2010

Multimodal Non-Rigid Registration for Image-Guided Head and Neck Surgery

Michael Lawrence Weissberger
College of William & Mary - Arts & Sciences

Multimodal Non-Rigid Registration for Image-Guided Head and Neck Surgery

Michael Lawrence Weissberger
Norfolk, Virginia

Bachelor of Science, The College of William and Mary, 2005

A Thesis presented to the Graduate Faculty
of the College of William and Mary in Candidacy for the Degree of
Master of Science

Department of Computer Science

The College of William and Mary
May 2010

APPROVAL PAGE

This Thesis is submitted in partial fulfillment of
the requirements for the degree of

Master of Science

Michael Lawrence Weissberger

Approved by the Committee, April 2010

Committee Chair
Professor Nikos Chrisochoides, Computer Science
The College of William and Mary

Associate Professor Weizhen Mao, Computer Science
The College of William and Mary

Assistant Professor Denys Poshyvanyk, Computer Science
The College of William and Mary

Associate Professor Joseph Han, Department of Otolaryngology
Eastern Virginia Medical School

ABSTRACT PAGE

Image guidance is a useful and enabling tool in the neurosurgical suite. The use of image guidance during brain tumor resection surgery provides the surgeon with detailed real-time anatomical information which has been shown to improve patient outcomes. Due to brain shift and surgical progression, the intra-operative anatomy may differ significantly from the pre-operatively acquired images. Therefore, intra-operative images are acquired, but must be "registered" or fused, to the pre-operatively acquired images. We previously investigated the performance and efficiency of our state-of-the-art bio-mechanical non-rigid registration algorithm to register intra-operative MRI to pre-operative MRI images in real time in a high-performance computing environment, and found that it improves registration accuracy between 3 and 8 times.

In this study, we investigate the use of intra-operative Computed Tomography (iCT) as an intra-operative modality for non-rigid registration to pre-operatively acquired CT. The goal of this study is to evaluate the accuracy of our non-rigid registration algorithm using iCT images. Successful application of this technique would yield an intra-operative pseudo-MRI at a fraction of the cost. We implement two additional similarity metrics and develop and employ a synthetic benchmark which we use to both select a similarity metric and automatically evaluate the performance of the algorithm. To this end, and in participation with our clinical partners, we create and use a database of six patient cases. We find that pre-operatively acquired CT has sufficient definition to be used in our non-rigid registration algorithm, showing an average improvement of 4.14 times over rigid registration alone using synthetically-generated intra-operative images. We find, however, that our iCT are not suitable as an intra-operative modality currently, showing an average improvement of only 19% over rigid registration alone. We conclude that either more sensitive similarity metrics or an improved iCT is required for successful application of our non-rigid registration algorithm.

for Mom and Dad

Table of Contents

ACKNOWLEDGMENTS

I am greatly indebted to the many people that made this thesis possible. Any errors which remain are completely and solely my own.

Most of all to my advisor Nikos who always impressed me with his attention to my wishes and goals: without your steady encouragement and guidance I would have been lost.

To my many wonderful friends: your companionship kept me focused at the right times, and distracted when I most needed it. I could not have asked for better support.

To the staff of the Computer Science department: Vanessa Godwin and Jacqulyn Johnson, and to the many staff that keep the building running, thank you.

This work would not have been possible without the patient and incredible cooperation of Dr. Joseph Han. I owe him much gratitude for exposing me to his fascinating practice, and sharing his invaluable time and expertise with me.

My deep gratitude goes to my father, Dr. Marshall Weissberger, for his medical expertise in general and his assistance identifying landmarks for the analysis of this work.

This work benefited from the use of the Insight Segmentation and Registration Toolkit (ITK), an open source software developed as an initiative of the U.S. National Library of Medicine and available at www.itk.org.

This work was performed [in part] using computational facilities at the College of William and Mary which were provided with the assistance of the National Science Foundation, the Virginia Port Authority, Sun Microsystems, and Virginia's Commonwealth Technology Research Fund. In addition, Tom Crockett provided me with very helpful support and guidance both in my office and in the field.

This work is supported in part by NSF grants: CCF-0916526, CCF-0833081, and CSI-719929 and by the John Simon Guggenheim foundation.

My thanks go to Ben Kennedy and Will Cline: this work is immeasurably more readable and accurate due to their careful and diligent attention.

Also to my roommate Ben goes my deep gratitude for his patient and warm-hearted friendship.

I cannot understate how important the Daily Grind has been for my academic and mental health. To the staff of the Daily Grind, and mostly to Scott Owen, your support, friendship, and coffee have provided for me immeasurably over the past 8 years.

Finally to my brother David and my parents, Marshall and Barbara: nothing I say can express my love for you and my gratitude for the many ways you have helped and supported me.

List of Tables

List of Figures

Multimodal Non-Rigid Registration for Image-Guided Head and Neck Surgery

Chapter 1

Introduction

1.1 Motivation

Cancer is the third-most-common cause of death in the United States. There are over 560,000 deaths every year in this country due to the disease, second only to heart disease. We concern ourselves here with cancer of the head and neck which is predicted to afflict around 22,000 patients in the United States in 2009, with a five-year survival rate of 35% [45]. The treatment of malignant neoplasms (tumors) is a primary research focus. Surgical intervention in the treatment of such neoplasms is a course chosen by many. The goals of surgical intervention are diverse, including primary treatment of a neoplasm by tissue reduction, diagnostic and investigative procedures, and preventative and palliative care. Should surgery be chosen, image-guidance is a useful and enabling tool. Image guidance in the OR suite provides the surgeon with detailed real-time anatomical information which can result in better surgical outcomes for the patient and better prognoses overall. All image-guidance techniques start with the acquisition of images of various modalities which capture different information about a patient

and her pathology. The description provided by a certain imaging modality reflects the image capture technology employed. For instance, Computed Tomography (CT) is a three-dimensional imaging technique which uses X-ray radiation to produce an image. The image that results from a CT session reflects the physical opacity of a subject to X-ray radiation. The more radio-opaque a subject is, the less radiation is absorbed by the detector. Bones appear whiter than soft tissue because more of the X-ray energy is absorbed by the dense bone tissue, and less reaches the detector, than in the soft tissue. Magnetic Resonance Imagery (MRI) on the other hand uses tuned radio-frequency fields to detect oscillations in the magnetic precession of certain molecules in the subject tissue. The effect is that tissues with higher concentrations of hydrogen (or water) have a different intensity in the final image than tissues with a low concentration of water. There are many other modalities (e.g. Positron Emission Tomography [PET], and ultrasound [US]) which are described by the same refrain: an imaging modality produces images which describe anatomical detail in a specific language. The language of CT describes radio-opacity. The language of MR describes water concentration. An important observation is that each imaging modality provides us with a different, complementary description of an underlying anatomical subject. A CT is a more beautiful and appropriate language to explain details of anatomy which are described by a difference between soft tissues and hard tissues. The vocabulary of MR is more appropriate to explain phenomena which appear in the soft-tissue range.

The use of functional modalities like PET and functional MRI (fMRI) provide different kinds of information than purely anatomical modalities like CT, MRI and US. PET and fMRI explain the metabolic activity of an image subject. That is, PET and fMRI distinguish areas of high metabolic activity (such as tumor tissue and hepatic

tissue) from areas of low metabolic activity. They may also be used to localize important functional areas of the brain, areas which should be avoided during surgery, and areas of relatively little import which may be sacrificed in the pursuit of maximal tumor resection.

These different imaging technologies, together with surgical navigation suites, are enabling technologies for image-guided neurosurgery (IGNS). IGNS refers to the capture and use of these images and special surgical navigation hardware to provide the surgeon with anatomical details which are otherwise unavailable and which may be of vital importance to the success of the procedure. This type of surgery has become commonplace in hospitals across the country, and image guidance tools are often available in ORs. Fried, Parikh and Sadoughi comment on IGNS, saying that it "is one of the most significant advances in the performance of endoscopic sinus surgery since the inception of the endoscopic approach in the mid-1980s [18]." Studies have shown that image guidance improves patient outcomes and makes complicated surgery more feasible, especially when anatomical landmarks are obscured by underlying disease . In addition, endoscopic surgery is a natural fit with image-guidance, since anatomical landmarks are obscured by tissue which would be retracted during open surgery [17]. These images may be used to perform procedures in minimally invasive ways or to perform traditional procedures with more precision. For instance: the boundary between tumor and non-tumor tissue in the brain is often indistinguishable to the naked eye. One of the goals of tumor resection is to provide symptomatic relief and curative treatment of a tumor. Post-operative prognoses are directly correlated with the amount of tumor tissue remaining [9, 23]. On the other hand, it is important to remove as little non-pathological tissue as possible to limit psychological and physiological impact. This balancing act is complicated by

the visual similarity of tumor tissue and non-tumor tissue. Images used during surgery can help the surgeon reduce the pathological tissue and preserve healthy tissue, since the boundary between the two is often more clear with an imaging modality like CT or MRI.

However, the use of pre-operative images for IGNS suffers from a systematic flaw. The problem of *brain shift* presents a challenge to the disruptive neurosurgeon. Brain shift refers to the movement of the brain during neurosurgery due to several unavoidable phenomena. Recent technological advances have allowed surgeons to track such changes during surgery with the use of intra-operative imaging modalities. Previous studies have shown both increased tumor-resection precision and better prognoses for the surgical patient when undergoing an image-guided procedure with intra-operative image acquisition [9, 23]. However, intra-operative images cannot substitute for pre-operative images. Intra-operative images are often noisier, and do not contain the same kinds of information as their pre-operative cousins. The use of intra-operative images to update pre-operative images requires a process known as *Image Registration.* We previously studied the problem of *non-rigid* image registration for pre- and intra-operatively acquired MRI images to address this problem [3].

1.2 Contribution of this Thesis

In this work, we describe our experience applying this tested and validated non-rigid registration algorithm to a novel application, and the methods and tools we developed to overcome the unique challenges encountered in this new application. The application we investigate is real-time non-rigid registration of medical images for image-guided

endoscopic head and neck surgery. The novel component of this application is its use of intra-operative CT in the place of intra-operative MRI.

The first contribution of this work is the development of the protocol and tools required to help accomplish the goal of using this new intra-operative modality. First we describe our experience using the existing protocol and software in this new application. We implemented and experimented with two additional mutual information-based similarity metrics to improve the performance of the registration software in these new circumstances. These new metrics are more robust than other metrics we have used previously when used across different modalities. We evaluate and compare these similarity metrics using a synthetic benchmark which we developed to measure their respective performance. Finally, we propose and evaluate a new protocol to improve the accuracy of the registration procedure.

The second major contribution of this work is the evaluation of this new protocol using prospective and retrospective cases prepared and analyzed especially for this study. We collaborated with our clinical partners, Dr. Joseph Han, Director of Endoscopic Sinus and Skull Base Surgery at the Department of Otolaryngology and Head and Neck Surgery at Eastern Virginia Medical School for the clinical evaluation portion of this study, and carried out several real-time non-rigid registrations from the operating suite. We also consulted with Dr. Marshall Weissberger a radiologist at Medical Center Radiologists, for portions of this thesis.

The overarching purpose of this study is to evaluate the feasibility of using intra-operative CT with our current non-rigid registration technique.

1.3 Structure of this Thesis

This thesis is structured as follows:

- Chapter 2 describes the necessary preliminaries in medicine, imaging technology and image registration, and describes the current state-of-the-art of non-rigid registration algorithms.
- Chapter 3 details our method for non-rigid registration of medical images using intra-operative MRI. In addition, we discuss the supporting technology and high-performance computing tools which enable the real-time use of this method in the OR.
- Chapter 4 explains the changes which we made to the method and clinical protocol to accommodate the change in intra-operative modality from MRI to CT.
- Chapter 5 presents the results of our new method using both synthetic and clinical case studies prepared in concert with our clinical partners.
- Chapter 6 discusses our findings and describes the direction we think future work should take.

Chapter 2

Preliminaries

This chapter discusses the necessary background information required to understand the rest of this work. We briefly describe images, rigid and non-rigid registration, the finite element method and the different types of medical images on which we base this work. We use the notation described by Hill throughout this chapter [25].

2.1 Images

A bounded, continuous set $\tilde{\Omega}$ defines the image space, and Ω defines the image domain for a particular image:

$$\Omega = \tilde{\Omega} \cap \Gamma_{\varsigma},$$

where Ω is a regular grid (an *anisotropic lattice*) on which samples are taken with spacing $\varsigma = \{\varsigma_x, \varsigma_y, \varsigma_z\}$. For our purposes, an image is simply a function which maps a point in this space onto the integers:

$$A : x \in \Omega \rightarrow A(x) \in \mathbb{Z}$$

which is represented by a grey value in medical images, and in **figure 2.1**.

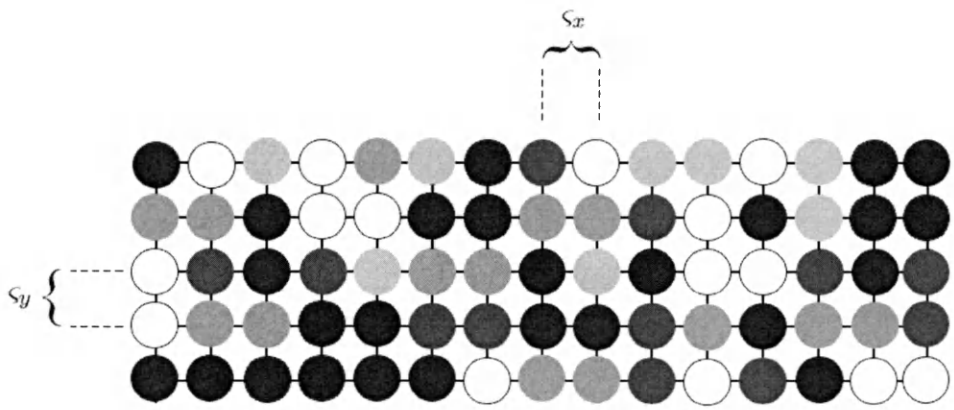

Figure 2.1: An illustration of an image with lattice spacing ς_x, ς_y. [1]

Each sample $A(x)$, sampled at position x is called a "volume element" or "voxel" if the image is in three dimensions (or a "pixel" if the image is in two dimensions), and is centered around a node of the lattice defined by Γ. This function's range is defined by the particular detector used to produce the image, and is generally different from detector to detector, and from modality to modality. It may be thought of as an intensity or a grey value. The images we use have spacing components ς between .5mm and 5mm, depending on the modality used and the settings used during acquisition. Intra-operative modalities generally are optimized for speed and have lower spatial resolution, while pre-operative modalities are optimized for fidelity and often have sub-pixel spatial resolutions [16].

Noise is present in images and has several sources. Martin et. al explain that the noise level in an image is related to the radiation dose delivered to the patient. They explain that as the number of photons used to create an image decreases, the background noise increases. This is because as the speed of radiographic film (or detectors) increases, fewer photons are required to produce an image or signal. This is desirable because this

delivers less radiation dose to the patient. However, as the number of photons per image area decreases, the resultant image is more vulnerable to random quantum fluctuations in the arrival of photons and to serendipitous photons not generated by the imaging device. This phenomenon is called *quantum mottle* [33]. We will see this is a concern since intra-operative detectors are optimized for speed of image acquisition, and as a result use a lower dose. Practically this makes intra-operative images more grainy and noisy, and reduces the signal-to-noise ratio.

