Improved seed-based MR and CT image registration for prostate brachytherapy

by

Elizabeth S. Kim

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2004

Author .....
Department of Electrical Engineering and Computer Science
May 20, 2004

Certified by .....
William M. Wells, III
Research Scientist, MIT CSAIL
Associate Professor of Radiology,
Harvard Medical School and Brigham and Women's Hospital
Thesis Supervisor

Accepted by
Arthur C. Smith
Chairman, Department Committee on Graduate Students

Improved seed-based MR and CT image registration for prostate brachytherapy

by

Elizabeth S. Kim

Submitted to the Department of Electrical Engineering and Computer Science
on May 28, 2004, in partial fulfillment of the
requirements for the degree of
Master of Engineering

Abstract

Prostate cancer's high incidence and high survivability motivate its treatment using tightly focused radiation therapy. Brachytherapy treatment, the implantation of radioactive seeds into the prostate, is increasing in popularity, spurred by advances in medical imaging techniques for prostate visualization. Successful brachytherapy requires precise positioning of implant seeds within the pelvic anatomy. Following implantation, precise localization of individual seeds is required to evaluate treatment, but this remains an open challenge. This thesis addresses the seed localization problem with contributions for improving seed-based registration of MR and CT post-implant images. A model for non-rigid, affine prostate motion is presented and demonstrated to improve on current techniques of rigid registration. Also, an evaluation of the benefit of using multiple, rather than a few, seeds is presented, along with a scheme for validating registrations using manually detected seeds in MR and CT volumes. Finally, a scheme for automatic seed-based MR and CT registration by aligning *all* seeds is suggested, with supporting algorithms for CT seed-finding and unmatched feature registration. A call for an MR seed-finder is issued, for this is the final component needed to achieve automatic and complete seed-based MR and CT registration.

Thesis Supervisor: William M. Wells, III
Title: Research Scientist, MIT CSAIL
Associate Professor of Radiology,
Harvard Medical School and Brigham and Women's Hospital

Acknowledgments

I thank Sandy for believing in me. He has generously shared his experience, ideas and community with me. He has guided my research and growth with an eye for good science and good fun. I have enjoyed working with him late at night and plumbing the depths of his bottomless store of patience.

I owe many thanks to Clare for the trust she has invested in me by giving me this project and funding my research with her NIH grant. She has assembled a lively team of sharp and friendly scientists, including Steve, whom I thank for listening, being honest and giving technical advice and tight code. I also thank Rob for his enthusiasm and answers to my many questions.

Thanks to MIT for embracing the hacking ethic. This has been a wonderful place to live and work. Long live athena!

Thanks to Wanmei for reaching out to me, feeding me and showing me the real meaning of hard work. Thanks to Dr. Demirdjian, Lilla, my officemates and the vgn neighborhood for being friendly. Cheers to the MEng boys for struggling with me.

Thanks to Veronica, Lucy, and Chosetec for taking living, fun and friendship seriously.

Thanks to Mom and Dad for their dedication to raising me, for music, and especially for paying for MIT! They have waited for me when I've wandered. Thanks to Catho and Anne for knowing what I mean and for feeling my pain and happiness. Komapsubnida to Auntie and Grandma for their wordless support.

Finally, I thank David for teaching me how to give myself a chance. I owe much of my happiness to him.

The NIH has provided the funding which supported my research (NIH R01 AG 19513). I extend my sincere gratitude to them for supporting humanity's collective growth through scientific advancement.

Contents

Chapter 1

Introduction

Medical image analysis has supported the development of diverse clinical treatments, including image-guided surgeries, biopsies, and radiation therapies, for the brain, breasts, pelvic and other regions in the body. This thesis presents image registration techniques for application to the treatment of prostate cancer.

This chapter describes prostate cancer, brachytherapy treatment for prostate cancer, the role of medical imaging in brachytherapy, and methods to improve post-implant evaluation of brachytherapy. It concludes with a description of the organization of the thesis.

1.1 Medical context and significance

1.1.1 Prostate cancer

Prostate cancer is the most common non-cutaneous malignancy among American men, with one in six American men expected to develop prostate cancer in his lifetime[5]. It is the second leading cause for death from cancer among American males.[5]. The high incidence of prostate cancer motivates significant research activity toward the development of treatments for this disease.

1.1.2 Motivation for brachytherapy treatment

For patients whose cancer is local to the prostate, three main treatment forms are available [5]. The first, radical prostatectomy, involves surgical removal of the prostate gland. The second, external beam radiotherapy, exposes the prostate to high-energy X-ray beams from outside the body. The third and most novel treatment option, brachytherapy, involves temporary or permanent implantation of radioactive *seeds* into the prostate. Both external beam radiotherapy and brachytherapy depend heavily on medical images for guiding and evaluating treatment.

The anatomy of the male pelvis presents challenges in the treatment of prostate cancer. The prostate sits immediately anterior to the rectum, and the urethra passes through the prostate. Therefore, all forms of prostate cancer therapy are known to damage, to varying degrees, healthy tissues which neighbor the prostate, adversely affecting urinary, rectal, and sexual functions [23, 7]. Figures 1-1 and 1-2 depict the prostate's anatomy, including substructures of the prostate, such as the *central gland* and the *peripheral zone*. The peripheral zone has been identified as the site where most prostate cancers arise, making it a region of special interest for prostate cancer treatment [13, 14].

Because 98% of all prostate cancer patients survive for five years after diagnosis, and 84% survive ten years, care providers strive to treat the prostate while preserving the urethra and rectum [5]. In other words, the ideal prostate cancer treatment maximizes tumor control while minimizing morbidity (adverse effects of treatment).

Brachytherapy implants deliver a high radiation dose over a small, focused area [12]. Thus, if precise and accurate implantation can be achieved, brachytherapy promises to target diseased tissue while preserving healthy tissues. Brachytherapy is less invasive and less expensive than prostatectomy, and requires fewer hospital visits than external beam radiotherapy. Due to its promise and practical advantages, brachytherapy is increasing in popularity [5].

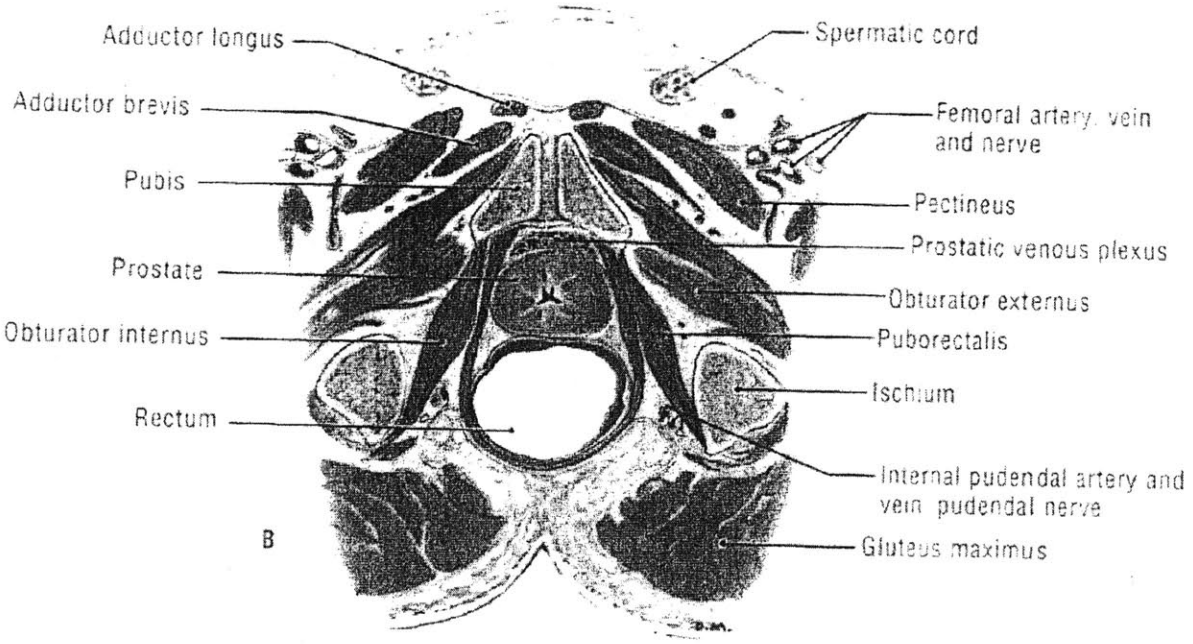

(a) Axial view of the male pelvis

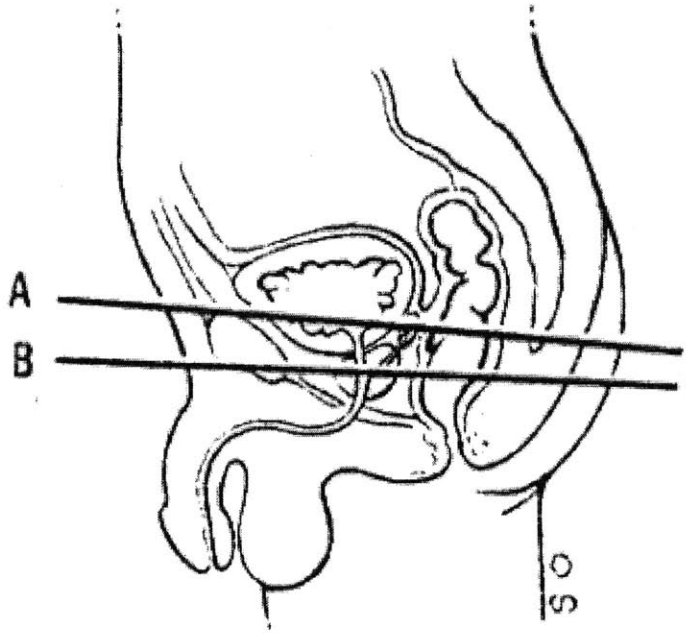

(b) Sagittal view of the male pelvis, indicating the location (lower of two cross-sections, marked "B") of the slice shown in (a).

Figure 1-1: Drawings showing axial (from the bottom) and sagittal (from the side) views of the prostate and surrounding pelvic anatomy. (a) shows the prostate's situation amidst the entire pelvis. (b) indicates the location of the slice shown in (a).

