Automatic Intermodal Image Registration for Alignment of Robotic Surgical Tools

AUTOMATIC INTERMODAL IMAGE REGISTRATION FOR ALIGNMENT OF ROBOTIC SURGICAL TOOLS

 $\mathbf{B}\mathbf{y}$

ETIENNE DE VILLIERS, B.Eng. (Engineering Physics) Queen's University Kingston, Ontarion, Canada

A Thesis

Submitted to the School of Graduate Studies in Partial Fulfilment of the Requirements for the Degree

Master's of Engineering

McMaster University

© Copyright by Etienne de Villiers, February, 2000

© Copyright 2000

by

Etienne de Villiers

MASTER'S OF ENGINEERING (2000)

McMaster University Hamilton, Ontario

(Mechanical Engineering)

TITLE:	Automatic Intermodal Image Registration
	for Alignment of Robotic Surgical Tools
AUTHOR:	Etienne de Villiers
	B.Eng. (Engineering Physics)
	Queen's University
	Kingston, Ontarion, Canada
SUPERVISOR:	Dr. M.A. Elbestawi
	Department of Mechanical Engineering
	McMaster University

NUMBER OF PAGES: 1,96

Abstract

This thesis outlines the development of an automatic image registration algorithm for matching 3D CT data to 2D fluoroscope X-ray images. The registration is required in order to calculate a transformation for measurements in the 2D image into the 3D representation. The algorithm ahieves the registration by generating digitally reconstructed radiographs from the CT data set. The radiographs are 2D projection images, and therefore may be compared with the 2D Fluoroscope images.

The X-ray and fluoroscope images were compared using the photometric-based registration algorithm, pseudocorrelation, with χ^2 as the distance metric. An automated search algorithm was implemented using the Downhill Simplex of Nelder and Meade. The algorithm was successful in locating the position and orientation of the CT data set for calculating a digitally reconstructed radiograph to match the fluoroscope image.

The CT data set was located with a maximum mean position error of 2.4 mm in xy, 4.4 mm in z, and xyz axial rotation within 0.5° . The standard deviation given 1800 random starting locations was 9.3 mm in x, 12.7 mm in y, 16.9 mm in z, xz axial rotation 2.5° , and y axial rotation of 1.9° .

The search algorithm was successful in handling gross misalignment, however there were difficulties in convergence once within the vicinity of the global minimum. It is suggested to implement a hybrid search technique, switching to a conjugate gradient search algorithm once in the vicinity of the global minimum. An additional refinement would be a possible change of the distant metric, or the registration algorithm, once within the vicinity of the global minimum.

Additional investigation needs to be directed towards testing the algorithm with live fluoroscope and CT data. This is required in order to assess registration performance when comparing different imaging modalities.

Acknowledgments

The author would like to thank his supervisor, Dr. Elbestawi, for the continued support and patience as the task of registering a robotic surgical tool evolved into a thesis on automatic image registration.

Technical programming and hardware assistance was generously given by Todd Pfaff, Richard Teltz, and Hugh Pasika. Without the aid of Todd Pfaff and his suggestion of memory leakage detection, the program would still be undergoing extensive debugging.

Thanks to Scott Hastie and Sarah Crothers for their Ideas expertise, and Paula Meyer for our instructive dialogue. The IMMRC lab in general was helpful in fostering a good working environment, and supplying suggestions and support on a variety of topics.

Contents

A	Abstract			iv	
A	ckno	wledgr	nents	vi	
1 Introduction					
		1.0.1	Revision Arthroplasty	2	
		1.0.2	Current Surgical Technique	3	
		1.0.3	Solution for Improved Surgical Outcome	4	
2	Bac	kgrou	nd and Literature Review	9	
	2.1	Bone	Growth and Implant Fixation	9	
		2.1.1	Bone Growth	10	
		2.1.2	Implant Fixation	11	
		2.1.3	Fixation Success	12	
	2.2	Comp	uter-Controlled Medical Systems	14	
		2.2.1	Radiotherapy	15	
	2.3	Comp	uter-Assisted Surgery	16	
		2.3.1	Computer Guidance	17	
		2.3.2	Robotic Surgery	19	
		2.3.3	Applications	19	

	2.4	Diagn	ostic Medical Imaging	22
		2.4.1	Computerised Tomography	22
		2.4.2	C-arm X-ray Fluoroscopy	24
3	Reg	istrati	on	27
	3.1	Patier	nt Registration	27
	3.2	Mecha	anical Registration Procedures	29
		3.2.1	Fiducial Registration	29
		3.2.2	Shape Registration	31
	3.3	Auton	natic Registration	35
		3.3.1	Intermodal X-ray Matching	36
	3.4	Regist	ration Through Synthetic X-rays	37
		3.4.1	Advantages of DRRs for Image Registration	39
	3.5	DRR	Calculation Complications	40
4	$\mathrm{Th}\epsilon$	eory		42
	4.1	Radia	tion Physics	42
		4.1.1	Diagnostic Radiology	43
		4.1.2	X-ray Interaction with Matter	45
		4.1.3	Calibration Constants	47
		4.1.4	Additional Considerations	48
	4.2	Featu	re Extraction	48
	4.3	Regist	ration Metrics	50
		4.3.1	Template Matching	51
		4.3.2	Photometrics	54
5	Ima	ige Re	gistration Algorithm	61

	5.1	Overview	61
	5.2	DRR Calculation	64
		5.2.1 Ray-Tracing	64
		5.2.2 Data Set Orientation and Position	65
		5.2.3 Numerical Integration	66
		5.2.4 Additional DRR Considerations	70
	5.3	Optimisation Search Algorithm	70
		5.3.1 Downhill Simplex - Nelder and Meade	71
6	Ros	ulte	76
U	nes 0 1		76
	0.1	Overview	10
	6.2	Search Algorithm Results	77
7	Disc	cussion	80
	7.1	Search Difficulties	80
		7.1.1 χ^2 Scaling	81
		7.1.2 Local Minima	82
	7.2	Performance of the Registration Metric χ^2	82
	7.3	Improving Overall Registration	83
		7.3.1 Improved Algorithm Structure	84
	7.4	Comparison with Other Studies	84
		7.4.1 Search Algorithm Performance	85
8	Con	nclusion	87
A	Para	allelised Code	89
	A.1	Overview	89

Bibliography		93
A.1.3	Relevance to Computer-Assisted Surgery	92
A.1.2	Overview of PVM System	91

List of Tables

.

6.1	Results for 180	0 Random	Starting Positions.													•	•	•	7	6
-----	-----------------	----------	---------------------	--	--	--	--	--	--	--	--	--	--	--	--	---	---	---	---	---

List of Figures

1.1	Illustration of femoral drilling and broaching.	4
1.2	Cross-section of manually broached and robotically machined femurs. [PBM+92] $$	5
2.3	Internal bone structure.[Cow89]	10
2.4	Intraoperative custom femoral component production. [LME+95] $\ .$	13
2.5	A sampling of possible radiation field geometries.[GvH93]	15
2.6	Typical surgical tools instrumented for navigation. $[DSJ^+98]$	18
2.7	Precision femoral drilling. [SH95]	20
2.8	CT slice displayed as an image, and voxel placement within the CT data set.	23
2.9	C-arm fluoroscope. [PDH94]	24
2.10	Typical X-ray spectra, and attenuation coefficient as a function of beam	
	energy.[JC83]	25
3.11	Fiducial registration.[Sim99]	29
3.12	Shape registration.[Sim99]	31
3.13	Physically probing a knee joint for cadaveric studies.[GSF99]	32
3.14	Matching a physically probed surface to a pre-defined model. [GSF99]	33
3.15	Template matching in computer-guided hip surgery. [DSJ+98] $\ .$	34
3.16	Perspective distortion creates the illusion that objects closer to the source	
	are larger than objects closer to the image plane.	39
4.17	Monochromatic X-ray attenuation.	43

4.18	Primary radiological image production.[JC83]	45
4.19	Edge detection on the reference image and the distance transform applied to	
	the initial guess.	49
4.20	Location of the center of a machined hole with the Hough Transform. \ldots	52
4.21	Linear correlation between two similar, and two dissimilar images	55
4.22	Pseudocorrelation χ^2 comparing two clinical images. [RRS94]	57
4.23	Variance of χ^2 as the data set is translated along the x co-ordinate axis	59
5.24	Setup geometry for DRR calculation	63
5.25	CT data set position, and direction vectors: $\vec{p}, \vec{a}, \text{ and } \vec{b}$	65
5.26	Two-dimensional bilinear interpolation.	69
5.27	Comparison of DRR calculated with and without interpolation	69
5.28	Possible steps taken by the Amoeba, here a 4 vertice figure searching in	
	3D.[PTVF94]	72
5.29	Clumping of a poor random number generator.	73
5.30	Variance of χ^2 as the CT Data Set is translated and rotated for 50 pixel points.	74
6.31	Sample DRR images of reference and initial guess.	78
6.32	The initial guess and mean solution images as feature templates superimposed	
	upon the reference image.	79
7.33	Comparison of variance of χ^2 along a single parameter line[0,0,1,0,0,0], and	
	several parameter lines.	81
7.34	Variance of χ^2 as the number of pixels sampled increases from 10 to 100	82
A.35	Illustration of the PVM computation model, and architecture schematic.[GBD+9	91 91

Chapter 1

Introduction

The Case for Computer-Assisted Surgery

There are over 16 000 hip replacements performed every year in Canada.[Can89] With an aging population, all indications point to an increase in the number of hip replacements, Total Hip Arthroplasty(THA), performed. This procedure greatly improves the quality of life for those afflicted with hip joint disease. A successful THA may restore full mobility, and permit the patient a return to their accustomed lifestyle. In addition, there is tremendous economic incentive as a wheelchair-bound individual incurs large costs associated with accessibility and long-term chronic care. Unfortunately this technique has yet to be perfected, with implants requiring revision on average after 15 years. Since a successful outcome is unlikely after a second revision, this places a minimum age at which a patient may receive a THA without becoming wheel-chair bound at a later point in life.

1.0.1 Revision Arthroplasty

Revision arthroplasty is a traumatic procedure as usually the implant has failed, or loosened within the femur. In the case of the former, extraction of the remaining portion of the implant may result in breaking the shaft of the femur. If the implant has failed due to loosening, often there is little solid bone left within which to place a new implant.

Studies have shown that "revision THA places inordinate demands on the resources of the hospital and requires significantly increased time, effort, stress, and liability from the surgeon".[Bar95] The operation takes longer, the outcome is never on par with the original replacement, and the patient usually suffers as a result. If a means to increase the longevity of hip replacements were possible, significant savings to the medical industry, improved productivity, and reduction of pain for the patient could be realised.

Causes of Revision Arthroplasty

Aside from outright mechanical failure of the prosthesis, the main cause of revision surgery is loosening of the implant within the femur. There are a variety of causative factors involved with implant loosening, but the common detail is a reduction in the strength and mass of the bone at the bone-implant interface. This reduction in bone density is known as resorption, and is caused by three main factors:

- Particulate Debris
- Size, Material Properties, and Surface Characteristics of the Prosthesis
- Natural Aging Process

While the third factor is dependent upon health, lifestyle, and age, the first two factors may be directly addressed with improvements in prosthetic design and surgical technique.[Pos92] The most advanced prostheses from a mechanical standpoint are anatomically-shaped to match the natural joint loading mechanics. Unfortunately these implants are extremely difficult to accurately locate in a precise cavity using current surgical technique. [PBM+92] The overall integrity of a joint replacement is highly dependent on the mechanical environment of the bone-implant interface.

Currently, THA relies on standard X-rays for assessing implant size and placement. The surgeon then shapes the femoral cavity using hand tools in order to fit the implant. There are two main types of implants, those fixated with surgical cement, and cementless implants which rely upon bone ingrowth for their fixation. While they both fail for different reasons, the end result remains the same.

Studies have indicated that for cementless prostheses, a stable fit requires bone ingrowth. However, the bonc-implant interface must be in close apposition $\sim 0.5 \, mm$. Too tight a fit causes swelling, pain, and poor healing. Too loose a fit leads to excessive implant motion, and scar tissue formation instead of a stable bone graft. Cemented prostheses are also highly dependent on an accurately shaped femoral cavity. One study of 836 cemented femoral prostheses identified four factors related to surgical technique which had a dramatic effect on improved long term outcome.[ESM⁺94]

- 2-5 mm Cement Mantle
- $< 2 \, mm$ Cancellous Bone Surrounding the Implant
- Implant Stem Fills More Than Half of the Medullary Canal
- Axial Alignment of Implant Neutral or Valgus

1.0.2 Current Surgical Technique

Manual surgical techniques are incapable of addressing these factors on a consistent basis. Current techniques utilise a broaching process that roughly chips the bone to implant



Figure 1.1: Illustration of femoral drilling and broaching.

shape.(Figure 1.1) The variations in the bone surface due to the cutting process itself are outside of the bounds stated in the literature for strong, long-term stable fixation. Figure 1.2 illustrates the improvement in accuracy and surface finish achievable by a machine tool. Finally, there have been studies suggesting that implant longevity is highly dependent upon the experience of the surgeon. While surgeons with several years of experience are able to consistently achieve implant survival rates around 15 years, inexperienced surgeons often have considerably less successful implant survival terms. While statistically poor outcomes are not the fault of these less experienced surgeons, the fact remains that a substantial learning curve exists in achieving successful femoral implant fixation.

1.0.3 Solution for Improved Surgical Outcome

Surgeons are currently limited to tools and techniques that have undergone very little substantive development since the American Civil War. Manual surgical techniques are



(a) Broached Femur



Figure 1.2: Cross-section of manually broached and robotically machined femurs.[PBM⁺92]

incapable of delivering fine tolerance in a repeatable manner, so research has looked to computer-controlled machining to attain the required precision. While a large change in thinking for medicine, such a revolution transformed the manufacturing and machining industry in the late 1940's through the 1960's with the advent of first NC, and the successor, CNC machine tools. Computers are able to control machine tools to an extremely fine tolerance and more importantly, repeat those accurate cuts indefinitely. In the case of THA, this requires a computer-guided surgical robot, the medical incarnation of the industrial CNC machine tool.

Current trends in the field of autonomous robotic surgery is limited to the implementation of pre-existing robotic arms such as the 5-axis Sankyo-Sankei known as ROBODOC [PBM⁺92], and the 6-axis Puma 260 Unimation.[MGS⁺93] The former is a femoral cavity preparation tool for THA, and the latter is designed as a saw guide for Total Knee Arthroplasty(TKA).

