

IRE : an image registration environment for volumetric medical images
by David Kenneth Lyle

A thesis submitted in partial fulfillment of the requirements for the degree of Master of Science in Computer Science
Montana State University
© Copyright by David Kenneth Lyle (2001)

Abstract:
Multi-modal volumetric medical image registration is accomplished by mapping the coordinate space of one three-dimensional image set to the coordinate space of another three-dimensional image set of the same or differing modality using some transformation function. In this thesis, the transformation function in this report is found by maximizing the mutual information between the image sets. A parabolic search function is used to optimize mutual information with respect to the multivariate transformation function. The results this search produces are presented. Advantages of implementing Parzen windows to estimate probability density are also tested.

The method's performance is demonstrated by registering magnetic resonance (MRI) images with positron emission tomography (PET) images and with computed tomography (CT) images. Additionally, experiments testing the algorithm's ability to handle small deformations of structural features in the images are presented.

IRE: AN IMAGE REGISTRATION ENVIRONMENT FOR VOLUMETRIC MEDICAL IMAGES

by

David Kenneth Lyle

A thesis submitted in partial fulfillment
of the requirements for the degree

of

Master of Science

in

Computer Science

MONTANA STATE UNIVERSITY
Bozeman, Montana

August 2001

N378
L9879

APPROVAL

of a thesis submitted by

David Kenneth Lyle

This thesis has been read by each member of the thesis committee and has been found to be satisfactory regarding content, English usage, format, citations, bibliographic style, and consistency, and is ready for submission to the College of Graduate Studies.

Denbigh Starkey — (Signature) — 8/30/01 — Date

Approved for the Department of Computer Science

Denbigh Starkey — (Signature) — 8/30/01 — Date

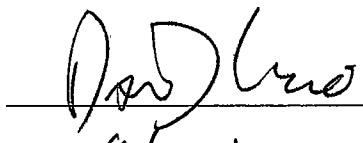
Approved for the College of Graduate Studies

Bruce McLeod — (Signature) — 8-30-01 — Date

STATEMENT OF PERMISSION TO USE

In presenting this thesis in partial fulfillment of the requirements for a master's degree at Montana State University, I agree that the Library shall make it available to borrowers under the rules of the Library.

If I have indicated my intention to copyright this thesis by including a copyright notice page, copying is allowable only for scholarly purposes, consistent with "fair use" as prescribed in the U.S. Copyright Law. Requests for permission for extended quotation from or reproduction of this thesis in whole or in parts may be granted only by the copyright holder.

Signature 

Date 8/30/2001

TABLE OF CONTENTS

LIST OF TABLES v

LIST OF FIGURES vi

ABSTRACT vii

1. INTRODUCTION 1

2. APPROACHES TO REGISTRATION 4
   Geometric Feature Registration 4
   Voxel Similarity Measure Registration 5

3. METHODS 7
   Image Sets 7
   Entropy 9
   Joint Entropy 10
   Mutual Information 12
   Approximating Probability Distributions Using Parzen Windows 12
   Transformation Function 16
   Quaternion Rotation 18
   Greedy Stochastic Gradient Ascent 19
   Viewing and Capturing the Results 20

4. RESULTS 21
   Note on Registration Accuracy 21
   MRI-MRI Registration 22
   CT-MRI Registration 23
   MRI-PET Registration 26
   MRI-MRI Time-Lapse 29
   Parzen Density Estimation versus Exhaustively Sampled Density Estimation 31

5. CONCLUSIONS 34
   Future Directions 34

APPENDICIES 36
   Appendix A – Quaternions and Rotation 36

BIBLIOGRAPHY 40

LIST OF TABLES

Table Page

1. CT-MRI image set composition 23
2. CT-MRI registration results 24
3. MRI-PET image set composition 26
4. MRI-PET registration results 28
5. MRI-MRI time lapse registration results 30
6. Comparison of registration times 32

LIST OF FIGURES

Figure Page

1. Medical image formats 1
2. Methods of interpolation 8
3. Entropy densities 10
4. Example images with grayscale values 11
5. Co-occurrence matrix 11
6. Depiction of kernel 14
7. Comparison of the probability distributions for individual voxel intensities 14
8. Sampling an image 15
9. Probability density histograms 16
10. CT-MRI prior to registration 24
11. CT-MRI post registration 25
12. MRI-PET prior to registration 27
13. MRI-PET post registration 28
14. MRI-MRI time-lapse prior to registration 30
15. MRI-MRI time lapse post registration 31

ABSTRACT

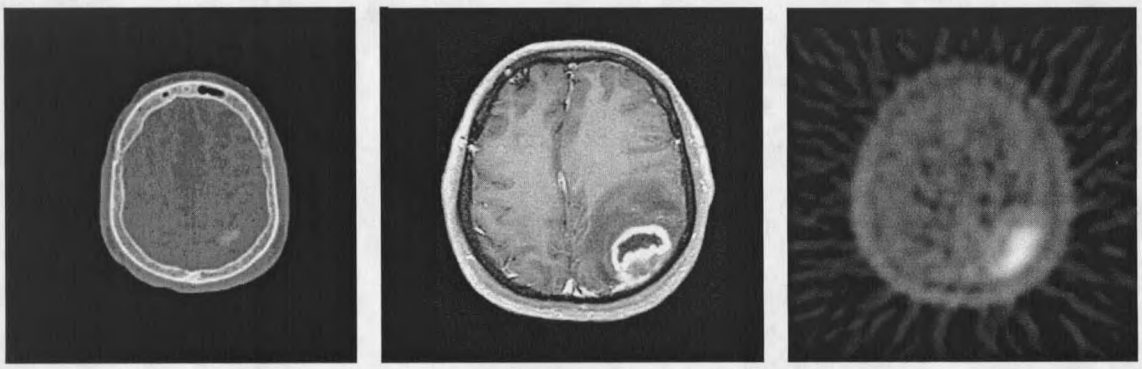
Multi-modal volumetric medical image registration is accomplished by mapping the coordinate space of one three-dimensional image set to the coordinate space of another three-dimensional image set of the same or differing modality using some transformation function. In this thesis, the transformation function in this report is found by maximizing the mutual information between the image sets. A parabolic search function is used to optimize mutual information with respect to the multivariate transformation function. The results this search produces are presented. Advantages of implementing Parzen windows to estimate probability density are also tested.

The method's performance is demonstrated by registering magnetic resonance (MRI) images with positron emission tomography (PET) images and with computed tomography (CT) images. Additionally, experiments testing the algorithm's ability to handle small deformations of structural features in the images are presented.

CHAPTER 1

INTRODUCTION

Medical images are routinely used for diagnosis, treatment planning, treatment guidance and tracking disease progression [1]. Different image modalities provide different information relevant to the treatment planning process. The imaging modalities considered in this thesis include computed tomography (CT), magnetic resonance imaging (MRI), and positron emission tomography (PET). CT imaging excels in providing detail in higher electron density features, such as bone, which provide the most consistent anatomic references. MRI imaging captures soft tissue information, which primarily permits the discernment of different types of soft tissue and anatomical reference. PET images record functional information [2] (Figure 1) and by combining these feature sets the treatment planner is able to take advantage of the complimentary information each modality provides.



(a) (b) (c)

Figure 1. Medical image formats (a) computed tomography (b) magnetic resonance (c) positron emission tomography

In order to obtain maximum benefit from multiple image modalities, the different images must be spatially aligned to each other. This process is called *registration*. The goal of image registration is to find a transformation function that can map the position of features in one image to the position of the same features in another image. Our focus will be on volumetric medical image registration, although our algorithm, described in Section 3, is also able to accommodate single image to single image registration. Registration is achieved by finding the transformation function that maximizes the mutual information between image sets [2, 3].

Mutual information measures the amount of information one image provides about another. One of the main advantages of using mutual information is its non-invasive nature, as fiducial markers [4] are not needed to accomplish accurate registration. Avoiding fiducial marker use is desirable because they can cause the patient considerable discomfort, there are no clear standards governing their use, and they cannot be used for alignment retrospectively.

Calculating mutual information is computationally expensive and the search space for volumetric registration is quite large, thus speed of the algorithm and tractability of the problem are important considerations. Composition of the transformation function can reduce search space complexity and sampling techniques can lessen the computational demands.

The product of this thesis is the development of IRE, an image registration environment for the TIRADE[1] software package developed for the INEEL Boron Neutron Capture Therapy (BNCT) project [5] which is a collaborative effort between

[1] Transport Independent RADiotherapy Environment

Montana State University and the Idaho National Engineering and Environmental Laboratory (INEEL) for the development of radiotherapy treatment planning software for use in BNCT.

CHAPTER 2

APPROACHES TO REGISTRATION

We can divide methods for image registration into two general types, registration using geometric features and registration based on voxel similarity measures. The focus in the medical imaging literature has been on rigid body registration where physical structures of interest neither appreciably deform nor distort between imaging sessions. Alignment based on geometric features is by its very nature a rigid body registration. Voxel similarity measures are still predominantly used in rigid body applications, but are extensible into non-rigid body registration [6].

Geometric Feature Registration

Alignment methods using physical features, specified manually or by an automated method, attempt to find the transformation function that aligns image features. Feature points can be either prominent naturally occurring physical structures or artificially introduced fiducial markers that are physically attached to or marked on the body.

Naturally occurring features present many difficulties. Generally, physical features are not viewable in all modalities and even when viewable may be challenging to locate precisely. Natural features may alter between imaging sessions and can appear differently in scans separated by a time interval. Distinct features may cover several pixels in diameter allowing for several pixel diameters of error. Markers occurring naturally are also problematic because the user interaction necessary to define the points

for comparison has the propensity to introduce error. SERA[2], the predecessor of TIRADE, implements this type of approach.