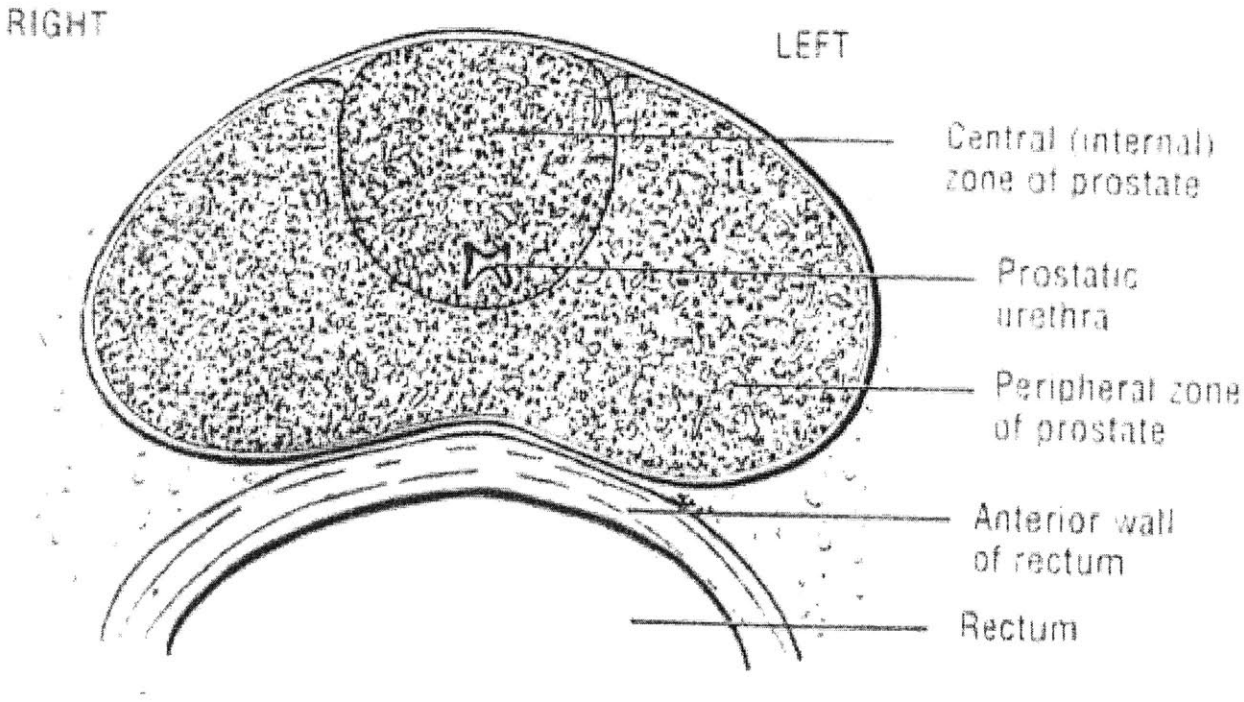

Figure 1-2: This drawing shows the prostate, rectum, prostatic urethra, and substructures of the prostate, including the central gland and the peripheral zone. Most cancers are found in the peripheral zone, making it a special target for brachytherapy.

1.1.3 History of prostate brachytherapy

In 1903 Alexander Graham Bell proposed treating prostate cancer by inserting radioactive sources into the prostate [22]. Within a decade, the radioisotope ^{226}Radium was being temporarily inserted into the prostatic urethra to treat cancer, and soon after Radium needles were inserted into the prostate, via the perineum, guided by "a finger in the rectum" [21]. Despite successful reduction of lumps using this technique, the medical community paid little attention to prostate brachytherapy until a small burst of development in the 1970s and 1980s. Only within the last decade has prostate brachytherapy attracted widespread research attention and clinical use, due to advances in medical imaging technology [22, 9, 21, 14].

1.1.4 Visualizing the prostate with medical images

Three image modalities are commonly used to visualize the prostate for a variety of clinical purposes, including brachytherapy treatment. Each modality offers unique advantages and suffers specific drawbacks.

Computed Tomography (CT)

Computed Tomography (CT) images are three-dimensional X-ray images. CT images are generated by taking multiple 2D X-ray images in different planes and solving the tomography (inverse projection) problem. Each discrete, 3D element, or *voxel* of a CT image, maps to a grayscale intensity, which corresponds to the density of the material imaged. Bones, implant seeds, and other dense materials appear in bright white in CT images. Various soft tissues, such as muscles and organs, all have approximately the same lower density, and appear in CT with duller intensity, in gray. CT provides excellent spatial resolution and excellent contrast among hard tissue, soft tissue, and air. Patients undergo CT scanning within an environment specialized for CT imaging. In general, the major drawback of CT imaging is its poor differentiation of soft tissues.

CT images visualize brachytherapy implants as crisp, bright spots, but neither clearly delineate the prostate from the neighboring rectum, nor differentiate the sub-

structures of the prostate. As of 2000, the American Brachytherapy Society asserted that CT is the best modality for imaging seeds [19]

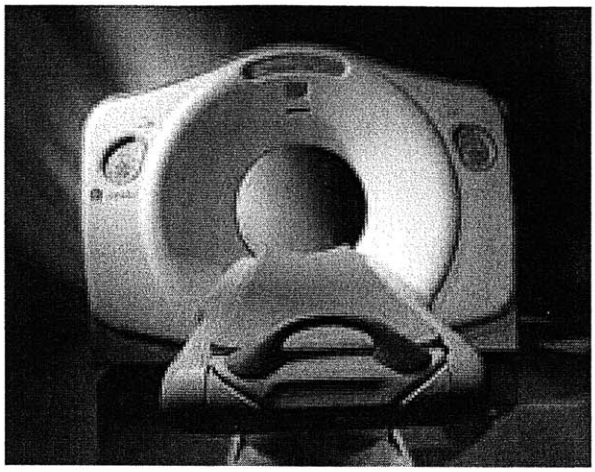

(a) CT scanner

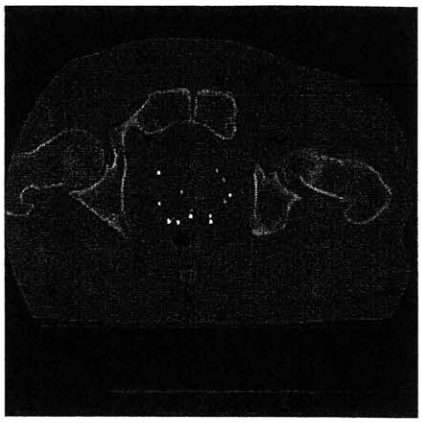

(b) CT scan of prostate

Figure 1-3: Brigham and Women's Hospital (BWH) uses the CT scanner (GE LightSpeed QX/i) shown in (a) to image prostate seeds. An axial slice from a CT volume of the post-implant prostate is shown in (b).

Ultrasound (US) Imaging

Ultrasound (US) is a nearly ubiquitous technology, best known for imaging fetuses in pregnant women. US generally produces two-dimensional images, with intensity values generated using sonar-based echo location. Three-dimensional US has recently become available, but is not yet popularly used. Soft tissue contrast in US is good, but air cavities and bones are imaged poorly. Another drawback of US imaging is limited spatial resolution and noisy image signal. US's primary advantages are its inexpensive operation and equipment cost, and its flexibility for use in a variety of environments, including interventional settings.

In the brachytherapy context, ultrasound is popularly used for intraoprative, real-time imaging feedback because it is inexpensive and easily applicable to the treatment setting. Unfortunately, US images seeds very poorly, and US signals are disrupted by the seeds [19]

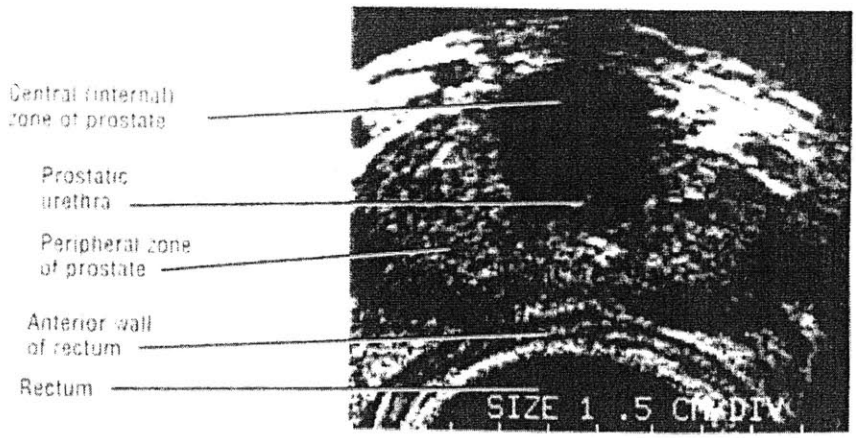

(b) US scan of prostate [6]

Figure 1-4: Ultrasound produces noisy images of the prostate.

Magnetic Resonance (MR) Imaging

Of all medical imaging modalities, magnetic resonance (MR) imaging provides the best contrast among different types of soft tissue. MR images are three-dimensional, with intensities corresponding to the nuclear magnetic resonance (NMR) properties of the materials being imaged. The main drawback of MR imaging is the often-prohibitively high cost of purchasing and operating an MR scanner. Purchase of an MR scanner costs millions of dollars. Another disadvantage of MR is that like CT, most MR scanners are incompatible with a surgical or interventional environment, and can be used only within specialized imaging settings.

Because of its superior soft tissue differentiation, MR is the imaging modality of choice for visualizing pelvic anatomy. MR not only cleanly delineates the rectum, urethra, and prostate, but also defines prostatic substructures. Among the substructures shown in MR is the *peripheral zone*, the region where most prostate cancers originate, and which clinicians are especially keen to target [13, 14]. MR is the best imaging modality for differentiating the peripheral zone from the central gland.

Unfortunately, seed visualization is problematic in MR. Seeds produce no MR signal, so MR images represent seeds as amorphous black spots, or signal voids [19]. Where two or more seeds lie close together, MR images present large black spots,

from which the number of seeds imaged cannot be determined. A variety MR imaging sequences, for example *T1-weighted*, *T2-weighted*, and *Spoil Gradient Acquisition in the Steady State (SPGR)*, can be employed to highlight specific tissue or material types. The distinctions between these sequences can be seen in Figure 1-5.

Interventional Magnetic Resonance (IMR) imaging

Generally patients' bodies are inaccessible during MR scanning. The patient lies on a table which is inserted into the bore of a very large, monolithic magnet.

Recently, interventional magnetic resonance (IMR) scanners have been developed, which split the scanner's magnet into two parts, separated by enough distance to a physician to access the patient during imaging. Because of its enormous expense, very few institutions operate IMR devices. IMR devices produce images which allow doctors to produce images with high quality soft tissue differentiation while performing surgery or other interventions.

1.1.5 Brachytherapy procedure

A brief description of the brachytherapy treatment procedure will be valuable to understanding the problem addressed by this thesis. The following is a brief overview of the brachytherapy implantation procedure used by collaborators at the Brigham and Women's Hospital (BWH) in Boston, MA. This description is summarized from the protocol published by D'Amico, Cormack, Tempany, et al. [14, 11]. Brachytherapy treatment is performed in fundamentally similar ways at most institutions. A note on key procedural variations follows this description of BWH's procedure.

Brachytherapy is performed on patients in the early stages of prostate cancer, at which point the cancer is local to the prostate. The entire brachytherapy treatment takes a few hours.

Prior to implantation, clinicians acquire a three-dimensional Magnetic Resonance (MR) image of the patient's pelvis in order to visualize and identify the target volume of tissue. Using this pre-operative image, clinicians plan where to implant the seeds.

