SEGMENTATION IN 3D VIRTUAL SPINE MODELING FOR ASSISTANCE IN SURGICAL PLANNING AND GUIDANCE

By

MATTHEW JAMES WILLIAMS

A THESIS PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF ENGINEERING

UNIVERSITY OF FLORIDA

2005

ACKNOWLEDGMENTS

There are many people that I am pleased to recognize who have had a very positive influence on my journey at the University of Florida. First, I would like to thank Dr. Frank Bova for his guidance and patience, and for making possible a very clinically interactive biomedical engineering experience. It is the kind of project I had hoped for when I began my studies at UF. Dr. Wesley Bolch and Dr. David Gilland both have my sincere appreciation for their commitment to come on board at such a late stage to help out. Literally, I could not have finished without them. Also, I would like to extend gratitude to Dr. Bernard Guiot, for introducing me to neurosurgery; Dr. Lionel Bouchet, for his tutelage and example of tireless dedication; Drs. Yunmei Chen and Bernard Mair, for lending their ears for mathematical reflection; and Drs. Christopher Batich and Anthony Brennan, for their continued support and words of motivation. In addition, I would like to acknowledge the financial contributions of the Department of Neurosurgery, the Department of Biomedical Engineering, and the Whitaker Foundation.

I have encountered and have had the pleasure of knowing many friendly personalities in the lab, in the operating room, in the classroom, and in various other nooks and crannies, and each has helped to make this a memorable experience; and for that I give thanks. Finally, I would like to extend a heartfelt appreciation to Alan Wineman, Jennifer Kadlowec, and Jeanette Clute for their belief in my abilities; my fraternity brothers, for their companionship; and my exceptional friends Matt, JP, Vishal, Christina, Chris, Tom, and Kent for their encouragement and distracting adventures. And

most of all, I am indebted to my brother, my parents, and the rest of my loving family as they are the foundation of my strength; and to Jennifer, the delightful blend of sunshine and spice that has helped me to believe in all of the wonderful things in life.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF OBJECTS

Abstract of Thesis Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Master of Engineering

SEGMENTATION IN 3D VIRTUAL SPINE MODELING FOR ASSISTANCE IN SURGICAL PLANNING AND GUIDANCE

By

Matthew James Williams

May 2005

Chair: Frank J. Bova
Major Department: Biomedical Engineering

Surgical procedures of the spine are technically demanding, requiring precise navigation to avoid critical vascular and nerve structures. Often, image guidance systems are employed to improve accuracy of surgical tool placement and increase the likelihood of a successful outcome. The most commonly used systems generate a patient-specific virtual 3D model. This model is used to create a surgical plan that, during the operation, safely guides the surgeon in placement of the surgical tools. Unfortunately, the current modeling technique limits the application of image-guided surgery.

The typical surgical model is created using the thresholding technique performed on a pre-operatively acquired diagnostic CT image. The result is a single rigid body model that appears as a surface-rendered outline of the involved bones. During the procedure, this model is registered or aligned with the patient. The drawback to the rigid model is its inability to account for the mechanical flexibility of the spine. Since the patient is likely to change positions between image acquisition and placement for

surgery, errors are introduced into the guidance. In order for the surgical model to properly deform, the virtual bones must be segmented and individually registered to the vertebrae. The current method of segmentation is very time-consuming and not clinically feasible on a routine basis. The unfortunate consequence is that, given accuracy restrictions, image-guided procedures of the spine are generally limited to one or two vertebral levels. It is the purpose of this research to develop a method that provides a segmented model for use in image-guided surgery. It is desired that this method be easy to work with, provide swift results, and require minimal user intervention.

The segmentation method developed here utilizes a simple region-growing scheme. It dilates seed regions that are uniquely assigned to the underlying bones. This allows the growing regions to retain their identity while expanding into the full bone profile. The method was tested on 31 high-resolution clinical scans. Over 80% of the resulting segmentations showed significant improvement over the standard manual methods. Success was gauged by the amount of time and effort required to achieve a segmentation using the method as compared to manual methods. The factors contributing to segmentation failures are attributed to poor bone resolution and inadequate starting information. This method showed marked success in the segmentation of the CT-derived models used in guided spine surgery. This method may be employed to significantly reduce segmentation time and facilitate surgical image-guided applications to multi-level spinal procedures.

CHAPTER 1
IMAGE-GUIDED SPINE SURGERY

Surgical procedures of the spine are technically demanding. They often involve placement of rigid instrumentation in close proximity to critical vascular and nerve structures and accordingly, demand a high degree of accuracy. In order to avoid complications, the neurosurgeon or orthopedic surgeon must accurately plan the trajectories for placement of this instrumentation. This requires visualization of the involved anatomy in three dimensions. Physical inspection allows the surgeon to make an initial assumption of specific anatomical orientation. However, imaging systems verify this assumption without unnecessary surgical exposure. In addition, they provide crucial information such as condition of vasculature or other hidden structures or extent of pathology such as fracture dimensions and tumor location.

Image-guided surgery, at its simplest, is the use of an imaging modality to facilitate surgical intervention. Systems dedicated to image guidance, though, actually allow non-invasive evaluation and planning through a virtual surgical environment. In addition, during surgery, these systems direct the surgeon's tools to the proper trajectories. The use of image guidance systems in surgery, through increased accuracy and precision, has reduced the risk of damage to critical areas and increased the ability of the surgeon to handle technically difficult surgeries. In addition to a reduction in the duration of the operation, these advancements have led to significant reductions in patient morbidity and mortality (Cleary, 1999).

The most common image guidance systems utilize a model of the vertebrae generated from a diagnostic Computed Tomography (CT) image. It appears as a surface-rendered image of the bone profile. Unfortunately, this model is not able to deform, or account for any of the natural motions that may occur with the spine. Since motion of the patient is likely to occur after image acquisition, the model of the vertebrae constructed may not adequately reflect the orientation of the vertebrae presented during surgery. This limits the accuracy of the model and its utility in guiding surgical procedures. In order to create a spine model that can properly reorient or deform each of the vertebrae must be isolated. The process of identifying these individual vertebrae in the model is called segmentation. The current method of segmenting spinal models is a labor-intensive, time-consuming procedure, which is infeasible on a routine clinical basis. The purpose of this research is to create a more user-friendly method for segmenting the models intended for image-guided surgery. The goal is to provide results in a more suitable time frame while requiring minimal user intervention.

Preoperative Image Guidance

The standard of care is the use of a presurgical scan for image-guided surgery. Typically, plain radiographs are supplemented with an additional modality such as myelography, computed tomography, or magnetic resonance imaging (MRI). Each modality provides another layer of information that the surgeon can utilize. However, the use of both CT and MRI is discouraged because of cost constraints (Cleary, 1999). The preoperative scans that are at the surgeon's disposal are generally used to mentally construct a 3D representation of the anatomy. This envisioned model is associated with the anatomy exposed during the operation. It allows the surgeon to proceed with more confidence without relying on any anatomical statistical convention. This is especially

true in cases of severe pathologic deformation. Overall, the greater the knowledge of anatomical positioning, the less invasive a procedure must be to assess the correct orientation of the involved elements (Lavallée et al., 1996). The accurate determination of vertebral alignment is essential for a successful surgical outcome. Of the modalities available for presurgical imagery, CT and MRI are typically used for spinal applications. They provide multiple 2D slices that can be mentally stacked to form a 3D model. CT has excellent bone and soft tissue contrast, yet requires ionizing radiation (Cleary, 1999). MRI, on the other hand, provides soft tissue contrast without the ionizing radiation and bone flare usually associated with CT. They are both susceptible to artifacts, such as non-uniformity distortions in MRI and starburst patterns in CT.

The primary drawback of preoperative imaging used for surgery is that it does not account for any anatomical alignment changes that occur after imaging. Any changes of positioning that occur detract from the accuracy of the representation. This change of positioning can be isolated from two sources. One is the movement of the patient between the presurgical scan and the operative positioning. The presurgical scan is taken in the supine position to minimize breathing artifacts. A process called "reduction" whereby the surgeon optimally orients the patient for the procedure determines the operative position. The second source of position change occurs during the operation. Throughout the procedure the vertebrae can experience significant surgeon-induced motion, especially in the case of spinal instability (Glossop & Hu, 1997). Patients with unstable spines may exhibit abrupt translations of spinal segments, and these movements are difficult to model and predict (Cleary, 1999). Accounting for these alignment changes is especially critical in cases where two vertebral bodies are to be fixed by the

Table 1-1. Change of orientation between the vertebrae during a surgical procedure. Presurgical and postsurgical CT scans were compared, and the change of the positions of C1 relative to C2 was recorded. The translation and rotation between the two vertebrae in orthogonal planes are shown for three patients.

	Patient 1			Patient 2			Patient 3		
	AP	Lat.	SI	AP	Lat.	SI	AP	Lat.	SI
Trans. (mm)	2.7	-6.1	3.7	-9.3	3.8	7.4	-6.7	0.7	10.0
Rot. (degree)	3.0	-5.4	-6.5	0.8	-14.0	9.7	-1.9	-18.0	2.4

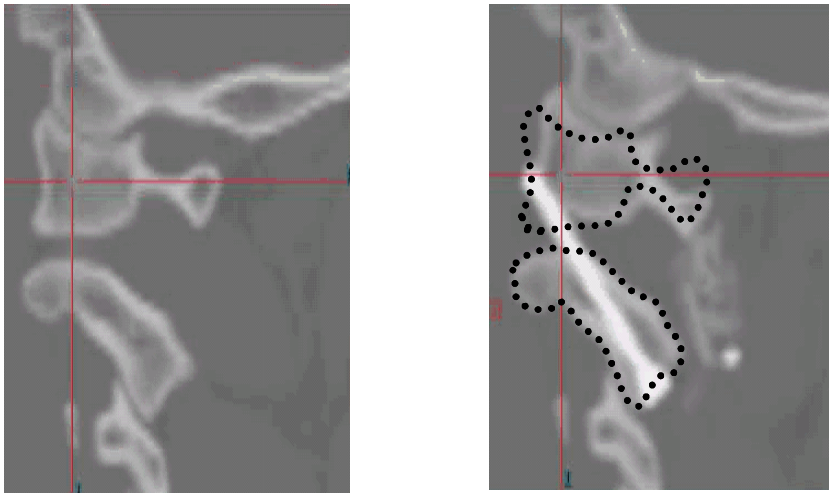

Figure 1-1. Change in anatomical orientation occurring during transarticular screw placement. A profile of vertebral positioning is taken from the presurgical CT scan (left) and overlaid onto the postsurgical CT scan (right). A shift in position can be seen between C1 and C2.

same screw (Figure 1-1). Table 1-1 quantifies the changes of position that occur after imaging by comparing preoperative and postoperative CT scans. Even though this particular example only involves two vertebral levels, it is apparent that the use of pre-operative imaging for image guidance has a limited accuracy. There are, however, two possible solutions for managing the problem of spinal motion after image acquisition. The first is intraoperative imaging and the second is the use of a guidance system that tracks each individual vertebral body.

Intraoperative Imaging

Imaging systems can be used to acquire scans during the operation and provide an instant indication of anatomical positioning and surgical tool location. Despite this advantage, however, there are certain considerations such as cost, intraoperative image quality, and impact on the operating room environment that must be considered before the investment in such a system.

Fluoroscopy

Fluoroscopy is the dominant intraprocedural imaging modality for spinal surgery. The unit is a common piece of equipment that easily fits around the table and provides minimum disruption to the operative field or operative technique. It is a low-cost, useful and familiar technology. It is also easily accessible and portable (Foley et al., 2001; Cleary, 1999). In the application of spinal surgery, fluoroscopy is commonly used to verify vertebral alignment and instrumentation and tool positioning. Fluoroscopy units provide high resolution, large field of view (FOV) scans in real time that have excellent bone to soft tissue contrast (Cleary, 1999). There are some drawbacks however, that limit the utility of x-ray fluoroscopy. The primary one is fluoroscopy only acquires a 2D image in one plane, and hence only gives one aspect of 3D positioning. It is impractical to gauge depth in the picture, as it is a display of overlapping tissues. In order to image in additional planes, the fluoroscope must be repositioned repeatedly (Foley et al., 2001). Also, fluoroscopic images have poor soft tissue discrimination, which makes the visualization of vascular and nerve structures difficult (Cleary, 1999). Finally, the capability of imaging in real time is at the expense of constant x-ray radiation exposure. In the case of spine surgery, where there may be long periods of imaging activity, there is

significant radiation exposure to the patient and operating team (Foley et al., 2001; Rampersaud et al., 2000; Cleary, 1999).

Computed Tomography

Computed tomography is an imaging modality in which a three dimensional image is constructed from a series of plane cross-sectional radiographic images made along an axis. The scans are high resolution and have excellent bone to soft tissue contrast. Intraoperatively, CT provides an excellent view of bony structures and their relationships, and can accurately localize the tip of interventional instruments (Cleary, 1999). However, the imaging is not real time and subjects the patient and those in the scanning field to ionizing radiation.

There are three main types of CT scanning systems that are used during surgery. The first is spiral CT. It is very fast with a large bore and excellent image quality, yet requires high capital and maintenance costs (Cleary, 1999). In addition, it is a large, fixed machine that has limited accessibility in a surgical environment. Mobile CT, on the other hand, is smaller, portable, and has a comparatively lower radiation dose. The disadvantages to this imaging system are slower acquisitions, decreased tube capacity, and lower image quality, which can lead to registration difficulties (Cleary, 1999). Mobile CT is costly as well. The last type of CT imaging systems uses a sweeping fluoroscope to generate images. This fluoro-CT has the advantage of quick reconstruction and display for 3D imaging as well as easy patient access and targeting. It offers a low cost, low patient dose alternative to spiral CT or mobile CT. However, fluoro-CT has comparatively poor tissue contrast which results in minimally acceptable bone reconstruction. Two examples of Fluoro-CT systems are the SIREMOBIL Iso-C^{3D} (Siemens Medical Solutions USA, Inc., Iselin, NJ) and FluoroCAT™ (G.E. Healthcare

Technologies, Waukesha, WI). Given these options, there is still one primary limiting factor in the use of intraoperative CT for spinal alignment determination. Scans from the typical surgical position are subject to breathing artifacts, and the resulting images are unsuitable for accurate image guidance.

Magnetic Resonance Imaging

Magnetic resonance imaging is a non-ionizing imaging modality that can be employed intraoperatively. The images are high-resolution with excellent soft tissue contrast. Safety considerations with the magnetic fringe field however, make it difficult to adapt to the surgical environment. The need for specialized tools and equipment that can operate in the intense magnetic field and radio frequency-rich environment make it a costly endeavor. This is in addition to the intrinsic cost of the device. Also, specific to guidance, there is insufficient tool tip viewing accuracy. MRI images have poor definition of bony structures, making edge interpretation difficult. The correct determination of bone boundaries is critical in spinal surgery. Finally, MRI has its own handful of potential artifacts that must be considered (Cleary, 1999). An example of an MR system optimized for surgical application is the Achieva I/T Interventional MR (Philips Medical Systems, Bothell, WA). Two other examples that allow intraoperative use are the Signa SP (G.E. Healthcare Technologies, Waukesha, WI) and the MAGNETOM Concerto (Siemens Medical Solutions USA, Inc., Iselin, NJ). Similar to computed tomography, this imaging modality is also susceptible to the breathing artifacts that plague scans performed in standard operative position.

Ultrasound

Ultrasound is an inexpensive, easily portable, real-time imaging modality that does not rely on ionizing radiation. However, the poor image quality, weak discrimination of

certain critical spinal tissues, and reliance on operator skill, does not make it a candidate for precise surgical visualization. Also, the use of ultrasound for surgery often requires increased intrasurgical access, which should be avoided if at all possible (Cleary, 1999). This modality is mentioned because it has utility in being used for image to patient registration, a process necessary for computer-assisted surgery.

Computer-assisted Surgery Systems

The problem of realizing spinal motion after pre-operative image acquisition is a primary concern for accurate surgical intervention. In response, computer-assisted surgical guidance systems have been developed which actively track spinal movement and relay real-time positioning of the vertebrae to the surgeon. Unlike the intraoperative imaging systems, which typically operate over discrete periods of time with some degree of surgical field interruption, these systems relay positioning information throughout the critical portion of the surgical intervention. Image-guided surgical systems accomplish their task by overlaying a virtual surgical field with the patient's operative field. The computer-generated virtual field is created preoperatively and includes a geometric model of the involved anatomy with visual indications of intended surgical pathways and targets of instrumentation. Once this virtual field is matched or registered intraoperatively to the physical surgical field, the surgeon can compare his surgical tool placement with the plan. The ability of the model to properly represent the spinal anatomy has a direct effect on the accuracy of the surgical guidance.

Virtual Surgical Model

The most preferred anatomical model used for spine surgery is one that best approximates the mechanical flexibility of the spine. The spine should be modeled as a series of rigid non-deformable bodies, vertebrae, connected by deformable tissue, muscle,

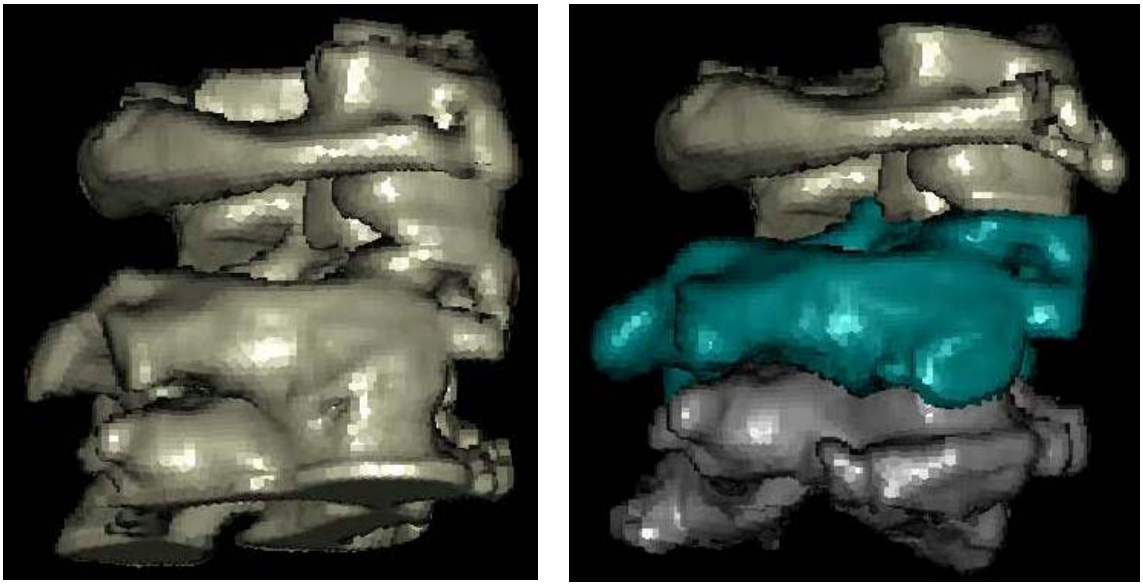

Figure 1-2. Segmentation of a surgical model. A rigid model of the cervical vertebrae is created using the thresholding technique (left). An applied segmentation method produces a model that properly represents the anatomy (right).

cartilage, and ligaments. The model used for guided surgery, however, only needs to include each individual vertebra as isolated bodies, so that they can slide, or translate and rotate relative to each other. Such a model best reflects the motion seen during the operation.

The computer-based anatomical model is typically created from a presurgically acquired CT image. As indicated by Peters (2000), a high-resolution, three-dimensional image is best suited to clearly represent the patient's anatomy. For example, a CT scan intended for use in spine surgery has a .6 to 1.2 mm slice thickness with a 0.2 – 0.3 mm in plane pixel size. These voxel dimensions (0.3 x 0.3 x 0.3 mm) equate to an approximate 18 cm field of view with 512 x 512 pixel images.

