A System for Surgical Planning and Guidance using Image Fusion and Interventional MR

by

David T. Gering

Submitted to the Department of Electrical Engineering and Computer Science

in partial fulfillment of the requirements for the degree of

Master of Science in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

December 1999

Author . . . .
Department of Electrical Engineering and Computer Science
May 7, 1999

Certified by . . . .
W. Eric L. Grimson
Bernard Gordon Professor of Medical Engineering
Thesis Supervisor

Certified by . . . .
Ron Kikinis
Associate Professor of Radiology, Harvard Medical School
Thesis Supervisor

Accepted by . . . .
Arthur C. Smith
Chairman, Departmental Committee on Graduate Students

A System for Surgical Planning and Guidance using Image Fusion and Interventional MR

by

David T. Gering

Submitted to the Department of Electrical Engineering and Computer Science
on May 7, 1999, in partial fulfillment of the
requirements for the degree of
Master of Science in Computer Science and Engineering

Abstract

In this thesis, we present a computerized surgical assistant whose core functionality is embodied in a software package we call the *3D Slicer*. We describe its system architecture and its novel integration with an interventional Magnetic Resonance (MR) scanner. We discuss the 3D Slicer's wide range of applications including guiding biopsies and craniotomies in the operating room, offering diagnostic visualization and surgical planning in the clinic, and facilitating research into brain shift and volumetric studies in the lab.

The 3D Slicer uniquely integrates several facets of image-guided medicine into a single environment. It provides capabilities for automatic registration, semi-automatic segmentation, surface model generation, 3D visualization, and quantitative analysis of various medical scans.

We formed the first system to augment intra-operative imaging performed with an open MR scanner with this full array of pre- operative data. The same analysis previously reserved for pre- operative data can now be applied to exploring the anatomical changes as the surgery progresses. Surgical instruments are tracked and used to drive the location of reformatted slices. Real-time scans are visualized as slices in the same 3D view along with the pre-operative slices and surface models. The system has been applied in over 20 neurosurgical cases at Brigham and Women's Hospital, and continues to be routinely used for 1-2 cases per week.

Thesis Supervisor: W. Eric L. Grimson
Title: Bernard Gordon Professor of Medical Engineering

Thesis Supervisor: Ron Kikinis
Title: Associate Professor of Radiology, Harvard Medical School

Acknowledgments

This section was the most rewarding portion of my thesis to write. None of this would have been possible without the amazingly strong encouragement I received from so many people at GE Medical Systems in Milwaukee. I begin with a chronological glimpse of those I owe for being here.

I thank Ken Accardi for his support in my development of the patented "cine fix" which first sparked my interest in research. I thank Art Larson and Armin Phoh for their guidance which kindled my desire to pursue a Ph.D. Steve Hushek really fanned the flames with his passion for image-guided surgery. His contagious zeal for the field first led me into GE's open-magnet program, and then his sharing of his fond (and also stressful) memories of his alma mater led me to choose MIT. I sincerely thank Joe Phelan for being the best mentor one could ask for. I certainly would not have pursued this without his unwavering belief that I could do this, and that GE would benefit. Many managers deserve credit for having the vision to see the synergistic benefits of supporting my work at MIT; Dave Mack, Morry Blumenfeld, Guy Rabbat, Jean-luc Duchmin, Steve Patscot, and especially, Madhavi Nunna, believed in me enough to invest in my education and bridge the gap between academic research and commercial products.

The relationships I've formed with the very nice, fun people of the AI Lab are a real blessing. My first contact was a phone call to Joe Smosky, in which he cemented my decision to apply by convincing me that this would be a fun place to be, and that his advisor, Eric Grimson, was outstanding. You were more than right on both counts, Joe, and I thank Eric not only for his great guidance, but his talent for baking delicious desserts. I thank my other advisor at MIT, Sandy Wells, for his devotion to helping with last-minute paper edits, his registration hacking, and his analytical skill in dissecting problems. My fellow grad students are a real joy to be with, and a couple deserve special thanks here. I thank my "boss", Tina Kapur, for her invaluable insights and enjoyable conversations on everything from classes and this thesis, to lab survival tips and internet stocks. I thank Neil Wiesen for the late nights working problem sets together (but not really together, Professor Horn). Furthermore, I just have to thank Anita Flynn for founding the AI Lab Olympics - the greatest event on earth.

In addition to working at the best engineering school in the world, I have been fortunate to have the privilege of developing and implementing my research clinically at the Surgical Planning Lab at Brigham and Women's Hospital. I thank my advisor, Ron Kikinis, for somehow managing to simultaneously be both my harshest critic, and most generous supporter. I thank Ferenc Jolesz for his vision for image-guided therapy and the 3D Slicer. I thank Arya Nabavi for teaching me everything I know about surgery, and mostly for transforming those late nights at work into fun and enjoyable memories.

I most gratefully thank the one person who has most supported me through it all. I understand the sacrifices my wife, Heidi, has made to make this possible, and I cannot adequately express the joy I receive from her loving companionship in sharing this experience.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In this thesis, we present a computerized surgical assistant whose core functionality is embodied in a software package we call the *3D Slicer*. We describe the system architecture and its novel integration with an interventional Magnetic Resonance (MR) scanner. We discuss the wide range of applications in which the 3D Slicer is currently being used at the Surgical Planning Lab at Brigham and Women's Hospital [SPL]. These applications include: guiding biopsies and craniotomies in the operating room, offering diagnostic visualization and surgical planning in the clinic, and facilitating research into brain shift and volumetric studies in the lab. A screen shot of the 3D Slicer is presented in Figure 1-1.

1.1 Motivations

Image-guided surgery seeks to augment the human eye with an enhanced vision of reality. An unassisted surgeon can see the surfaces of exposed tissues, but not internal structures. Image-guided surgery provides X-ray vision of what lies beyond the exposed surfaces, what types of tissue are seen, and what functions the tissues serve. Different types of tissue may be difficult to distinguish with the eye alone, but appear markedly different on certain medical imaging scans. Similarly, tissues that handle critical functions, such as voluntary movements, speech, or vision, appear identical to companion tissue, but can be highlighted by a functional exam.

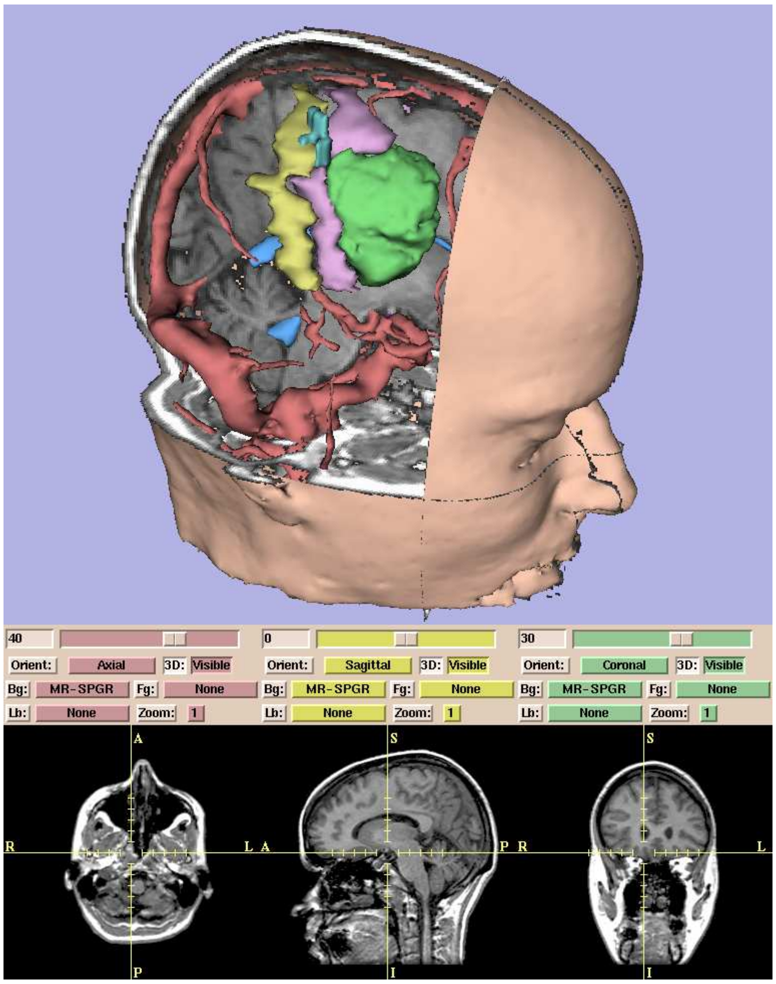

Figure 1-1: The 3D Slicer features a 3D view where 2D slices are integrated with 3D surface models. The 2D slices are also shown in 2D image windows below the 3D view.

We believe better pre-operative visualization leads to better surgical planning. It allows the surgeon to cut smaller access holes, which leads to quicker patient recoveries and shorter hospital stays. In addition, pre-operative guidance allows the surgeon to select an approach trajectory that more safely invades the tissue. This can reduce the cost of the surgery by minimizing the time the surgeon spends planning in the operating room. To complement *pre*-operative planning, *intra*-operative guidance facilitates more complete removal, or resection, of tumor tissue while reducing the chances of harming healthy tissue that surrounds the area under treatment. This suggests an overall improvement in patient outcomes.

1.1.1 Conventional Image-Guided Surgery

An overview of image-guided medicine at some of the leading research institutions can be found at the following sites on the World Wide Web [ETH, MAYO, Montreal, SPL, MGH, JHU, Shadyside, Lueven, UCLA, UCSF, Guys, INRIA, GE, MERL] Conventional image-guided surgery systems present the surgeon with data that was gathered prior to surgery, track surgical instruments within the operating field, and render the tracked devices along with the data.

Data sets are aligned, or registered, with the patient's position at the time of surgery to allow the surgeon to establish a correspondence between the patient lying on the operating table and the images rendered on a nearby computer screen. Image-guided surgery has progressed from attaching cumbersome reference frames to the patient [Galloway1993], to frameless stereotaxy that integrates registration and tracking [Grimson et al.1998]. Registration of pre-operative imagery to patient position is commonly achieved by attaching fiducial markers to skin or bone [Maciumas et al.1993], with drawbacks in patient comfort and accuracy [Peters et al.1996]. Surface-based registration offers complicated, yet flexible and accurate, alternatives using techniques such as probing points on the skin [Ryan et al.1996], video alignment [Colchester et al.1996], or laser scanning [Grimson et al.1996b].

For more difficult surgeries, the goal is to present the surgeon with not just one diagnostic scan, but all the useful information available. An array of information can be derived from fusing data sets about morphology, cortical function, and metabolic activity. For example, bone is most practically imaged with Computed Tomography (CT), soft tissues are usually best discriminated with T1-weighted and T2-weighted Magnetic Resonance Imaging (MRI), and vasculature can be derived from MR Angiography (MRA). Functional physiologic data is available from transcranial magnetic stimulation (TMS), or functional MRI (fMRI), and is invaluable for executing surgical resection without sacrificing critical brain functions. Metabolic information is scanned using single photon emission computed tomography (SPECT) or positron emission tomography (PET) which distinguishes metabolically active tumor parts from necrotic areas. These varied data sets are acquired in different coordinate systems and need to be registered to a common framework before that framework is in turn registered to the patient.

1.1.2 Limitations of Conventional Image-Guided Surgery

The accuracy provided by both frame-based and frameless stereotaxy decreases as surgery progresses. The positions of anatomical structures change as the brain swells when the skull is opened, cerebrospinal fluid (CSF) egresses, or tumor tissue is removed. Greater dynamics accompany the resection of larger or deeper tumors.

Intra-operative imaging is being developed to address these concerns and also to detect unexpected complications during surgery such as hemorrhage. Interventional imaging is typically performed with X-ray fluoroscopy or ultrasound. Fluoroscopy is often used to guide catheters in cardiac procedures and needle placement in spinal punctures. Ultrasound is most widely used in guiding biopsies in the abdominal cavity, and preliminary experiments have included its use in guiding localization in neurosurgery [Auer1990, Bucholtz et al.1997]. More recently, technological developments have lead to high-resolution, high-contrast 3D interventional imaging, such as interventional CT [Lunsford, Parrish and Albright1984], and MRI [Gronemeyer et al.1995, Schenk et al.1995]. This thesis focuses on computer software for augmenting the functionality of interventional MRI. Compared to the other named interventional imaging modalities, an open MR system provides the advantages of high soft-tissue contrast, a lack of radiation exposure to patient and surgeon, a clear definition of resection boundaries, and continuous access to the operative field [Jolesz1997].

1.1.3 Augmenting Interventional MRI with the 3D Slicer

In order to amplify the benefits of interventional MRI, we propose augmenting the scanning component with computer software that maximizes the interactivity of an image-guided therapy system through focusing on the following five areas:

Image Resolution — Some anatomical structures are difficult to distinguish on interventional MR images, but are clearer on conventional, diagnostic MRI that benefit from a higher magnetic field and longer imaging times.

Imaging Time — For surgical guidance to be interactive, images must be acquired quickly enough to be utilized without disrupting or slowing down the procedure. Fast imaging techniques are being developed [Speilman, Pauly and Metyer1995], but in general, faster imaging brings lower image quality.

Multi-modal Fusion — Functional and metabolic data that is acquired pre-operatively could deliver increased benefit if integrated with intra-operative, anatomical information.

Faster Localization — Interventional MR provides the capability of planning approach trajectories by maneuvering a tracked wand and collecting images at the rate of 6-20 seconds per image. While this is a significant accomplishment, an ideally interactive system needs an update rate of 10 frames per second.

3D Visualization — Interventional images are presently two-dimensional — requiring the surgeon to mentally map the 2D images seen on a computer screen to the 3D operating field.

The 3D Slicer is a software package that addresses the aforementioned areas. Image resolution, imaging time, and localization are improved by performing real-time re-slicing of both pre-operative and intra-operative data sets, and displaying them for simultaneous review. Multi-modal information is incorporated through automatic registration. The 3D Slicer features a computer graphics display that offers the flexibility to see the situation from viewpoints not physically possible. It has the ability to “fly through” virtual data to facilitate the understanding of complex situations, and aid in avoiding damage to healthy tissue.

1.2 Contributions of this Thesis

Aside from the ongoing benefits reaped from its applications, the 3D Slicer itself contributes to the field of image-guided surgery in the following four ways.

1.2.1 Highly Integrated Analysis and Visualization

The 3D Slicer uniquely integrates several facets of image-guided medicine into a single environment. Interleaving the results from the available assortment of image-guided tools can be cumbersome and time-consuming. The 3D Slicer provides an end-to-end solution that bundles different aspects of analysis into a single visualization framework. It allows medical scans containing anatomical or functional data to be automatically registered [1] and semi-automatically segmented [2]. 3D computer models of key structures such as skin, brain vessels, tumor, and motor cortex can be generated and visualized in a 3D scene along with reformatted images. Reformatting presents cross-sectional views through medical scan data on slices of any position and orientation. Feedback from clinical users of the 3D Slicer suggest that they most appreciate its ability to help them make navigate a complicated assembly of information.

1.2.2 Integration with interventional MR

The 3D Slicer is the first software system to augment an open-MR scanner with this full array of pre-operative data. The analysis previously reserved for pre-operative data [ANALYZE] can now be applied to exploring anatomical changes that occur as surgery progresses. Surgical instruments are tracked and drive the location of the reformatted slices. Real-time scans are visualized as slices in the same 3D view as the pre-operative slices and graphical models.

[1] Aligned.

[2] Segmentation refers to the process of assigning a label to each data point to classify it according to its tissue type.

1.2.3 Medical Reality Modeling Language

We created a novel file format, the Medical Reality Modeling Language (MRML), for conveniently describing 3D scenes composed of a heterogeneous mixture of scanned volumes, 3D surface models, and the coordinate transforms that express the geometric relationships between them. This format specifies where the data is stored, how it relates hierarchically to other data in the 3D scene, and what display attributes are for rendering the data.

1.2.4 Freely Available, Open-Source Software

We have made the 3D Slicer freely available as a tool to clinicians and scientists [SPL]. It has a modular, easily extendable design for other researchers to create new capabilities. We developed the 3D Slicer on top of the OpenGL graphics library [OpenGL] using the Visualization Toolkit (VTK) [Schroeder, Martin and Lorensen1996] for processing, and the Tcl/Tk scripting language [Ousterhout1994] for the user interface. The development of the 3D Slicer required adding many new tools to VTK, some of which have been released to the worldwide community of VTK users.

1.3 Applications

This section briefly introduces three applications for which the 3D Slicer is currently being used at the Surgical Planing Lab (SPL) at Brigham and Women's Hospital: surgical planning, surgical guidance, and volumetric measurements. Later, we elaborate on each of these applications in individual chapters.

1.3.1 Visualization or Pre-Operative Planning

The 3D Slicer was chiefly designed to serve as an intra-operative companion to neurosurgery in an open magnet, but it has also found routine use outside of the operating room and neurosurgery. Doctors and scientists at the SPL have applied the 3D Slicer for visualization and planning in procedures concerning the spine, the pelvis, the

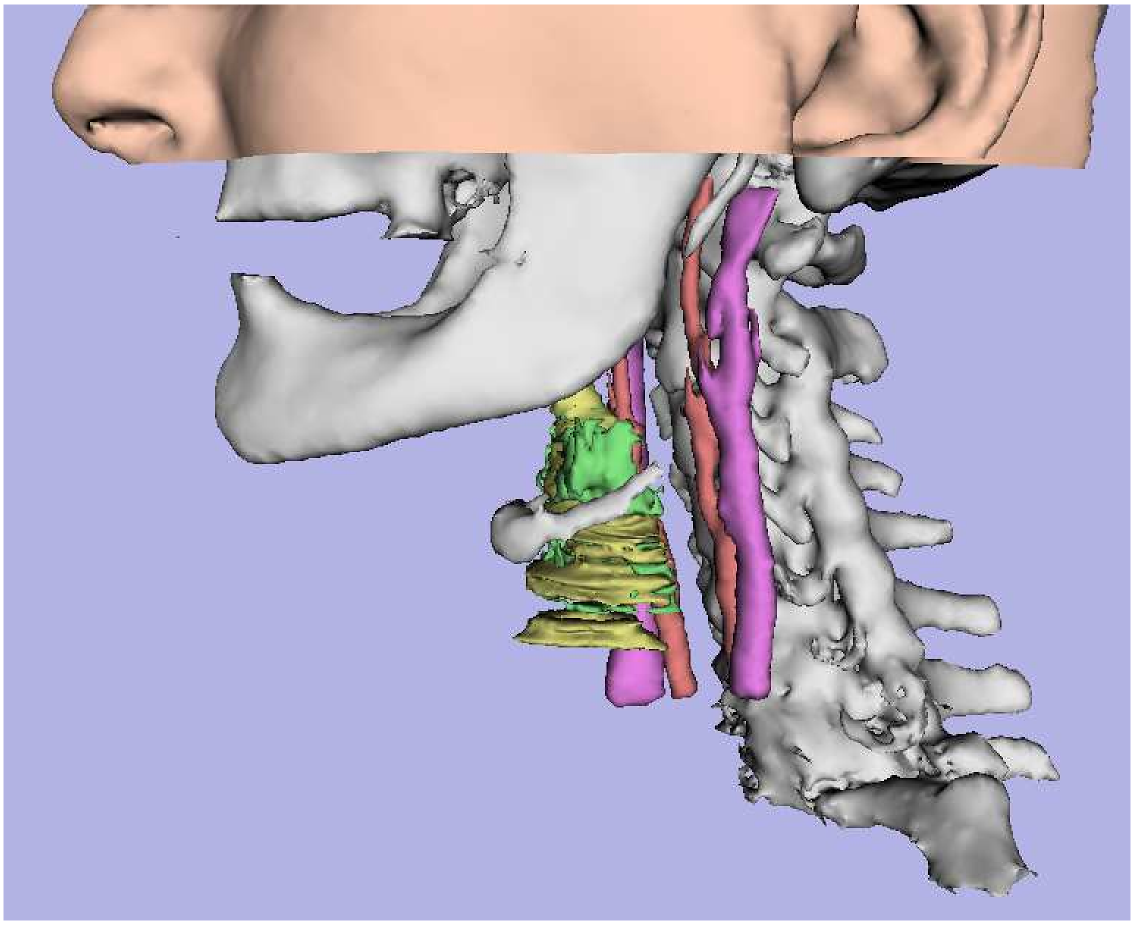

Figure 1-2: This figure shows the use of the 3D Slicer for ENT surgical planning. A computerized reconstruction of the jaw and neck was integrated from from multiple exams such as MR and CT (used to segment the bone) into a single view. This allows the character of the tumor (green) to be clearly explored.

bladder, the heart, and the throat. Figure 1-2 shows the use of the 3D Slicer for ENT surgical planning.

Prior to surgery, pre-operative data sets (T1- weighted MRI, T2-weighted MRI, MRA, fMRI, CT, PET, or SPECT) are fused with an automatic rigid registration. This merged data representation can be visualized by reformatting up to three slices at once, and rendering them in an interactive 3D graphics environment. Each slice can display multiple volumes blended with adjustable opacities. For example, functional information can be overlaid in color on graylevel, anatomical data. Clinicians appreciate the facility of being able to view available data on a single screen because this makes diagnosis more straightforward.

Volumetric data can be segmented semi-automatically using the 3D Slicer's collection of editing tools. Effects such as thresholding, connectivity, morphological operations, and free-hand drawing can be applied to the data on a 3D or slice-by-slice basis while being visualized interactively in 3D. The output label maps are converted to 3D surface models using the Marching Cubes triangulation algorithm [Lorensen and Cline1987]. The surface models can then be visualized in the 3D view along with the reformatted slices, and the slices can selectively clip away portions of some models, such as the skin, to reveal other unclipped models beneath, such as a tumor. Distances, angles, surface areas, and volumes of structures can be measured quantitatively.

Figure 1-3 displays a screen shot of the 3D Slicer being used for trajectory-planning for a craniotomy. A key component of neurosurgical planning is plotting an approach trajectory that avoids critical structures such as the blood vessels and the motor cortex. This is vitally important, because tumors can either infiltrate functional tissue or push it aside, and a tumor that intricately invades eloquent cortex can be considered inoperable for the sake of preserving the quality of life rather than longevity. In the case shown in the figure, a functional MRI exam was registered to an anatomical exam. Surface models of the skin and tumor were constructed from the anatomical exam, and models of the motor cortex, and areas for auditory verb generation and visual verb generation were constructed from the fMRI. The same three slices that are visible in the 3D view are displayed as cross-sections the bottom. The volumetric form of the fMRI results are overlaid in color on the grayscale, anatomical images.

This patient had a tumor in Broca's area on the left side of the brain where 96% of speech is generally processed. A functional MRI exam was administered to ascertain whether the tumor tissue had infiltrated the eloquent cortex. The 3D Slicer's integrated 3D view clearly demonstrated that the speech activity migrated to the right side, proving the operability of this lesion. It is interesting to note that although the 3D Slicer was only scheduled for use in pre-operative planning for this patient, the surgeon came out of the operating room halfway through the surgery to

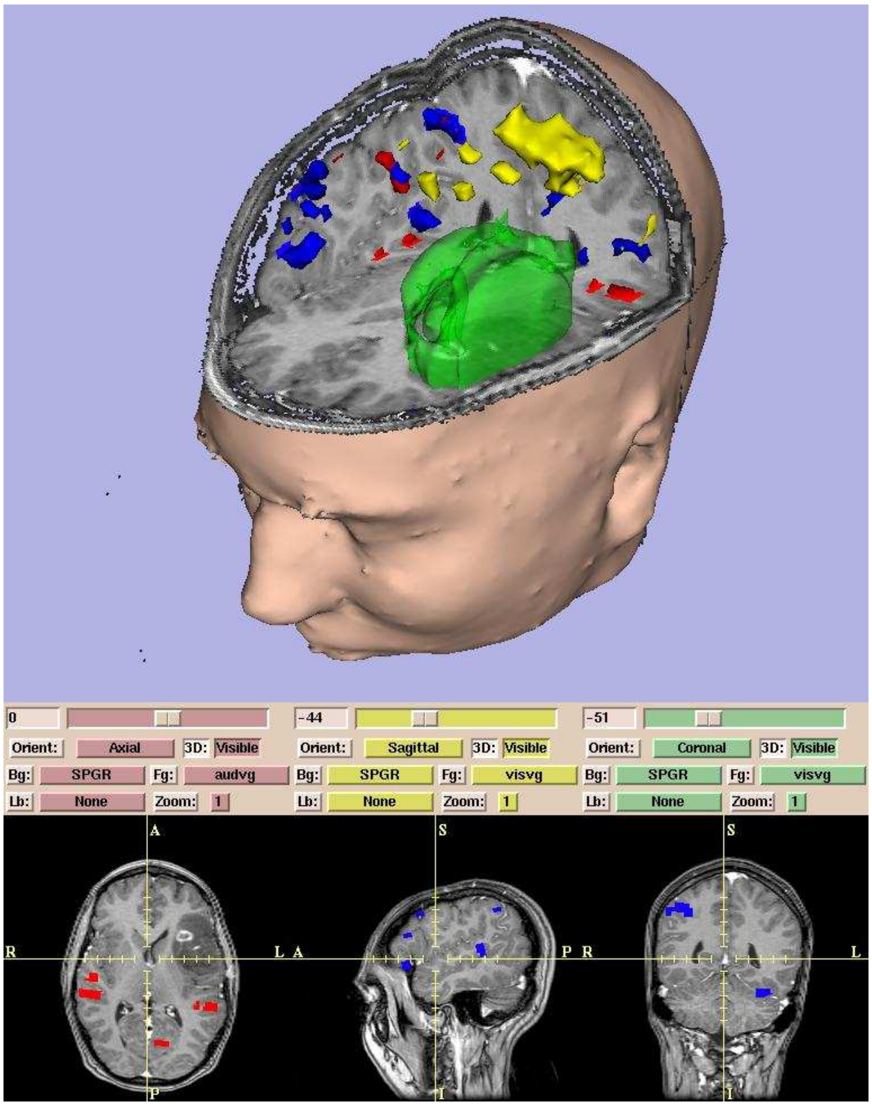

Figure 1-3: Surface models of skin, tumor (green), motor cortex (yellow),auditory verb generation (red), and visual verb generation (blue),are integrated with slices.

view the 3D Slicer again.

1.3.2 Surgical Guidance

A fused, pre-operative data set can be used to augment interventional MR imaging, which is limited by a weaker magnetic field (0.5 vs. 1.5 Tesla) and shorter imaging times (to not prolong the operation). A fast volume scan is collected while the patient is being prepped in the operating room to create a second data representation. Then rigid registration using mutual information [Wells III et al.1996] relates

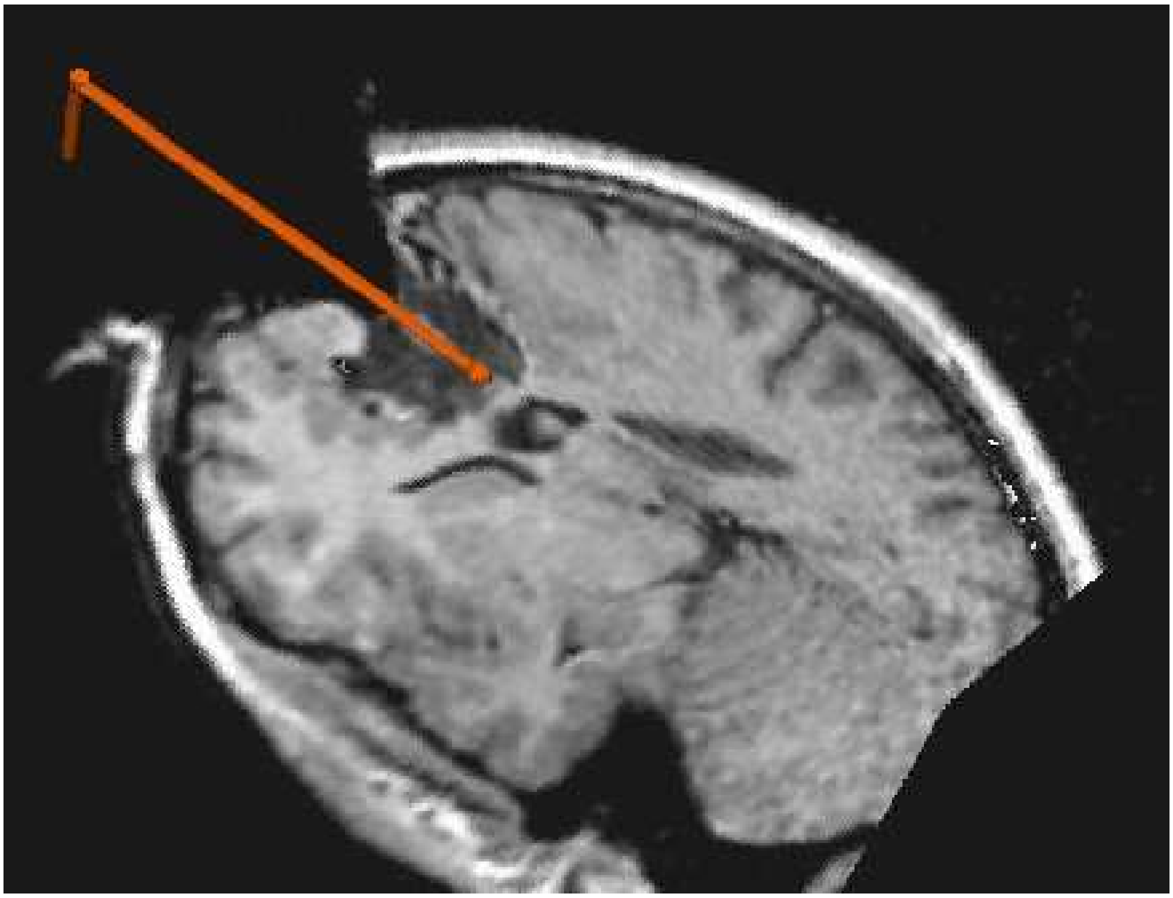

Figure 1-4: The surgical instrument (red) is tracked and 2 orthogonal planes are shown reformatted relative to the instrument tip's position.

the pre-operative data sets to the coordinate frame of the interventional MR scanner. As a tracked, surgical instrument is moved within the surgical field, it is rendered in the 3D view, and the reformatted slice planes follow its position, sweeping through the volumes as a form of virtual real-time imaging. Figure 1-4 displays the two reformatted image planes generated at the location of the tracked instrument, rendered as a red cylinder. The neurosurgeon observes the 3D graphics display on an LCD panel in the scanner gantry, as shown in Figure 1-5.