Tool Registration

A CNC machine tool, or a robotic arm requires its position and orientation relative to the workpiece; the robot needs to be registered. Within the context of a manufacturing environment, registration is a fairly straightforward procedure. When machining a new piece of material, the first set of cuts creates a flat, machined surface, which provides a registration or datum point. If the part to be registered must be located by previously machined surfaces a touch probe is utilised. The touch probe is brought up to the part surface, and when just in contact with the surface, the part is considered referenced to the machine tool.

The difficulty for robotic surgery is there are no accurately locateable datum positions with which to precisely locate the machine tool. The ROBODOC registration method originally relied upon the implantation of radio-opaque pins into the femur. These pins are actually fairly large screws which require an incision into the leg to prepare the site, and then are screwed into the femur a week prior to the surgery. The screws provide a series of accurate, fixed-point reference points on the femur. The patient leg is then imaged using a CT scanner, which generates a 3D density map of the femur. The registration pins may be located within the CT data by standard image processing techniques. The CT scan is then used in a pre-operative planning session at which the surgeon fits and places a model of the implant into the CT image of the patient femur. This model, subtracted from the CT femur data, defines the bone removal volume for the robot. A computer then converts the defined removal volume into a set of machine instructions, G-code. The registration pins provide a set of identifiable datum points between the femur, and the intended tool path.

During the operation, the surgeon fixates the femur with respect to the robot using two large cortical bone screws. The pin wounds are reopened, and the surgeon guides a touch probe to the pin head sites. Since the registration pins are fixed relative to the bone, and the touch probe is fixed to the bone via the bone screws, the femur may now be registered to the robot. That is, the robot knows the exact orientation and position of the femur in space, since the pin locations were accurately defined in the CT scan.[PBM⁺92]

The ROBODOC pin registration method requires an operative procedure a week prior to surgery in order to place the pins. This, in conjunction with the associated pain and trauma of inserting bone screws is considered excessive. It was proposed to develop a taskspecific machine tool to replace the robotic arm, custom designed to minimise the intrusion of the robot on the operating environment. As a part of this project, a specific need was defined: eliminate the pins to expedite the registration.

This thesis concerns the development of a pin-less registration system for computerassisted surgery. The initial scope considered only the registration of a machine tool to the proximal end of the femur. Once registered the machine tool would conduct robotic surgery for THA. After conducting the literature review however, it soon became evident that there are a variety of surgical applications that would benefit from a pin-less, non-invasive registration system. While recent innovations have eliminated the reliance on pin-based registration techniques, the replacement procedures generally require additional time and consideration on the part of the surgeon during the operation. This suggested an additional mandate for this study, minimise the role played by the surgeon in registering the patient.

As well, the field of computer-assisted surgery has subsequently divided into two main areas: guidance, and robotics. While both show promise, computer-guided surgery offers substantial improvement to surgical technique both through improved accuracy, and reduced radiography requirements while still using established surgical techniques and tools. A computer-guided operation requires a smaller financial commitment, less change to current surgical technique, and a smaller leap of faith on the part of the surgeon and the patient than robotic surgery. This suggests the adoption of computer-guided surgical techniques at an increased rate over robotic surgery. This led to consideration of the requirements for computer-guided surgical applications. For these reasons the focus of this thesis has been broadened to include a variety of anatomical situations, while still considering the limitations associated with robotic surgery applied to THA.

Chapter 2

Background and Literature Review

The difficulty faced with a thesis addressing a topic such as image registration for computer-assisted surgery is the necessity to simultaneously consider a range of diverse, seemingly unrelated, disciplines. In order to lay the groundwork for an examination of the requirements for a registration system, it is necessary to outline some of the issues and relevant details surrounding the implantation of biomedical devices for reconstructive surgery.

2.1 Bone Growth and Implant Fixation

The main difference from normal practise when applying standard engineering principles to the design of biomedical devices is the regenerative capabilities of the human body. Unlike generic engineering materials, bone is dynamic and actually grows stronger in response to heavier loads. This reactionary remodeling of bone was defined in Wolff's oft-quoted book 'The Law of Transformation of Bone': "change in the function of a bone is followed by



Diagram of arrangement of cells associated with growing bone (cross-section of mid-shaft of femur of young rabbh)

Figure 2.3: Internal bone structure. [Cow89]

certain definite changes in the internal architecture and external conformation in accordance with mathematical laws." [Tre81] This characteristic has great relevance as the converse remodeling process, bone resorption, in response to reduced loads is also found to be true.

2.1.1 Bone Growth

Bone is composed of inorganic minerals, collagen, water, and 1% by weight cellular material. The collagen fibers resist tensile loads, while the binding mineral matrix provides the high compressive strength. The cellular material, composed of: osteoblasts, osteocytes, and osteoclasts, regulate the formation and resorption of bone. Essentially, osteoblasts generate, or synthesise, the precursor to collagen, tropocollagen. The tropocollagen is created on the surface of the bone, and evolves into collagen fibers. The collagen becomes mineralised to form mature bone tissue. Osteoclasts on the other hand are responsible for the converse process, bone resorption. Osteocytes are osteoblasts trapped within the bone tissue, and are responsible for maintaining the mineral composition of the bone. When a bone is under high stress, osteoblasts are called upon to lay down new tropocollagen, and the bone becomes extremely dense in these regions. If on the other hand, the bone does not experience heavy loads, osteocytes resorb the bone. This leads to a decrease in the bone's density, and a corresponding decrease in the yield strength of the bone. When an injured person is given an implant, or function is affected, the structure of their anatomy will change to best manage the new biomechanics and loading patterns.[Cow89][McR94]

2.1.2 Implant Fixation

The use of THA in restoring functionality has been successfully implemented since the 1960's. The first widely successful system, developed by Charnley, utilised a straight, smooth femoral implant that was fixated within the femoral canal with bone cement (PMMApolymethlmethacrylate). The success of Charnley's system owed much to the simple shape of the implant, and his demand that practitioners must attend an educational clinic at his orthopaedic center; ensuring exacting femoral cavity preparation on the part of the surgeon. Since then many attempts have been made to perfect both the prosthesis, and the implantation procedure, with varied success.

Cemented Femoral Implants

In the short to medium term implant success was high, however after 8 years the implants generally failed due to a breakdown and loosening of the bone cement. Several steps were taken to improve implant longevity such as improving the composition of PMMA, pressurised injection of the cement, careful femoral canal preparation, and plugging the canal distally to contain the cement. While implant lifetime has improved to ~ 15 years, there appears to be a limit to what degree advanced technique can extend the lifetime of a cemented implant. As well, there are health risks associated with the cement such as pulmonary emboli and osteolysis('cement disease'). These concerns were weighed against the main benefit of cemented prostheses, namely that there is an increased margin of error in the quality and precision of the femoral cavity preparation. Though optimal results require a consistent cement mantle, 1-2 mm thick, the use of cement does allow for an uneven fit to some degree.

Cementless Femoral Implants

The noted drawbacks led to the development of a cementless implant. These implants rely on the regenerative properties of bone for stable fixation. The basic design philosophy centers around creating a porous network on the surface of the implant, permitting the bone to mechanically interdigitate with the network, securing the implant. A variety of methods have been attempted for creating a hospitable environment for bone ingrowth from texturing the surface of the implant to adding a porous coating. In recent years the standard has become a porous coating created by sintering a metallic powder onto the surface of the implant.

2.1.3 Fixation Success

While average implant lifetime has improved as surgical technique and implant design has been modified, it is clear from various studies that the critical aspect of implant longevity is the quality of implant fit within the femoral canal. It has been found that not only must the implant fit snugly within the canal, but that primary fixation must occur at the proximal(top) end of the femur. The cause of this is thanks to Wolff's law. If an implant is fixated at the bottom end, then all loads are transmitted through the implant to the fixation site. The bypassed bone is resorbed, the fixation interface weakens, and the implant fails. This consideration led to the concept of Customised Cementless Femoral Components.

Customised Femoral Components

The concept of a custom implant makes inherent sense. No two human anatomies are identical, so why would one assume that a generic implant would be applicable across the spectra of anatomical types. This led to the development of a system capable of custom machining the femoral component of a hip replacement based on a custom CAD/CAM model derived from a mold, or CT scan of the patient. Subsequent studies found that the initial surgical outcome was favorable.[Bar89][SSW89][MMB⁺89] In contrast, a smooth-



Figure 2.4: Intraoperative custom femoral component production.[LME⁺95]

surface custom implant, the Identifit Hip, did not fare as well. An independent study reported a 28% failure rate on a trial of 68 patients within a range of 1-44 months. This is considerably worse than the failure rate for standardised femoral components. The authors noted that no postoperative measure was predictive of early failure.[LME⁺95]

As a biomedical engineer designs an implant both the functional properties of the device, and its affect on the surrounding bone must be considered. It has been demonstrated for instance, that wear particles from abrasion of articulating surfaces, such as in knee or hip implants, are directly responsible for resorption of bone around the implant. This has led engineers to search for bio-compatible materials that are wear-resistant. The Identifit Hip introduced a combination of new features that may have led to its high failure rate. A smooth surface with longitudinal grooves provided no area for bone ingrowth.

While the structural design of an implant has great impact on its functional success, its orientation and fit within the bone are just as important for ensuring a successful outcome. With standard engineering materials, a loose fitting part may work its way free under cyclic loading. In biomedical applications the bone will actually be resorbed around the loose part, accelerating the loosening process. It is this aspect of the environment that drives biomedical engineers to develop robotic surgical tools. It is impossible to design and develop biomedical devices with any level of sophistication without a valid performance assessment. Currently there is no adequate control over the position, orientation, or quality of fit of an implant within the bone. Unless a device is an outright failure like the Identifit Hip, researchers must wait 10-15 years for statistical failure rates in order to deduce the viability of an implant design. The use of robotic tools for site preparation will enable objective evaluation of clinical outcome, and hopefully augment the optimisation of implant design.

2.2 Computer-Controlled Medical Systems

There are currently a variety of active medical research topics that involve computerassistance.[PDH94][MGS⁺93][RKT⁺97][BMBH95][PBM⁺92] The introduction of computers into medicine is a natural development given the requirements for accurate, timely analysis of a variety of diagnostic tests. Often this information must be gathered from a variety of sensors, or tests, and conveyed to the surgeon in a meaningful, quantitative manner. In addition to the fields of computer-assisted surgery, computers are being utilised for assessing patient alignment in radiotherapy treatment.

2.2.1 Radiotherapy

Radiotherapy is the act of treating a cancer by irradiating the area with a focussed beam of high energy electrons or photons. The concept relies on the radiation dose killing the cancer cells, without causing irreparable harm to the patient. In order to achieve this, the patient is exposed to a series of low dose treatments over the course of weeks or months. At each visit, the patient must be aligned with the radiotherapy machine in order to ensure that the treatment beam is reapplied to the cancerous tissue, and exposure to healthy tissue be minimised. The alignment of the patient with the therapy machine is achieved by acquiring



Figure 2.5: A sampling of possible radiation field geometries.[GvH93]

patient images with portal X-ray cameras attached to the machine. These cameras function in a similar manner to a standard X-ray device, but instead of a diagnostic imaging beam the treatment beam of the machine is used for exposure. This results in an image with low contrast and poor resolution. This *portal* image is then compared to a *reference* image which outlines the relevant anatomical features in relation to the prescribed radiation field(the area to be irradiated). The reference image could be a DRR, a plain X-ray, or a prior portal image. The radiation field is a user-defined geometry created to isolate the cancerous cells from the healthy tissue. This template is superimposed upon the portal or X-ray image, defining the area to be treated by the beam. Examples of some sample radiation fields are illustrated in Figure 2.5. Deviation in the current alignment of the patient with respect to the therapy machine as compared with previous visits will appear as discrepancies in the position and orientation of anatomical structures relative to the radiation field.

Automatic Patient Alignment

There have been a series of attempts to develop an automatic patient alignment scheme for radiation therapy.[WF94][GvH93][CTT⁺95][GDT⁺96] The three main techniques involving DRRs are: single reference image transformation, interpolation between a set of reference DRRs, and iterative DRR projection. Of these, only the third is capable of consistently measuring out-of-plane rotations.

Current research is focusing on iterative techniques, and specifically methods to permit real-time registration. One specific application for real-time automatic registration is the use of automated radiation therapy machines like the Cyberknife. This machine is capable of maintaining the therapy beam alignment dynamically, correcting for patient motion during therapy.[Mur97] This feature maximises radiation exposure to the target, and minimises exposure to neighboring healthy tissue.

2.3 Computer-Assisted Surgery

There is currently a limited selection of computer-assisted surgical systems in full clinical use throughout the world.[DSJ⁺98][PBM⁺92] It is important to note that computer-assisted surgery is quite different from tele-surgery, or remote surgery. In tele-surgery a camera system is positioned to present a view of a particular anatomical system. A local surgeon

prepares the site, and locates a robotic manipulator within the region being treated. The operation itself, or aspects requiring particular skills or training, is conducted by a surgeon viewing the operation by way of the cameras from a remote site. The robotic manipulator follows the instructions of the specialist, and reproduces the instructed cuts. There is no interpretation or representation of the surgical scene. The camera displays the patient, and the robot reproduces the surgeon's motions. These systems have also been used for a variety of procedures from minimally invasive to micro-surgery. In the former, the surgeon operates through a small incision, viewing the cuts via a television camera. Micro-surgery is similar, but the operation is conducted through a microscope, and the surgeon's actions are scaled by the manipulator in relation to the magnification of the microscope.

A computer-assisted surgical system utilises the computational power of a computer to enhance, or augment, the information available to the surgeon. There are two main fields of computer-assisted surgery: guidance or surgical navigation, and robotic surgery. The former tracks instruments and the patient, with the computer synthesising the information and displaying relative graphics to the surgeon in real-time. The graphics, superimposed upon an X-ray image of the patient, elucidate the actual live position and orientation of a tool with respect to anatomical structure that is obstructed from direct view. Robotic surgery in comparison takes the next step. The computer not only tracks the patient and instruments, but institutes the cuts itself within the constraints of a pre-planned surgical model. The surgical tool is responsible for enhancing the motor function of the surgeon, or including information in the cutting process that is unavailable to the surgeon.

