interfere greatly with methods such as IIB and XES, because of their dependence on low-level image information.

6.3 DISCUSSION OF RESULTS : ALGORITHMS

6.3.1 PERFORMANCE AND FUTURE DIRECTIONS

The success of each of the three algorithms is now discussed in detail, in relation to the results in Chapter Five. Where relevant, suggestions for future improvements are given.

6.3.1.1 Image intensity based approach

As seen in Chapter Five, the image intensity based routine can be effective at isolating and highlighting the vertebral body outline. We noted that the contrast in the image affects how much the vertical sides of the vertebral body appear in the bi-level image. These sides are lower in intensity than the horizontal ridges, and can often be within the grey level range of the centrum. While the thresholding step desirably eliminates the edge-filtered signals from the centrum, it can also easily eliminate some of the vertical vertebral body sides as well. This leads to gaps in the vertebral body boundary, and a loss of connectivity of the contour. However, depending on the search routine, we may not require all four sides to be well-connected. Secondly, we noted that the mode filtering reduced the speckles remaining in the centrum after thresholding. The speckles may be caused by signals in the centrum arising from variations in image intensity noise in the centrum.

Currently, the success of this algorithm depends largely on the threshold value, T . If too high a value is chosen, greater gaps in the boundary are created. If too low a threshold, other signals in the surrounding regions become visible. A suggestion for dealing with this challenge is to aim for more localized processing. Instead of processing the 400x400 image clip, we could take smaller 9 x 9 neighbourhoods and process them serially. Doing so may eliminate the problem where the interior of the vertebral body is almost as bright and dominating as the edges are. Neighbourhood thresholding may be especially effective.

An advantage of the minimal user-interaction involved in this algorithm is in the flexibility it gives the user for specifying the vertebrae to be analyzed. Also, the robustness requirement on the search algorithm is reduced by giving it a starting point. This may save a significant amount of time, and may prove to be more accurate than a completely automatic routine. An alternate algorithm based on the similar approaches of IIB and XES can perhaps begin with the user selecting only the vertebral body centres in the images. The angle of the axis from one centre to the next is then calculated. For each vertebra, we can take its axis of orientation as the average of the axes from its centre to the centre above and the centre below. The axis of orientation can then be used to adjust the angle criteria and edge orientations.

Instead of 400 x 400 rectangular image clips, we could try processing trapezoids with parallel sides oriented parallel to the axis of orientation of each vertebra. The disadvantage of rectangular clips is that the high intensity ridges of adjacent vertebrae may be present in the image clip. We want to eliminate these 'artifacts' and have only the desired vertebra in the image clip, and with no residues from other vertebrae. The challenge in processing trapezoids would be in registering the image clip for rectangular filter operations.

6.3.1.2 Gradient search and cross-correlation detection (XES)

This method proves promising for semi-automatic detection and measurement of vertebrae. Its search algorithm takes advantage of the fact that the superior and inferior ridges of the vertebral body are of higher intensity than the vertically oriented edges.

Problems arose with the cross-correlation detection of vertebrae that were further down away from the initial starting vertebra (Fig. 5.5). Some of this misdetection can be attributed to changes in image contrast and the presence of the iliac crest in that region of the image. It may be possible to improve this weakness by controlling the maximum number of iterations of the cross-correlation detector.

As we saw with image AG (Fig. 5.10), the XES algorithm is very dependent on the successful corner detection of the initial vertebra, since the recentred box around that vertebra is used for localizing the next vertebra. If the localization is inaccurate or not well-centred on the vertebra, the corners will subsequently not be detected correctly, and the error will be propagated down to the following vertebrae. As well, the first box must be chosen by the user to be no smaller than any box around each of the vertebrae in the image. If the first box is small (perhaps because of a smaller sized starting vertebra), then all other vertebrae will not fit in this box size (especially if they require larger rectangular boxes to account for their rotated orientation), and their corners will be incorrectly located. Thus this algorithm is very dependent on the success of the vertebra localization. It may be desirable to separate the two stages of the algorithm and allow for user correction of the vertebra localization stage before the corner detection stage is executed.