Currently, thresholding is the scheme most commonly used to create the anatomic virtual model used in spine surgery. It is a convenient procedure that takes little of the clinician's time. Thresholding is a process that creates a boundary at disjoint intensity ranges such as between bone and soft tissue on CT images. It is easy to indicate the

proper bone boundary using this method, however, it often leaves adjacent bones as one contiguous mass (Kikinis et al., 1998). The resulting single rigid body model, once registered to the patient, is ill suited to account for any movement of the spine that may occur during the surgical procedure, which potentially introduces errors into the surgical guidance. This challenges the efficacy of the guidance system and presents the need for a segmented model. The single rigid body model created through thresholding can be converted into a model that allows for flexibility. This process is called segmentation, or the division of the single spinal bone model into its constituent vertebrae (Figure 1-2). The resulting composite model, if properly matched to the patient, will allow the orientation of the vertebrae to be tracked intraoperatively. Currently, segmentation is a time-consuming process in which the clinician must use basic processing tools to manually contour the boundaries of each vertebra; there is to date no automatic, reliable, and robust method of doing so. Therefore, the development of real-time image processing techniques for model creation is of primary importance for image-guided surgery (Cleary, 1999).

Surgical Planning

Once an adequate virtual model is created, it can be used to aid treatment selection (Cleary, 1999). The surgeon will use the model to visualize inside structures, and establish optimal surgical pathways. Preparation involves the placement of virtual tools, which represent the actual instruments used in surgery. The surgeon can plan the incision, define resection margins, and determine the appropriate orientation for instrumentation (Welch et al., 1997). This includes identifying working corridors that provide adequate access while minimizing the risk of damage to fragile tissues (Cleary, 1999). An example of tool placement is shown in Figure 1-3. After the planning phase

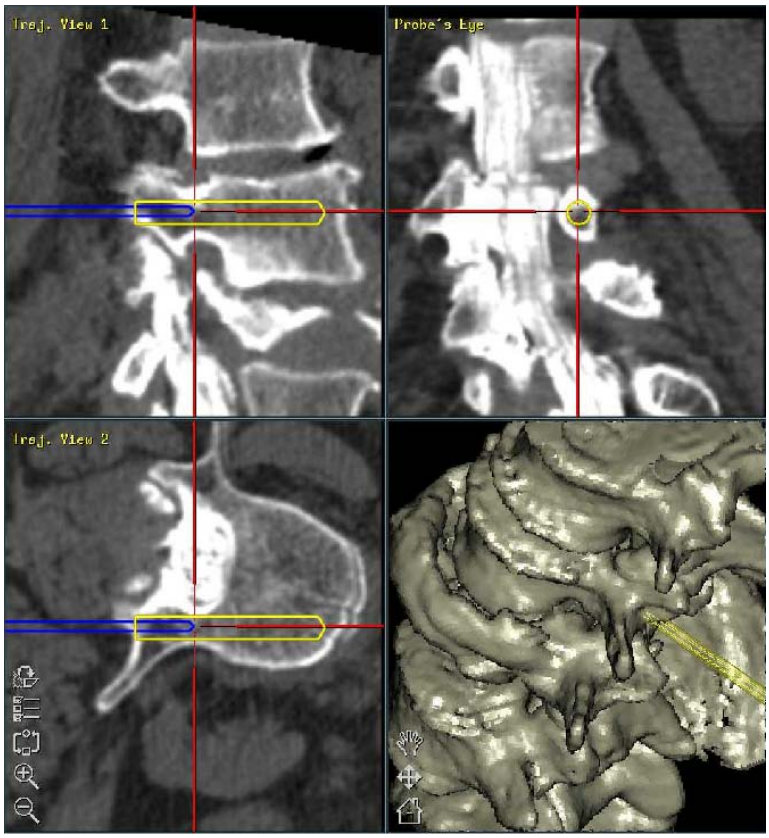

Figure 1-3. Surgical plan. Pedicle screw trajectory is planned using a virtual tool. Note the surgical model in the lower-right corner.

of surgical treatment, the virtual model must be matched, or registered to the patient to allow for guidance.

Registration and Tracking

Registration is the mapping of coordinates between any two spaces. In a guided surgery application, it is the process that links the coordinate frame of the virtual computer field to the surgical field. Registration is done so the displayed computer model of the involved anatomy and overlaid virtual surgical tools accurately represents the physical operation. A popular system at the University of Florida accomplishes registration using a stereotactic, optical camera system. It tracks precise positioning of the anatomy and surgical tools in three-dimensional space using reference frames. These frames have attached light emitting diodes (LEDs) that allow the infrared (IR) camera to

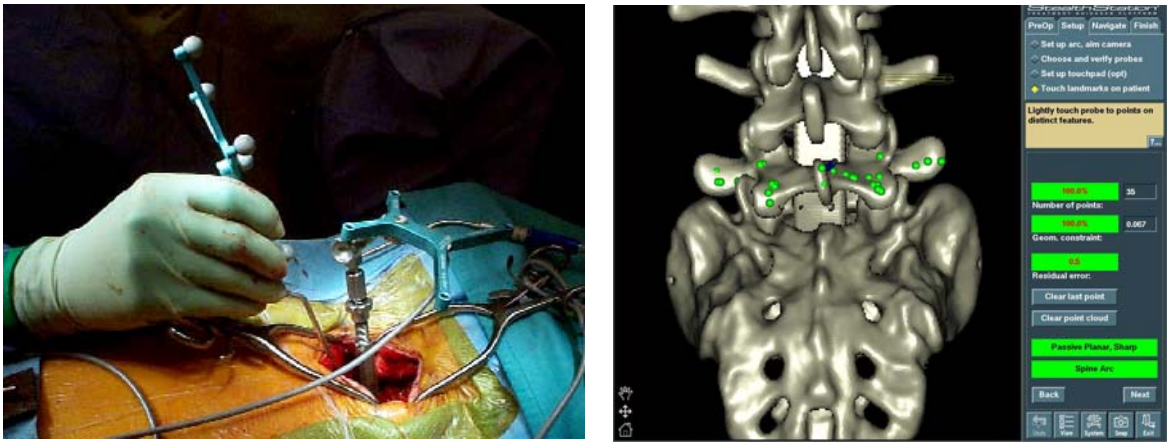

Figure 1-4. Registration of the surgical model to the patient. The reference frames attached to the probe and vertebra are used by the IR camera system to track stereotactic positioning. The surgeon uses a probe to touch various points on an exposed vertebra for surface registration (left). The corresponding points are reflected on the virtual model (right).

track their position and pose. There are several steps taken to perform registration for this system. First, a reference frame is attached to an exposed spinal process and to each of the tools involved in the guided surgery. Since the thresholding process creates a single rigid-body model, only one vertebra can be registered. Therefore, for best utilization, the model is registered to the most significant or important vertebral level. Next, the computer-built virtual anatomic model must be registered or geometrically associated to the patient's selected vertebra. Registration is required because the computer is unaware of the actual position of the anatomy relative to the attached reference frame. Registration for the tracked tools is much simpler because of predefined attachment sites for the reference arrays.

The primary registration method for the spinal system is point matching. A set of points, usually landmarks, are identified on the virtual vertebrae and then matched to corresponding points on the patient's vertebrae using a tracked probe. Once point matching is completed, and an estimate of the association between the surgical model and exposed vertebrae is achieved, an additional registration step, surface matching, may be

employed. In order to surface match, the surgeon touches a multitude of points on the surface of the exposed vertebrae (Figure 1-4). These points are used to create a surface profile, which is then aligned to the surface of the virtual model. This step may improve the accuracy of the registration. Once registration is complete, whichever vertebra the reference frame is attached to is tracked dynamically to account for changes in the patient's vertebral orientation. Registration assures that the virtual surgical tools, which correspond to the real tools, are properly overlaid upon the spine model as the operation is performed.

Registration is usually limited to one vertebral level. This is due to both current modeling restrictions and the ability of the tracking system to easily recognize multiple rigid bodies. Therefore, to preserve accuracy, the use of surgical navigation is generally reserved for those procedures that only involve a couple of vertebral levels. An increase in availability of segmented models, however, will encourage registration and tracking of multiple vertebral levels. This will allow the benefits of surgical guidance to be applied to a broader scope of procedures.

Alternate techniques exist for tracking and registration of the spine. One noteworthy tracking system uses electromagnetic (EM) markers to locate anatomy. These point localizers minimize signal blockage commonly associated with the cumbersome IR reference frames, but may be susceptible to EM interference and the presence of magnetic or ferrous objects. Registration, alternatively, can be accomplished with a localized imaging system. Specifically, ultrasound can be used to create a surface profile for use in the surface-matching method of registration. One novel development presented by Medtronic is the FluoroMerge® software (Medtronic Surgical Navigation

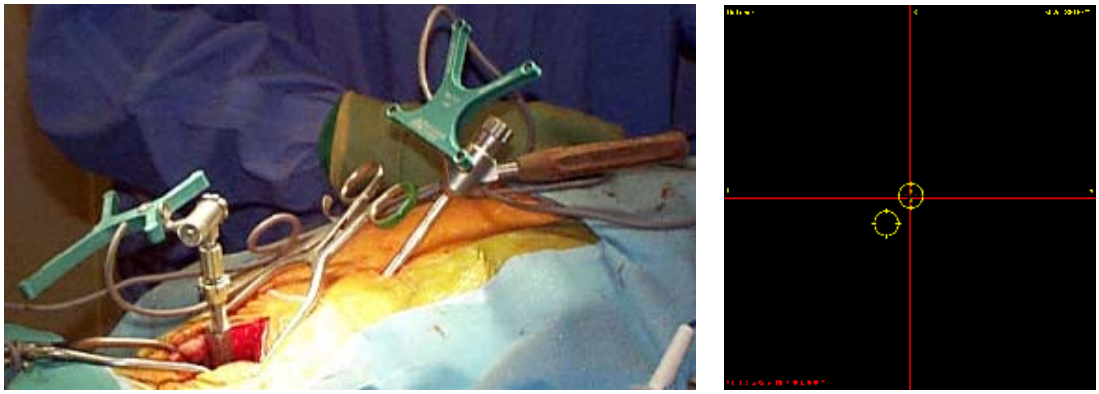

Figure 1-5. Active surgical guidance. A targeting screen (right) guides the surgeon to the proper trajectory of the surgical drill guide (left).

Technologies, Louisville, CO). It can perform the registration of a preoperative CT image using only two fluoroscopic images. This not only provides a quick, "hands-free" registration method, but may also be useful for multiple rigid body registration.

Surgical Guidance

Once all of the steps of preparation (model creation, planning, registration) have been completed, the system is ready to provide the surgeon with the feedback necessary to perform a successful operation. A visual terminal in the operating room allows the surgeon to simultaneously view surgical tool placement relative to both the virtual model and the patient. Also included is a targeting screen that compares the trajectories of tools relative to the surgical plan (Figure 1-5). This feedback assures proper placement of the instrumentation during the surgical procedure.

Surgical Application

There are several benefits to the use of image guidance in surgery. First, it provides a multidimensional view of anatomic relationships in the operative field, including extent of bone and soft tissue resection (Welch et al., 1997). It is especially useful when traditional surgical landmarks are obscured or altered, as may be the case with pathology or bony fusion (Austin et al., 2002). Image guidance reduces the need for

exposure and increases the confidence of the surgeon during the procedure (Welch et al., 1997). It also improves accuracy of the surgical tool placement, which leads to the reduced risk of structurally significant violations such as neurovascular injury. Ultimately, image guidance lessens the technical difficulty of spine surgery and allows for more safe and effective outcomes (Assaker et al., 2001; Ohmori et al., 2001; Youkilis et al., 2001; Welch et al., 1997).

Currently, there are three general classifications of spinal surgery that take advantage of image-guidance. The first is decompression, or removal of anatomical pressure on the spinal cord. This category has the highest volume of cases. Stabilization is the category with the next highest volume of procedures. Surgical intervention is almost always required when instability occurs below the level of C2. The final group is deformity correction, which has the highest risk of undesirable outcomes (Cleary, 1999). The hard tissues manipulated in image-guided surgical procedures often include the vertebral body, facet joints, iliosacral joint, and intervertebral disc. The pathology addressed during the procedure varies from fracture to inflammation and may involve the spinal cord and other adjacent soft tissues (Cleary, 1999). The use of computer-assisted surgery has especially seen use in cervical procedures, given the need for a high level of accuracy. Vaccaro & Singh (2001) discusses various applications in cervical spine surgery. He states the primary use for computer-assisted surgery is placement of C2-C1 transarticular screws for atlantoaxial fusion in order to minimize the risk of injury to the vertebral artery. Other infrequent uses include transoral odontoid resections, cervical subaxial pedicle screw placement, and anterior cervical corpectomies.

Spinal surgery often includes placement of instrumentation such as screws, rods, hooks and wires. Screws can be used alone to repair fracture or instability as in C1-C2 transarticular screw fixation, or as an anchor for further instrumentation as in the case of pedicle screw fixation. Pedicle screws are often used in the treatment of pathological conditions such as arthritic deformity (spondylosis), fractures (iliosacral or vertebral), and can be used to support bony fusion (arthrodesis) (Cleary, 1999). In certain levels of the spine, extremely high accuracy is needed for the placement of pedicle screws to avoid perforation of the pedicle wall (Rampersaud et al., 2001). In many cases, misplacement of the screw can result in vertebral artery injury (Weidner et al., 2000). Cervical procedures are even more technically difficult as the screw trajectory is in very close proximity to the spinal canal, vertebral artery, and spinal nerve root. Image-guided surgery has been shown to improve the placement of pedicle screws and reduce the risk of screw misplacement (Austin et al., 2002; Youkilis et al., 2001; Weidner et al., 2000; Amiot et al., 2000; Henderson et al., 1996).

There are a few considerations in the use of image-guidance for spinal surgery. First, navigation systems eliminate the need for repetitive intraoperative fluoroscopy for tool placement, dramatically reducing radiation exposure (Foley et al., 2001; Welch et al., 1997). Conversely, the use of fluoroscopy or other intraoperative modality can be used to verify a CT-based image guidance system and avoid complications resulting from registration errors, modeling errors, shifting of the reference frame, or untracked intraoperative shifting of anatomy (Dickman, 2000). As far as how guided-surgery effects the overall duration of the surgery, there are arguments for no appreciable effect

and increased time consumption (Assaker et al., 2001; Weidner et al., 2000; Henderson et al., 1996).

Systems

There are a number of image guidance systems that are used in spine surgery. Each of them relies on the data from an imaging modality and a tracking system to relay vertebral orientation to the surgeon. Medtronic offers three configurations that use IR tracking, yet utilize different imaging modalities (Medtronic Surgical Navigation Technologies, Louisville, CO). The first, the StealthStation®, uses a traditional CT-derived model to perform navigation. The second system uses 3D fluoro-CT images by interfacing the StealthStation® with a Siemens SIREMOBIL® Iso-C^{3D}, an isocentric, automated fluoroscopy system (Siemens Medical Solutions USA, Inc., Iselin, NJ). Fluoroscopic-CT images generally have excellent spatial accuracy, yet their poor image contrast results in a 3D image reconstruction with poor tissue resolution and differentiation (Foley et al., 2001). These qualities hinder the construction of suitable virtual models for use in surgical guidance. The third pertinent system from Medtronic is the FluoroNav® virtual fluoroscopy system. It facilitates real-time navigation using C-arm fluoroscopy. This system employs a two-plane display, yet the drawback is similar to standard fluoroscopy in that 2D views do not give an accurate appreciation of 3D anatomy. G.E. Healthcare Technologies (Waukesha, WI) offers surgical guidance system configurations that use EM tracking, such as the OEC 9800 FluoroTrak™, which couples a high-resolution fluoroscopic imaging system with surgical navigation. A software upgrade to the FluoroTrak™ Surgical Navigation allows the use of 3D Fluoroscopic-CT images.

In addition to CT and Fluoro-CT, there are 3D image-based guidance systems that use MR technology. An example is the PoleStar™ Intraoperative MR Image Guidance System (Odin Medical Technologies, Inc., Newton, MA). It uses an infrared navigation system similar to CT-based systems; however, this system allows intraoperative image acquisition. The primary drawbacks to use of this system are very high cost and small size which prohibits spinal operations on the trunk.

Considering the options for tracking intraoperative spinal motion during surgical procedures, the cost of intraoperative MR and CT, the difficulty in resolving breathing artifacts, the lack of dimensionality of fluoroscopy, and poor 3D image construction from fluoro-CT, it is most clinically appealing to perform surgical guidance using a navigation system that utilizes preoperatively acquired CT images. These images provide a suitable platform from which to create an accurate anatomical model, plan the surgery, and program the navigation. However, the primary obstacle is that multi-level spinal procedures require the registration of a segmented model to be accurate.

Segmentation Introduction

Segmentation, as it is considered in medical image application, is the division of an image into anatomically labeled sections. A clinician experienced in the modality can easily recognize and outline the pertinent anatomy, however, a similar computational recognition or identification scheme is challenging to develop. Algorithms that automatically segment images have the potential to significantly reduce involvement of the clinician, encouraging the beneficial application of segmented models. These models can not only be used for visualization, surgical planning, and surgical navigation, but as a medium to study structural information particular to the modality.

In order for a segmentation algorithm to be successful, many factors must be considered. First, the time of the clinician is at a premium, so the algorithm must operate quickly (in accordance with current computational technology) and with minimum user input. It also must be flexible to adjustment; as such models will always be subject to the fine-tuning of a trained eye. Also, the algorithm must be accurate and robust despite the varied pathologies presented such as crushed vertebrae, fractures, scoliosis, and involvement of tumors. Artifacts are problematic to any imaging modality, and a segmentation algorithm must be tolerant to these instances. In the specific case of CT scans, instrumentation can cause such artifacts. Finally, the performance of the algorithm must be consistent and reliable, optimally eliminating the variability introduced through human error. These parameters, combined with a thorough review of methodology, will appropriately guide research of automatic segmentation algorithms.

Image segmentation is acknowledged as the most difficult and prohibitive step in the modeling of anatomical data. Despite this impediment, segmentation of images has a clear benefit to spinal surgery. The proper registration of an extended-level surgical plan to a patient diminishes the inaccuracies introduced during intraoperative movement. It is the focus of this research to develop and test methods to automate the process of vertebral segmentation of 3D spine models constructed from CT image slices. This will encourage image-guided spine surgery involving multiple vertebral levels, ultimately improving patient outcomes.

CHAPTER 2
SEGMENTATION BACKGROUND

Segmentation is the division of a whole image into a subset of connected pixel regions that have some common property. Segmentation in medical image processing, however, includes the implied step of anatomical labeling. That is the association of the image information to the appropriate anatomy for analysis. Computer algorithms that accurately segment medical images can be challenging to develop. However, there are many potential benefits such as enhanced visualization and ability to perform complex surgical planning and navigation procedures.