Since the positions of anatomical structures change during the course of operation, new slabs of data are acquired several times during surgery for reformatting. The surgeon can then observe changes in the position of anatomy on the intra-operative images, while also benefiting from the higher resolution, higher contrast, and functional information of the pre-operative images. During biopsies and laser ablations,

Figure 1-5: The surgeon stands within the gap of the interventional MR magnet and observes the display output of the 3D-Slicer for cues.

2D images are scanned in real time and rendered in the 3D display so spatial relationships can be clearly seen.

After several clinical trials during development, we have used the Slicer as a navigational tool in more than 20 cases since January 1999. This initial feasibility phase proved the Slicer to be a stable and reliable application. As we use the Slicer on a routine basis, a clinical evaluation of its influence on surgical decision making and resection control is being conducted.

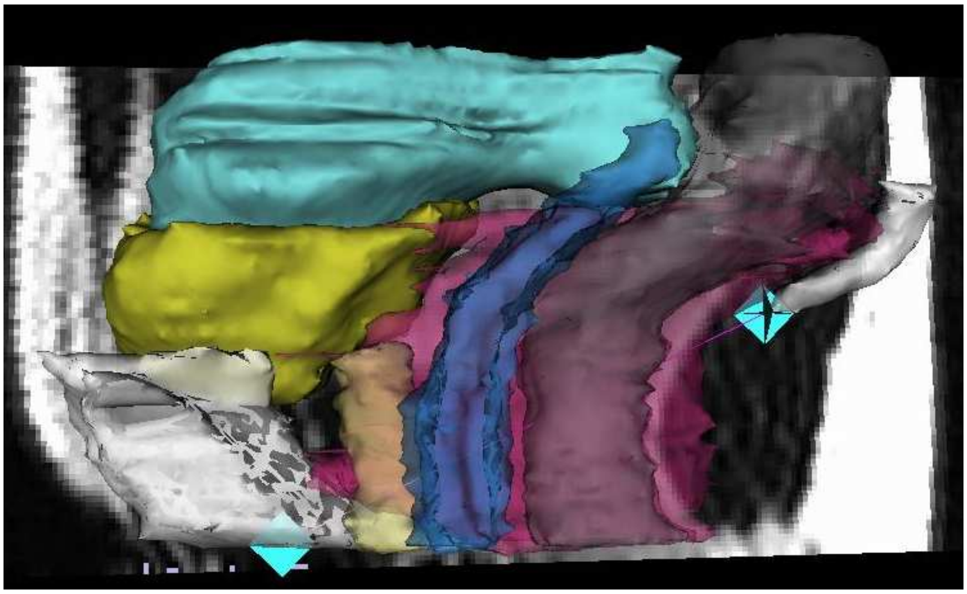

Figure 1-6: 3D models of the female pelvis such as bones (white), bladder/urethra (yellow), vagina (blue), uterus (green), rectum (gray), and the levator ani muscle (pink) can be visualized along with gray-scale images and quantified. The purple line between the two blue markers is measuring the distance of the pubococcygeal line (level of the pelvic floor, and minimum width of the birth canal).

1.3.3 Volumetric Analysis and Studies of Dynamics

The 3D Slicer is used outside of surgery for quantitative studies that require making measurements on source images and 3D models simultaneously. Applications to date include modeling the female pelvic floor, taking quantitative measurements of muscle mass, and conducting orthopedic range of motion studies. Figure 1-6 is derived from a study that relied on the 3D Slicer to quantify the appearance of the normal female pelvic floor to better understand the causes of incontinence.

The open MR presents unprecedented opportunities for scientifically studying the dynamics that accompany intervention. The advantage of the open MR is that the patient is not repositioned between operating and imaging. We are using the 3D Slicer to analyze brain shift in the attempt to reach a level of understanding of the biomechanics of the brain that includes predictive power. Then certain image-guided surgical techniques used in the open MR could be applied to conventional operating rooms. Our findings thus far have revealed quite inhomogeneous distortions of the ventricles and unexpected shifts in the brain's midline of up to 1.0 cm. These

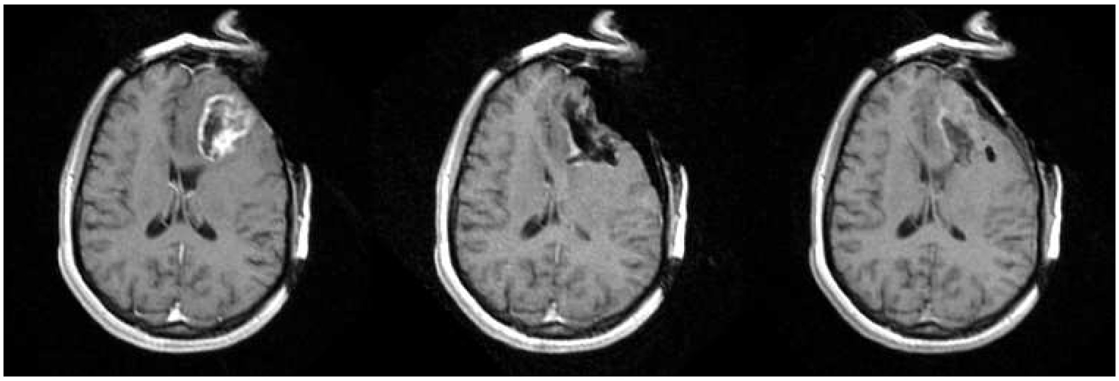

Figure 1-7: From left to right, after opening of the brain's protective dura mater, after tumor resection, and after dural closure.

structures were previously believed to be stationary and suitable for use as landmarks during intra-operative imaging. We have also documented large dynamics during dural closure (shown in Figure 1-7) that cannot be detected by most other approaches.

1.4 Roadmap

We present the system architecture in the next chapter, and we explore each application in the three chapters that follow.

Chapter 2: System Architecture — descriptions of related work, the evolution of the architecture, processing of reformatted images, creation of surface models, multi-modal image registration, integration with MR-guided surgery, the Medical Reality Modeling Language, coordinate systems and transforms, and the user interface.

Chapter 3: Surgical Planning — illustrations of the 3D Slicer's value for aiding diagnosis from multiple scans, trajectory planning, orthopedic surgical planning, visualizing parametrically colored surface models, and constructing and applying anatomy atlases.

Chapter 4: Surgical Guidance — descriptions of operating room protocol beginning with clinical setup and proceeding with descriptions of image registration, trajectory planning, tissue identification, incorporating intra-operative imaging updates, and performing post-resection validation.

Chapter 5: 3D Analysis — descriptions of analyzing intra-operative brain deformations and performing quantitative studies that require making measurements on source images and 3D models simultaneously.

Chapter 6: Conclusion

Chapter 2

System Architecture

This chapter describes the architecture of the system organized into the following sections:

Related Work — comparisons of the 3D Slicer to other systems in the field.

Evolution of the Architecture — a look at the phases of development of the system's architecture and the lessons learned along the way.

Processing Pipeline — a description of the image processing pipeline that produces reformatted slices.

Surface Models — generating and visualizing 3D polygonal models of important structures.

Multi-Model Registration — rigid registration between data sets generated from various imaging modalities (i.e.: MR, CT, SPECT).

Interventional MR-Guided Surgery — the integration of the 3D Slicer with the interventional MR system to form a surgical assistant.

Medical Reality Modeling Language — a novel file format for expressing scenes composed of volumes, surface models, and the coordinate transforms between them.

Coordinate Systems and Transforms — the mathematics behind slicing though multiple data sets simultaneously.

User Interface — a brief description of the user interface.

2.1 Related Work

Image-guided surgery systems strive to achieve one or more of the following:

- Data analysis
- Surgical planning
- Surgical guidance
- Surgical guidance with intra-operative updates

Interleaving the results of several of the currently available systems can be cumbersome and time-consuming — making it impractical for clinical applications. Our system is uniquely valuable in its integration of all of these facets into a single environment. Such cohesion is required for seamless integration of pre-operative and intra-operative data. The same level of analysis previously reserved for pre-operative data can now be applied to exploring the anatomical changes as the intervention progresses.

2.1.1 Systems for Analysis and Planning

Several existing software packages present a more extensive collection of tools for data analysis and planning, but do not support surgical guidance. [ANALYZE] has been developed at the Mayo Clinic since the early 1970's, and it offers a vast set of tools for 2D and 3D visualization, image processing, registration, segmentation, and quantitative analysis. [MEDx] and [MNI] have followed with a similar scope as ANALYZE, and are not designed provide any surgical guidance.

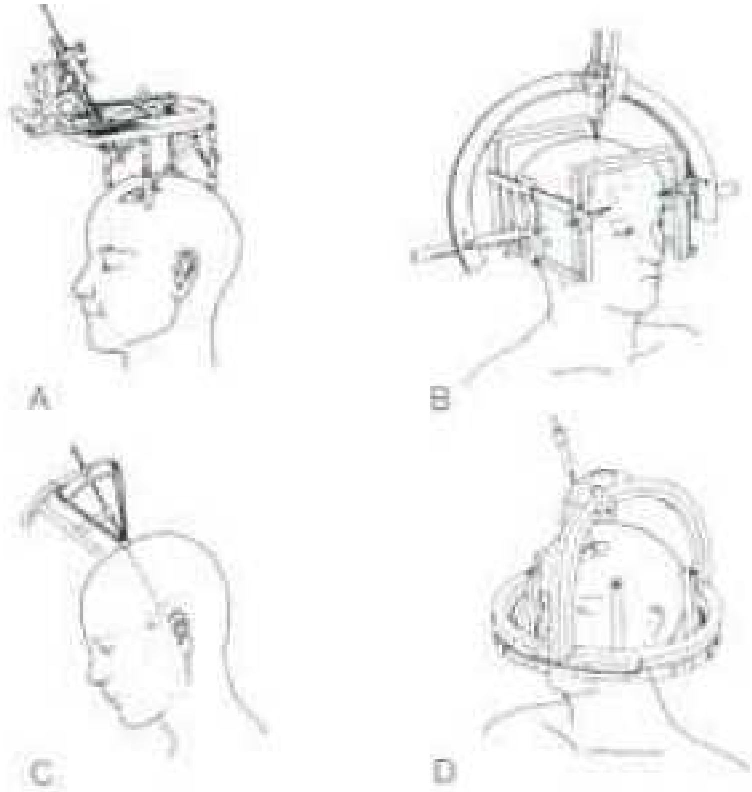

Figure 2-1: Reference frames were attached to the patient's skull to establish correspondence between imaging and surgery, which were sometimes several days apart, yet the frame had to remain intact in the interim.

2.1.2 Systems for Surgical Guidance

Several systems have been developed to facilitate surgical guidance, but they feature leaner support for analysis and planning, and they are not designed to incorporate intra-operative updates. These systems provide a means of establishing a correspondence between fused, pre-operative data and the patient as positioned on the operating table. Image guided surgery has progressed from attaching cumbersome reference frames to the patient [Galloway1993] (Figure 2-1), to frameless stereotaxy that integrates registration and tracking [Grimson et al.1998].

Registration of pre-operative imagery to patient position is commonly achieved by attaching fiducial markers to skin or bone [Maciumas et al.1993], with drawbacks in patient comfort and accuracy [Peters et al.1996]. Surface-based registration offers complicated, yet flexible and accurate, alternatives using techniques such as probing

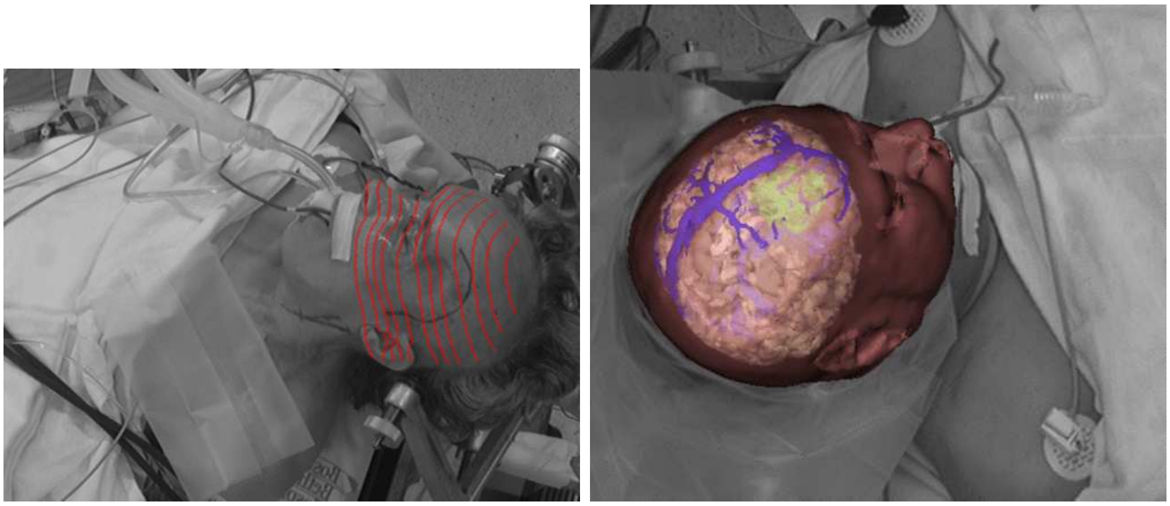

Figure 2-2: LEFT: Topography identified by a laser scanner. RIGHT: Video overlay of pre-operative data.

points on the skin [Ryan et al.1996], video alignment [Colchester et al.1996], or laser scanning [Grimson et al.1996b] (Figure 2-2).

Tracking of surgical instruments has been achieved through a variety of methods. The Radionics System (Burlington, MA) reads angles from encoders on the joints of an articulated arm to compute the position of the tip. Flashpoint (Image-Guided Technology, Boulder, CO) and Optotrak (Northern Digital Inc., Waterloo, Ontario) use multiple cameras to view flashing LEDs on the instrument and triangulate its position. The VISLAN system [Colchester et al.1996] passively extracts a pattern on a wireless instrument from two stereo images. Fastrack (Polhemus, Inc., Colchester, VT) and Flock of Birds (Ascension Technology Corporation, Burlington, VT) utilize electro-magnetic digitizers that alleviate the need for a clear line of site between the instrument and the tracking equipment.

Intra-operative imaging is being performed with X-ray fluoroscopy, CT [Lunsford, Parrish and Albrig ultrasound [Bucholtz et al.1997], and MRI [Gronemeyer et al.1995, Schenk et al.1995]. Our system is the first to augment an open-MR scanner with various pre-operative data sets with information on morphology (MR, CT, MR angiography), cortical function (fMRI), and metabolic activity (PET, SPECT) [Alexander et al.1993].

2.2 Evolution of the Architecture

A system for registering pre-operative and intra-operative images for MR-guided therapy has been undergoing development for the last two and a half years. The basic concepts were proven in a clinical trial where each slice could just reformat a single volume [Hata et al.1997] and surface models could be seen, but not constructed. The system was then redesigned from scratch.

We developed the 3D Slicer on top of the OpenGL graphics library using the Visualization Toolkit (VTK) [Schroeder, Martin and Lorensen1996] for processing, and the Tcl/Tk scripting language [Ousterhout1994] for the user interface. VTK provides a set of objects written in C++ that can be chained together to form a data-processing pipeline. Pipelined processing maximizes interactivity because the output of each stage is stored in memory, and any update from user-interface controls triggers a change at a minimal depth into the pipeline.

Although the system was built on VTK and Tcl from the beginning, it has grown to more fully leverage VTK as well as contribute to VTK's development. We added several classes to VTK by deriving them from existing, documented classes, which results in well-defined inputs and outputs. This is vitally important because other researchers will continue to add to the 3D Slicer's development as the field grows to meet needs in a variety of specific application areas where researchers at various sites have more expertise. Therefore, extendibility and modularity are key attributes of the design. Consequently, each module was first prototyped to achieve correct functionality, and then rewritten to meet these criteria. During the earlier work of this thesis, the 3D Slicer's image processing pipelines had an ad hoc implementation utilizing VTK's 3D graphics pipeline. As VTK developed a separate imaging pipeline more appropriate for the operations we perform on the reformatted slices, the 3D Slicer was completely rewritten to take advantage of, and contribute to, this new architecture.

Appropriately coded VTK classes handle multiple data types, handle data with multiple components per pixel, support multi-threading, and run on multiple platforms. We operate the 3D Slicer on PCs running Windows and Sun workstations running Solaris.

2.3 Processing Pipeline

This section describes the image processing pipeline that generates each reformatted slice. The pipeline begins with extracting a 2D slice through a 3D volume, and proceeds to color the data and merge it with other layers to form a composite slice.

2.3.1 Reformatting

Medical imaging scanners typically output data in the form of 3D arrays of volume elements appropriately called *voxels*. Each voxel is a scalar value, often an integer, representative of a physical property that the scanner is sensitive to. For example, CT images are a map of X-ray attenuation coefficients produced by sending X-ray beams through the body at all angles and measuring the output radiation on the flip side.

MR imaging exploits the principle that any nucleus with an odd number of protons has a dipole moment and therefore behaves as a miniature magnet. When the body is placed in a magnetic field of 1.5 Tesla (stronger than a automobile salvage-yard magnet), a minute fraction of the hydrogen atoms are reoriented to align with the field. Imaging is performed by momentarily generating a different magnetic field and measuring the radio signal emitted by the dipoles as they relax back to their equilibrium positions. More details can be found in [Westbrook and Kaut1993].

Volume data is stored as a stack of 2D images as displayed in Figure 2-3. The 3D Slicer enables one to better visualize volume data through Multi-Plane Reformatting (MPR). A reformatted image is derived by arbitrarily orienting a plane in 3D space, and assigning values to each 2D pixel of the plane by interpolating the 3D voxels of the volume data intersected by the plane. We reformat up to three slices at once with independent orientations, in real time. Slices can be arbitrarily oblique, or orthogonal. Orthogonal slices can be oriented relative to either the coordinate frame

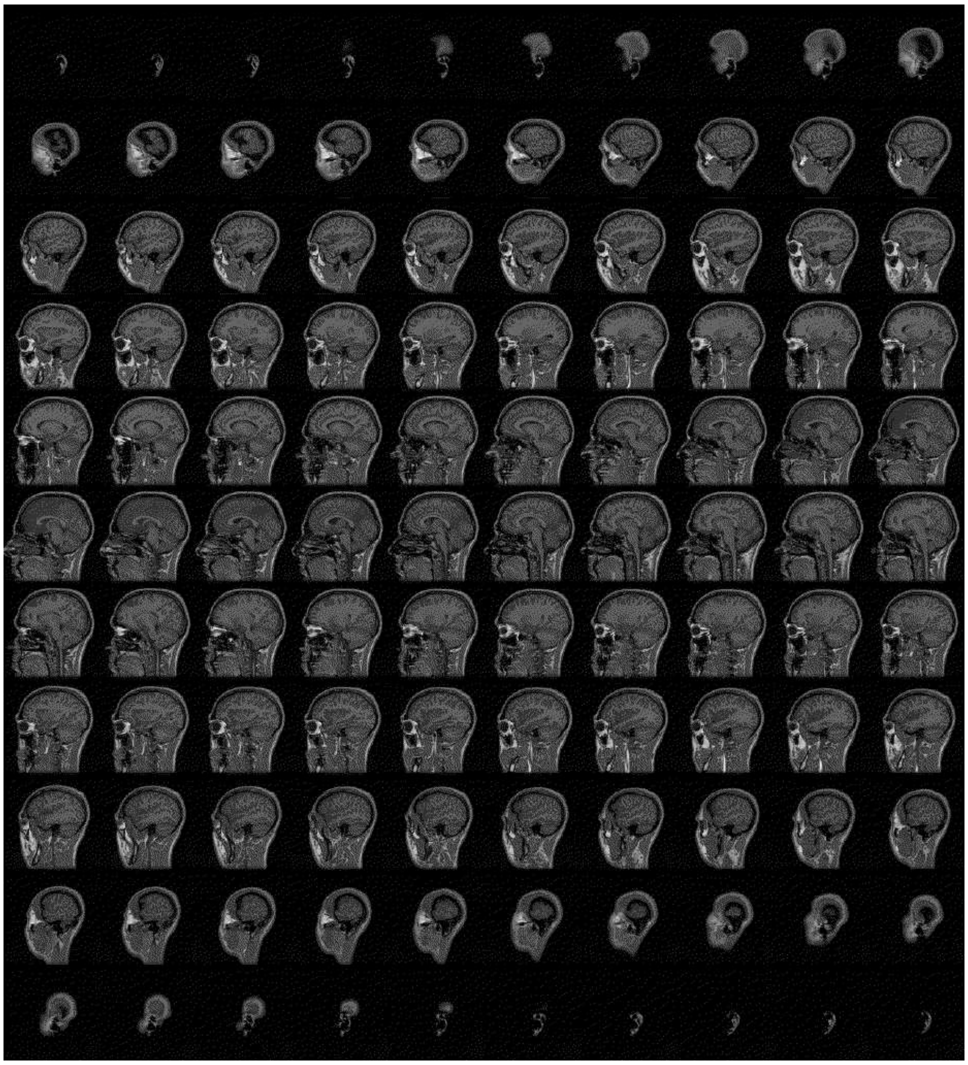

Figure 2-3: A 3D volume is stored as a set of 2D images.

of the scanner or the tracked pointing device. Figure 2-4 demonstrates these two, different cases. Sliders on the user interface allow the slice locations to be slid along the axis perpendicular to the image plane.

Some radiological applications insist on minimum interpolation, and so we provided additional orientation options that generate orthogonal slices relative to the data itself rather than the coordinate frame it has been registered to. In this case, the sliders on the user interface set the slice position in terms of voxels rather than millimeters.

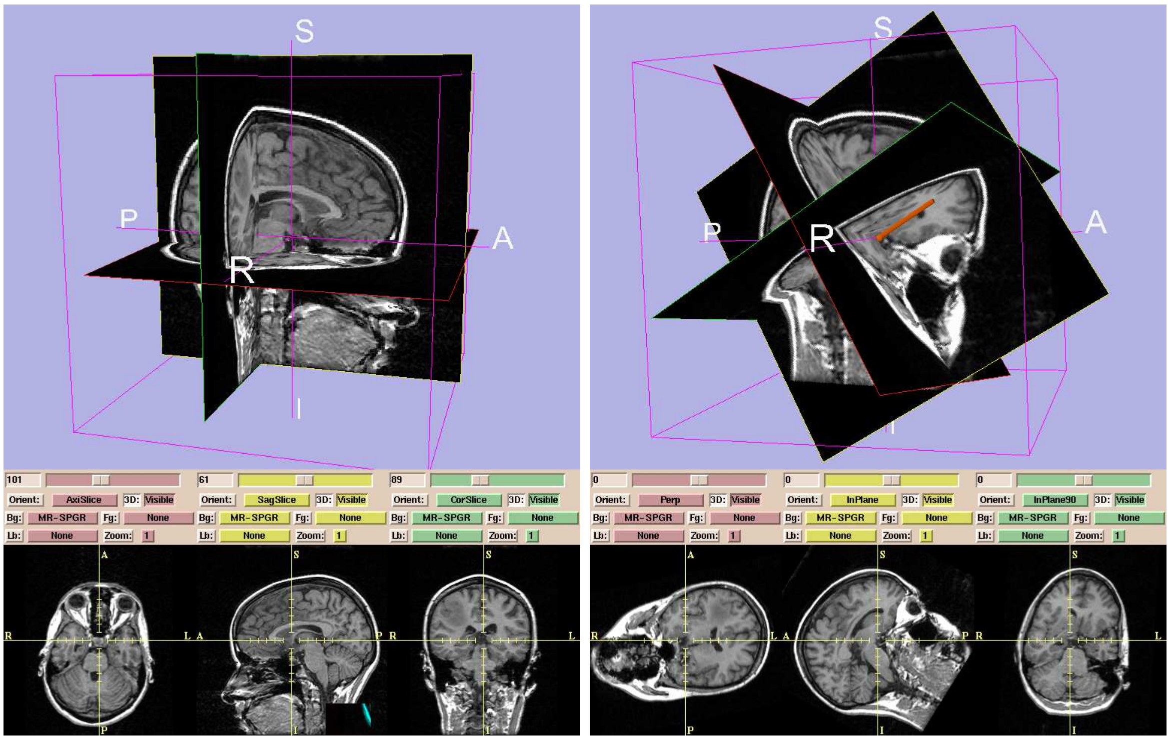

Figure 2-4: Left: slices oriented relative to the reference coordinate frame. Right: slices oriented relative to the pointing device (orange).

2.3.2 Multiple Volumes

Each reformatted slice may be generated from any of the available data sets. For example, one slice could be reformatted from pre-operative SPGR, a second slice could be reformatted from pre-operative T2- weighted MRI, and a third slice could display the images that are being scanned in real time. Figure 3-1 is an example.

2.3.3 Multiple Volumes on the Same Slice

We extended reformatting to slice through both anatomical and functional volumes of data simultaneously. Each slice is the composite of a background layer, foreground layer, and label layer. The background image is typically gray-scaled anatomical data, and the optional foreground can be colored, functional information. The foreground layer is overlaid on the background layer with an adjustable opacity to form the image that is displayed on the slice. As for the third layer, the output of image segmentation is a labelmap where each voxel is assigned an integer label according to the segmented structure it belongs to. The boundaries of each structure can optionally be drawn in color as the third layer. For example, the segmented tumor is outlined in green in Figure 2-5.

There are different ways of blending images with respective opacity values. We lay down the first image, and then blend in the second as follows:

```
new = first * (1-opacity) + second * (opacity)
```

But if the second data set is mostly zeros (as in fMRI), then we are mostly just dimming the first image. Consequently, we have a second mode that leaves the first image's pixel alone whenever the second image's pixel has an opacity value of 0.

```
if (second has per-pixel opacity=0)
    new = first
else
    new = first * (1-opacity) + second * (opacity)
```

Figure 2-5 describes these techniques pictorially. A vascular scan (phase-contrast MRA) is colored with a *Hot Iron* palette (range from black to red to yellow to white) and a threshold is applied to remove the noise floor. The uniform blending technique (bottom left image of 2-5) works well for registering images or observing differences in images gathered at different time points. The method of selective overlay (bottom right image of 2-5) is best for displaying the location of small, vitally important features, such as vessels or functional areas.

2.3.4 Mapping Scalars to Colors

Medical scanners (MR, CT) generate single-component, scalar values with a dynamic range that far exceeds the number of shades of gray distinguishable with the human eye. For display on a computer screen, values are transformed into data with four components (red, green, blue, and opacity) through spreading a palette of colors (such as 256 shades of gray) over a *window* of values centered around a *level* value.