Presently, the search routine begins from four starting points, located at $1/3$ distance away, in both x and y directions, from each box corner. In a few cases where a vertebra has a highly rotated orientation, $1/3$ distance away might be either very close to the vertebral edge, or in fact beyond it, so that the search will head out outside the vertebra. This is potentially a problem for severely biconcave deformed vertebrae. On the other hand, if a larger distance away from the box corners is used to locate starting points for the search, the search may detect strong edges within the centrum and terminate prematurely at these edges. This is especially conceivable at the posterior edge, where the neural arch extends out. We also saw instances where strong rib edges lay in the path of the gradient search and

caused misdetection of the corners, and other instances where the rib shadows do not mislead the algorithm because the $1/3$ distance just misses them.

On some of the vertebrae (for example, L4 and L5 of RH (Fig. 5.15)), corners were identified outside of the vertebral box identified by the cross-correlation detector. This effect should be modified by a constraint on the x-direction gradient search distance. In addition, different search masks could be attempted.

At present, the amount of image pre-processing for the method is minimal (median filtering and linear stretching only); however, it may be beneficial to contrast enhance local image regions for the cross-correlation detector. The resulting effects may or may not be beneficial.

On the whole, this method has been developed to a stage where it is able to perform with some success. Further improvements will be on the small scale, involving some of the ideas presented above. It is difficult yet to determine how robust this method is, but it definitely faces a challenge in striving for robustness.

6.3.1.3 Snakes contouring

The Snakes contouring method was the least successful of the three algorithms. There was minimal convergence onto the vertebral body boundaries (Figs. 5.17 and 5.18). The resulting snakes tend to spread out their points, and there is little connectivity of the contours. When the maximum number of iterations was small, the snake did not have time to converge onto the edge information; when large, the snake continued to stretch beyond the desired boundaries. A constraint should be placed on the continuity term, E_{cont} , so that the adjacent snake point spacing is allowed to become larger than the average spacing, but not greater than a certain factor of the average spacing.

At this stage in the customization of the algorithm for our application, the results are not yet promising. Improvements to this algorithm require further customization of the curvature and image energy terms to the features of vertebral body boundaries. As well, it will probably prove useful to impose additional customized constraints on the iterations of the snake contouring. To make Snakes more functional, we may also need to consider adding other energy functionals that will attract the snake to salient features in the image (Kass, Witkin and Terzopoulos, 1987).

6.3.2 COMPARISON OF SYSTEM AND USER REQUIREMENTS

An algorithm's success depends on how well it performs its task. In addition to the end results, it is important to consider also the system and user requirements.

User interaction

All three algorithms at present require user interaction. IIB requires the selection of one centre point per vertebral body; XES requires two corners of a box to be selected; and Snakes requires that a contour be initialized by dragging the mouse. (XES could easily be implemented with a box drag, where a resizable box can be dragged by the user, to facilitate selecting opposite corners of the box.)

Runtime

Of the three algorithms we examined, Snakes has a significantly longer runtime (about 2-3 minutes per vertebral body contour versus about 1 minute for both IIB and XES) in part because it is an iterative algorithm utilizing neighbourhood computations at each snake point in each iteration. The cross-correlation detector of XES is also iterative, but it performs computations on the image clip as a whole, and not in small neighbourhoods.

Space

The disk and memory requirements of each algorithm is different. IIB and XES employ several image buffers, while Snakes stores mainly linked lists of snake nodes, and thus requires only a minimum of memory. The bigger the image size that we work with, the bigger this difference in memory requirements will become.

Robustness

Since Snakes is given an initial contour to start with, it should be more robust in terms of accurate localization of the vertebral body contour. When an algorithm is robust, it has a high chance of success under diverse conditions. Given the nature of X-ray images of the thoracolumbar spine, we are working with considerably diverse images, in terms of contrast, artifacts, pathology of the vertebrae and other considerations, and therefore would need an extremely robust algorithm to function in these conditions.