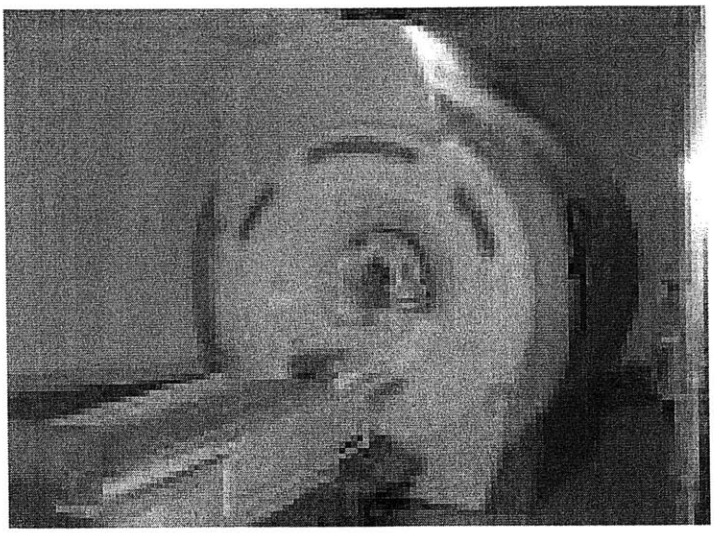

(a) MR traditional scanner

(b) T1-weighted MR axial slice.

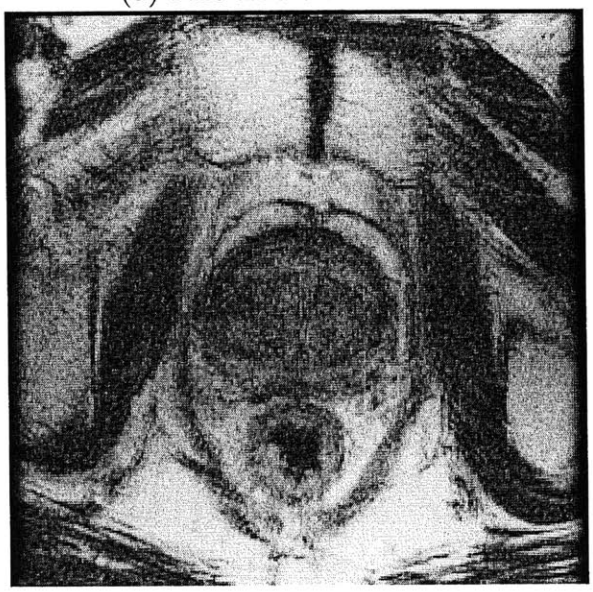

(c) T2-weighted MR axial slice

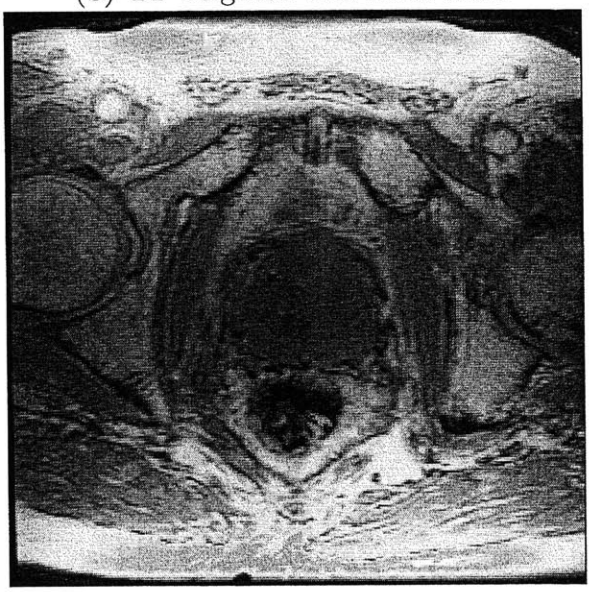

(d) SPGR-weighted MR axial slice.

Figure 1-5: MR images are captured using a large magnet as shown in (a). Different sequences highlight different tissues, as shown in images (b) of T1-weighted MR, (c) of T2-weighted MR, and (d), of SPGR sequence MR. T1- and T2-weighted images produce superior soft tissue differentiation but poor seed differentiation, while SPGR sequences resolve seeds as black signal voids but do not differentiate soft tissue well. These three MR scans were taken in one session, less than an hour long, without moving the patient between scans. The consistency in patient positioning throughout these three scans introduces the possibility of easily registering these images to achieve MR's promised soft tissue differentiation with suggested seed positions.

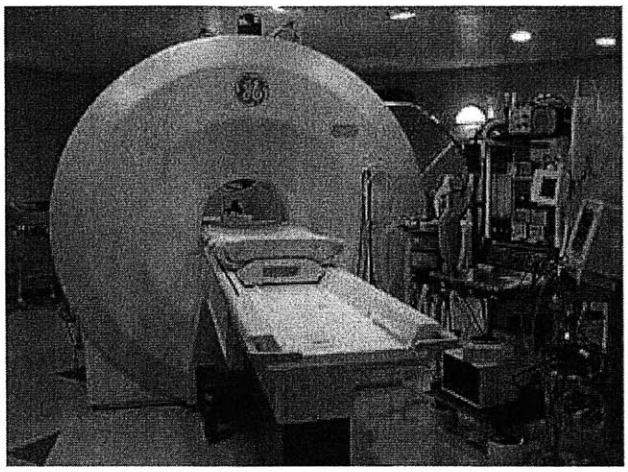

(a) Front view of interventional MR scanner.

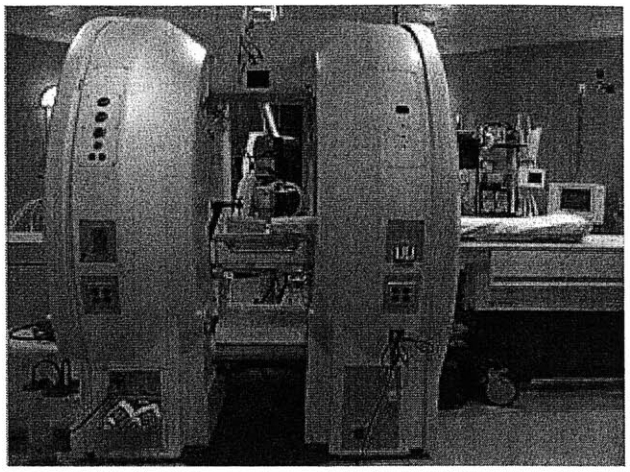

(b) Side view of IMR scanner.

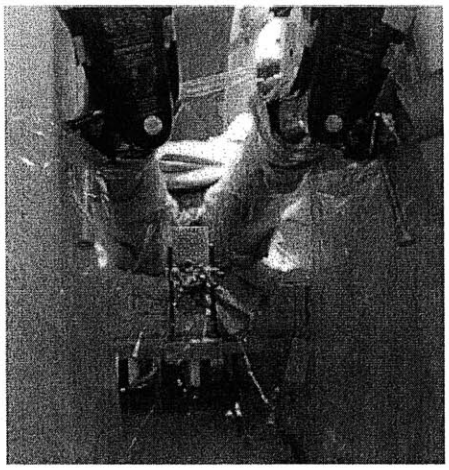

(c) Brachytherapy patient in IMR scanner.

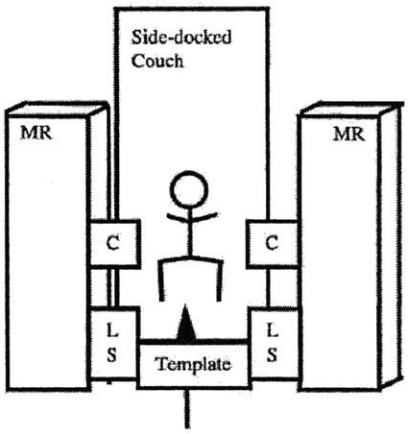

(d) Patient's position in IMR scanner.

Figure 1-6: Compared with a traditional scanner (see Figure 1-5), an interventional MR scanner's magnet is cut into two pieces. Shown above, in (a) and (b), is BWH's interventional MR scanner (GE Signa SP), configured for loading a patient through th bore of the magnet. Photograph (b) shows why IMR scanners are often called "open magnets." Photograph (c) shows a prostate brachytherapy patient loaded from the side, in BWH's IMR scanner, for prostate brachytherapy. The patient's legs are raised in the lithotomy position, and a needle guidance template is situated before the patient's perineum [11]. The drawing in (c) offers another perspective of the patient's position within the IMR scanner [11].

The planning of seed positions must be precise and accurate because seeds deliver a high radiation dose to immediately neighboring tissues, and the dose falls rapidly with distance [12].

Implantation

During implantation, the patient lies on a table, in the lithotomy position, with legs raised and spread to provide access to the perineum. Large needles are loaded with radioactive seeds. The physician then inserts needles via the perineum into the prostate. Following the insertion of each needle, an MR image of the patient's pelvis is captured within the surgical setting. The resulting intraoperative image provides real-time feedback about the needle's, and thus the seeds', actual position, which clinicians compare with planned positions. Discrepancies between the planned and actual seed positions can be corrected at this time.

When the needle position has been adjusted to fit the plan, the physician deposits the seeds and removes the needle. This process is repeated until all planned seeds have been implanted.

Typically 40-120, or an average of about 80, seeds are implanted. Seeds are titanium cases filled with a radioactive isotope. At BWH and many other institutions the radioactive isotope used is ^{125}Iodine, with a half-life of 60 days [12, 21]. Other institutions use the radioactive isotope ^{103}Palladium, with a half-life of 17 days[21]. Seeds deliver a clinically effective dose of radiation, enough to kill surrounding tissue faster than the tissue can regenerate, for one to two half-lives [10]. At 4.5 mm in length and 0.8 mm in diameter, each seed is just smaller than a grain of rice, as can be seen in Figure 1-7 [21]. At BWH, the seeds are implanted permanently.

Intraoperative image guidance

Intraoperative imaging provides *real-time dosimetric feedback*, giving clinicians immediate feedback about the dose distribution over pelvic volumes. BWH performs brachytherapy under interventional MR guidance, a technology described in Section 1.1.4. The IMR scanner's split magnet is a very recently developed, very expensive

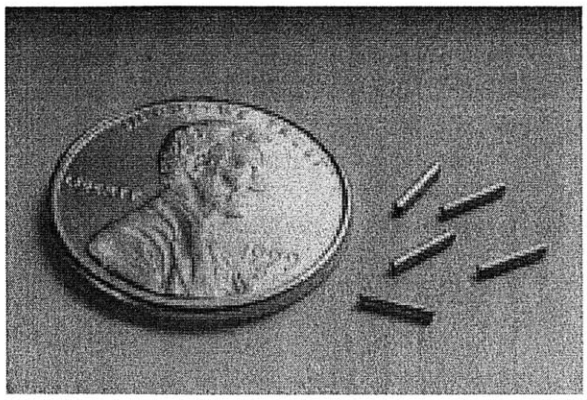

Figure 1-7: Radioactive seeds, shown here next to a penny, are just smaller than a grain of rice.[1]

technology, used by very few institutions. More popularly, trans-rectal ultrasound (TRUS) is used to provide real-time dosimetric feedback during brachytherapy. As described in Section 1.1.4, MR imaging provides significantly higher quality images and delineation of the prostate and its substructures than does US, making it a much more powerful intraoperative guidance tool. The advantages of IMR over TRUS-guidance are discussed by D'Amico, Cormack, Tempany, et al. [14, 11, 12]. CT intraoperative guidance is also possible, but is very rarely used [19].

