

**THE REGISTRATION OF MULTIMODALITY
MEDICAL SCANS**

By

Torre Dana Zuk

B. Sc. (Computer Science) University of Alberta

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF

THE REQUIREMENTS FOR THE DEGREE OF

MASTERS OF SCIENCE

in

THE FACULTY OF GRADUATE STUDIES

COMPUTER SCIENCE

We accept this thesis as conforming
to the required standard

.....
.....
.....



THE UNIVERSITY OF BRITISH COLUMBIA

OCTOBER 1993

© TORRE DANA ZUK, 1994

In presenting this thesis in partial fulfilment of the requirements for an advanced degree at the University of British Columbia, I agree that the Library shall make it freely available for reference and study. I further agree that permission for extensive copying of this thesis for scholarly purposes may be granted by the head of my department or by his or her representatives. It is understood that copying or publication of this thesis for financial gain shall not be allowed without my written permission.

Computer Science
The University of British Columbia
2075 Wesbrook Place
Vancouver, Canada
V6T 1Z1

Date:

Abstract

The registration, or alignment, of multiple medical scans can provide a plethora of new information about a patient beyond what can be gained from studying the individual scans in isolation. This added information can increase the accuracy in specifying regions of interest, help in radiation therapy planning, and aid in a myriad of other treatment and diagnostic techniques. However, the registration of the scans is very complex, and no panacea exists for adjusting for the numerous variations and distortions that can occur between scans.

This thesis briefly summarizes the current technology of MRI, CT, PET, and SPECT imaging. Possible uses and benefits of using registered data sets are then explained in this context. This is followed by a structured description of the general problem of data registration and a framework that can be used to solve it. Current methods that have been utilized by various researchers are surveyed. The implementation of a surface matching method and the adaptations and innovations that were utilized is described in detail. These are analyzed under diverse test conditions involving both real and simulated data. The surface matching methods that were developed provide an accurate means for the registration of rigid body transformations, although many problems remain to be solved before a robust solution exists.

Table of Contents

Abstract	ii
List of Tables	viii
List of Figures	ix
Acknowledgment	xii
1 An Introduction to Registration	1
1.1 Reader's Guide	1
1.2 Variations Between Scans	2
1.2.1 Patient Positioning	2
1.2.2 Sampling	3
1.2.3 Noise	3
1.2.4 Imaging Configuration	3
1.2.5 Time-series Studies	3
1.2.6 Different Modalities	4
1.3 The Goals of Registration	4
1.3.1 Increasing the Precision of Regions of Interest	4
1.3.2 Radiation Treatment Planning	4
1.3.3 Localization of Functional Data	5
1.3.4 Accurate Comparison of Regions Over Time	5
1.3.5 Comparison of Patient Data to Statistical Norms	5
1.3.6 Removal of Distortions	6
1.4 Method of Registration	6

1.4.1	Segmentation of the Data	6
1.4.2	Matching of the Features	7
1.4.3	Verification of the Solution	7
1.5	Overview of the Thesis	8
2	Imaging Physics and Technology	10
2.1	Electromagnetics	10
2.2	Volumetric Representations	11
2.3	CT	13
2.3.1	The Physical Nature of CT	13
2.3.2	Current Technology and Trends	14
2.4	MRI	15
2.4.1	The Physical Nature of MRI	15
2.4.2	Current Technology and Trends	17
2.5	SPECT	17
2.5.1	The Physical Nature of SPECT	18
2.5.2	Current Technology and Trends	18
2.6	PET	19
2.6.1	The Physical Nature of PET	19
2.6.2	Current Technology and Trends	20
3	The Theory of Registration	21
3.1	Definitions	21
3.2	Types of Transformations	23
3.2.1	Rigid Body Motion	23
3.2.2	Affine	25
3.2.3	Projective	25
3.2.4	Non-linear	25

3.3	Variations in Data	26
3.3.1	Corrected Distortions	27
3.3.2	Uncorrected Distortions	28
3.3.3	Variations of Interest	30
3.3.4	Summary of Variations	30
3.4	Conditions for Finding a Solution	31
4	The Registration Process	32
4.1	Segmentation	32
4.1.1	Feature Types	33
4.1.2	Choice for Features	34
4.1.3	Techniques	36
4.2	Matching	37
4.2.1	Features	38
4.2.2	Comparison Function	39
4.2.3	Heuristics	43
4.3	Verification	46
4.3.1	Comparison Function Residual	46
4.3.2	Simulated Data	46
4.3.3	Patient Data	48
4.4	Case Studies	51
4.4.1	Iterative Manual Alignment	51
4.4.2	Solving Equations Formed From Corresponding Points	52
4.4.3	Iterative Surface Alignment	53
5	Implementations and Innovations	55
5.1	Platform	55
5.2	Direct Point Matching	57

5.2.1	Segmentation	57
5.2.2	Matching	58
5.2.3	Verification	59
5.3	Iterative Surface Matching	60
5.3.1	Segmentation	60
5.3.2	Matching	62
5.3.3	Verification	71
6	Analysis of Surface Matching	76
6.1	Accuracy of Segmentation	76
6.1.1	Types of Segmentation Error	77
6.1.2	Segmentation and Transformations	78
6.2	Theoretical Analysis	79
6.2.1	Comparison Function	79
6.2.2	Bounds of the Parameter Space	83
6.2.3	The Nature of the Search Space	83
6.3	Simulations: Transformed Real Features	87
6.3.1	Methods	88
6.3.2	Thresholds	90
6.3.3	Hierarchies of Scale and Thresholding	92
6.3.4	Distribution of Starting Points	95
6.3.5	Multiple Searches	96
6.4	Simulations: Transforming Real Data	97
6.5	Patient Data Studies	100
6.5.1	Registering to Norms	100
6.5.2	MRI-MRI Registrations	101
7	Conclusions and Future Work	107

7.1	Experimental Conclusions	107
7.2	Extensions and Future Development	108
7.3	Applications of the Implemented Techniques	109
7.4	Summary	110
	Bibliography	111
	Glossary	119

List of Tables

2.1	Examples of SPECT Radionuclides	18
2.2	Examples of PET Radionuclides	20
4.3	A List of Feature Types	35
4.4	Comparison of the Choices for Intrinsic Features	36
4.5	Features and Appropriate Comparison Functions	40
4.6	Overview of Pelizzari et al.'s Results	54
6.7	Distance Transform Errors	80
6.8	Errors in an Anisotropic Euclidean Chamfer Distance Transform	82
6.9	Comparison of Search Methods with a Noise Free Search Space	90
6.10	Comparison of Distance Transform Thresholds on Matching	92
6.11	Comparison of Thresholding in Conjunction with Multiple Scales	93
6.12	Results of Using Multiple Searches	96
6.13	Simulation Results with Transformed Real Data	100
6.14	Analysis of Interpolation for Reducing the Effects of Anisotropy	103

List of Figures

1.1	The Registration Process	7
2.1	The Electromagnetic Spectrum	11
2.2	Sample Types	12
2.3	Basic Forms of the 3D Scans	12
2.4	A CT Slice	14
2.5	An MRI Slice	16
2.6	A SPECT Slice	18
2.7	A PET Slice	19
3.1	The Default Labeling of Co-ordinate Axis	22
3.2	An Example of Registration (Rigid Body)	22
3.3	Categories of Transformation Elasticity	23
4.1	A Taxonomy of Feature Types	34
4.2	The Formation of a 2D Distance Transformation	42
4.3	A Comparison of Features	42
5.1	A WiT Igraph for Registration	57
5.2	PET/SPECT Segmentation	62
5.3	PET Data	63
5.4	PET Segmentation	63
5.5	MRI Data	64
5.6	MRI Segmentation	64
5.7	Actual Distance	67

5.8	DT Function	67
5.9	DT with Threshold at 10	68
5.10	DT with Threshold at 2	68
5.11	The “Colorwashing” of an Image	73
5.12	A WiT Igraph for “Colormerging”	74
5.13	A T1 Weighted MRI Scan	74
5.14	A T2 Weighted MRI Scan	74
5.15	Alternate Scan Line Merging of T1 and T2 Weighted MRI	75
6.1	Minor Error in Different Segmentations of a PET Scan	78
6.2	A 3D Chamfer Euclidean Distance Transform (101x101x101)	81
6.3	Comparison Function Over X Rotations Only ($c(F_s, f(rot_x, F_d))$)	85
6.4	$c(F_s, f(rot_y, F_d))$	86
6.5	$c(F_s, f(rot_z, F_d))$	86
6.6	$c(F_s, f(trans_x, F_d))$	87
6.7	$c(F_s, f(trans_y, F_d))$	87
6.8	Comparison Function Over Z Translation Only ($c(F_s, f(trans_z, F_d))$)	88
6.9	A Floating Point Anisotropic Chamfer Distance Transform	89
6.10	The Effect of Thresholding with Minor Error	91
6.11	Accuracy of Method A	94
6.12	Accuracy of Method B	94
6.13	Accuracy of Method C	95
6.14	Accuracy of Method D	95
6.15	Histogram of Actual Error Testing Method B ($\bar{d}=2.05, \sigma_d=1.51$)	96
6.16	Histogram of the Comparison Function Testing Method B	97
6.17	Search Sensitivity to Initial Parameters	98
6.18	A WiT Igraph for the Registration of Simulated Data	99
6.19	A Simulated Scan	101

6.20	Segmented (Dynamic) Surface	101
6.21	Result of Registration	102
6.22	Difference of Static and Registered Sets	102
6.23	PET Scan “X”	103
6.24	“X” Registered to Normal	103
6.25	PET Scan “Y”	104
6.26	“Y” Registered to Normal	104
6.27	T1 Weighted MRI Data	105
6.28	T1 Weighted MRI Segmentation	105
6.29	T2 Weighted MRI Data	105
6.30	T2 Weighted MRI Segmentation	106

Acknowledgment

The following people and groups should be acknowledged for their help in the direction and/or development of this thesis and the research behind it. Their contribution was greatly appreciated.

I would like to thank: my parents for their continuing support throughout, Brian Johnston (UBC¹) for his help with various aspects of the thesis, Dr. Stella Atkins (SFU²) and Dr. Kellogg Booth (UBC) for supervision and help in making revisions, Dr. Alain Fournier (UBC) for supplying constructive criticism, Terry Arden (Logical Vision Ltd.) for adapting the software package WiT to help in our research, Dr. Anna Celler (Dept. of Nuclear Medicine, VGH³) and Dr. Ken Poskitt, M.D. (B.C.'s Children's Hospital) for providing practical information, The PET Group (TRIUMF⁴) for the informative seminars that provided much needed background knowledge, NSERC⁵ for their Post Graduate Scholarship and funding under grant #8020, and Larry Palazzi (UBC) for his help in the layout and formatting of this thesis.

I would also like to thank all my friends and colleagues here at UBC for the assistance they have provided and for making my time here so enjoyable.

¹University of British Columbia

²Simon Fraser University

³Vancouver General Hospital

⁴Tri-University Meson Facility

⁵Natural Sciences and Engineering Research Council of Canada

Chapter 1

An Introduction to Registration

Human knowledge and human power meet in one; for where the cause is not known the effect cannot be produced. Nature to be commanded must be obeyed; and that which in contemplation is as the cause is in operation as the rule.

- F. Bacon, *Novum Organum*, 1620, First Book of Aphorisms, III

Registration is the process of finding the condition for correct alignment or proper relative positioning of one medical scan with reference to another. In this context the alignment is of one imaged object relative to another imaged objects' reference frame. Any two medical scans may be registered to compare and contrast the unique information present in each. This correlation can provide important insights into both data sets.

Registration requires the removal of certain variations, while ignoring others. When done correctly it can facilitate the process of diagnosis and qualitative analysis. The process can be structured into the phases of segmentation, matching, and verification. This chapter will clarify the what, why, and how of registration and then provide an overview of the rest of the thesis.

1.1 Reader's Guide

Due to the interdisciplinary nature of this thesis it is impossible to present all the relevant material. Therefore the reader is encouraged to read the provided references in order to gain a deeper understanding of the concepts. For this reason when multiple articles are referenced on a single topic they are listed from general to specific. Words that are marked with a “†” are defined in the glossary at the end of the thesis. Chapters 2, 3, and 4 are intended to provide the general background necessary for understanding the specific implementations and analyses in Chapters 5 and 6. A reader unfamiliar with medical imaging is advised to read Chapter 2 first

and then return to Chapter 1, while those with the background knowledge may skip Chapter 2 with no loss of continuity.

1.2 Variations Between Scans

There can be many types of variations between the data sets. The registration process must be able to handle each one in an acceptable manner. Voxels[†] in different modalities that represent corresponding physical spaces may have no direct functional relation to each other. Even two scans of a single modality will never be identical. Within the same scan voxel values will be dependent on the patient's position and the configuration parameters of the individual scanning machine. In almost every case the two scans to be registered are not identical and may differ as a result of a variety of known or unknown factors.

1.2.1 Patient Positioning

The position of the patient in the imaging device is the major variation for which registration will try to compensate. During successive scans the patient may be situated at a different orientation and location in the machine, or the patient may be in a different machine altogether.

It is possible to fix the patient into a common 3D reference frame for each scan, but this usually requires a device be physically attached to the patient. This may be traumatic, especially if the head is clamped into a non-moving position for an extended period of time. Another more flexible way of establishing the patients' reference frame is through the use of fiducial markers[†]. They are limited by the accuracy of their placement, they can be costly, and they cannot be introduced into a scan retrospectively. In another technique, by Saeed and Durrani [77], a previous MRI scan is used to realign the patient into the same position during a preliminary scanning procedure.

1.2.2 Sampling

The voxel values calculated by the imaging device are representative of a discrete volume. The exact spatial domains of these voxels are almost never duplicated from scan to scan due to differences in sampling rates and voxel dimensions. The exact spatial resolution will even vary with position in a single scan.

1.2.3 Noise

Every scan will be subject to a variable amount of noise. The noise may be due to the physics of the imaging technique, external sources, or patient movement, all of which can result in noise and artifacts[†] of various frequencies and amplitudes. Often this noise may be reduced by filtering or similar techniques, but it cannot be eliminated entirely.

1.2.4 Imaging Configuration

The specific imaging device parameters in effect during a scan can introduce large variations in corresponding voxels of the same modality. For example, an MRI can be specially configured to precisely image the brain's ventricles[†], but then may not provide good detail of other structures. With SPECT and PET the dose of the radiopharmaceutical is yet another variable that affects the scans.

1.2.5 Time-series Studies

In many cases a single patient will require multiple scans over varied periods of time. The dynamics that occur may then be structural and/or physiological depending on treatments, disease progression, trauma, or other changes. One example is the response of a tumour to radiotherapy. The arbitrary variations of a region of interest (ROI) in each scan may create a situation that must be accommodated on a case-by-case basis.

1.2.6 Different Modalities

The different imaging modalities have no intrinsic correspondence between voxels that represent the same physical space. The physical property being measured usually has little relation to that used in the other modalities. The only common features are often the boundaries between internal structures which manifest themselves as gradients.

1.3 The Goals of Registration

Numerous insights that can be attained from registered data sets are pushing its development. The diverse advantages provided by registration lie in the increased ability to quantify regions and monitor changes. Multimodality registration of functional[†] and structural[†] data provides a more complete understanding of the patient, which is essential for an optimal diagnosis.

1.3.1 Increasing the Precision of Regions of Interest

The complementary information provided in the different imaging modalities will often allow a more accurate delineation of an arbitrary ROI. The boundary of a tumour may be more clearly observed in CT used in conjunction with a contrasting agent than in MRI [55], while MRI may better contrast tumour and edema to gray and white matter [15]. In a limited study by Hill [34] in which two surgeons and a radiologist assessed registered CT and MRI from seven patients, it was shown that the combined data provided significant new information clarifying the ROIs in each case.

1.3.2 Radiation Treatment Planning

Radiation treatment involves the irradiation of a specific ROI with high energy photons (X-rays). A CT volume is required for radiation treatment planning to provide the electron density function for calculation of the photon beam specifications [48]. An MRI scan may be used to more accurately segment the normal critical structures and identify the treatment volume; these

can then be transferred to the CT after registration [48]. Similarly PET or SPECT ROIs can be transferred to CT for beam attenuation calculations.

1.3.3 Localization of Functional Data

The registration of functional (SPECT or PET) and anatomical[†] data (CT or MRI) allows direct correlation between the structure and the observed functionality. The comparison of the low resolution functional data directly to corresponding structures allows for a more accurate diagnosis, especially in patients with large anatomic variations (congenital[†] or disease related) [73, 47, 19]. The localization[†] of the functionality can also guide the surgeon in deciding on a biopsy site [38].

1.3.4 Accurate Comparison of Regions Over Time

For the analysis of a treatment's effect on a certain patient, it is important to spatially correlate the changes within a ROI [68]. One example is that a tumour's size and position will be studied over time to see how a patient is responding to radiation therapy. The size of the tumour can be calculated in each volume of the time series, but this does not provide an accurate spatial-dependent response. With registered data sets ROIs can be superimposed over each other for direct comparison. This increases the precision of treatment monitoring [55].

1.3.5 Comparison of Patient Data to Statistical Norms

A new idea which has not yet gained widespread use is the direct comparison of the patients' data to accumulated statistical norms. By registering patient data to a uniform average data set, abnormalities may be more easily assessed. This registration is different from the other types in that the intra-patient variability must be accommodated by the registration process, enabling a meaningful comparison of the patient's data to the norms.

1.3.6 Removal of Distortions

The registration of a distorted data set to one free of the distortion may allow some distortions to be eliminated. MRI, which suffers from many non-linearities, could be registered to CT to remove them, or a SPECT scan could be reconstructed[†] with fewer artifacts if registered to a CT scan in order to utilize attenuation[†] values for improved scatter correction.

1.4 Method of Registration

The goal of registration is to find a transformation that enables an ROI to be compared directly and accurately to the same region in another scan. The process of reaching this goal can be broken down into three phases:

1. Segmentation - features that will be compared are selected
2. Matching - calculation of the registering transformation
3. Verification - a visual or numerical check on the accuracy of the alignment

The general flow of the registration process through these three phases is shown in Figure 1.1.

1.4.1 Segmentation of the Data

Segmentation is usually employed to classify data for analysis, but registration requires segmentation to isolate features that exist in both data sets for matching purposes. This allows the data to be matched in a more controlled and precise way. The user's actions and knowledge required for segmentation vary greatly, from completely manual methods, which are time consuming and work intensive, all the way up to fully automatic ones, which require no operator control at all.

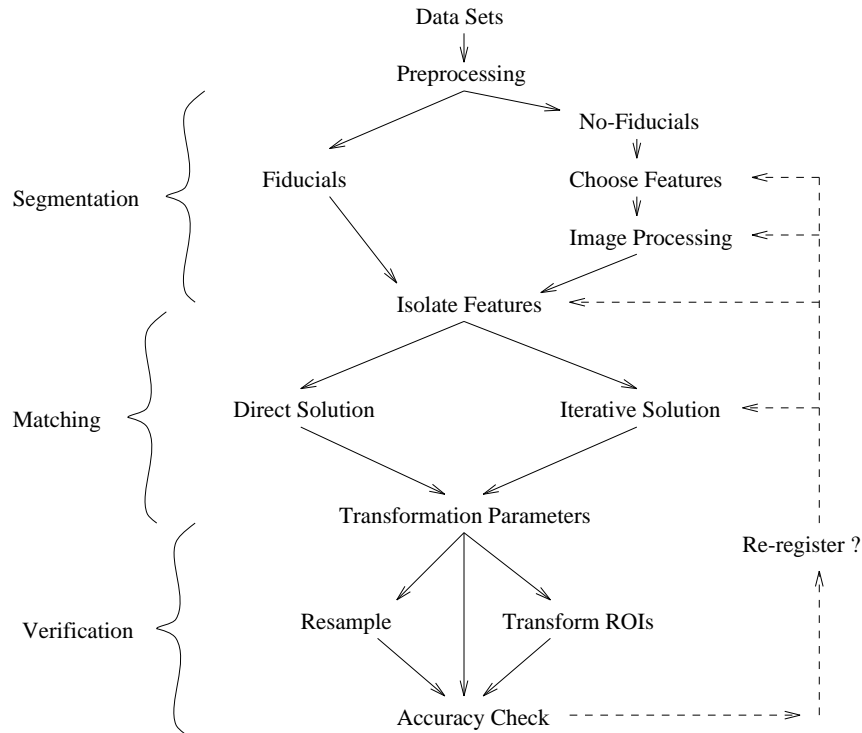


Figure 1.1: The Registration Process

1.4.2 Matching of the Features

Matching is the process of discovering what transformation can align the segmented features, and thus the entire scans. Initially a comparison function must be chosen for matching the two sets of features, and it will be highly dependent on the types of features chosen for segmentation. Based on the function, a heuristic for finding the transformation will attempt to determine the optimal match. Numerous different heuristics exist for solving for the transformation directly or iteratively converging to a solution. The large number of variations that can occur make this a very difficult part of the process.

1.4.3 Verification of the Solution

The validity of a solution can be checked by using an appropriate display technique to inspect the match visually, although this may require a radiologist's skills. Another way to check

the solution can be the proximity of the comparison function to the “optimal” value; this theoretical verification is often incorporated into iterative heuristics as a stopping criteria. In real applications the exact accuracy of the registration is unknown, as all of the variations cannot be removed.

1.5 Overview of the Thesis

The next chapter provides a brief summary of the background material needed to place subsequent chapters in the proper context. It touches on the physics utilized for the creation of the three dimensional data sets, and introduces much of the vocabulary that will be assumed later. A reader familiar with the physics of CT, MRI, PET, and SPECT may skip this chapter with no loss of continuity.

Chapter 3 provides the theoretical foundation for an accurate formulation of the problems of registration. It provides a categorization of the possible registering transformations, as well as the types of distortions that exist within and across modalities. Preconditions for an exact solution are discussed. It is not meant to be an in-depth analysis, but contains references to more thorough treatments of the material within the literature.

A detailed description of the method of registration is provided in Chapter 4. This contains both theoretical and application-driven aspects of the process, and covers the implementation details from the general to the specific. The basic constructs for segmentation, matching, and verification are detailed in this chapter, as well as a brief discussion of three registration implementations typical of the major genres (manual, point mapping, and surface matching).

Chapter 5 describes the implementations that form the main thrust of the thesis. The registration process was initially based on point features, and then modified to use surfaces. This required the construction of segmentation, matching, and verification tools, all of which were constructed in a data flow environment for image processing. A complete set of data flow tools for the registration of multi-modality data was created.

An investigative analysis of the implemented ideas is summarized in Chapter 6. A formal

declaration of the registration parameters, and the search space is given in order to clarify the results. Simulated and real data were utilized in this process to quantify the theoretical and practical performance of the various implemented techniques.

The final chapter presents conclusions and provides the direction that future developmental and analytic work may take. The algorithms developed provided good results, and many of the design decisions are in agreement with the recent Computer Vision in Radiology (COVIRA¹) registration project [2, 1]. The testing of the applications by research groups will further exercise the implementations, and provide useful knowledge for future performance improvements.

¹Computer Vision in Radiology, an Advanced Informatics in Medicine (AIM) Programme of the European Communities Commission

Chapter 2

Imaging Physics and Technology

Now to descend to the particular tenets of vain philosophy derived to the Universities, and thence into the church, partly from Aristotle, partly from blindness of understanding; I shall first consider their principles. There is a certain *Philosophia Prima*, on which all other philosophy ought to depend; and consisteth principally, in right limiting of the significations of such appellations, or names, as are of all others the most universal: which limitations serve to avoid ambiguity, and equivocation in reasoning; and are commonly called definitions; such as are the definitions of body, time, place, matter, form, essence, subject, substance, accident, power, act, finite, infinite, quantity, quality, motion, action, passion, and divers others, necessary to the explaining of a man's conceptions concerning the nature and generation of bodies.

- T. Hobbes, *Leviathan*, 1651, Part IV, "The Kingdom of Darkness", Chapter 46

Each modality reveals various unique physical properties of the regions it scans. In order to measure and localize these diverse qualities scanning devices utilize a wide range of electromagnetic (EM) radiation. Different modalities sample volumetric regions in distinctive ways and amounts. This chapter introduces the physical principles underlying each modality and briefly describes the technology that implements it.

