Stereopsis-Guided Brain Shift Compensation

A Thesis Submitted to the Faculty
in Partial Fulfillment of the Requirements for
the Degree of Doctor of Philosophy

by

Hai Sun

Thayer School of Engineering, Dartmouth College
Hanover, New Hampshire

May 2004

Examining Committee:

Keith D. Paulsen, Ph.D., Chairman

Hany Farid, Ph.D.

David W. Roberts, M.D.

Alex Hartov, Ph.D.

James S. Duncan, Ph.D.

Dean of Graduate Studies

Signature of Author

Abstract

Brain deformation models have proven to be a powerful tool in compensating for soft tissue deformation during image-guided neurosurgery. The accuracy of these models can be improved by incorporating intraoperative measurements of brain motion. We have designed and implemented a passive intraoperative stereo vision (iSV) system capable of estimating the 3-D shape of the surgical scene in near real-time. The motion of the estimated cortical surface is then tracked over time. This cortical motion is used to guide a full brain model, which subsequently updates a preoperative MR volume. We have found that the iSV system is accurate to within approximately 1mm. Executing on a 1.1 GHz Pentium machine, the 3-D estimation from a stereo pair of 1024×768 resolution images requires approximately 60 seconds of computation. Based on data from representative clinical cases, we show that stereopsis guidance improves the accuracy of brain shift compensation both at and below the cortical surface. Specifically, we report an average error reduction of 3.7mm relative to the model estimate without using intraoperative data. In addition, advantages of our iSV system include its fast acquisition rates, predominantly automated steps, and overall accuracy. These characteristics allow us to capture other important cortical behavior such as the tissue resection and the pulsatile motion of the cortical surface.

Acknowledgments

I would like to thank my current and former teachers and mentors for their guidance, my colleagues and friends for their support, and my family for their love.

This work is supported by grant R01-NS33900 awarded by the National Institute of Neurological Disorders and Stroke.

The real voyage of discovery consists of not in seeking new landscapes but in having new eyes.

- Marcel Proust

Contents

List of Tables

List of Figures

Contributions

Below are what we consider the major contributions of this work:

- We have designed and implemented an intraoperative stereo vision (iSV) system capable of estimating the shape of the cortical surface. Some advantages of this approach are:

    1. The system utilizes the binocular optics the operating microscope.
    2. The system acquires the stereo images at 15 pairs per second.
    3. We have selected and implemented three matching constraints that have reliably solved the correspondence problem of stereopsis for the reconstruction of the cortical surface.
    4. The reconstruction of the cortical shape from an image pair takes less than 60 seconds on a 1.2 GHz PC.
    5. Through multiple experiments, we have found that the iSV system is accurate to within approximately 1mm.

- Based on the estimates of the 3-D cortical shape over time, the system tracks the cortical surface displacement.

    1. The tracking of the cortical surface does not interrupt the surgery.
    2. A variety of cortical motions, including soft tissue deformation, tissue retraction and resection, and the cortical pulsation, have been recovered.
    3. The tracking accuracy using the iSV system is comparable to that of an optical stylus.

- Under the guidance of the cortical motion, a poroelastic brain model predicts the full volume deformation, which is then used to update the preoperative MR.

1. Based on data from representative clinical cases, we show that stereopsis guidance improves the accuracy of brain shift compensation both at and below the cortical surface.
2. The versatility of the iSV system may aid the brain model to compensate for more complex tissue deformation occurring during retraction and resection.

- We have implemented and tested all the proposed steps as a complete end-to-end system.

List of Abbreviations

Imaging Modalities

CT	computed tomography
iCT	intraoperative computed tomography
iMRI	intraoperative magnetic resonance imaging
iSV	intraoperative stereo vision
MRI	magnetic resonance imaging
MRA	magnetic resonance angiogram
pMR	preoperative magnetic resonance
SV	stereo vision
uMR	updated magnetic resonance

Others

CSF	cerebrospinal fluid
FE	finite element
GS	gravitational sag
ICP	iterative closest point
IRED	inferred emitting diode
OR	operating room
PPM	perspective projection matrix

Chapter 1

Introduction

The value of medical images in providing navigational guidance during surgical interventions was recognized immediately after the techniques of generating these images were invented. Less than three months after the discovery of x-rays, a bromide print was used to aid the surgical removal of a needle from a woman's hand [93]. The advent of advanced imaging modalities such as magnetic resonance (MR) and computed tomography (CT) imaging has dramatically enhanced the safety and effectiveness of surgical procedures. Images generated from these techniques provide a detailed description of the patient's anatomy and can be used as a surgical atlas. The goal is to make surgical procedures safer, less invasive, and more effective.

We know from experience that, in order to utilize a map, one needs to establish a correlation between the map and the physical space through which one tries to navigate. A publicly displayed map at an international airport or a crowded shopping mall clearly points out the location of its readers on the map ("YOU ARE HERE!"). This information, along with other useful cues such as knowing the direction of north both in the physical space and on the map, helps to correlate the map and the physical space. Based on this correlation, travellers utilize the map to guide their journey. In order to use medical images to guide a surgery, a similar correspondence between these images

and the patient in the operating room (OR) needs to be established. With this correspondence, the position of a surgical target (e.g., a tumor or a herniation in the spine), clearly visible in the MR/CT images, can be more easily visualized by the surgeon in the OR. A suitable surgical trajectory leading to the target is then planned.

1.1 Image-guided Neurosurgery

1.1.1 Current Implementations

In neurosurgery, MR images of the patient's head acquired prior to the procedure are commonly used as the atlas to help surgeons navigate. In order to establish image guidance, the coordinate system of the preoperative MR (pMR) volume is *rigidly* rotated and translated to align with the coordinate system of the patient in the OR. Under homogeneous coordinates [1], this transformation takes the form:

$$\begin{pmatrix} X_W \\ Y_W \\ Z_W \\ 1 \end{pmatrix} = \begin{pmatrix} & \vdots & & \vdots \\ \dots & R & \dots & T \\ & \vdots & & \vdots \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} X_P \\ Y_P \\ Z_P \\ 1 \end{pmatrix}$$
$$= {}^{W}T_P \begin{pmatrix} X_P \\ Y_P \\ Z_P \\ 1 \end{pmatrix}, \tag{1.1}$$

where a point in the preoperative image volume, (X_P, Y_P, Z_P), is related to a point in the world coordinate, (X_W, Y_W, Z_W), through the 3×3 rigid-body rotation matrix, R

[1] Note that by using homogeneous coordinates, rotation *and* translation can be performed by multiplying with a *single* transformation matrix; that is, a rotation in the homogeneous 4-D space is equivalent to a rotation and translation in the underlying 3-D space.

(generated from 3 Euler angles, roll, pitch, and yaw), and the 3×1 translation vector, T. R and T are embodied into the single 4×4 matrix ${}^{W}T_{P}$, which relates the preoperative image frame with the world coordinate frame. Whereas the preoperative coordinate frame is defined by the image voxel grid, the world coordinate frame is usually defined by a 3-D positioning device used to track the position of the patient and surgical devices in the OR.

This transformation between the two coordinate systems, ${}^{W}T_{P}$, is computed based on the locations of a set of features both in the image volume and in the OR. The commonly used features for patient registration are scalp-attached fiducials, as shown in Figure 1.1(a). They are constructed of materials that will appear hyperintense on MR imaging so that they will be readily visible on preoperative imaging studies.

Given the position of n points in the world $(X_{Wi}, Y_{Wi}, Z_{Wi}), 1 \leq i \leq n)$ and their coordinates in the pMR volume $(X_{Pi}, Y_{Pi}, Z_{Pi}), 1 \leq i \leq n)$, ${}^{W}T_{P}$ can be determined by solving the following equation using standard linear algebra:

$$\begin{pmatrix} X_{W1} & X_{W2} & \dots & X_{Wn} \\ Y_{W1} & Y_{W2} & \dots & Y_{Wn} \\ Z_{W1} & Z_{W2} & \dots & Z_{Wn} \\ 1 & 1 & \dots & 1 \end{pmatrix} = \begin{pmatrix} & \vdots & & \vdots \\ \dots & R & \dots & T \\ & \vdots & & \vdots \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} X_{P1} & X_{P2} & \dots & X_{Pn} \\ Y_{P1} & Y_{P2} & \dots & Y_{Pn} \\ Z_{P1} & Z_{P2} & \dots & Z_{Pn} \\ 1 & 1 & \dots & 1 \end{pmatrix} \quad (1.2)$$

The process of obtaining the transformation between the pMR volume and the world coordinates is often referred to as *registration*. Here, we use the term registration to describe the establishment of a spatial correspondence between preoperative images and a physical space, i.e., the OR. As we will see later, the term registration can also mean aligning different medical images or aligning images with a computational model. To avoid confusion, we refer to the process of aligning the preoperative images with the patient in the OR as *patient* registration.

The estimated transformation expressed by Equation (1.1) allows any point in the

surgical field to be viewed in the context of patient anatomy, which is represented by the preoperative images. An example of the procedural process and outcome of establishing the correspondence between the operating field and the pMR is shown in Figure 1.1. Since the coordinate system in the OR is defined by the positioning system capable of tracking other surgical tools, registration can be expanded to include other surgical devices such as an operating microscope.

1.1.2 The Stereoscopic Operating Microscope

Since its introduction in 1962 [39], the stereoscopic operating microscope has been an integral part of a neurosurgical OR. Surgeons use the microscope to locate cortical structures and deeper features exposed during operation. As an extension of neurosurgical image-guidance, an operating microscope can also be employed to assist the surgical navigation. For example, at Dartmouth-Hitchcock Medical Center, a stereoscopic operating microscope (Model M695, Leica USA, Rockleigh NJ) is ceiling-mounted on a robotic platform, as shown in Figure 1.2. The microscope is continuously tracked by a positioning device (Polaris System, Northern Digital Inc. Ontario, CANADA) that employs infrared sensors. Focusing of the microscope is accomplished by aligning two laser beams propagated through each of the binocular paths of the microscope and which converge at the focal point. Through the transformation established during patient registration, this focal point can be viewed in the context of 3-D and patient-specific image studies. In Figure 1.1(c), the location of the focal point is depicted as the crosshairs in sagittal, axial, and trajectory views. By registering an important element of the surgeon's visual field with the preoperative images, the navigational system offers more direct assistance to the surgeon in selection of the trajectory through or around cerebral structures.

Although it is an important element of the surgeon's view through the operating microscope, the focal point is insufficient to accurately convey the geometrical relationship between the current surgical scene and the patient's anatomy. It is desirable to establish

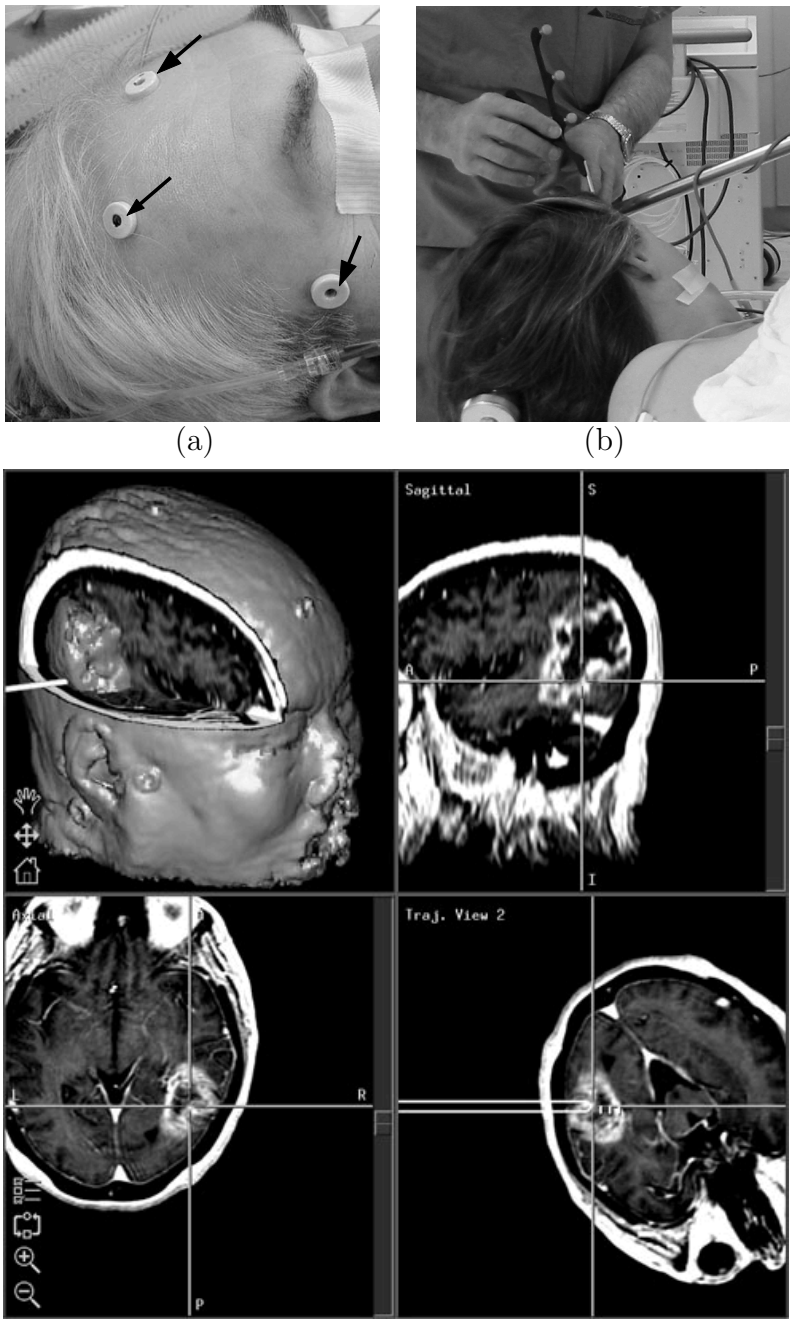

Figure 1.1: Pictorial description of image-guided surgery: (a) fiducials (arrows) are attached to a patient's scalp to correlate the patient's surface anatomy with the pre-operatively obtained image data set; (b) these fiducials are localized in the OR prior to surgery. (c) After patient registration, three-dimensional reconstruction, sagittal, axial, and trajectory views demonstrate a left parieto-occipital tumor. This display aids surgical planning.

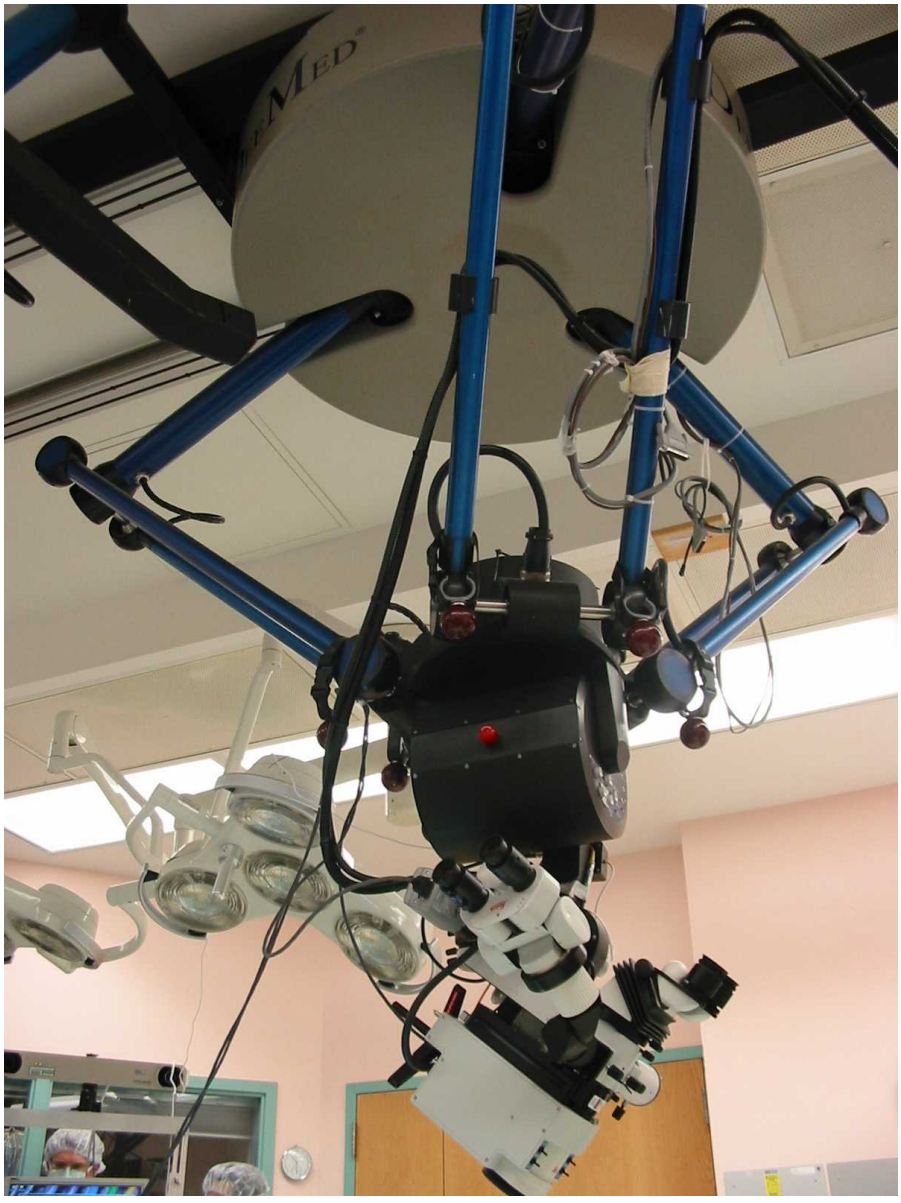

Figure 1.2: Stereoscopic operating microscope ceiling-mounted on a robotic platform.

an accurate correspondence between the surgeon's view through the operating microscope and the pre-operative 3-D MR data. The registration between preoperative data and the real-time microscope image allows the surgeon to access a surgical locus deep in the brain parenchyma. Projections calculated from 3-D pre-operative data can be overlaid with corresponding microscope views. Such image overlays ensure the surgeon a more accurate surgical trajectory.

In as early as 1985, Roberts et al. used an ultrasonic ranger-finder to track the operating microscope and displayed CT image data with the co-registered surgical scene in the correct perspective without requiring a stereotactic frame [71]. In their work on microscope-assisted guided interventions (MAGI), Edwards et al. carefully calibrated the imaging parameters of the microscope optics and experimented with injecting the projection of anatomical landmarks from the preoperative scans to both eyepieces [18]. The details of the implementation of this system will be discussed further in Chapter 2.

So far, the neuro-navigational systems discussed here are based on rigid registration, where any two coordinate systems are related to each other through a rotation and a translation. With rigid registration between the preoperative images, the patient coordinate system, and even the 2-D surgical scene under the microscope, the preoperative images help the surgeon plan the surgical trajectory, evaluate the extent of a lesion, and avoid critical brain structures. However, it is well recognized that brain tissue deforms non-rigidly during surgery. Neuro-navigational systems based solely on rigid registration are subject to errors due to brain deformation.

1.2 Brain Deformation

There are several factors that cause the brain to deform during open-cranial surgery. Parenchymal deformation can take place from the drainage of cerebrospinal fluid (CSF) which reduces the buoyancy forces acting on the brain. Hyperosmotic drugs such as

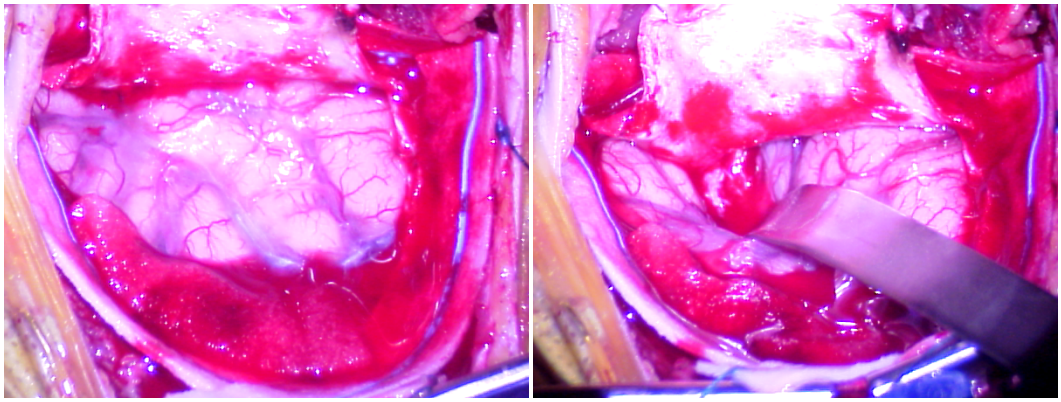

Figure 1.3: The cortical surface of the same patient viewed through the operating microscope immediately after the opening of the dura (left) and during the tissue retraction (right). Note the significant tissue deformation induced by the retractor blade.

mannitol can also cause shift by pulling fluid from the parenchyma and thus reducing brain volume. Brain motion can also result from tissue retraction and resection or the drainage of a cystic lesion. An example of brain deformation during neurosurgery is shown in Figure 1.3.

Many groups have investigated brain shift and the resulting compromise in accuracy of registration between patient and guidance system [85, 10, 38, 67, 36, 72, 57, 76, 35, 32]. Here, a few examples of these studies are given to show the general magnitude of these tissue motions. Let us begin at the cortical surface. Hill et al. employed intraoperative functional mapping to measure the post dura-removal location of approximately 60 cortical points in 21 cases, and compared them to their position in the preoperative images, reporting mean displacements of 4.4-5.6mm in the direction corresponding to gravity [36]. Through tracking the cortical vasculature using the focal point of the operating microscope, Roberts et al. measured a 1cm mean surface displacement, also dominantly in the direction of gravity [72]. Subsurface shift has been quantified in a study of ten cases by Dickhaus et al. [15] using an open MRI (0.2 Tesla). They measured a 6-7mm displacement for internal structures, namely the interhemispheric fissure and ventricles, as well as up to 2cm shift on the brain surface. Maurer et al. [57] completed

a qualitative study of intraoperative deformation in six cases using interventional MR (1.5 Tesla), noting that changes in ventricular volume greater than 1ml can indicate the occurrence of deformation. From these studies, we can see that brain shift causes significant changes in a patient's anatomy during open-cranial surgery. These changes may result in an error up to a few centimeters in registration between the patient's brain in the OR and its pMR images. This error in misregistration is incongruent with the goal of image-guided neurosurgery, i.e., to accurately target the surgical focus and at the same time avoid any critical brain structure. To compensate for brain shift, the preoperative images, which are static and cannot reflect tissue changes during surgery, need to be updated.

1.3 Brain Shift Compensation

As stated in the previous section, in order to maintain the fidelity between the patient's dynamic anatomy and its image presentation, the static preoperative scan needs to be updated. An attractive method for updating preoperative images is to image the full brain volume using intraoperative MRI (iMRI) or intraoperative CT (iCT) [96, 87, 8, 37, 86, 31]. These imaging devices offer multiplanal, high resolution volume data for accurate definition of most boundaries of tumor, vessels and heterogenic tissue structures during surgery. They are also helpful in monitoring potential intraoperative complications [8]. Despite these advantages, iMRI/iCT impose significant financial and logistical constraints. In addition, updates can only be obtained a few times during a case because the imaging process is lengthy and interrupts the surgery.

An alternative method for providing a full volume image update is to use computational models to predict brain deformation. The goal of brain modeling is to produce a displacement field over the patient's brain that accurately represents the current deformation in the operating room (OR). This vector field provides a *non-rigid* transformation that relates the preoperative image data with the patient anatomy during surgery, and

hence, can be used to update the pMR volume. Typical brain deformation models combine the bio-mechanical properties of the brain and boundary and volumetric conditions to predict this displacement field. Proposed bio-mechanical models of the brain utilize solid linear elastic [24, 75], solid nonlinear elastic [13], linear viscoelastic [62], or biphasic elastic equations [64, 29, 17]. Boundary and volumetric conditions include the direction of gravity, intracranial pressure, the location and size of the craniotomy, the level of cerebral spinal fluids, and the constraint of the rigid skull.

Compared to iMRI/iCT, brain deformation models provide an inexpensive and significantly less intrusive means to compensate for soft tissue deformation during image-guided neurosurgery. But this approach to brain shift compensation faces several challenges. First, recent reports have indicted that the pattern of brain shift is complex. Hartkens et al., using an iMRI, have determined that the direction of gravity alone is unreliable in predicting brain shift [33]. The direction of shift is likely to be a result of a complex interplay between gravity, the condition of the current surgical field (e.g., resected region), fluid pressure, and other forces. Given the complex tissue deformation pattern, brain models that make use of very limited intraoperative information (e.g., just the direction of gravity) are unlikely to accurately predict the full volume deformation.

Another challenge facing brain modeling lies in the uncertainties among model parameters. Model parameters include the mechanical properties of the brain and the boundary and volumetric forcing conditions under which the model is deployed. Limited work has been done to test mechanical properties *in vivo* and quantified results vary significantly [61, 4, 63, 92, 20]. The values we employed for our brain model are based on reports in the literature, as well as those deduced through optimization in our own experiments [61]. The brain material properties may in fact vary among patients due to differences in age and tissue pathology and our parameters do not reflect this variability. The uncertainty in the boundary and volumetric forcing conditions arise from the difficulty in estimating these conditions in the OR. For example, although point-based sensing devices are available to measure the intracranial pressure, accurate

full-volume pressure data is difficult to obtain; The level of CSF inside the cranium is usually assigned qualitatively based on the patient's position in the OR.

The complexity of brain deformation patterns and uncertainty in model parameters pose challenges for brain models to accurately predict the full volume displacement field. Several groups have proposed that the accuracy of these models can be improved when measurements of brain motion acquired during surgery are used to guide the model estimates [1, 24, 75, 49, 80].

1.4 Intraoperative Data

Imaging modalities used for measuring brain deformation in the operating room (OR) include iMRI [56, 7, 28, 30, 73, 96], iCT [52], and intraoperative ultrasound (iUS) [49, 12, 9, 41, 44, 90]. In an attempt to reduce the amount of intraoperative imaging, Ferrant et al. have combined a linear elastic brain model with MR-estimated motions of cortical surface and lateral ventricles to predict tissue motions at other parts of the brain [24]. Although iMRI offers rich intraoperative data, as stated earlier, its high cost and stringent implementation requirements hinder its acceptance. Intraoperative ultrasound is considerably less cumbersome in the OR and is gradually becoming accepted as an alternative to iMRI/iCT. Unsgaard et al. demonstrated that iUS was capable of successfully localizing brain lesions in 112 out of 114 clinical cases [90]. By tracking 1mm metal beads during their porcine studies [51], Lunn et al. showed that a co-registered iUS has a mean target localization error of 1.5mm and a mean vector error for displacement of 1.1mm. Despite these encouraging results, iUS still faces some challenges when employed to track structures in the brain. First, spatial and temporal variation in image quality limits its ability to identify homologous features in the brain for tracking purposes. Second, iUS requires tissue contact for imaging, which may induce further brain deformation and hence compromise tracking accuracy.

The underlying goal of these techniques is to acquire intraoperative data to guide brain deformation models. Along these lines, but in contrast to these methods, several groups have employed non-contacting techniques for estimating the 3-D cortical topology. For example, in their phantom studies, Audette et al. [1] used a commercial 3-D laser range sensor to scan the cortical surface and proposed to use surface motion to guide a brain model. Miga et al. employed a similar scanner to digitize cortical features which were used to register the patient in the OR with the pMR volume [60]. According to their experience, it took approximately 1.5 minutes to scan the exposed cortical surface. Skrinjar et al. utilized a pair of stereo cameras to guide their bio-mechanical brain model [77]. In order to avoid the stereo correspondence problem, they coupled the topological reconstruction with their brain model. As a result, only estimates that satisfied the model were produced. Because of a limited camera calibration and a restrictive Lambertian assumption, the accuracy of estimating the shape of the cortical surface using their method has not been fully evaluated.

1.5 Proposed Solutions

So far, we have seen that the conventional neuro-navigational system based on rigid registration is subject to inaccuracies caused by brain deformations. Factors that cause such deformations include gravity, drug administration, loss of cerebrospinal fluid (CSF), and tissue retraction and removal. As a result of these changes, image studies acquired prior to surgery are no longer an accurate representation of the patient's anatomy during surgery. Strategies to update the static preoperative images include full-volume intraoperative scanning techniques such as iMRI/iCT and computational brain models that develop estimates of the full volume deformation. Using computational models to compensate for brain shift is attractive because of its low cost and some initial success has been reported by several groups using this strategy [13, 24, 29, 46, 61, 64, 91]. Although brain deformation models offer an inexpensive alternative, but their accuracy warrants

concerns due to the complexity of tissue motion and the uncertainty in model parameters. To contend with these challenges, we have designed and implemented a brain shift compensation system that uses intraoperatively measured brain motions to guide a brain model for updating the pMR image [81, 80, 83, 82]. Our goal is to create an affordable and accurate solution to brain deformation compensation for incorporation into neuro-navigational image-guidance system.

This system is based on both rigid and non-rigid registrations. Rigid registrations provide transformations between all the coordinate systems, such as the pMR volume, the patient in the OR, and the operating microscope field of view. Non-rigid registrations consist of the estimation of the cortical motion, brain modeling and the pMR update.

Central to this system is the estimation of brain motion during surgery. Similarly to the studies using laser range scanners and stereo vision, we have also designed a non-contacting strategy that measures the 3-D cortical topology using stereo vision (SV). Our SV system is incorporated into the standard stereoscopic operating microscope. Our decision to utilize the stereoscopic operating microscope for the estimation of the cortical motion is based on the following considerations:

1. The conventional operating microscope is readily available in most neurosurgical suites, therefore, the system does not require a new imaging device.

2. The coordinate systems of the microscope and its image plane are registered with the patient in the OR and the preoperative image volume (Section 1.1.2). This registration offers convenience in surface tracking and data assimilation in the brain model.

3. The microscope is equipped with binocular optics. By attaching a pair of CCD cameras to the optics, we can perform stereopsis.

4. Stereopsis is an attractive strategy for estimating the shape of the cortical surface because it requires no tissue contact and does not interrupt the surgery.

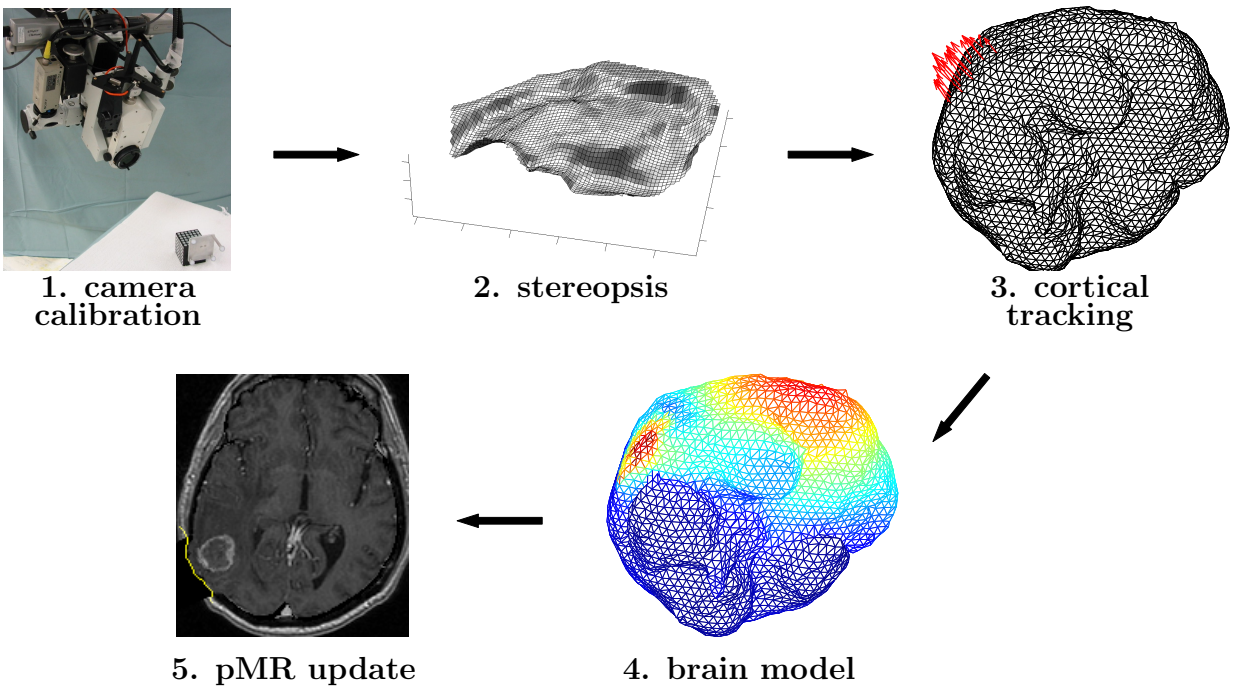

Figure 1.4: Flow diagram illustrating the steps in stereopsis-guided brain shift compensation.

Based on estimates of the 3-D cortical shape over time, the system tracks the cortical surface displacement. Under guidance from the cortical motion, a poroelastic brain model predicts the full volume deformation, which is then used to update the pMR. We have implemented and tested all of the proposed steps as a complete end-to-end system.

Next, we will give an overview of this system which consists of several steps leading toward our goal of updating pMR during surgery. The methodology developed in this thesis combines the rigid image-guided techniques, stereopsis, with a nonrigid brain model to correct for intraoperative brain movement. We begin by presenting the steps involved in stereopsis-guided brain shift compensation. As depicted in Figure 1.4, our technique involves five basic steps:

1. perform the system calibration to register the preoperative image volume with the patient in the OR and the 2-D microscope images,

2. estimate the shape of the 3-D cortical surface using stereopsis,

3. estimate the displacement of the 3-D cortical surface over time,

4. incorporate the estimated motion of the cortical surface into a computational model to produce a full volume description of brain deformations,

5. and finally, update the pMR images.

These five steps utilize both rigid and nonrigid registration strategies introduced in the previous sections. Step 1 employs rigid registration to establish the conventional image-guided neurosurgery. Steps 2 and 3 estimate intraoperative data through intraoperative imaging. Steps 4-5 employ nonrigid registration to update the pMR volume. As we will see later, the rigid registration step serves as the basis for non-rigid registration steps. In the next five chapters, each of the five steps will be discussed in detail. In each chapter, methods utilized in the corresponding step are first presented and then followed by the relevant results.

Chapter 2

Rigid Registration

In this chapter, we describe the steps required to implement conventional Microscope-based Image-guided Neurosurgery. Most of this work has been discussed in [79]. The image guidance employed here is based solely on rigid registration. To establish Microscope-based Image-guided Neurosurgery, the following three steps are required: a rigid-body registration between the preoperative and OR spaces, calibration of the microscope optics, and tracking the patient and microscope throughout the surgery. These three steps establish a series of coordinate transformations as shown in Figure 2.1. In this figure the common (world) coordinate system is given by the Polaris, an optical 3-D positioning system [1]. The final overall transformation from the preoperative data [2] to the real-time microscope image is determined through a series of coordinate transformations: (1) the transformation from 3-D microscope frame to 2-D microscope image is denoted as ${}^{I}T_{M}$ [3]; (2) the transformation from world to 3-D microscope frame is denoted as

[1] Hybrid Polaris, Northern Digital, Waterloo, Ontario.

[2] Prior to surgery, a standard volumetric MRI (typically a $256 \times 256 \times 124$ matrix with the voxel size $0.97 \times 0.97 \times 1.5mm^3$) is obtained for the patient using 1.5 Tesla GE MR scanner.

[3] In order to preserve the linearity of all transformations, we use only the pinhole perspective projection to describe the imaging process. The full camera model which includes the nonlinear lens distortion compensation will be presented in the later section.

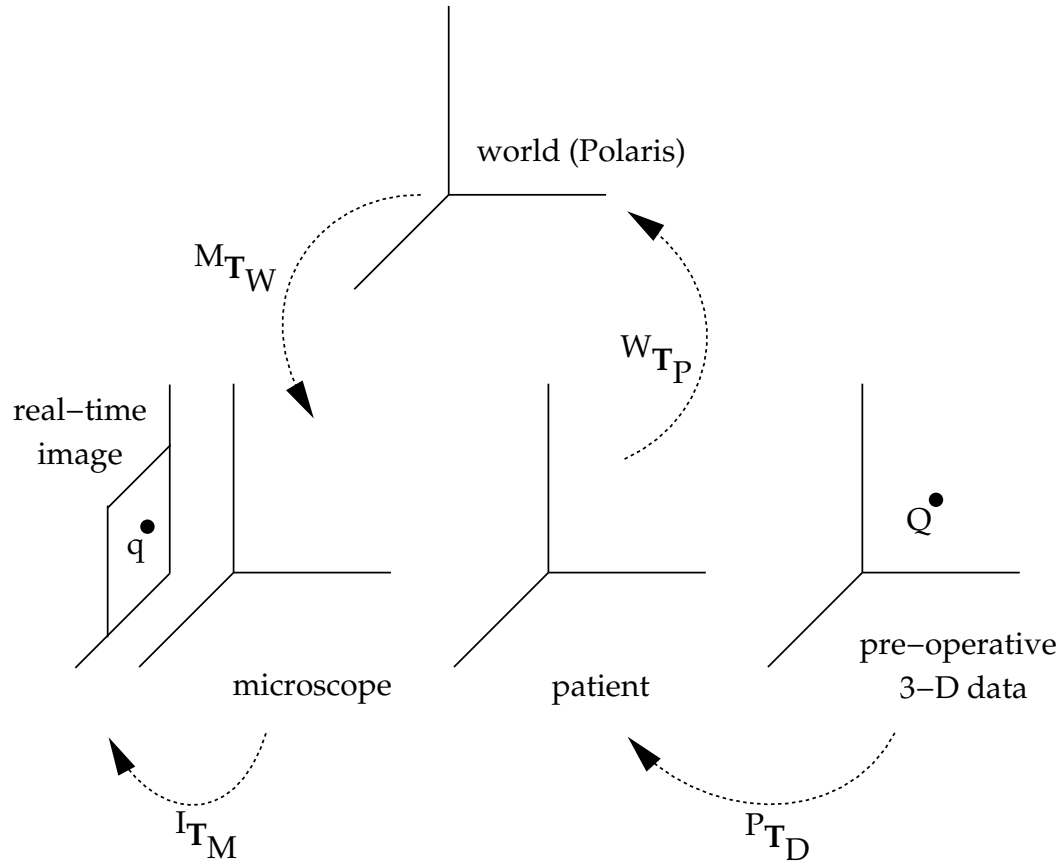

Figure 2.1: Coordinate transformations from the 3-D pre-operative data to the real-time 2-D microscope image.

${}^{M}T_{W}$; (3) the transformation from patient to world is denoted as ${}^{W}T_{P}$; (4) and finally, the pre-operative image to patient transformation is symbolized as ${}^{P}T_{D}$. The complete transformation is simply a concatenation of individual transformations:

$$\vec{q} \;=\; \left({}^{I}T_{M}\ {}^{M}T_{W}\ {}^{W}T_{P}\ {}^{P}T_{D}\right)\ \vec{Q}, \tag{2.1}$$

where $\vec{Q}$ is a point in the 3-D pre-operative coordinate space and $\vec{q}$ is the same point in the 2-D real-time microscope image coordinate system. We next describe the procedure of system calibration for obtaining each of the four transformations in Equation (2.1). From left to right in this equation, our discussion begins with imaging parameters, ${}^{I}T_{M}$.

2.1 Imaging Parameters: ${}^{I}T_{M}$ and others

The calibration of imaging parameters is not only important for registering the preoperative data with the 2-D microscope image but also, as we will see later, crucial for the implementation of stereopsis. The process of image formation corresponds a ray of points in the scene to a pixel in the image plane. Prior to surgery, a pair of stereo cameras are calibrated to obtain their intrinsic and extrinsic parameters using the method developed by Tsai [89]. These parameters are used to relate the points in the scene to the pixels in the image plane. Other calibration methods can also be employed. For example, the orthocenter property and the calibration of the image center through vanishing points by Caprile and Torre [11], direct calibration of the projection matrix described by serval groups including Faugeras and Toscani [22], and a method for calibrating a stereo system from arbitrary stereo images developed by Thacker and Mayhew [84]. We choose Tsai's method to calibrate the microscope optics due to its relatively simple camera model and its accuracy in capturing the image formation process.

Shown in Figure 2.2 is our operating microscope with two CCD cameras [4] optically coupled. The operating microscope is tracked throughout the surgery using infrared emitting diodes (dashed arrow). We begin by presenting the camera model employed by Tsai's camera calibration strategy.