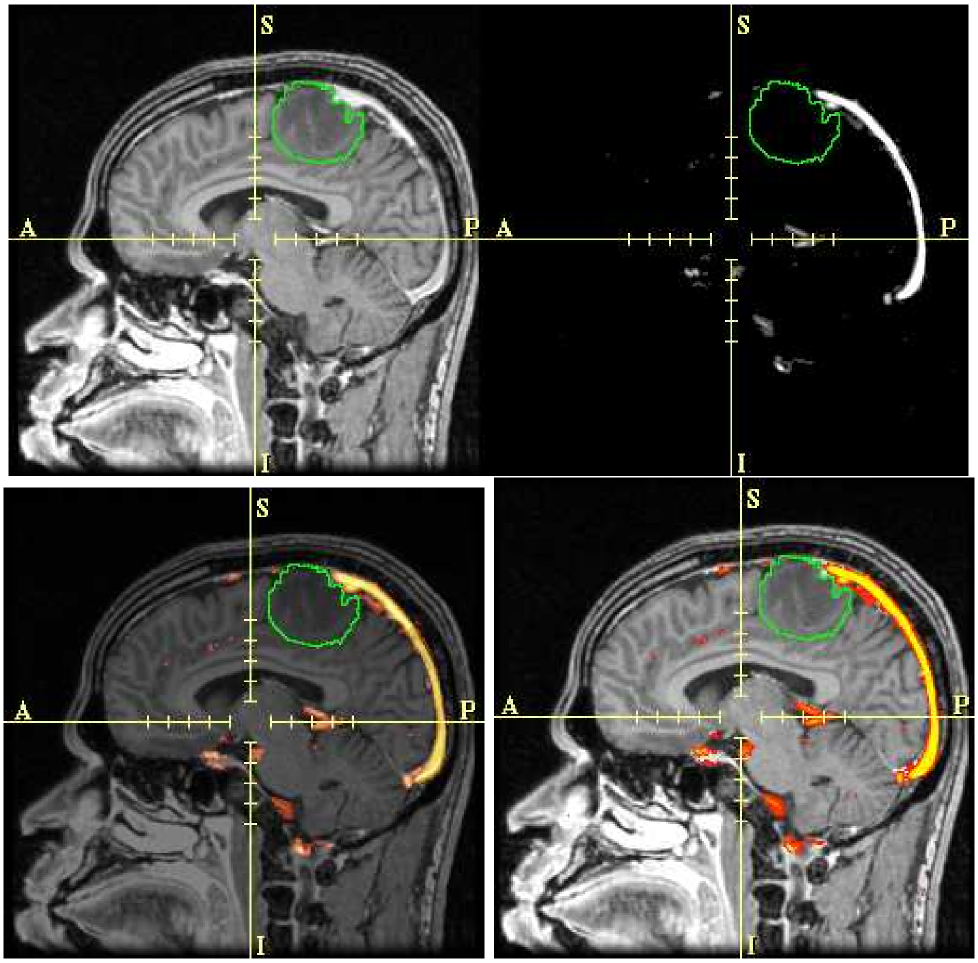

Figure 2-5: TOP: the outline of the tumor segmentation is drawn in green on an anatomical image on the left, and on a vascular image on the right. BOTTOM: The vascular image is fused with the anatomical image using uniform blending on the left, but selective overlay on the right.

We provide several color palettes in addition to a gray palette to facilitate overlaying one data set in color on another in gray.

We take advantage of the opacity component to implement thresholding. We threshold functional data to hide pixels with insignificant probability of being functional tissue. We threshold anatomical data because the 2D images are texture-mapped onto planes in 3D space. The images have black space outside the body and it is undesirable to have plain black planes occluding the view of the rest of the 3D scene. Each color palette has a color at index 0 that is transparent black. All other colors have full opacity. This way, pixels that fail to meet the thresholds can be as-

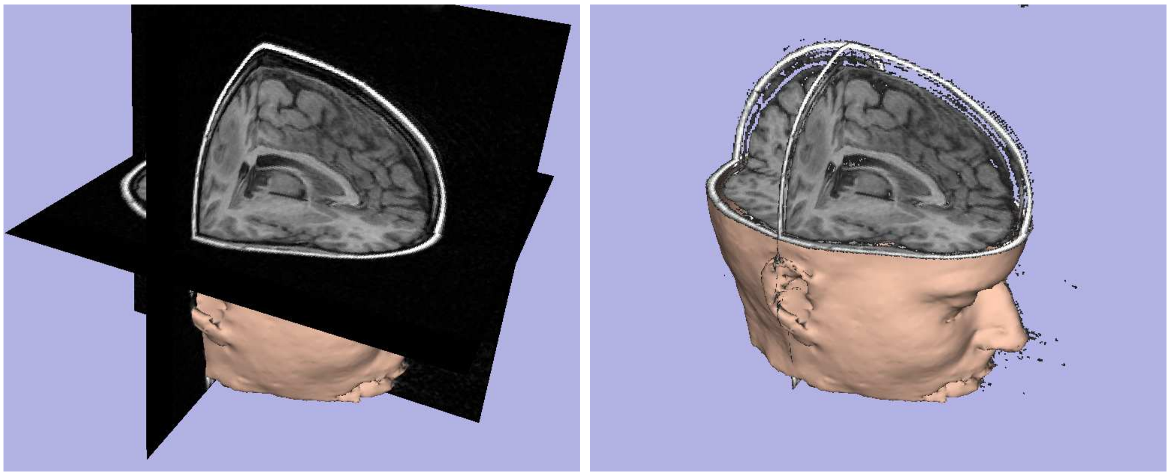

Figure 2-6: With (left) and without (right) applying a threshold to remove the noise floor of the slices.

signed color #0 and they will be texture-mapped transparently in 3D, but displayed black on a 2D window. Figure 2-6 demonstrates that the view is less occluded when uninteresting values are transparent.

The 3D Slicer can automatically apply this threshold after computing the image histogram, which is a plot of the frequency of each voxel value. The threshold is set to dissect the histogram at the intersection of the probability densities for noise and signal. This occurs at the bottom of the first valley of the histogram. The 3D Slicer also implements automatic window and level by computing the centroid (center of mass) of the histogram, and setting the level to the value of the centroid, and the window to twice this quantity (a fast approximation of the width of the signal's distribution).

2.3.5 Reformatted Slice Location

Besides using the sliders on the GUI to position reformatted slices, users may simply click on a location to change the center of all slices (the focal point) to that point in 3D space. Figure 2-7 is a screen shot taken after clicking on the center of the tumor. In this case, all slices were set to have sagittal orientations and different zoom factors. Note the rulers with centimeter markings.

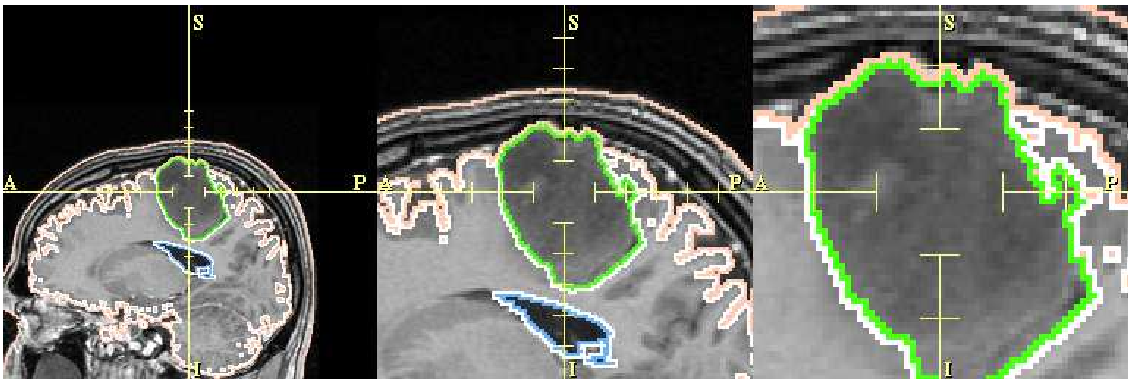

Figure 2-7: Viewing segmented structures under various magnifications.

2.4 Surface Models

Reformatted images provide cross-sectional views through soft tissues that are similar to what physician is accustomed to viewing. A surgeon, however, also needs to instantly visualize precisely where a tumor is or is not, or exactly where a winding blood vessel lies in relationship to the scalpel. These spatial relationships become more readily apparent when the structures are extracted, or segmented, from the volume data [Grimson et al.1996a].

The output of this segmentation process is a set of label maps, where each voxel is assigned a value according to its membership in an anatomical structure such as skin, brain, vessel, or tumor. A label map is an intermediate step in creating a surface model, which is a set of joined triangles arranged in 3D space to form a closed surface. A surface model is made from a label map by creating a mesh of triangles that bound all voxels of a common label. The resulting triangle mesh is the 3D graphics equivalent of wrapping a fishnet around a tumor.

2.4.1 Volume Editor

Volumetric data can be semi-automatically segmented using the 3D Slicer's suite of editing tools. Effects such as thresholding, morphological operations (erosion, dilation) [Serra1982, Brummer et al.1993], island-removal (erasing small groupings of similar pixels), measuring the size of islands, cropping, and free-hand drawing of poly-

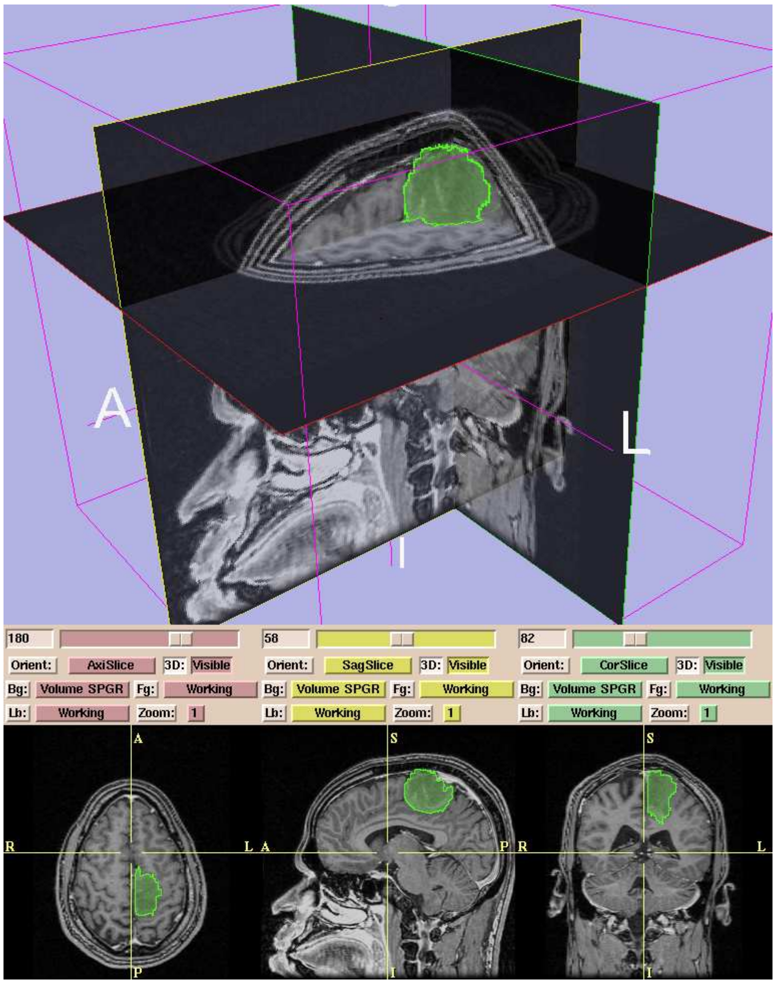

Figure 2-8: While a tumor is being interactively edited, it is displayed on both the 2D slices and in the 3D view. The segmented tumor is shown as translucent green overlaid on the original gray-scale data. Opaque green clearly delineates the tumor boundary.

gons, lines, or points can be applied to the data. Each effect can be administered on either a 3D or slice-by-slice basis. We found it very helpful to users to allow the slice-by-slice editing to be administered on either axial, coronal, or sagittal slices merely by clicking on the appropriate slice. Another strength of our system is that effects can be visualized by overlaying the working volume translucently on the original volume and explored in the 3D view, as shown in Figure 2-8.

Each effect may be applied to either the *original* volume, or a *working* volume. Multiple editing effects may be applied to the working volume, and when finished, the working volume may be merged with a *composite* volume to overlay smaller structures

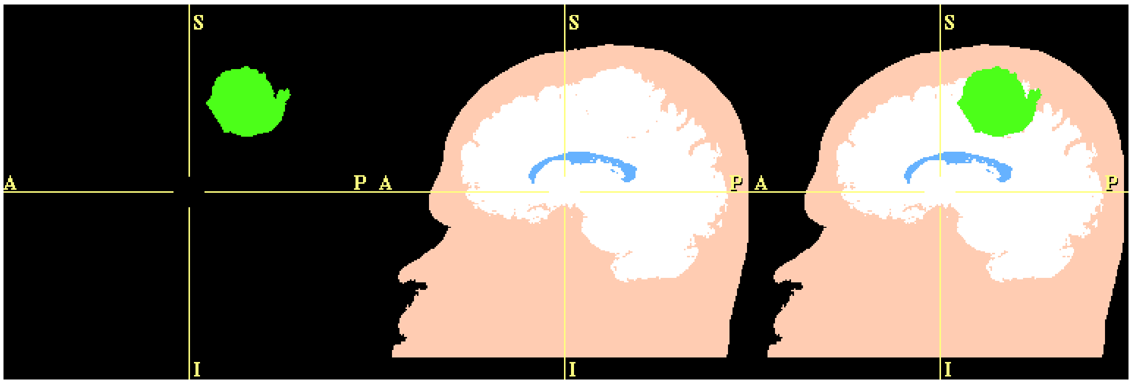

Figure 2-9: The original volume is edited to produce a tumor in the *Working* volume (left). This working volume is then written over the *Composite* volume (center) to form a new composite volume (right).

onto larger ones. For example, we segment skin automatically from an MR scan by applying a threshold at the first trough of the histogram (where the noise and signal lobes overlap) and storing the binary output in the working volume. We then remove islands in the working volume, and finish with an erosion and dilation to smooth the edges. This working volume is then copied to the composite volume. Next, segmentation of the brain can be performed in the working volume, and then all non-zero voxels of the working volume can overwrite the composite volume to form a combination of skin and brain. Figure 2-9 illustrates this process.

Through a modular architecture that exploits C++ inheritance, other researchers can easily add new segmentation tools that increase the level of automation, such as [Kapur1999].

2.4.2 Model Generation

The bounding surfaces of the label maps are extracted and represented as a collection of triangles using the Marching Cubes algorithm [Lorensen and Cline1987]. Decimation reduces the number of triangles to a quantity that can be more quickly rendered with little observable loss in detail [Schroeder, Zarge and Lorensen1992]. For example, a typical brain surface is reduced from approximately 500,000 triangles to 150,000. Figure 2-10 illustrates how models originate from label maps.

We found it helpful to have the 3D Slicer call a separate process to generate

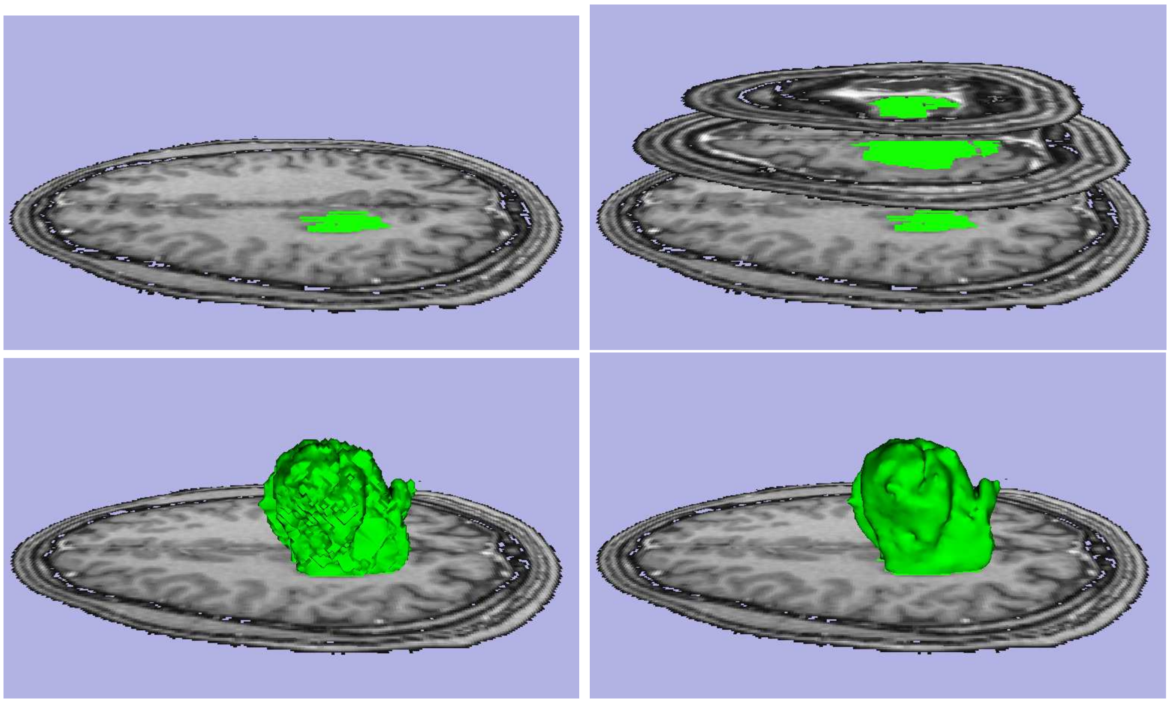

Figure 2-10: The label map of a tumor is created by segmenting slices (top). Next, a polygonal surface is created to encompass the segmentation (bottom-left). This surface is smoothed (bottom-right) to remove the digitization artifacts.

the models so that sets of many models could be performed as a batch job in the background while the user continues to use the 3D Slicer for other tasks. Model generation can be run on another machine with more memory and the user is emailed when the job is complete so the models can be viewed in the 3D Slicer. This feature in conjunction with an interface designed to minimize mouse clicks and movement has approximately halved the 3-hour time required to manually segment a neurosurgical case.

2.4.3 Model Visualization

Surface models of key anatomical structures can be visualized in the 3D view along with the reformatted slices. Our surgical colleagues favor viewing skin as a landmark. Therefore, we allow for the slice planes to selectively clip away the skin model to reveal other unclipped models beneath, such as a tumor or physiologically critical structures like blood vessels, as well as the respective image planes. Figure 2-11 demonstrates

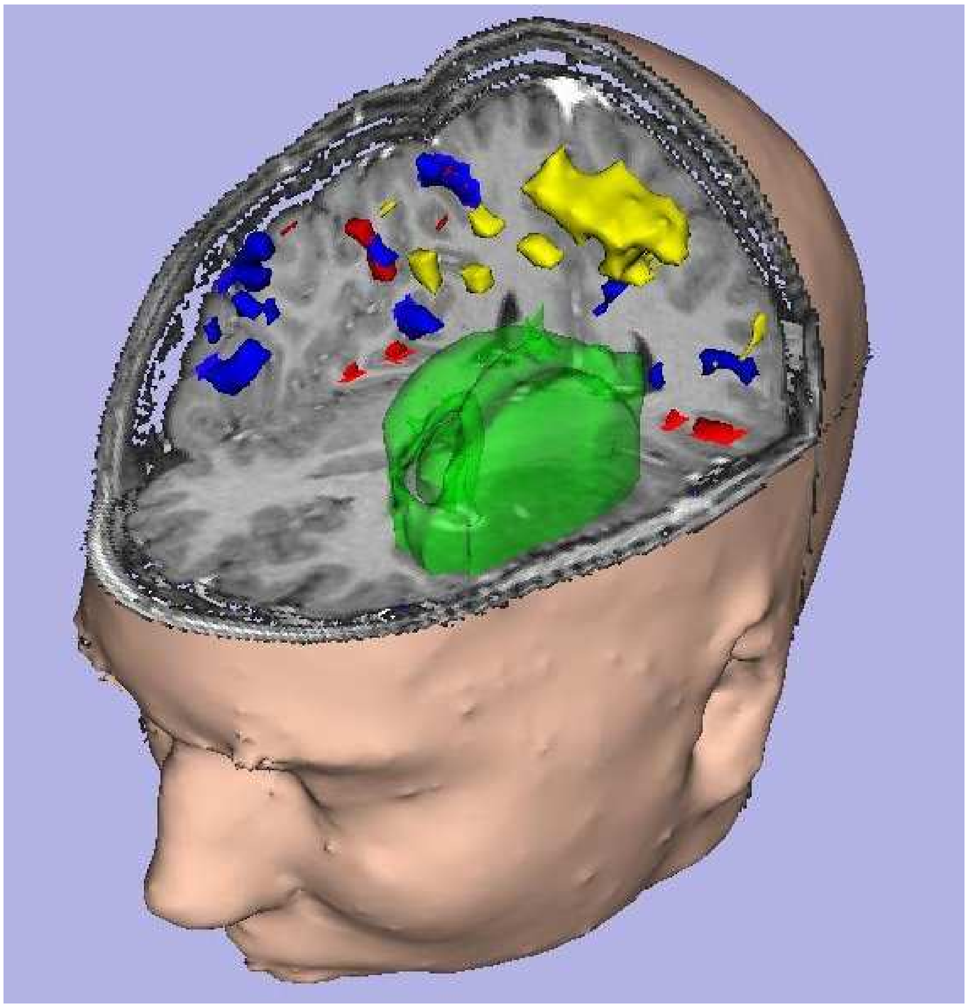

Figure 2-11: Reformatted planes clip the surface model of the skin to reveal models beneath.

this technique.

In cases where the tumor is near the eloquent cortex, such as centers of speech or motor processing, functional models are included to illustrate their proximity to tumor tissue. These models display the 3D structure of a detected correlation of functional response. Each model is colored differently (but consistently between cases), and rendered with adjustable opacity (such as the translucent tumor in Figure 2-11). We also support applications that require varying a model's color along its surface to visualize a property, such as bladder wall thickness, for example.

2.5 Multi-Model Registration

The various volumes being segmented and reformatted in the 3D Slicer are acquired in different coordinate systems and need to be registered to a common framework in order to correctly view all components in a single scene.

The 3D Slicer supports rigid, manual registration by allowing the user to specify which volume to move, and then translate and rotate that data set by clicking and dragging on the images in the 2D windows. The reference volume can be displayed as the background layer of the slices, and the registered volume can be displayed translucently as the foreground layer.

Furthermore, with one button click, the result of manual registration can be used as an initial pose for automatic registration by maximization of mutual information [Viola and Wells III1995, Wells III et al.1996]. This method is of general utility, and other implementations of it have performed well in an NIH-sponsored test [J. West1997]. When registered by mutual information, knowing something from one data set allows predicting something about the corresponding location in the other data set. This method is more robust than conventional correlation techniques that use a mean squared error. For example, when registering T1-weighted MRI to T2-weighted MRI, tissue that appears hypointense in one image can be hyperintense in the other, and therefore, correlation is an incorrect metric.

The registration capability is implemented as a separate process that communicates with the 3D Slicer through the use of MRML files as described later. Figure 2-12 illustrates the automatic registration of two, 34-slice data sets in about one minute.

2.6 Integration with MR-Guided Surgery

2.6.1 The Open MR System

We integrated the 3D Slicer with the interventional MR system (Signa SP, GE Medical Systems, Milwaukee, WI) to form a computerized surgical assistant as envisioned by Jolesz as the *operating room of the future* [Jolesz1997, Black et al.1997]. A system

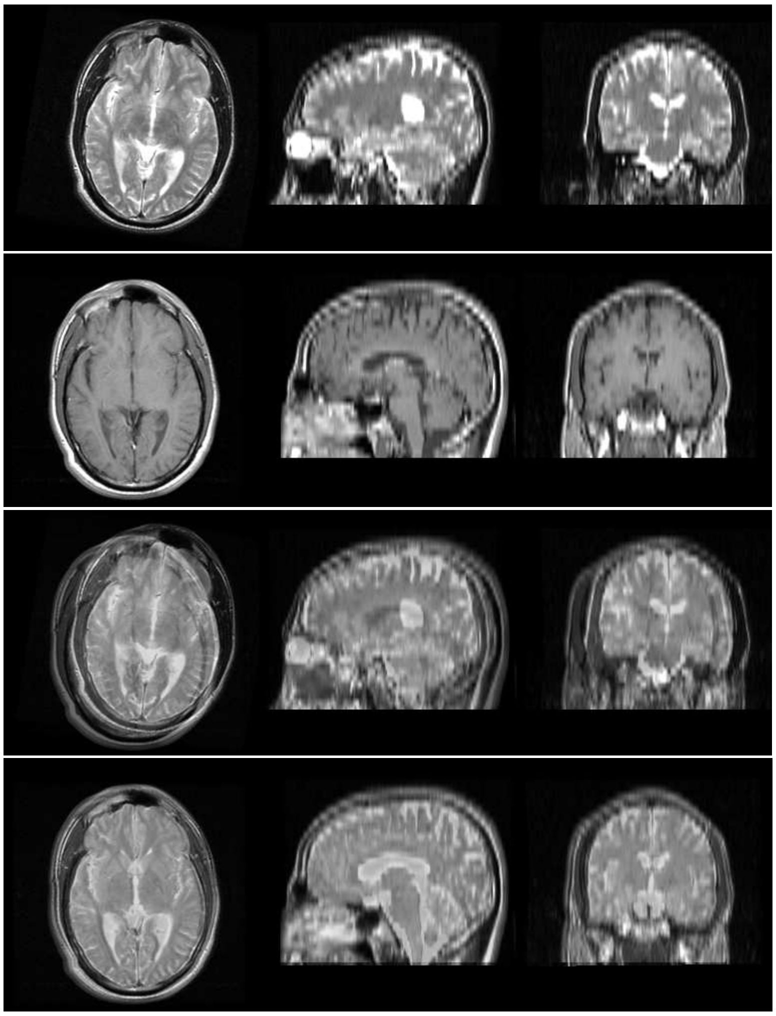

Figure 2-12: T1-weighted MR (top row) and T2- weighted MR (second row) are shown overlaid prior to registration (third row) and after (bottom).

diagram is shown in Figure 2-13, and a view of the operating room is given in Figure 2-14.

The location of the imaging plane (the position of the plane's center, and the orientation of the plane in space) can be specified with an optical tracking system (Image Guided Technologies, Boulder, CO), which we will refer to as the *locator* or *pointer*. Light-emitting diodes affixed to the locator are tracked by 3 CCD cameras that are attached to a rail above the operating field. The spatial relationship of the instrument relative to the scanner is reported as both a position and an orientation with an update rate of approximately 10 Hz. The locator is shown in Figure 2-15 with an LED on each prong of the star-shaped handle.

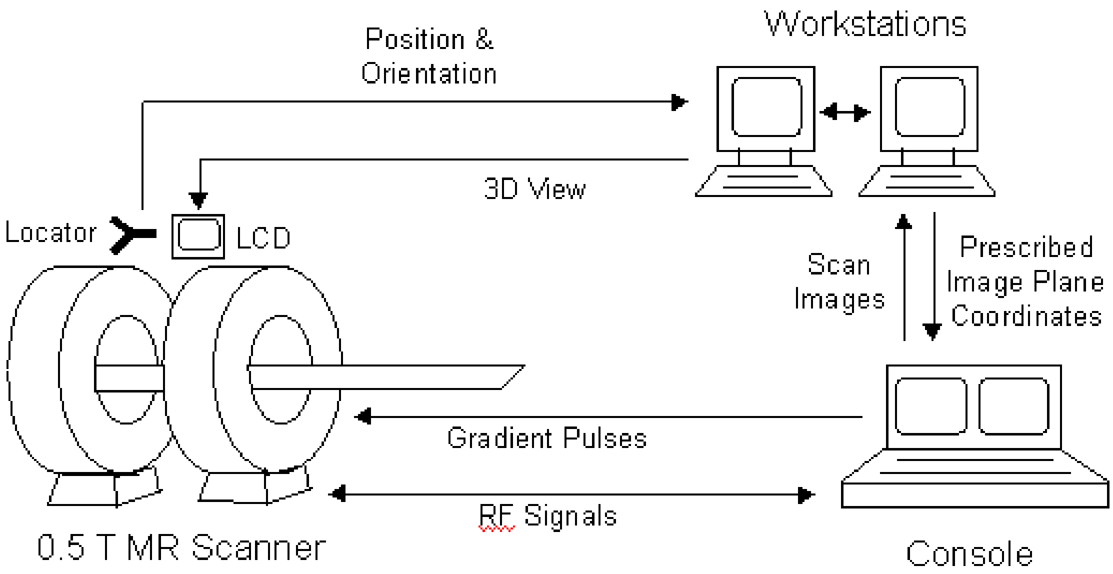

Figure 2-13: We added a workstation that runs the 3D Slicer, reads the position and orientation of the locator, receives real-time images, and displays output to the surgeon

The Signa SP console and Signa SP imaging workstation have a network connection with a TCP/IP interface by which we added a visualization workstation (Ultra 30, Sun Microsystems, Mountain View, CA) running the 3D Slicer. The 3D Slicer connects through a network socket to a server we created on the SP imaging workstation. Whenever the locator's position or orientation changes, or a new image is scanned, the server sends the new data. The peak transfer rate between the two workstations is presently 10MB/s, which transports a 256x256x16 bit image in 0.72 seconds. The server is concurrent so that when the 3D Slicer opens a connection to it, the server spawns a child process to perform the actual serving, and the parent server is said to sleep. This is important because if the server crashes, the 3D Slicer merely re-connects to the sleeping parent, and the surgery can continue uninterrupted.