2.1 Electromagnetics

CT, MRI, SPECT, and PET are non-invasive[†] methods. Internal information is carried externally for measurement by EM waves. This look into the inner nature of the body is dependent on accurately recording and quantifying these EM signals. The position in the EM spectrum of the waves used by the different modalities is shown in Figure 2.1 [30].

The EM-radiation carries energy (the signal) in a wave of alternating electric and magnetic fields whose energy is expressed in Joules by the formula:

$$E = h\nu$$




Figure 2.1: The Electromagnetic Spectrum

where E is the energy, h is Planck's constant $\approx 6.63 \times 10^{-34} J \cdot s$, and ν is the frequency. Therefore the relative energy of the EM used for each modality is shown as increasing from left to right in Figure 2.1.

2.2 Volumetric Representations

As opposed to traditional images that represent 2D data (pixels), the scanning methods of MRI, CT, SPECT, and PET create 3D data (voxels). This data is often viewed as a two-dimensional plane slicing through the volume, but the elements still represent discrete three-dimensional samples.

Samples can be taken from an arbitrary variety of different configurations which then form the voxels of the data set. The samples represent a function of the actual data over their spatial domain. With adjacent voxels the sampling functions commonly overlap. This can be seen in Figure 2.2. The distance between the centres of the functions (d) represent the inter-voxel distances and may vary in all three dimensions. The rectangular voxel is a gross abstraction and commonly represents only a tessellation of the sampling space.

The most prevalent data format contains planes (slices) of approximately adjacent voxels (Case 2) with gaps (Case 3) between the "slices". This is depicted as type A in Figure 2.3, with

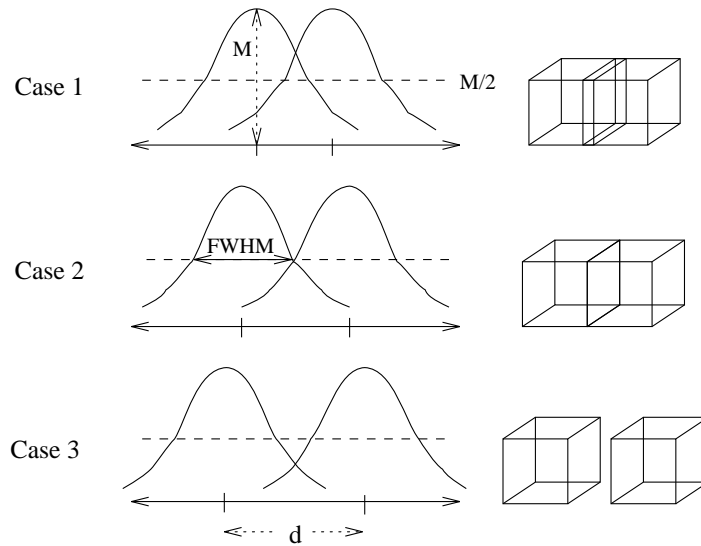


Figure 2.2: Sample Types

type B being contiguous 3D samples, which could possibly be formed by resampling type A. References to data resolution ($|V_x| \times |V_y| \times |V_z|$) indicate first the number of planar voxels (interslice) and then the number of slices (4x2x3 for Figure 2.3). The effective spatial resolution, or volume a voxel represents, is more complex and it is stated as the full width of the point spread function[†] at half its maximum value (FWHM).




Figure 2.3: Basic Forms of the 3D Scans

Due to different spatial resolutions the data is commonly resampled into cubic form (Figure 2.3.B) for the simplification of manipulation and visualization techniques. The resampling

process is based on a filter which determines the type of interpolation that is performed. The resampled data will not contain any information that was not inherent in the original data, but may contain less due to the approximations used for interpolation.

2.3 CT

Computer Tomography (CT) was the first 3D non-invasive medical imaging technique to be developed. By measuring the attenuation of X-rays through the body it provides the highest in-plane resolution of all the medical imaging techniques. It is the most widely used technique, and the least expensive. An example of the data created by CT is shown in Figure 2.4 which illustrates a single 512x512 axial[†] slice through the top of the brain. The smoothly curved structure around the bottom of the image is the external head holder used with the scanning device. Please note that all images of medical data included in this thesis have been inverted (black \Leftrightarrow white or 0 \Leftrightarrow 255) for better reproduction with the printer's greyscale gamut.

2.3.1 The Physical Nature of CT

CT utilizes an earlier imaging technology which provided 2D X-ray attenuation values. X-ray film records the amount of high energy electromagnetic radiation that passes through the exposed object. If numerous X-ray images are projected at different angles it becomes possible to solve for the volumetric (3D or voxel) representation of the attenuation values.

In CT imaging the 2D film is replaced by a 1D array of photon detectors. The X-ray beam is spread out in fan-like form for a single exposure. The source is then rotated in a circle so that numerous single X-rays are taken at different angles around the object. This rotation delineates the plane (slice) that is imaged. Numerous slices then can be acquired by repeating the process after moving the relative position of the camera.

The CT voxels represent a mathematically formed approximation of attenuation values for a discrete slice in 3-space. The slice is commonly computed (or reconstructed) from the radial samples by projecting lines back to the source, accumulating the counts for each voxel, and




Figure 2.4: A CT Slice

then filtering to account for artifacts. This technique is called “filtered back-projection” and it is the most popular. Another technique for reconstruction is the use of Fourier transforms according to the Fourier projection-slice theorem [57].

2.3.2 Current Technology and Trends

Current CT commonly provides an intra-slice resolution of 512×512 voxels while the slice resolution is dependent on the number of times the planar process is repeated. Each voxel commonly has around a 0.5mm in-plane spacing and inter-slice spacing of 1 to 10mm [91]. The number of slices varies greatly depending on what is to be imaged, but commonly ranges from 10 to 40.

The fine detail and relatively few distortions present in CT enables very detailed structures to be imaged. Contrasting agents may also be used to change the attenuation characteristics of regions of interest, thereby creating sharper contrast between different structures.

The latest development is the spiral CT. With scanner a 2D array of sensors is rotated in a spiral (translated while rotating) which allows the generation of arbitrarily positioned slices. Other current work is on the creation of novel contrasting agents. More detailed background descriptions are provided by Bates [8], Vannier [91], Lancaster [53], and Kreel [52].

2.4 MRI

Magnetic Resonance Imaging (MRI) has recently become a widely accepted imaging technique. The use of low-energy EM radiation provides a safety advantage over and above its imaging capabilities, but future research may discover that the strong magnetic fields that are utilized have some negative effects. It also provides superior flexibility in the many possible imaging configurations. A typical slice of MRI data (256x256) is shown in Figure 2.5. It represents a sagittal[†] slice through a head revealing a large tumor below the posterior corpus collosum[†].

2.4.1 The Physical Nature of MRI

MRI reveals the number of atoms with a chosen odd atomic number within a 3D sample. The physical basis for MRI is that the resonant frequency (Larmor frequency) of an atom in a magnetic field is dependent on the atomic nuclei and the strength of the applied magnetic field. Atoms are put in a strong magnetic field to align their magnetic moments[†] and then a selected EM pulse in the radio frequency (RF) range is used to excite them. The resonating atoms produce another RF signal and in this way the concentration of a selected atom may be measured. The nature of the resonance is in two planes, each of which can be used for the measurement of atom concentration. This provides scans labelled T1 or T2, depending on the relative weighting of the two planes. If the resonation is measured as a function of the time it takes the magnetization to return to 63% of its equilibrium value (before excitation) it is called



Figure 2.5: An MRI Slice

a T1 weighted scan. A T2 weighted scan is measured as time it takes for the signal to decay to 37% of its maximum value. The different weightings provide slightly different information because they represent different magnetic properties that vary with tissue type [3].

3D spatial encoding of the excited resonant frequencies is performed by the use of a magnetic gradient, and frequency/phase variations of the pulse. The signals created are in the frequency domain and so the slices are reconstructed by the use of an inverse Fourier transform[†]. The fact that the EM triggering pulses are highly controllable allows arbitrary oblique slices to be reconstructed.

2.4.2 Current Technology and Trends

Currently only the concentration of 1H (Hydrogen) is measured because it has the strongest magnetic moment and thus provides the strongest signal. This provides very high resolution data, commonly around $256 \times 256 \times 100$ with inter-voxel spacing of $1\text{mm} \times 1\text{mm} \times 2\text{mm}$. The use of different signal parameters as well as the T1 and T2 weightings allows a greater control over the imaging process than is possible with other modalities.

A result of the push for faster scanning times is the Fast Low Angle Shot (FLASH) imaging method, which can provide images in the time frame of a second rather than a minute, but at a cost of a lower signal to noise ratio (SNR). Future developments may make the measurement of ^{31}P (Potassium), ^{23}Na (Sodium), ^{19}F (Fluorine), and ^{13}C (Carbon) more common, allowing a greater amount of functional data retrieval. Further information can be found in articles by Jones [44], Hinshaw [35], Kreel [52], and Riederer [76].

2.5 SPECT

Single Photon Emission Computed Tomography (SPECT) provides a means to view the internal physiology of a patient. *In vivo* functional data may be measured by imaging specific metabolically related molecules. The molecules are made radioactive (labeled) and then their emissions can be observed externally. A single SPECT slice (64×64) is shown in Figure 2.6. It is taken from an axial Technetium series of the head.

2.5.1 The Physical Nature of SPECT

As opposed to CT, the source of EM radiation for SPECT is inside the body and the EM radiation is in a higher band of the spectrum. When radioactive source elements (radionuclides) are contained within specific molecules they form radiopharmaceuticals, which are injected into the body. These release radiation in the form of gamma rays (photons) which are recorded by external detectors. The direction of the incident photon is determined by using a device (called a collimator) that only allows perpendicular rays to enter the detector. Although there are many possible configurations of detectors, they typically form a 2D array, and so a single 360 degree rotation of the array around the body can enable a 3D volume reconstruction.




Figure 2.6: A SPECT Slice

The various radiopharmaceuticals that can be created add a new dimension to medical imaging. The level of emissions in a region reveals the uptake of the radioactive substance. Due to the use of high energy radioactivity, the process becomes dependent on the decay times of the radionuclides as well as the amount of relative uptake throughout the body.

2.5.2 Current Technology and Trends

The SPECT data is commonly of resolution $64 \times 64 \times 64$ with voxel spacing from 6mm to 15mm . They have low SNRs resulting in effective spatial resolutions of 9 to 18mm FWHM. Some current radionuclides are listed in Table 2.1 [21].

Radionuclide	Half-life (h)	Main Emission (keV)	Typical Use
^{99m}Tc (Technetium)	6.0	140	brain perfusion [†]
^{123}I (Iodine)	13.2	159	brain perfusion
^{111}In (Indium)	67.3	171, 245	inflammation
^{201}Tl (Thallium)	73.1	68-82	myocardial perfusion
^{133}Xe (Xenon)	126.0	81	lung scintigraphy [†]

Table 2.1: SPECT Radionuclides

The spatial resolution is limited mainly by the collimator. Future progress will provide higher resolution data, higher SNRs, and new radionuclides for monitoring various metabolic functions. For more information on SPECT, see discussions by Jaszczak [40], Peters [72], and Ell [21].

2.6 PET

Positron Emission Tomography (PET) records functional data in a similar manner to SPECT, but provides higher resolution images. PET is the most recent imaging technology and currently the most expensive to own and to operate. An illustration of a PET slice (100x100) is given in Figure 2.7, showing an axial slice of the head using a fluorine radionuclide.



Figure 2.7: A PET Slice

2.6.1 The Physical Nature of PET

PET, akin to SPECT, uses radiopharmaceuticals, but higher energy ones that emit positrons as they decay. The positrons travel a short distance until they hit an electron and then annihilate releasing two photons (gamma rays) in opposite directions. These pairs are recorded by detectors and subsequent computer processing determines the direction of the incident rays.

The counts of matching rays are then used to calculate slice data. Commonly a ring or rings of detectors are used to catch co-incident pairs in the plane(s) being imaged.

The use of higher energy radiation provides more emissions and thus more statistically accurate approximations can be made of the 3D distributions of the radionuclides. PET offers a more versatile range of radiopharmaceuticals than does SPECT. A typical example is ^{18}F labeled deoxy-glucose, which can reveal important details of sugar metabolism.

2.6.2 Current Technology and Trends

The resolution of PET is around $100 \times 100 \times 20$, with the planar (in-slice) voxel dimensions of approximately 5mm FWHM and slice thickness of 10mm FWHM. Examples of PET radionuclides are listed in Table 2.2 [40, 72].

Radionuclide	Half-life (m)	Emission (keV)	Typical Use
^{82}Rb (Rubidium)	1.3	511	perfusion
^{15}O (Oxygen)	2.0	511	tissue blood flow
^{13}N (Nitrogen)	10.0	511	oxygen utilization
^{11}C (Carbon)	20.0	511	drug labelling
^{68}Ga (Gallium)	68.0	511	toxoplasmosis [†] vs. lymphoma [†]
^{18}F (Fluorine)	111.0	511	glucose metabolism

Table 2.2: PET Radionuclides

Future research is increasing the scanning resolution and the range of radiopharmaceuticals. PET is inherently limited by the fact that a positron travels a variable distance before annihilating; therefore a spatial resolution of 2mm FWHM is a theoretical minimum [45]. For more information on PET see papers by Jaszczak [40], Peters [72], and Ell [21].

Chapter 3

The Theory of Registration

Philosophy is written in that vast book which stands forever open before our eyes, I mean the universe; but it cannot be read until we have learnt the language and become familiar with the characters in which it is written. It is written in mathematical language, and the letters are triangles, circles and other geometrical figures, without which means it is humanly impossible to understand a single word.

- G. Galileo, *Il Saggiatore*, Question 6

The underlying process of registration is the formation of a transformation to map one scan onto another. The type of transformation required depends on the variations between the two scans. By using an appropriate transformation that correctly models the variations present, it is possible to properly align the data.

The simplest registration transformation will require only translation and rotation to correct for patient positioning, and possibly scaling to correct for different sampling rates. Some modalities may contain other variations that must be counteracted for a successful registration. The ability to find a solution and accurately quantify the error can be related to how well an individual case satisfies certain conditions.

3.1 Definitions

Volumetric registration between two medical scans can be defined as a spatial mapping. Let p_{vi} specify a function, such that given a scan (V_i) and a position in that scan (x, y, z), it returns the 3D positions in the *real* or physical coordinate system defined by the “imaged object”. The registration is then the mapping between two scans V_1 and V_2 by a function f , given in equation form as,

$$p_{V_1}(x, y, z) = p_{V_2}(f(x, y, z)),$$

where f is a 3D spatial-coordinate transformation. Therefore $f = p_{V_2}^{-1}(x, y, z) \cdot p_{V_1}$. In layman's terms, it is the alignment of two scanned objects into the same physical space. The default labelling of the co-ordinate axis is provided in Figure 3.1. When the transformation maps set B onto set A, B will be called the dynamic set, and A the static set. An example registration is shown in Figure 3.2.

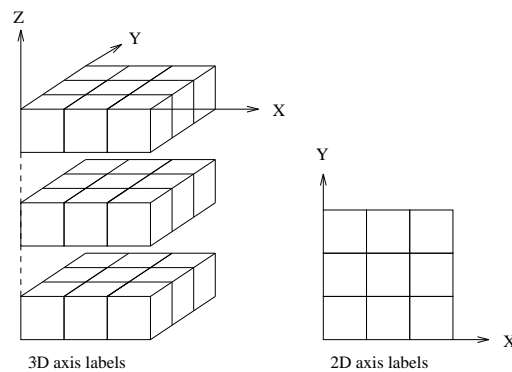


Figure 3.1: The Default Labeling of Co-ordinate Axis

An exact transformation function f will not be discovered for most cases. The types of variations between V_1 and V_2 will determine the complexity of the function f . When dealing with actual data subject to noise and other distortions, a good approximation for f is all that is feasible.




Figure 3.2: An Example of Registration (Rigid Body)

3.2 Types of Transformations

The transformation from one scan to another can be arbitrarily complex and can be characterized in a large number of ways. The transformation may be global and have a domain of the entire volume, or be local and affect only sub-regions of the volume. Global changes are commonly due to patient positioning, while local ones are usually more complex and can be caused by object deformation, joint movement, or other distortions. Local transformations may be handled by segmenting the data into smaller regions each of which can then be handled individually, or with complex non-linear transformations which can affect each region differently. A major problem with local transformations is the discontinuities that may appear between local regions.

Geometric transformations have been grouped into four basic categories of elasticity as rigid, affine, projective, and non-linear by van den Elsen et al. [88]. Each category's class of transformations is theoretically contained in the next, hence curved transformations form the superset. A visual example of a global transformation of each type is shown in Figure 3.3.




Figure 3.3: Categories of Transformation Elasticity

3.2.1 Rigid Body Motion

The most common and simplest transformation is rigid body motion, which Giancoli [30] defines as: "... By a *rigid body* we mean a body that has definite shape which doesn't change, so that the particles that make it up stay in fixed positions relative to each other.". This results in motion that can be analyzed by rotation and translation about the center of mass. While the rotation and translation can correct for different patient positioning, the addition of scaling

factors is necessary to account for the arbitrary sampling rates of voxels. Although an object may be manipulated about an arbitrary axis, the equivalent transformation may always be performed by three separate rotations about each of the x, y and z axis through the center of mass, and then a single translation in 3D. Baraff provides an indepth presentation [6] of the various aspects of rigid body motion.

Rigid body motion can be represented in matrix form as a composition of scaling, rotations about each axis, and finally translation, expressed in homogeneous coordinates [25]:

$$\begin{bmatrix} x_2 \\ y_2 \\ z_2 \\ 1 \end{bmatrix} = T R_z R_y R_x S \begin{bmatrix} x_1 \\ y_1 \\ z_1 \\ 1 \end{bmatrix}$$

where

$$T = \begin{bmatrix} 1 & 0 & 0 & t_x \\ 0 & 1 & 0 & t_y \\ 0 & 0 & 1 & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix}, R_z = \begin{bmatrix} \cos(\theta_z) & -\sin(\theta_z) & 0 & 0 \\ \sin(\theta_z) & \cos(\theta_z) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

$$R_y = \begin{bmatrix} \cos(\theta_y) & 0 & \sin(\theta_y) & 0 \\ 0 & 1 & 0 & 0 \\ -\sin(\theta_y) & 0 & \cos(\theta_y) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, R_x = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta_x) & -\sin(\theta_x) & 0 \\ 0 & \sin(\theta_x) & \cos(\theta_x) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

$$S = \begin{bmatrix} s_x & 0 & 0 & 0 \\ 0 & s_y & 0 & 0 \\ 0 & 0 & s_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

3.2.2 Affine

The next more general transformation is the affine, which allows for scaling and shearing in each dimension independently. It can be expressed in matrix form as:

$$\begin{bmatrix} x_2 \\ y_2 \\ z_2 \\ 1 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} & t_x \\ a_{21} & a_{22} & a_{23} & t_y \\ a_{31} & a_{32} & a_{33} & t_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ z_1 \\ 1 \end{bmatrix}$$

An affine transformation preserves straight lines and parallel lines remain parallel after transformation.

3.2.3 Projective

After a projective transformation parallel lines do not necessarily remain parallel. The expansion of the affine transformation to provide a projection (perspective) requires the full power of the homogeneous matrix representation. However the result will not be in regular 3-space, but this can be formed by dividing the x'_2, y'_2, z'_2 coordinates by the weighting factor w'_2 [25].

$$\begin{bmatrix} x'_2 \\ y'_2 \\ z'_2 \\ w'_2 \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \\ a_{41} & a_{42} & a_{43} & a_{44} \end{bmatrix} \begin{bmatrix} x_1 \\ y_1 \\ z_1 \\ w_1 = 1 \end{bmatrix}, \begin{bmatrix} x_2 \\ y_2 \\ z_2 \\ 1 \end{bmatrix} = \begin{bmatrix} x'_2/w'_2 \\ y'_2/w'_2 \\ z'_2/w'_2 \\ w'_2/w'_2 \end{bmatrix}$$

Projective transformations are commonly used to map between 2D X-ray images and volume data [89]. These types of projections are commonly used in graphics applications to create realistic 2D images which portray the 3D nature of the data.

3.2.4 Non-linear

Non-linear transformations allow the mapping of lines to curves, and do not necessarily preserve any lines. There are numerous types of functions that map to curves, but polynomial are perhaps

the simplest. An arbitrary polynomial transformation may be performed by using functions of the form:

$$\begin{aligned}
 x_{v2} &= \sum_{k=0}^n \sum_{j=0}^{(n-k)} \sum_{i=0}^{(n-k-j)} a_{ijk} x_{v1}^i y_{v1}^j z_{v1}^k \\
 y_{v2} &= \sum_{k=0}^n \sum_{j=0}^{(n-k)} \sum_{i=0}^{(n-k-j)} b_{ijk} x_{v1}^i y_{v1}^j z_{v1}^k \\
 z_{v2} &= \sum_{k=0}^n \sum_{j=0}^{(n-k)} \sum_{i=0}^{(n-k-j)} c_{ijk} x_{v1}^i y_{v1}^j z_{v1}^k
 \end{aligned}$$

Although more than an infinite number of other non-linear transformations exist, many are too complex for practical use.

Often with complex variations between scans a single transformation will not properly transform the entire data set. This necessitates the use of local transformations that may provide increased control. Thin-plate splines are a family that provide a means for piecing together local transformations without discontinuities.

3.3 Variations in Data

Registration requires the removal of or compensation for distortions between data sets, in order to facilitate an accurate comparison. Distortions can be defined as the variations in the data that cause misregistration [14]. The removal of variations may also provide insight into other differences between the two scans. All variations can be broadly grouped into those which are inherent within a single scan (intrinsic), and those that occur between different scans of the same or differing modality (extrinsic). Intrinsic variations are noise, physical distortions of the imaging technique and scanner calibration. Some extrinsic variations are patient positioning, physiological changes, and the differing physical nature of each modality.

Brown, in her survey of image registration techniques [14], describes three categories for the variations: corrected distortions, uncorrected distortions, and variations of interest. The definition of these variations and examples in their respective category or categories are provided

in the next three subsections. Many other corrected and uncorrected distortions exist due to the construction and associated physical limitations of the scanning machines.

3.3.1 Corrected Distortions

Corrected distortions are distortions that can be modelled. The model hopefully provides a means for correction of the distortion in the registration process.

1. All modalities

- different voxel sizes and sampling rates (intrinsic)

Multi-modality data sets will usually have different spatial resolutions and extents. In order to directly compare the data sets they must be mapped to a single viewing space. The scale transformation may be performed using *a priori* knowledge provided by the imaging device, or by including scale deformations in the transformation.

- body motion (intrinsic/extrinsic)

Rigid body motion can be corrected by the use of rotation, translation, and scaling to align the data. While the rigid body motion assumption often holds for the skull, it will not for many other regions that contain joints and soft tissue, and thus are more flexible. Motion within a single scan creates spatial artifacts, and when periodic can sometimes be isolated and corrected in the Fourier domain.

2. CT

- beam hardening (intrinsic)

The beam that passes through the center of a scanned object is usually of higher average energy than the beams passing through the periphery [3]. This is due to a greater amount of lower energy photons being absorbed, and results in higher relative attenuation values near the center of a scanned object (also known as cupping).

Various reconstruction techniques can provide some corrective measures.

3. MRI

- object independent signal inhomogeneities (intrinsic)

Main-field inhomogeneities can be caused by shimming[†] in the electromagnets, abnormal magnetic gradients can occur due to improper calibration, and non-linear variations can be created by electromagnetic eddies[†]. These are related to the MRI configuration, and can be corrected if quantified correctly [20]. Various regularly shaped phantoms[†] can reveal stable object-independent distortion patterns.

- object inherent signal inhomogeneities (intrinsic)

These distortions can be due to the variation of the magnetic susceptibility of materials, chemical shifts, or flow shifts. With extra work, they can usually be removed satisfactorily using various techniques [83].

4. CT, SPECT, PET

- photon attenuation (intrinsic)

As electromagnetic waves travel they may be absorbed and refracted by various mediums.

- Compton scattering (intrinsic)

The photons can be deflected by collisions with electrons in the medium they pass through. This will result in lower energy photons travelling in different directions than the original photons.

3.3.2 Uncorrected Distortions

Uncorrected distortions are those that are difficult to model and so are often ignored, or handled separately. It is hoped that their effect will be on the same order of magnitude as the residual error from correcting the other distortions.

1. All Modalities

- partial volume effects (intrinsic)

The volume of a single sample will often encompass more than a single contiguous structure. In these cases the resulting voxel value represents a composition of all the tissue and/or bone types present within the sample.