No matter the implantation protocol or choice of image-guidance modality, in all cases, after implantation an evaluation is necessary to determine the actual dose distributed to the pelvic tissues.

1.2 Post-implant dosimetry and seed localization

Knowledge of the seeds' final resting positions allows doctors to evaluate the success of the implantation. Precise seed localization with respect to anatomical structures can indicate whether adequate doses of radiation are being delivered to the targeted tissue, and whether potentially dangerous doses are being delievered to healthy tissues.

Post-implant dosimetry is the problem of measuring actual dose distributions with respect to the patient's prostate and neighboring tissues. A necessary precursor to post-implant dosimetry is precise *seed localization*, with respect to the prostate, rec-

tum, and urethra. Some clinicians are also interested in knowing seed positions with respect to the substructures of the prostate, specifically the peripheral zone [14].

During and after implantation, the prostate and surrounding tissues swell, possibly causing seeds to shift within the patient's body. Therefore, despite knowing the locations of seed-bearing needles at implantation time, post-implant seed localization is not straightforward.

Precise seed localization for post-implant dosimetry remains a major challenge in prostate brachytherapy. This problem is the clinical motivation for this thesis.

1.2.1 MR and CT registration for seed localization

Given the complementary strengths of MR and CT modalities for visualizing soft tissue anatomy and implant seeds, MR and CT image registration has emerged as a popular approach to seed localization. This general technique, also called *fusion*, involves reformatting either image to align with the other image to allow for combination of the visual information in the two images. In the context of seed localization, the goal of MR and CT image registration is the alignment of seeds.

Early methods in post-implant multi-modal registration between MR and CT or US and CT aligned anatomical surfaces such as contours of the urethra, bladder, or rectum between the two images [20] [8]. Similar techniques have been used to solve registrations for external beam radiotherapy treatment applications. Anatomical contour-based regitration techniques may roughly align the prostate capsule or the prostatic urethra between two images, but they do not specifically align individual seeds. At best they allow localization with respect to the prostate capsule as a whole.

A seed-based alignment was first introduced by Dubois [16], and this style of registration has since become the standard for MR- and CT-based dosimetry [10]. Commercial vendors now produce software to perform rigid MR and CT alignment based on manually selected corresponding seeds.

1.2.2 Limitations of existing seed-based MR and CT registration methods

The prostate's deformability is a well-known and frequently observed phenomenon. Van Herk showed that the prostate is subject to deformation due to slight movements of the legs and filling of the bladder and rectum [18]. Existing post-implant dosimetry registrations lack a model which captures the prostate's deformability. A non-rigid model is needed to better represent the movements of the prostate between post-operative image scans.

Also, in practice, existing seed-based registration methods use only three pairs of seeds, rather than the 80 or so pairs of seeds which are implanted and which must be precisely aligned. This effective "three-seed limit" is due primarily to the inconvenience of seed-finding, which is a manual process. This limitation speaks to a need for automatic seed registration, and an implied need for automatic seed detectors.

Finally, there is a lack of methods for validating alignments of seeds.

1.3 Contribution

This thesis addresses the need for a non-rigid model of motion, more automatic registration, and validation of seed localization.

First, we verify that improved registration accuracy results from the use of a large number of seeds. Second, we establishe that seed-based registration is improved using a non-rigid model of prostate motion. Third, we contribute a method for registering automatically detected seeds. Finally we suggest methods for automatic seed detection in MR and CT post-implant images. The combination of the above four contributions culminates in a scheme for validation of post-implant registration techniques, and suggests an automatic validation scheme.

In short, this thesis immediately meets the first need, by contributing improved model of non-rigid prostate movement, and addresses the third need by demonstrating

a manual validation scheme. The second need, for increased automation of post-implant seed localization presents further research challenges in the area of automatic seed detection.

1.3.1 Immediate usability in clinical settings

Improved registration with non-rigid affine modeling can be immediately applied in the clinical setting, to replace the current standard three-point rigid registration algorithm. The concept of validation by multiple seeds can also be immediately applied to evaluate the the quality of existing registrations.

In the future, if a reliable MR seed-finder can be designed, then this work provides a registration scheme and a CT seed-finder which will perform automatic, seed-based non-rigid registration of the prostate, using all seeds. The process will require an initial registration which can be generated using intensity-based or current three-point techniques. This technique promises excellent seed localization because all seeds will be individually aligned.

1.4 Organization of Thesis

This thesis begins by presenting detailed background information which will facilitate the reading of the remainder of the thesis. An automatic seed detection algorithm, which works reliably in CT and problematically in MR, is presented in Chapter 3. Chapters 4, 5, 6, and 7 describe the design and performance of rigid and affine registration techniques for matched and unmatched features. Experiments in Chapter 7 demonstrate the need for multiple-seed alignment, the utility of a non-rigid affine model, and the feasibility of registration on unmatched seeds. This thesis concludes in Chapter 8 with a review of contributions made by this thesis and a survey of directions for further research.

Chapter 2

Background

A background on image acquisition, representation, and registration will facilitate an understanding of the problem of seed localization by MR and CT registration, and the present research. This chapter provides this background information as well as a detailed view of current, widely used approaches to post-implant dosimetry.

2.1 Post-implant image acquisition

All data which motivated, and which was used to evaluate, this research was collected, with appropriate consent, at BWH from patients undergoing prostate brachytherapy treatment, as described in the Introduction (see 1.1.5). Six weeks following implantation, CT and MR scans were acquired on the same day at separate imaging suites within the hospital. Time between scans was allowed to span the entire day.

CT images were using a General Electric Medical Systems (Milwaukee, WI) LightSpeed scanner, typically with 320 mm field of view and 1.25 mm slice thickness; matrix size 512. This corresponds to voxels of dimension 0.625 x 0.625 x 1.25 mm.

Post-implant MR images were scanned using a conventional MR scanner (Signa 1.5T, General Electric, Milwaukee, WI), typically with 200 mm field of view and 2 mm slice thickness; matrix size 256. For each patient a series of different MR scans were collected in the same session, including T1-weighted, T2-weighted, and Spoil Gradient Acquisition in the Steady State (SPGR) sequences. Each of these MR

imaging techniques produces high contrast of a specific type of tissue or material. T2-weighted images provide the best delineation of the prostate, rectum, urethra and prostatic substructures. Meanwhile SPGR produces the highest contrast, among various MR techniques, between seeds and soft tissues, while also providing effective, if not ideal, differentiation of the prostate from neighboring anatomy. Despite its leading the MR techniques in seed visualization, SPGR seed imaging is nevertheless problematic in the ways described in 1.1.4, necessitating registration with CT images.

Since no single post-implant MR technique provides ideal visualization of both anatomical structures and seeds, the T2 and SPGR datasets can be registered to produce a dataset featuring excellent soft tissue contrast and slightly problematic seed visualization. This technique of registering CT with SPGR sequences, and then registering SPGR with T2-weighted MR is not yet used, but is expected to provide optimal seed localization.

2.2 Image representation

Of the various MR scans collected postoperatively, the SPGR scan is selected for registration with CT because of its superior resolution of seeds. In general, this thesis will use *post-implant MR* to refer specifically to the SPGR image.

Post-implant CT and MR images are three-dimensional, or volumetric, images. They are represented as three-dimensional matrices. Each matrix element, called a *voxel*, maps to a grayscale intensity (brightness) value. We can think of volumetric images as functions mapping 3D coordinates to intensities.

Consider an arbitrary image defined over a coordinate space S in R^3. The image is an intensity function $f : S \rightarrow intensity$ over the space. Let each point in the image $p \in S$ be represented as a column vector in R^3, $p = \begin{bmatrix} p_x \\ p_y \\ p_z \end{bmatrix}$. The grayscale intensity of the image at any point p in the image is $f(p)$.

2.2.1 Homogeneous Coordinates

Often in image processing, and specifically in the present research, we will find it convenient to use homogeneous coordinates to represent image point coordinates. A three-dimensional point $p = \begin{bmatrix} p_x \\ p_y \\ p_z \end{bmatrix}$ can be represented in homogeneous coordinates as

$$p_h = \begin{bmatrix} p \\ 1 \end{bmatrix} \text{ or } p_h = \begin{bmatrix} p_x \\ p_y \\ p_z \\ 1 \end{bmatrix}.$$

The additional, fixed, unitary parameter facilitates compact representations of rigid and affine transformations. The utility of this representation will become apparent in 4.1.1 and 5.1.

2.3 The general image registration problem

Image registration is the problem of reformatting one image, which we will call the *floating image*, into the coordinate space of another image, which we will call the *fixed image*. Image registration is solved by finding the floating image's position, rotational orientation, and possible deformations in the fixed image's coordinate space. The solution is given as the parameters of a *coordinate transformation* which maps locations in the floating image to their corresponding locations in the fixed image's space.

A coordinate transformation maps a coordinate p from the floating image to a coordinate p' in the fixed image, according to a transformation T. In other words, $p' = T(p)$.

If the floating image moves into the fixed space as a rigid body, then the transformation represents changes only in position and rotational orientation, and is called the floating image's *pose* in the fixed space. In this thesis, we will slightly abuse this terminology and let *pose* refer to slight non-rigid deformations as well as rigid

movements.

2.4 General solution sructure for image registration

As just explained, the goal of an image registration algorithm is to find the transformation parameters which move the floating image into optimal alignment with the fixed image.

An image registration algorithm is comprised of three key components:

- A *coordinate transformation*, or model of motion or deformation.

- A *objective function of similarity*, or measurement of the quality of the alignment. This should be based on comparisons of the two images' intensities or features.

- An *optimization strategy* or method for searching for the optimal transformation parameters, as measured according to the the similarity objective.

In somewhat rare instances, optimization strategies can be implemented as simple linear systems. More generally, iterative process are employed to search for an optimal alignment.

The process of iteratively solving a registration is a loop over the following steps:

1. Produce an alignment; that is, assign values to the parameters of the coordinate transformation which will align the floating image into the fixed image's space.

2. Reformat the floating image according to the new alignment. Measure the similarity between the aligned images, in terms of the similarity objective.

3. Use the optimization strategy to guide the next iteration of the process toward a better similaritiy.

The result of this iterative process is a set of transformation parameters, also sometimes called a *pose estimate*, which describe the optimal alignment, as measured by the specified similarity objective.