2.3.1 Computer Guidance

In surgical navigation, a computer tracks both the patient, and hand-held surgical tools within the operating theatre. Current tracking technology utilises infrared markers on a template. The templates are rigidly fixed to bones of interest, and any surgical tools that



Figure 2.6: Typical surgical tools instrumented for navigation. [DSJ+98]

require guidance. After taking a referenced X-ray of the patient in the operating room (OR), the surgeon is able to direct surgical cuts by observing a real-time computer model of the tool super-imposed upon the X-ray by the computer.

In comparison, current technique requires the surgeon to insert a guide or drill wire, take a radiograph, reposition the wire, take another exposure, and continue until satisfied with the alignment. This commonly leads to several minutes of radiation exposure to both the surgeon and the patient.

Since the computer follows patient and tool positioning, there is no need in computerguided surgery to take additional radiographs as the surgery progresses. Instead, the tool position is updated on the original reference radiographs in real-time, saving both surgeon and patient from unnecessary radiation exposure.[DSJ⁺98]

2.3.2 Robotic Surgery

Robotic surgery arose out of a recognised need to improve the accuracy and precision of surgical cuts. As an additional benefit, well-designed robotic surgical tools often require less invasive incisions. Unlike the surgeon, the tool requires only minimal access to the surgical site. There are numerous benefits to robotic surgical tools: cutting force monitoring, repeatable surgical cuts, control over feeds and speeds, and no learning curve required for accurate bone cuts. The robot will never replace the surgeon, however its accuracy and control over the cutting process promise immediate improvements in surgical outcome.

The biggest drawbacks associated with robotic surgery center around adaptability to the surgical environment, and acceptance by both the medical community and the general public.[PBM⁺92]

2.3.3 Applications

In addition to THA and TKA, there are a variety of surgical procedures that could benefit from the assistance of computer-guidance and control. There have been limited studies examining the efficacy of computer-guided surgical tools applied to a variety of surgical procedures. The two most obvious cases involve spinal instrumentation and femoral drilling.(Figure 2.7) These techniques generally require the accurate placement of a drill wire within a small region of dense bone. In the case of spinal surgery, if the wire strays outside the target, nerve and spinal cord damage may result. Femoral drilling requires accurate wire placement in order to achieve stable fixation of the femoral head. The success of both surgeries is directly related to the accuracy of screw placement. Since the drill wire creates locating holes for the screws, its accurate placement is crucial. Real-time tracking techniques give the surgeon accurate feedback while the procedure is underway, improving outcome and patient assessment. It is likely that such procedures will be adapted for robotic



Figure 2.7: Precision femoral drilling. [SH95]

surgical technique in the years to come.[SH95]

The striking advantage of computer-assisted surgery is the reduction in X-ray exposure to both the patient, and the surgeon. This factor cannot be stressed enough, as surgeons may receive quite high levels of radiation exposure as case after case requires extensive fluoroscope imaging. Computer-guided surgery requires only two or three fluoroscope images, one from each viewing angle. The surgeon is then able to place holes more accurately, and in less time than with current, unguided, techniques.[DSJ+98]

There are additional benefits to both the patient and surgeon through the use of computer-assisted robotic surgery. The accuracy of bone cuts, both in precision of cut shape and direction, and in the placement of the cuts relative to the bone is far greater for a surgical tool, than for an unaided surgeon. In addition, the machine tool is capable of repeating the cuts with the same accuracy for each trial. This is especially important for investigative studies concerning implant longevity. Current studies suffer due to the difficulty in assessing the impact of implant design on lifetime when there is variance in placement between implants. If a robotic surgical tool is utilised however, the effect of over-sizing, or under-sizing the femoral cavity to the implant may be examined in a statistically significant manner. It is currently accepted that bone requires a close interference fit for successful bone ingrowth in cementless arthroplasty. However, due to variation in manual techniques there is no way to objectively determine at what degree of interference one achieves optimal ingrowth. The same dilemma follow for cemented arthroplasty in determining the ideal thickness of the cement mantle for maximising implant lifetime.

The main technical difficulty in implementing both forms of computer-assistance is the registration of the patient to pre-operative imaging data. While a surgeon can conduct a computer-assisted surgery without the benefit of pre-operative data, there promises to be great improvement in surgical outcome when such data is included in the surgical procedure. Robotic surgery on the other hand depends entirely upon the registration of the patient to

a pre-operative model. It is this model that the surgeon uses for analysis, and planning the surgical procedure.

2.4 Diagnostic Medical Imaging

An alternative to X-ray based imaging, ultrasound, was rejected after preliminary examination. Ultrasound relies on timing the travel of sound waves through matter in order to obtain a distance measurement. The accuracy of the measurement is dependent upon the wavelength of the sound wave. However, as frequency increases to improve the resolution, so the signal attenuates faster in the medium. In the case of solid body registration such as the femur, an ultrasound imaging unit cannot supply adequate information to ensure an accurate determination of position and orientation.

This thesis relies on two standard forms of diagnostic medical imaging: Computerised Tomographic(CT), and C-arm X-ray Fluoroscopy. Both are industry standard imaging techniques, and available at modern surgical facilities.

2.4.1 Computerised Tomography

A CT scan is actually a 3D density mapping of an object. The CT scanner scans repeatedly through 180° arcs at a series of levels, or positions along an object. The resulting scan is a series of 2D density slices through the object. An individual volume element in a slice is referred to as a voxel, as each slice is interpreted to have a depth. Generally the voxels have square xy dimensions, while the depth, or z direction, is dependent upon the spacing of slices selected by the operator. In this study a standard voxel depth of 3 mmwas chosen as it was the spacing of the CT slices in the male visible human data set.[oM] Increasing the number of slices leads to an increase in the effective absorbed radiation dose, so there is definitely an incentive to make do with the minimum information required to


Figure 2.8. CT slice displayed as an image, and voxel placement within the CT data set

register the tool. Investigation determining the minimum number of CT slices required for accurate registration would be useful, but was not a component of this study. Other studies have utilised CT slices taken in the range of $1-3\,mm$ spacing, so the choice of $3\,mm$ is a reasonable start.



Figure 2.9: C-arm fluoroscope. [PDH94]

2.4.2 C-arm X-ray Fluoroscopy

The C-arm fluoroscope as illustrated in Figure 2.9 is a portable imaging device used extensively in precision surgical procedures. The fluoroscope is the surgeon's eyes through the patient's body for guiding drill wires, placing screws, and assessing alignment in the operation. There has been extensive study in the past two years regarding fluoroscopy for computer-assisted surgery. This technique links operating tools, and the fluoroscope through an infrared motion tracking system. After the surgeon takes a single series of fluoroscope images, the computer tracks patient and tool position and provides a real-time update of the tool position relative to the patient on the initial fluoroscope reference images. Since only one set of fluoroscope images taken at the beginning of the operation are required, this has been shown to significantly reduce the required radiation exposure to both surgeon and patient.



(a) X-ray Spectra



Figure 2.10: Typical X-ray spectra, and attenuation coefficient as a function of beam energy.[JC83]

C-arm as a Precision Measurement Tool

The final accuracy of a registration system is limited by the resolution of the measurement devices. In the case of the C-arm, current generation fluoroscopes are generally inadequate for precision measurements. These instruments were intended to supply qualitative radiographic images, assisting the surgeon's actions. Their main limitations are due to both the quality of X-ray production and detection, and the physical construction of the device itself.

The rate of X-ray attenuation varies both with object density, and X-ray energy. Standard fluoroscopes tend to produce X-rays in a broad band, or distribution, of energy (Figure 2.10). This leads to variances in image production between fluoroscopes for a given object. The detection unit, a fluorescent screen, is curved which leads to distortions in the image. Finally, the C-arm that holds source and detector is prone to sag with prolonged use, leading to misalignment between source and detector.

There are various calibration regimes developed to standardise the information collected

by different C-arm Fluoroscopes. However, most of the variation will be corrected with the advent of modern Digital Fluoroscopes. These devices are not fluoroscopes at all, their detection arrays are in fact silicon-based photon sensitive arrays. Similar to CCD cameras, these detectors may be tuned for sensitivity to a narrow range of photon energy, are quite flat, and weigh considerably less than a conventional fluoroscope. The reduction in weight minimises potential deflection in the C-arm. While not standard equipment in ORs yet, future imaging systems will likely be solely based on digital technology.[MZHR99]

There has been research into high and dual energy fluoroscopy, however these techniques are not currently applied in general. They are important to note however, as the accuracy of the registration process may become limited by the resolution of the fluoroscope. Essentially high energy fluoroscopy, \sim 1MeV, improves resolution producing a high contrast image. Dual-energy fluoroscopy implements sensor fusion, combining two X-ray images taken at different energies. The additional information produces a more accurate image with greater detail.

Chapter 3

Registration

3.1 Patient Registration

In order to successfully conduct computer-assisted surgery, it is necessary to locate the position and orientation of the patient relative to a datum point. In the case of robotic surgery the position and orientation of a specific bone is required relative to the machine tool, or a fixed reference point. This point is referenced to the origin of a co-ordinate system that is defined for the computer. Radiation therapy requires the position and orientation of a tumor relative to the fixed co-ordinate system of the radiation therapy machine.

The accurate CT data is used as a reference model defining the patient's anatomy for all types of registration methods. Prior to the operation, or procedure, the surgeon develops a pre-surgical plan on the computer within the CT data set. For THA, this involves virtually placing and orienting the implant within the CT data. Robotic surgical techniques in general require the computer to calculate a corresponding set of machining instructions, Gcode, in order for the machine tool to carry out its operation. The computer references the G-code relative to the origin of the defined co-ordinate system. Radiation therapy follows a similar procedure in which the surgeon identifies the treatment volume, and the computer calculates the required therapy beam path to achieve the desired exposure.

During treatment, the relationship between either the bone and the datum, or the tumor and the therapy machine is required. Since the therapy machine is a gantry device capable of accurately targeted radiation exposure, it is intuitively obvious to utilise the imaging capabilities of the therapy machine to capture orthogonal views of the patient in situ. These orthogonal views were originally used to manually define the treatment volume, and locate the volume relative to the therapy machine. Recent developments have led to automated procedures for matching structures within these 2D images with their corresponding representation in the 3D CT data set.

Radiation Exposure

One aspect of computer-assisted surgery that causes some practitioners to express concern is the reliance on CT imaging. This sentiment is especially strong in cases such as THA, in which previous and current technique have only required 1-2 plain X-rays. The question is raised as to what level must treatment outcome be improved to justify the use of CT and its corresponding increased radiation exposure. There is currently no definitive measure to compare anticipated outcome improvements versus higher radiation exposure.

It is important to minimise radiation exposure to the patient, and current computerassisted surgical techniques have actually demonstrated a dramatic decrease in the fluoroscope time required for a procedure. This notwithstanding, diagnostic X-rays are a necessary part of surgical intervention, and a short exposure en route to improved surgical outcome is generally preferred to no exposure, and inferior results. In the case of THA, the promise of increased implant longevity would lead to fewer revision cases. Since revision surgery is highly traumatic, and can also involve extensive use of fluoroscopy, limited use in the initial procedure could save excessive use in subsequent operations.

3.2 Mechanical Registration Procedures

Computer-assisted surgical techniques have thus far relied upon mechanical registration procedures as there are no standard imaging devices prevalent in OR use for capturing orthogonal X-rays. The mechanical registration techniques may so far be divided into two main groups: fiducial registration, and pin-less registration.

3.2.1 Fiducial Registration



Figure 3.11: Fiducial registration.[Sim99]

Fiducial registration procedures, locating pre-implanted bone screws, rely upon manually guiding a robotic arm to the screw sites. The computer records these datum positions, and refers to a pre-operative CT scan which also contains the screws. By identifying the position of the screws in CT space, an inverse transformation may be derived to link the CT co-ordinates with the datum point positions in physical space. As long as the robotic arm remains fixed relative to the bone, registration is maintained throughout the surgery. This technique has been proven mechanically effective, and extremely robust throughout the ROBODOC trials. The procedure is broken down into two main components: the pre-operative and the operative.

Pre-operative

The pre-operative component involves inserting surgical pins, bone screws, into the anatomical structure to be tracked during a minor pre-operative surgical procedure. With pins implanted in the bone, the patient is imaged by a CT scanner and the image data transferred to a computer workstation. The workstation must run object recognition and registration code in order to identify and locate the surgical pins within the CT data. A 3D model of the bone is created from the CT data, and the surgeon performs a virtual operation on this model. The computer records the cuts required by the operative plan and relates them to the locations of the bone screws in the CT data. Once the position of the bone screws has been determined, the robot's tool path code may be transformed to consider the position of the bone screws relative to the machine tool.

Operative

The operative aspect concerns the mechanical registration of the pins. After initial preparation, the bone is fixed relative to the computer's local co-ordinate system. Fixation is achieved by securing the bone using two bone screws. Given that the bone is fixed in space, the surgeon manually guides a touch probe to each of the registration pin sites. This permits the computer to determine the position and orientation of the anatomical structure relative to the machine tool.

The main drawback of the pin-based registration system is pain associated with the implantation of the registration pins. In the case of THA, some pins must be located

CHAPTER 3. REGISTRATION

at the distal end of the femur, near the knee joint. Patients complain of post-operative pain at the registration pin sites, and aching in the knee joint up to six months after the surgery. Surgeons familiar with the operation complain of complications resulting from the screw implantation including: long-lasting bone pain, poor knee functionality, and screw site infection. Since these complications are due to procedures not directly related to the corrective surgery, proper medical ethics led researchers to replace the method with a pinless registration method.[Wie99a]

Procedure Plan Medical Images Patient Preoperative Shape Extraction With It Spatially Reconcile Shape Plan with Execution Registration Intraoperative Shape Extraction Patient min min Procedure Execution

3.2.2 Shape Registration

Figure 3.12: Shape registration. [Sim99]

The only current alternative to pin-based fiducial registration is surface template matching or shape registration.[RKT⁺97] This technique was developed for computer-aided guidance procedures, specifically concerning: pedicle screw insertion, hip guide wire insertion, reconstructive hip surgery, and THA performed by ROBODOC. This pin-less registration

CHAPTER 3. REGISTRATION

procedure does away with the reference pins by matching the surface of the patient's femur in the CT data to the physically defined surface of the femur in the OR. The femur's surface is located by touching a series of reference points on the patient's bone with a tracked stylus. The position and orientation of the stylus at each of the points is combined by the computer to generate a model of the bone's surface. This surface is then matched to the femur representation within the CT data to calculate the position and orientation of the bone with respect to the stylus tracking system.