2.1.1 Camera Model

The schematic drawing of the pinhole camera model followed by lens distortion compensation is shown in Figure 2.3. In order to determine the 2-D image coordinates of a point specified by its 3-D position in a world coordinate system, a series of transformations and projections must be performed:

1. rigid body transformation from world coordinates (X, Y, Z) to camera coordinates

[4] Sony DFW-x700, Resolution $1024(H) \times 768(V)$, Sony Corp., New York, NY, USA

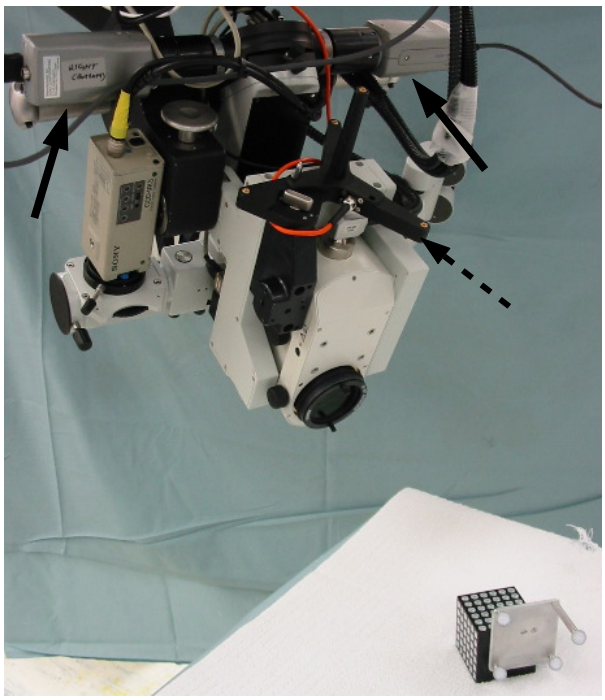

Figure 2.2: Two CCD cameras (solid arrows) have been attached to the binocular optics of the operating microscope. The intrinsic and extrinsic camera parameters are calibrated prior to surgery using the calibration target shown in the lower right. The operating microscope is tracked throughout the surgery using infrared emitting diodes (dashed arrow).

(X_c, Y_c, Z_c)

2. perspective projection from camera coordinates (X_c, Y_c, Z_c) to undistorted sensor coordinates (x, y)

3. adjustment from undistorted sensor coordinates (x, y) to distorted sensor coordinates (x_d, x_d)

4. unit conversion (mm to pixels) from distorted sensor coordinates (x_d, x_d) to image coordinates (x_f, y_f)

The individual transformations and projections between each of these coordinate systems are given below.

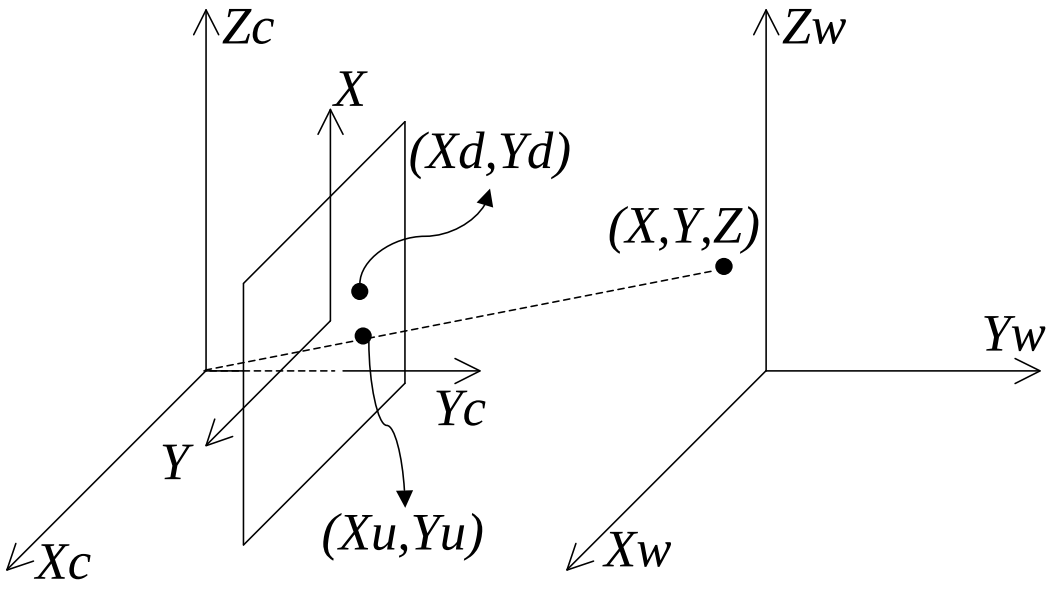

Figure 2.3: Camera Model: perspective projection and lens distortion.

1. The transformation from world coordinates (X, Y, Z) to camera coordinates (X_c, Y_c, Z_c) in homogeneous coordinates is:

$$\begin{pmatrix} X_c \\ Y_c \\ Z_c \\ 1 \end{pmatrix} = \begin{pmatrix} & \vdots & & \vdots \\ \dots & R & \dots & T \\ & \vdots & & \vdots \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix}, \tag{2.2}$$

where the rotation matrix, R, (generated from 3 Euler angles, roll, pitch, and yaw) and the translation matrix, T, are given by:

$$R = \begin{pmatrix} r_1 & r_2 & r_3 \\ r_4 & r_5 & r_6 \\ r_7 & r_8 & r_9 \end{pmatrix} \quad \text{and} \quad T = \begin{pmatrix} t_x \\ t_y \\ t_z \end{pmatrix}. \tag{2.3}$$

The world coordinate frame for camera calibration can be arbitrarily defined. During a surgical procedure, the microscope is moving through a ceiling-mounted robotic arm (Surgiscope, Elekta AB, Stockholm), as shown in Figure 1.2. In order to avoid updating the camera parameters upon movement of the microscope,

we choose a coordinate frame which is rigidly attached to the microscope as the world coordinate frame for camera calibration. This coordinate system is defined by an optical tracker consisting of infrared emitting diodes tracked by the Polaris system, as indicated by the dotted arrow in Figure 2.2. Note that this world coordinate system is not the same as the world coordinate frame in Figure 2.1, which is defined by the Polaris.

2. Undistorted sensor coordinates (x, y) are computed from camera coordinates (X_c, Y_c, Z_c) using the following equations:

$$x = \frac{fX_c}{Z_c} \quad \text{and} \quad y = \frac{fY_c}{Z_c}, \tag{2.4}$$

where f is the camera focal length [5].

So far, steps 1 and 2 constitute the simple pinhole camera model. Under this model, a point in the world coordinates, (X, Y, Z), is projected onto a point in the 2-D imaging plane, (x, y). Combining equations 2.2 and 2.4, we have :

$$\begin{aligned} \begin{pmatrix} x_s \\ y_s \\ s \end{pmatrix} &= \begin{pmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} & \vdots & & \vdots \\ \dots & R & \dots & T \\ & \vdots & & \vdots \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix} \\ &= P \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix}, \end{aligned} \tag{2.5}$$

where the projected image coordinates are $(x, y) = (x_s/s, y_s/s)$. The intrinsic camera parameters are the focal length, f, and the extrinsic parameters are the 3×3 rigid-body rotation matrix, R, and the 3×1 translation vector, T. All of

[5]The focal length of the lens is the distance from the lens where the image of an object that is infinitely far away is focused.

these parameters are embodied into the single 3×4 matrix P, referred to as the perspective projection matrix (PPM). ${}^{I}T_{M}$ in Equation (2.1) is a PPM.

Next, we apply the radial distortion model and then represent the process by which digital images are actually formed (i.e., the conversion of the sensor coordinates to the final image coordinates).

3. The conversion of undistorted sensor coordinates (x, y) to distorted sensor coordinates (x_d, y_d) takes the form:

$$\begin{aligned} x_d &= x - x_d(\kappa_1 r^2 + \kappa_2 r^4 + ... + \kappa_n r^{2n}) \\ y_d &= y - y_d(\kappa_1 r^2 + \kappa_2 r^4 + ... + \kappa_n r^{2n}), \end{aligned} \tag{2.6}$$

where $r = \sqrt{x_d^2 + y_d^2}$ and $(\kappa_1, \kappa_2, ..., \kappa_n)$ are the lens distortion parameters.

We have found that a first-order model is sufficient to compensate for lens distortion in our case. Equation (2.6) is then simplified to:

$$x_d = x - \kappa x_d r^2 \quad \text{and} \quad y_d = y - \kappa y_d r^2, \tag{2.7}$$

where κ is the lens distortion coefficient and $r = \sqrt{x_d^2 + y_d^2}$.

4. The conversion of distorted sensor coordinates (x_d, y_d) to image coordinates (x_f, y_f) takes the form:

$$\begin{aligned} x_f &= S_x \frac{N_{fx}}{N_{cx} d_x} x_d + C_x \\ y_f &= \frac{1}{d_y} x_d + C_y, \end{aligned} \tag{2.8}$$

where N_{fx} is the number of pixels in a line sample by the computer; N_{cx} is the number of sensor elements in the X direction; d_x and d_y are the center to center distance between adjacent sensor elements in X and Y directions, respectively; (these parameters may be obtained from the manufacturer supplied camera and frame grabber specifications); S_x is a scale factor arising from digitizer-to-computer memory translation and (C_x, C_y) is the image center.

Table 2.1: Camera parameters to be calibrated

parameter	description
intrinsic	
f	focal length
κ	lens distortion coefficient
S_x	non-square pixel scalar
C_x, C_y	camera center
extrinsic	
$R(\theta_x, \theta_y, \theta_z)$	rigid-body rotation
$T(t_x, t_y, t_z)$	rigid-body translation

Note that we now have a camera model that projects a point in the world to its image coordinates, the next step is to determine (i.e., calibrate) several unknown parameters among the equations presented above. In particular, the extrinsic camera parameters to be calibrated are the rotation and translation matrices (R, T) and the intrinsic parameters are the focal length (f), lens distortion coefficient κ, scale factor (S_x), and image center (C_x, C_y). These parameters are listed in Table 2.1.

2.1.2 Camera Calibration

Principles

Equations (2.2) - (2.8) specify how to project from world to image coordinates assuming perspective projection and a radial lens distortion model. As noted, the extrinsic and intrinsic camera parameters need to be calibrated in order to perform such a projection. Since the camera parameters consist of 11 unknowns (i.e., 3 rotation angles $(\theta_x, \theta_y, \theta_z)$, 3 translation components (t_x, t_y, t_z), the focal length (f), the lens distortion parameter(κ_1; assuming a first-order approximation), the scale factor (S_x), and the 2 components of the image center (C_x, C_y)), 11 equations with known world and image coordinates should be sufficient to solve for the 11 camera parameters. The transformation from world to image coordinates can be expressed, in homogeneous coordinates, as:

$$\begin{pmatrix} x_s \\ y_s \\ s \end{pmatrix} = \underbrace{\begin{pmatrix} \alpha_x & 0 & C_x & 0 \\ 0 & \alpha_y & C_y & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}}_{A} \underbrace{\begin{pmatrix} r_1 & r_2 & r_3 & t_x \\ r_4 & r_5 & r_6 & t_y \\ r_7 & r_8 & r_9 & t_z \\ 0 & 0 & 0 & 1 \end{pmatrix}}_{B} \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix}, \tag{2.9}$$

where the projected image coordinates are $(x_f, y_f) = (x_s/s, y_s/s)$. Note that matrix B is the rigid body transformation matrix (Equation (2.2)) and matrix A incorporates the perspective projection from camera to sensor coordinates (Equation (2.4) and the transformation from sensor to image coordinates (Equation (2.8)). For exposition, the lens distortion parameters $(\kappa_1, \kappa_2, ..., \kappa_n)$ are assumed to be zero. This restriction will be relaxed later when the full calibration procedure is presented. The extrinsic and intrinsic camera parameters can be solved for by setting up a set of at least 10 equations of the form in Equation (2.9):

$$\begin{pmatrix} x_{s1} & x_{s2} & \dots & x_{sn} \\ y_{s1} & y_{s2} & \dots & y_{sn} \\ s_1 & s_2 & \dots & s_n \end{pmatrix} = \begin{pmatrix} \alpha_x & 0 & C_x & 0 \\ 0 & \alpha_y & C_y & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \cdot \begin{pmatrix} r_1 & r_2 & r_3 & t_x \\ r_4 & r_5 & r_6 & t_y \\ r_7 & r_8 & r_9 & t_z \\ 0 & 0 & 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} X_1 & X_2 & \dots & X_n \\ Y_1 & Y_2 & \dots & Y_n \\ Z_1 & Z_2 & \dots & Z_n \\ 1 & 1 & \dots & 1 \end{pmatrix}, \tag{2.10}$$

Given the position of n points in the world $(X_i, Y_i, Z_i), 1 \leq i \leq n)$ and in image coordinates $(x_i, y_i, 1 \leq i \leq n)$, the camera parameters can be determined *within a scale factor*, using standard linear algebra. The scale factor is s in the left hand side of Equation (2.9). This scale factor comes from the non-linear perspective projection equations, $x = \frac{fX}{Z}$, that is, for a fixed x image coordinate and a fixed X world coordinate, there exist several solutions for f and Z which satisfy $x = \frac{fX}{Z}$. This non-linearity and the added non-linearity of a radial lens distortion model complicates the camera calibration procedure. In particular, in addition to solving linear equations, a non-linear

optimization (e.g., gradient descent) on the focal length f, translation component t_z, and lens distortion parameters, $(\kappa_1, \kappa_2, ..., \kappa_n)$ is needed in order to recover the *exact* camera parameters.

Procedure

In the previous section, we have shown that a set of at least 11 points with known world and image coordinates is needed in order to determine a camera's extrinsic and intrinsic parameters. The points in world coordinates can be coplanar (i.e., lie in a single plane) or non-coplanar (i.e., lie in several planes). Accordingly, as outlined by Tsai [89], there are two forms of calibration. Since non-coplanar calibration yields considerably more accurate results than coplanar calibration, we have decided to perform this form of calibration.

In order to estimate the intrinsic and extrinsic parameters we first manufactured a 3-D calibration target with 108 fiducial markers having known relative positions, as shown in Figure 2.4. This target is placed in the microscope field of view, as shown in Figure 2.2, and an image is digitized from each camera of the stereo pair. The centroid of each calibration marker is automatically estimated and its spatial coordinates, (x_f, y_f), along with its known 3-D positions, (X, Y, Z) are recorded. These image and world coordinates are then used to estimate each camera's intrinsic and extrinsic parameters [95]. The details of how to compute these parameters can be found in Appendix A.

Although each camera is independently calibrated, they are all rigidly attached and calibrated to a common world coordinate system, therefore the camera positions are known relative to one another. As we will see later, this pairwise calibration strategy is important for computing stereo correspondence.

Having obtained ${}^{I}T_{M}$ in Equation (2.1) and other necessary camera parameters, we now describe the calibration of the next transformation matrix ${}^{M}T_{W}$.

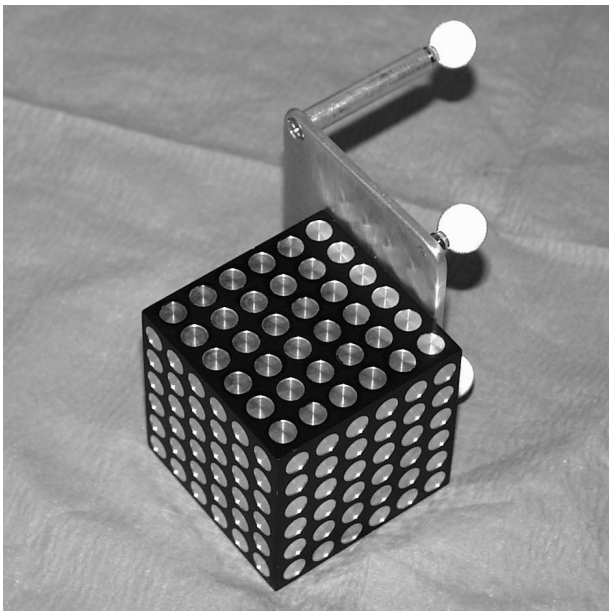

Figure 2.4: The calibration target with the attached infrared reflecting beads for localization of the target in the world coordinate reference frame.

2.2 Microscope Tracking: ${}^{M}T_{W}$

As we have mentioned in the previous section, the microscope is moving through a ceiling-mounted robotic arm to accommodate the surgical needs during a procedure. Camera calibration is performed with respect to a coordinate system defined by an optical tracker which is rigidly attached to the microscope. The optical tracker consists of infrared emitting diodes (IRED), as indicated by arrows in Figure 2.5. Through tracking these markers, the Polaris records the current position and orientation of the microscope with respect to the true world coordinate system defined by the Polaris. These data are used to construct the matrix ${}^{M}T_{W}$ in Equation (2.1), which relates the coordinate system of camera calibration with the true world coordinate system defined by the Polaris.

2.3 Patient Tracking and Registration: ${}^{W}T_{P}$ and ${}^{P}T_{D}$

The final calibration step requires a link between the patient and world and also between the patient and 3-D pre-operative data to be established. In studies conducted by other groups (e.g., [18, 55, 94]), bone-implanted markers were used to register and/or re-register the patient to the pre-operative data. As stated in Section 1.1, we attach small fiducials to the patient's scalp in order to minimize the invasiveness to the patient. These points are localized in the preoperative data space and on the patient's head prior to surgery [25]. From these corresponding points, the transformation matrix ${}^{P}T_{D}$ is estimated using Equation (1.2) [34].

The position of the patient with respect to the world system is determined by tracking (throughout the surgery) the location of IREDs rigidly attached to the patient's head clamp. A patient tracker consisting of 4 IREDs is shown in Figure 2.5. As with the tracking of the microscope, the tracking of the patient's head clamp provides the final transformation matrix ${}^{W}T_{P}$.

This completes the estimations of all the transformations in Equation (2.1) needed to facilitate the projection of a 3-D point in the preoperative image volume to the 2-D microscope image plane. After the projection calculation, a lens distortion model and a digital image conversion are applied to obtain the desired 2-D pixel coordinate. These steps establish the non-invertible rigid registration between the 3-D preoperative volumetric data and the 2-D microscope image plane. With this registration, useful information from preoperative images can be presented to the surgeon. Next, we describe some image display strategies.

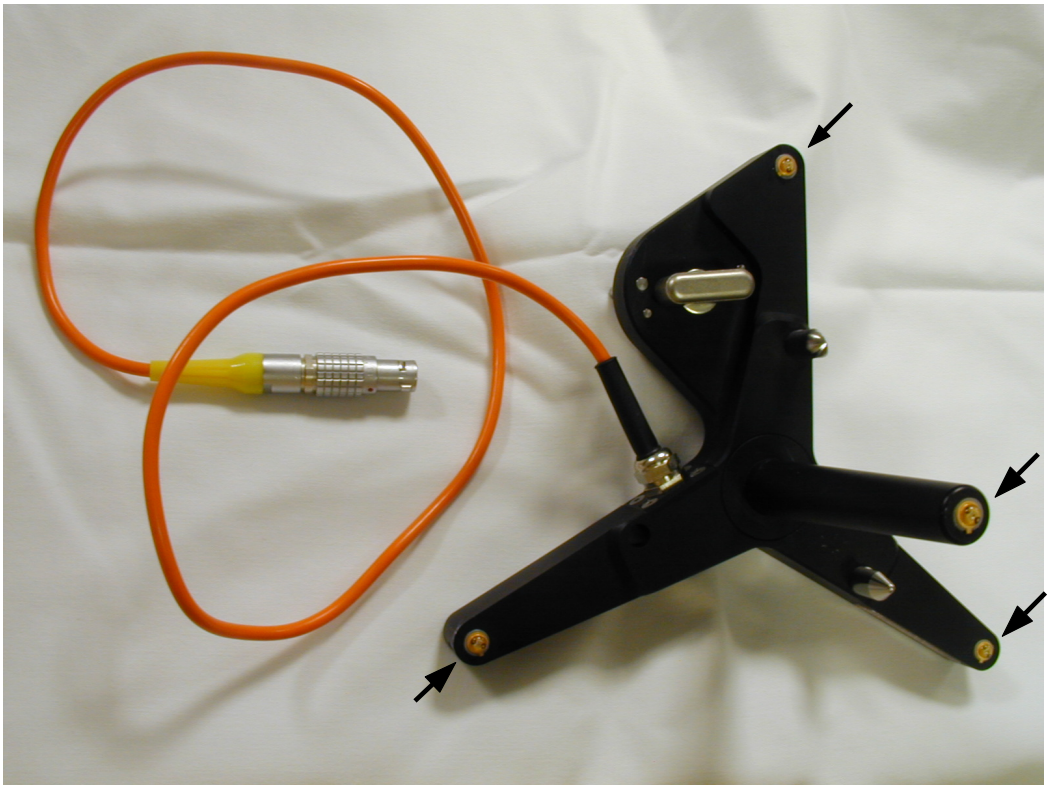

Figure 2.5: The tracking device which consists of infrared emitting diodes (IREDs, arrows) has been rigidly attached to the microscope for monitoring its location and orientation, and to the patient head clamp for updating the patient position with respect to the world system.

2.4 Image Display

We employ two strategies to overlay the information from the preoperative image volume with the current surgical scene.

1. Segmented Feature Overlay. Salient features such as cortical vessels and tumors can be segmented from pMR volume. The 3-D coordinates of these features can be projected onto the 2-D microscope images.

2. Plane Intersection Overlay. The second strategy avoids brain segmentation and utilizes the focal plane of the operating microscope. As described earlier, we can focus the microscope by aligning two laser beams projected through each of the

Surgery No.	Age (yr)/ Gender	Lesion/Pathology	Location	Procedure
1	27/F	Ganglioglioneurocytoma	Right Frontal	Resection
2	55/F	Metastatic carcinoma	Left Frontal	Resection
3	39/F	Metastatic carcinoma	Right Frontal	Resection

Table 2.2: Patient Characteristics

binocular paths of the microscope and which converge at the focal point. This focal point and one optical ray (there are two optical rays due to stereo) determines a focal plane in the surgical field, which is perpendicular to the optical ray and passes through the focal point. The same focal plane can be formed in the coregistered pMR volume. All pMR voxels on this focal plane can be projected and overlaid with the microscope image of the current surgical scene. Another plane parallel with the focal plane but located deeper within the brain can be formed and all voxels associated with this plane can again be projected and overlaid with the microscope image. The overlay of a deeper brain structure with the current surgical scene helps the surgeon evaluate the surgical trajectory.

2.5 Results

The system calibration establishes the rigid registration between the 3-D pMR volume and 2-D microscope images. This system has been implemented by Rick et al. using the Graphical User Interface (GUI) platform offered by MATLAB [68]. Three patient cases are chosen to illustrate the accuracy of this rigid registration. The patient characteristics are listed in Table 2.2. The microscope optics were calibrated before the patients' arrival in the OR. Figure 2.6 shows an example of the camera calibration results. In this figure, an red point indicates the true position of a feature point and an blue open circle denotes its estimated location by projecting a centroid in the world coordinates to the image

Table 2.3: Camera Calibration Errors for Surgery 1, 2, and 3

	Mean (pixel)	Max (pixel)	Min (pixel)
Surgery 1			
left	0.53	0.97	0.14
right	0.45	0.88	0.16
Surgery 2			
left	0.27	0.54	0.11
right	0.21	0.33	0.09
Surgery 3			
left	0.36	0.78	0.13
right	0.37	0.82	0.18

plane using the camera parameters. To assess the accuracy of these calibrations, we compared the 2-D image coordinates of the centroid of each calibration marker with the estimated image coordinates. The RMS differences between the true image coordinates of these centroids and their projected coordinates are given in Table 2.3. These error values indicate that Tsai's camera model captures, with sufficient accuracy, the image formation process by the microscope optics.

After the patient was anesthetized in the OR, patient registration was performed. Two metrics, Fiducial Registration Error and Target Registration Error, were used to quantify the registration accuracy [26]. Results from these three cases are presented in Table 2.4. The values of FRE and TRE [6] in these three cases are comparable to results reported by other groups using scalp-attached fiducials for patient registration.

We employed the first image display strategy, Segmented Feature Overlay, to display the coregistered pMR volume with the current surgical scene under the microscope during the first clinical case. Figure 2.7(a) shows the microscope image of the surgical field. The cortical vasculature segmented from the patient's magnetic resonance angiogram (MRA) and the tumor boundary segmented from the T-1 pMR volume were projected on the microscope image plane, as shown in Figure 2.7(b) and (c), respectively. Figure 2.7(d) contains a RGB image with the surgical scene, cortical vasculature,

[6] The target for computing TRE was generally defined in the vicinity of the surgical locus in the brain. In all three cases reported here, the target was a point in the tumor region.

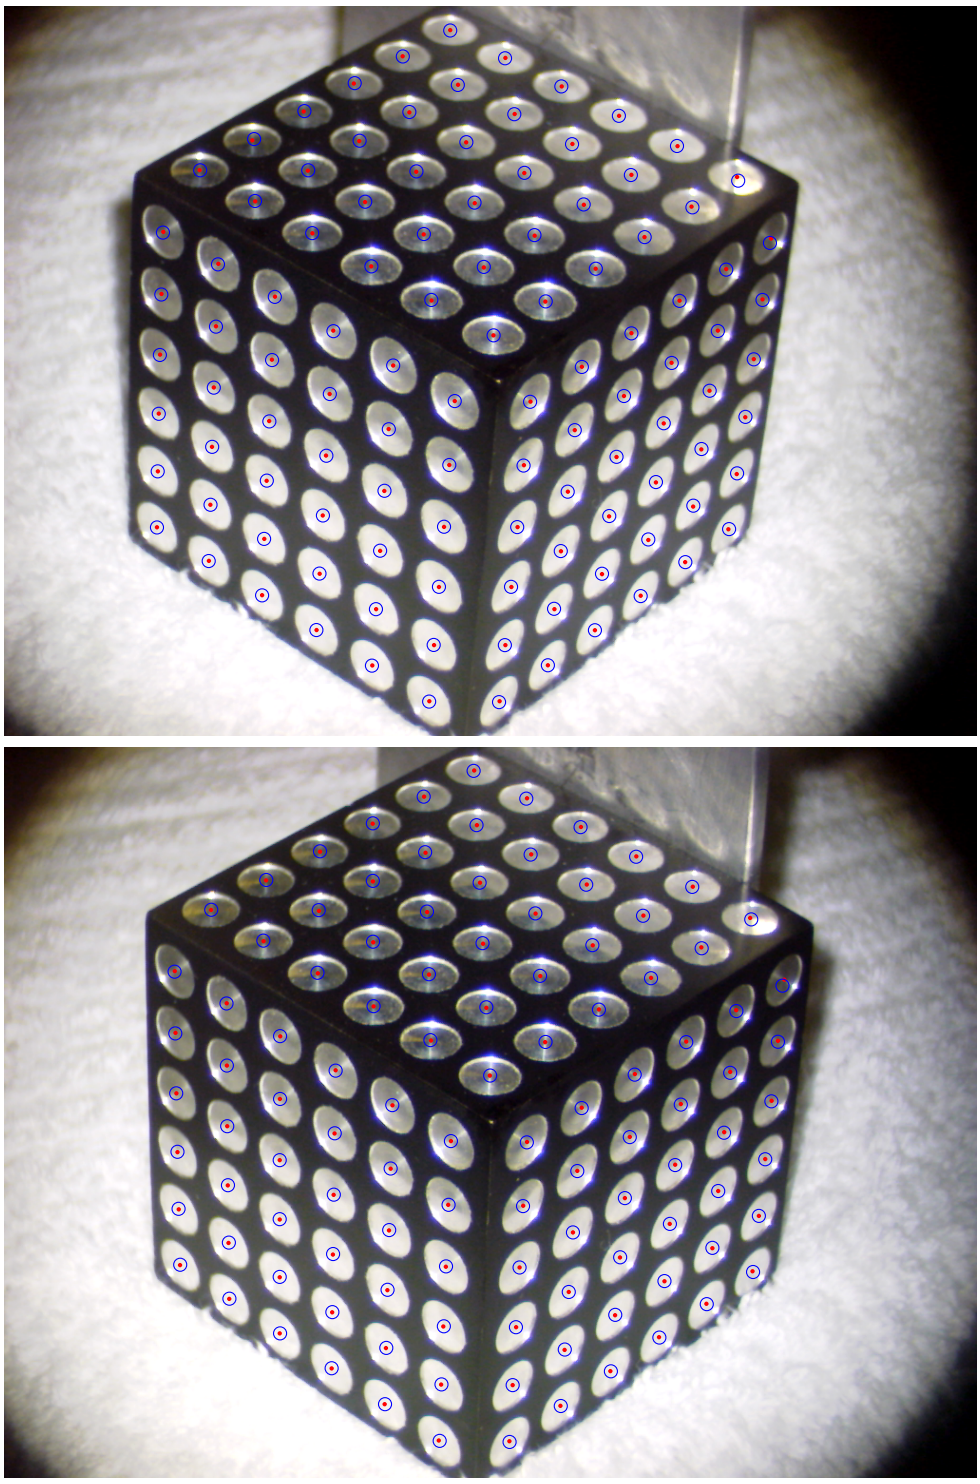

Figure 2.6: Shown are the camera calibration results of both left (top) and right (bottom) cameras prior to surgery 1; open circles denote the true positions of the centroids of all circular markers in 2-D images, and dots denote estimated positions by projecting the 3-D locations of centroids to the 2-D image planes using the camera parameters.

Table 2.4: Patient registration results for surgery 1, 2, 3

	$Num.of Fidus$	FRE (mm)	TRE (mm)
Surgery 1	10	2.17	1.67
Surgery 2	8	1.97	1.54
Surgery 3	7	2.32	1.87

and segmented tumor boundary displayed through the red, green, and blue channels, respectively.

We employed the second image display strategy, Plane Intersection Overlay, to display the projection of the coregistered pMR volume with the current surgical scene during the second and third clinical cases. During the second surgery, an image of the surgical scene under the microscope was acquired right after the dural opening, as shown in Figure 2.8(a). Using the second image display technique, all the voxels in the coregistered MR volume intersected by the focal plane were projected onto the image plane, as shown in Figure 2.8(b) and overlaid with the current surgical scene, as shown in Figure 2.8(c). Note that the curved cortical vessel aligned well in the overlay. With the microscope in the same position, all the voxels in the MR volume intersected by a plane parallel with the focal plane but 8mm deeper into the brain were projected onto the image plane, as shown in Figure 2.8(e) and overlaid with the current surgical scene , as presented in Figure 2.8(f). This overlay showed the relative position of the tumor seen in the MR image with respect to the current surgical scene.

In the third surgery, an image of the surgical scene under the microscope was acquired right after the dural opening, as shown in Figure 2.9(a). All voxels in the coregistered MR volume intersected by the focal plane were again projected onto the image plane (see Figure 2.9(b)) and overlaid with the current surgical scene (see Figure 2.9(c)). Note that the bifurcating cortical vessel aligned well in the overlay. With the microscope in the same position, all the voxels in the MR volume intersected by a plane parallel with the focal plane but 16mm deeper into the brain were projected onto the image plane, as shown in Figure 2.9(e), and overlaid with the current surgical scene, as presented in

Figure 2.9(f). This overlay showed the relative position of the tumor seen in MR image with respect to the current surgical scene. The three cases shown here demonstrated the navigational value of the registration between the 3-D pMR volume and 2-D surgical scene.

Finally, while the calibration of this system is relatively straightforward, a number of factors conspire to corrupt the overall process. These factors include, for example, motion of the microscope between calibration and surgery, the long duration of the surgery (approximately 5 hours), or inadvertent jostling of the Polaris. An adaptive procedure is needed to respond to accruing registration error [79]. We have developed a real-time correction scheme to maintain the accuracy of the rigid registration throughout surgery. Interested readers can refer to Appendix B for details.

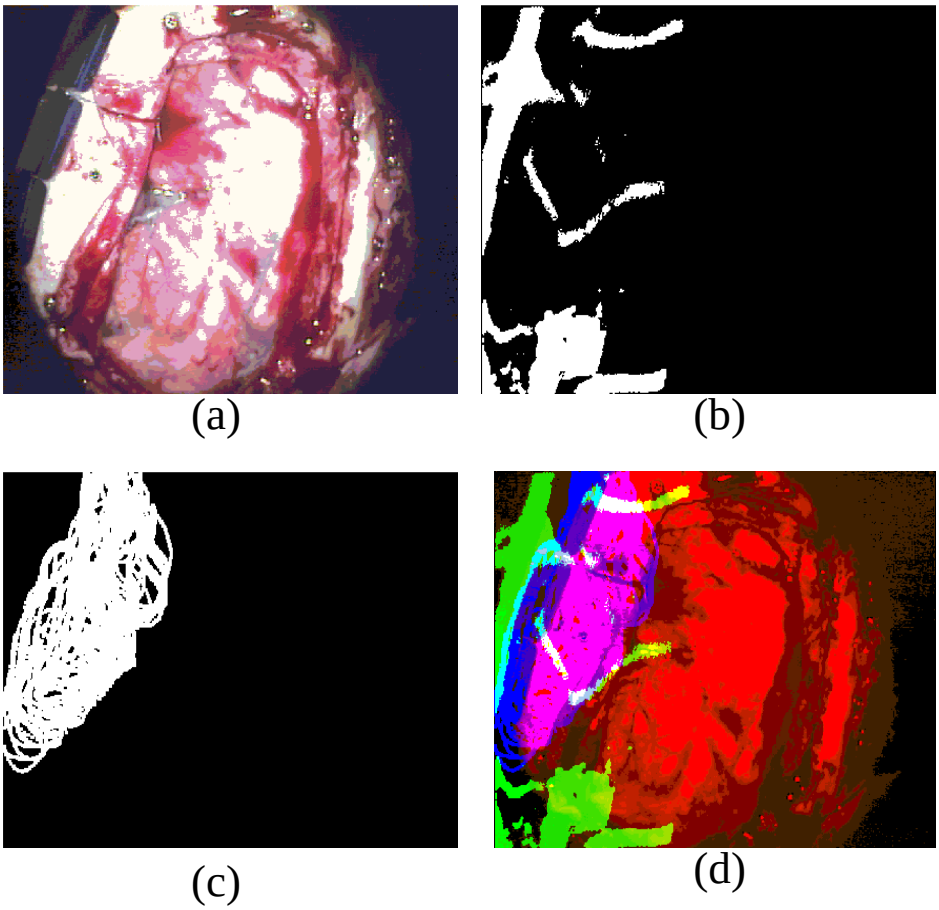

Figure 2.7: Salient structures segmented from co-registered MRA and MRI (surgery 1) were projected to microscope image planes: Image (a) is the surgical scene under the microscope; image (b) is the projected image of the segmented brain surface vasculature from a co-registered phase-contrast MR angiogram; image (c) is the projected image of all tumor boundaries outlined in all slices in the co-registered pMR; the RGB image (d) is the overlay of images (a), (b), and (c), in the red, green, and blue channels, respectively.

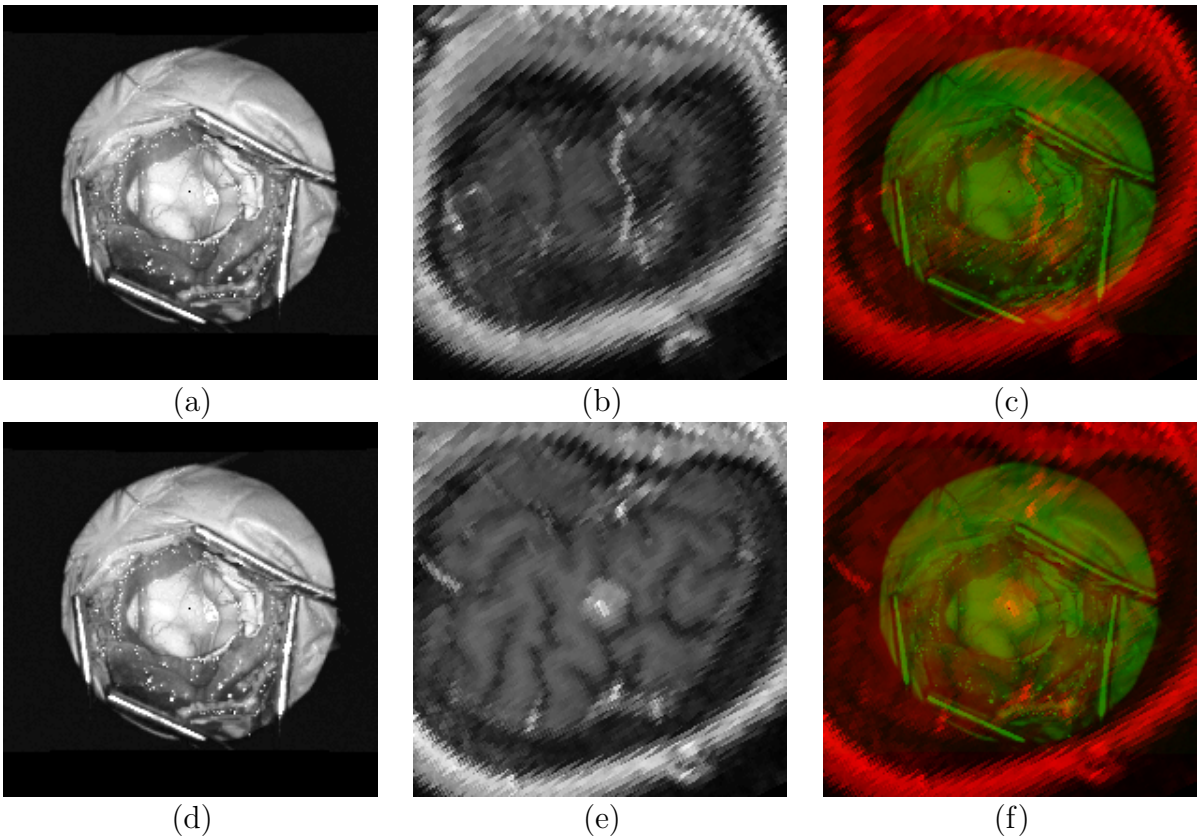

Figure 2.8: Shown in the first row are (a) the surgical scene (surgery 2), (b) the projected MR voxels intersected by the focal plane, and (c) the overlay of (a) and (b), in green and red, respectively; in the second row are (d) the same surgical scene , (e) the projected voxels intersected by a plane parallel with the focal plane but 8mm deeper into the brain intersecting the tumor, and (f) the overlay of (d) and (e).

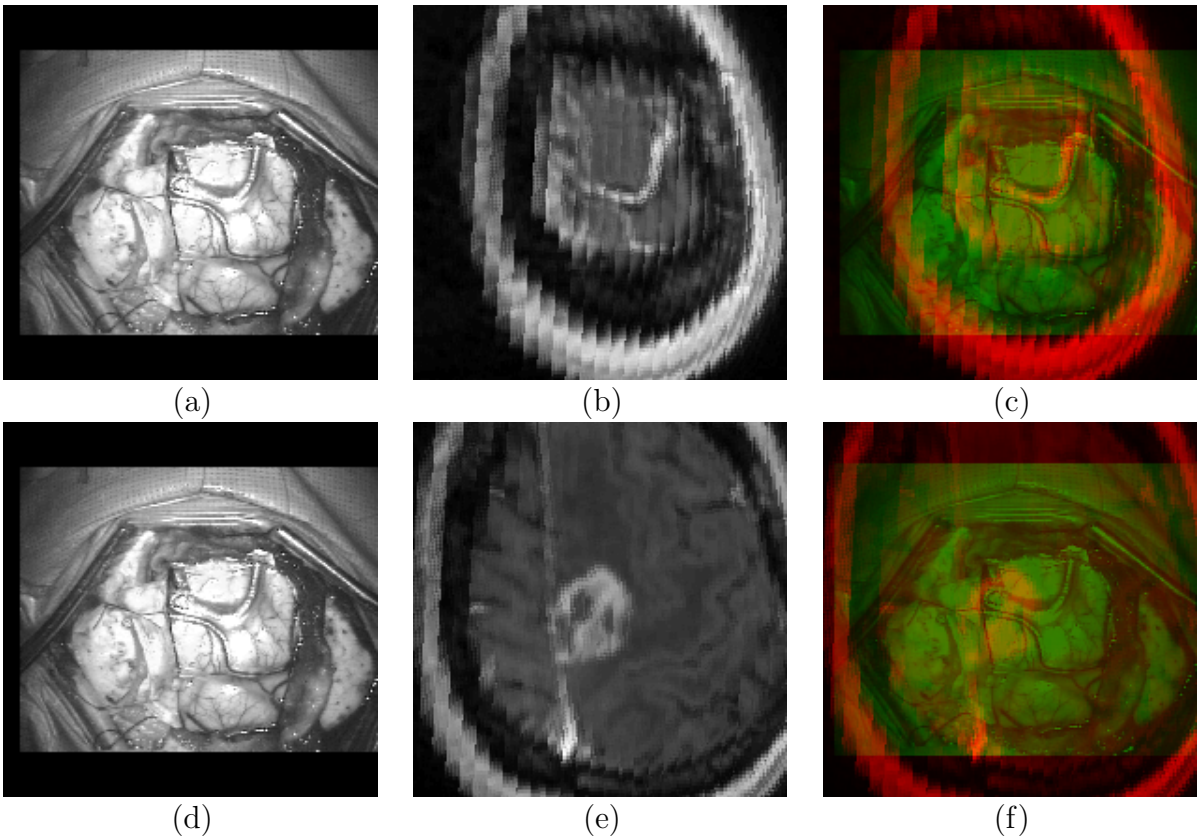

Figure 2.9: Shown in the first row are (a) the surgical scene (surgery 3), (b) the projected MR voxels intersected by the focal plane, and (c) the overlay of (a) and (b), in green and red, respectively; in the second row are (d) the same surgical scene , (e) the projected voxels intersected by a plane parallel with the focal plane but 16mm deeper into the brain intersecting the tumor, and (f) the overlay of (d) and (e).

Chapter 3

Stereopsis

So far, we have established the rigid registration between the 3-D pMR volume with the 2-D microscope images. This registration is soon subject to inaccuracy after craniotomy due to soft tissue deformation. The next step is to use the stereoscopic microscope to estimate the shape of the cortical surface during surgery, and in doing so prepare to track the cortical motion.