2. Multi-modality

- different imaging physics (extrinsic)

Because the different physical properties in each modality are independent there is no requirement of identical features existing in both.

3. MRI

- object dependent signal inhomogeneities (intrinsic)

Cannot be removed completely, although recent correction methods provide approximate results (see Corrected Distortions)¹. The variations may be from non-linearities in the gradient magnetic fields, or EM-eddies, or loss of signal due to foreign objects in the body.

4. SPECT and PET

- Radionuclide decay and uptake (intrinsic)

With decay, the time elapsed from radionuclide creation to imaging will greatly affect the scale of imaged values and SNR. The relative uptake of the radionuclide can also vary drastically from scan to scan.

5. PET

- positron wander (intrinsic)

¹It is often assumed that nonisotropic distortions are negligible in MRI and PET [47].

The positron created in the decay wanders a variable distance before annihilating, and thereby emitting a gamma ray. Thus the source of the radiation can only be determined approximately, which lowers the spatial resolution.

3.3.3 Variations of Interest

Variations of interest are the changes or differences in the data that the researcher or radiologist wishes to visualize. These variations should not be removed, and it is the goal of registration to reveal them more accurately.

1. Single Modality

- temporal differences (extrinsic)

The anatomical or physiological structure may change over time. One example is a tumor's progress that can be monitored by its change in shape and volume.

2. Multi-modality

- different information (extrinsic)

Two data sets may provide complementary representations of a physical structure (CT/MRI). A functional data example would be PET/MRI registration. In this case, the dysfunction of an anatomical region without structural problems could be revealed by the metabolic variations.

3.3.4 Summary of Variations

The act of correcting the distortions represents the registration process. Although the uncorrected distortions and variations of interest may be ignored in the registration model, they will help determine the features and heuristics used to find the optimal transformation. A good strategy for registration is to remove the corrected distortions independently if they can be modelled correctly and isolated. This will provide a more controlled and flexible process.

3.4 Conditions for Finding a Solution

Exact registration of the data requires two conditions:

1. The regions in each set must be correctly segmented to represent the same features with respect to the object reference frame.
2. There must exist a one-to-one transformation that maps one set of features onto the corresponding features exactly.

If these requirements are fulfilled, then an exact solution may be found. Both requirements depend highly on the segmentation of the data because this forms the basis of the solution. The derived solution is based on the features, and not the physical space. For the registration to be valid, the segmentation must be a true mapping between them.

In real cases, the conditions for exact registration will not hold completely. An approximate solution will depend on how well the conditions are met, and on the ability of heuristics to find the correct mapping. The ability to measure the accuracy of the solution is also dependent on the degree of fulfillment of the conditions.

Each corrected distortion will have its own necessary conditions to properly remove it. For example, with rigid body motion there must be at least three matching pairs of features (points) that are not collinear in order to provide a sufficient basis for solving the transformation. Other variations further compound the problem of finding an exact solution.

Chapter 4

The Registration Process

It is manifest that if we could find characters or signs appropriate to the expression of all our thoughts as definitely and as exactly as numbers are expressed by arithmetic or lines by geometrical analysis, we could in all subjects, in so far as they are amenable to reasoning, accomplish what is done in Arithmetic and Geometry. ... Moreover, we should be able to convince the world of what we had discovered or inferred, since it would be easy to verify the calculation either by doing it again or by trying tests similar to that of casting out nines in arithmetic.

- G. Leibniz, *Preface to the General Science*, 1677

The process of registration takes the raw data from two or more scans and transforms one onto the other, providing an estimate of the resulting accuracy. The entire process can be performed in three phases:

1. Segmentation
2. Matching
3. Verification

The series of actions may be tightly or loosely coupled (overlapping or disjunct phases) depending on the method used. Loose coupling allows each of the phases to be performed by methods specialized for individual cases, thus providing flexibility.

4.1 Segmentation

Segmentation is the process of reducing the data into a more compact representation (called a feature) that may be compared more easily and effectively. It can be stated as a function (s) acting on both the static (V_s) and dynamic (V_d) scans:

$$s(V_s) = F_s$$

$$s(V_d) = F_d,$$

resulting in the set of static (F_s) and dynamic features (F_d). It is important that the features selected represent the same physical entities in both data sets. This is especially true with the registration of functional data. In these data sets the lack of pronounced distinctions between different anatomical structures is compounded by the lower spatial resolution. It requires expert knowledge to select appropriate features, as there is great variability within functional data of the same modality due to metabolic variations. This is probably clearest in SPECT, where both radioisotope uptake rates and delivered dose amounts can vary.

Segmentation may be performed manually, semi-automatically, or automatically, although the diverse variations in the data dictate that the segmentation can almost never be fully automatic. Pre-processing of the data may facilitate the segmentation process. It is usually performed by various filtering operations. The goal of segmentation is to eliminate all of the data except those features that will be used for the matching process. The features chosen should be representative of the transformation needed to register the entire data.

4.1.1 Feature Types

There are numerous ways to represent the features of three-dimensional data. A taxonomy (as they are derived) of the different types of features is given in Figure 4.1. Points are the simplest feature, and are the underlying constructs for most other features. Their compact representation allows operations to be performed efficiently. Edges are the next basis and are most commonly the representations of boundaries. A boundary may be based on an equal value, an equal gradient, edge-detection[†] operations, classified regions, or any other distinguishing trait. Volumes are different in that they are 3D space filling, but they utilize surfaces as external boundaries. Hence they are often formed from very similar techniques to those used to create surfaces. Statistical features are formed from numerical analyses of the data, which group the data into various salient components. They are often formed from a subset (a rough segmentation) of the data. There are also many other representations based on various levels

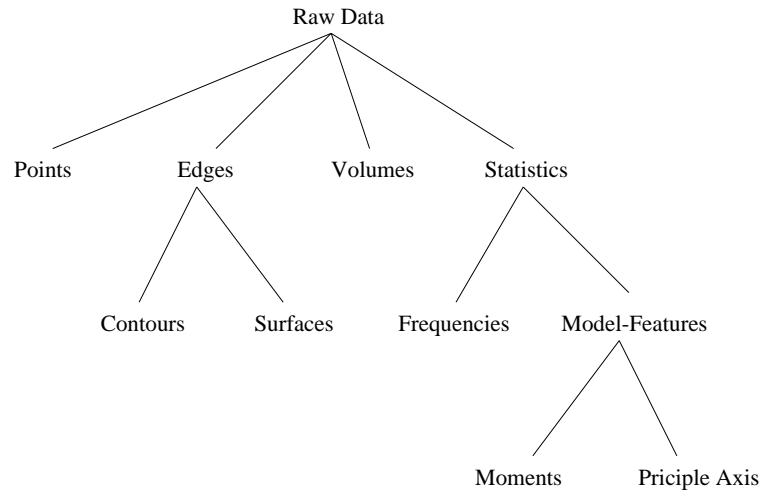


Figure 4.1: A Taxonomy of Feature Types

of abstraction and structures.

Often features are composed into a model of the object being scanned. The model can be formed from statistical norms or strictly theory, and a group of these models for a specific region is commonly called an “atlas”. Hohne et al. [36] have shown that an atlas can be a powerful visualization and educational tool. Incorporating knowledge of the data into a feature creates a model, the simplest of which may be the interhemispheric fissure of the brain when represented as a plane. It has been used by Gershon et al. [29, 28, 26, 27] and Kapouleas [47] for registration, and Gershon also discusses incorporating various other information into a representation. Chinzei et al. [16] presents the use of models for segmentation and matching, Peifer et al. [67] utilize a model of the coronary arterial tree to aid in registration, and Kovacic et al. [50] register PET and CT by mapping them to a normal (model) position. The most common features are listed in Table 4.3.

4.1.2 Choice for Features

The region of the data to be selected for a feature may vary for each patient. It will depend on the types of scans used and the regions being imaged. Van den Elsen [89] has labelled all features as either intrinsic or extrinsic. The intrinsic features are only dependent on the patient, while

Feature	Description	Representation
points	3D discrete positions or voxel volumes	tuples
contours	stacks of 2D lines	lines
surfaces	iso-surfaces/gradient	polygons
volumes	volumetric region	closed-surfaces
frequencies	transform into Fourier domain	complex numbers
moments	centroids, ...	tuples
principal axes	compound moments	vectors
atlas	model of average object	surfaces & volumes

Table 4.3: Feature Types

extrinsic are the result of artificial objects present in the scan. Extrinsic features are commonly referred to as fiducial markers, as they are intended to provide a stable object-defined reference frame.

Intrinsic features are usually chosen as the boundaries of physical structures. Bone structures are often the most rigid and so are good choices if they exist in both scans. With functional and anatomical scans of the head, the brain surface is often used [73, 15] but the epidermis of the head can be obtained from PET transmission scans¹ and may provide a sharper feature [55, 9]. Intrinsic features should not be deformable objects, unless these alone are the objects to be registered. A comparison of some of the possible choices for features of the head was performed by Wowra [98]. This is summarized² in Table 4.4.

There are many different types of extrinsic features (fiducials) that allow the delineation of the object's reference frame. Fiducial markers (commonly points or lines) are constructed to facilitate easy segmentation. They are limited by the accuracy of repeated positioning and segmentation. The major limitations of extrinsic markers are the inability to be used retrospectively and the discomfort inflicted on the patient.

¹A transmission scan is similar to a CT scan, and is done before each PET slice to obtain attenuation values to aid in the reconstruction.

²Wowra's survey did not include SPECT, but it can be approximately derived as a degradation of the PET values.

Feature	CT	MRI	PET/ SPECT	Total for Structure ($X/9$)
mid plane	3	3	3	9
skull	2	0	0	2
skull base	3	0	0	3
ventricles	3	3	2	8
cortical relief	2	3	2	7
vessel architecture	1	2	0	3
Total for Modality ($X/18$)	14	11	7	
Legend: Detectability Scale 0 - not, 1 - in principle, 2 - sufficient quality, 3 - optimal				

Table 4.4: Comparison of the Choices for Intrinsic Features in the Head

4.1.3 Techniques

The techniques are often tailored for specific modalities. Comparisons of manual tracing, count-based methods, static thresholds, adaptive thresholds, and 2 and 3D edge-detection methods are provided by Long et al. [60]. Lee and Chung [54], Hohne and Hanson [37] and Andreasen et al. [4] present various other techniques. For more detail, see Johnston's thesis [43] for a discussion of segmentation.

Segmentation techniques require varied amounts of user interaction and time, and this provides one of the means for classifying them. Current techniques are most often image based operations in which each slice is treated as an independent image and the 2D results combined. 3D segmentation exists but the huge amount of data makes most routines impractical. In each case they will reduce the data into features for use in the matching process.

- Manual Segmentation

Manual segmentation is the process of direct human interaction to select the various features. The class can be defined as those methods in which the inclusion or exclusion of all voxels is directly decided by the user. Two examples are the selection of single 3D points by a user, or the manual tracing of an ROI, both of which require interaction through the use of a mouse, trackball, or other user interface. These methods can be very

accurate (although they are not always so) as they can utilize the highly adapted visual system of an expert. They also are very time intensive for the person performing the task.

- Semi-Automatic Segmentation

Most methods of segmentation fall into the class of semi-automatic methods. These require less human interaction because they utilize some form of knowledge about the data in order to perform the decision making for including or excluding voxels. Threshold methods allow the selection of voxels in various ranges. They may be combined with seed points to form region or volume growing techniques. Auto-contouring is another simple one of these methods, and is used in two dimensions to create a line according to certain boundary criteria. A three-dimensional counterpart is isosurface generation from a seed value or point. With increased knowledge, a probabilistic method can utilize statistical techniques both *a priori* and “on the fly” to help make decisions.

- Automatic Segmentation

Automatic methods can be created for very specific tasks. When variations in the data are highly predictable, that knowledge can be incorporated into a model that can allow the segmentation to proceed automatically. When the knowledge is explicit it can be incorporated into an expert system, and it will perform well to the degree the assumptions underlying the rule base hold. These types of segmentation usually cannot tolerate too much variance, but they may be necessary to feasibly perform the large number of segmentations required for a comprehensive statistical analysis.

4.2 Matching

Each matching method can be broken down into its three main components:

1. Feature Space - type and choice of features that will be compared
2. Comparison Function - quantification of how closely the features match

3. Heuristic - an approach based on the comparison function to find a solution

The great diversity of the different images to be registered precludes the existence of a single best method for all cases. The huge transformation parameter space necessitates a strategy for finding the solution, as any brute-force method would be infeasible. The solution may be the result of a direct calculation or it may involve some form of iterative process to optimize the match.

4.2.1 Features

The feature type that will be selected is linked directly to the matching method to be used, although this usually does not restrict the method used to segment the features. The features should be chosen to minimize the effects of uncorrected distortions and variations of interest, as their inclusion would only lead to a less accurate registration. Depending on the differences between the data sets, a certain level of abstraction (feature segmentation) may provide a useful and more efficient method for handling the matching process. The more similar the data sets are, the less abstraction is necessary. Simple features can require more computationally expensive methods; complex features often inherently contain more relevant information in a compact representation and thus require less.

The types of features selected allow the desirable properties of the object being imaged to be isolated. Structures or regions that are less vulnerable to physiological variations can be chosen. This allows the removal of some variances due to growth, motion, or dynamic density. Boundaries of structures are more easily delineated and often represent more stable positions. The differing physical properties imaged by the different modalities require that some segmentation into features be performed. Again, the high contrast that often occurs on structure boundaries is often the most suitable property consistent across different modalities. Only in registering data sets of the same modality would it be appropriate to use the entire data set without some form of segmentation.

4.2.2 Comparison Function

When any two non-identical sets of features are to be matched, an appropriate criterion must be chosen to reveal the relative success of the match. The comparison function quantifies how closely the features match, and therefore is a relation between two sets of features. It can be formulated as a function. In equation form it is:

$$c(F_s, f(F_d)) = residual$$

where c is the comparison function, F_s and F_d are the static and dynamic sets of features, f is a transformation on the dynamic set of features, and the *residual* is the function value. The registration process attempts to find the transformation f in order to minimize the residual. A single numerical value can usually be produced from the chosen function. Thus when the function has a zero value it represents no discernible difference between the two sets of features.

The comparison function can range from the simplistic to the complex. After the dynamic features undergo a 3D transformation the resulting features will be in a continuous³ domain. Therefore the comparison function may be discrete or continuous in nature. When using a discrete function the features are commonly projected onto the slice planes of the static data set, or rounded to integer values.

It is highly desirable to use a function which is a metric, and therefore has the following properties:

1. The global minimum is reached only when the two sets of features are identical
2. The function is equal regardless of the order of comparison in that $c(F_1, F_2) = c(F_2, F_1)$.
3. The triangle inequality holds between individual features. $c(i, j) + c(j, k) \geq c(i, k) \forall i, j, k \in A, B$

Another useful property is that the overall function be monotonic; as the difference between the features increases the function is non-decreasing.

³Continuous in the sense of being real number coordinates not restricted to the original slice planes.

If the sets (F_s & F_d) are composed of more than one feature, the residual of the comparison function is based on some method for the combining of all the values for the individual features. This method will be dependent on the types of features. Any comparison function should be able to tolerate variations in order to allow the heuristic to find an optimal match. The use of more features usually provides increased robustness. A brief summary of features and common comparison functions is presented in Table 4.5.

Feature	Comparison Function	Comments
raw data	differences of intensity cross-correlation	few variations (same modality) raw data/bitmap, expensive
points	distance	any distance function or approx.
contours surfaces	distance, perpendicular distance	computationally expensive
volumes	differences of volumes	non-overlapping volume or (\cap/\cup)
frequencies	Fourier correlation	computationally expensive
moments	distance	simple
principal axes	rotation angles & distance	simple

Table 4.5: A Summary of Features and Corresponding Comparison Functions

Differences are the simplest function and are usually calculated by arithmetic subtraction, although the binary AND, OR, or XOR operations could be used. While normalized cross-correlation is still too numerically expensive for 3D, it can be used on stacks of 2D images if the registration process can be reduced to two dimensions. Distances are the most common comparison function and will be discussed in more detail later. A large amount of theoretical and practical work along lines similar to registration with surfaces has occurred in the area of computer vision in the area of *attitude determination*. This work on the registration of surfaces has developed numerous comparison functions based on curvature and other surface properties, but has not been applied to very noisy or complex surfaces present in medical applications. See Li [58] and Little [59] for further information. 3D Fourier correlation methods are not utilized for similar reasons as those for cross-correlation. Sometimes the comparison function may be used directly to calculate the optimal transformation, as when the angle between principal axes is used to align them.

Distances and Their Computation

The distance between features is the most prevalent form of comparison. When the features allow a distance function to be used, it is most commonly Euclidean distance⁴:

$$D[(x_1, y_1, z_1), (x_2, y_2, z_2)] = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2}$$

but other less computationally demanding functions such as the class of non-negatively weighted absolute functions:

$$D[(x_1, y_1, z_1), (x_2, y_2, z_2)] = c_0|x_1 - x_2| + c_1|y_1 - y_2| + c_2|z_1 - z_2|$$

(this is just Manhattan distance when $c_0 = c_1 = c_2 = 1$) are often implemented.

When an iterative solution requires repetitive calculation of the distance functions, an approximation may be necessary. One way this can be done is with a distance transform (DT). “A distance transformation converts a binary digital image, consisting of feature and non-feature pixels, into an image where all non-feature pixels have a value corresponding to the distance to the nearest feature pixel [11].” This is essentially the precomputation of distances over a discrete region surrounding a feature, and can be applied in 2D or 3D.

One type of DT, called chamfering, can be efficiently formed by a two pass filtering operation on a data set where features are set to 0 and the rest to ∞ . This is illustrated in a simple 2D example shown in Figure 4.2. As the mask is moved across the feature image, if the sum of any mask value and the pixel below it is less than the center pixel value, the center value is replaced. After two passes of the image represents the distance transform of the features, where each pixel represents the “distance” to the nearest feature pixel. Using the mask values in this simple example results in the computation of what is called the *City block* distance. 3D chamfer Distance Transforms were incorporated in the surface matching techniques that were implemented. They are discussed further in Section 5.3.2 and analyzed in Section 6.2.1.

Other DT approximations exist such as the n-neighbor and octagonal methods, but they provide less accurate approximations [10]. The actual Euclidean distances can be calculated

⁴The L-norm, $L_p = \sqrt[p]{\sum_{i=1}^n \|x_i - y_i\|^p}$, for $p = 2$




Figure 4.2: The Formation of a 2D Distance Transformation

with a similar but expanded chamfer DT algorithm. The increased computation and memory requirements necessary for this are not reasonable, although Paglieroni's 2D algorithm [66], which reduces the increased time costs, may be expandable to 3D. The chamfer DT described is good enough for most applications and is efficient to calculate. It is used for a hierarchical 2D edge matching by Borgefors [12], and has been adopted for 3D registration of surfaces by Jiang [41] and van den Elsen et al. [88].

When more than one feature is present there are different ways in which an overall function value can be formed. Borgefors' [12] tests with 2D chamfer matching have shown that of median, arithmetic sum, root-mean-square (rms^\dagger), and maximum, the rms value provided the best results. This analysis is applicable to medical imaging and the two dimensional case is illustrated in Figure 4.3.

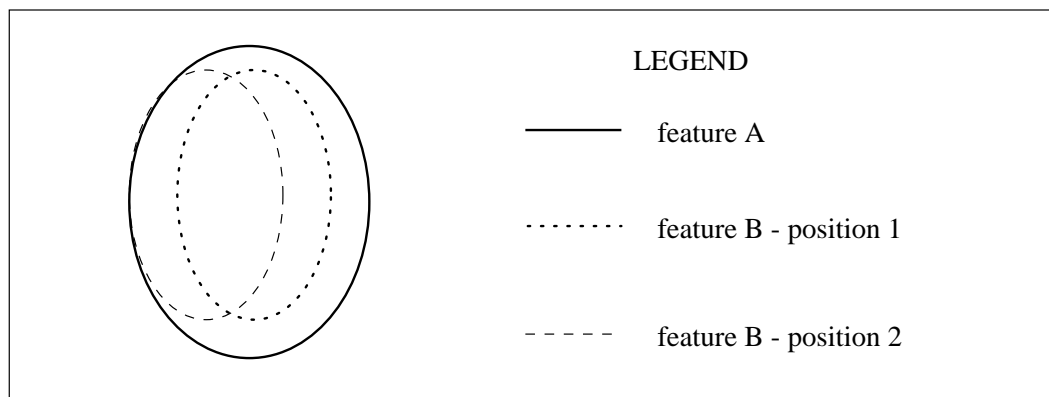


Figure 4.3: A Comparison of Features

If feature B is being matched to feature A, then the *rms* value of position 1 would be smaller than position 2, while the mean value would be the same for both. It is probably more appropriate to have the more consistent error present in position 1, rather than the large variations in error with position 2. This type of closed feature is commonly segmented in medical scans, in which a slight error may result in one feature being contained within the other.

4.2.3 Heuristics

Once the data has been segmented into features, a transformation must be found that will map one set to the other. The type of data, as well as the nature of the changes that have occurred, will determine the complexity of the heuristics needed to find a correct mapping. A minimization of the comparison function will usually drive the heuristic. There are two major classes of heuristics: those that calculate the solution directly, and those that iteratively refine a approximation until it converges.

Direct Solutions

Direct heuristics are those that calculate the optimal transformation in one step. A single residual value from the comparison function must be used to define the registering transformation. Although the techniques used to solve systems of equations and correlation methods have an iterative nature, they are grouped in this class because there is no large distinction between steps. Direct methods commonly contain numerically expensive algorithms at their core, and so are not feasible for an iterative process. These techniques require enough knowledge of the transformation so that they do not require a search of the different possibilities.

One of the most commonly used 2D methods is cross-correlation,

$$a \otimes b = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} a(x - \xi, y - \eta) b(\xi, \eta) d\xi d\eta$$

and the selection of the maximal value for x and y [39]. In practice it may utilize the form of

a normalized discrete convolution, which has the form in three dimensions of:

$$C(t, u, v) = \frac{\sum_x \sum_y \sum_z V_1(x, y, z) V_2(x - t, y - u, z - v)}{\sqrt{\sum_x \sum_y \sum_z V_1^2(x - t, y - u, z - v)}}$$

calculated for an arbitrary translational range for t, u and v for x, y and z -axis respectively. The convolution can be seen to be identical to the correlation except volume V_1 is inverted in all dimensions, while the normalization provides voxel magnitude invariance. The global maximum of this function represents the optimal translational coefficients.

By resampling (mapping) the data into polar co-ordinates, rotations about the chosen polar origin can be found by a translation [17, 33]. Hibbard [33] does this with the autocorrelation ($a \otimes a$) function to avoid dependence on an arbitrarily chosen sampling origin, while Cideciyan [17] maps the Fourier transform of the data into log-polar form to find rotation and scaling from the translations.

A major problem is that the computation required for these correlation methods is too great for use in three dimensions. The results of these methods will not produce a one step solution for both rotation and translation, and so will usually have to be iterative to be exact. The techniques are still useful. Maguire et al. [62, 61] utilized 2D cross-correlation in segmentation to improve localization of their point features. Future development of parallel processing or faster computers will make these methods more feasible.

Moment based methods use basic statistical values derived from the shape of the data to directly calculate a transformation. Centers of mass can be aligned to correct for translation, and principal axes (if they are representative and stable) can be aligned for rotation correction. These methods are very dependent on the low frequency nature of the features being similar. They require that the variations in the data be uniformly distributed so that the averaging performed with a moment calculation can accurately reveal the optimal transformation. The use of moments and principal axes were studied by Faber and Stokely [24] to find the orientation of cardiac blood pools, while Bartoo and Hanson [7] and Toennies et al. [86] utilized moments for registration. The use of moments are also inherently utilized in most model-based registrations.