2.2 Image Registration

Image Registration is the process of discovering the relationship, correlation, or matching between two images that differ in some way. The images may differ in time, orientation, subject or in sensor modality. Brown describes image registration as finding, for each of the points in one image, the corresponding point in another image [8]. Image registration is an important task because different images may contain complementary information, and the fusion of two or more images may multiply the usefulness of any individual image. Functional images like fMRI and PET are commonly registered to anatomical images like CT and MRI in order to better localize functional metabolic activity to a specific anatomical location in the image. In the case of image-guided neurosurgical tumor surgery, it is helpful to very precisely localize functional information, provided by a functional modality like fMRI, to anatomical information provided by an anatomical modality like T1 MRI in order to more precisely target tumor and non-tumor tissue. In addition, this information is useful to better predict where non-tumor tissue may be safely removed to maximize tumor tissue removal margins, and where non-tumor tissue

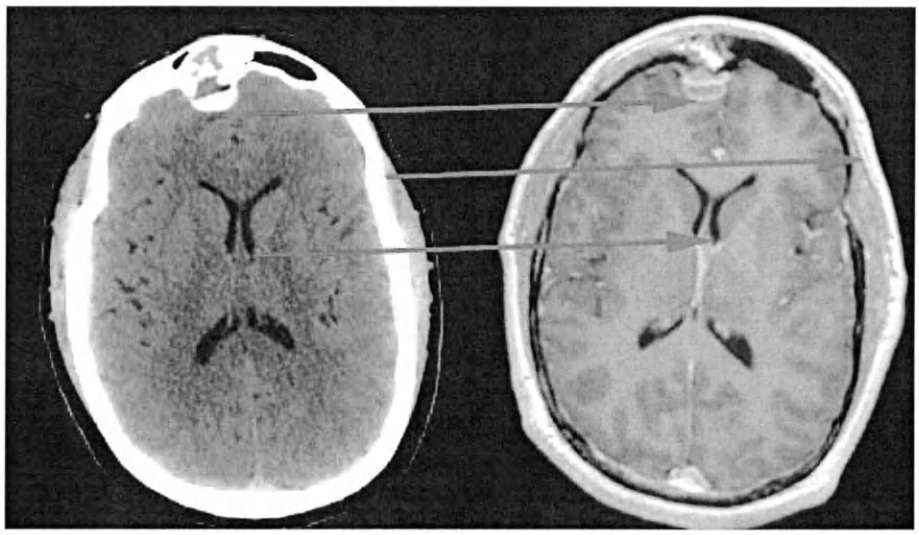

Figure 2.2: Example of a rigid registration between a CT and an MRI. Because these two images were captured with different devices, and with different patient positioning, the two images need to be registered to align their anatomical features. Since both were obtained pre-operatively (and are of the same subject), we may assume safely that the two images are related in a rigid manner. This means that a single, global rigid transformation defines the difference between the two images.

is critical to maintain physiological and mental function. In order to accomplish these tasks the two images must be registered together.

We will now discuss the terminology of image registration. When registering two images, one is called the *floating* or *moving* image, and one is called the *fixed, target* or *reference* image. They are so called because we wish to deform or transform the floating or moving image (thus it will be *moved*) so that its features match the fixed, target or reference image. Image registration is divided into two broad categories: *Rigid* and *Non-rigid* registration.

2.2.1 Rigid Registration

Rigid registration describes the differences introduced between two images as a global rigid-body transformation. These global rigid transformations may involve a rotation

factor, a scale factor, and a translation factor and may be described by a single global function which maps image coordinates in the first image to image coordinates in the second image. These transformations are called rigid because the relationship between different parts of the two images is the same across the entire image. These transformations in two dimensions may be described as in [8] by three parameters, a global rotation θ, a translation t_x, t_y and a scaling s.

$$\begin{pmatrix} x_2 \\ y_2 \end{pmatrix} = \begin{pmatrix} t_x \\ t_y \end{pmatrix} + s \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} x_1 \\ y_1 \end{pmatrix}$$

In general, rigid registration problems are limited to resolving image distinctions introduced due to differing image acquisition times, positions of the camera and subject, or different imaging modalities of the same, internally consistent subject. For example: registering a CT and an MR of the same patient, of an organ that is largely static internally, is a problem well-suited to rigid registration. Rigid registration is a generally solved problem with robust, automatic and accurate solutions. These solutions include feature-based methods such as paired landmark registration, contour registration, and voxel-based methods. Landmark and contour methods use paired points to calculate a rigid transformation between two images [25]. These points and contours may be manually selected and matched by hand, or automatically selected and matched. Automatic feature selection may be divided by feature type: points, lines, regions and corners as well as other geometric shapes. Edges and lines are covered in [13] and [56]. Region features are described in the segmentation literature [36]. Point feature selection has been understood as line-intersection [51], identification of high variance points [4, 14], and various corner-detection algorithms, among others. Automatic feature matching is commonly done with the iterative closest point method [55], or area-based methods

like maximization of some similarity metric [57]. The resulting global transformation is generally calculated using a least-squares approximation which minimizes the residuals of the target registration polynomial at the corresponding points [57]. Both point-based and contour-based rigid registration are extensively covered in the literature [8, 32].

A completely different method involves maximization of a global similarity metric between the two images being registered. Wells et. al, contemporaneously with Maes et. al, use gradient descent to maximize the mutual information between pairs of images in many different modalities with great success [31, 53]. Škerl, et. al. compare five different performance attributes of nine different similarity criteria for rigid registration of multi-modal images [58]. Currently, rigid registration methods are the only methods commercially available to our clinical colleagues in their surgical navigation systems.

2.2.2 Non-rigid registration

Non-rigid registration refers to registration between two images which may have undergone an elastic deformation. Pairs of images which display non-rigid transformations include images in which the underlying subject matter has changed internal configuration, not just changed in relation to the sensor. Non-rigid registration attempts to recover deformations in a subject which are local, not global, and which may be non-uniform across the entire image. Non-rigid registration is generally much more computationally intensive, but reduces overall registration error when correctly applied to these kinds of problems.

There have been various approaches to non-rigid (also called *elastic* or *deformable*) registration. Rueckert et. al. model non-rigid movement of the breast using MR images by altering local transformation parameters in order to maximize a global similarity

measure via gradient descent [41]. Hata and Maintz do the same using synthesized and real CT and MRI images and mutual information as the similarity metric [22, 32]. A deformable model using attractive forces modeled after Maxwell's demons was also successfully employed to register MRI of the brain and cardiac motion [49]. A segmentation-based uni-modal non-rigid registration algorithm was investigated by Collins, and revisited by Gaens to use mutual information. In these segmentation-based methods, an atlas is used to identify structures in the target images, and these structures are transformed to maximize a similarity criteria [12, 19]. Likar and Pernuš introduced a hierarchical model for registration of skeletal muscle using mutual information. They use local mutual information maximizations to generate an elastic deformation at regular grid points. They generate successively finer details of deformation by iteratively traversing down a hierarchy of region sizes, from large to small. In addition, they compare how incorporation of prior information into the joint probability density estimate, histogram binning and random sub-sampling can improve the performance of mutual information in small regions [28]. Andronache et. al. extended this hierarchical model to three dimensions. They also use a spatial autocorrelation measure to exclude structureless patches from the registration hierarchy, while switching between two different similarity criteria (mutual information and cross correlation). They also introduce intensity mapping to create a pseudo-modality in-between the two modalities to be registered. The goal of this pseudo-modality is to generate an image which contains only the anatomical features shared by the two target modalities [1, 2].

As far as our application is concerned, non-rigid registration is a better choice than rigid registration to model the changes that the brain undergoes during tumor resection. During resection, the elastic and non-rigid movement of the brain creates complex and

local deformations to occur. Uncovering these local deformations is the goal of a non-rigid registration solution. In later sections we will describe our approach to non-rigid registration for image-our guided neurosurgery.

2.3 Finite-element analysis

The study of real-world dynamic systems such as the one we examine in this thesis sometimes require the evaluation of partial differential equations. Solving these partial differential equations (PDEs) is often an intractable problem when dealing with real-world systems and their accordingly non-ideal geometries and boundary conditions. In our case, the geometry of the brain is complex and makes direct evaluation of our PDE infeasible. Finite-element analysis is a numerical technique which allows us to find an approximate solution to such PDEs. This technique requires us to generate a finite-element discretization, or mesh. We will use this mesh to convert our PDE into a system of linear equations which we may solve directly. The generation and optimization of this mesh is an open problem. In this study we rely on previous work to generate a high-quality mesh from an image segmentation. This meshing approach has previously been evaluated extensively and is well-suited to our chosen application. The particular meshing approach is described in [16].

The PDE whose solution we approximate over this mesh describes the tradeoff between several different forces acting in the biomechanical model we use during neurosurgery. We discuss the meshing approach and its use in **Chapter 3**.

2.4 Similarity Metrics

The non-rigid registration procedure described below relies on a notion of "similarity" between two regions of an image. In this section we describe what "similarity" means and how to measure it. We discuss three different similarity metrics: the *Normalized Correlation Coefficient (NCC)*, the *Normalized Mutual Information (NMI)* and the *Regional Mutual Information (RMI)*. Each of these metrics relates the intensity distributions of the two images, as characterized by the example joint histograms in **figure 2.3**. We will assume we have two image regions I and J, with intensities $I(x)$ and $J(y)$ and mean $\overline{I}$ and $\overline{J}$ respectively, where x and y are voxels in the two images: $x \in I$ and $y \in J$.

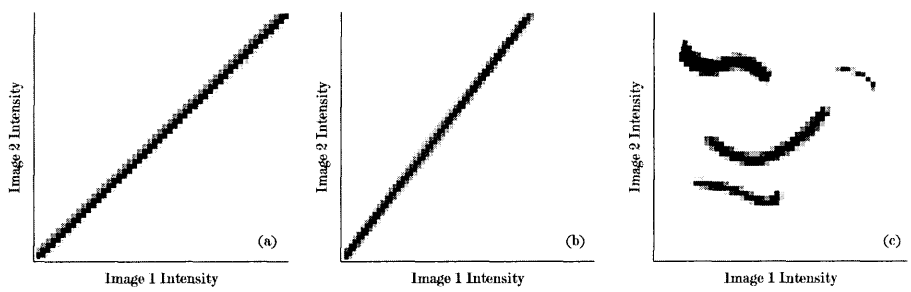

Figure 2.3: Three cartoon joint histograms. (a) displays an identity relationship between the two image intensities, (b) displays an affine relationship, and (c) displays a statistical relationship. The metrics in this section measure the degree to which corresponding image intensities follow the given relationship.

2.5 Normalized Correlation Coefficient

The simplest metric which we consider and use is the Normalized Correlation Coefficient. The NCC is defined as the ratio of the covariance between I and J to the product of the standard deviation of I and J:

$$NCC = \frac{Cov(I, J)}{\sqrt{Var(I)} * \sqrt{Var(J)}}$$

The NCC is useful for registering images from the same modality. This measure is valuable for comparing two images whose probability density functions are related in an affine manner, as is suggested by Roche [39]. In addition, the value of the NCC is normalized between zero and one, with higher values meaning two regions' intensity distributions are more similar, and lower values meaning they are less similar.

2.6 Mutual Information-Based Metrics

When using the NCC, we make the assumption of an affine relationship between the image intensity distributions. However, a more robust similarity metric does not require this assumption. In addition, noise may exist in one of the two images in the form of random background noise or artifacts from motion, partial volume, or beam-hardening[2] This noise may be more effectively compensated for with a more flexible and robust similarity metric.

Mutual information is a statistical measure of the amount of "information" one image tells us about another image. In practice we may use mutual information-based measures to compare images which may be related in a non-functional, statistical manner, such as in multi-modal image registration. Mutual information, and all derived similarity measures, are born out of the field of information theory.

[2]The back-propagation algorithm for tomographic reconstruction assumes a monochromatic X-ray beam. In practice this is not true, and artifacts may result[7].

The images I and J may be thought of as random variables with individual outcomes or values I_i and J_j. Shannon defines the *self-information* Ψ of an individual outcome I_i as:

$$\Psi(I_i) = -\log p(I_i)$$

where p(·) is the probability of outcome I_i [44]. This may be thought of as the amount of information associated with an outcome. Outcomes which are less common have a higher information content than outcomes which are more common and are more expected. A similar and interchangeable term is an outcome's surprisal, or, the surprise which we experience when observing a particular outcome. The information is defined in terms of a unit of information, which manifests itself in the base of the logarithm in the previous definition. This unit is commonly the "bit" resulting in the use of the log_2.

For example, if we flip a fair coin, we expect the coin to land heads up with probability .5. This makes the surprisal or self-information of this outcome

$$\Psi(I_{heads}) = -\log_2 .5 = 1.$$

Alternatively, if we have an unfair coin which lands on its head with probability .75, then

$$\Psi(I_{heads}) = -\log_2 .75 \approx .416$$

Intuitively, if an outcome is more likely (the unfair coin landing on heads), then the information contained in that event is lower. The Shannon (marginal) entropy $H(I)$ of a given variable is defined as the expected value (mean) of the self-information of a variable over all possible outcomes:

$$H(I) = E(\Psi(I)) = -\sum_{i=1}^{n} p(I_i) * \log_2 p(I_i)$$

In our unfair coin example:

$$H(I) = -p(I_{heads}) * \log_2 p(I_{heads}) - p(I_{tails}) * \log_2 p(I_{tails}) = .75 * .416 + .25 * 2 = .812$$

which is less than the entropy of a fair coin ($.5 * 1 + .5 * 1 = 1$). In other words, the uncertainty of a flip of a fair coin is higher than the uncertainty of a flip of an unfair coin.

The conditional entropy of two variables $H(I|J)$ may be understood as the uncertainty remaining in a random variable I once we know the outcome of a second variable J. It may be calculated as:

$$H(I|J) = -\sum_{i,j} p(I_i, J_j) \log p(I_i|J_j).$$

Mutual information relates the marginal and conditional entropies of the target images. It may be thought of as the amount one variable's uncertainty is reduced by the knowledge of a second variable:

$$MI(I, J) = H(I) - H(I|J).$$

Equivalently, it may be thought of as the information content "shared" by two variables. Mutual information of discrete variables may be calculated as:

$$MI(I, J) = \sum_{i \in I} \sum_{j \in J} p(i, j) * \log \frac{p(i, j)}{p_1(i) p_2(j)},$$

where $p_1(i)$ and $p_2(j)$ are the marginal probability distribution functions for variables I and J and $p(i, j)$ is the joint probability distribution function between the same. Intuitively, mutual information of two independent variables is zero, since $H(I) = H(I|J)$

for independent variables. In other words, if I and J are independent, knowing the value of variable J tells us nothing about the variable I, so the conditional entropy $H(I|J)$ is the same as the marginal entropy $H(I)$. If, on the other hand, I is perfectly dependent on J, then the conditional entropy $H(I|J) = 0$, and the mutual information is equal to the marginal entropy of I [3].