Real-time images are incorporated into the 3D Slicer by setting one of the reformatted planes to automatically assume the same orientation in 3D space as the real-time scan. One of the slice's three layers (background, foreground, or label) can be set to display the real-time image, while the other layers can overlay function information, segmented structures, or high-quality, pre-operative anatomical data.

Figure 2-14: The surgeon can observe the 3D Slicer's graphical information while operating from within the interventional MR gantry.

The visualization workstation contains two Sun Creator3D graphics accelerator cards. One drives the 20-inch display placed in the control area of the surgical suite, and the other outputs only the 3D view and no control windows. Its signal is converted to NTSC TV and displayed on a color LCD panel in the scanner gantry.

2.6.2 Trajectory Assistance

In addition to slices and models, the locator is also rendered as a cylinder in 3D space at its correct location and orientation. The software can also provide means for assisted trajectory planning and guidance. Once a low-contrast target is identified using near real- time imaging, it is sometimes difficult to quickly find it again. This can be simply avoided by marking the probe's position with a bright red sphere on the 3D display, so the location can be easily returned to.

Trajectory assistance is achieved through marking the point of entry with one sphere, and the target tissue with another sphere, and rendering a cylinder between the two points. Thus the surgeon has only to align the rendered locator with the target in a video-game- like exercise to find the planned approach.

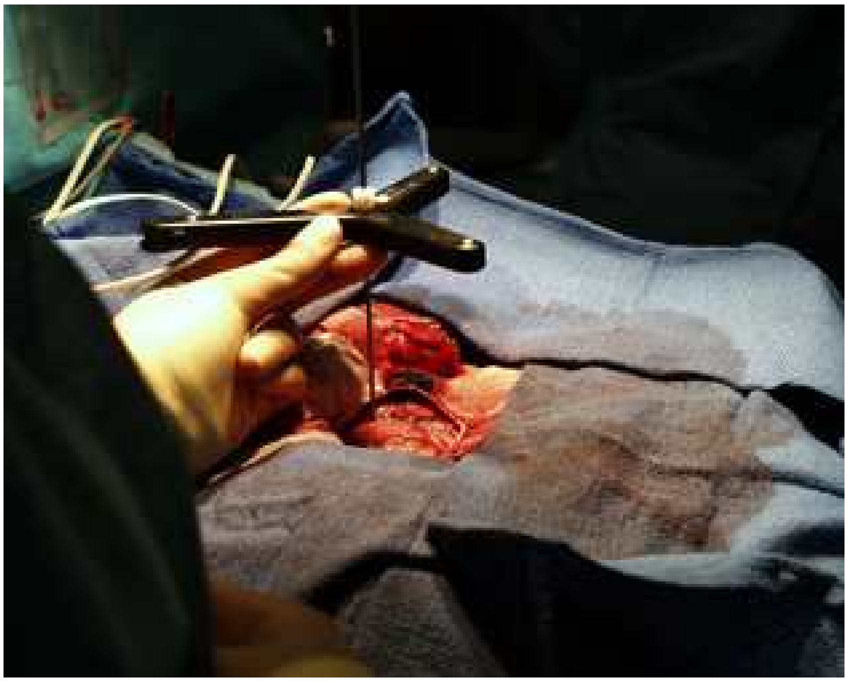

Figure 2-15: The optically-tracked handpiece has an LED on each prong of its handle. Various surgical needles may be inserted in the middle and the software is notified of the needle's length.

2.7 Medical Reality Modeling Language

Visualizing medical data involves combining various data sets into a single scene, and exploring the scene interactively. The usage of the 3D Slicer typically involves the creation of a scene from a variety of volume data sets, surface models derived from those volumes, and transformations derived from 3D registrations of both the volumes and models. We have found that the proper coordination of these items is easiest to obtain by the use of a hierarchical modeling paradigm as exemplified by the modeling systems and languages of graphics and CAD/CAM.

Toward this end, we created a novel file format for expressing 3D scenes composed of volumes, surface models, and the coordinate transforms between them. The Medical Reality Modeling Language (MRML) presents a format for describing scenes that consist of various types of data sets collected in various geometric locations. MRML files are not a copy of the data in another format. Instead, a MRML file describes where the data is stored so the data can remain in its original format and location.

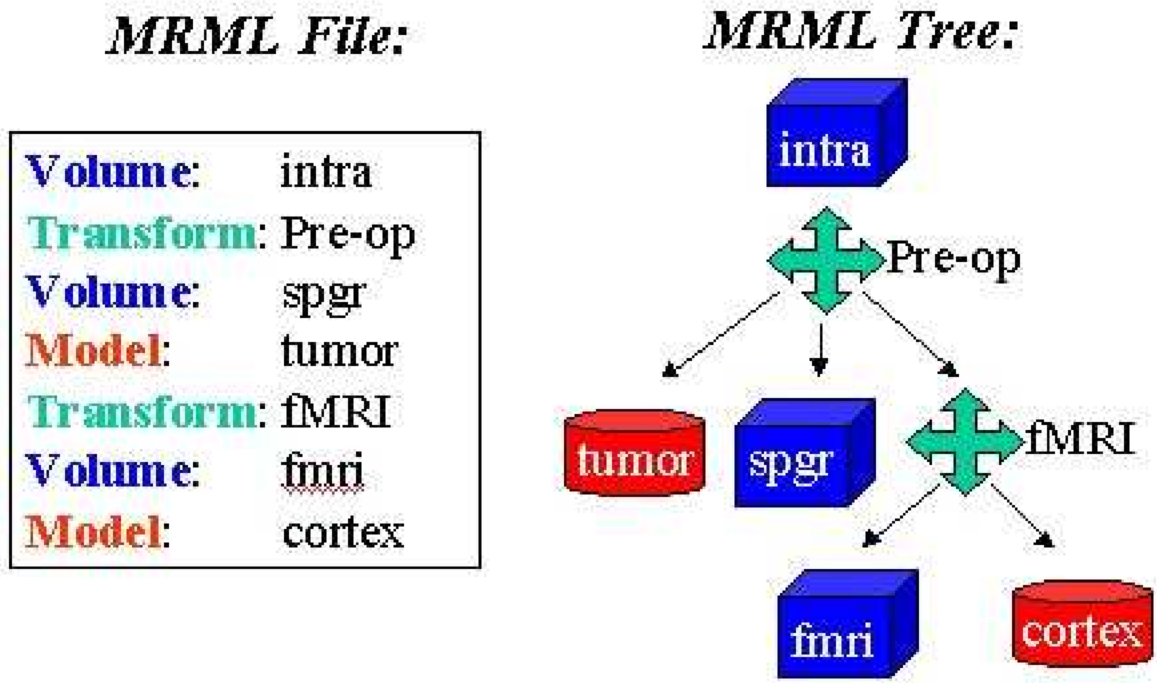

Figure 2-16: Pre-operative volumes named "spgr" and "fmri" were registered, and the "fMRI" transform was computed to align the "fmri" volume to the "spgr". Then, the "Pre-op" transform was computed to align this fused, pre-operative data to the intra-operative volume called "intra". Thus, the transform applied to the "fmri" volume is the concatenation of both transform nodes in the MRML file.

Second, a MRML file describes how to position the data sets relative to each other in 3D space. Third, a MRML file describes how to display the data by specifying parameters for rendering and coloring.

A scene is represented as a tree-like graph where volumes, models, and other items are the nodes in the graph. Each node has attributes for specifying its data. Transforms along each branch are concatenated to form a composite transform applied to the leaf node. Figure 2-16 displays an example MRML file (which has been simplified for illustrative purposes) and the tree-like graph it represents. The order of nodes in the file defines the tree structure; each transform is a parent node, and all the nodes following it, yet before the next transform, are its children. In the figure, the transform for the last volume in the file is the concatenation of all the transforms in the file. More complex trees, such as where multiple siblings have children, can be created using the separator construct.

The actual text of the first portion of this example MRML file is displayed below. Nodes possess more attributes than those listed in the file, because attributes only appear when their values differ from the defaults.

```
Volume (
   name intra
   filePrefix data/mri/I
   imageRange 1 14
   spacing 1.4 1.4 6.0
   rasToIjkMatrix -0.71 0 0 128 0 0 -0.71 128 0 0.167 0 12.5 0 0 1
)
Transform (
   matrix 1 0 0 9.5 0 1 0 5 0 0 1 0 0 0 0 1
)
Volume (
   name spgr
   filePrefix data/spgr/I
   rasToIjkMatrix -0.8 0 0 128 0 0 0.8 0 128 0 0 0.125 15 0 0 0 1
)
Model (
   name tumor
   fileName structures/glioma.vtk
   colorName Tumor
   opacity 0.7
)
```

An application programming interface (API) allows programs to easily parse and write MRML files.

Each type of node is briefly described below. More documentation is available from the SPL [SPL] in the guide named MRMLaid.

2.7.1 Volume Node

Volume nodes describe data sets that can be thought of as stacks of 2D images that form a 3D volume. Volume nodes describe where the images are stored on disk, how to render the data (window and level), and how to read the files. This information is extracted from the image headers (if they exist) at the time the MRML file is generated. Consequently, MRML files isolate MRML browsers from understanding how to read the myriad of file formats for medical data.

2.7.2 Model Node

Model nodes describe polygonal data. They indicate where the model is stored on disk, and how to render it (color, opacity, etc). Models are assumed to have been constructed with the orientation and voxel dimensions of the original segmented volume.

2.7.3 Transform Node

The output of a rigid-body registration is a rotation and translation expressed mathematically as a transformation matrix. These transforms can be inserted into MRML files as transform nodes. Each transform affects volumes and models that appear below it in the MRML file. Multiple transforms can be concatenated together.

2.7.4 Separator

A separator is not a node with attributes, but a construct for building MRML files. A separator encapsulates the nodes inside it such that they are invisible to nodes outside the separator. Separators can be used to achieve two types of effects:

- Any transform node inside a separator will only affect volumes and models that are positioned both inside the separator, and after the transform.
- Any model node inside a separator will be interpreted to belong to the volume node that is the most immediately above it and also inside the separator. Then,

registration of a volume will also move its models.

2.7.5 URL Node

A URL node can serve one of two purposes:

- State the path of another MRML file to effectively place the contents of the other file into this one.
- State a root path from which relative paths (in volume and model nodes) are based.

2.7.6 Config Node

Config nodes allow for a MRML browser, such as the 3D Slicer, to be configured at startup.

2.7.7 Color Node

Color nodes define colors by describing not only the actual color value, but also its name and a list of label values. One attribute of model nodes is the name of its color. When the 3D Slicer displays label maps, it colors each voxel by looking up the color associated with that label value. Thus, when label maps are displayed on reformatted slices, their colors automatically match the corresponding surface models in the 3D view.

2.8 Coordinate Systems and Transforms

This section describes the various coordinate systems used by the 3D Slicer and the coordinate system transformations used to drive reformatting.

2.8.1 The RAS Coordinate System

The RAS coordinate system defines a right-handed, three-dimensional space with millimeter units and orthogonal axes relative to the patient. The R axis extends from the patient's left to Right side, the A axis runs from the posterior to the Anterior, and the S axis runs from the inferior to the Superior (feet to head). Every scanner has its own coordinate system, and patients can also assume different positions within the same scanner. By working in RAS space, the 3D Slicer becomes scanner independent, and automatic registration has an immediate head start because the patient's nose will always roughly point in the direction of the A axis.

2.8.2 The IJK Coordinate System

The IJK coordinate system refers to a framework for accessing the 3D array of voxels in a volume scan. The axes extend in the order that the voxels are stored on disk. Typically, this implies that the I axis runs across an image from left to right, the J axis runs down an image from top to bottom, and the K axis extends from the first slice acquired to the last. The coordinates of a particular voxel equal the indices into a 3D array that a computer would use to access the voxel.

2.8.3 The XYZ Coordinate System

The position of the tracked locating device is reported with respect to a calibration point in the center of the open MR imaging volume. These XYZ coordinates are in millimeter units, and conversion to RAS is merely a matter of permuting the axes. The permutation is a function of the patient entry (feet-first or head-first), patient position (supine, prone, left-decubitus, or right-decubitus), and the docking position of the table (front or side). RAS coordinates equal the scanner's XYZ coordinates when when the patient is lying head-first, supine on a table docked at the front of the magnet.

2.8.4 Converting from RAS to IJK

Reformatting is the process of identifying which RAS points lie on the desired image plane, and then transforming those points to IJK space to obtain indices into the array of voxels stored in computer memory. This conversion from RAS to IJK space is stored in an RAS-to-IJK matrix in the MRML file. The matrix can be computed from the corner points of the volume, when available,or from the voxel size, scan orientation (axial, sagittal, or coronal), and slice order (such as left-to-right, or right-to-left for sagittal scans).

2.8.5 Handling Rigid Registration

Rigid registrations that allow 6 degrees of freedom (3 translation, 3 rotation) can be represented by a 4x4 transformation matrix. Registrations between volumes are stored as heirarchies of transform nodes in the MRML file. These transforms are concatenated when the MRML file is read to form a REF-to-RAS transform, where *REF* refers to the *ref*erence volume that the data was registered to. If the 3D Slicer is used to perform registration, it computes a WLD- to-REF transform which relates the 3D Slicier's view of the world (the RAS coordinate system for its 3D view window) to the reference volume. At bootup, the WLD-to-REF transform for each volume is initialized to the identity matrix.

For the purpose of reformatting, points in the 3D Slicer's RAS world space are converted to the IJK space of the specific volume being reformatted This is achieved by multiplying the points by a transformation matrix, $\underline{M}$, that is generated by concatenating the RAS-to-IJK, REF-to-RAS and WLD-to-REF transforms. In the equation below, WLD is a point from world space, and IJK is a point from IJK space.

$$\underline{M} = RAStoIJK \cdot REFtoRAS \cdot WLDtoREF \tag{2.1}$$

$$\underline{M} \cdot WLD = IJK \tag{2.2}$$

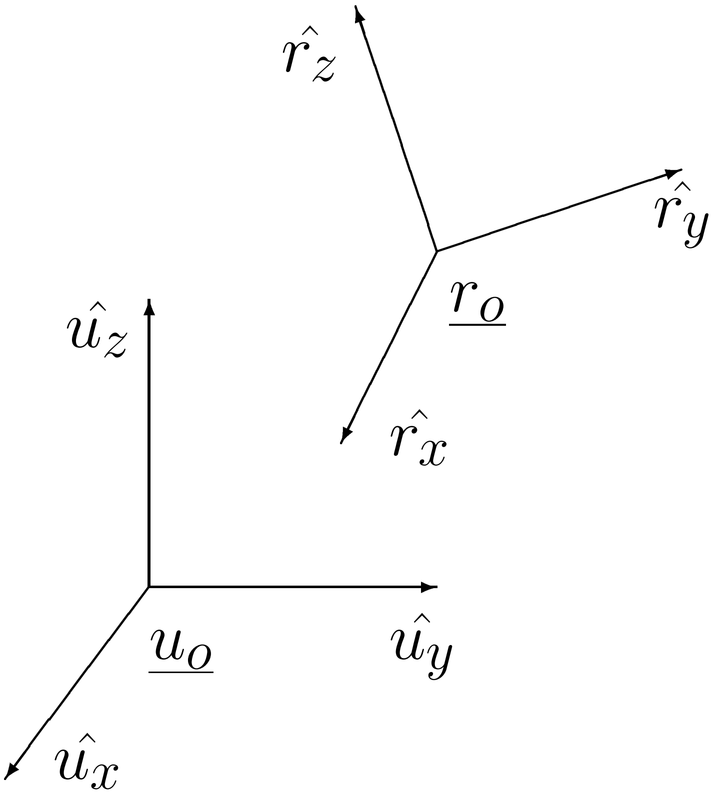

Figure 2-17: The location of a reformatted plane is expressed as a matrix that transforms the base coordinate system, u, to the coordinate system, r, of the reformatted plane.

2.8.6 Computing the Reformat Matrix

The orientation and position of a reformatted plane are specified by what we call the reformat matrix. This is the matrix that would transform a plane initially centered at the origin, and lying in the RA-plane, to the desired location of the reformatted slice. Figure 2.8.6 shows the coordinate system of the desired reformatted slice (defined by the axes $\hat{r_x}$, $\hat{r_y}$, and $\hat{r_z}$) has been rotated and translated with respect to the base coordinate system defined by the axes $\hat{u_x}$, $\hat{u_y}$, and $\hat{u_z}$.

The 4x4 matrix $\underline{R}$ that rotates the r coordinates system to the u system will take the form of equation 2.3.

$$\underline{R} = \begin{bmatrix} r_{11} & r_{12} & r_{13} & 0 \\ r_{21} & r_{22} & r_{23} & 0 \\ r_{31} & r_{32} & r_{33} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{2.3}$$

The upper left 3x3 submatrix of $\underline{R}$ is orthogonal, which means that the rows form a set of orthogonal unit vectors. When $\underline{R}$ is multiplied by $\hat{u_x}$, it forms the $\hat{r_x}$ axis of the reformatted system as shown below. (See [Hearn and Baker1997] for more details.)

$$\underline{R} \cdot \hat{u_x} = \hat{r_x} \tag{2.4}$$

$$\underline{R} \cdot \begin{bmatrix} 1 \\ 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} r_{11} \\ r_{21} \\ r_{31} \\ 1 \end{bmatrix} \tag{2.5}$$

If the reformatted system is translated from the original system by a translation $\underline{r}_o = (t_x, t_y, t_z)$, then the translational component of the offset may be specified as:

$$\underline{T} \cdot \underline{u}_o = \underline{r}_o \tag{2.6}$$

$$\underline{T} = \begin{bmatrix} 1 & 0 & 0 & t_x \\ 0 & 1 & 0 & t_y \\ 0 & 0 & 1 & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{2.7}$$

The complete reformat matrix, $\underline{M}$, that transforms any points between the two systems is the concatenation of the $\underline{R}$ and $\underline{M}$ transforms:

$$\underline{M} = \underline{R} \cdot \underline{T} \tag{2.8}$$

$$\underline{T} = \begin{bmatrix} r_{11} & r_{12} & r_{13} & t_x \\ r_{21} & r_{22} & r_{23} & t_y \\ r_{31} & r_{32} & r_{33} & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{2.9}$$

This reformat matrix, M, is computed by the 3D Slicer in one of two different ways. When the tracked locator device is driving the location of the reformatted slices, then the elements of M are filled in by using the position of the locator's tip as (t_x, t_y, t_z), and the normal direction of the locator (from handle to needle tip) is used as $\hat{r_z}$, and the transverse axis (from handle to cord) is used as $\hat{r_y}$, and $\hat{r_x}$ can be computed by their cross product.

When the user, rather than the locator, is driving the location of the reformatted planes, and the slices are oblique, then they may be generated relative to the viewing direction instead of the locator. Here, the focal point (center of rotation of the 3D dislay) is used as the translation, the view direction is used as $\hat{r_z}$, and the "up" direction is used as $\hat{r_y}$.

2.9 User Interface

This section outlines the methodology behind the user interface design. Readers interested in details or instructions on how to operate the 3D Slicer are encouraged to consult the 3D Slicer User's Guide available from the Surgical Planning Lab [SPL].

2.9.1 Efficiency through Simplicity

Though often overlooked by researchers, the user interface becomes more important in a surgical setting. The opportunities for operator error must be minimized along with the time it takes for a clinician to respond to a surgeon's request. Consequently, we present the user with only two, non-overlapping windows: one for controls, and one for viewing the data. This prevents the screen from becoming cluttered with confusing, overlapping, pop-ups. Figure 2-18 presents a screen shot of a session in

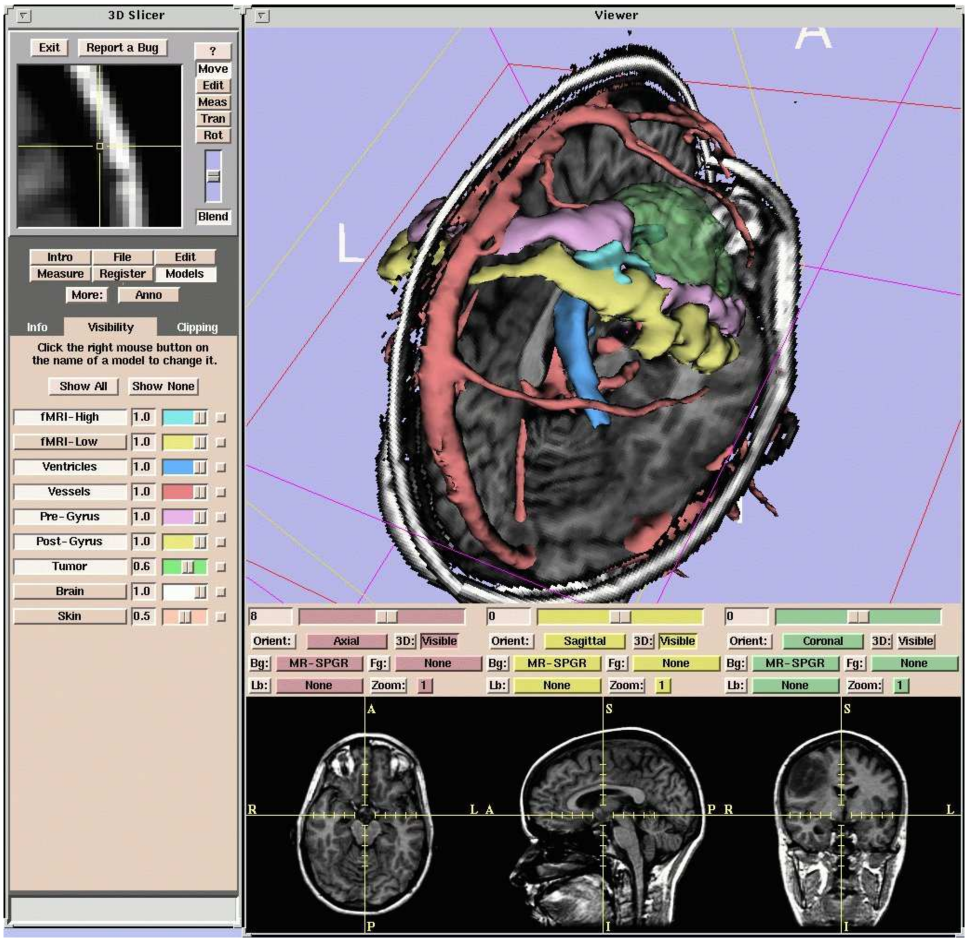

Figure 2-18: The system interface features a control window on the left, and a re-sizeable view window on the right.

progress.

2.9.2 Display Windows

The reformatted images are displayed in 2D image windows (one for each slice) and texture-mapped onto graphics and rendered in a 3D view. The planes are positioned in 3D space in their correct location; that is, where the reformatted voxels exist in the volume. We found that users often need to vary their focus between the 2D images and the 3D rendering. We provided three different methods in attempting to create a convenient interface for all users. First, there is a foveal view in the form of a 2D window that provides a magnified version of the area around the cursor. Second, each 2D image may be zoomed independently as shown in Figure 2-7. Third, the display mode may be set to show a huge 3D view, a large 3D view and the three 256x256 2D images (as shown in Figure 2-18), or all 4 windows set to the same size of 512x512.

2.9.3 Modular Design

The control window contains a main menu of buttons. There is a button for topics such as Volumes, Models, Registration, Measurements, Guidance, etc. When a button is pressed, the window displays a set of virtual index tabs corresponding with that topic. The first tab in each set displays a help page for that topic.

The internal architecture of the 3D Slicer is oriented around the same, modular paradigm of the user interface. Adding a new module (for interfacing with a robot, for example) only consists of adding one new Tcl file that codes the tabbed pages, and editing a configuration file that adds a button to the main menu.

2.9.4 Protecting Against Crashes

The software state, including the view parameters and condition of every control item, can be named and saved to facilitate rapid toggling between various states. Additionally, the current program state is automatically written to a file every minute. Then, in the event of a program crash, the system can be restarted and quickly loaded into its last known configuration to continue operation with minimal interruption to the surgery.

2.10 Chapter Summary

This chapter has explored the evolution of the system architecture into an environment for performing registration and segmentation, integrating with interventional imaging, and visualizing all available information in a common display. The next three chapters highlight the systems application domains.

Chapter 3

Surgical Planning

Outside of the operating room, the 3D Slicer is routinely relied upon as a stand-alone software tool. This chapter explores its value for a variety of applications related to surgical planning and visualization.

3.1 Diagnosis from Multiple Scans

Diagnostic MR scanning protocols often involve collecting multiple scans with various parameters to be sensitive to different tissue properties. Recall from Chapter 2 that MR imaging is performed by measuring the radio signal emitted by magnetic dipoles (hydrogen nuclei) as they relax back to their equilibrium position following excitation by a momentarily-applied magnetic field. The dipoles cannot merely align themselves with the magnetic field as little bar magnets would, because the laws of quantum physics restrict these dipoles to be in one of two states. They precess like spinning tops, and the "tops" can make one of two angles with the axis of rotation. The applied magnetic field excites about one in a million of the dipoles to flip states, and the total sum of all these miniature magnets is a magnetization that decays once the field ceases to be applied.

This decay has two separate components referred to as T1 and T2 relaxation. T1 relaxation occurs as the dipoles return their orientation to the equilibrium position, and T2 relaxation results from the precession of the dipoles falling out of phase with each other. The rate of T1 and T2 decay varies depending on the molecular chemistry of the tissue inhabited by the hydrogen nuclei. Scanning parameters can be set so that the source of image contrast (light and dark regions) is weighted more toward either the T1 or T2 relaxations.

In many instances, physicians acquire both T1 and T2 weighted MRI. For example, extracting a well-defined tumor boundary from diagnostic images may be hindered by surrounding edema. Edema, or liquid diffused between cells, spreads finger-like into the white matter, while avoiding the gray matter and cortex whose cell packing is too dense to harbor as much fluid. The extra-cellular fluid of edema and increased intra-cellular fluid of tumors can be confused when attempting to ascertaining the tumor/tissue interface. Ambiguity can be diminished by having both T2-weighted MR images and T1-weighted MR images with contrast. A contrast medium (liquid that appears bright on MRI) is administered to the patient, and taken up more by the areas of active tumor tissue. The contrast agent forms a hyperintense region on MRI where the agent leaks out of vasculature into tissue. This occurs where the blood-brain barrier breaks down, and is thus an indication of a high grade, rather than a low grade, glioma (a mass created in the brain by the growth of abnormal cells, or the uncontrolled proliferation of cells).

A direct diagnosis can be better attained by presenting all available data (both T1 and T2-weighted MRI) to the physician for simultaneous review. Figure 3-1 is a screen shot of the 3D Slicer displaying side-by-side slices at similar locations, although from different acquisitions. This particular case is a prime example of requiring multiple, complementing scans to arrive at a correct diagnosis. The T1-weighted MRI with contrast reveals a small hyperintense region inside a large hypointense area. This suggests there is either a large low grade glioma with a small high grade core, or a small metastasis surrounded by edema. It is imperative to disambiguate this diagnosis because the large hypointense region extends over the central sulcus — home of the motor cortex which controls voluntary movement and dramatically impacts operative strategy. While both edema and glioma appear hypointense on T1-weighted MRI, glioma is strongly hyperintense on T2-weighted MRI. The T2-weighted

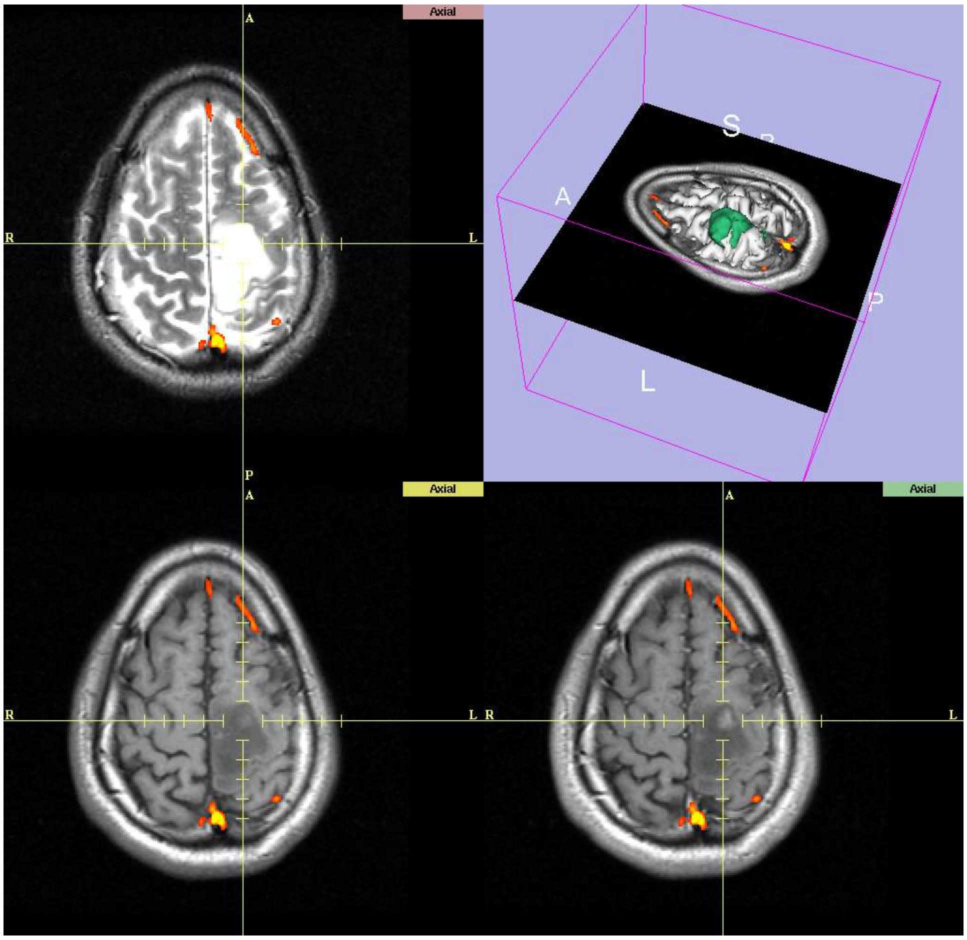

Figure 3-1: The 3D Slicer aids diagnosis by simultaneously presenting T2-weighed MRI (top left), T1-weighted MRI (bottom left), and T1-weighted MRI with contrast (bottom right). The T2-weighted MRI clearly delineates the tumor boundary from edema, while the T1-weighted MRI with contrast details the tumor's internal character.