There are several approaches to the problem of medical image segmentation. Each of them relies on different information and can give somewhat different results, depending on the application. The first, voxel classification, uses globally defined characteristics to determine segmentation. For example, the use of an intensity threshold to identify bony anatomy qualifies as voxel classification. The next approach to segmentation involves the creation of a boundary concept that is used to mark the division between regions. This concept relies on the identification of discontinuities between regions. The contrast to this edge-based segmentation is region-based segmentation. This approach uses regional characteristics such as common intensity patterns to identify clusters. The final approach to image segmentation is the use of deformable atlases. The atlas is a general model that is associated to a particular anatomic structure. To perform segmentation, the general atlas is deformed to match an

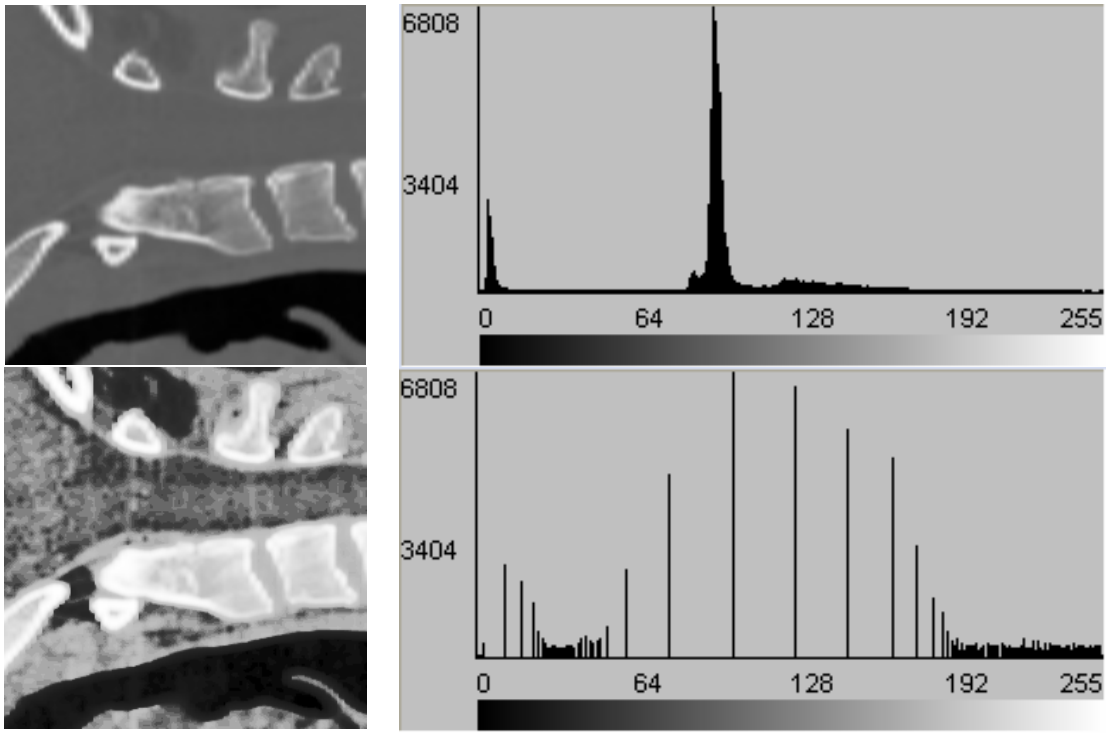

Figure 2-1. Intensity histograms and histogram equalization. An intensity histogram (upper right) is calculated from the sample image (upper left). The horizontal axis shows grayscale value and the vertical axis shows frequency. Histogram equalization improves the contrast over the most densely populated intensity ranges. A histogram equalization (lower right) is performed on the sample image (lower left).

individual image. Often, these techniques for medical image segmentation can be integrated in an algorithm to produce a more favorable result.

Image Processing

Certain methods can be used to extract information from an image for use in a segmentation algorithm or to enhance an image to better suit a segmentation algorithm. Examples include histogram equalization, boundary detection, and filtering.

An intensity histogram of an image is a frequency distribution of pixel intensities. The resultant graph indicates which intensity ranges have the most pixels (Figure 2-1). One use of the information provided in a histogram is to allow contrast normalization. This normalization process, called histogram equalization (Figure 2-1), maps pixels to

new intensity values to approximate a flat histogram. This process creates no new intensity values.

Boundary detection is a feature detection method used to identify boundary pixels in an image. These boundary pixels correspond to the region between different tissue types and are usually defined by a high intensity gradient (Roberts, 1965; Sobel, 1970; Prewitt, 1970). The gradient of a three-dimensional image with image intensity I is defined as:

$$\nabla I(x,y,z) \hat{=} \vec{i}\,\frac{\partial}{\partial x} I(x,y,z) + \vec{j}\,\frac{\partial}{\partial y} I(x,y,z) + \vec{k}\,\frac{\partial}{\partial z} I(x,y,z)$$

Various tracking systems can use the gradient information to make an edge trace. However, ambiguous or discontinuous edge data can produce errors. A drawback to the use of gradient information is that the image often must be smoothed because gradient calculations are susceptible to noise.

Filters are used to enhance the qualities of an image in accordance with some desired characteristic. In one instance, multi-dimensional adaptive filters are used to resample the image data to reduce partial volume effects and noise. They also handle the low off-plane resolution of CT images (Westin et al., 2000). Another filter of note is specifically designed to allow for a more robust image segmentation for use in guided surgery. It proposes to enhance separation of joint spaces in a CT scan, while allowing the retention of important edge information (Westin et al., 1998).

Mathematical Morphology

Mathematical morphology is a geometrical approach to signal processing (Matheron, 1975; Serra, 1982). It performs many image processing tasks using object quantification and easily deals with attributes such as shape and size, connectivity, and

contrast. Also, edge information is preserved during boundary manipulation. Common applications include noise reduction, texture analysis, and shape changing such as thickening, pruning, or skeletonization. These characteristics make mathematical morphology well suited to the task of image segmentation.

Mathematical morphology uses set theory as a foundation for many of its functions. In accordance, its primary operations use binary structures which are defined with an object and complementary background:

Object: $A = \{a | property(a) == TRUE\}$

Background: $A^C = \{a | a \notin A\}$

A common technique for the creation of such a structure is the threshold mask. The threshold mask T_M is a binary overlay that indicates which voxels in a three-dimensional image *I* are included within a prescribed intensity range $[t_0, t_1]$ as shown:

$$T_M = \{(i, j, k, 1) | (i, j, k, g) \in I, g \in [t_0, t_1]\}$$

Simple operations such as reflection and translation form the basis of more complex set functions, and can be applied to this binary image structure. Reflection of set *A* is indicated by $\hat{A}$:

$$\hat{A} = \{w | w = -a, \forall a \in A\}$$

In other words, $\hat{A}$ is the set of elements, *w*, such that *w* is formed by multiplying each of the coordinates of the elements, *a*, of set *A* by –1. Note that the elements or voxels in an image are considered vectors, so in a three-dimensional image in Z^3 space $a = (a_1, a_2, a_3)$. Another basic set operation is translation:

$$(A_x) = \{c | c = a + x, \forall a \in A\}$$

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

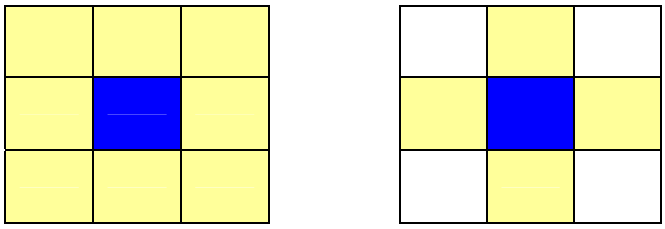

Figure 2-2. Morphological structuring elements and neighborhood configurations. A flat 3x3 structuring element (FSE) (left) is equivalent to a pixel with its eight nearest neighbors (8NN configuration) (middle). On the right is a structuring element composed of a pixel and its four nearest neighbors (4NN).

The equation says that the shifting of set A by vector x is the set of pixels c such that c is equal to a+x for all pixels that are a member of A.

Binary Operations

Two operations, dilation and erosion, are part of a core of image processing algorithms used in mathematical morphology. They produce a result by passing a structuring element over the image. This structuring element is analogous to the convolution kernel used in linear filter theory and it must have the same dimension as the image. Technically, in this application, the structuring element can be considered an image. For example, in an algorithm for a two-dimensional image, a flat structuring element (FSE) is commonly used. It has each element in the structuring array set to "true" or "1" (Figure 2-2). Other types of structuring elements can be defined according to their neighborhood configuration. For example, equivalent to the 3x3 FSE is a structure consisting of a pixel and its adjacent horizontal and diagonal pixels. These adjacent pixels are known as the eight nearest neighbors (8NN). Another common structure is a central pixel with its adjacent orthogonal pixels or four nearest neighbors (4NN). These neighborhood configurations are shown in Figure 2-2.

Nearest neighbors are defined by their connectivity patterns; the arrangements that determine if adjoining pixels are part of the same object. In the 4NN structure example

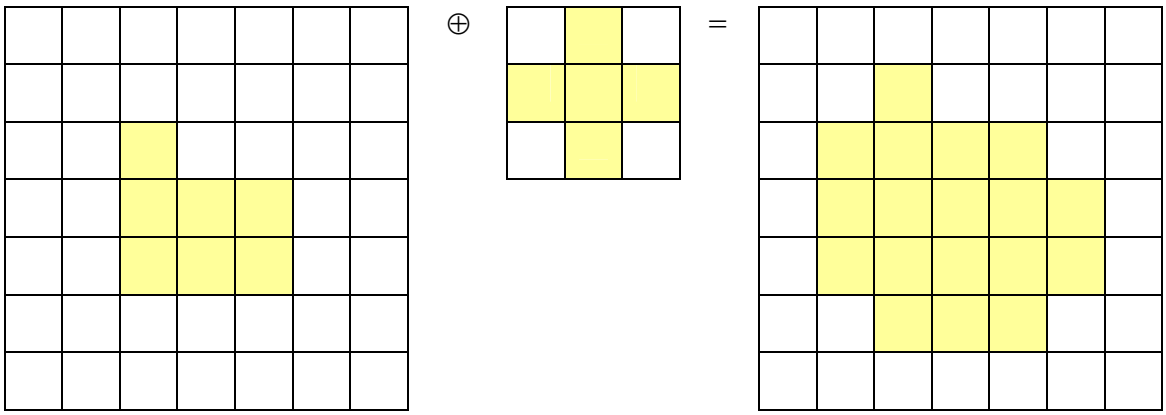

Figure 2-3. Binary morphological dilation. A binary image (left) dilated by a structure (middle) results in an expanded object (right)

above, all of the four nearest neighbors are 4-connected to the central pixel. On the other hand, pixels are 8-connected if they are connected in the orthogonal or diagonal direction, as in the 8NN structure. Connectivity arrangements for three-dimensional structures include:

1. 6-connected: voxels are connected by their faces
2. 18-connected: voxels are connected by their faces or edges
3. 26-connected: voxels are connected by their faces, edges, or corners

One of the simplest morphological functions to implement is dilation. It serves to expand the binary image structure and is analogous to convolution (Figure 2-3). In typical notation, A is the image set and B is the structuring element. The set theory formula for dilation is given by:

$$A \oplus B = \{z | (\hat{B})_z \cap A \subseteq A\}$$

That is, all coordinates *z* where the translation of the reflected structure *B* by *z* intersects with the binary image *A*. The implementation of dilation on a computer, also known as Minkowski addition, is in a slightly different form. It is the union of the sets where the

binary image is translated according to each of the points in the dilating structure:

$$A \oplus B = \bigcup_{b \in B} A_b$$

Erosion, another simple morphological image processing operation has the effect of shrinking a binary object. The set operation for erosion as well as Minkowski subtraction is given by:

$$A \ \theta \ B = \{x : (B)_x \subseteq A\} = \bigcap_{x \in B} A_x$$

The set equation simply states that the erosion of *A* by *B* is the set of points *x* such that *B* translated by *x* is contained in *A*. The Minkowski variant computes the erosion by taking the intersection of all of the sets from the result of *A* translated by each element of *B*. The two simple morphological operations dilation and erosion can be combined in series to form the compound morphological operations *open* and *close*. The open operation serves several functions such as smoothing of object contours, breaking of narrow isthmuses, eliminating thin protrusions, and acts as a filter to remove background noise:

$$A \circ B = (A \ \theta \ B) \oplus B$$

The close operation, on the other hand, fills gaps in contours, fuses breaks, eliminates small holes, and acts as a filter to remove foreground noise:

$$A \bullet B = (A \oplus B) \ \theta \ B$$

Opening and closing have the advantage of being idempotent, which means that repeated applications will not further change the signal. Examples of opening and closing are shown in Figure 2-4.

There are a handful of other binary morphological operations that are useful for segmentation. One is boundary extraction. A boundary is extracted from a binary object

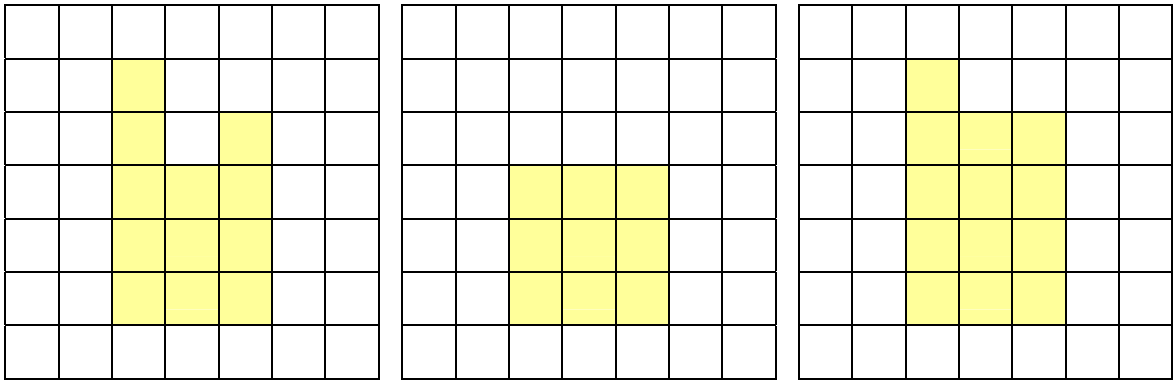

Figure 2-4. Compound morphological operations. Comparison of open (middle) and close (right) morphological operations performed on a binary image (left) using a 3x3 Flat Structuring Element.

by subtracting the result of an erosion operation performed on that object. The connectivity of this boundary is determined by the structure used for the erosion. Region filling is another useful operation. It utilizes the background or complement of a binary image as the area to be filled. A point X_0 is selected in the background and repeatedly dilated until it achieves the desired fill result:

$$X_k = \left(X_{k-1} \oplus B\right) \cap A^c$$

This technique is especially successful for filling internal gaps. Finally, distance transforms can be integrated into mathematical morphology (Cuisenaire, 1999). Distance transforms create distance maps from binary image objects (Figure 2-5). Distance maps are images in which the intensity of a pixel *p* is an indication of proximity or nearest distance to the object O:

$$D(p) = \min\left\{dist_M(p,q), q \in O\right\}$$

Distance maps have a smooth surface and an even gradient making them desirable as shape representations. In addition, distance maps have an interesting property. The result of an object dilated by a spherical structure can be expressed as the threshold of a distance map of that same object. *B* is a spherical structure with radius *d*, created by the

0	0	0	0	0	0
0	0	0	0	0	0
0	1	1	0	0	0
0	1	1	0	0	0
0	0	1	0	0	0
0	0	0	0	0	0

2.2	2	2	2.2	2.8	3.6
1.4	1	1	1.4	2.2	3.2
1	0	0	1	2	3
1	0	0	1	2	3
1.4	1	0	1	2	3
2.2	1.4	1	1.4	2.2	3.2

Figure 2-5. Morphological distance transforms. A Euclidean distance transform is used to form a distance map (right) from a binary image object (left). The pixel intensity, as indicated by the numerical value, relays the distance of the pixel to the object.

selection of all points that are less than or equal to a certain distance d away from point (0,0,0):

$$B = \{b | dist_M(b,(0,0,0)) \le d\}$$

And an object *X* dilated by the structure *B* is equivalent to the threshold of the distance transform *DT(x)* of the object at distance *d*:

$$X \oplus B = \{x | DT(x) \le d\}$$

Grayscale Operations

Binary morphological image processing methods can be extended to grayscale (Bangham & Marshall, 1998). Dilation and erosion become useful filtering functions. Dilation extends an object by using the maximum filter to remove low-valued regions. Accordingly, erosion contracts an object by using the minimum filter to remove high valued regions. These operations both have the effect of smoothing an image. Grayscale morphological dilation assigns to each pixel the maximum of the sum of the local region and the structuring element. In the special case that this structuring element contains all

zeroes, this operation is equivalent to the maximum filter:

$$A \oplus B = \max_{(i,j,k)\in B} \left(A_{x+i,y+j,z+k} \right)$$

And the dual operation for erosion using the minimum operator:

$$A \theta B = \min_{(i,j,k)\in B} \left(A_{x+i,y+j,z+k} \right)$$

These operations, as with their binary analogues, can also be combined to form the compound operations open and close. Opening acts as a high intensity point filter and closing acts as a low intensity point filter. Another useful grayscale operator is the morphological gradient. Given an image I, the morphological gradient is given by the difference between the respective dilation and erosion:

$$g(I) = (I \oplus B) - (I \theta B)$$

This intensity gradient is very useful for boundary extraction. The advantage of using these morphological operations is that they provide useful image-processing features that are easy to implement, and integrate into a segmentation routine.

Voxel-based Segmentation Methods

The simplest of the segmentation methods utilizes global image information to assign the memberships of voxels into particular anatomical regions. Boundaries of regions are then implicitly determined from a complete label map. Given a set of anatomical structures contained in an image $\{\omega_1 \ldots \omega_k\}$, a label map $L(x)$ overlays the image $I(x)$ and $L(x) = \omega_i$ where ω_i represents the anatomical structure at $I(x)$. The assignment of voxels takes into account such factors as intensity value, neighboring pixel classification, and relative distance of neighboring pixels. However, if the chosen factors result in an overlap of neighboring anatomical structures, the global representation makes

additional helpful information such as shape and geometric relationships difficult to incorporate. On the other hand, voxel methods are quickest in speed and do not require a cumbersome training model to achieve a result.

Thresholding

Thresholding is a simple, extensively used image processing technique that isolates a region of an image based solely on intensity criteria. It is a computationally inexpensive and fast low-level technique. In addition, the result may be used as an input to a higher-level segmentation model. It is defined as:

$$G(i, j, k) = \begin{cases} 1 & \text{for } I(i, j, k) \geq T \\ 0 & \text{for } I(i, j, k) < T \end{cases}$$

G is the resultant three-dimensional binary image of a grayscale image I thresholded at intensity value T (Figure 2-6). Semithresholding is a similar technique which masks out the image background leaving gray level information present in the objects:

$$G(i, j, k) = \begin{cases} I(i, j, k) & \text{for } I(i, j, k) \geq T \\ 0 & \text{for } I(i, j, k) < T \end{cases}$$

An upper bound may also be incorporated into these thresholding forms. It serves to isolate a specific intensity band. Overall, this class of voxel methods is typically applied in cases where particular anatomy or tissue can be identified within a certain intensity range. The most apparent example is the segmentation of bone, where all included voxels are for the most part at a higher intensity than adjoining tissues. The selection of a threshold value can be done through manual inspection or probability estimation. To perform the estimation, a tissue model is created that predicts which voxels belong to which structures based on probability distributions of intensities. The probability of

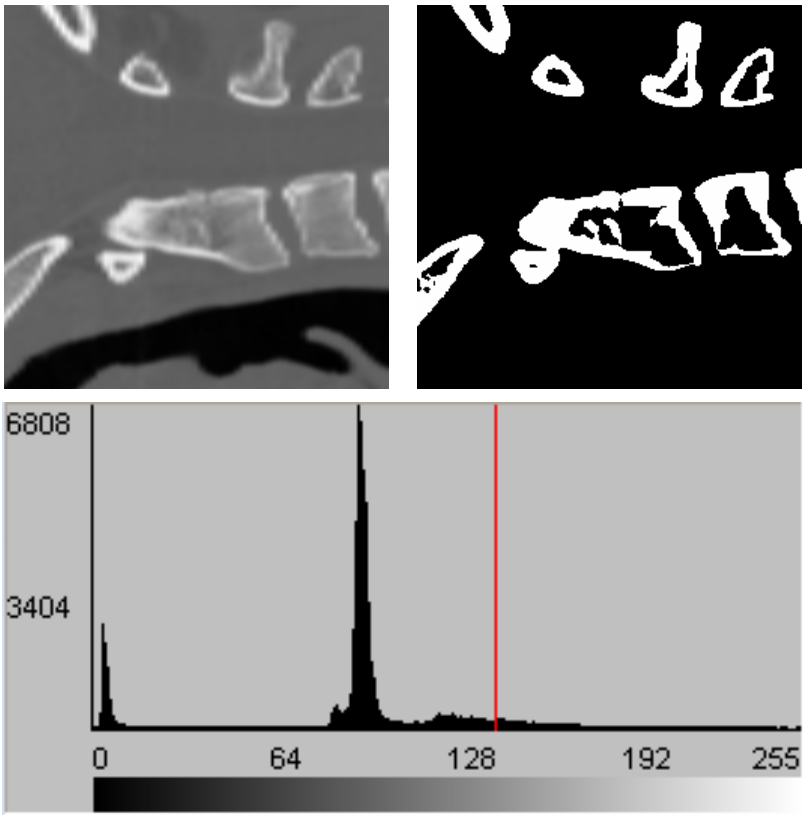

Figure 2-6. Threshold modeling. A binary image (upper right) is created by thresholding the sample image (upper left). The red line on the histogram of the sample image (bottom) indicates the threshold intensity value. Each of the pixels with an intensity value higher than the threshold value is highlighted in the binary image.

intensity value x, estimated at each tissue class ω_I is based on a set of training data $P(x|\omega_i)$. Histograms are commonly used to collect training data for this technique. Care must be taken in application of these voxel methods if intensity ranges are not disjoint. For example, in spinal CT images, it is not uncommon to have certain intensity values coinciding with both bone and cartilage or ligament. Also, thresholding or tissue class estimation is susceptible to imaging noise and artifacts that cause intensity overlap and unclear boundaries between tissues. A popular method for probability estimation is called Expected Maximization (Dempster et al., 1977). It is especially effective at handling incomplete data.