In many cases, the optimization strategy may be a local method, which is subject to becoming trapped in local extrema. Local strategies also require initialization to rough estimates of the optimal alignment. Without such an initilaization, local strategies sometimes converge on false, locally optimal solutions.

2.5 State of the art in dosimetry: three-point rigid (3PR) registration

Currently BWH performs post-implant dosimetry using commercial MR and CT registration software.

SPGR sequence MR images are selected for post-implant registration because of superior resolution of seeds. These MR images are registered with CT using Advantage Fusion software (General Electric Medical Systems, Milwaukee, WI). Advantage Fusion software produces rigid registrations. It suggests an intensity-based initial rigid registration, within which the user selects three or more pairs of matching points in the two images. Advantage Fusion then corrects the registration to the best rigid alignment of the user-defined control points [?].

At BWH, the physicist, who plans and evaluates seed positions, produces three alignments using Advantage Fusion software. The physicist first selects one pair of matching control points to initialize the software's initial, intensity-based registration. He then refines the suggested rigid alignment by selecting three spatially extreme, unambiguous seed pairs, as rigid registration control points. The physicist reports that these seed pairs generally lie outside the prostate organ, within highly deformable, surrounding tissue. Because these seeds are prone to move in relation to the prostate organ, the planning physicist deselects these extra-prostatic seed pairs once the software has used them to update its rigid alignment. Given the updated

alignment, the planning physicist selects three new, unambiguous seed pairs from the spatial extremes of the prostate organ itself. Using this second set of three control points, Advantage Fusion software produces its final, rigid alignment, a three-point rigid registration (3PR). By personal communication, the BWH planning physicist reports that although commercial software allows alignment by more than three points, clinicians typically use only the minimum three, as a matter of convenience [10]. Clinicians currently perform post-implant dosimetry using a fused dataset, produced by overlaying high intensity voxels from the CT image over the 3PR-reformatted MR image. These fused images align seeds well enough to show whether or not they cover the prostate. However, the 3PR registration can be prone to significant error. Experiments in registration in Chapter 7 illustrate the limitations of 3PR registration.

Chapter 3

Seed Detection

The goal of post-implant MR and CT registration is to locate seeds in relation to soft tissue anatomy in order to facilitate seed localization for dosimetry. Therefore, we are concerned with improving seed-based registration. Seeds must first be detected before they can be used to control and evaluate image registration. This chapter describes the design and performance of a seed detector, whose function is to semi-automatically find seeds within MR and CT volumes.

3.1 A need for automatic seed detection

The widely used 3PR registration algorithm uses only three seeds to define a rigid registration and makes no attempt to align each of the remaining seeds. However, if each seed is to be precisely located within soft tissue anatomy, then *all* seeds should be aligned. Manual detection of seeds can be time-consuming and tedious, enough so that clinicians currently stop at finding a minimal three seed pairs out of convenience. Registration of all seeds is therefore unfeasible if it depends on the manual detection of all 40-120 seeds implanted in each patient. Therefore, automated seed detection is necessary for full, seed-based registration.

3.2 CT seed-finding

This section describes a simple, semi-automatic seed-finding algorithm. The present algorithm relies on interactive visualization of axial and sagittal views of the image volume in question. This algorithm was designed for CT seed-finding and is presented, and evaluated first in the CT context.

3.2.1 Intensity thresholds

CT image intensity is very high for voxels representing seeds. Generally, only bones appear with comparably high intensity in CT images. Voxel intensities are generally the highest in the center of the seed and are slightly lower on the periphery of the seed. In an image with 0.625mm pixel size and 1.25mm slice thickness, seeds are typically 2 pixels by 2 pixels by 1 slice at their brightest, with a slightly less bright border around the center measuring 1-2 pixels or 1 slice in thickness. In other words, including their less bright periphery, seeds are roughly spherical, with diameter about 3.2-4.8 mm. A closeup view of a typical CT seed is shown in Figure 3-1.

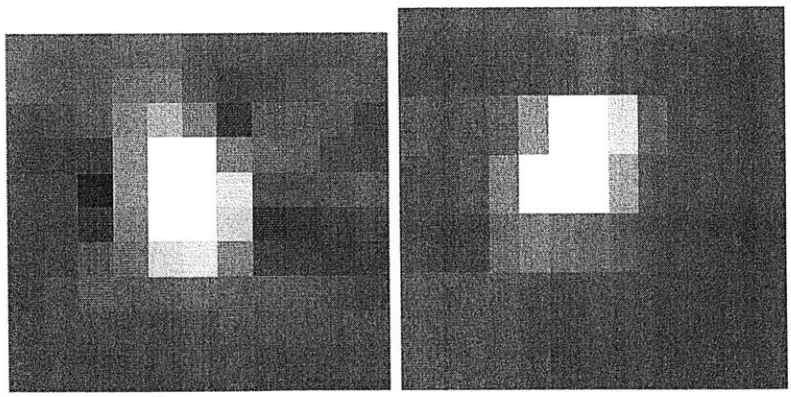

Figure 3-1: An axial (left) and sagittal (right) closeup view of a CT seed. The sagittal view has been stretched in the vertical axis according to the slice:pixel ratio to form an isometric view. Each pixel width (and half the vertical pixel height in the sagittal view) corresponds to 0.625 mm.

To begin the CT detection, the user first defines a region of interest immediately around the prostate, and excluding bone structures if possible.

On any slice containing seeds, the user interactively finds an initial threshold intensity by selecting high intensity seed-center regions. The minimum value from these selections is used as an initial estimation for CT seed intensity threshold.

From here, the user can view the thresholded regions and adjust the threshold to capture only seed-like regions.

3.2.2 Connected components as candidate seeds

A comparison of the entire CT volume against the threshold produces a binary volume, which can be thought of as an indicator function mapping voxels to an indication of whether or not the voxel has seed-like intensity.

"Connected component analysis" is performed on the threshold-binary to produce clusters of adjoining seed-colored voxels. Each cluster is a *candidate seed.*

In an effort to minimize the clustering of multiple seeds into a single candidate seed, strong connection, specifically 6-connectedness (three-dimensional connection by full faces), is chosen as criterion for two high intensity voxels to be considered connected. This eliminates single-point or single-edge connections, both of which are assumed to be too tenuous to represent cohesion within a concave, cylindrical seed.

3.2.3 Representing seeds

After connected component analysis, further shape or intensity analysis can be performed on candidate seeds to filter out bones, non-seed artifacts, and . However, at the writing of this thesis, all candidates seeds are simply assumed to be seeds. Extended techniques are discussed in Section 3.6.

Each seed is represented as a single, three-dimensional point, given by the coordinates of its centroid in the space of the image volume.

3.3 Performance of the semi-automatic CT seed-finder

The CT seed-finder finds 89 seeds for the volume shown with automatically detected seeds in Figure 3-2. A scan through axial and sagittal slices of the volume with found seeds overlaid reveals that the CT seed-finder detects one false positive, a hollow, high-intensity structure, about three seeds in diameter. The CT seed-finder counts 3 pairs of closely neighboring seeds as single seeds, and misses three high intensity regions, which are visibly duller and either significantly larger or smaller than seeds, which are generally extremely consistent in shape and size. In summary, out of 89 seeds detected, the CT seed-finder produces no certain misses and one false positive, with three clustering errors, for very successful detection.

User interaction for the CT seed-finder is minimal, as threshold adjustment is fast and straightforward, and connected component analysis is typically completed within 20 seconds on a high-end workstation.

3.4 Semi-automatic MR seed-finding

In contrast with crisp their crisp appearance in CT, seed visualization in MR is less than ideal. Figure 3-3 shows a closeup view of a typical MR seed, whose intensity pattern is darkest in the center, with a gradual increase in intensity with distance from the center of the seed. At its vaguely defined periphery, a typical MR seed measures 3.9-5.1 mm in roughly spherical diameter.

Unfortunately, air and some anatomical tissues produce low MR signal, which can be confused with seed intensities. MR seed visualization is further complicated because MR signal strengths are known to vary across a single volume. In other words, whereas seed intensities are roughly constant in CT, this is not so in MR.

For simplicity, the same approach to CT seed-finding is applied to the MR seed-finding problem. Because seeds produce no MR signal (see 1.1.4), MR intensities are very low at seed voxels. Therefore, the MR seed-finder's intensity threshold is selected

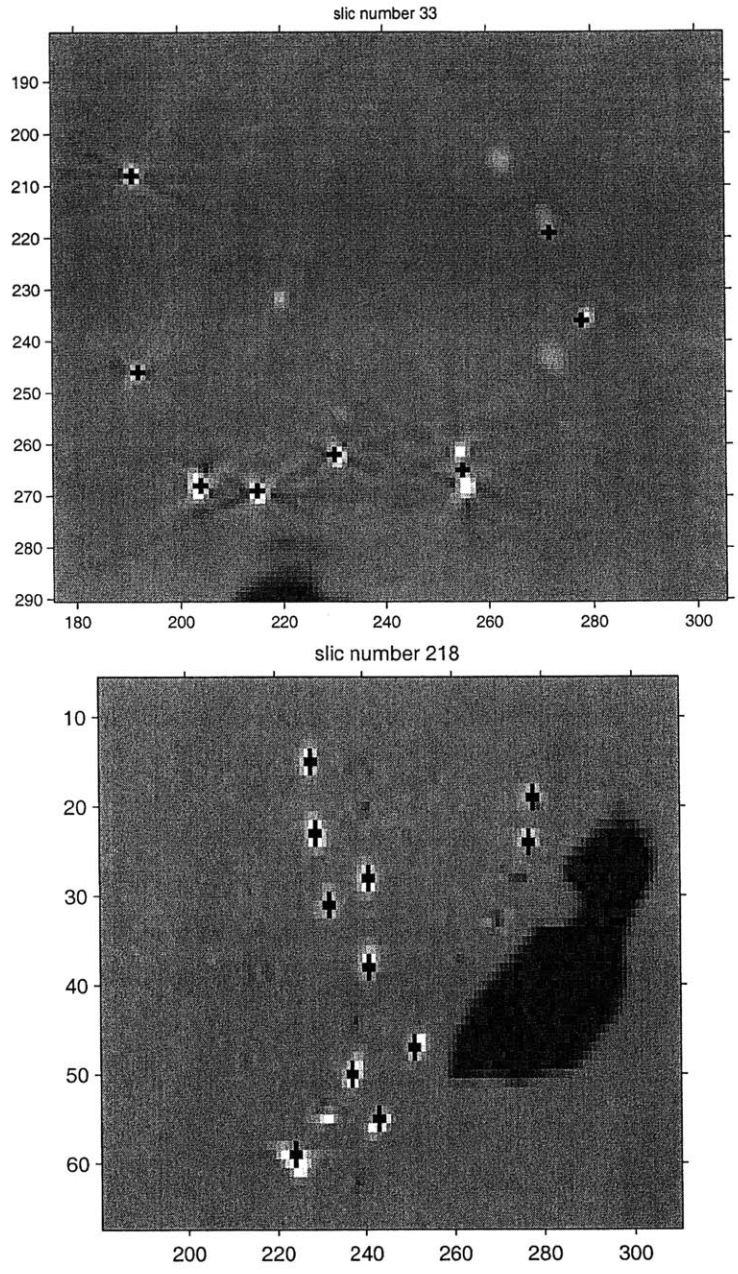

Figure 3-2: Seeds detected in CT volume, overlaid on the CT image in an axial (top) and sagittal (bottom) view. At their centers, seeds appear as bright white spots in the CT image. Black crosses are overlaid at the centroids of semi-automatically detected seeds. Some smaller or duller spots appear in the axial view. These are seeds whose centers are in the preceding or following slices, and are detected in those slices (not shown). Some dark spots appear in the roughly linear seed patterns in the sagittal view. These are artifacts which surround the bright white seed spots. The artifacts can be seen in the axial view. One pair of seeds in the axial view can be seen to have produced a single detected candidate seed. Corrections for this error are addressed in Section 3.6