The use of corresponding point features (commonly extracted fiducial markers) provides a basis for solving a system of equations. Once a transformation type is chosen and matching points are selected, any of the numerous methods for solving systems of equations can be used. A comparison of algorithms that can be used to solve for the transformation matrix is provided by Timmens [85]. Pelizzari et al. [69] found that solving for an affine transformation provided reasonable results when enough corresponding points could be chosen, and Henri et al. [32] also solve equations for a transformation matrix. Polynomial transformation (commonly called warping [95]) is one of the curved transformation types that uses the matching points to derive polynomial functions which can map one set onto the other. This may be done by either approximation or interpolation, although the use of higher order polynomials, needed for complex warping, can have problems with large oscillations in the solution⁵. Approximation attempts to minimize the variance (least squares) from all the control points, while interpolation requires that the transformation map all the corresponding points exactly. For distortion correction Schad et al. [78] utilize fourth order 2D-polynomials. Maguire et al. [62, 61] and Kramer and Noz [51] use first and second order polynomials for the mapping of points in 2D, and Schiers et al. [79] use them for 3D for registration.

Iterative Solutions

Iterative methods utilize successive approximations to “zero in” on the optimal parameters, and so are often more flexible in accommodating variations than direct methods may be. The comparison function must be simpler than those used for direct methods due to their repeated calculation. One exception is the novel technique discussed by Woods et al. [96], in which the variance of the voxel-to-voxel ratios of two PET scans is iteratively minimized. In a similar but more complex method [97] Woods performs MRI-PET registrations. The process of iteration often allows errors in initial approximations to be corrected later. The formulation of the heuristic into discrete steps allows for the integration of a hierarchical (often pyramidal)

⁵This is called the Rungian effect. When this occurs the solution may have only small errors near the control points, but elsewhere the error will be very large.

structure into the process. This can provide an efficiency gain by initially optimizing on high levels of abstraction and then fine tuning the parameters on the lower levels.

Gradient search methods can be utilized when the comparison function provides a generally convex shape over the parameter space. This allows a series of small refinements of the parameters to guide the search from an initial position to the minimum. The presence of local minimum, as opposed to one global minimum, is the major stumbling block for these methods. The comparison space is very complex and so these local minimums inevitably exist for 3D rotation and translation. Jiang [41, 42] deals with this problem by using a large number of starting points spread over the parameter space.

4.3 Verification

The result of the matching process will be a transformation to register the data. How accurately the data are aligned must be determined before results can be used. This is difficult to do in real applications, and the residual from the matching heuristic and visual inspection are the only two techniques. Therefore the heuristic should be tested in simulations which allow the error to be accurately quantified and the robustness tested.

4.3.1 Comparison Function Residual

The residual of the comparison function is often used to describe the similarity between features after a calculated transformation. This residual should be a good indicator of the relative success, however, it is dependent on the segmentation and the approximations of the comparison function and so should not be used as the final indicator. Often the residual can only be used to quickly discover which results are poor.

4.3.2 Simulated Data

Simulated data is just about the only way to arrive at a truly quantitative measure of verification. A simulation is the artificial creation of two scans that are to be registered. With real

data there are often too many unknowns and other distortions to perform a highly accurate check on the mathematical soundness and completeness of the results.

With most simulations the actual correspondence between 3D points in the features can be calculated. The correspondence of static and dynamic feature points can be used to calculate the Euclidean distances between the points after the registration. Depending on the position of the features, relative to the other data, this may provide a bound for the misalignment error of the entire data set. When using resampled real data the segmentation error can be excluded from the resulting alignment error. If the simulation transformation is also used on the static set of features, then it too can be registered and compared to the identical unmodified set.

With the simulated data, the registering transformation can be compared to the one used to create the simulation. When the transformations are in matrix form, the product of the two matrices should result in the identity matrix.

$$M_{reg} * M_{sim} = \begin{bmatrix} 1.0 & 0.0 & 0.0 & 0.0 \\ 0.0 & 1.0 & 0.0 & 0.0 \\ 0.0 & 0.0 & 1.0 & 0.0 \\ 0.0 & 0.0 & 0.0 & 1.0 \end{bmatrix}$$

Although this results in a quantitative measure of error, it is not very intuitive. The matrix can be then used to transform a single set of features to determine the average distance between corresponding points mentioned previously.

Mathematically Created

Simulated data can be based on various levels of abstraction:

- mathematically created data and/or features

In this way highly controlled shapes can be created and transformed often with no uncertainty in orientation and position. Ellipsoids or other more complex shapes can be used to give bench marks of the heuristics with regards to well-behaved comparison functions.

If segmentation methods require testing then original data can be created, otherwise the features can be formed directly from equations.

- numerically transformed real features

Using real features allows the heuristic to be tested on the actual comparison space that will exist in real cases. A single set of features can be transformed, distorted, or modified in a completely controllable manner, or else two sets of matching features can be transformed accordingly. This can provide data sets which test aspects of the heuristic that relate to specific variations.

- resampling real data

Resampling at different rates after an arbitrary transformation creates data that mimics the results of multiple scans, with the only variation being the position and orientation. This also allows the segmentation to be tested along with the matching method. Added noise and other modifications can be used to replicate the differences occurring from various scanning device parameters, or possibly other variations, enabling the creation of highly realistic simulations.

Phantom Data

A phantom is an object that is imaged for the purpose of simulating the effects of imaging a real object. It can be scanned after various controlled physical transformations. The exact composition of the phantom's volume is known, so the scanned volume can be checked. Often the phantom contains various imaging substances that can be exactly controlled.

4.3.3 Patient Data

When actual data has been registered, it can be visually inspected to verify the result. This will, of course, be a subjective evaluation. It may be aided by the presence of fiducial or reference

markers in the images, which can be easily segmented and compared. Current analysis of the data is almost always performed on 2D images, and so it is most likely that one data set will be resampled along the same image planes as the other data set. This enables any numerical analysis or algorithms to be performed on identical regions. While these operations are normally done to utilize the added information from the registered data set they can also be used to discover errors.

Various visualization techniques allow the comparison of the data sets for verification. These can enable the user to detect various signs of misregistration.

- **Linked Cursor**

The simplest visual inspection of registered data is the use of a linked cursor. In this case the movement of a cursor in one data set will move another cursor to the spatially equivalent position in the other data set. For 2D viewing, this is commonly performed after an initial step of resampling identical slice planes using the registering transformation. This simplifies the process of calculating the matching cursor position, but more importantly it preserves directional motion across the data sets (e.g. horizontal motion of the cursor in one set results in horizontal motion in the other).

- **Transferring ROIs**

With the registering transformation any ROI may be transferred from one set to the other. The ROI may be of any type from simple polygons to the output from an edge-detection method. With a 2D ROI in set A, the data in set B may be resampled so as to lie in the plane of the ROI. The ROI from set A can then simply be overlaid on the other set. While this also can be done with 3D ROIs, another alternative is that the ROI can be transformed over to set B and then intersected with the appropriate planes.

- **Image Merging**

It is often beneficial to resample one entire data set so that it lies in the same planes as the other. The equivalent slices can be shown side by side, or can be integrated using

various visualization techniques. This may involve averaging, color-washing, selective segmentation, or other compositing techniques.

- Volume Rendering

Three-dimensional display techniques have only recently been utilized to provide a better understanding of the data. Diagnosis is still mainly performed by looking at the 2D slices, but volume visualization techniques are quickly gaining wide-spread acceptance. The process converts the 3D data into 2D images that retain some form of 3D information. There are two broad classes of volume visualization. One utilizes surfaces and the other various types of projections. The first fits surfaces onto the data and then uses simulated lighting and hidden surface removal to simulate 3D viewing effects, while the other projects rays through the data onto the 2D image and accumulates desired characteristics.

For all 3D display techniques the data requires some segmentation to provide suitable results. The segmentation may be inaccurate, compounded by the volume rendering approximations, this may create a misleading 3D image. This is a major concern with the use of volume visualization, and a reason why 2D slices are the primary diagnostic platform. Another disadvantage of volume visualization is the long time necessary for creating the rendered images. A 3D object is often rotated to provide different viewpoints, and so this multiplies the number of images to be created and further increases the time needed for creating the visualization. One method to accelerate image creation is the use of an octree representation of the data [94], which eliminates the traversal of empty space, or in surface rendering techniques, a reduction in the number of polygons [87, 80]. Ultimately, parallel computation and faster processors may provide the only true solution for real-time volume visualization. See Mealha et al. [64] for a discussion of various data structures and optimizations. Different volume visualization techniques are surveyed by Elvins [22], while Wallis and Miller [93] and Tiede et al. [84] describe them in a medical context. A survey of the hardware involved is given by Stytz et al. [82], and two recent techniques are explained by Yoo et al. [100] and Levoy [56].

4.4 Case Studies

Three specific registration methods will be discussed here in further detail. These include typical examples from the three most common techniques represented in the literature: manual methods, corresponding point mapping, and the iterative alignment of surfaces.

4.4.1 Iterative Manual Alignment

A completely interactive method for the 3D alignment of functional and morphological scans is provided by Uwe Pietrzyk, Karl Herholz, and Wolf-Dieter Heiss [73]. Another manual method is described by Kapouleas et al. [47].

Feature extraction

The feature extraction is formed by 2D preprocessing of the data. This involves convolution with various filters for edge-detection and smoothing. The output of this preprocessing is hopefully the set of contours which represent the boundaries of various internal structures. The segmentation is only performed on one set, usually the higher resolution data set.

Registration

The contours obtained from the feature extraction are overlaid on the other data sets to provide a visual means of comparison. The data sets are interactively and arbitrarily re-sliced to find the correct transformation. When an expert is satisfied that there is enough correspondence, the registration process is complete. This typically required only 5-10 min on standard computing equipment.

Results

Pietrzyk et al. discovered that it is visually easy to detect a variation in rotation of 3 degrees and a 2 pixel translation. Therefore they concluded that an accuracy to within ± 1 voxels was

obtained. Their matching process was able to register to a point where the greatest error was below the FWHM of the line spread function in the lower resolution imaging device.

4.4.2 Solving Equations Formed From Corresponding Points

An approach that uses numerous corresponding 3D points to define a polynomial function to transform and register the data is presented by Schiers et al. [79]. A similar approach in 2D was implemented by Maguire et al. [62, 61].

Feature extraction

The features chosen are single points (landmarks) in each data set that represent the same 3D point. To accomplish this, Schiers created a tool for the interactive selection of 3D points. This was commonly done with a projection of the outer skin surface (CT/MR), which could be manipulated to provide different views.

Registration

Once the corresponding points are selected by an expert, the calculation of a polynomial which approximates or interpolates them must be made. Linear polynomial interpolation requires at least 4 points, and if distortions are present higher polynomials are needed. From these points, Schiers solves three equations for the respective x , y and z transformations. The coefficients are calculated from the minimization of a least squares residual.

Results

The accuracy of results were not discussed specifically, but were stated as being sufficient for rigid objects. It was noted that other distortions required higher degree polynomials, but then problems with oscillations occur, so many more points are needed.

4.4.3 Iterative Surface Alignment

Pelizzari et al. [68, 55, 15] has devised a well-known “hat” and “head” surface matching method. Other methods of this type are Jiang et al. [41, 42], Malandain and Rocchisani [63], van den Elsen et al. [88], Collignon [18], and van Herk [90].

Feature extraction

Pelizzari’s method involves the extraction of a surface from each data set. This is done by manually selecting a feature (contours) in each set. The user will go through all the slices segmenting the contours in 2D. This can be aided by edge-detection and contour-tracing programs which provide a faster, semi-automatic procedure. This process will yield a stack of contours that can be used to represent a surface for both the static data (which they call the head) and the data set to be transformed (the hat). These surface features are then used for the registration process.

Registration

The features are initially transformed according to the machine setup parameters. The transformation is assumed to be rigid body motion with scaling used for configuration error and MR distortions, and therefore 9 parameters are needed to encompass all rotation, translation, and scaling.

Powell’s minimization technique [74] is used to find the transformation function parameters that minimize the residual; with the residual being the sum of an approximated distance for each “hat” point to the closest point on the “head”. The minimization method is based on the changing of one parameter at a time and is a conjugate-gradient approach. The comparison function is the distance between the “hat” and “head” surfaces calculated by a ray travelling from the “hat” point to the centroid of the “head”. Points that are not within the range of the “head” slices are ignored with the intent of allowing a better match of features that only partially overlap. Typical surface features consist of ≈ 1500 points for the “head” and ≈ 250

for the “hat”. The minimization process required 5-10 min on a VAX-11/750.

Results

In phantom studies (CT to CT/MR/PET and MR to PET) the method found transformations that had residual values (*rms*) that were less than the largest pixel sizes [68]. In-plane errors were estimated from the deviations of the centroids of cylinders imaged in both modalities. The overall analysis is shown in Table 4.6 [68].

Modalities	Type	Voxel Spacing (<i>mm</i>)	Residual (<i>rms mm</i>)	In-plane Error (<i>mm</i>)
CT/CT	Phantom	1.15x1.15x4 / 1.15x1.15x4	0.73	0.7±0.48
CT/MR	Phantom	1.15x1.15x4 / 1.95x1.95x10.5	1.16	1.36±0.45
CT/PET	Phantom	1.15x1.15x4 / 2.50x2.50x14	1.80	1.38±0.65
MR/PET	Phantom	1.95x1.95x10.5 / 2.50x2.50x14	1.83	2.48±1.09
PET/MR	Clinical	2.50x2.50x14 / 1.17x1.17x10.5	1.6	n/a

Table 4.6: Overview of Pelizzari et al.’s Results

Chapter 5

Implementations and Innovations

But as soon as I had acquired some general notions in physics, and beginning to test them in various particular difficulties, had observed how far they can carry us, and how much they differ from the principles that have been employed up to the present time, I believed that I could not keep them concealed without offending gravely against the law by which we are bound to promote, as far as in us lies, the general good of mankind.

- R. Descartes, *Discourse on Method*, 1637

As part of the research for this thesis, numerous different methods and techniques for each phase of registration were implemented and tested. The main thrust of the work involved a thorough examination of *surface matching* methods. The programs used to perform the registration were created in an object-oriented data flow environment. In this environment images are objects which can flow down a pipeline of processing modules.

The registration was restricted to rigid body transformations because they are simple to perform, they maintain the relation of intra-data distances, and they are the most widely used. They are appropriate for the scans of the head that were the focus of testing. This chapter describes the major characteristics of the various methods and provides an in-depth description of the techniques used.

5.1 Platform

The registration process was performed in a data flow environment created for multi-purpose image processing. This was provided by the software package WiT [5] (Object Oriented Imaging Technology) running on SUN workstations. The WiT data flow environment provides a large library of I/O, image processing, and control modules. These modules are visually connected together with links to describe the data flow algorithm.

A high-level toolkit was used to decrease the cost of developing a complete registration system. This eliminated the need for the programming to start from scratch, because low-level image operators and graphical user interfaces are provided by WiT. It did, however, lock the programmer into a framework in which the programmer is biased towards using the provided tools, whereas accuracy and efficiency gains could be achieved if specialized functions were coded manually. With WiT, user defined modules can be easily incorporated into the system. In this way a large number of different algorithms can be prototyped quickly and efficiently, simply by connecting together different procedures.

An example of a data flow diagram (called an Igraph) produced directly from WiT is shown in Figure 5.1. WiT is a 2D image processing system and so pseudo 3D methods were performed by utilizing a vector of images and an external data structure to represent a scan volume. The *scan info* operators on the far left of the diagram represent scan or volume information objects. They contain the dimensions of the data set as well as the inter-voxel distances, all of which are shown below the operator as the *dim* and *scale* values. Each operator can have its parameters manipulated by clicking on it's icon, which provides a pop-up control box. The *read raw data* operators load the data into a vector of images, which is then sent to a vector function operator. The operators with solid black circles in the top right corners represent hierarchical operators which contain arbitrary Igraphs within them. These can be expanded, and are constructed the same as any other Igraph. Each element in the input vector is sent to the *auto segment* operator connected to it. This operator (which will be explained later) segments each image independently. The resulting vector of segmented data, representing a surface, is passed down to the *surface to points* operator. The magnifying glass seen on this link can be placed on any link, and allows the user to examine the data that is flowing through it. The *surface to points* operator converts the images to points which will be the primitive used by the registration process. The *registration* operator (which has many other parameters that have not been selected to appear below the icon) then registers the two sets of features and outputs the registering transformation that minimizes the comparison function, as well as

the final residual value. The centre of mass of the features is used in the registration process. It is also used in the next operator that resamples the dynamic set of data to extract slices that should match with the static set. The resampled data, and residual are displayed, as well as the aligning rigid body transformation (the RBT seen through the magnifying glass).

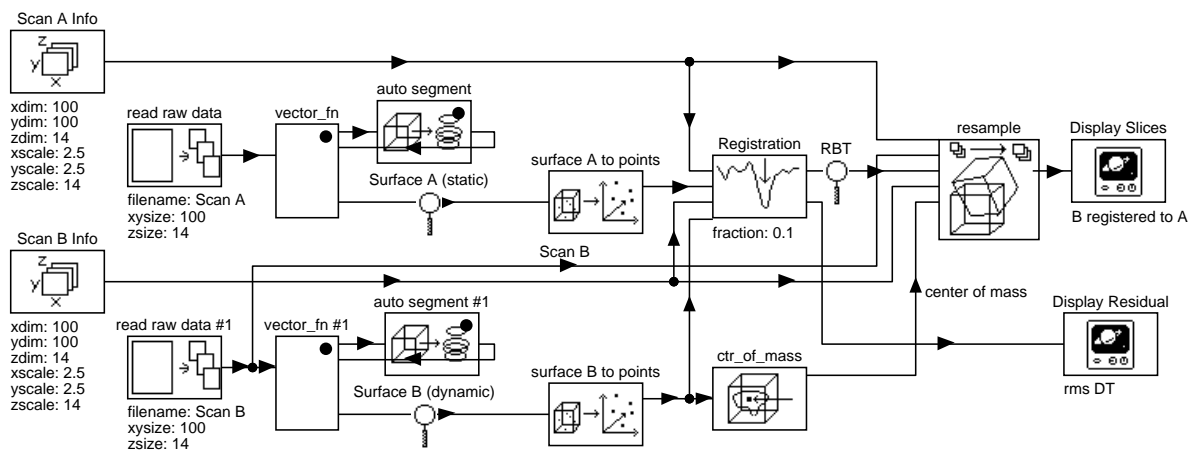


Figure 5.1: A WiT Igraph for Registration

5.2 Direct Point Matching

An initial implementation utilized individually selected point features, which in their simplicity allowed a brute force matching method to be used. This served as a foundation for the implementation and testing of various segmentation and verification subroutines. The simple nature of this method precludes it from use for highly distorted data sets, but it provided a useful starting point for further development. Because of its labour intensive nature, no detailed analyses of this algorithm were performed.

5.2.1 Segmentation

The features chosen for this registration method were corresponding 3D points, representing the same relative position within each data set. These points, commonly called “landmarks”,

are often the center of local extrema or a discontinuity in a boundary. In each data set the points were selected by pointing the mouse at the desired pixels in a 2D slice of the data and clicking, while a slider was provided to allow the user to find the desired slice.

In this way corresponding points were selected from two data sets as representing the same physical position in space. The correspondence between points that were selected could not conveniently be recorded (in WiT), although the user's selections were almost always entered in matching pairs. This segmentation provided a small number of relatively accurate points, the accuracy of each point being dependent on the voxel dimensions, the expertise of the user, and the display and interaction techniques.

5.2.2 Matching

The comparison function for point matching is commonly the distance to the closest point (*rms* for > 1 point). The initial implementation used a brute force method to test all possible transformations within a given range. For two sets of N points the closest distance calculation requires $O(N^2)$ operations. Hence, a less computationally expensive approximation is essential for large N and brute force methods. This was done by precomputing the distances for a small region surrounding each static point and putting the values in a 3D lookup table, similar to that used for a distance transform. After transforming the dynamic set of points, each point would be rounded to get a lookup table distance entry, and the squared distances totalled. The square root of the average point value was then used for the comparison function.

The WiT interface did not maintain the order of selected points, so corresponding points could not be recorded easily. Therefore a system of equations could not be set up ¹. Instead an initial guess of the transformation was used to restrict the parameter space, and a complete search was carried out. The scanning devices' configuration values (sampling distributions) were taken as accurate and used to perform any necessary scaling between data spaces. This

¹Although these selected points would naturally provide a basis for a mathematical least squares techniques such as Procruste's point cloud matching, not knowing corresponding points turns the order $O(N^3)$ matrix elimination into a order $O(N!N^3)$ problem due to the choices for matching pairs.

left six parameters for rotation and translation, with the rotation performed about the centre of the data set. The algorithm then iterated by looping through all parameter ranges in discrete steps.

The minimum value of the comparison function (within the tested range) was taken as the true registration point. If the optimal value occurred when a parameter is at the limit of its tested range, then the search would be repeated with extended ranges for all such parameters. Increased accuracy can be obtained simply by decreasing the step sizes. The brute force method ensures that the correct solution (an approximation) will be found if it is within the bounds of the parameter search, and if the other assumptions hold.

5.2.3 Verification

This basic algorithm was tested on PET-PET and MRI-MRI registrations. The tests were in no way comprehensive, but meant only to initiate the development of the verification phase. Simulated data requiring registration was created by resampling a patient data set after a mathematically performed rigid body transformation. In this way the exact transformation is known. The simulated data was then resampled again, along identical slice planes as the original, using the result of the registration. These registered slices were compared with the original static set, both visually and by subtraction, in order to reveal inconsistencies in the result.

A linked cursor was also simulated by allowing point ROIs to be selected in one set and identical points shown in the other. This allows selected misalignment to be seen more clearly. Identical slice planes can be extracted from the original and registered data and averaged to provide insight into the errors. These techniques, along with the actual numerical values of the result, provided a reasonable and accurate verification of the correctness.

5.3 Iterative Surface Matching

Currently the most successful general purpose registration techniques utilize surface matching. From the “hat” and “head” matching used by Pelizzari [70] in conjunction with Powell’s numerical optimization [74], to a recent physically based rigid body model which incorporated torque [63], these methods have provided good results by using a large number of surface points. The large surface features can be segmented with much less expertise than specific points, especially when dealing with low-resolution data or multi-modalities. Therefore multiple variations of surface based registration were developed and tested.

5.3.1 Segmentation

The major benefit of a surface is that by using a large number of points it allows small uncertainties in segmentation to be averaged out. A surface feature inherently provides more information such as curvature, but it is not often used due to the increased computation and storage requirements. The chosen surface will usually represent the boundary of a single anatomical structure.

There are a great number of different methods available to segment a surface. Most common are two dimensional approaches that operate on each slice independently, and then combine the resulting contours in stacks. Manual segmentation, in which all contours are drawn by hand, was performed but it is extremely time consuming. The 2D iso-value contour was also extended into 3D with an algorithm [99] to provide a single iso-surface. Simplistic heuristics such as thresholding in conjunction with morphological operations were also used to provide a surface feature.

PET/SPECT segmentation

With PET and SPECT the data has a low SNR and so most iso-value and iso-gradient based contour creation techniques do not work well. Although PET transmission scans can be segmented more easily, they were not readily available and so could not be used. The basic segmentation heuristic (for the brain surface) utilized a single thresholding value, region smoothing, and outlining.

The variability in reconstruction techniques does not necessarily preserve relative values between slices, so a single threshold value for all slices could not always be used. Therefore the average value in each slice was calculated as a base value to which a certain proportion of the standard deviation was added to arrive at a threshold value. Although highly dependent on the variance of the noise, this approach provided an approximate segmentation of the brain surface. The WiT Igraph for this technique is shown in Figure 5.2 along with typical data and results in Figures 5.3 and 5.4.

This Igraph is the expanded hierarchical operator *auto segment* from Figure 5.1. The slice (image) enters in the top left corner and proceeds to the *threshold* and *image stats* operators. The *image stats* operator calculates the average (\bar{x}) and standard deviation (σ) which are later used. The *proportion factor* operator provides a constant which is linked to the *memory* operator because it will be used repeatedly for all the images that come through. The *calc* operators calculate $0.8 * \sigma + \bar{x}$ which is then used as the lower threshold. The *closeBin* operator performs 5 iterations of the morphological operations: dilation followed by a erosion, that removes small holes from the segmentation. The *outline* operator then outputs the contours that form the segmented surface.