Operative Procedure



Figure 3.13: Physically probing a knee joint for cadaveric studies.[GSF99]

The basic procedure involves fixing the bone with respect to a mechanical reference point, or a tracking marker. This is achieved by securing the femur with bone screws to



Figure 3.14: Matching a physically probed surface to a pre-defined model.[GSF99]

either the computer's local co-ordinate system, or to a tracking marker. The tracking marker permits the computer to follow movements of the femur with respect to the computer's local co-ordinate system. A series of bone landmarks are then mapped with the use of a touch probe or stylus.(Figure 3.2.2) This procedure defines the position and orientation of the femur with respect to the computer's co-ordinate system. The touch probe may either be mechanically linked to a reference point, or may itself tracked by an infrared motion tracking system.[DJB⁺98] In general some 30 surface points on the relevant anatomy are defined, and used to generate a 3D surface topography. This topography is then matched to the CT data and an inverse transformation calculated to map points from CT space to physical space in relation to the computer's co-ordinate system. The problem is one of over-constraint, 30 points are more than required to match the surface, with the extra information utilised to minimise error and verify the alignment. The surface matching technique does not require an invasive procedure prior to the patient receiving the CT scan. However it does suffer from two significant drawbacks.

CHAPTER 3. REGISTRATION

Shape Registration Drawbacks

The first drawback, is that the computer requires a relatively large area of bone to be exposed to the stylus. The accuracy of the registration procedure is directly proportional to the distance spanned by the located points. This is not a problem for procedures like reconstructive hip surgery in which a fairly large area of bone is exposed as a matter of routine in the standard procedure. In the case of THA, a small incision is made at the hip joint site, and the majority of the femur remains hidden from view. This requires a series of incisions to be made at the distal end of the femur to permit stylus access. The use of additional incision sites during the procedure is undesirable as each incision causes additional bone to be exposed, increasing the possibility of infection.

An additional issue for surface matching is the time it takes the surgeon to reference the points. Current procedure utilises approximately 30 points to define the surface, taking up to 5 minutes for the surgeon to locate all of the points. While not an inordinate amount of time, information from preliminary studies suggest that surgeons would be hesitant to perform the procedure multiple times throughout the operation to verify the registration procedure.



Figure 3.15: Template matching in computer-guided hip surgery.[DSJ+98]

While fairly accurate the procedure requires additional surgical time, and surgeons hesitate to verify registration by repeating the touch probe procedure at later points in the operation. Additionally, in cases such as multiple pedicle screw insertion over a series of vertebrae, surgeons believe registering one vertebrae to be adequate. This is far from an optimal solution, and ideally registration should be verified at various points along the spinal column throughout the surgery.[Jar99][DJB⁺98][Sim99][DSJ⁺98][GSF99][Wie99a][Wie99b]

3.3 Automatic Registration

These shortcomings led to the search for a registration procedure that is: accurate, reproducible, and requires minimal surgical time. Since X-ray fluoroscopes are routinely engaged to assess alignment in surgical procedures, a method was sought that would combine its ease and versatility with the accuracy of the CT scan.

The basic concept is simple, take a fluoroscope image during the operation with the machine tool in place. From this image, a measurement may be taken to determine the position of the machine tool relative to the femur. However, these measurements must be calibrated or scaled with the actual dimensions of the femur. The CT scan provides an accurate model of the femur from which to calculate the reference dimensions. This turns the registration process into a matching procedure, between a two-dimensional (2D) plain X-ray and a three-dimensional (3D) CT scan.

Since computer-assisted and robotic surgery is being implemented to minimise invasive surgery, it seems counter-productive to introduce new registration techniques that require invasive cuts. In addition, a registration system should minimally infringe upon an operation. A surgeon should be encouraged to verify registration at various points throughout the surgery. Time consuming measures like surface matching lead a surgeon to attempt as few registration procedures as possible. Finally, some procedures require extreme exposure just to provide a surgeon with partial access. In these cases the benefits of computer-assistance still exist and are perhaps heightened, though physical access is too limited for touch probe registration.

These concerns have fortified the need for an imaging-based automatic registration procedure. Such a system does not require any pre-operative invasive procedures to install pins, nor any bone to be exposed for surface definition. The quality of registration may be verified at any point by re-imaging the patient with the fluoroscope. An additional benefit is that the system would only require hardware that would normally be a part of the operation. In the case of computer-guided surgery, the fluoroscope is an integral part of the operation, currently providing the reference images that the surgeon uses to direct the instrumentation.

By automating the registration procedure it is hoped the surgeon will be able to devote more energy and attention to the actual operation. In addition, if the process is transparent to operating practise, it is more likely to be incorporated across the widest possible range of surgical and medical procedures.

3.3.1 Intermodal X-ray Matching

The basic premise for this registration scheme is to match objects in the 2D fluoroscope image with the corresponding object in the 3D CT data set. The 3D CT data is preoperative, and is the model from which the surgeon plans the operation. The 2D fluoroscope image is taken during the operation, and illustrates the relative position of the surgical tool, or a reference marker, with respect to the bone. The act of identifying, measuring, and placing an object in a 2D image is fairly well documented in the literature.[SFH92][Bro92] Likewise, there are various schemes for identifying objects within a 3D data set.[BM96] The problem remains however to determine the transformation for resolving an object in the 2D image with it's subsequent 3D representation. Various techniques have been developed for matching 2D projections to 3D data sets, and some specifically for 2D X-ray matching with 3D CT data. However, most of these techniques rely on segmentation, and contour matching. These techniques tend to be fairly time consuming, and are highly dependent on being able to clearly identify object outlines in the 2D X-ray image. In the case of the femur, the only real outline feature for axial rotation is the neck and head of the femur. Unfortunately it is this structure that is usually first removed in preparing the patient for a THA. The remaining possibility is to use internal image structure, features caused by bony landmarks, to align the femur with the CT data.

Internal image structure arises as a result of the heterogeneous density characteristics of bone. As the bone is rotated, these lines will move in the X-ray image, even though the outline or exterior contours of the bone appear unchanged. By identifying bone landmarks on the 2D X-ray with pattern recognition algorithms, corresponding internal density structures may be located within the CT scan. Preliminary research into these techniques, notably among vascular reconstruction, and other 3D registration cases suggested that identifying these features would be a time consuming approach.

A compromise may be reached between the utility of relying upon internal structure and the performance penalty associated with complex recognition algorithms. There are various registration algorithms recently developed for intermodal registration that are able to register images in a time efficient manner. The common theme among these algorithms is to minimise the information required to achieve a successful registration.

3.4 Registration Through Synthetic X-rays

An alternative solution was found based on registration techniques used in radiation therapy. When a patient requires irradiation of a tumor or cancer the process is usually carried out over several visits to the clinic. At each session the oncologist must verify the position and orientation of the tumor relative to the radiation therapy machine. The machine is then programmed to irradiate the patient following such a path that a succession of therapy beams pass through the tumor, while minimising exposure to healthy tissue. Thus treatment success relies upon accurately identifying the treatment volume consistently at each visit.

Identification of the treatment volume, and registering its position relative to the radiation therapy machine is achieved by creating portal images. These are poor quality X-rays obtained by exposing X-ray film to an attenuated treatment beam. The oncologist then visually matches the portal image to a previous reference image which could be another portal image, a plain X-ray, or a Digitally Reconstructed Radiograph (DRR). DRRs are computer-generated synthetic X-rays produced by ray tracing through the CT data set, and calculating the line integral along the ray. The resultant sum corresponds to the total linear density through which the ray passes, which yields the average amount the X-rays falling on that particular detector element will have attenuated. The important feature to note is that internal bone structure is included in the resultant 2D X-ray image. Thus matching a portal image to a DRR has the effect of matching the internal bone structure without the high computational cost.

In recent years efforts have been made to automate the process with some success. The problem is fairly well constrained as the patient is lying flat on a table, and generally out of plane rotations are ignored. The problem is then one of four degrees of freedom(DoF): magnification, rotation, and x-y translation, in order to match a 2D reference image with a 2D portal image. Generally, single pixel accuracy is sufficient for a successful outcome.

In order to register a 2D image to a 3D data set, it is proposed to calculate 2D DRRs from the 3D CT data, and then match the resultant 2D DRRs to the 2D fluoroscope image. Therefore in order to find a transformation mapping objects in the 2D fluoroscope image back to the 3D CT data set, it is actually necessary to determine the parameters that will

CHAPTER 3. REGISTRATION

permit the successful calculation of a DRR to match the fluoroscope image. Unlike the four DoF radiation therapy situation this is a six DoF problem, requiring special attention to computational speed while maintaining registration accuracy.

Once the correct DRR calculation parameters have been determined the orientation and position of the machine tool may be extracted from the fluoroscope image using standard 2D image analysis techniques. The DRR parameters define a mapping to place the extracted machine tool within the CT data set. It is likely that two or more fluoroscope images will be required to locate the machine tool within the CT data. The second image is required to minimise errors due to out-of-plane rotations and translations. Additional images could be used to verify registration accuracy, and provide the surgeon with a measure of fit.



Figure 3.16: Perspective distortion creates the illusion that objects closer to the source are larger than objects closer to the image plane.

3.4.1 Advantages of DRRs for Image Registration

The advantage of utilising a DRR for the 2D to 3D registration procedure, is that the inherent projection distortions generated while capturing the fluoroscope image will be recreated from the non-distorted CT data. "This approach [using DRRs] seems the most logical one, because it simulates the imaging process directly, thereby making optimal use of all geometrical information in the CT data and in the transmission images." [GDT⁺96]

Distortions arise in two main ways. First, the X-rays diverge from the source leading to a perspective distortion (Figure 3.4). Objects close to the source are magnified to a greater extent than those far from the source (close to the image plane). This leads to difficulties in being able to match contours and features directly with the CT data set. The second distortion is caused by the physical construction of the fluoroscope unit. Current generation fluoroscopes rely on photo-multiplier tubes which have a curved imaging surface. This leads to so-called 'pin-cushion' distortion. Essentially, pixel-spacing near the edges of the field of view are both closer together, and no longer on the flat image plane. This leads to magnification and sampling, quantisation, distortions.

The second distortion is not of great concern for two main reasons: first, it may be reproduced in the DRR by modeling the curved image plane when ray-tracing the CT data set. That is the instead of tracing rays onto a flat image plane with equal pixel spacing, they are traced to points corresponding to the outline of a curved image plane.[FMH97] This additional consideration is actually irrelevant as next generation fluoroscopes will utilise flatpanel technology. Instead of photo-multiplier tubes, photo-diodes will be used to measure the incident photon energy. The first distortion, perspective or magnification distortion is recreated when calculating the DRR. Since the DRR is calculated by modeling the actual physics of an X-ray it contains distortions inherent in the imaging process.

3.5 DRR Calculation Complications

There are a variety of issues that would need to be addressed in a commercial system. Issues such as calibration of the fluoroscope, and compensating for differences bwtween the fluoroscope image, and the DRR would need to be examined. However, most of these considerations are fine-tuning issues as opposed to major process related problems. The most important issue centers around verification of bone registration. The surface matching registration procedure utilises additional mapped surface points to verify the success of the registration. It is proposed that additional fluoroscope views be taken to verify the success of the registration algorithm. This thesis focussed solely on the act of correlating a reference X-ray to DRRs calculated from a CT data set. However it is important to note additional algorithm requirements.

Chapter 4

Theory

4.1 Radiation Physics

In order to compute the DRRs, it is necessary to take into consideration the underlying physics present in capturing an image with the C-arm fluoroscope. The question of image distortion with respect to the C-arm has been briefly examined, however there are additional considerations that need to be addressed before the C-arm may be presented as a precision measuring device. While the process of collecting the CT data does not need to be examined, the actual nature and definition of the data itself is important if one wishes to numerically reproduce images similar to the images that the C-arm will capture. This chapter relies on information described in "The Physics of Radiology" [JC83], and "Principles of Computerized Tomographic Imaging". [AS88]



Figure 4.17: Monochromatic X-ray attenuation.

4.1.1 Diagnostic Radiology

X-ray Attenuation

As X-rays propagate from a source, they interact with all of the matter in their path. Thus if an incident beam of I_o photons impinge upon a thin absorber of thickness Δx , the number of photons transmitted through the layer, $T = \Delta I_1$ is proportional to I_o and Δx . (Figure 4.17)

The rate of their attenuation is dependent upon an attenuation factor, μ_{xyz} , that varies with the density of the material.

$$T = \Delta I = -\mu I_o \,\Delta x \tag{4.1}$$

Which holds for some small Δx and ΔI . In the limit, this equation becomes:

$$dI = -\mu I_o \, dx \tag{4.2}$$

$$\frac{dI}{I_o} = -\mu \, dx \tag{4.3}$$

integrating and simplifying yields a standard exponential decay relationship:

$$T(x) = I_0 e^{-\mu x} \tag{4.4}$$

This simple relationship only holds where μ is constant over the interval of integration. In order to expand this concept for a non-homogeneous material, like the human body, we must consider μ as a function of position. Thus:

$$T = I_o \exp\left[-\int_{ray} \mu_{xyz} \, ds\right] \tag{4.5}$$

Again, this is a simplification of the physical event, as this formula assumes that the X-rays are produced from a monochromatic source. That is, that all of the X-rays have the same energy, which is not the case for medical imaging devices. A commercial system might require calibration of the code to match the X-ray spectra produced by a particular fluoroscope; the code would require an effective energy measurement of the fluoroscope to compare with the provided effective energy of the CT scanner. These issues, while highly relevant, were not considered in this study as attention was directed towards building the intermodal registration code.

Projection Image

Generalising from the previous section it is evident that an X-ray image, in which an object is placed between a photon-sensitive medium and some form of X-ray source, is actually a form of shadow image. The detail is a result of the attenuation of X-rays. As illustrated in Figure 4.18 the image is a record of the density through which the rays pass from source to detection medium. The primary radiological image is the result of computing a series of density line integrals from the source to the detection medium as defined in Equation 4.5. Bone landmarks appear in an X-ray image as light features, regions where



Figure 4.18: Primary radiological image production.[JC83]

the X-rays have attenuated to a greater extent as they passed through denser bone.