Artificially introduced fiducial markers are physically placed markers, which are designed to be highly visible and to facilitate automated registration. The use of fiducial markers provides accurate registration and is frequently the standard used to measure error [7], but they introduce complications. Fiducial markers [4] cannot be applied in retrospect, so images taken before introduction of the markers cannot be registered by the same method. Fiducial markers can be invasive, and for non-surgical patients they may not be a desirable option.

Voxel Similarity Measure Registration

Registration using voxel similarity measures attempts to use some metric calculated from the voxel intensities directly to derive the transformation function. Several methods exist based on voxel characteristics for intramodality alignment. Intramodality registration is used to compare images for detection of any pathological or anatomical changes over time assisting in diagnosis and determining the efficacy of implemented treatments.

Using intensity similarities allows for application of intuitive means of measuring registration accuracy. Examples of intuitive metrics used to register intramodality images are the sum of squared intensity differences between images and the correlation coefficient. These methods are limited to images of the same modality and intermodality registration is one of our main goals.

[2] Simulation Environment for Radiotherapy Applications.

Intermodality alignment is more complex because there is no simple relationship to measure misregistration. Many existing methods specialize in mapping one image modality to another utilizing modality specific mapping relationships. Intensity re-mapping of CT images by inverting low and high intensities to more closely approximate MRI intensities allows for cross-correlation as a measure of misregistration. This method relies on *a priori* knowledge of the problem space and is limited to a problem space of MRI to CT registration.

Our desire is to implement and optimize a general metric that performs well without knowledge of the problem space. Mutual information proves to be a versatile measure of registration. Comparisons with other voxel similarity measures for registration of MRI and CT images [8] as well as MRI and PET images [9] found mutual information a more flexible and robust metric. This metric also behaves excellently in intramodality alignment. We will show that mutual information successfully measures registration of mild deformations. The TIRADE IRE implements an algorithm that optimizes alignment based on this information measure.

CHAPTER 3

METHODS

Image Sets

Input to the registration algorithm is two sets of images. Each set is an ordered collection of incremental planar images along a common axis where the axes of the two sets are generally not identical. Information about the distance between planes and the dimensions of each pixel is provided for each set.

A planar slice is composed of a matrix of pixels, where each pixel is a grayscale value that, for the purpose of this thesis, is normalized to an eight bit value between 0 and 255. Pixels are generally square in medical imaging, and an imported image set with rectangular pixels is scaled to provide square pixels. Intensity values for the new, scaled pixels are linearly interpolated from the neighboring pixels in the original image. The second image set is then scaled to match the dimensionality of the first set loaded ensuring that each pixel represents an equivalent area. By scaling the image sets in this manner, scaling need not be considered a parameter in the search.

A voxel is a three-dimensional pixel or volume element. We will use a stricter definition where all three dimensions have equal magnitude and voxels are cubic. Cubic volumes ensure that distances are preserved, so the need to scale and interpolate new voxel intensity values for each rotation or translation is eliminated. This enhances the performance of the algorithm significantly.

The cubic voxels are created from the image slices by interpolating from neighboring slices. The method of interpolation can significantly affect the results of mutual information optimization [10], which is based on pixel values. TIRADE IRE implements a modification of tri-linear interpolation that takes advantage of the square structure of the planar images (Figure 2).

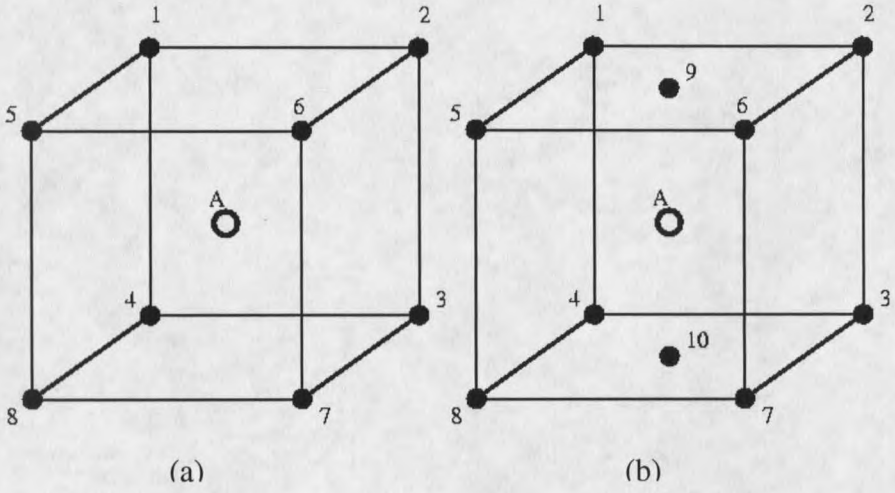


Figure 2. Methods of interpolation (a) tri-linear interpolation (b) modified tri-linear interpolation

In part (a) of Figure 2, the intensities of voxels 1-8 are weighted inversely proportional to their distance from *A* and averaged to determine the intensity of voxel *A* performing tri-linear interpolation. Given the square grid of the plane above and below the desired voxel *A*, modified tri-linear interpolation, as seen in (b), is accomplished when the intensities of voxels 9 and 10 are added to the traditional eight voxels considered in tri-linear interpolation and all weighted inversely proportional to their

distance from *A*. Such an interpolation produces a slightly blurred result, but blurring has demonstrated better results in mutual information searches (see [10] for discussion).

If the dimensionality of the resulting volume is not cubic, then pixels with an intensity value of zero are added equally to both sides of any shorter dimension to create a cubic volume. A cubic volumetric image enables rotation and translation without having to change the dimensions of the image.

In this discussion, the resulting cubic volume image sets will be referred to as the *reference set* and *test set*. The reference set will remain unaltered by the alignment process, while the test set is transformed to maximize mutual information.

Entropy

Entropy is a measure of systemic randomness that is defined for discrete variables as [11]:

$$H(A) = -\sum p_i \log p_i \qquad (1)$$

Where, A is an image or image set, i is a grayscale value and p_i is the probability of that grayscale value occurring in the image (Figure 3). The probabilities are calculated from the image intensity histogram. Entropy is maximized when all voxel intensities in an image have the same probability of occurring, 1/n (Figure 3 part a). The value of entropy is maximized to zero if all voxels in the image share the same grayscale value with a probability of 1. Only when the outcome is known does the entropy value disappear.

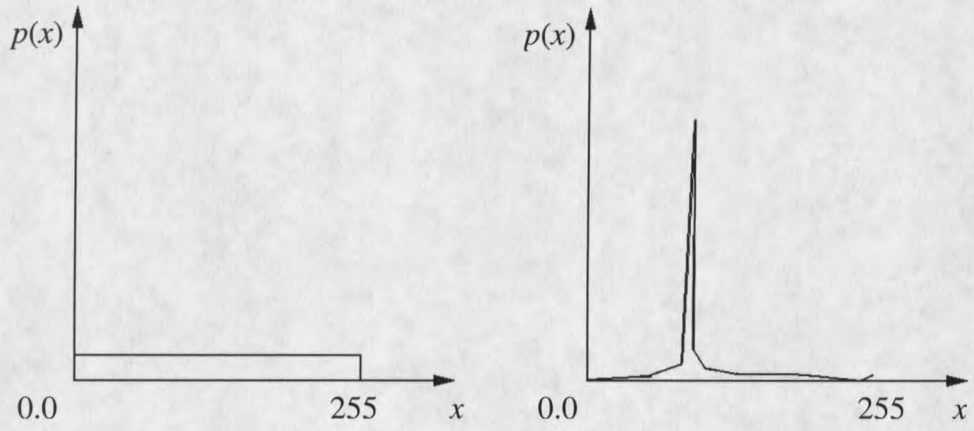


Figure 3. Entropy densities (a) high (b) low

Joint Entropy

Joint entropy measures how well the voxel intensity of a specific location in one image predicts the voxel intensity in the other image. Joint entropy is defined as [11]:

$$H(A,B) = -\sum p(i,j) \log p(i,j) \quad (2)$$

Where $p(i,j)$ is the joint probability that i is the grayscale value at a specific coordinate location in image A, and that j is the grayscale value at the same location in B. Probabilities are calculated from the co-occurrence histogram of A and B. If A and B are completely unrelated, then their joint entropy is the sum of the entropies of A and B individually. If two images are identical, the more uniform the joint probability matrix and the lower the joint entropy.

The calculation of joint entropy is demonstrated in the following example. Consider two 4x4 images, with grayscale values ranging from [0-3] in Figure 4 parts a and b.

2	2	3	2	
1	0	3	3	*y*
1	2	2	3	
0	0	1	2	

x

(a)

1	2	3	2	
0	1	2	3	*y*
2	2	3	2	
0	1	1	1	

x

(b)

Figure 4. Example images with grayscale values (a) image 1 (b) image 2.

Calculating the joint entropy of the two images requires generation of the co-occurrence histogram. The co-occurrence histogram represents how many times a grayscale value *i* in image 1 at location *x*,y is matched with a grayscale value *j* in image 2 at location *x*,*y*. The resulting co-occurrence matrix for image 1 and image 2 is given in Figure 5.

3	0	0	2	2
2	0	2	3	1
1	1	1	1	0
0	1	2	0	0
	0	1	2	3

image 1

image 2

Figure 5. Co-occurrence matrix. Grayscale value co-occurrences between images.

By dividing by the number of combinations (16 in the example), the $p(x,y)$ values in Equation 2 can be found. For this example, the joint entropy is:

$$5 \times \frac{1}{16} + 3 \times \frac{2}{16} + 2 \times \frac{3}{16} = 1.0625$$

Mutual Information

Image registration can be thought of as attempting to maximize the information in common between the two images [2, 3]. In medical images, if two images are aligned, then the resulting combined image will not show the same structures in different spaces. One means of characterizing the amount of information shared by two images is mutual information. Mutual information is defined as [11]:

$$I(A,B) = H(A) + H(B) - H(A,B) \tag{3}$$

The advantage of this metric over joint entropy alone is that should the transformation of B translate all non-zero intensity voxels out of the coordinate space of the search, then the entropy value of B will be zero resulting in a lower information score regardless of the increase in the joint entropy of A and B. It also places the base line at zero when A and B are independent.