We are primarily interested in the implementation of a passive depth estimation method known as stereopsis. This technique enables us to infer information on the 3-D structure and distance of a scene from two or more images taken from different viewpoints. This is a depth estimation method with which we are all familiar, and one can learn a great deal regarding the basic principles of a stereo system through a simple experiment. Hold one thumb at arm's length and close the right and left eye alternatively. What do you expect to see? With presumably little surprise you find that the relative position of thumb and background appears to change, depending on which eye is open or closed. It is precisely this difference in retinal location that is used by the brain to reconstruct a 3-D representation of what we see.

3.1 Non-Invertibility of Image Formation

Before we discuss the method of stereopsis, it is useful to explain why a pair of cameras is needed for the depth estimation. As we have learned previously, image formation is a 3-D (world coordinates) to 2-D transformation (image coordinates). Inherent to such a transformation is a loss of information: the information lost is the distance to the objects in the world, i.e., range. Specifically, under a perspective projection model all points of the form (cX, cY, cZ), for any $c \in \mathcal{R}$, are projected to the same point (x, y). The projection is *not* one-to-one thus *not* invertible.

The non-invertibility of image formation can also be modeled via linear algebraic equations. In particular, we have seen that the image formation equations may be written in matrix form as, $\vec{q} = M_{n\times m}\vec{Q}$, where $n < m$ (Equations (2.5) and (2.10)). Since the projection is from a higher dimensional space to a lower dimensional space, the matrix M is not invertible and, thus, the projection is not invertible. Since image formation is not an one-to-one transformation and thus not invertible, assuming no prior information and a standard imaging system [1], the full 3-D structure of the world *cannot* be recovered from a single 2-D image.

The information lost in a 3-D to 2-D transformation can be recovered, however, using at least two images of the same scene taken from two different vantage points. Theories on stereopsis can be found in [21, 88] and [14] provides a general survey of some important technical implementations. Our goal is to use stereopsis to estimate the shape of the cortical surface during neurosurgery. To facilitate this process, a stereo pair is constructed using the stereoscopic operating microscope.

[1] There are many methods of recovering the lost information using one camera plus other cues, see [14] for reviews.

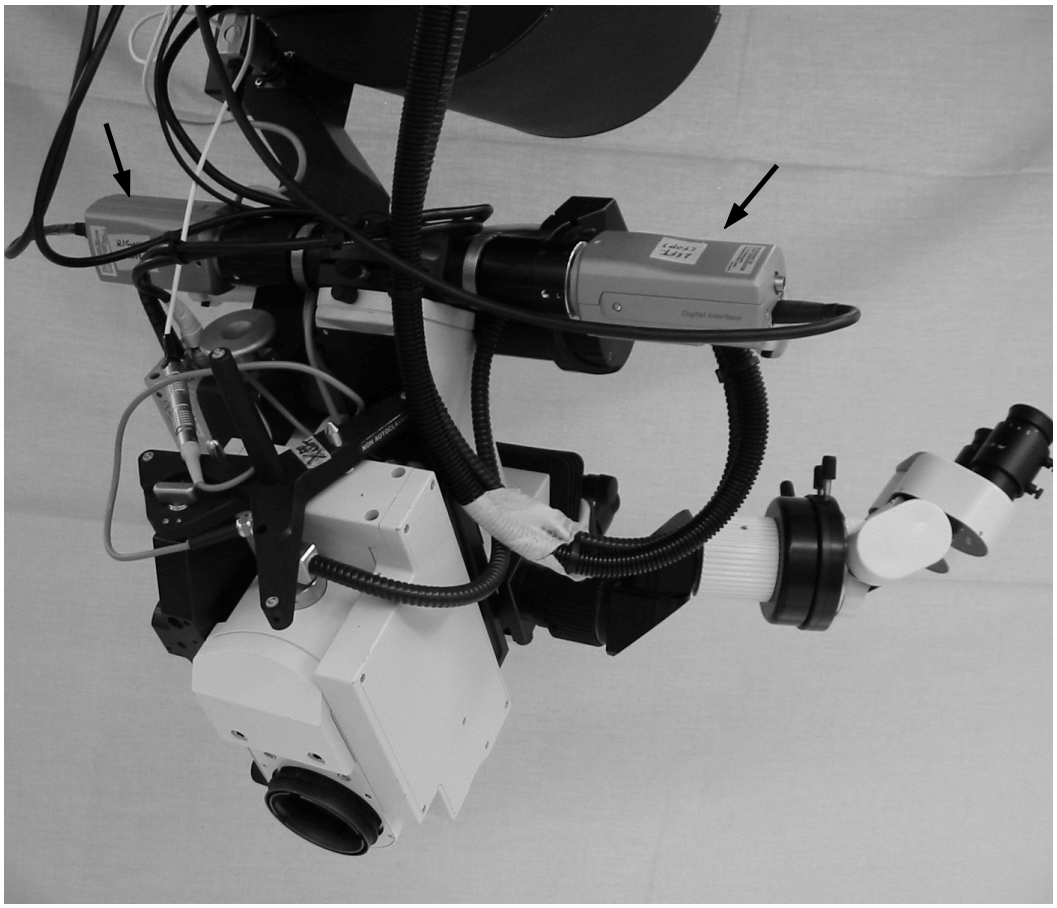

Figure 3.1: Two CCD cameras (black arrows) have been attached to the binocular optics of the operating microscope. The resulting stereo vision system is used to estimate the shape of the cortical surface during surgery.

3.2 The Stereo Pair

Shown in Figure 3.1 is an operating microscope to which two CCD cameras have been attached in order to estimate the shape of the cortical surface. A beamspliter has been added to the operating microscope to accommodate the two CCD cameras and to avoid obstructing the surgical team's visual access to the surgical field. Two optical couplers with a c-mount [2], as shown in Figure 3.2, are used to connect the CCD cameras to the left and right beamsplitter ports. The two cameras are controlled by a PC via a pair

[2]Microstereopsis, Ijamsville, Maryland, USA

Figure 3.2: Shown are the CCD camera (left) and the optical coupler with c-mount (right) used to connect the camera to the microscope beamsplitter port.

of firewire cables. Stereo image pairs are acquired at 15 frames per second. The stereo images can be displayed on a computer monitor and fused using shuttered glasses to provide a real-time 3-D view of the surgical scene, as illustrated in Figure 3.3.

3.3 Stereo Matching

Having constructed an intraoperative stereo vision (iSV) system by attaching two CCD cameras to the stereoscopic operating microscope, the next task is to estimate the shape of the cortical surface using stereo images. From the thumb experiment, we have learned that our vivid 3-D perception of the world is due to the interpretation that the brain gives of the computed difference in retinal position, named disparity, between corresponding items. Following the wisdom of Nature, the basic concept of stereo vision is illustrated in the left panel of Figure 3.4: a point (X, Y, Z) in the world reference frame is projected, through two pinhole cameras centered at C_l and C_r, onto the image sensors at locations

Figure 3.3: The surgical team can view the operating scene in 3-D through the shuttered glasses.

(x_l, y_l) and (x_r, y_r). Given these image coordinates, the goal is to estimate the location of the point (X, Y, Z). Beginning with an image projection onto, say, the left image, (x_l, y_l), we must first compute the corresponding projection onto the right image (x_r, y_r). This is the first step in stereopsis, which is often referred to as stereo matching. The displacement between the corresponding points, (x_l, y_l) and (x_r, y_r), is termed disparity. Combining the camera calibration (described above) with the established correspondence yields the desired 3-D coordinates. This is the second step in stereopsis, which is called triangulation.

It turns out that the first step of stereopsis is more difficult because the correspondence problem is ambiguous. Given a point (x_l, y_l) in the left image, it may *a priori* be put in correspondence with any point in the right image. To overcome this difficulty, we must constrain the number of potential matches for any given point. Literature regarding stereo matching is immense. Two classic matching algorithms have been presented by Marr and Poggio [53, 54]. Examples of some other significant matching strategies can

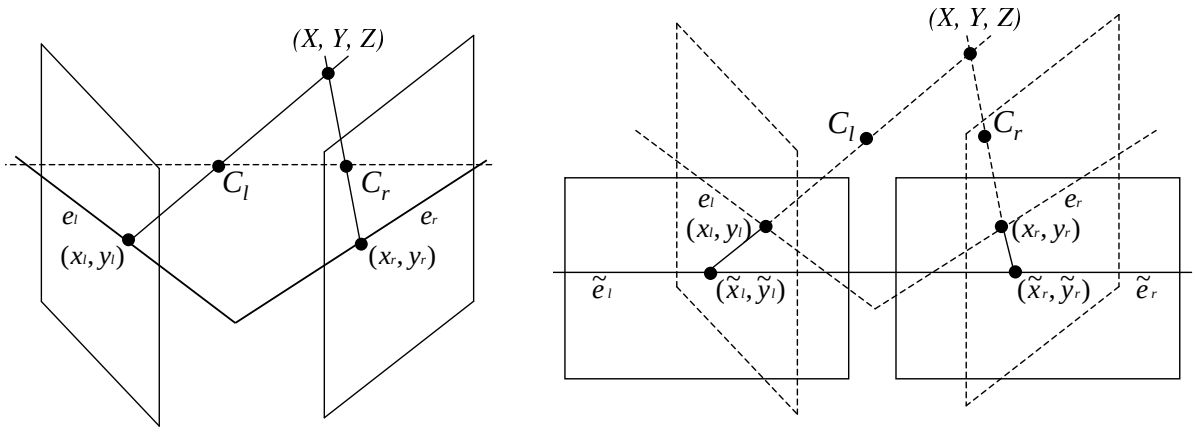

Figure 3.4: The 3-D vision problem and the epipolar constraint: shown on the left are two pinhole cameras, with centers C_l and C_r, forming the image of a single point. The projection onto the left image, (x_l, y_l), has a corresponding match in the right image, (x_r, y_r), lying on the epipolar line e_r. Shown on the right is the rectified image so that the epipolar lines are coincident with image scan lines (solid lines).

be found in [3, 42, 47, 58]. Building on this literature, we have implemented a number of matching constraints tailored to the estimation of the shape of the cortical surface.

1. Epipolar Constraint

   This constraint is derived from the fact that the stereo camera pair is calibrated with respect to a common coordinate system and therefore the camera positions are known relative to one another. As a result, the corresponding point (x_r, y_r) of the point (x_l, y_l) is constrained to lie on the epipolar line. This line, e_r, is defined to be the projection of the optical ray $\langle (x_l, y_l), C_l \rangle$ through the camera center C_r, as shown in Figure 3.4. The epipolar constraint is symmetric and hence the corresponding point (x_l, y_l) of the point (x_r, y_r) is constrained to lie on the epipolar line, e_l, which is constructed similarly. This epipolar constraint restricts correspondence to a single line rather than an entire image.

   The epipolar line is often oblique, and it is more convenient to rectify the images so that the epipolar lines are coincident with image raster lines (as shown in the right panel of Figure 3.4) – the rectified image coordinates are denoted as $(\tilde{x}_l, \tilde{y}_l)$

and $(\tilde{x}_r, \tilde{y}_r)$.

The basic idea of rectification is straightforward: if the left and right image planes are coplanar and the horizontal axes are collinear (no rotation about the optical axes), then the image rows are epipolar lines and stereo correspondences can be found by searching for matches along corresponding rows. Rectification is discussed by Ayache [2], Hartley et al. [33], and Faugeras [21]. Camera parameters are not required for rectification [69], but the use of calibrated stereo pairs gives rise to a simpler technique.

Given a pair of stereo images and a calibrated stereo pair, rectification determines a transformation of each image plane such that pairs of epipolar lines become collinear and parallel to one of the image axes, as seen in Figure 3.4 (right). For each camera, this transformation takes the form:

$$\begin{pmatrix} \tilde{x_s} \\ \tilde{y_s} \\ s \end{pmatrix} = R \begin{pmatrix} x \\ y \\ 1 \end{pmatrix}, \tag{3.1}$$

where the coordinates in the rectified image are $(\tilde{x}, \tilde{y} = (\tilde{x_s}/s, \tilde{y_s}/s)$, and (x, y) are the coordinates in the original image pair. When the two cameras are rigidly attached to each other and are calibrated with respect to a common coordinate frame, R is a 3×3 matrix and:

$$R = P_r P^{-1}, \tag{3.2}$$

where P is the PPM of a camera in Equation (2.5) and is obtained via camera calibration. P_r is a 3×4 matrix similar to P. Since it projects a point in the world coordinate, (X, Y, Z), to a point in the rectified image plane, $(\tilde{x}, \tilde{y})$:

$$\begin{pmatrix} \tilde{x_s} \\ \tilde{y_s} \\ s \end{pmatrix} = P_r \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix}, \tag{3.3}$$

P_r is often referred to as the camera's rectifying PPM.

Knowing P from camera calibration, we can determine P_r by applying several constraints imposed by the characteristics of the desired rectification: (1) the rectified image pair must share the same focal plane; (2) the camera centers of the rectified perspective projections must be the same as those of the original projections; and (3) all pairs of epipolar lines are parallel to one another and parallel to one of the pixel rows (usually the horizontal one). The mathematical derivation of $\tilde{P}_r$ can be found in Appendix C. With P_r and P, R can be computed using Equation (3.2).

Given all the camera and rectification parameters, the rectification of the original image pair consists of three steps. First, the projected images coordinates, (x, y), are computed from the original (distorted) image coordinates, (x_d, y_d), using Equation (2.7). Secondly, the rectified image coordinates, (x_r, y_r), are computed from the projected image coordinates, (x, y), using Equation (3.1). Finally, to create the final rectified image pair, the rectified image coordinates, (x_r, y_r), are used to linearly interpolate the same sampling lattice of the original image pair.

2. Intensity Correlation Constraint

   Given the rectified stereo pair of images, an intensity-based correlation metric is used to solve the correspondence problem. In the correlation-based matching strategy, the elements matched are image windows of certain size, and the similarity criterion is a measure of the correlation between windows in two images. The corresponding element is given by the window that maximizes the similarity criterion within a search region. Clearly, only those features which satisfy the epipolar constraint are considered for matching.

   Specifically, the horizontal displacement (or disparity) between corresponding points $(\tilde{x}_l, \tilde{y}_l)$ and $(\tilde{x}_r, \tilde{y}_r)$ is given by a scalar value d (there is no vertical displacement due to the rectification). In order to estimate this disparity, the following error

function is minimized:

$$E(d) = \sum_{i=-W}^{W} \sum_{j=-W}^{W} [I_l(\tilde{x}_l + i, \tilde{y}_l + j) - I_r(\tilde{x}_l + d + i, \tilde{y}_l + j)]^2, \tag{3.4}$$

where $I_l(\cdot)$ and $I_r(\cdot)$ are the left and right images, and W is a parameter defining the size of the correlation neighborhood. Kanade and Okutomi describe a method for adapting the shape and size of the correlation neighborhood to accommodate different parts of the same image [43]. Give the size of our rectified stereo pairs (approximately 1400×800), we use a square window size of W equal to 20 pixels. This non-linear error function, $E(d)$, is minimized by using a search strategy over a fixed range of integer-valued disparities. The disparity is estimated with a sub-pixel resolution by fitting $E(d)$ with a second-order polynomial and analytically computing the minimum. This entire process is repeated for all, or for a subset, of the image pixels.

To reduce computational costs, we constrain the disparity search to a relatively small region as follows. The disparities are computed for 50 equally spaced image pixels in the left image – a large search region of 100 pixels is used. The search region for the entire image is then constrained to consist of the pixel with the median of these disparities plus its nearest 20 neighboring pixels on the epipolar line, thus greatly reducing the computational effort.

3. Smoothness Constraint

In most craniotomy sites, the depths of scene points corresponding to nearby feature points are likely to be close to one another. We can impose a smoothness constraint on our estimated disparities in order to exploit this heuristic. The disparities of all image points are used to form a so-called disparity map, to which a 3×3 median filter is applied to the estimated disparity map. The median filtering has the effect of forcing pixels with distinct disparities to be more like their neighbors, thus it helps to eliminate spurious disparities, and is generally reasonable given the smoothly varying nature of cortical surfaces.

Given a pair of images of a surgical scene, we avoid computing the correspondence between all pixels. Instead, only a subset of the image pixels were selected for matching. Specifically, the boundary of the cortical surface was outlined, and approximately 1000 uniformly spaced pixels [3] within this region were chosen as matching tokens. This strategy offers two advantages: first, since the craniotomy site is usually centered in the microscope view and visible by both cameras, the matching algorithm need not determine the image parts that should *not* be matched; second, computational cost is greatly reduced with fewer matching tokens. One concern when estimating the shape of the cortical surface from only a subset of pixels is that the resulting surface will be spatially coarse. This concern will be addressed later with experimental data.

3.4 Triangulation

Given an estimated disparity map, triangulation is used to compute the 3-D coordinates, (X, Y, Z), from each pair of matched image points, $(\tilde{x}_l, \tilde{y}_l)$ and $(\tilde{x}_r, \tilde{y}_r)$, as shown in Figure 3.4. First, these image coordinates are transformed back into the original image planes, (x_l, y_l) and (x_r, y_r). Then, these image coordinates and the camera centers, C_l and C_r, are transformed from their 2-D image space to 3-D camera space using the intrinsic camera parameters. The extrinsic parameters map these 3-D coordinates into a common world coordinate system. The set of four 3-D coordinates, defining two optical rays, are used to estimate the final 3-D coordinates. In the absence of noise, this is performed by simply finding the intersection of the rays. In the presence of noise, however, the rays may not intersect and therefore the estimated intersection is taken to be the point of minimal and equal distance from both rays. Repeating this process for each pair of matched coordinates yields an estimated 3-D cortical structure. It is worth pointing out that it is *not* essential to have both the intrinsic and extrinsic parameters of the stereo system for triangulation. With only the intrinsic parameters of the cameras,

[3] Smoothness constraint is applied to the downsampled disparity map.

the 3-D coordinates of an object in the world can be reconstructed up to a scale factor. In our case, the exact location of the cortical surface in the world coordinates is important for registering this surface with the pMR volume, as we will see later.

Next, we present our experimental results in the estimation of the shape of the cortical surface. We begin with a simulation probing the effect of the matching results on the accuracy of the final reconstructed shape. Then, the reconstructed shapes of the cortical surfaces from eleven clinical cases are given. Finally, to demonstrate the accuracy of the estimated cortical shape, error analysis from both phantom and clinical studies is presented.

3.5 Results

The previous section demonstrates that estimating the 3-D surface from a stereo pair of images involves two main steps: (1) for each point in the left stereo image, the corresponding point in the right image must be found (i.e., stereo matching), and (2) from this correspondence, triangulation is then used to determine the 3-D coordinates of the object.

3.5.1 Sensitivity to Matching

In order to quantify the performance of the first step of stereopsis, i.e., stereo matching, a numerical simulation was designed and executed to probe the error in stereo range estimation as a result of error in solving the correspondence problem. We used stereo images of the camera calibration target, as shown in Figure 2.4, for the simulation. The target was tracked, so that the 3-D location of the centroids of each marker was known. At the beginning of the experiment, the camera was calibrated. The camera calibration results and microscope tracking data were used to rectify the stereo image pairs. The centroids of markers ($n = 108$) in both the left and right images were

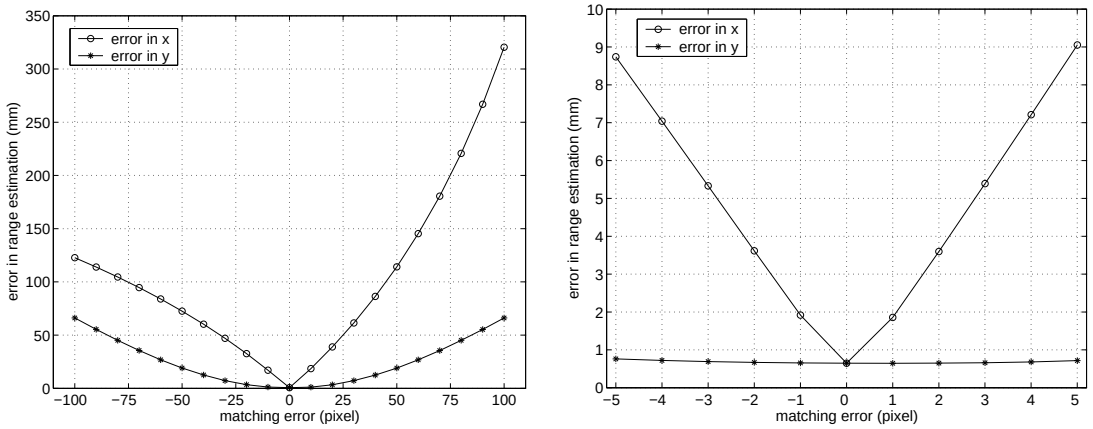

Figure 3.5: The 3-D coordinates of centroids of markers on the calibration target were first computed based on the manually estimated matching results. Errors in either the x (the directions along the epipolar line) or y (the directions perpendicular to the epipolar line) image axis were subsequently added to this best case matching results. RMS errors of range estimation based on matching with and without added errors were computed by comparing the estimated 3-D coordinates with the locations of these centroids recorded by the Polaris. The plot on the right is a subplot of the one on the left. Note that a matching error in x is more significant in causing an error increase in range estimation than the corresponding error in y.

estimated via segmentation, and each corresponding pair was matched manually. Then the triangulation produced the 3-D coordinates of these centroids, and the RMS error was estimated by comparing the results with the locations of these centroids recorded by the Polaris system. Since the homologous points were estimated via segmentation and matched manually, we consider that there is no matching error so far. This minimal RMS error of range estimation (0.65mm) was determined solely by errors from camera calibration and tracking. Next, we simulated matching error (in pixel) in either the x (the directions along the epipolar line) or y (the directions perpendicular to the epipolar line) image axis by altering the estimated locations of centroids in the right image. The RMS error of the range estimation based on altered homologous pairs was computed. Figure 3.5 shows the relationships between RMS errors in range estimation and simulated matching errors. The plot on the left, based on simulations over a bigger range of matching error values, depicts general quadratic relationships, and whereas the plot on the right, a subplot of the plot of the left, shows more details in the neighborhood of

the true correspondences.

These simulated results have revealed that a matching error in x is more significant in causing an increase of error in range estimation than the same error in y. Our interpretation of this result is that in order to achieve more accurate range estimation, matching constraints must be able to solve correspondence with subpixel accuracy in the x direction; while in the y direction, the performance of these constraints can be allowed some tolerance. In order to achieve this requirement, two of the three matching constraints that were implemented, i.e., the intensity correlation constraint and smoothness constraint, were designed to reinforce the matching accuracy in the x direction. Due to the tolerance of the matching error in y, the search for a potential match was generally widened to include the two lines of pixels immediately adjacent to the epipolar lines to counter potential matching errors in the x direction that may be a result of inaccurate rectification.

3.5.2 Clinical Cases

We next present the results of estimating the shape of the cortical surface using stereopsis. Most of these results have been discussed in a paper submitted to Neurosugery [83]. An exemplary image pair from a clinical case is used to illustrate both the incremental results from different steps involved in this estimation and the final estimated shape of the cortical surface. The final results from nine other clinical cases are then used to demonstrate the success of the iSV system in capturing a variety of shapes of the cortical surfaces. The characteristics of these ten patients are listed in Table 3.1.

Our exemplary clinic case involved a 48-year-old male patient with intractable epilepsy. During surgery, craniotomy was performed at the left anterior temporal region of the skull for the placement of electrode implants. The reconstruction of the surgical scene following craniotomy began with the acquisition of a stereo pair, as shown in Figure 3.6(top). The RGB image pair was first converted to a grayscale image pair, and

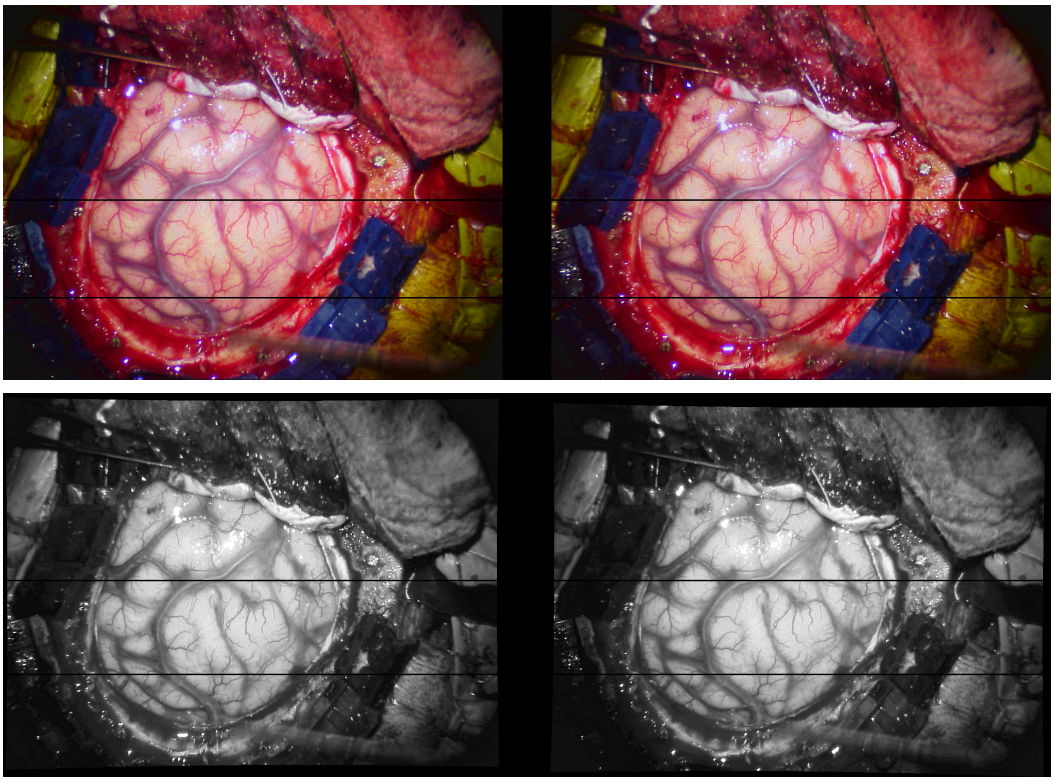

Figure 3.6: Left and right images of a typical early surgical scene (top) and the rectified image pair (bottom). Note that the homologous features from both left and right images align with the same pixel row after rectification.

rectification was applied using the method described in Appendix C. As a result of rectification, epipolar lines became collinear and parallel to the image row, as shown in Figure 3.6(bottom).

On the left image of the rectified pair, the boundary of the craniotomy site was manually outlined and about 1000 evenly spaced pixels within this region were selected for matching. The corresponding points for these pixels were found by applying the Intensity Correlation Constraint, our second matching constraint. The disparity maps of these pixels are shown in the top row of Figure 3.7. A 5×5 median filter was then applied to the disparity maps to eliminate spurious disparities, as shown in the bottom row of Figure 3.7.

Combining the estimated correspondence with the camera parameters, triangulation

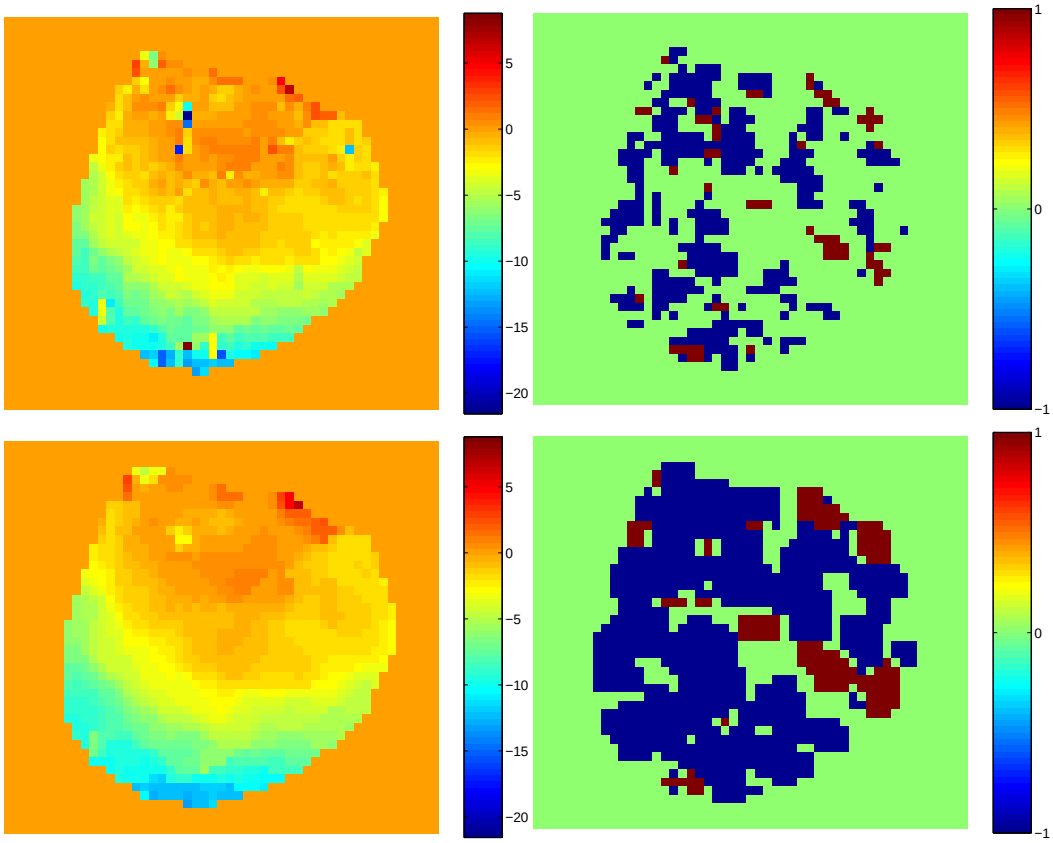

Figure 3.7: Shown in the top row are the disparity maps in the x direction (left) and y direction (right) before the smoothing using a median filter. The smoothed disparity maps are shown in the bottom row.

yields the desired 3-D shape of the cortical surface. The estimated 3-D surface is shown both in wire frame and in texture mapping [4] modes in Figure 3.8. On a 1.1 GHz PC, the reconstruction of a surface for a stereo image pair using 1000 matching tokens requires approximately 60 seconds of computation. The reconstructed cortical surfaces from nine other surgical cases are shown in Figure 3.9, and the relevant patient information is given in Table 3.1.

[4]For texture mapping, after interpolating the reconstructed surface onto the same sampling lattice as the left image of the stereo pair, each reconstructed point was assigned with its corresponding pixel intensity.

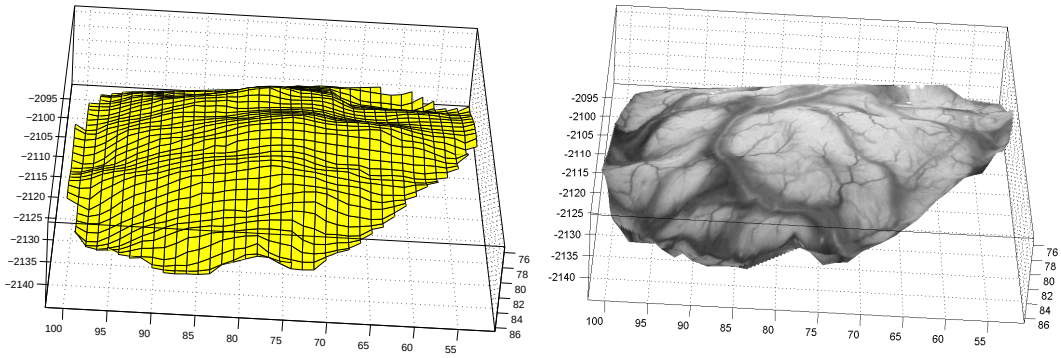

Figure 3.8: The reconstructed 3-D cortical surface (surgery 1), shown in both wire frame (left) and texture mapping mode (right) (units on axes are in millimeters; also see Figure 3.9).

3.5.3 Reconstruction Accuracy I

In order to quantify the accuracy of our iSV system, the localization of points in 3-D were compared to an optical stylus probe, a well quantified 3-D point digitizer. The stylus, as shown in Figure 3.10, is tracked by the Polaris system, which defines the world (OR) coordinates. The RMS tracking error of a single marker on a tool such as the optical stylus is reported by the manufacturer to be 0.35mm in either active or passive mode [66]. The accuracy of locating a known point using the stylus depends on the number of the markers on the stylus and the geometry of the stylus. We have tested the precision of our stylus. Our results show that, when returning repeatedly to a certain set of coordinates in space, the device consistently displays the same set of coordinates within 0.87mm.

In our first approach, we used the stylus probe to measure the 3-D location of several micro-screws [5] placed on the bony rim of the cranial opening as fiducial markers. These fiducials are denoted MS1,..., MS4 in Figure 3.11(left). The 3-D locations of the micro-screws were estimated using stereopsis; the center of each micro-screw was manually segmented in the microscope images and their 3-D positions were computed as the

[5]Leibinger E-Z Flap Plating System, Styker Cooperation, Kalamazoo, MI., USA

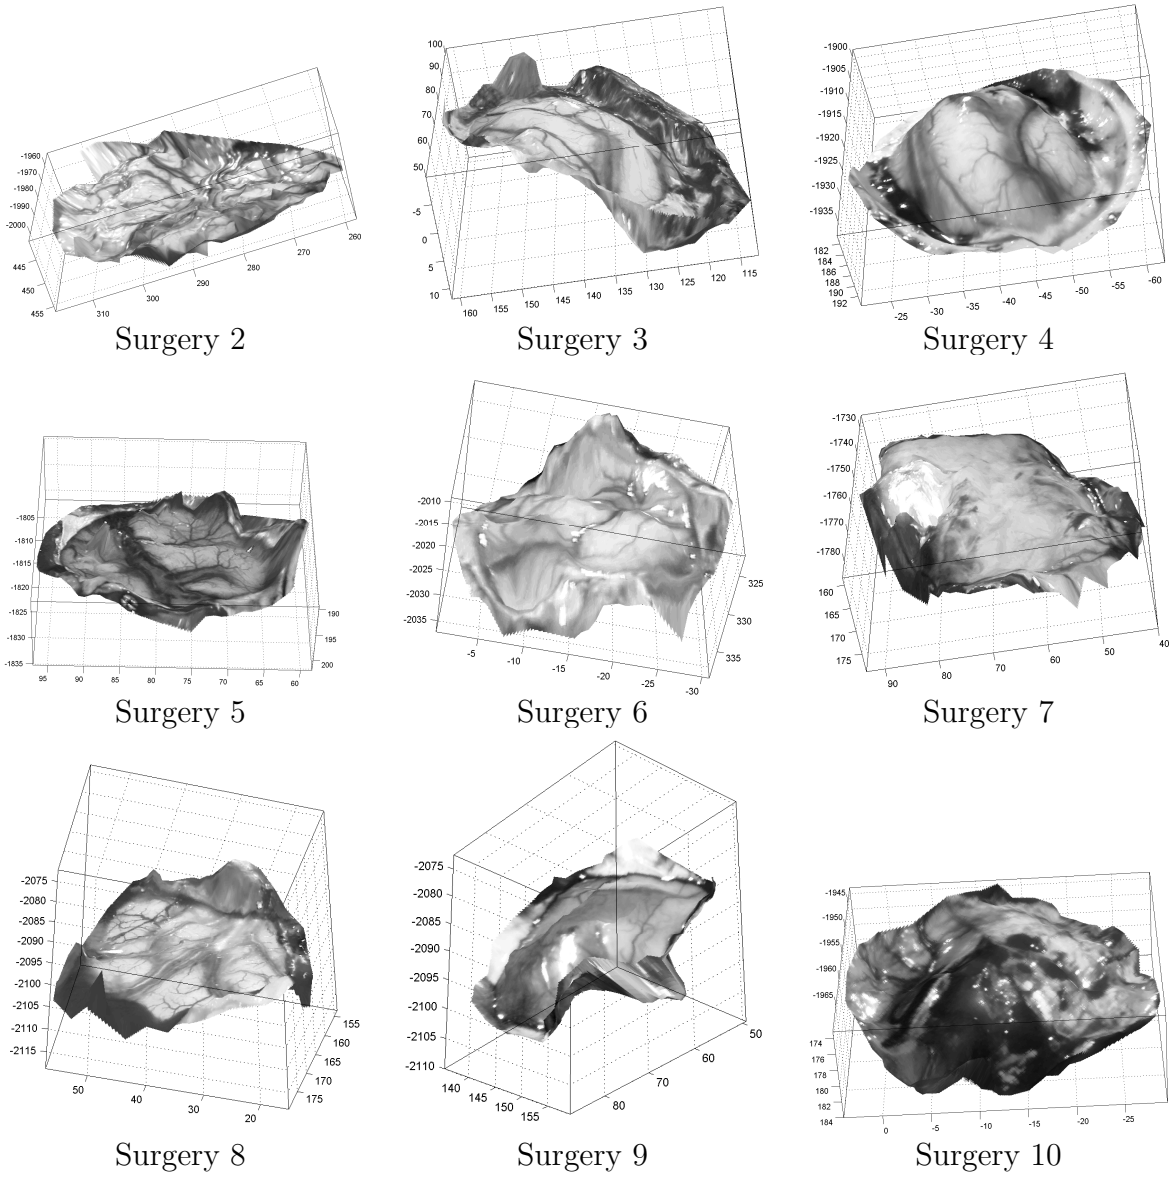

Figure 3.9: Initial cortical surface reconstructions with texture mapping from nine other craniotomies (units on axes are in millimeters; also see Figure 3.8).

Surgery No.	Age (yr)/ Gender	Lesion/Pathology	Location	Procedure
1	39/F	Intractable epilepsy	Left anterior temporal	Electrode implant
2	63/F	Hemangioma and right MCA aneurysm	Right temporal	Resection
3	62/M	Renal cell carcinoma	Right posterior temporal	Resection
4	50/M	Solitary metastasis	Left frontal	Resection
5	27/F	Seizure Disorder	Right medial temporal	Amygdala removal
6	39/M	Cystic Hamartoma	Left Temporal	Resection
7	52/M	Recurrent glioblastoma	Right Temporal	Resection
8	48/M	Intractable epilepsy	Left temporal	Lobectomy
9	44/F	Metastatic carcinoma	Left parietal	Resection
10	50/F	Meningioma	extra-axial parasagittal	Resection

Table 3.1: Patient Characteristics

mean of the estimated 3-D coordinates of a small neighborhood of pixels (about 15 pixels) centered about this position. Stylus acquisitions were then compared to their corresponding stereopsis estimates. During ten clinical cases, 38 fiducials were implanted and digitized using both methods. Each case employed an average of 3.8 fiducials, with a maximum of 5 and a minimum of 3. The RMS difference between these estimates with respect to the number of matching tokens used to reconstruct the surgical scene is shown in Figure 3.11(right). The absolute difference between these estimates decreases as the number of matching tokens increases from 100 to 600, then it levels off at around 1mm when the number of matching tokens is greater than 700. With 1000 matching tokens applied, the difference is 1.02mm with a standard error of 0.07mm. The maximum error is 1.35mm, with a minimum error of 0.61mm. This number of matching tokens is more than sufficient to capture the details of different shapes of the surgical scene during the

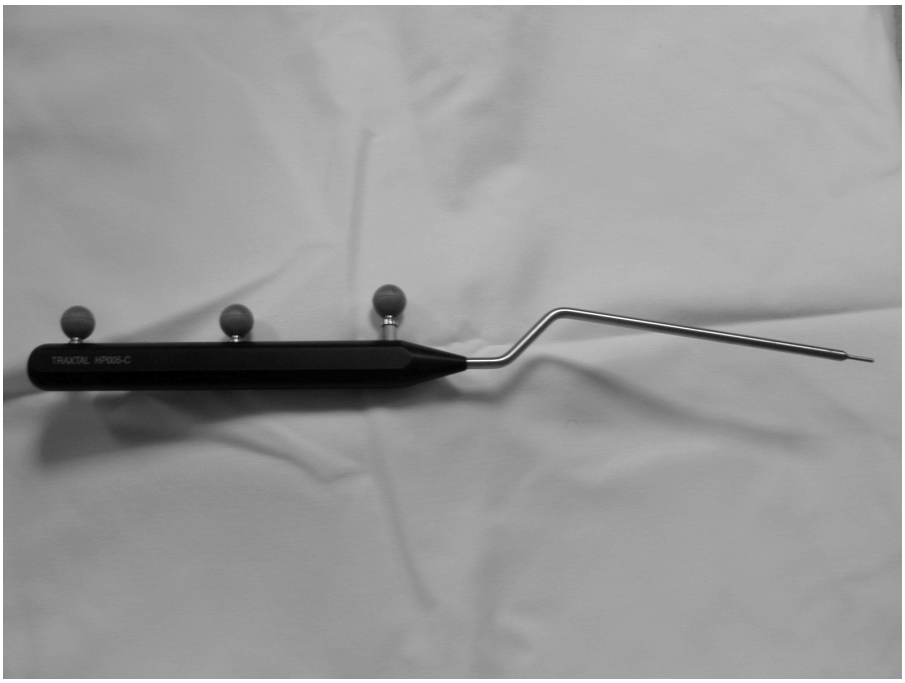

Figure 3.10: An optical stylus used to localize 3-D points during surgery.

course of surgery.