6.3.3 WHAT ELSE DO WE NEED?

In order to better test and evaluate each algorithm, the need for i) a gold standard, ii) a diseased image set, and iii) a larger image set, must be considered.

Gold standard

A gold standard is necessary for comparing results. To this end, we wrote a tool for digitally measuring lengths on the images. Using MUMC Display, the image is first displayed on the computer screen. The user then selects four corners of a vertebral body. The routine computes and displays the four side lengths, the locations of the four midpoints, the midheight, the midwidth, and the angles at which each of the six line segments are oriented, relative to horizontal and vertical axes (Fig. 6.1). An advantage of this manual method is that the user can decide whether he/she needs to perform some image enhancement on each particular image (easily done with the threshold and window features of MUMC Display) to visually aid the selection of the vertebral body corner points. Such enhancement is adjusted for the particular image, and need not be a generalized procedure. With this tool, the user can collect data satisfying his/her subjective interpretations of the vertebral boundary, and compare them with the results collected from the semi-automated routines.

Osteoporotic spine images

In our testing of the algorithms, normal spines made up the test set. While these presented many inherent problems, we need osteoporotic spine images on which to further test and evaluate the performance of each algorithm. The nature of the disease of osteoporosis means that diseased bones are likely to be less bone dense, and thus less X-ray



Instructions:
 Click corner points
 in this order
 Results:
 pixel distance/
 deg. rel. to -- or. |

- Top 232.4/-13
- Mid 235.7/-13
- Bottom 239.1/-14
- | Left 190.7/ 16
- Mid 191.8/ 15
- Right 193.1/ 14

- Top 237.2/ -5
- Mid 237.6/ -8
- Bottom 238.5/-11
- | Left 181.6/ 8
- Mid 192.6/ 8
- Right 203.7/ 8

- Top 236.0/ 0
- Mid 238.2/ -2
- Bottom 240.8/ -4
- | Left 180.4/ 4
- Mid 190.3/ 4
- Right 200.2/ 3

- Top 243.0/ 5
- Mid 243.3/ 3
- Bottom 244.0/ 0
- | Left 178.0/ -1
- Mid 187.0/ -1
- Right 196.0/ -2

- Top 242.6/ 15
- Mid 240.4/ 17
- Bottom 238.5/ 20
- | Left 181.9/-15
- Mid 172.0/-14
- Right 162.0/-13

R

MUMC BK1 Sparc 10

Vertebral Deformity Menu

Manual identify

IIB preprocess

XES x-corr/gradient

Snakes contour

Return to previous page

view a study

Change view mode->

Pan at full resolution

Clear screen

ROI operations->

Pan image

Image manipulation->

Plot a slice

3d resectioning ops->

Cine mode

Vertebral deformity->

Close to the ICON

AV- to select command

■□□- to execute command

□□□- help for this command

□□□- to previous page

Fig. 6.1 Digital manual measurement routine

attenuating, therefore yielding less contrast between the intensity of bone and the intensity of the soft tissues in an image. As well, with severe vertebral deformities, the vertebrae may be of various different shapes and angles, and no longer rectangular or properly aligned along the spine (Fig. 6.2). These characteristics of osteoporotic spines will present greater challenges to the algorithms.

Larger image set

A larger image set is necessary for testing and identifying problems, and to increase the robustness of the algorithms.