CT/MRI segmentation

Manual segmentation was used to generate reasonably accurate surfaces. It was performed by a WiT operator that provided mouse controlled contour drawing for each image. This method was extremely slow because no special aids were developed. The thresholding techniques used

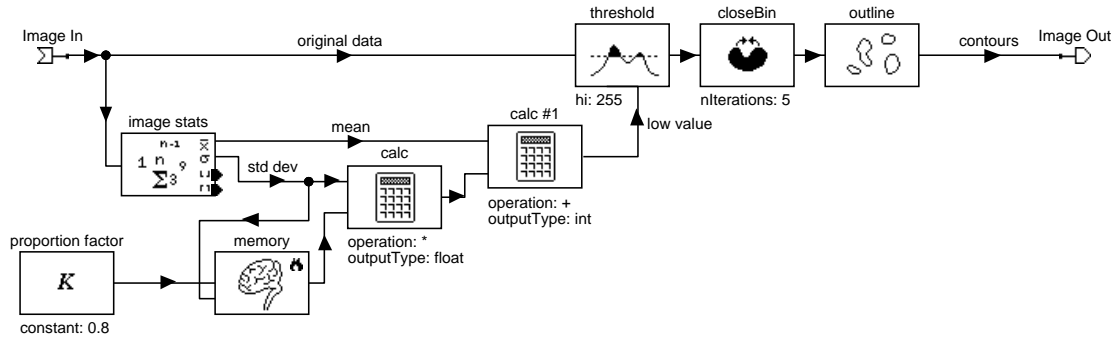


Figure 5.2: PET/SPECT Segmentation

for SPECT and PET were also utilized.

Wyvill's iso-surface growing algorithm [99] was also implemented to segment surfaces. The algorithm tags the voxels and then triangulates the intersections of the surface and the voxel. However the large amount of data in a polygonal surface representation was not needed, so only the tagging of voxels was performed. This was initiated by selecting a seed point from the data, which began the growth of the iso-value surface. An example MRI data set and a resulting segmentation are shown in Figures 5.5 and 5.6.

5.3.2 Matching

The surface was treated only as a set of 3D points which the heuristics could manipulate to find the best match. Numerous comparison functions were studied for their suitability with surfaces. Based on these comparison functions, various iterative strategies were compared to study their behavior under diverse conditions.

Comparison Functions

The comparison function was chosen as the *rms* value of the distance from each point on the dynamic surface to the closest point on the static surface. The precomputation of local distances used in the initial registration method was originally utilized, but it was too inefficient for the




Figure 5.3: PET Data

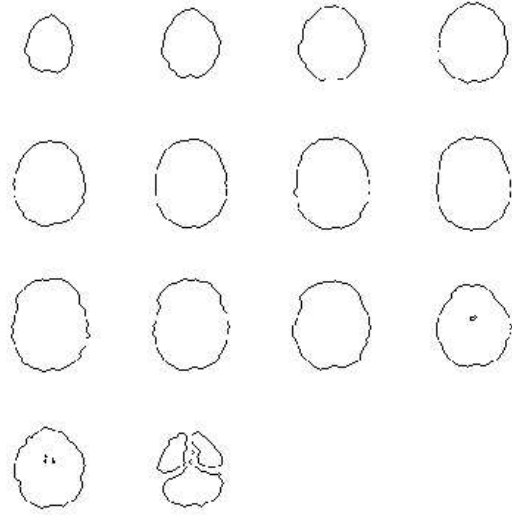


Figure 5.4: PET Segmentation

calculation of large distances. The great amount of overlap when points are close together on a surface made most of the computation redundant. Therefore a distance transform was used.

The process of creating a 3D chamfer distance transform initially requires a volume in which the features voxels are set to 0 and the other voxels to ∞ . Then two passes with a 3x3x3 filter mask will create the DT. The process is provided here in pseudocode:

```
/* 3D Chamfer Distance Transform */
```

```
/* Mask for Forward Pass */
```

$$MF[-1] = \begin{bmatrix} +d_3 & +d_2 & +d_3 \\ +d_2 & +d_1 & +d_2 \\ +d_3 & +d_2 & +d_3 \end{bmatrix}; MF[0] = \begin{bmatrix} +d_2 & +d_1 & +d_2 \\ +d_1 & 0 & \infty \\ \infty & \infty & \infty \end{bmatrix}; MF[1] = [\infty];$$

```
/* Mask for Backward Pass */
```

$$MB[-1] = [\infty]; MB[0] = \begin{bmatrix} \infty & \infty & \infty \\ \infty & 0 & +d_1 \\ +d_2 & +d_1 & +d_2 \end{bmatrix}; MB[1] = \begin{bmatrix} +d_3 & +d_2 & +d_3 \\ +d_2 & +d_1 & +d_2 \\ +d_3 & +d_2 & +d_3 \end{bmatrix};$$




Figure 5.5: MRI Data

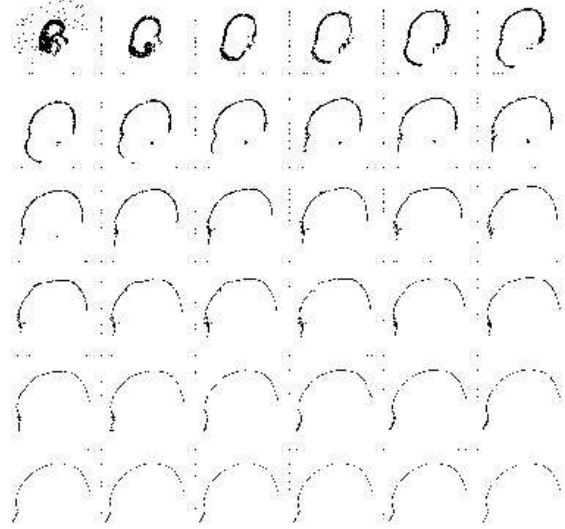


Figure 5.6: MRI Segmentation

```

/* Forward Pass */
FOR (k = 0; k < slices; k++)
    FOR (j = 0; j < rows; j++)
        FOR (i = 0; i < columns; i++)
            DT[k][j][i] = min(a,b,c)∈MF(DT[k+a][j+b][i+c] + MF[a][b][c])

/* Backward pass */
FOR (k = slices-1; k >= 0; k--)
    FOR (j = rows-1; j >= 0; j--)
        FOR (i = columns-1; i >= 0; i--)
            DT[k][j][i] = min(a,b,c)∈MB(DT[k+a][j+b][i+c] + MF[a][b][c])

```

The comparison function provided by a chamfer DT results in a more efficient calculation of approximate Euclidean distances [11]. Anisotropic data was kept in its original form (not interpolated to cubes) and so required an adaptation of the basic chamfer distance transform (DT) that could best be performed using floating point mask values. Although this allows anisotropic DTs to be calculated it may require more memory than is feasible. A 256x265x70

MRI scan would require slightly more than 18MB for the DT. This can be circumnavigated by forming integer approximations that may be stored in a 0 to 255 valued `char` (1 byte) versus the arbitrary valued `float` (4 bytes), or by forming the DT from the lower resolution modality when possible. The first of these two optimizations has recently been utilized in the European COVIRA project on surface registration [2, 1].

Abstractions

A few different search strategies were implemented incorporating hierarchical abstractions of various types. One hierarchy is based on the representation of the data on multiple spatial scales as Borgefors [12] utilized for 2D images. This allows the lower frequency features to be matched first (for initial approximations) and then the higher frequency features can be utilized for fine tuning the parameters. At each hierarchical level, more and more detail of the complete segmentation would be used. Beyond the efficiency gain provided by using a reduced amount of data, the hierarchy may provide more accurate results by ignoring the fine details when gross alignments are being made.

Jiang et al. [41] created different levels of this type by scaling both data sets down to create smaller sets for each level. This required the precomputation of a DT for each scale, and a filtering step to provide accurate features at the lower resolutions. We created the different scale levels by sampling the dynamic set of features (points) at various rates. With a large number of points this probabilistic approach can attain the same hierarchy of scales. Our approach of sampling the dynamic set is more computationally efficient than Jiang's, and allows a single DT of the static set to be used for all levels of the hierarchy.

A hierarchical threshold approach [42] was also incorporated. It requires only the establishment of a limiting value on each DT look up; this will restrict the maximum distance any point will contribute to the comparison function. The focus of the matching process can then be controlled by restricting the matching to localized regions.

The strategy usually starts with a large threshold (maximum distance) allowing gross differences and thus gross alignment, and then decreases the threshold, so that for fine adjustments of the parameters only the regions that match very closely make a difference. The thresholding also allows variations in segmentation beyond a certain limit to be ignored. The basic effect of thresholding is illustrated in the comparison function shown in Figures 5.7 - 5.10. The surface shown in Figure 5.4 was used for both static and dynamic surfaces with intra-slice voxel separation of 1.0mm and inter-slice separation of 2.0mm . The matching surfaces were identical with the exception that the dynamic surface had approximately 25% of its surface translated by -15mm , and another 25% translated by -6.5mm , both in the x direction. The actual distance between corresponding points and the DT comparison function are shown as the dynamic surface is translated by x in Figures 5.7 and 5.8 respectively. In these figures the actual distance between corresponding points can be calculated because the same surface was used to form both dynamic and static sets. The surfaces are registered and so the global minimum should occur when there is no translation in x . In the comparison function without thresholds (Figure 5.8) it can be seen that the error biases the minimum. The threshold value of 10 (Figure 5.9) moves the minimum closer to its optimal position by keeping the effect of the 15mm error constant around the 0 point. The threshold value of 2 (Figure 5.10) provides the global minimum at the correct position of zero, but it can also be seen to create local minimum at the erroneous segmentation locations. Where the errors in segmentation are larger than the threshold value, the effect of these errors will remain constant, and the thresholding will move the global function minimum to a more accurate registration point.

Search Methodology

The surface dependent nature of the search space is too complex for an in-depth analysis. Therefore an iterative gradient search approach was initially adopted. It consisted of discrete steps taken independently for each of the six parameters. The step size was modified and controlled in various manners. Rotational parameters were initialized to zero and rotations

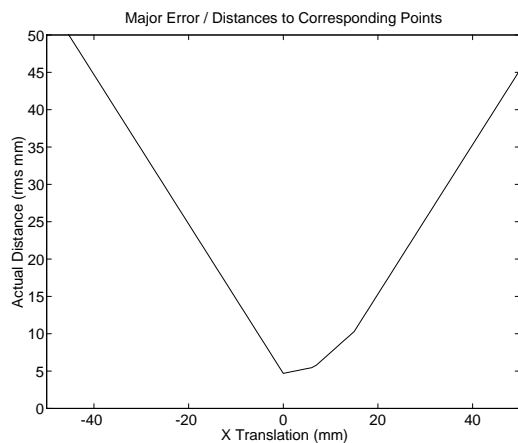


Figure 5.7: Actual Distance

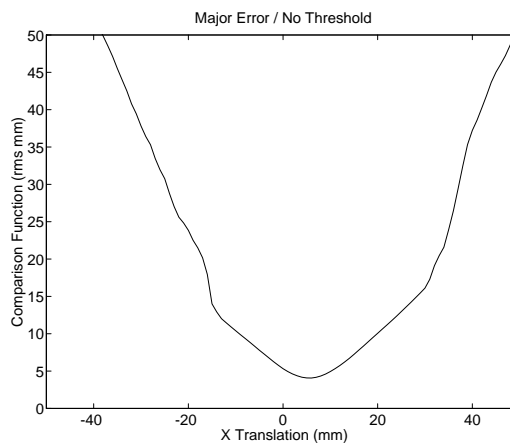


Figure 5.8: DT Function

were performed about the centre of mass so that identical step sizes would not bias the search. With highly irregular objects this may still lead to inadequate rotational steps for some axes. The initial translational parameters were set to align the centres of mass. This was intended to quicken the search process and increase the accuracy, but this result is dependent on the features being globally similar.

The highly complex nature of the search space led to the use of univariate relaxation techniques [65]. Given the parameter state s after iteration k , univariate relaxation completes a next state that can be written as:

$$s^{k+1} = s^k - \alpha_k p^k, k = 0, 1, \dots,$$

where α_k is the step size scaling factor, $p^k = \pm \Delta_{k \bmod 6}$, and basis $\Delta = \{x\text{-axis rotation } (xr), y\text{-axis rotation } (yr), z\text{-axis rotation } (zr), x\text{-axis translation } (xt), y\text{-axis translation } (yt), z\text{-axis translation } (zt)\}$. The gradients were calculated from finite differences using the entire step size. With only descending steps being taken the comparison function (c) satisfies:

$$c(F_s, f(s^{k+1}, F_d)) \leq c(F_s, f(s^k, F_d)), k = 0, 1, \dots,$$

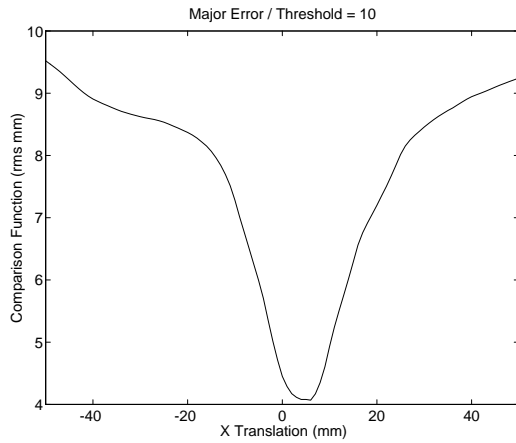


Figure 5.9: DT with Threshold at 10

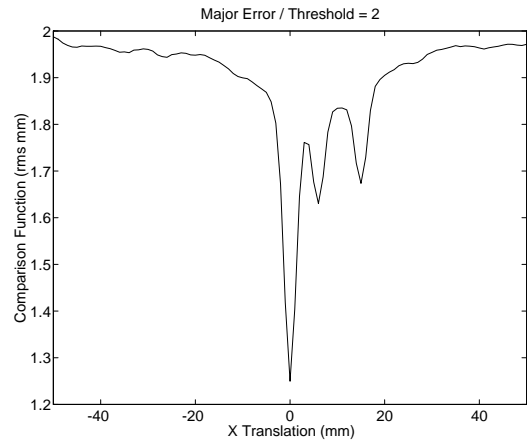


Figure 5.10: DT with Threshold at 2

We formed hierarchical levels in the search after either a local minimum was found, or when a single loop through the parameters resulted in no change in parameters. A level was also formed after a predetermined maximum number of iterations in order to stop slowly converging or oscillating searches. The α parameter was also varied to study its effect and interaction along with the scale and threshold levels. For each parameter the step sizes (Δ_i) were set to an initial value, and at the end of each hierarchical level they were halved.

The basic matching heuristic is presented here in pseudocode:

```

/* Hierarchical Univariate Relaxation */

xt = initial x translation; yt = initial y translation; zt = initial z translation;
xr = initial x rotation; yr = initial y rotation; zr = initial z rotation;
Δxr = x rotation step size; Δyr = y rotation step size; Δzr = z rotation step size;
Δxt = x translation step size; Δyt = y translation step size; Δzt = z translation step size;
αxr = αyr = αzr = αxt = αyt = αzt = step size scaling factor; /* decay term */
FOR (i = 0; i < levels in hierarchy; i++)
{
  Fd = subset(scale(i), Fdynamic); /* sample dynamic features at multiple scales */
  threshold = DTthreshold(i); /* refine maximum DT value */
  δxr = Δxr; δyr = Δyr; δzr = Δzr; /* instantiate step sizes */
  δxt = Δxt; δyt = Δyt; δzt = Δzt;
  converged = false; /* minimum not found */
  FOR (j = 0; j < maximum iterations AND NOT(converged); j++)
  {
    /* x rotation (gradient calculated by finite difference of step size) */
    r0 = c(threshold, Fs, f(xr, yr, zr, xt, yt, zt, Fd));
    r1 = c(threshold, Fs, f(xr + δxr, yr, zr, xt, yt, zt, Fd));
     $\frac{\delta c}{\delta x_r} = r_1 - r_0$ ; /* gradient */
    IF ( $\frac{\delta c}{\delta x_r} > 0$ ) THEN δxr = -αxrδxr /* reverse & decrease step size */
    ELSE xr = xr + δxr; /* step taken */
    /* similarly for y and z rotations */
    ...

    /* x translation */
    r0 = c(threshold, Fs, f(xr, yr, zr, xt, yt, zt, Fd));
    r1 = c(threshold, Fs, f(xr, yr, zr, xt + δxt, yt, zt, Fd));
     $\frac{\delta c}{\delta x_t} = r_1 - r_0$ ;
    IF ( $\frac{\delta c}{\delta x_t} > 0$ ) THEN δxt = -αxtδxt
    ELSE xt = xt + δxt;
    /* similarly for y and z translations */
    ...

    IF (no further steps possible) converged = true;
  } /* one loop through all parameters */

  /* halve the step sizes */
  Δxr = Δxr/2.0; Δyr = Δyr/2.0; Δzr = Δzr/2.0;
  Δxt = Δxt/2.0; Δyt = Δyt/2.0; Δzt = Δzt/2.0;
} /* one level completed */

/* xr, yr, zr, xt, yt, zt now contain the final rigid body parameters */

```

Four basic methods of this type were created, with various modifications:

- Method A

The initial method utilized an independent α for each parameter. It would decrease and invert the step size whenever the gradient was positive. The decay provided by α allowed the iterations to quickly converge to a minimum. A major problem with this method was the existence of many local minima which occurred with various features. In order to help avoid this, the step sizes were reinitialized to half their initial value that after each level. This erases the knowledge that a minimum was found and may allow the search to escape a local minimum.

- Method B

Method B utilized a constant magnitude α ($\alpha_k = 1.0, \forall k$). Therefore the step lengths were only reduced after each level to provide finer and finer adjustments. The initial step sizes were intended to be large enough to enable the search to jump over the local minima. Each level was terminated when no step had been taken during one loop through the parameters. This will often be earlier than when it is actually stuck in a minimum. This was intended to prevent the search from getting to the extrema of a local minimum at the initial levels.

- Method C

In the third technique, a simulated annealing strategy [49] was adopted to deal with the local minimum problem. With this procedure the step size was decreased (by α) with each reversal of direction similarly to method A. At each level in the hierarchy the *reheating* process would occur, whereby the step sizes would be reinstated, and the parameters would be given a random offset. This offset, in addition to the increased step sizes, was intended to solve the difficult problem of allowing the search to escape the concavity of local minima, while not leaving the valley of a global minimum. Similarly to method B, a level was completed when no step had been taken on a single loop through the parameters.

- Method D

Method D incorporated the same technique as method B with the exception of the level termination criterion. It continued the search at each level until it no further decreasing steps could be made (stuck at a relative minimum).

Multiple Searches

Another approach for coping with local minima is the use of multiple starting points. This is not a different method, but is the repetition of one method to avoid sensitivity to initial conditions. This brute force ideology is based on the assumption that if there are enough searches performed in a disjunct set of search extents then the global minimum will be found.

Multiple starting points can also provide a quick approximation of the entire search space. After only a few iterations the relative minima from the different searches can be used to form decision criteria on how to narrow the search bounds. It must be noted that with all hierarchical search strategies a few iterations can only provide an approximation of the search space, and so there exists the danger that the region containing the global minimum may be excluded from further consideration at the early levels.

5.3.3 Verification

Any verification can only represent the approximate overall behavior of the methods. Because the search space is highly dependent on the features segmented, it is possible that the algorithms presented will lead to erroneous alignments in some cases. This can only be overcome by large amounts of simulated and real testing. The accuracy of each implemented registration technique was rigorously tested and compared. The testing was intended to examine the behavior and effect of various adaptations of the matching methods.

In simulation studies the exact misregistration was examined in a large variety of tests. In this way each of the different modifications and methods was tested and compared thoroughly. Then the methods were further tested on actual patient data. The results of all these analyses

are discussed in the next chapter.

Simulation Studies

Rigid body simulation studies were performed on a single surface composed from actual data contours. In this way the correspondence of points was known, allowing the exact precision of the registration to be measured. This simulated transformation was performed on the centre of the data set, while the registration transformations operate on the centre of mass. This fact, as well as the registering rotations being performed in the same order (not the reverse), meant the registering parameters would not be the inverse of those that created the simulation. The complete transformation matrix would, however, be the inverse of the original one (in principle).

Simulations were made more realistic by the use of artificial noise and other distortions. Noise was introduced into the simulation process because the points composing the transformed features were rounded to integers, and also by the addition of random offsets to each point on the dynamic surface. Non-overlapping surface regions were created by using subsets of both the dynamic and static surfaces.

Patient Data Studies

When the true registration transformation was not known, as occurs with patient data, numerous 2D visualization techniques can be used to inspect the results. When the same modalities are registered, the subtraction of identical slice planes provides a good indication of the registration. Other techniques for 2D viewing of multi-modality data are “colorwashing” and “color merging”, which can be easily implemented in WiT. With “colorwashing”, the greyscale values of one image are modified (colored) by the values of the other image. Perault et al. [71] utilize SPECT to red “wash” greyscale CT. A problem with this is the resulting loss of information [23].

The WiT Igraph for “colorwashing” is shown in Figure 5.11. At the top left image A is loaded and sent to the *ALU Operation* (arithmetic logic unit) operator which can perform the pixel addition, multiplication, subtraction, division, maximum, or minimum operations between

images. Image B goes into the *ALU Operation* operator and the red and blue channel inputs of the *rgbMerge* operator. An image formed by choosing the maximum value of corresponding pixels in A and B (Output Image $_{i,j} = \max(A_{i,j}, B_{i,j})$) is output from *ALU Operation*, and this becomes the green channel in the RGB image. Therefore the output image is an RGB greyscale of image B except where image A's pixels (or voxels) exceeds its value, where A will green "wash" B proportionate to A's greater intensity. This can be performed on all the slices in two registered data sets.

A variation on this is the merging of the different images by putting them on different color (RGB) channels. This results in no relative information loss, but may not be as visually informative. Figure 5.12 contains the WiT Igraph for this procedure. This is a similar Igraph to that of "colorwashing", but all the data from each image is used. The *copy clear* operator simply creates a zero valued image the size of the input image. Therefore image A is shown on the green channel and image B on blue. Where the data sets contain equal values a shade of purple results. The choice of channels is somewhat arbitrary, as the red channel could be utilized to replace green or blue. In another method [81], called alternate merging, the two images may be interleaved on alternating horizontal lines allowing comparison. Figures 5.13 and 5.14 represent registered T1 and T2 weighted MRI slices through the head. Figure 5.15 illustrates an alternate merging the two scans in which details from each image can be accurately compared. All these techniques were constructed quite easily in the WiT environment.

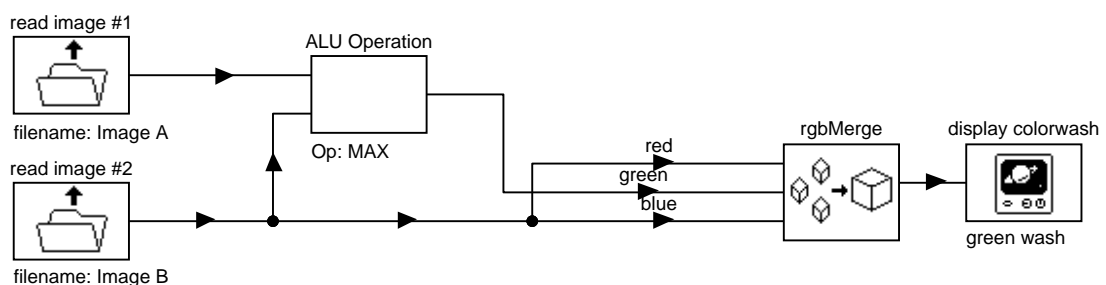


Figure 5.11: The "Colorwashing" of an Image

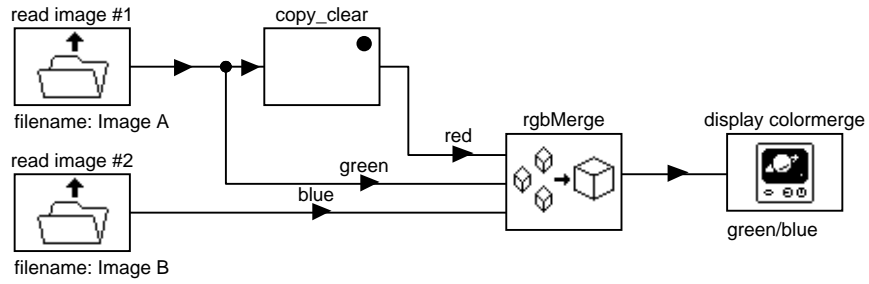


Figure 5.12: A WiT Igraph for “Colormerging”




Figure 5.13: A T1 Weighted MRI Scan




Figure 5.14: A T2 Weighted MRI Scan



Figure 5.15: Alternate Scan Line Merging of T1 and T2 Weighted MRI

Chapter 6

Analysis of Surface Matching

I may now turn my attention to what is the most important subject of all, namely, to the correction of the understanding and to the means of making it able to understand things in such a way as is necessary to the attainment of our end. To bring this about, the natural order we observe demands that I should recapitulate all the modes of perception which I have used thus far for the indubitable affirmation or negation of anything, so that I may choose the best of all, and at the same time begin to know my powers and nature which I wish to perfect.

