AUTOMATIC PLANNING OF A SAFE DRILLING PATH TO BE USED IN

COCHLEAR IMPLANTATION SURGERY USING IMAGE REGISTRATION

TECHNIQUES

By

Hasan Mohammed Al-Marzouqi

Thesis

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

MASTER OF SCIENCE

in

Electrical Engineering

August, 2006

Nashville, Tennessee

Approved:

Professor Benoit Dawant

Professor Michael Fitzpatrick

ACKNOWLEDGMENTS

I should start by thanking my academic advisor, Dr. Benoit Dawant, for his guidance and help throughout this project. I would also like to thank Dr. Michael Fitzpatrick for reviewing this thesis. Acknowledgements must also go to Dr. Robert Labadie for introducing me to the medical background necessary to work through this project. Further thanks go to Dr. Frank Warren for drawing the outlines on the scans used in this project and for verifying the safety of the computed paths. I would also like to thank Rui Li for the time he spent on modifying some of the tools used in the MIP laboratory to suite the demands of this project. Finally, I'd like to express my gratitude to my dear parents for their invaluable continuous support and encouragement.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER I

INTRODUCTION

Image registration is the process of finding the geometric relationship between corresponding points in two images of the same scene. The two images can be taken at different times, with different acquisition methods, or from different views. This makes image registration a fundamental step when trying to integrate information taken from different images. A wide range of applications exist for image registration covering different fields such as remote sensing, computer vision, and medical imaging. Over the last decade significant technical developments occurred in the area of medical imaging that led to an increase in its use and importance. Nowadays medical images are routinely used in diagnosis, guidance of therapy, and monitoring of disease progression. More recently it has been used for image guided surgery (IGS) in the operating room. Consequently, the use and importance of image registration have risen significantly. For example, in many applications images from different modalities are acquired and the need arises for automatic and accurate ways of combining them. It is also common to acquire a series of images dynamically in a short period of time. Relating those images to each other is essential to extract the required relevant information. In this work image registration is used in atlas based segmentation where relevant structures in one image are labeled and used to identify the same structures in other images. An algorithm is developed using this idea to identify the centroid of two structures in patient images. The path connecting them should be a safe path that can be used in cochlear implantation

surgery. Success of this project would facilitate the surgical procedure currently used for cochlear implantation. In the following chapter, clinical information about the cochlear implantation surgery and a description of the problem are presented. The two registration methods that are used in this work are discussed in Chapters three and four. Chapter five presents details of the developed algorithm followed by the results obtained in chapter six. Conclusions and a view of future work are in chapter seven.

CLINICAL BACKGROUND AND PROBLEM DESCRIPTION

Cochlear implantation surgeries are usually performed on patients having profound to total hearing loss. In a normal ear sound waves travel down the ear canal and hit the ear drum in the middle ear. This causes the ear drum to vibrate. Those vibrations are amplified by three bones in the middle ear: the malleus, incus, and the stapes (Figure 1) before reaching the cochlea in the inner ear. The cochlea is filled with liquid that carries those vibrations to thousands of hair cells launching an electric signal that travels through the auditory nerve to the brain for processing.
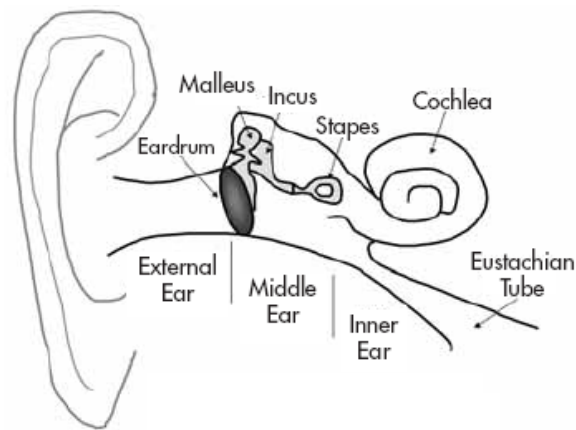
Figure 1: Ear anatomy [1]

Different problems can arise affecting the functionality of normal ears. Approximately 90% of hearing loss is caused by damage to the inner ear (the cochlea). Hearing aids are

typically used in those cases to treat mild and moderate levels of hearing loss. However, when the hearing loss is severe, cochlear implantation becomes the preferred treatment.

Cochlear Implants

A cochlear implant bypasses the functionality of the inner ear. It stimulates the auditory nerve directly by an electrode array that is threaded into the cochlea. This electrode array is connected to an internal receiver placed beneath the skin that is magnetically linked to an external receiver/processor. The external receiver converts sound waves into electric pulses and sends the appropriate signal to the cochlea through the internal receiver.
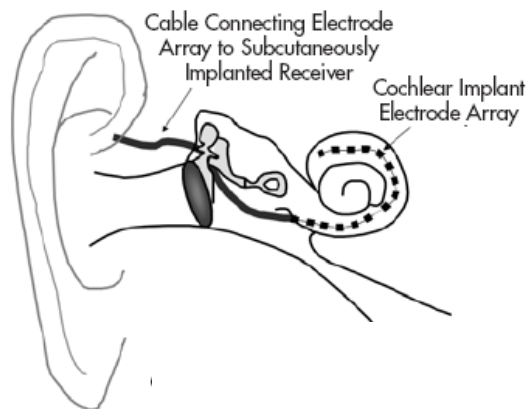
Figure 2: Cochlear implantation [1]

Access To The Cochlea

In the current surgical procedure access to the cochlea is based on a wide surgical exposure (approximately 6 cm incision) to identify anatomic landmarks. Those landmarks are used to avoid injury to any critical structures such as the facial nerve. Entrance to the middle ear is made through the *facial recess*; a region bounded by the facial nerve and its branch the chorda tympani (Figure 3). The cochlea is identified after

4

opening the recess and a hole is made into its basal turn. The electrode array is then

inserted into the cochlea, the internal receiver is secured to the skull, and the incision is

closed. More information about the procedures can be found in [2]-[4].
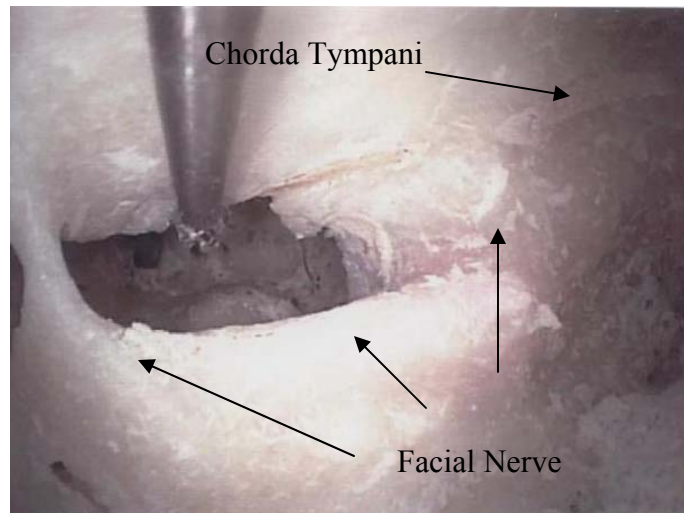
Figure 3: The surgical drill going through the facial recess - shown drilled out -
on its way to the cochlea. Critical structures surrounding it are shown on the image

The New Surgical Technique

Improvements to this surgical procedure can be made if the drill's path was determined

before the surgery. The drill can then be constrained to pass in the pre-determined

trajectory using a drill guide. This guide is custom built based on the path found on CT

images for each patient. This would allow access to the cochlea without the need for a

wide opening. The new procedure would reduce the time and the cost of the surgery

(down to 30 minutes instead of about an hour). It will also standardize CI placement by

removing the surgical technique as a variable of outcome and allow more surgeons to

safely place CI's. Activation of implants could happen on the same day as the surgery

(currently 1-3 weeks are required). But, determining a safe path requires finding two

points through which the path must pass. Paths accessing the basal turn of the cochlea through the facial recess are considered safe and appropriate. By locating a point on the facial recess and a point on the basal turn on CT scans a suitable path can be found. Unfortunately, a majority of surgeons performing CI surgeries find it difficult to locate such a path on CT images. An algorithm helping them to do it would thus be useful and important. In this project an automatic way of finding such a path is developed, using techniques of atlas based segmentation.

CHAPTER III

IMAGE REGISTRATION

Geometrical Transformations

Rigid transformations

The goal of registration is to find a transformation that relates every point in an image with its corresponding point in the second image. For objects that are rigid or approximately rigid such as bones or the head, possible transformations are limited and are dependent on few parameters. In 3D the transformation can be completely described by three translations t_x, t_y, and t_z and three rotations $theta_x$, $theta_y$, and $theta_z$. To apply a shift by an amount of a in the x-direction, b in the y-direction and c in the z-direction to a point **X** one can use the following equation:

$$\mathbf{X'} = \mathbf{X} + \begin{bmatrix} a \\ b \\ c \end{bmatrix}$$

Rotations require multiplying the coordinates of the point by a 3 by 3 matrix that depends on the rotation axis and the rotation angle:

To rotate around the x-axis: $\mathbf{X'} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta) & \sin(\theta) \\ 0 & -\sin(\theta) & \cos(\theta) \end{bmatrix} * \mathbf{X}$

To rotate around the y-axis: $\mathbf{X'} = \begin{bmatrix} \cos(\theta) & 0 & \sin(\theta) \\ 0 & 1 & 0 \\ -\sin(\theta) & 0 & \cos(\theta) \end{bmatrix} * \mathbf{X}$

7

To rotate around the z-axis: $\mathbf{X'} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix} * \mathbf{X}$

Multiplication of the above three rotation matrices will generate a matrix representing different possible combinations of rotations around the three axes:

$$\mathbf{R} = \begin{bmatrix} \cos(\beta)\cos(\gamma) & -\cos(\alpha)\sin(\gamma)+\sin(\alpha)\sin(\beta)\cos(\gamma) & \sin(\alpha)\sin(\gamma)+\cos(\alpha)\sin(\beta)\cos(\gamma) \\ \cos(\beta)\sin(\gamma) & \cos(\alpha)\cos(\gamma)+\sin(\alpha)\sin(\beta)\sin(\gamma) & -\sin(\alpha)\cos(\gamma)+\cos(\alpha)\sin(\beta)\sin(\gamma) \\ -\sin(\beta) & \sin(\alpha)\cos(\beta) & \cos(\alpha)\cos(\beta) \end{bmatrix}$$

Where α is the angle of rotation around the x-axis, β is the angle of rotation around the y-axis, γ is the angle of rotation around the z-axis. The above representation is the "Euler angles" parameterization of rotation. Other representation systems that are also used are the *axis-and-angle* specification and the *quaternion* specification [5]. Combining rotation and translation in one equation yields the following equation:

$$\mathbf{X'} = \mathbf{RX} + \mathbf{t}$$

Where $\mathbf{t}$ is a 3 by 1 translation vector and $\mathbf{R}$ is the above rotation matrix.