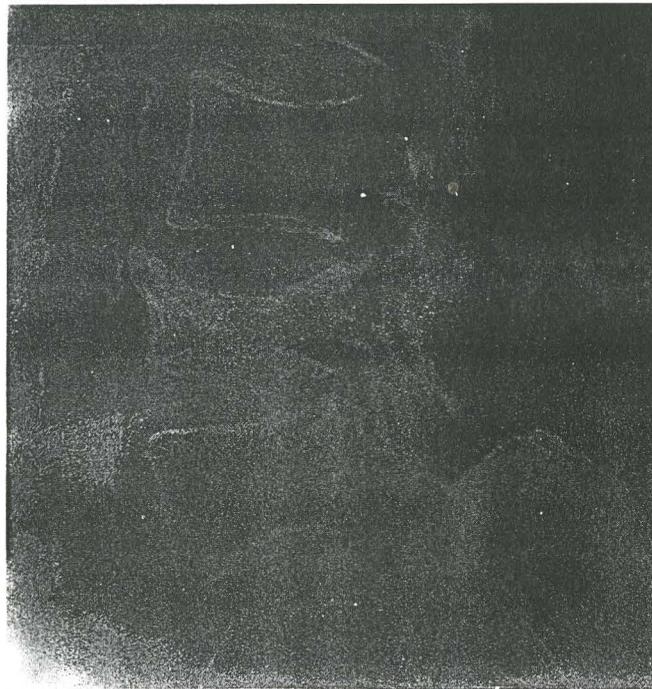


Fig. 6.2 Osteoporotic vertebrae.

6.3.4 WHICH ALGORITHM IS BEST?

The three algorithms are at various stages of development. They all show potential applicability to the task of automating vertebral body detection and measurement. It was hoped that with an accurate contour obtained using Snakes, we would have the information on the entire boundary available for any type of measurement. This information may be useful for vertebral deformity identification and classification using criteria other than ratios of heights and widths. The IIB and XES approaches are more closely related, as they use low-level decisions based on the pixel values in the image. Many of the steps in the IIB process can, and should be, incorporated into the XES method. At this point, it would be useful to continue developing each of the algorithms to the point of full functionality. Then their results may be compared more appropriately. As well, it would be useful to test the three methods on images of a spine phantom, under various ideal conditions such as having no ribs attached, or without the surrounding tissue. Such tests would provide an idea of how well the algorithms work under ideal image conditions, and would indicate on what the continued development of the algorithms should focus.

6.4 CHAPTER SUMMARY

In this chapter, we studied the performance of the three algorithms for detecting vertebral boundaries. Because of factors such as patient size, radiation risk factors, and variations in image acquisition technique and patient positioning, the contrast levels in the

various images were very different, and the number of vertebrae in the images differed as well. From the digital image processing viewpoint, the lateral spine images presented a challenge in the presence of ribs extending down from the thoracic vertebrae across the view of one or more lumbar vertebrae. The iliac region was also problematic since it is a region of high attenuation (hip bones, pelvic bones), which skews the intensity range of the image and often presents it as being 'washed out'. The high average intensity in this region lowers the visibility of the fifth lumbar vertebra, L5. High edge gradients can now belong not only to the vertebra, but also to the ilium.

The algorithms were discussed in detail and improvements suggested. System and user requirements of each algorithm were compared, and the need for gold standards, magnification correction, and a wider range of images, was described.

With all this in mind, it appears worthwhile to invest in further development of all three methods for computer-aided detection and measurement of vertebral body dimensions. When these methods become fully functional, the practicality and usefulness of each can be fully determined.

CHAPTER SIX

REFERENCES

Kass M, Witkin A, Terzopoulos D, 1988. Snakes: active contour models. *Int J Computer Vision* 1:321-331.

CHAPTER SEVEN

CONCLUSION

7.1 OVERVIEW

At present, the lack of a standardized, objective and accurate method for classifying and quantifying vertebral body deformities in osteoporotic spines indicates the need for such a procedure. A goal of this work was to develop computer automated software for the morphometric assessment of vertebral body deformities. In this thesis, we firstly considered the selection of morphometric criteria for classifying and quantifying deformity. Secondly, we explored methods for improving the quality of digital radiographic images of the spine. Finally, we developed and evaluated three algorithms for automating the measurement of vertebrae from digital Computed Radiography images.

After considering morphometric criteria developed by various other groups, we developed a set of criteria, based on ratios of anterior, mid, and posterior heights, and inferior width. For each of the three deformity types (wedge, biconcave, and compression), normal values of these dimensions for deformed vertebrae are interpolated by a polynomial fit to the data from undeformed vertebrae of the same spine. Deviations from the predicted norms are then summed to give a deformity index for each vertebral level. The sum of these indices then yields a Total Vertebral Deformity Index (TVDI) for that particular spine. Such a score may prove useful to clinicians in making recommendations for treatments which may depend upon the severity of the deformity. As well, a single deformity score like the TVDI may be useful for judging the progression of deformity between followup assessments of osteoporotic patients.