- B. Spinoza, *The Treatise on the Correction of Understanding*

The techniques and methods implemented in the previous chapter were studied and their results analyzed. Experimental values for the robustness and accuracy of the methods were obtained by using simulations. Patient scans were used as the basis of the simulations to provide a realistic framework. Actual patient data registrations were also performed as a test of true application performance. This chapter presents the details and results of these analyses.

6.1 Accuracy of Segmentation

With surface based approaches the accuracy of the segmentation is a major factor contributing to the success or failure of the registration. Segmentation is, however, another extremely difficult problem, and so could not be thoroughly investigated. Simple segmentation techniques that could be automated were used. This was necessary in order to allow a large number of registrations to be performed. The resulting accuracy of the registrations were reduced, but this was inevitable in the time frame available. The segmentation issue was avoided in many simulations by utilizing the same segmentation for both surfaces, and modifying one or both to mimic different conditions.

6.1.1 Types of Segmentation Error

Errors in segmentation may be categorized in two major groups:

1. Minor Error

Minor error will be used to indicate those errors where the correct structure is segmented, but the boundary between it and surrounding regions is slightly off. Therefore there is an approximate mapping from the region of erroneous segmentation to the correct one.

2. Major Error

Major error exists when a certain region is erroneously classified as another region type or when the magnitude of any error is very large. This type of segmentation can usually not be mapped to the correct one. When this type of error occurs there will not be a corresponding feature in the other segmented data set.

Although these two classes represent different aspects of the same problem, distinguishing between the two will be useful in future analysis. When comparing two segmentations with minor errors there will usually be corresponding features in each segmentation. Major errors are more of a problem because there is often no similar feature to match between sets. Commonly minor errors occur when a single threshold is used to perform binary classification in low resolution images. It is difficult to determine the proper threshold; slight variations of the threshold value will cause a continuous variation in the surface. The global errors are more visible in MRI and CT when an automatic segmentation program has been used to extract finely detailed structures. In this case some regions are either included or excluded erroneously. Minor error is illustrated in Figure 6.1. This PET example is a variation of the segmentation presented earlier in Figure 5.4, which is shown as the lighter line. The minor error was created by changing the threshold value from 95 to 70. This change in the threshold value increased the size of the segmented brain as more background noise was included.

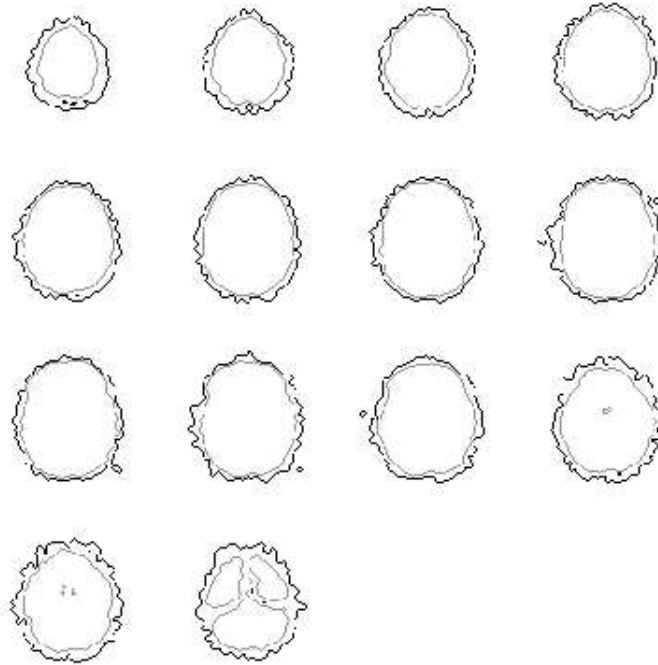


Figure 6.1: Minor Error in Different Segmentations of a PET Scan

6.1.2 Segmentation and Transformations

Unless the segmentation is highly accurate it will introduce biases in the registering transformation. Although this cannot be helped, it can be understood, and it may be used to help decide which type of transformations to search for. Scaling is commonly used in the registration process to correct for errors in imaging device calibration. This is probably the transformation most susceptible to error, because minor errors are very prevalent and they affect the perceived scale. A surface may be segmented consistently larger or smaller than the identical surface in the other data set, resulting in a scaling factor that corrects the segmentation error, but one not based on true spatial scale differences. For this reason scaling factors should be corrected elsewhere, and not included in the transformation.

6.2 Theoretical Analysis

The analyses performed depend to a great extent on the segmentations. In many cases this cannot be avoided. When possible, more general tests were performed. Most of the simulation testing was performed using the typical PET data set shown before. A single segmentation was used for both dynamic and static surfaces (along with various levels of noise and truncation) to enable exact error calculation.

6.2.1 Comparison Function

Each surface was represented as a set of points. The comparison function was chosen to be the *rms* value of the distances to the nearest neighbour. A 1-to-1 mapping of points is not utilized as errors in segmentation as well variations in voxel size violate the properties required for its accurate use. Given two features with N points in each, the distance from a point in one feature to the closest point in the other feature requires $O(N^2)$ operations¹. The segmentations usually contained from around 2000 points in a PET scan to over 50000 in a MRI scan. Therefore a less computationally expensive approximation may be essential. This was initially performed by precomputing the distances for small regions around each static point.

The three-dimensional look up table for a maximum local distance of d_{max} requires pre-computations of $O(N(2d_{max} + 1)^3)$, but then provides an $O(1)$ calculation of the distance to the nearest point. The lookup table will then provide exact Euclidean distance within the local regions. This precomputation is only feasible for small distances of up to a d_{max} of about 10 voxels. These precomputed distances are worthwhile when there are a very large number of points (more than 2000) in each set and when the heuristic requires a great number of iterations (> 50).

This need for more efficient pre-computation of larger distances led us to the Chamfer Distance Transform. Chamfering creates a lookup table (of approximate distances) the entire

¹In a naive approach each point (of N) must be compared to all the points (N) in the other set to find the minimum when a 1-to-many mapping is allowed

size of the data set (M^3), versus the small distances (at most ≈ 10) that were feasible with the brute force method. The algorithm was given previously in Section 5.3.2 and requires $O(M^3)$ operations on a $M \times M \times M$ data set. However, it can calculate exact distances only up to the size of the filter mask used to create the transform. The mask values (d_i) are variants of the local distances. The distance from a voxel to a feature that is beyond the size of the mask will be approximated by the minimum path of 3D-diagonals ($m_3 \& d_3$), 2D-diagonals ($m_2 \& d_2$) and then vertical and horizontal steps ($m_1 \& d_1$). If we assign each coordinate variable to m_i so that $m_3 \leq m_2 \leq m_1$ then for isotropic voxel spacing the error from the actual Euclidean distance is [10]:

$$error = \sqrt{m_1^2 + m_2^2 + m_3^2} - m_3 d_3 - (m_2 - m_3) d_2 - (m_1 - m_2) d_1$$

This expression can be differentiated to find extrema. In Table 6.7 Borgefors summarizes the accuracies of various choices for the d_i values when using a $(M + 1)$ dimensioned cube. However, the data is commonly not isotropic and so must be interpolated to use these values appropriately.

Distance Function	d_1	d_2	d_3	Upper Bound on Error
Chamfer Euclidean	1	$\sqrt{2}$	$\sqrt{3}$	-0.147M
Chamfer Optimal	1	≈ 1.314	≈ 1.628	$\pm 0.104M$
Chamfer (3,4,5)	3	4	5	-0.118M

Table 6.7: Distance Transform Approximations

The other possibility is to use anisotropic Euclidean values for the DT mask. The errors for various anisotropic values were analyzed by performing simulations using a single point feature in the center of a 101 cubic volume and calculating the corresponding DT. An example of the 3D distance transform based on a single point feature and isotropic inter-voxel distances is provided in Figure 6.2. Distances are mapped to a grayscale in which darker pixels indicate smaller distances. The single point feature is at the center of the volume, and the variation from perfect circles reveal the approximation errors.

The errors resulting from using various anisotropic mask values are summarized in Table 6.8.




Figure 6.2: A 3D Chamfer Euclidean Distance Transform (101x101x101)

In these cases the d_i values forming the mask are the exact anisotropic Euclidean distances. The maximum and average relative errors are dependent on the ratio of the intra-slice and inter-slice dimensional units, and not on their absolute value. Therefore the ratios 2:2, 3:3, ... result in the same errors as shown for the 1:1 ratio and the entries are not repeated. The maximum error as a fraction of the actual distance was found by comparing all values in the computed DT lookup table to the actual distance to the center point, and similarly the average error reflects all the values in the entire DT.

From this table it can be seen that the use of the exact anisotropic Euclidean distances does not degrade the performance of the chamfer DT to a large extent. Although the maximum relative error grows large (up to 35%) for highly disproportionate voxel spacing (1:7), the average relative errors remain quite small (around 9%). These values were calculated from a single point feature. It has been shown that the maximum error occurs at relatively large

Dimensional Units		Max	Avg
Intra-slice	Inter-slice	<i>Error/ActualDist.</i>	<i>Error/ActualDist.</i>
1.00	1.00	0.128093	0.081525
1.00	2.00	0.199794	0.087848
1.00	3.00	0.254452	0.094127
1.00	4.00	0.292891	0.096605
1.00	5.00	0.320485	0.096638
1.00	6.00	0.341011	0.095288
1.00	7.00	0.356817	0.092987
2.00	3.00	0.165241	0.083803
2.00	4.00	0.199794	0.087848
2.00	5.00	0.229514	0.091466
2.00	6.00	0.254452	0.094127
2.00	7.00	0.275336	0.095694
3.00	4.00	0.152918	0.082638
3.00	5.00	0.177242	0.085087
3.00	6.00	0.199794	0.087848
3.00	7.00	0.220171	0.090387
4.00	5.00	0.146692	0.082184
4.00	6.00	0.165241	0.083803
4.00	7.00	0.183068	0.085794
5.00	6.00	0.142946	0.081973
5.00	7.00	0.157888	0.083082
6.00	7.00	0.140448	0.081839
All values calculated from a single point feature at the center of a 101x101x101 dimension volume.			

Table 6.8: An Error Analysis of the Anisotropic Euclidean Chamfer DT

distances [10, 11]; therefore when a large number of point features are distributed throughout the volume the errors will be smaller.

In an application where real numbers result, the distance values can be formed from the anisotropic transform by scaling each point coordinate into DT space and then rounding off to form the lookup table index. This rounding is an approximation that will introduce error, but when the points lie continuously between voxels (when rotated out of the DT planes) it should statistically provide the same results as if the DT, features or data were interpolated into smaller units. With highly disproportionate dimensions the rounding error may have to be reduced by interpolation of some type, and this will be discussed later.

Other existing techniques could be utilized to calculate the distances for the comparison function. One way in which “on the fly” distances calculations could be improved is by the use of an octree partitioning of the points. This would result in small groups of points that would not suffer drastically from the $O(N^2)$ search. It would also not be limited by the finite nature of the precomputed comparison function. Bouts [13] presents an optimized Euclidean distance transform based on a single propagation of a “distance wavefront” through the volume. This is an area for possible future investigation.

6.2.2 Bounds of the Parameter Space

The parameter space is the set of all possible transformation parameters in which the registering transformation may exist. By assuming rigid body motion we have 6 parameters: 3 for rotation, and 3 for translation. The parameters for rotation of the dynamic surface about each axis may be completely arbitrary in theory, anywhere in the range of 0 to 360 degrees. This can in all practicality be reduced to a small range of values by visually inspecting the data sets, unless the registration is to be performed completely automatically. The implementations created are intended to be semi-automatic. They rely on the assumption that the parameter space has been reduced. A similar case can be presented for the translation parameters which must be contained in $[-XSize, XSize]$, $[-YSize, YSize]$, $[-ZSize, ZSize]$ of the larger data set for each respective dimension, while the average translations needed for registration will usually be much smaller, on the order of 1/20 of these ranges.

6.2.3 The Nature of the Search Space

The search space is the comparison function value over the entire parameter space. It is dependent on the features and is highly complex. Its six dimensional nature makes it too large to be completely analyzed, even for simple features. With identical static and dynamic features the space can be shown to have a global minimum at the true registration point, unless there are symmetries in a feature, which will then introduce periodicities into the search space. Although

perfect symmetries will almost never exist, approximate ones may produce similar effects.

The epidermis of the head and brain surface are approximately ellipsoidal, and this results in an approximately periodic search space. Using the PET surface shown previously in Figure 5.4 (with distances of 2.0mm inter-slice and 1.0mm intra-slice) this phenomenon is revealed in the one dimensional plot of the search space shown in Figures 6.3 - 6.8. The function is calculated using the single surface for both dynamic and static (the DT) surfaces. These figures represent the function when all other parameters are zero. Thus when the rotation parameter is zero or 360 degrees, or when the translation parameter is zero, they are exactly aligned. Because the static and dynamic surfaces are the same, corresponding points are known and can be used to determine the actual distances. In all these figures the dashed line represents the actual *rms* distance between corresponding points and the solid line is the distance transform approximation used for the comparison function. The actual distance lines are sinusoidal for rotation as the points follow the rotational path up to the furthest point (where the data set is completely flipped about one axis) and then return back to the original position. The translational actual distance lines are linear with the displacement, as expected.

A periodic search space poses a large problem to the gradient search methods as they produce the deep local minimum which are difficult to distinguish from the true minimum. Although the chamfer distance transform over estimates the actual distance for a single point, it can consistently underestimate the true distance for corresponding points without a 1-to-1 mapping of points. This can be seen in the 0-45, 135-235, and 315-360 degree ranges of Figure 6.3. In the other ranges the comparison function reveals discontinuities where the function begins rapidly increasing towards peaks. With rotation about the X axis beyond 50 degrees, the large Y dimension of the surface comes into play as a large region of the surface leaves the region of the finite DT. As the points leave the finite region of the look up table they are assigned the maximum value of the entire distance transform. This shows a major limitation of using a distance transform, or any other finite function. The function is biased towards keeping the surfaces contained within the look up table, and therefore is only applicable when the dynamic

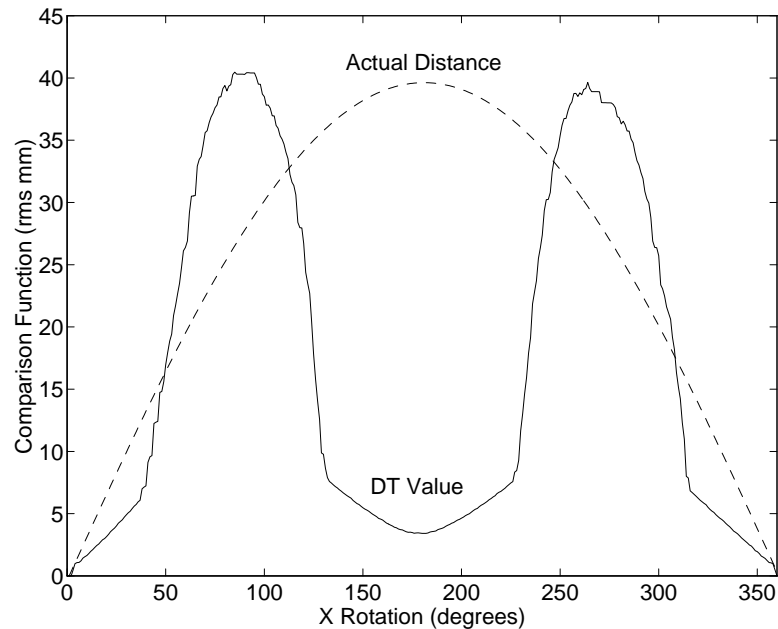




Figure 6.3: Comparison Function Over X Rotations Only ($c(F_s, f(rot_x, F_d))$)

surface is contained within it upon registration.



Figures 6.4 and 6.5 represent plots of the Y and Z rotational aspects. With Y rotation the function suffers slightly from the effect seen in the X rotation plot, but is significantly less dramatic because it is the major axis of the ellipsoid, and so the points move less during rotation. The Z rotation plot reveals none of the effects of the finite approximation. The X and Y translational plots in Figures 6.6 and 6.7 reveal the finite nature of the transform after $20mm$, however these large translations would almost never be needed. Within the Z translation plot of Figure 6.8, the finite effects are apparent before $20mm$, and are magnified due to the lower resolution of this dimension ($14 \cdot 2.0mm$ compared to $100 \cdot 1.0mm$ for X and Y), but again these effects appear beyond the normal working ranges needed.

Another problem, which is visible in the Z translation plot, is the large inter-slice distance. When the step size ($1.0mm$ in these examples) is below the size of the DT units ($2.0mm$ inter-slice) there are flat regions in the search space. For Z translation this effect will only be large

Figure 6.4: $c(F_s, f(rot_y, F_d))$ Figure 6.5: $c(F_s, f(rot_z, F_d))$

when the dynamic features are parallel with the DT planes (no X or Y rotations), so all dynamic points in a plane are rounded to the same plane. This can introduce variance in the solution of up to half the inter-slice distance. There are two methods to deal with this, both involving interpolation. The first method is to interpolate the data, or segmentation [41], into cubic or at least less disproportionate distances. The other is to use interpolation on the anisotropic DT [1]. The interpolation of the DT requires added computation on each calculation of the comparison function, while the interpolation of the data requires greater precomputation. Either of these techniques will eliminate the problem, when there is no rotation. The DT was initially utilized without either of these modifications in order to test its performance, but the registration of two data sets both of which have large inter-slice separation and little rotation will require the use of some form of interpolation. These forms of interpolation are analyzed in Section 6.5.2.

These examples have shown some of the weakness inherent in the distance transform, but do not reveal the local minimum problem that may occur with various features. The local minimum problem is discussed by Jiang et al. [41]. Local minima often appear when major segmentation errors create false matches and when the features contain periodic shapes which can be aligned in multiple positions.

Figure 6.6: $c(F_s, f(trans_x, F_d))$ Figure 6.7: $c(F_s, f(trans_y, F_d))$

6.3 Simulations: Transformed Real Features

Arbitrary rotation and translations were applied to surface data to create simulations of data requiring registration. These were then varied by using only subsets of the surface and adding random noise in a manner similar to how Jiang et al. [41] created their simulated data. A single surface was used as the basis for both the static and dynamic surfaces; this allows the exact misalignment of identical points to be calculated. This is the most quantifiable method available for testing the process. In this way, a highly controlled and flexible test bed was made, in which the correct result is known to great accuracy.

The surface was formed from a single PET scan of a normal patient. It consists of 14 slices of 100x100 dimensional data. The inter voxel distances were taken to be 1.0mm intra-slice and 2.0mm inter-slice to simplify and clarify the results. The 1:2 ratio is a good approximation of the actual ratio of 4:7.14 and provided a realistic brain surface when viewed in an orthographic projection. The data and the extracted features were shown in two dimensions in Figures 5.3 and 5.4. The contours, which consist of 2290 points, were used for the following analysis. These original contours were used to form the floating point chamfer DT shown as an array of 30 slices




Figure 6.8: Comparison Function Over Z Translation Only ($c(F_s, f(trans_z, F_d))$)

in Figure 6.9. The addition of an eight voxel border in all dimensions was utilized to create a more stable search environment (116x116x30). This is done so that the finite nature of the DT does not bias the search when a transformed surface extends slightly beyond a boundary of the imaging data. In this gray-scale representation of the 3D lookup table the darkest regions are closest to the static features. The 14 slices in the middle contain zero-valued representations of the surface in 5.4, but they are complicated as the distances from neighboring slices are propagated throughout the DT formation.

6.3.1 Methods

The four different search methods described in Section 5.3.2 were compared in their basic forms and with modifications. Each method was first compared with identical step lengths, a constant 0.5 fraction of the dynamic contour points, and no thresholding. The parameters were individually tuned to provide appropriate responses for each heuristic. All methods were based




Figure 6.9: A Floating Point Anisotropic Chamfer Distance Transform

on five hierarchical levels of gross step length adaptation, to keep as much constant as possible across techniques.

Initial testing was performed without random noise and both the surfaces were exactly the same, except for the round-off after the simulated rigid body transformation, which averages a quarter voxel in each dimension for each point. The limit of the registration accuracy would be expected to be the error for round-off, which is $0.612mm$ rms. The averages are compiled from 729 tests, which consist of rotations from -40 to $+40$ degrees with 10 degree increments for all axis. Any simulated translations would be removed by the initial alignment of the centres of mass, but the search still included the translation parameters. The searches for all of the 729 transformations began with initial rotation values of zero, and the centre of mass alignment provided the starting values for translation. The results are analyzed in the Table 6.9. It shows the averages (\bar{d}) of all the misalignment values (d_i) and the standard deviation (σ_d) of these

values. This misalignment was calculated as the *rms* distance between all corresponding points on both surfaces after registration. Because the distribution was not normal, the percentage of all the tests in which the distance (d_i) was below some stated values is also provided.

Method	\bar{d} (rms mm)	σ_d (rms mm)	%d < 1.0	%d < 2.0	%d < 5.0	%d > 10.0
A	0.73	1.15	99.4	99.7	99.9	0.1
B	0.90	0.82	86.3	97.0	99.2	0.1
C	0.94	1.42	89.2	96.0	99.2	0.6
D	0.71	0.19	97.3	99.5	100.0	0.0
Averages computed from 729 trials.						

Table 6.9: Comparison of Basic Methods Without Noise

The results of this testing show that with very little noise in the dynamic set, each method provided good results. The results show that methods A and D, which continue until they get stuck, provide better results. This was expected with the “error-free” data, but this may lead to the finding of more false minima when features have errors. The difference between the \bar{d} values of method A and method D is not significant enough to make any general conclusions, but method D was much more consistent, probably due to its constant step size (α) value. These statements are also applicable with regards to methods B and C, with method B being the more consistent for similar reasons. The percentages reveal that the standard deviations are not very important in regards to accuracy, as A has the most solutions below 1.0mm and 2.0mm yet the second highest standard deviation. Thus the standard deviation should only be used as a measure of the consistency of the approach.

6.3.2 Thresholds

For the analysis of different thresholds, method B was utilized. It provides natural stages for refining the threshold after each hierarchical level. The more simplistic nature of method B, as compared to the dynamic step sizes of method A and the random offsets of simulated annealing, provides a more stable means for the independent judgment of the thresholds.

The types of segmentation error will influence the results obtained by using thresholds in the search process. Global errors can be best accommodated by the thresholding process. With a small threshold, outliers and large errors can be ignored. This allows closely matching features to be used optimally for fine tuning. This has been examined thoroughly elsewhere by Jiang et al. [41]. However, a problem may occur at the transition between the minor and global errors. If minor errors dominate, the threshold value must remain larger than the magnitude of these errors, otherwise the thresholding will be throwing away relevant information. In other words, the threshold should remain above the mean error in the segmentation. Although this cannot usually be calculated, it can often be estimated. This places a lower bound on a suitable threshold value.

The effect of a threshold below the mean segmentation error is shown here in Figure 6.10, illustrating a Z rotation plot of the search space. The static surface is the example segmentation (from Figure 5.4) and the dynamic set is the example of minor error presented in Figure 6.1. For




Figure 6.10: The Effect of Thresholding with Minor Error

this plot the three threshold values of 7, 5, and 4 reveal the effects. As the threshold decreases below the magnitude of the minor error (between 5 and 4) the comparison function becomes less accurate, to the extreme where the true registration point becomes close to a maximum.

For a general assessment of threshold values only random noise was added to each point in

the dynamic surface. This provided minor error, but due to its random nature, thresholding would not eliminate too much data. On each of the 5 levels of the hierarchy various thresholds were compared. Testing was performed with a dynamic surface created by adding uniformly distributed random noise of $\pm 2.0mm$ to the static surface, which was also affected by round-off errors after the simulated rigid body transformation. The results are analyzed in Table 6.10 with the 729 simulated transformations identical to those in Table 6.9. The resulting actual misalignment in the registered surfaces was calculated by transforming the noise-free surface by the simulated transformation, then using the registering transformation to align them, and finally computing the *rms* distance between corresponding points. These results indicate that

Thresholds	\bar{d} (rms mm)	σ_d (rms mm)	%d < 1.0	%d < 2.0	%d < 5.0	%d > 10.0
8 6 5 3 2	0.90	1.13	88.8	97.5	98.9	0.55
constant 7	0.90	0.82	86.3	97.0	99.2	0.14
8 7 6 5 4	0.94	1.42	89.2	96.0	99.2	0.55
max 12 8 3 2	1.01	1.46	88.5	94.9	97.7	0.67
max 10 4 2 $\sqrt{2}$	1.09	2.04	88.9	95.5	96.8	1.51
none	1.48	2.08	56.8	88.6	96.2	1.23
Averages computed from 729 trials.						