Approximating Probability Distributions Using Parzen Windows

The calculation of entropy based on the entire image is computationally expensive. The necessary amount of calculation is reduced by estimating the probability distribution from a relatively small sample rather than calculating probabilities from the histogram created from the entire image. For a single 128x128x128 voxel image set, 128^3 voxels are examined to determine the image histogram for a 100% sample, referred to here after as an *exhaustive sample*. To calculate the co-occurrence histogram exhaustively, this number is squared. Sampling only a portion of this information reduces the calculations necessary. We employ the non-parametric method called *Parzen*

Windows [12, 13] to estimate the probabilities. The general form of Parzen window density estimation is:

$$p_n(x) = 1/n \sum R(x - x_i) \tag{4}$$

where n is the number of voxels sampled and R is the kernel smoothing function. We have empirically optimized the size of n to 10 percent of the total number of voxels when calculating entropy for an image. Fitness of a sample size is determined by how well it estimates the probability distribution calculated from an exhaustive sample.

A Gaussian density function:

$$R(x - x_i) = \frac{1}{\sigma(2\pi)^{0.5}} e^{\frac{1}{2}\frac{(x-x_i)^2}{\sigma^2}} \tag{5}$$

is used for R, which is commonly referred to as the *kernel*. The kernel has been empirically tested to be seven in width and a bin size of 1 (Figure 6). The goal of these optimizations is to reproduce the probabilities calculated from the entire image, or pair of images, as accurately as possible while reducing computational expense. To best approximate the results obtained from exhaustive sampling, we found that having a kernel, as shown in Figure 6, with a large variance was optimal. The larger variance assists in smoothing the sample data.

Figure 7 compares the sample-derived probability distribution and the exhaustive sample based probability distribution. The exhaustive sample plot charts the probability distribution calculated using all pixel intensities in the image. The four and ten percent plots show probability distributions based on the relative sample size. Samples are taken by using the intensities of evenly spaced clusters of twenty voxels in the image set

(Figure 8). The probability distribution from the sample is then smoothed by the Parzen windowing method. Figure 7 illustrates that the ten percent sample still follows the features of the exhaustive sample curve while a smaller sample deviates from the exhaustive sample curve. Registration performance was poor with sample sizes smaller than 10 percent.

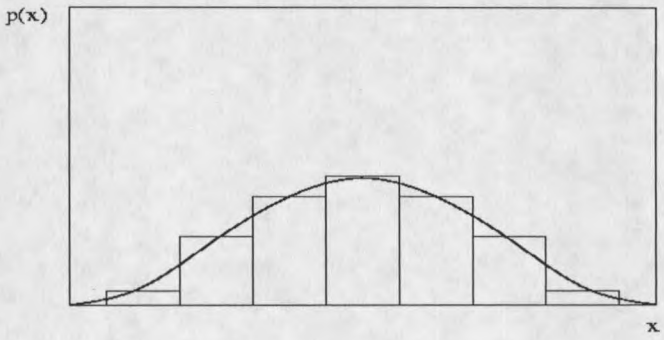


Figure 6. Plot of Parzen kernel. Each box represents one intensity level.

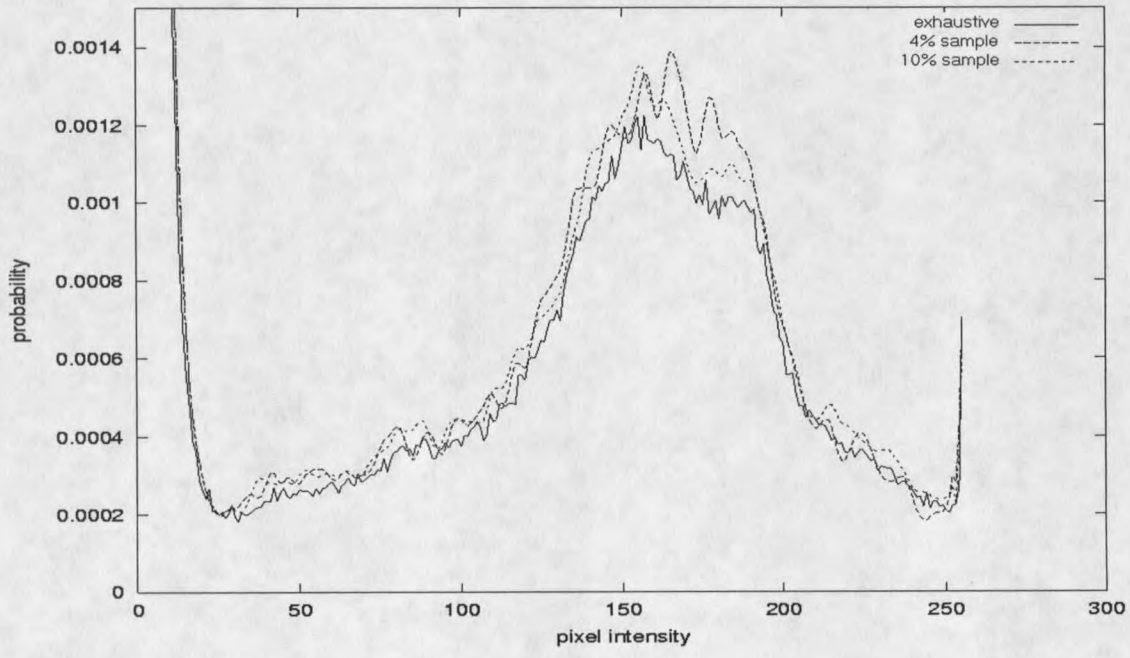


Figure 7. Comparison of the probability distributions for individual voxel intensities.

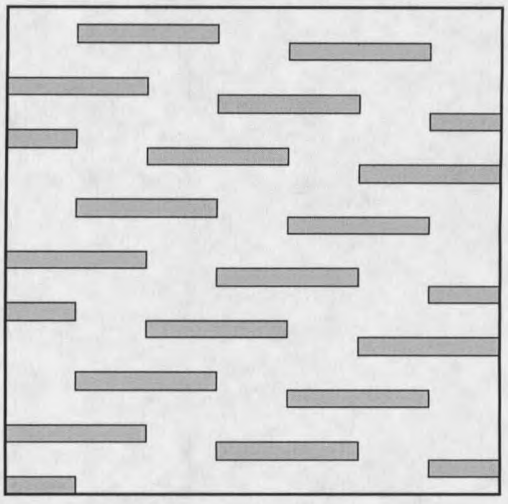
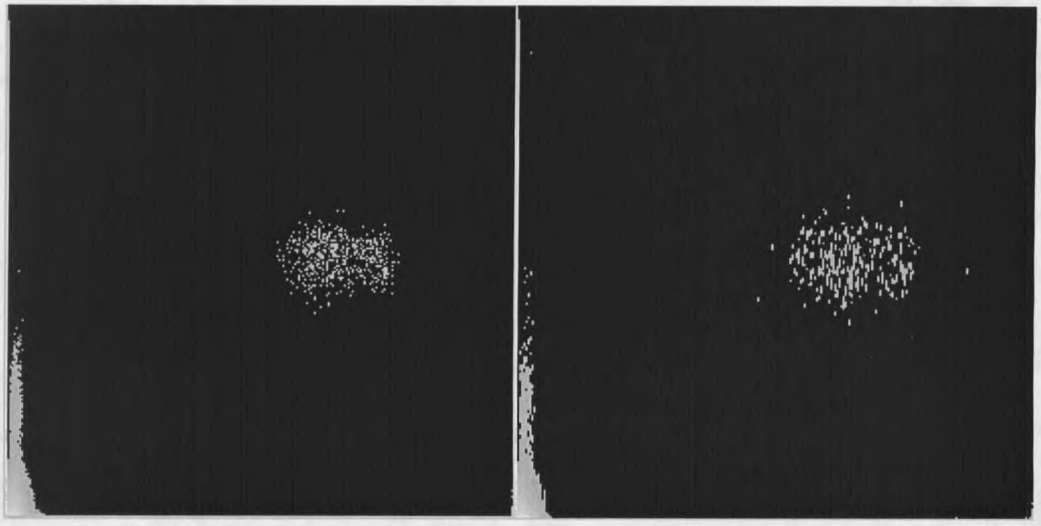
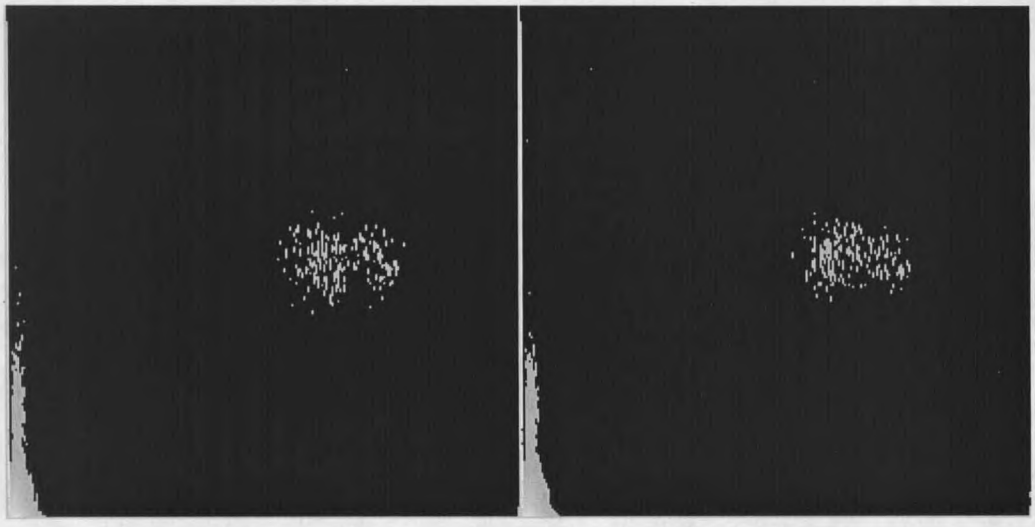


Figure 8. Sampling an image. Dark areas represent samples taken.