Three digital image processing algorithms were developed with the aim to automate the quantification process. The three methods produced varying results. The first method was an image enhancement and thresholding process with the goal of obtaining a bi-level image containing only the vertebral body boundary. This method was fairly effective, but suffered from its dependence on the thresholding step. The second method automated the detection of the vertebrae, given a starting vertebra, by computing cross-correlation values. Then, edge-gradient tracking traced the boundary of each detected vertebra. Results of this method seemed promising, although the cross-correlation detector failed to accurately locate vertebrae that were further away from the starting vertebra. The third method, an active contouring algorithm called Snakes,

represents the vertebral body contour by a pliable spline, which seeks to minimize an energy functional consisting of curvature, continuity, and image energy terms. This method is limited by the need for user initialization of a starting contour. Present results displayed minimal convergence onto actual vertebral boundaries. Future work in further customization of the energy terms should improve the performance of this algorithm. In addition to the characteristics of lateral spine images (such as high intensity vertebral edges) which are used to advantage by the algorithms, all three methods encountered problematic characteristics such as the presence of the high intensity ribs and ilium, bowel gas, lung shadows, and vertebral body orientation.

7.2 FUTURE DIRECTIONS

Each of the three algorithms developed in this thesis presented differences in approach and in the amount of user-interaction required. The desirability of automation, besides its presumed savings of human hours, is in achieving measurement reproducibility. User-interaction involves decisions which are often subjective in nature (for example, "Where precisely should I place this point?"), and this subjectiveness invariably leads to inter- and intra-observer variations in the measurements. However, where these variations are small enough not to interfere with clinical assessment in general, the desirability of complete automation must be weighed against practical considerations such as computational time, space, and algorithmic complexity.

Once satisfactory results are obtained from automated or semi-automated measurement of the vertebral body boundaries, then the morphometric criteria of Chapter Two may be tested. The results of assessing vertebral body deformity according to these morphometric criteria must be evaluated against assessments made by clinicians. Such evaluation is typically accomplished by ROC analysis to determine the sensitivity and specificity of measured data. An important issue here is the absence of a gold standard in assessments of the extent of vertebral deformity.

To further the scope of this project, neural network methods may be introduced in considering both the image processing routines and the development of prognosis models for predicting risk of fracture. Morphometric data from the former can be combined with qualitative data from a patient's health record to form a data vector for input into a neural network. The result of the neural network can be the present extent of vertebral deformity, as well as a prediction of risk of further deformity.

The cost of osteoporosis-related fractures to society as a whole and to individuals suffering from the debilitating effects of a crumbling skeleton is high. Not only do vertebral fractures lead to further complications involving the spinal cord, but they may also lead to death arising from the trauma of unpredicted fracture. With the ability to manipulate digital images efficiently and economically, the development of a comprehensive diagnostic system as presented in part in this thesis will prove invaluable to the osteoporosis clinical and research communities.

APPENDIX A

PROGRAM MODULES

The program modules for the algorithms were coded in C and integrated into the Sunview based software MUMC Display (see Figure 6.1). Functions named in plain *italics* were existing functions in MUMC Display.

General Display Routines

Appendx, Iwrite, Image_coords_to_display_coords, Shift_coords_back, Spine

General Mathematical Routine

Sqr

Image Intensity Based (IIB) Routines

Spine_pp0, Spine_pp1, Medianf, Modef, Conenh, Edgefilt, Linear_con, Threshold

Cross-correlation Detection and Edge Gradient Search (XES) Routines

Segment_vertebra, Segment_continued, Segment_continued2, Subimage, Walker, Mean, Multiply, Variance, X_deriv, Y_deriv

Snakes Routines

Spine_contour, Snake_0_drag, Snake_1_drag, Snake_2_drag, Snake_fini, Snake, E_cont, Adj_spacing, E_curv, E_image, Draw_Snake, Spine_fini

Manual Measurement Routines

Vert_man1, Vert_man2, Spine_comp_correct_1r, Spine_clear_point_0, Spine_comp_correct_fini, Spine_comp_correct_fini_r, Vert_man3, Vert_man_display

General Display Routines**Appendx(fname) :**

Description : Appends a blank image to the current study

Parent : Spine_pp1

Child : *Image_num_to_database*

Function : Appends a blank image to the current study and returns the name of the image file.