Tissue class estimators can actually use criteria other than intensity to perform segmentation of tissues. One example, which makes a classification according to regular neighborhood intensity patterns, uses a technique called statistical clustering (Leahy et al., 1991). Texture patterns are typically measured though a co-occurrence of distance and intensity.

Morphological Segmentation

Methods in mathematical morphology can be used to segment an image (Dougherty, 1993). And although morphology is a language for shape representation and manipulation, its basic segmentation methods can be considered voxel-based. The first concept in morphological segmentation is seeded region growing. A seed is a point voxel or a group of voxels. The seed is expanded by checking to see if neighboring boundary voxels are within specified criteria. A common criterion is if the absolute value of the intensity difference of a seed and its neighboring voxel is beneath a threshold. A very simple version of region growing is found in the watershed transform (Vincent & Soille, 1991; Beucher & Meyer, 1992).

The watershed transform is a segmentation method for grayscale images. It interprets the topology of an image and assigns watershed lines to boundaries between catchment basins. Catchment basins are the areas of local minima on the topographical map and are analogous to the depressions that would collect drops of water. The watershed lines are the crests between these basins. The common input to the watershed transform is a morphological gradient of an image so that watershed lines would correspond to the areas of strong edge evidence and divide the original image into homogonous regions. In application, however, simply taking local minima can result in oversegmentation, especially if there is a large noise contribution from false minima.

Figure 2-7. Binary image segmentation. The distance transform can be used to create a topographical image (center) from a binary image (left). The binary image is a silhouette of three overlapping shapes. Utilizing the topographical image, the watershed transform can then be called to approximate a segmentation (right).

The problem of oversegmentation can be handled in a couple of ways. The first is to apply a merging scheme, which combines adjacent areas according to some guideline. This guideline is usually based on statistical gray level properties. The second way to deal with an oversegmented image is to use marker selection. First, a new analogy should be introduced, and that is the gradual flooding of the image topography using a rising water table. Given this, marker selection is the decision of which local minima or seed regions will flood (Beucher & Meyer, 1992; Meyer & Beucher, 1990). The result is regions that are selected will flood the regions that are not, eliminating oversegmentation. Caution must be used in attempting to use smoothing operators or other such filters to reduce the false minima from noise contributions as they have the potential to remove areas of strong edge information. In modern application, the watershed transform has been used to segment out maxillofacial bone in CT (Böhm et al., 1999). The cited method also uses a tissue classification scheme to label the segmented regions. The watershed transform has applications on binary images (Figure 2-7). It can use a

topographical image created by the distance transform to separate overlapping objects (Cuisenaire, 1999).

The last morphological medical image segmentation method of note uses Recursive Erosion (RE) and Geodesic Influence (GI) (Kaneko et al., 2000). The recursive erosion is used generate candidate seeds while the geodesic reconstruction recovers the separated organs.

Edge-based Segmentation

Some segmentation problems can be solved using boundary localization. This concept involves the creation of an edge or surface model that is designed to converge on an object boundary. An accurate convergence will then describe the border of the segmented object. Given a three-dimensional image, closed surfaces are defined $\{S_1 \ldots S_k\}$ with all points inside surface S_i corresponding to an anatomical structure ω_I. Also, all points that represent ω_i are contained by S_i. The information for the model is gathered by relating the edge representation to its associated image information. This includes image gradient, texture discontinuities, or any other useful measure that can be geometrically associated with the boundary. One indirect scheme for the segmentation of 3D objects involves the unification of multiple segmentations in two dimensions. An example is the segmentation of bone in CT images. Once the proper segmentation of one image slice is accomplished, it can be used to direct the result of adjacent slices. The combination of all of these individual slices results in a segmented three-dimensional image.

Active contours are dynamic deformable models used for edge-based segmentation (Blake & Isard, 1998; Caselles et al., 1993; Malladi et al., 1995; Tek & Kimia, 1995).

One such model is called the snake (Kass et al., 1987). It makes use of an energy function to guide the evolution of the boundary. The snake itself is a spline function parameterized by a set of node points. The cubic polynomial is a common choice for the spline function:

$$x(u) = a_x u^3 + b_x u^2 + c_x u + d_x$$

$$y(u) = a_y u^3 + b_y u^2 + c_y u + d_y$$

The energy equation is then used to direct the movement of the nodes. The terms in the equation are a balance of forces that pull the spline toward the desired edge features. A proper segmentation is achieved through the minimization of this energy equation:

$$E_{total} = \int_{\omega} [E_{\text{int}\,ernal} + E_{image} + E_{constra\,\text{int}}]$$

This equation includes a term for internal energy, image energy and constraint energy. Internal energy is solely dependent on the shape of the spline. It includes parameters for stretching and flexing, which are optimized according to the geometric knowledge of the object to be segmented. The image energy is based on the image values along the path of the spline and can include attraction to strong gradients or regions of light or dark intensity. Constraint energy is intended to capture the higher level knowledge about the image and features. Example constraint energy terms are manually defined attraction or repulsion fields such as volcanoes and springs. A specific snake model designed for image segmentation was introduced by Kass et al. (1987):

$$E(C) = \beta \int |C'(q)|^2 dq - \lambda \int |\nabla I(C(q))| dq$$

It includes internal and external energy terms for the curve *C*. Internal energy is a regulating force that keeps the curve smooth by penalizing high curvature. This also

makes the curve robust against noise. External energy is image energy and is designed to act as an attraction force to high gradients. The coefficients of these terms are empirically adjusted according to the specific application. One consideration in the optimization of snakes is the dynamic adjustment of nodes. If the number of nodes remains consistent, then snake expansion will result in decreased resolution and snake contraction will result in unnecessary computational expense. Algorithms that effectively adjust number and spacing of nodes will maximize accuracy and speed.

Snakes have the advantage of being conveniently autonomous in their search for a minimal energy state. Also, since the integral operator is an inherent noise filter, they are insensitive to noise and other image ambiguities. These ambiguities include spatial aliasing and sampling artifacts that can cause boundaries to be indistinct and disconnected (McInernery & Terzopoulos, 1996). In contrast, snakes often overlook minute features in the process of minimizing the energy over the entire path of their contours.

Initiation of the snake is critical to the success of the algorithm and accordingly there are a few points to consider. First, the snake must be initialized close to the expected boundary for good performance since they can get stuck in lower minima states. This is a concern since contours are generally difficult to initialize around the region of interest. For best results, the snake should be initiated depending on the uniformity of the intensity distribution, either inside or outside the boundary. Caution must be observed if the curve is initialized on the inside of the expected boundary in a place where the image information has little influence as the curvature penalty will cause the spline to shrink to a point. The addition of an outward pointing force addresses the shrinking problem

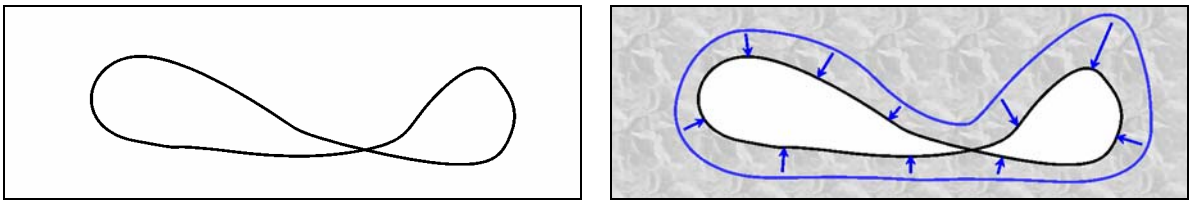

Figure 2-8. Singularity in contour evolution. A self-intersecting curve creates a singularity which can be described as in the process of either splitting or merging. Grass burning is one method used to solve this ambiguity (right). It determines the state by marking the path (texture) traveled by the curve (blue). In this case, the contour is in the process of splitting.

inherent in the internal energy term of the snake. This edge-based active contour is called a balloon (Cohen, 1991). The balloon function includes stability aids such as the normalizing of image force to restrict boundary movement and interpolation to avoid discretization errors.

A particular concern with the evolution of contours occurs when image topography causes the contour to self-intersect. The intersection point can be in one of two states, splitting or merging. The nodes can be reparameterized to adjust, as in establishing two separate contours in the case of splitting (Figure 2-8). However, the state must be established. Several solutions exist to tackle this problem. One is to use a grass-burning assumption (Sethian, 1996a). This creates an entropy condition where the evolving front leaves a “burnt” or irreversibly marked path as it travels. This will determine what direction the contour was moving. Another solution is to use T-snakes, which use a triangulation of the embedded space to determine new node points (McInerney & Terzopoulos, 1997). The Level Set Method is yet another method for tracking the evolution of interfaces (Osher & Sethian, 1998).

Level set methods offer highly robust and accurate methods for tracking interfaces moving under complex motions (Osher & Sethian, 1988). They handle topologically breaking and merging naturally as in the creation of channels in a surface contour.

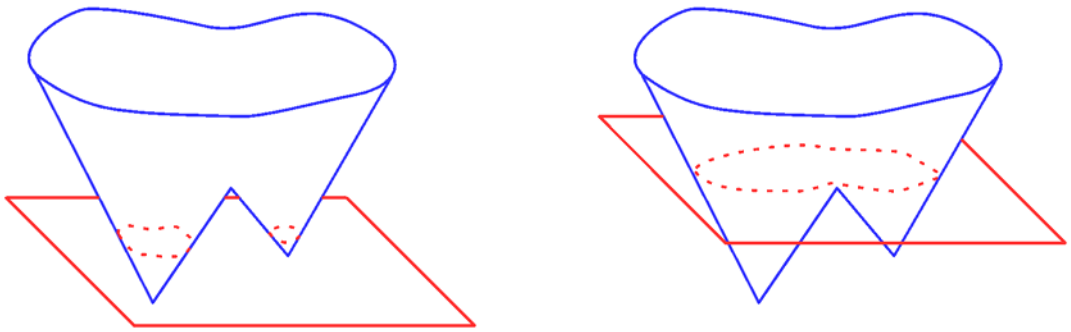

Figure 2-9. Level set application for contour evolution. The contour (red) is embedded as a zero level set of a higher dimensional surface (blue). This allows the smooth adaptation of the contour to topographical changes.

Details about theory, implementation, and application of level set methods are in Sethian's *Level Set Methods* (1996a). It relies on two central embeddings. One, the contour is embedded as the zero level set of a higher dimensional function (Figure 2-9). Two, the velocity of the contour is integrated into the higher level function. This is called the extension velocity, and is curvature dependent. A companion technique to Level Set Methods is Fast Marching Methods (Sethian, 1999, 1996b). They are numerically efficient methods for evolving a front traveling in one direction. They can be used to construct a distance map and provide appropriate extension velocities for level sets (Adalsteinsson & Sethian, 1999). There are several associated active contour methods of note that are used for image segmentation. The first uses a geodesic or minimal path formulation for active contours in the level set approach and applies it to segmentation of medical imagery data (Caselles et al., 1997). The second uses gradient flows to direct the curve evolution (Kichenassamy et al., 1995). The last integrates other geometric techniques for image segmentation (Malladi et al., 1995).

Much of the current research in active contour models deals with generalizing the form of the contours and overcoming the convergence and stability problems encountered during the energy minimization process (Davatzikos & Prince, 1992).

Region-based Segmentation

In applications such as medical imaging, adequate edge information such as strong gradients may not always be present for an edge-based segmentation approach to be consistently accurate. In these cases region-based models may be employed. In the region-based segmentation approach, the boundaries of the deformable model are determined by statistics inside and outside a region or cluster of voxels. These statistics measure properties such as unique texture patterns, homogeneity of intensity, or some other pixel-based statistic. The goal in evolving the region is that variation of these properties is less inside a region then between regions. Similar to edge-based methods, region based methods are evolved through the minimization of an energy term. This global energy term, however, is defined for the entire area of the region rather than only the boundary. There are a variety of region-based segmentation methods that include Bayesian segmentation (Geman & Geman, 1984), piecewise constant energy (Mumford & Shah, 1989), region competition with balloons (Zhu & Yuille, 1996), and an energy-based watershed (Bleau & Leon, 2000; Nguyen et al., 2003). One technique uses the level set formulation to evolve the boundary of a region based on texture statistics (Yezzi et al., 1999).

Region-based segmentation algorithms have several advantages. First, they have a greater capture range and are not as dependent on initialization as edge-based methods. Also, they are not reliant on high frequency information and are not as susceptible to

noise (Chan & Vese, 1999). Care must be taken in application, though, as the determination of regional parameters can be computationally expensive.

As a final note, efforts have been made to integrate region-based and edge-based modeling techniques (Chakraborty et al., 1996; Paragios & Deriche, 1999). The addition of a boundary regulation term is used to control smoothness in a region-based segmentation model. Also, geometric information, such as a shape prior, is much more easily incorporated into a boundary formulation. However, the tradeoff of boundary regularization comes at a cost of strength of region description.

Shape-based Segmentation

The final approach to medical image segmentation is to incorporate shape information in a deformable modeling algorithm. This is accomplished through the construction of an atlas. An atlas is a collection of prior information used to direct the deformation of the algorithm. It can be a set of points, finite elements, flow fields, or other similar parameter, but is most commonly a shape descriptor. Once defined, the atlas or standard template molds itself to the target image, imprinting an inherent label map. The data associated with the atlas can also be used to direct the mentioned edge or region-based segmentation strategies.

Atlas Warping

The classic warp consists of a representative image scan that has an assigned label map $L^*(\mathrm{x})$. Given a new image, some transformation T must be computed that deforms the representative scan $I^*(T(\mathrm{x}))$ to correspond with $I(\mathrm{x})$. The same transformation applied to the label map $L^*(T(\mathrm{x}))$ will then describe the label map for the new image $L(\mathrm{x})$. There are several examples of this template-driven segmentation (Pentland & Sclaroff, 1991; Sclaroff & Pentland, 1995; Wang & Staib, 1998b; Warfield et al., 1998).

One specific method computes the deformation field in a way that allows a course to fine solution (Christensen et al., 1996; Christensen, 1999). Another method begins with a rigid transformation and iteratively progress in plasticity to find a solution (Miller et al., 1997). A technique especially useful in the atlas-based segmentation of composite boney structures incorporates the rigidity of the tissue class in the deformation model (Little et al., 1997). More examples of atlas warping include works by Cootes and Taylor (1992a), Jones and Poggio (1998), and Pichumani (1997).

Modeling

There are several factors to consider in the development of a shape-based model. First, for automatic interpretation, it is essential to have a model that not only describes the size, shape, location and orientation of the target object but that also permits expected variations in these characteristics. In order to properly account for this object variability seen in application, a statistical analysis must be performed (Dryden & Mardia, 1998; Neumann & Lorenz, 1998). Casting the fitting process of deformable modeling into a probabilistic framework allows incorporation of prior statistics as well as an inherent measure of uncertainty (McInerney & Terzopoulos, 1996). These statistical models have the advantage of being flexible to ambiguity and noise. Since the development of a model and its associated training data set can be data intensive and computationally laborious, techniques are used to minimize the number of statistically significant parameters. One technique in particular is Principle Component Analysis (PCA). This process identifies the primary modes of variation and reduces the dimensionality of the data set, optimizing the framework (Golland et al., 2000; Lorenz & Krahnstover, 1999).

Multiple strategies for shape description have been developed. Feature detection is a common approach. A feature is a geometric property of the object to be segmented that

can be easily identified by an algorithm. The algorithm performs a segmentation by matching the set of features that is extracted from the unlabeled image to the training set. Depending on the application, the association of these features should be defined with certain invariances such as translation, rotation, or homogeneity. One method of feature extraction uses medial axis or skeletonization. The geometric features produces are scale-invariant, which allow detection at different resolutions (Dougherty, 1993). A drawback to the use of skeletons is that original size and edge information are lost. An improvement on skeletonization uses fixed topology skeletons. They have the advantage of being robust to the noise and quantization errors that traditional skeletons are susceptible to (Golland et al., 1999). Mathematical morphology can be used to build a shape classification strategy as well (Dougherty, 1993). The binary object of interest is probed with an array of simple shape primitives, from which a statistically appropriate feature set can be collected. In another method, distance maps are used as a shape descriptor (Golland et al., 2000). Distance maps provide smooth and wide minima for a matching algorithm.

Methods

There are a variety of shape-based segmentation algorithms. One particular group uses Active Shape Models (ASM) (Cootes & Taylor, 1992a; Cootes et al., 1992b; Cootes et al., 1999b). The ASM is a parametric deformable model that is represented in the image as an n-point polygon. The algorithm then deforms the points of the polygon to match the landmark points in the target image object. The points are evolved according to a point density model (PDM), which is constructed using PCA. The PDM provides a statistical distribution of global shape variation. Once the model is deformed, an accompanying label atlas identifies the segmented object. The ASM is significant

because it is fast and accurate and can describe complex shapes with minimal parameters. In application, models have been developed to account for both Gaussian and non-Gaussian data (Cootes et al., 1995; Cootes & Taylor, 1999a). In one instance, ASMs are incorporated into a Bayesian formulation for boundary location (Wang & Staib, 1999a). ASMs have also been extended to accommodate grayscale information as Active Appearance Models (Cootes & Taylor, 1998).

Shape-based segmentation approaches have been integrated into edge and region-based algorithms to retain the benefits of those approaches. For example, contours are well suited to handle shape information, and such information can easily be incorporated as priors. One method uses statistical priors on the Fourier coefficients of the contour to represent shape (Staib & Duncan, 1992). A related method combines snakes with a shape-based Fourier parameterization method (Szekely et al., 1996). Chen et al. (2001), uses shape priors with a level-set evolved geometric active contour model. Leventon (2000) tests a method that incorporates distance-intensity profiles and shape contour with a probability-based segmentation approach. It uses a level set implementation of geodesic active contours for shape evolution. Also, PCA is implemented to reduce the dimensionality of the training set. A couple of region-based shape methods are also of note. Cremers et al. (2002) incorporate statistical shape knowledge into the Mumford-Shah functional to perform segmentation. In application of cervical spine segmentation, Pichumani (1997) uses a finite-element shape model integrated into a region-based approach.

CHAPTER 3
SEGMENTATION

Currently, to create a segmented model appropriate for guided surgery, the clinician follows a few basic steps. First, he reviews the volumetric CT scan and selects an intensity value that defines the edge of the bone. Isocontours are displayed on the terminal to assist in this task. Typically, voxels of bone have much higher intensity values than adjacent soft tissues, and the boundary value can be easily identified. The selected intensity is used to define a threshold mask. The mask is a binary overlay which identifies all voxels that are associated with bone (Figure 3-1). The highlighted voxels describe an object which can usually be seen with accompanied 3D visualization software.