Finally, if we treat an image as a random variable measuring intensity, then we may use mutual information as a measure of how related two image intensity distributions are. Mutual information was simultaneously introduced in this manner for image registration by Viola et. al. and Maes et. al. Various methods have been explored to find the optimal registration using the mutual information, including gradient descent [31, 53], simulated annealing [38], and hill climbing [46]. A normalized variant of mutual information was also introduced by Studholme et al., which proved to be more robust than other variants [47]. Mutual information has been proven to be a robust and reliable automatic similarity measure for use in both mono- and multi-modal rigid registration. Studies have shown success using MI among and between CT, MRI, PET, SPECT, US and fMRI [31, 53, 58, 37].

2.7 Regional Mutual Information

One limitation of mutual information stems from its use of the probability density function (PDF) for images to be registered. Since we do not generally know this function, (and since the PDF is unique for each patient, device, and even image acquisition) we

[3] Another equivalent formulation of the mutual information is $MI(A, B) = H(A) + H(B) - H(A, B)$. Using this formulation, we see that maximizing the mutual information between two variables (say, by registering their corresponding images) is equivalent to minimizing their joint entropy $H(A, B)$. Minimizing two variable's joint entropy has the effect of clustering the corresponding variables' probability density functions.

must estimate it. This is most simply done with either a discrete histogram measurement or using Parzen windowing [46, 48]. As the number of spatial samples decreases, any method of PDF estimation becomes less accurate. For instance, with very few samples, the accuracy of a discrete histogram method suffers from the noisy and sparse available data. To solve this problem, we may decrease the number of bins in the histogram. This leads to a "fuzzy" and lossy PDF estimate and many registration artifacts. (Improvements have been realized through adaptive histogram binning, K-means clustering, and robust maximum *a posteriori* (MAP) estimation of joint histograms [42, 50].)

A direct implication of this problem is that traditional MI metrics do not perform well with small image sizes or image regions. The predictive power of the joint and marginal PDFs in such cases is simply too low to provide a good measurement of similarity. Likar and Pernus elegantly illustrate this problem by showing that the mutual information measurement between a small image and the same small image spatially translated has multiple local maxima which may obscure the global maximum and prevent proper registration. They show that the mutual information may not reach its global maximum at the correct registration for images of less than 64x64 pixels [28]. The same effect is less pronounced but still experienced when using three dimensional images.

It has been suggested that by incorporating some prior information about the probability density function into the estimate of the floating PDF, we may increase the measure's performance for small image sizes. Various methods exist to do this, and we chose to implement the Regional mutual information (RMI)[40].

Rogelj et. al examine and derive a point-based similarity metric from the mutual information which allows us to measure the similarity of arbitrarily sized regions of an image. They derive it by noting that the definition of mutual information may be

rewritten in terms of image voxels v, a member of the cartesian product of image voxels in I and J: $v = \{(x, y)|x \in I \text{ and } y \in J\}$ and intensities i and j in the corresponding intensity distributions I_p and J_p:

$$MI(I, J) = \sum_{i \in I_p} \sum_{j \in J_p} \frac{N_{(i,j)}}{N} \log \frac{p(i, j)}{p_1(i)p_2(j)} = \frac{1}{N} \sum_{v \in I \times J} \log \frac{p(v)}{p_1(I(v_I))p_2(J(v_J))},$$

where N is the number of overlapping voxels in the image, $N_{(i,j)}$ is the number of voxels with intensity i in image I and intensity j in image J, and v are the overlapping voxel intensity pairs between the two images. They then remark that this definition of mutual information may be understood as an average of what they call point similarities $S_{MI}(v)$.

$$MI = \frac{1}{N} \sum_{v} S_{MI}(v),$$

where

$$S_{MI} = \log \frac{p(v)}{p_1(I(v_I))p_2(J(v_J))}.$$

They specify that S_{MI} is an *estimate* of the intensity dependence between the images when they are correctly registered, and as such is only usable if we expect the final transformation relating the two images to be relatively small.

They then use this notion of a point similarity measure to define the similarity of an image region S_R:

$$S_R = \frac{1}{N_R} \sum_{v \in R} S_{MI}(v).$$

The key difference between the regional mutual information and the canonical mutual information applied to a small image region is that the S_{MI} uses a joint PDF which is calculated from the entire images before registration, while MI is calculated using the joint PDF of just the image region, during registration. A consequence of this is that if

the images are grossly misaligned, the point similarity function will be a poor estimate. However, when the two images are relatively well-registered and we may be confident that the prior PDF of the two images resembles the floating PDF of the image regions to be registered, we may expect the estimate to be a good one. By incorporating these global statistics, we increase the predictive power of the PDF estimate, and reduce the interpolation and binning problems with vanilla mutual information based measures.

2.8 The Problem of Brain Shift

The assumption of internal rigidity or internal consistency required for rigid registration is generally inapplicable. Human subjects are not rigid, and their internal compositions change depending on the internal and external forces applied to them. This is especially true in the special case of neurosurgery. Neurosurgeons must deal with a phenomenon known as "brain shift", in which the brain changes internal shape as the dura is punctured and cerebrospinal fluid (CSF) escapes. The brain is essentially floating in an inertial buffer of CSF. During neurosurgery the envelope enclosing the brain and CSF is cut, causing a loss of this fluid. The brain floating in CSF experiences very different external pressures and forces than the brain resting in the absence of CSF. This causes the (relatively elastic) brain tissue to deform up to 25mm in some cases [4] [35]. In addition, tumor tissue resection cuts connections between brain tissues, and causes the brain to change shape in a non-rigid manner as more and more tissue is removed [3].

These two factors combine to render images acquired pre-operatively obsolete. Images which are acquired before surgery describe the anatomical configuration of the

[4]The adult brain is typically around 1500 ccm, so a 25mm deformation represents a deformation of nearly 20% of the brain size.

brain before surgery (and before any resultant brain-shift) which may be markedly different than the configuration during surgery due to these phenomena. These shifts also render functional information less useful, since this functional information is registered to the now-obsolete anatomical information. This is a problem because, as previously described, incomplete information during surgery can result in a less than successful procedure where tumor tissue remains, or large amounts of functional non-tumor tissue is removed. In addition, accurate anatomical and functional images are a prerequisite for image-guided surgery. Since the surgeon is relying partially on the image guidance system to localize his instruments inside the patient, brain shift and the resultant non-rigid movement of the brain is a serious problem. A solution to this problem is the introduction of intra-operative imaging procedures, a protocol which has been used for some time to provide updated information to surgeons about the state of the brain during surgery [34].

During image-guided neurosurgery, the surgical procedure is paused while an updated intra-operative image is acquired. This image provides the surgeon with an updated landscape of the brain. This intra-operative image may be used by the surgeon to direct the remainder of the procedure. The problem of brain shift is essentially solved by acquiring and using intra-operative images. These intra-operative images generally fall into two modal categories, both anatomical: CT and MR.

There is a tradeoff in the use of intra-operative images during surgery. The anesthetized patient must be physically placed in the imaging device while the procedure is put on hold. Since the surgery cannot be paused for long in this manner, the image acquisition times must necessarily be short. In addition, since the surgeon must have access to the patient, an intra-operative MR must use an "open magnet" or "open

bore" configuration. These open magnets are generally of a lower magnetic field (for instance, .5 T instead of 3 T) [6], and thus have a lower signal-to-noise ratio. In addition, the long acquisition times required for functional modalities like fMRI are impractical during surgery. In order to overcome these limitations, image registration has been investigated as a solution. By registering the pre-operative images to the intra-operative images, it is possible to "update" the high-detail, pre-operative images, including those containing functional information, to more accurately reflect the configuration of the brain during surgery. To do this, we can use the intra-operative images as a reference image, and the pre-operative images as the floating image to be registered. Rigidly registering pre-operative images with intra-operative images is insufficient, especially around the area of tumor resection and maximal brain shift. Rigid registration yields registration errors exceeding 50 percent of the actual deformation. We previously used a robust non-rigid registration method in order to register pre-operative MRI and fMRI to intra-operative MRI. We have shown this method to be between three to eight times as accurate as rigid registration alone [3]. A major contribution of our previous work is a computational framework to provide real-time non-rigid registration during an ongoing procedure. Using this technique, we have successfully provided high-quality MR and functional MR to neurosurgeons during surgery. We summarize our previous work on MRI-MRI image registration for image guided neurosurgery in **Chapter 3**.

Unfortunately, complicating the applicability of this procedure is the rarity of intra-operative MRI devices in-clinic today. Only a handful exist in the country, and they are extremely expensive to purchase and operate. In addition to acquiring an intra-operative MRI, an entire OR suite must be tailored to its use [5]. Intra-operative CT

[5]For example, all surgical equipment must be made of non-ferrous materials to be compatible with

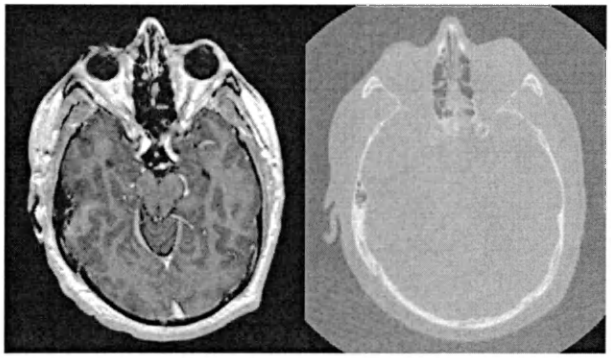

(a) MR of the head (b) CT of the head

Figure 2.4: a) MR of the head with superior soft tissue differentiation. b) CT of the same patient with inferior soft tissue differentiation. The differences in soft tissue differentiation arises from the kind of sensor which produces the image. This patient is included in our case studies below.

scanners, on the other hand, are relatively inexpensive and are ubiquitous in OR suites across the country. One drawback of using intra-operative CT as an image-guidance modality is that it has far inferior soft-tissue resolution as compared with MRI. **Figure 2.4** shows an example of this difference in a CT and MR of the head. We will discuss the implications of this difference in **Chapter 4**.

Our previous work was focused on image-guided neurosurgery for intracranial gliomas, a type of brain tumor. A contribution of this work is its application of our method to a novel set of tumor sites. In this study, we will examine the performance of our non-rigid registration algorithm during endoscopic, and open head-and-neck and skull base surgery in coordination with our clinical partners at Eastern Virginia Medical School. To this end, we developed a set of image collection protocols for this study. Finally, we present our method for non-rigid registration with intra-operative CT for head and neck surgery, also in **Chapter 4**.

---

the MR scanner's high magnetic field.

2.9 Imaging Modalities

Medical imaging modalities vary widely. In general medical images are used to solve an inverse problem: inferring pathological cause by examining the effect, the observed signal. Different detector and probe types yield different resulting images. In this section, we will present a brief overview of medical imaging technologies. We divide medical imaging modalities into two classes: anatomical and functional. This distinction is a fuzzy one, however. Anatomical images do convey functional information, however–in general–we define a functional image as one that describes phenomena that are not merely configurational. That is, a functional image shows details of the subject that may not be reflected in an underlying physical difference. We describe the two anatomical modalities we use in this thesis.

2.9.1 X-ray Computed Tomagraphy (CT)

It is apparent what the probe is for this very common modality: X-ray radiation. An X-ray CT is a three dimensional image which describes a subject's radio-opacity. X-rays are electromagnetic radiation with wavelength of between 10 and .01 nanometers that are produced using an *X-ray tube*. An X-ray tube contains a cathode (a source of electrons) and an anode (a target for electrons) enclosed in a vacuum. The cathode is generally a heated fiber of tungsten which emits electrons when heated. A high voltage is applied across the cathode and anode, which accelerates the elecrons towards the anode. Electrons striking the anode are abruptly slowed down, which releases their energy. Approximately 1 percent of this released energy is converted to emitted photons in the X-ray range, with the rest converted to heat. These photons are then projected through the patient, where some fraction of them are absorbed. The fraction which is transmitted

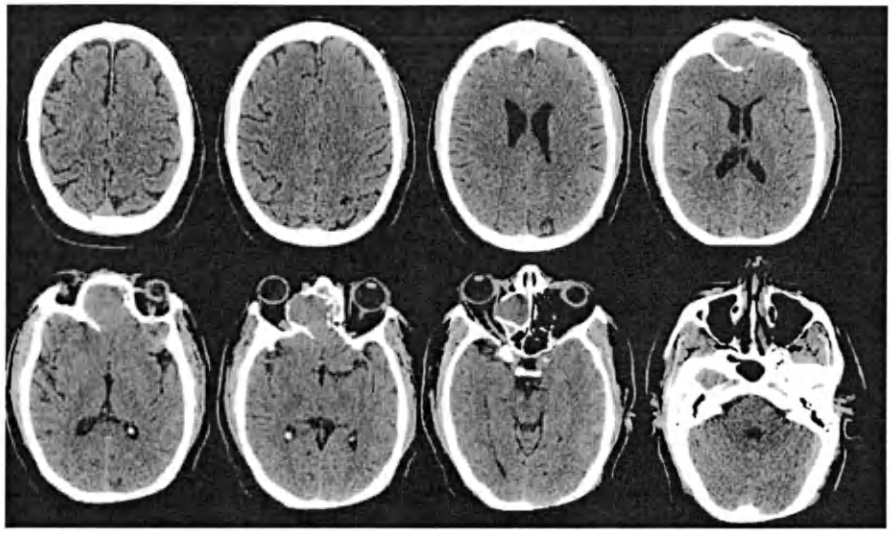

Figure 2.5: Example CT of the brain, clockwise from anterior to posterior. Visible are the folds (*sulci*), ventricles, contrast agent, slight differentiation between the brain soft tissues, and a frontal and ethmoidal sinus neoplasm. This patient is included in our case studies below.

through the material is then detected by a sensor. Many measurements are taken in a circular or helical fashion as the patient is slowly moved through the device. These measurements are then used to create a three dimensional reconstruction representing the radio-opacity of the subject in a process called tomographic reconstruction [27]. Contrast agents, which are very radio-opaque, are sometimes used to highlight certain anatomical structures in an image.

CT acquisition of the brain takes a few minutes, and is relatively inexpensive as compared with the other modalities below. The resulting greyscale image has voxel intensities which are proportional to the attenuation of X-rays in the image, as described theoretically by Beers law:

$$I = I_0 e^{(-\mu\chi)}$$

with μ representing the attenuation constant for a material, χ representing the length of

the transmission path through that material, I representing the measured x-ray intensity and I_0 representing the initial x-ray intensity. Many measurements of intensity are taken and a resulting greyscale image is produced which reflects the attenuation constant $\mu_(x, y, z)$ at each voxel (x, y, z) in the image. Image reconstruction is performed in most modern CT units by a process known as back-propagation. In this technique, the attenuation coefficient measurement μ is "smeared" backwards along the direction of ray-propagation. This process is repeated for each intensity measurement, and an image is produced by the convolution of these back-propagated intensities.

2.9.2 Magnetic Resonance Imagery (MRI)

MRI makes use of a novel feature of nuclear physics called Nuclear Magnetic Resonance (NMR). NMR refers to the tendency of magnetic nuclei in a magnetic field, in the presence of an applied electromagnetic pulse, to radiate energy outwards predictably. In the presence of a large magnetic field, the magnetic moments of a target material's nuclei largely "line up", either along or in direct opposition to the applied field. An applied electromagnetic pulse at a specific ("Larmour") frequency, causes these magnetic moments to precess and eventually return to their initial (lower energy) equilibrium state (or "relax"). This relaxation produces a signal, which the detector uses to create an image [52].