When calculating joint entropy, only 20 percent of the co-occurrence histogram needs to be derived. Figure 9 illustrates the exhaustive sample and partial sample based joint probability density histograms. The joint probability density histogram that is generated from creating the entire co-occurrence histogram is shown in part (a). Parts (b), (c) and (d) are derived from calculating samples of ten, twenty and thirty percent, respectively, of the co-occurrence histogram. As with estimating entropy, samples are taken by using the intensities of evenly spaced clusters of twenty voxels in the image set (Figure 8). The probability densities from the sample are then smoothed by the Parzen windowing method. Comparing (b) with (a), it is apparent that a ten percent sample is noisy and adds misinformation that complicates the search process. The twenty percent sample is sufficient to accurately emulate the joint probability histogram derived from an exhaustive sample.



(a) (b)



(c) (d)

Figure 9. Probability density histograms derived from (a) entire co-occurrence histogram (b) 10 percent sample (c) 20 percent sample (d) 30 percent sample

Transformation Function

In image registration, the objective is to find the transformation, T, from the coordinate frame of the test set to the reference set. A voxel in the test set is denoted $v(x)$

and $u(x)$ is a voxel in the reference set. So, $v(T(x))$ is the voxel from the test set that is associated with the reference set voxel $u(x)$. The mutual information equation (Equation 3) can be expressed in terms of T as:

$$I(u(x),v(T(x))) = h(u(x)) + h(v(T(x))) - h(u(x),v(T(x))) \quad (6)$$

Where ultimately, T will be the transformation that provides the maximum value for I.

There are many types of transformations applicable in three-dimensional medical image registration. The most common are rotation and translation. Scaling and skewing are considered in some algorithms. The registration algorithm presented here does not consider scaling as part of the optimization problem. Scaling is handled on image importation as discussed in Section 3.1. Skewing does not provide adequate improvement in the rigid body registration result to warrant the additional computational complexity [1]. T is a rigid body transformation where all distances are preserved. We define T as being composed of six variables, translation in x, y and z, and rotation about x, y and z. The variables of T are then converted to a transformation matrix that is applied to all points in the test set.

Limiting the number of design variables is important in increasing the speed of the algorithm. Each design variable adds to the complexity of the search space. We have set the number of variables in the transformation function to six by scaling the images prior to the search and eliminating skewing as a necessary variable. This improves the speed of the algorithm and reduces the complexity of the search space.

Quaternion Rotation

Three-dimensional rotation can be represented by the combination of three right-hand rotations about the x, y and z axes. This method of rotation is called Euler rotation [14]. Rotations completed in this fashion are non-commutative and must follow strict ordering conventions. This makes the search space discontinuous.

Quaternions (Appendix A) are numbers with one real part and three imaginary parts. A quaternion is written as:

$$\mathbf{q} = w + xi + yj + zk \tag{7}$$

The coefficients x, y and z in the imaginary terms are real numbers that represent elements of a vector in three-dimensional space which can represent an arbitrary axis of rotation. The real number w is a scalar that can represent the magnitude of the rotation. The quantities i, j and k are three mutually orthogonal unit quaternions. The use of quaternions eliminates the non-commutative nature of rotation about an arbitrary axis.

Quaternions by definition have four components (w, x, y, z), but we can represent w as the magnitude of the direction vector:

$$w = \sqrt{x^2+y^2+z^2} \tag{8}$$

The vector (x, y, z) is then normalized by the value of w. Thus, the number of variables involved in rotation remains three and the laws of quaternion rotation are preserved. Since w is the magnitude of the vector (x, y, z) any change in x, y or z will result in a change in the amount of rotation.

Greedy Stochastic Gradient Ascent

Three rotations and three translations form the independent variables of our optimization routine. The objective is finding the transformation T that maximizes mutual information as rapidly as possible. We use a variation of Powell's [15] algorithm for our gradient-based optimization method.

Powell's algorithm is an iterative optimization method that loops through each of the design variables finding the maximum in the dimension of the search space that the design variable represents. Each dimension is maximized in turn until the exit criterion is reached. So, for x, y and z:

```
repeat until exit criterion satisfied
        maximize mutual information in x
        maximize mutual information in y
        maximize mutual information in z
```

Our method differs by selecting the dimension that promises the maximum gain based on the gradient and pursues optimality in that dimension. This process is repeated until all gradients are near zero.

```
repeat until max information gain is less than minimum
        find gradients in all 6 dimensions
        do for all dimensions
                if information gain is greater than 0 and greater than max gain
                        max gain equals information gain
        find maximum in dimension of max gain
```

Gradient-based methods perform well in situations where the function is differentiable. Entropy and joint entropy both meet this criterion [11, 2]. This method explores the search space, testing the gradient in six directions from the current location and determining the direction of greatest ascent.

Stochastic approximation is used to determine the gradients at each location rather than calculating the true gradients. Estimations are accomplished by evaluating the mutual information function in close proximity to the current position in both the positive and negative direction. We found that sampling the function in only one direction did not provide an accurate enough estimate of the true gradient. Although this is more computationally expensive for each gradient calculation, the more accurate gradient results in fewer steps to find the maximum.

Having found the gradients, the search now finds the maximum in one dimension. The dimension searched is that of the variable with the greatest gradient. Brent's [15] method is used to locate this one-dimensional maximum by bracketing and then finding the best parabolic fit. The algorithm is prevented from searching in the same dimension in successive iterations, searching until no further improvement can be found. Testing is performed attempting to escape small local maxima by calculating information gain for five iterations at increasing distances. When no further information gain is found, the mutual information is maximized.

Viewing and Capturing the Results

With TIRADE IRE, the user has the option to choose the resolution at which the images should be registered. Acceptable resolutions are 64, 96, 128, 256 or 512 pixels per image plane. This feature allows for greater speed when aligning images, as fine resolution is not always necessary for the registration process.

CHAPTER 4

RESULTS

Experiments were conducted registering image sets using several different resolutions. The speed of the algorithm is dependent on the number of voxels in the image space that must be operated on, so when a lower resolution is used, registration times are shorter. Accuracy can be sacrificed for speed. For example, at the lowest resolution tested, 64x64x64, each voxel roughly represents a volume of $64mm^3$. This allows for a large margin for error. If a rough alignment is desired quickly, then for most registration problems the IRE requires between two and four minutes. Testing on multiple resolutions demonstrated that 128x128x128 provided the best blend of speed and accuracy. The following experiments use a resolution of 128x128x128.

Note on Registration Accuracy

Accuracy of the following registration trials was determined visually, as we did not have access to sets of images where the absolute registration parameters are known. The accuracy of intramodal and CT-MRI registration are relatively easy to determine visually. MRI-PET registration accuracy is more difficult to gauge because anatomical structures and borders in PET images are not clearly represented. When results proved inconclusive we counted this as registration failure.

MRI-MRI Registration

In this section, we describe the results of a series of experiments using the TIRADE IRE for registering one MRI image and a copy of that MRI image set transformed in some manner. These experiments are used as a base case for the registration algorithm, and provide a basis for evaluating the accuracy for the algorithm. The alignment of an MRI image set that has been rotated, translated or both to the same MRI image set performs accurately and consistently. Certain limitations were observed. Consistent registration using the IRE is possible for any combination of rotations under forty-five degrees, about the x, y and z axes, and any combination of translations in x, y and z that still allow for some overlap between the two copies of the image set. These limitations will hold for all intramodality and intermodality registrations in IRE.

The limitations of IRE can be explained by the nature of the optimization algorithm. Using mutual information requires an overlap between image sets where some physical structure in one image occupies the same coordinate space as any physical structure in the other image. Since the search samples mutual information values in close proximity to the current position, a gain in mutual information will not be found without image overlap.

Our algorithm is susceptible to stopping at local maxima rather than the global maximum. For rotations over 45 degrees, we observed that the search direction of the algorithm becomes unpredictable and a local maximum 180 degrees may be pursued rather than the true maximum. The affinity for a 180 degrees rotation is due to the oval shape of the human head. Rotational differences this large are easy to recognize visually.

We designed the IRE to allow the user to approximate alignment of the two image sets manually before running the optimization routine. Manual alignment speeds up the search process and steers the algorithm away from several local maxima.

CT-MRI Registration

Experiments were performed to gauge the ability of the IRE to register CT images and MRI images from the same patient. We used a single pair of MRI and CT images provided courtesy of the Visible Human Project [16]. The MRI image set was used as the test set and the CT image set as the reference set. Mutual information is a symmetrical measure and will produce equivalent results using the CT image set as the test set. Dimensions of image sets can be seen in Table 1. All planar slices are composed of 256x256 pixels.

SET	TYPE	SLICES	PIXEL WIDTH	PIXEL HEIGHT	SLICE SPACING
1	CT	64	1.0	1.0	4.0
1	MRI	33	1.0	1.0	5.0

Table 1. CT-MRI image set composition. All measurements are in millimeters.

Table 2 summarizes the experiments. Multiple initial positions of the MRI image set are tested. The initial parameters cover all trials. The distance of translation from the registered position in the x, y and z dimensions are averaged and given in millimeters. The rotation in degrees about the x, y and z axes are averaged for all trials. The final parameters represent the average deviation from the parameters of the registered position

for only the successful trials. The visual success rate represents percentage of samples that were classified as correctly registered by visual inspection.

INITIAL		FINAL						VISUAL SUCCESS RATE
Average ΔT	Average $\Delta\theta$	Translation in mm			Rotation in degrees			
mm	°	σ_x	σ_y	σ_z	σ_α	σ_β	σ_γ	%
13.1528	2.6068	1.1217	0.5756	1.6850	0.2416	0.4833	0.1656	85.7

Table 2. CT-MRI registration results.