There are a few considerations in the creation of a threshold model. First, it is crucial that the edge of the bone is accurately defined, as many of the registration techniques use surface points to do the matching. A model created with a threshold value that is low will include non-bony soft tissue such as cartilage and ligaments. A model created with a threshold value that is high will diminish the volume of the bone, shrinking the edge and withering the interior. The process of using a threshold to describe an area of the image can also be utilized to isolate an object from high intensity artifacts or noise. Instead of a just a single lower intensity threshold, an upper intensity bound is used as well.

The model created by thresholding rarely results in a structure that is segmented or appropriately sectioned to match the underlying anatomy. Therefore, the model must be

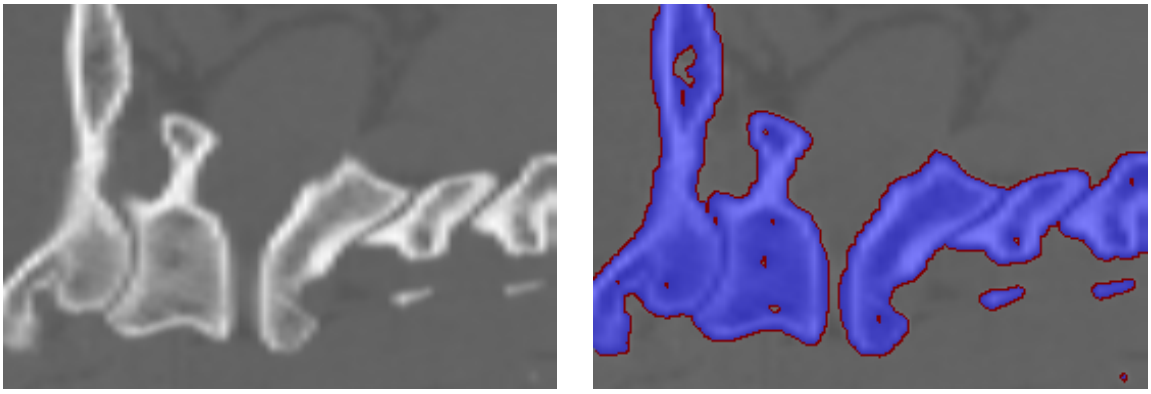

Figure 3-1. Building a threshold model. On the left is a sagittal slice of a cervical CT image. On the right is an isocontour (red) scribed at the boundary of the bone, with its implied threshold mask area highlighted in blue. It can be seen that several of the bones are inappropriately grouped together.

manually segmented. This is a tedious and time-consuming process whereby the clinician flips through each of the series of CT image slices and uses simple pixel highlighting tools to identify each bone. The result is a segmented model that can be used for guided surgery (Figure 3-2). However, even though this method produces an acceptable result, the time commitment involved in this manual process restricts the routine use of guided-surgery for multi-level applications. Therefore, the focus of this research is to develop methods to automate the segmenting process and facilitate these procedures.

Problem Evaluation

Several guidelines were established to direct this segmentation research. First, the time spent by the clinician to prepare a segmented model must be minimized. This includes increasing simplicity and minimizing user input. In contrast, however, there must be enough prior knowledge provided by the user to direct the algorithm to a correct solution. Another contributor to the involved time of the clinician is the speed or computational efficiency of the algorithm, but this optimization is secondary at this stage of the research. Second, there must be an accurate approximation of bone edge for the

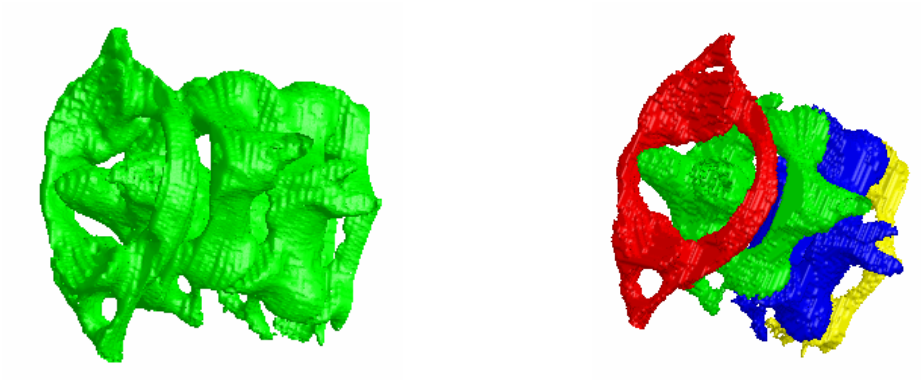

Figure 3-2. Manual segmentation and display. The image on the left is a 3-D view of a contiguous model formed using the thresholding method. The image on the right is the same model that has been manually segmented. The goal of this research is to reduce the involved time to produce segmentation.

registration methods to be successful. Third, the segmentation process must be robust to the assorted image variations. These include pathologies and artifacts, such as those induced with instrumentation.

The initial step in development of the segmentation algorithm is to assess a random sample of images used for image guided surgery. Currently, most of the image-guided surgical procedures are implemented on the cervical spine. This is because there is a very small tolerance available for tool placement to avoid critical vascular and nerve structures. A series of the high-resolution cervical CT scans used for these surgeries were evaluated. Upon analysis, it was apparent that the general shape and intensity distribution of the vertebrae were similar. Some patients had an overall higher bone density, with clear distinction of boundaries. Others had bone density that was close to that of soft tissue. In some of the most severe cases, the individual bones were difficult to recognize. Some of the pathologies observed include fractures and fusion, each which showed unique discontenuities and intensity patterns. Instrumentation, such as rod and screw assemblies or wires used to stabilize bones, produce unique artifact patterns in the

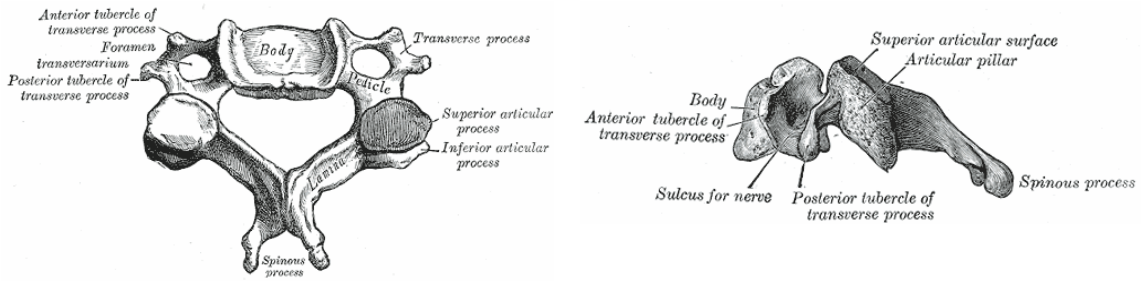

Figure 3-3. Surface characteristics of vertebrae. The typical cervical vertebra has very complex topology, with areas of high curvature and channels. These properties make it a troublesome candidate for edge-based or shape-based algorithms. (Left) Superior view. (Right) Lateral view. (Gray's Anatomy).

CT scan and disrupt the normal intensity patterns. A careful review of image variations was helpful in selecting the approach to algorithm design for this research.

Method Selection

Several avenues exist for the development of a segmentation algorithm. One common approach is to use boundary localization techniques. These involve the creation of a boundary concept that will deform until it reaches a minimum energy state at an appropriate edge. They have the capacity to handle small gaps in information, but also have the potential to arrive at false minima. These are erroneous minimal energy states that locate an incorrect boundary. A primary difficulty in the use of boundary localization techniques is that the typical cervical vertebra has very complex topological characteristics including areas of high curvature and channels. This makes it a difficult surface to evolve (Figure 3-3). In addition, the peculiar shape would exacerbate the initialization problems that have the potential to plague these methods.

Another avenue for image segmentation is to use region-based methods. These methods have the advantage of not being as dependent on initialization as edge-based methods and are less affected by the gaps in information or high frequency errors that may cause boundary methods to reach a false minima. A drawback, however, is that

region-based parameters are often computationally expensive. In addition, region-based methods may not be best suited to handle the complex boundaries and topology presented with a typical vertebra. One advantage to a region-based approach, however, is that shape information can be more easily incorporated. An atlas, if properly developed, could capture variance across a population and be used not only to direct segmentation, but also to identify abnormality or pathology. The incorporation of shape information, however, can be costly. First, reliable features must be identified. And then, a comprehensive database must be put together and maintained. Given the number of pathologic conditions to account for, this could be a very time consuming process.

The third avenue to develop a segmentation algorithm is to use voxel-based techniques, which are independent of any boundary or region components. This significantly reduces complexity, yet has the drawback of not being able to easily account for geometric information. Thresholding is a proven, if inefficient, voxel-based modeling technique, and makes a good approximation of the bone surface. Therefore, it provides a convenient platform from which to launch other voxel-based modeling techniques. Given that this is the simplest approach, it shall be investigated first. To facilitate this work, morphological operators will be used, as they are computationally efficient, work in 2D and 3D and preserve edge information.

Research Tools

A programming language and accompanying toolset were selected to facilitate the segmentation research. IDL (Research Systems Inc., Boulder, CO) was chosen because of its strong image processing toolset, assortment of visualization methods, a GUI interface, and ability to process large arrays. First, a GUI was created to visualize volumetric CT images. Through it, any orthogonal image slice could be viewed and

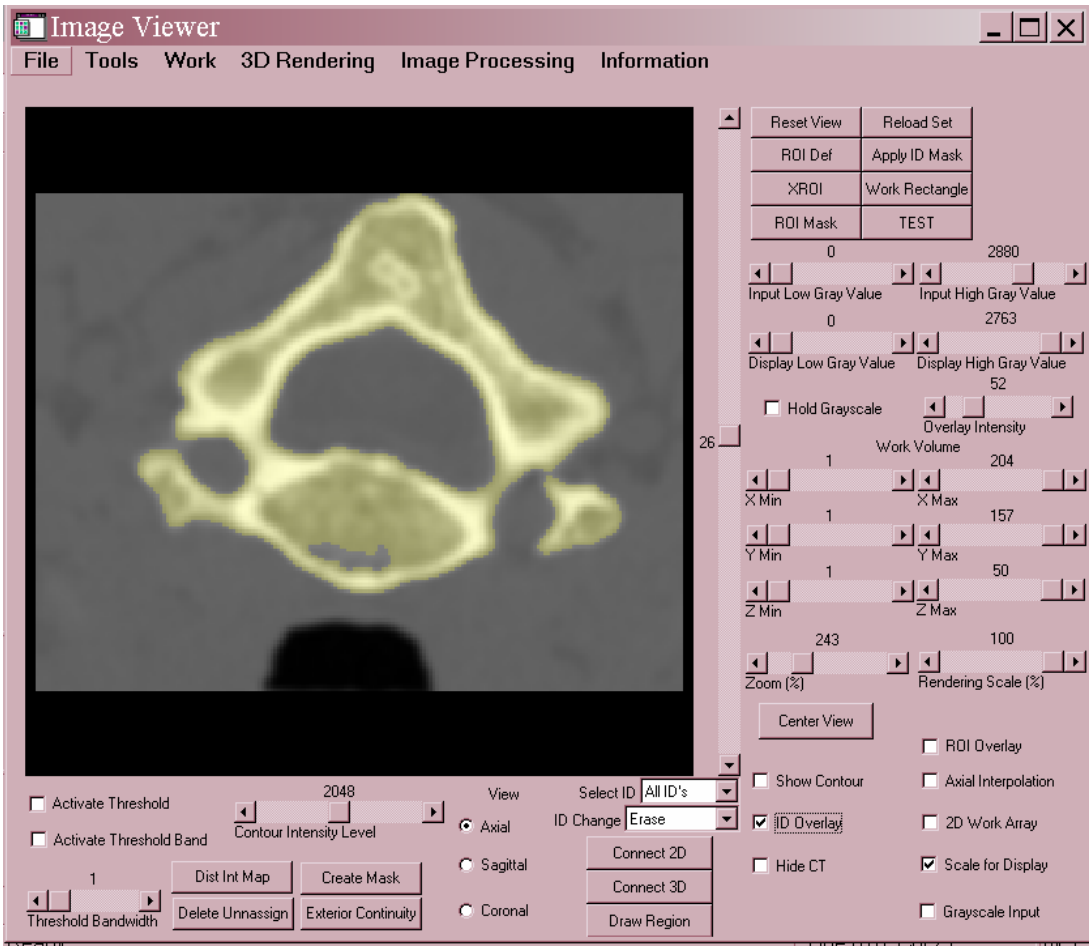

Figure 3-4. Toolset GUI. Multiple functions are easily arranged to facilitate segmentation. Shown is an axial CT slice with an overlaid indexed array (yellow highlight), indicating the segmentation boundary and corresponding anatomical association.

manipulated. Plus, a function for overlaying intensity contours to aid in the selection of threshold values was included. Since a 2D view does not give an accurate account of 3D connectivity, several 3D visualization schemes were also setup. Once a viewing system was constructed, multiple groups of 2D and 3D image processing functions were added to the menu structure, such as the morphological functions open, close, and gradient.

The next step in toolset development was the addition of functions to the program environment that mimicked the manual segmentation process. Before this could be implemented, however, a data representation of the segmentation model had to be

chosen. We decided on an indexed overlay. This is an integer data array equal in size to the image array with an index number assigned to each voxel. This index identifies that voxel as being associated to a specific bone. Once created, this array could be blended with the original picture for visualization (Figure 3-4). This data representation for segmentation is certainly inefficient, yet is easy to associate with any voxel-based manipulation. Next, several options were added to the GUI to allow the creation of the index array based on threshold techniques. Several tools were created to assign or change index assignments. They include 2D single pixel and polygon and 3D sphere and rectangular prism assignment options. In addition, continuity functions in 2D and 3D were also added. They can be used to "paint" interior spaces or to fill holes. Throughout the research, functionality of the interface was expanded accordingly.

Cursory Examination

In order to test the toolset interface, a few simple, preliminary examinations were made. The primary failure of segmentation by thresholding is that the process does not properly isolate each of the bones, leaving one contiguous clump of bone. These initial experiments, then, attempt to use a few simple techniques to accomplish improved separation. There are two approaches to test. One is to cut the threads that leave the bones connected. The other is to identify the boundary and subtract it from the contiguous model, leaving separated regions.

The morphological open operator is a simple function that can be used to effect a break of thin connections on a binary object. Hypothetically, this could be used to separate bones on a threshold model. Several tests were performed using the open function on sample images. Symmetric structuring elements were applied in an order of increasing width to see what size would effectively break the connections. The result is

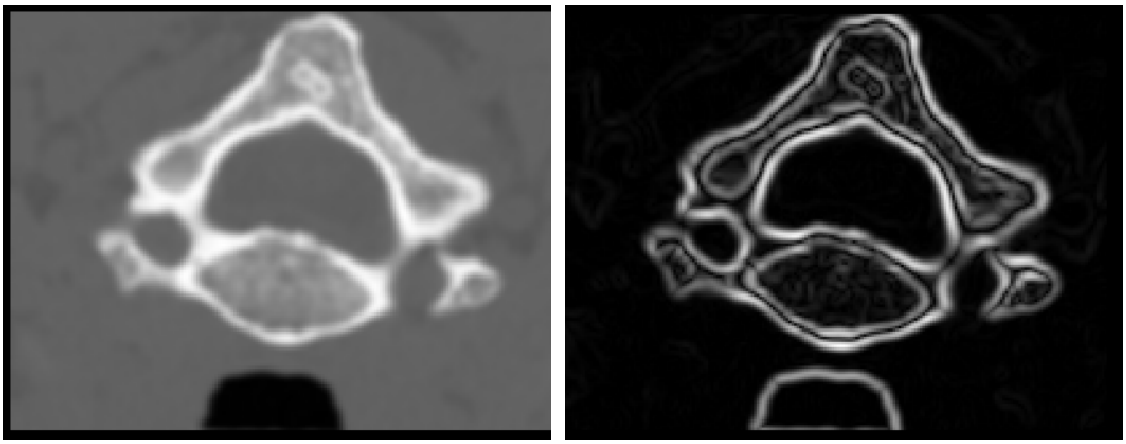

Figure 3-5. 3D intensity gradient. The morphological gradient is constructed from a CT volume. On the left is an axial slice of a cervical vertebra. On the right is an equivalent slice from the gradient map. Notice the interior areas of strong edge-information.

that the size of structure required was thicker than some of the natural boney structures, and the application of which significantly altered the surface of the vertebra.

The intensity gradient is a very commonly used indicator of edge strength. Several 3D gradient intensity maps were calculated from sample CT scans and analyzed. A majority of the boundaries were highlighted in these trials; however, there were multiple gaps or areas of low gradient. Therefore, without further modification, the gradient map could not be directly used to make a boundary subtraction and foster segmentation. However, an interesting property of the vertebrae is that there is a strong gradient on the interior of the bone (Figure 3-5). One observation is that, given these circumstances, an edge-based algorithm solely dependent on gradient information would have difficulty localizing the correct boundary. The interior areas of strong gradient would tempt the curve during evolution.

As a result of these preliminary experiments, a rectangular region of interest (ROI) tool was setup to allow the application of image processing algorithms on a defined area. It was theorized that this would allow unique image faults to be addressed without

affecting the whole image. For example, it could be used to address specific gaps in the gradient map. After a few trials, it was apparent that ROI definition in accordance with a specific region in 3D was a time-consuming process subject to burdensome iterations and is something that should be avoided.

CHAPTER 4
REGION GROWING

Kernel Definition

The fundamental problem in using a threshold model created at the edge defining value for bone is that it often connects adjacent anatomy as one contiguous model. However, it has been observed that a regression of the isocontour to a higher threshold value results in a disjoint model that can be uniquely associated with the underlying anatomy. It is hypothesized that these regions can be used to create a successful segmentation.

A method has been developed to create and label this kernel-based threshold model using IDL. First, an isocontour is selected that allows the creation of disjoint regions. This is an iterative process that involves using the visualization software to check for contiguous regions. Once an intensity level is selected, a binary threshold is created in the label array. The proper IDs are assigned through point selection and continuity labeling tools (Figure 4-1). Since the goal is to minimize the intervention of the clinician, not every unassigned region is labeled. It is assumed that the user will use knowledge of the segmentation method to label the regions that will have the most profound effect on the results.

Shape Correlation

It is observed in the creation of the disjoint threshold model that the shapes of the kernel regions bear a resemblance to their associated boundaries. It is reasoned that if

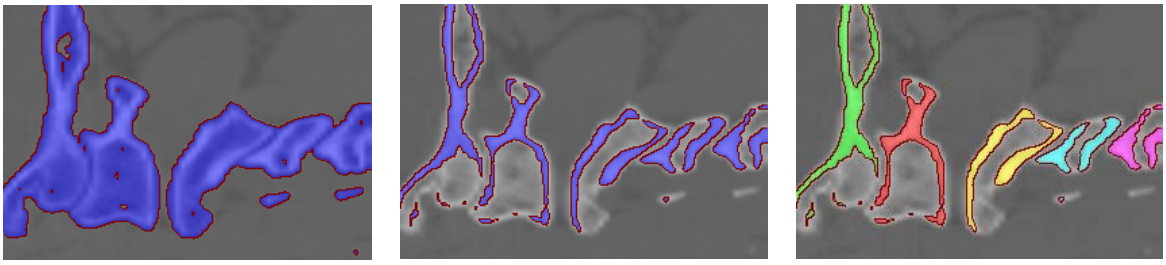

Figure 4-1. Manual kernel selection. A threshold model created at the bone-defining intensity value of a sagittal slice of a cervical spine CT image (left). The threshold value is raised, resulting in a model with multiple discontinuous bony elements (middle). Each area is labeled to correspond to the underlying anatomy (right).

enough boundary information is preserved in the shape of the kernels, then they can be used for segmentation.