In general, the image produced from an MRI describes the relative hydrogen concentration of tissues in a subject. There are several different flavors of MR. These include FLAIR, diffusion-weighted, T1, T2, spin-density, among others.

Chapter 3

Previous Work: Biomechanical Model for Non-rigid Registration of Intra-operative MRI

3.1 Overview

The use of intra-operative MRI (iMRI) to produce an image of an *in situ* brain during surgery has revolutionized neurosurgery. iMRI is a relatively new tool, developed by General Electric during the 1990s. During an image-guided procedure, a series of intra-operative images are acquired using an open-bore MRI in a specially designed operating room suite. Often, these images are incorporated in a surgical navigation tool which allows a surgeon to precisely localize their progress through the procedure. Serial image acquisitions have been shown to accurately image surface and subsurface deformations which result from gravity and subsidence of the brain near the resected volume [34]. Nimsky et. al. describe the impact which iMRI has on image guidance by noting that

surgical navigation systems which rely on pre-operative images cannot be used during the "critical steps of the surgical procedure, e.g., identification of the deep tumor margin." [35]

iMRI does not completely replace pre-operative imaging. These intra-operative images have several drawbacks. First – in general – iMRI have a relatively lower signal-to-noise ratio than their pre-operative cousins. This results from the more complicated open-bore configuration required to provide access to the patient during surgery, and its corresponding magnet with a lower field strength. Pre-operative MRI have field strengths as high as 3T, while most pre-operative MRI machines have a field strength of less than 1T. As a result, iMRI images have a lower spatial resolution and produce a noisier image than more traditional MRI devices. In addition, it is possible in pre-operative imagery to capture functional images which may, for example, be useful to more precisely localize tumor margins. In contrast, it is in general not possible to capture functional images with an iMRI device due to the specialized requirements of such imaging procedures.[1]

As a result of these drawbacks, we previously introduced a robust algorithm for non-rigid registration of pre-operative and intra-operative MRI in order to solve the problem of brain shift during image-guided neurosurgery. Utilizing updated information present in serially acquired iMRI, we non-rigidly registered a set of pre-operative images to each captured intra-operative image. Pre-operative images included higher-quality T1 and T2-weighted MRI, fMRI and DTI. A major barrier to performing non-rigid registration during surgery is its high computational cost. Our previous work also addresses this

[1]Recently intraoperative fMRI have started to come online. However, pre-operative fMRI is still generally regarded as the gold standard [20].

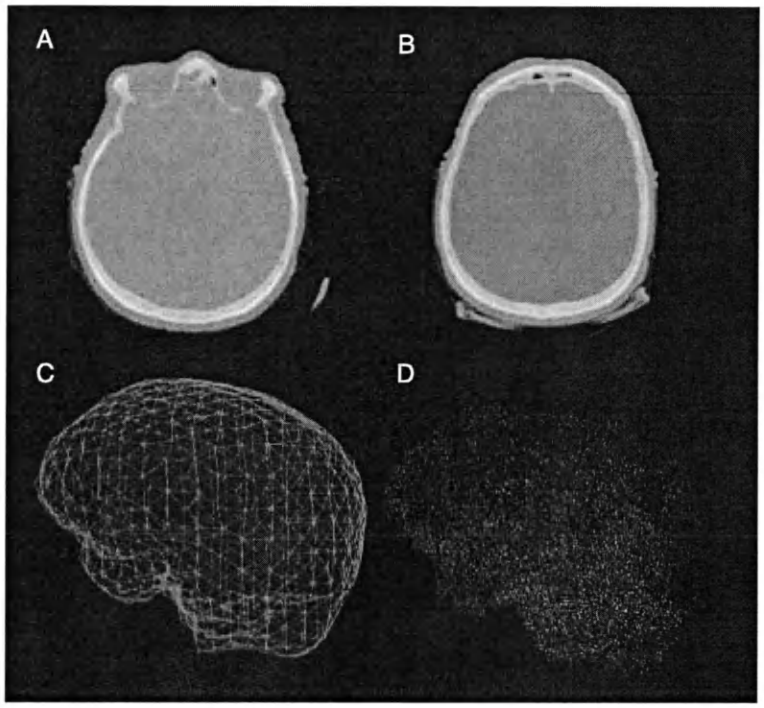

Figure 3.1: An illustration of the pre-operative steps. A. Pre-operative CT. B. Segmented pre-operative CT. C. Finite element Mesh generated with RGMesh. D. Feature points selected from highest-variance blocks.

issue by leveraging high performance computing environments (including cluster and grid computing resources) to provide real-time non-rigid registration for image-guided neurosurgery for the first time. In this section, we briefly describe our method. For a more detailed treatment, see [3, 11].

We may divide our approach into pre-operative and intra-operative processing stages. We attempt to perform as much calculation as possible in the pre-operative stage, since the intra-operative stage must be done in near real-time. The first step in the pre-operative stage of our NRR algorithm is the physical acquisition of the pre-operative image. The second step is the segmentation of the pre-operative MR to obtain a binary mask of the intra-cranial cavity (ICC). This segmentation is then used to generate a

patient-specific model which consists of a tetrahedral volume mesh of this binary segmented image. The penultimate step in the pre-operative stage consists of the selection of optimal candidates for "block matching" in the point (or feature) selection step. The last step of pre-operative processing is the distribution of these images to the various computational sites to be used in the intra-operative phase.

The intra-operative stage is further divided into two stages, the pre-resection stage and the intra-resection (time critical) stage. Recall that the goal is to update the pre-operative images. Once the first intra-operative image arrives, which is generally taken before resection begins, the first step is to rigidly register that image to the pre-operative MR. This is done because the patient is physically affixed to the MRI machine, so any future images will be in the same coordinate space as the first intra-operative image. We use this transformation to register the pre-operative data (including the segmentation, floating images and mesh) to the soon-to-follow intra-operative images. The time between the first pre-resection image and the first intra-resection image is generally long enough that the computational requirements of these steps are easily satisfied.

Once the first intra-resection image is taken, our time constraints are much more restrictive. The first step is to generate a sparse deformation field by finding a block matching between the pre-operative floating and intra-operative fixed images. The second (and last) step is to approximate a dense deformation field by employing a novel iterative approximation and interpolation heuristic which converges on a solution while rejecting outliers from block matching. Using this dense deformation field, we may update the pre-operative images and deliver these to the surgeon.

3.2 Pre-operative processing

3.2.1 Image Acquisition and Segmentation

Pre-operative image acquisition typically occurs in the week before the procedure. In general, segmentation refers to the demarcation or separation of an image into pieces or regions. Segmentation can be used to separate an image into constituent organ systems or tissues, or to isolate a region over which our model is defined, as in our case [54]. After acquisition, we must separate the image into two regions in order to isolate the intra-cranial cavity (ICC). This segmentation is required to restrict the elastic motion of the brain model to ensure it does not deform outside of the cranial cavity, which would be physically unrealistic but otherwise allowed without the use of this patient-specific model in the mechanical energy term of our non-rigid registration formulation. In addition, we must segment the brain in order to create the mesh which will be the basis of this patient-specific biomechanical model. We describe the mesh in more detail in the next section.

Segmenting the intra-cranial cavity is a semi-automatic process. Segmenting a two dimensional image is relatively straightforward. Segmenting a three dimensional image, however, is more difficult and time-consuming. Automatic segmentation is typically performed with thresholding, region growing and connectivity operators, which select regions based on similar voxel intensities which may be in continuous regions. For example, the ICC on a bony-window CT has a nearly uniform intensity, which allows relatively simple region growing to correctly segment. In contrast, an MRI has much higher tissue specificity and superior soft-tissue resolution, which makes this process more difficult and time consuming. We employed a combination of automatic operators,

like region growing and level-set filters, with slice-by-slice manual segmentation to correct any erroneously included regions. We used Insight's Slicer3D software, and ITK's SNAP, parts of the Segmentation and Registration Toolkit [26], as well as in one case a novel model-based segmentation method [30].

3.2.2 Mesh Generation

Our NRR uses a patient-specific model to tailor the computation to an individual patient's anatomy. This model uses a tetrahedral mesh generated from the segmentation produced in the previous step. Not all meshes are equal, and this mesh has several properties that make it well suited to our problem [15].

The mesh should conform to the boundary of the binary image. This is important so that the model of the patients brain accurately describes the actual shape of the brain. Second, we strive for an equidistribution of registration points with respect to the mesh vertices. This affects the magnitude of the displacement error, as well as the numerical condition of the problem. Next, the gradation of the mesh is an important property that may reduce interpolation error by having elements of smaller size. However, more numerous small elements may lead to a longer solution time. Thus it is important to be able to adaptively refine the mesh in selected regions of interest. Next, the shape of mesh elements is important, as small angles may increase the condition number of the stiffness matrix, and lead to slower convergence in the numerical solver. Next, mesh generation time is important, as short times allow real-time mesh refinement. Such refinement is desirable because as outlier registration points are discarded in the iterative solver previously mentioned, the first requirement (equidistribution of registration points) may be violated. Finally, the mesh must provide a reasonably close approxima-

tion of the surface of the meshed object. A mesh which conforms perfectly to the image boundary would provide great fidelity, but at the cost of an excessive level of detail and an overwhelming number of mesh elements. Extensive previous work was done to refine our meshing procedure to produce a mesh best suited to our application of image-guided neurosurgery [15]. We use the red-green mesh algorithm presented in [16].

The resulting mesh, finely tuned to our application of image-guided neurosurgery, is used as the basis of the non-rigid registration method. It is used for the approximate numerical solution of the partial differential equation which we must solve to find the best trade-off between empirically measured block deformation and the elastic properties of the real tissue, described later.

3.2.3 Salient Feature Point Selection

The goal of the registration method is to recover the movement of the brain anatomy between the acquisition of the pre-op image and the intra-op image. The heart of the algorithm is a window-bounded block-matching algorithm that measures the sparse initial deformation field. We perform block-matching between the pre-operative (or floating) image (MRI or CT) and the intra-operative (or target) image (CT).

Before we can do this, we must target block matching to areas of the image which have a reasonable degree of structural information. We do this because block matching between structured sub-image regions will be less error-prone than block matching between regions with relatively less structural information. In other words, floating image regions which are random noise, or which have little structural definition, will be harder to identify in the fixed image. We call this stage point selection, and is defined over several parameters. These parameters are described in table 3.1. We begin by placing a

point (hereafter called a *registration point*) at the centers of all voxels inside the region of interest (ROI) of the floating image. This ROI is the intersection of the floating image with the tetrahedral mesh. The registration points define the centers of sub-regions called blocks. The variance of the voxel intensities of each block is calculated, and a fraction of the blocks with the lowest variance is rejected. The remaining blocks are hypothesized to be better candidates for block-matching with the intra-operative fixed image, which occurs in the next step.

Parameters		
Parameter	Space	Default value
Block half-size	$\mathbb{Z}^3$	(3,3,3)
Voxel connectivity	$\mathbb{Z}$	26
Rejection Fraction	$\mathbb{R}$	.95

Table 3.1: Point selection parameter space

3.3 Intra-operative processing

The intra-operative stage of our non-rigid registration comprises four parts: rigid registration, mesh generation, sparse deformation field measurement during block-matching, dense deformation field estimation using a finite element solver and deformation of the floating image. These stages are in the time-critical path of this procedure. Image-guided therapy, specifically intra-operative image acquisition, is a major disruption in the surgical process. As such, a major requirement for our technique is that it cause minimal down-time during the surgical procedure. Practically, all of these intra-operative stages must be completed in a matter of minutes to be useful to the neurosurgeon during surgery.

3.3.1 Rigid Registration

The first intra-operative stage is the rigid registration of the floating image and the just-acquired intra-operative fixed image. This step is necessary to remove the global rigid difference between the two images. The floating and fixed images are related to each other by a rigid transformation which results from differences in patient and image positioning. In other words, the rigid registration process is necessary to put the two images into the same spatial coordinate system. The floating image (I_1) and the reference image (I_2) are related by a transformation which comprises a rigid transformation (τ_r) and a non-rigid transformation τ_n:

$$I_1 = \tau_n(\tau_r(I_2)).$$

The estimation of τ_n is made more tractable by the prior transformation of τ_r, since the rigid part is globally applied to the entire image and is easier to calculate. This rigid registration is performed using ITK's implementation of affine registration by maximization of normalized mutual information [26]. This discovered rigid registration is applied to the intra-operatively acquired images, obviating the need for intra-operative ICC segmentation and mesh generation, which is impractical and too time-consuming.

3.3.2 Block Matching

In the block matching phase we begin to estimate the difference between the two images. The inputs to block matching are the pre-operative and intra-operative images to be registered, and the set of registration points selected during the pre-operative point selection phase. We may make the assumption that all of the registration points (p_i^o) in the pre-operative image have homologs (p_i^*) in the intra-operative image. p_i^* is the

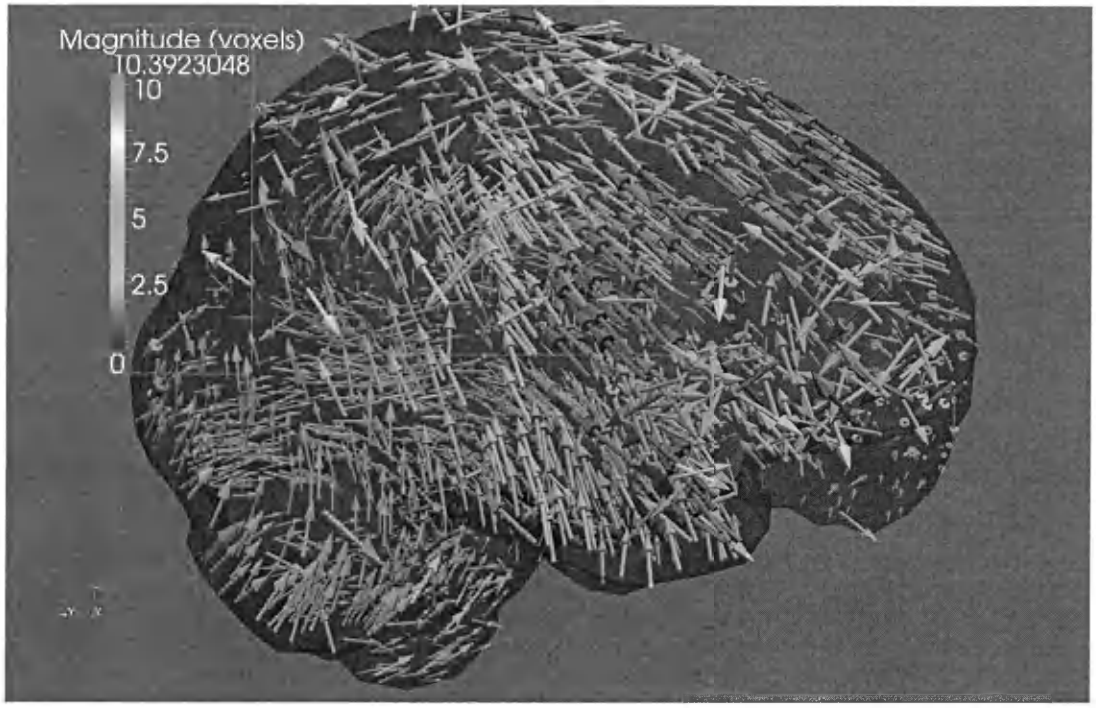

Figure 3.2: A completed block matching showing a small subset of the recovered matches. Each arrow represents the deformation which maximizes the similarity of the subject (origin) block to the pointed-to (target) block. The magnitude of the displacement is represented by the size and color of the arrow. Note that there are outliers, and that the displacement field is *sparse*.

hypothetical intra-operative position of point p_i^o. In fact, the difference between the location of the homologous points $p_i^* - p_i^o = \delta_i$ is the deformation we are searching for. In order to find this deformation, we choose to perform a bounded exhaustive search near p_i^o. We choose a search window size W and block size B as parameters to the block matching algorithm. We first define a neighborhood β_i^o around each p_i^o in the pre-operative image. This neighborhood (called a block) is a cube of dimension B. We then define cube-shaped blocks $\{\beta_j^* : \mathbf{j} \in W_i\}$ of dimension B, centered at $\mathbf{j}$ in the intra-operative image, where W_i is the neighborhood (of dimension W) of voxels around p_i^o.