Table 6.10: General Comparison of Threshold Values With Noise

the use of thresholding can allow the search to more precisely determine the correct registering transformation. The constant threshold (7) provided close to the lowest misalignment error but provided fewer solutions within $1.0mm$ than all other threshold options. These results are feature dependent and only indicate that multi-level thresholding can increase the accuracy.

6.3.3 Hierarchies of Scale and Thresholding

An analysis of thresholding and the use of multiple scales was performed by running the different algorithms with and without the use of each. The data used was again the standard PET data surface used throughout this section. The results for each method are provided in Table 6.11. These tests were run with uniform random noise of $\pm 3.0mm$ in all dimensions in addition to

the round-off error mentioned previously on the whole surface. Only 60% of the two matching surfaces was identical. This was done by taking the first 70% of the points (taken top to bottom and left to right from the data set) from the base surface for the dynamic set and excluding the first 10% of the surface for the static set. This data was again formed from 729 trials with the same transformations as used previously. The accuracies are further compared in Figures 6.11,

Method	Scales	Thresholds			
	Multiple Fractions	Yes \bar{d} (mm)	No \bar{d} (mm)	Yes σ_d (mm)	No σ_d (mm)
A	Yes	1.78	2.14	1.14	1.21
	No	1.72	2.49	0.54	1.49
B	Yes	2.05	2.43	1.51	1.13
	No	1.97	3.00	0.75	2.14
C	Yes	2.37	2.63	1.43	1.19
	No	2.24	3.16	0.80	2.02
D	Yes	1.81	2.21	1.11	0.83
	No	1.75	2.59	0.59	1.70
Averages computed from 729 trials. Scales: Yes = {0.1, 0.2, 0.4, 0.8, 1.0} No = 0.5 Thresholds: Yes = {7, 6, 5, 3, 2} No = none					

Table 6.11: Comparison of Thresholding and Hierarchies

6.12, 6.13, and 6.14, showing the percentage of trials that converge to within specified ranges of the optimal solution. The solid line represents the use of thresholding and multiple scales, the dashed only thresholding, the dotted only multiple scales, and the dotted and dashed no thresholding or multiple scales.

The histograms for method B utilizing thresholding in conjunction with multiple scales are shown in Figures 6.15 and 6.16. The non-normal distribution of actual errors is revealed in this histogram of d_i values. The histogram in Figure 6.16 shows the minimum values of the comparison function and reveals a somewhat random distribution of final values, indicating that the differences between a good and poor registration in the comparison function are not very distinct. Therefore it would be difficult to determine from the comparison function alone whether the final solution is very accurate.




Figure 6.11: Accuracy of Method A




Figure 6.12: Accuracy of Method B

The results of this table and the figures indicate strongly that the use of thresholding increases the accuracy of the found solution. With every method the use of thresholds reduced average distance error in the solutions. The use of thresholding reduced the standard deviation when multiple scales were not used, but increased it for methods B, C, and D when they were used. This may be due to the fact that the changes between levels, in both scale and threshold, destabilize the search. The results are dependent on the single set of surfaces, but similar results were observed with various levels of random noise. It must be pointed out that random noise is different from the minor error most commonly found in segmentations, in that even small thresholds will not exclude accurate data. Therefore in real applications care must be taken to set the smallest threshold above the mean segmentation error. Major errors are usually handled appropriately by thresholding [41].

The analysis of using multiple scales does not provide a clear judgment. The features did not contain many frequency levels, which may be why there was no increase in accuracy. Further analysis with other features would be required to make a decision on this question. The use of the scales does provide much faster searches at the initial levels (due to less points). This is very useful if a large number of searches will be carried out. If most of the searches are




Figure 6.13: Accuracy of Method C




Figure 6.14: Accuracy of Method D

terminated early, there will be significant time savings from the use of multiple scales.

6.3.4 Distribution of Starting Points

The search methods are all initiated at specific starting points. The human operator may provide the starting point, or some automatic strategy may be used. The proximity of the initial transformation estimate to the optimal solution will affect the searches' success.

For this reason, an analysis of various starting positions was carried out with method B. The results are provided in Figure 6.17. For each of the data points five hundred searches were performed with rotational starting positions placed randomly within the stated range (0-X). This was performed with the same two surfaces as used previously, having 60% overlapping regions, and $\pm 2.0\text{mm}$ added random noise.

The percentage of these starts that converged to within actual point correspondence of 1.0 and 2.0mm rms are charted. This reveals the increased chances of being trapped in local minima the further away from the true registration point you begin. These effects will be amplified with more distorted features.

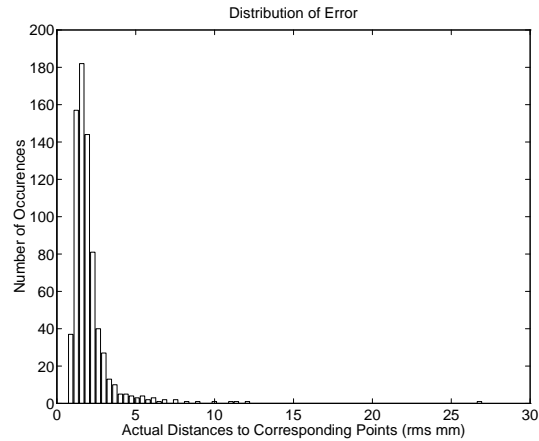


Figure 6.15: Histogram of Actual Error Testing Method B ($\bar{d}=2.05$, $\sigma_d=1.51$)

6.3.5 Multiple Searches

Multiple starting points can be incorporated in many ways. The simplest is the repetition of complete searches started at different positions. From the results provided in the previous subsection, this has a high probability of increasing the accuracy of the result. This is because by using more starting points you increase the odds of one being closer to the true solution, and so there is less chance of getting stuck in a local minima.

A sample test of this was performed for 27 searches using Method B for each of the 729 simulated registrations. The searches were started with the initial rotations of -20 , 0 , and 20 degrees for each axis, and with no variation in translational starting values. The surfaces had $\pm 2.0\text{mm}$ random noise in each direction and 60% overlapping regions. Results are given in Table 6.12.

Number of Searches	\bar{d} (rms mm)	σ_d (rms mm)	%d < 1.0	%d < 2.0
1	2.05	1.51	4.12	64.75
27	1.66	0.52	5.90	78.33
Averages computed from 729 trials.				

Table 6.12: Analysis of Multiple Searches

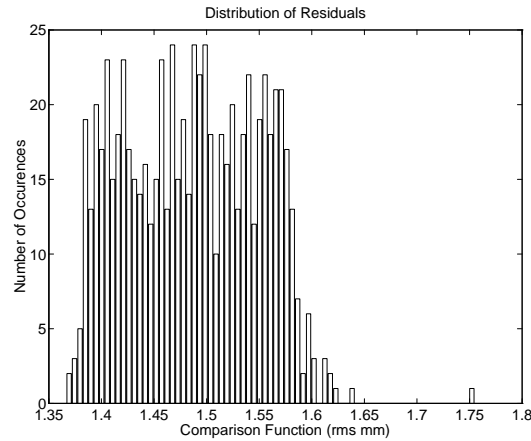


Figure 6.16: Histogram of the Comparison Function Testing Method B

The table indicates that the use of multiple searches is an easy way to increase accuracy. The price is the multiplication of the search time. This could be reduced by terminating some of the worst searches at early hierarchical levels. In this way enough searches can be started throughout the search space so that no user-specified starting position or estimate is required. The computation will be wasted if the searches converge to the same basins, but this is difficult to know beforehand as the space is feature dependent.

6.4 Simulations: Transforming Real Data

In this section we examine simulations that transformed an entire data set and then resampled it. This was done to create new data sets to be registered. The segmentation phase was then performed on each data set, as would be the case with real data.

Using the registration transformation the dynamic data is then resampled along identical slice planes as the static data. This allowed visual verification techniques to be utilized. The WiT Igraph for this process is shown in Figure 6.18.

For these simulations method D was used in conjunction with multiple thresholding and scale levels. It was chosen because of its superior performance in many of the previous studies. The thresholds values were the ones used earlier: 7, 6, 5, 4, 2, and the scales were 0.1, 0.2, 0.4,




Figure 6.17: Search Sensitivity to Initial Parameters

0.8, 1.0 for the respective levels. The results of some of these simulations are summarized in Table 6.13.

Due to the nature of the simulations, any large rotations about the x and y axis result in a clipping of the data which composes the surface used for registration. For this reason the rotations were kept small for these variables. The results of these tests show that the method performed quite well considering the circumstances. The clipping of the surface just mentioned creates major segmentation errors, and this can be seen in Figure 6.20. An accuracy, whereby the surfaces are off by an average of one voxel per point, results in the average error of $2.45mm$ and for an average of two voxels this becomes an error of $4.90mm$. These values provide a context for judging the results, although they do not provide a clear cut evaluation of whether the method's performance was acceptable.

To provide a better picture of these simulations the data representing the first of the simulated trials and its segmentation are provided. The static data is the PET scan shown previously

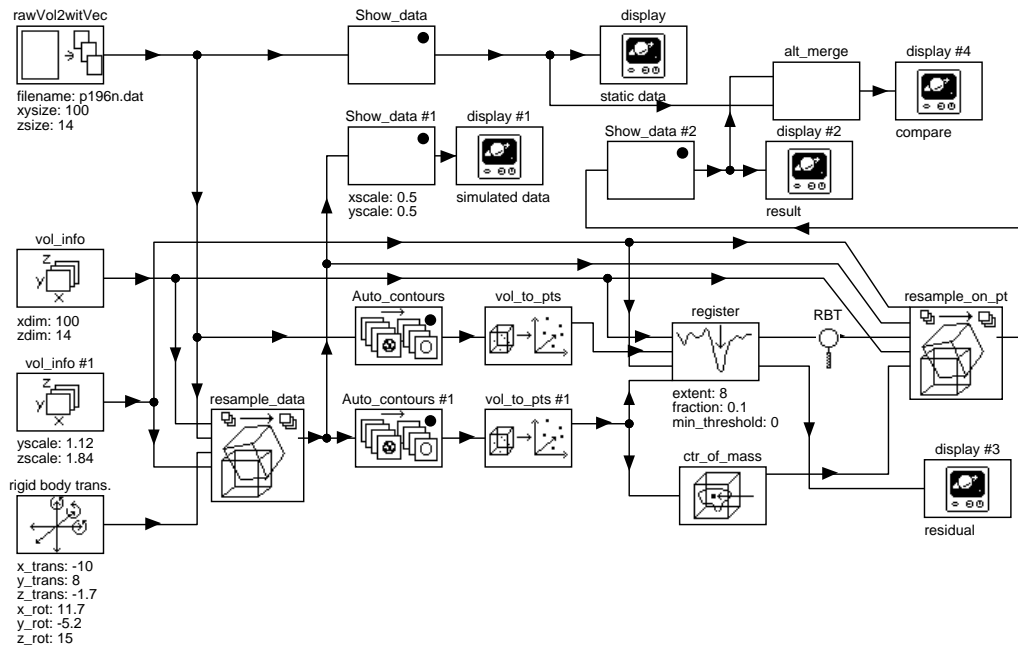


Figure 6.18: A WiT Igraph for the Registration of Simulated Data

in Figure 5.3 and its segmentation in Figure 5.4. The simulated data and its segmentation are shown in Figures 6.19 and 6.20.

The results of registration are shown in which the simulated data has been resampled along identical slice planes as the original using the registering transformation in Figure 6.21. An image in which the registered and static data sets are subtracted to reveal differences is provided in Figure 6.22. The subtracted image is a signed image so that the bright and dark regions represent the most error, but also note the relative differences apparent in a circle outside the brain. This is created by the variations in the inherent background noise alone. The image values have also been scaled to amplify the differences for visual clarity, and the highly visible differences are due to the resampling and clipping of the dynamic data set during this simulation.

Transformation						Actual Distance		Comparison Metric
Translation (<i>mm</i>)			Rotation ($^{\circ}$)			d	σ_d	
x	y	z	x	y	z	(rms <i>mm</i>)	(rms <i>mm</i>)	(rms <i>mm</i>)
-4	-8	-1.7	8	0	15	2.96	0.54	1.28
5	-9	-2.3	3	-6	10	2.18	0.29	0.91
3	12	3.8	-8	2	-6	2.58	0.20	0.87
-7.4	-15	3.1	-2	9	-10	3.80	0.76	0.89
-2.3	4.6	-1.7	1	-3	5	1.85	0.64	0.78
Dynamic Scan Created by Resampling Static Scan								

Table 6.13: Analysis of Simulations Using Transformed Real Data

6.5 Patient Data Studies

The methods were tried on a number of patient data sets requiring registration. In these cases the exact transformation is not known. Therefore the results had to be visually checked. The results of these studies can only be verified to have results to be within approximately 3 degrees for rotation and 2 pixels for translation as revealed visually [73].

6.5.1 Registering to Norms

The registration of an individual patient's data to a "normal" patient can provide an improved juxtaposition when comparing the two scans. This can help to reveal abnormal variations and improve the visualization of relative positions. For these reasons the registration of PET data of different patients was performed.

The data registration was performed on numerous data sets. The results of these tests were promising. The variation of each patient precluded accurately quantifying the error. The patients' brain sizes, as well as various levels of radiopharmaceutical uptake provided both minor and major errors. When the segmentations accurately delineated the brain surface, it was possible to register the data into a common position.

For these tests method D was utilized. The spatial scales: 0.1, 0.2, 0.4, 0.8, 1.0, were used with thresholds from a maximum of 7*mm* down to a minimum of 4*mm*. The larger minimum value was used to compensate for the larger anatomical differences present between people. A




Figure 6.19: A Simulated Scan

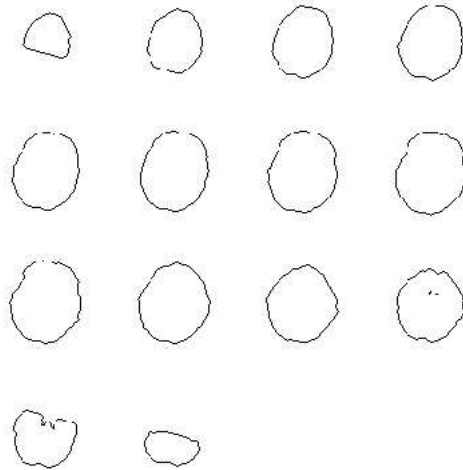


Figure 6.20: Segmented (Dynamic) Surface

step size of 8° was used for all rotations, and $8mm$ for all translations. The use of an 8 voxel border was also incorporated.

The results of two cases are provided in Figures 6.23 & 6.24 and 6.25 & 6.26. These two data sets were registered to the normal patient used throughout the testing process (see Figures 5.3 and 5.4). For the first data set the optimal registering transformation required was $-2, 0.5, 3.5$ for the $x, y,$ and z translations respectively, and no rotations. The second data set required $-2.5, -7.5, -4$ for translation, and -16 degrees rotation about the x -axis. In comparing the normal data to the registered ones, the ventricles and mid plane can be seen to be in close alignment, despite the anatomical variations.

6.5.2 MRI-MRI Registrations

The registration of 2 MRI scans was performed. The two scans were already registered: the T1 and T2 weightings of a single MRI scanning session. Because of this the resulting transformation should be the identity (0 translation and 0 rotation), at least if an accurate segmentation of




Figure 6.21: Result of Registration




Figure 6.22: Difference of Static and Registered Sets

each scan had been provided.

The first scan and its segmentation are shown in Figures 6.27 and 6.28, and the second scan and its segmentation in Figures 6.29 and 6.30. The scans are axial views through the top of the head. The 20 slices have inter-voxel distances roughly of $1mm$ in-plane and $8mm$ inter-slice. This illustrates the problem with largely disproportionate spacing of voxels discussed in Section 6.2.3 on the nature of the search space. This is revealed in the resulting registration in which various solutions to the problem were compared. They all represent the same search configurations of method D that were used with the registrations in Section 6.5.1. The only variation is that a minimum threshold value of $2.0mm$ was used. The results are summarized in Table 6.14.

The interpolation of the data required added time for calculation as well as segmentation. Feature interpolation is possible but it was not available. The interpolation of the static surface doubles the size of the DT, and so is computationally expensive. Similarly the interpolated dynamic surface increased search time, although it would be possible to use reduced fractions of




Figure 6.23: PET Scan "X"




Figure 6.24: "X" Registered to Normal

points to eliminate this added time. The interpolation of the DT itself for only the z dimension requires little added computation, and no increase in pre-computation. For this reason, as well as the fact that it provided the most accurate registration, it is probably the best solution. It can be utilized for increased accuracy even when the data is isotropic.

DT Technique	Inter-slice Distance		Transformation					
	Static (mm)	Dynamic (mm)	Translation (mm)			Rotation ($^{\circ}$)		
			x	y	z	x	y	z
normal	8.0	8.0	0.5	0.0	2.75	0	0	0
normal	4.0	8.0	-0.5	0.5	0.0	-0.5	0	-0.5
normal	8.0	4.0	0.0	1.0	-2.0	0	0	0
normal	4.0	4.0	0.0	1.0	-2.0	0	0	0
interpolated	8.0	8.0	0.0	1.0	0.0	0	0	0
Inter-slice distances of 4.0mm were created by interpolating the data and then segmenting.								

Table 6.14: Analysis of Interpolation Techniques




Figure 6.25: PET Scan "Y"




Figure 6.26: "Y" Registered to Normal




Figure 6.27: T1 Weighted MRI Data




Figure 6.28: T1 Weighted MRI Segmentation




Figure 6.29: T2 Weighted MRI Data




Figure 6.30: T2 Weighted MRI Segmentation

Chapter 7

Conclusions and Future Work

Roseau pensant. - Ce n'est point de l'espace que je dois chercher ma dignité, mais c'est du règlement de ma pensée. Je n'aurai pas d'avantage en possédant des terres: par l'espace, l'univers me comprend et m'engloutit comme un point; par la pensée, je le comprends.

- B. Pascal, *Pensées*, Appendix to Chapter XXIII

The analyzed implementations provided a good method for the registration of multimodality medical images. Surface matching allows the data sets, in which no highly discernible corresponding points exist, to be matched. Tolerance of small errors from the averaging of a large number of points can allow segmentation without the expert knowledge of a radiologist. These results were experimental, and a large amount of work remains in extending and optimizing the implemented procedures.

7.1 Experimental Conclusions

The segmented surface features can be compared efficiently and simply with the use of an anisotropic chamfer distance transform. This provides accurate results as long as the transformed feature remains entirely within the discrete distance transform. With this comparison function, and the implemented gradient descent methods, the registering transformation can be calculated.

These results were checked for accuracy through numerous experiments using both simulated data and actual patient data. The simulation experiments used numerically based verification, whereas various visual tools were used for the patient data verification. The fully automatic procedure usually resulted in success below the visual threshold for error.

The use of multiple starting points was not examined thoroughly enough to determine precise guidelines for their use. The feature-dependent nature of the search space necessitates the use

of some higher level techniques for a more accurate understanding. Sparse approximations of the search space can be calculated. If the space is well behaved, a case-specific distribution of starting points can be determined.

Analyses of multiple spatial scales and thresholds were not rigorous enough to provide all-encompassing generalizations. Abstractions of this type are very case dependent and should be used with caution. Multiple thresholds resulted in more accurate registrations in the simulations. This should also occur with the real data. The effects of using multiple scales were not as obvious, but this does provide an efficient way to manage a large number of searches. The initial stages of a hierarchical search can be performed on the lower resolution data efficiently with little or no degradation in accuracy. It may be possible to analyze the selected features to provide appropriate calibration of the threshold and scale levels.

7.2 Extensions and Future Development

Although the search methods provided good results, there exist many areas for improvements and other work. These are possible routes that future work might take:

1. Numerous other techniques could possibly be used to incorporate knowledge of the problem into the search heuristic.
2. It would be possible to modify the implementation to provide a more efficient method in terms of computer time and memory requirements. Accuracy reduction in exchange for faster performance would require justification that should be studied in further detail.
3. A more quantitative analysis of the discussed gradient search methods and modifications requires the establishment of a performance benchmark. This could possibly be provided by the use of a numerical analysis minimization procedure such as the FORTRAN FSQP [101] routine. This mathematical library subroutine solves constrained non-linear systems. It would provide a good comparison for the optimization routines discussed here.

4. A rigorous comparison of surface matching to other current techniques could be performed. This should provide guidelines for the determination of which method or methods are possible for individual cases, and their chance of success. The use of different methods will also provide a very good check on the accuracy, which is still extremely difficult to measure in actual cases.

7.3 Applications of the Implemented Techniques

Future work will lie in the practical application of these methods in a medical research environment. This will provide further testing with real data. More diverse applications will also provide insights into the optimization and modifications needed to promote these techniques in a clinical setting.

The tools and methods used in these implementations can be utilized for further research. Various research groups are interested in registration for aiding and improving their work. The simplicity of the techniques and the use of the WiT environment should allow easy application development.

One application is the registration of SPECT and CT to help in the quantitative analysis of SPECT data. This would facilitate the work of a research group at Vancouver General Hospital. The added information from the registered data could aid in forming more accurate reconstructions and in segmentation of various regions.

Another application is the registration of PET scans of different patients to a normal patient to improve the analysis and visualization of variations. Research in averaging normal patients to further develop criteria for detecting abnormal variations is being performed. This could benefit from the registration of different normal patients into one “standard” position to reduce the variance in the average data set.

After future progress with the research applications, the registration may be utilized in a clinical setting. The registration of CT, MRI, SPECT, or PET could aid in a variety of ways. British Columbia’s Children’s Hospital is one site where the use of registered data could provide

help in patient diagnosis and data analysis.

7.4 Summary

The registration techniques developed here provide a toolkit for registering multimodality data, but the segmentation tools required for this type of application must still be developed further. The limitations of surface matching were analyzed as well as the performance of various heuristic modifications. This resulted in a specification of criteria for setting the various parameters associated with the registration process and a more thorough understanding of their effects. The registration problem is still far from being solved. This thesis provides a simple and robust method for registration that can be expanded for use in many different areas.

Bibliography

- [1] AIM, P. First Versions of SW Tools for Multimodality Image Registration: WP 1.2.1 Software Tool for Registration by Stereotactic Fiducial Markers, WP 1.2.2 Software Tool for Landmark Based Registration, WP 1.2.3 Basic SW Tool for Registration of crisp Surfaces. Technical report, Commission of the European Communities DG XIII, 1992. Project A2003, COVIRA, Deliverable D2/1.2T.
- [2] AIM, P. Specification Reports for Multimodality Image Registration. Technical report, Commission of the European Communities DG XIII, June 1992. Project A2003, COVIRA, Deliverable D1/1.2.
- [3] ANDREASEN, N., editor. *Brain Imaging: Applications in Psychiatry*. American Psychiatric Press, Inc., Washinton, DC, 1989.
- [4] ANDREASEN, N., COHEN, G., HARRIS, G., CIZADLO, T., PARKKINEN, J., REZAI, K., AND SWAYZE, V. Image Processing for the Study of Brain Structure and Function: Problems and Programs. *Journal of Neuropsychiatry*, 4(2), pp. 125–133, 1992.
- [5] ARDEN, T. AND POON, J. *WiT User's Guide*. Logical Vision Ltd., Suite 108, 3700 Gilmore Way Burnaby, B.C., Canada, 1993.
- [6] BARAFF, D. Rigid Body Simulation. *Lecture Notes for SIGGRAPH '92 Course - An Introduction to Physically Based Modeling*, pp. H1–H30, July 1992.
- [7] BARTOO, G. AND HANSON, W. A. Multi-Modality Image Registration Using Centroid Mapping. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 11(2), pp. 550–551, 1989.
- [8] BATES, R., GARDEN, K., AND PETERS, T. Overview of Computerized Tomography with Emphasis on Future Developments. *Proceedings of the IEEE*, 71(3), pp. 356–372, March 1983.
- [9] BIDAUT, L. Accurate Registration of Various Medical Imaging Modalities: An Application to PET and MRI for Anatomical Localization and Kinetic Modelling. *IEEE Melecon*, pp. 1233–1237, 1991.
- [10] BORGEFORS, G. Distance Transformations in Arbitrary Dimensions. *Computer Vision, Graphics, and Image Processing*, 27, pp. 321–345, 1984.
- [11] BORGEFORS, G. Distance Transformations in Digital Images. *Computer Vision, Graphics, and Image Processing*, 34, pp. 344–371, 1986.