One simple voxel-based method to test this theory uses both the labeled region kernel model and the contiguous boundary threshold model. They are compared, and each of the highlighted voxels in the boundary model are assigned according to the closest labeled kernel region. The assumption is that there is a high likelihood that an unassigned voxel has the same label as the closest pre-labeled region. This can be accomplished through a system of constrained morphological dilations. Specifically, the kernels of adjacent anatomies are dilated equally until they touch. The voxel IDs assigned in this process cannot be overwritten, so the boundary created is fixed. The process of repeated dilation continues until all of the voxels in the boundary threshold model are assigned. Expanding on this idea, distance maps may be used to approximate this process. A distance map is an image whereby each of the voxels is assigned an intensity value corresponding to the Euclidean distance to an object. This value, as it turns out, is approximately equal to the number of unit dilations necessary to reach that voxel from the object. Further, if a distance map is created for two sets of objects, the

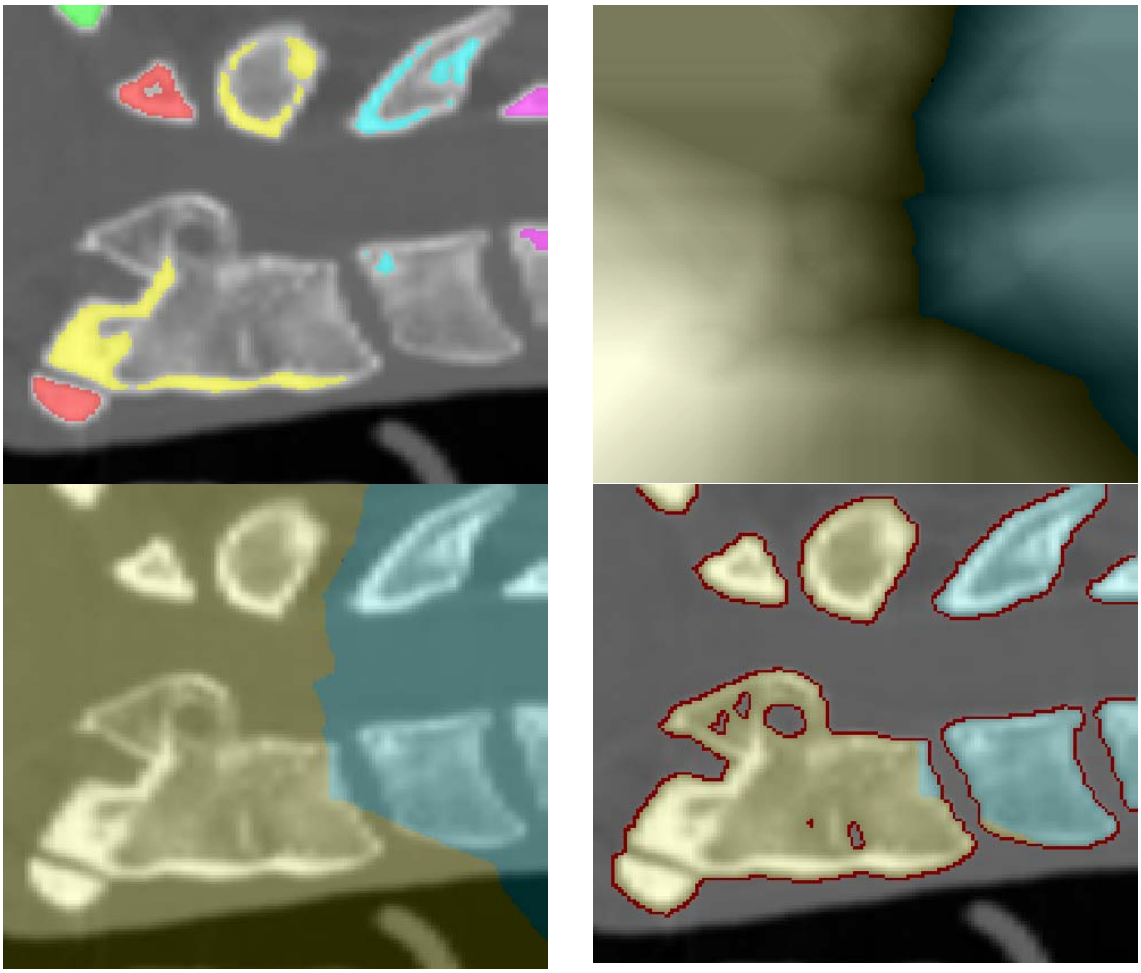

Figure 4-2. Segmentation using distance maps. Unique, discontinuous kernel regions are created from a cervical CT scan (upper left). Then, two distance maps are created. One represents the kernels associated with C2 (yellow) and the other represents the kernels associated with C3 (Cyan). The difference of the resulting maps (upper right) approximates a boundary between C2 and C3 (lower left). A segmentation results when a boundary threshold mask is applied (lower right).

difference of the two maps scribes a zero line exactly equidistant between the object sets. Figure 4-2 illustrates this process enacted upon two sets of adjacent kernels.

This segmentation method is a good way to quickly approximate a boundary. However, there is one failing that made further pursuit inadvisable. The amount of shape information contained in the labeled kernels rapidly deteriorates as the threshold level is

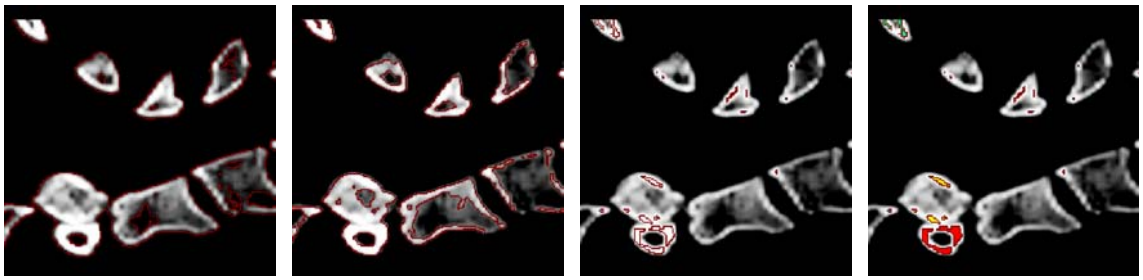

Figure 4-3. High threshold values result in diminished kernels. The threshold contour (red) is increased from the boundary defining value (left) to one that allows proper label associations. The resulting high threshold and label (right) shows kernels that retain very little boundary shape information.

Object 4-1 .High threshold movie (object1-highthreshold.mpg, 27.6 MB)

increased (Figure 4-3). Therefore, circumstances that warrant a high threshold to achieve isolated kernel regions produced the least accurate boundaries.

Hysteresis

The previous study showed that the shapes of the segmentation kernels are not infallibly correlated to boundaries. More information is needed to direct the growth of these kernels to arrive at the appropriate segmented boundary. In the kernel labeling process, an intensity threshold is used to withdraw the isocontour to the interior high intensity areas of the bone. This threshold information, if appropriately utilized, can assist the kernels to revert back to the bone edge.

An algorithm was developed using a combination of morphologic dilation and thresholding functions to progress labeled kernel regions to a fully segmented model. The inputs to this algorithm are the CT image, the user-determined bone boundary threshold intensity value, and the kernel region dataset. Each iteration of the region-growing algorithm involves three basic actions (Figure 4-4):

1. The step to a lower threshold and resulting isocontour.
2. The constrained orthogonal dilation of the seed regions.
3. The restriction of the dilated voxels to the area outlined by the isocontour.

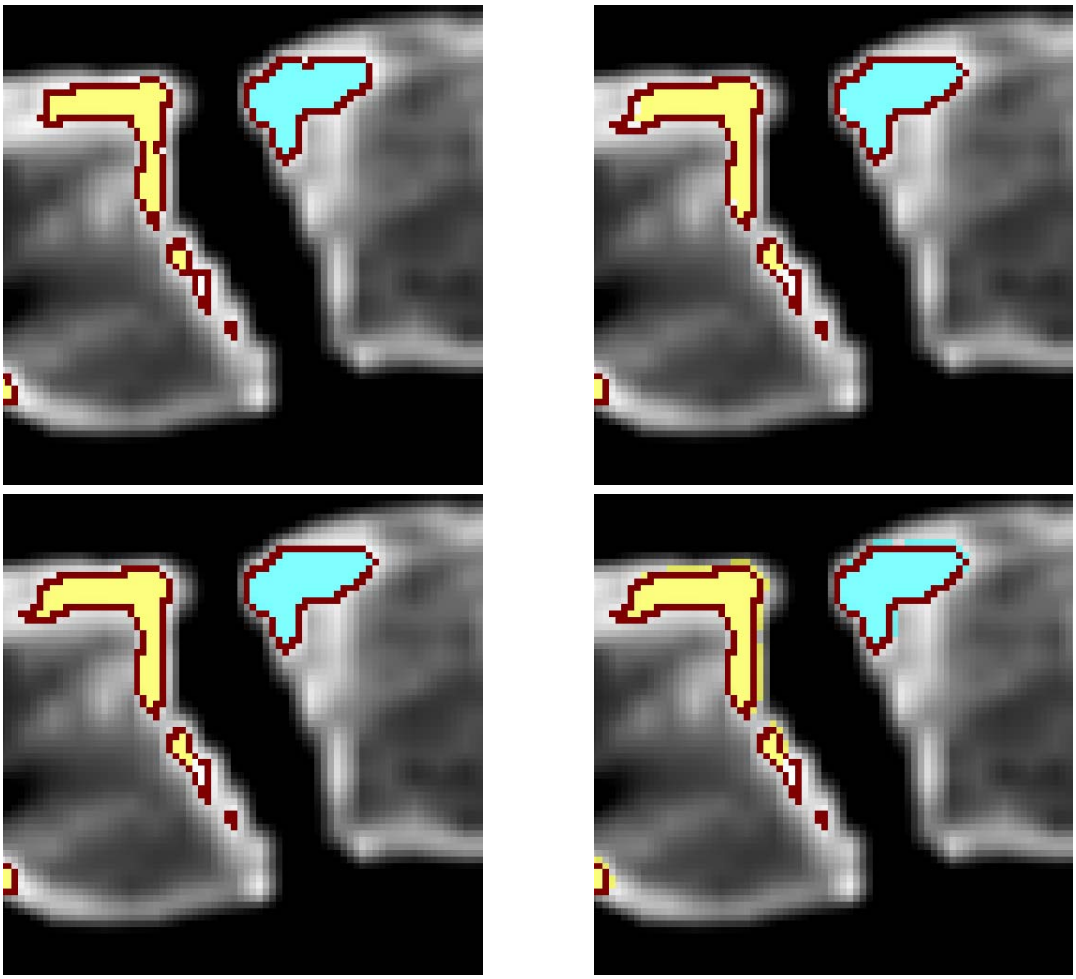

Figure 4-4. Region growing algorithm iteration. Shown in the upper left are two adjacent bone segments with corresponding seed regions. The first step progresses the threshold contour to a lower value (upper right). Then the seed regions are grown through morphological dilation (lower right). Finally, any highlighted voxels outside the threshold contour are removed (lower left).

This basic cycle is initiated at the kernel-defining threshold value and repeated until all of the volume outlined by the lower threshold mark is filled.

A couple parameters can be adjusted to discover the most effective sequence for region growing. The first is the number of threshold divisions, or the number of times that the isocontour value will step in order to reach the edge defining value. The height of the step will depend on the number of threshold divisions and the difference between the two boundary threshold values. The second definable parameter is the number of unit

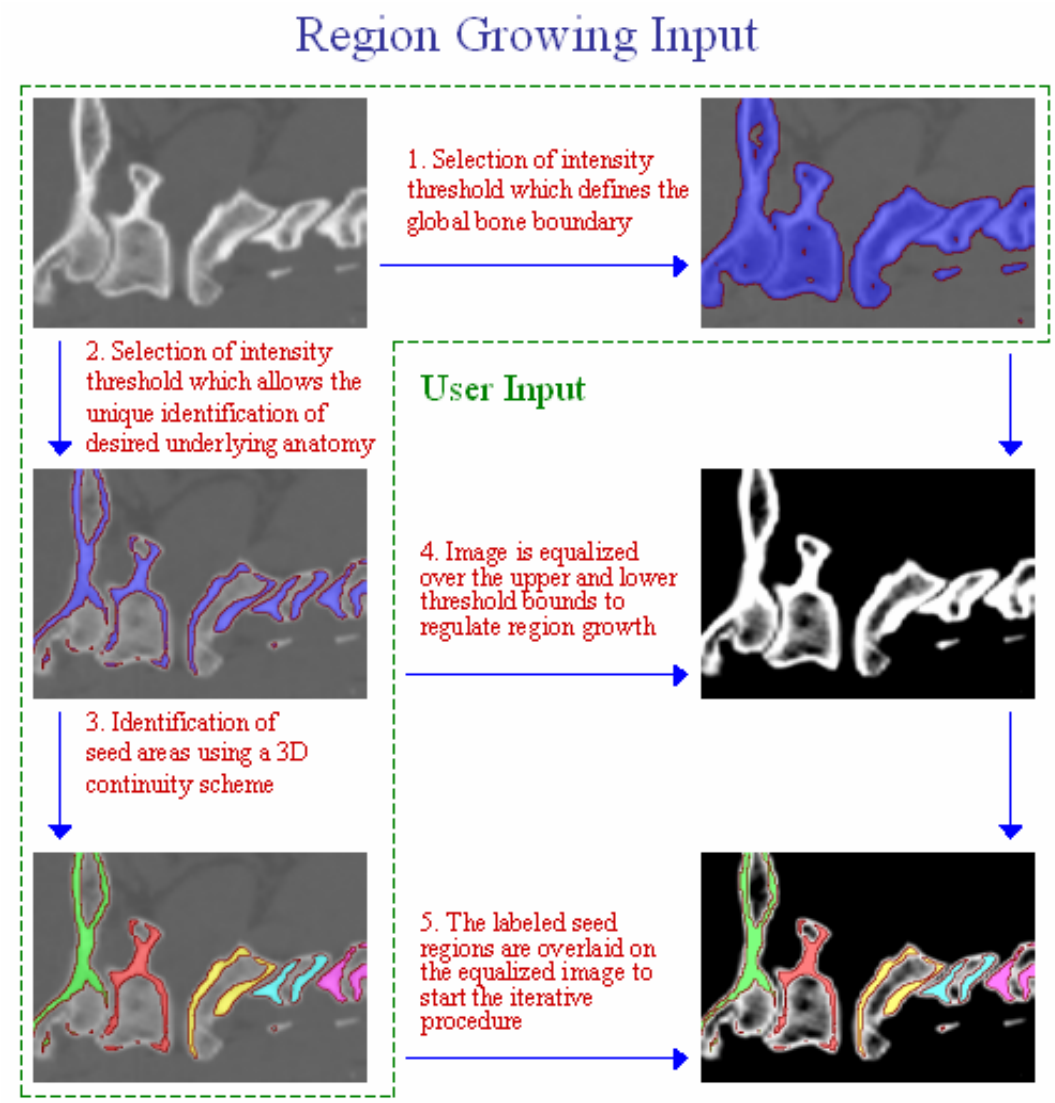

Figure 4-5. Algorithm input processing sequence. The steps requiring user input are outlined in green. The input image is histogram equalized over the boundary threshold ranges to regulate growth while preserving the geometric boundary conditions.

dilations per threshold step. This determines the rate at which the available threshold-bound area will fill. Through repeated trials, a couple of basic trends were noted. Generally, a higher number of threshold divisions produced a more accurate segmentation. Secondly, the number of dilations required to properly fill a threshold step varied widely. In order to regulate this growth, an analysis was performed. It was determined that an intensity histogram could be used to determine the number of voxels

Figure 4-6. Region growing. Once the input image is prepared, the algorithm instructs the labeled kernel regions to grow incrementally until the bone-defining threshold is filled.

Object 4-2 .Region growing movie (object2-success.mpg, 25.2 MB)

available to be filled per threshold step. Accordingly, the histogram can be equalized and applied to the input image to even the growth. This step has been added to the algorithm input processing sequence as shown in Figure 4-5. Once the preparation steps have been accomplished, the iterative region-growing algorithm is activated and continues until the total bone volume is assigned a label (Figure 4-6). The completed label set then defines the segmentation for the original image.

Before the segmented dataset is used for a surgical application, it must be verified. This involves overlaying the segmented dataset on the original CT image for visual comparison. Manual post-processing steps may be necessary if the segmentation does not succeed in the proper isolation of adjacent vertebral bodies. The same tools mentioned for manual segmentation are used for this post-processing step.

Algorithm Enhancement

A few avenues were investigated to improve this region-growing method of image segmentation. The first was an examination of growth characteristics, and how they could be better controlled to suit the application. The second involves the modification of image information to produce more auspicious growth patterns. The third tests the

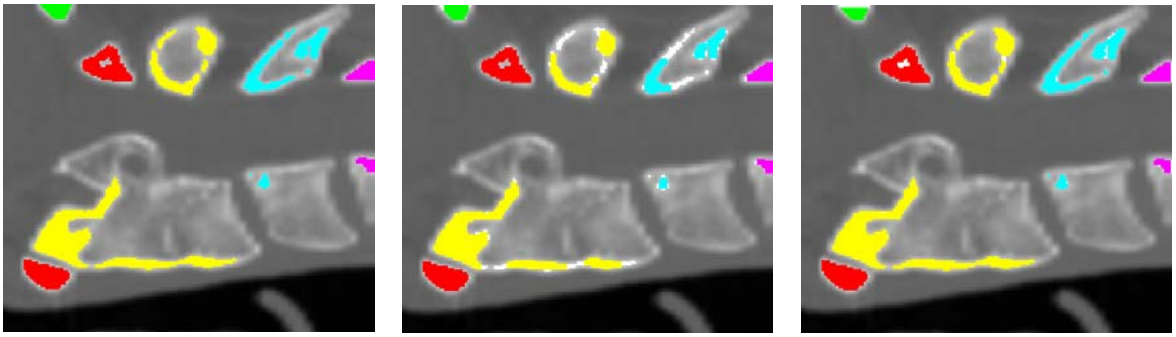

Figure 4-7. Open and close operators are used to alter the evolution of the growth algorithm. On the left is the original image. In the center is an image with the open operator applied. The white regions highlight the areas removed by the operation. On the right is the image with a close operator applied. The white regions are the areas filled by the operation.

inclusion of user-defined input information in addition to the initial kernel-laden segmentation dataset.

During the region growing operations, the kernels tend to expand in finger-like fashion until they merge with other anatomically self-similar sections. The problem arises when one of these extends to a piece of unrelated anatomy, essentially allowing it to falsely populate. Two shape filters were inserted into the algorithm to modify growth behavior. One, the morphological open operator, was used to curb the extension of thin tendrils to false anatomy (Figure 4-7). The other, the morphological close operator, was used to facilitate growth into self-same anatomy. Orthogonal (6-connected) and diagonal (26-connected) structures were both tested. The trials with an orthogonal open operator applied tended toward the more favorable result by restricting finger growth, even though it also inhibited movement between self-same regions. The close operator, however, created too many false bridges to be effective.

Image Variations

Several variations of the input image were presented to the region-growing algorithm to determine if they could be used to improve the quality of the segmentation.

In each of these trials, the kernels were defined based on the altered image and allowed to progress using the iterative scheme. However, the bone-defining threshold value selected using the original image was retained and acts as a barrier for growth. This was done because the intensity threshold is a good approximation of bone edge and these image variations would undoubtedly result in less appealing threshold-based boundary approximations.

First, it was thought that a smoothing filter applied to the original CT image would weaken the intensity of anatomical cross-links and promote a more uniform, centralized geometry from which to evolve. The resulting anatomical edges were very clean and smooth, yet inappropriately labeled. This was attributed to smeared boundaries which lessened the inhibition of the kernels to leak into false anatomy.