We next iteratively compare the putative block in the intra-operative image with the subject block in the pre-operative image using some measure of similarity $\{\sigma(\beta_i^o, \beta_j^*) : \beta_i^o \otimes \beta_j^* \rightarrow \mathbf{R}\}$. We choose the block $\hat{\beta}_j^*$ which maximizes the chosen similarity measure σ. The choice of σ is an important one, and the optimal metric depends on many factors. The resulting deformation $i - \hat{j}$ is our preliminary estimation of the deformation measured between the pre-operative and intra-operative images at the ith registration point. The block selection procedure is repeated for each registration point in the pre-operative image. The resulting sparse deformation field is used as a rough estimation of the displacement between the two images [2].

The optimal similarity measure is a function of several factors. In previous studies we used the Normalized Cross-Correlation [3]. NCC is appropriate for mono-modal image registration, or image registration between images of the same modality. Registering iMRI to MRI is a problem well-suited to NCC, as there is a linear or affine relationship between the two images' joint probability density functions. In other words: we can expect the intensity of any voxel in one image to be an affine function of the value of the corresponding voxel in the other image.

This assumption does not necessarily hold for images from different modalities. The intensity of a voxel in an MRI is not an affine function of the corresponding voxel in a CT image. In fact, the relationship between these two images isn't even functional — it is statistical [24]. A more complicated relationship often requires a more sophisticated similarity measure, such as the Correlation Ratio or one of many measures derived from information theory, like normalized mutual information (NMI). The most appropriate

[2] We call this a *sparse* deformation field because it is only defined at the registration points previously selected. We will later use this sparse field to generate a dense deformation field.

measure, however, is dictated by many inputs: the modality of the images, and the specifics of the images such as image subject, detector position, and lighting conditions [39]. We have compared the performance of several different similarity measures in order to choose an optimal measure for our purposes of image-guided head and neck surgery. To do this, we developed a set of synthetic benchmarks which are used to select the optimal similarity measure for two given images.

3.3.3 Estimation of the dense deformation field

The sparse deformation field measured with block matching suffers from two problems. First, it is sparse and discontinuous (defined only at the registration points). We wish to have a deformation which is smooth and defined everywhere. Secondly, it is noisy because the block matching algorithm is imperfect. It may not always find the correct block matching because the numerical similarity between two homologous blocks may happen to be lower than between two non-homologous blocks. In other words, two blocks may be similar with respect to the chosen similarity metric, but not with respect to the anatomy. The block matching measure may "choose" the wrong putative block, simply because it happens to be numerically "closer" in terms of the pre-defined similarity metric to the target block. This happens because of noise in the input images, differences in structure between the two image blocks, and because some of our image similarity metrics do not take structural information into account. This is especially true in areas homogenous tissue intensity. The registration points with such erroneous block matchings are called "outliers." The first approach to this problem we consider is interpolation.

If we model the total energy of the mesh as W, and the matching energy as $W_{matching}$, then

$$W_{matching} = (HU - D)^T S(HU - D)$$

where H is the linear interpolation matrix between the displacements recovered by block matching and those at the mesh vertices, U are the (unknown) displacements at the mesh vertices, D are the measured displacements from block matching, and S is the block matching stiffness matrix (matches with higher confidence are assigned higher weights). Minimization of this matching energy (with respect to the unknown displacements U) will minimize the error between the measured displacements and the solution mesh vertex displacements. However, as just discussed, the interpolation formulation performs poorly in the presence of noise. To correct for this, we use an approximation approach. An elastic energy which resists the deformation of the mesh is added, $W_{mechanical}$:

$$W_{mechanical} = U^T KU$$

where K is the mechanical stiffness matrix, which describes the elastic properties of the mesh and reflects an estimation of the physical properties of the brain tissue. The total energy in this formulation is therefore:

$$W_{total} = W_{mechanical} + W_{matching} = U^T KU + (HU - D)^T S(HU - D)$$

This mechanical energy may be understood as an approximation of tissue's tendency to resist deformation. A further rationale for this formulation is that the sparse deformations measured by block matching do not, in and of themselves, take into account the elastic properties of real tissue. This is important because we want to estimate a deformation which is physically realistic. In order to provide a registration that is

physically realistic, free of outliers, and continuous, we choose to model this problem as an energy minimization between the resistive (mechanical) energy, and the deformative (matching) energy.

In this formulation, the final deformation is determined by the competition between the mechanical energy, which is inversely related to displacement, and the matching energy, which is directly related to displacement. This equilibrium state is found classically by differentiating the total energy with respect to U, and setting equal to zero:

$$\frac{\partial W}{\partial U} = [K + H^T SH]U - H^T SD = 0$$

where $H^T SD$ is essentially the weighted displacement at each mesh vertex, estimated from the displacements at the surrounding registration points. This is called the approximation formulation, where solving for U gives us our solution. This system of linear equations is solved using the biconjugate gradient stabilized method (implemented within the Gmm++ library with the diagonal preconditioner). A major problem with the approximation formulation, however, is that it suffers from a systematic error. The approximation formulation, by definition, may not pass through all of the measured displacements. In the presence of outliers, this is undesirable, since block displacements which are correct will be equally weighted with displacements which are incorrect.

To solve this problem, we wish to identify the likely outliers, and seek to remove them from our final solution. We introduce an external force F to the approximation formulation, which counter-balances the internal mesh stress:

$$[K + H^T SH]U = H^T SD + F$$

This force is then iteratively updated to balance the internal force KU, at each iteration i:

$$F_i \leftarrow KU_i$$

$$U_{i+1} \leftarrow [K + H^T SH]^{-1}[H^T SD + F_i]$$

The tetrahedral finite element mesh produced by the image-to-mesh conversion has a dual role in the above formulation. First, it is used in the mechanical energy ($U^T KU$) of the system to model deformation of the brain as a physical body. This is used to discover and discard the outlier registration points. At the end of each iteration a certain fraction of outlier block matches are discarded. The discarded block matchings are those which were found to have low numerical confidence during block matching, or which disagree with their neighbors. Second, the mesh is also used to regularize (or smooth) the displacements estimated from the minimization of the matching energy ($(HU - D)^T S(HU - D)$) from block matching. Once the solution converges, the algorithm is completed, and we have calculated a dense volumetric deformation that conforms to the measured block matchings. For a more precise treatment of this algorithm, see [10, 11].

3.4 Computational Framework

In order for these results to be useful during surgery they must be available in real-time. In practical terms, this means a matter of minutes. There are a variety of constraints during surgery, and pausing the procedure for more than a few minutes poses unacceptable risk to the patient. The computational requirements for the software described previously are large, and can take days to run in serial computation. The two dif-

ferent phases, pre-operative and intra-operative, have different time constraints. The pre-operative stage (includes image preparation, segmentation, mesh generation, and feature block/point selection) are not subject to rigid time constraints, and are successfully executed on a single hot processing node. The intra-operative phase consists of rigid registration, block matching and solving. This phase is subject to rigid time constraints. Block matching is the slowest of the intra-operative steps. Our parallel version of block matching, which optionally employs real-time load balancing, allows execution of the block-matching stage in a matter of a few minutes in a cluster environment with 36-52 processors. See [10] for a detailed treatment of the computational framework employed. We also investigate using a heterogeneous cooperative architecture using the GPU to execute block matching [29, 30].

Chapter 4

An Updated Protocol: Intra-operative CT

Unfortunately, intra-operative MR facilities are rare across the US today. The cost of an intraoperative MRI machine is prohibitive: between US$1 million and $6 million for a single machine. In addition, the OR suite must be configured for the specific use of the machine. This can require a special layout, specialized tools, and specialized training for technicians. This can cost an additional $700,000 to $1 million. In addition, maintenance of the machine is between 5% and 10% of the initial cost outlay, annually [43, 21]. This scarcity restricts the usefulness and deployment of our registration software, since intra-operative MR are required during the registration process. The goal of this work is to increase the usefulness and broaden the availability of this protocol. To do this, we needed to identify a suitable replacement for the intra-operative MRI which provide the updated reference images during surgery. To this end, we investigate intra-operative CT as a candidate modality.

Interventional CT scanners are common in surgical suites around the country. CT

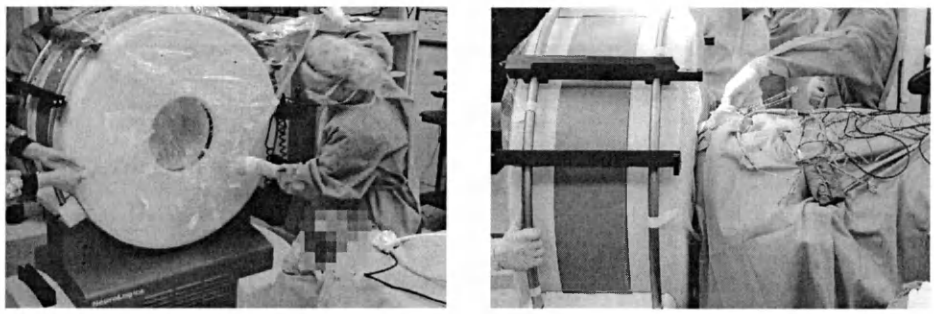

(a) (b)

Figure 4.1: The intra-operative CT machine used in this study.

is a much less expensive image to acquire than MR. A CT machine costs a fraction of an MR machine, and the surgical suite does not need to be tailored for the specific use of these machines. As a result, many hospitals have a portable CT machine which may be used during surgery, and many more may easily and affordably acquire one. The actual machine is very compact, in order to be maneuverable and portable enough for use during surgery with an anesthetized patient. The intra-operative machine in this study is the Neurologica CereTom, and is shown in figure 4.1. It is approximately 4'x5'x1', and is capable of capturing CT images with an in-plane resolution of 512x512 voxels, and a slice thickness of down to 1.25mm. The main contribution of this thesis is that we extend our previous non-rigid registration algorithm to use intra-operative images acquired in the CT modality.

4.1 Overview

The use of intra-operative CT introduces several challenges to the existing protocol. These challenges spring from the difference between CT and MRI: what each measures, and how. As mentioned in Chapter 2, CT and MRI measure very different physical

properties, stemming from differences in the probe. While MRI uses nuclear magnetic resonance to measure the concentration of water in a sample, CT uses radiation in the X-ray range to measure the opacity of a sample to that radiation. It turns out that many different soft tissues in the human body are similar in radio-opacity, and that the resulting CT measures similar attenuation constants for tissues that many be very different. The result can be seen in figure 4.2. On the other hand, varying tissues are often very distinctly differentiated in an MRI, owing to the differences in water concentration in them. For instance, fatty tissue has a different water content than muscle tissue. The different tissues in the brain have different concentrations of water, making MRI images of the brain very articulated. On the other hand, and important for our research, different tissues in the brain, including normal and pathological tumor tissue, can be very similarly radio-opaque. This leads to relatively little definition between them in a CT.

This difference has immediate consequences for our software, which uses the intra-operative images and relies on their precision. In this section, we describe the changes we made to the protocol, the real-time and retroactive clinical cases we studied, and describe how we evaluated the modified protocol.

4.2 Description of the revised protocol

Both the original protocol and the revised protocol call for segmenting the pre-operative brain. In the iMRI protocol, this segmentation was performed using an MRI rather than a CT. In addition, the segmentation was performed by the cooperative neurosurgeon. In this study, we investigated a variety of automatic and manual tools to segment

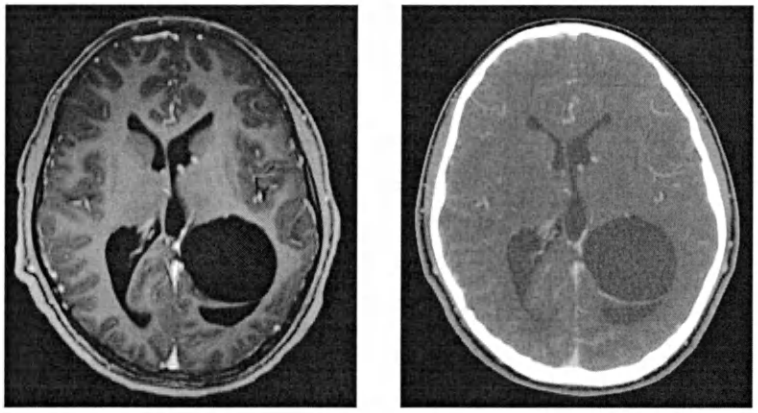

(a) (b)

Figure 4.2: A comparison of the imaging qualities of MRI (a) and CT (b) of the brain. Notice the higher soft-tissue fidelity in the brain, where CSF appears black, contrast appears white, and grey matter and white matter are highly discriminated. The folds (sulci) are well-distinguished. Compare with the CT where the bone is white, the CSF darker grey, and the white and grey matter are middle-grey. The sulci are much more difficult to locate.

the volume — including the open source Slicer, and itk-SNAP. We had repetitive success segmenting our pre-operative images with ITK-SNAP, which merges an automatic component with a manual component [54]. We chose to use the pre-operative CT instead of the pre-operative MRI for segmentation, because of the superior reproduction of the soft-tissue/bony-tissue barrier. Since we segment the entire intra-cranial cavity, segmenting the CT will suffice. Our cooperating neurosurgeon verified several of our segmentations.

Another change was made to the intra-operative protocol. Recall that we first rigidly register the pre-op image with the intra-op image before we model the non-rigid movement of the brain. This rigid registration is a required step to make the non-rigid registration block-matching step more tractable. In the previous protocol, the patient's head is affixed to a stereotactic frame which is not moved during surgery. The first intra-operative image is acquired with the patient attached to this frame, but before

resection begins. This allows us to properly rigidly register the pre-operative and intra-operative MRI once, and ensure that the transformation mapping pre-operative space to intra-operative space does not change during the procedure. In this protocol, the medical procedure does not use a stereotactic frame, and does not acquire this first intra-operative image. To overcome this challenge, we instead perform a rigid registration each time an intra-operative image is acquired. This change adds a small amount of pre-processing time to the overall process. A bigger concern is that as the resection progresses, depending on the size of the resected volume, the automatic rigid registration methods we employ may fail. While we did not encounter this problem, we suggest it would be wise to prepare several rigid registration procedures, in case one of them fails to provide a satisfactory registration due to tissue loss. Our automatic rigid registrations were verified by our cooperating neurosurgeon.