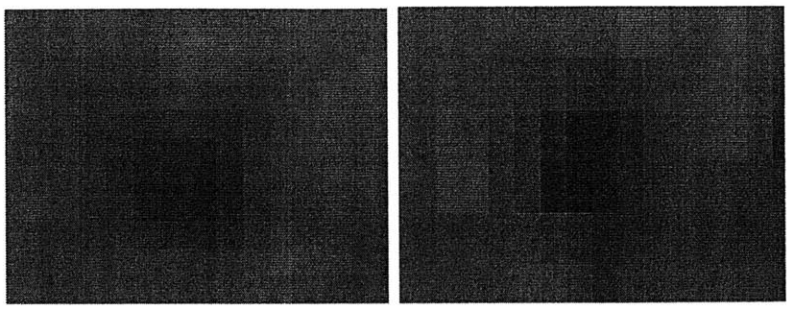

Figure 3-3: An axial (left) and sagittal (right) closeup view of an MR seed, which has been reformatted into CT space using 3PR registration. The sagittal view has been stretched in the vertical axis according to the slice:pixel ratio to form an isometric view. Each pixel width (and half the vertical pixel height in the sagittal view) corresponds to 0.625 mm.

to accept only very low intensities. As before, the intensity threshold is adjusted interactively using visual feedback of the voxels captured by the threshold. As with the CT seed-finder, strong connected component analysis is used to form candidate seeds, which are at this stage declared to be seeds. This intensity thresholding and connected component approach is very similar to the CT seed-finder reported by Brinkmann and Kline, which was validated in comparison with manual seed-finding in two-dimensional radiographs [?]

3.5 Performance of the MR seed-finder

A constant number of seeds appears in the paired MR and CT volumes used to test the seed-finders. The CT seed-finder found approximately 90 visually validated seeds. Given the user-adjusted threshold, the MR seed-finder found 189 candidate seeds.

A scan through the volume reveals that roughly one quarter of these candidate seeds are false positives, corresponding to regions in bones and the rectum, and other artifacts which are not shaped like seeds. Another error results from the significant number of single seeds which produce several "candidate seeds" because their lowest intensity regions are not connected. Despite the high number of candidate seeds produced, at least 10 misses also occurred. The results of the MR seed-finder can be

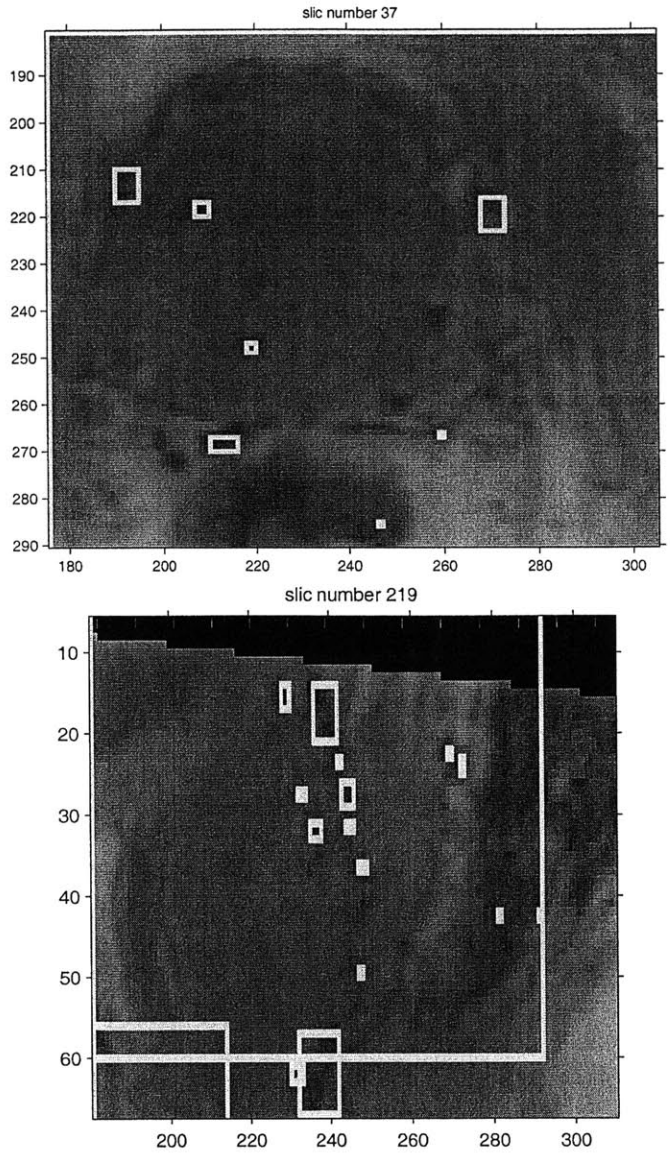

Figure 3-4: Candidate seeds' bounding boxes generated by the MR seed-finder are overlaid on the MR image. In the axial view (top), the seed-finder successfully finds 6 seeds in the large, central, circular prostate region. The lower-most seed is most likely a false positive corresponding to air in the rectum. The seed-finder glaringly misses a few seeds, which should be completely outlined by their bounding boxes. In the sagittal view, a vertical swath of about ten seeds, and falsely connected seed pairs, is successfully detected. It is not clear whether the candidates detected in the upper right are false positives. The candidates in the dark region are probably false positives, lying in the rectum. The very large bounding boxes are caused by poorly defined region of interest, which includes zero signal areas from 3PR reformatting.

visualized in 3-4.

A visualization of found candidate seeds' bounding boxes, overlaid on the MR volume indicates the reason for the high number.

3.6 Future directions for automatic seed-finding

The simple seed-finder presented identifies seeds as connected clusters of voxels above or below fixed intensity thresholds. To improve both MR and CT seed-finding, the following techniques have been considered but not yet implemented and tested. This section describes directions for further research in automatic or semi-automatic seed-finding.

3.6.1 Separating adjacent seeds and excluding large shapes

First of all, the shape of seeds can be used to improve seed detection. Seeds are consistently small, and round-shaped in both MR and CT images.

To further minimize multiple-seed clustering, connected components with large bounding boxes can be eroded using a simple, 2-dimensional square-shaped structuring element in an effort to eliminate connections between adjacent seeds. Erosion is a morphological (shape) operation which reduces the extent of shapes to fit the shape of the designated structuring element [?].

After erosion, tightly neighboring seeds should have been separated. Therefore, all remaining large components, which exceed 5 x 5 pixels x 3 slices, can be excluded on the assumption that they represent bone, or other non-seed objects.

Following erosion and large-size-filtering, all remaining connected components are assumed to be seeds. This assumption can be confirmed visually by overlaying the connected components' (that is, the seeds') bounding boxes over the CT image volume from which they have been found.

3.6.2 Shading gradients

If the MR seed intensity is variable, MR seeds nevertheless feature an invariant pattern of low-intensity centers with higher intensity peripheries. Rather than using a simple intensity threshold to find seeds, intensity gradients can be searched for ball shapes with low intensities in their centers.

3.7 Conclusion on semi-automatic seed-finding

Semi-automatic seed-finding is an important tool for seed-based registration. This work verifies that CT seed-finding is straightforward and requires only minimal user interaction to set the threshold.

On the other hand, simple intensity thresholding has been shown to be problematic for MR seed-finding. If a semi-automatic MR seed-finder can be achieved, then given automatically found seeds, full seed-based registration can be automated. Therefore, automatic seed-finding should receive further research attention.

Chapter 4

A rigid registration algorithm for matched seeds

The development of a rigid registration algorithm allows us to eventually answer two questions. The first asks, can three seeds produce a good alignment of all seeds in the prostate, or should more seeds be used? The second asks, is it appropriate to perform registration using a rigid model of the prostate? These questions will be answered in Chapter 7. The present chapter describes a rigid registration algorithm based on corresponding seeds.

4.1 The rigid model of motion

As described in Chapter 2, an image registration algorithm consists of a coordinate transformation, an objective or measure of similarity, and an optimization scheme. The solution to a registration problem is an assignment of values to the transformation's parameters.

A rigid registration uses a rigid transformation, which describes the physical motion of rigid bodies. The characteristic feature of rigid bodies is that the distances and angles between any two parts of the object remain fixed through any motion of the object. Likewise, on an image, a rigid transformation preserves distances and angles between all pairs of coordinates in the image.

Rigid motion encompasses translation and rotation. In three dimensions, a rigid transformation is defined by three degrees of freedom for translation, and three for rotation, for a total of six parameters.

4.1.1 Representing the rigid transformation

First, we recall our representation, introduced in 2.2, of an image as an intensity function over a coordinate spaces, and we recall our representation of a coordinate point as a column vector $p = \begin{bmatrix} p_x \\ p_y \\ p_z \end{bmatrix}$. Now, if we think of the rigid transformation T_{Rigid} as a function which maps the 3D coordinate point p to a new 3D coordinate p', then we can express the rigid transformation as the composition of a rotation $T_{Rotation}$ and a translation $T_{Translation}$ on p, as follows

$$p' = T_{Rigid}(p),$$

where

$$T_{Rigid}(p) = T_{Translation} \circ T_{Rotation}(p). \tag{4.1}$$

Translation

Translation $T_{Translation}(p)$ of a point p is usually represented as the addition of the point with a column vector $D = \begin{bmatrix} d_1 \\ d_2 \\ d_3 \end{bmatrix}$ which represents the distance and direction traveled. Translation of a 3D point p to a new 3D coordinate p' can be expressed as

$$p' = T_{Translation}(p),$$

where

$$T_{Translation}(p) = p + D \tag{4.2}$$

$$= \begin{bmatrix} p_x \\ p_y \\ p_z \end{bmatrix} + \begin{bmatrix} d_1 \\ d_2 \\ d_3 \end{bmatrix}$$

Rotation

A rotation is a linear transformation on a coordinate vector. A proper rotation matrix only rotates coordinates, and does not scale, translate or distort relative distances or angles. Rotation transformations are constrained to be orthonormal with determininant 1 [27].

