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

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

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

LIST OF FIGURES

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.

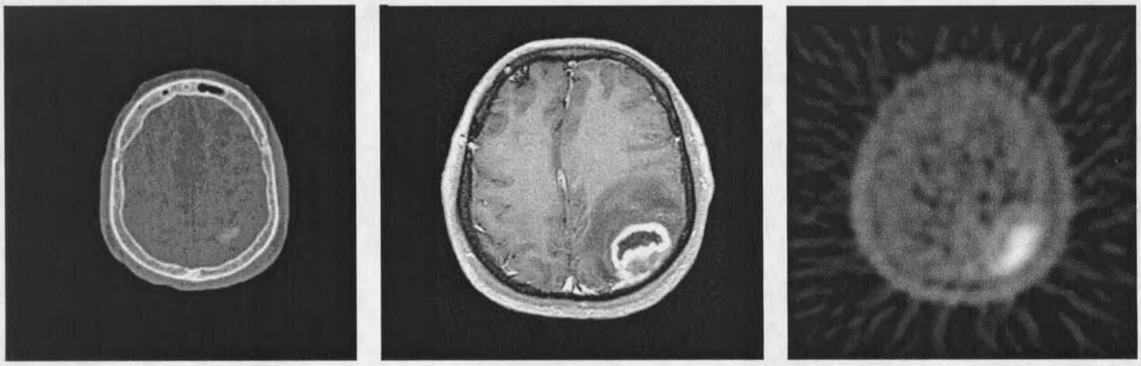

(a) (b) (c)

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).

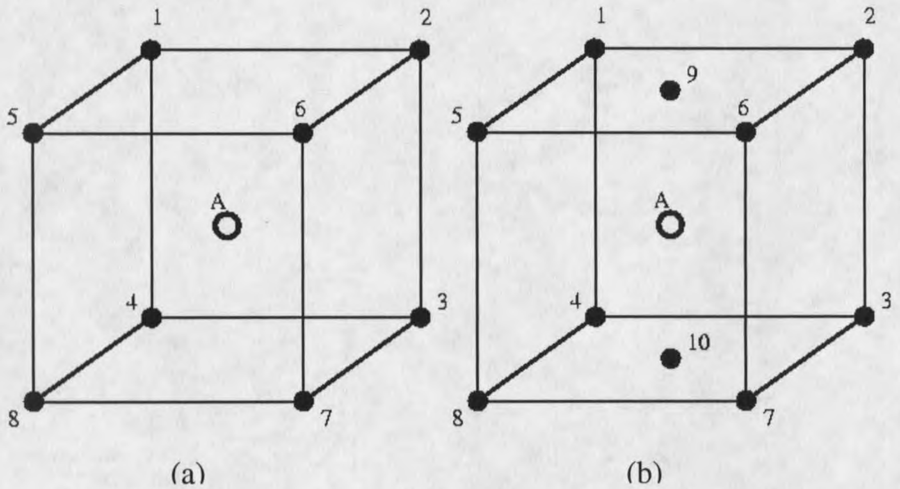

Figure 2. Methods of interpolation (a) tri-linear interpolation (b) modified tri-linear interpolation

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.

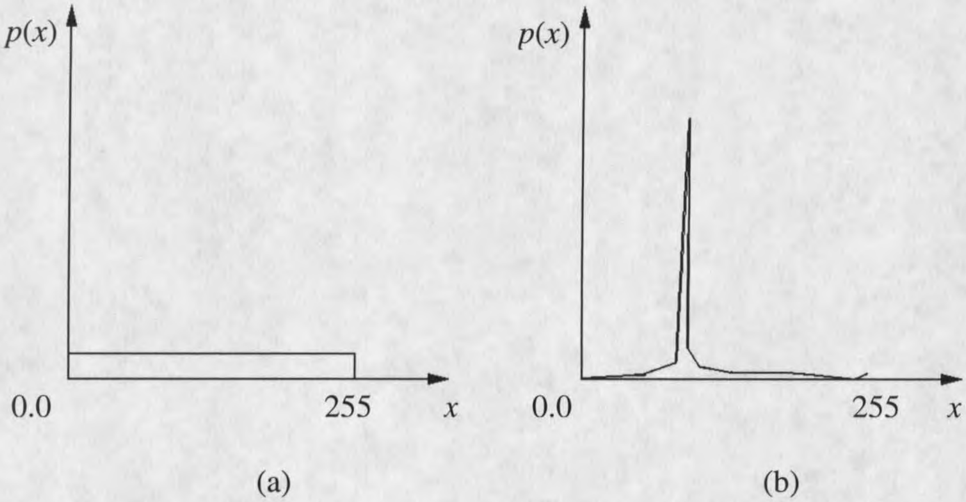

Figure 3. Entropy densities (a) high (b) low

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]:

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

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.

(a)

y \ *x*				
	2	2	3	2
	1	0	3	3
	1	2	2	3
	0	0	1	2

(b)

y \ *x*				
	1	2	3	2
	0	1	2	3
	2	2	3	2
	0	1	1	1

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.

image 1 \ *image 2*	0	1	2	3
3	0	0	2	2
2	0	2	3	1
1	1	1	1	0
0	1	2	0	0

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:

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

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*