The end goal of this study is to produce an updated intra-operative pseudo-MRI. To do this, we first investigated multi-modal non-rigid registration between the pre-operative MRI and the intra-operative CT. We implemented two additional similarity measures, related to the Mutual Information. We implemented the Normalized Mutual Information, and the Regional Mutual Information [31, 40, 53]. We suspected that performing a nonrigid registration between the pre-operative MRI and the intra-operative CT would produce unreliable results, primarily because of our block-matching strategy. Our blocks are relatively small regions: less than 10 voxels in each dimension. Mutual Information based measures are generally unreliable as they are applied to smaller and smaller regions [2], and our blocks are indeed very small for this purpose. Since multi-modal block matching is a much more complicated process, and since we also have a pre-operative CT available in this medical protocol, we decided to change the protocol to

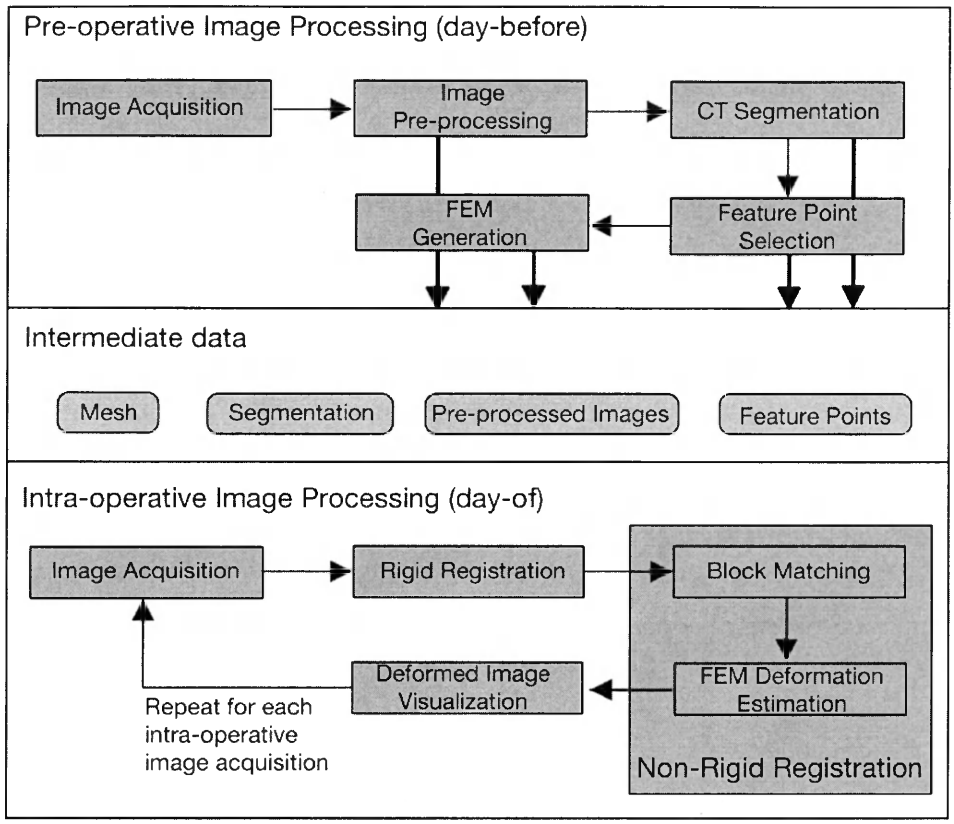

Figure 4.3: An illustration of the clinical protocol in this study.

perform a non-rigid registration between the pre-operative CT and the intra-operative CT. Once this transformation is discovered, we then apply the same transformation to the pre-operative MRI, which has been first rigidly registered to the pre-operative CT. We suspect that the error introduced by this additional rigid registration (between the pre-operative CT and the pre-operative MRI) to be generally less than the error introduced by a deficient block-matching.

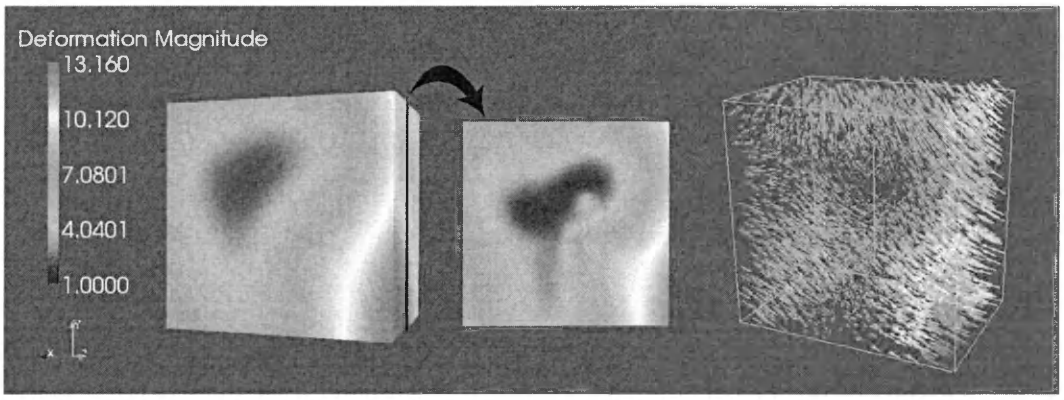

Figure 4.4: An example deformation field volume. The deformation field shown is a three dimensional deformation volume, on the left, with magnitude of deformation shown as color. A two dimensional slice is extracted in the center. The field vectors for the same deformation field are shown on the right. All deformations are in units of voxels.

4.3 Synthetic Deformation Benchmarks

As previously described, the optimal block matching metric $\alpha(\cdot)$ depends on many variables. Roche et. al. describe that the normalized correlation coefficient performs well for mono-modal image registration [39], and that the information theoretical metrics perform more robustly with multi-modal image registration. We used a synthetic deformation benchmark which we used to select the proper block matching metric from among these three available metrics. This synthetic benchmark involves synthetically deforming the input CT in such a way that we know the real deformation ("ground truth" deformation), recovering that deformation using our non-rigid registration algorithm, then calculating the error between the recovered deformation and the actual deformation. We repeat this process for each available similarity metric, and are able to compare the absolute error each suffers from.

We use software described in [5] to generate the synthetic deformation fields. The software follows a method described by Rogelj et al in [40]. We first create an isotropic

lattice (grid), and assign a deformation vector at each node of the lattice. This deformation vector is drawn from a Gaussian distribution with parameters mean $\mu = 0$ and variance σ [5 – 25], where higher σ results in larger deformation. These deformation vectors at the nodes of the lattice form a sparse and discontinuous deformation field. In order to obtain a more realistic continuous field, and to generate a dense deformation field, we then regularize these deformation vectors using thin-plate splines across the entire image volume.

The dense deformation field can then be applied to the input CT, resulting in a synthetically deformed CT which may play the role of the intra-operative CT during image registration.

Chapter 5

Results

We include two studies in this work. The first uses pre-operative data exclusively. We use these pre-operative CT images to answer a single question: *does CT have sufficient definition in the soft tissue regions to be a good candidate for intra-operative use as a fixed image during non-rigid registration?* To answer this question, we require a source of ground truth with which we will assess the performance of the registration algorithm. We use the previously described method to generate a synthetic deformation which serves as our ground truth. We then deform each pre-operative CT according to a synthetic deformation, and treat this new deformed CT as the intra-operative image. We perform segmentation, mesh generation, point selection, block matching and dense deformation field estimation as previously described. We perform the final two steps three times, once for each similarity metric: NMI, PMI and CC. Finally, we compare the recovered deformation with the ground truth deformation in several ways.

We calculate the **matching error** of block matching by subtracting the recovered displacement from the true displacement for each measured displacement vector, located at the registration points. This measure is reported in millimeters. We calculate the

Table 5.1: A summary of the image characteristics for the synthetic study

Case	Image Size (vox)	σ	ς	Applied Deformation (mm)		
				max	avg	std
1	$(512, 512, 57)$	10	50	9.40	4.36	1.89
2	$(512, 512, 190)$	20	50	12.98	6.02	2.22
3	$(512, 512, 144)$	10	20	10.34	3.85	1.27
4	$(512, 512, 144)$	30	50	6.00	2.45	1.14
5	$(512, 512, 136)$	5	50	7.09	3.43	1.46
6	$(512, 512, 75)$	10	50	13.37	5.98	2.68

registration error by subtracting the final dense deformation calculated at the end of each registration with the true deformation, for each mesh vertex of the finite element mesh on which the dense deformation is defined. This measure is also reported in millimeters. The **symmetric Hausdorff distance** (SHD) is a measure of the distance between two point-sets as previously described. Finally, **mutual information** is a measure of image intensity similarity. It may also be used to assess the quality of a registration, as it measures how much information is contained in the relationship between the two intensity distributions. As two images are more tightly registered, the mutual information between them increases. Alternatively, as two images are less well-registered, the information approaches zero[1]. Here we use the normalized mutual information variant.

The second study performed uses "actual" intra-operative data. These images are captured intra-operatively with the previously described intra-operative CT scanner. We present three clinical cases which correspond to pre-operative synthetic cases 4 – 6. The registration method is the same as for the synthetic case, however the intra-operative data are not synthetically deformed in any way. In this study, we attempt to recover the

[1]The mutual information of two independently distributed random variables is also zero

Table 5.2

Case	Metric	Matching error		Registration error		SHD	NMI Ratio	Avg
		max	avg	max	avg			Improvement
1	CC	**17.92**	**0.92**	**6.34**	**1.02**	1.49	8.84	4.27
	NMI	18.97	5.05	10.16	2.96	1.91	1.81	1.47
	PMI	**12.85**	**0.66**	**6.38**	**1.16**	1.51	8.60	3.76
2	CC	**20.10**	**1.68**	**6.94**	**1.56**	1.66	9.33	3.86
	NMI	20.64	6.59	12.98	3.71	1.93	1.56	1.62
	PMI	**18.60**	**2.66**	**12.71**	**1.65**	1.58	7.56	3.65
3	CC	**15.03**	**1.12**	**8.49**	**2.62**	1.94	3.16	1.47
	NMI	15.78	5.28	9.59	3.04	2.02	1.23	1.26
	PMI	**16.91**	**1.87**	**8.79**	**2.77**	1.94	2.84	1.39
4	CC	**11.75**	**0.60**	**4.80**	**1.36**	2.12	1.60	1.80
	NMI	15.03	4.03	9.77	2.06	2.32	1.25	1.19
	PMI	**11.70**	**0.37**	**5.04**	**1.52**	1.94	1.55	1.61
5	CC	**11.70**	**0.52**	**2.74**	**0.58**	2.50	3.80	8.58
	NMI	16.03	4.29	5.68	1.13	2.59	2.50	3.54
	PMI	**10.05**	**0.40**	**2.83**	**0.67**	2.49	3.65	7.15
6	CC	**22.23**	**2.73**	**5.76**	**1.23**	1.62	1.72	4.86
	NMI	22.69	6.26	15.91	4.34	1.98	1.44	2.01
	PMI	**21.98**	**2.59**	**7.96**	**2.89**	1.61	1.71	3.32

real deformation caused by brain shift and mass effect during surgery. This presents a challenge for accuracy assessment and validation. In the synthetic case, we have a readily available source of ground truth: the synthetic deformation. In these clinical cases, there is no synthetic ground truth. In order to assess the accuracy of this registration method using real intra-operative CTs, we collaborate with our clinical partners to derive observed ground truth. For each case, our two experts selected five points in the pre-operative image, and found the same five homologous points in the intra-operative images. We then treat this deformation as the true deformation and compare it to the

deformation recovered at the same points with this non-rigid registration algorithm. The selection of these specific points was left to our clinical partners, and they represent a range of anatomical landmarks, some near the resection margin, and some farther from the resection margin in deep brain tissues.

5.1 Synthetic Ground Truth

We have six cases with pre-operative images. For all of the cases we have both pre-operative MRI and pre-operative CT. We generate a synthetic deformation for each of the cases using the previously described tool. We vary the two input parameters to the synthetic deformation, the variance of the Gaussian distribution σ, and the spacing of the deformation lattice ς to simulate different clinical situations. We measure the resulting deformation only in the region of interest, and report these statistics in **table 5.1**. For each case we also report:

(i) the maximum and average of the block matching error distribution

(ii) the maximum and average of the overall registration error distribution

(iii) the ratio of the Hausdorff distance between the two images when non-rigidly registered to that of the images rigidly registered

(iv) the ratio of the normalized mutual information between the two images when non-rigidly registered to that of the images rigidly registered

in **table 5.2**.

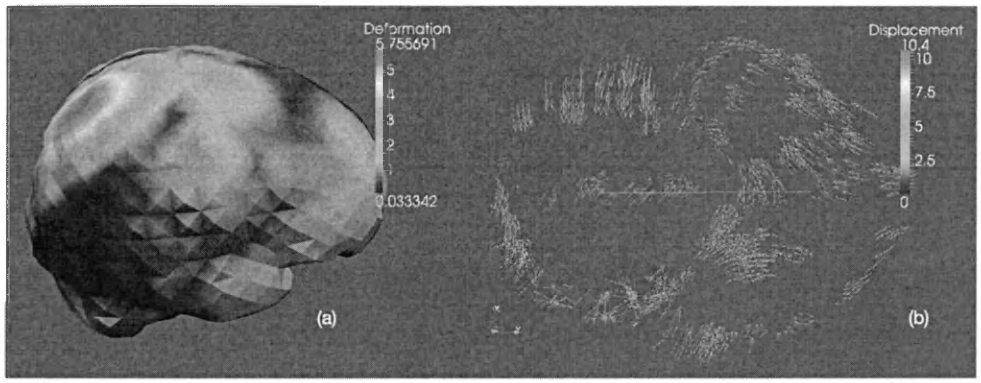

Figure 5.1: An example of the final registration error (a) and sample block matchings (b) from a thin slice of case 6. Note the presence of outliers in the block matchings, and the relatively large regions without any registration points.

5.1.1 Problems with synthetically-derived ground truth

Synthetic data in the form of generated deformations which produce synthetic intra-operative images constitute an important and standard way of assessing the performance of a registration method. This method is used repeatedly in the literature. Synthetic cases, however, are not sufficient simulators of clinical experience. The deformations we generate and apply in these simulations do not accurately reflect the forces which act on the brain during surgery. Real tissue has varying elastic properties (and thus different reactions to the same applied force) which are not simulated or compensated for with the described deformation generation technique. Synthetic ground truth is an important, but incomplete, method of validation for these reasons.

5.2 Expert-produced Ground Truth

Real intra-operative brain shift is very different than the best synthetically derived deformations. The problem of producing good ground truth from real intra-operative

Table 5.3: A summary of the results for the clinical study including the magnitude of expert-derived deformation, the rigid registration landmark error, and the non-rigid registration landmark error.