Scaling transformations

The six degrees of freedom transformation discussed in the previous section is not appropriate when the registration is to be done between objects of different sizes. While intra-patient registration problems can usually be solved using rigid transformation, inter-patient registration will require the addition of more parameters. Scaling transformations use nine degrees of freedom. Three scaling factors, one for each of the coordinate axes,

are used in addition to the six parameters used in rigid registration. The transformation takes the form of:

$$\mathbf{X'} = \mathbf{RSX} + \mathbf{t} \quad \text{or} \quad \mathbf{X'} = \mathbf{SRX} + \mathbf{t}$$

Where $\mathbf{S} = \begin{bmatrix} sx & 0 & 0 \\ 0 & sy & 0 \\ 0 & 0 & sz \end{bmatrix}$, and sx, sy, and sz are constants representing scaling factors.

Since $\mathbf{RS}$ is not necessarily equal to $\mathbf{SR}$, two representations are necessary.

Affine transformations

The elements of the rotation matrix are so far determined by the use of trigonometric functions. Three other parameters can be added to the transformation by removing that constraint and letting all the parameters be independent. The resulting 12 degrees of freedom transformation is called *affine transformation*. This transformation allows for skewing in the direction of the three coordinate axes and can be represented by the following equation:

$$\begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} & t_1 \\ a_{21} & a_{22} & a_{23} & t_2 \\ a_{31} & a_{32} & a_{33} & t_3 \end{bmatrix} * \begin{bmatrix} x \\ y \\ z \end{bmatrix}$$

Where (x,y,z) are coordinates of the point in the original image and (x',y',z') are the new coordinates of the points. The above 4 by 4 matrix combines translation parameters with the other parameters in one matrix, which simplifies mathematical operations. This transformation preserves the straightness of lines but allows the angles between them to change. The task of an affine registration algorithm is to find values for those twelve

parameters that are best in aligning corresponding points in the given two images. The following figures show examples of the above transformations applied on a 2-D image.

Figure 4: Examples of translation and rotation applied to the original image to the left
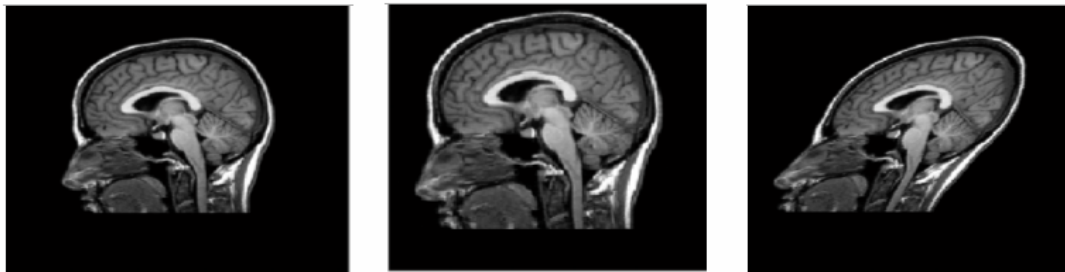
Figure 5: Examples of scaling and skewing applied to the original image on the left

It is obvious that rigid and affine transformations have global effects over all voxels of the image. *Curved non rigid transformations* are used in applications where it is necessary to use more flexible transformations with higher degrees of freedom. Such transformations allow for local deformations. The following figure shows an example of a non rigid registration transformation:
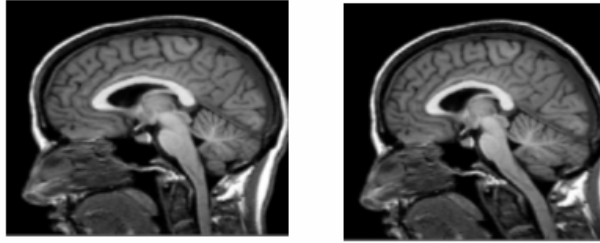
Figure 6: Non rigid transformation example, original image is to the left

Applications where the use of those transformations becomes a necessity include registering pre and post surgery images, inter-patient registration, and atlas based segmentation. In this project due to variability in the structures of interest between patients, non rigid transformations are used in deforming the facial recess and the target window in the atlas to match them with their correspondent structures in the unknown images.

Intensity Based Registration

Introduction

A wide variety of techniques have been developed for registering medical images, which can be classified in many ways. Based on the technique used for registration we can classify most of them into three categories: *point-based* methods, *surface-based* methods and *intensity-based* methods. In points-based methods a set of corresponding points are identified in the two images to be registered. After that the transformation required to bring them into alignment is computed. Those points, called *fiducials*, can be identified visually by an experienced user who looks at the images to localize well defined points at the resolution level of the scan in the images to be registered. Alternatively those points

can be localized by markers that are attached to the anatomy and are designed to be accurately localizable [6]. If the transformation is confined to be rigid then the mathematical solution that calculates the necessary transformation has been known for many years and the algorithm used to calculate it is straight forward [7]. Surface based methods approach the problem by using sets of points representing surfaces of interest in the two images. Different algorithms are then employed to find the transformation that best aligns these surfaces. The Head and Hat approach developed by Pelizzari *et al*. [8] and iterative closest point (ICP) approach [9] are examples of such algorithms. The third category covers intensity-based methods. Unlike point and surface based methods those methods rely only on the intensity of the images. Because of their automatic nature and the good results that were obtained by them in recent years these methods are now widely used for a variety of applications. In those methods a *similarity measure* is used to quantitatively measure how well the images are registered. Obviously if the images are well registered they will look similar to one another. The parameters of the geometrical transformations used to align the two images are optimized by using the similarity measure. An optimization function applies different transformations to an image B and the value of the similarity measure computed on the other image A and the transformed B is used to predict better values for the parameters until the optimization algorithm converges. The work done in this project relies primarily on these methods. Hence they will be discussed in detail.

Similarity measures

Sum of Squared Difference (SSD) is one of the simplest similarity measures. SSD is calculated over the set of N overlapping voxels *i* between images A and transformed image B which is referred to as B^T in the following equation:

$$SSD = \frac{1}{N}\sum_i [A(i) - B^T(i)]^2$$

This measure is the optimum measure when working with images that differ only by Gaussian noise [10] and is widely used for serial MR registration [11]. Correlation Coefficient CC requires a less strict assumption. It can generate optimum results if the relationship between the intensities of the two images is linear

$$CC = \frac{1}{N}\frac{\sum_i (A(i) - \overline{A})\times(B^T(i) - \overline{B^T})}{\{\sum_i (A(i) - \overline{A})^2 \times \sum_i (B^T(i) - \overline{B^T})^2\}^{1/2}}$$

Where $\overline{A}$ and $\overline{B}$ are the mean intensity values of voxels in images A and transformed B, N is the number of overlapping voxels *i*. Possible values of CC range from 0 to 1. A value of 1 indicates that the two images are identical. Since the assumption of linearity is not generally true, researchers proposed other measures that use other indicators of alignment. In 1992 a new similarity measure, the *Ratio Image Uniformity* (RIU), was introduced by Woods [12]. In this algorithm a ratio image is calculated from images A and B and the uniformity of the ratio image (its normalized standard deviation) was used as a measure of similarity. The ratio image should have a uniform constant value if the images are well registered. Deviations in the ratio image indicate misalignments. The *Partitioned Intensity Uniformity* (PIU) developed also by Woods shortly after his RIU

algorithm [13] sums the normalized standard deviations of voxel values in an image for each intensity value/partition in the other image. Minimizing PIU is the goal of the optimization algorithm. The underlying assumption used is that voxels corresponding to an intensity value in an image A should have intensities similar to each other since they correspond to the same tissue type. By plotting 2-D histograms of the intensities between the images to be registered one can notice that the more scattered the points on the histogram are the higher the misalignment. Researchers recognized that *Mutual information* (MI), a basic concept from information theory, can be used as a function to measure this scattering [14-15] and results started to appear showing success of mutual information as a similarity measure. Since its introduction MI has appeared in numerous publications achieving very good results that made it one of the best and most commonly used measures of similarity.

Mutual Information

Introduction

A graphical description of the relationship between intensities in two images can be obtained by plotting their 2-D histogram. The axes of the histogram are the intensities in each image. The value of a point (a,b) on the histogram is equal to the number of voxels that have the intensity value a in the first image and the intensity value b in the second image. Obviously sets of intensity partitions can also be used instead of the individual intensities. Plotting this histogram for identical images will generate a straight line. As one image is spatially transformed with respect to the other the points on the

14

histogram will start to scatter. This scattering effect will also be observed if the images are from different modalities. Figure 2 shows examples of this effect

(a)

(b)
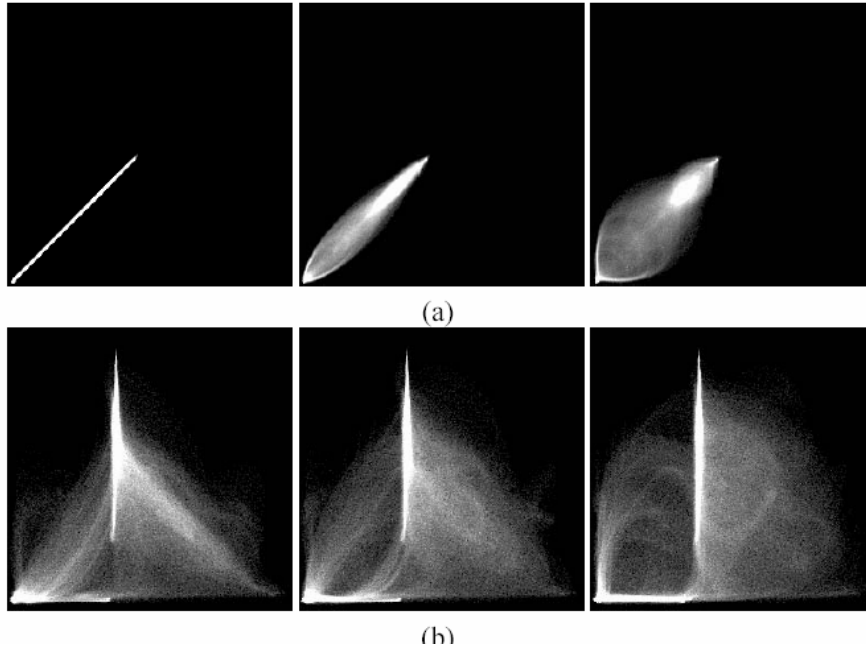
Figure 7: Examples of 2-D histograms from Hill *et al* [16]. In (a) both images are identical. In (b) a CT image and an MR image of the brain are used. Images on the left are correctly registered. In the center a translation of 2mm is applied and a translation of 5 mm is applied to the image on the right.

Using information theory this clustering/scattering effect is quantitatively described by measuring the *Entropy* and the *Mutual Information* of those joint histograms.