Next, the region-growing algorithm was tested on the gradient of the input image. The gradient is a common image filter that is used to highlight boundaries (Figure 4-8). A first note is that gradient kernels are much more disconnect at higher thresholds which makes the index assignment phase more laborious. The resulting segmentations showed similar success to those of the unaltered image, however, the borders were choppy and the areas of bleed-through were relocated. One noticeable property of the gradient image is that there are high intensity rings around the periphery of the vertebrae. This leaves the interior relatively uniform at a low intensity value. Therefore, an inverted gradient image would leave a strong centralized bulk kernel to feed into the iterative algorithm (Figure 4-8). The kernel assignment phase for this image variation, as predicted, was simple. The large internal areas were easy to label. The resulting segmentation, however, left

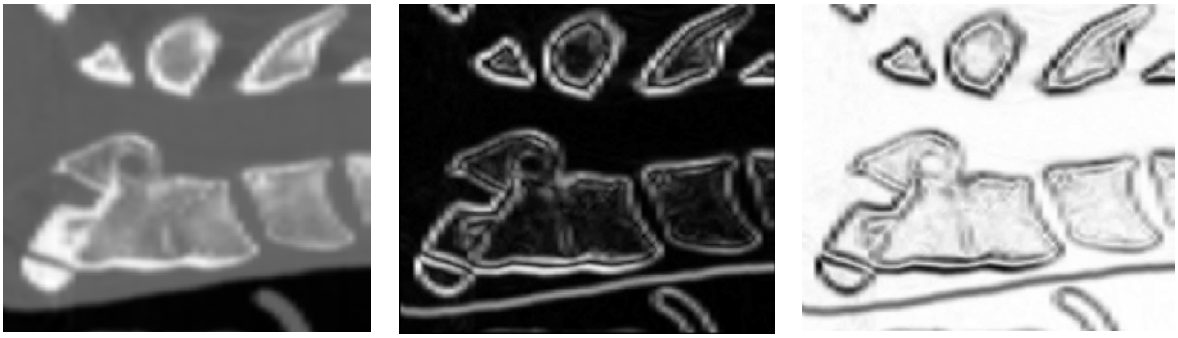

Figure 4-8. Unique properties of the input CT image may be isolated or combined to produce a more favorable segmentation. On the left is the original CT scan. In the center, the gradient of the image highlights boundaries. The inverse of the gradient (right), in contrast, highlights the interior of the image.

unpredictable, jagged boundaries. Example segmentations of the mentioned image variations are shown in Figure 4-9.

The previous image variation tests, individually, did not provide wholly beneficial results. However, they each revealed advantageous properties which could be utilized to improve segmentations. Since this region-growing scheme is based on global intensity patterns, the various images were blended or combined to form a composite image with more suitable qualities. The most prosperous combination was a blend of original CT and inverse gradient images. The inverse gradient image served to add a more centralized interior which would quickly unify distant pieces of the vertebra. The original CT image contribution served to temper the erratic boundary formation that the inverse gradient image produced. This method had the potential to improve the quality of the segmentation, however, the optimal blending ratios were scan specific and tedious to determine. Also, boundary defects still existed which would foil the segmentation.

Prelabeling

The kernel assignment phase of the region growing process has a strong influence on the success of the segmentation. When circumstances require a high threshold level to

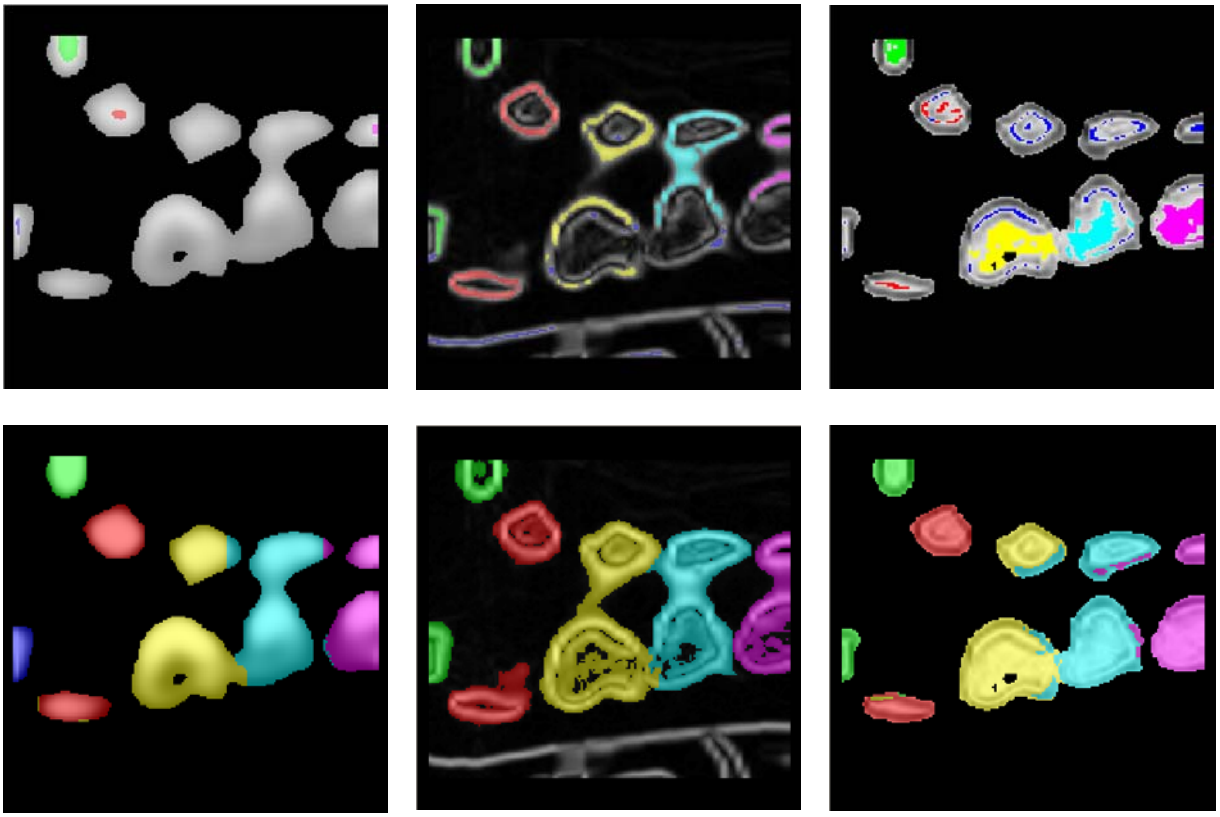

Figure 4-9. Segmentation results using variations on the input CT image. The top row indicates the input kernels and the bottom row displays the result of the algorithm. Examples of smoothed image (left column), image gradient (center column), and inverse gradient (right column) are shown.

create unique kernel elements, the algorithm may be left starved. Specifically, only a limited percentage of voxels in any one piece of anatomy may be available to seed growth. The consequence is a much smaller chance of the correct label reaching all other areas of the anatomy before borders are assigned. The third avenue pursued to improve the region-growing segmentation method is the inclusion of additional user-defined labels. These labels, set during the kernel assignment phase, force the assignment of unassigned voxels that are exposed during region growth. It is expected that these pre-labeled voxels will appropriately spread to adjacent areas and result in a better segmentation. The updated iterative action consists of these steps:

1. The step to a lower threshold and resulting isocontour.
2. The pre-label of unassigned voxels.
3. The constrained orthogonal dilation of the labeled regions.
4. The restriction of the labeled voxels to the area outlined by the isocontour.

The goal is to find a specific labeling technique that makes a strong contribution to segmentation while requiring a minimal amount of user input. The two methods tested were 2D image slice and sphere labeling.

The CT image slice was decided as a good candidate for pre-labeling because it is a straightforward manual segmentation process. Sagittal image slices were chosen for the pre-labeling study because the vertebrae are fairly easy to visualize and a medial slice is likely to contain all of the boney sections. Between one and five slices were tested. Generally, a sparse kernel model or disruptive pathology forces the user to seek additional slices to label. The first problem noted is the pre-labeled slices tend to have little influence on adjacent axial growth patterns unless they happen to cross the labeled sagittal plane. The response was to attempt to position the pre-labeled slice over critical contact areas, such as the transarticular facets. This moderately improved the output; however, the determination of the most effective slice is a tedious process that often requires frequent references to adjacent sagittal image slices.

In response to the inability of the 2D slice pre-labeling to effect growth in the axial direction, sphere labeling was developed. This is simply the creation of spherical regions given a user-defined point, radius, and label. Remarkably, this method showed little improvement in segmentation quality. The reason is that kernels tended to develop in a ring fashion a bit inside the edge of the vertebra before progressing to the border. Therefore, in order to effect kernel growth, the sphere must extend very close to the vertebral boundary. The definition of a sphere in this manner, however, requires strict

attention to adjacent slices to assure that there is no violation into adjacent anatomy by the sphere. Any such occurrence would be very defeating to the overall result. The time-consuming verification process for sphere placement makes this avenue discouraging and impractical.

Instrumentation

The segmentation methods discussed rely on intensity profiles of the bones. When surgical instrumentation such as rods, screws, and wires are implanted into the patient, though, this profile is severely disrupted. Metallic implants cause auras of high intensity where they are positioned and serve as bridges incorrectly linking anatomy. Fortunately, the intensities produced are much higher than any occurring in the natural bone profile. This allows an intensity threshold mask to subtract the instrumentation out of the region growing algorithm. The iterative sequence is similar to pre-labeling; void regions are created in which the dilating kernels cannot travel. After a couple of tests, this method of instrumentation treatment proved very successful. The algorithm was able to navigate around the instrumentation and produce a segmentation (Figure 4-10). The only modification that was appended was an additional dilation or two of the instrumentation region. The first reason is that partial volume effect alters bone profile slightly outside the actual geometry of the instrumentation. The second reason is that implants that have resided in the body for a lengthy period of time will likely have bone growth that alters the natural bone profile.

Results

The process of segmentation with the region-growing algorithm was tested on a total of 31 clinical, high-resolution cervical CT scans intended for spinal surgery. Success was gauged by the amount of effort necessary to achieve an accurate

Figure 4-10. Instrumentation subtraction. The instrumented area (purple) was successfully removed from the region growing algorithm, severing the false bridge and allowing the region-growing algorithm to naturally progress.

visualization of the segmented vertebrae. Note that this may be more stringent than what is required for any specific registration method. Primarily, success was measured by the amount of image slices needing corrective manual post-processing alterations, with an awareness of input effort and algorithm runtime. There were 19 (61%) of the scans that required no or minimal post-processing. There were 7 (23%) that required manual processing, yet still remained a significant improvement over manual methods. The segmentations that showed no significant improvement over manual segmentation amounted to 5 (16%). The algorithm run time was roughly 3-5 minutes (Pentium 4 at 1.4 MHz) for 3-5 vertebral levels. The array size was a very significant factor in runtime since the routines operate on full rectangular matrices. Optimization schemes would significantly reduce processing time.

There are a couple of key factors that affected the outcome of this segmentation method. First, the quality of the kernel regions had a significant contribution. If the seed

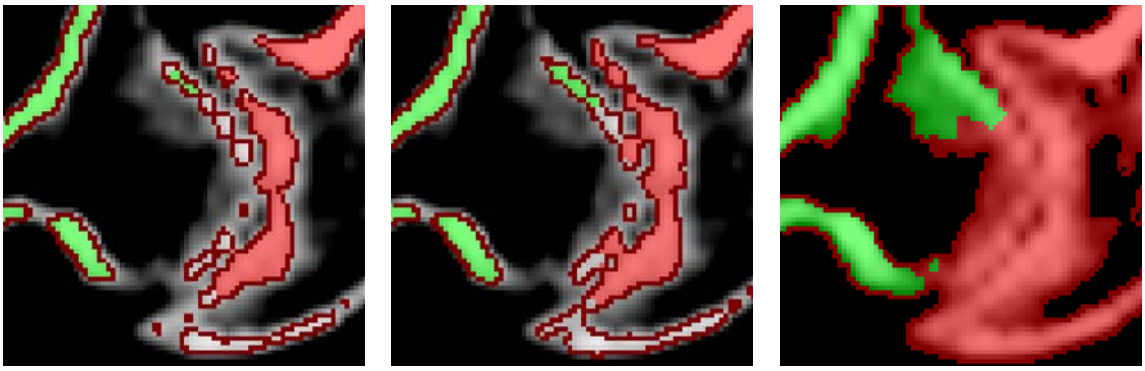

Figure 4-11. Segmentation failure: boundary leakage. Seed regions defined distant to an anatomical boundary must carry the accurate label to the boundary before a competing label invades. The center image depicts the boundary leakage with the final segmentation on the right.

Object 4-3. Segmentation failure: boundary leakage movie (object3-failureleakage.mpg, 25.8 MB).

regions were small and disconnected, then considerable weight was placed on the ability to branch and interconnect properly. The failure of those actions leads to boundary leakage (Figure 4-11). Threshold models that are created at a very high intensity produce these low quality kernels. Areas of high intensity contact will force the seed region selection to a high threshold. Severe posturing can change the proximity of the bones and bring high intensity areas into contact, usually at the transarticular surfaces. The second primary component to success of the segmentation process was image resolution. The cervical boundary profile requires a high level of detail. A low-resolution image will increase partial volume effects yielding incomplete intensity patterns and increase opportunity for failure. A similar effect is seen when posturing compresses the space between vertebrae. The last main ingredient to success of the segmentation was the condition of the bone intensity profile. Some patients had comparatively lower bone contrast with soft tissue, which created threshold models with significant gaps. These gaps effected the ability of the anatomic labels to spread. The intensity profile was also

disrupted with pathology such as fusion. The most severe result observed was kernels that dilated with little distinction between correct anatomical locations (Figure 4-12).

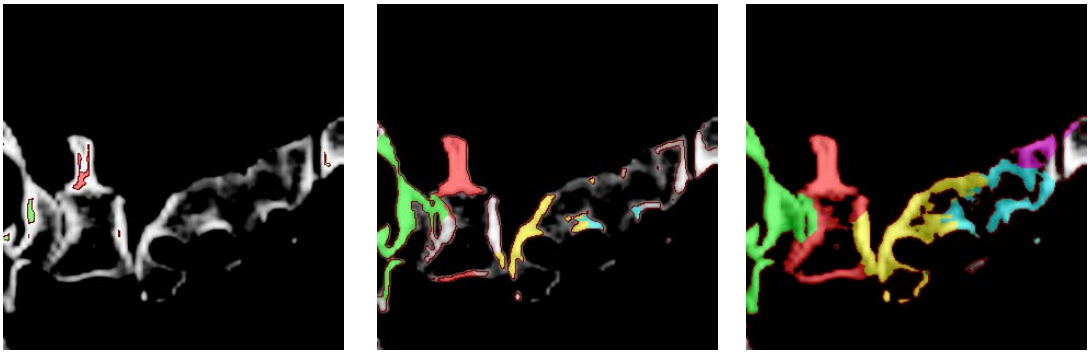

Figure 4-12. Segmentation failure: weak bone intensity profile. A CT scan that has low bone to soft tissue contrast or pathology that disrupts the clarity of the bone profile does not allow the dilating kernel regions to properly meet at the vertebral boundaries.

Object 4-4. Segmentation failure: weak bone intensity profile movie (object4-failureweak.mpg, 28.6 MB).

CHAPTER 5
CONCLUSIONS AND FUTURE WORK

A semi-automatic method has been developed to expedite the production of segmented CT images beyond that of manual methods. It promotes accurate visualization of bony components and encourages the use of image guidance with surgeries involving multiple levels. This region-growing hysteresis method has a couple of notable benefits. First, it is a relatively simple procedure to implement. It has no pre-defined datasets to build or maintain. It works well with fracture pathology. Plus, instrumentation can be easily subtracted without significantly affecting the performance of the algorithm.

In contrast, this segmentation process has a few limitations. One is the potential for lengthy user intervention, such as during kernel assignment or post-processing, if required. Second, the threshold model does not completely define what is considered the interior of the anatomy. Some of the interior regions have the same intensity as the surrounding soft tissue and so are not included in the model. Region filling methods may be used to fill the interior; however, this is not a thorough solution. Note that image guidance registration may not require a segmentation with complete internal labeling to be successful. Third, the algorithm is susceptible to small intensity defects, especially those that bridge adjacent bones.

A couple of directions exist which have the potential to improve upon the developed segmentation method, remedying some of its shortcomings. The region-growing hysteresis method relies strictly on global voxel-based operations. Additional

elements, such as boundary or region-based components, if integrated, may provide the additional information necessary to improve the segmentation. For example, an elastic boundary concept and internal region statistics could be used to make the method more robust to noise and point defects. Also, any statistical geometric profile could be used to guide the evolving boundary to its correct location. One cost to these options, however, is the maintenance of a potentially cumbersome dataset. For the addition of these parameters to a method to be successful, it must be considered how well-suited the algorithm is to incorporate the information. For example, in the current algorithm, choosing a penalty term to curvature for a boundary-based representation may be difficult with regard to the chaotic nature of how the regions intersect. After review, we decided to pursue a method that was perhaps more suited to the integration of the mentioned voxel, boundary, and region-based components.

Leventon (2000) presents a noteworthy approach to segmentation by calling upon curvature models, prior shape models, and statistical image-surface relationships to perform segmentations. One distinct advantage to this method is it blends these techniques into a level set framework, allowing the boundary to evolve over complex topography. This particular algorithm uses a signed distance map as the level set surface representation. It has several advantages. First, boundary and region-based representations are implicitly contained in the signed distance function. Second, the notion of "holes" in the boundary is seamlessly contained in the representation; a small "hole" in the object does not create an entire new set of features in the object. Third, it is robust to slight misalignment of features. One interesting aspect of Leventon's method is the use of intensity/distance-to-boundary distributions. These profiles are incorporated as

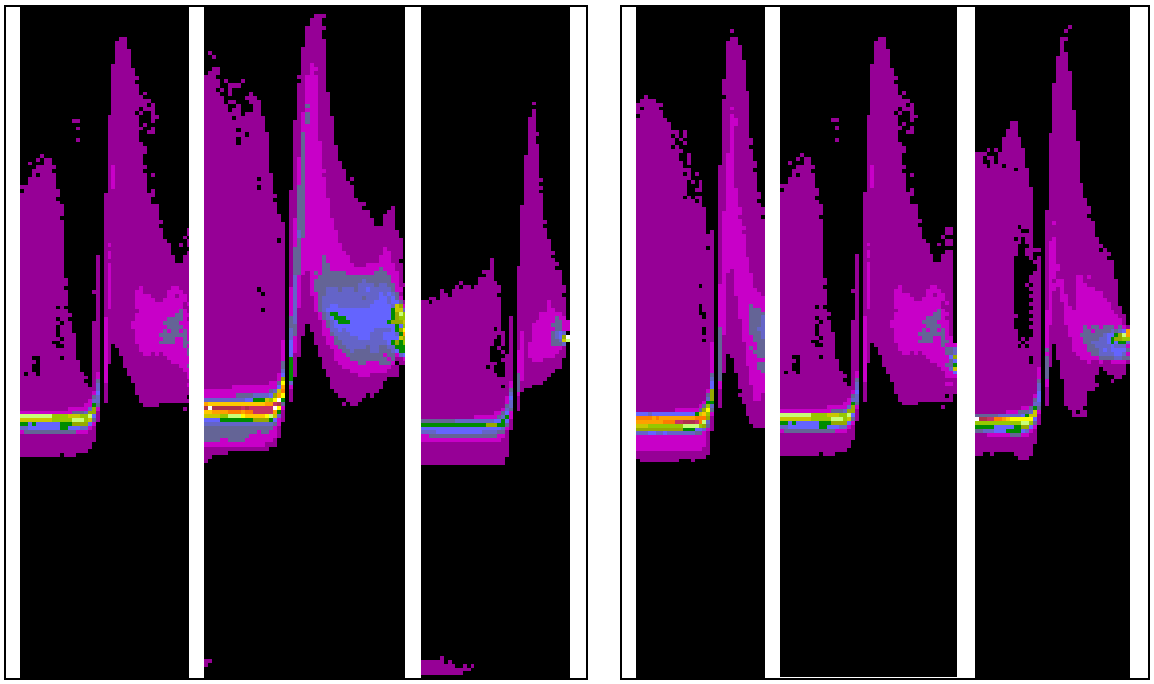

Figure 5-1. Intensity/distance-to-boundary profiles. The vertical axis is grayscale value, with the top having the highest intensity. The horizontal axis is distance from object boundary (vertical black stripe). Right indicates interior distance from the boundary and left indicates the exterior distance from the boundary. Color indicates frequency of pairs. The first three images on the left show the profile of vertebra C2 from three random individuals. The three images on the right show three different levels, C1, C2, and C3 (center to right) from one individual.

training sets into a model that makes boundary estimations. Specifically, with this particular adaptation of the level set framework, distance maps are used to describe the higher-dimensional surface. The distribution of surface-intensity pairs, then, is a voxel match of the morphological distance map of the segmented object, and it's associated CT image value. A cursory examination of intensity/distance-to-boundary distributions of spinal images was performed to determine if these profiles would be useful for segmentation. Six cervical spine CT images were manually segmented, and vertebrae from C1 to C4 were individually isolated. From that, 31 distinct distance-intensity maps were created (Figure 5-1). The images were interpolated to equivalent square voxels

before the creation of the maps, so that the distance maps were accurate, and could be compared. Also, no images were selected that contained significant intensity artifacts.