Case	Point	True Deformation (mm)	RR Error	NRR Error (mm)		
				CC	NMI	PMI
4	1	(-0.43,-0.87, 1.25)	1.59	0.89	1.81	0.71
	2	(-0.43,6.54, -2.50)	7.02	7.39	7.71	9.59
	3	(-0.87, 0.00,-0.50)	0.87	3.64	4.72	2.56
	4	(-3.49, -3.06, -3.75)	5.96	3.47	2.45	5.99
	5	(-1.75,-0.45, 0.00)	1.80	2.10	0.79	2.40
	6	(-6.55, -0.87, 0.00)	6.60	6.81	5.41	7.26
	7	(0.43, 0.87, 0.00)	0.98	1.11	2.48	2.56
	8	(0.44, 1.30, 0.00)	1.38	3.55	6.06	3.49
	9	(-0.44, 1.31, 1.25)	1.86	4.33	6.66	3.67
5	1	(-1.31, 0.87, -1.00)	1.86	1.82	1.60	2.95
	3	(-0.44,1.74, 3.00)	3.50	1.39	4.57	1.73
	4	(-0.44,0.43, 3.00)	3.06	0.88	2.53	1.39
	5	(-3.49, 0.87,1.00)	3.73	4.46	4.20	5.25
	6	(-0.87, 4.37, 2.00)	4.88	4.23	4.83	4.39
	7	(2.18, 4.36, 2.00)	5.27	4.22	5.12	3.76
	8	(0.44, 3.49, 0.00)	3.52	3.29	4.56	3.21
	9	(0.00,4.80, 0.00)	4.80	4.98	6.28	4.62
	10	(-0.87, 3.06, 0.00)	3.18	2.53	2.99	2.72
	11	(0.44, 5.23, 1.00)	5.35	5.20	6.09	5.12
6	1	(-3.93 2.18,-5.00)	6.72	5.92	4.58	6.04
	2	(1.74,3.05, 2.50)	4.32	2.66	6.26	4.36
	3	(2.18,-0.44, 2.50)	3.35	2.51	8.04	5.74
	4	(1.75,0.00, 2.50)	3.05	2.19	4.91	3.29
	5	(-1.75, 5.67, 5.00)	7.76	6.74	10.69	7.88
	6	(1.75, 4.36, 5.00)	6.86	4.60	7.17	4.79
	7	(0.00,-1.31, -2.50)	2.82	3.44	2.43	3.38
	8	(1.74, 3.49, 0.00)	3.90	2.98	5.92	2.94
	9	(5.24, 2.62, 0.00)	5.85	6.25	6.80	7.61

Table 5.4: A summary of the results for the clinical study. A blank entry indicates no improvement, or a negative improvement.

Case	Avg. RR Error	NRR Improvement		
		CC	NMI	PMI
1	3.04	-	-	-
2	3.56	18%	-	10%
3	4.96	20%	-	-

images, however, is a serious impediment to the evaluation and validation of registration techniques such as ours. As previously described, this section details the challenges we encountered generating and using this ground truth, and the results of three clinical cases using real intra-operative, and in one case post-operative, data.

5.2.1 Problems with expert-produced ground truth

Finding specific anatomical landmarks in medical images is a difficult and onerous task in and of itself. Our task is made even more challenging owing to the nature of our medical images. Our pre-operative CT are generally detailed enough and have sufficient spatial resolution to identify — with sub-millimeter accuracy — a number of anatomical landmarks. The intra-operative images we use, however, present a real challenge. Images acquired during surgery suffer from several undesirable phenomena. The intra-operative images we received had a significantly lower signal-to-noise ratio than the pre-operative images of the corresponding patient. This is due to the lower power used by the X-ray tube in the intra-operative CT machine, which results in the previously described quantum mottling. A second source of difficulty, which exacerbates the first, is the presence of movement artifacts due to the (unconscious) patient's breathing motion in the device. This creates motion artifacts which manifest as streaks or concentric circular

shapes. These two phenomena combine to make the reliable identification of homologous anatomical landmarks between intra-operative and pre-operative images difficult, error-prone and sometimes impossible. Our experts captured ten anatomical landmarks on each set of images. The result of this registration method as compared with their anatomical landmarks is listed in table 5.3 and table 5.4.

Chapter 6

Discussion

In this thesis we have performed several practical feasibility studies of an FEM-based non-rigid registration algorithm using intra-operative CT. The studies involved several steps. We have worked with our clinical partners to acquire and catalogue a database of six cases, comprising multiple images of real patients. This includes six sets of pre-operative CT and MRI, of which half are supplemented by intra-operative or post-operative CT follow-ups. We used these data to rigorously evaluate the non-rigid registration method. To this end we developed two mutual-information-based metrics to supplement the existing similarity metric, and created a synthetic benchmark to evaluate the performance of these metrics applied to all six cases.

The synthetic benchmarks show that the normalized cross-correlation-based block matching and registration algorithm out-performed the others, and achieved an average accuracy improvement of **4.14** times. This is compared to **1.84** times improvement for Normalized Mutual Information, and a **3.48** times average improvement for Regional Mutual Information. This result is not surprising, as this mono-modal image registration problem (CT-to-CT) allows us to be confident in an assumption of an affine relationship

between the intensity distributions of the fixed and floating images. This assumption leads us naturally to the NCC as a similarity metric. The NMI clearly performs poorly, most likely due to the extreme sensitivity of the joint histogram calculation to a dearth of sample data as we have in our small blocks. That said, our results show that the PMI metric is nearly as accurate as NCC, and clearly does not suffer from the shortcomings of the NMI metric, being on average **1.89** times as accurate as the NMI. We suggest that the PMI be evaluated further as a multi-modal block matching metric in this framework.

One of the questions we attempt to answer in this study is whether CT has enough soft-tissue definition to successfully complete our relatively sparse displacement-recovering block-matching approach. We believe that we have shown that, given a sufficiently high resolution pre-operative CT, we can recover a deformation to within **1.39 mm** of the actual synthetic value on average, which is much better than rigid registration alone (**4.35 mm** average error). This study suggests strongly that, given an intra-operative image of similar articulation, we may expect similar accuracy. We conclude that this synthetic study implies that block-matching between a pre-operative and intra-operative CT is possible using either the Normalized Correlation Coefficient or the Pointwise Mutual Information.

We next explored the three cases which present clinically-acquired intra-operative CT. The three cases involved correspond to cases 4 – 6 of the synthetic study. That is, the pre-operative CT for cases 4 – 6 in the synthetic study are the same images as the pre-operative CT for cases 1 – 3 in the clinical study. Cases 1 and 2 involve cancerous neoplasms of the frontal and ethmoidal sinus cavities, while case 3 involves a defect of the skull base, resulting in an nasoethmoidal encephalocele. In cases 1 and 2, we attempt to register the pre-operative CT with a CT acquired intra-operatively. In case 3, the fixed

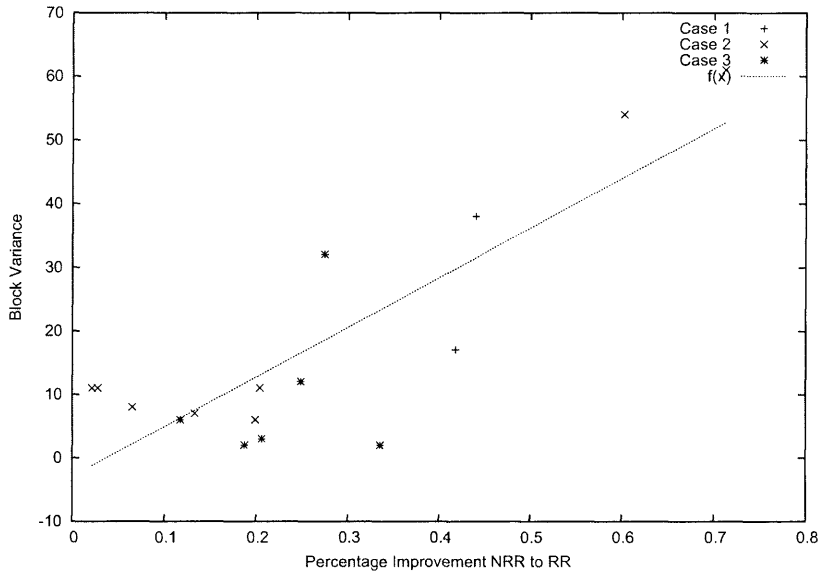

Figure 6.1: A positive relationship between registration improvement (NRR-to-RR) and block variance at registration landmarks in the clinical study is exhibited. The OLS line of best fit has slope = 78.33 variance units for every 100% improvement in accuracy.

image is a post-operative CT acquired on a similar scanner as the pre-operative CT. There was no CT contrast present in the images, with the exception of the pre-operative case 3 image. We first note that the pre-operative images in all cases have superior fidelity than the the intra- (and post-) operative images. This manifests itself in several ways: the intra-operative images of cases 1 and 2 have a relatively low signal-to-noise ratio and are visually very noisy and grainy. The post-operative scan available for case 3 has a visually similar noise profile to the pre-operative scans, but has a lower soft-tissue articulation. In addition, the post-operative scan has fewer slices in the Z-direction. These are all undesirable properties in an image registration problem. If the noise in an image overpowers the features of the image, the similarity metric we use will perform poorly. The registration is also sensitive to the resolution of the input images, as it will

be unable to detect block matchings which reside in between the available slices. In practice, we increase the resolution of the registered image by interpolating at points between the available data. This interpolation will obviously lead to inaccuracies as the interpolation nodes (available image slices) have a larger slice spacing. The fidelity of a medical image such as those we deal with depends on many scan parameters — and setting those parameters is a problem of mutual accommodation of sometimes competing goals.

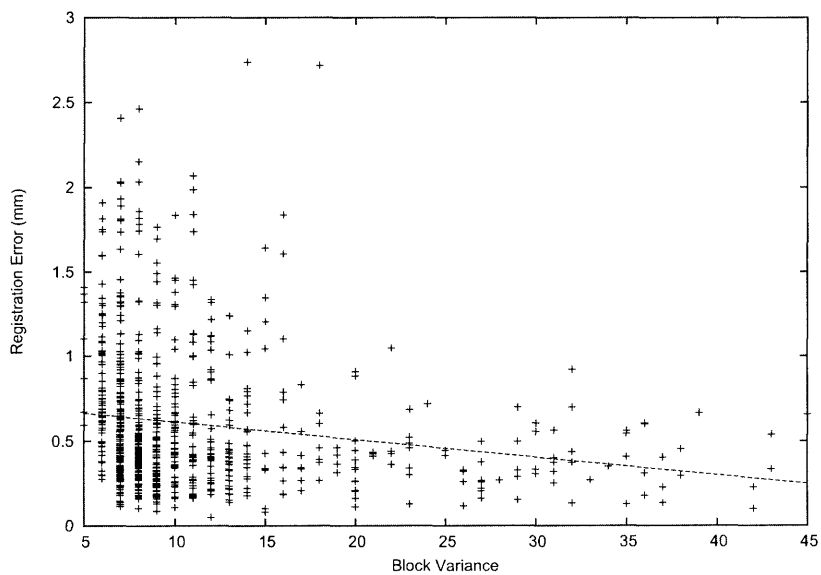

Figure 6.2: A positive relationship between registration improvement (NRR-to-RR) and block variance at registration mesh points in the synthetic study is exhibited.

The landmark acquisition for the three clinical cases was difficult. Both of our physician experts commented that the quality of the intra-operative scans made finding reliable landmark pairings a hard task. Some of the landmarks captured have a relatively high confidence: the ventricles are an easy feature to discern, even in a very bad image, and some of the landmarks (particularly near the tumor of interest) lie in relatively CT-

homogenous tissue. That said, the difficulty of landmark matching in CT is an elocution of the motivation for our study. Nevertheless, expert-acquired landmarks are the best ground truth available to us.

Our results for the clinical study show that more work remains to be done to use this registration method with intra-operative CT. The improvement in individual landmark error was between **3.5 times** better and **4.18 times** worse than rigid registration, with an average improvement of only **1.19** times, or **19%**. That said, there is a correlation between the block variance around an individual landmark and the registration accuracy at that landmark. The higher the block variance is at a landmark, the higher the relative improvement is in registration accuracy. This relationship is described in **figure 6.1**. We ran a similar study of a synthetic deformation, which show the same trend in **figure 6.2**: increased block variance yields lower registration error.

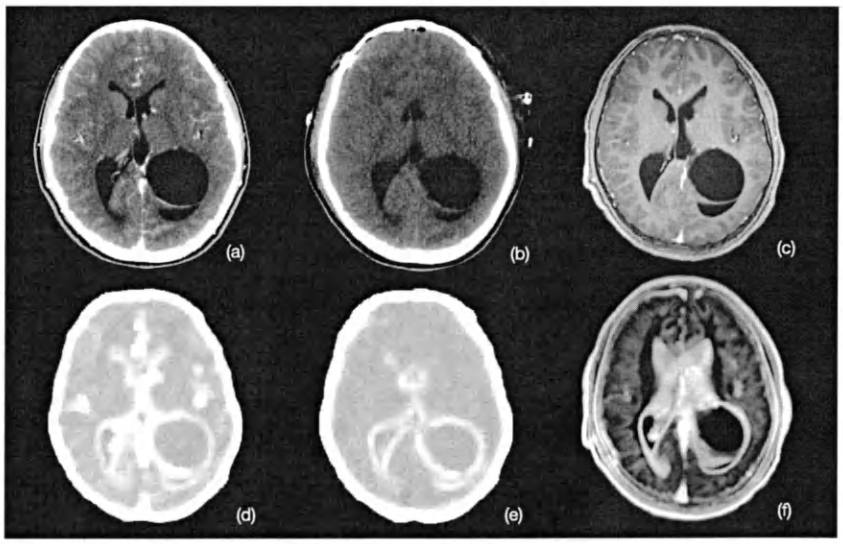

Figure 6.3: Case 3 pre-op (a), pre-op block variance (d), intra-op (b), intra-op block variance (e), pre-op MRI (c), and pre-op MRI block variance (f).

There are two relevant sources of the variance in our images. First is the variance

of a block due to the presence of an anatomical or functional feature in the image. We call this source the "feature variance." The second source of variance in a block is noise, either random quantum mottle or artifact noise, as previously described. We call this "noise variance." Refer to **Figure 6.3** for the pre-op CT, pre-op MRI, post-op CT, and block variance images for each of these images from case 6. These block variance images are produced by calculating the variance of a three dimensional isotropic block centered at each voxel coordinate. We see that in the pre-operative image (a), with contrast, the block variance (d) is largely coincident with the edge features in the image. The ventricles, large cyst, sulci, and midline are all displayed in the block variance image. In addition, regions which are largely structurally void in the CT are largely homogenous in the block variance image. This is a characteristic conducive to high registration accuracy. On the other hand, in the intra-operative image (b), the intra-operative block variance image (e) shows less feature variance and more noise variance than the pre-operative images. In other words, the variance pattern does not coincide with structural edges or features in the image, but rather is a result of noise. This strongly suggests a hypothesis for the relatively poor performance of the clinical study: the intra-operative images used have poor articulation of soft tissues and low signal-to-noise ratio, which causes block matching to generate a large number of erroneous matches, or outliers. Contrast this with the pre-operative MRI (c) and its block variance image (f), which shows a much larger range of feature variance, and less homogenous and noisy variance regions. The same is illustrated in **figure 6.4**, where we see the distribution of block variance is much broader in MRI than CT. We draw the conclusion that the major limiting factor in the improvement of this non-rigid registration method is the acquisition of superior quality intra-operative CT scans. The dual of this problem is

the development of noise-robust similarity metrics, and similarity metrics which perform well in the absence of highly articulated soft tissue. Alternatively, it would be worth investigating to extend Andronache et. al.'s method of detecting structureless patches to choose registration points, rather than simply choosing the highest variance blocks. Andronache uses Moran's spatial autocorrelation coefficient to do exactly this task [2]. Alternatively, it may be beneficial to use edge detection images to weight possible feature points by their inclusion or proximity to a detected edge.