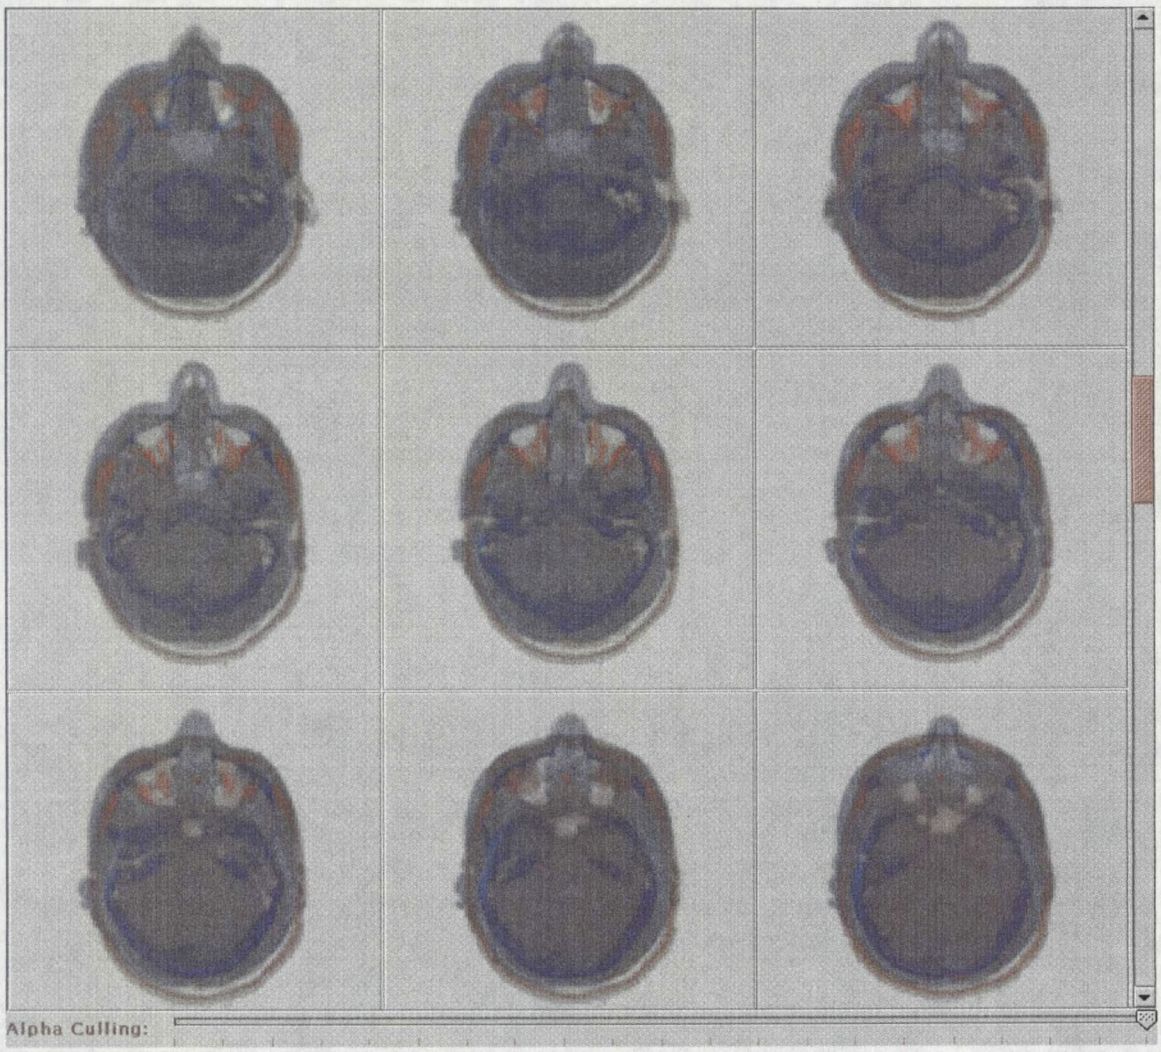


Figure 10. CT-MRI prior to registration.

Alignment accuracy for CT-MRI image sets is high. See Figures 10 and 11 for registration results. CT and MRI images contain well-defined complimentary structures that assist in alignment. Accuracy decreased only as the initial position varied increasingly from the aligned position.

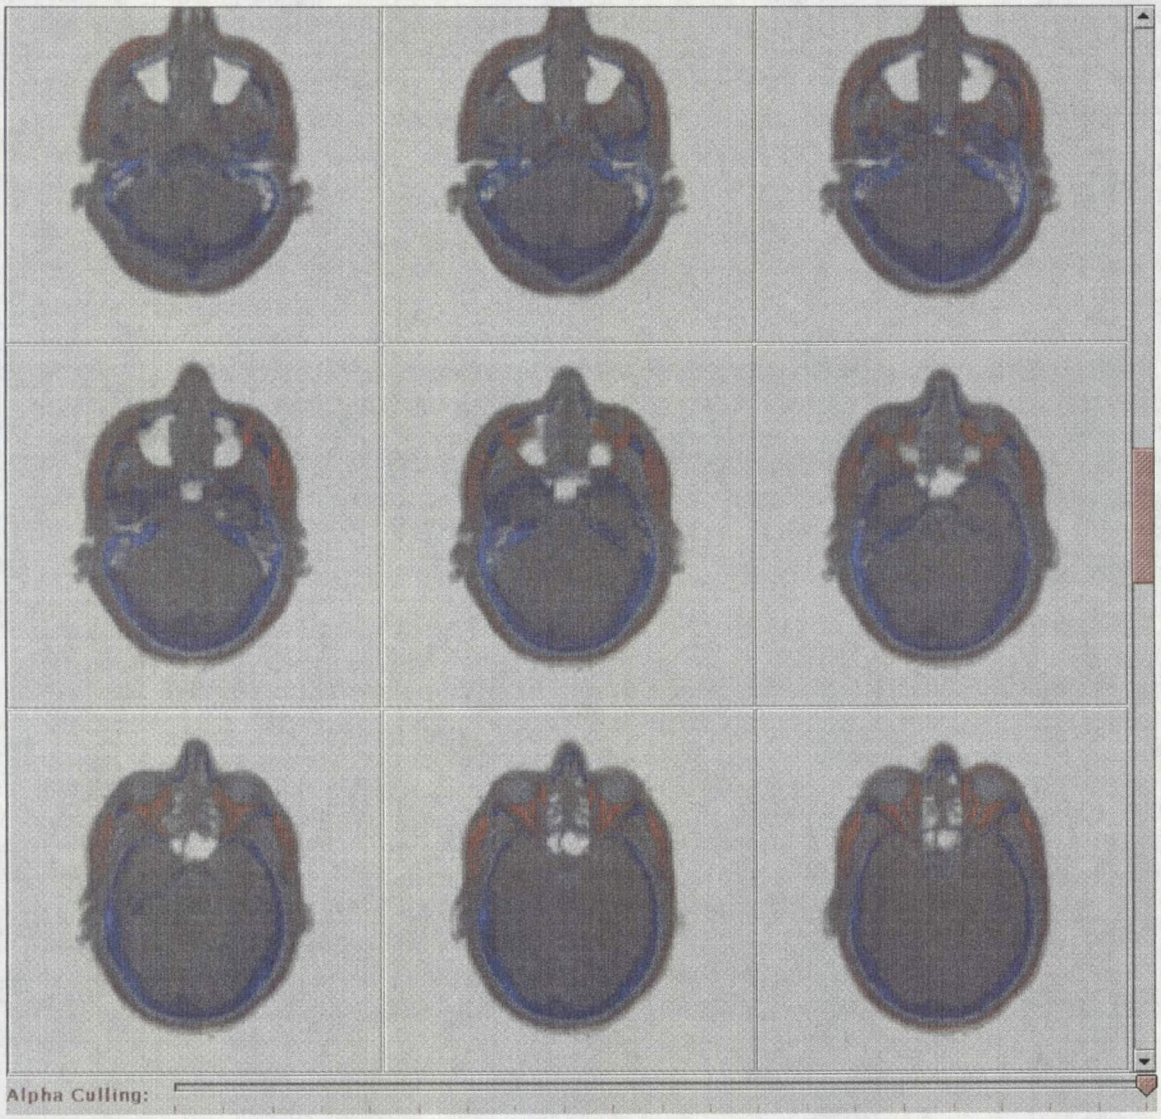


Figure 11. CT-MRI post registration.

MRI-PET Registration

This section describes experiments performed to investigate the ability of the IRE to register PET and MRI image sets. The PET image sets were used as the test sets and the MRI image sets as the reference sets. Four pairs of MRI and PET image sets were analyzed. Dimensions of the original image sets can be seen in Table 3. All planar slices are composed of 256x256 pixels.

SET	TYPE	SLICES	PIXEL WIDTH	PIXEL HEIGHT	SLICE SPACING
1	MRI	20	0.89375	0.89375	7.5
1	PET	47	1.05470	1.05470	3.375
2	MRI	20	0.89375	0.89375	7.5
2	PET	47	1.05470	1.05470	3.375
3	MRI	20	0.89375	0.89375	7.5
3	PET	44	1.05470	1.05470	3.375
4	MRI	19	0.89375	0.89375	8.0
4	PET	22	1.05470	1.05470	7.0

Table 3: MRI-PET image set composition. All measurements are in millimeters.

MRI and PET image set registration poses a more difficult problem than CT and MRI image set registration. PET is a functional representation and MRI a structural representation; this presents difficulties as different portions of the same structural region may have varied functional representations. Where MRI and CT image set features compliment each other, PET and MRI image set features can interfere with each other. This makes accurate alignment difficult.

Table 4 represents the experiments performed. Multiple initial geometric positions of the PET image set were attempted. Registration accuracy was largely dependent on the starting position. MRI and PET image registration presents numerous

local maxima in the mutual information search space. Close proximity to the correct alignment, within a few degrees, is necessary. If the user begins with a reasonably close alignment, then the IRE is capable of accurate registration. See Figures 12 and 13 for first pair of image set registration results.