Entropy

Entropy H is a measure of the amount of information in a given signal. It was developed by Shannon the 1940s [17].

$$H = - \sum_{s} p(s) \log p(s)$$

where s is the set of possible symbols in a message, p(s) the probability of occurrence of that symbol. H will take a minimum value of 0, if the message has only one possible

symbol. The probability of that symbol would be 1 and the message will contain no information since we know for sure that this symbol is occurring. In this respect entropy can also be used as a measure of *uncertainty*. As the number of different messages that can be formed increases, the uncertainty about which one of them is received increases. When all symbols have equal probability, p(s) will be equal to 1/n where n is the number of different symbols and the entropy takes its maximum value:

$$\text{Max. } H = -\log\left(\frac{1}{n}\right) = \log(n)$$

The uncertainty is also at a maximum value because one is completely not sure which symbol will be received. If one of the symbols has a higher probability this uncertainty and therefore the entropy decreases. By viewing intensity value occurrences in different pixels/voxels in images as probabilities, the concept of entropy can be applied to images. In an image that is formed by only a single intensity value (an empty black image for example) entropy is equal to its minimum value 0. When all different possible intensities are uniformly distributed in an image the entropy will have its maximum value log (n) where n is the number of different intensities. As the distribution of intensities tends to favor a small number of intensities, thus giving them high probabilities, the entropy decreases. Based on this discussion, entropy can be used to measure the amount of clustering/scattering in the 2-D histograms presented earlier. As registration accuracy increases, bright areas in the histogram start to cluster and become less dispersed. This clustering indicates that a small number of intensities are having high probabilities. Therefore the increase in clustering is accompanied by decrease in entropy. The *Joint entropy* of images A and B is calculated from the intensities of images A and B using the following equation:

$$H(A,B) = - \sum_{i,j} p_{AB}(i,j) \log p_{AB}(i,j)$$

where i and j are the intensities or intensity partitions in the region of overlap between images A and B. The following figure presents an example showing the relationship between intensity histograms and the value of the joint entropy between two images:

(a) H=5.9   (b) H=10.5   (c) H=11.2   (d) H=12.1
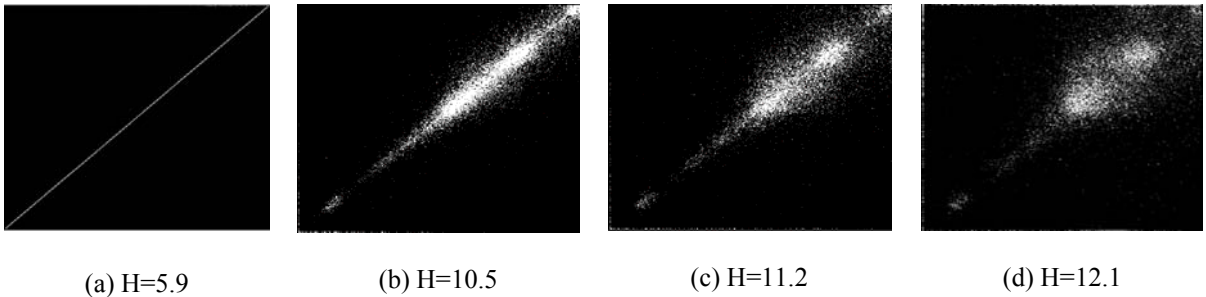
Figure 8: Join histogram and joint entropy are shown between (a) image A and itself. (b), (c) and (d) image A and image A rotated by 1, 3 and 10 degrees respectively.

The use of Entropy as a similarity measure is tempting, but the joint entropy calculation depends only on the region of overlap between the two images. This makes it very sensitive and dependent on this region of overlap. The optimization algorithm might, for example, minimize the joint entropy *H(A,B)* by maximizing the amount of air (which normally surrounds the images) in the region of overlap. The final position of the images in this case could be such that it would contain only air because it leads to a minimum in the joint entropy as shown in the next figure.
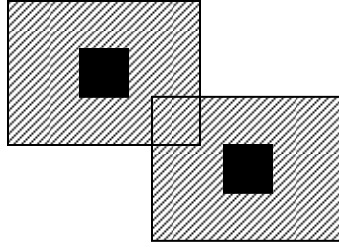
Figure 9: The two images are not registered but their joint entropy will be at a minimum.

This problem can be solved by considering the information in each of the two images along with the joint entropy.

Definition of Mutual Information

Mutual Information (MI) overcomes the problems of joint entropy by taking into account the entropy of the two images used to compute the joint entropy. The mutual information of two images A and B is given by

$$MI(A,B) = H(A) + H(B) - H(A,B)$$

where H(A) and H(B) are the entropies of images A and B computed over the region of overlap :

$$H(A) = - \sum_i p_A(i) \log p_A(i)$$

$$H(B) = - \sum_i p_B(i) \log p_B(i)$$

where $p_A(i)$ and $p_B(i)$ are the probabilities of occurrence for each intensity (or intensity partition) i in the area of overlap in images A and B, respectively. To maximize MI the optimization algorithm should look for the transformation that yields maximum entropies from image A and B and at the same time a minimum joint entropy value. The small joint

entropy value that results from aligning air surrounding images is now penalized by the small individual entropies of the two images forming that alignment. MI can also be expressed in terms of marginal distributions H(A|B) and H(B|A) as:

$$MI(A,B) = H(A) - H(A|B)$$

$$= H(B) - H(B|A)$$

where

$$H(A|B) = - \sum_{i,j} p_{AB}(i,j) \log p_{A|B=j}(i)$$

with $p_{A|B=b}(i)$ the conditional probability of A given B = b. Finally MI can also defined by:

$$MI(A,B) = - \sum_{i} \sum_{j} p_{AB}(i,j) \log \frac{p_{AB}(i,j)}{p_A(i) p_B(j)}$$

Like the joint entropy, MI does not make any assumptions about the nature of the relationship between the image intensities of corresponding voxels in two images. This can explain its importance as a similarity measure. Since its introduction MI was used successfully in a variety of different applications covering both intramodality [18] and intermodality [19] classes of registration problems. It was tested with different types of geometrical transformations including rigid [20], affine [21], and non rigid [22]. In this project mutual information was used successfully in intramodality affine and non rigid registration. The accuracy of Mutual Information compared to other registration techniques was tested as part of the Retrospective Registration Evaluation Project (RREP) at Vanderbilt University. Using fiducials implanted on sets of patient images and point based registration algorithms, transformations that registered the images were

obtained and were used as a gold standard to evaluate the different retrospective registration methods. Different registration techniques were then applied to the same set of images and the transformations generated by those methods were compared to the gold standard to determine their accuracy. Results showed considerable advantages for Mutual Information algorithms over the other retrospective voxel, surface and, points based measures [23]. Table 1 shows some of the results obtained by registering MR to CT images in this study. The median error in millimeters is reported for 14 different methods of registration over six types of MR images (T1, T2 and PD weighted before and after correction for geometrical distortion). Two different implementations of Mutual information appear in this table under the column CO [24] and HI [25] and their superiority is obvious.

Table 1: Median error in mm for different MR- CT registration methods as reported in [23], rect. indicates that the MR image was corrected for geometrical distortion before registration

MR Modality	Technique														N
	BA	CO	EL	HA	HE	HI‡	MAI	MAL	NO†	PE	RO1	RO2	RO3	RO4	
T1	1.6	1.5	1.6	3.4	1.4	1.2	5.1	4.3	3.3	2.7	4.2	5.2	5.7	5.4	7
PD	1.9	1.5	2.0	3.1	2.4	1.9	4.1	4.0	7.8	1.9	4.5	5.5	4.9	4.8	7
T2	2.5	1.5	1.6	4.2	4.7	1.5	3.9	5.0	3.9	2.5	4.5	4.5	5.4	4.7	7
T1 rect.	1.4*	0.7	0.9	3.3	1.0	0.7	4.9	5.4	3.4	2.2	5.9	5.9	6.3	5.9	6
PD rect.	1.7*	0.8	1.1	3.0	1.7	0.7	3.0	4.0	4.6	2.1	5.9	5.7	5.5*	5.5*	7
T2 rect.	2.1*	0.8	1.6	3.5	1.6	0.8	4.3	5.3	4.2	2.9	5.5	5.3	5.3	5.3	7

Other techniques listed in the above table are:

BA: chamfer matching (surface based method) [26], EL: intensity correlation [27], HA: head – had [8], HE: head – hat [28], MA: edgeness correlation (voxel based) [29], MAL: potential field (surface based) [30], NO: elastic warp (point based) [31], PE: head – hat, RO's: chamfer matching (surface based) [32].

Normalized Mutual Information

Another solution that addresses the problem with the joint entropy H(A,B) is a similarity measure developed by Studholme *el al.* [33]. In their method these authors use the ratio between the sum of the individual entropies and the joint entropy as a similarity measure.

$$NMI(A,B) = \frac{H(A) + H(B)}{H(A,B)}$$

They found that maximizing NMI is superior to MI if the region of overlap between the images is relatively small at the correct registration solution. NMI was shown to be less sensitive to the changes in the region of overlap than MI.

Interpolation Schemes

Registration by maximization of mutual information involves iteratively transforming an image B with respect to image A until an optimum value for their mutual information is found. Transformed points in B will generally not coincide with a grid point in A and interpolation will be required. Nearest Neighbor interpolation (NN) is the simplest possible interpolation method. The transformed point B will take the intensity value of the nearest voxel lying on the grid as shown in the following figure. Point x will take the intensity value of point N4 since the distance from x to N4 is the smallest when compared to distances to the other grid points.
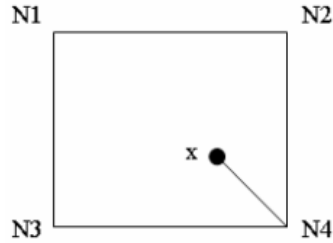
Figure 10: Using Nearest Neighbor interpolation
x takes the intensity of N4

Nearest Neighbor is insensitive to translations up to 1 voxel and the resulting similarity measure is going to be discontinuous. A better interpolation method is bilinear interpolation. In this method four weights are assigned to each surrounding voxel on the grid based on distances from the interpolated point to the grid. The interpolated voxel value is assigned based on those weights and the intensities of the surrounding voxels. For example with bilinear interpolation the intensity of point x in figure 6 will be computed as a weighted average of the intensities of the surrounding voxels. Because of the small distance between x and N4, the algorithm would make the effect of point N4 on the intensity computed for point x bigger.
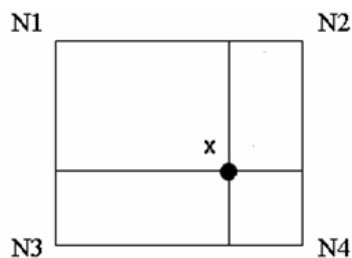
Figure 11: Illustration of bilinear interpolation

With this approach the interpolated function will be changing continuously. An implementation of this interpolation method is given in [34]. Extending the ideas of bilinear interpolation to three dimensions requires eight different weights and is usually