MRI clearly demarcates the tumor boundary but fails to detail its internal character. In this case, the T2-weighted MRI completes the diagnosis as a large low grade glioma without edema. In cases like this one, 3D Slicer's capability for simultaneous display of multiple scans aids diagnosis through convenient comparative evaluation of all available data.

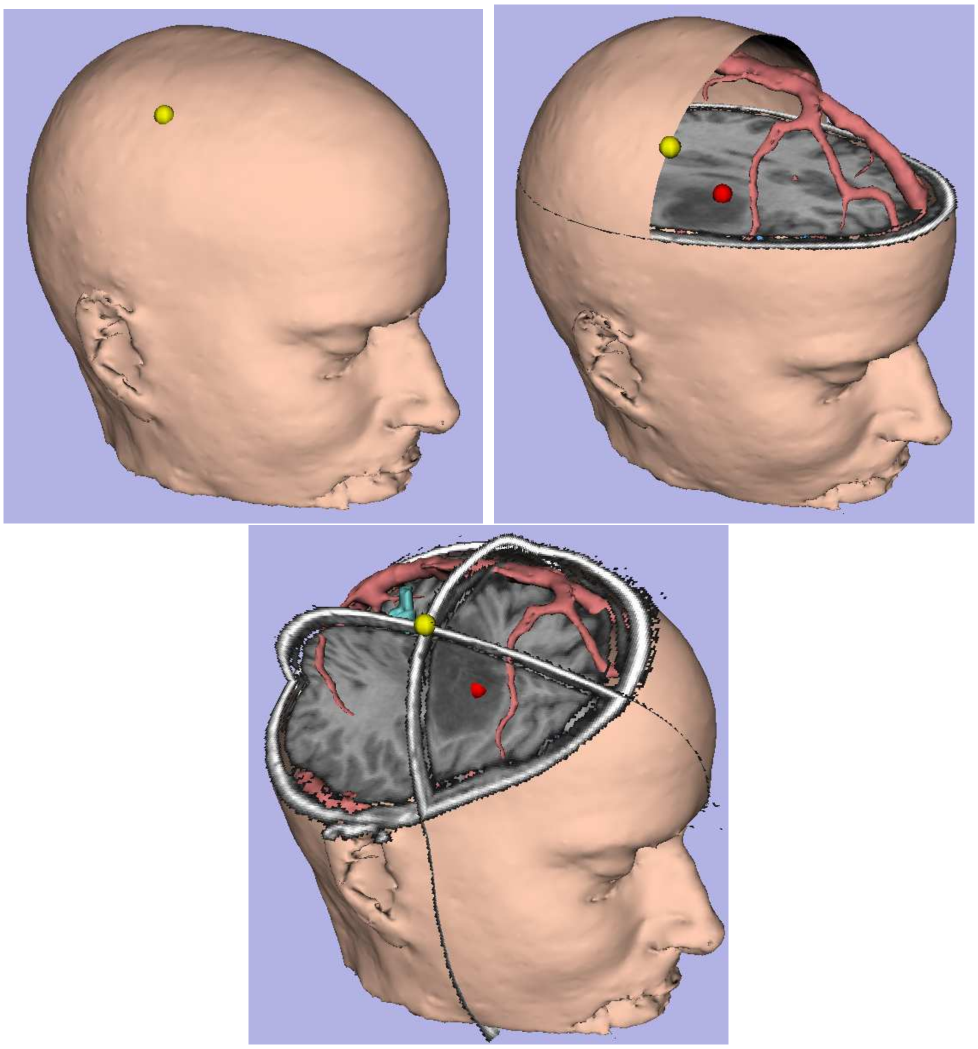

Figure 3-2: Trajectory planning is performed by positioning points for entry (yellow)and target (red), and examining the reformatted planes oriented relative to the connecting path.

3.2 Trajectory Planning

A key component of neurosurgical planning is plotting an approach trajectory that avoids critical structures such as blood vessels or the motor cortex. The 3D Slicer facilitates trajectory planning with the three-step process shown in Figure 3-2. A yellow sphere is placed on the skin model at the entry point under investigation. A red sphere is positioned on the target anatomy. Then, with the click of a button, the 3D view is set to align the viewing angle along the two points — from entry to target. The three reformatted planes become oriented relative to this trajectory.

The plane perpendicular to the candidate trajectory may be slid along its path to visualize the structures that will be encountered en route. Additionally, the other two orthogonal planes allow views through the tissue lining the approach.

3.3 Orthopedic Surgical Planning

While the 3D Slicer is routinely applied to neurosurgery, it has also provided assistance in planning for an orthopedic surgery. And this trend should grow in the future. As an example case, a twelve- year-old girl was diagnosed with a benign pelvic lesion four years earlier. However, a current follow-up exam revealed significant growth, so the 3D reconstruction shown in Figure 3-3 was created to aid the doctor's decision for partial or total removal.

This case highlights another feature of the 3D Slicer. These orthopedic surgeons demanded high-resolution data with minimum interpolation. A very high-resolution MR volume can take a long time to acquire, so we modified the typical scanning protocol for optimum visualization in the 3D Slicer. Instead of collecting many slices with narrow separation, a few slices with wide separation were acquired in three orthogonal directions. Then, one slice from each volume was displayed simultaneously in the 3D Slicer. We used the reformatting mode that avoids interpolation by restricting the planes to be positioned in 3D space in coincidence with the actual slices.

3.4 Parametrically-Colored Surface Models

It is difficult to accurately visualize the thickness of hollow, 3D structures, such as the bladder wall. For example, one approach we could try would represent the outer and inner walls as different surface models, and render the outer wall semi-transparently to offer a view of the inner wall, yet the thickness is very difficult to ascertain accurately from such a view. We could also use the 3D Slicer's ability to clip away models with a reformatting plane so that the bladder's cross-section can be viewed straight-on, but then a feel for the thickness along the entire surface would not be acquired

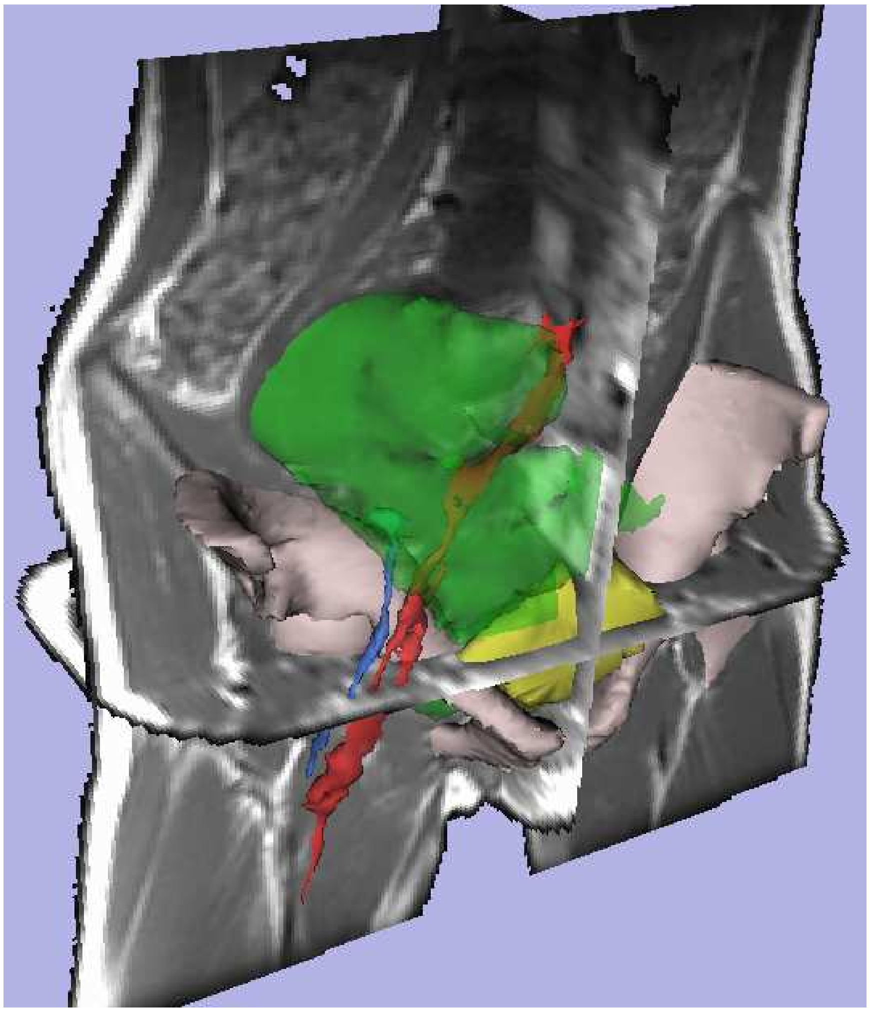

Figure 3-3: Three orthogonal volumes are visualized simultaneously with one slice through each volume. This high-resolution 3D scene aided the doctor's decision making. Color scheme: bone (cream), tumor (green),bladder (yellow), vessels (red), nerves (blue).

simultaneously.

Therefore, software was written in our lab to measure the thickness of the bladder wall at every point along the inner wall. The result is a scalar value for each vertex of the polygonal model. The 3D Slicer can color the surface of the model by applying a color lookup table to the scalar value at each vertex. The resulting display, shown in Figure 3-4, allows a doctor to instantly assess the thinning portions at the bottom of the bladder. This experiment demonstrates the ability to interface other medical analysis software with the 3D Slicer's capabilities for integrated visualization.

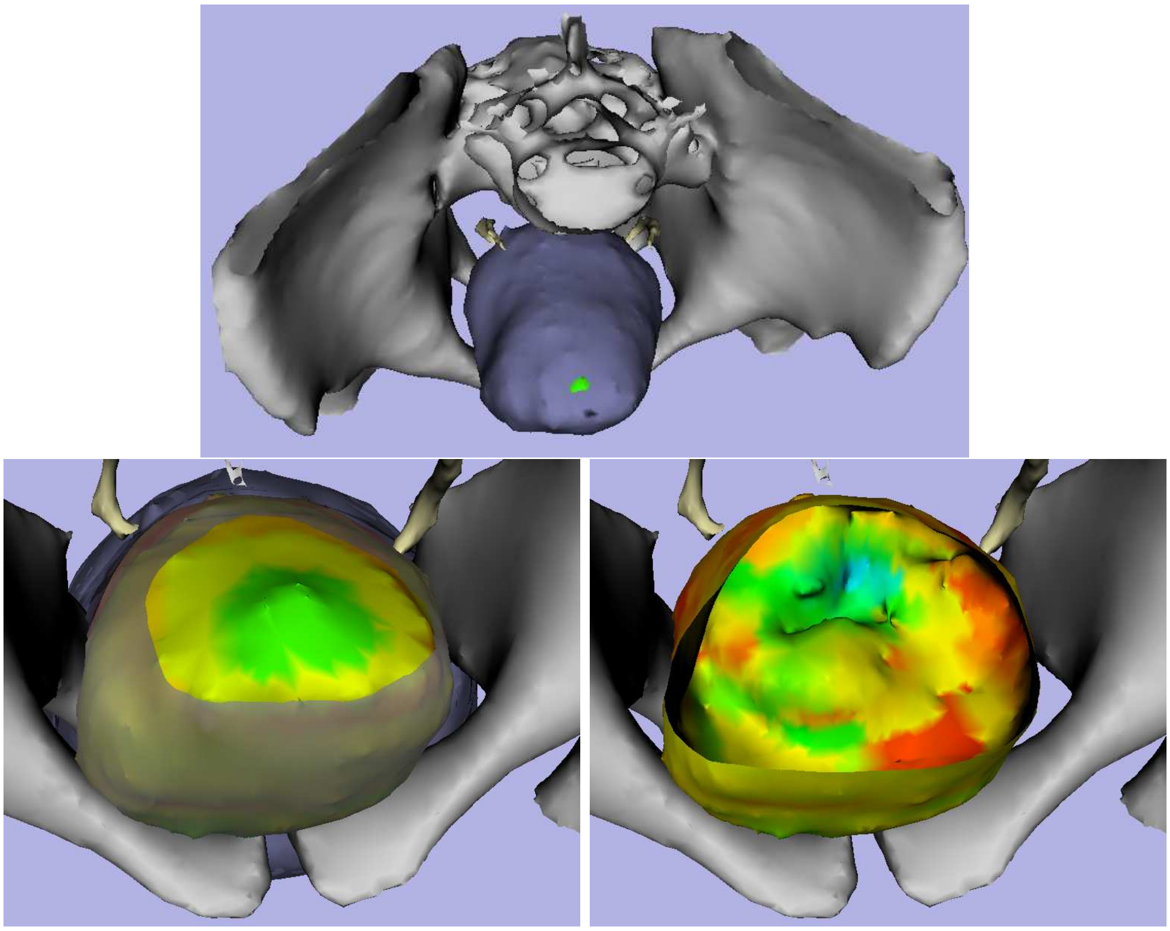

Figure 3-4: TOP: The bladder (blue) positioned within the pelvis. BOTTOM-LEFT: The bladder's outer wall (blue) is semi-transparent and also partially cut away to reveal the inner wall (color-coded). BOTTOM-RIGHT: The inner wall is also cut away to peer inside the bladder at the color-coded thickness of the bottom wall.

3.5 Construction and Application of Anatomy Atlases

The advent of high-resolution medical scanning combined with painstakingly detailed segmentation by experts has led to the creation of anatomy atlases [Höhne1992, Golland et al.1998, AnatomyBrowser]. Such atlases provide collections of intricately segmented structures for use in educating medical students and supplying the templates in template-driven segmentation [Warfield et al.1998]. Atlases can be created and visualized using the 3D Slicer as exemplified by the spine atlas depicted in Figure 3-5.

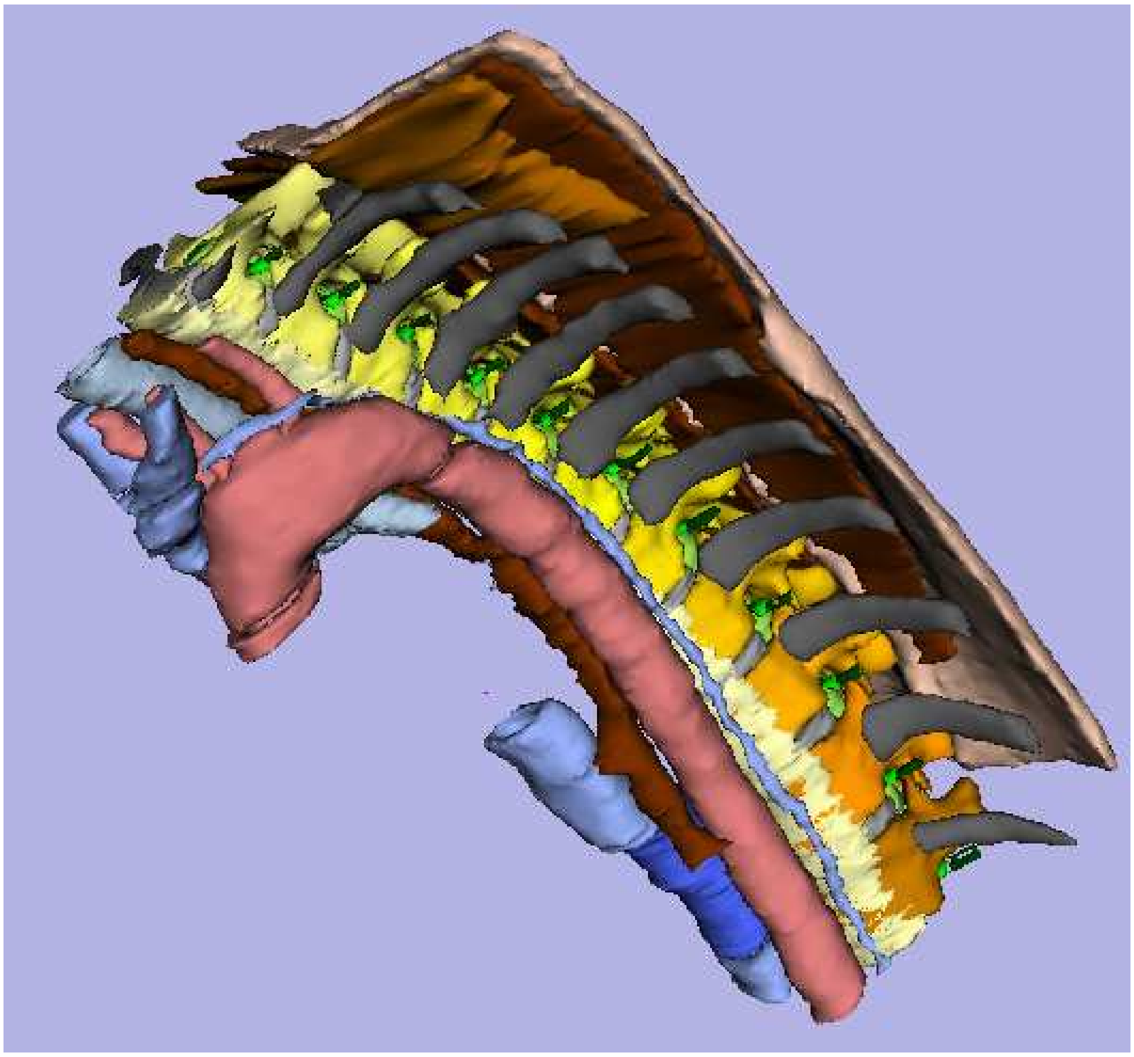

Figure 3-5: A spine atlas created and visualized using the 3D Slicer

An entire atlas can be morphed to match a scan of a certain patient. The registration is driven by comparing the prominent features of the patient-specific data with the corresponding data in the atlas, which is a subset of the entire atlas. Once registered, any structure may be pulled from the atlas and rendered along with the patient's data for surgical planning in the 3D Slicer.

This thesis has, in essence, described augmented-reality, in which data present in medical scans is enhanced or combined with data from other scans of the patient. The method here deviates from that theme by not just *augmenting* information, but rather *adding* information that is not clearly present in any of a patient's medical scans. Figure 3-6 displays an example application where the corticospinal tract was added to a patient's data despite not appearing on the scan, or at least appearing too vague to segment. The tract's model from the axis, on the other hand, was constructed by experts segmenting ultra high-resolution data.

3.6 Summary of the 3D Slicer's Role in Surgical Planning

This chapter demonstrated the 3D Slicer's varied applicability to surgical planning. The ability to register and reformat slices through multiple data sets simultaneously aids diagnosis through providing a convenient display of heterogeneous information. The ability to quickly align the reformatted slices along various approach trajectories enables investigation, and a form of simulation, of the surgical tactic. As researchers develop new analysis tools, such as the parametrically-labeled bladder wall presented here, the results can be readily visualized with other surface models and slices in the 3D Slicer's integrated visualization environment. Lastly we saw that atlases may be constructed using the 3D Slicer's manual segmentation suite. Once other software warps the atlas to register with a patient's scan, selected models from the atlas may be visualized in the 3D Slicer accompanied by the patient's actual data.

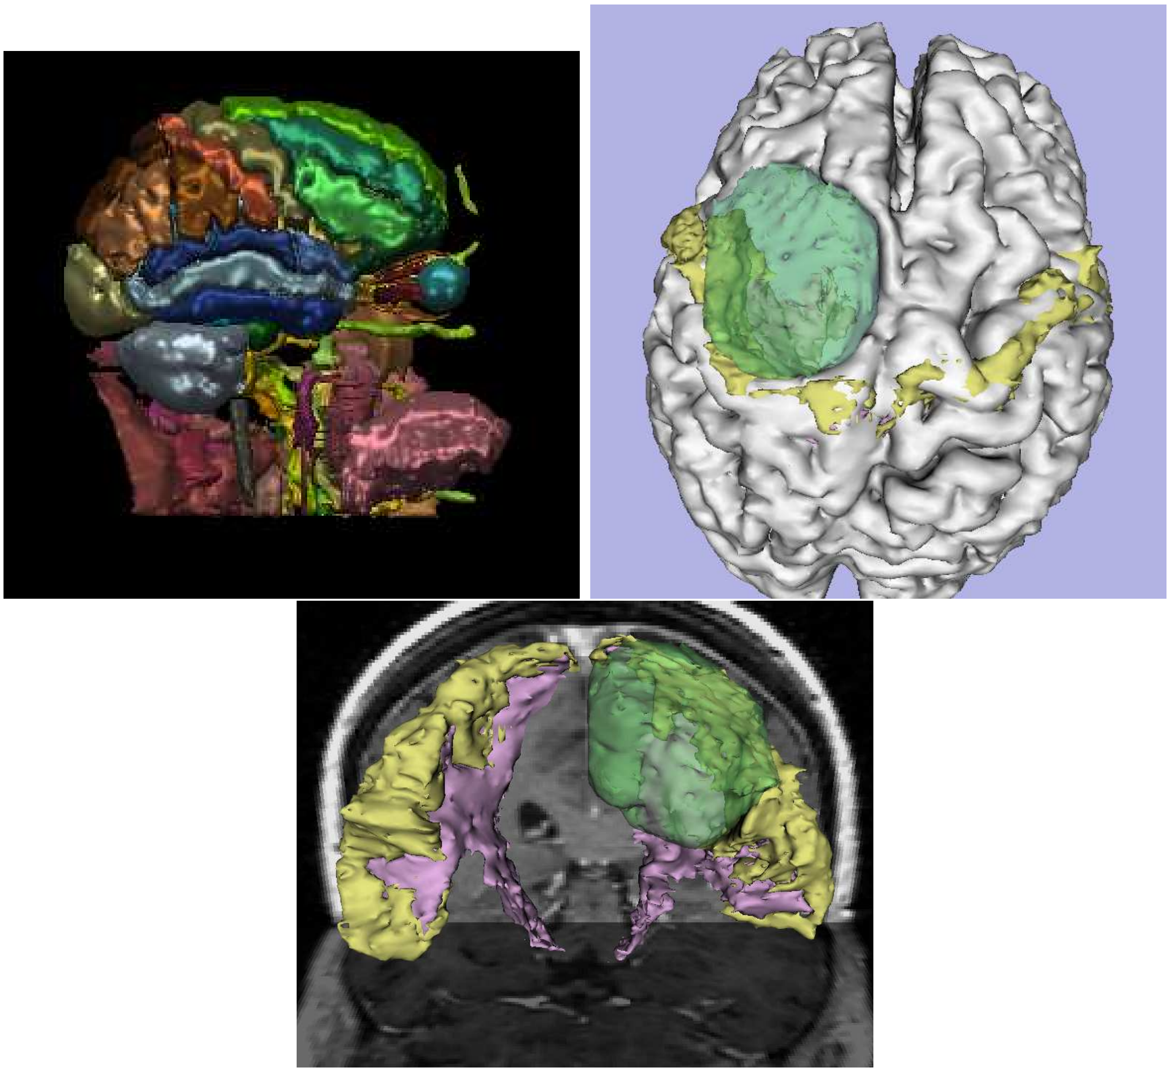

Figure 3-6: TOP-LEFT: an example brain atlas. TOP-RIGHT: tumor (green) near the post-central gyrus (yellow). BOTTOM: corticospinal tract (pink) from an atlas is added to the patient's data in the top-right image.

Chapter 4

Surgical Guidance

After several clinical trials during development, we have used the 3D Slicer as a navigational tool in 22 cases between January and May of 1999. This initial feasibility phase proved the 3D Slicer to be a stable and reliable application. As we apply the 3D Slicer on a routine basis, a clinical evaluation of its influence on surgical decision making and resection control is being conducted.

This chapter describes the utility of the 3D Slicer in the intra-operative setting. The chapter is organized in an order representative of an actual surgery: we begin with clinical setup and proceed with descriptions of image registration, trajectory planning, tissue identification, incorporating intra-operative imaging updates, and performing post-resection validation.

4.1 Clinical Setup

The 3D Slicer is setup and tested for a couple of minutes before the patient enters the MR. This involves checking that the Slicer's output is properly displayed on the monitors in the operating room, and verifying that the locator is being tracked by the Slicer.

After the patient is positioned within the interventional MRI, an initial 3D volume is acquired. Another check is made to verify that this volume, as rendered in the Slicer, is aligned with the patient's actual position by localizing anatomical landmarks. For example, we typically point the locator at the patient's nose, chin, and ear, and verify that the real-time graphics rendering agrees.

4.2 Image Registration

4.2.1 Merging Pre-operative and Intra-operative Data

For cases in which pre-operative data sets have been fused using rigid registration, a second rigid registration using mutual information is performed to the intra-operative setting. The pre-operative data sets are thus aligned with the coordinate frame of the interventional MR scanner by relating them to the newly scanned volume. The relation is expressed as a transformation matrix that is inserted into the MRML file for the case. We treat the MR scanner as the reference frame and choose to transform the pre-operative data to preserve the highest accuracy for the measurements of intra-operative movements relative to earlier intra-operative data.

As an example of a case, refer to Figure 4-1, where a pre-operative SPGR [1] volume is registered to an intra-operative T2-weighted volume. Surface models generated from the pre-operative data are shown in conjunction with the intra-operative images in Figures 4-2.

4.2.2 The Challenge of Registration in the Open MR

One caveat of registering intra-operative MR data stems from magnetic field inhomogeneity. The high cost of MR scanners is driven not by the technology needed to generate a strong magnetic field, but by the requirement that the field strength be constant to within 1-13 parts per million throughout the entire imaging volume. Unlike conventional MR scanners, the vertically open MR maintains such homogeneity in a small volume only the size of a basketball (30cm diameter).

Patients are positioned in the magnet to optimize the surgeon's access to the area of interest, which sometimes leaves features near the extremities, such as the tip of

[1] Spoiled Gradient Recalled Acquisition in the Steady State

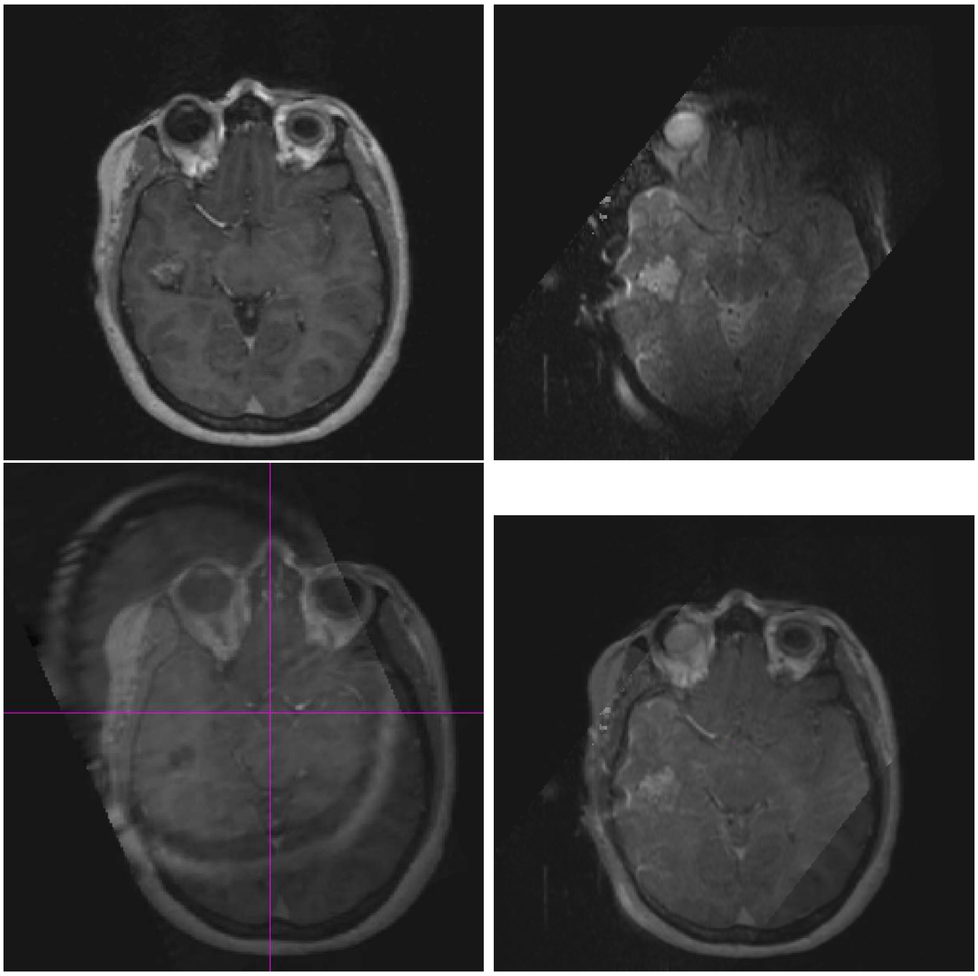

Figure 4-1: TOP-LEFT: Pre-operative SPGR. TOP-RIGHT: Intra-operative T2-weighted. BOTTOM-LEFT: Overlaid before registration. BOTTOM-RIGHT: After registration

the skull, near the boundary for homogeneity. This can present a problem for fully automatic registration algorithms that rely significantly on the strong signal present in the skin. In particular, we found that ill-positioned patients experience a flattening of the head in the interventional images. This usually does not present a problem for accurately tracking the locator since the tumor tends to be positioned near the center of the magnetic field. We scanned a healthy volunteer in a position that illustrates a worst-case scenario, and the results are presented in Figure 4-3. Proper patient positioning is presently the preferred practice for preventing this problem. Antoni Raimondi, a great, Italian pediatric neurosurgeon said, "In love, war, and surgery, positioning is the most important thing."