4.1.2 X-ray Interaction with Matter

A full description of how ionising radiation interacts with matter is beyond the scope of this thesis. However, it is necessary to minimally cover the topic in defining exactly which processes are considered. The design of the DRR generating algorithm assumed that the X-ray images to be matched would be generated by a fluoroscope producing keV range X-rays. It is important to note this, as the attenuation coefficient, μ , is actually the sum of the cross-sections of all four interaction mechanisms. These mechanisms each dominate the total attenuation coefficient at different X-ray energies.

$$\mu_{total} = \tau_{photoelectric} + \sigma_{coherent} + \sigma_{incoherent} + \kappa_{pair} \tag{4.6}$$

CHAPTER 4. THEORY

In coherent scattering, the incident electromagnetic wave induces the atom's electrons to vibrate. The oscillation of the electrons causes emission of radiation at the same wavelength as the incident wave. Though no energy is transferred to kinetic energy, the secondary emission leads to a broadening of the photon beam, and a loss in effective energy.

Incoherent scattering, or compton scattering involves the release of an electron(kinetic energy E), the annihilation of the incident photon (energy $h\nu$), and the production of a new photon (energy $h\nu'$). By the conservation of energy:

$$h\nu = h\nu' + E \tag{4.7}$$

Thus the created photon has less energy than the incident, resulting in an attenuation of the beam.

Pair production only occurs when the incident photon has energy in excess of 1.02MeV. In this situation a photon passing near the nucleus may disappear, and result in the production of an electron-positron pair $(h\nu - 1.022 = E_+ + E_-)$. This transformation derives from the mass energy equivalence, and so the kinetic energies of the two particles results from any incident energy above the threshold for pair production.

Given current fluoroscope technology and the standard energy being utilised, only the fourth mechanism, photo-electric absorption, is considered in calculating the DRRs. In photo-electric absorption the photon is annihilated, and an electron liberated with some kinetic energy, E. This process occurs with the highest probability if the photon energy is just higher than the binding energy yielding an electron with small E.

These considerations are important, as previously developed DRR algorithms were implemented for radiotherapy, in which a number of interaction mechanisms are prevalent both singly, and in combination. Modeling the wrong interaction mechanism will result in a DRR that will not match the fluoroscope image.

4.1.3 Calibration Constants

The CT data as utilised for this project is idealised from the actual Hounsfield units supplied by commercial CT imaging systems. The data is shifted to start at zero, and a simple threshold correction was made for material attenuation characteristics. In order to accurately model the interaction between X-rays and organic matter, there are a number of calibration constants required to specify the true nature of the interaction. For the purposes of this investigation dealt solely with numerical calculation and correlation. Therefore nonbone was eliminated by applying a factor of zero, while bone was modified by a factor of one. The effect of leaving out the actual mass attenuation coefficients had no observable effect on the registration algorithm, and is consistent with other DRR calculation schemes. As a point of interest the difference in the attenuation constant between bone and non-bone is roughly an order of magnitude.

The sample DRRs included in the results section (Figure 6.31) were generated with calibration constants assuming a beam energy of 60keV, and rescaled to fall within a 128 greyscale. While a computer has no problem discerning between 4000 shades of grey, the human eye registers somewhere in the range of 128-256 greyscales. If the DRRs were left as 4000 greyscale images they would appear as muddy images with low contrast. This is only due to our limited range of recognition, and is not a factor when the images are analysed by the computer.

When the registration code is integrated with a real C-arm fluoroscope, it will be necessary to take the actual calibration constants for the current beam energy into account. This is required in order to generate DRRs that are equivalent to the live fluoroscope image. The most efficient way to do this would be to pre-process the CT data. As previously stated, the CT data is actually listed in Hounsfield units, which under ideal conditions are calculated by:

$$H = \frac{\mu - \mu_{water}}{\mu_{water}} \times 1000 \tag{4.8}$$

Thus the CT data points must be converted to attenuation factors, given the machine specific value for μ_{water} , and then modified by the mass attenuation coefficients for a given beam energy. The CT data set in Hounsfield units would be replaced with an attenuation factor data set so that the ray tracing algorithm would not need to calculate the attenuation factor at each step.

4.1.4 Additional Considerations

The accuracy of the DRR code in exactly reproducing the fluoroscope image is not necessarily a requirement for accurate registration. There is however a need to accurately reproduce the features upon which a particular registration algorithm relies in order to match a set of images. This investigation utilised a photometric-based comparison, which is a comparison of relative pixel intensities. Since relative pixel intensities were considered, the algorithm is actually comparing structural details. Thus the DRR image, and the reference image may exhibit different average intensities and still constitute a match as long as they contain the same contextual information. The choice of metric, that is the matching criterion, has a direct impact on the quality of information required from the DRR calculation engine.

4.2 Feature Extraction

In order to generate a shape template of an image highlighting salient features, detail information must be extracted from the image. The features of an image are obtained by







Figure 4.19: Edge detection on the reference image and the distance transform applied to the initial guess.

edge detecting the image, and only selecting edges with a strong gradient. Edge detection is accomplished by passing a gradient filter such as the Sobel filter across the image.

Sobel Edge Detection

The Sobel operator is a method of gradient calculation that yields a magnitude, independent of orientation. Essentially, the pixel gradient is calculated in two orthogonal directions, and summed in quadrature.(Equation 4.9) In addition the gradient direction may be computed as listed in Equation 4.10.[Rus92]

Sobel Magnitude =
$$\sqrt{\frac{\partial B^2}{\partial x} + \frac{\partial B^2}{\partial y}}$$
 (4.9)

Sobel Gradient =
$$\arctan\left(\frac{\frac{\partial B}{\partial y}}{\frac{\partial B}{\partial x}}\right)$$
 (4.10)

Figure 4.19a illustrates the resultant image after the reference image has been operated upon by the edge detection algorithm. By applying a threshold operation across the edge detected image, the strong edges may be extracted.

Once the edges are defined a distance transform is computed across the image, and a maxima search yields the feature lines.(Figure 4.19b) These features could be used in a template-based registration algorithm such as the Hough Transform, or the Hausdorff Distance.

4.3 Registration Metrics

There are two main classes of registration metrics: photometric, and template matching. Photometric analysis considers only the intensity of individual pixel elements without regard to their relationship to neighboring pixels. Template matching involves some form of edge detection, and extraction of basic object features within the image. Pixel intensity is only considered so far as the intensity defines an edge or feature on the object. There has been research involving both methods in medical imaging, with each method finding relative success in a particular application.

A variety of registration metrics were examined in the course of designing the registration algorithm. Since processing time was considered a factor, consideration was paid both to the cost of computing a particular metric, and the information required to compute the metric. In this particular case, the calculation of the DRR may be considered the most costly part of the registration process. Therefore any method which minimises DRR computation will enjoy the greatest time savings.

4.3.1 Template Matching

Template matching operates by extracting some form of object features from the reference image. The most commonly used features are edges and contours. A template is created once the features have been identified. A distance metric is then used to quantify the degree to which the template matches a particular image.

The template matching regimes generally require the entire DRR to be traced in order to identify and fit models to the image. This would suggest a greater computational cost to successfully register the reference image. One study selectively traces landmark features to create a DRR that contains only feature information. These features may then be registered to the reference image.[GDT⁺96] The advantage is that the selective tracing faces a far smaller computational burden than tracing the full CT data set. The drawback is increased registration complexity, and pre-processing of the CT data. Two forms of template matching were examined and subsequently discarded for registering the DRR images.

Hough Transform

The first template matching scheme, the Hough Transform, is illustrated in Figure 4.20. In the example image (a) illustrates the original grey scale image being examined, while image (b) shows the metal block after an edge detection algorithm has been applied. Once an edge detected image has been obtained, the features may be identified. In this example the center of the machined circle is the feature of interest.

In order to identify a feature one requires a template. This template may either be supplied by a reference image, or be described by a mathematical equation. A circle may simply be described by:

$$x^2 + y^2 = r^2 \tag{4.11}$$

CHAPTER 4. THEORY



(a) Original Image of Block



(b) Edge Detected Block



(c) Detection of Circle Center in Hough Space



(d) Center of Circle Located on Original Image

Figure 4.20: Location of the center of a machined hole with the Hough Transform.

If an edge detected point lies on a circle, it will satisfy this equation. The first step in the Hough Transform is to mark in parameter, or Hough, space all of the possible circles on which this point could lie. Each possible circle is noted as a 'vote' in Hough Space.

The physical points which lie on the actual circle will all cast one of their votes to the same point in Hough Space. That is, the actual circle is represented by the equation that receives the most votes in Hough Space.

While useful for a simple geometry for a circle, the Hough Transform would not be useful for complex geometries if one was required to know the parametric equation that described the figure. This is where the reference template would come into play. In the case of image correlation, edge detected points may be referenced to an arbitrary 'center' point. By calculating a direction, or gradient, in combination with the Sobel magnitude, a template may be described by each edge point's distance from the center and corresponding gradient.

A matched pair of images will both exhibit the most votes at the same center point. The main advantage of the Hough Transform is that one can describe fairly complex geometries by the vote pattern in parameter space.[SFH92] Other template matching routines usually require some form of model fitting to match a polynomial to identified features.

Hausdorff Distance

The second template matching routine that was investigated was the Hausdorff Distance. Given two sets of points representing a model and an image:

"The Hausdorff distance between these two point sets is small exactly when every point in the model is close to some point in the image, and every point in the image is close to some point in the model." [Ruc96]

Essentially the procedure involves defining a skeleton model, which in this case would be an edge detected fluoroscope image. This reference is then analysed to derive a feature set of line segments and points. A DRR would then be calculated, edge detected, and feature points determined. The Hausdorff distance could then be calculated between the model or reference points, and the DRR generated points. It is noted that the Hausdorff distance measure has advantages over the Hough Transform in that with the Transform

"every model feature is paired with a large number of image features, and vice versa, ... the total number of counts due to incorrect pairings is much larger than the number of counts due to correct pairings. The chance of a large peak being formed by a random combination of incorrect pairings is thus quite large." [Ruc96]

These two techniques were not investigated further, as both required considerable analysis beyond the DRR calculation. It seemed logical following the principle of parsimony to first attempt registration with a simpler metric that didn't involve image manipulation and analysis. An additional consideration is that the fluoroscope image is also often cluttered with surgical instruments and other physical artifacts. While both of these template-based techniques are extremely robust in dealing with selective regions of interest(ROI), there was concern that obscuring a dominant feature might foil the registration process.

4.3.2 Photometrics

The simplest form of image comparison deals with individual pixel intensities. Surprisingly, there are quite a range of methods to perform this single task. While all are functional the goal of minimal computational cost and high reliability ruled several out. The metric chosen was the method of pseudocorrelation as described by Radcliffe.[RRS94]

Pseudocorrelation was developed in order to measure the alignment of images produced by Video Electronic Portal Imaging Devices (VEPID) as used in radiotherapy. These images are acquired in near real-time and compared with simulator or portal images taken at a



Figure 4.21: Linear correlation between two similar, and two dissimilar images.

previous treatment. The benefit of pseudocorrelation is that a successful measure of fit may be calculated by only sampling a few pixels out of the entire image. Radcliffe suggests 25 pixels as a useful minimum, however investigation using DRRs during this thesis points to 100 as a good choice. The difference has minimal effect on algorithm performance given that other techniques like FFT, and template matching correlation schemes require all of the pixels in an image, necessitating the tracing of the entire DRR when measuring the fit with a reference image. Even a small image, 128×128 pixels, contains a total of 16 384 pixels. The 75 additional pixels in the pseudocorrelation scheme is not even a row worth of pixels in a full image.

The basic strategy of the pseudocorrelation algorithm is to plot a particular pixel intensity from the reference image versus its corresponding intensity from the image to be matched. By collecting a number of such pixel pairs at random locations throughout the image the linear correlation between the grey levels of the two images may be calculated. Figure 4.21 illustrates the procedure with two very similar images, and two completely different images (simulated by generating a random set of pixel intensities). A line of best fit is drawn on both, and Pearson's r coefficient calculated as a measure of linear fit. The similar images display a slope and r^2 near unity. The dissimilar images also manage a slope near one, however r^2 identifies the fit as poor.

The exact grey level values in each image are unimportant, rather the relationship between the relative pixel intensities. An image that is uniformly darker or lighter corresponds to a shift in the pixel line of fit. A progressive darkening or lightening in one image relative to another will lead to a change in the slope of the line of fit. As long as the pixels' relative intensity remains consistent the line of fit will remain straight, and the images will be considered to match.

One nice feature of the pseudocorrelation algorithm is that regions of interest (ROIs) can be defined as required either restricting pixel sampling inside the region, or as regions of exclusion. This procedure could be implemented in an automatic fashion, excluding pixels containing metal for instance. Since this is a photometric correlation measure, a pixel's relation to its neighbors is ignored, there is less risk of eliminating a key feature and subsequently failing the registration.

Search Space Topography

The mechanics of the search operation involve rotating and translating the CT data cube in order to optimise the match between the reference X-ray image, and a numerically computed DRR. The metric, or measure of fit utilised by the search algorithm is based on the principle of pseudocorrelation, as described by Radcliffe.[RRS94]

As suggested by Radcliffe, the χ^2 statistic was used to describe the linear correlation as opposed to Pearson's r, or some other linear measure of fit. Thus χ^2 plays the role of a distance measure. A pair of perfectly matched images will return a χ^2 of zero. As the misalignment of the images increases, so χ^2 increases. The calculation of χ^2 for binned data is listed in Equation 4.12, where N_i represents the pixel intensity in the calculated DRR,



Figure 4.22: Pseudocorrelation χ^2 comparing two clinical images. [RRS94]

and n_i the pixel intensity in the reference image.

$$\chi^2 = \sum_{i} \frac{(N_i - n_i)^2}{n_i} \tag{4.12}$$

According to Radcliffe χ^2 is a good choice for a number of reasons, but the strongest feature is that it grows large for poorly matched images, providing 'absolute probability

of good alignment'. [RRS94] Pearson's r coefficient is found lacking in that there is an inconsistency regarding the size of r and the degree of correlation between two images.

Unfortunately though χ^2 does not vary smoothly with the geometric transformation parameters.(Figure 4.22) Radcliffe utilised a grid search method for locating the global minimum to counter this irregularity of the χ^2 metric space. While a possible solution in 4 DoF, in 6 DoF the grid search method quickly becomes untenable. For instance, sampling at just 3 transformation locations (+1mm, 0mm, -1mm) in 4 DoF requires 3⁴ evaluations as opposed to 3⁶ in 6 DoF. This corresponds to 81 evaluations instead of 729. Since only measuring at 3 locations is unrealistic, a more realistic value being around 20 (20⁶ - 64,000,000), some form of optimisation search algorithm is essential, as opposed to a blind grid search.