22

referred to in the literature as *Trilinear Interpolation*. Linear interpolation methods usually generate new intensity values that were not occurring in the original images. Collignon suggested the use of *Partial Volume* interpolation [35]. Using this method the weights are still calculated using the same procedure for trilinear interpolation; however instead of introducing new intensity values the algorithm updates several histogram entries. The histogram entry for each neighboring voxel is updated by its associated weight w_i. Both of those methods generate artifacts in the registration process. The behavior of mutual information function at grid aligning positions will be different from its behavior at positions where interpolation is needed. Ideally we are looking for a smooth monotonic function that reaches its peak value at the correct registration position. Trilinear and partial volume interpolation will introduce local minima and maxima that can deteriorate the robustness of registration. With trilinear interpolation the neighboring voxels are averaged which makes the resulting intensities more similar to each other. When calculating the joint entropy lots of those similar intensities will merge, thereby reducing joint entropy. The entropy of the interpolated image also decreases but by a smaller amount. These changes increase the value of mutual information, which results in local maxima at non grid aligning positions where interpolation is needed. On the other hand partial volume interpolation updates different histogram entries for each interpolated voxel. This reduces mutual information by increasing the joint entropy of the two images. The entropy of the interpolated image also increases but by a smaller amount. These changes decrease the value of mutual information. At grid aligning positions no interpolation is needed and local maxima appear. In [36] Maes *et al.* showed that traces of MI for Nearest Neighbor and Trilinear interpolation are noisy and show

many local maxima. Partial volume interpolation was shown to be less sensitive to this artifact. Pluim *et al.* showed that a small resampling of one of the images such that the voxel sizes of both images are no longer equal will largely remove interpolation artifacts in both trilinear and partial volume interpolation [37]. Resampling establishes a difference in voxel sizes and grid alignments therefore cannot occur.
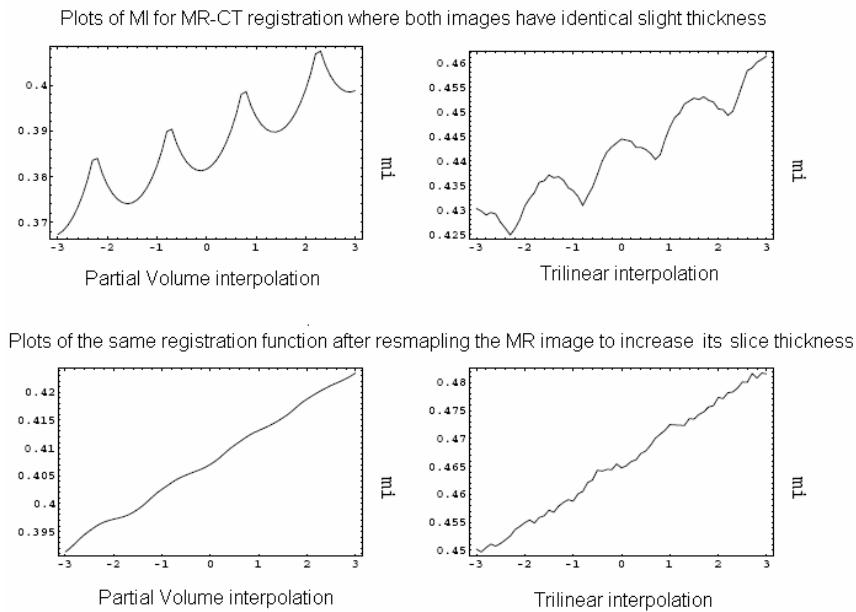
Figure 12: Effect of resampling on the registration function. The first row plots the registration function for MR-CT images with identical slight thickness. In the second row the MR image is resampled to have a slightly higher thickness [37]

Smoothing the images before registration was also found to help in reducing artifacts formed by trilinear interpolation since it reduces the dispersion of the joint histogram. However it has a very limited effect on partial volume artifacts [38].

Optimization methods

A successful search strategy is important for the registration algorithm to work. Different optimization algorithms have been used to maximize mutual information. Two of those, which are widely used, are the Simplex and Powell's direction set methods. Implementations of these methods in C are given in [34]. Since Powell's method finds the optimum by going through a sequence of one-dimensional optimizations, the order in which these optimizations are carried out is important. In [39] Maes *et al.* suggest starting the search with the in-plane degrees of freedom before optimizing the out of plane degrees of freedom. As the algorithm progresses it may introduce other optimization directions and change the order in which the parameters are considered. An important property of the Simplex and Powell's method is that they do not need information about the gradient of the function being optimized. Several authors tested developing analytical expressions for the gradient of MI and using them to introduce gradient base search strategies to this registration method [40-42]. In many applications a large number of degrees of freedom have to be optimized. The affine registration problem, for example, involves 12 parameters. Combining this with the large number of voxels in a high resolution image makes the registration problem computationally expensive. It has been shown that in those cases one can downsample the high resolution image without affecting the robustness of the algorithm. Important speed-ups can be achieved by a multiresolution optimization strategy. The optimization algorithm starts it's iterations with a downsampled version of the images until it converges to a local minimum. After

that results of registration at this level are used as a starting point in optimizing the images at their original resolution. Depending on the size and resolution of the images, the algorithm can be implemented to work through different resolution levels achieving speedups without reducing the precision or the robustness of the algorithm [43].

Number of bins

In calculating the joint entropy it is important to set an appropriate number for the intensity partitions or 'bins' to be used:

$$H(A,B) = - \sum_{i,j} p_{AB}(i,j) \log p_{AB}(i,j)$$

Usually each value on i or j represents an intensity partition. Those bins have to be set to values appropriate to the range of intensities in images A and B. For example, with MR and CT images having intensities ranging from 0 to 4096 (12 bits), the use of 32 to 256 bins is common.

CHAPTER IV

THE ADAPTIVE BASES ALGORITHM

Introduction

In many applications rigid and affine registration are not sufficient to describe the

relationship between the two images to be registered. With inter-subject registration for

example non rigid registration is needed to accommodate the variability in anatomic

structures across patients. Such applications will demand the use of a higher number of

degrees of freedom. A number of methods have been proposed to solve this problem. For

instance, An elastic approach is used by [44]. A model based on finite element methods

has been proposed by Edwards *et al* [45]. Christensen developed algorithms based on

viscous fluids [46]. Following the success that mutual information based methods had

with rigid registration, different methods appeared using it as a similarity measure driving

the deformations. In intensity-based, non-rigid registration, a source image B(x) is

deformed to best match a target image A(x). A similarity function is used to measure the

disparity between the two images and to guide the search algorithm to the optimal

solution. Mathematically, in three dimensions this can be described as:

$$\arg\max_{x'} F(\ B(\mathbf{x'}),\ A(\mathbf{x}),\ \mathbf{x'}\ )$$

in which

$$\mathbf{x'} = \mathbf{x} + \mathbf{v(x)}$$

and

$$\mathbf{v(x)} = (\ v_x(\mathbf{x})\ ,\ v_y(\mathbf{x})\ ,\ v_z(\mathbf{x})\ )$$

The function F is the similarity measure to be optimized. Various methods have been proposed to compute the value of v(x). Mayer uses a technique based on thin plate splines. After starting with an affine registration between the two images to be registered, the algorithm optimizes the location of the control points (initially location is selected by the user) using mutual information as a similarity measure. Rueckert and Studholme use a similar approach using B-splines [47-48]. Different authors suggested the use of linear combination of B-splines placed on a regular grid [47-50]. In many applications structures of interest are on the order of millimeters. In this project for example the facial recess is approximately 2x2x5 mm in dimension. This would require the regular grid algorithms to place a basis function every couple of millimeters. Finding the optimal solution in such cases is computationally expensive and difficult due to the effects of local minima or maxima. Rohde *et al.* developed the Adaptive Bases Algorithm (ABA) that overcomes this problem by using an adaptive grid of basis functions [51]. The rest of this chapter goes over this algorithm in detail.

Bases Functions Used

The deformation field is modeled by a linear combination of basis function spaced irregularly over the image domain. The deformation vector at point **x** is calculated using:

$$\mathbf{v}(\mathbf{x}) = \sum_{i=1}^{N} \mathbf{c}_i \Phi\left(\frac{\|\mathbf{x} - \mathbf{x}_i\|}{s}\right)$$

in which

$$\mathbf{c}_i = (c_i^x, c_i^y, c_i^z)$$

$$\Phi(r) = (1 - r)_+^4 + (3r^3 + 12r^2 + 16r + 4) \text{ for } r \geq 0$$

where vector $\mathbf{x}_i$ has the coordinates of control point i, N is the total number of control points used, $(1 − r)_+ = \max(1 − r, 0)$, s is a predetermined scale factor, and $\|\mathbf{x}\|$ is the Euclidian norm for vector $\mathbf{x}$. The next figure shows a plot of the basis function $\Phi(r)$ in one and two dimensions.
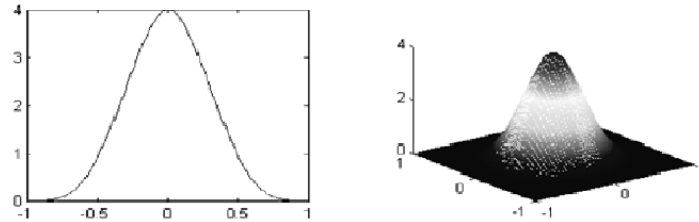
Figure 13: Plots of the basis function $\Phi(r)$ in one and two dimensions

The algorithm approaches the final deformation field by going through a series of resolutions and scales. Images to be registered are first downsampled to the lowest resolution level. The algorithm starts with few control points with a large scale value. As the algorithm approaches finer resolutions, the scale of the basis function is lowered and a higher number of control points are used. The algorithm computes the deformation field for each level (a particular combination of scale and resolution) and the total deformation field is computed by the following equation:

$$\mathbf{v}(\mathbf{x}) = \mathbf{v}_1(\mathbf{x}) + \mathbf{v}_2(\mathbf{x}) + \ldots + \mathbf{v}_M(\mathbf{x})$$

where M is the total number of levels.

Identifying Regions of Misregistration

For each level the algorithm places a regular grid of basis function and models the deformation field as:

$$d(\mathbf{x}) = \mathbf{x} + \sum_{k=1}^{m-1} \mathbf{v}_k(\mathbf{x}) + \sum_{i=1}^{N} \mathbf{c}_i \Phi(\frac{\|\mathbf{x} - \mathbf{x}_i\|}{s})$$

where $\mathbf{x}_i$ is the position of the basis function, s is the scale factor and $\sum_{k=1}^{m-1} \mathbf{v}_k(\mathbf{x})$ is the sum

of the deformation fields obtained up to level m-1. The gradient of the normalized mutual

information NMI(A($\mathbf{x}$), B(d($\mathbf{x}$))) is computed with respect to the coefficients $\mathbf{c}_i$ and is used

to determine areas that are misregistered between images A($\mathbf{x}$) and B($\mathbf{x} + \sum_{k=1}^{m-1} \mathbf{v}_k(\mathbf{x})$).

Areas where the magnitude of the gradient is larger than a selected threshold are

considered areas of mismatch. Areas adjacent to each other are removed from the list of

regions of mismatch to prevent overlap (This is discussed further in the next section).

Local Optimization

The coefficients $\mathbf{c}_i$ for each region of mismatch are optimized independently. Eight

locations surrounding the centre of each region are picked as centers for basis functions.

Those locations form a cube around each area of mismatch found in the previous section.