Iwrite(image,SIZEx,SIZEy,fname) :

Description : Writes the current image to the file

Parent : Spine_pp1

Child : N/A

Function : Takes the filename from Appendx() and writes the result image to this file.

Image_coords_to_display_coords(x, y) :

Description : Converts pixel coordinates

Parents : all screen display routines

Child : N/A

Function : Converts a pixel's coordinates from an image matrix reference to a screen display reference.

Shift_coords_back(x,y) :

Description : Offsets pixel locations for display

Parents : all screen display routines

Child : *Shift_factor*

Function : Offsets the pixel coordinates for display, depending on the image display format.

Spine() :

Description : Main control module for vertebral deformity routines

Parent : Main MUMC Display menu

Child : *Push_panel_num, Init_command_window*

Function : First level MUMC Display menu item spawning submenu of the vertebral deformity routines (see Figure 6.1 for menus).

General Mathematical Routine**sqr(a) :**

Description : Computes square of a number

Parent : Adj_spacing

Child : N/A

Function : Returns the double precision floating point result of squaring a double precision floating point number.

(i) Image Intensity Based (IIB) Routines**Spine_pp0() :**

Description : Main control module for IIB method

Parent : Spine

Child : Spine_pp1

Function : When selected, assigns a left mouse button click event to Spine_pp1.

Spine_pp1(pos,x,y) :

Description : IIB algorithm

Parent : Spine_pp0

Children : Medianf, Linear_con, Conenh, Edgefilt, Threshold, Modef, Appendx, Iwrite

Function : Given the point selected by the user to indicate the centre of a vertebra, a 400x400 image clip, centred on this point, is formed. The clip is processed according to the IIB method, and the resulting image is stored and appended to the current study. The left mouse button event is reset to nothing.

Medianf(image,SIZEx,SIZEy)

Description : Median filtering to reduce noise in image

Parent : Spine_pp1

Child : N/A

Function : A 5x5 median filter is passed through the image clip matrix, and the result (a matrix of median values within a pixel's 5x5 neighbourhood) replaces the original image clip matrix.

Modef(image,SIZEx,SIZEy) :

Description : Mode filtering to remove speckles in image

Parent : Spine_pp1

Child : N/A

Function : Given the image clip matrix, the mode pixel value of a pixel's 5x5 neighbourhood replaces that pixel's value.

Conenh(image,SIZEx,SIZEy) :

Description : Piecewise linear contrast enhancement to enhance bone over soft tissue

Parent : Spine_pp1

Child : N/A

Function : Piecewise linear contrast enhancement with grey level breakpoints at 400 and 700 is performed on the image clip.

Edgefilt(image,SIZEx,SIZEy) :

Description : Edge enhancement

Parent : Spine_pp1

Child : N/A

Function : The 3x3 Sobel edge masks are convolved throughout the image clip.

Source : S.K. Yu, 1994.

Linear_con(image,SIZEx,SIZEy) :

Description : Linear histogram stretching to full dynamic range

Parent : Spine_pp1

Child : N/A

Function : The maximum and minimum pixel values of the current image clip are found, and this pixel value range is linearly transformed to full range (1024 levels).

Threshold(image, SIZEx, SIZEy) :

Description : Bi-level thresholding of the image

Parent : Spine_pp1

Child : N/A

Function : The average pixel value of the image clip is found, and the image clip is thresholded at twice this average value, so that pixel values below this threshold are set to black (0), and pixel values greater than or equal to this threshold are set to white (1023).