A more recent study by M.J. Murphy examined 6 DoF image registration between CT and radiographs for radiosurgery. Murphy utilised a combination of orthogonal radiograph views which were masked to isolate ROIs. In this study Murphy also used χ^2 but restricted the search space to small ROIs around the edge of the skull.[Mur97] The photometric was used to identify changes in the position of the skull outline. This effectively makes the photometric-based registration algorithm perform like a template matching scheme. It was considered to follow a similar technique for registering the femur, however such an approach would require the surgeon to identify ROIs in the fluoroscope image since the size and position of the femur in the image would depend upon the position and orientation of the fluoroscope. This seemed counter-productive to the goal of fully automated registration, suggesting investigation into registration without identified ROIs.

It is important to note however that in Murphy's study a hybrid search algorithm is utilised. This algorithm begins with a gradient search, and progresses to a fitting function using Taylor Series Approximation near the global minimum. Murphy calculates the derivative required for this technique by building a lookup table of derivatives across a range of


Figure 4.23: Variance of χ^2 as the data set is translated along the x co-ordinate axis. the six parameters prior to patient treatment. During the procedure a specific derivative is calculated by finite differencing between lookup entries.

It was hoped for this study that an accurate registration could be achieved without the use of derivatives. Murphy notes that

"the hybrid gradient search algorithm typically reaches a good minimum in four iterations, thus requiring five pairs of DRRs ... For comparison a downhill simplex search would require seven pairs of DRRs to initialise, and is notoriously slow to converge." [Mur97]

The speed of convergence was an issue in selecting the Downhill Simplex, however there was concern that a gradient search would become trapped within the local minima found throughout the search space. Murphy managed to avoid this problem by restricting the search range to within +/-8 mm along each co-ordinate axis and $+/-5^{o}$ around the rotation axes.

This restriction localises the search to a region close to the reference point, and exhibits a fairly smooth search space topography. This point is illustrated in Figure 4.23 which plots the variance in χ^2 as the data set is translated along the x co-ordinate axis. It is important to note that with less than 10 mm translation, χ^2 remains fairly well behaved. However, to restrict the possible range of transformations within 8mm and 5° is not very realistic for the clinical setting. Thus a range of $(+/-20 \text{ mm } xy, +/-30 \text{ mm } z, +/-15^{\circ})$ in the transformation between the reference and initial trial images was used in this study. This magnitude of initial error is more representative of the type of variance expected in real trials.

Since χ^2 appears to exhibit large local minima in which a search algorithm may become trapped, the ideal routine must exhibit robust performance and the ability to find its way out of local minima en route to the global minimum.

Chapter 5

Image Registration Algorithm

5.1 Overview

The scope of this thesis concerns the development of an automatic registration algorithm to match an arbitrary X-ray image with a DRR produced from a CT scan of the object. When it was first determined that DRRs would be used for registration, a literature review was conducted to identify current DRR generation engines. These algorithms are essentially modified ray-tracing code (Figure 5.24) in which a line is drawn from a source, through space, and falls on an image plane. The pixel on which the ray falls will take its value dependent upon what objects the ray intersects along its path. In the case of DRRs, the total density through which the ray passes is recorded and is represented as the shade of the pixel on the image plane. Different views of a patient's anatomy are calculated by rotating the source and image plane about the CT data set.[CFP+85][CTT+95] A line integral is calculated along each ray through the CT data set, in order to compute the magnitude of attenuation (Equation 4.5).

Additional Considerations

For the purposes of this thesis, it was necessary to consider both the speed of calculation, and the flexibility of the DRR algorithms. The implementations cited above were developed for radiotherapy patient alignment techniques. While the matching procedure is similar, radiotherapy is afforded the luxury of a well-positioned patient and the time to move and orient the patient for good initial alignment with the machine. In computer-assisted surgery on the other hand, the patient is prepared on the operating table, and positioned to allow optimal access for surgery, not imaging. This may result in gross misalignment of the patient with respect to the fluoroscope and the algorithm's initial guess. Thus consideration must be made to deal with this large initial misalignment.

Algorithm Efficiency

In order to conduct automatic registration a search algorithm would need to manipulate the orientation of the data set with respect to the X-ray source in an efficient manner. It is most likely that any implementation of this algorithm would make use of hardware acceleration, or be parallelised in order to complete registration and verification as quickly as possible. Thus while attention was paid to keeping DRR calculation as efficient as possible, it was not required to ensure absolute optimal speed of calculation. Rather, the algorithm was structured to permit simple replacement of individual components, and permit parallelisation of the entire program. As a research tool it is not desire-able to rely on hardware optimisation as that ties the research to a single vendor, and often resources cannot provide for multiple installations at a single institution. A parallelised algorithm on the other hand may, if scalable, be run on a variety of platforms from a cheap PC clone to a MPP workstation. A framework for parallelising the code is discussed in Appendix: A.



Figure 5.24: Setup geometry for DRR calculation.

Geometric Setup

The next step after building a DRR raytracing engine was to develop an imaging analysis toolkit. The toolkit is responsible for manipulating the CT data set in order to generate DRRs from any position or orientation of the CT data with respect to the source position.

An automatic search algorithm must interface with the data set manipulation routines to search the transformation space, and reproduce the DRR that matches the reference image. The search algorithm draws measure of fit information from a registration metric function. This information must be collected, assessed, and used to direct the search algorithm's next step.

5.2 DRR Calculation

5.2.1 Ray-Tracing

The ray-tracing engine must be versatile and fast, quickly tracing both individual rays, and entire images. The toolkit requires some form of correlation measure in order to quantitatively assess the match between a calculated image, and the reference image. The toolkit must then be able to act on this information, and search the parameter space to find the DRR that matches the reference image. At this stage it is not important which method the toolkit utilises to succeed with each of these steps. Once the overall algorithm is in place, it is a simple process to replace individual algorithms. For instance, the switch from a photometric based measure of fit to template-matching requires different treatment of the data, however the basic framework will be in place.

The algorithm was developed from principles and basic functions used for computer graphics calculations. The basic setup follows the geometry of Figure 5.24, however the global co-ordinate system (GCS) is defined in terms of the X-ray source position, and not the center of the CT data set as used by previous DRR generation code. The CT data set is treated as a six-sided polygon placed between the X-ray source and the image plane. The X-ray source is fixed at the origin of the GCS which is measured in millimeters. The image plane is centered at [0, 0, -1099.31], and rays are defined using the *point-vector* definition of a line.

$$\vec{r} = \vec{x_1} + s\vec{x_2} \tag{5.13}$$

The separation distance was selected from an arbitrary commercial C-arm X-ray unit, but requires calibration on a unit-by-unit basis.

5.2.2 Data Set Orientation and Position

The data set orientation and position is defined by a set of three vectors: \vec{p} , \vec{a} , and \vec{b} as shown in Figure 5.25. The three vectors define the position of the top left corner and the direction and length of the CT slice rows and columns. This arrangement saves additional calculation when determining the intersection of rays with the data set. The \vec{p} vector defines the position of the top left corner of the top of the cube. This location corresponds to the top left corner of the first CT slice, or ct[0][0][0]. The \vec{a} vector defines the direction of the columns across the first slice, while the \vec{b} defines the direction of the rows down the first slice.



Figure 5.25: CT data set position, and direction vectors: \vec{p} , \vec{a} , and \vec{b} .

Manipulation of the Data Set

In order to generate DRRs from different views of the patient's anatomy, it is necessary to alter the position and orientation of the CT data set with respect to the X-ray source and image plane. This was accomplished by applying translation and rotation transformations to the three data set vectors. Translation was achieved by adding the translation matrix to the data set position vector, \vec{p} . Vectors \vec{a} and \vec{b} are not translated as they are direction vectors describing the orientation of the data cube relative to \vec{p} , the position vector. Rotation required that the three vectors first be referenced as points relative to the center of the CT data set. The transformation matrix is applied, and then the new points related back in terms of their relation to the X-ray source (GCS origin).

5.2.3 Numerical Integration

Integration Technique

In order to calculate a DRR, the ray-tracing algorithm must become a numerical integration scheme once a ray has successfully penetrated the CT data set. There are several successful techniques for calculating line integrals through CT data sets, however for the purposes of this study it was determined to generalise from the fastest equations.[KU94][Sid85] The reason for generalising is that these algorithms make assumptions such as: the requirements of a square data set, or interpolation taking place at every voxel edge intersected by the ray. In order to ensure fast ray-tracing, while still maintaining control over integration accuracy it is necessary to permit varied integration step sizes. This way the quality of the DRR may be tuned to match the requirements of the registration metric.

Voxel Dimensions

In order to integrate the data set, the computer must make a change of basis from the global co-ordinate system (GCS) in mm to the CT local co-ordinate system (LCS) which steps in voxel dimensions. The xy voxel dimensions may differ between machines, and the z voxel dimension depends upon the slice spacing selected by the radiologist when the CT image was recorded. This research utilised the visible human data set, and a standard slice separation of 3 mm as used for the male data set was selected.

There has been research in sub-sampling CT data sets in order to generate a 3D square data set from a non-square original. However this was not deemed necessary for the purposes of this study. The xy voxel dimensions were 0.8 mm as defined in the CT Header files. Thus incremental steps within the CT data set LCS proceed at 0.8 mm in the xy direction, and 3 mm in the z direction.

Changing Basis

The change of basis from the GCS to the CT LCS is achieved by calculating the point of intersection by the ray on the sides of the CT data polygon. Each side of the polygon is defined by three vectors: \vec{p}_k , \vec{a}_k , and \vec{b}_k . These vectors are calculated by simple transformations from the original data set basis vectors. \vec{a}_k and \vec{b}_k span the width and length of each face, defining the bounds for intersection. The ray is tested for intersection with each of the six polygon faces. In the case of an intersection, the intersection point is identified by two scalar multipliers, one for each vector. Since each vector spans the face the conversion from mm to voxels (LCS to GCS) is trivial, obtained by multiplying the scalar by a conversion constant.

Integration

The length of the ray segment passing through the data set is calculated, and the number of required steps is determined based upon a user-defined step size. A direction vector is calculated between intersection points, and used to step along the line in step size increments. At each step the line integral is calculated by interpolating with the 6-nearest neighbor voxels. The accumulated sum is the total density through which that ray passes.

The reason for the step size parameter is to permit selection of coarse or fine calculation of the DRR pixels. A coarse calculation would result in a lower quality DRR, but at a decreased computational cost. The minimum integration resolution required for consistent, accurate, registration would have to be determined by experimentation.

Interpolation

There are a wide range of interpolation techniques available for determining the density at each step along the ray. As a balance between speed and accuracy a 3D bilinear interpolation algorithm was generalised from 2D by performing two successive 2D bilinear interpolations, and then performing a 1D linear interpolation between the resulting density values (Equation 5.15).[PTVF94] Figure 5.26 illustrates the method in 2D.

$$y(x_1, x_2) = (1 - a_{x_1})(1 - a_{x_2})y(x_{n_1}, x_{n_2}) + a_{x_1}(1 - a_{x_2})y(x_{n_1+1}, x_{n_2})$$

$$+ a_{x_1}a_{x_2}y(x_{n_1+1}, x_{n_2+1}) + (1 - a_{x_1})a_{x_2}y(x_{n_1}, x_{n_2+1})$$
(5.14)

where:

$$a_{x_1} = (x_1 - x_{n_1})/(x_{n_1} - x_{n_1+1})$$

$$a_{x_2} = (x_2 - x_{n_2})/(x_{n_2} - x_{n_2+1})$$
(5.15)

CHAPTER 5. IMAGE REGISTRATION ALGORITHM



Figure 5.26: Two-dimensional bilinear interpolation.

Sampling by simply taking the local voxel density at each step location was found inadequate for producing an acceptable DRR. This is mainly due to the large data spacing in the z-direction (3 mm). The interpolated DRR in contrast appears quite clear and accurate.



(a) Interpolation



(b) No Interpolation



5.2.4 Additional DRR Considerations

There were two main factors involved in shaping the development of the DRR algorithm. The first considered the ability of the registration metric to match the images. Since the desired accuracy for computer-assisted surgery is greater than any previous medical imaging registration scheme (*sub mm* vs. 1-2mm) the DRR algorithm must be flexible enough to permit later refinements. The second factor concerns the speed of calculation. The registration procedure occurs during the operation, requiring the calculation time to be minimised.

DRR Accuracy

As alluded in the theory section, the DRR calculation algorithm is a simplification of the actual projection X-ray process. The fluoroscope is essentially a collection device, that quantises the continuous distribution of X-rays that fall on the detector surface. The digital fluoroscope limits the image to a greater extent as it is more selective, only photons containing at least E_{min} energy are counted. The DRR model assumes each pixel in the image represents the path along a single ray through the data set. In reality a pixel is actually the sum of all photons that land within the active area of the detector. Thus it might be necessary to actually calculate a series of rays falling on each pixel, and taking their average as the pixel value.

5.3 Optimisation Search Algorithm

There are a wide variety of optimisation strategies for multi-dimensional search spaces. For the purposes of this investigation it was important to implement a robust strategy that wouldn't be caught up in local minima. Radcliffe's work with Pseudocorrelation suggests that any optimisation strategy must be able to cope with numerous local minima along each parameter line in order to be successful. However a subsequent examination of the 6 DoF DRR search space illustrated a fairly smooth variance in comparison to Radcliffe's study (Figure 4.22), with one or two local minima along individual parameter lines. This would suggest a search algorithm may have more success than the grid search method. The difference in performance of the χ^2 statistic may be attributed to the variance in image quality between VEPID and DRR images. Although not a clinical setting, speed remains a consideration due to the 6 DoF search space. For these reasons the simplex optimisation algorithm of Nelder and Meade was implemented based on the algorithm provided in Numerical Recipes. As will be discussed in the Analysis Chapter, this algorithm exhibited some weaknesses that question its utility as a final solution.

5.3.1 Downhill Simplex - Nelder and Meade

There are numerous multi-dimensional optimisation search strategies, however the most efficient also tend to be sensitive to irregularities in the search space. One of the most robust, but sometimes inefficient methods, is the Downhill Simplex optimisation algorithm(Figure 5.28).

This strategy is likened to a multi-dimensional amoeba that stretches and contracts, oozing its way down to a minimum. As long as the basic step size is larger than any local minima, the amoeba will often step out of local minima and not become trapped like a garden variety gradient search strategy.