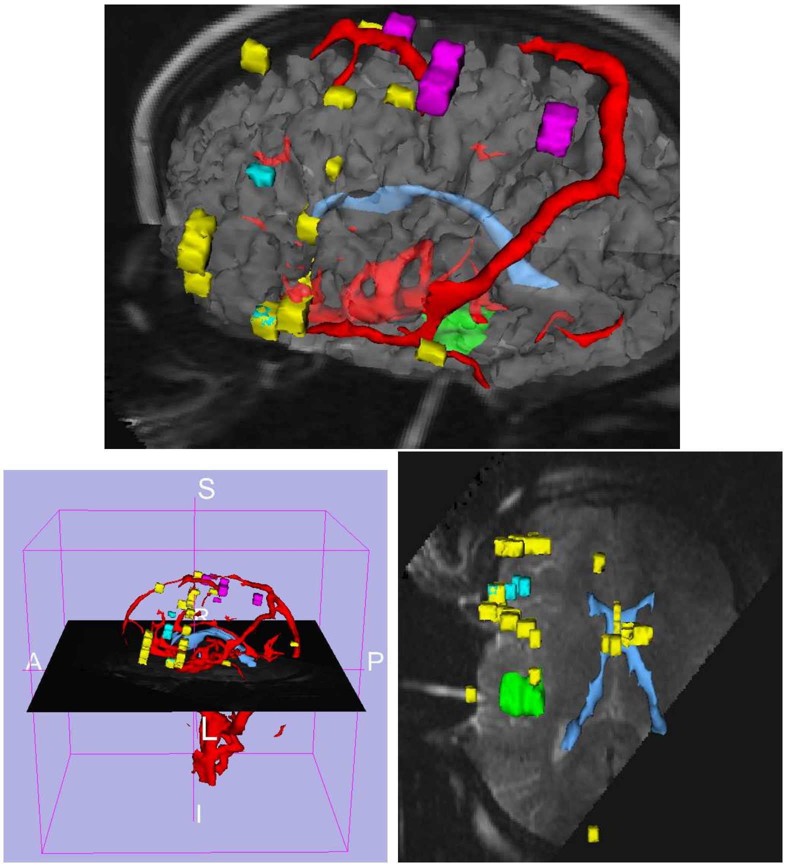

Figure 4-2: These images shows the capacity of displaying intra-operative information with registered, multi-modal, pre-operative information for updated guidance. The pre-operative MRI (sagittal plane) and the models (brain = white, ventricle = blue, lesion = green) are fused with the models of the separate MR phase contrast angiography and the functional studies (fMRI = magenta, movement tasks = yellow, complex speech tasks = blue) and the intra- operative T2.

4.3 Trajectory Planning

4.3.1 Peering Inside before Going Inside

In many cases, the optimal trajectory from the skin surface to the tumor is obvious from the 2D images. The 3D Slicer becomes valuable in cases where the tumor is deeply situated or dangerously near critical functional tissue or vasculature. This is where generating a surface model of the skin from either the pre-operative or intra-operative data comes into play. Prior to the craniotomy (opening of the skull), the surgeon maneuvers the locator to point at his intended craniotomy at various angles.

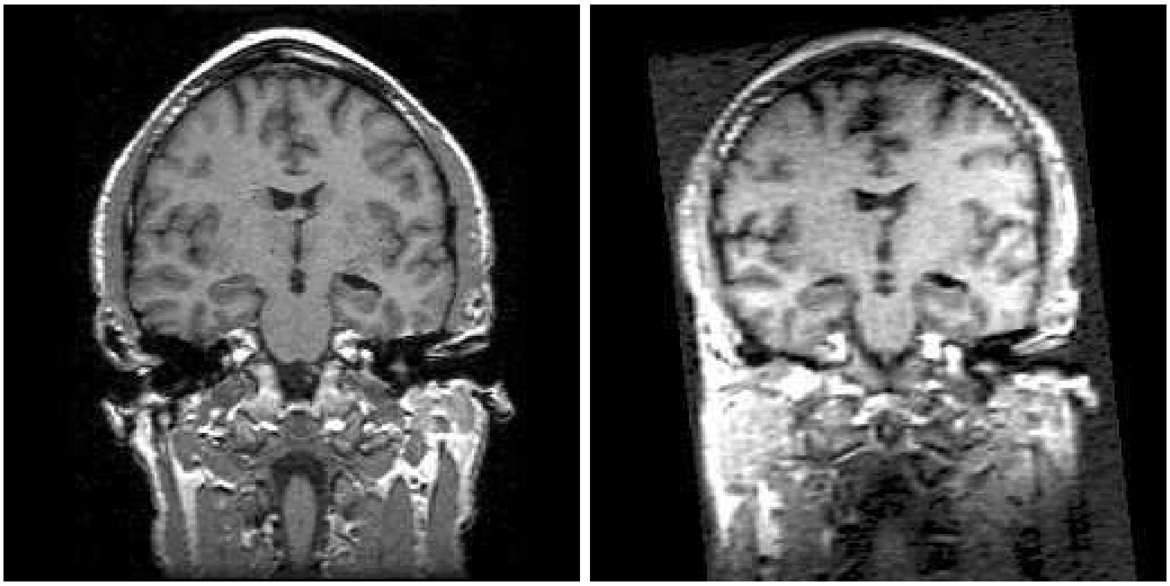

Figure 4-3: LEFT: Pre-operative SPGR. RIGHT: Intra-operative SPGR appears to have a flattened head.

The pointer and its trajectory are rendered with the anatomy in a view that only a computer can provide.

As the tracked pointer moves within the surgical field, the reformatted slice planes follow its position, sweeping through the volumes. The surgeon verifies, and if necessary, alters, the planned approach by visualizing it on the display relative to all the surface models of critical structures.

Figure 4-4 illustrates the locator (green) negotiating a subtemporal approach (beneath the brain's temporal lobe) prior to craniotomy for a low-grade, pediatric glioma (tumor) deeply situated in the brainstem.

4.3.2 Biopsies

Besides resections (tumor removal), the 3D Slicer has also seen clinical experience in two biopsies where the goal is to obtain a tissue sample for histopathological testing. A needle is inserted though the brain to the deep site for tissue extraction. If the needle's trajectory is only slightly incorrect, the needle's tip veers off target. During penetration, images are collected for the attending radiologist to assess whether the needle rests inside the suspect tissue. Positioning is critical because negative

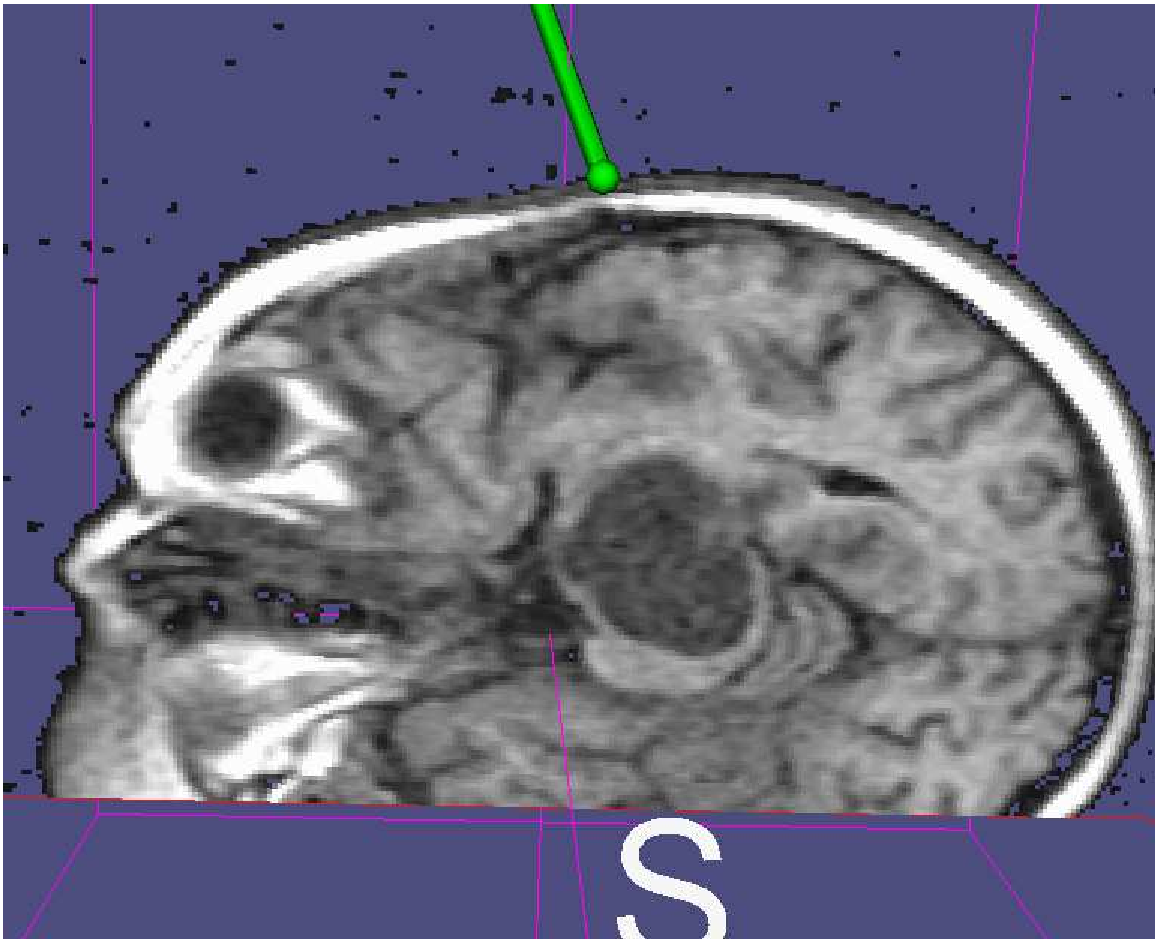

Figure 4-4: Locator (green) tries a trajectory for reaching a deep tumor.

histopathological results can indicate either that the tumor is benign, or that the needle missed the mark. The surgeon watches the 3D Slicer's display to help prevent the latter while simultaneously avoiding repeated needle repositioning. In this way, the importance of the 3D Slicer increases with the deeper tumors since the risks associated with tissue damage rise with repeated penetrations and maneuvering.

The first clinical case involved a six year-old girl with a tumor in the brainstem. It was from discussions with the surgeon following this particular case, that we added the capability to set an entry point on the skin, and a target, and render these points as well as the path between them with graphical cues. Reformatted images can then be automatically generated along this path as described in Chapter 3.

The second biopsy case is an example of the surgeon altering the approach following examination with the 3D Slicer. According to a conventional examination of MRI

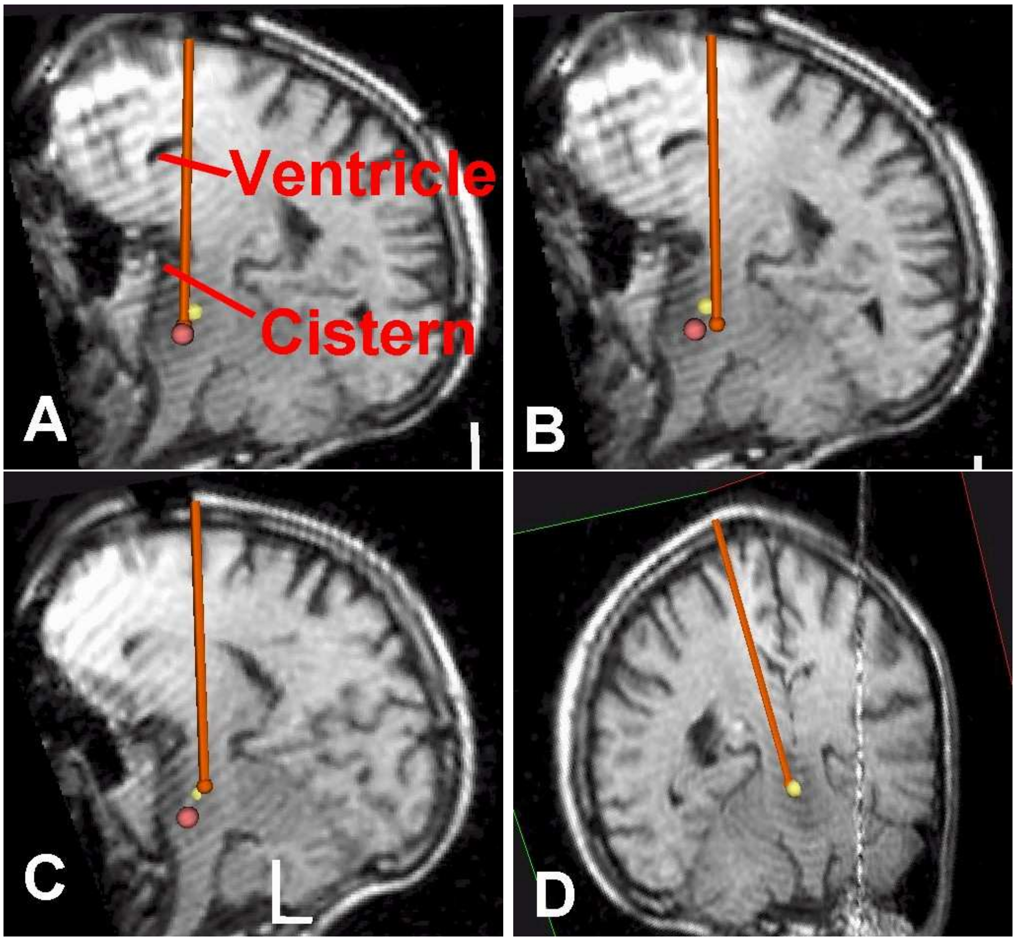

Figure 4-5: (A) Trajectory planning using the original hole, but hitting the cistern. (B) Avoiding the cistern and ventricle, but missing the tumor (marked with red and yellow spheres). (C) Trajectory planning using a new hole that allows reaching the tumor while missing both the cistern and ventricle. (D) Actual biopsy with two reformatted slices.

on a light box outside the OR, an access hole was created near the top of the head to reach a deep tumor in the brain stem. Then the locator was held at the site of the craniotomy by a Buckwalter clamp and oriented to try various trajectories. The locator — normally a handpiece with a needle of known length protruding from its center — was used *without* the needle in order to visualize *virtual* penetration in the 3D Slicer. The 3D Slicer was used to reformat images in the plane of the locator, and the biopsy targets were clearly marked using red and yellow spheres. The surgeon needed to reach the tumor while avoiding the ventricle (which if penetrated, would cause brain shift and bleeding) and the peripeduncular cistern with the Basilar artery (deadly if penetrated). The 3D Slicer revealed that the surgeon would pierce either the ventricle or cistern (Figure 4-5-A). Avoiding both hazards would require missing the tumor as well (Figure 4-5-B).

Therefore, a new access hole was drilled in the skull. Figure 4-5-C shows the same planning exercise being executed with the new hole. With the 3D Slicer's guidance, the tumor may now be reached while avoiding damage to both the ventricle and cistern. Figure 4-5-D shows the actual biopsy with a real needle in place. Here, two slices are being reformatted in the plane of the locator, yet perpendicular to each other so that both hazards may be seen simultaneously. Note that a single 2D view cannot reveal all hazards, and thus the 3D Slicer is an advantageous tool for trajectory planning compared to a conventional light box.

4.3.3 Cavernomas

One very important application of the 3D Slicer is in cases of cavernomas. A cavernoma is a very small vascular lesion that can bleed repeatedly. Repeated bleeding sessions grow layers like an onion — eventually leading to a larger mass that can cause seizures or other symptoms dependent on the lesion's location and the area of the brain that is pressured during bleeding. The key in this application is that the lesion is benign, small, and difficult to find. Prior to imaging, such lesions could be found only through following the blood after an episode of bleeding. Even with imaging, experienced surgeons have trouble pinpointing the lesion. When inserting the surgical tool, they pass it, pull back, take another stab, and repeat. Invasion of the brain always poses serious threats of side-effects such as paresis or bleeding. Such complications may be unavoidable in cases of malignant tumors that must be removed completely for survival. But cavernomas are benign, which reduces the tolerance of error. The 3D Slicer can significantly reduce risk of damage by guiding the surgeon more directly toward small lesions.

4.4 Tissue Identification

One challenge of neurosurgery is distinguishing the diseased tissue from its surroundings when there may be little or no differences visible to the human eye. To compound the difficulty, there are also no visible indications of which tissue is responsible for critical functionality such as motor control or speech. This section begins by hanging a backdrop of existing techniques, and then unveils how the 3D Slicer can provide assistance. Refer to [Osborn1994, Higgins, Hricak and Helms1997] for more details.

4.4.1 Vascularization

The experienced surgeons sometimes identify diseased tissue from the manner in which it bleeds. Tumors grow their own, more dense, yet more fragile, vasculature to meet the nutritional needs of their rapidly growing tissue.

4.4.2 Ultrasonic Aspirator

Another method employs a modern device called an ultrasonic aspirator. The aspirator employs ultrasonic energy waves to disintegrate tumor tissue while leaving healthy structures mostly intact. It takes advantage of the less-adhesive quality of tumor tissue in contrast with non-diseased tissue. The aspirator is used in conjunction with a sucking tube that extracts the debris. Experienced surgeons actually listen for the subtle differences in the sucking sound the tube makes depending on the type of tissue in its mouth.

4.4.3 Functional Areas

Functional areas responsible for speech can be protected by keeping the patient awake rather than anesthetized. The surgeon talks to the patient throughout resection while noting speech impediments as clues for when to back off. Similarly, surgeons pause to identify the motor cortex by asking the patient to move their hands and feet.

4.4.4 Imaging

The aforementioned methods are not fully sufficient, and identifying a boundary by crossing over it is certainly not a desired method. Imaging offers the opportunity for boundary validation without violation. The more dense cell lattice of tumors is depicted as hyperintense regions on T2-weighted MR images, and as hypointense areas on T1-weighted MRI. The fastest growing tumors generate a mixture of both bright and dark areas on MRI that appear almost anarchical in contrast to the structured surroundings. These tumors are called multiforme because they outgrow the very vasculature that provides their nutrition — leaving portions of the tumor dead — and creating their inhomogenous character.

4.4.5 Contrast Agents

When differences appear subtle even on MR images, an agent is injected into the bloodstream that produces more contrast on MR images in the more densely packed tumor. However, as the surgery progresses, the contrast diffuses more into the surrounding tissue, blurring the boundary.

Therefore, the 3D Slicer's capability to fuse earlier scans with present images augments the clear border definition even in the cases with washed out contrast. By "fuse", we refer to the capability to overlay foreground and background images on the same reformatted slice, and interactively adjust the opacity of the foreground image so that it fades in and out with respect to the background.

4.4.6 Using the 3D Slicer for Assistance

The 3D Slicer can aid the surgeon's decision making when used in conjunction with any of these conventional methods. The Surgeon sees the location of the locator's tip on the display, giving him a direct correlation of his visual impression of the tissue and the MR definition of the tissue. This is especially useful for areas that appear normal to the eye but display abnormal signal intensities in the applied sequences (T1 or T2).

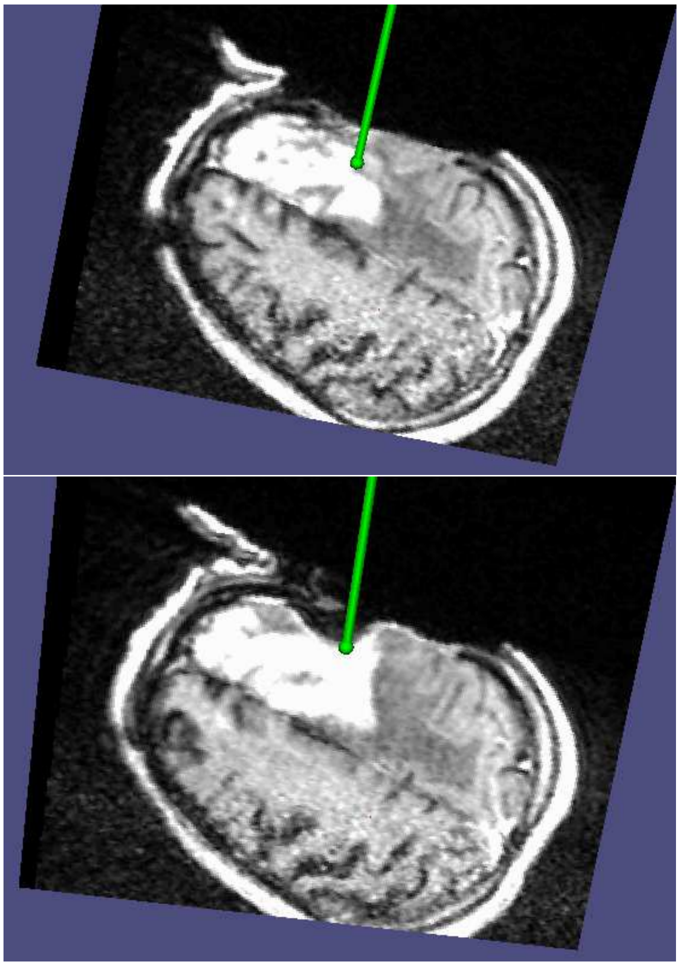

Figure 4-6: The surgeon leads a probe through the surgical cavity to observe the tumor boundary as amplified by an administered contrast agent. Great care must be taken in this case because the posterior border of the tumor reaches the motor area.

Figure 4-6 illustrates a case with a level 4 glioma (high grade brain tumor caused by erratic cell growth) that appears obviously bright under contrast-enhancement. The 3D Slicer is being used to create real-time, reformatted images through an intra-operatively acquired slab of images. The surgeon steers the locator around the perimeter of the cavity vacated through resection and notes the boundary between the non-enhanced normal tissue and the enhanced diseased tissue after resection of the posterior portion of the lesion.

The 3D Slicer invaluably blends data sets together to merge functional, anatomical, and information from contrast agents into a single view. Additionally, we have experimented with affixing the locator's handle to the ultrasonic aspirator so the sur-

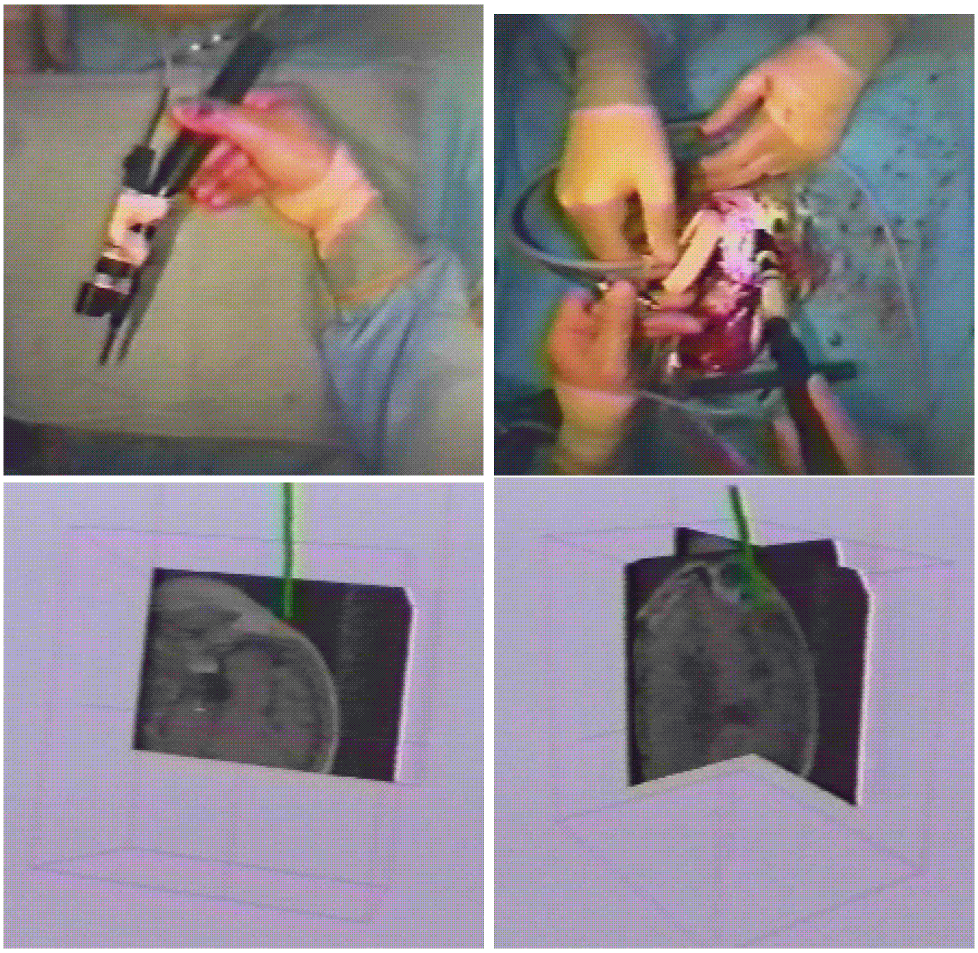

Figure 4-7: TOP-LEFT: The locator's handle attached to the ultrasonic aspirator. TOP-RIGHT: Applying the novel device. BOTTOM: Following the new tool's position on the higher definition, pre-operative, reformatted images with one (left) or two (right) reformatted slices. (Note: the poor quality is a result of extracting these images from video, but the surgeon's view is clear.)

geon can directly correlate the aspirator's position to the information present in the dynamically reformatted images. Figure 4-7 demonstrates this technique.

4.5 Intra-operative Imaging Updates

4.5.1 Progressive Updates

Cerebrospinal fluid (CSF) leakage, as well as tumor removal itself, causes intra-operative brain deformation. Subsequent acquisitions at various times during the surgery allow adjustments to intra-operative changes. According to the magnitude of the intra-operative changes these acquisitions are repeated (on average 3-5 times during the case) as demanded by the surgeon. We typically acquire a slab of 10-14 slices with 5 mm thickness in 2-2.5 minutes.

The 3D Slicer's capability of displaying intra-operative T1 and T2 imaging, as well as preoperative multi-modal information, allows the comparative evaluation and the subsequent modification of the surgical tactic. The surgeon can observe the changes in position of anatomy on the intra-operative images, while also benefiting from the higher resolution, higher contrast, functional information, and administered contrast agents of the pre-operative images.

Figure 4-8 depicts a SPGR scan with contrast enhancement scanned after craniotomy, but prior to resection, of a fronto-temporal glioma. The bottom image is an overlay of a slab of lower-resolution (acquired fast), T1-weighted images scanned during resection (note the hole in the tumor) and overlaid semi-transparently on the SPGR (top image). The overlay demonstrates the ability to relate the intra-operative changes to the higher-resolution definition of tumor location.

Figure 4-9 displays a reformatted axial slice of the same patient in 4-8. The bottom image depicts the top image overlaid on a reformatted slice from the same location prior to resection. The collapse of brain tissue during the resection is quite pronounced by the shadow. (Observant readers who note the movement of the skin flap may believe the patient has moved, but in fact, the skin flap loses rigidity and falls during the procedure.)