As Horn suggests, a multitude of representations for rotation exist. Rodrigues' formula for rotation, provides a straightforward representation of rotation, written in terms of a three-dimensional unit vector ω representing axis of rotation, and a scalar angle θ of rotation about the axis [17]. Rodrigues' formula is written as:

$$T_{Rotation}(p) = (cos\theta)p + (sin\theta)\omega \times p + (1 - cos\theta)(\omega \cdot p)\omega \tag{4.3}$$

As shown in [26], Rodrigues' formula can be expanded into matrix form:

$$T_{Rotation}(p) = R \begin{bmatrix} p_x \\ p_y \\ p_z \end{bmatrix}, \tag{4.4}$$

where

$$R = \begin{bmatrix} cos\theta + w_x^2(1 - cos\theta) & w_x w_y(1 - cos\theta) - w_z sin\theta & w_y sin\theta + w_x w_z(1 - cos\theta) \\ w_z sin\theta + w_x w_y(1 - cos\theta) & cos\theta + w_y^2(1 - cos\theta) & -w_x sin\theta + w_y w_z(1 - cos\theta) \\ -w_y sin\theta + w_x w_z(1 - cos\theta) & w_z sin\theta + w_y w_z(1 - cos\theta) & cos\theta + w_z^2(1 - cos\theta) \end{bmatrix} \tag{4.5}$$

We now compose rotation and translation into a rigid transformation. For a point p defined as a three-dimensional column vector, equation 4.1 can be rewritten as

$$T_{Rigid}(p) = Rp + D.$$

Using homogeneous coordinates, we can represent the rigid transformation of a 3D point p to a new 3D coordinate p' as as a matrix multiplication

$$T_{Rigid}(p) = T \begin{bmatrix} p \\ 1 \end{bmatrix}$$

with

$$T \begin{bmatrix} p \\ 1 \end{bmatrix} = Rp + D,$$

where

$$T = \begin{bmatrix} r_{11} & r_{12} & r_{13} & d_1 \\ r_{21} & r_{22} & r_{23} & d_2 \\ r_{31} & r_{32} & r_{33} & d_3 \end{bmatrix}.$$

A note on the non-linearity of rigid transformation

Because $T_{Rigid}(p)$ can be represented as a matrix multipliciation with a column vector addition, rigid transformation may at first appear to be linear in its parameters. Unfortunately, the rotation matrix R's parameters are never simply its elements. Instead, as we have discussed, R is constrained to be orthogonal. Rodrigues' formula for rotation is obviously non-linear in its parameters θ and ω. While other parametrizations of rotation are possible, no known parametrization of rotation is linear in its parameters. This observation will be relevant to the optimization of the transformation, discussed in 4.3.

4.2 Seed-based similarity criterion

Since our goal is to align seeds, we use seed positions to measure the similarity between aligned images. A popular similarity objective for a variety of approximation and estimation problems is the Sum-Squared Error (SSE) criterion. We will use sum-

squared distance between corresponding seeds as our objective function. In general, for a transformation $T(p)$ with parameters collected in a column vector β, we can write T as a function of both p and β, or $T(\beta, p)$. We then define our objective function as

$$J(\beta) = \sum_i ||u_i - T(\beta, v_i)||^2, \tag{4.6}$$

where u_i and v_i are the coordinates of a pair of corresponding seeds in the two images. According to the objective function $J(\beta)$, the best rigid transformation is the one whose parameters β minimize the sum-squared distance between seeds. This can be expressed as

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}}\, J(\beta),$$

or equivalently

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \sum_i ||u_i - T_{Rigid}(\beta, v_i)||^2. \tag{4.7}$$

$\hat{\beta}$ is the rigid transformation which minimizes the sum-squared distance between seeds.

4.2.1 Representating transformation parameters in the context of optimization

For a rigid transformation $T_{Rotation}$, β is a column vector representing three degrees of rotational and translational freedom. For convenience in calculating the rigid transformation, we previously characterized the rotation matrix R using Rodrigues' formula, in terms of a three-dimensional axis of rotation ω and a scalar angle of rotation θ. Rodrigues' formula represents three rotational degrees of freedom using four parameters (three in ω and one in θ), though ω is constrained to be a unit vector.

The constraint makes it difficult to search among values for the parameters for the solution which optimize our objective function. Another view of the inconvenience

of optimizing Rodrigues' representation is to consider the shape of the space we are trying to search. We are essentially trying to find the optimal point on a unit sphere, which corresponds to the axis of rotation, and some angle, which is difficult to characterize. We are generally used to maximizing and minimizing functions in R or R^2 or R^3, to optmize a function on a unit sphere with some angle is more complicated.

To make optimization or search for parameter values easier, we can select new rotational parameters which are convenient to search, and then relate them to Rodrigues' parameters. We define our new rotation parameters as the elements of a three-dimensional column vector Q which is related to Rodrigues' ω and θ as follows.

$$Q = \omega\theta \tag{4.8}$$

or

$$\begin{bmatrix} q_1 \\ q_2 \\ q_3 \end{bmatrix} = \begin{bmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{bmatrix} \theta$$

While ω and θ have the physical meaning of a unit axis and associated angle of rotation, Q does not have an actual, physical meaning. However, Q can be geometrically understood as a vector in the direction of ω with length θ, or alternatively as a point in R^3. Unlike ω, Q is not restricted to lie on the unit sphere. Also, whereas ω previously lacked a clear representation in a searchable space, its value now describes the length of Q. Q spans R^3, a space over which optimization is very well-developed.

We now have two parametrizations for rotation. The first, Rodrigues', is convenient for evaluating the coordinate transformation T_{Rigid}. The second, Q, is convenient for optimizing search. Equation 4.8 describes how to convert Rodrigues' representation into Q. To convert Q back to Rodrigues' representation, we can do

$$\omega = \frac{Q}{|Q|}$$

and

$$\theta = |Q|$$

We now have an easily searchable trio Q of rotation parameters spanning R^3. Our translation parameters D also span R^3. Now satisfied that we have six parameters which we can optimize using existing techniques, we collect the parameters into a single, six-dimensional column vector, $\beta = [QD]'$.

4.3 Nonlinear optimization of rigid sum-squared error

We must now search a six-dimensional Euclidean space spanned by the parameters collected in β for the values which optimize our objective function. In general, optimizing an objective criterion is equivalent to maximizing or minimizing the objective function.

Recall our solution $\hat{\beta}$, as characterized by 4.7. To find $\hat{\beta}$, we might consider trying to solve

$$\nabla_\beta J(\hat{\beta}) = 0. \tag{4.9}$$

Because $J(\beta)$ is quadratic in T_{Rigid}, 4.9 is linear in T_{Rigid}. At first 4.9 may seem at to be a linear equation. Unfortunately, as we saw in 4.1.1, $T_{Rigid}(\beta)$ is not linear in β, and so 4.9 may be difficult to solve.

To optimize β we can use any of several standard techniques of nonlinear optimization. For the experiments discussed in Chapter 7, the Downhill Simplex algorithm was selected as the nonlinear opimization algorithm.

The Downhill Simplex optimization method, developed by Nelder and Mead, searches the parameter space by evaluating the objective function at the vertices of an object which moves about the space according to well-defined rules which direct the shape toward optimal values of the objective function. In other words, the objective function is evaluated at points in the search space which eventually converge upon the optimal points in the search space. The points of evaluation define the

vertices of the object, called a *simplex*, which specifically has one more vertex than the dimensionality of the space. The simplex traverses the search space by reflecting over one of its edges, by contracting, or by expanding [24].

As a local method for optimization, Downhill Simplex requires the user to provide an initial, rough pose estimate.

4.4 Rigid registration summary

This chapter has presented a rigid registration algorithm which optimizes a rigid transformation to minimize the sum-squared distance between corresponding seeds in the post-implant MR and CT images. Optimization is conducted using the Nelder-Mead Downhill Simplex method.

The clinically used commercial solutions for rigid registration likely use an algorithm very similar to the one presented in this chapter. The software used at BWH requires initialization by intensity-based alignment, and can perform registration over three or multiple seeds. However, in practice, only three seeds are used.

This rigid registration scheme requires knowledge of correspondences between seeds. It has been provided in this thesis primarily to allow for comparison between rigid and affine models of prostate movement, and for comparison between three-point and many-point registrations. These comparisons are made in Chapter 7.

Chapter 5

An Affine Registration Algorithm For Matched Seeds

Van Herk reported that the prostate changes shape between image scans, even when taken with the patient in the same anatomical position [18]. Although patients lie in the supine position for both post-implant CT and MR image acquisitions, these scans are performed at different imaging suites, sometimes hours apart. Thus, MR and CT image acquisition allows for variations in patient positioning, and in bladder and rectal filling, both of which are factors identified by Van Herk as affecting prostate shape and position.

Despite van Herk's evidence of non-rigid prostate motion, rigid registration remains the clinical standard [10].

In this chapter, we explore a model fe assume that the post-implant MR and CT imaging procedures allow the prostate to move non-rigidly betwen scans.

The prostate's shape change motivates the use of a non-rigid transformation for MR and CT registration. The simplest non-rigid model is the affine transformation, which generalizes rigid movement to include shears and non-isotropic scaling. Affine transformations preserve parallel lines.

5.1 Non-rigid model of motion using affine transformation

Affine motion generalizes rigid motion to include shearing (slanting) and non-isotropic scaling (that is, stretching along one but not all dimensions). Affine transformations do not preserve relative distances and angles but they do preserve parallel lines.

The affine transformation of a three-dimensional coordinate point p to a new 3D coordinate p' is given as $p' = T_{Affine}(p)$.

As in the rigid case, the affine transformation can be thought of as the composition of two transformations, a generalized rotation transformation and a translation transformation. $T_{Affine}(p) = T_{Translation} \circ T_{GeneralizedRotation}(p)$.

Translation can be represented just as it was in the rigid case, as $T_{Translation}(p) = p + D$.

In contrast to the complicated rigid rotation matrix, the generalized rotation transformation is simply an unconstrained 3×3 matrix. That is,

$$T_{GeneralizedRotation}(p) = Mp$$

$$= \begin{bmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{bmatrix} \begin{bmatrix} p_x \\ p_y \\ p_z \end{bmatrix}$$

Composing rotation and translation into the affine transformation, we have

$$T_{Affine}(p) = Mp + D.$$

Using homogeneous coordinates, the affine transformation $T_{Affine}(p)$, which maps p to a new 3D coordinate p', can be written as a matrix multiplication with the affine matrix A, as

$$T_{Affine}(p) = A \begin{bmatrix} p \\ 1 \end{bmatrix}$$

with

$$A \begin{bmatrix} p \\ 1 \end{bmatrix} = Mp + D,$$

where

$$A = \begin{bmatrix} m_{11} & m_{12} & m_{13} & d_1 \\ m_{21} & m_{22} & m_{23} & d_2 \\ m_{31} & m_{32} & m_{33} & d_3 \end{bmatrix}$$

The affine transformation A is essentially an unconstrained 3×4 matrix. Whereas the rigid transformation was nonlinear in its parameters because of constraings on the rotation matrix R, the affine's generalized rotation matrix M is unconstrained. Therefore, the affine transformation's parameters are its elements themselves, and it is linear in those parameters.