(ii) Cross-Correlation Detection and Edge Gradient Search (XES) Routines

(developed by J.S. Sloka and modified by L. Tan)

Segment_vertebra() :

Description : Main control module for XES algorithm

Parent : Spine

Child : Segment_continued

Function : Assigns left mouse button click to store coordinates of a user-selected point.

Segment_continued(pos,x,y) :

Description : User point initialization

Parent : Segment_vertebra

Child : Segment_continued2

Function : Upon left mouse button click, the coordinates of the first user-selected point are stored. This point is the upper left corner of a rectangular region.

Segment_continued2(pos,x,y) :

Description : User point initialization

Parent : Segment_continued

Child : Subimage

Function : Upon left mouse button click, the coordinates of the second user-selected point are stored. This point is the lower right corner of a rectangular region.

Subimage() :

Description : Cross-correlation vertebra detector

Parent : Segment_continued

Child : Walker

Function : User-initialized box for first vertebra is used to locate the second vertebra according to the highest correspondence by cross-correlation between the two image matrices. Each newly detected vertebra is used similarly to detect the next vertebra.

Walker(sub_row, sub_col, row, col, nrows, ncols, buffer) :

Description : Vertebral edge and corner detection

Parent : Subimage

Child : Linear_con, X_deriv, Y_deriv

Function : Tracks the superior and inferior vertebral body edges using edge gradients, and detects corners using angle criteria.

Mean(buffer,nrows,ncols) :

Description : Calculate the mean pixel value in an image matrix

Parent : Subimage

Child : N/A

Function : Returns the mean value in a matrix.

Multiply(buffer1,buffer2,prodbuffer,ncols) :

Description : Multiply two matrices (element by element)

Parent : Subimage

Child : N/A

Function : Given two matrices of the same dimensions, multiply each pair of corresponding elements and store result in a third matrix.

Variance(buffer,ncols) :

Description : Calculate the variance of the values in a matrix

Parent : Subimage

Child : N/A

Function : Returns the variance of the values in a matrix.

X_deriv(inbuffer,outbuffer,n,m) :

Description : Horizontal edge filtering

Parent : Walker

Child : N/A

Function : Returns a matrix of the horizontally edge filtered image region.

Y_deriv(inbuffer,outbuffer,n,m) :

Description : Vertical edge filtering

Parent : Walker

Child : N/A

Function : Returns a matrix of the vertically edge filtered image region.

(iii) Snakes Routines

(based on routines developed by W.N. Street, 1993)

Spine_contour() :

Description : Main control module for the Snakes algorithm

Parent : Spine

Children : Snake_0_drag, Snake_1_drag, Snake_2_drag

Function : Allocates memory for the first Snake node and initializes the left mouse button click, drag, and release events to the child functions.

Snake_0_drag(pos,x,y) :

Description : Contour initialization (1st node)

Parent : Spine_contour

Child : N/A

Function : Store first Snake point upon left mouse button click.

Snake_1_drag(x,y) :

Description : Contour initialization

Parent : Spine_contour

Child : N/A

Function : Each point picked up by mouse dragging is displayed on the screen, allocated and stored in a linked list data structure containing each point's x coordinate, y coordinate, pixel value, and memory address of the next point (when next point is registered). With each new point, a variable counting the number of Snake points is incremented.

Snake_2_drag(x,y) :

Description : Contour initialization

Parent : Spine_contour

Children : Snake_fini, Snake

Function : On release of the left mouse button, this function does some cleaning up of the user-initialized Snake contour. It links the last node to the first node, thereby closing the contour's linked list. Then it creates a matrix to check for repeated points, which are then deleted from the Snake linked list.

Snake_fini() :

Description : Finish up Snakes contour initialization

Parent : Snake_2_drag

Child : N/A

Function : Resets mouse button event functions when the user has finished initializing the contour.

Snake() :

Description : Main Snakes energy minimization module

Parent : Snake_2_drag

Children : E_cont, E_curv, E_image, Adj_spacing, Draw_Snake, Spine_fini

Function : Performs the Snakes energy functional minimization according to the algorithm described in Chapter 4, Section 4.5.3.2.