4.5.2 Augmenting Intra-operative Images with Pre-operative Images

The need to augment quickly-acquired, low-resolution, low-contrast, intra-operative slabs of slices with high-quality images scanned prior to resection is made apparent in Figure 4-10. Reformatted images that are relatively parallel to the other images in the slab look fine, but the image shown is oriented perpendicular to those in the

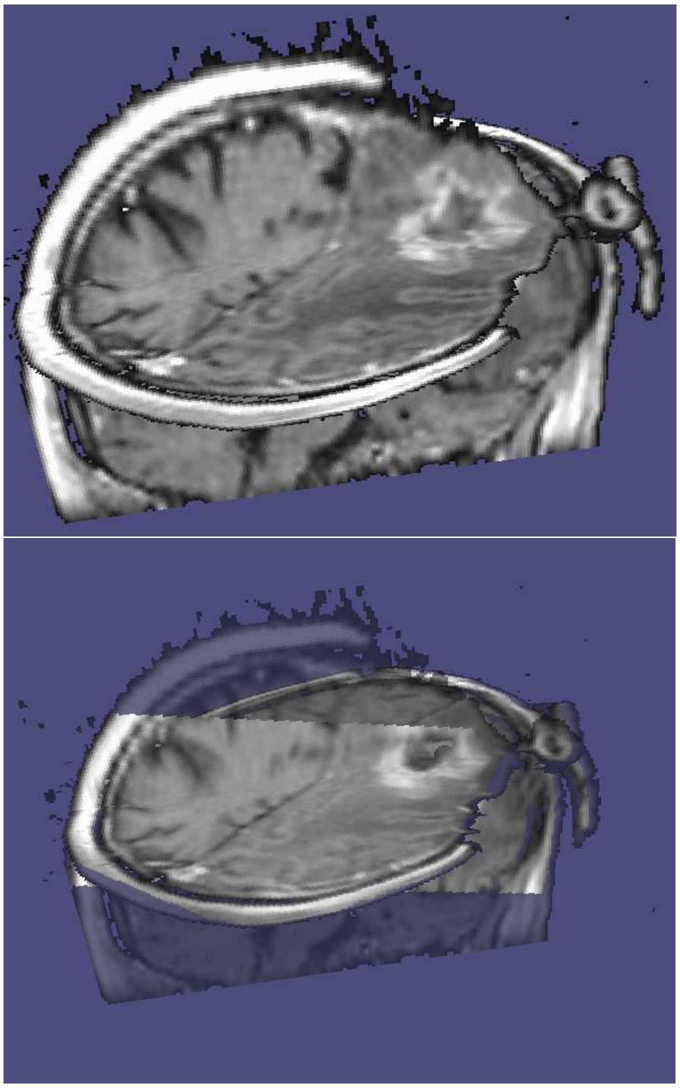

Figure 4-8: Top: Pre-resection data. Bottom: Intra-resection data overlaid on pre-resection data.

slab. Even with interpolation, the stair-step jaggedness is a testament to the large spacing between slices (7 mm rather than the customary 1.5 mm).

4.5.3 Intra-operative Deformable Registration

Since model deformation has yet to achieve a sufficiently dependable accuracy for surgeons, we often stop using pre-operative models once resection is well underway. Deformable registration is being developed [Hata et al.1998] that holds promise in being able to warp pre-operative data to match the intra-operative changes. Existing

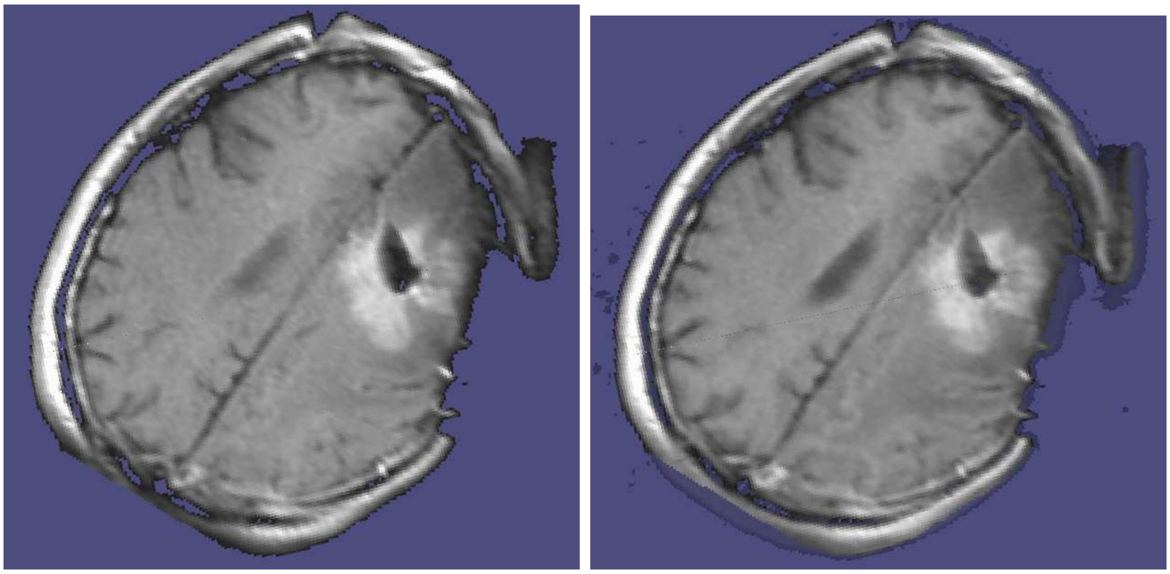

Figure 4-9: Top: Intra-resection image. Bottom: Intra-resection data overlaid on pre-resection data to illustrate brain shift.

algorithms are not yet accurate enough to be relied upon by surgeons, but when ready, such tools can be easily inserted into the 3D Slicer.

4.6 Post-resection Validation

Figure 4-11 illustrates the surgeon exploring the cavity remaining from resection. The attending neuroradiologist examines the reformatted images for evidence of residual, diseased tissure to be removed. The surgeon frequently pauses for similar examinations during the course of resection.

4.7 Summary of Clinical Procedure

This section summarizes the clinical procedure for using the 3D Slicer for neurosurgical guidance:

1. Collect diagnositic, anatomical scans (MR, CT, angiography).

2. For cases near the eloquent cortex, conduct functional exams (fMRI).

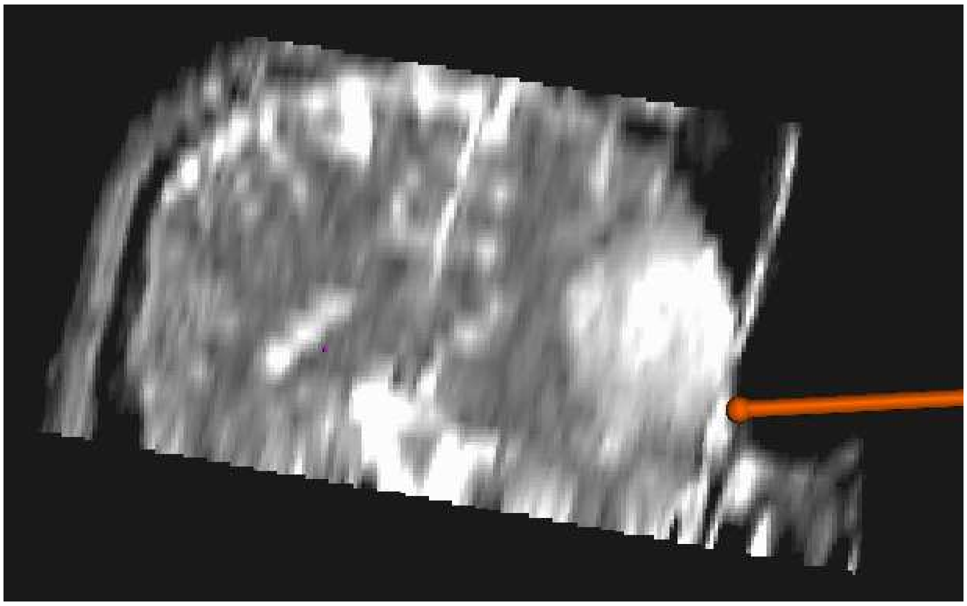

Figure 4-10: Reformatted images have poor resolution when oriented perpendicuar to thick slices.

3. Use the 3D Slicer to segment the scans and generate 3D surface models of key structures – both anatomical and functional.

4. Automatically register all pre-operative data together using the 3D Slicer.

5. Position patient in the interventional MR, collect a volume scan, and setup the 3D Slicer.

6. Automatically register all pre-operative data to the intra-operative data using the 3D Slicer.

7. Use the combination of all data to plan trajectory and guide the initial resection by aiming the locator while seeing with X-ray vision by viewing the 3D Slicer's display on the in-bore [2] monitor.

8. Periodically perform intra-operative imaging as surgery progresses and visualize the slabs with reformatted slices in the 3D Slicer. Drive the location of the slices using the tracked locator device.

[2]Where the surgeon stands between the double-donut shaped magnets

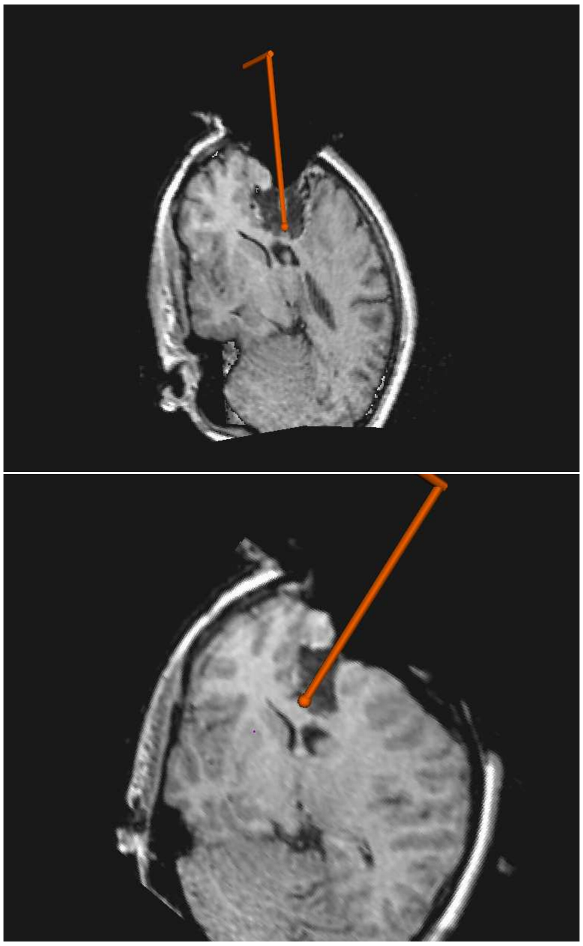

Figure 4-11: The surgeon leads a probe through the surgical cavity to check for tumor remnants on the reformatted images. ABOVE: two reformatted planes. BELOW: one plane.

4.8 Proven Validity

We conclude this chapter with an assessment of the 3D Slicer's validity and accuracy. Its validity for planning and guiding resection is highlighted by the histopathology and location of the clinical cases for which it was used (See Table 4.1). Ten tumors were low-grade gliomas, for which total resection is correlated with a prolonged survival time and presumably even with total cure. Of these, six were located in the patients' dominant left hemisphere, three in the immediate vicinity of the motor cortex, and two close to speech areas. In these cases, precise margin definition is of utmost importance. Visually differentiating tissue in these tumors can be difficult to impossible, whereas T2-weighted images give a good estimate of the tumor extent. The display of these

Table 4.1: Neurosurgical Guidance Cases

Case	Side	Lobe	Histopathology
1	Left	Parietal	Low grade glioma
2	Left	Parietal	Low grade glioma
3	Left	Basal ganglia	Low grade glioma
4	Left	Parietal	Low grade glioma
5	Right	Frontal	Low grade glioma
6	Left	Frontal	Low grade glioma
7	Right	Frontal	Low grade glioma
8	Right	Parietal	High grade glioma
9	Right	Parietal-occipital	High grade glioma
10	Right	Parietal	High grade glioma
11	Left	Parietal	High grade glioma
12	Right	Frontal	High grade glioma
13	Right	Frontal	High grade glioma
14	Left	Parietal	High grade glioma
15	Left	Frontal	Cavernoma
16	Left	Frontal	Metastasis
17	Right	Basal ganglia	Low grade glioma
18	Right	Frontal	High grade glioma
19	Right	Frontal	Low grade glioma
20	Left	Parietal	Low grade glioma
21	Right	Brain stem	Biopsy
22	Right	Brain stem	Biopsy

images allows a more careful control of total tumor resection. Of the remaining seven cases with high-grade (astrocytoma III and glioblastoma multiforme) tumors, four were adjacent to the motor cortex, thus making updated navigation a valuable tool to define the surgical goal and prevent morbidity.

4.9 Accuracy Testing

We have conducted experiments to verify that reformatted images generated by the 3D Slicer relative to the locator appear in the same location as the images scanned by the MR scanner at the coordinates of the locator. Figure 4-12 demonstrates testing relative to the superior, posterior, and anterior borders of a tumor.

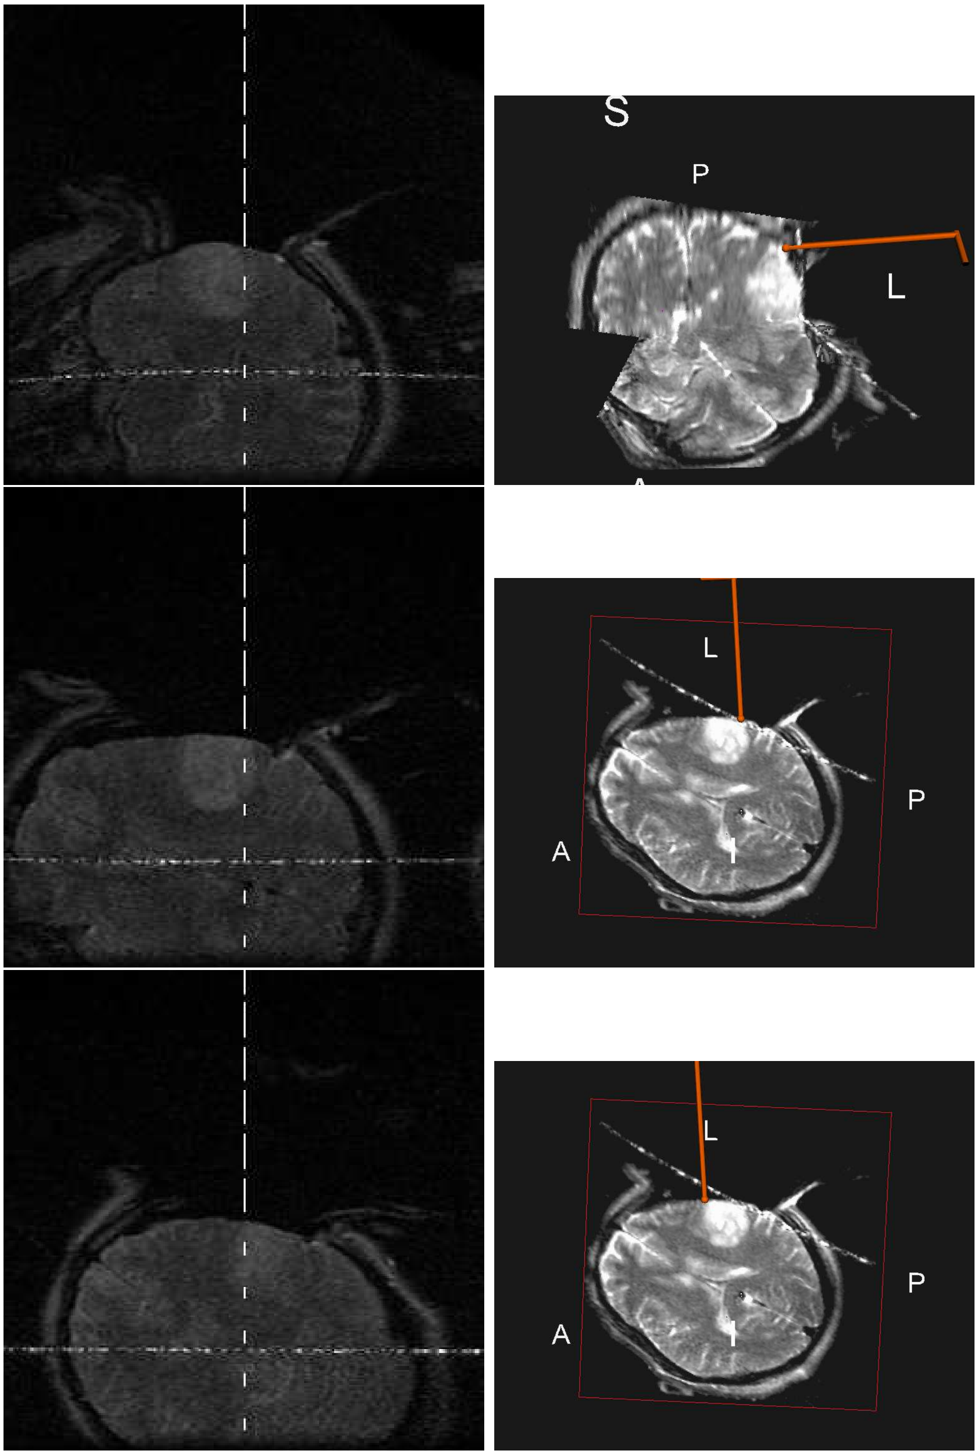

Figure 4-12: The images on the left are generated by the MR scanner, and the images at on the right are generated by the 3D Slicer relative to the locator.

Chapter 5

3D Analysis

The 3D Slicer is used outside of surgery for performing analysis that requires a 3D environment. In this chapter, we describe how our group is using the combination of the 3D Slicer and the open MR to analyze intra-operative brain deformations, and we examine quantitative studies that require making measurements on source images and 3D models simultaneously.

5.1 Studying Brain Shift

The precision offered by typical neuronavigation systems diminishes as the brain deforms during the course of surgery. Current attempts to measure and correct for brain shift involve using navigation systems alone, intra-operative MR or CT, intra-operative ultrasound, or intra-operative ultrasound combined with a navigation system. Our system of intra-operative MR combined with a navigation system offers several advantages to the prior approaches. Our disadvantage is currently cost, but the discoveries made with our system will hopefully enable lower-cost systems to adequately handle deformations. At present, interventional imaging works well to update navigation, before we have the capability to model the dynamics.

We study the behavior of the brain during surgery to further the construction of models that simulate the brain's biomechanical properties. These are determined by the brain's tissue properties, the internal fiber system, the external supporting structures, the arterial and venous mesh, and the pathology [Nabavi et al.].

5.1.1 Other Methods for Detecting Brain Shift

This section introduces the progress made in the field to date, and outlines the disadvantages of other systems that are overcome by the combination of the 3D Slicer with the open MR unit.

Navigation Systems Alone

Modern navigation systems have been applied unaided by intra-operative imaging to measure deformations. [Hill et al.1997] studied the influence of exposing the brain to gravity and the subsequent sinking of the surface during an open craniotomy. [Roberts et al.1998] extended this study to collect surface measurements during actual surgery. [Dorwand et al.1998] attempted estimate depth movement beyond the exposed surface, but clearly, accounting for all subsurface movement invasively is impossible. The surgery itself impairs repeated measurements, and there is rare opportunity for the definition of accurate subsurface landmarks. All these approaches are limited by the small area exposed by the craniotomy. The goal of minimum invasiveness, which is one purpose of the navigational systems, demands that only a small region is opened and is thus accessible to measurements. This does not yield enough information to deduce deformation and predict the course of brain deformation with precision and reducibility, especially for subsurface structures. Figure 5-1 demonstrates how surface measurements can be made using a navigation system, which in this case is MIT's X-plan system [Leventon1997, Grimson et al.1998, Black and Chaberie1998]. As an optically-tracked point device swabbed the tumor's exposed surface, the positions were marked and rendered with spheres on the 3D graphics display.

Intra-operative Ultrasound

MR is the superior imaging modality for delivering soft-tissue contrast as needed for neurosurgery. Ultrasound can detect surface boundaries, so there are synergies

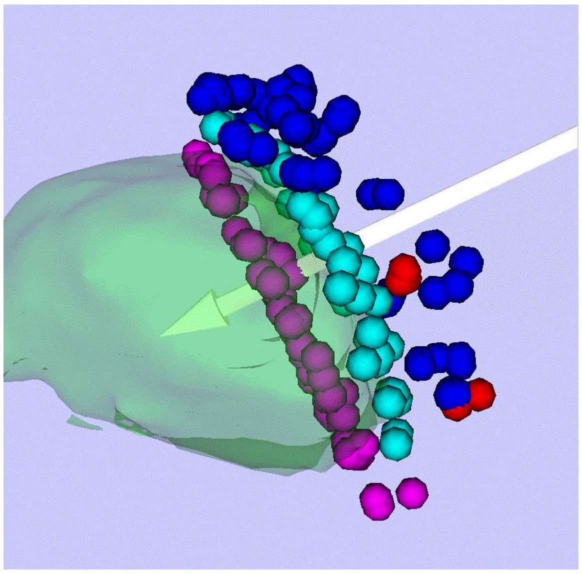

Figure 5-1: Color spheres mark surface locations touched by a tracked pointer before dural opening (green), after dural opening (blue — note the bulging brain), and after tumor removal (pink - note the sinking of the surface).

in teaming pre-operative MR with intra-operative ultrasound. Ultrasound has been shown to delineate a lesion when operated at 6MHz which limits the insonation depth to reaching the ipsilateral ventricle [Jodicke et al.1998]. [Bucholtz et al.1997] first combined intra-operative ultrasound with a neuronavigation system. Pre-operative CT/MR were registered with the intra-operative ultrasound, and then deformed to correct for intra-operative changes. Their registration relied on stable midline structures — an assumption challenged by findings made using our system! Their landmarks may be obscured by pronounced midline shift or changes in the appearance of midline structures. While we agree with Bucholz' descriptions of various shift probabilities, we found there can be remarkable exceptions which will lead to a misinterpretation of the ultrasound updates [Nabavi et al.].

Other Intra-operative MR or CT Systems

The Magnetom Open Scanner (Siemens AG, Erlangen, Germany) is situated outside the OR, requiring patient transport to and from the scanner. Bone screws facilitate neuronavigation, but require an hour of setup time. The Phillips ACS-NT

(Phillips Medical Systems, Best, The Netherlands) is a 1.5 Tesla system within the OR, but the patient must be moved to the scan location and back. With either system, repeated scanning is cumbersome, time-consuming, and endangers sterility. Consequently, scanning is primarily pre-operative and post-resection, and misses the dynamic components in-between.

5.1.2 Results with Our System

The specific results of our studies are beyond the scope of this thesis and are presented elsewhere [Nabavi et al.]. Nonetheless, we present an overview of our findings here to establish the utility of our system for this application.

The advantage of the Signa SP is that the patient is not repositioned between operating and imaging. We performed repeated scanning to provide enough temporal resolution to uncover all phases of deformation. Scanning time, and therefore image quality, is limited by surgical demands, and we found a sufficient compromise in routinely acquiring a volume comprised of 60 slices with 2.5 mm thickness in 4 minutes. The images were used for navigation and later post-processed in the 3D Slicer for measuring brain shift.

Figure 5-2 demonstrates that the craniotomy causes swelling. Consequently, the brain bulges out after the opening of the dura mater — the brain's tough, protective membrane. The surrounding area becomes edematous (fluid accumulates in the tissues), and the ventricle becomes elongated towards the opening. It is important to note that the swelling and subsequent distortion of the ventricle is not homogeneous.

Figure 5-3 surprisingly reveals that the midline can not be relied upon as a rigid landmark. Following resection, the cortical displacement is 1.5 cm, but there is also a significant midline shift of almost 1 cm. The left ventricle is compressed, and even the brainstem itself has shifted.

Figure 5-4 shows another case with views of the same location after dural opening, after tumor resection, and after dural closure. The stages are characterized by different extents of intra-operative displacement and the different development of the various components. The initial phase is a sinking of the tissue in the direction of

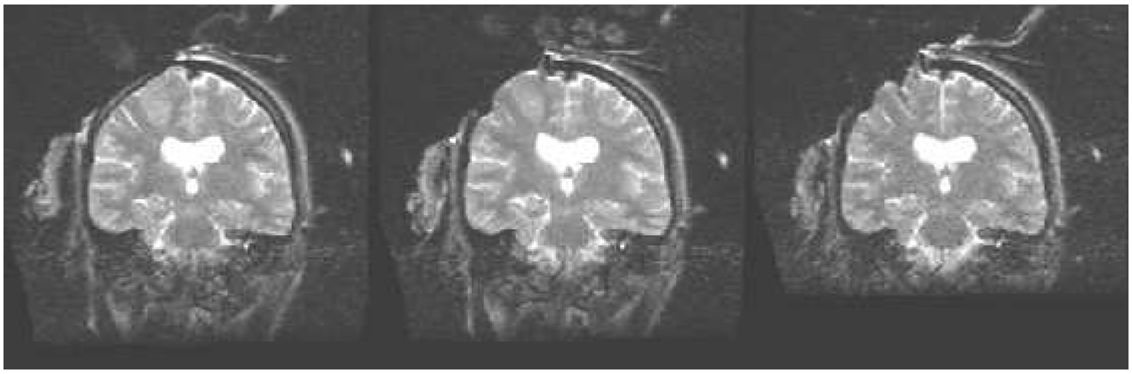

Figure 5-2: The brain swells out after dural opening in a non-homogeneous manner, and the ventricle becomes elongated.

gravity, which is barely apparent on the first image, but the cortical surface falls into the resection cavity on the second. Additionally, the remote surface sinks in by almost a centimeter. The midline structures change completely as the left ventricle is almost completely absent, and the right ventricle is slightly compressed.

Following dural closure, in the last image, the midline structures have regained their basic shape, the resection cavity is shrunken, and the surrounding regions have increased their volume. This case provides evidence that suggests other studies conducted by wheeling the patient into a second room for imaging after dural closure visualize a less pronounced shift than actually occurred during the operation. Without the data acquisition during surgery, the shift in this case would have been greatly underestimated. This even holds true for structures further away from the surgical site, as for instance the temporal lobe. Displacement here is as pronounced, and in the same direction, as in the previous description in the other areas.

We also experimented with visualizing the displacement of the temporal lobe by segmenting and building 3D surface models from the intra-operatively acquired volumes, as shown in Figure 5-5.

The conclusions we can draw from these experiments are that the intra-operative distortions are not linear or homogenous, can affect the midline considerably, and will be difficult to predict [Nabavi et al.].

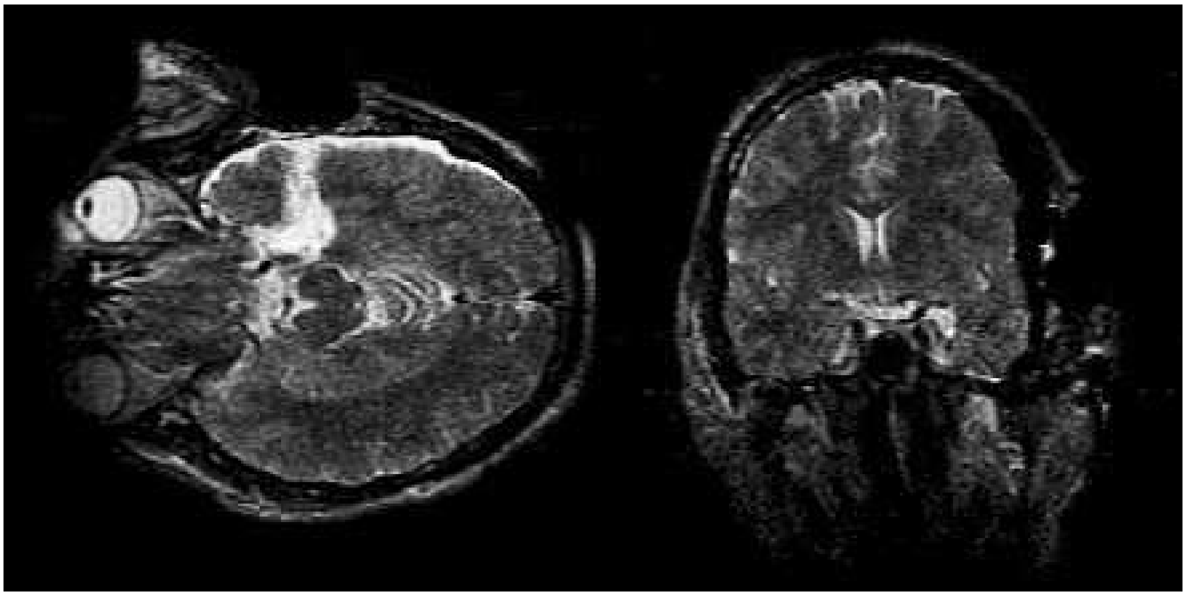

Figure 5-3: Contrary to common belief, our results show that the midline can shift, even up to a centimeter.