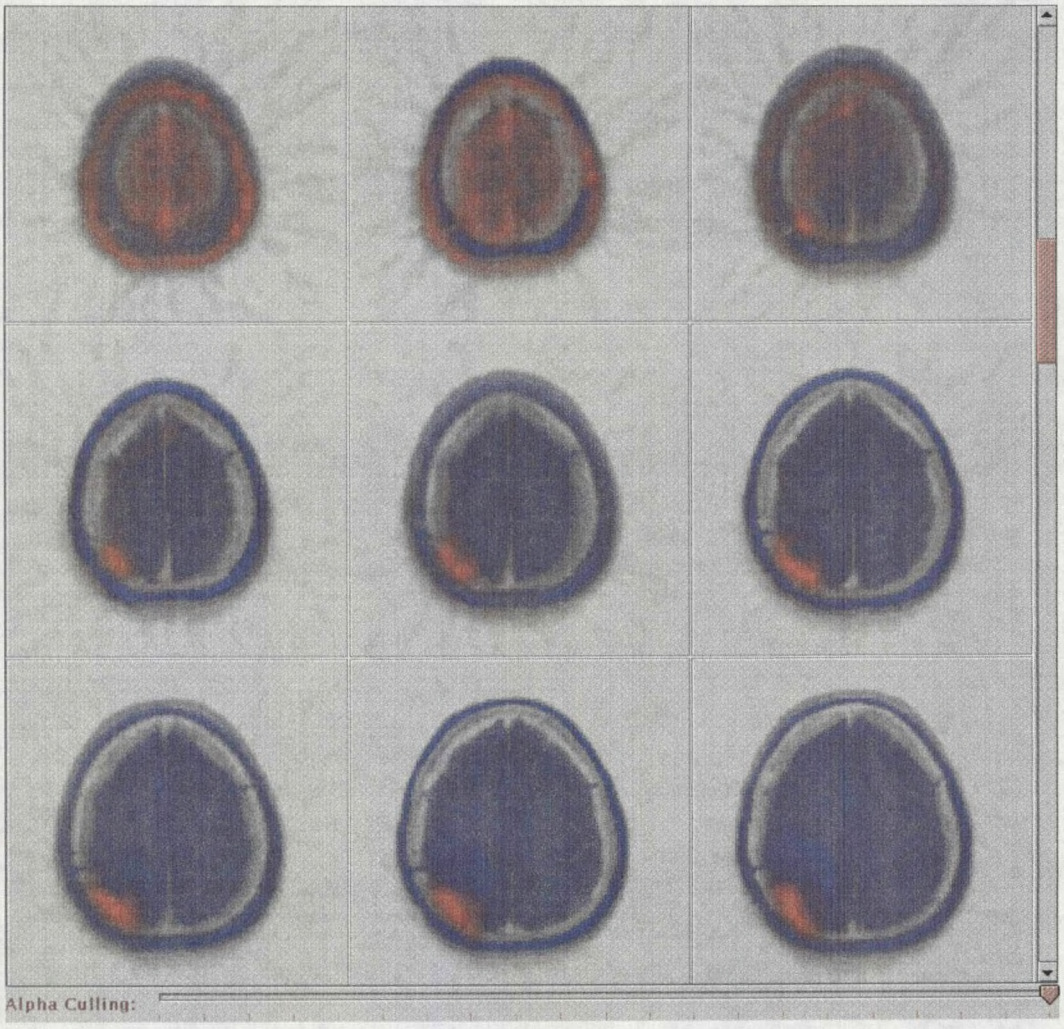


Figure 12. MRI-PET prior to registration.

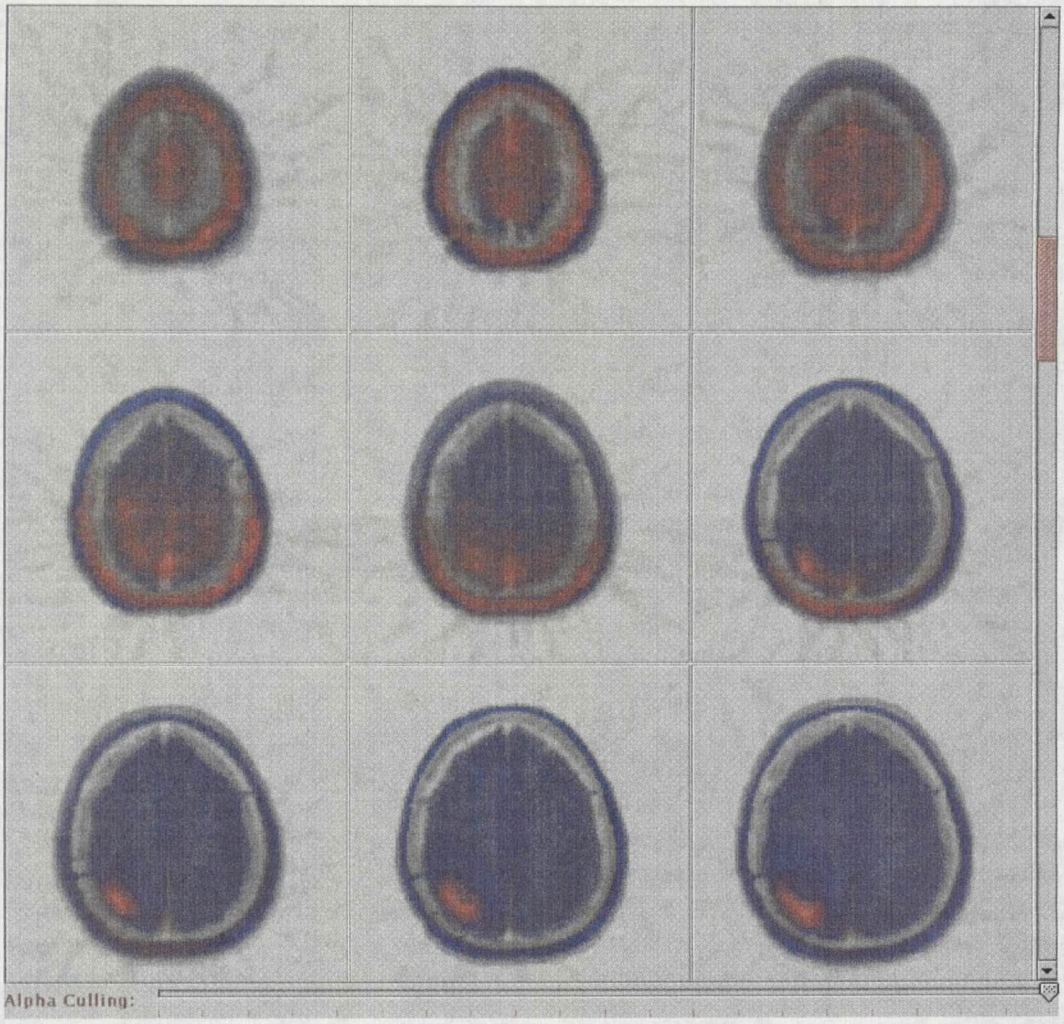


Figure 13. MRI-PET post registration.

SET	INITIAL		FINAL						VISUAL SUCCESS RATE
	Average ΔT	Average $\Delta\theta$	Translation in mm			Rotation in degrees			
	mm	°	σ_x	σ_y	σ_z	σ_α	σ_β	σ_γ	%
1	16.3620	3.9416	3.1818	13.1449	11.5886	5.0214	0.9781	1.1147	77.78
2	16.0975	3.3293	4.4015	5.5452	7.4432	2.0657	0.1452	1.8524	33.33
3	7.4414	2.1372	6.5705	9.5202	10.3409	4.7969	1.6652	2.6211	62.50
4	12.7920	3.4888	3.9784	7.4831	5.4907	2.9509	0.3977	0.7449	66.67

Table 4. MRI-PET registration results.

MRI-MRI Time-Lapse

The experiments described in this section test the algorithm's strength on registration of two sets of MRI images for the same patient with a time lapse between scans. In this case, we were limited to a single pair of image sets.

The MRI image sets used in these experiments represent pre and post-surgical scans, which presents an interesting problem in image alignment. Structural changes are present between scans and inflammation is prominent. The mass of the tumor has been removed between scans surgically through the skull, locally changing the structure of the brain and skull. Such deformations force the mutual information search to favor matching regions over structurally altered regions. Outside of a rigid body, this is more difficult as the volume of altered regions may be greater than the volume of unaltered regions.

Both MRI image sets contain slices of 44 planar images at 256x256 pixels. The pixels are square with a side length of 1.015625mm. The planar images are evenly spaced with 5.0mm between images.

Several starting configurations were attempted. Table 5 summarizes the results. Even though there are slight structural deformations between scans and superficial inflammation, the IRE registers the two image sets accurately. See Figures 14 and 15 for registration results. Misregistration occurs only as the starting position varies largely from the aligned position.

INITIAL		FINAL						VISUAL SUCCESS RATE
Average ΔT	Average $\Delta\theta$	Translation in mm			Rotation in degrees			
mm	°	σ_x	σ_y	σ_z	σ_α	σ_β	σ_γ	%
8.6271	4.8478	0.5632	0.6418	0.9806	0.2085	0.2410	0.1487	85.7

Table 5: MRI-MRI time-lapse registration results.