A steepest gradient descent algorithm combined with the quadratic interpolation four-

point bracketing update method of line minimization is then applied to the 24 coefficients

of the cube of basis functions under the following cost function:

$$F(A(\mathbf{x}),B(\mathbf{x'}),\mathbf{x'})$$

in which

$$\mathbf{x'} = \mathbf{x} + \sum_{k=1}^{m-1} \mathbf{v}_k(\mathbf{x}) + \mathbf{v}_m(\mathbf{x})$$

where

$$\mathbf{v}_m(\mathbf{x}) = \sum_{j=1}^{J}\sum_{i=1}^{8}\mathbf{c}_i^{j}\Phi\left(\frac{\left\|\mathbf{x}-\mathbf{x}_i^{j}\right\|}{s}\right)$$

J represents the number of regions of mismatch while $\mathbf{x}_i^{j}$ are the coordinates of the cube mentioned earlier. Because the regions were picked so that they don't overlap this optimization procedure can be done on each region independently. To ensure consistent topological deformations constraints have to be applied to the values of $\mathbf{c}_i$. The algorithm includes a parameter λ that sets an upper limit on the absolute differences allowed between the values of the coefficients $\mathbf{c}_i$ in adjacent basis functions. Smaller values of λ yield smoother deformation fields. More detailed information about the ABA algorithm can be found in [51-52].

CHAPTER V

THE PROPOSED ALGORITHM

Initial Preparation

The path of the drill can be determined by identifying a voxel in the facial recess and another voxel in the area on the basal turn of the cochlea called the "*target window*". The proposed algorithm uses labeled regions in an image referred to as the *atlas* to identify the same regions in the patient images. The atlas was picked by an experienced surgeon so that it represents a typical normal ear. The *facial recess* and the *target window* were labeled in each atlas manually. Two atlases were used in this study, one for each ear since none of the available images were such that both ears appear good enough to be used as an atlas. Two binary volumes having the same dimensions as the atlas images were created for each atlas. In each volume voxels within the region of interest were given an intensity of 255 and the intensities of all other voxels were set to zero. Those volumes mask the regions of interest in the atlases and will be used to identify their correspondents in the patient images. The following figures show the outlines that were used to generate the masks.

Axial view

Sagittal view

Coronal view
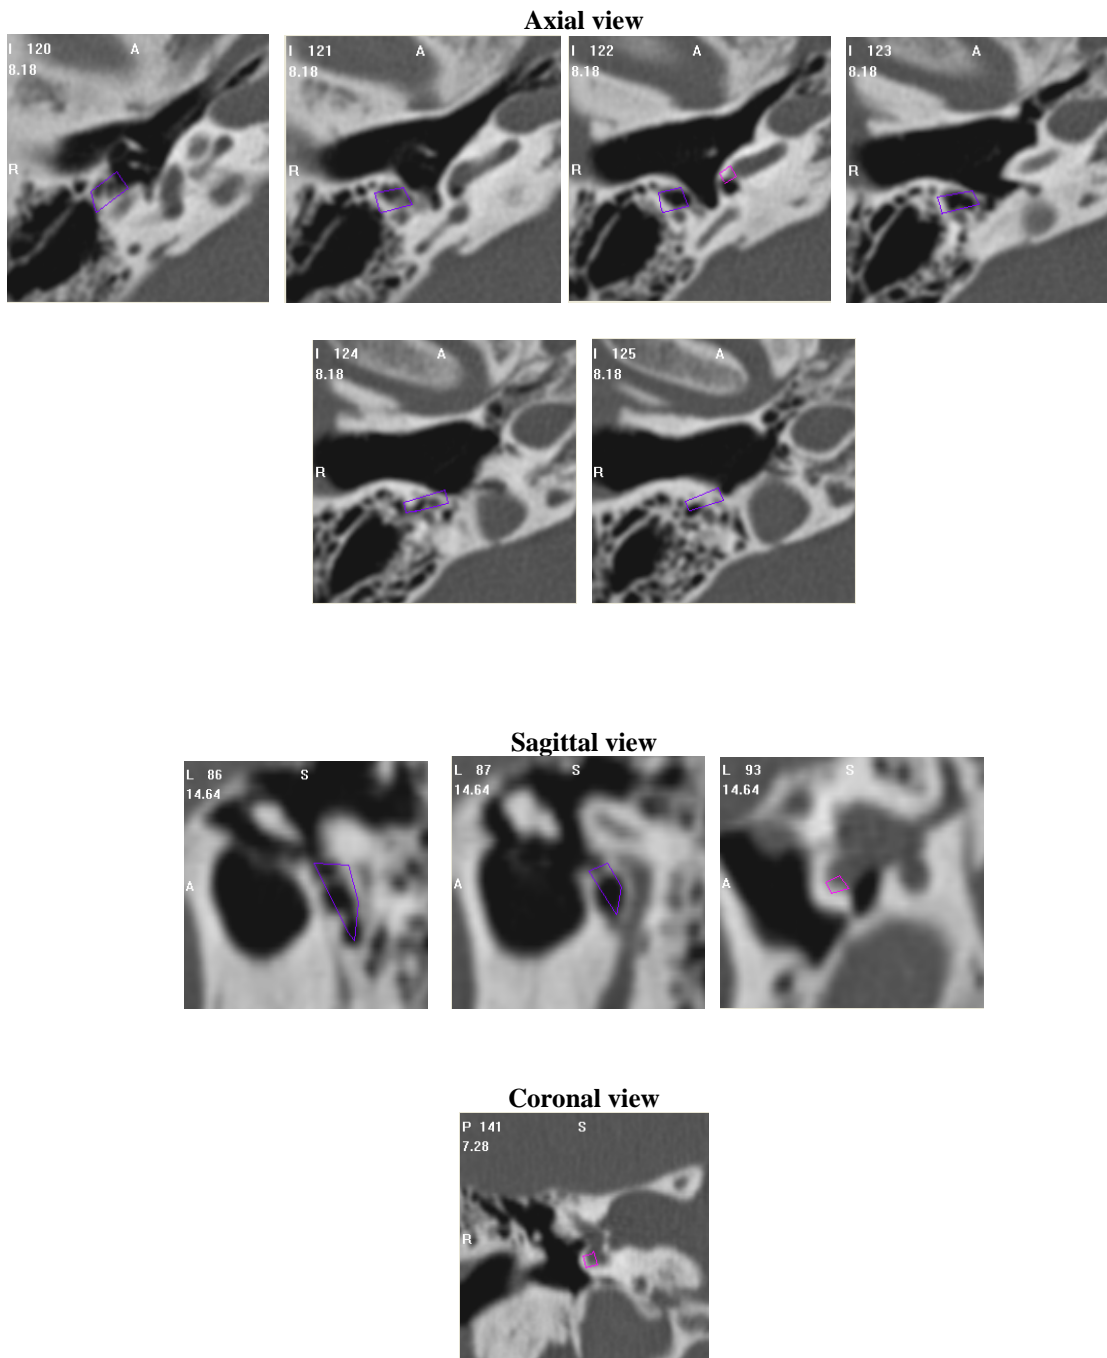
Figure 14: Outlines used to create masks for the right ear's atlas. The purple outlines label the facial recess. The red outlines are labeling the target window
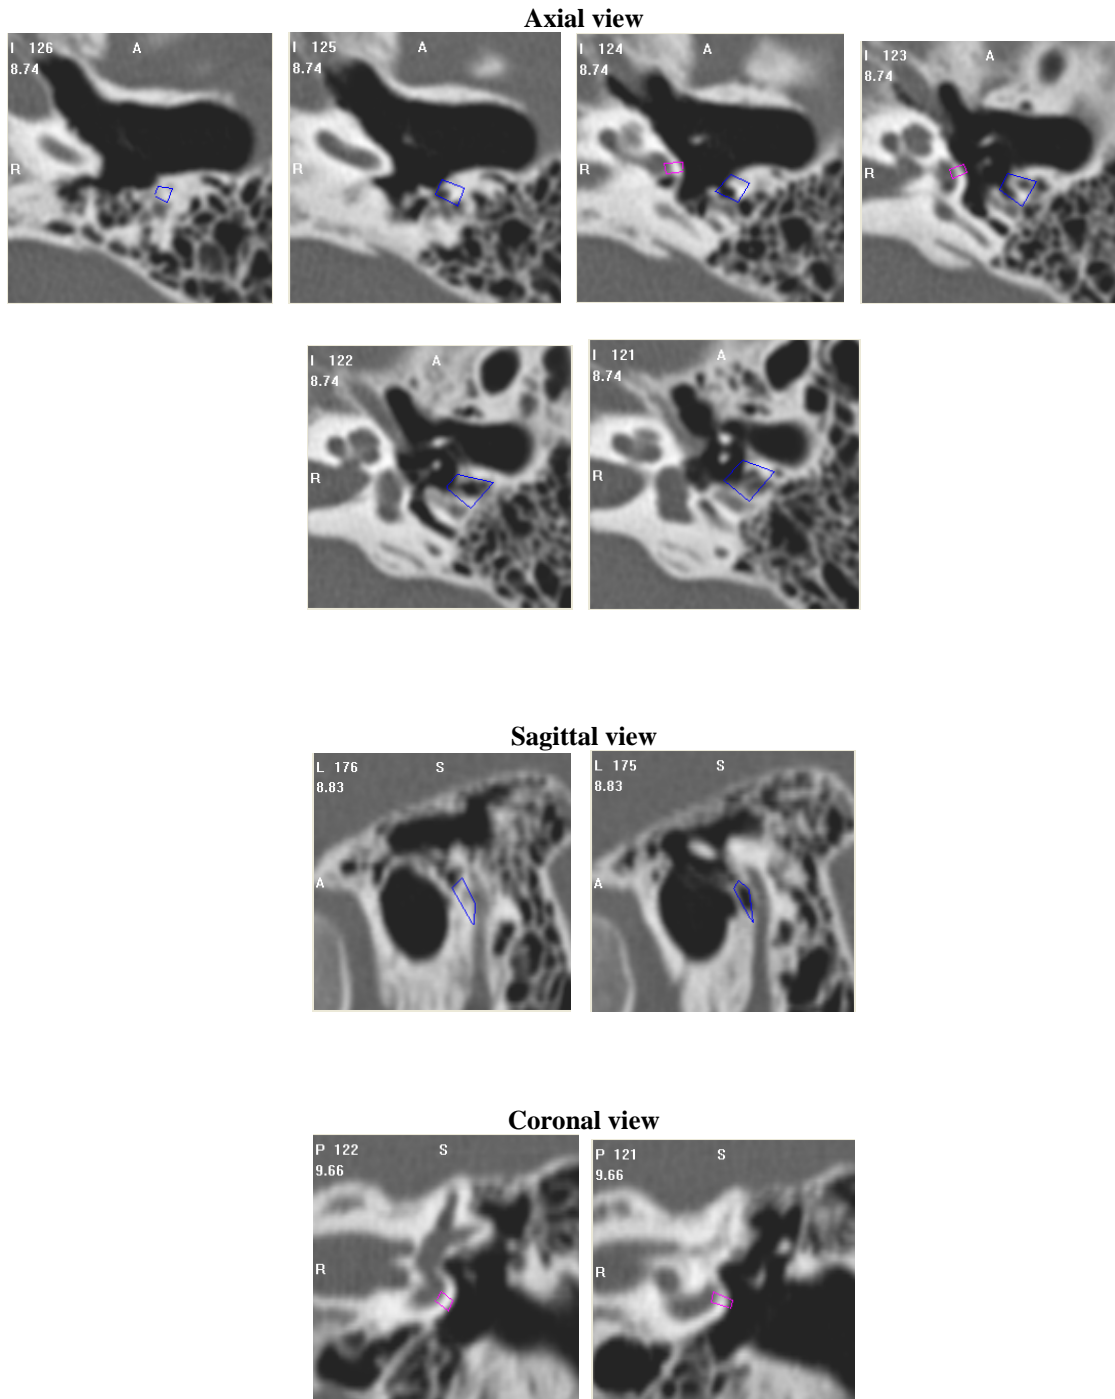
Figure 15: Outlines used to create masks for the left ear's atlas. The purple outlines label the facial recess. The red outlines are labeling the target window