A visual comparison of the distance-intensity profiles shows many similarities. It is anticipated that with some training, an algorithm could use this information to properly locate a boundary. The next step would be to formulate a probability density model and implement it into a segmentation routine, most likely that presented by Leventon (2000).

LIST OF REFERENCES

D. Adalsteinsson & J. Sethian, "The fast construction of extension velocities in level set methods," *The Journal of Computational Physics*, 148:2-22, 1999.

L.P. Amiot, K. Lang, M. Putzier, H. Zippel & H. Labelle, "Comparative results between conventional and computer-assisted pedicle screw installation in the thoracic, lumbar, and sacral spine," *Spine*, 25(5):606-614, 2000.

R. Assaker, N. Reyns, M. Vinchon, X. Demondion & E. Louis, "Transpedicular screw placement: image-guided versus lateral-view fluoroscopy: in vitro simulation," *Spine*, 26(19):2160-2164, 2001.

M.S. Austin, A.R. Vaccaro, B. Brislin, R. Nachwalter, A.S. Hilibrand & T.J. Albert, "Image-guided spine surgery: a cadaver study comparing conventional open laminoforaminotomy and two image-guided techniques for pedicle screw placement in posterolateral fusion and nonfusion models," *Spine*, 27(22):2503-2508, 2002.

J. Bangham & S. Marshall, "Image and signal processing with mathematical morphology," *Electronics & Communication Engineering Journal*, 10:117-128, 1998.

S. Beucher & F. Meyer. "The morphological approach of segmentation: the watershed transformation," *Mathematical Morphology in Image Processing*, E. Dougherty, Ed., chapter 12, 433-481, Marcel Dekker, New York, 1992.

A. Blake & M. Isard, *Active Contours*, Springer-Verlag, London, 1998.

A. Bleau & L. Leon, "Watershed segmentation and region merging," *Comp. Vision and Image Understanding*, 77(3):317-370, 2000.

G. Boehm, C.J. Knoll, V.G. Colomer, M. Alcañiz-Raya & S.E. Albalat, "Three dimensional segmentation of bone structures in CT images," *In Proc. SPIE, Medical Imaging*, 3661:277-286, 1999.

V. Caselles, F. Catte, T. Coll & F. Dibos, "A geometric model for active contours in image processing," *Numerische Mathematik*, 66:1-31, 1993.

V. Casellas, R. Kimmel & G. Sapiro, "Geodesic active contours," *Int. Journal of Computer Vision*, 22(1):61-79, 1997.

A. Chakraborty, L. Staib & J. Duncan, "Deformable boundary finding in medical images by integrating gradient and region information," *IEEE Trans. Medical Imaging*, 15(6):859-870, 1996, http://noodle.med.yale.edu/staib/papers/chakrabtmi.pdf, last accessed March 22, 2005.

T. Chan & L. Vese, "An active contour model without edges," *Scale-Space Theories in Computer Vision*, LNCS 1682:141-151, Springer-Verlag, London, 1999.

Y. Chen, S. Thiruvenkadam, H. Tagare, F. Huang, D. Wilson & E. Geiser, "On the incorporation of shape priors into geometric active contours," *In Proc. 1st IEEE Workshop on Variational and Level Set methods in Computer Vision*, pp.145-152, IEEE, Los Alamitos, 2001.

G. Christensen, R. Rabbitt & M. Miller, "Deformable templates using large deformation kinematics," *IEEE Trans. on Image Processing*, 5(10):1435-1447, 1996.

G. Christensen, "Consistent linear-elastic transformations for image matching," *In Proc. Information Processing in Medical Imaging*, 1613:224-237, Springer-Verlag, London, 1999.

K. Cleary, "Technical Requirements for Image-Guided Spinal Procedures", Workshop Report, April 1999, ISIS Center, Department of Radiology, Georgetown University Medical Center, Washington, D.C., http://www.visualization.georgetown.edu/Spine_Workshop/report/report.html, last accessed March 22, 2005.

L. Cohen, "On active contour models and balloons," *CVGIP: Image Understanding*, 53(2):211-218, 1991.

T. Cootes & C. Taylor, "Active shape models – 'smart snakes'," *In Proc. British Machine Vision Conference*, pp. 266-275, Springer-Verlag, London, 1992.

T. Cootes, C. Taylor, D. Cooper & J. Graham, "Training models of shape from sets of examples," *In Proc. British Machine Vision Conference*, pp. 9-18, Springer-Verlag, London, 1992.

T. Cootes, C. Taylor, D. Cooper & J. Graham, "Active shape models - their training and application," *Computer Vision and Image Understanding*, 61:38-59, 1995.

T. Cootes & C. Taylor, "A mixture model for representing shape variation," *Image and Vision Computing,* 17(8):567-574, 1999.

T. Cootes, C. Beeston, G. Edwards & C. Taylor, "A unified framework for atlas matching using active appearance models," *In Proc. Information Processing in Medical Imaging*, 1613:322-333, Springer-Verlag, London, 1999.

T. Cootes, G. Edwards & C. Taylor, "Active appearance models," *5th European Conference on Computer Vision*, 2:484-498, Springer-Verlag, Berlin, 1998.

D. Cremers, F. Tischhäuser, J. Weickert & C. Schnörr, "Diffusion snakes: introducing statistical shape knowledge into the Mumford-Shah functional," *International Journal of Computer Vision*, 50(3): 295-313, 2002.

O. Cuisenaire, "Distance transformations: fast algorithms and applications to medical image processing," Ph.D. thesis, Université cotholique de Louvain, Laboratoire de Telecommunications et Teledetection, 1999, http://www.tele.ucl.ac.be/PEOPLE/OC/these.pdf, last accessed March 22, 2005.

C. Davatzikos & J. Prince, "Convergence analysis of the active contour model with applications to medical images," *In Proc. SPIE*, *Visual Communications and Image Processing*, 1818:1244-1255, 1992.

A. Dempster, N. Laird & D. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Royal Statistical Soc.*, Ser. B, 39(1):1-38, 1977.

C.A. Dickman, "Point of view," *Spine*, 25:2674, 2000.

E. Dougherty, *Mathematical Morphology in Image Processing*, M. Dekker, New York, 1993.

I. Dryden & K. Mardia, *Statistical Shape Analysis*, Wiley, Chichester, 1998.

K.T. Foley, D.A. Simon & Y.R. Rampersaud, "Virtual fluoroscopy: computer-assisted fluoroscopic navigation," *Spine*, 26:347-351, 2001.

G.E. Healthcare Technologies, Waukesha, WI, U.S.A., www.gehealthcare.com/index.html, 2004, last accessed March 22, 2005.

S. Geman & D. Geman, "Stochastic relaxation, Gibbs, and the Bayesian restoration of images," *IEEE Trans. on PAMI*, 6(6):721-741, November, 1984.

N. Glossop & R. Hu, "Assessment of vertebral body motion during spine surgery," *Spine*, 22:903-909, 1997.

P. Golland, W.E.L. Grimson & R. Kikinis, "Statistical shape analysis using fixed topology skeletons: corpus callosum study," *In Proc. Inf. Proc. in Med. Imaging* (IPMI), 1613:382-387, Springer-Verlag, 1999.

P. Golland, W.E.L. Grimson, M. Shenton & R. Kikinis, "Small sample size learning for shape analysis of anatomical structures," *In Proc. MICCAI' 2000: Medical Image Computing and Computer-Assisted Intervention*, LNCS 1935:72-82, Springer-Verlag, 2000.

Gray's Anatomy, http://education.yahoo.com/reference/gray/, last accessed March 22, 2005.

M. Jones & T. Poggio, "Multidimensional morphable models: a framework for representing and matching object classes," *International Journal of Computer Vision*, 29(2):107-131, 1998.

T. Kaneko, L. Gu & H. Fujimoto, "Abdominal organ recognition using 3D mathematical morphology," *In Proc. 15th International Conference on Pattern Recognition*, 2:263-266, 2000.

M. Kass, A. Witkin & D. Terzopoulos, "Snakes: active contour models," *Int. Journal on Computer Vision*, 1:321-331, 1987.

A. Kichenassamy, A. Kumar, P. Olver, A. Tannenbaum & A. Yezzi, "Gradient flows and geometric active contour model," *In Proc. IEEE Int. Conf. Computer Vision*, pp. 810-815, IEEE, Los Alamitos, 1995.

R. Kikinis, S.K. Warfield & C.F. Westin, "High performance computing (HPC) in medical image analysis (MIA) at the surgical planning laboratory (SPL)," *In Proc. High Performance Computing Asia*, pp. 290-297, Singapore, 1998, http://splweb.bwh.harvard.edu:8000/pages/papers/kikinis/hpc/HPCfinal.pdf, last accessed March 31, 2005.

S. Lavallée, J. Troccaz, P. Sautot, B. Mazier, P. Cinquin, P. Merloz & J.P. Chirossel, "Computer-assisted spinal surgery using anatomy-based registration," *Computer-Integrated Surgery: Technology and Clinical Applications*, R. Taylor, S. Lavallee, G. Burdea & R. Mosges, Eds., pp. 425-449, MIT Press, Cambridge, 1996.

R. Leahy, T. Herbert & R. Lee, "Application of Markov random fields in medical imaging," *Information Processing in Medical Imaging*, J. Llacer & D. Ortendahl, Eds., pp. 1-14, Wiley-Liss, NY, 1991.

M. Leventon, "Statistical models for medical image analysis," Ph.D. Thesis, Massachusetts Institute of Technology, 2000.

J.A. Little, D.L.G. Hill & D.J. Hawkes, "Deformations incorporating rigid structures," *Computer Vision and Image Understanding*, 66(2):223-232, 1997.

C. Lorenz & N. Krahnstover, "3D statistical shape models for medical image segmentation," *In Proc. Int. Conf. 3-D Imaging and Modeling*, pp. 414-423, IEEE, Los Alamitos, 1999.

R. Malladi, J. Sethian & B. Vemuri, "Shape modeling with front propagation: a level set approach," *IEEE Trans. Pattern Anal. Machine Intell.*, 17:158-175, 1995.

G. Matheron, *Random Sets and Integral Geometry*, John Wiley & Sons, New York, 1975.

T. McInernery & D. Terzopoulos, "Medical image segmentation using topologically adaptable surfaces," *In Proc. Conf. Computer Vision, Virtual Reality, and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery (CVRMed-MRCAS)*, 1205:23-32, Springer-Verlag, Berlin, 1997.

T. McInerney & D. Terzopoulos, "Deformable models in medical image analysis: A survey," *Medical Image Analysis*, 1(2):91-108, 1996, http://mrl.nyu.edu/~dt/papers/mia96/mia96.pdf, last accessed March 22, 2005.

Medtronic Surgical Navigation Technologies, StealthStation®, Louisville, CO, www.stealthstation.com, last accessed March 22, 2005.

F. Meyer & S. Beucher, "Morphological segmentation," *Journal of Visual Comm. and Image Rep.*, 1(1):21-46, 1990.

M. Miller, A. Banerjee, G. Christensen, S. Joshi, N. Khadeja, U. Grenander & L. Matejic, "Statistical methods in computational anatomy," *Statistical Methods in Medical Research*, 6(3):267-99, 1997.

D. Mumford & J. Shah, "Optimal approximation by piecewise smooth functions and associated variational problems," *Comm. Pure and Applied Math.*, 42:577-684, 1989.

A. Neumann & C. Lorenz, "Statistical shape model based segmentation of medical images," *Computerized Medical Imaging and Graphics*, 22:133-141, 1998.

H. Nguyen, M. Worring & R. van den Boomgaard, "Watersnakes: energy-driven watershed segmentation," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 25(3): 330-342, 2003.

Odin Medical Technologies, Inc., Newton, MA, www.odinmed.com, last accessed March 22, 2005.

K. Ohmori, Y. Kawaguchi, M. Kanamori, H. Ishihara, H. Takagi & T. Kimura, "Image-guided anterior thoracolumbar corpectomy: a report of three cases," *Spine*, 26:1197-1201, 2001.

S. Osher & J. Sethian, "Fronts propagating with curvature dependent speed: Algorithms based on Hamilton-Jacobi formulations," *Journal of Computational Physics*, 72:12-49, 1988.

N. Paragios & R. Deriche, "Geodesic active regions for supervised texture segmentation," *In Proc. IEEE Int. Conf. Comp. Vision*, pp. 926-932, IEEE, Los Alamitos, 1999.

A. Pentland & S. Sclaroff, "Closed-form solution for physically based shape modeling and recognition," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(7):715-729, 1991.

T.M. Peters, "Image-guided surgery: from x-rays to virtual reality," *Comput. Methods Biomech. Biomed. Engin.*, 4(1):27-57, 2000.

R. Pichumani, "Construction of 3D geometric models for segmentation and visualization of cervical spine images," Ph.D. thesis, Stanford University, November 1997.

Philips Medical Systems, Achieva I/T Interventional MR, Bothell, WA, http://www.medical.philips.com, last accessed March 22, 2005.

J. Prewitt, "Object enhancement and extraction," *Picture Processing and Psychopictorics*, B. Lipkin & A. Rosenfeld, Eds., pp. 75-149, Academic Press, New York, 1970.

Y.R. Rampersaud, K.T. Foley, A.C. Shen, S. Williams & M. Solomito, "Radiation exposure to the spine surgeon during fluoroscopically assisted pedicle screw insertion," *Spine*, 25:2637-2645, 2000.

Y.R. Rampersaud, D.A. Simon & K.T. Foley, "Accuracy requirements for image-guided spinal pedicle screw placement," S*pine*, 26:352-359, 2001.

L. Roberts, "Machine perception of three-dimensional solids," *Optical and Electro-optical Information Processing*, J. Tippett, Ed., pp. 159-197, MIT Press, Cambridge, 1965.

S. Sclaroff & A. Pentland, "Modal matching for correspondence and recognition," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 17(6):545-561, 1995.

J. Serra, *Image Analysis and Mathematical Morphology*, Academic Press, New York, 1982.

J. Sethian, *Level Set Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision, and Material Science*, Cambridge University Press, New York, NY, 1996.

J. Sethian, "A fast marching level set method for monotonically advancing fronts," Proc. Nat. Acad. Sci., 93(4):1591-1595, 1996.

J. Sethian, "Fast marching methods," *SIAM Review*, 41(2):199-235, 1999.

Siemens Medical Solutions USA, Siemens SIREMOBIL Iso-C^{3D}, Inc., Iselin, NJ, http://energized.esiemenshealthcare.com/SIREMOBIL/iso3d.htm, last accessed March 22, 2005.

I. Sobel, "Camera models and machine perception," Technical Report AIM-21, Stanford Artificial Intelligence Lab, Palo Alto, California, May 1970.

L. Staib & J. Duncan, "Boundary finding with parametrically deformable models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14:161-175, 1992.

G. Szekely, A. Kelemen, C. Brechbuhler & G. Gerig, "Segmentation of 2-D and 3-D objects from MRI volume data using constrained elastic deformations of flexible Fourier contour and surface models," *Med Image Anal.*, 1(1):19-34, 1996.

H. Tek & B. Kimia, "Image segmentation by reaction diffusion bubbles," *Proc. Int. Conf. Computer Vision*, pp. 156-162, IEEE, Los Alamitos, 1995.

A.R. Vaccaro & K. Singh, "Computer-assisted image-guided surgery: application in cervical spine surgery," *Current Opinion in Orthopaedics*, 12:245-250, 2001.

L. Vincent & P. Soille, "Watershed in digital spaces: an efficient algorithm based on immersion simulations," *IEEE Trans. on PAMI*, 13(6),583-598, 1991.

Y. Wang & L. Staib, "Boundary finding with correspondence using statistical shape models," *In Proc. IEEE Conf. Comp. Vision and Patt. Recog.*, pp. 338-345, IEEE, Los Alamitos, 1998.

Y. Wang & L. Staib, "Elastic model based non-rigid registration incorporating statistical shape information," *In Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, LNCS 1496:1162-1173, Springer-Verlag, London, 1998.

S. Warfield, A. Robatino, J. Dengler, F. Jolesz & R. Kikinis, "Nonlinear registration and template driven segmentation," In *Brain Warping*, Arthur W. Toga, Ed., Ch.4:67-84, Elsevier Science & Technology Books, St. Louis, 1998.

A. Weidner, M. Wähler, S.T. Chiu & C.G. Ullrich, "Modification of C1-C2 transarticular screw fixation by image-guided surgery," *Spine*, 25:2668-2674, 2000.

W.C. Welch, B.R. Subach, I.F. Pollack & G.B. Jacobs, "Frameless stereotactic guidance for surgery of the upper cervical spine," *Neurosurgery*, 40:958-964, 1997.

C.F. Westin, S. Warfield, A. Bhalerao, L. Mui, J. Richolt & R. Kikinis, "Tensor controlled local structure enhancement of CT images for bone segmentation," *Int. Conf on Medical Image Computing and Computer-Assisted Intervention (MICCAI'98)*, LNCS 1496:1205-1212, Springer-Verlag, London, 1998.

C.F. Westin, J. Richolt, V. Moharir & R. Kikinis, "Affine adaptive filtering of CT data," *Medical Image Analysis*, 4:161-177, 2000, http://splweb.bwh.harvard.edu:8000/pages/papers/westin/media2000/westinMedia2000.pdf, last accessed March 22, 2005.

A. Yezzi, A. Tsai & A. Willisky, "A statistical approach to image segmentation for bimodal and trimodal imagery," *In Proc Int. Conf. Comp. Vision*, pp. 898-903, IEEE, Los Alamitos, 1999.

A.S. Youkilis, D.J. Quint, J.E. McGillicuddy & S.M. Papadopoulos, "Stereotactic navigation for placement of pedicle screws in the thoracic spine," *Neurosurgery*, 48:771-779, 2001.

S.C. Zhu & A. Yuille, "Region competition: unifying snakes, region growing, and Bayes/MDL for multiband image segmentation," *IEEE Trans. on PAMI*, 18(9): 884-900, 1996.

BIOGRAPHICAL SKETCH

Matthew James Williams, born April 13, 1975, received a Bachelor of Science in Engineering in mechanical engineering, with concentrations in design and biomedical applications, from the University of Michigan in 1998. During his undergraduate years, he was blessed with many rewarding experiences and opportunities that have helped to shape his life. Some of these include internships at a variety of manufacturing and technical facilities, capped with a cooperative at the Denso Tech Center in 1997. After graduation, Matthew was hired as a project management consultant at ACM, Inc. Throughout his career, he has been cementing his resolve to explore engineering opportunities in medicine. This led him to the University of Florida Biomedical Engineering Program; and in 2000, he migrated to Gainesville.

Now that he has finished the requirements for his master's degree, Matthew has elected to continue taking advantage of the cornucopia of opportunities at the University of Florida through the study of the regenerative medicine aspect of biomedical engineering. Also, Matthew continues the pursuit of the performing arts, nourishing creativity in all aspects of his life.