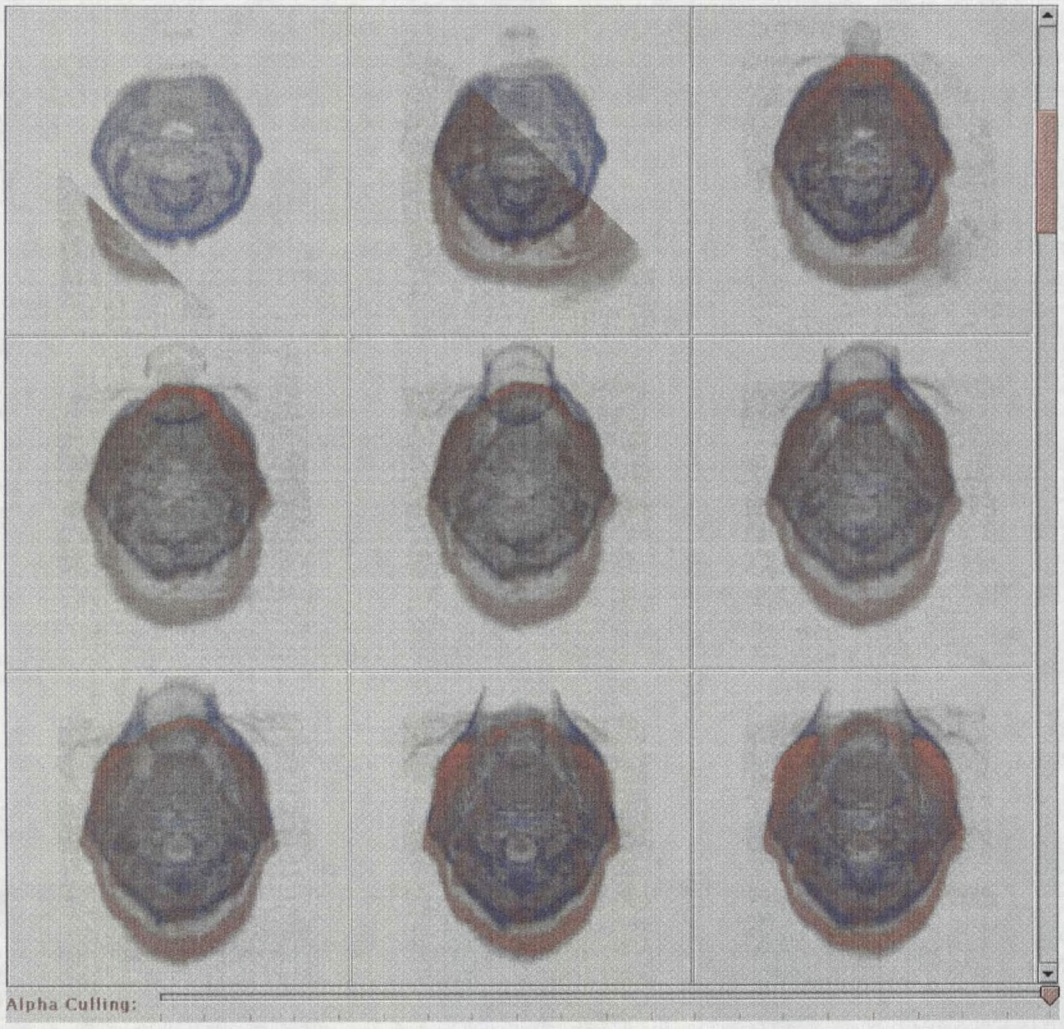


Figure 14. MRI-MRI time-lapse prior to registration.

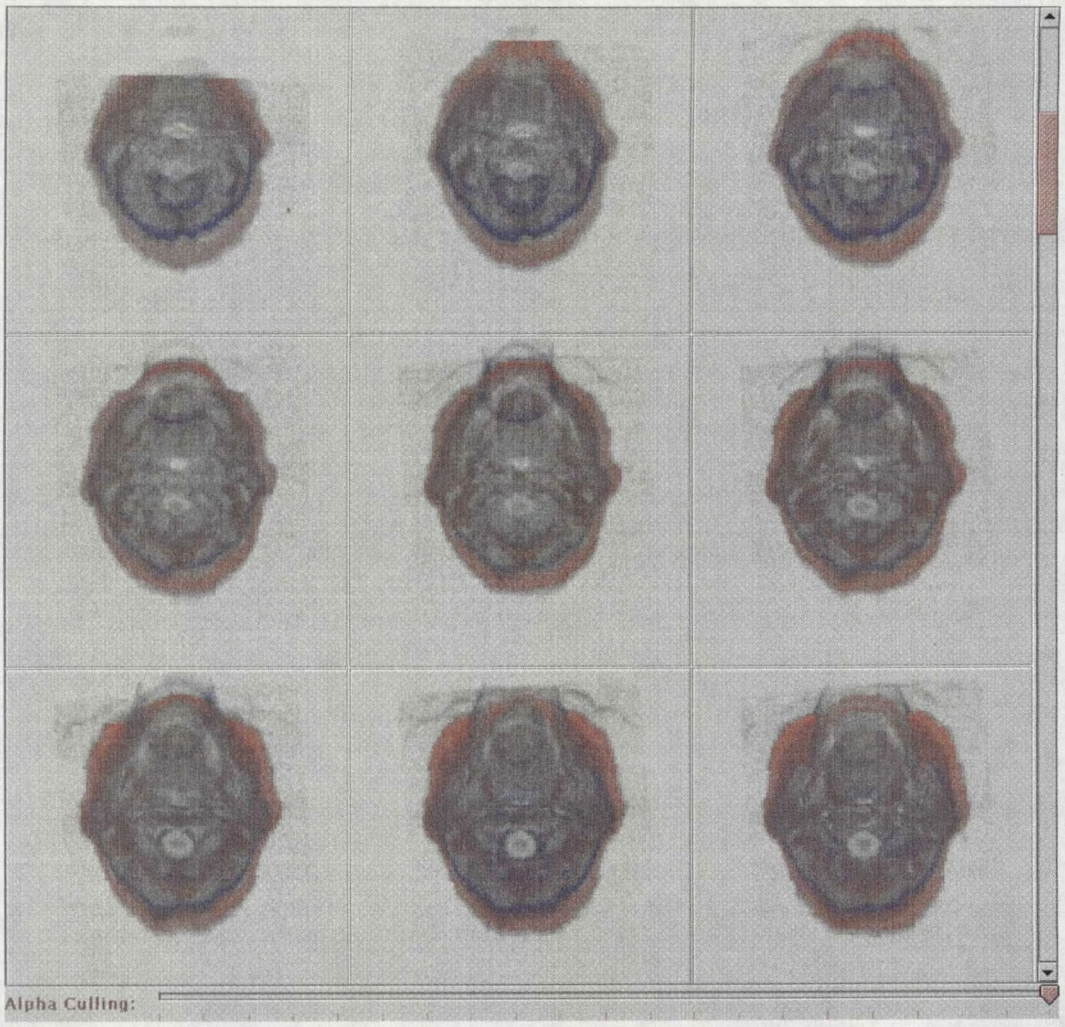


Figure 15. MRI-MRI time-lapse post registration.

Parzen Density Estimation versus Exhaustively Sampled Density Estimation

Experiments were performed comparing the mutual information optimization performance using Parzen windows to estimate probability densities and calculating probability densities from an exhaustive sample. For all of the experiments described in Sections 4.3, 4.4 and 4.5, optimization was conducted using both types of probability density estimation. All experiments were performed on a machine with a Pentium III

700Mhz. The results of these experiments are presented in Table 6. All times are seconds. Times represent averages of registration times for all successful trials.

	CT-MRI AVERAGE TIME	MRI-PET AVERAGE TIME	TIME-LAPSE AVERAGE TIME	AVERAGE
Parzen Density Estimation	941	875	1045	954
Exhaustive Density Estimation	1047	1281	1680	1336
% Gain	10.12	31.69	37.80	28.59

Table 6. Comparison of registration times using Parzen windows for density estimation and density estimation calculated from an exhaustive sample.

A decided reduction in registration times was observed when the optimization algorithm implemented Parzen windows. It is important that the densities estimated using sampling and smoothed by the Parzen windows method closely approximate the density estimations calculated from an exhaustive sample. It was observed that if the sampled densities did not closely follow the densities derived from exhaustive sampling, misinformation could lead to misregistration, as erroneous probabilities corrupt both the entropy and joint entropy calculations.

In contrast, it was observed that a smoother probability density function leads to quicker transversal to the global maximum. Smoothing of the probability density function results in fewer local maxima where the mutual information may become temporarily trapped. Smoothing out these maxima reduces the number of iterations necessary to reach the maximum. As each iteration requires at least twenty mutual information calculations, reducing the number of iterations is desirable.

Implementation of Parzen windows results in an average speed up of approximately 29%. The use of Parzen windows in the search algorithm did not adversely affect accuracy. Anomalous failures were observed with both types of probability density estimation, but success rates were equivalent.

CHAPTER 5

CONCLUSIONS

Mutual information is an adaptive and robust metric for volumetric medical image registration and the IRE provides an accurate image registration environment for the TIRADE radiotherapy treatment planning process. Intramodality and intermodality image alignment are now possible in the TIRADE treatment planning software package. Small deformations in temporally distinct images are also within the registration capabilities of the IRE.

The IRE presents a tractable solution to the image registration problem. Times required to accurately align volumetric image sets can be as little as two minutes and are generally less than thirty minutes on a Pentium III 700Mhz. The implementation of voxel intensity sampling followed by smoothing with Parzen windows provides a substantial search time reduction without sacrificing accuracy. Thus, the IRE represents a quick, easy to use, versatile environment.

Future Directions

Modification of the algorithm could allow for a greater speed increase. The major factor limiting faster running times is the need to transform the entire image volume for each calculation of mutual information. One method of pursuing this could be based on knowledge of the sampling procedure for the Parzen windows density estimation technique. Knowing the locations of the voxels to be sampled could facilitate use of the

inverse of the transformation matrix to back calculate the source of those pixels without having to transform the entire image volume. The expense of calculating the inverse of the 3x3 transformation matrix should be considerably less than transforming all voxels.

Another extension to pursue is the use of normalized mutual information to allow for more than small deformations. One of the goals of the TIRADE is to be able to help in the treatment process for all radiation modalities all areas of the body including such deformable regions as the chest cavity and abdomen. Normalized mutual information has been proposed [6] as a method to deal with this type of deformation.