Because we are only interested in the ears and the non rigid registration algorithm is computationally expensive it is necessary to crop the images, picking only the region

containing and surrounding the ear of interest. However to automate this step the patient image (unknown) has to be first rigidly registered to the atlas image. This enables the algorithm to automatically pick the region of interest in the patient image based on the alignment of the ears between the atlas and the patient. Obviously this region of interest has to be picked in each atlas before running the algorithm. Two points were manually picked in each atlas to define a 3D box containing the region of interest. The coordinates of the points were picked so that the distance between the edge of the picked region and the closest structure of interest (target window or facial recess) contains at least thirty voxels (6 – 12mm). Experiments showed that this surrounding area is necessary to obtain good deformations by the ABA algorithm. This area does also account for voxels lost due to rotation and skewing as explained in the next section. The use of larger surrounding areas did not improve the results. With the two atlases, four masks, and two sets of input points the algorithm is now ready to start.

<div align="center">Affine And Non Rigid Registration</div>

The algorithm starts by performing intensity based affine registration with 12 degrees of freedom between the atlas and the patient image. Mutual information is used as a similarity measure and Partial volume is used as the interpolation method. Images are smoothed before the registration starts (3×3 box filtering), and the Powell's method is used as an optimization algorithm. Images used in this study were downsampled by a factor of 2 in the x and y directions since they had a large number of voxels (768 x 768). A multiresolution approach using two levels was then applied to the images. The order of optimization in Powell's method starts with the x and y directions before moving into

rotation around the z axes. Rotations around the other axes are optimized later followed by optimizing scaling and skewing in the three directions. The resulting transformation matrix is then saved with the *registered atlas image* that results from applying the transformation matrix to the atlas image. After that the algorithm cuts two 3-D boxes representing the region of interest in the registered atlas image and the patient image. The coordinates of the 3D boxes are calculated by applying the transformation matrix to the input points defined in the algorithm. Since the two images are aligned at this point those coordinates should correspond to approximately the same structures in both images. Because of skewing, the initial box defined before running the algorithm can transform into a 3-D parallelogram. This will change the originally defined area by adding or subtracting few voxels. This change in the covered area can also occur because of rotation. The following figures illustrate the effects of rotation and skewing to the area covered by the initial box.
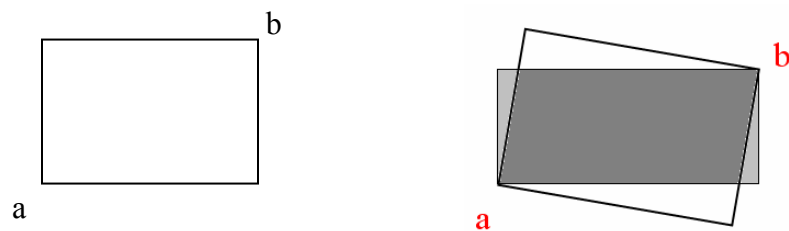

Figure 16: Original rectangle is shown on the left. On the right the effect of rotation on the box defined by points *a* and *b*. The new box is presented by the gray area (light and dark). White areas are lost and light gray areas are added to the area covered by the original rectangle.

Figure 17: Original rectangle is shown on the left. On the right effect of skewing on the box defined by points *a* and *b*. The new box is presented by the gray area. White areas are lost.

Since the degrees of rotation or skewing required are usually small and the initial box contained a safe region surrounding the structures of interest, the effect of this change can be ignored. However if the images to be registered required high degrees of rotation or skewing then increasing the surrounding area in the initial box should solve the problem of lost voxels. The two extracted boxes are then saved as two images representing the *atlas ear* and the *patient ear*. The Adaptive Bases Algorithm (ABA) is then applied on the two images. The ear images are registered by going through a series of 15 levels, which differ by the number of control points used in each one of them. The algorithm starts with two control points in each slice direction (x, y, and z) and goes through the different levels until it terminates at the last level where it uses a grid of 50 control points in each direction. The Jacobian threshold was set at a value of 0.3. The algorithm is used to generate a *deformation field* which is a map of 3D vectors representing the direction and magnitude of the optimal deformation that was found for each voxel in the atlas ear image. Now the overall deformation field is computed by combining the transformation matrix generated by the affine registration algorithm with the deformation field from the ABA algorithm. Since the deformation field was calculated using the smaller ear images

appropriate conversions need to be preformed. The following figure illustrates the steps

discussed so far.
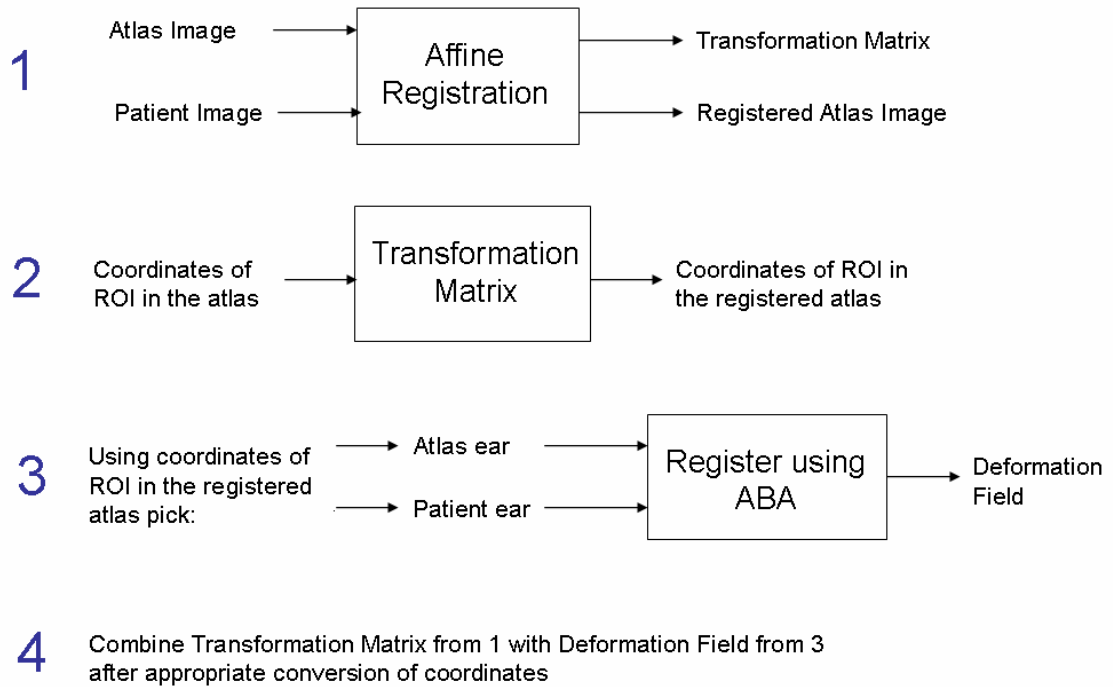
Figure 18: Steps used to find the overall transformation field that registers the atlas ear to patient ears

Deforming Masks And Finding Centroids

The algorithm now applies the deformation field obtained at the previous step to the

facial recess and the target window masks mentioned earlier in this chapter. The

deformed masks are computed by creating an empty image that has the same dimensions

as the patient image. The overall deformation field is then used to find the intensities at

each voxel in the deformed masks. Grabbing the intensities from the original masks will

most likely involve the use of interpolation and therefore the resulting deformed masks

are not going to be binary images. The intensity centroid is then computed for each

output mask and the line connecting the calculated centroids is drawn using parametric

38

equations. Ideally this line should be an appropriate drilling path that can be used in accessing the cochlea.

CHAPTER VI

RESULTS

Volumes Used

Six different CT scans were available for this study. The scans covered the ears and

neighboring areas in the head, and sometimes in the neck. A qualified surgeon selected

two scans, one for each ear[1], to be used as atlases. Voxel sizes were taken into

consideration while selecting the scans. The two atlases had a high resolution compared

with most of the other images used as unknowns. The following table lists the dimensions

and voxel sizes for the six scans on which the algorithm was tested.

Table 2: Volumes used in this study

| Scan	Dimensions	Voxel Sizes in mm	
A	768 x 768 x 145	[0.218750 0.218750 0.4]	left ear used as atlas
B	768 x 768 x 147	[0.230469 0.230469 0.4]	right ear used as atlas
C	768 x 768 x 333	[0.358073 0.358073 0.4]	right ear was not used
D	768 x 768 x 300	[0.384115 0.384115 0.4]	
E	768 x 768 x 331	[0.339844 0.339844 0.4]	
F	768 x 768 x 84	[0.365890 0.365890 1.0]	

---

[1] A scan where the two ears are could be used as atlases could not be found

The right ear in scan "A" and the left ear in scan "B" were used as unknowns. The right ear in scan "C" was not used because it went through a surgery that obliterated its facial recess.