The studies performed in this thesis are a good first step toward the improvement of this state-of-the-art non-rigid registration algorithm when used with intra-operative CT. The objective of achieving superior performance with intra-operative CT is important for reducing the cost and increasing the accessibility of real-time non-rigid registration for image guided neurosurgery. We believe the studies included here, particularly the synthetic studies, show the feasibility of this goal.

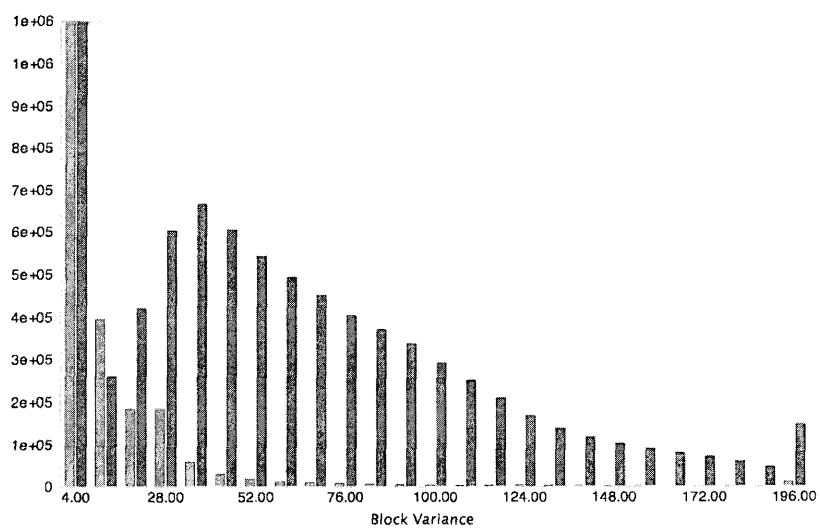

Figure 6.4: A histogram of the block variance levels in a CT (green) and an MRI (red).

Bibliography

[1] A. Andronache, P. Cattin, and G. Szekely. Adaptive subdivision for hierarchical non-rigid registration of multi-modal images using mutual information. In *Medical Image Computing and Computer-Assisted Intervention MICCAI 2005*, Lecture Notes in Computer Science, pages 976–983. Jan 2005.

[2] A. Andronache, M. von Siebenthal, and G. Szkely. Non-rigid registration of multi-modal images using both mutual information and cross-correlation. *Medical Image Analysis*, 12(1):3 – 15, Jan 2008.

[3] Neculai Archip, Olivier Clatz, Andriy Fedorov, Andriy Kot, Stephen Whalen, Dan Kacher, Nikos Chrisochoides, Ferenc Jolesz, Alexandra Golby, Peter Black, and Simon K. Warfield. Non-rigid alignment of pre-operative mri, fmri, dt-mri, with intra-operative mri for enhanced visualization and navigation in image-guided neurosurgery. *Neuroimage*, 35(2), 2007.

[4] Neculai Archip, Olivier Clatz, Stephen Whalen, Dan Kacher, Andriy Fedorov, Andriy Kot, Nikos Chrisochoides, Ferenc Jolesz, Alexandra Golby, Peter M. Black, and Simon K. Warfield. Non-rigid alignment of pre-operative mri, fmri, and dt-mri with intra-operative mri for enhanced visualization and navigation in image-guided neurosurgery. *NeuroImage*, 35(2):609 – 624, 2007.

[5] Eric Billet, Andriy Fedorov, and Nikos Chrisochoides. The use of robust local hausdorff distances in accuracy assessment for image alignment of brain mri. *The Insight Journal*, January 2008.

[6] P. Black, T. Moriarty, E. Alexander III, and P. Stieg. Development and implementation of intraoperative magnetic resonance imaging and its neurosurgical application. *Neurosurgery*, 41(4), Jan 1997.

[7] R. A. Brooks and G. Di Chiro. Beam hardening in x-ray reconstructive tomography. *Physics in Medicine and Biology*, 21(3):390–398, 1976.

[8] Lisa Gottesfeld Brown. A survey of image registration techniques. *ACM Comput. Surv.*, 24(4):325–376, 1992.

[9] Eric C. Burton and Michael D. Prados. Malignant gliomas. *Current Treatment Options in Oncology*, 1(5):459 – 68, 2000.

[10] Nikos Chrisochoides, Andriy Fedorov, Andriy Kot, Neculai Archip, Peter Black, Olivier Clatz, Alexandra Golby, Ron Kikinis, and Simon K. Warfield. Toward real-time, image guided neurosurgery using distributed

and grid computing. In *IEEE/ACM International Conference for High Performance Computing, Networking, Storage and Analysis (Supercomputing)*, 2006.

[11] O. Clatz, H. Delingette, I.F. Talos, A.J. Golby, R. Kikinis, F.A. Jolesz, N. Ayache, and S.K. Warfield. Robust non-rigid registration to capture brain shift from intra-operative mri. In *IEEE Transactions on Medical Imaging*, volume 24, pages 1417–1427, 2005.

[12] D. Louis Collins. *3D Model-based segmentation of individual brain structures from magnetic resonance imaging data.* PhD thesis, McGill University, 1994.

[13] L. Davis. A survey of edge detection techniques. *Computer Graphics and Image Processing*, 4:248–270, 1975a.

[14] M. Ehlers. Region-based matching for image registration in remote sensing databases. *Proceedings of the International Geoscience and Remote Sensing Symposium*, pages 2231–2234, 1991.

[15] Andriy Fedorov and Nikos Chrisochoides. Tetrahedral mesh generation for non-rigid registration of brain mri: Analysis of the requirements and evaluation of solutions. In *International Meshing Roundtable*, number 17th. Springer Verlag, October 2008.

[16] Andriy Y. Fedorov. *Enabling Technology for Non-Rigid Image Registration during Image-Guided Neurosurgery.* PhD thesis, The College of William & Mary, 2009.

[17] W. Freysinger, A. Gunkel, and W. Thumfart. Image-guided endoscopic ent surgery. *European Archives of Otolaryngology*, Jan 1997.

[18] Marvin P. Fried, Sanjay R. Parikh, and Babak Sadoughi. Image-guidance for endoscopic sinus surgery. *The Laryngoscope*, 118(7):1287–1292, 2008.

[19] T. Gaens, F. Maes, D. Vandermeulen, and P. Suetens. Non-rigid multimodal image registration using mutual information. In *MICCAI '98: Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention*, Lecture Notes in Computer Science. Jan 1998.

[20] Thomas Gasser, Oliver Ganslandt, Erol Sandalcioglu, Dietmar Stolke, Rudolf Fahlbusch, and Christopher Nimsky. Intraoperative functional mri: implementation and preliminary experience. *Neuroimage*, 26(3):685–93, Jul 2005.

[21] Joseph K. Han. Personal Correspondence, September 2009.

[22] Nobuhiko Hata, Takeyoshi Dohi, Simon K. Warfield, William M. Wells III, Ron Kikinis, and Ferenc A. Jolesz. Multimodality deformable registration of pre- and intraoperative images for mri-guided brain surgery. In *MICCAI '98: Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 1067–1074, London, UK, 1998. Springer-Verlag.

[23] Stephen J. Hentschel and Raymond Sawaya. Optimizing outcomes with maximal surgical resection of malignant gliomas. *Cancer Control*, 10(2), 2003.

[24] Gerardo Hermosillo, Christophe Chefd'Hotel, and Olivier Faugeras. Variational methods for multimodal image matching. *International Journal of Computer Vision*, 50(3):329–343, July 2002.

[25] Derek L. G. Hill, Philipp G. Batchelor, Mark Holden, and David J. Hawkes. Medical image registration. *Physics in Medicine and Biology*, 46(3):R1–R45, 2001.

[26] L. Ibanez, W. Schroeder, L. Ng, and J. Cates. *The ITK Software Guide*. Kitware, Inc. ISBN 1-930934-15-7, http://www.itk.org/ItkSoftwareGuide.pdf, second edition, 2005.

[27] Slaney M. Kak, A.C. *Principles of computerized tomographic imaging*. IEEE Service Center, 1988.

[28] B. Likar and F. Pernus. A hierarchical approach to elastic registration based on mutual information. *Image and Vision Computing*, 19(1-2):33 – 44, 2001.

[29] Yixun Liu, Andriy Fedorov, Ron Kikinis, and Nikos Chrisochoides. Real-time non-rigid registration of medical images on a cooperative parallel architecture. In *IEEE International Conference on Bioinformatics Biomedicine*, November 2009.

[30] Yixun Liu, Andriy Fedorov, Ron Kikinis, and Nikos Chrisochoides. A data parallel non-rigid medical image registration method on cooperative architecture. *Parallel Computing*, preprint.

[31] F. Maes, A. Collignon, and D. Vandermeulen. Multi-modality image registration maximization of mutual information. *Proceedings of MMBIA*, Jan 1996.

[32] J. B. Antoine Maintz, Erik H. W. Meijering, and Max A. Viergever. *General Multimodal Elastic Registration Based on Mutual Information*, pages 144–154. SPIE Press, Jan 1998.

[33] C. J. Martin, D. G. Sutton, and P. F. Sharp. Balancing patient dose and image quality. *Applied Radiation and Isotopes*, 50(1):1 – 19, 1999.

[34] A. Nabavi, P. M. Black, D. T. Gering, C. F. Westin, V. Mehta, R. S. Pergolizzi, M. Ferrant, S. K. Warfield, N. Hata, R. B. Schwartz, W. M. Wells, R. Kikinis, and F. A. Jolesz. Serial intraoperative magnetic resonance imaging of brain shift. *Neurosurgery*, 48(4):787–97; discussion 797–8, Apr 2001.

[35] C. Nimsky, O. Ganslandt, S. Cerny, P. Hastreiter, G. Greiner, and R. Fahlbusch. Quantification of, visualization of, and compensation for brain shift using intraoperative magnetic resonance imaging. *Neurosurgery*, 47(5):1070–9; discussion 1079–80, Nov 2000.

[36] Nikhil R. Pal and Sankar K. Pal. A review on image segmentation techniques. *Pattern Recognition*, 26(9):1277 – 1294, 1993.

[37] J. Pluim, J. Maintz, and M. Viergever. Mutual-information-based registration of medical images: a survey. *IEEE transactions on medical imaging*, 22(8), Jan 2003.

[38] Nicola Ritter, Robyn Owens, K. Yogesan, and P. van Saarloos. The application of mutual information to the registration of stereo and temporal images of the retina. In *Advanced Topics in Artificial Intelligence*, Lecture Notes in Computer Science, pages 67–76. 1997.

[39] Alexis Roche, Grgoire Malandain, Xavier Pennec, and Nicholas Ayache. The correlation ratio as a new similarity measure for multimodal image registration. In *Medical Image Computing and Computer-Assisted Interventation MICCAI98*, Lecture Notes in Computer Science. 1998.

[40] P. Rogelj and S. Kovacic. *Point similarity measure based on mutual information.* Lecture Notes in Computer Science. Jan 2003.

[41] D. Rueckert, C. Hayes, and C. Studholme. Non-rigid registration of breast mr images using mutual information. In *Medical Image Computing and Computer-Assisted Interventation MICCAI98*, volume 1496 of *Lecture Notes in Computer Science*. Jan 1998.

[42] D. B. Russakoff, C. Tomasi, T. Rohlfing, and C. R. Maurer Jr. Image similarity using mutual information of regions. In *Computer Vision - ECCV 2004*, volume 3023 of *Lecture Notes in Computer Science*, pages 596–607. 2004.

[43] Ann Scott. *Interventional and Intraoperative Magnetic Resonance Imaging.* Alberta Heritage Foundation for Medical Research,, March 2004.

[44] C. E. Shannon. A mathematical theory of communication. *Bell system technical journal*, 27, 1948.

[45] American Cancer Society. *Cancer Facts & Figures.* Atlanta: American Cancer Society, 2009.

[46] C. Studholme, D.L.G. Hill, and D.J. Hawkes. Automated 3-d registration of mr and ct images of the head. *Medical Image Analysis*, 1(2):163 – 175, 1996.

[47] C. Studholme, D.L.G. Hill, and D.J. Hawkes. An overlap invariant entropy measure of 3d medical image alignment. *Pattern Recognition*, 32(1):71–86, 1999.

[48] Philippe Thevenaz and Michael Unser. Optimization of mutual information for multiresolution image registration. *IEEE Transactions on Image Processing*, 9:2083–2099, 2000.

[49] J.P. Thirion. Image matching as a diffusion process: an analogy with maxwell's demons. *Medical Image Analysis*, 2(3):243 – 260, 1998.

[50] M. Toews, D. Collins, and T. Arbel. Maximum a posteriori local histogram estimation for image registration. In *Medical Image Computing and Computer-Assisted Intervention MICCAI 2005*, volume 3750 of *Lecture Notes in Computer Science*. Jan 2005.

[51] A. S. VASILEISKY AND M. BERGER. Automated co-registration of multi-sensor images on the basis of linear feature recognition for subsequent data fusion. *Proc. Conf. on Fusion of Earth Data*, pages 59–65, 1998.

[52] MARSHALL A. WEISSBERGER. Personal Correspondence, 2009-2010.

[53] W. WELLS, P. VIOLA, H. ATSUMI, AND S. NAKAJIMA. Multi-modal volume registration by maximization of mutual information. *Medical Image Analysis*, Jan 1996.

[54] PAUL A. YUSHKEVICH, JOSEPH PIVEN, HEATHER CODY HAZLETT, RACHEL GIMPEL SMITH, SEAN HO, JAMES C. GEE, AND GUIDO GERIG. User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage*, 31(3):1116–1128, 2006.

[55] ZHENGYOU ZHANG. Iterative point matching for registration of free-form curves and surfaces. *Int. J. Comput. Vision*, 13(2):119–152, 1994.

[56] DJEMEL ZIOU AND SALVATORE TABBONE. Edge detection techniques - an overview. *International Journal of Pattern Recognition and Image Analysis*, 8:537–559, 1998.

[57] BARBARA ZITOV AND JAN FLUSSER. Image registration methods: a survey. *Image and Vision Computing*, 21(11):977 – 1000, 2003.

[58] DARKO KERL, BOTJAN LIKAR, AND FRANJO PERNU. A protocol for evaluation of similarity measures for rigid registration. *IEEE Transactions on Medical Imaging*, 25(6):779–791, June 2006.