In order to pick a good starting point for the Simplex, 1000 randomly selected starting points (position and orientation) for the CT data cube were selected using the 'minimal standard random number generator' from Numerical Recipes. It is important to select a good random number generator to avoid clumping and an uneven distribution of points throughout the 6 DoF search space. The pixel intensities listed in Figure 5.29 were generated by a supposedly normally distributed random number generator resident in a spreadsheet.



Figure 5.28: Possible steps taken by the Amoeba, here a 4 vertice figure searching in 3D.[PTVF94]



Figure 5.29: Clumping of a poor random number generator.

As is evident from the figure, the values are actually found to clump in a few localised areas along the line. If this 'random' number generator were used to sample the 6 DoF transformation space, large areas of the space would remain unsampled, and localised areas would be oversampled. The starting points are generated by selecting translation and rotation parameters from within a 40 mm, 20° range about some arbitrary starting position. The CT data cube is transformed using these parameters, the pixels are traced and χ^2 is calculated. The 1000 starting points are then sorted according to ascending order of χ^2 , and the best 5 points are used as starting points for the Simplex.

At each starting point the Simplex is run with a large step size, a large negative step size, and then two intermediate step sizes. The Simplex requires a large negative step in case the first step happens to be in the opposite direction from the global minimum. In this situation the amoeba seems to become totally lost in the wilderness of transformation space. The final result is generally the best result the Simplex can generate given a particular starting



Figure 5.30: Variance of χ^2 as the CT Data Set is translated and rotated for 50 pixel points.

point. The best point of the 5 runs is used as a center for another 1000 random points with a smaller parameter range $(10mm, 2^o, 20mm z$ -direction).

The four Simplex restarts are repeated on the best 5 sampled points from the new 1000 point distribution. The reason for all of the restarts is that the Simplex sometimes folds up on itself, and fails to resolve down to the global minimum. Part of the problem is likely the extremely small gradient in the z-direction. The lack of sensitivity to translation along the z-axis is demonstrated in Figure 5.30. Ideally a change of optimisation strategy is called for after the first Simplex run, minimising along just the z-axis. This is discussed in detail in the Discussion section.

Chapter 6

Results

6.1 Overview

The preliminary results are promising, as outlined in Table 6.1. It appears possible to search the 6 DoF parameter space, and return the data set co-ordinates required to recreate the reference image. An extensive trial was conducted of 1800 runs with random starting positions located within a $(+/-20 mm, +/-20 mm, +/-30 mm, +/-15^{o}, +/-15^{o}, +/-15^{o})$ transformation space hypercube about the reference image. The algorithm was able to consistently find its way to the location and orientation of the data set used to generate the reference image. The three data set vectors were located with a mean error of less than 6 mm in each axial direction.

	p_x	p_y	p_z	a_x	a_y	a_z	b_x	b_y	b_z
Actual	-229.6	120	-1099.3	459.1	0	0	0	0	459.1
Mean	-232.0	119.1	-1094.9	458.4	6	-2.5	2.5	4.1	458.4
Std. Dev.	9.3	12.7	16.9	1.9	19.5	15.2	15.2	20.1	2.1

Data Set Position and Orientation (mm)

Table 6.1: Results for 1800 Random Starting Positions.

The standard deviations in the direction vectors convey the error in angular position of the CT data set. The rotation angles of the CT data set about the x-axis, α , is derived from the error in b_y . The y-axis, β , and z-axis, γ , are derived from a_z or b_x , and a_y respectively. A 20mm standard deviation is consistent with an error of ~ 2.5°, and 15mm represents an error of 1.9°.

6.2 Search Algorithm Results

The search algorithm was able to locate the position and orientation of the CT data set with fair accuracy for a Simplex. The mean χ^2 value was 0.012 with a standard deviation of 0.018. To place this result in perspective, an initial random starting location would exhibit a χ^2 of ~ 500-1000.

In order to obtain these results it was necessary to restart the Simplex multiple times, using the co-ordinates for the previously computed minimum, as outlined in Section 5.3. A typical scenario would be from an initial $\chi^2 \sim 500$, the first run through the simplex would improve the result to a χ^2 of ~ 100 . After resampling the parameter space, the restarted amoebas would generally locate the data set with a χ^2 of less than 1. However, improvement at this stage depended upon how smooth the search space topography was at each point close to the minimum.

The main problem is that there is considerable variance in accuracy between runs. The standard deviation for position in x and y was 9.3 mm and 12.7 mm respectively. The z component exhibited a particularly poor standard deviation, 16.9 mm due to the reduced change in the resultant image for magnification.

The error in the direction vectors, a and b, appears large however these vectors actually represent the angular orientation of the data cube. By appropriate trigonometry the resultant angular standard deviation: $\sim 2.5^{o}$ and 1.9^{o} are actually quite reasonable and close to

CHAPTER 6. RESULTS

the results of other studies.

DRR Calculation



(a) Reference Image

(b) Initial Guess

Figure 6.31 Sample DRR images of reference and initial guess.

Figure 6.31 illustrates the operation of the DRR algorithm in calculating the reference image, and a hypothetical initial guess by the registration algorithm, transformed from the reference location by $(+20 mm, +20 mm, +30 mm, +15^{o}, +15^{o}, +15^{o})$ This initial guess is intended to demonstrate gross misalignment, and is not indicative of the initial manual alignment performed by the clinician.

DRR Registration

The images in Figure 6.32 illustrate the relative shift and change in composition between the initial guess and the mean solution obtained by the search algorithm, relative to the reference image. The initial guess from Figure 6.31(b) is displayed as a feature template superimposed over the reference image in Figure 6.32(a) An image generated from the mean solution deviation, as listed in Table 6.1, is superimposed upon the reference in

CHAPTER 6. RESULTS



(a) Initial Guess



Figure 6.32. The initial guess and mean solution images as feature templates superimposed upon the reference image.

Figure 6.32(b) These feature templates were not used in the registration process, and are included solely as a visual aid for comparison between the unregistered and registered images.

As is evident from the figure, the algorithm is successful at gross registration of the images. Given a random starting location, the algorithm does on average achieve a close fit to the reference image. As far as its utility as a registration tool for robotic surgery, it is clear that there is definitely room for substantial improvement

Chapter 7

Discussion

The search algorithm is capable of achieving registration, however currently the variance between runs is too high for clinical application. It is important to consider that while there may be a 1-2 mm error in registration, that the machine tool will still perform the cuts to high precision. Currently, surgeons will achieve position and orientation accuracy of the order of 1-2 mm, and 1-2°. However the quality of fit between implant and cavity will be poor in comparison with a machine prepared site.

7.1 Search Difficulties

A step-by-step examination of the Simplex in operation yielded insight into its limitations. By the second sampling of the search space, the Simplex was generally very close to the reference position. At this stage it would often have all of the parameters nearly exact, except for p_z which could be as much as 30 mm off. On subsequent iterations the amoeba would often step off the exact xy positions in an attempt to reduce χ^2 .

7.1.1 χ^2 Scaling



(a) χ^2 Along X axis Parameter Line (1,0,0,0,0,0)

(b) χ^2 Along Multiple Parameter Lines

Figure 7.33. Comparison of variance of χ^2 along a single parameter line[0,0,1,0,0,0], and several parameter lines.

One particular problem faced by the search algorithm is the lack of even scaling by the distance measure along all parameter lines. For instance, the χ^2 gradient is so slight along the z axis, varying from zero to 1 2 for some 30 mm in translation, that the search algorithm will make small movements in xy for a large improvement instead of large movements in z for a small improvement. This was illustrated in Figure 5.30(a) Unfortunately, this reduced sensitivity along the z-axis is a function of the setup geometry, and not a result of the choice of registration metric.

The solution is to switch to a 1D optimisation strategy along the z axis after the Simplex has found its perceived global minimum. The search space is very smooth along the z-axis so it is likely that a 1D search algorithm would be successful in minimising despite the gradual gradient.

7.1.2 Local Minima

Another difficulty encountered by the search algorithm is illustrated in Figure 7.33. The gradient of a registration metric may be quite smooth along a single DoF, such as only x-translation of the image. The search algorithm however must contend with a complex topography containing many local minima as it searches in 6 DoF. While the Simplex is inherently robust, it is still fairly easy for the algorithm to be trapped in local minima near the global minimum. The ability of the amoeba to step out of local minima is extremely useful when far from the global minimum. When the Simplex has made its way close to the minimum however, it seems to step over and around the minimum into small pockets of local minima. Once in a local minima the Simplex will contract in about itself, and lose sight of the global minimum. The solution is to improve the ability of the search algorithm to handle local minima, or switch to a registration metric that exhibits a smoother topography.

7.2 Performance of the Registration Metric χ^2



Figure 7.34: Variance of χ^2 as the number of pixels sampled increases from 10 to 100.

There are a several aspects of the registration metric which could be modified to improve the match with the CT data. The simplest improvement is to simply increase the number of pixels sampled from the DRR. This effect is illustrated in Figure 7.34. Sampling only 10 pixels produces a χ^2 surface that is quite flat in comparison to a sampling of 100 pixels. All results listed in this section were derived from 100 pixel sampling. However trial runs using a larger number of pixels than 100 did not appear to realise an improvement in the final registration accuracy.

The quality of χ^2 as a registration metric could be a limiting factor in overall registration accuracy. Radcliffe opted to use a grid-based search method due to the complexity of the search space topography, and he was only operating in 4 DoF.[RRS94] One of the major factors in selecting χ^2 was the simplicity of its calculation, while still exhibiting characteristics of a good distance measure. It might be possible to improve the search space topography by switching to a different measure such as energy or moments. These metrics are slightly more complex to calculate, and carry an associated increase in computational burden. A proposed solution would be to utilise χ^2 until the algorithm is close to the global minimum, and then switch to a more complex measure.

Since this is a numerical simulation χ^2 resolves down to a global minimum of zero along all parameter lines when the images match. This however will not be the case for an actual image registration procedure conducted between CT data and a live fluoroscope image. In this situation it is unlikely the images will match exactly, and so when the images match there will be a global minimum, but this is not necessarily, and most unlikely to be zero.

7.3 Improving Overall Registration

The best solution to these problems would be to introduce a different optimisation strategy close to the global minimum. The Downhill Simplex, while robust, appears to have difficulty settling to the global minimum. The answer would be to switch to a standard gradient-based search algorithm once the likelihood of running into a local minima has been eliminated.

7.3.1 Improved Algorithm Structure

Hybrid Search Algorithm

A possible improvement on the algorithm would be to proceed as before with the random sampling of the search space, and permitting the amoeba to run through a trial. The subsequent minima would then, assuming χ^2 is below some threshold, be passed to a conjugate gradient method for minimisation down to the global minimum. The drawback of such a method is the requirement of calculating the gradient in all 6 DoF prior to each function evaluation. Such a method is really only viable if one is sure the global minimum is nearby. The same consideration would apply for simulated annealing methods, or other advanced minimisation strategies.

Change of Registration Metric

An alternative to simply adapting the search algorithm would be to alter the registration inetric. One example would be the introduction of a model or template-based metric after the pseudocorrelation algorithm had successfully located the local vicinity of the global minimum.

7.4 Comparison with Other Studies

The results from other studies are quite different from those listed here. The closest in philosophy centers around developing an automatic registration algorithm for image-guided radiosurgery by M.J. Murphy.[Mur97] In this study the goal is to register the patient's skull to a radiosurgery system, the Cyberknife. The author utilised a photometric measure, χ^2 , in conjunction with a masked DRR. The masking, isolated ROIs along the skull edge, serves to increase the consideration of features by the photometric measure. Murphy was able to locate the phantom with a standard deviation of 0.5-1 mm translation and 0.6°-1.3° rotation. It is important to note however that this study utilises orthogonal pairs of radiographs captured by digital fluoroscopes. Thus translations and rotations that are outof-plane in one radiograph will appear in-plane in the second radiograph. This permits a marked improvement in alignment accuracy.

The other radiation therapy studies also make use of orthogonal reference images. Their matching techniques may vary, but generally orthogonal views are relied upon to detect out of plane rotations. These studies are able to rely on the orthogonality condition due to the presence of the therapy machines. These gantry-type systems are designed to be capable of accurate, reproducible orthogonal images. Unfortunately, C-arm Fluoroscopes are not currently capable of accurate, reproducible orthogonal views. Thus in the general case it is not possible to assume orthogonality of the reference fluoroscope images. In future work it is intended to integrate a second fluoroscope image in improving registration accuracy.

7.4.1 Search Algorithm Performance

Improving the performance of the search algorithm will most likely rely on adapting both the registration metric, and the search algorithm as the matching process progresses. Ideally a simple metric like χ^2 would be used with a robust search algorithm in the initial phase of registration. Once the position and orientation of the CT data set has been localised and the search algorithm is within a few mm of an ideal match, the strategy should adjust both the metric and the search strategy to take advantage of a smoother metric topography. It does not make sense to rely solely upon either a search algorithm that takes a long time to converge upon a minimum, or a fast, accurate algorithm that gets lost if moderately far from a pair of matching images.

One additional point to consider is that in testing the performance of this registration algorithm, no attempt was made to accept or reject the registration results on the basis of the magnitude of χ^2 . One of the key selling points of χ^2 as a measure is that poorly aligned images result in a large value for χ^2 . This investigation has tried to evaluate the raw performance of a registration algorithm. An actual implementation would include a threshold for acceptable χ^2 , indicating a successful match. Utilising such a threshold would markedly improve the posted mean and standard deviation of the results. A threshold was not utilised in order to promote optimisation and refinement of the algorithm.

Chapter 8

Conclusion

There has been a flurry of recent development in computer-assisted surgery in the past few years. As computational power has increased, while the cost of computers has drastically fallen, so the ability to perform real-time image-based decision making has become a reality. One of the difficulties in conducting research in a field that is so rapidly developing is the reality that one's work may be outdated by technological development. Many of the issues surrounding this thesis grew from the technical reality of 1996. At that time Pentium class computers operating under 200 MHz were the norm, and incapable of the demands of realtime imaging. The main goal at that point in time was to develop an algorithm that would be capable of accurate registration in a timely manner. Since that time processor speeds have increased to speeds over 500 MHz, and more importantly bus and RAM access times have decreased considerably. It still remains quite difficult however to quickly match intermodal images without relying on specialised hardware or graphics processing computers.