Performance Of Registration

The registration accuracy was analyzed by comparing the deformed atlas ear with the given unknown ear. In all cases the algorithm successfully deformed the atlas ear generating new image ears that are very similar to the unknown ears. An Example of the results generated are shown in the following figure. In this figure the registered atlas ear is in the top row while the unknown ear is in the second row. The rows are showing the same region in the inner ear in three different views: sagittal, axial, and coronal. The red circle indicates the location of the target window, while the green square is showing the location of the facial recess. In the top rows the circles were picked manually and their coordinates were used to draw the circles in the unknown ears. It is clear that the algorithm was able to align the structures of interest. In the coronal views the cochlea can easily be noticed by its spiral shape in both the registered atlas and unknown ears.
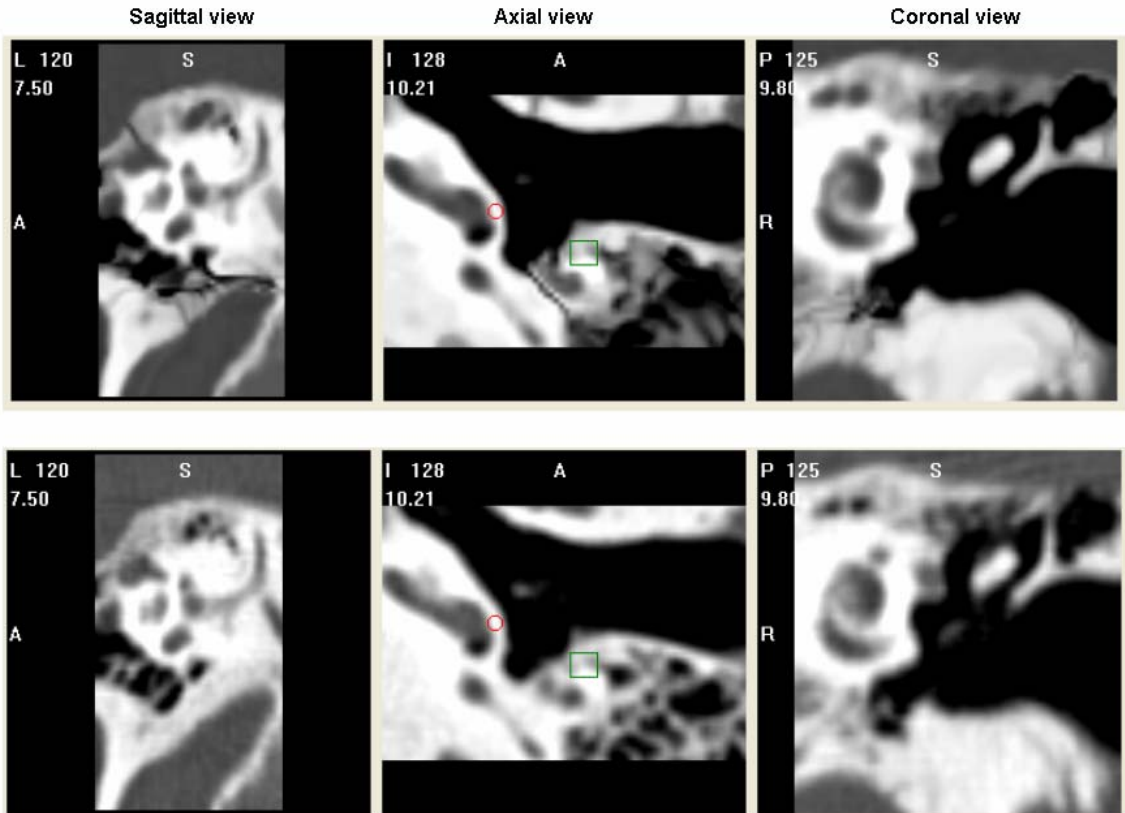
Figure 19: A slice of a deformed atlas ear is shown using three views on the top row, while the unknown ear is on the bottom row. The square is indicating the location of the facial recess. The circle is showing the location of the target window.

Suitability Of Paths

The 3-D paths generated by the algorithm were plotted on the scans and evaluated by a qualified surgeon. The algorithm was found to be successful in finding a safe path in all of the nine ears. All the paths accessed the basal turn of the cochlea by going through the facial recess, thus avoiding injuring any important structures. The next figure shows an example of such a path using the three different views tracking its progress to the cochlea on different slices.

Axial view

1 **2** **3**

Coronal view

1 **2** **3** **4**

5 **6** **7** **8**

Sagittal view

1 **2** **3** **4**
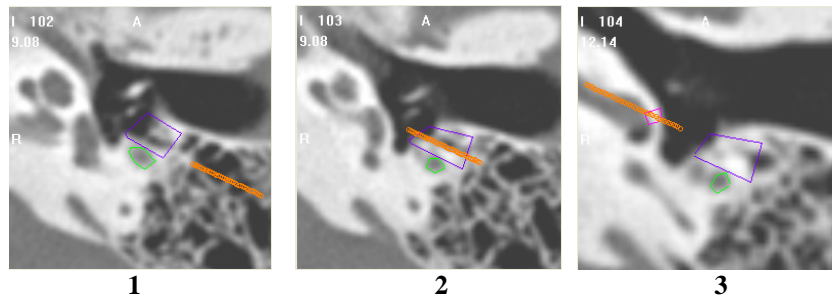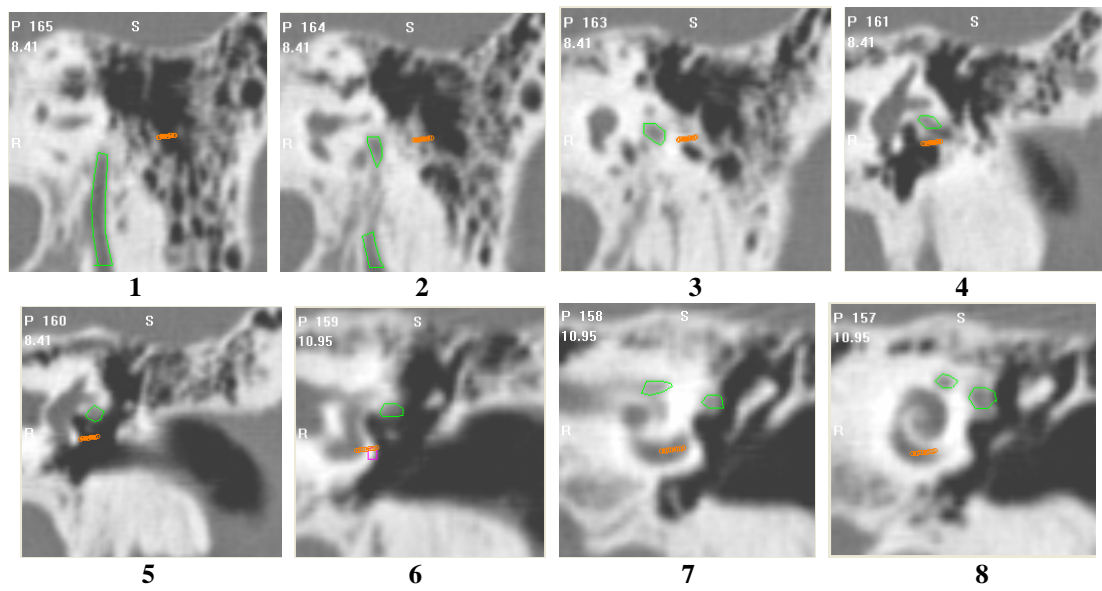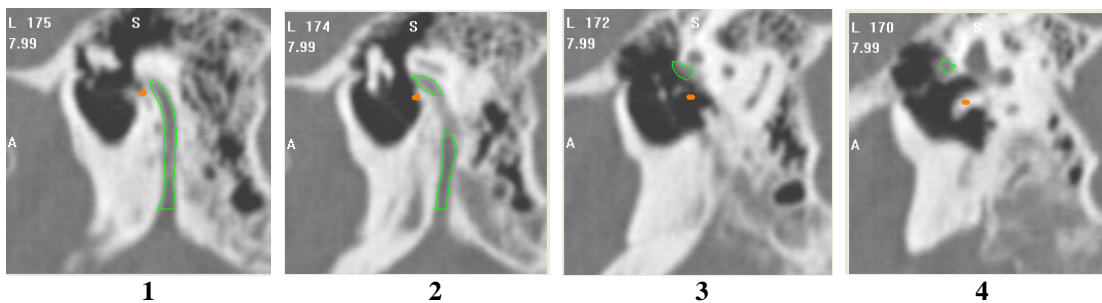
Figure 20: The path found by the algorithm for the left ear in scan 'C'. It reaches the cochlea by going through the target window outlined in red through the facial recess outlined in purple. The facial nerve is outlined in green and the path succeeds in avoiding it

43

Accuracy In Locating Centroids

The facial recess and the target window were manually labeled on the unknowns by the same surgeon who drew the outlines on the atlas images. Distances between the centroids of the manually labeled structures and the centroids calculated by the algorithm were computed and used as an error measurement. The results are shown in the following tables.

Table 3: Centroid error on the left ears used in this study

Scan	Facial recess error in mm	Target error in mm
B	1.343	1.205
C	0.599	0.495
D	2.379	0.233
E	1.343	0.683
F	1.618	0.416
Average error	1.456	0.602

Table 4: Centroid error on the right ears used in this study

Scan	Facial recess error in mm	Target error in mm
A	1.259	0.611
D	0.965	0.972
E	0.932	1.339
F	1.365	1.323
Average Error	1.130	1.062

CHAPTER VII

CONCLUSIONS AND VIEW OF FUTURE WORK

An algorithm to plan an appropriate drilling path to be used in cochlear implantation surgery was developed and tested in this thesis. The path was found by locating the centroids of two structures through which it must pass. The algorithm demonstrated success of atlas based segmentation methods in localizing the basal turn of the cochlea and the facial recess. Improvements can be made to the running time of the algorithm. Currently the algorithm takes between 1-3 hours depending on the size of the patient image. Cropping the head's and atlas images so that they only have the side of the head where the ear of interest is can achieve speedups by a factor close to two. It would be interesting to test using only the area of interest around the inner ear as an atlas image. Further studies can be made into the relationship between the number of control points used by the ABA algorithm and the size of the area of interest picked when forming ear images with the goal of finding the optimum of those two parameters with respect to speed savings and accuracy.

REFERENCES

[1]   Courtesy of Dr. Robert F. Labadie

[2]   House WF. Surgical considerations in cochlear implantation Ann Otol Rhinol Laryngol Supp 1982;91(2pt3):15-20

[3]   Kroenenberg J, Migriov L, Dagan T. Suprameatal approach: new surgical approach for cochlear implantation. J Laryngol Otol 2001;115:283-5

[4]   Kronenberg J, Baumgartner W, Migirov L, Dagan T, Hildesheimer M. The suprameatal approach: an atleranative approach to cochlear implantation. Otol Neurotol 2004;25:41-5

[5]   Sonka and Fitzpatrick, "The Handbook of Medical Imaging, Volume 2: Medical Image Processing and Analysis", ISBN 0-8194-3622-4, SPIE , Belligham, WA, 2000,

[6]   Calvin R. Maurer, Jr., Member, J. Michael Fitzpatrick, Matthew Y. Wang, Robert L. Galloway, Jr., Robert J. Maciunas, and George S. Allen, "Registration of Head Volume Images Using Implantable Fiducial Markers", *IEEE Transactions on Medical Imaging*, Vol. 16, pp. 447-461, 1997.

[7]   B. F. Green, "The orthogonal approximation of an oblique structure in factor analysis," Psychometrika, vol. 17, pp. 429–440, 1952.

[8]   C. A. Pelizzari, G. T. Y. Chen, D. R. Spelbring, R. R. Weichselbaum, and C.-T. Chen, "Accurate three-dimensional registration of CT, PET, and/or MR images of the brain," J. Comput. Assist. Tomogr., vol. 13, pp. 20–26, 1989.

[9]   P. J. Besl and N. D. McKay, "A method for registration of 3-D shapes," IEEE Trans.Pattern Anal. Mach. Intell., vol. 14, pp. 239–256, 1992.