-
- [12] BORGEFORS, G. Hierarchical Chamfer Matching: A Parametric Edge Matching Algorithm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(6), pp. 849–865, November 1988.
- [13] BOUTS, E. A Fast, Error Free, Squared, Euclidean Distance Transform. In *VIP '93 International Conference on Volume Image Processing*, pp. 47–49, Utrecht, The Netherlands, June 1993.
- [14] BROWN, L. A Survey of Image Registration Techniques. *ACM Computing Surveys*, 24(4), pp. 325–376, December 1992.
- [15] CHEN, G., PELIZZARI, C., AND LEVIN, D. Image Correlation in Oncology. In DEVITA, V., HELLMAN, S., AND ROSENBERG, S., editors, *Important Advances in Oncology 1990*, pp. 131–141. J.B. Lippincott Co., Philadelphia, U.S.A., 1990.
- [16] CHINZEI, K., DOHI, T., HORIUCHI, T., OHTA, Y., SUZUKI, M., YAMAUCHI, Y., HASHIMOTO, D., AND TSUZUKI, M. Quantitative Integration of Multimodality Medical Images. In *SPIE Proceedings - Visualization in Biomedical Computing 1992*, volume 1808, pp. 187–195, Chapel Hill, North Carolina, USA, 1992.
- [17] CIDECIAYAN, A., JACOBSON, S., KEMP, C., KNIGHTON, R., AND NAGEL, J. Registration of High Resolution Images of the Retina. In *SPIE - Medical Imaging VI: Image Processing*, volume 1652, pp. 310–322, 1992.
- [18] COLLIGNON, A., VANDERMEULEN, D., SUTENS, P., AND MARCHAL, G. Registration of 3D Multi-Modality Medical Images Using Surfaces and Point Landmarks. In *VIP '93 International Conference on Volume Image Processing*, pp. 27–30, Utrecht, The Netherlands, June 1993.
- [19] DANN, R., HOFORD, J., KOVACIC, S., REIVICH, M., AND BAJCSY, R. Evaluation of an Elastic Matching System for Anatomic (CT/MR) and Functional PET Cerebral Images. *Journal of Computer Assisted Tomography*, 13(4), pp. 603–611, July/Aug 1989.
- [20] EHRICKE, H., DAIBER, G., SONNTAG, R., WOLFGANG, S., LOCHNER, M., SCHAD, L., AND LORENZ, W. Interactive 3D-Graphics Workstations in Stereotaxy: Clinical Requirements, Algorithms and Solutions. In *SPIE Proceedings - Visualization in Biomedical Computing 1992*, volume 1808, pp. 548–558, Chapel Hill, North Carolina, USA, 1992.
- [21] ELL, P. Nuclear Medicine. *Postgraduate Medical Journal*, 68, pp. 82–105, 1992.
- [22] ELVINS, T. A Survey of Algorithms for Volume Visualization. *Computer Graphics*, 26(3), pp. 194–201, August 1992.
- [23] EVANS, N. Combining Imaging Techniques. *Clinical Physics and Physiological Measurement*, 11(A), pp. 97–102, 1990.
- [24] FABER, T. AND STOKELY, E. Orientation of 3-D Structures in Medical Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(5), pp. 626–633, September 1988.

-
- [25] FOLEY, J., VAN DAM, A., FEINER, S., AND HUGHES, J. *Computer Graphics Principles And Practice*. Addison-Wesley Publishing Company, Inc., New York, 1990.
- [26] GERSHON, N. Image Understanding and Registration of Positron Emission Tomography (PET) Images and Magnetic Resonance (MR) Images. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 13(3), pp. 1321–1322, 1991.
- [27] GERSHON, N. Registration and Visualization of Three-Dimensional Magnetic Resonance (MR) and Positron Emission Tomography (PET) Images. *Annals of Biomedical Engineering*, 19(5), pp. 611, 1991.
- [28] GERSHON, N. Three-Dimensional Registration of Magnetic Resonance (MR) and Positron Emission Tomography (PET) Images. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 13(1), pp. 217–218, 1991.
- [29] GERSHON, N., CAPPELLETTI, J., HINDS, S., AND GLENN, M. Image Understanding, Visualization, and Registration of Magnetic Resonance (MR) and Positron Emission Tomography (PET) Images. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 12(3), pp. 217–218, 1990.
- [30] GIANCOLI, D. *General Physics*. Prentice-Hall, Englewood Cliffs, NJ, 1984.
- [31] HAMPSHIRE, S. *The Age of Reason*. Penguin Books Canada Ltd., Markham, Ontario, 1956.
- [32] HENRI, C., CUKIERT, A., COLLINS, D., OLIVIER, A., AND PETERS, T. Towards Frameless Stereotaxy: Anatomical-Vascular Correlation and Registration. In *SPIE Proceedings - Visualization in Biomedical Computing 1992*, volume 1808, pp. 214–224, Chapel Hill, North Carolina, USA, 1992.
- [33] HIBBARD, L., ARNICAR-SULZE, T., DOVEY-HARTMAN, B., AND PAGE, R. Computed Alignment of Dissimilar Images for Three-dimensional Reconstructions. *Journal of Neuroscience Methods*, 41(2), pp. 133–152, 1992.
- [34] HILL, D., GREEN, S., CROSSMAN, J., HAWKES, D., ROBINSON, G., RUFF, C., COX, T., STRONG, A., AND GLEESON, M. Visualisation of Multi-modal Images for the Planning of Skull Base Surgery. In *SPIE Proceedings - Visualization in Biomedical Computing 1992*, volume 1808, pp. 564–573, Chapel Hill, North Carolina, USA, 1992.
- [35] HINSHAW, W. AND LENT, A. An Introduction to NMR Imaging: from the Bloch Equation to the Imaging Equation. *Proceedings of the IEEE*, 71(3), pp. 338–350, March 1983.
- [36] HOHNE, K., BOMANS, M., RIEMER, M., SCHUBERT, R., TIEDE, U., AND LIERSE, W. A Volume-based Anatomical Atlas. *IEEE Computer Graphics & Applications*, pp. 72–78, July 1992.
- [37] HOHNE, K. AND HANSON, W. Interactive 3D Segmentation of MRI and CT Volumes Using Morphological Operations. *Journal of Computer Assisted Tomography*, 16(2), pp. 285–294, 1992.

- [38] HOLMAN, B., ZIMMERMAN, R., JOHNSON, K., CARVALHO, P., SCHWARTZ, R., LOEFLER, J., ALEXANDER, E., PELIZZARI, C., AND CHEN, G. Computer-Assisted Superimposition of Magnetic Resonance and High-Resolution Technetium-99m-HMPAO and Thallium-201 SPECT Images of the Brain. *Journal of Nuclear Medicine*, 32(8), pp. 1478–1484, 1991.
- [39] HORN, B. *Robot Vision*. The MIT Press, Cambridge, Massachusetts, 1986.
- [40] JASZCZAK, R. Tomographic Radiopharmaceutical Imaging. *Proceedings of the IEEE*, 76(9), pp. 1079–94, September 1988.
- [41] JIANG, H. AND ROBB, R. Image Registration of Multimodality 3D Medical Images by Chamfer Matching. In *SPIE - Biomedical Image Processing and Three-Dimensional Microscopy*, volume 1660-35, pp. 356–366, San Jose, CA, USA, 1992.
- [42] JIANG, H., ROBB, R., AND HOLTON, K. A New Approach to 3-D Registration of Multimodality Medical Images by Surface Matching. In *SPIE Proceedings - Visualization in Biomedical Computing 1992*, volume 1808, pp. 196–213, Chapel Hill, North Carolina, USA, 1992.
- [43] JOHNSTON, B. The Segmentation of Medical Data. Master's thesis, The University of British Columbia, Department of Computer Science, 1993.
- [44] JONES, R. AND MACFALL, J. Computers in Magnetic Resonance Imaging. *Computers in Physics*, 2(5), pp. 25–30, Sept-Oct 1988.
- [45] JONES, T. Positron Emission Tomography. *Clinical Physics and Physiological Measurement*, 11(A), pp. 27–36, 1990.
- [46] KAPOULEAS, I. Error Analysis of 3D Registration of Brain MR and PET Images Using the Interhemispheric Fissure. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 13(1), pp. 227–228, 1991.
- [47] KAPOULEAS, I., ALAVI, A., ALVES, W., GUR, R., AND WEISS, D. Registration of Three-dimensional MR and PET Images of the Human Brain Without Markers. *Radiology*, 181(3), pp. 731–739, 1991.
- [48] KESSLER, M., PITLUCK, S., PETTI, P., AND CASTRO, J. Integration of Multimodality Imaging Data for Radiotherapy Treatment Planning. *International Journal of Radiation Oncology, Biology, Physics*, 21(6), pp. 1653–67, Nov 1991.
- [49] KIRKPATRICK, S., GELATT JR., C., AND VECCHI, M. Optimization by Simulated Annealing. *Science*, 220(4598), pp. 671–680, May 1983.
- [50] KOVACIC, S., GEE, J., CHING, W., REIVICH, M., AND BAJCSY, R. Three-Dimensional Registration of PET and CT Images. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 11(2), pp. 548–549, 1989.

-
- [51] KRAMER, E. AND NOZ, M. CT-SPECT Fusion for Analysis of Radiolabeled Antibodies: Applications in Gastrointestinal and Lung Carcinoma. *Nuclear Medicine and Biology*, 18(1), pp. 27–42, 1991.
- [52] KREEL, L. Medical Imaging. *Postgraduate Medical Journal*, 67, pp. 334–36, 1991.
- [53] LANCASTER, J. AND FULLERTON, G. Physics and Medicine: Imaging the Body. *Computers in Physics*, 2(5), pp. 16–22, Sept-Oct 1988.
- [54] LEE, S. AND CHUNG, S. A Comparative Performance Study of Several Global Thresholding Techniques for Segmentation. *Computer Vision, Graphics, and Image Processing*, 52, pp. 171–190, 1990.
- [55] LEVIN, D., PELIZZARI, C., CHEN, G., CHEN, C., AND COOPER, M. Retrospective Geometric Correlation of MR, CT, and PET Images. *Radiology*, 169, pp. 817–823, 1988.
- [56] LEVOY, M. A Hybrid Ray Tracer for Rendering Polygon and Volume Data. *IEEE Computer Graphics & Applications*, pp. 33–40, March 1990.
- [57] LEVOY, M. Volume Rendering Using the Fourier Projection-Slice Theorem. In *Proceedings Graphics Interface '92*, pp. 61–69, Vancouver, BC, CAN, May 1992.
- [58] LI, Y. *Orientation-Based Representations of Shape and Attitude Determination*. PhD thesis, The University of British Columbia, Department of Computer Science, April 1993.
- [59] LITTLE, J. *Recovering Shape and Determining Attitude from Extended Gaussian Images*. PhD thesis, The University of British Columbia, Department of Computer Science, April 1985.
- [60] LONG, D., KING, M., AND PENNEY, B. Comparative Evaluation of Image Segmentation Methods for Volume Quantitation in SPECT. *Medical Physics*, 19(2), pp. 483–489, March-April 1999.
- [61] MAGUIRE JR., G., NOZ, M., E.M., L., AND SCHIMPF, J. Correlation Methods for Tomographic Images Using Two- and Three-Dimensional Techniques. In BACHARACH, S., editor, *Information Processing in Medical Imaging*, pp. 266–279. Martinus Nijhoff, Dordrecht, The Netherlands, January 1986.
- [62] MAGUIRE JR., G., NOZ, M., RUSINEK, H., JAEGER, J., KRAMER, E., SANGER, J., AND SMITH, G. Graphics Applied to Medical Image Registration. *IEEE Computer Graphics & Applications*, 11(2), pp. 20–8, March 1991.
- [63] MALANDAIN, G. AND ROCCHISANI, J. Matching of 3D Medical Images with a Potential Based Method. *Rapports de Recherche - Unité de recherche inria-sophia antipolis*, Programme 4 - Robotique, Image et Vision(1890), pp. 1–36, Mars 1993.
- [64] MEALHA, O., PEREIRA, A., AND SANTOS, B. Data Structures for Multimodality Imaging: Concepts and Implementation. In *SPIE - Biomedical Image Processing and Three-Dimensional Microscopy*, volume 1660-37, pp. 367–374, San Jose, CA, USA, 1992.

-
- [65] ORTEGA, J. AND RHEINBOLDT, W. *Iterative Solution of Nonlinear Equations in Several Variables*. Academic Press, New York, New York, 1970.
- [66] PAGLIERONI, D. A Unified Distance Transform Algorithm and Architecture. *Machine Vision and Applications*, 5, pp. 47–55, 1992.
- [67] PEIFER, J., GARCIA, E., COOKE, C., KLEIN, J., FOLKS, R., AND N.F., E. 3D Registration and Visualization of Reconstructed Coronary Arterial Trees on Myocardial Perfusion Distributions. In *SPIE Proceedings - Visualization in Biomedical Computing 1992*, volume 1808, pp. 225–234, Chapel Hill, North Carolina, USA, 1992.
- [68] PELIZZARI, C., CHEN, G., SPELBRING, D., WEICHELBAUM, R., AND CHEN, C. Accurate Three-Dimensional Registration of CT, PET, and/or MR Images of the Brain. *Journal of Computer Assisted Tomography*, 13(1), pp. 20–26, January/February 1989.
- [69] PELIZZARI, C., EVANS, A., NEELIN, P., CHEN, C., AND MARRETT, S. Comparison of Two Methods for 3D Registration of PET and MRI Images. *Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, 13(1), pp. 221–223, 1991.
- [70] PELLIZZARI, C. AND CHEN, G. Registration of Multiple Diagnostic Imaging Scans Using Surface Fitting. In ET AL, I. B., editor, *The Use of Computers in Radiation Therapy*, pp. 0–0. Elsevier Science Publishers B.V., North-Holland, 1987.
- [71] PÉRAULT, C., LOBOGUERRERO, A., LIEHN, J., AND BATTEUX, F. Automatic superposition of CT and SPET immunoscintigraphic images in the pelvis. In *SPIE Proceedings - Visualization in Biomedical Computing 1992*, volume 1808, pp. 235–240, Chapel Hill, North Carolina, USA, 1992.
- [72] PETERS, A. Recent Advances and Future Projections in Clinical Radionuclide Imaging. *The British Journal of Radiology*, 63(750), pp. 411–429, June 1990.
- [73] PIETRZYK, U., HERHOLZ, K., AND HEISS, W. Three-dimensional Alignment of Functional and Morphological Tomograms. *Journal of Computer Assisted Tomography*, 14(1), pp. 51–59, 1990.
- [74] POWELL, M. An Efficient Method for Finding the Minimum of a Function of Several Variables Without Calculating Derivatives. *The Computer Journal*, 7, pp. 155–162, 1964.
- [75] RATIB, O., BIDAUT, L., SCHELBERT, H., AND PHELPS, M. A New Technique for Elastic Registration of Tomographic Images. In *SPIE: Medical Imaging II*, volume 914, pp. 452–455, 1988.
- [76] RIEDERER, S. Recent Advances in Magnetic Resonance Imaging. *Proceedings of the IEEE*, 76(9), pp. 1095–1105, September 1989.
- [77] SAEED, N. AND DURRANI, T. A New MRI Rotation Algorithm for the Registration of Temporal Images. *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 3, pp. 1496–1499, May 1989.

- [78] SCHAD, L., BOESECKE, R., SCHLEGEL, W., AND ET AL. Three Dimensional Image Correlation of CT, MR, and PET Studies in Radiotherapy Treatment Planning of Brain Tumors. *Journal of Computer Assisted Tomography*, 11, pp. 948–54, 1987.
- [79] SCHIERS, C., TIEDE, U., AND HÖHNE, K. Interactive 3D Registration of Image Volumes from Different Sources. In *Proceedings of the International Symposium CAR '89 - Computer Assisted Radiology*, pp. 667–670, Berlin (West), Federal Republic of Germany, 1989.
- [80] SCHROEDER, W., ZARGE, J., AND LORENSEN, W. Decimation of Triangle Meshes. Technical report, Corporate Research and Development, General Electric Company & ConSolve Inc., Schenectady, NY & Lexington, MA, 1992.
- [81] SPETSIERIS, P., DHAWAN, V., TAKIKAWA, S., MARGOULEFF, D., AND EIDELBERG, D. Imaging Cerebral Function. *IEEE Computer Graphics & Applications*, pp. 15–26, January 1993.
- [82] STYTZ, M., FRIEDER, G., AND FRIEDER, O. Three-Dimensional Medical Imaging: Algorithms and Computer Systems. *ACM Computing Surveys*, 23(4), pp. 421–499, December 1991.
- [83] SUMANAWEEERA, T., GLOVER, G., BINFORD, T., AND ADLER, J. MR Susceptibility Distortion Quantification and Correction for Stereotaxy. In *SPIE - Biomedical Image Processing and Three-Dimensional Microscopy*, volume 1660-04, pp. 30–43, San Jose, CA, USA, 1992.
- [84] TIEDE, U., HOEHNE, K., BOMANS, M., POMMERT, A., RIEMER, M., AND WIEBECKE, G. Investigation of Medical 3D-Rendering Algorithms. *IEEE Computer Graphics & Applications*, pp. 41–53, March 1990.
- [85] TIMMENS, G., VAN DEN ELSSEN, P., APPELMAN, F., AND VIERGEVER, M. Landmark Based Registration of Images: A Comparison of Algorithms. In *VIP '93 International Conference on Volume Image Processing*, pp. 145–148, Utrecht, The Netherlands, June 1993.
- [86] TOENNIES, K., HERMAN, G., AND UDUPA, J. Surface Registration for the Segmentation of Implanted Bone Grafts. In *Proceedings of the International Symposium CAR '89 - Computer Assisted Radiology*, pp. 381–385, Berlin (West), Federal Republic of Germany, 1989.
- [87] TURK, G. TR006-92: Re-Tiling of Polygonal Surface. Technical report, University of North Carolina at Chapel Hill, Department of Computer Science, 1992.
- [88] VAN DEN ELSSEN, P., MAINTZ, J., POL, E., AND VIERGEVER, M. Image Fusion Using Geometrical Features. In *SPIE Proceedings - Visualization in Biomedical Computing 1992*, volume 1808, pp. 172–186, Chapel Hill, North Carolina, USA, 1992.

-
- [89] VAN DEN ELSEN, P., POL, E., AND VIERGEVER, M. Medical Image Matching - A review with classification. *IEEE Eng. Med. Biol.*, 0(0), pp. 000–000, 1992. (in press).
- [90] VAN HERK, M. AND KOOY, H. Three-Dimensional CT-CT, CT-MR, and CT-SPECT Correlation Using Chamfer Matching. In *VIP '93 International Conference on Volume Image Processing*, pp. 25–26, Utrecht, The Netherlands, June 1993.
- [91] VANNIER, M. Computers in Computer Axial Tomography. *Computers in Physics*, 2(5), pp. 39–43, Sept-Oct 1988.
- [92] VANNIER, M. AND GAYOU, D. Automated Registration of Multimodality Images. *Radiology*, 169(3), pp. 860–861, 1988.
- [93] WALLIS, J. AND MILLER, T. Three-Dimensional Display in Nuclear Medicine and Radiology. *Journal of Nuclear Medicine*, 32(3), pp. 534–546, 1991.
- [94] WILHELMS, J. AND VAN GELDER, A. Octrees for Faster Isosurface Generation. *ACM Transactions on Graphics*, 11(3), pp. 201–227, July 1992.
- [95] WOLBERG, G. *Digital Image Warping*. IEEE Computer Society Press, Los Alamitos, CA, 1990.
- [96] WOODS, R., CHERRY, S., AND MAZZIOTTA, J. Rapid Automated Algorithm for Aligning and Reslicing PET Images. *Journal of Computer Assisted Tomography*, 16(4), pp. 620–633, July/August 1992.
- [97] WOODS, R., MAZZIOTTA, J., AND CHERRY, S. MRI-PET Registration with Automated Algorithm. *Journal of Computer Assisted Tomography*, 17(4), pp. 536–546, July/August 1992.
- [98] WOWRA, B., SCHAD, L., SCHLEGEL, W., AND KUNZE, S. Scopes of Computer Application in Stereotactic Neurosurgery. In *Proceedings of the International Symposium CAR '89 - Computer Assisted Radiology*, pp. 297–302, Berlin (West), Federal Republic of Germany, 1989.
- [99] WYVILL, G., MCPHEETERS, C., AND WYVILL, B. Data Structure for Soft Objects. *The Visual Computer*, 2, pp. 227–234, 1992.
- [100] YOO, T., NEUMANN, U., FUCHS, H., PIZER, M., CULLIP, T., RHOADES, J., AND WHITAKER, R. Direct Visualization of Volume Data. *IEEE Computer Graphics & Applications*, pp. 63–71, July 1992.
- [101] ZHOU, J. AND TITS, A. *User's Guide for FSQP Version 3.2: A FORTRAN Code for Solving Constrained Nonlinear (Minimax) Optimization Problems, Generating Iterates Satisfying All Inequality and Linear Constraints*. Electrical Engineering Department and Institute for Systems Research, University of Maryland, College Park, MD, USA., March 1993.

Glossary

- anatomical data** - Data which represents physical, or structural quantities. This is often the characteristics of tissue or bone. CT and MRI provide this type of data.
- artifact** - An artificial or erroneous region of the data that is often created as a result of under-sampling.
- attenuation** - The reduction in the number of photons passing through a object.
- axial** - Planes perpendicular to the long axis of the body (spine).
- congenital** - Any state which has existed from birth.
- coronal** - Frontal planes passing through the long axis of the body.
- corpus callosum** - A section of the brain uniting the cerebral hemispheres.
- edge-detection** - This term is used to indicate the various methods for the isolation of gradients. An exemplar case being the zero-crossings in a convolution with the Laplacian of a Gaussian.
- electromagnetic eddies** - EM flow induced by a changing magnetic flux acting on a conductor or by a moving conductor in a static magnetic field.
- fiducial markers** - These markers are externally placed objects which show up in the scans and can be placed in designated positions relative to the object being imaged. In this way they establish a common reference frame.
- functional data** - Data which represents metabolic or physiological quantities. SPECT and PET provide data of this type. The functionality that is revealed is dependent on the metabolic attributes of the radiopharmaceutical that is used.
- Fourier transform** - A transformation that maps data from a spatial representation into a frequency representation.
- localization** - The correlation of a region to its specific anatomic position.
- lymphoma** - A tumour of lymphoid tissue.
- magnetic moment** - The net magnetic dipole of an atom represented as a vector.
- non-invasive** - Not involving penetration of the skin by surgery or hypodermic needle.
- perfusion** - Blood flow through an organ or tissue.
- phantom** - A scanned object (used in simulations) for which the exact composition is known.

point spread function - If an infinitely precise point was scanned the imaged point would be blurred or spread out to a finite area. The function that describes this spreading is called the point spread function [39].

reconstruction - The process of transforming the raw scanning device information into the 3D voxel representation.

region of interest - The general terminology for any region of a medical data set which has been delineated for specific examination.

rms - see root-mean-square.

ROI - see region of interest.

root-mean-square - The root-mean-square (*rms*) value of a set $x_i, i = 0 \dots n - 1$ is defined as: $rms(x) = \sqrt{\frac{\sum_{i=0}^{n-1} x_i^2}{n}}$.

sagittal - Parallel to the plane which bisects a bilateral animal into right and left halves.

scintigraphy - A 2D imaging technique that utilizes an internal radiation source.

shimming - Inconsistent control of the magnetic field over the scanning volume.

structural data - see anatomical data.

toxoplasmosis - An infection caused by a parasitic microorganism (of the *Toxoplasma* genus) that affects the central nervous system.

ventricles - Fluid filled cavities in the brain that are connected to the central canal of the spinal cord.

voxel - A voxel is a single unit of the sample space which represents a discrete 3D volume.