3.5.4 Reconstruction Accuracy II

One concern regarding the use of a micro-screw to simulate a point in 3-D is its finite size. Although it is small (1mm in diameter), the finite size of each screw may introduce uncertainty in whether the same point was digitized by both the stylus probe and the iSV system. In order to better simulate a point in 3-D, we constructed a planar geometrical pattern using a CAD drawing program and printed it on a 600 dpi laser printer, as shown in Figure 3.12. Within this pattern, the location of each circle, line, and bounding rectangle are analytically known. An optical stylus, tracked by the Polaris, was used to determine the 3-D coordinates of the four corners of the bounding rectangle. For each corner, the acquisition was repeated 100 times, and averaged to generate the 3-D position of each corner point. The intersection between the various geometric shapes

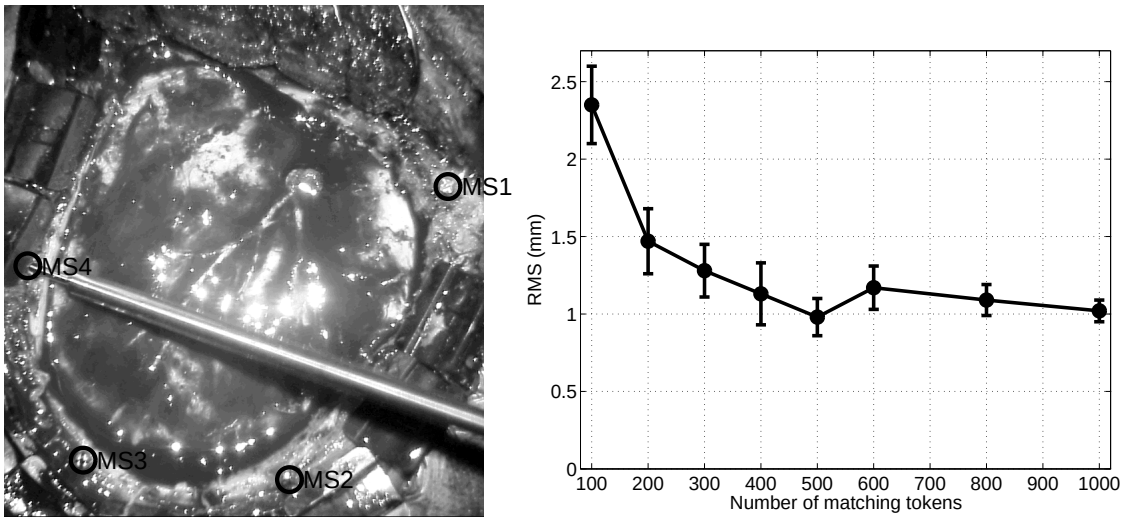

Figure 3.11: Shown on the left is a camera view of a craniotomy site with 4 implanted micro-screws (MS1,...,MS4) and the optical stylus used to digitize their locations. Shown on the right is the RMS difference between the measured positions of the micro-screws using the optical stylus and our iSV system as a function of the number of matching tokens involved in the stereoscopic surface reconstruction.

(lines and circles) were analytically determined from these coordinates.

The 3-D coordinates of these same intersections were also estimated using stereopsis. All intersection points were located in the left stereo image by manually fitting lines and circles to each geometric shape. The coordinates of each intersection in this image were computed from these analytic fits. The stereo correspondences and 3-D coordinates for each intersection were then estimated. From the 98 intersections, the RMS difference between the stereopsis and stylus estimates was 0.93mm with a standard deviation of 0.41mm. The maximum error was 1.61mm, with a minimum error of 0.20mm. The mean error of 0.93mm is close to our observed accuracy for the optical stylus of 0.87mm.

3.5.5 Reconstruction Accuracy III

So far, we have learned that the accuracy of our iSV system in localizing a point in 3-D is comparable to the conventional optical 3-D point digitizer. Since we are primarily interested in estimating the shape of the cortical surface, we need to probe the accuracy

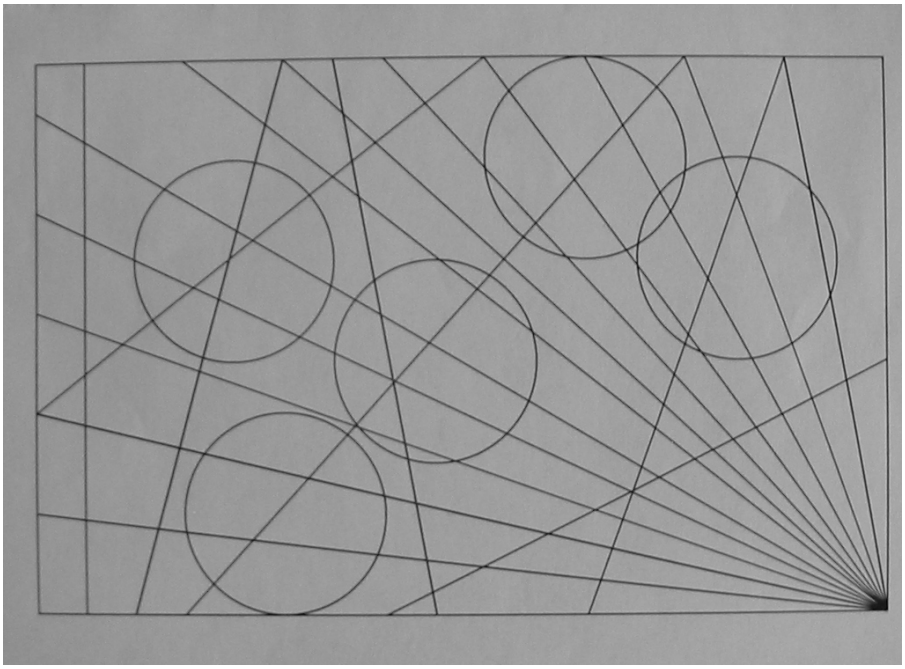

Figure 3.12: Pattern used to test the accuracy of the stereo reconstructions.

of our stereo system in shape estimation. In order to further confirm the viability of our approach, we compared the accuracy of our iSV system against a CT reconstruction. We first constructed a brain-shaped phantom from agar gel, as shown in Figure 3.13. Copper wires were embedded into the gel and placed close to the surface to simulate cortical vasculature. A CT scan of the phantom was then obtained, from which the shape of the cortical surface was manually computed. The phantom was then brought to the OR, and its surface estimated using stereopsis. Shown in Figure 3.14(a) are the stereo pair, and shown in Figure 3.14(b) is the estimated surface. Figure 3.14(c) reports the structure obtained from the CT scan whereas Figure 3.14(d) contains the difference between the stereopsis- and CT-based estimates. [6] The mean distance between these structures was 1.16mm with a maximum of 2.96mm and a minimum of 0.08mm.

[6]In order to compare the stereopsis- and CT-based estimates, we applied the ICP algorithm [5] to bring these two surfaces, which were previously in different coordinates systems, into alignment.

Figure 3.13: Shown is the brain phantom used to quantitatively evaluate the accuracy of the iSV system in capturing the shape of the cortical surface. Electric wires were embedded in the phantom to simulate cortical vasculature.

3.5.6 Reconstruction Accuracy IV

One concern regarding the estimation of the cortical surface from only a subset of pixels is that the resulting surface will be spatially coarse. To address this concern, we measured the density of the reconstructed surface by computing distances between all neighboring pairs of matching tokens. Table 3.2 contains the measured distances between neighboring pairs of matching tokens in ten reconstructed surfaces shown in both Figure 3.8 and Figure 3.9. From these ten estimations, the point clouds representing the final estimated

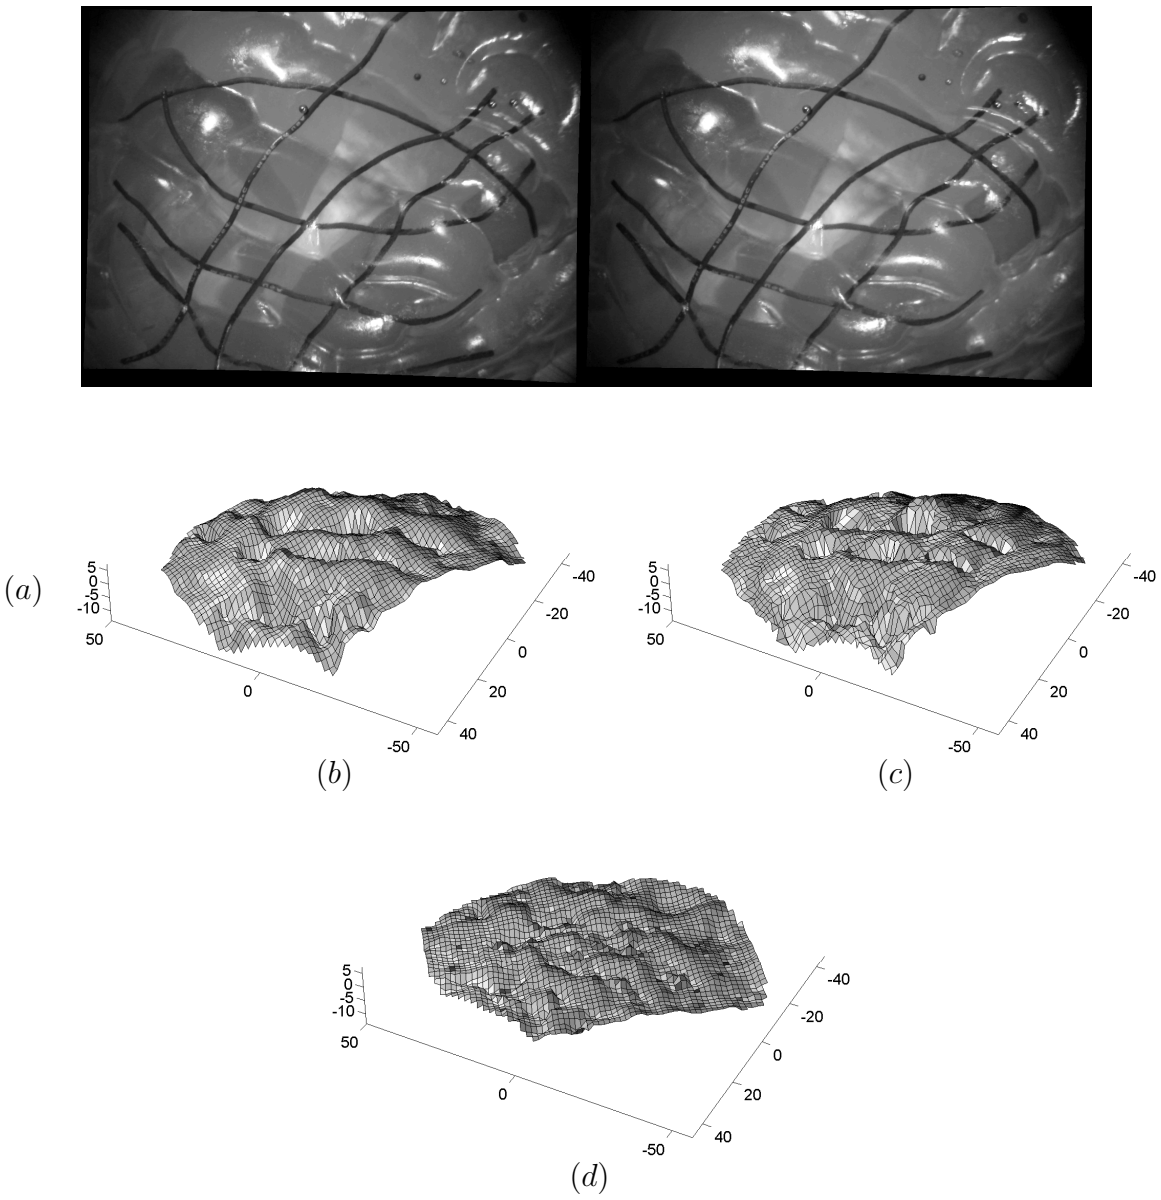

Figure 3.14: Shown are (a) rectified stereo image pair of the phantom brain; (b) the stereopsis-based estimate of the surface; (c) the CT-based estimate; (d) and the difference between the stereopsis- and CT-based estimates. Units are in millimeters.

Surgery	Size of Point Cloud	Mean (mm)	Min (mm)	Max (mm)
1	1038	1.44	1.17	2.17
2	1073	0.76	0.56	1.29
3	1023	1.56	1.12	2.08
4	1123	0.98	0.67	1.37
5	1067	1.25	0.98	1.87
6	1231	0.56	0.34	0.86
7	1114	1.34	0.95	1.73
8	1056	1.35	1.09	2.11
9	1221	0.71	0.39	1.22
10	1038	1.01	0.76	1.25
mean	1098	1.10	0.80	1.60

Table 3.2: Measured Distances between neighboring pairs of matching tokens in reconstructed surfaces

surfaces contained a mean of 1098 points. The mean distance between neighboring pairs of matching tokens was 1.1mm. Based on these studies, we found that approximately 1000 matching tokens is sufficient to capture the topology of the cortical surface.

In summary, we have implemented an intraoperative stereo vision (iSV) system capable of capturing the 3-D shape of the cortical surface in a variety of clinical cases. The final shape of the cortical surface is represented by a point cloud that consists of approximately 1000 points. Compared to an optical stylus, our iSV system is capable of localizing a point in 3-D with an average error of approximately 1mm. Compared to CT, our system is capable of capturing the shape of the cortical surface with an average error less than 1.2mm. Next, we show how to track the motion of the cortical surface in the OR.

Chapter 4

Cortical Tracking

As described in the previous chapter, the 3-D shape of the cortical surface can be estimated with sufficient accuracy using stereopsis during surgery. Our goal is to gain intraoperative data on the cortical motion to guide brain deformation models to recover the full volume description of tissue shift.

In our implementation, the tracking of the cortical surface can be divided into two phases: (1), in order to establish the baseline motion of the cortical surface, the first stereoscopically estimated cortical surface in the OR is compared to the cortical surface in the pMR volume; and (2), the cortical surface is continuously tracked in the OR by comparing the current shape of the cortical surface with the previously estimated shape.

4.1 Phase I: Cortical Motion Relative to pMR

In anticipation of using the cortical motion to guide a brain model for updating the pMR volume, it is important to establish the baseline motion of the cortical surface relative to the same surface in the pMR volume. In order to compare the cortical surface with the pMR volume we must (1) generate a finite element mesh from the pMR volume; (2) register the 2-D microscope image with the pMR mesh; (3) locate the exposed cortical

surface in the pMR mesh; and (4) compute the displacement between the 3-D cortical surface in the real-time surgical scene and the pMR volume.

Prior to surgery, a standard volumetric MRI (typically a $256 \times 256 \times 124$ matrix where the voxel size was $0.97 \times 0.97 \times 1.5mm^3$) is obtained for the patient using a 1.5 Tesla GE MR scanner. This MRI scan is then imported into the software ANALYZE [1] and a 3-D representation is constructed from the MR slices. The cranium is removed through a semi-automatic segmentation strategy implemented in ANALYZE. Based on segmented brain, a marching cubes algorithm [74] is used to generate a surface description of the brain consisting of triangular patches. To maintain an reasonable number of patches, appropriate levels of surface decimation are applied. The segmented brain is now rendered as a surface wire frame composed of surface nodes. A nodal averaging scheme is usually applied to the initial wire frame to smooth the surface representation. Based on the newly smoothed boundary description, a full 3-D tetrahedral volume mesh can be generated with designated areas of refinement. Details on the mesh generation are described by Sullivan et al. [78].

The coordinates of the finite element mesh (the preoperative space) are registered to the 3-D operating room and then the 2-D operating microscope (intraoperative space). This process requires a rigid-body transformation between the pMR volume and OR space, calibration of the microscope optics, and tracking of the patient and microscope throughout the surgery, details of which have been described in Chapter 2.

With the pMR mesh registered to the microscope image, we identify the portion of the mesh that corresponds to the exposed cortical surface after craniotomy. To do so, the surface nodes are first projected into the microscope image. To avoid the inherent ambiguity resulting from this 3-D to 2-D projection, the surface nodes from only the hemisphere closest to the craniotomy are projected. The boundary of the craniotomy site is manually outlined in the microscope image. The surface nodes that fall within

[1]The software ANALYZE was provided by the Mayo Foundation, Rochester, MN, USA.

this boundary are identified in the microscope image and hence in the original mesh - we will refer to this subset of nodes as the craniotomy nodes.

Given the preoperative and intraoperative cortical surfaces, we next calculate the displacement between these surfaces. The displacement would be easy to compute if the correspondence between the surfaces were known. Lacking this information, however, we employ the iterative closest point (ICP) algorithm [5] to simultaneously establish this correspondence and estimate the displacement. Given the two surfaces, each represented by the point clouds, the ICP algorithm iteratively determines the correspondence between point clouds and the registration matrix (rotation and translation) that yields the best alignment (in terms of RMS Euclidean distance). The resulting displacement is represented as a set of 3-D translation vectors between each craniotomy node from the pMR volume and its corresponding point on the estimated cortical surface.

4.2 Phase II: Intraoperative Tracking

Having established the baseline cortical motion relative to the pMR volume, we are able to continuously track the cortical motion in the OR by comparing the stereoscopically estimated cortical shapes. During surgery, the cortical surface is re-estimated approximately every minute. Although the shape of the craniotomy site, including the cortical surface and surrounding structures such as the bony skull, can be estimated using stereopsis, only the cortical surface itself is selected for tracking.

In order to quantify the amount of cortical motion over time, it is necessary to track the estimated cortical structure between times t and $t + \Delta t$. This process involves first estimating the cortical surface at time t and $t + \Delta t$. Then, for points on the surface at time t, the iterative closest point algorithm is employed to find corresponding points on the surface at time $t + \Delta t$ such that the overall Euclidean distances between the two surfaces, represented as respective point clouds, is minimized. The displacement is the

vector that connects the point being tracked to its corresponding point at the next time increment.

In the case of missing cortical surface due to resection, surface tracking is then performed in the reverse direction to avoid any erroneous tracking. Namely, for points on the surface at time $t + \Delta t$ (a smaller cortical surface after resection), the iterative closest point algorithm is used to find correspondence points on the surface at time t (a larger cortical surface prior to the resection). The tracking can be executed with as fine of a temporal resolution as desired (subject to image acquisition limits and the computational load).

4.3 Results

The previous section explains that the tracking of the cortical surface is conducted in two phases. In phase I, in order to establish its baseline motion, the stereoscopically estimated shape of the cortical surface is compared with the same surface on the preoperative FE mesh. In phase II, the cortical motion is continuously tracked by comparing the stereoscopically estimated cortical shape over time. In this section, we present results from tracking cortical surfaces during three clinical cases. The accuracy of our tracking algorithm is then compared with that of tracking the features on the cortical surface using an optical stylus.

4.3.1 Phase I

We have selected three neurosurgical procedures (surgeries 2, 3, and 4) from the ten cases listed in Table 3.1 to illustrate the phase I cortical tracking. As we have stated earlier, this stage of cortical tracking is performed by comparing the stereoscopically estimated cortical surface to the cortical surface in the pMR volume.

The phase I tracking of the cortical surface in Surgery 2 began with the acquisition of

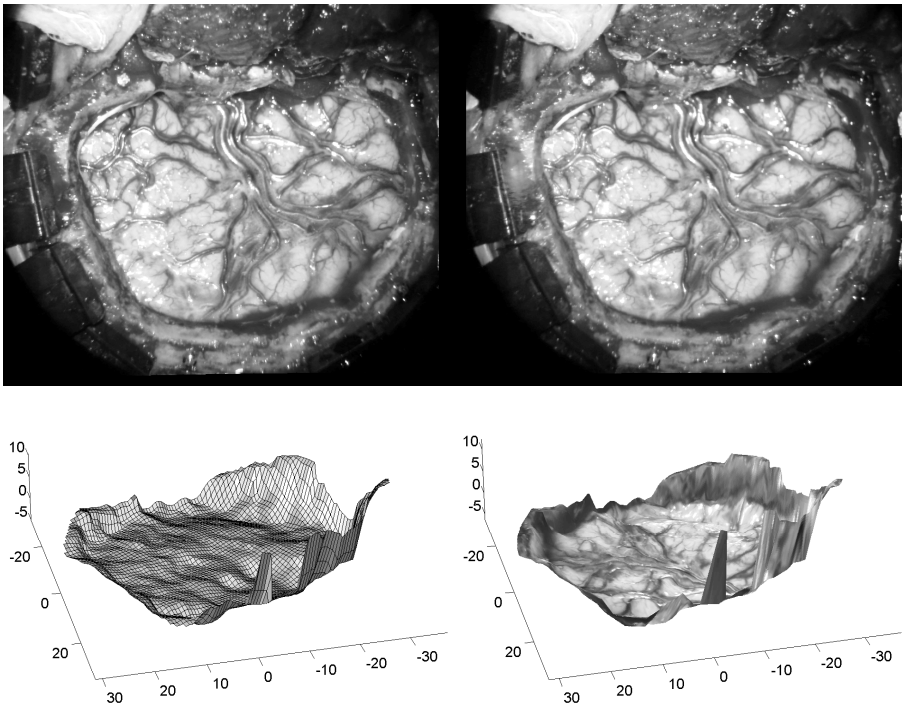

Figure 4.1: Rectified stereo pair (top row) and the estimated cortical surface (bottom row) presented in a wire frame view (left) and with texture mapping (right) – units are millimeters.

the stereo pair shown in the top row of Figure 4.1. Stereopsis was then used to estimate the shape of the cortical surface. Similar results for the third and fourth surgeries are presented in Figure 4.2 and Figure 4.3, respectively. During all surgeries, the bony rims of the craniotomy sites were also reconstructed to illustrate the motion of the cortical surfaces after dural opening. Note the collapsing cortical surface of the second patient and the distended surface of the third and fourth subjects.

The iSV-based estimates of the cortical surface, Figures 4.1, as shown in Figure 4.2, and 4.3, were compared with the pMR volumes to obtain intraoperative displacements. During each surgery, the coordinate space of the preoperative mesh was first rigidly registered to the 3-D operating room using scalp-attached fiducials. In the second surgery the fiducial and target registration errors (FRE/TRE) for 8 fiducials were 1.82mm and 1.45mm, respectively. For the third surgery, six fiducials had a FRE/TRE of 2.60mm and 1.92mm. For the fourth surgery, seven fiducials had a FRE/TRE of 2.14mm and 1.62mm.

Table 4.1: Cortical tracking results for surgery 1 and 2

Motion	Num. of Nodes	max (mm)	min (mm)	mean (mm)
Surgery 2				
Distension	62	5.64	0.12	3.56
Collapsing	112	8.66	1.57	4.76
Surgery 3				
Distension	44	9.03	2.95	6.29
Surgery 4				
Distension	48	17.82	10.45	14.35

Shown in Figure 4.4 is the overlay of the preoperative mesh with the surgical scene after registration – the craniotomy nodes are shown as asterisks, and their corresponding region in the preoperative mesh is shown in gray.

Figure 4.5 contains the results of overlaying the estimated cortical surface with the preoperative mesh. The result on the top left panel came from the second surgery – note that the anterior portion of the exposed cortical surface collapsed, while the posterior portion distended. Shown on the top right is the result from the third surgery, where the entire surface distended. The results from the fourth surgery are shown in the second row of the same figure. The intraoperative displacement was estimated using the ICP algorithm. For the second surgery, of the 174 craniotomy nodes, 112 collapsed and 62 distended. The average displacement of collapsing nodes was 4.76mm, with a maximum of 8.66mm and a minimum of 1.57mm. The average displacement of distending nodes was 3.54mm, with a maximum of 5.64mm and a minimum of 0.12mm. For the third surgery, all 44 craniotomy nodes distended. The average displacement was 6.29mm, with a maximum of 9.03mm and a minimum of 2.95mm. For the fourth surgery, all 48 craniotomy nodes distended, with an average magnitude of 14.35mm, a maximum of 17.82mm and a minimum of 10.45mm. These results are also presented in Table 4.1.

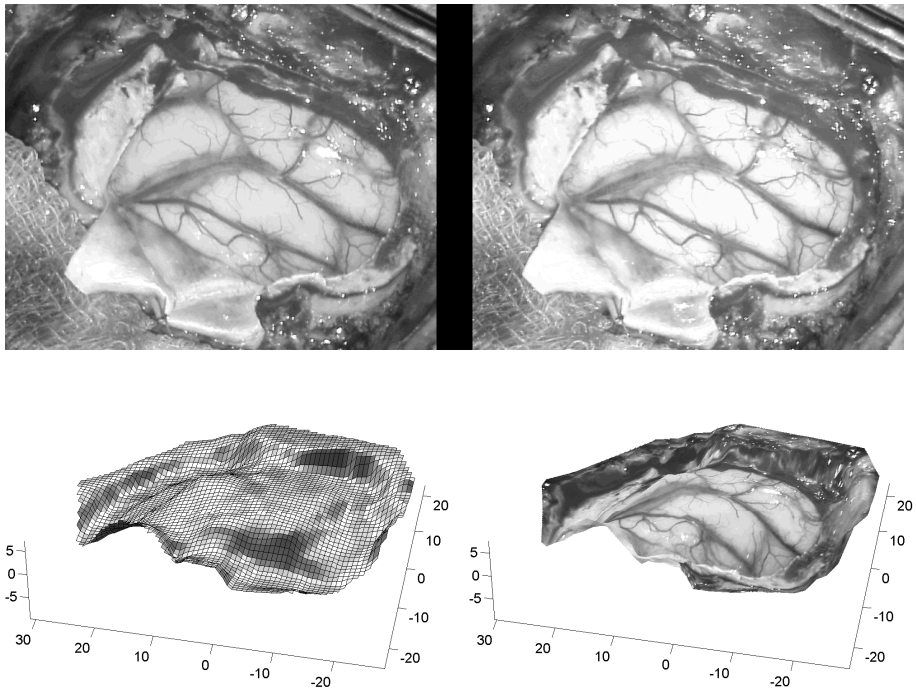

Figure 4.2: Rectified stereo pair (top row) and the estimated cortical surface (bottom row) presented in a wire frame view (left) and with texture mapping (right) – units are millimeters.

4.3.2 Phase II

We used one patient case (surgery 3 from Table 3.1) to illustrate how cortical surface motion is tracked during surgery and then present the results for the remaining nine cases. This surgical procedure involved a 62-year-old male with a cystic renal cell carcinoma metastasis in the right posterior temporal region of the brain. Stereo pairs of the cortical surface were acquired at three stages of the surgery: before and after the opening of the dura, and after partial cyst drainage. Three corresponding stereo pairs are shown temporally from top to bottom in Figure 4.6. Locations of the cortical surface were then estimated using stereopsis. Prior to surgery, a standard volumetric MRI was acquired for each patient. Using the estimated rigid registration between the OR space and the pMR volume, the reconstructed surfaces were overlaid with the pMR slices, as shown in Figure 4.7. Three colored lines were created by calculating the intersections between the reconstructed surfaces and MR slices. Green lines correspond to the surface

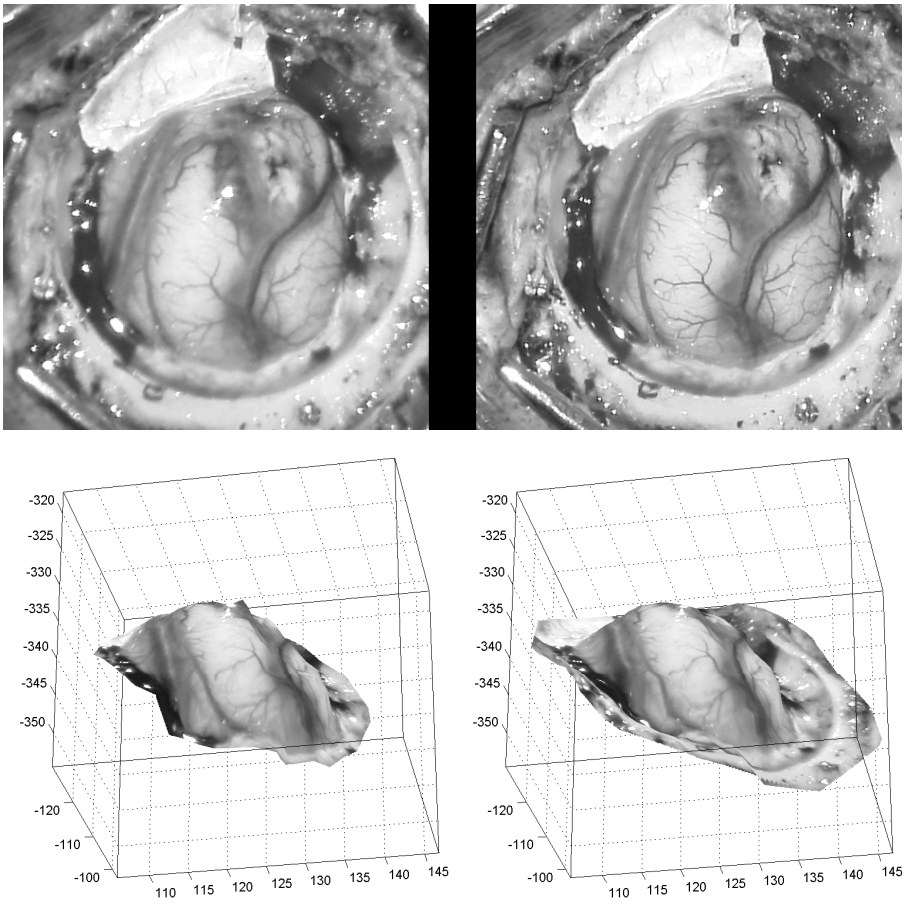

Figure 4.3: Shown are the rectified stereo pair of the current surgical scene in surgery 2(top); and the reconstructed cortical surface in 3-D, both with (bottom right) and without (bottom left) the surrounding bony skull. Units on axes are millimeters. Note the bulging motion of the cortical surface relative to the skull.

before the opening of the dura, yellow lines to the surface after the opening of the dura, and magenta lines to the surface after partial cyst drainage. To aid the visualization, the portion of the skull at the site of craniotomy was removed from the pMR volume based on the registered stereo view of the surgical opening. In Figure 4.7, prior to the removal of the dura, brain deformation is minimal (green lines) and the tracked cortical surface is well aligned with the rigidly registered pMR. The cortical surface (yellow lines) bulged after dural opening, due to elevated intracranial pressure. When the surgeon opened and partially drained the cyst cavity, the cortical surface (magenta lines) sagged from

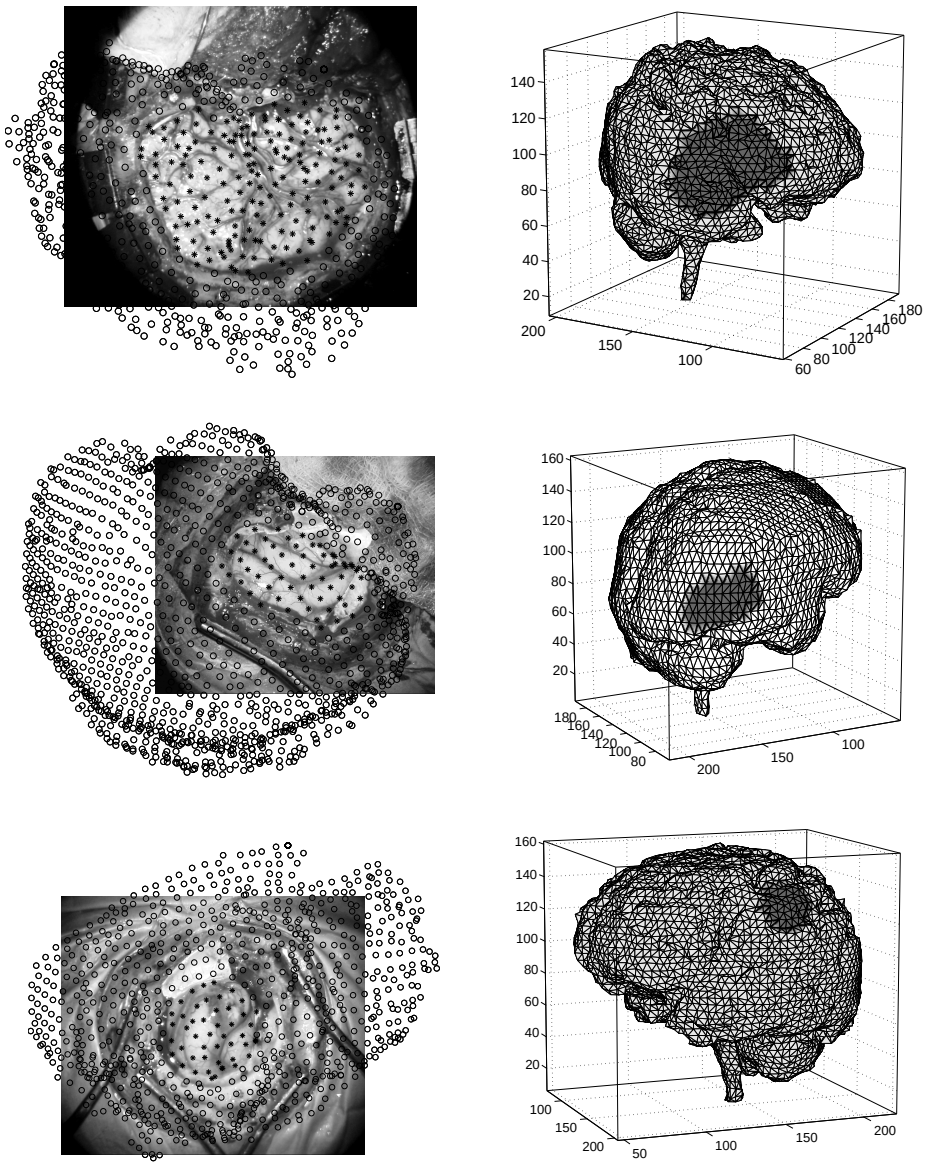

Figure 4.4: Shown on the left is the overlay of the preoperative mesh with the surgical scene after rigid registration, with the craniotomy nodes shown as asterisks. Shown on the right is the preoperative mesh with the craniotomy nodes highlighted. The top, middle, and bottom row corresponds to the first, second, and third surgeries, respectively.

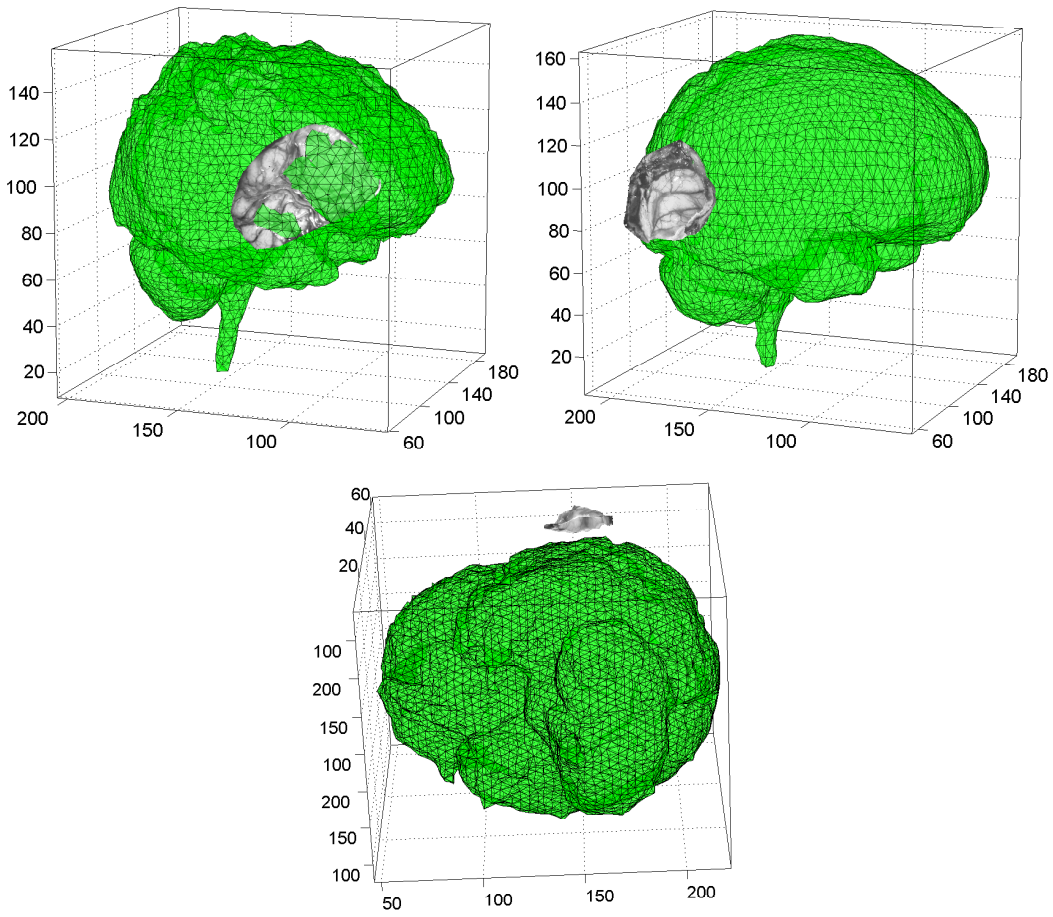

Figure 4.5: Relative position of the cortical surface to the preoperative mesh for the first (top left), second (top right), and third (bottom) surgery.

its previous position (yellow lines). Note that in this figure, although the cortical surface underwent a multiplicity of motions, the reconstructed bony edges of the craniotomy site from the three stereo pairs remained fused with the skull in the pMR slices.

Approximately one thousand points from the reconstructed cortical surface (green) prior to dura opening were selected for tracking. The iterative closest point algorithm was employed to determine distances moved when the surface later bulged (yellow). The average magnitude of the bulge among the points tracked was 5.00mm with a maximum of 7.79mm and a minimum of 2.03mm. The iterative closest point algorithm was again used to determine how these points moved from the bulged surface (yellow) to the sagged

surface (magenta) after partial cyst drainage. Since a portion of the cortical surface was resected to open the cyst cavity, the reversed iterative closest point algorithm was used for tracking. The average magnitude of the sag was 5.77mm with a maximum of 8.60mm and a minimum of 2.35mm.

By tracking these 3-D points, we estimated the cortical displacement with respect to the cortical surface prior to any deformations. In most cases, the cortical surface under the dura was regarded as the state of no deformations. When a stereo pair of the cortical surface with dura intact is unavailable, the cortical surface segmented from the co-registered pMR volume is used instead. Shown in Table 4.2, the topology of the cortical surface was captured with approximately one thousand point elements during each of the ten clinical cases. Minima, maxima and means of the estimated displacements are listed. A negative value indicates a clasping motion, while a positive value indicates a distending motion, relative to the closed cranium position. When the estimated motion is a combination of bulge and sag, the minimum displacement represents the largest sag, and the maximum displacement represents the largest bulge.

So far, we have seen that the ICP algorithm is capable of establishing the correspondence between two estimated shapes of the cortical surface, from which the displacement was then computed. Next, we compare our tracking strategies with the optical tracking system.

4.4 Tracking Accuracy

To quantify the accuracy in tracking of the cortical surface, we again compare results obtained through stereopsis with optical stylus measurements. This is done during surgery 6 in Table 3.1. Shown in Figure 4.8 are three cortical features (CF1, CF2, and CF3). Their locations were digitized using the stylus before and after cyst drainage. These locations were compared to the positions of the same features on the estimated

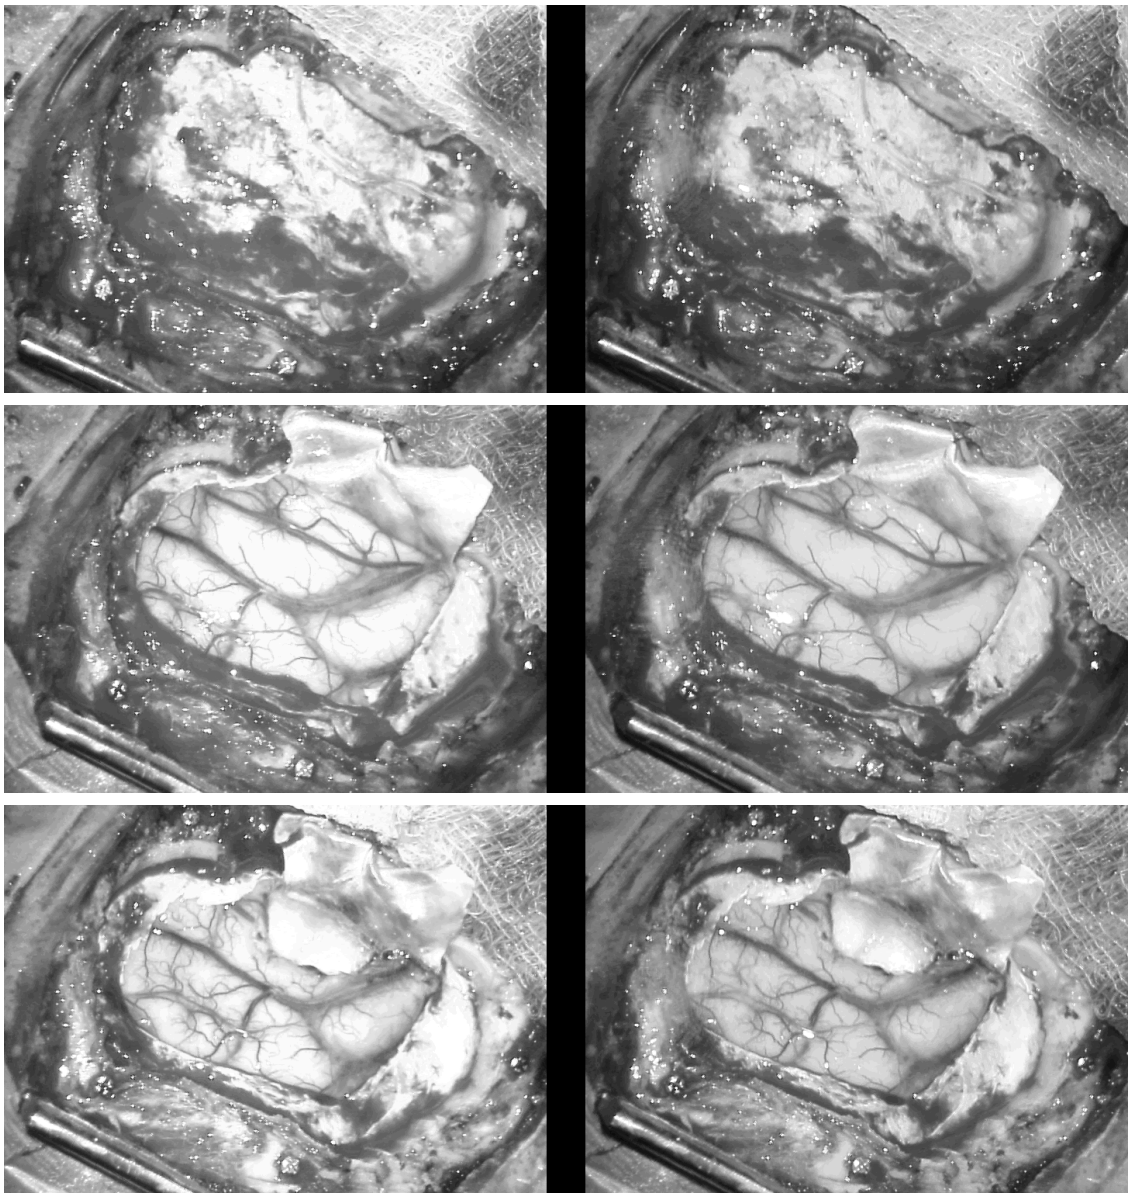

Figure 4.6: Stereo pairs acquired at different stages of the surgery for patient 3. Shown, from top to bottom, are stereo pairs acquired before dural opening, after dural opening, and after partial cyst drainage.