[10]   P. A. Viola, "Alignment by maximization of mutual information". Ph.D. thesis, Massachusetts Institute of Technology, 1995

[11]   J.V. Hajnal, N. Saeed, A. Oatridge, E.J. Williams, I.R. Young, and G.M. Bydder, Detection of subtle brain changes using subvoxel registration and subtraction of serial MR images., J. Comput. Assist. Tomogr. , vol 19, pp. 677-691, 1995

[12]   R.P. Woods, S. R. Cherry, and J.C. Mazziotta, Rapid automated algorithm for aligning and reslicing PET images, J. Comput. Assist. Tomogr., vol 16, pp. 620-633, 1992

[13] R.P. Woods, J.C. Maziotta, and S.R. Cherry, MRI-PET registration with Bautomated algorithm, J. Comput. Assist. Tomogr., vol 22, pp. 139-152, 1998

[14] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *IEEE Trans. Med. Imag.*, vol. 16, pp. 187–198, Apr. 1997

[15] W. M. Wells III, P. Viola, H. Atsumi, S. Nakajima, and R. Kikinis, "Multimodal volume registration by maximization of mutual information," *Med. Image Anal.*, vol. 1, no. 1, pp. 35–51, Mar. 1996.

[16] D. L. G. Hill and D. J. Hawkes, "Voxel similarity measures for automated image registration", Visualization in Biomedical Computing 1994, vol. Proc. SPIE 2359, pp. 205-216, 1994

[17] C. E. Shannon, Communication in the presence of noise, Proc. IRE, vol. 37, pp. 10-21, 1949. Reprinted in Proc. IEEE 86:447-457, 1998

[18] M. Holden, D. L. G. Hill, E. R. E. Denton, J. M. Jarosz, T. C. S. Cox, T. Rohlfing, J. Goodey, and D. J. Hawkes, "Voxel similarity measures for 3-D serial MR brain image registration," IEEE Trans. Med. Imag., vol. 19, pp. 94–102, July 2000

[19] M. Debois, R. Oyen, F. Maes, G. Verswijvel, G. Gatti, H. Bosmans, M. Feron, E. Bellon, G. Kutcher, H. V. Poppel, and L. Vanuytsel, "The contribution of magnetic resonance imaging to the three-dimensional treatment planning of localized prostate cancer," Int. J. Radiat. Oncol. Biol. Phys, vol. 45, no. 4, pp. 857–865, Nov. 1999

[20] T.Rohlfing, J. B. West, J. Beier, T. Liebig, C. A. Taschner, and U.-W. Thomale, "Registration of functional and anatomical MRI: accuracy assessment and application in navigated neurosurgery," Computer Aided Surgery, vol. 5, no. 6, pp. 414–425, 2000.

[21] M. Jenkinson and S. Smith, "A global optimisation method for robust affine registration of brain images," Medical Image Analysis, vol. 5, no. 2, pp. 143–156, 2001.

[22] C. R. Meyer, J. L. Boes, B. Kim, P. H. Bland, K. R. Zasadny, P. V. Kison, K. Koral, K. A. Frey, and R. L. Wahl, "Demonstration of accuracy and clinical versatility of mutual information for automatic multimodality image fusion using affine and thin-plate spline warped geometric deformations," Medical Image Analysis, vol. 1, no. 3, pp. 195–206, 1997.

[23] J.West, J. M. Fitzpatrick, M. Y. Wang, B. M. Dawant, C. R. Maurer Jr., R. M. Kessler, R. J. Maciunas, C. Barillot, D. Lemoine, A. Collignon, F. Maes, P. Suetens, D.Vandermeulen, P. A. van den Elsen, S. Napel, T. S. Sumanaweera, B. Harkness, P. F. Hemler, D. L. G. Hill, D. J. Hawkes, C. Studholme, J. B. A. Maintz, M. A. Viergever, G. Malandain, X. Pennec, M. E. Noz, G. Q. Maguire

Jr., M. Pollack, C. A. Pellizari, R. A. Robb, D. Hanson, and R. P.Woods, "Comparison and evaluation of retrospective intermodality brain image registration techniques," J. Comput. Assist. Tomogr., vol. 21, pp. 544–566, 1997.

[24]	A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal , "Automated multi-modality image registration based on information theory", in Information Processing in Medical Imaging 1995, Y. Bizais, C. Barillot, and R. Di Paola, Eds., pp. 263-274. Kluwer Academic, Dordrecht, The Netherlands, 1995.

[25]	C. Studholme, D. L. G. Hill and D. J. Hawkes, "Automated 3D registration of MR and CT images of the head", Med. Image Anal., Vol. 1, 1996

[26]	D. Lemoine, D. Liegeard, E. Lussot, and C. Barillot, "Multimodal registration system for the fuson of MRI, CT, MEG, and 3D or steretactic angiographic data", Medical Imaging 1994: Image Caputre, Formatting and Display, vol. Proc. SPIE 2164, pp. 46-56, 1994 .

[27]	P. A. van den Elsen, E.-J. D. Pol, T. S. Sumanaweera, P. F. Hamler, S. Napel, and J. R. Adler, "grey value correlation techniques used for automatic matching of CT and MR brain and spine images", Visualization in Biomedical Computing 1994, vol. Proc. SPIE 2359, pp. 227-237, 1994.

[28]	P. F. Hemler, T. S. Sumanaweera, P. A. van den Elsen, S. Napel and J. R. Adler, "A versatile system for multimodality image fusion", J. Image Guid. Surg., vol. 1, pp. 35-45, 1995

[29]	J. B. A. Maintz, P. A. van den Elsen, and M. A. Viergever, "Comparison of feature-based matching of CT and MR brain images", in Computer vision, Virtual Reality, and Robotics in Medicine 1995, N. Ayache, Ed., pp. 219-228. Springer-Verlag, Berlin, 1995

[30]	G. Malandain, S. Fernandez-Vidal, and J. M. Rocchisani, "Rigid registration of 3-D objects by motion analysis", in Proc. 12th Int. Conf. Pattern Recognition, 1994, pp.579-581.

[31]	G. Q. Maguire, Jr., M. E. Noz, H. Rusinek, J. Jaeger, E. L. Kramer, J. J. Sanger, and G. Smith, "Graphics applied to medical image registraion", IEEE Comput. Graph. Appl., vol 11, pp. 20-29, March 1991.

[32]	H. Jiang, K. S. Holton, and R. A. Robb, "Image registration of multimodality 3-D medical images by chamfer matching", Biomedical Image Proecssing and Three-Dimensional Microscopy 1992, vol. Proc. SPIE 1660, pp. 356-366, 1992.

[33] C. Studholme, D. L. G. Hill and D. J. Hawkes, "An overlap invariant entropy measure of 3-D medical image alignment," *Pattern* Recognit., vol. 32, no. 1, pp. 71–86, 1999.

[34] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, Numerical Recipes in C, Second ed. Cambridge, U.K.: Cambridge Univ. Press, 1992.

[35] A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal, "Automated multimodality medical image registration using information theory," in Proc. 14th Int. Conf. Information Processing in Medical Imaging (IPMI'95), vol. 3, Computational Imaging and Vision, Y. Bizais, C. Barillot, and R. Di Paola, Eds., Ile de Berder, France, June 1995, pp. 263–274.

[36] F. Maes, "Segmentation and registration of multimodal medical images: From theory, implementation and validation to a useful tool in clinical practice ," Ph.D. dissertation, Dept. Elect. Eng. (ESAT/PSI), KU Leuven, Leuven, Belgium, 1998.

[37] J. P. W. Pluim, "Interpolation artefacts in mutual information-based image registration," *Comput. Vis. Image Understand.*, vol. 77, no. 2, pp. 211–232, 2000.

[38] Rui Li, "Automatic Placement of regions of interest in medical images using image registration", Master's thesis, Vanderbilt University 2001

[39] F. Maes, D. Vandermeulen, and P. Suetens, "Comparative evaluation of multiresolution optimization strategies for multimodality image registration by maximization of mutual information," *Med. Image Anal.*, vol. 3, no. 4, pp. 373–386, 1999.

[40] D. R. Holmes III, J. J. Camp, and R. A. Robb, "Evaluation of search strategies and cost functions in optimizing voxelbased image registration," in Medical Imaging: Image Display, Y. Kim, Ed. 1996, vol. 2707 of Proc. SPIE, pp. 554–562, SPIE Press, Bellingham, WA.

[41] F. Maes, D. Vandermeulen, and P. Suetens, "Comparative evaluation of multiresolution optimization strategies for multimodality image registration by maximization of mutual information," Medical Image Analysis, vol. 3, no. 4, pp. 373–386, 1999.

[42] M. J. Clarkson, D. Rueckert, D. L. G. Hill, and D. J. Hawkes, "Registration of multiple video images to pre-operative CT for image guided surgery," in Medical Imaging: Image Processing, K. M. Hanson, Ed., SPIE Press, Bellingham, WA, 1999, vol. 3661 of Proc. SPIE, pp. 14–23.

[43] P. Thévenaz and M. Unser, "Optimization of mutual information for multiresolution image registration," IEEE Trans. Image Process., vol. 9, pp. 2083–2099, Dec. 2000.

[44] J. C. Gee, M. Reivich, and R. Bajcsy, "Elastically deforming 3D atlas to match anatomical brain images," J. Comput. Assist. Tomogr., vol. 17, no. 2, pp. 225–236, 1993.

[45] P.J. Edwards, D.L.G. Hill, J.A. Little, and D.J.Hawkes. A three-component deformation model for image-guided surgery. Med. Image. Anal. 2(4):355-367,1998.

[46] G. E. Christensen, M. I. Miller, and M. Vannier, "3D brain mapping using a deformable neuroanatomy," *Phys. Med. Biol.*, vol. 39, pp. 609–618, 1994.

[47] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes, "Nonrigid registration using free-form deformations: Application to breast MR images," *IEEE Trans. Med. Imag.*, vol. 18, pp. 712–721, Aug. 1999.

[48] C. Studholme, R. T. Constable, and J. Duncan, "Accurate alignment of functional EPI data toanatomical MRI using a physics-based distortion model," *IEEE Trans. Med. Imag.*, vol. 19, pp. 1115–1127, Nov. 2000.

[49] J. Kybic, P. Thévenaz, A. Nirkko, and M. Unser, "Unwarping of unidirectionally distorted EPI images," *IEEE Trans. Med. Imag.*, vol. 19, pp. 80–93, Feb. 2000.

[50] T. Rohlfing, C. R. Maurer Jr, W. G. O'Dell, M. C. Schell, and J. Zhong, "Modeling liver motion and deformation during the respiratory cycle using intensity-based free-form registration of gated MR images," *Proc. SPIE (Medical Imaging 2001: Image Processing)*, vol. 4319, pp. 337–348.

[51] Gustavo K. Rohde, Akram Aldroubi, and Benoit M. Dawant, "The adaptive bases algorithm for intensity-based nonrigid image registration", *IEEE Trans. Med. Imag.*, vol. 22, pp. 1470-1479, Nov. 2003

[52] G. K. Rohde, "The Adaptive grid registration algorithm: A new spline modeling approach for automatic intensity-based nonrigid registration," masters thesis, Vanderbilt Univ, Nashville, TN, 2001.