E_cont(current, previous, average_spacing, cont_array) :

Description : Computes continuity energy for the neighbourhood points

Parent : Snake

Child : Adj_spacing

Function : Stores in a neighbourhood matrix called cont_array the difference between the distance from the current Snake node to a candidate in its neighbourhood and the average distance between all adjacent points in the Snake.

Adj_spacing(xx1, yy1, xx2, yy2) :

Description : Calculates the Euclidean distance between two points

Parents : Snake, E_cont

Child : N/A

Function : Returns a floating point value for the Euclidean distance between any two points.

E_curv(current, previous, next, centre, curv_array) :

Description : Calculates the curvature energy for the neighbourhood points

Parent : Snake

Child : Adj_spacing

Function : Calculates and stores in a matrix called curv_array the values of curvature energy for each point in the neighbourhood of the current Snake point.

E_image(curr, centre, image_forces_array) :

Description : Calculates the edge energy for the neighbourhood points

Parent : Snake

Child : N/A

Function : Calculates and stores in a matrix called image_forces_array the values of the edge gradients at each neighbourhood point.

Draw_Snake() :

Description : Displays result of Snakes

Parent : Snake

Children : Shift_coords_back, Image_coords_to_display_coords

Function : Puts a coloured dot on the screen for every point in the converged Snake contour.

Spine_fini() :

Description : Finish up Snakes

Parent : Snake

Child : N/A

Function : Frees all Snake node memory locations, reinitializes global Snake variables, and resets mouse button functions to nothing.

(iv) Manual Measurement Routines (see Figure 6.1)**Vert_man1() :**

Description : Main control module for manual measurement routine

Parent : Spine

Children : Vert_man2, Spine_comp_correct_1r, Spine_clear_point_0,
Spine_comp_correct_fini, Spine_comp_correct_fini_r

Function : Initializes the mouse button events, initializes global variables, and displays the diagram of instructions on the selection order of the points.

Vert_man2(pos,x,y) :

Description : Collects user-selected points

Parent : Vert_man1

Child : N/A

Function : On left mouse button click, stores the coordinates of a user-selected point, displays the point on the screen, and increments a point counter.

Spine_comp_correct_1r(x,y) :

Description : Warns the user that a maximum number of points has been selected

Parent : Vert_man1

Child : N/A

Function : On left mouse button release, checks the point counter for maximum limit, and alerts user if the maximum number of points has been reached.

Spine_clear_point_0(pos,x,y) :

Description : Erase previous point

Parent : Vert_man1

Child : N/A

Function : On middle mouse button click, the last point in the point buffer is removed from the screen, and the point counter decremented.

Spine_comp_correct_fini(pos,x,y) :

Description : Finish selecting points for measurement

Parent : Vert_man1

Child : Vert_man3

Function : On right mouse button click, the buffer of stored points is converted from the display coordinate reference to the image matrix coordinate reference and the coordinates written out to a text file. The global variables are reinitialized.

Spine_comp_correct_fini_r(x,y) :

Description : Warns the user that the minimum number of points has not been selected

Parent : Vert_man1

Child : N/A

Function : On right mouse button release, check for enough points in the buffer. If not enough, alert user to continue selecting points.

Vert_man3() :

Description : Calculates vertebral dimensions from user-selected points

Parent : Spine_comp_correct_fini

Child : Adj_spacing, Vert_man_display

Function : Opens file of user-selected point coordinates and calculates lengths, widths, midpoints, and angles of the line segments connecting the points.

Vert_man_display(C1,C2,C3,C4,LCy,Top,Left,Bottom,Right,Midw,Midh,Topmp, Bottommp,Leftmp,Rightmp,Topor,Leftor,Bottomor,Rightor,Midwor,Midhor) :

Description : Displays measurement results

Parent : Vert_man3

Child : N/A

Function : Takes the calculated results from the parent module and displays the line segments, midpoints, and numerical results on the screen.