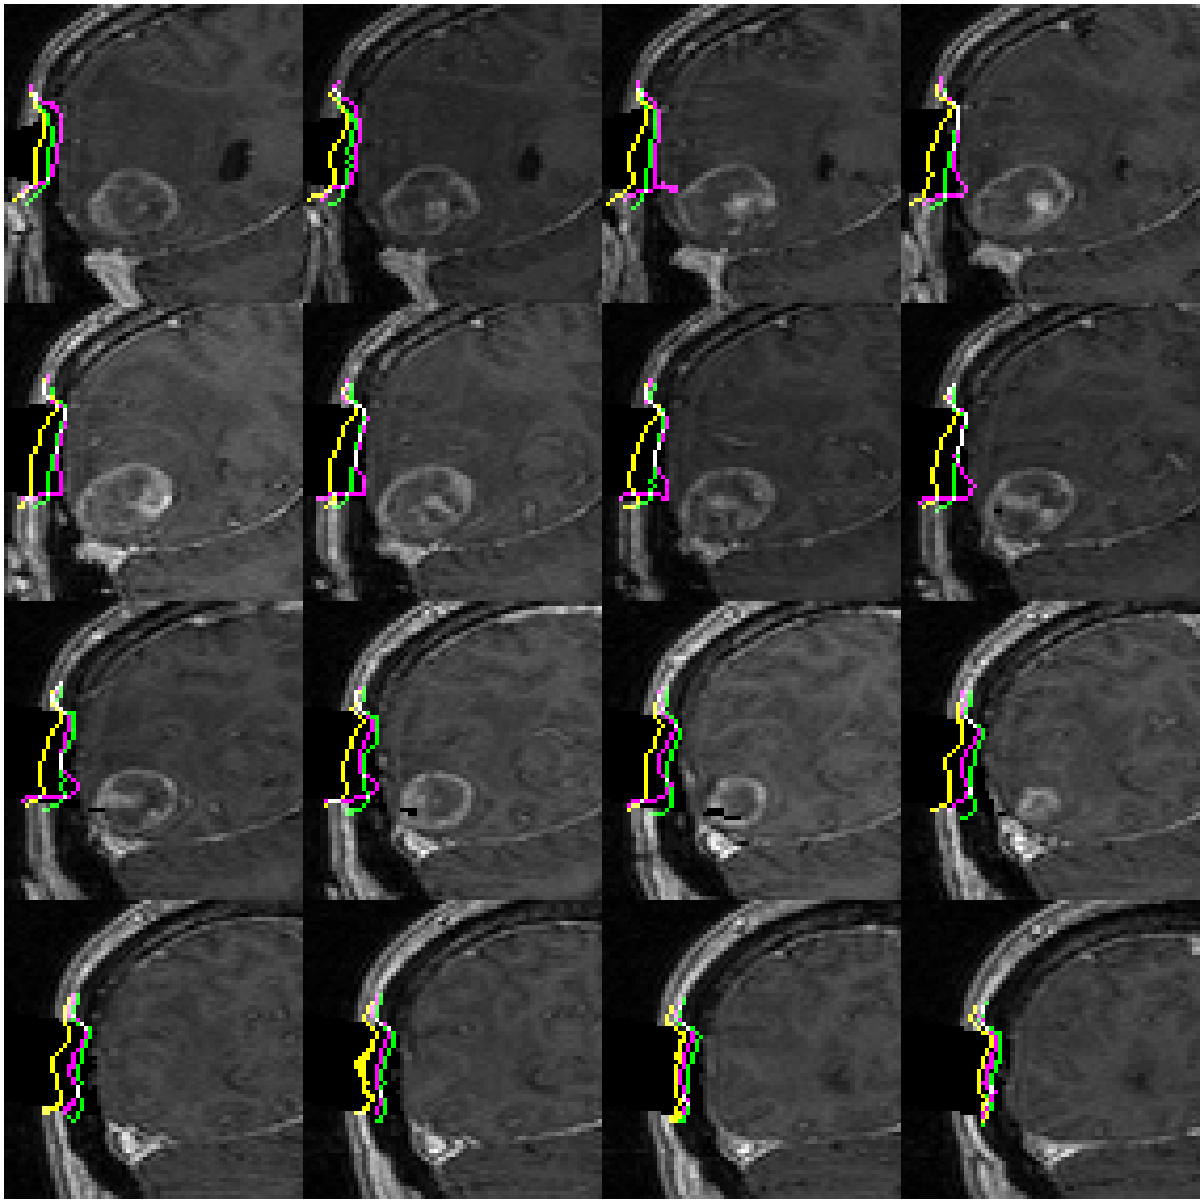

Figure 4.7: Three cortical surfaces (color lines) reconstructed from the stereo pairs in Figure 4.6 overlayed on the rigidly registered coronal pMR slices. Green, yellow, and magenta lines correspond to surfaces before dural opening, after dural opening, and after partial cyst drainage, respectively. Note that the cortical surface bulged when the dura was opened, and then it sagged once the cyst was partially drained.

Patient No.	Surgical stage	Number points reconstructed	Minimum disp. (mm)	Maximum disp. (mm)	Average disp. (mm)
1	After electrode implant	1057	2.15	5.92	3.30
2	After dural opening	1023	-4.51	3.97	1.22
3	After dural opening	1098	-3.22	4.30	0.73
4	After dural opening	1002	6.82	14.97	9.06
5	After dural opening	1004	2.03	7.79	5.00
6	After cyst drainage	1038	-3.43	-12.42	-9.53
7	After partial resection	1019	-3.43	-8.35	-6.45
8	After partial resection	1011	-2.56	-6.52	-4.59
9	After dural opening	1029	-11.15	8.64	-1.83
10	After partial resection	1101	-1.10	-9.31	-6.17

Table 4.2: Measured Cortical Displacements for Ten Individual Cases

cortical surface. When using stereopsis for tracking these cortical points, we manually locate them on the reconstructed surface before cyst drainage. The displacement of each feature was then calculated using ICP, as described in Section 4.2, and using a simple gravitational sagging model (GS). Tracking using GS is performed by intersecting each node of the surface at time t with the surface at time $t+\Delta t$ along the direction of gravity. The GS model, we believe, is a suitable way to establish the correspondence between two shapes, when the motion is induced in the direction of gravity such as in the case of the cyst drainage. Table 4.3 contains the estimated displacements for all three techniques. Note that, in this case at least, the ICP approach under-estimates and the GS method over-estimates the displacement measurements acquired from the stylus probe. While the stylus probe defines the world (OR) coordinate system, which is regarded as "ground truth," it also contributes an error component although if well calibrated the probe error

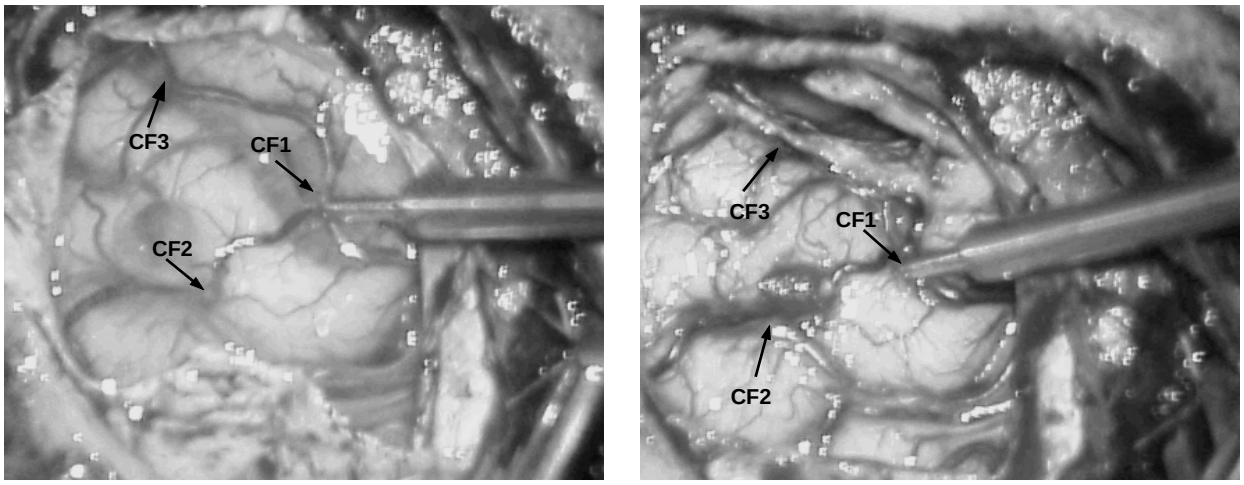

Figure 4.8: The cortical features (CF1, CF2, and CF3) were digitized using the optical stylus before (left) and after (right) cyst drainage.

	CF1	CF2	CF3
stylus	11.6	8.9	11.8
ICP	9.3	7.5	9.7
GS	13.3	9.8	12.1

Table 4.3: The estimated displacements (mm) of these features from the image-based tracking methods (ICP and GS) relative to the stylus method.

is expected to be small. This error, however, may be exacerbated due to the inaccuracies in locating precisely the same position on the cortical surface before and after cyst drainage.

In summary, we have presented a method of tracking the cortical surface motion from the stereoscopically estimated shapes. Our tracking method is divided into two phases. During phase I, the baseline cortical motion is established by comparing the stereoscopically estimated cortical surface with the surface in the pMR volume. During phase II, the cortical surface is continuously tracked by comparing the current shape with the previous estimation. We tested our method in ten clinical cases during which a variety

of cortical motions have quantitatively estimated. Having obtained the intraoperative data of cortical motion, we now conclude our discussion on stereopsis and move on with a presentation of how to use the stereoscopically estimated intraoperative data to guide brain models to compensate for brain shift.

Chapter 5

Brain Modeling

Brain modeling offers a less expensive alternative to full-volume intraoperative imaging techniques such as iMRI for the compensation of the brain shift. The goal of brain modeling is to produce a displacement field over the patient's brain that accurately represents the current deformation in the OR . This vector field provides a *non-rigid* transformation that relates the preoperative image data with the patient anatomy during surgery, and hence can be used to update the pMR volume. Typical brain deformation models combine the bio-mechanical properties of the brain and boundary and volumetric forcing conditions to predict how tissue moves under these conditions [64, 29, 17, 24, 75, 13]. Due to the complexity of brain shift patterns [32] and uncertainties among forcing conditions, several groups have suggested that the accuracy of these models can be improved when measurements of brain motion acquired during sugery are used to guide the model estimates [24, 75, 49, 80]. In previous chapters, we have employed stereopsis to estimate the shape of the cortical surface and track its motion over time. In this chapter, the intraoperative displacement is used to guide a brain model to estimate the motion of other parts of the brain. We begin our discussion with a general description of the brain model employed for this task.

5.1 Dartmouth Brain Model

We have adopted a poroelastic model, using Biot's consolidation formulation [6] to represent the brain as a deformable porous medium. Biot's theory of consolidation describes the continuum as biphasic in nature. More specifically, the brain volume contains a solid matrix with the interstitial spaces filled with an incompressible fluids (such as CSF and saline). The gravity and hydraulic body forces as well as surgical forces are included in the model. The volumetric force due to gravity is based on the difference between the density of the brain tissue and the surrounding fluid. The hydraulic body force is created by the fluid in the interstitial space. This model implies that tissue deformation is characterized by two distinct stages. When the brain is subjected to a load there is an instantaneous deformation at the contact area followed by a second stage of additional deformation from exiting interstitial fluid. Gradients of interstitial fluid pressure are distributed volumetrically, hence, the net body force acting on a given volume is the interstitial pressure gradient integrated spatially over that volume. The above concepts can be described mathematically using Equation (5.1). Although the interstitial fluid is incompressible, the interstitial space is not necessarily saturated with the fluid, i.e., gaseous voids may be present. This can be conceptualized as a net compressibility which acts to delay the distribution of pressure and create a more diffusion-like behavior of the pressure rather than the instantaneous propagation associated with incompressibility, Equation (5.2):

$$\nabla \cdot G\nabla \mathbf{u} + \nabla \frac{G}{1-2\nu}\nabla \cdot \mathbf{u} - \alpha \nabla p + (\rho_t - \rho_f)\mathbf{g} = 0, \tag{5.1}$$

$$\alpha \frac{\partial}{\partial t}(\nabla \cdot \mathbf{u}) + \frac{1}{S}\frac{\partial p}{\partial t} - \nabla \cdot k\nabla p = \Psi. \tag{5.2}$$

In these equations, the functions $\mathbf{u}$ and p are the displacement vector and pore fluid pressure to be computed, and the remaining parameters are defined in Table 5.1. Finite element treatment of equations 5.1 and 5.2 in the pore fluid pressure leads to a modification of the discrete matrix system for time advancement of the solution reported in [64]. The solution to equations 5.1 and 5.2 is obtained by applying known boundary

Table 5.1: Parameters used in the brain model equations

parameter	description	typical value
G	shear modulus	2100Pa
ν	Poisson's ratio	0.46
ρ_t	density of brain tissue	$1000\frac{kg}{m^3}$
ρ_f	density of surrounding fluid	$1000\frac{kg}{m^3}$
k	hydraulic conductivity	$1 \times 10^{-7}\frac{m^3 s}{kg}$
α	ratio of fluid vol. extracted to change in solid vol.	1
$1/S$	fluid that can be forced into the tissue under constant vol.	0
Ψ	pressure source strength	0Pa/s

and volumetric forcing conditions and solving for $\mathbf{u}$ and p. The material parameters in this equation are based on values in the literature, as well as those deduced through optimization in our own experiments [61].

These boundary and volumetric forcing conditions include fluid saturation in the lower portion of the brain, the gravitational direction acquired during surgery, stress-free conditions at the craniotomy and at the highest elevation of the brain, zero normal displacement and zero tangential traction beneath the remainder of the cranium, free flow of fluid at the craniotomy site, and no flow of fluid beneath the rest of the cranium. With the constraints of these conditions, the brain model produces a prediction of deformation over the entire patient's brain.

The estimated intraoperative displacement is incorporated into the brain model as follows. For each craniotomy node for which a displacement was estimated, the corresponding finite element equations are precisely enforced by this displacement [75, 24]. The output of this brain model is a displacement field over the entire patient brain, which is then used to update the entire pMR volume.

5.2 Intraoperative Update

The brain model guided by the intraoperative data can now be used to update the pMR volume to reflect the changes occurring during surgery. First, the bony skull corresponding to the craniotomy site is removed from the pMR volume. This region is identified by rigidly registering the estimated cortical surface in the OR space to the pMR volume. Once overlayed, the estimated cortical surface defines the boundaries of the region of the pMR volume to be removed. Second, the brain model displacement field is applied to the entire brain in the pMR volume to update the internal tissue. The updated MR (uMR) volume is a combination of the remodeled skull and the deformed brain.

5.3 Results

In the previous chapter, three neurosurgical procedures (surgeries 2, 3, and 4 from Table 3.1) were used to illustrate how the baseline cortical deformation was estimated. In this chapter, we use these estimated intraoperative data to guide our poroelastic model for recovering full brain deformations. Some of these results have appeared in [82]. Here the patients's characteristics and the estimated cortical motion relative to pMR volumes are stated again. Surgery 2 in Table 3.1 involved a 63-year-old female with a right temporal hemangioma and MCA aneurysm. Surgery 3 involved a 62-year-old male with a renal cell carcinoma in the right posterior temporal region of the brain. Surgery 4 involved a 62-year-old male with renal cell carcinoma in the right posterior temporal region of the brain. For surgery 2, of the 174 craniotomy nodes, 112 collapsed and 62 distended. The average displacement of collapsing nodes was 4.76mm, with a maximum of 8.66mm and a minimum of 1.57mm. The average displacement of distending nodes was 3.54mm, with a maximum of 5.64mm and a minimum of 0.12mm. For surgery 3, all 44 craniotomy nodes distended. The average displacement was 6.29mm, with a

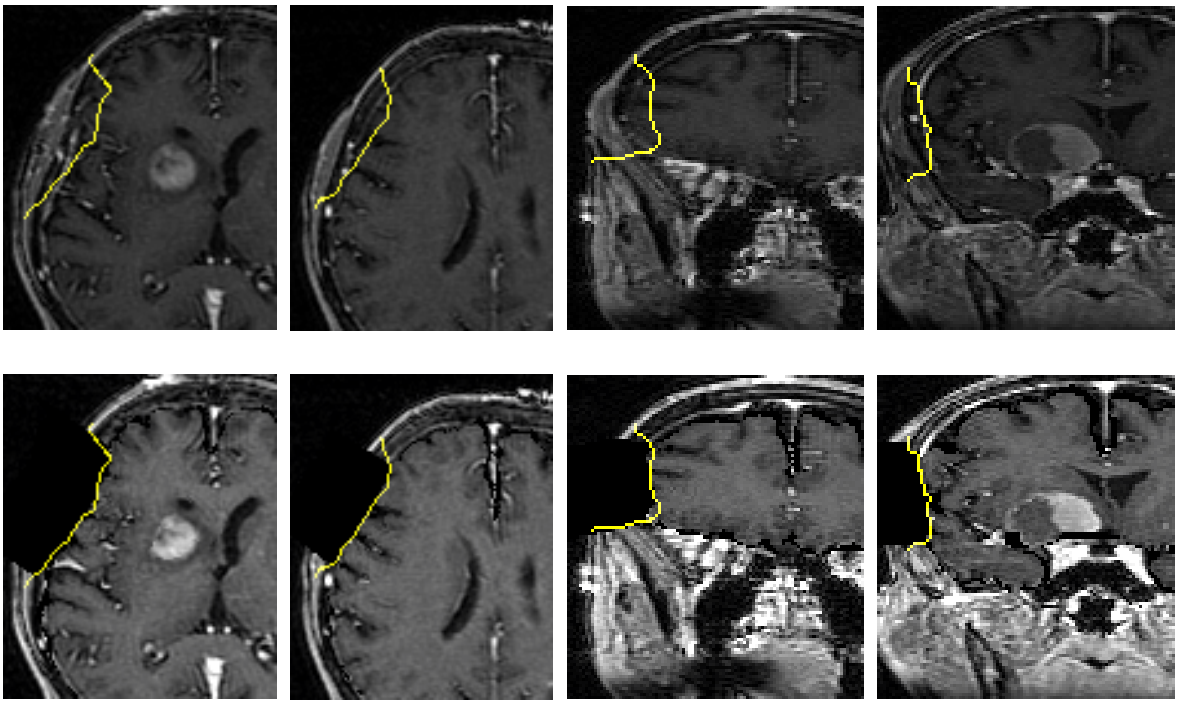

Figure 5.1: The intraoperative cortical surface (yellow) overlayed with two axial slices and two coronal slices from the pMR volume (top row) and the same slices from uMR volume (bottom row).

maximum of 9.03mm and a minimum of 2.95mm. For surgery 4, all 48 craniotomy nodes distended, with an average magnitude of 14.35mm, a maximum of 17.82mm and a minimum of 10.45mm. These results are also presented in Figure 4.5 and Table 4.1 in the previous chapter.

5.3.1 Intraoperative Update

The estimated cortical displacement (previous chapter) was used to guide the biomechanical model, as described in Section 5.1. This model generated a displacement field which was then applied to the pMR volume. Figures 5.1, 5.2 and 5.3 present the results from the three surgeries. The top row in each figure consists of the pMR slices (two axial slices and two coronal slices) with the reconstructed cortical surface overlayed in yellow. Note the misalignment between the intraoperative and preoperative surfaces.

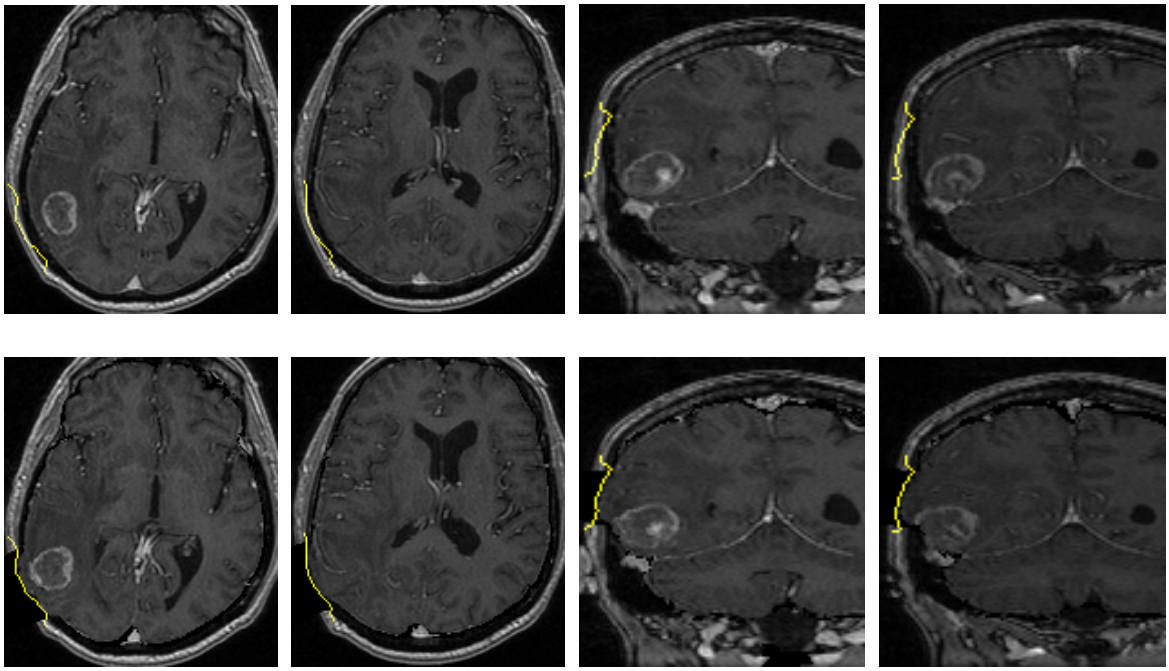

Figure 5.2: The intraoperative cortical surface (yellow) overlayed with two axial slices and two coronal slices from the pMR volume (top row) and the same slices from uMR volume (bottom row).

The bottom row displays the intraoperative surface overlayed with the updated MR (uMR). Note that these surfaces are now in good alignment.

5.3.2 Validation Using Co-registered iUS

In the previous section we showed how the incorporation of measured cortical motion into brain models can more accurately represent the surface of the brain. Here, we demonstrate how this correction can more accurately compensate for tissue motions in the interior regions of the brain as well.

Three MR volumes from both surgery 2 and surgery 3 were selected for comparison: the pMR volumes, and the two uMR volumes generated by the brain model with and without stereopsis guidance. In the model calculation guided by stereopsis, the boundary condition at the craniotomy site was of type I, constraining the cortical surface to the

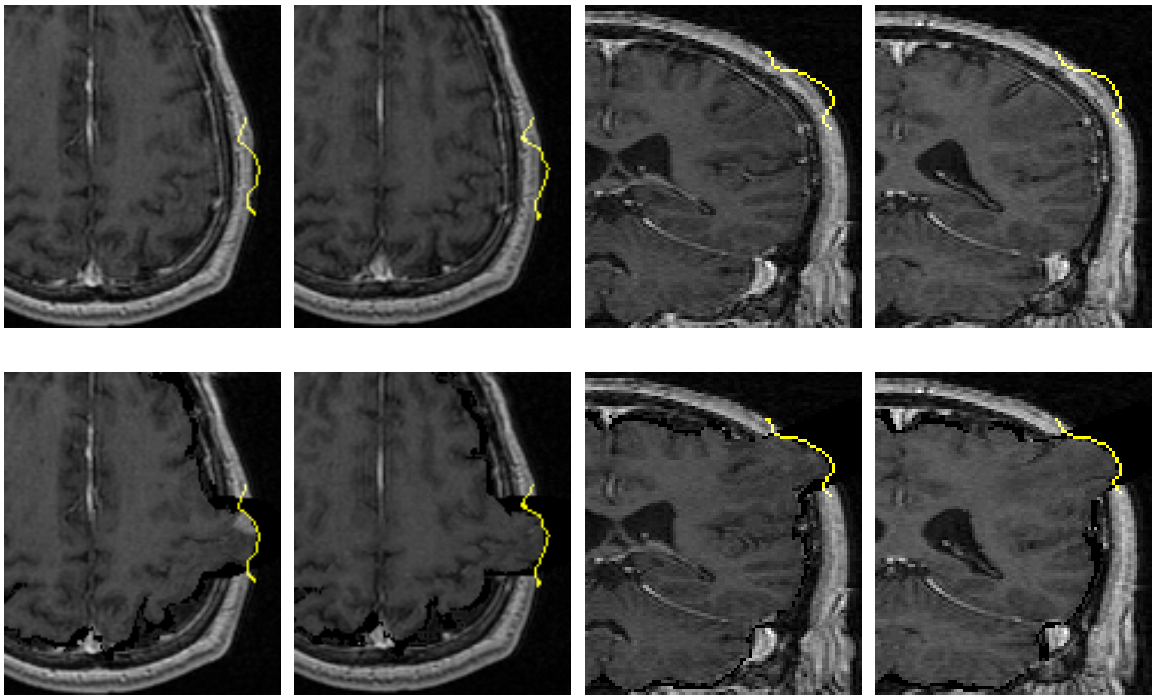

Figure 5.3: The intraoperative cortical surface (yellow) overlayed with two axial slices and two coronal slices from the pMR volume (top row) and the same slices from uMR volume (bottom row).

observed deformation, whereas for the model without guidance, the boundary condition on the exposed surface was a type II, stress free condition. All other forcing conditions were identical for the two model calculations used to generate the uMR volumes. The location of a sub-surface tumor was compared in all three of these MR volumes with intraoperative ultrasound (iUS) [1]. Prior to surgery, the iUS was calibrated and its 2-D image plane was rigidly registered to the world frame defined by the Polaris System and then to the pMR volume [49]. For each iUS image, a corresponding oblique image slice from the MR volume was formed by computing the intersection between the co-registered iUS image plane and 3-D MR volume. The sampling lattice of the oblique

[1] A 2-D iUS was employed (SONOLINE Sienna digital ultrasound, Siemens Medical Systems, Elmwood Park, NJ, USA) with a 5MHz curved array. A 6cm depth scale was used, and the iUS transducer was rigidly attached with a 3-D spatial tracking device (Model TT002-C, Traxtal Technologies, Toronto, Ontario, Canada).

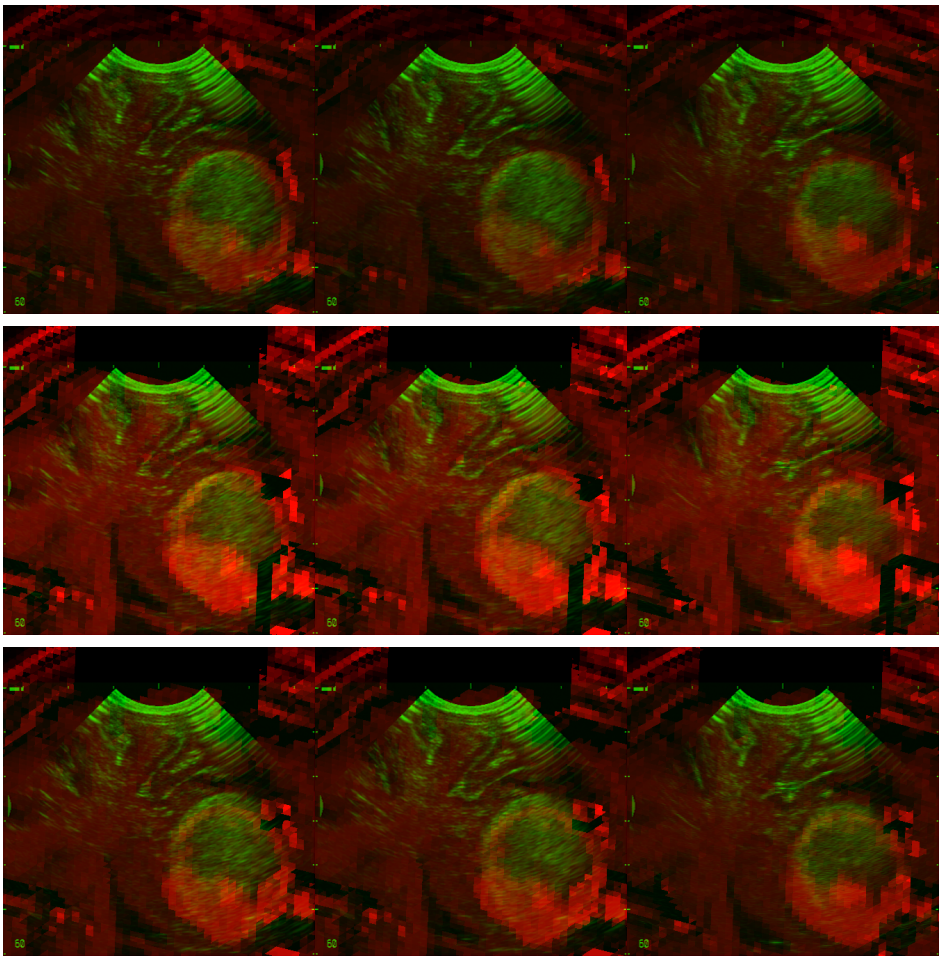

Figure 5.4: MR (red) and US (green) overlays from surgery 2. Three different iUS images and their corresponding MR slices are shown in each column. The top row is the pMR volume, the middle row is the uMR with stereopsis guidance, and the bottom row is the uMR without stereopsis guidance.

MR image was chosen to be the same as the iUS image and a color image of the overlay between the two was formed with the iUS image in green and the corresponding oblique MR slice in red. Figures 5.4 and 5.5 are representative examples from the first and second surgery, respectively.

Shown in each column of each of the two figures are three different iUS images with their corresponding MR slices. The top row is the pMR volume, the middle row is the uMR with stereopsis guidance, and shown in the bottom row is the uMR without

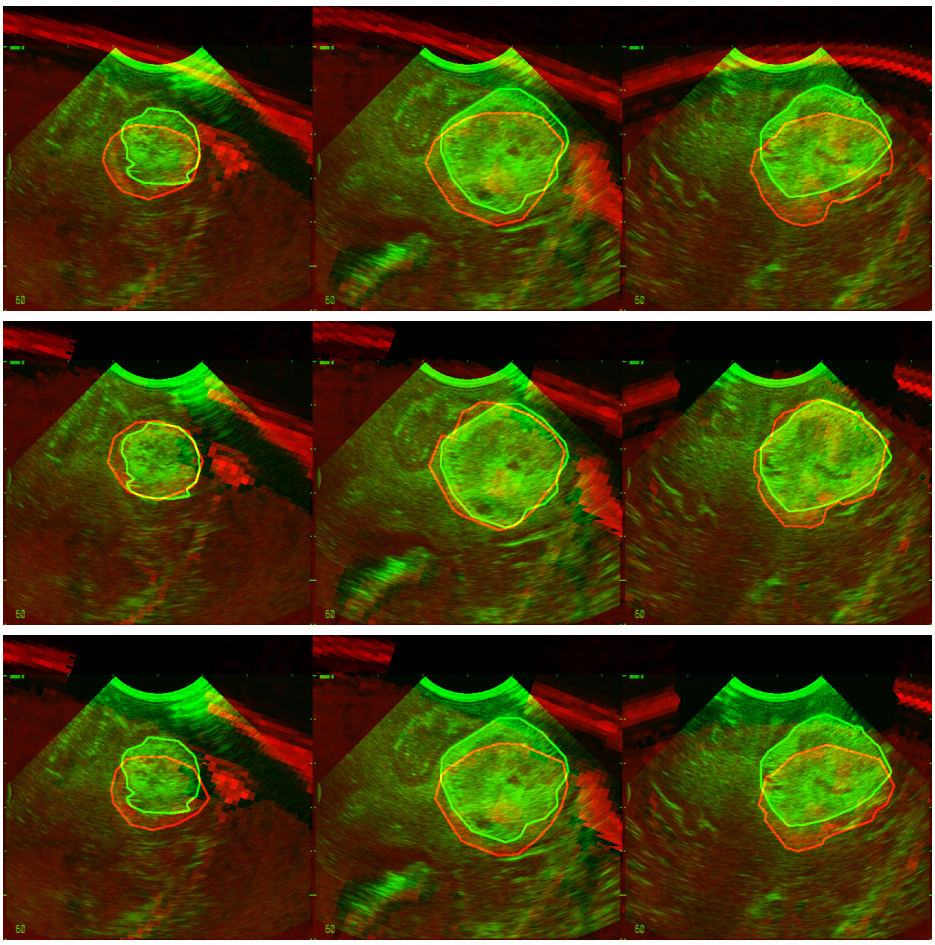

Figure 5.5: MR (red) and US (green) overlays from surgery 3. Three different iUS images and their corresponding MR slices are shown in each column. The top row is the pMR volume, the middle row is the uMR with stereopsis guidance, and the bottom row is the uMR without stereopsis guidance.

stereopsis guidance. Note that, in both Figures 5.4 and 5.5, the sub-surface tumors in the iUS are in good alignment with the uMRs with stereopsis guidance, but are generally in poor alignment in the other volumes. Note also, in both Figures 5.4 and 5.5, that the surfaces of the iUS transducer in the middle rows are well aligned with the cortical surfaces.

To quantify this improvement, we manually outlined the tumor in the pMR volume, in the stereopsis-guided uMR volume, in the non-guided uMR volume, and in the corresponding iUS images from surgery 3. These contours were aligned using the ICP

Table 5.2: Comparison of tumor alignment between MR volumes and co-registered iUS

MR Volume	max (mm)	min (mm)	mean (mm)
pMR	8.7	1.9	5.2
stereo-guided uMR	5.3	0.4	2.8
non-guided uMR	10.5	2.6	6.5

algorithm, as described in Chapter 4. The alignment errors between these MR volumes and iUS are presented in Table 5.2. The mean alignment error was reduced from 5.2mm between pMR and iUS to 2.8mm between stereo-guided uMR and iUS. The same error, however, increased from 5.2mm to 6.5mm when the brain model was not deployed under stereo guidance. This was due to the fact that, without intraoperative data, the model predicted a collapsing tissue motion in response to the craniotomy and dural opening. The example illustrates that motion estimation at the surface is capable of improving the accuracy of the entire updated image volume of the brain, including sub-surface regions.

In summary, we have designed and implemented an intraoperative stereo vision (iSV) system capable of estimating the 3-D shape of the surgical scene in near real-time. In three clinical examples, we have estimated the deformation of the cortical surface upon the opening of the dura. The intraoperative cortical displacement is then used to guide a full brain deformation model, which updates the pMR volume. We have shown that the estimated cortical shape and motion is able to guide a full brain deformation model, allowing it to accurately compensate for significant shifts both at and below the cortical surface. In an exemplary case, we report an average error reduction of 3.7mm relative to the model estimate without using intraoperative data. In the next chapter, we extend our modeling method to a different stage of a surgical procedure, where tissue retraction

is modeled with the guidance of the stereoscopic data.

Chapter 6

Tissue Retraction

During neurosurgery, the surgeon often employs retractors to intentionally deform the tissue in order to gain access to deeper structures. Although the retraction of brain parenchyma during surgery is common, detailed studies of the effects of retraction on tissue are few [32]. Existing simulations are largely limited to qualitative descriptions of the retraction process [45] and are not capable of producing an accurate estimate of the mechanical effect on the parenchyma.

As we have seen in the previous chapter, computational brain models have proven to be powerful in compensating for brain shift, but their success has been largely limited to the initial tissue response to the craniotomy and dural opening. Miga et al. proposed a strategy for digitizing the motion of the retractor and the volume of the resection cavity in order to incorporate these changes into a brain model [59]. In their initial attempt, the intraoperative data were only qualitatively estimated based on preoperative images due to the difficulty in tracking surgical tools and the brain surface during surgery. Using porcine subjects, Platenik et al. [65] demonstrated that their brain deformation model is capable of capturing 75% - 80% of the tissue motion generated during interhemispheric retraction.

In order to accurately model the retraction of a human brain, we employ the iSV

system to capture the motion of a retractor and continuously monitor the cortical motion during tissue retraction. The measured motion of the retractor is incorporated into the brain model to produce full volume deformation estimates, which are then used to update the pMR volume. Using this approach, we have modeled tissue changes both during retraction and after the release of the retractor in a clinical case.

6.1 A Clinical Case

The modeling of tissue resection follows the same strategy as presented in the last section of Chapter 1. A clinical case study is presented to illustrate how we use the motion of the retractor to model the tissue retraction process. The patient was a 52-year-old female with a right sphenoid wing meningioma. [1] Because of the location of the tumor, the surgeon performed a craniotomy at the right temporal region of the skull and retracted the brain tissue at the craniotomy site in order to gain access to the tumor. In order to capture the motion of both brain and surgical tools, the iSV system was used to record the surgical field every ten seconds during the entire procedure.

Prior to surgery, the iSV system was calibrated to obtain the intrinsic and the extrinsic camera parameters. After the patient was anesthetized in the OR, patient registration was performed using seven scalp-attached fiducials. FRE and TRE of the patient registration were 1.98mm and 1.78mm respectively.

Prior to surgery, a standard volumetric MRI was obtained for the patient using 1.5 Tesla GE MR scanner, and the brain was segmented. A finite element mesh was generated from the resulting segmented brain and registered to the real-time microscope images acquired at the start of surgery, as shown in Figure 6.1(left and middle). The surface nodes that correspond to regions of the craniotomy site (lighter gray) and the tissue underneath the retractor (darker gray) are identified using the rigid registration

[1] This surgery is not listed in Table 3.1.

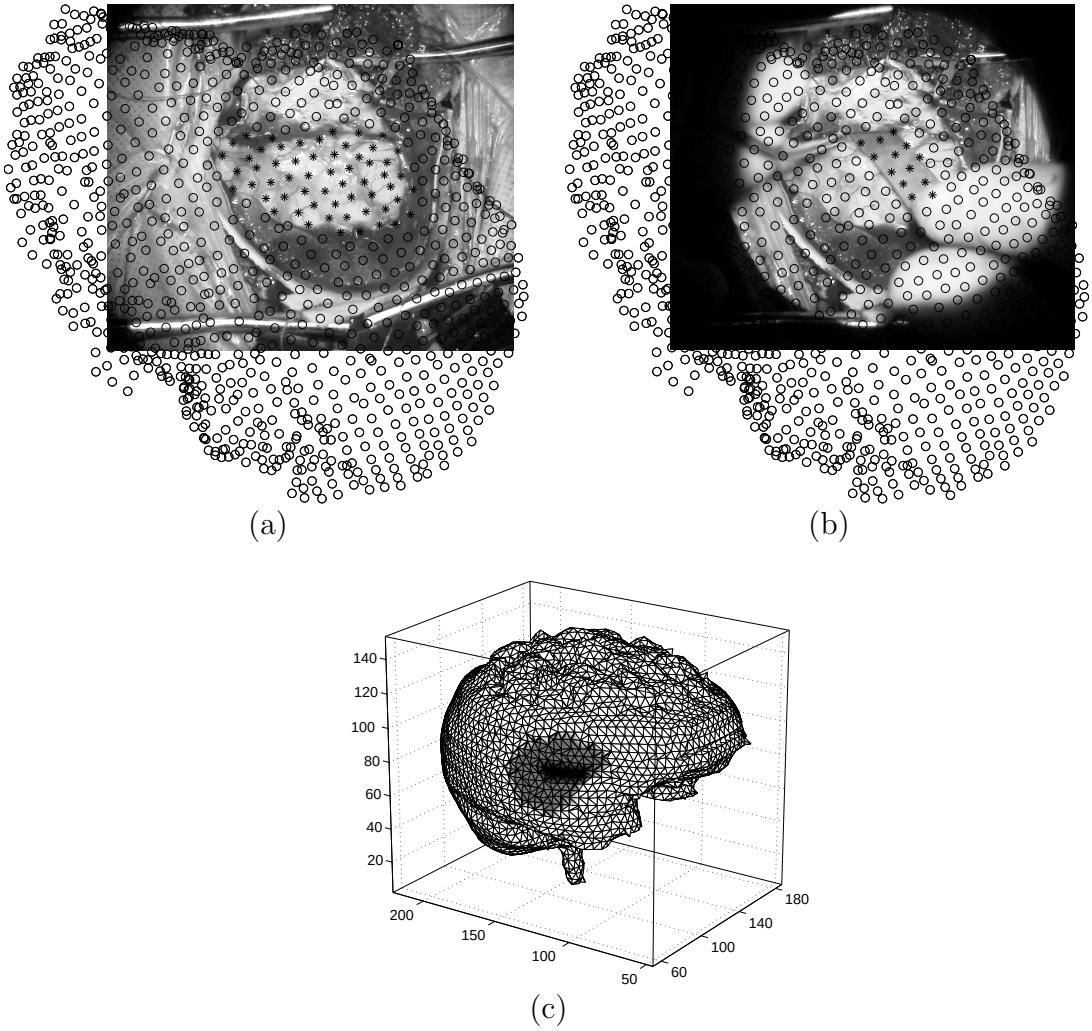

Figure 6.1: Shown, from left to right, are (a) the results of projecting the mesh surface nodes onto the surgical scene, where the nodes within the craniotomy boundary are plotted as asterisks and the others as open circles; (b) a second projection of the mesh surface nodes, where those corresponding to brain tissue underneath the retractor are plotted as asterisks and the others as open circles; and (c) the mesh with the craniotomy nodes colored in lighter gray and the retractor nodes colored in darker gray.

between the microscope view and the pMR volume, as shown in Figure 6.1(right).