APPENDIX A

QUATERNIONS AND ROTATION

Sir William Hamilton invented quaternions in 1843 to extend complex numbers into four dimensions. Quaternions can be used to represent a point in four-dimensional space. Quaternions are numbers with one real part and three imaginary parts. A quaternion is written as:

$$\mathbf{q} = w + xi + yj + zk$$

The coefficients x, y and z in the imaginary terms are real numbers that represent elements of a vector in three-dimensional space which represents the arbitrary axis of rotation. The real number w is a scalar that represents the magnitude of rotation. The quantities i, j and k are three mutually orthogonal unit quaternions that have three zero elements and one positive element defined as [17]:

$$i = (0, 1, 0, 0)^{\mathrm{T}},\ j = (0, 0, 1, 0)^{\mathrm{T}},\ k = (0, 0, 0, 1)^{\mathrm{T}}$$

where

$$i^2 = j^2 = k^2 = ijk = -1,\ ij = -ji = k,\ jk = -kj = i,\ ki = -ik = j$$

Note, ij is not equal to ji. If ij were equal to ji, then there would be some quaternion **a** not equal to zero and some quaternion **b** not equal to zero where **a**·**b** does equal zero. An algebra where **a**·**b** equals zero if and only if **a** equals zero or **b** equals zero is called a division algebra. It turns out that there can only be three division algebras with associative multiplication, real numbers, complex numbers and quaternions [18].

Finding the conjugate and magnitude of a quaternion is similar to finding the conjugate and magnitude of complex numbers.

$$\mathbf{q}' = w - xi - yj - zk$$

$$\|\mathbf{q}\| = (\mathbf{q} * \mathbf{q}')^{0.5} = (w^2 + x^2 + y^2 + z^2)^{0.5}$$

A unit quaternion has a magnitude of one and satisfies the relationship:

$$\|\mathbf{q}\| = w^2 + x^2 + y^2 + z^2 = 1$$

The inverse of a quaternion is defined as:

$$\mathbf{q}^{-1} = (w - xi - yj - zk) / (w^2 + x^2 + y^2 + z^2) \quad \text{or} \quad \mathbf{q}^{-1} = \mathbf{q}' / \|\mathbf{q}\|$$

For a unit quaternion its inverse is equivalent to its conjugate.

$$\|\mathbf{q}\| = 1 \Rightarrow \mathbf{q}^{-1} = \mathbf{q}'$$

A vector in Euclidean three-space can be expressed as a quaternion with the scalar w equal to zero. Addition, subtraction of quaternions is done by adding scalars to scalars and vectors to vectors. Scalar multiplication is achieved by multiplying the terms w, x, y and z by the desired scalar. We can represent a quaternion as $(s,\mathbf{v})$ where s is the scalar w and $\mathbf{v}$ is the vector composed of elements x, y and z. The product of two quaternions is given by:

$$\mathbf{q}_1 * \mathbf{q}_2 = (s_1 s_2 - \mathbf{v}_1 \mathbf{v}_2, s_1 \mathbf{v}_2 + s_2 \mathbf{v}_1 + \mathbf{v}_1 \times \mathbf{v}_2)$$

Note, quaternions are not commutative,

$$\mathbf{q}_1 * \mathbf{q}_2 \neq \mathbf{q}_2 * \mathbf{q}_1$$

Shortly after Hamilton's formulation of quaternions, Arthur Cayley published a paper showing that rigid body rotations could be represented using quaternions [19]. When a quaternion $\mathbf{q}$ is a unit quaternion then it is also a rotation quaternion and represents a rotation in three-dimensional space. The unit quaternion that rotates points in an object through an angle θ, around a unit vector $[a,b,c]$ is defined as:

$$\mathbf{q} = (\cos(\theta/2), \sin(\theta/2)[a,b,c])$$

and thus

$w = \cos(\theta/2),\ x = (a)\sin(\theta/2)\ \ y = (b)\sin(\theta/2),\ \ z = (c)\sin(\theta/2)$

In four-dimensional space, the point (x,y,z) is a point on the unit sphere called the *pole of the rotation* [20], the unit vector from the origin to this point is the axis of rotation and w is the amount of rotation. All rotations are performed counter-clockwise around the *pole of the rotation* from the perspective of looking down on the point (x,y,z) from outside the unit sphere.

Cayley described [21] using the vector $\mathbf{p}$ to represent a point with the coordinates (x,y,z) treated as the pure quaternion $\mathbf{P}$, $(0, \mathbf{p})$, we can compute the desired rotation of $\mathbf{P}$, $\mathbf{P}'$ as follows:

$$\mathbf{P}' = \mathbf{q}\mathbf{P}\mathbf{q}^{-1}$$

The result will always be a quaternion with zero scalar component, $(0, \mathbf{p}')$.

And the rotation back to the initial point is:

$$\mathbf{P} = \mathbf{q}^{-1}\mathbf{P}'\mathbf{q}$$

Following Shoemake's derivation [22] we can obtain the general rotation matrix for quaternion rotation defined as:

$$\mathbf{Q} = \begin{bmatrix} 1-2(y^2+z^2) & 2(xy-wz) & 2(xz+wy) & 0 \\ 2(xy+wz) & 1-2(x^2+z^2) & 2(yz-wx) & 0 \\ 2(xz-wy) & 2(yz+wx) & 1-2(x^2+y^2) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Thus, for any point $\mathbf{p}$, $\mathbf{p}' = \mathbf{Q}\mathbf{p}$ for any unit quaternion.

BIBLIOGRAPHY

[1] D. Hill, P. Batchelor, M. Holden and D. Hawkes. Medical Image Registration. *Physics in Medicine and Biology*, 46:R1-R45, 2001.

[2] W. Wells, P. Viola, H. Atsumi, S. Nakajima and R. Kikinis. Multi-Modal Volume Registration by Maximization of Mutual Information. *Medial Image Analysis*, 1:35-51, 1996.

[3] P. Viola. Alignment by Maximization of Mutual Information. *PhD Thesis*, Massachusetts Institute of Technology, 1995.

[4] V. Mandava, J. Fitzpatrick, C. Maurer, Jr., R. Maciunas, G. Allen. Registration of Multimodal Volume Head Images via Attached Markers. *Medical Imaging VI: Image Processing*, Proc. SPIE 1652:271-282, 1992.

[5] WWW URL. http://www.cs.montana.edu/~bnct. INEEL, MSU Boron Neutron Capture Therapy home page.

[6] D. Rueckert, M. Clarkson, D. Hill, D. Hawkes. Non-rigid registration using higher-order mutual information. *Medical Imaging 2000: Image Processing*, Proc. SPIE 3979:438-447, 2000.

[7] J. West, J. Fitzpatrick, M. Wang, B. Dawant, C. Maurer, Jr., R. Kessler, R. Maciunas, C. Barillot, D. Lemoine, A. Collignon, F. Maes, P. Suetens, D. Vandermeulen, P. van den Elsen, S. Napel, T. Sumanaweera, B. Harkness, P. Hemler, D. Hill, D. Hawkes, C. Studholme, J. Maintz, M. Viergever, G. Malandain, X. Pennec, M. Noz, G. Maguire, Jr., M. Pollack, C. Pelizzari, R. Robb, D. Hanson, R. Woods. Comparison and Evaluation of Retrospective Intermodality Brain Image Registration Techniques. *Journal of Computer Assisted Tomography*, 21(4):554-566, 1997.

[8] C. Studholme, D. Hill, D. Hawkes. Automated 3D Registration of MR and CT Images of the Head. *Medical Image Analysis*, 1:163-175, 1996.

[9] C. Studholme, D. Hill, D. Hawkes. Automated 3D Registration of MR and PET Brain Images by Multi-resolution Optimization of Voxel Similarity Measures. *Medical Physics*, 24:25-35, 1997.

[10] J. Pluim, J, Maintz, M. Viergever. Interpolation artifacts in mutual information based image registration. *Computer Vision Image Understanding*, 77:211-232, 2000.

[11] C. Shannon. A mathematical theory of communication (parts 1 and 2). *Bell Systems Technical Journal*, 27:379-423 and 623-656, 1948. (reprint available: http://cm.bell-labs.com/cm/ms/what/shannonday/shannon1948.pdf)

[12] R. Duda, P. Hart. *Pattern Classification and Scene Analysis*. John Wiley and Sons, 1973.

[13] S. Theodoridis, K. Koutroumbas. *Pattern Recognition*. Academic Press, 1999.

[14] D. Hearn, M. Baker. Computer Graphics. Prentice Hall, 2nd edition, 1997.

[15] W. Press, B. Flannery, S. Teukolsky, W. Vetterling. Numerical Recipes in C: The Art of Scientific Computing. Cambridge University Press, 2nd edition, 1999.

[16] The National Library of Medicine's Visible Human Project. http://www.nlm.nih.gov/research/visible/visible_human.html.

[17] W. Hamilton. *Elements of Quaternions*. Chelsea, 3rd Edition, 1969.

[18] I. Kantor, A. Solodovnikov. *Hypercomplex Numbers*. Springer-Verlag, 1989.

[19] A. Cayley. On certain results relating to quaternions. *Philosophical Magazine*, 26:141-145, 1845. (Reprint available: *The collected mathematical papers*. Cambridge, 1963.)

[20] H. Smith. *Quaternions for the masses*. http://pw1.netcom.com/~hjsmith/Quatdoc/page01.html.

[21] A. Cayley. On the application of quaternions to the theory of rotation. *Philosophical Magazine*, 33:196-200, 1848. (Reprint available: *The collected mathematical papers*. Cambridge, 1963.)

[22] K. Shoemake. *Quaternions*. ftp://ftp.cis.upenn.edu/pub/graphics/shoemake/quatut.ps.Z

MONTANA STATE UNIVERSITY - BOZEMAN
3 1762 10353666 8