5.2 Volumetric Measurements

The 3D Slicer's tools for creating label maps, generating surface models, and performing operations on models allow volumetric analysis of 3D structures. Although cadavers are used to assess anatomy, these results are not patient-specific and cannot be applied to quantification of muscle mass. There are a vast number of these analytic applications, and we mention just a few here to demonstrate the 3D Slicer's utility in this domain.

Orthopedic Experiments

One example of a patient-specific experiment is to maneuver bone models prior to orthopedic surgery to determine which course of action delivers the optimum post-operative range of motion [Richolt et al.1998]. At the time of this writing, another researcher is adding collision detection of 3D surface models to the 3D Slicer [Everett].

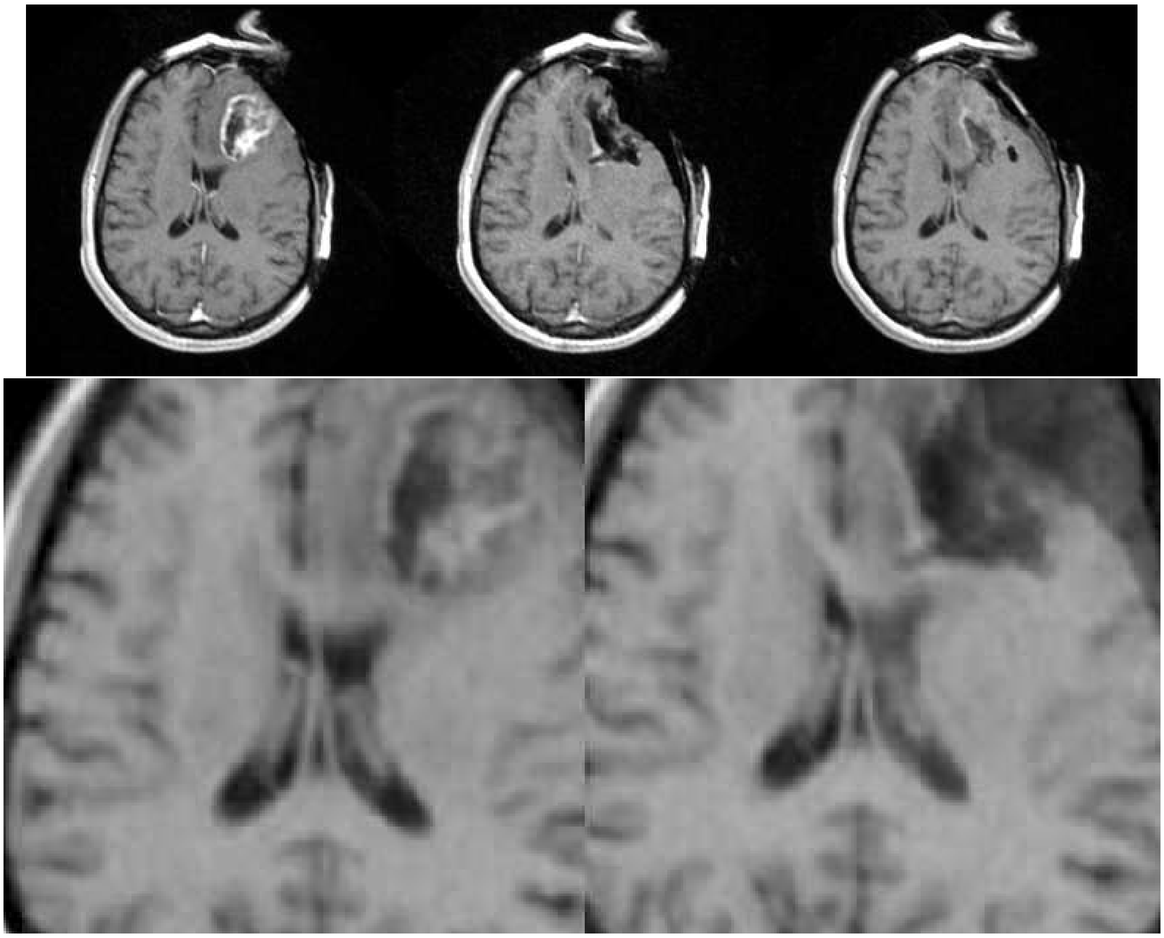

Figure 5-4: Brain structures regain their shape following dural closure, which proves that post-resection imaging studies are inaccurate. TOP: From left to right, after dural opening, after resection, and after dural closure. BOTTOM: Close-up of the top left image blended with the top right

Quantification of Muscle Mass

Ten million American women are affected by stress urinary incontinence and pelvic organ prolapse at an annual cost of nearly 10 billion dollars [Hu1990]. Our limited knowledge of the anatomy of the pelvic floor abets the cause of incontinence to remain elusive.

We used the 3D Slicer to quantify the normal appearance of the female pelvic floor to better understand the structures responsible for maintaining support. Clinical studies of this nature could routinely guide treatment in the future. For example, it seems reasonable that women with intact pelvic floor support structures would

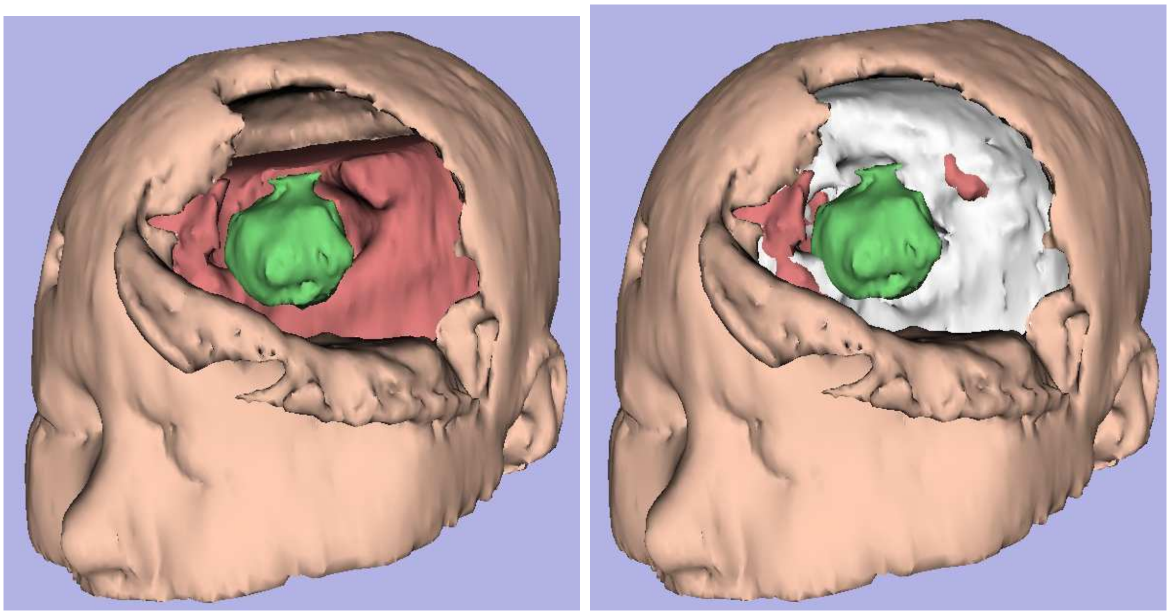

Figure 5-5: The brain after dural closure (white) covers the brain model made from data scanned before closure (pink) showing the displacement more intuitively than the 2D images convey.

respond well to behavior modification techniques and estrogen replacement therapy while those with disruption of the levator ani (support muscle) would benefit from surgery.

We recently conducted a study in which ten volunteers underwent MR scanning. We then used the 3D Slicer to measure distances, angles, and volumes from the 3D structures. The detailed results are reported in [Fielding et al.1999]. Figures 5-6 and 5-7 display 3D models of the normal female pelvis with superimposed gray-scale images.

Measuring Response to Therapy

The 3D Slicer's combined capabilities for segmentation, visualization, and volumetric measurements comprise a tool for measuring a patient's response to therapy. For example, the present clinical practice for following progress of cervical cancer treatments is to measure the diameter of the lesion on a 2D image — clearly an inaccurate metric. An upcoming study at Brigham and Women's Hospital will investigate the advantages of using the 3D Slicer to measure the true volume in under ten minutes of a clinician's time.

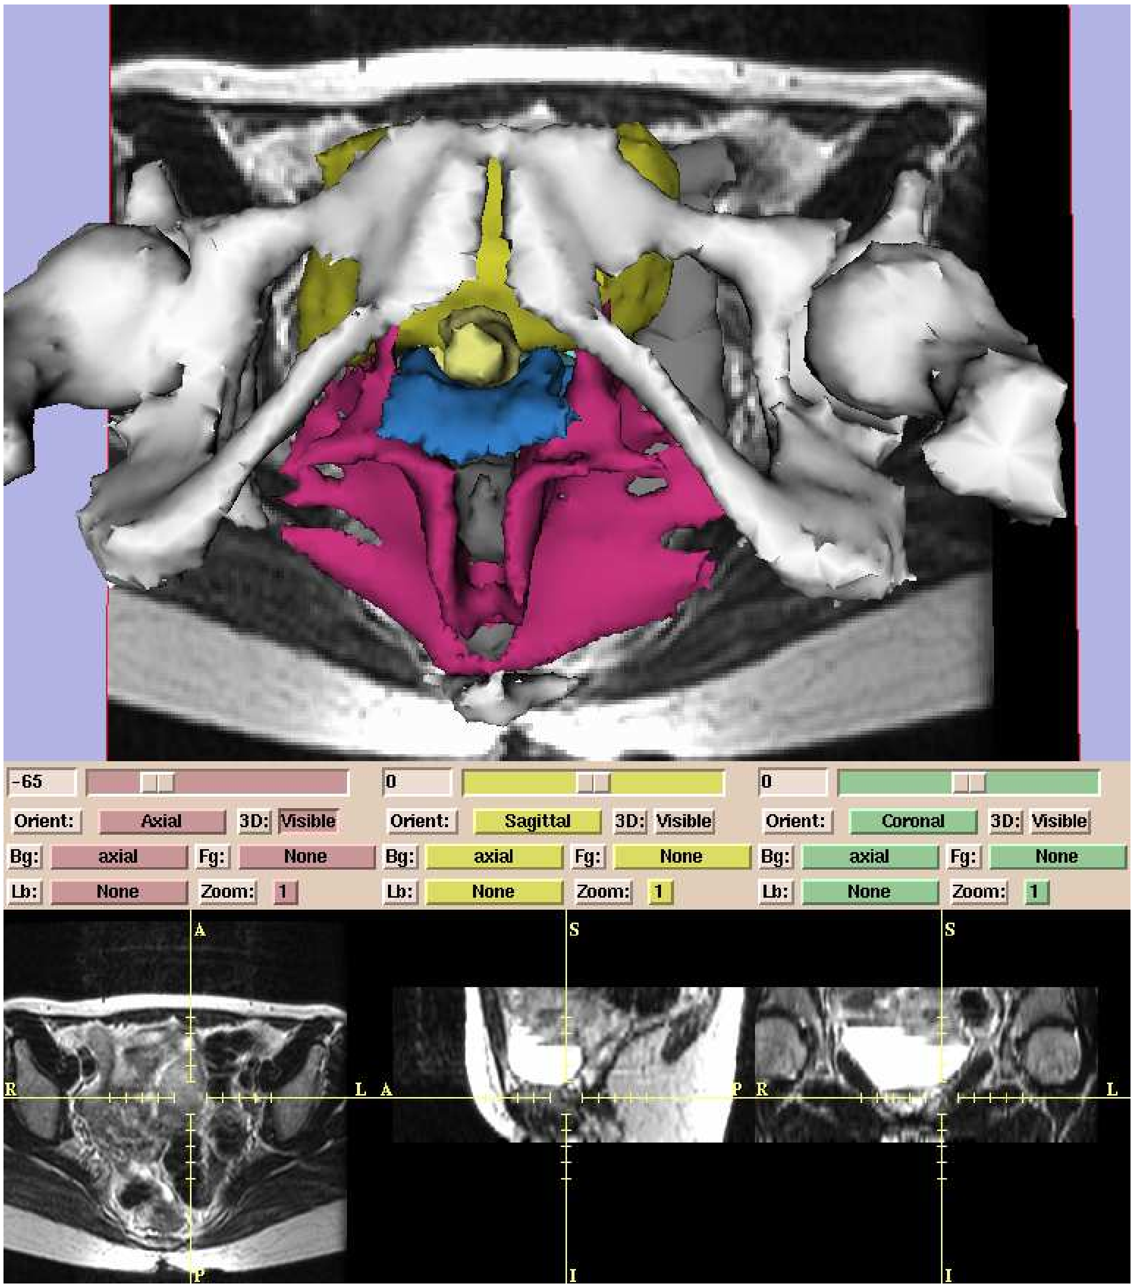

Figure 5-6: Bones (white), bladder/urethra (yellow), vagina (blue), uterus (green), rectum (gray), and levator ani (pink) can be visualized and quantified.

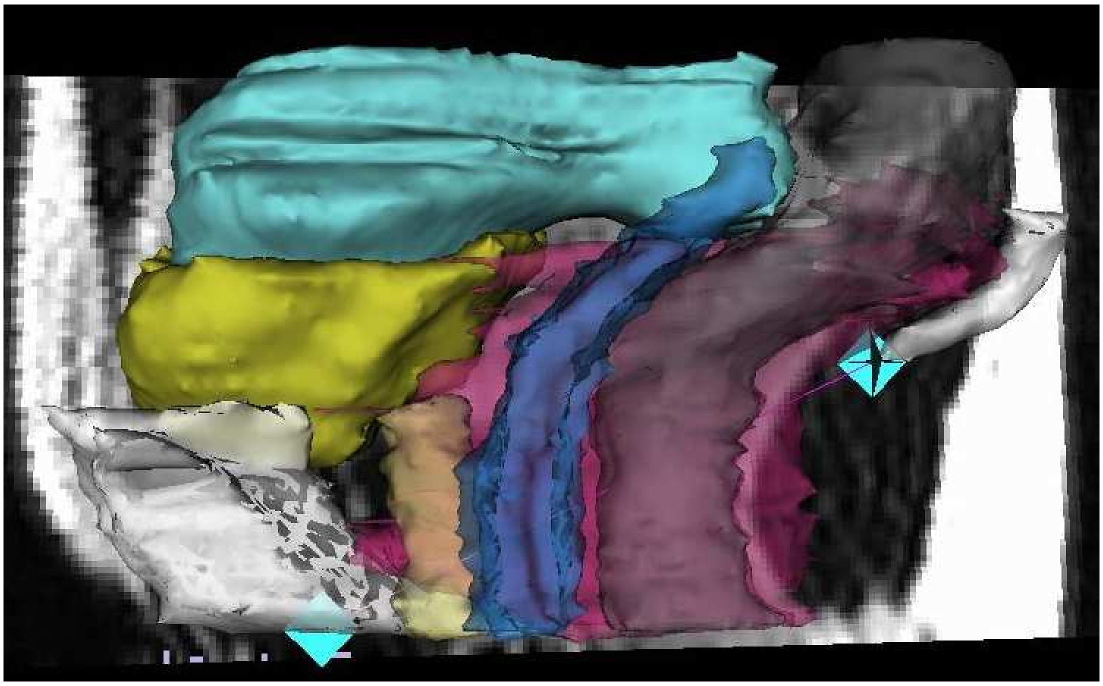

Figure 5-7: The purple line between the two blue markers is measuring the distance of the pubococcygeal line (level of the pelvic floor, and minimum width of the birth canal).

Chapter 6

Conclusion

This thesis describes a uniquely valuable system for registration, segmentation, visualization, volumetric analysis, instrument-tracking, surgical planning, surgical guidance, and seamless integration of pre-operative data with intra-operative imagery.

The 3D Slicer has been used as a navigational tool in more than 20 cases, and the feasibility phase proved it to be a stable and reliable application. As we apply the 3D Slicer on a routine basis in 1-2 cases per week, a clinical evaluation of its influence on surgical decision making and resection control is being conducted. The system has been used in the most difficult cases where precise margin definition is of utmost importance, and where updated navigation is a valuable tool to define the surgical goal and prevent morbidity.

The system's extendible design, and free availability to researchers, will allow it to expand its value as a highly integrated suite for analysis and visualization.

Bibliography

[Alexander et al.1993] Alexander, E., Loeffler, J., Schwartz, R., Johnson, K., Carvalho, P., Garada, B., Zimmerman, R., and Holman, B.1993. Thallium-201 technetium-99m HMPAO single photon emission computed tomography (SPECT) imaging for guiding stereotactic craniotomies in heavily irradiated glioma patients. *Acta Neruochirugica*, 122:215–217.

[ANALYZE] *ANALYZE Software.* Mayo Clinic. http://www.mayo.edu/bir/analyze/ANALYZE_Main.html.

[AnatomyBrowser] AnatomyBrowser. http://www.ai.mit.edu/projects/anatomy_browser.

[Auer1990] Auer1990. *Ultrasound in Neurosurgery.* Springer.

[Black and Chaberie1998] Black, P. and Chaberie, A.1998. Clinical experience with XPlan. In [Wells, Colchester and Delp1998].

[Black et al.1997] Black, P., Moriarty, T., Alexander III, E., Stieg, P., Woodard, E., Gleason, P., Martin, C., Kikinis, R., Schwartz, R., and Jolesz, F.1997. The development and implementation of intraoperative MRI and its neurosurgical applications. *Neurosurgery*, 41(4):831–842.

[Brummer et al.1993] Brummer, M., Mersereau, R., Eisner, R., and Lewine, R.1993. Automatic detection of brain contours in MRI data sets. *IEEE Transactions on Medical Imaging*, 12(2):153–166.

[Bucholtz et al.1997] Bucholtz, R., Yah, D., Trobaugh, J., McDurmont, L., Sturm, C., Baumann, C., Henderson, J., Levy, A., and Kessman, P.1997. The correction of stereotactic inaccuracy caused by brain shift using an interaoperative ultrasound device. In *Proceedings First Joint CVRMED/MRCAS*, Grenoble France.

[Colchester et al.1996] Colchester, A., Zhao, J., Holton-Tainter, K., Henri, C., Maitland, N., Roberts, P., Harris, C., and Evans, R.1996. Development and preliminary evaluation of VISLAN, a surgical planning and guidance system using intra-operative video imaging. *Medical Image Analysis*, 191:73–90.

[Dorwand et al.1998] Dorwand, N., Alberti, O., Velani, B., Gerritsen, F., Harkness, W., Kitchen, N., and Thomas, D.1998. Postimaging brain distortion: Magnitude, correlates, and impact on neuronavigation. *Neurosurgery*, 88:656–662.

[ETH] *Image Science Group*. ETH Zürich. http://www.vision.ee.ethz.ch.

[Everett] Everett, P. Collision detection for orthopedic simulations using 3D surface models. Personal Communication.

[Fielding et al.1999] Fielding, J., Dumanli, H., Schreyer, A., Okuda, S., Gering, D., Zou, K., Kikinis, R., and Jolesz, F.Press, 1999. MR based three-dimensional modeling of the normal female pelvic floor with quantification of muscle mass. *American Journal of Roentgenology*.

[Galloway1993] Galloway, R.1993. Stereotactic frame systems and intraoperative localization devices. In Maciunas, R., editor, *Interactive Image-Guided Neurosurgery*. American. Association Neurological Surgeons.

[GE] *Three Dimensional Medical Reconstruction*. GE Corporate Research & Development. http://www.crd.ge.com/esl/cgsp/projects/medical.

[Golland et al.1998] Golland, P., Kikinis, R., Umas, C., Halle, M., Shenton, M., and Richolt, J.1998. Anatomybrowser: A framework for integration of medical information. In [Wells, Colchester and Delp1998].

[Grimson et al.1996a] Grimson, W., Ettinger, G., Kapur, T., Leventon, M., Wells III, W., and Kikinis, R.1996a. Utilizing segmented MRI data in image-guided surgery. *IJPRAI*.

[Grimson et al.1996b] Grimson, W., Ettinger, G., White, S., Lozano-Perez, T., III, W. W., and Kikinis, R.1996b. An automatic registration method for frameless stereotaxy, image guided surgery, and enhanced reality visualization. *IEEE Transactions on Medical Imaging*, 15(2).

[Grimson et al.1998] Grimson, W., Leventon, M., Ettinger, G., Chabrerie, A., Ozlen, F., Nakajima, S., Atsumi, H., Kikinis, R., and Black, P.1998. Clinical experience with a high precision image-guided neurousurgery system. In [Wells, Colchester and Delp1998].

[Gronemeyer et al.1995] Gronemeyer, D., Seibel, R., Melzer, A., Schmidt, A., Deli, M., Friebe, M., and Busch, M.1995. Future of advanced guidance techniques by interventional CT and MRI. *Minimally Invasive Therapy*, 4:251–259.

[Guys] *Computational Image Science Group.* Guy's and St. Thomas' Hospital. http://www-ipg.umds.ac.uk/cisg/indexn.htm.

[Hata et al.1997] Hata, N., Dohi, T., Kikinis, R., Jolesz, F., and Wells III, W.1997. Computer assisted intra-operative MR-guided therapy: Pre and intra-operative image registraion, enhanced three-dimensional display, deformable registration. In *7th Annual meeting of Japan Society of Computer Aided Surgery*, pages 119–120.

[Hata et al.1998] Hata, N., Dohi, T., Warfield, S., III, W. W., Kikinis, R., and Jolesz, F.1998. Multimodality deformable registration of pre- and intraoperative images for MRI-guided brain surgery. In [Wells, Colchester and Delp1998], pages 1067–1074.

[Hearn and Baker1997] Hearn, D. and Baker, M.1997. *Computer Graphics.* Prentice Hall.

[Higgins, Hricak and Helms1997] Higgins, H., Hricak, H., and Helms, C.1997. *Magnetic Resonance Imaging of the Body.* Lippincott-Raven, Philidelphia.

[Hill et al.1997] Hill, D., Maurer Jr., C., Wang, M., Maciunas, R., Barwise, J., and Fitzpatrick, J.1997. Estimation of intra-operative brain surface movement. In *Proceedings First Joint CVRMED/MRCAS*, Genoble France.

[Höhne1992] Höhne, K.1992. A 3D anatomical atlas based on a volume model. *IEEE Computer Graphics Applications*, 2:72–78.

[Hu1990] Hu, T.1990. Impact of urinary incontinence on healthcare costs. *J. Am. Geriatr. Soc.*, 38:292–295.

[INRIA] *Epidaure Research Project.* INRIA. http://www-sop.inria.fr/epidaure/Epidaure-eng.html.

[J. West1997] J. West, J. F.1997. Comparison and evaluation of retrospective intermodality brain image registration techniques. *Journal of Computer Assisted Tomography*, 21:554–566.

[JHU] *Computer-Integrated Surgical Systems and Technology.* Johns Hopkins University. http://cisstweb.cs.jhu.edu.

[Jodicke et al.1998] Jodicke, A., Deinsberger, W., Erbe, H., Kriete, A., and Boker, D.1998. Intraoperative three-dimensional ultrasonography: An approach to register brain shift using multidimensional image processing. *Minimally Invasive Neurosurgery*, 41:13–19.

[Jolesz1997] Jolesz, F.1997. Image-guided procedures and the operating room of the future. *Radiology*, 204:601–612.

[Kapur1999] Kapur, T.1999. *Model Based Three Dimensional Medical Image Segmentation.* PhD thesis, Massachusetts Institute of Technology.

[Leventon1997] Leventon, M.1997. A registration, tracking, and visualization system for image- guided surgery. Master's thesis, Massachusetts Institute of Technology.

[Lorensen and Cline1987] Lorensen, W. and Cline, H.1987. Marching cube: A high resolution 3-D surface construction algorithm. *Computer Graphics*, 21(3):163–169.

[Lueven] *Medical Image Computing.* Leuven. http://www.esat.kuleuven.ac.be/mi2/mic.html.

[Lunsford, Parrish and Albright1984] Lunsford, L., Parrish, R., and Albright, L.1984. Intraoperative imaging with a therapeutic computed tomographic scanner. *Neurosurgery*, 15:559–561.

[Maciumas et al.1993] Maciumas, R., Fitzpatrick, J., Galloway, R., and Allen, G.1993. Beyond stereotaxy: Extreme levels of application accuracy are provided by implantable fiducial markers for interactive image-guided neurosurgery. In Maciunas, R., editor, *Interactive Image-Guided Neurosurgery.* American. Association Neurological Surgeons.

[MAYO] *Biomedical Image Resource.* Mayo Clinic. http://www.mayo.edu/bir/.

[MEDx] *MEDx Software.* Sensor. http://www.sensor.com/medx_info/medx_docs.html.

[MERL] *Medical Applications Research.* Mitsubishi Electric Research Laboratory. http://www.crd.ge.com/esl/cgsp/projects/medical.

[MGH] *http://research.neurosurgery.mgh.harvard.edu.* Massachusetts General Hospital and Harvard Medical School.

[MNI] *MNI Software.* MNI. http://www.bic.mni.mcgill.ca/software/.

[Montreal] *Image Guided Neurosurgery Laboratory.* Montreal Neurological Institute and the University of Western Ontario. http://www.irus.rri.on.ca/igns.

[Nabavi et al.] Nabavi, A., Gering, D., Pergolizzi, R., Hata, N., Kikinis, R., Jolesz, F., and Black, P. Shift happens. Submitted for publication.

[OpenGL] *Sun OpenGL for Solaris.* Sun. http://www.sun.com/software/graphics/OpenGL.

[Osborn1994] Osborn, A.1994. *Diagnostic Neuroradiology.* Mosby, St. Louis.

[Ousterhout1994] Ousterhout, J.1994. *Tcl and the Tk Toolkit.* Addison-Wesley, Massachusetts.

[Peters et al.1996] Peters, T., Davey, B., Munger, P., Comeau, R., Evans, and Olivier1996. Three-dimensional multimodal image-guidance for neurosurgery. *IEEE Transactions on Medical Imaging*, 15(2):121–128.

[Richolt et al.1998] Richolt, J., Teschner, M., Everett, P., Girod, B., Millis, M., and Kikinis, R.1998. Planning and evaluation of reorienting osteotomies of the proximal femur in cases of scfe using virtual three-dimensional models. In [Wells, Colchester and Delp1998].

[Roberts et al.1998] Roberts, D., Hartov, A., Kennedy, F., Miga, M., and Paulsen, K.1998. Intraoperative brain shift and deformation: a quantitative analysis of cortical displacement in 28 cases. *Neurosurgery*, 43:739–747.

[Ryan et al.1996] Ryan, M., Erickson, R., Levin, D., Pelizzari, C., Macdonald, R., and Dohrmann, G.1996. Frameless stereotaxy with real-time tracking of patient head movement and retrospective patient-image registration. *Journal of Neurosurgery*, 85:287–292.

[Schenk et al.1995] Schenk, J., Jolesz, F., Roemer, P., et al.1995. Superconducting open-configuration MR imaging system for image-guided therapy. *Radiology*, 195:805–814.

[Schroeder, Martin and Lorensen1996] Schroeder, W., Martin, K., and Lorensen, W.1996. *The Visualization Toolkit: An Object-Oriented Approach to 3- D Graphics.* Prentice Hall, NJ.

[Schroeder, Zarge and Lorensen1992] Schroeder, W., Zarge, J., and Lorensen, W.1992. Decimation of triangle meshes. *Computer Graphics.*

[Serra1982] Serra, J.1982. *Image Analysis and Mathematical Morphology.* London Academic.

[Shadyside] *Medical Robotics and Computer Assisted Surgery.* Shadyside Hospital and Carnegie Mellon University. http://www.mrcas.ri.cmu.edu.

[Speilman, Pauly and Metyer1995] Speilman, D., Pauly, J., and Metyer, C.1995. Nonuniform k-space sampling techniques: Interleaved sprials. *Magn Res Med*, 33:326–336.

[SPL] *Surgical Planning Lab.* Brigham & Women's Hospital and Harvard Medical School. http://www.splweb.bwh.harvard.edu:8000.

[UCLA] *Laboratory of Neuro Imaging.* UCLA. http://www.loni.ucla.edu.

[UCSF] *Department of Radiology.* UCSF. http://www.radiology.ucsf.edu.

[Viola and Wells III1995] Viola, P. and Wells III, W.1995. Alignment by maximization of mutual information. In *International Conference on Computer Vision.*

[Warfield et al.1998] Warfield, S., Kaus, M., Jolesz, F., and Kikinis, R.1998. Adaptive template moderated spatially varying statistical classification. In [Wells, Colchester and Delp1998], pages 431–438.

[Wells, Colchester and Delp1998] Wells, W., Colchester, A., and Delp, S., editors1998. *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'98)*, Boston. Springer-Verlag.

[Wells III et al.1996] Wells III, W., Viola, P., Atsumi, H., Nakajima, S., and Kikinis, R.1996. Multi-modalvolume registration by maximization of mutual information. *Medical Image Analysis*, 1(1):35–51.

[Westbrook and Kaut1993] Westbrook, C. and Kaut, C.1993. *MRI in Practice.* Blackwell Science.