Shown in Figure 6.2(top) is a stereo pair of the retracted brain tissue. From these images, shapes of the retractor and the cortical surface were computed. These shapes were overlaid with the preoperative mesh to illustrate the tissue deformation, as shown in Figure 6.2(bottom). As a result of tissue retraction, the anterior portion of the exposed cortical surface collapsed and the posterior portion of the cortical surface distended. From the retractor nodes in Figure 6.1 and the shape of the blade in Figure 6.2, the motion of the retractor was estimated using the ICP algorithm. Of the 13 retractor nodes, 8 collapsed and 5 distended. The average displacement of collapsing nodes was 9.23mm, with a maximum of 12.34mm and a minimum of 4.34mm. The average displacement of distending nodes was 5.42mm, with a maximum of 6.34mm and a minimum of 2.45mm.

The estimated retractor motion was used to guide the brain deformation model, as described in Chapter 5. This model produced a displacement field which was then applied to the pMR volume. Figure 6.3 presents the result from this update. The images on the left are the axial slices in the pMR volume with the reconstructed cortical surface (white curves) overlaid. Note the misalignment between the intraoperative and preoperative surfaces. The images on the right display the intraoperative surface overlaid with the updated MR (uMR). Note that these surfaces are now in better alignment.

6.2 Validation

In order to confirm the viability of our approach in modeling tissue retraction, we compared the model-estimated cortical motion with the same motion captured via stereopsis. Given the accuracy of the iSV system at approximately 1mm [82], we consider the latter estimate as close to the ground truth. The cortical surface surrounding the retractor was used as an independent evaluative landmark.

To begin, the retractor nodes were subtracted from the craniotomy nodes. The re-

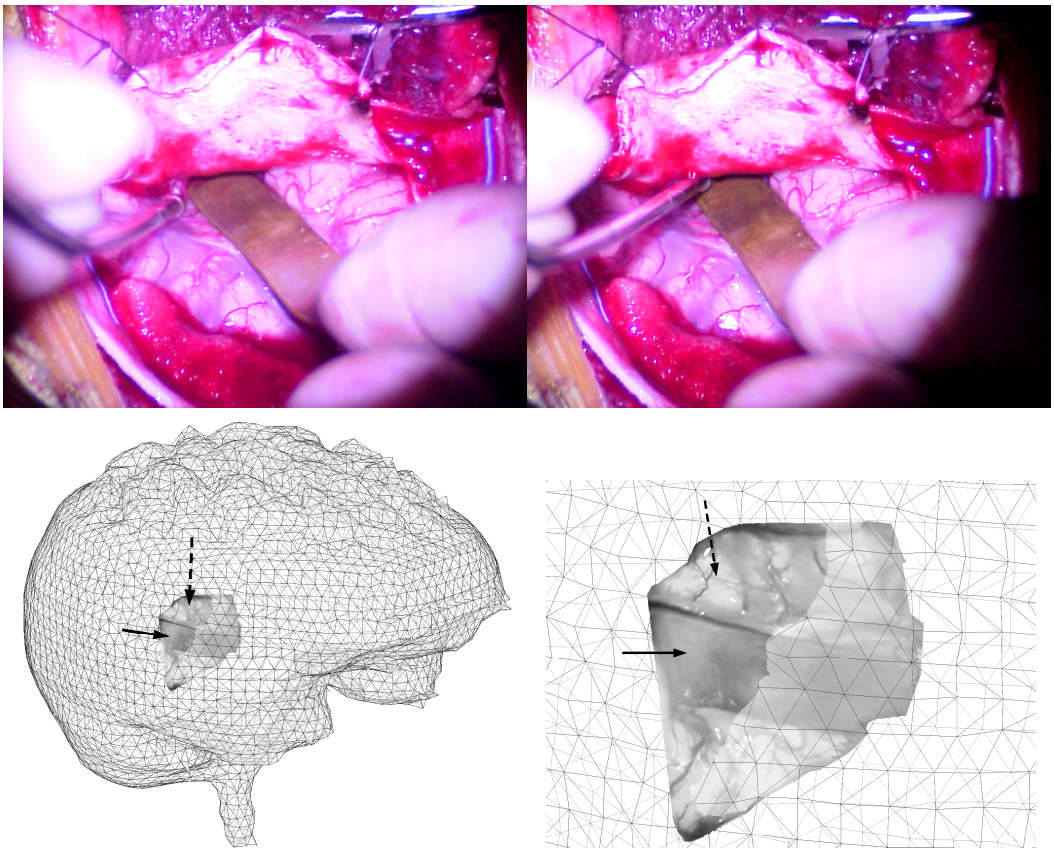

Figure 6.2: Shown are: the stereo pair of the retracted brain (top); relative position of the retractor (black arrow) and the deformed cortical surface (dotted arrow) to the preoperative mesh (bottom left) and a close-up view of the same figure (bottom right).

maining craniotomy nodes, corresponding to the cortical surface surrounding the retractor, were selected for comparison. This set of surface nodes was defined as the *cortical* nodes. As described previously, the brain model had predicted a displacement for each cortical node. From the stereo pair shown in Figure 6.2, the shape of the cortical surface surrounding the retractor was estimated. Using the method described in Chapter 4, the stereo-based estimate of the exposed cortical surface was compared with the corresponding mesh surface, i.e., the cortical nodes, to obtain their displacement. Of 24 cortical nodes, the stereo-based strategy estimated that, 15 of those nodes collapsed and 9 distended during retraction. The average displacement of collapsing nodes was 8.76mm,

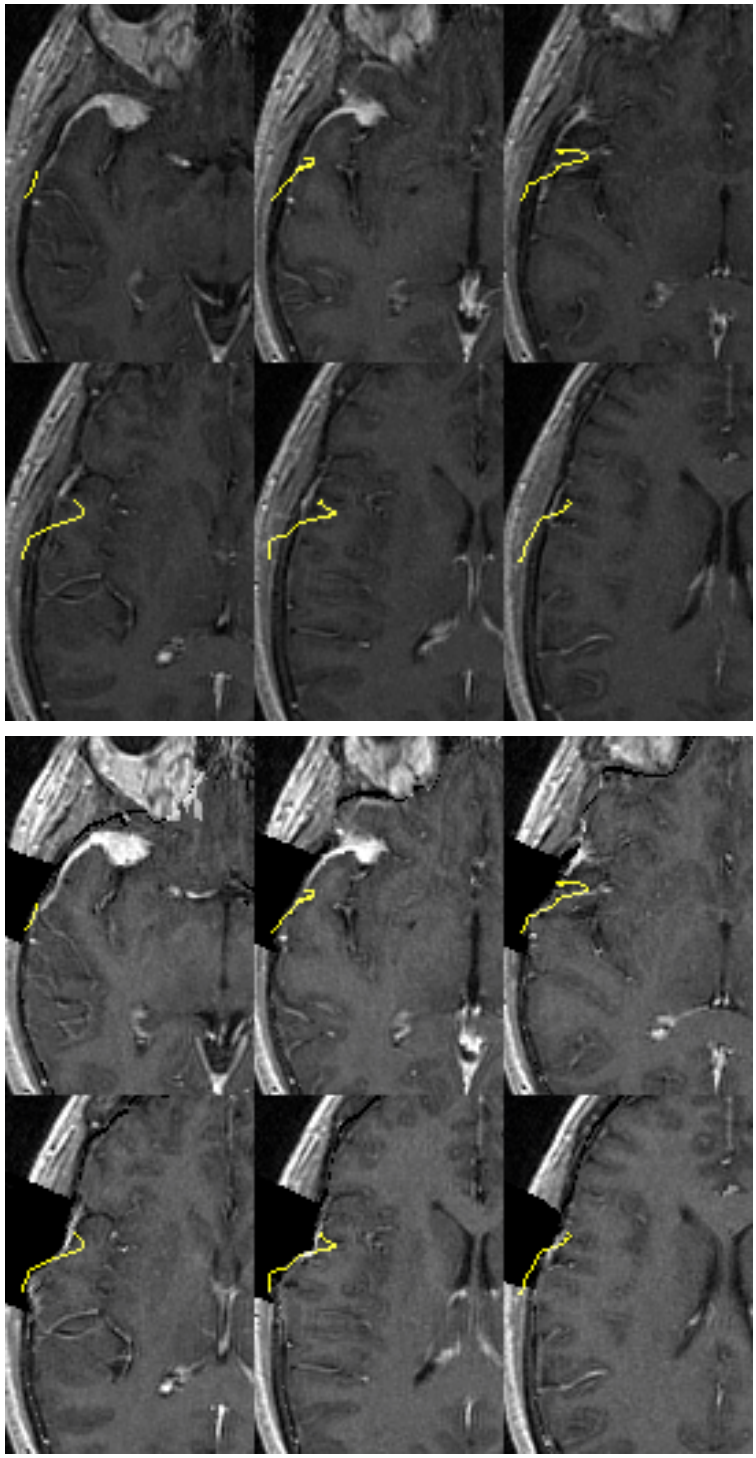

Figure 6.3: The retracted cortical surface (white curves) overlaid with axial slices from the pMR volume (top) and the same slices from uMR volume (bottom).

with a maximum of 10.65mm and a minimum of 2.46mm. The average displacement of distending nodes was 4.64mm, with a maximum of 5.64mm and a minimum of 1.83mm.

We next computed, for each cortical node, the absolute difference between the stereo- and model-estimated displacement. The relative difference was determined as the absolute difference divided by the amount of displacement estimated via stereopsis. The final error estimates were cast as the percent capture of tissue motion by subtracting each relative difference from 100 percent. These error estimates were computed in all Cartesian directions (X, Y, Z) and in overall magnitude, Table 6.1. In this particular case, we have found that our model was capable of capturing approximately 75% of the cortical motion.

In order to further confirm the viability of our approach, we have employed the same strategy to model tissue changes after the release of the retractor. Thirty seconds after the stereo pair in Figure 6.2 was acquired, the iSV system recorded another stereo pair shown in Figure 6.4 (top). The shape of the cortical surface was estimated using stereopsis and overlaid with the uMR generated during tissue retraction, as shown in Figure 6.4(bottom left). Note that the cortical surface did not fully recover from the induced tissue deformation. The motion of the retractor nodes was tracked during the recovery phase after blade removal. Of 13 cortical nodes, four collapsed and nine distended after the tissue was released. The average displacement of collapsing nodes was 2.43mm, with a maximum of 4.26mm and a minimum of 1.02mm. The average displacement of distending nodes was 7.32mm, with a maximum of 9.18mm and a minimum of 4.69mm. These retractor node displacements were used to guide the brain deformation model for updating the uMR generated during retraction. The images on the right display the current cortical surface overlaid with the uMR generated after 30 seconds during the recovery phase. Note that, compared to the images on the left, the alignment between the intraoperative cortical surface and the surface in the uMR improved as a result of this update. For this modeling step, we also compared the motion of the cortical nodes predicted by the model with the stereo estimation. The model was capable of

Table 6.1: Percent Capture of Deformation During Retraction

	mean	max	min
X	71.6	97.4	56.7
Y	66.3	89.9	44.3
Z	82.1	98.9	64.4
magnitude	74.8	97.2	53.1

Table 6.2: Percent Capture of Deformation After Retraction

	mean	max	min
X	69.2	94.3	49.2
Y	78.1	98.4	55.4
Z	65.3	92.1	45.1
magnitude	72.3	96.8	47.9

capturing approximately 72% of the cortical motion during the recovery phase, as shown in Table 6.2.

In summary, we have presented a technique for modeling tissue retraction. By monitoring changes in the surgical field every ten seconds, we have estimated the motion of a retractor and its concomitant cortical surface movement. The retractor motion is then used to drive brain deformation models for updating the pMR volume. Using the cortical surface surrounding the retractor as an independent evaluative landmark, we show that our approach can recover on average approximately 75% of the tissue deformation. So far, we have completed our discussion on using stereoscopically estimated intraoperative data to guide the brain model to compensate for brain shift. Next, we present some other cortical motions captured by the iSV system during surgery.

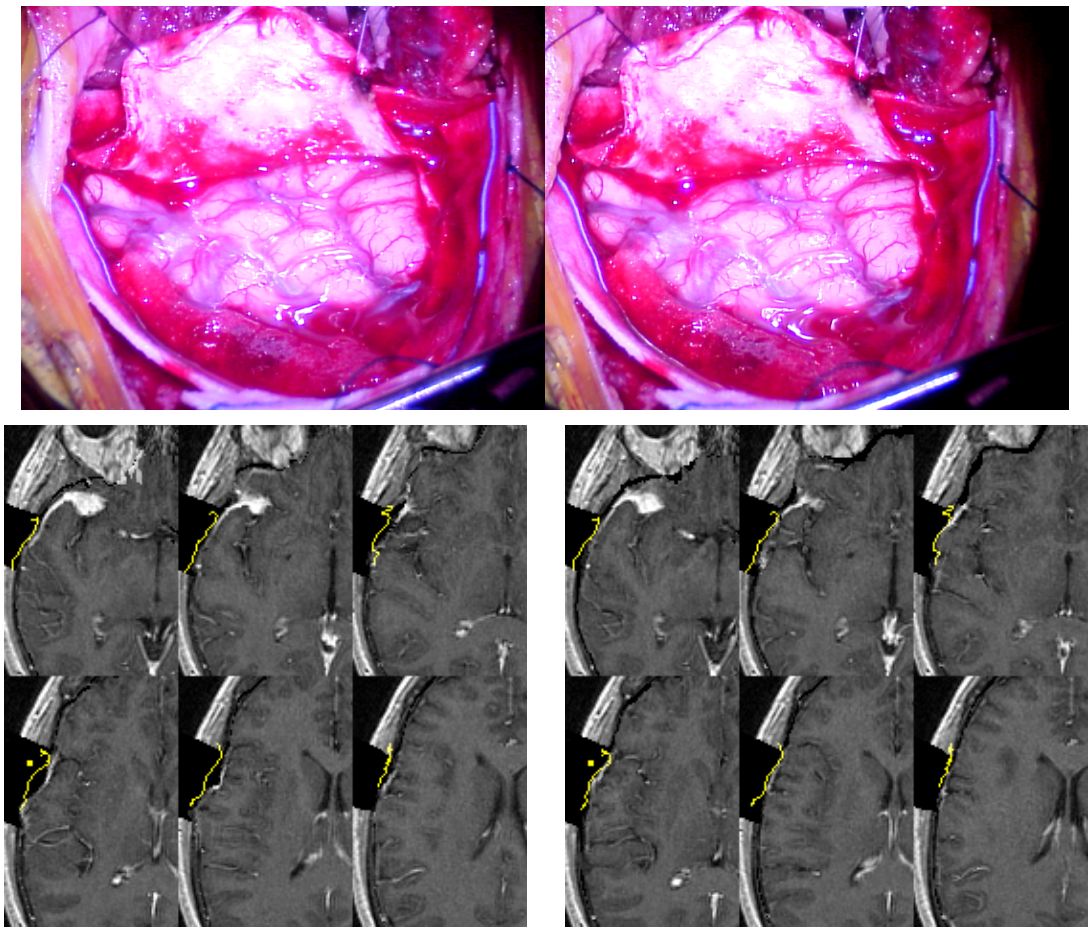

Figure 6.4: Stereo pair of the surgical scene taken 30 seconds after the retractor was released (top); The intraoperative cortical surface (white curves) overlaid with axial slices from the uMR volume generated during retraction (bottom left) and the same slices from uMR volume generated after retraction (bottom right).

Chapter 7

Other Types of Cortical Motion

Advantages of our iSV system include its fast acquisition rates, predominantly automated steps, and overall accuracy. As described below, these characteristics allow us to capture other important cortical behavior that enhances image-guidance and overall clinical decision-making during surgical procedures.

7.1 Resection Cavity

Tissue resection is the major objective in many neurosurgical procedures. Since the pMR images cannot reflect brain deformations, there is a need to update the information regarding the location of a target lesion. To date, intraoperative MRI/CT and intraoperative US are the only modalities employed for intraoperative imaging. While iMRI is costly, the void in the resection cavity poses challenge for ultrasound.

We have employed our iSV system to provide a real-time update of the tissue topology at the craniotomy site referenced to the initially registered pMR. We use surgery 2 in Table 3.1 as an example to illustrate how the shape of the resection cavity is captured. This operation involved a 63-year-old female with a right temporal hemangioma and MCA aneurysm. During surgery, the tumor in the right parietal region was resected

and a resection cavity was created. A stereo pair acquired at this stage is shown in Figure 7.1. The shape of the cortical surface, including the resection cavity, was reconstructed, as shown in Figure 7.1 (bottom) from two different vantage points. Using the techniques described in the previous section, the reconstructed cortical surface and resection cavity are presented as yellow lines in the pMR space, as shown in Figure 7.2. The ability to provide surgical guidance was further enhanced through the illustration of the progression of tissue deformation during surgery. Here, the cortical surface, after the opening of the dura, was also plotted (green) in the preoperative image volume. Note that the resection cavity extended into the tumor region in the pMR and the cortical surface deformed significantly after tissue retraction and resection. Note that the resection cavity did not encompass the tumor. This is likely due to the subcortical tissue deformation resulting from the resection.

Tissue retraction and resection can also be mathematically modeled for updating the pMR images. Miga et al. proposed a strategy to model the effect of tissue retraction and resection [59]. Their technique requires the identification of the resected tissue volume. In their initial attempt, the resected tissue volume was identified manually based on the pMR images of the expected tumor margin. Using the iSV system, we can perform this task intraoperatively and obtain much more accurate results. Shown in Figure 7.3 are the pMR images with the resected area painted in green. The resected area consists of the resection cavity reconstructed using stereopsis and the tumor volume manually segmented from the pMR volume. Our iSV system offers the ability to continuously and intraoperatively monitor the shape of the cortical surface during resection. It is only able to continuously estimate the shape of the cortical surface during resection. The ICP algorithm is not suitable for tracking tissue motion resulting from resection due to the complex and nonrigid tissue deformation pattern.

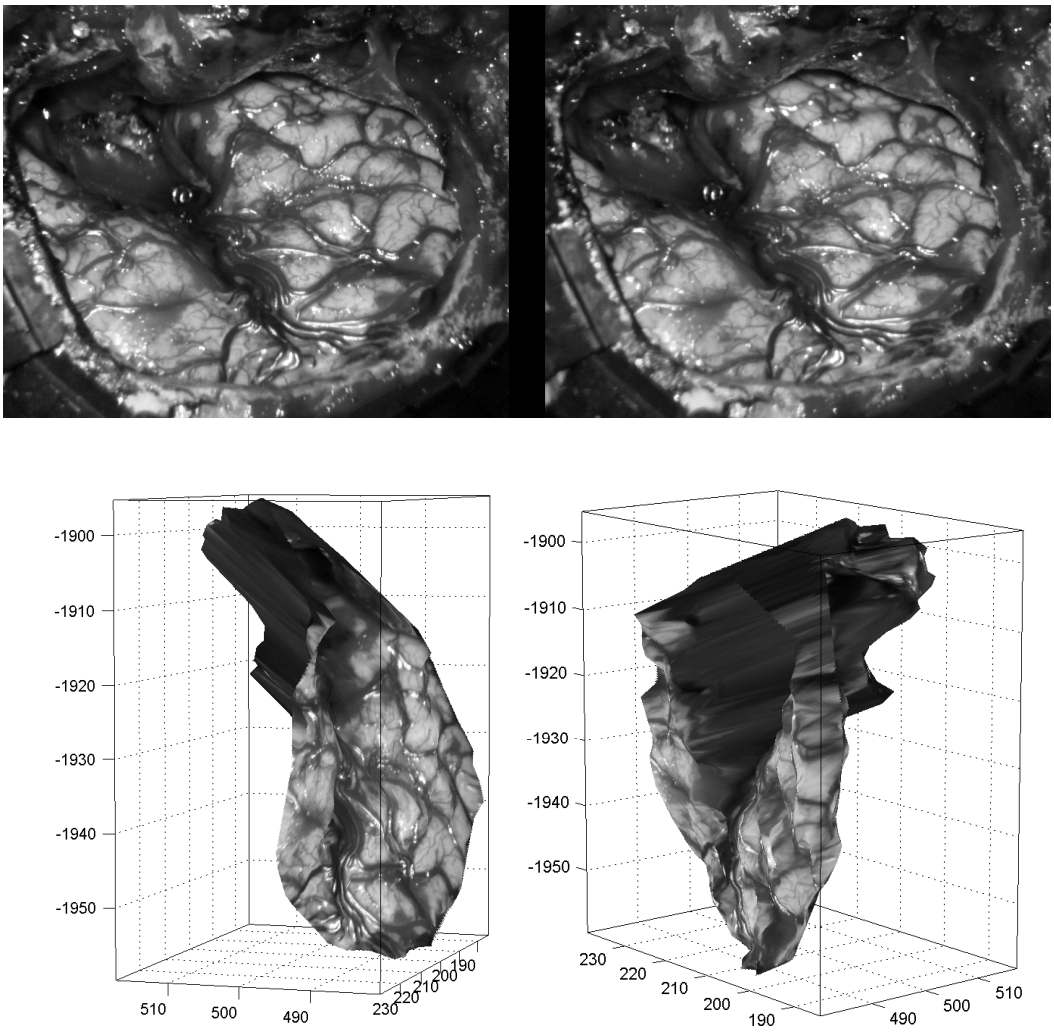

Figure 7.1: Stereo pair of the surgical scene (top) after partial tumor resection and reconstructed cortical surface (bottom, units on axes are in millimeters.) with the resection cavity viewed from two different vantage points.

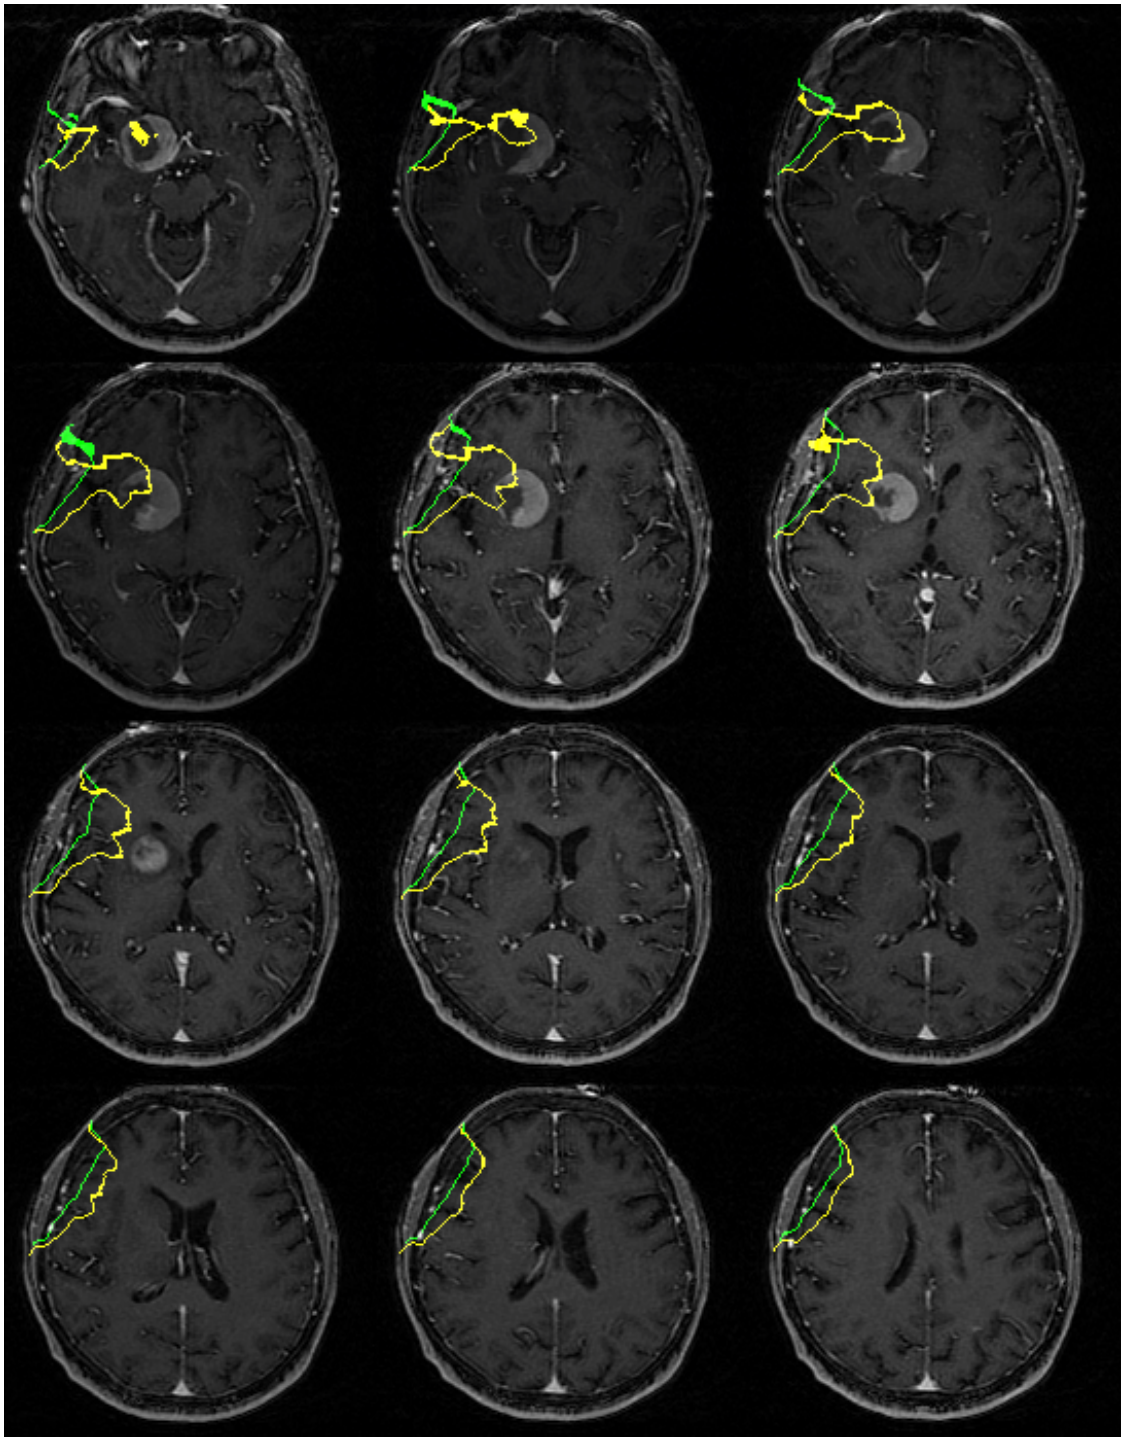

Figure 7.2: Reconstructed cortical surface with the resection cavity (yellow) plotted in registered pMR slices. Note that the cavity extends to the tumor region in the preoperative volume. The cortical surface after dural opening is also shown (green) for comparison. Note the significant tissue motion after tumor resection and that the reconstructed bony edge of the craniotomy, as expected, remains fused with the skull in these images.

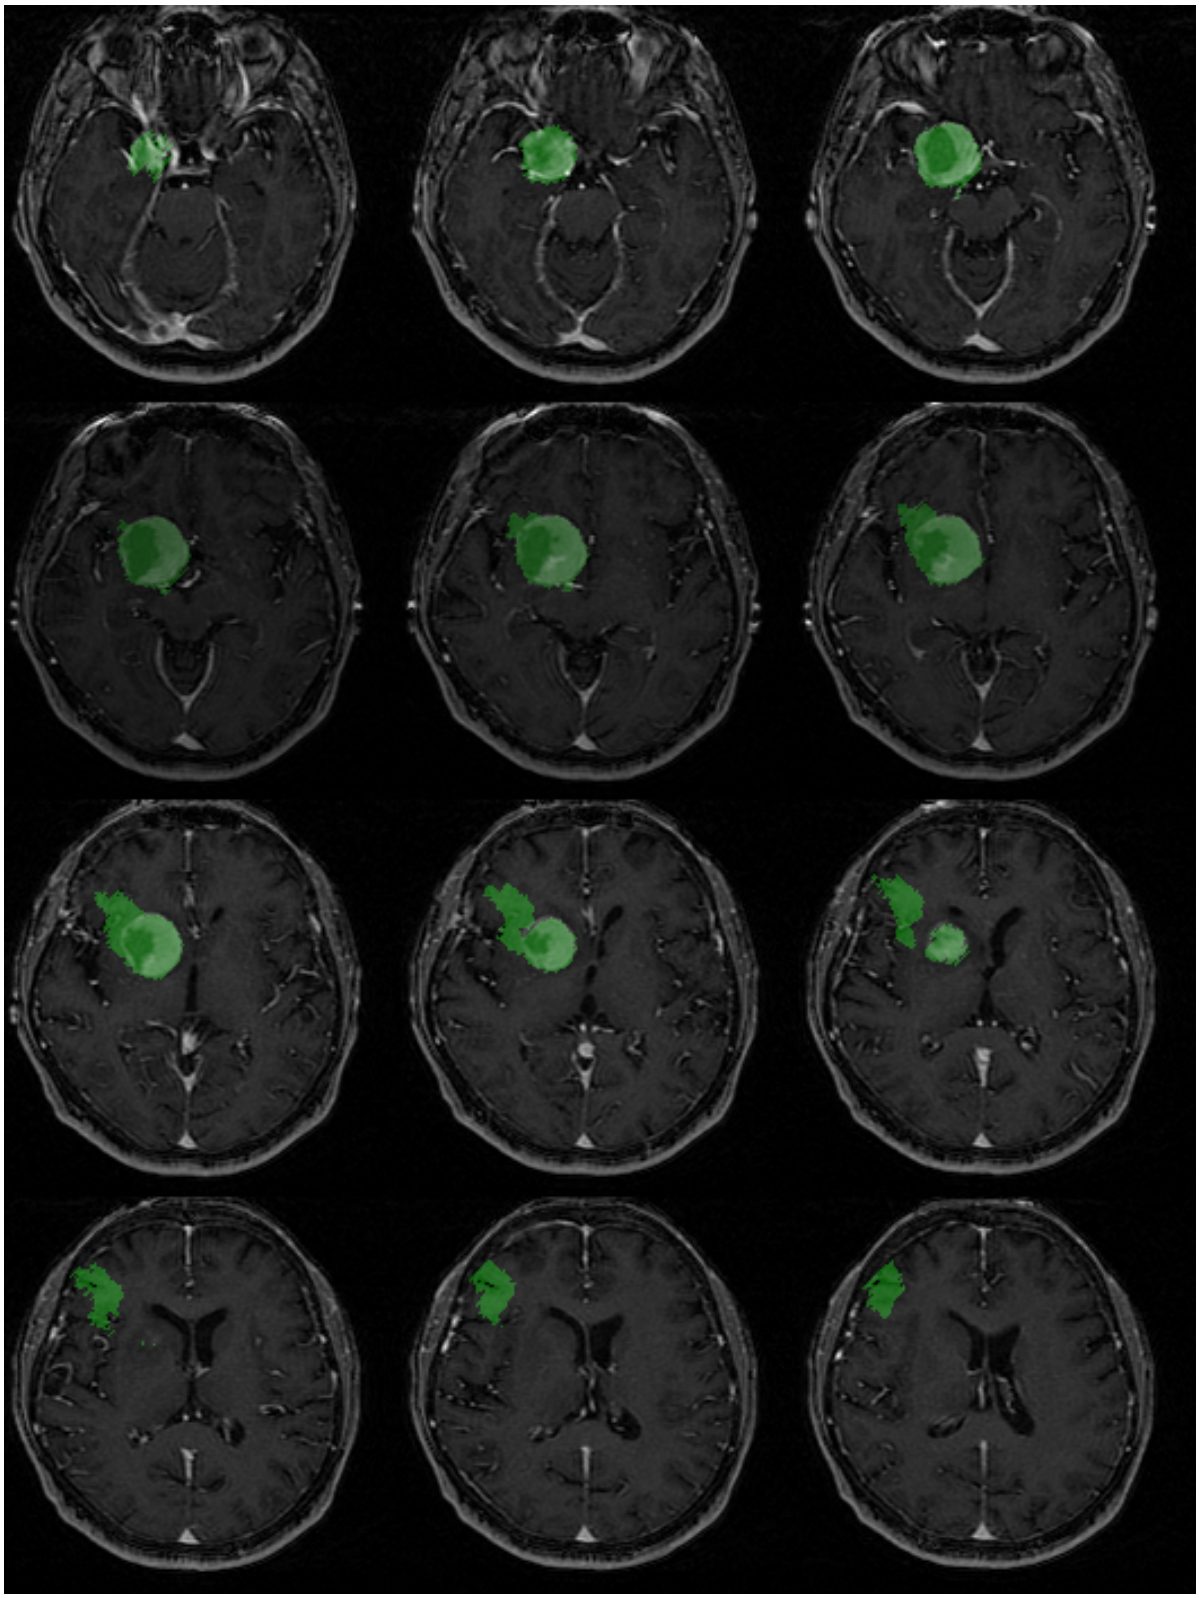

Figure 7.3: Shown are the pMR images with the resected tissue volume painted in green.

7.2 Brain Pulsation

In 1987, Feinberg et al. described what they believed to be the first observations and measurements of human brain motion, which occurs in extensive internal regions (particularly the diencephalon and brainstem) and is synchronous with cardiac systole [23]. This physiologic motion of the brain reflects the response of the brain parenchyma, spinal chord, and CSF to changes in arterial and venous pressures and volumes during the cardiac cycle [27, 19]. The pulsatile motion can also be observed at the cortical surface at the site of craniotomy, but the degree of these motions varies among patients.

To estimate the pulsation of the cortical surface, a video of the surgical scene is acquired at 15 frames per second. Each frame consists of a stereo image pair. The operating microscope is kept stationary during the image acquisition to minimize the tracking error. From the first image pair in the video, several cortical features were chosen to be tracked. As described earlier, the motions of these features were estimated over the duration of several seconds.

Shown in Figure 7.4 are the cortical surfaces captured during surgery 2 and surgery 6 (left and right respectively). Surgery 2 involved a 63-year-old female with a right temporal hemangioma and MCA aneurysm. Surgery 6 involved a 38-year-old male with a cystic glioma in the left temporal lobe of the brain. Three points on each of the cortical surfaces (indicated by a circle, an asterisk and a triangle) were selected for tracking for a duration of approximately two cardiac cycles. The displacements of these features were measured relative to their locations at the beginning of tracking. These displacements are plotted with respect to time in Figure 7.4. The patient in surgery 6 shows a pulsation with a maximum amplitude of nearly 2mm, while the patient in surgery 2 shows practically no motion (an absolute displacement smaller than 0.5mm is likely due to noise in the reconstruction). The factors that influence the pulsation of the brain may include patient's sex, age, tissue pathology, cardiac output, and drug administration. With the ability to track the cortical motion, we now can study the

relationship between these factors and the pulsation of the brain. The variation in brain pulsation may be correlated to the physiological characteristics; thus it may be valuable to quantify this motion.

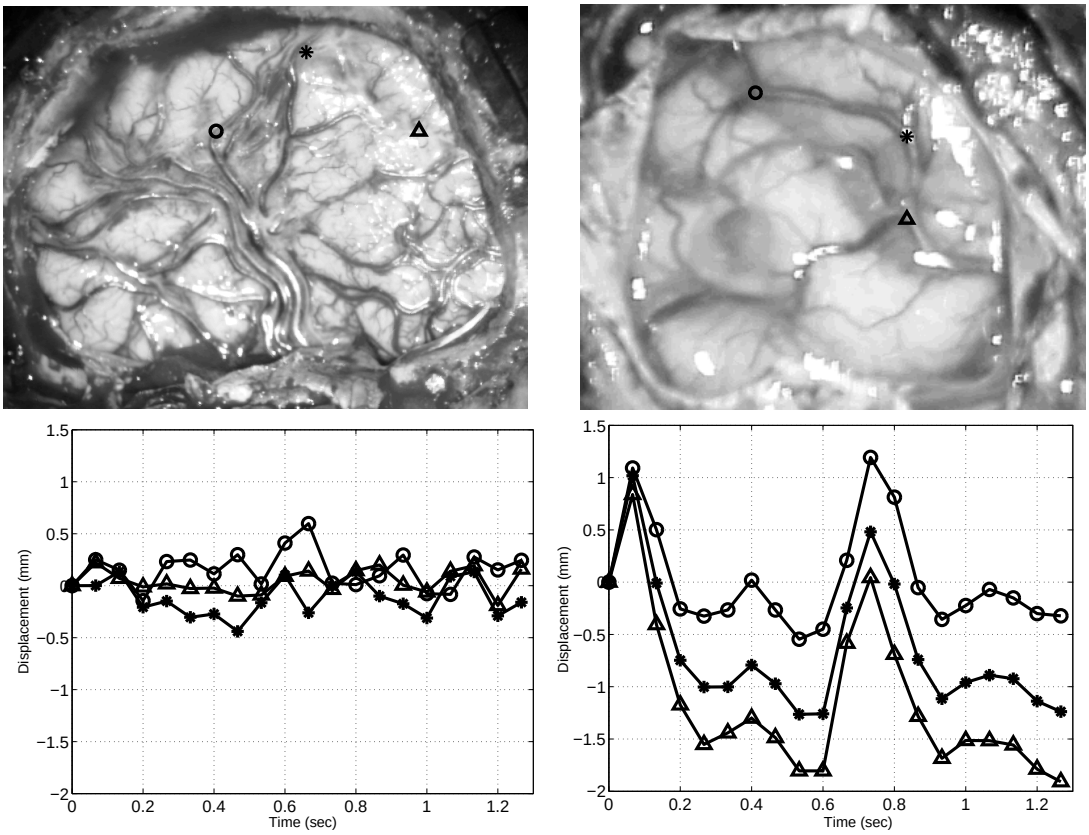

Figure 7.4: Comparison of cortical pulsations between the patient from surgery 2 (upper left) and the patient from surgery 6 (upper right). Three points, indicated by a circle, an asterisk, and a triangle, were tracked over a duration of approximately two cardiac cycles. The pulsatile motion of the patient in surgery 6 (lower right) is much more pronounced (approximately 2mm) than that of the patient in surgery 2 (lower left).

Chapter 8

Discussion

The work described in this thesis is motivated by the concern over the misregistration between the preoperative image studies and the patient anatomy during surgery caused by brain tissue deformation in an open cranium procedure. Due to this misregistration, the preoperative image volume needs to be updated to reflect these changes. A promising strategy in compensating for brain shift combines biomechanical brain deformation models with intraoperative data to predict the full-volume tissue motion. Aiming at gaining the cortical displacement data, we have designed and implemented an intraoperative stereo vision (iSV) system capable of estimating the 3-D shape of the surgical scene in near real-time. In this chapter, we summarize the results from both phantom experiments and clinical studies that support the conclusions that (1) stereopsis is an accurate and effective way to estimate cortical motions and (2) stereopsis-driven deformation models are an accurate and effective means of compensating for brain shift. Following these conclusions, future works to improve the current system are also suggested.

8.1 Stereopsis

Central to the strategy of stereopsis-guided brain shift compensation is the estimation of brain motion during surgery. Similar to other studies using laser range scanners [1, 60] and stereopsis [77], we have also designed a non-contacting strategy that measures the 3-D cortical shape using stereopsis. Our stereo system is incorporated into the standard stereoscopic operating microscope. This system acquires stereo images at 15 frames per second without interrupting the surgical procedure. Given a stereo image pair, the system automatically reconstructs the 3-D shape of the cortical surface every minute (running on a 1.1GHz computer). The motion of the estimated cortical surface is then tracked over time by comparing the different shapes of these surfaces. The estimated cortical motion is used to guide a full brain deformation model, which updates the pMR volume.