The philosophy of this project has also been modified based on developments within the field of computer-assisted surgery. As computer-guidance has strongly established itself as the future of surgical technique so the applicability of computer-assistance has widened to

CHAPTER 8. CONCLUSION

a variety of anatomical locations. This development has led to the desire to develop a registration system that is not so feature-based that particular ROIs must be defined. A flexible system that can handle restricted fields of view, and various image paradigms. Computerassisted surgical guidance relies on providing a qualitative real-time representation of the location and orientation of surgical instrumentation relative to the patient. Robotic surgery however requires the additional quantitative measure of the relative position and orientation of the objects of interest.

The introduction of computer assisted techniques promises an increased number of favorable surgical outcomes. The crucial factor involved in implementing computer assistance is accessing, and acting on, accurate and timely diagnostic information during a surgical intervention. The implementation of an efficient intermodal registration algorithm will support the adoption of computer assisted techniques through a wide range of surgical applications.

These concerns have led to the development of a multi-modal image registration framework. The algorithm is modular, and designed to consider parallelisation at a later stage. The code does not require specialised or proprietary hardware, and can run on a variety of platforms. Unlike some work which reduces calculation time by pre-processing the CT data to consider only bony features, this registration scheme keeps all available information in case different applications require the detail. The intended use of the code is as a development platform. It will usually be more efficient to devise application specific algorithms. However in evaluating particular registration strategies, it makes sense to build from a flexible framework that minimises the development time required to build test each algorithms effectiveness at registration.

Appendix A

Parallelised Code

A.1 Overview

The raytracing code, and the registration algorithm both suffer from severe performance limitations under standard serial CPU computers. Most of the current clinical implementations of patient registration for computer-assisted surgery utilise advanced Silicon Graphics(SGI) hardware, and still require a couple of minutes to complete the registration process. The obvious solution is to turn to the power of parallel processing to ease the bottleneck.

In Parallel processing, a large task is divided into many small tasks and distributed among a number of computation engines. Often a performance increase results not only as a result of concurrent calculation, but also due to reduced memory access and cache flushing.

The algorithms and functions in this program were structured to facilitate conversion to parallel code. The reasons for this are two-fold. Firstly without specialised graphics hardware it is unlikely the current, or even next generation of PC's would be capable of real-time image registration in a clinical setting. The second concern is that in order to improve registration accuracy beyond +/-1 mm, $+/-1^{\circ}$, it is necessary to incorporate more sophisticated image analysis and registration techniques. The current generation of registration methods, which tend to follow a minimalist approach, are limited by the inherent accuracy of the data, and the limitations of the algorithms. More advanced registration techniques, such as utilising 3D model fitting, are far too computationally intensive to currently consider for real-time registration applications.

A.1.1 Parallel Virtual Machine - PVM

In keeping with the overall design philosophy of a platform independent calculation tool, the Parallel Virtual Machine (PVM) system was used as a development model.[GBD+94] PVM is a message passing protocol that permits a heterogeneous cluster of networked UNIX computers to appear and function as a single, virtual, parallel computer.

PVM was developed by Oak Ridge National Laboratory in collaboration with: the University of Tennessee, Emory University, and Carnegie Mellon University. The nice feature of this system is that it doesn't place any requirements on what the system components are, or even whether or not they are the same. This means that a cluster containing PC's, Pentium III's, and DEC Alpha's may all tackle the same problem, and computational load is distributed to maximise the use of resources. At the same time, the code could be run on a single massively parallel processor (MPP) without any modifications.

With the cost of PC's dropping rapidly, a high-powered multi-processor PC may be purchased for \sim \$5000. By implementing the PVM system one may achieve efficient computation in an extremely cost effective manner. This is especially important in medical imaging applications where research and clinical use occurs at the same facility. If an affordable hardware solution is presented, hospital imaging facilities can afford to have both clinical and research systems available to its staff. As an added bonus this situation precludes dependence upon a single hardware vendor.



(a) PVM Computation Model



Figure A.35: Illustration of the PVM computation model, and architecture schematic.[GBD⁺94]

A.1.2 Overview of PVM System

The PVM system consists of a set of software tools and libraries that enable a networked cluster of varied architecture computers to function as a single multiple processor computer. The PVM paradigm is broad enough to permit operation as either a parallel, or a concurrent solution computation engine.

PVM Structure

The basic PVM structure involves a host pool of computers that is defined by the user. This host pool may consist of a variety of computational architectures, and may be altered by adding or deleting machines during a calculation. The user may either choose to assign a task to the entire pool, or target specific machines within the pool.

The main program is initiated as a *parent* routine. The parent may spawn *child* tasks as required to solve a particular problem. The *child* tasks can all be executed on a single processor, or distributed throughout the entire pool. This flexibility within the PVM structure makes it an ideal system for medical imaging. Within a single image generation and registration process different tasks will require different levels of hardware support. Tasks may be spawned either to maximise computational efficiency on a dedicated clinical system, or on a development platform they may be spawned to share resources with other applications.

A.1.3 Relevance to Computer-Assisted Surgery

The benefits of parallel processing may be delivered throughout the field of computerassisted surgery. The entire field, from registration to real-time tracking and controls requires intensive computation in a timely manner. By implementing an algorithm modification like parallelisation, these tasks may be performed on standard computational platforms and end the reliance on high-end specialised hardware systems. This is a timely observation given the current financial constraints felt by the public health care system. One may argue the benefits of a sophisticated medical system, but if the initial hardware costs are prohibitive it will never be given the opportunity to demonstrate its merit.

For more information regarding PVM, and scientific parallel computing consult PVM: Parallel Virtual Machine. A Users' Guide and Tutorial for Networked Parallel Computing.[GBD+94]

References

- [AS88] C.K. Avinash and M. Slaney. Principles of Computerized Tomographic Imaging. IEEE Press, 1988.
- [Bar89] W.L. Bargar. Shape the implant to the patient. *Clinical Orthopaedics*, 249:73–78, December 1989.
- [Bar95] R.L. Barrack. Economics of revision total hip arthroplasty. *Clinical Orthopaedics*, 319:214, October 1995.
- [BM96] P.H. Bland and C.R. Meyer. Robust three dimensional object definition in ct and mri. *Medical Physics*, 23:1:99–107, January 1996.
- [BMBH95] K. Bouazza-Marouf, I. Browbank, and J.R. Hewit. Robotic-assisted internal fixation of femoral fractures. Proceedings of the Institution of Mechanical Engineers, 209:51–58, 1995.
- [Bro92] L.G. Brown. A survey of image registration techniques. ACM Computing Surveys, 24:325-376, December 1992.
- [Can89] Health Canada. Canada health reports, 1989.
- [CFP+85] E. Chaney, S. Fuchs, S. Pixer, J. Rosenman, G.W. Sherouse, E. Staab, and M. Varia. A three-dimensional imaging system for radiotherapy treatment planning. *Medical and Biomedical Engineering and Computing*, Proceedings of the XIV International Conference on Medical and Biomedical Engineering and the VII International Conference on Medical Physics:23:951-952, 1985.
- [Cow89] S.C. Cowin. Bone Mechanics. CRC Press Inc, 1989.
- [CTT⁺95] E.L Chaney, J.S. Thorn, G. Tracton, T. Cullip, J.G. Rosenman, and J.E. Tepper. A portable software tool for computing drrs. *International Journal of Radiation Oncology, Biology, Physics*, 33:1287–1300, May 1995.

- [DJB⁺98] A.M. DiGioia, B. Jaramaz, M. Blackwell, D.A. Simon, F. Morgan, J.E. Moody,
 C. Nikou, B.D. Colgan, C.A. Aston, R.S. Labarca, E. Kischell, and T. Kanade.
 Image guided navigation system to measure intraoperatively acetabular implant
 alignment. *Clinical Orthopaedics*, 355:8–22, October 1998.
- [DSJ⁺98] A.M. DiGioia, D.A. Simon, B. Jaramaz, M. Blackwell, F. Morgan, R.V. O'Toole, B. Colgan, and E. Kischell. Hipnav: Pre-operative planning and intra-operative navigational guidance for acetabular implant placement in total hip replacement surgery. technical report http://ortho.cor.ssh.edu/projects/hipnav/paper.html, Center for Orthopaedic Research, Shadyside Hospital, Pittsburgh, PA, 1998.
- [ESM⁺94] E. Eramzadeh, A. Sarmiento, H.A. McKellop, A. Llina, and W.G. Gogan. The cement mantle in total hip arthroplasty. *Journal of Bone and Joint*, 76A:77, January 1994.
- [FMH97] R. Fahrig, M. Moreau, and D.W. Holdsworth. Three-dimensional computed tomographic reconstruction using a c-arm mounted xrii: Correction of image intensifier distortion. *Medical Physics*, 24(7):1097–1106, July 1997.
- [GBD+94] A. Geist, A. Beguelin, J. Dongarra, J. Weicheng, R. Manchek, and V. Sunderam. PVM: Parallel Virtual Machine. A Users' Guide and Tutorial for Networked Parallel Computing. The MIT Press, 1994.
- [GDT⁺96] G.A. Gilhuijs, K. Drukker, A. Touw, P.J.H. van den Ven, and M. van Herk. Interactive three dimensional inspection of patient setup in radiation therapy using digital portal images and computed tomography data. Int. Journal of Radiation Oncology, Biology, Physics, 34:4:873–885, 1996.
- [GSF99] D. Glozman, M. Shohman, and A. Fischer. Efficient registration of 3-d objects in robotic-assisted surgery. Presented at CAOS, June 1999.
- [GvH93] K.G.A. Gilhuijs and M. van Herk. Automatic on-line inspection of patient setup in radiation therapy using digital portal images. *Medical Physics*, 20(3):667– 677, May/June 1993.
- [Jar99] B. Jaramaz. Xalign: Computer assisted radiographic measurement of cup placement using "synthetic" xrays. Presented at CAOS, June 1999.
- [JC83] H.E. Johns and J.R. Cunningham. The Physics of Radiology. Charles C. Thomas, 1983.
- [KU94] A.R. Krommer and C.W. Ueberhuber. *Numerical Integration*. Springer-Verlag, 1994.
- [LME⁺95] A.V. JR. Lombardi, T.H. Mallory, R.W. Eberle, M.B. Mitchell, M.S. Lefkowitz, and J.R. Williams. Failure of intraoperatively customized non-porous femoral components inserted without cement in total hip arthroplasty. *The Journal of Bone and Joint Surgery*, 77-A:1836–1844, December 1995.
- [McR94] R. McRae. Practical Fracture Treatment. Churchill Livingstone, 1994.
- [MGS⁺93] F.A. Matsen, J.L. Garbini, B. Sidles, J.A.and Pratt, D. Baumgarten, and R. Kaiura. Robotic assistance in orthopaedic surgery. *Clinical Orthopaedics*, 296:178–186, November 1993.
- [MMB⁺89] J.C. Mulier, M. Mulier, L.P. Brady, H. Steenhoudt, Y. Cauwe, M. Goossens, and M. Elloy. A new system to produce intraoperatively custom femoral prosthesis from measurements taken during the surgical procedure. *Clinical Orthopaedics*, 249:97–112, December 1989.
- [Mur97] M.J. Murphy. An automatic six-degree-of-freedom image registration algorithm for image-guided frameless stereotaxic radiosurgery. *Medical Physics*, 24:6:857– 867, June 1997.
- [MZHR99] N. Matsuura, W. Zhao, Z. Huang, and J.A. Rowlands. Digital radiology using active matrix readout: Amplified pixel detector array for fluoroscopy. *Medical Physics*, 26:5:672–680, May 1999.
- [oM] National Library of Medicine. Visible human project ct data set.
- [PBM+92] H.A. Paul, W.L. Bargar, B. Mittlestadt, B. Musits, R.H. Taylor, P. Kazanides, and B. Zuhars, J.and Williamson. Development of a surgical robot for cementless total hip arthroplasty. *Clinical Orthopaedics*, 285:57–66, December 1992.
- [PDH94] P. Potamianos, B.L. Davies, and R.D. Hibberd. A robotic system for minimal access surgery. *Journal of Engineering in Medicine*, 208H:119-126, 1994.
- [Pos92] R. Poss. Natural factors that affect the shape and strength of the aging human femur. *Clinical Orthopaedics*, 274:194–201, January 1992.
- [PTVF94] W.H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Flannery. Numerical Recipes in C. Cambridge University Press, 1994.
- [RKT+97] A.L. Redhead, A.C.W. Kotcheff, C.J. Taylor, M.L. Porter, and D.W.L. Hukins. Robotic-assisted internal fixation of femoral fractures. *Proceedings of the Insti*tution of Mechanical Engineers, 211:145–154, 1997.

REFERENCES

- [RRS94] T. Radcliffe, R. Rajapakshe, and S. Shalev. Pseudocorrelation: A fast, robust, absolute, grey-level image alignment algorithm. *Medical Physics*, 21:761–769, June 1994.
- [Ruc96] W. Rucklidge. Efficient Visual Recognition Using the Hausdorff Distance. Springer, 1996.
- [Rus92] J.C. Russ. The Image Processing Handbook. CRC Press, 1992.
- [SFH92] P. Suetens, P. Fua, and A.J. Hanson. Computational strategies for object recognition. *ACM Computing Surveys*, 24:5–61, March 1992.
- [SH95] I. Shuaib and M. Hillery. Forces generated in guide-wires when drilling human bone. Proceedings of the Institution of Mechanical Engineers, Part H: JOurnal of Engineering in Medicine, 209:157–162, 1995.
- [Sid85] R.L. Siddon. Fast calculation of the exact radiological path for a three dimensional ct array. *Medical Physics*, 12:2, March 1985.
- [Sim99] D.A. Simon. The tools for caos: Medical imaging, registration, and tracking devices. technical report, Medtronic Sofamor Danek, Broomfield, CO, USA, 1999. Presented at CAOS June 17-19, 1999.
- [SSW89] S.D. Stulberg, B.N. Stulberg, and R.L. Wixson. The rationale, design characteristics, and preliminary results of a primary custom total hip prosthesis. *Clinical Orthopaedics*, 249:79–96, December 1989.
- [Tre81] R.W. Treharne. Review of wolff's law. Orthopaedic Review, 10:35, 1981.
- [WF94] H. Wang and B.G. Fallone. A robust morphological algorithm for automatic radiation field extraction and correlation of portal images. *Medical Physics*, 21:2:237–244, February 1994.
- [Wie99a] U. Wiesel. Robodoc at bgu frankfurt experiences with the pinless system. Presented at CAOS, June 1999.
- [Wie99b] U. Wiesel. Using the robodoc system. Presented at CAOS, June 1999.