As compared to laser range scanners, our system is faster and requires no additional imaging devices. As compared to iMRI or iCT, it is less expensive, much easier to implement, and does not interrupt the surgery during acquisition. As compared to iUS, it is more reliable in finding the same feature for tracking and does not require tissue contact; hence, it does not introduce additional deformations.

To assess the accuracy of the iSV system, we compared the localization of points in 3-D with an optical stylus probe, a well quantified point digitizer. In our first approach, we measured, using the stylus probe, the 3-D location of several micro-screws used as fiducial markers placed on the bony rim of the cranial opening. During ten clinical cases, 38 fiducials were implanted and digitized using both methods. Each case employed an average of 3.8 fiducials, with a maximum of 5 and a minimum of 3. With 1000 matching tokens applied, the difference between the stereoscopically and stylus- estimated 3-D locations of the markers was 1.02mm with a standard error of 0.07mm.

In our second approach, a planar geometrical pattern was constructed to better simulate point objects. The 3-D coordinates of the intersection between the various

geometrical shapes on the pattern were estimated using both stereopsis and the optical stylus. From the 98 intersections, the RMS difference between the stereopsis and stylus estimates was 0.93mm with a standard deviation of 0.41mm. The maximum error was 1.61mm, with a minimum error of 0.20mm. The mean errors observed both in this experiment and the experiments involving micro-screws (0.93mm and 1.02mm, respectively) are close to the accuracy for the optical stylus of 0.87mm.

In order to confirm the viability of our system in estimating a 3-D shape, we compared the accuracy of our stereo system against a CT reconstruction. A brain-shaped phantom was constructed to facilitate this process. A CT scan of the phantom was then obtained, from which the shape of the cortical surface was manually extracted. The phantom was then brought to the OR, and its surface estimated using stereopsis. The mean difference between the shape estimated using stereo and the one detected by the CT was 1.16mm with a maximum of 2.96mm and a minimum of 0.08mm.

In order to address the concern using only a subset of matching tokens for the stereo reconstruction, we measured the density of the reconstructed surface by computing the distance between all neighboring pairs of matching tokens. From ten clinical cases, the point clouds representing the final estimated surfaces contained a mean of 1098 points. The mean distance between neighboring pairs of matching tokens was 1.1mm. Based on these studies, we found that approximately 1000 matching tokens is sufficient to capture the topology of the cortical surface.

From these series of experimental results, we conclude that our iSV system is accurate to within approximately 1mm. This accuracy serves as a solid basis for the tracking of cortical motion, which is performed by comparing the shape of the cortical surface over time.

8.2 Cortical Tracking

The tracking of the cortical surface involves two phases: (1) the first stereoscopically estimated cortical surface in the OR is compared with the cortical surface in the pMR volume to establish the baseline cortical motion; and then (2), the cortical surface is continuously tracked in the OR by comparing the current shape of the cortical surface with the previously estimated shape. In both phases of motion tracking, we employ the iterative closest point (ICP) algorithm to simultaneously establish this correspondence between different shapes of the cortical surface and estimate the displacement.

To quantify the accuracy in tracking of the cortical surface, we again compare results obtained through stereopsis with optical stylus measurements. In a clinical case involving a cyst drainage, three cortical features were digitized using the stylus before and after the cortical deformation and the displacement vectors associated with these features were then computed based on the digitization data. When using stereopsis for tracking these cortical points, we manually located them on the reconstructed surface before cyst drainage. The displacement of each feature was then calculated using ICP and using a simple gravitational sagging model (GS). In this particular case, compared to the stylus estimation, ICP underestimated the cortical displacement by approximately 1.9mm and GS overestimated the cortical displacement by approximately 1.0mm. The quantification of the tracking error is especially challenging because of the lack of ground truth [40]. While the stylus probe defines the world (OR) coordinate system, which is regarded as "ground truth", it also contributes an error component although if well calibrated the probe error is expected to be small. This error, however, may be exacerbated due to the inaccuracies in locating precisely the same position on the cortical surface before and after cyst drainage.

We recognize that the error in tracking the cortical surface in the OR may be higher than either that associated with localizing points in 3-D or that obtained during phantom studies of the cortical shape. In particular, the correspondences between two recon-

structed cortical shapes were established by minimizing the overall Euclidean distance between the two shapes, i.e., the iterative closest point algorithm. The true characteristics of tissue motion may in fact be more complex than that represented by this assumption. For example, the tissue motion induced by resection, as shown in Figure 7.2, is largely non-rigid. Therefore, the ICP algorithm may not be suitable for tracking tissue motion during resection.

8.3 Brain Shift Compensation

The most significant benefit of our system is that the estimated cortical shape and motion is able to guide a full brain deformation model, allowing it to accurately compensate for significant shifts both at and below the cortical surface. For each craniotomy node for which a displacement was estimated, the corresponding finite element equations are precisely enforced by this displacement. Under the guidance of the cortical displacement, the poroelastic brain model generated a displacement field which was then applied to the pMR volume.

In three clinical cases, we show that the model was able to accurately compensate for significant shift at the cortical surface. In two of these clinical cases, intraoperative ultrasound was employed to demonstrate how this correction more accurately compensated for tissue motions in the interior regions of the brain. To quantify this improvement, we manually outlined the tumor in the pMR volume, in the stereopsis-guided uMR volume, in the non-guided uMR volume, and in the corresponding iUS images of one surgery. The mean alignment error of tumor margins between MR and iUS was reduced from 5.2mm between pMR and iUS to 2.8mm between stereo-guided uMR and iUS. The same error, however, increased from 5.2mm to 6.5mm when the brain model was not under stereo guidance. In this particular case, an average error reduction of 3.7mm was observed relative to the model estimate without using intraoperative data.

We also applied the stereopsis-guided brain shift compensation strategy to the deformation induced by retraction. In another clinical example, by monitoring changes in the surgical field every ten seconds, we have estimated the motion of a retractor and its concomitant cortical surface movement. The retractor motion is then used to drive the brain deformation model to update the pMR volume. Using the cortical surface surrounding the retractor as independent evaluative landmarks, we show that our approach can recover on average approximately 75% of the tissue deformation.

From these clinical examples, we conclude that cortical data driven deformation models do hold promise as cost-effecitive and efficient means of providing updates to the preoperative image volume so that the accuracy of an image-guided system can be maintained during surgery.

8.4 Other Types of Cortical Motion

As we have seen from both phantom studies and clinical cases, advantages of our iSV system include its fast acquisition rates, predominantly automated steps, and overall accuracy. These characteristics allow us to capture other important cortical behavior that enhances image-guidance and overall clinical decision-making during surgical procedures.

One of these important cortical changes is the resection cavity. We have employed our iSV system to provide a real-time update of the tissue topology at the craniotomy site referenced to the initially registered pMR. Although it is challenging to recover the tissue motion in response to the resection, the estimated resected volume can be used for the modeling of the tissue resection, a strategy proposed by Miga et al. [59]. Our iSV system is capable of estimating the resected volume intraoperatively.

Another utility of the iSV system is its ability to measure brain pulsatility, whose relationship to both normal and pathological patient conditions warrants further inves-

tigation.

8.5 Future Works

In light of these successes, we are considering several extensions to the system. First, we plan to expedite the process of stereoscopic reconstruction. With some refinement of the code and a faster computer, the stereo reconstruction can be executed at essentially the image acquisition rate of 15 frames per second. The current cortical surface tracking also involves some manual steps, which we plan to eliminate.

There are two ways in which we would like to improve the cortical surface tracking. First, during the first phase stage of cortical tracking, the preoperative mesh is used as the starting point of tracking. The tracking result at this stage relies on the accuracy of both the brain segmentation and registration between the pMR volume and the OR space. Due to the limited MR image resolution and the uncertainty in brain segmentation, the finite element mesh surface may not correspond to the true cortical surface in the pMR volume. Since the brain deformation is minimal prior to opening of the dura, the shape of the cortical surface estimated via stereopsis may be used to improve the tracking accuracy during Phase I. Second, the ICP algorithm is deployed to establish the correspondences between the preoperative mesh and the estimated cortical surface. This registration is constrained to a rigid transformation, and will not fully capture the often non-rigid tissue motion. We are therefore investigating more sophisticated and flexible tracking algorithms that allow us to correlate features in the microscope image with the reconstructed cortical surface.

In the modeling step, the estimated cortical surface displacement is used directly as type I boundary conditions to guide the brain model. This strategy offers a convenient way of constraining the model estimation with intraoperative data and achieves a perfect match between the data and the model prediction at the points of constraint. Studies

conducted by Lunn et al. [50, 48], however, have shown that an exact fit to the displacement data may come at the expense of introducing unrealistic forces at the measurement sites. Furthermore, the inevitable presence of measurement error suggests that enforcing an exact fit to the data may not lead to a correct full-volume displacement estimate. We are evaluating several alternative means of incorporating intraoperative data into the model in order to avoid these limitations. We also plan to combine estimated cortical displacement data with interior brain motion obtained using iUS.

In the modeling of retraction, we plan to further validate our approach through imaging subsurface regions of the brain using co-registered intraoperative ultrasound. In addition, in order to improve the modeling results, we plan to test other more sophisticated boundary conditions at the retraction site, in order to produce more realistic tissue motion such as traction along the retractor blade. Finally, we plan to investigate the role of the material properties of the brain in the finite element model. The ability to continuously monitor tissue behavior during the retraction process has the potential of generating more realistic and patient-specific model parameters, which may improve the accuracy of the model estimate.

We would like to model tissue resection by incorporating the stereoscopically estimated resection volume into the brain model. Note that, in Figure 7.2, the resection cavity (yellow) did not encompass the entire tumor. This is likely due to subcortical tissue deformation in response to the resection, itself, which was not quantified in this case and deserves more attention in the future. We are also working on the synchronization of the brain pulsatile motion with the patient's cardiac cycle and other physiological parameters to investigate its relationship to both normal and pathological patient conditions.

In summary, we have designed, implemented, and tested an end-to-end system for measuring intraoperative cortical motion, and have shown that these measurements can

be used to more accurately compensate for brain shift.

Appendix A

Non-coplanar Calibration

Assuming the camera model described in Section 2.1.1, we need to estimate the camera parameters consist of 11 unknowns (3 rotation angles $(\theta_x, \theta_y, \theta_z)$, 3 translation components (t_x, t_y, t_z), the focal length (f), the lens distortion parameter(κ_1); assuming a first-order approximation, the scale factor (S_x), and the 2 components of the image center (C_x, C_y). A calibration target shown in Figure 2.4 is employed for the calibration. This target is placed in the microscope field of view, as shown in Figure 2.2, and an image is digitized from each camera of the stereo pair. The centroid of each calibration marker is automatically estimated and its spatial coordinates, (x_f, y_f), along with its known 3-D positions, (X, Y, Z) are recorded. The calibration markers are located on three different planes and hence require the non-coplanar calibration scheme. Below is the ten-step non-coplanar calibration procedure developed by Tsai [89].

1. For each calibration point, i, with computer image coordinates (x_{fi}, y_{fi}), determine the distorted sensor coordinates (x_{di}, y_{di}):

$$\begin{aligned} x_{di} &= \frac{N_{cx} d_x}{N_{fx} S_x}(x_{fi} - C_x) \\ y_{di} &= d_y(y_{fi} - C_y), \end{aligned} \tag{A.1}$$

   letting, $S_x = 1$, and (C_x, C_y) be the center pixel of the computer image. These

equations are obtained by solving for (x_{di}, y_{di}) in Equation (2.8) in Section 2.1.1.

2. For each calibration point, i, solve the set of linear equations:

$$\begin{pmatrix} y_{di}X_i \\ y_{di}Y_i \\ y_{di}Z_i \\ y_{di} \\ -x_{di}X_i \\ -x_{di}Y_i \\ -x_{di}Z_i \end{pmatrix}^T \cdot \begin{pmatrix} t_y^{-1}S_x r_1 \\ t_y^{-1}S_x r_2 \\ t_y^{-1}S_x r_3 \\ t_y^{-1}S_x t_x \\ t_y^{-1}r_4 \\ t_y^{-1}r_5 \\ t_y^{-1}r_6 \end{pmatrix} = x_{di} \tag{A.2}$$

Equation (A.2) is a set of linear equations with, seven unknowns $(t_y^{-1}S_x r_1, t_y^{-1}S_x r_2, t_y^{-1}S_x r_3, t_y^{-1}S_x t_x, t_y^{-1}r_4, t_y^{-1}r_5, t_y^{-1}r_6)$. Given seven known points (X_i, Y_i, Z_i) and (x_{di}, y_{di}), the set of linear equations may be solved by inverting the square 7×7 matrix on the left hand side and multiplying by the 7×1 matrix on the right hand side. Given n calibration point $(n > 7)$, the set of over-constrained linear equations can be solved by multiplying the right hand side by the pseudo-inverse of the $n \times 7$ matrix.

3. Solve for the absolute value of t_y:

$$|t_y| = ((t_y^{-1}r_4)^2 + (t_y^{-1}r_5)^2 + (t_y^{-1}r_6)^2)^{-1/2} \tag{A.3}$$

Equation (A.3) stems from the fact that R is a rotation matrix, as such the sum of squares of each column or row is equal to 1 ($[r_4, r_5, r_6]$ constitute the second row of the rotation matrix).

4. Determine the sign of t_y:

   (a) Select a calibration point with image coordinates (x_{fi}, y_{fi}) and world coordinates (X_i, Y_i, Z_i) whose image coordinates are *not* near the center (C_x, C_y).

(b) Let the sign of t_y be $+1$.

(c) Calculate the values x_t and y_t as given in Equation (A.4).

(d) If x_t and x_{fi} have the same sign *and* y_t and Y_i have the same sign, then the sign of t_y is $+1$, otherwise, the sign of t_y is -1.

$$\begin{aligned} x_t &= r_1 X_i + r_2 Y_i + r_3 Z_i + t_x \\ y_t &= r_4 X_i + r_5 Y_i + r_6 Z_i + |t_y|, \end{aligned} \tag{A.4}$$

where,

$$\begin{aligned} r_1 &= (t_y^{-1} S_x r_1)|t_y| \\ r_2 &= (t_y^{-1} S_x r_2)|t_y| \\ r_3 &= (t_y^{-1} S_x r_3)|t_y| \\ r_4 &= (t_y^{-1} r_4)|t_y| \\ r_5 &= (t_y^{-1} r_5)|t_y| \\ r_6 &= (t_y^{-1} r_6)|t_y| \\ t_x &= (t_y^{-1} S_x t_x)|t_y| \end{aligned}$$

5. Solve for S_x:

$$S_x = \sqrt{(t_y^{-1} S_x r_1)^2 + (t_y^{-1} S_x r_2)^2 + (t_y^{-1} S_x r_3)^2}|t_y| \tag{A.5}$$

Note that once again, the sum of square of $[r_1, r_2, r_3]$ is equal to 1.

6. Solve for the rotation matrix R:

$$\begin{aligned} r_1 &= (t_y^{-1} S_x r_1) t_y / S_x \\ r_2 &= (t_y^{-1} S_x r_2) t_y / S_x \\ r_3 &= (t_y^{-1} S_x r_3) t_y / S_x \\ r_4 &= (t_y^{-1} r_4) t_y \end{aligned}$$

$$r_5 = (t_y^{-1} r_5) t_y$$
$$r_6 = (t_y^{-1} r_6) t_y \tag{A.6}$$

These equations come directly from the unknown in the set of linear equations in Equation (A.2), where, t_y and S_x are given in Equations(A.3) and (A.5), respectively. The elements r_1 through r_6 constitute the first tow rows of the rotation matrix. Using the orthonormal property of R, the third row $[r_7, r_8, r_9]$ can be computed as the cross product of the first two rows, $[r_1, r_2, r_3]$ and $[r_4, r_5, r_6]$:

$$r_7 = r_2 r_6 - r_3 r_5$$
$$r_8 = r_3 r_4 - r_1 r_6$$
$$r_9 = r_1 r_5 - r_2 r_4 \tag{A.7}$$

7. Solve for t_x:

$$t_x = (t_y^{-1} S_x t_x) t_y / S_x, \tag{A.8}$$

where, t_y and S_x are given in Equations (A.3) and (A.5), respectively.

8. Determine an approximation for f and t_z:

$$(y_{ti} - d_y x_{di}) \begin{pmatrix} f \\ t_z \end{pmatrix} = w_i d_y y_{di}, \tag{A.9}$$

where,

$$y_{ti} = r_4 X_i + r_5 Y_i + r_6 Z_i + t_y$$
$$w_i = r_7 X_i + r_8 Y_i + r_9 Z_i$$

Equation (A.9) stems from the perspective projection equations, where lens distortion (κ) is assumed to be zero. Similar to Equation (A.2) is a set of linear equations with, two unknowns (f, T_z). Given two known points (X_i, Y_i, Z_i) and (x_{di}, y_{di}), the set of linear equations may be solved by inverting the square 2×2 matrix on the left hand side and multiplying by the 2×1 matrix on the right

hand side. Given n calibration points ($n > 2$), the set of over-constrained linear equations can be solved by multiplying the right hand side by the pseudo-inverse of the $n \times 2$ matrix.

9. Compute exact solution for f, T_z, and κ_1: Using the approximation of f and T_z from the previous step, and $\kappa_1 = 0$ as initial values, minimize the following error function, over n calibration points, using a gradient descent optimization:

$$\sum i = 1n(x_{1i} - x_{2i})^2 + (y_{1i} - y_{2i})^2, \tag{A.10}$$

where, (x_{1i}, y_{1i}) are calculated using the inverse transformation: image to distorted sensor and distorted sensor to undistorted sensor; (x_{2i}, y_{2i}) are calculated using the transformations: world to camera and camera to undistorted sensor.

10. Compute image center (C_x, C_y): Using the approximation of image center as the center pixel of the computer image, minimize the error function given by Equation (A.10).

Appendix B

Real-time Correction Scheme

In Chapter 2, we described in detail the calibration steps for establishing the rigid registration between the 3-D pMR volume and the 2-D microscope view. It might appear that the calibration required for the registration is strictly an off-line process. This system, however, is inherently dynamic in nature. We have found that, a one-time calibration is simply insufficient to provide the necessary accuracy over the duration of a surgery. Shown in Table B.1 is a comparison of the registration errors from a phantom and clinical study. The phantom study consists of a mannequin head with 20 metal beads implanted on its surface. A CT scan of this phantom is first obtained. The phantom is then brought into the surgical room where the position of the beads are determined with the Polaris. A real-time 2-D image is then acquired through the microscope and the position of each bead is measured by hand. Shown in column one of Table B.1 is the average error in camera calibration. Shown in column two is the average error between the 2-D position of the metal beads and the result of passing the 3-D points in patient space through the transformation ${}^{I}T_{M}\ {}^{M}T_{W}\ {}^{W}T_{P}$. Shown in column three is a similar error, but where the points from the 3-D pre-operative space are passed through the transformation ${}^{I}T_{M}\ {}^{M}T_{W}\ {}^{W}T_{P}\ {}^{P}T_{D}$. The clinical study consists of a similar procedure but with a patient undergoing surgery. The degradation in accuracy in the clinical case

Study	Camera Calibration RMS Error (pixels)	Registration with World RMS Error (mm)	Registration with 3D data RMS Error (mm)
Phantom	0.023±0.005	0.39±0.15	1.37±1.31
Clinical	0.035±0.012	1.767±0.366	4.56±2.17

Table B.1: Registration errors for phantom and clinical studies.

post camera calibration is evident in Table B.1.

Sample images from a phantom and a clinical study are presented in Figure B.1. In this figure, an asterisk indicates the true position of a feature point and the open circle denotes the estimated location. Figure B.1(a) shows the camera calibration results for a phantom study. Figure B.1(b) shows the overlay of the true position of the metal beads with the result of passing the 3-D points in patient space through the transformation ${}^{I}T_{M}\ {}^{M}T_{W}\ {}^{W}T_{P}$. Figure B.1(c) shows the camera calibration outcomes for a clinical case. Figure B.1(d) shows the corresponding overlay of the true position of surface feature points and the result of passing the 3-D points in patient space through the same 3-D to 2-D transformation. Note that while the camera calibration errors between the phantom and clinical studies are comparable, the clinical registration errors are significantly larger. This increase is due to a number of factors, for example, motion of the microscope between calibration and surgery, the longer duration of the surgery (approximately 5 hours), or inadvertent jostling of the Polaris. In order to contend with this system drift in the calibration we have developed an adaptive procedure to re-estimate, in real-time, the overall system calibration.

We have developed an adaptive procedure to contend with the calibration errors that accrue during a long surgical procedure. To begin, a few 1mm micro-screws are

Study	Bone Markers RMS Error (mm)	Cortical Features RMS Error (mm)	TRE (mm)
Before Correction	1.767±0.366	1.823±0.655	3.223±0.989
After Correction	0.146±0.072	0.513±0.224	0.623±0.264

Table B.2: Comparison of errors before and after correction scheme. All the errors are from the registration between the patient space and the real-time image.

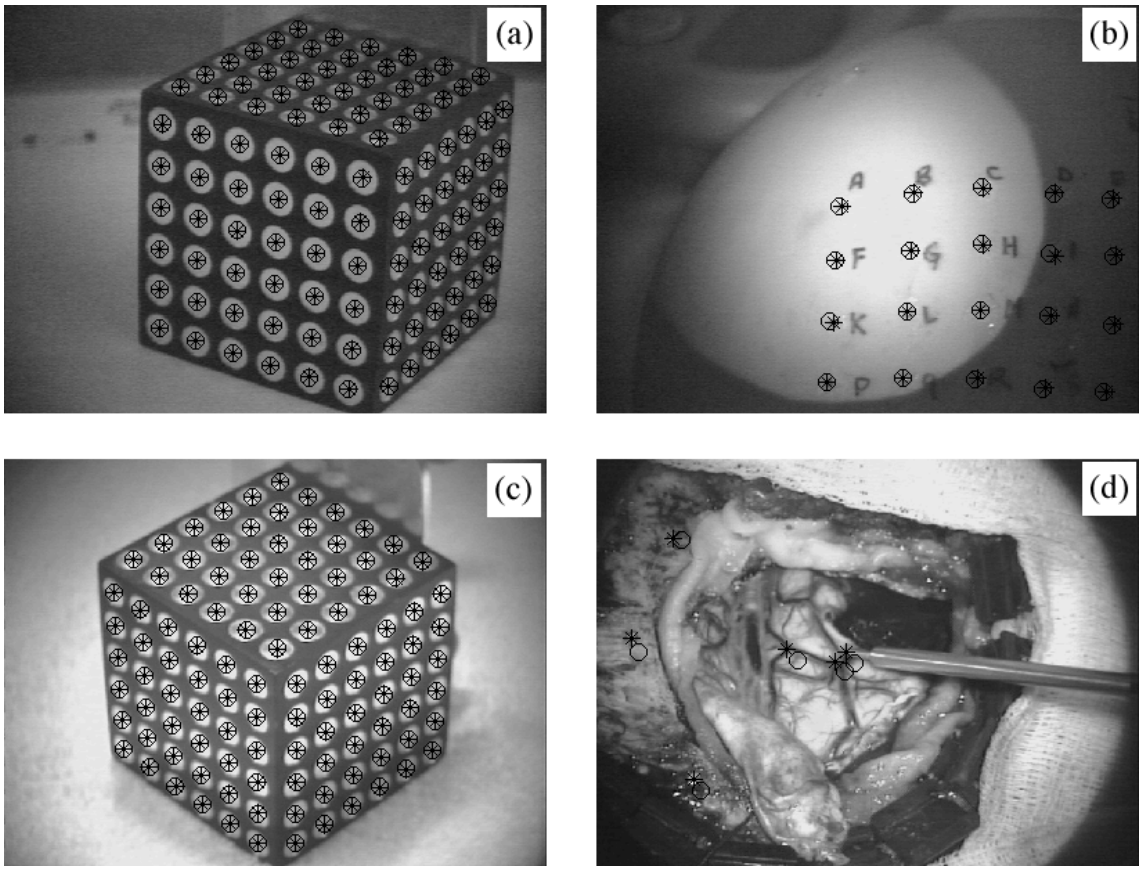

Figure B.1: Comparison between phantom (a)-(b) and clinical (c)-(d) studies. Asterisks denote the true position of feature points, and open circles denote estimated positions. Note that while the calibration errors are small for both the phantom and clinical study, (a) and (c), the registration errors are significantly higher in the clinical compared to the phantom study, (d) and (b), respectively.

implanted on the bony rim of the craniotomy site. These screws act as fiducial markers that are always in the microscope view. We note that since brain tissue undergoes significant deformation during surgery, cortical surface features are typically unreliable fiducial markers for rigid registration.

The 3-D position of the markers are continuously and automatically tracked with the Polaris and their 2-D position in the real-time microscope image is determined manually. The matrix ${}^{M}T_{W}$ is updated by minimizing the RMS error between the projected position of the 3-D markers and their measured image location. This update is performed several times throughout the course of the surgery. As we are interested

in minimizing the system drift, we assume that we can correct all parts of the system by updating the single matrix ${}^{M}T_{W}$. The minimization of this matrix is inherently non-linear, and a standard gradient descent minimization is employed. At the end of surgery, the implanted micro-screws are easily removed from the patient's skull.

Several cortical features on the surface of the brain, usually junctions of cortical veins, are chosen to test the performance of the correction scheme. The location of these cortical features are determined and tracked (in world coordinates) using the Polaris [34]. These points are then passed through the entire system transformation with and without the dynamically updated calibration. Shown in Table B.2 is a comparison between the registration errors before and after correction. In this table the errors for the fiducial markers and cortical features are compared separately. There is an average RMS reduction of 92% (from 1.767 to 0.146 mm) for the markers, and 72% (from 1.823 to 0.513 mm) for the cortical features. Also shown in this table is the target registration error (TRE) [26]. This metric is used here to show that the RMS error for cortical features is smaller than the calculated TRE with respect to the fiducials. Namely, the improvement in accuracy after correction is better than our least expectation. Shown in Figure B.2 are two sample microscope images with several virtual points overlayed. In this figure, asterisks denote the true position of feature points, triangles the estimated position before correction, and circles the estimated position after correction. After correction, the positions are nearly in perfect alignment.

It is worth pointing out that our goal is not to obtain an accurate *physical* account of the overall system drift. Rather, it is to reduce the final registration error and increase the robustness of microscope-guided neurosurgery. To this end, we account for overall system drift through the single transformation matrix ${}^{M}T_{W}$ (specifically, the extrinsic camera calibration portion of this matrix). This matrix is updated periodically throughout surgery by minimizing the positioning error of known fiducial markers. This straight-forward approach has proven quite effective in several lengthy surgeries, as illustrated in Figure B.2 and Table B.2.

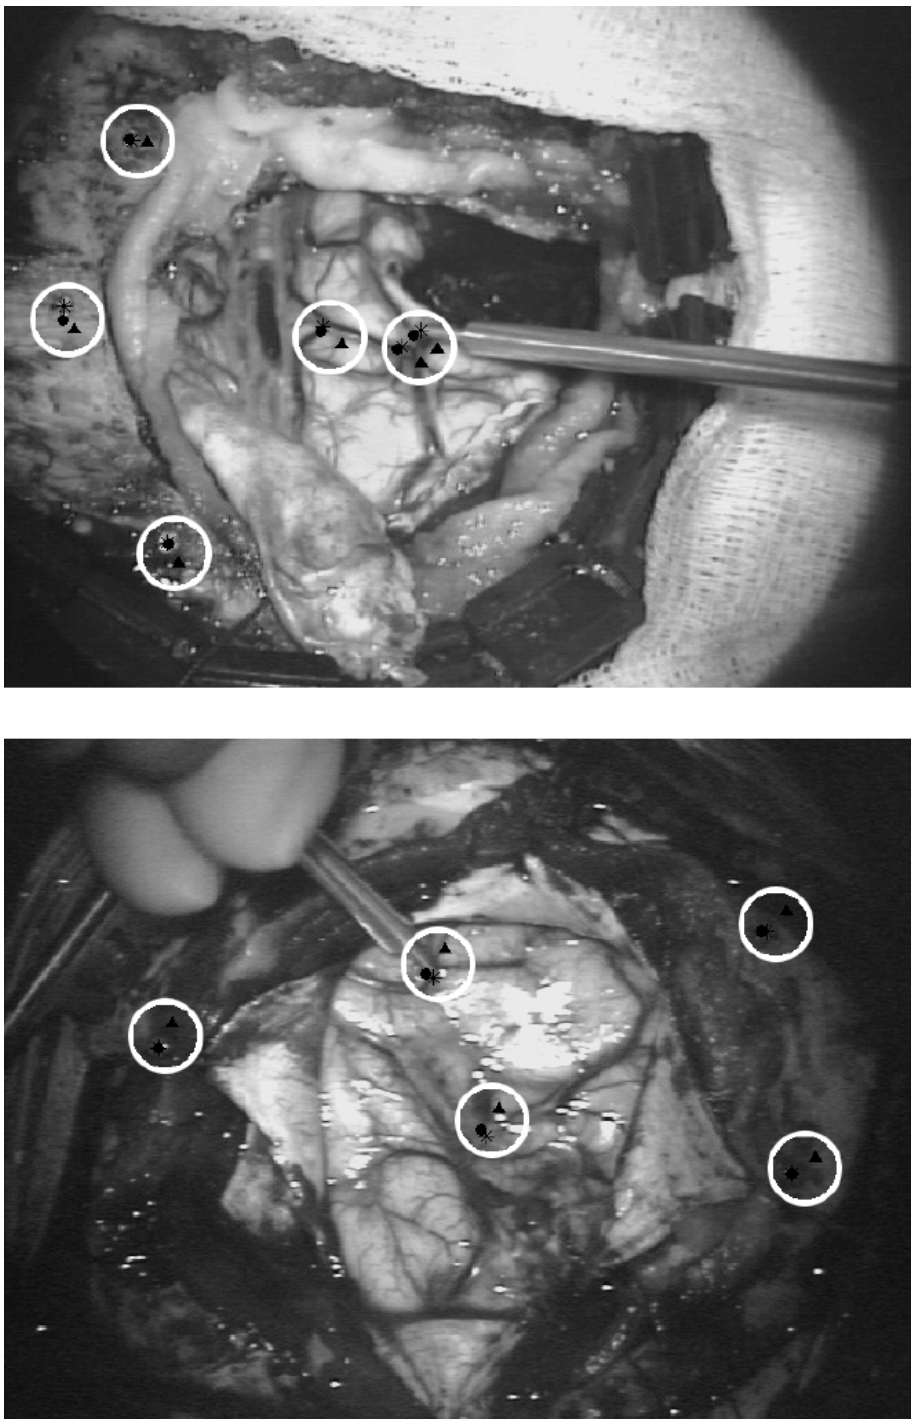

Figure B.2: Sample microscope views with an overlay of virtual points. Asterisks denote the true position of feature points, triangles the estimated position before correction, and circles the estimated position after correction.

In addition to providing more accurate registration, this technique has applications in evaluating non-rigid brain tissue deformation. Clinical conditions (gravity, loss of CSF, etc.) can cause significant non-rigid shifts in the brain [16]. For example, gravitational shifts alone can account for an average of 10mm of intraoperative displacement [70]. This tissue deformation compromises the alignment accuracy of microscope-guided neurosurgery systems. Edwards et al. [18] have shown that when studies move from phantoms to patients, registration errors increase from 0.3-0.5 mm to 3-4 mm. Although intraoperative tissue movement probably contributes to this increase, it is almost certainly not the only component of the increase in error account for. To model non-rigid deformation, it is essential to have an accurate rigid calibration. Since this calibration drift can be substantial, it is critical to update, in real-time, the initial system parameters.

Appendix C

Rectification

The goal of rectification of a stereo pair is to warp the stereo images so that the epipolar lines correspond to one of the image axes, typically image rows. The importance of rectification is that the identification of the epipolar lines becomes a trivial task after the image pair are rectified. An image pair can be rectified without knowing the camera parameters, but the calibration of a stereo pair leads a more straight forward solution to the rectification. In our case, the stereo pair is calibrated with respect to a common coordinate frame, therefore the camera positions are known relative to one another.

Let us first review the formation of the original image pair. As shown in Figure C.1(left), a point (X, Y, Z) in the world reference frame is projected, through two pinhole cameras centered at C_l and C_r, onto the image sensors at locations (x_l, y_l) and (x_r, y_r). The image formation of both left an right camera can be described by two perspective projection matrices (PPMs), P_l and P_r, respectively. It is important to point out that the nonlinear portion of the camera model including the radial lens distortion and the unit conversion of image axes is not relevant to the discussion of the rectification. This is due to the fact that the epipolar geometry is solely based on the pinhole projection model. As a result of this assumption, the projected images coordinates, (x, y), are computed from the original (distorted) image coordinates, (x_d, y_d) before the

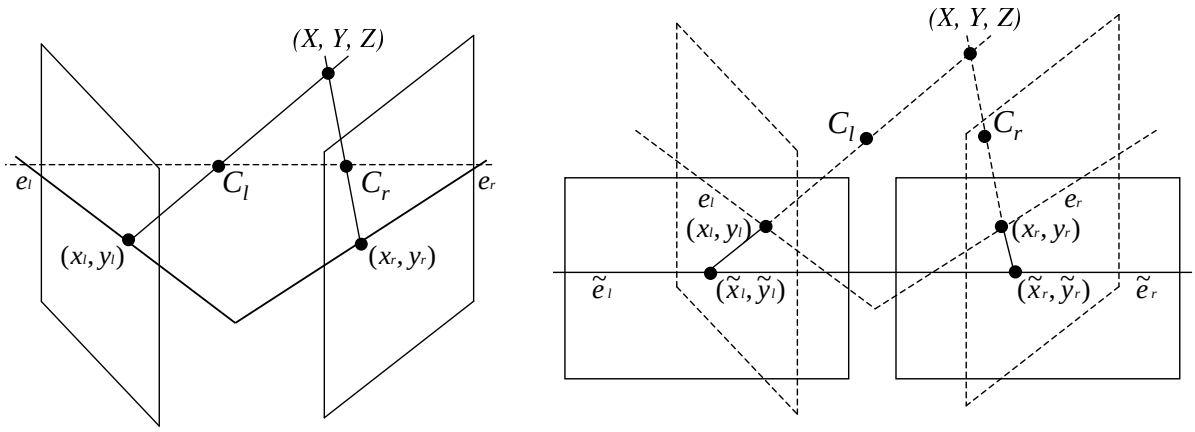

Figure C.1: The 3-D vision problem and the epipolar constraint: shown on the left are two pinhole cameras, with centers C_l and C_r, forming the image of a single point. The projection onto the left image, (x_l, y_l), has a corresponding match in the right image, (x_r, y_r), lying on the epipolar line e_r. Shown on the right is the rectified image so that the epipolar lines are coincident with image scan lines (solid lines).

rectification takes place. After rectification, the epipolar lines, e_l and e_r, are aligned with the image raster lines, as shown in Figure C.1(right). The goal of rectification is to compute two new PPMs, P_{lr} and P_{rr}, that describes the the formation of the rectified image pair, and hence, these new PPMs are often referred to as the rectifying PPMs.

Knowing P_l and P_r from camera calibration, we can determine the rectifying PPM P_lr and P_rr by applying several constraints imposed by the characteristics of the final rectified image pair. Before we delve into the derivation of the rectifying PPMs, let us first use the original PPM of an pinhole camera to depict its imaging geometry. This portion of our discussion is applicable to both left and right cameras, and hence we omit the subscription which indicates either the left or right camera. As described in Section 2.1.1, the PPM of a pinhole camera, P, is a 3×4 matrix:

$$P = \begin{pmatrix} p_{11} & p_{12} & p_{13} & p_{14} \\ p_{21} & p_{22} & p_{23} & p_{24} \\ p_{31} & p_{32} & p_{33} & p_{34} \end{pmatrix} \tag{C.1}$$

Using P, we can derive the mathematical expressions of imaging geometry associated with this camera. For example, any point (X, Y, Z) on the focal plane must satisfy the

following equation:

$$\begin{pmatrix} p_{31} & p_{32} & p_{33} & p_{34} \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix} = 0, \tag{C.2}$$

since a point on the focal plane is projected to infinity ($s = 0$). The optical center C, (X_c, Y_c, Z_c), is the intersection of three planes: the focal plane, the two planes that intersect the image plane in the vertical (x = 0) and horizontal (y = 0) axis of the image coordinates, respectively; therefore,

$$P \begin{pmatrix} X_c \\ Y_c \\ Z_c \\ 1 \end{pmatrix} = 0, \tag{C.3}$$

and

$$\begin{pmatrix} X_c \\ Y_c \\ Z_c \end{pmatrix} = - \begin{pmatrix} p_{11} & p_{12} & p_{13} & p_{14} \\ p_{21} & p_{22} & p_{23} & p_{24} \\ p_{31} & p_{32} & p_{33} & p_{34} \end{pmatrix}^{-1} \begin{pmatrix} p_{14} \\ p_{24} \\ p_{34} \end{pmatrix}. \tag{C.4}$$

The optical ray associated to an image point (x, y) is the line $\langle C, (x, y) \rangle$. Any point on this ray, (X, Y, Z) must satisfy the following:

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} X_c \\ Y_c \\ Z_c \end{pmatrix} + \lambda P^{-1} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix}, \tag{C.5}$$

where λ is an arbitrary real number. With these expressions, we next derive the epipolar geometry.

To do so, we return to the stereo pair. The PPMs of the left and the right cameras are denoted as P_l and P_r, respectively and,

$$P_l = \begin{pmatrix} l_{11} & l_{12} & l_{13} & l_{14} \\ l_{21} & l_{22} & l_{23} & l_{24} \\ l_{31} & l_{32} & l_{33} & l_{34} \end{pmatrix},$$

$$P_r = \begin{pmatrix} r_{11} & r_{12} & r_{13} & r_{14} \\ r_{21} & r_{22} & r_{23} & r_{24} \\ r_{31} & r_{32} & r_{33} & r_{34} \end{pmatrix}. \tag{C.6}$$

The optical center C_l and C_r are:

$$\begin{pmatrix} X_{cl} \\ Y_{cl} \\ Z_{cl} \end{pmatrix} = - \begin{pmatrix} l_{11} & l_{12} & l_{13} & l_{14} \\ l_{21} & l_{22} & l_{23} & l_{24} \\ l_{31} & l_{32} & l_{33} & l_{34} \end{pmatrix}^{-1} \begin{pmatrix} l_{14} \\ l_{24} \\ l_{34} \end{pmatrix}$$
$$\begin{pmatrix} X_{cr} \\ Y_{cr} \\ Z_{cr} \end{pmatrix} = - \begin{pmatrix} r_{11} & r_{12} & r_{13} & r_{14} \\ r_{21} & r_{22} & r_{23} & r_{24} \\ r_{31} & r_{32} & r_{33} & r_{34} \end{pmatrix}^{-1} \begin{pmatrix} r_{14} \\ r_{24} \\ r_{34} \end{pmatrix}, \tag{C.7}$$

respectively. For an image point (x_l, y_l) in the left image, its associated optical ray, $\langle C_l, (x_l, y_l) \rangle$, can be expressed as:

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} X_{cl} \\ Y_{cl} \\ Z_{cl} \end{pmatrix} + \lambda P^{-1} \begin{pmatrix} x_l \\ y_l \\ 1 \end{pmatrix}, \tag{C.8}$$

where λ is an arbitrary real number. Similarly, the optical ray, $\langle C_r, (x_r, y_r) \rangle$, can be expressed as:

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \begin{pmatrix} X_{cr} \\ Y_{cr} \\ Z_{cr} \end{pmatrix} + \lambda P^{-1} \begin{pmatrix} x_r \\ y_r \\ 1 \end{pmatrix}, \tag{C.9}$$

where λ is another arbitrary real number. The two optical rays are used to define epipolar lines.

As shown in Figure 3.4, the corresponding point (x_r, y_r) of the point (x_l, y_l) is constrained to lie on the epipolar line. This line, $\langle e_r, (x_r, y_r) \rangle$, is defined to be the projection of the optical ray $\langle (x_l, y_l), C_l \rangle$ through the camera center C_r. Its conjugate epipolar line, $\langle e_l, (x_l, y_l) \rangle$, is the projection of the optical ray $\langle (x_r, y_r), C_r \rangle$ through the camera center C_l.

Furthermore, all the epipolar lines in one image plane (e.g., in the right image) pass through a common point called the epipole (e_r), which is the projection of the conjugate

optical center (the projection of C_l through C_r). Mathematically,

$$\begin{pmatrix} sx_{er} \\ sy_{er} \\ s \end{pmatrix} = P_r \begin{pmatrix} X_{cl} \\ Y_{cl} \\ Z_{cl} \\ 1 \end{pmatrix},$$
$$\begin{pmatrix} sx_{el} \\ sy_{el} \\ s \end{pmatrix} = P_l \begin{pmatrix} X_{cr} \\ Y_{cr} \\ Z_{cr} \\ 1 \end{pmatrix}. \tag{C.10}$$

The coordinates of the two camera centers, (X_{cl}, Y_{cl}, Z_{cl}) and (X_{cr}, Y_{cr}, Z_{cr}), are given by Equation (C.4). For a point in the left image, (x_l, y_l), knowing its optical ray, $\langle (x_l, y_l), C_l \rangle$ (Equation (C.8)), and the epipole on the right image, e_r (Equation (C.10)), we can write the parametric equation of its epiploar line, $\langle e_r, (x_r, y_r) \rangle$, as follows:

$$\begin{pmatrix} x_s \\ y_s \\ s \end{pmatrix} = \begin{pmatrix} x_{er} \\ y_{er} \\ 1 \end{pmatrix} + \lambda P_r P_l^{-1} \begin{pmatrix} x_l \\ y_l \\ 1 \end{pmatrix}. \tag{C.11}$$

Similarly, the conjugate epiploar line, $\langle e_l, (x_l, y_l) \rangle$, can be expressed as follows:

$$\begin{pmatrix} x_s \\ y_s \\ s \end{pmatrix} = \begin{pmatrix} x_{el} \\ y_{el} \\ 1 \end{pmatrix} + \lambda P_l P_r^{-1} \begin{pmatrix} x_r \\ y_r \\ 1 \end{pmatrix}. \tag{C.12}$$

Next, we will compute the rectifying PPMs, P_{lr} and P_{rr} as the solutions of a system of equations, which express the constraints arising from rectifying requirements plus other constraints necessary to ensure a unique solution. Let

$$P_{lr} = \begin{pmatrix} lr_{11} & lr_{12} & lr_{13} & lr_{14} \\ lr_{21} & lr_{22} & lr_{23} & lr_{24} \\ lr_{31} & lr_{32} & lr_{33} & lr_{34} \end{pmatrix},$$
$$P_{rr} = \begin{pmatrix} rr_{11} & rr_{12} & rr_{13} & rr_{14} \\ rr_{21} & rr_{22} & rr_{23} & rr_{24} \\ rr_{31} & rr_{32} & rr_{33} & rr_{34} \end{pmatrix}. \tag{C.13}$$

They will be computed by applying the following requirements and constraints:

1. Common Focal plane. If cameras share the same focal plane, the common retinal plane is constrained to be parallel to the baseline and epipolar lines are parallel. The two rectifying PPM have the same focal plane if only if:

$$\begin{aligned} lr_{31} &= rr_{31} \\ lr_{32} &= rr_{32} \\ lr_{33} &= rr_{33} \\ lr_{34} &= rr_{34} \end{aligned} \tag{C.14}$$

2. Position of the optical centers. The optical centers of the rectifying perspective projections must be the same as those of the original projections:

$$P_{lr} \begin{pmatrix} X_{cl} \\ Y_{cl} \\ Z_{cl} \\ 1 \end{pmatrix} = 0 \quad \text{and} \quad P_{rr} \begin{pmatrix} X_{cr} \\ Y_{cr} \\ Z_{cr} \\ 1 \end{pmatrix} = 0, \tag{C.15}$$

Combining the equation above with Equation (C.7) and (C.13) yields six linear constraints:

$$\begin{aligned} lr_{11}X_{cl} + lr_{12}Y_{cl} + lr_{13}Z_{cl} + lr_{14} &= 0 \\ lr_{21}X_{cl} + lr_{22}Y_{cl} + lr_{23}Z_{cl} + lr_{24} &= 0 \\ lr_{31}X_{cl} + lr_{32}Y_{cl} + lr_{33}Z_{cl} + lr_{34} &= 0 \\ rr_{11}X_{cr} + rr_{12}Y_{cr} + rr_{13}Z_{cr} + rr_{14} &= 0 \\ rr_{21}X_{cr} + rr_{22}Y_{cr} + rr_{23}Z_{cr} + rr_{24} &= 0 \\ rr_{31}X_{cr} + rr_{32}Y_{cr} + rr_{33}Z_{cr} + rr_{34} &= 0 \end{aligned} \tag{C.16}$$

3. Alignment of conjugate epipolar lines. The vertical coordinate of the projection of a 3-D point onto the rectified image plane must be the same in both image:

$$\frac{lr_{21}X + lr_{22}Y + lr_{23}Z + lr_{24}}{lr_{31}X + lr_{32}Y + lr_{33}Z + lr_{34}} = \frac{rr_{21}X + rr_{22}Y + rr_{23}Z + rr_{24}}{rr_{31}X + rr_{32}Y + rr_{33}Z + rr_{34}}. \tag{C.17}$$

Combining the equation above with the constraints in Equation (C.14) yields:

$$\begin{aligned} lr_{21} &= rr_{21} \\ lr_{22} &= rr_{22} \\ lr_{23} &= rr_{23} \\ lr_{24} &= rr_{24} \end{aligned} \tag{C.18}$$

The equations written to this point are sufficient to guarantee rectification, but the orientation of the rectified image plane has still one degree of freedom. In other words, the constraints written up to now are not enough to obtain a unique set of PPMs. Next, we will introduce several other arbitrary but reasonable constraints to ensure a unique solution:

1. Orientation of the rectified image plane. We choose the rectifying focal planes to be parallel to be intersection of the two original focal planes, i.e.,

$$\begin{pmatrix} lr_{31} & lr_{32} & lr_{33} \end{pmatrix} \left[\begin{pmatrix} l_{31} & l_{32} & l_{33} \end{pmatrix}^\top \times \begin{pmatrix} r_{31} & r_{32} & r_{33} \end{pmatrix}^\top \right] = 0, \tag{C.19}$$

   where $\begin{pmatrix} l_{31} & l_{32} & l_{33} \end{pmatrix}$ and $\begin{pmatrix} r_{31} & r_{32} & r_{33} \end{pmatrix}$ are the third rows of P_l and P_r respectively. Notice that the corresponding equation,

$$\begin{pmatrix} rr_{31} & rr_{32} & rr_{33} \end{pmatrix} \left[\begin{pmatrix} l_{31} & l_{32} & l_{33} \end{pmatrix}^\top \times \begin{pmatrix} r_{31} & r_{32} & r_{33} \end{pmatrix}^\top \right] = 0,$$

   is redundant because of Equation (C.14).

2. Orthogonality of the rectifying reference frames. The intersections of the image plane with the planes:

$$\begin{pmatrix} lr_{11} & lr_{12} & lr_{13} & lr_{14} \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix} = 0 \quad \text{and} \quad \begin{pmatrix} lr_{21} & lr_{22} & lr_{23} & lr_{24} \end{pmatrix} \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix} = 0, \tag{C.20}$$

correspond to the two image axis of the rectified image plane. In order for this reference frame to be orthogonal, the planes must be perpendicular to each other. With equations in C.18, we have,

$$(\, lr_{11} \quad lr_{12} \quad lr_{13} \quad lr_{14} \,) \begin{pmatrix} lr_{21} \\ lr_{22} \\ lr_{23} \\ lr_{24} \end{pmatrix} = 0 \quad \text{and} \quad (\, rr_{11} \quad rr_{12} \quad rr_{13} \quad rr_{14} \,) \begin{pmatrix} lr_{21} \\ lr_{22} \\ lr_{23} \\ lr_{24} \end{pmatrix} = 0, \tag{C.21}$$

3. Image Centers. The image centers (x_0, y_0) is given by:

$$.x_0 = (\, lr_{11} \quad lr_{12} \quad lr_{13} \,) \begin{pmatrix} lr_{31} \\ lr_{32} \\ lr_{33} \end{pmatrix} \quad \text{and} \quad y_0 = (\, lr_{21} \quad lr_{22} \quad lr_{23} \,) \begin{pmatrix} lr_{31} \\ lr_{32} \\ lr_{33} \end{pmatrix}. \tag{C.22}$$

We set the two image centers at $(0, 0)$ and use Equations (C.14) and (C.18) to obtain the constraints:

$$\begin{aligned} (\, lr_{11} \quad lr_{12} \quad lr_{13} \,) \begin{pmatrix} lr_{31} \\ lr_{32} \\ lr_{33} \end{pmatrix} &= 0 \\ (\, lr_{21} \quad lr_{22} \quad lr_{23} \,) \begin{pmatrix} lr_{31} \\ lr_{32} \\ lr_{33} \end{pmatrix} &= 0 \\ (\, rr_{11} \quad rr_{12} \quad rr_{13} \,) \begin{pmatrix} lr_{31} \\ lr_{32} \\ lr_{33} \end{pmatrix} &= 0 \end{aligned} \tag{C.23}$$

4. Focal lengths in pixels. The horizontal and vertical focal lengths in pixels are given by:

$$\begin{aligned} f_x &= \| \, (\, lr_{11} \quad lr_{12} \quad lr_{13} \,)^\top \times (\, lr_{31} \quad lr_{32} \quad lr_{33} \,)^\top \, \| \\ f_y &= \| \, (\, lr_{21} \quad lr_{22} \quad lr_{23} \,)^\top \times (\, lr_{31} \quad lr_{32} \quad lr_{33} \,)^\top \, \|. \end{aligned} \tag{C.24}$$

By setting arbitrarily the values of f_x and f_y, we obtain the constraints:

$$\begin{aligned}
\| \begin{pmatrix} lr_{11} & lr_{12} & lr_{13} \end{pmatrix}^\top \times \begin{pmatrix} lr_{31} & lr_{32} & lr_{33} \end{pmatrix}^\top \| &= f_r^2 \\
\| \begin{pmatrix} lr_{21} & lr_{22} & lr_{23} \end{pmatrix}^\top \times \begin{pmatrix} lr_{31} & lr_{32} & lr_{33} \end{pmatrix}^\top \| &= f_r^2 \\
\| \begin{pmatrix} rr_{11} & rr_{12} & rr_{13} \end{pmatrix}^\top \times \begin{pmatrix} lr_{31} & lr_{32} & lr_{33} \end{pmatrix}^\top \| &= f_r^2.
\end{aligned} \tag{C.25}$$

By definitions of vector norm and vector cross product and Equations (C.23), we can rewrite the equation above as:

$$\begin{aligned}
\| \begin{pmatrix} lr_{11} & lr_{12} & lr_{13} \end{pmatrix} \|^2 \| \begin{pmatrix} lr_{31} & lr_{32} & lr_{33} \end{pmatrix} \|^2 &= f_r^2 \\
\| \begin{pmatrix} lr_{21} & lr_{22} & lr_{23} \end{pmatrix} \|^2 \| \begin{pmatrix} lr_{31} & lr_{32} & lr_{33} \end{pmatrix} \|^2 &= f_r^2 \\
\| \begin{pmatrix} rr_{11} & rr_{12} & rr_{13} \end{pmatrix} \|^2 \| \begin{pmatrix} lr_{31} & lr_{32} & lr_{33} \end{pmatrix} \|^2 &= f_r^2.
\end{aligned} \tag{C.26}$$

5. The scale factor. PPMs are defined up to a scale factor, and a common choice to block the latter is to set

$$\| \begin{pmatrix} lr_{31} & lr_{32} & lr_{33} \end{pmatrix} \| = 1 \quad \text{and} \quad \| \begin{pmatrix} rr_{31} & rr_{32} & rr_{33} \end{pmatrix} \| = 1 \tag{C.27}$$

Now, we have established enough constraints that guarantee a unique solution for each of the two rectifying PPMs. To capitulate, we organize the constraints introduced here in the following four systems:

$$\begin{cases}
\begin{pmatrix} lr_{31} & lr_{32} & lr_{33} & lr_{34} \end{pmatrix} \begin{pmatrix} X_{cl} \\ Y_{cl} \\ Z_{cl} \\ 1 \end{pmatrix} = 0 \\
\begin{pmatrix} lr_{31} & lr_{32} & lr_{33} & lr_{34} \end{pmatrix} \begin{pmatrix} X_{cr} \\ Y_{cr} \\ Z_{cr} \\ 1 \end{pmatrix} = 0 \\
\begin{pmatrix} lr_{31} & lr_{32} & lr_{33} \end{pmatrix} \left[\begin{pmatrix} l_{31} \\ l_{32} \\ l_{33} \end{pmatrix} \times \begin{pmatrix} r_{31} \\ r_{32} \\ r_{33} \end{pmatrix} \right] = 0 \\
\| \begin{pmatrix} lr_{31} & lr_{32} & lr_{33} \end{pmatrix} \| = 1
\end{cases} \tag{C.28}$$

$$
\begin{cases}
\begin{pmatrix} lr_{21} & lr_{22} & lr_{23} & lr_{24} \end{pmatrix} \begin{pmatrix} X_{cl} \\ Y_{cl} \\ Z_{cl} \\ 1 \end{pmatrix} = 0 \\
\begin{pmatrix} lr_{21} & lr_{22} & lr_{23} & lr_{24} \end{pmatrix} \begin{pmatrix} X_{cr} \\ Y_{cr} \\ Z_{cr} \\ 1 \end{pmatrix} = 0 \\
\begin{pmatrix} lr_{21} & lr_{22} & lr_{23} \end{pmatrix} \begin{pmatrix} lr_{31} \\ lr_{32} \\ lr_{33} \end{pmatrix} = 0 \\
\left\| \begin{pmatrix} lr_{31} & lr_{32} & lr_{33} \end{pmatrix} \right\| = f_r
\end{cases}
\tag{C.29}
$$

$$
\begin{cases}
\begin{pmatrix} lr_{11} & lr_{12} & lr_{13} & lr_{14} \end{pmatrix} \begin{pmatrix} X_{cl} \\ Y_{cl} \\ Z_{cl} \\ 1 \end{pmatrix} = 0 \\
\begin{pmatrix} lr_{11} & lr_{12} & lr_{13} \end{pmatrix} \begin{pmatrix} lr_{21} \\ lr_{22} \\ lr_{23} \end{pmatrix} = 0 \\
\begin{pmatrix} lr_{11} & lr_{12} & lr_{13} \end{pmatrix} \begin{pmatrix} lr_{31} \\ lr_{32} \\ lr_{33} \end{pmatrix} = 0 \\
\left\| \begin{pmatrix} lr_{11} & lr_{12} & lr_{13} \end{pmatrix} \right\| = f_r
\end{cases}
\tag{C.30}
$$

$$\begin{cases} \begin{pmatrix} rr_{11} & rr_{12} & rr_{13} & rr_{14} \end{pmatrix} \begin{pmatrix} X_{cr} \\ Y_{cr} \\ Z_{cr} \\ 1 \end{pmatrix} = 0 \\ \begin{pmatrix} rr_{11} & rr_{12} & rr_{13} \end{pmatrix} \begin{pmatrix} lr_{21} \\ lr_{22} \\ lr_{23} \end{pmatrix} = 0 \\ \begin{pmatrix} rr_{11} & rr_{12} & rr_{13} \end{pmatrix} \begin{pmatrix} lr_{31} \\ lr_{32} \\ lr_{33} \end{pmatrix} = 0 \\ \| \begin{pmatrix} rr_{11} & rr_{12} & rr_{13} \end{pmatrix} \| = f_r \end{cases} \tag{C.31}$$

plus

$$\begin{cases} lr_{21} = rr_{21} \\ lr_{22} = rr_{22} \\ lr_{23} = rr_{23} \\ lr_{24} = rr_{24} \\ lr_{31} = rr_{31} \\ lr_{32} = rr_{32} \\ lr_{33} = rr_{33} \\ lr_{34} = rr_{34}. \end{cases} \tag{C.32}$$

The first four systems have all the same structure, each one being a 3×4 linear homogeneous system subject to a quadratic constraint, that is,

$$\begin{cases} Ax = 0 \\ \|x'\| = k, \end{cases} \tag{C.33}$$

where x' is a vector composed by the first three components of x, and k is a real number. The four systems above are solved in sequence, top to bottom. The solution of each system is obtained by first computing (e.g., by SVD factorization) a one-parameter family of solutions to $Ax = 0$ of the form $x = \alpha x_0$, where x_0 is a nontrivial solution and α is an arbitrary real number, and then letting α equal to $\frac{k}{\|x'_0\|}$.

Appendix D

Line Intersection in 3-D

In 3-D, two lines often do not intersect and are not parallel. Shown in Figure D.1 are two lines in 3-D, $\langle P_1, P_2 \rangle$ and $\langle P_3, P_4 \rangle$, which are non-intersecting and non-parallel. In this case, "the point of intersection" between these two lines is defined as the point at equal and minimum distance to both lines. This point, P_0, is found by: first, constructing $\langle P_a, P_b \rangle$ so that it intersects and is perpendicular to both $\langle P_1, P_2 \rangle$ and $\langle P_3, P_4 \rangle$; and then, computing the middle point between P_a and P_B. With known P_1 and P_2,

$$P_a = P_1 + \lambda_a (P_2 - P_1), \tag{C-1}$$

and similarly, with known P_3 and P_4,

$$P_b = P_3 + \lambda_b (P_4 - P_3). \tag{C-2}$$

To determine P_a and P_b, we only need to determine λ_a and λ_b, respectively. Since $\langle P_a, P_b \rangle$ is perpendicular to both $\langle P_1, P_2 \rangle$ and $\langle P_3, P_4 \rangle$, the points on these lines satisfy the following equations:

$$(P_a - P_b) \cdot (P_2 - P_1) = 0$$
$$(P_a - P_b) \cdot (P_3 - P_4) = 0$$

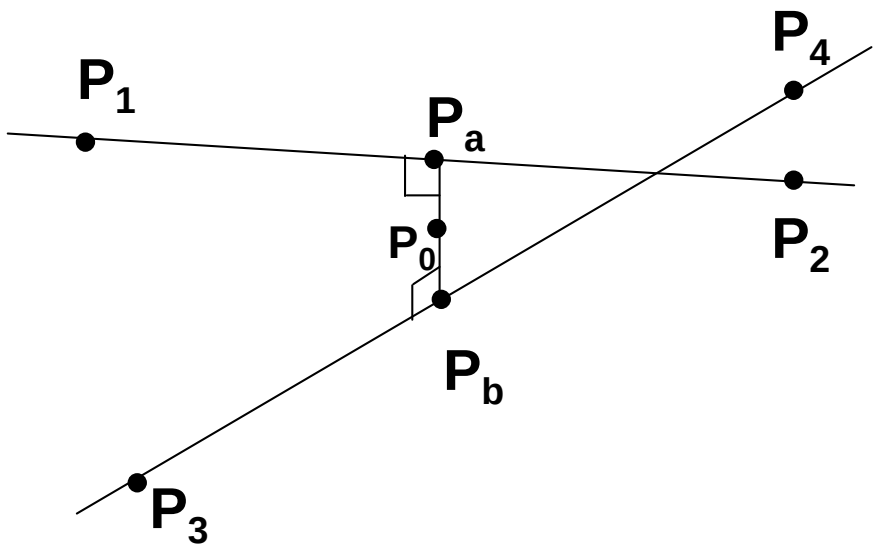

Figure D.1: The two lines, $\langle P_1, P_2 \rangle$ and $\langle P_3, P_4 \rangle$, do not intersect and are not parallel. The point of intersection between these two lines, P_0, is defined as the point at equal and minimum distance to both lines.

After substitute P_a and P_b with expression (C-1) and expression (C-2), respectively, the two equations above become:

$$(P_1 - P_3 + \lambda_a(P_2 - P_1) - \lambda_b(P_4 - P_3)) \cdot (P_2 - P_1) = 0$$
$$(P_1 - P_3 + \lambda_a(P_2 - P_1) - \lambda_b(P_4 - P_3)) \cdot (P_4 - P_3) = 0$$

The two unknowns, λ_a and λ_b, can then be solved. With λ_a and λ_b, compute P_a and P_b using expression (C-1) and expression (C-2). Finally, P_0, the point of minimum distance to two lines, $\langle P_1, P_2 \rangle$ and $\langle P_3, P_4 \rangle$, is then calculated by the following expression:

$$P_0 = \frac{P_a + P_b}{2}.$$

Bibliography

[1] Audette, M., Siddiqi, K., and Peters, T. Level-set surface segmentation and fast cortical range image tracking for computing intrasurgical deformations. **Medical Image Computing and Computer-Assisted Intervention**, pages 788–797, 1999.

[2] Ayache, N. **Artificial Vision for mobile robots: stereo vision and multisensory perception**. MIT Press, Cambridge, MA, USA, 1991.

[3] Ayache, N., and Faverjon, B. Efficient registration of stereo images by matching graph descriptions of edge segments. **International Journal of Computer Vision**, 1(2), 1987.

[4] Basser, P. J. Interstitial pressure, volume, and flow during infusion into brain tissue. **Microvasc. Res.**, 44:143–165, 1992.

[5] Besl, P., and McKay, N. A method for registration of 3-D shapes. **IEEE Transactions on Pattern Analysis and Machine Intelligence**, 14:239–256, 1992.

[6] Biot, M. General theory of three-dimensional consolidation. **Journal of Applied Physics**, 12:155–164, 1941.

[7] Black, P., E. Alexander, I., Martin, C., Moriarty, T., Nabavi, A., Wong, T., Schwartz, R., and Jolesz, F. Craniotomy for tumor treatment in an intraoperative magnetic resonance imaging unit. **Neurosurgery**, 45:423–433, 1999.

[8] Black, P. M., Moriarty, T., Alexander, E., Stieg, P., Woodard, E. J., Gleason, P. L., Martin, C. H., Kikinis, R., Schwartz, R. B., and Jolesz, F. A. Development and implementation of intraoperative magnetic resonance imaging and its neurosurgical applications. **Neurosurgery**, 41(4):831–842, 1997.

[9] Bonsanto, M. M., Staubert, A., Wirtz, C. R., Tronnier, V., and Kunze, S. Initial experience with an ultrasound-integrated single-rack neuronavigation system. **Acta Neurochirurgica**, 143(11):1127–1132, 2001.

[10] Bucholtz, R. D., Yeh, D. D., Trobaugh, J., McDurmont, L. L., Sturm, C. D., Baumann, C., Henderson, J. M., Levy, A., and Kessman, P. The correction of stereotactic inaccuracy caused by brain shift using an intraoperative ultrasound device. **CVRMed-MRCAS'97**, pages 459–466, 1997.

[11] Caprile, B., and Torre, V. Using vanishing points for camera calibration. **International Journal of Computer Vision**, 3:127–140, 1991.

[12] Comeau, R., Sadikot, A., Fenster, A., and Peters, T. Intraoperative ultrasound for guidance and tissue shift correction in image-guided neurosurgery. **Medical Physics**, 27:787–800, 2000.

[13] Davatzikos, C., Shen, D., Mohamed, A., and Kyriacou, S. K. A framework for predictive modeling of anatomical deformations. **IEEE Transactions on Medical Imaging**, 20(8):836–843, August 2001.

[14] Dhond, U., and Aggarwal, J. Structure from stereo - a review. **IEEE Transactions on Systems, Man and Cybernetics**, 19:1489–1510, 1989.

[15] Dickhaus, H., Ganser, K., Staubert, A., Bonsanto, M. M., Wirtz, C. R., Tronnier, V. M., and Kunze, S. Quantification of brain shift effects by MR-imaging. **Proc. An. Int. Conf. IEEE Eng. Med. Biology Soc.**, pages 491–494, 1997.

[16] Dorward, N., Alberti, O., Velani, B., Gerritsen, F., adn N.D. Kitchen, W. H., and Thomas, D. Postimaging brain distortion: Magnitude, correlates, and impact on neuronavigation. **Journal of Neurosurgery**, 88:656–662, 1998.

[17] Edwards, P., Hill, D., Little, J., and Hawkes, D. Deformation for image-guided interventions using a three-component tissue model. **15th International Converence for Information Processing in Medical Imaging, ed. Duncan, Gindi: Lecture Notes in Computer Science**, 1230:218–231, 1997.

[18] Edwards, P., King, A., Jr., C. M., Hawkes, D., Hill, D., Gaston, R., Fenlon, M., Jusczyzck, A., Strong, A., Chandler, C., and Gleeson, M. Design and evaluation of a system for microscope-assisted guided interventions (MAGI). **IEEE Transactions on Medical Imaging**, 19:1082–1093, 2000.

[19] Egnor, M., Rosiello, A., and Zheng, L. A model of intracranial pulsations. **Pediatric Neurosurgery**, 35:284–298, 2001.

[20] Fallenstein, G. T., and Huke, V. D. Dynamic mechanical properties of brain tissue. **J. Biomechanics**, 2:217–226, 1969.

[21] Faugeras, O. **Three-dimensional Computer Vision**. MIT Press, Cambridge, MA, 1993.

[22] Faugeras, O., and Toscani, G. The calibration problem in stereo. **Proceedings of IEEE Conference on Computer Vision and Pattern Recognition**, pages 15–20, 1986.

[23] Feinberg, D., and Mark, A. Human brain motion and cerebrospinal fluid circulation demonstrated with MR velocity imaging. **Radiology**, 163:793–799, 1987.

[24] Ferrant, M., Nabavi, A., Macq, B., Jolesz, F., Kikinis, R., and Warfield, S. Registration of 3-d intraoperative MR images of the brain using a finite-element biome-

chanical model. **IEEE Transactions on Medical Imaging**, 20(12):1384–1397, December 2001.

[25] Fitzpatrick, J., D.L.G.Hill, and Jr., C. **The Handbook of Medical Imaging, Volume 2: Medical Image Processing and Analysis**, chapter 8:Registration, pages 447–513. SPIE Press, Bellingham, WA, 2000.

[26] Fitzpatrick, J. M., West, J. B., and C. R. Maurer, J. Predicting error in rigid-body point-based registration. **IEEE Transactions on Med. Imageing**, 17:694–702, 1998.

[27] Greitz, D., Wirestam, R., Franck, A., Nordell, B., Thomsen, C., and Stahlberg, F. Pulsatile brain movement and associated hydrodynamics studied by magnetic-resonance phase imaging - the Monro-Kellie Doctrine revisited. **Neuroradiology**, 34:370–380, 1992.

[28] Hadani, M., Spiegelmen, R., Feldman, Z., Berkenstadt, H., and Ram, Z. Novel, compact, intraoperative magnetic resonance imaging-guided system for conventional neurosurgical operating rooms. **Neurosurgery**, 48:799–809, 2001.

[29] Hagemann, A., Rohr, K., Stiehl, H., Spetzger, U., and Gilsbach, J. M. Biomechanical modeling of the human head for physically based, nonrigid image registration. **IEEE Transactions on Medical Imaging**, 18(10):875–884, October 1999.

[30] Hall, W., Liu, H., Martin, A., Pozza, C., Maxwell, R., and Truwit, C. Safety, efficacy, and functionality of high-field strength interventional magnetic resonance imaging for neurosurgery. **Neurosurgery**, 46:632–642, 2000.

[31] Hall, W. A., Luis, H., Martin, A. J., Pozza, C. H., Maxwell, R. E., and Truwit, C. L. Safety, efficacy, and functionality of high-field strength interventional magnetic resonance imaging for neurosurgery. **Neurosurgery**, 46:632–642, 2000.

[32] Hartkens, T., Hill, D., Castellano-Smith, A., Hawkes, D., Maurer, C., Martin, A., Hall, W., Liu, H., and Truwit, C. Measurement and analysis of brain deformation during neurosurgery. **IEEE Transactions on Medical Imaging**, 22(1):82–92, 2003.

[33] Hartley, R., and Gupta, R. Computing matched-epipolar projections. **Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition**, pages 549–555, 1993.

[34] Hartov, A., Eisner, S., Roberts, D., K.D.Paulsen, Platenik, L., and Miga, M. Error analysis for a free-hand three-dimensional ultrasound system for neuronavigation. **Neurosurgical Focus**, 6(3), 1999. http://www.neurosurgery.org/focus/mar99/6-3-5.html.

[35] Hata, N., Nabavi, A., Warfield, S., Wells, W., Kikinis, R., and Jolesz, F. A. A volumetric optical flow method for measurement of brain deformation from intraoperative magnetic resonance images. **Lecture Notes in Computer Science: Medical Image Computing and Computer-Assisted Intervention - MICCAI'99**, 1679:928–935, 1999.

[36] Hill, D. L. G., Maurer, C. R., Maciunas, R. J., Barwise, J. A., Fitzpatrick, J. M., and Wang, M. Y. Measurement of intraoperative brain surface deformation under a craniotomy. **Neurosurgery**, 43(3):514–528, 1998.

[37] Hill, D. L. G., Maurer, C. R., Martin, A. J., Sabanathan, S., Hall, W. A., Hawkes, D. J., Rueckert, D., and Truwit, C. L. Assessment of intraoperative brain deformation using interventional mr imaging. **Lecture Notes in Computer Science: Medical Image Computing and Computer-Assisted Intervention - MICCAI'99**, 1679:910–919, 1999.

[38] Hill, D. L. G., Maurer, C. R., Wang, M. Y., Maciunas, R. J., Barwise, J. A., and Fitzpatrick, J. M. Estimation of intraoperative brain surface movement. **Proc.**

CVRMed-MRCAS'97, Springer Lecture Notes in Computer Science, 1205:449–458, 1997.

[39] Jacobson, J., Wallman, L., Schumacher, G., Flanagan, M., Suarez, E., and Donaghy, R. Microsurgery as an aid to middle cerebral artery endarterectomy. **J Neurosurg**, 19:109–115, 1962.

[40] Jannin, P., Fitzpatrick, J., Hawkes, D., Pennec, X., Shahidi, R., and Vannier, M. Validation of medical image processing in image-guided therapy. **IEEE Transaction on Medical Imaging**, 21(12):1445–1449, 2002.

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

[42] Jones, D., and Malik, J. Determining 3-D shape from orientation and spatial frequency disparities. **Proc. 2nd European Conference on Computer Vision**, pages 661–669, 1992.

[43] Kanade, T., and Okutomi, M. A stereo matching algorithm with an adaptive window: Theory and experiments. **IEEE Transactions on Pattern Analysis and Machine Intelligence**, 16:920–932, 1994.

[44] Koivukangas, J., Louhisalmi, Y., Alakuijala, J., and Okarinenen, J. Ultrasound-controlled neuronavigator-guided brain surgery. **Journal of Neurosurgery**, 79:36–42, 1993.

[45] Koyama, T., Okudera, H., and Kobayashi, S. Computer-generated surgical simulation of morpholocial changes in microstructures: Concepts of 'virtual retractor'. **Neurosurgery**, 46(1):118–135, 2000.

[46] Kyriacou, S., and Davatzigkos, C. A biomechanical model of soft tissue deformation, with applications to non-rigid registration of brain images with tumor pathology. **Lecture Notes in Computer Science: Medical Image Computing and Computer-Assisted Intervention - MICCAI '98**, 1496:531–538, 1998.

[47] Lucas, B., and Kanade, T. An iterative image registration technique with an application to stereo vision. **Proc. International Joint Conference on Artificial Intelligence**, pages 674–679, 1981.

[48] Lunn, K. **Data Assimilation in Brain deformation modeling**. Phd thesis, Thayer School of Engineering, Dartmouth College, Hanover, NH, December 2003.

[49] Lunn, K., Hartov, A., Roberts, D., Sun, H., and Paulsen, K. Extracting displacement data from co-registered ultrasound for brain modeling. **Proceedings of SPIE Medical Imaging, Visualization, Display, and Image-Guided Procedures**, 4681:30–38, 2002.

[50] Lunn, K., Paulsen, K., Roberts, D., Kennedy, F., and Hartov, A. Inverse technique for combined model and sparse data estimates of brain motion. **Proceedings of SPIE Medical Imaging, Visualization, Display, and Image-Guided Procedures**, 4(20):87–97, 2003.

[51] Lunn, K., Paulsen, K., Roberts, D., Kennedy, F., Hartov, A., and West, J. Displacement estimation with co-registered ultrasound for image guided neurosurgery: A quantitative *in vivo* porcine study. **IEEE Transactions on Medical Imaging**, 22(11):1358–1368, 2003.

[52] Lunsford, L., Kondziolka, D., and Bissonette, D. Intraoperative imaging of the brain. **Stereotactic and Functional Neurosurgery**, 66:58–64, 1996.

[53] Marr, D., and Poggio, T. Cooperative computation of stereo disparity. **Science**, 194:283–287, 1976.

[54] Marr, D., and Poggio, T. A computational theory of human stereo vision. **Proc. R. Soc. Lond.**, B-204:301–328, 1979.

[55] Maurer, C., Fitzpatrick, M., Wang, M., R.L. Galloway, J., Maciunas, R., and Allen, G. Registration of head volume images using implantable fiducial markers. **IEEE Transactions on Medical Imaging**, 16:447–462, 1997.

[56] Maurer, C., Hill, D., Martin, A., Liu, H., McCue, M., Rueckert, D., Lloret, D., Hall, W., Maxwell, R., Hawkes, D., and Truwit, C. Measurement of intraoperative brain deformation usinga 1.5 tesla interventional MR system: Preliminary results. **IEEE Transactions on Medical Imaging**, 17:817–825, 1998.

[57] Maurer, C. R., Hill, D. L. G., Martin, A. J., Liu, H., McCue, M., Rueckert, D., Lloret, D., Hall, W. A., Maxwell, R. E., Hawkes, D. J., and Truwit, C. L. Investigation of interoperative brain deformation using a 1.5-t interventional mr system: Preliminary results. **IEEE Trans. Med. Imaging**, 17(5):817–825, 1998.

[58] Mayhew, J., and Frisby, J. Psychophysical and computational studies towards a theory of human stereopsis. **Artificial Intelligence**, 17:349–385, 1981.

[59] Miga, M., Roberts, D., Kennedy, F., Platenik, L., Hartov, A., , Lunn, K., and Paulsen, K. Modeling of retraction and resection for intraoperative updating of images. **Neurosurgery**, 49(1):75–85, July 2001.

[60] Miga, M., Sinha, T., Cash, D., Galloway, R., and Weil, R. Cortical surface registration for image-guided neurosurgery using laser-range scanning. **IEEE Transactions on Medical Imaging**, 22:973–985, 2003.

[61] Miga, M. I. **Development and Quantification of a 3D Brain Deformation Model for Model-Updated Image-Guided Stereotactic Neurosurgery**. Ph.D. thesis, Dartmouth College, Thayer School of Engineering, Hanover, N.H., 1998.

[62] Miller, K., Chinzei, K., Orssengo, G., and Bednarz, P. Mechanical properties of brain tissue in-vivo: experiment and computer simulation. **Jornal of Biomechanics**, 33:1369–1376, May 2000.

[63] Nagashima, T., Shirakuni, T., and Rapoport, S. I. A two-dimensional, finite element analysis of vasogenic brain edema. **Neurol. Med. Chir.**, 30:1–9, 1990.

[64] Paulsen, K., Miga, M., Kennedy, F., Hoopes, P., Hartov, A., and Roberts, D. A computational model for tracking subsurface tissue deformation during stereotactic neurosurgery. **IEEE Transactions on Biomedical Engineering**, 46:213–225, 1999.

[65] Platenik, L., Miga, M., Roberts, D., Lunn, K., Kennedy, F., Hartov, A., and Paulsen, K. In vivo quantification of retraction deformation modeling for updated image-guidance during neurosurgery. **IEEE Transactions on Biomedical Engineering**, 49(8):823–835, 2001.

[66] Northern Digital Inc. Products of NDI: Polaris. Webpage, Dec 2000. http://www.ndigital.com/polaris.html.

[67] Reinges, M. H. T., Krombach, G., Nguyen, H., Spetzger, U., Küker, W., Thron, A., and Gilsbach, J. M. Assessment of intraoperative brain tissue movements by frameless neuronavigation. **Comput. Aided Surgery**, 2:218–230, 1997.

[68] Rick, K., Hartov, A., Roberts, D., Lunn, K., Sun, H., and Paulsen, K. Graphical user interface for intraoperative neuroimage updating. **Proceedings of SPIE Medical Imaging, Visulization, Display, and Image-Guided Procedures**, 5029:210–221, 2003.

[69] Robert, L., Zeller, C., Faugeras, O., and Hebert, M. Applications of non-metric vision to some visually guided robotic tasks. **Technical Report 2584, INRIA (available from http://www.inria.fr)**, 1995.

[70] Roberts, D., Hartov, A., Kennedy, F., Miga, M., and Paulsen, K. Intraoperative brain shift and deformation: A quantitative analysis of cortical displacement in 28 cases. **Neurosurgery**, 43(4):749–760, Oct 1998.

[71] Roberts, D., Strohbehn, J., Hatch, F., Murray, W., and Kettenberger, H. A frameless stereotaxic integration of computerized tomographic imaging and the operating microsope. **Journal of Neurosurgery**, 65:545–549, 1986.

[72] Roberts, D. W., Hartov, A., Kennedy, F. E., Miga, M. I., and Paulsen, K. D. Intraoperative brain shift and deformation: a quantitative clinical analysis of cortical displacements in 28 cases. **Neurosurgery**, 43(4):749–760, 1998.

[73] Rubino, G., Farahani, K., McGill, D., de Wiele, B. V., Villablanca, J., and Wang-Mathieson, A. Magnetic resonance imaging-guided neurosurgery in the magnetic fringe fields: The next step in neuronavigation. **Neurosurgery**, 46:643–654, 2000.

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

[75] Skrinjar, O., Nabavi, A., and Duncan, J. Model-driven brain shift compensation. **Medical Image Analysis**, 6:361–373, 2002.

[76] Skrinjar, O., Spencer, D., and Duncan, J. Brain shift modeling for use in neurosurgery. **Lecture Notes in Computer Science: Medical Image Computing and Computer-Assisted Intervention - MICCAI'98**, 1496:1067–1074, 1998.

[77] Skrinjar, O., Tagare, H., and Duncan, J. Surface growing from stereo images. **IEEE Computer Society Conference on Computer Vision and Pattern Recognition**, pages 571–576, 2000.

[78] Sullivan, J., Charron, G., and Paulsen, K. A three-dimensional mesh generator for arbitrary multiple material domains. **Finite Elements in Analysis and Design**, 25:219–241, 1997.

[79] Sun, H., Farid, H., Hartov, A., Lunn, K., Roberts, D., and Paulsen, K. Real-time correction scheme for calibration and implementation of microscope-based image-guided neurosurgery. **Proceedings of SPIE Medical Imaging, Visualization, Display, and Image-Guided Procedures**, 4681:47–54, 2002.

[80] Sun, H., Farid, H., Rick, K., Hartov, A., Roberts, D., and Paulsen, K. Estimating cortical surface motion using stereopsis for brain deformation models. **Medical Image Computing and Computer-Assisted Intervention**, 2878:794–801, 2003.

[81] Sun, H., Farid, H., Roberts, D., Rick, K., Hartov, A., and Paulsen, K. A non-contacting 3-D digitizer for use in image-guided neurosurgery. **Stereotactic and Functional Neurosurgery**, 80(1-4):120–124, 2003.

[82] Sun, H., Lunn, K., Farid, H., Wu, Z., Roberts, D., Hartov, A., and Paulsen, K. Stereopsis-driven brain shift compensation. Submitted to:IEEE Transactions on Medical Imaging, February 2004.

[83] Sun, H., Roberts, D., Farid, H., Wu, Z., Hartov, A., and Paulsen, K. Intraoperative brain imaging using an operating microscope. **Neurosurgery (in review)**, 2004.

[84] Thacker, N., and Mayhew, J. Optimal combination of stereo camera calibration from arbitrary stereo images. **Image and Vision Computing**, 9:27–32, 1991.

[85] Trobaugh, J. W., Richard, W. D., Smith, K. R., and Bucholz, R. D. Frameless stereotactic ultrasonography: Methods and applications. **Computerized Medical Imaging and Graphics**, 18(4):235–246, 1994.

[86] Tronnier, V. M. Craniotomy for tumor treatment in an intraoperative magnetic resonance imaging unit - COMMENT. **Neurosurgery**, 45(3):431–432, 1999.

[87] Tronnier, V. M., Wirtz, C. R., Knauth, M., Lenz, G., Pastyr, O., Bonsanto, M. M., Albert, F. K., Kuth, R., Staubert, A., Schlegel, W., Sartor, K., and Kunze, S.

Intraoperative diagnostic and interventional magnetic resonance imaging in neurosurgery. **Neurosurgery**, 40(5):891–900, 1997.

[88] Trucco, E., and Verri, A. **Introductory Techniques for 3-D Computer Vision**. Prentice Hall, Upper Saddle River, NJ, 1998.

[89] Tsai, R. Y. A versatile camera calibration technique for high-accuracy 3-D machine vision metrology using off-the-shelf TV cameras and lenses. **IEEE Journal of Robotics and Automation**, 3:323–344, 1987.

[90] Unsgaard, G., Gronningsaeter, A., Ommedal, S., and Hernes, T. Brain operations guided by real-time two-dimensional ultrasound: New possibilities as a result of improved image quality. **Neurosurgery**, 51(2):402–412, 2002.

[91] Škrinjar, O. **Deformable Models in Image-Guided Neurosurgery**. PhD thesis, Yale University, New Haven, CT, 2002.

[92] Walsh, E. K., and Schettini, A. Calculation of brain elastic parameters in vivo. **Am. J. Physiol.**, 247:R693–R700, 1984.

[93] Webb, W. **The Physics of Medical Imaging**. IOP Publishing Ltd., Bristol, 1988.

[94] West, J., Fitzpatrick, J., Wang, M., B.M. Dawant, C.R. Maurer, J., Kessler, R., and Maciunas, R. Retrospective intermodality registration techniques for images of the head: Surface-based versus volume-based. **IEEE Transactions on Medical Imaging**, 18:144–150, 1999.

[95] Willson, R. Computer code for Tsai's camera calibration algorithm. webpage, Oct. 1995. http://www-2.cs.cmu.edu/afs/cs.cmu.edu/user/rgw/www/TsaiCode.html.

[96] Wirtz, C. R., Bonsanto, M. M., Knauth, M., Tronnier, V. M., Albert, F. K., Staubert, A., and Kunze, S. Intraoperative magnetic resonance imaging to up-

date interactive navigation in neurosurgery: Method and preliminary experience. **Computer Aided Surgery**, 2:172–179, 1997.