5.2 Seed-based objective criterion

Just as we did for rigid registration, we still desire to align seeds, and so we use the same sum-squared distance criterion as defined in equation 4.6 and reiterated below.

$$J(\beta) = \sum_i ||u_i - T(\beta, v_i)||^2$$

As before, the optimal affine parametrization $\hat{\beta}$ is given as

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \sum_i ||u_i - T_{Affine}(\beta, v_i)||^2.$$

5.2.1 Linear Least Squares optimization

Linear least squares is the problem of minimizing sum-squared error of a function which is linear in its parameters.

Since the affine transformation $T_{Affine}(\beta, v_i)$ is linear in its parameters β, and since the objective function is quadratic in the affine transformation, this problem lends itself to a Linear Least Squares (LLS) optimization. The solution to the problem is to set the gradient of the objective function equal to zero and then to solve the resulting linear system.

We have not yet said how the parameters of T_{Affine} are represented in β. Earlier, for comparison with the rigid transformation, we represented $T_{Affine}(v_i)$ as Av_i, where the parameters of T_{Affine} are the elements of A. Now to facilitate the solution of these parameters, we rearrange them into the column vector

$$\beta = [a_{11}a_{12}a_{13}a_{14}a_{21}a_{22}a_{23}a_{24}a_{31}a_{32}a_{33}a_{34}]'.$$

We must now find an equivalent formulation for Av_i in terms of β. We rearrange the matrix multiplication Av_i as $V_i\beta$ by reformatting v_i as the matrix V_i, where

$$V_i = \begin{bmatrix} v_{i,x} & v_i, y & v_{i,z} & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & v_{i,x} & v_i, y & v_{i,z} & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & v_{i,x} & v_i, y & v_{i,z} & 1 \end{bmatrix}$$

We can now rewrite

$$T_{Affine}(v_i) = Av_i$$

as

$$T_{Affine}(\beta, V_i) = V_i\beta$$

Our sum-squared distance objective function

$$\hat{\beta} = \underset{\beta}{\operatorname{argmin}} \sum_i ||u_i - Tv_i^2||$$

We optimize it by equating the gradient to 0.

First, we rewrite the objective function in equation 4.6 to reflect our new repre-

sentation of the affine transformation

$$J(\beta) = \sum_i ||u_i - V_i\beta||^2$$

and expand to get

$$J(\beta) = \sum_i [u_i - V_i\beta]^T [u_i - V_i\beta]$$

which can be rearranged as

$$J(\beta) = \sum_i [u_i^T u_i - 2\beta^T V_i^T u_i + \beta^T V_i^T V_i \beta] \tag{5.1}$$

Now, to minimize, we take the gradient and equate it to zero.

$$0 = \nabla_\beta J(\hat{\beta}) \tag{5.2}$$

Substituting equation 5.1 we have

$$0 = \nabla_\beta \sum_i [u_i^T u_i - 2\beta^T V_i^T u_i + \beta^T V_i^T V_i \beta]|_{\hat{\beta}} \tag{5.3}$$

We take the gradient and evaluate at $\hat{\beta}$ and get

$$0 = \sum_i [-2V_i^T u_i + 2V_i^T V_i \hat{\beta}] \tag{5.4}$$

which we can rearrange as

$$(\sum_i V_i^T V_i)\hat{\beta} = \sum_i V_i^T u_i \tag{5.5}$$

If we define the following two constructs

$$X = \sum_i V_i^T u_i$$

and

$$Y = \sum_i V_i^T V_i$$

then we can simplify equation 5.5 as

$$Y\beta = X$$

and solve

$$\beta = Y^{-1}X. \tag{5.6}$$

Since β represents twelve parameters, at least four three-dimensional seed coordinate pairs must be defined to solve equation 5.6.

5.3 Affine registration summary

Given known correspondences among seeds, we now have an algorithm for solving for an affine transformation which minimizes the sum-squared distance between corresponding pairs of seeds. The algorithm uses Linear Least Squares to estimate the parameters of the affine transformation.

This algorithm allows us to solve for the optimal non-rigid, affine transformation which can bring an unconstrained number of matched seeds into alignment.

Chapter 6

Iterated Affine Registration for Unmatched Seeds

In this chapter, the affine registration algorithm (see Chapter 5) is generalized to allow two collections of seeds to be aligned, even if the correspondences among these seeds is not known. The rigid and affine registrations presented earlier require that correspondences be established between pairs of seeds in the two images. In our application, requiring known correspondences reduces to requiring a user to establish matches by hand.

Because one aim of this thesis is automatic seed-based registration, we require a registration algorithm which can take automatically detected seeds as input. As we saw in Chapter 3, seeds are detected in each image independently. Therefore, automatic seed-detectors cannot establish correspondence. This chapter describes an affine registration algorithm for unmatched seeds, such as those which might be produced by automatic seed-finders.

6.1 Dual pose and correspondence problems for unmatched features

In feature-based registration problems, there is frequently a duality among knowing pose, or knowing feature correspondences. If one is known, it is usually straightforward to estimate the other. A fully automatic solution to the seed-based registration problem could use automatic seed finders to detect the features in the input images, but their correspondences would at that stage be unknown.

If we do not know which features in one image correspond with which features in the next, the problem of using these features to measure the similarity of the two images becomes much more challenging. The problem of estimating a transformation becomes compounded with the problem of estimating correspondences between features.

We solve for an affine transformation by first generalizing our pose estimation to a *weigted least squares* approach, that entertains all possible correspondences, and weights them according to an estimate of their validity, as

$$J(\beta) = \sum_{ij} w_{ij} ||(u_i - T(\beta, V_j))||^2 \tag{6.1}$$

where w_{ij} are ideally one if features u_i and v_j correspond, and zero otherwise.

Wells [28, 29] described an algorithm, Posterior Marginal Pose Estimation (PMPE), that uses the Expectation Maximization (EM) alogorithm [15] to simultaneously estimate pose and correspondences. This approach iterates, in alternating fashion, pose estimation via equation 6.1, and estimation of the correspondence weights, using

$$w_{ij} = \frac{G_{\psi_{ij}}(u_i - T(V_j, \beta))}{\frac{B_i}{1-B_i}\frac{m}{Z} + \sum_j G_{\psi ij}(u_i - T(V_j, \beta))} \tag{6.2}$$

This equation estimates the correspondence weights w_{ij} by their posterior probabilities based on the current pose estimate β, and on a Gaussian model of feature detection error

$$G_{\psi_{ij}}(x) \equiv (2\pi)^{-\frac{3}{2}} |\Psi_{ij}|^{-\frac{1}{2}} exp(-\frac{1}{2} x^T \Psi_{ij}^{-1} x)$$

where Ψ_{ij} describes the covariance of the feature detection error along the three dimensions of the volume. We assume that for our application, feature detection errors in each direction are independent of other directions. Thus, we can represent feature error with a covariance that is a scaled identity matrix, where the scaling factor corresponds to variability in the distance between the detected location and the true location of each seed. Equation 6.2 also depends on the fixed probability B_i that some seed from the first image has no counterpart in the second image, on m, the number of features in the second image, and on Z, the volume of the second image's space.

This algorithm is similar to the popular Iterated Closest Point (ICP) [2] algorithm, but it uses a softer estimate of correspondences, and explicitly allows for missing features.

As an instance of the EM algorithm, this iteration is guaranteed to converge to a local maximum of the posterior probability of pose. Being a local method, good initial values are needed. The algorithm may be started with either an initial pose estimate, or an estimate of the correspondence weights.

In our experiments, we have initialized the algorithm with the pose that was determined by three-point rigid registration, and run the algorithm to convergence, typically in 10-30 iterations, within a few seconds on a high end workstation, for collections of 30 unmatched seeds.

6.2 Conclusion on iterative affine registration

This chapter has described an affine registration algorithm which does not require prior knowledge of correspondences between seeds. This algorithm is a local method which requires initialization.

The iterative affine registration method can be used to align automatically detected seeds, for which correspondences are not known. Thus, if we can design automatic seed detectors, use their output as input to this umatched-feature-based registration algorithm, which would result in an almost fully automated seed-based

registration algorithm.

Chapter 7

Experiments and Results

In this chapter, we demonstrate and compare the quality of alignments produced by the rigid, affine, and iterated affine registration algorithms previously developed. We evaluated these algorithms on five patients' post-implant MR and CT images. This chapter poses three questions which are the focus of this research and answers them in based on three comparisons of the registration algorithms.

7.1 Questions addressed by experiments

Registration experiments were conducted to answer three questions:

1. What is the effect of aligning many seeds, instead of three?

2. What is the effect of modeling the prostate's motion with a non-rigid, affine model, instead of a rigid model?

3. If we do not know seed correspondences, how well can we align seeds?

The experiments described in this chapter were conducted on post-implant MR (SPGR) and CT image pairs, which were taken within hours of each other, six weeks following brachytherapy implantation, as described in Section 2.1. Image pairs from five patients were used to conduct the experiments. Because the semi-automatic MR seedfinder was highly prone to error, automatically detected seed sets were not used

to evaluate the registration algorithms. Instead, test data was produced by manually selecting matching seed pairs in paired volumes.

To facilitate manual seed-finding, the MR and CT volumes were aligned using an initial 3PR rigid alignment, found using commercial software by the method described in Section 2.5. Therefore, the MR images used for testing are actually reformatted images, aligned using a commercial 3PR implementation. Within 3PR-aligned MR and CT volumes, between 27 and 32 pairs of corresponding seeds were manually selected from the five patients' MR and CT volumes, using interactive display of axial and sagittal views of the 3PR-aligned volumes. Manual selection took approximately 20 minutes per dataset. The number thirty was chosen to maximize the number of easily identifiable matches and to maxmize convenience. Seeds were required to have unambiguous correspondences in order to be selected . The coordinates of the seeds' centers were used to represent the seeds.

Three registrations were conducted on each set of 30 seed pairs:

- Multi-Point Rigid (MPR) registration
- Multi-Point Affine (MPA) registration
- Iterated Multi-Point Affine (IMPA) registration on unmatched seeds

Both the MPR and MPA techniques require that the two input seed sets U and V have correspondences established between their elements. One simple way to represent the correspondences is by putting the seeds in ordered lists with corresponding seeds having equal indices in their respective lists.

Although a 3PR registration was used to initialize the test data, the MPA algorithm is unaffected by initial registration. However, MPR registration requires a good initial alignment before it can use a local nonlinear optimization technique to refine the alignment. The IMPA registration algorithm can solve an alignment on unmatched seeds, but it requires an initial alignment, which we provide using the results of 3PR alignment.

7.2 Measuring performance

Registration performance was quantified for each of the three multi-point registrations, and also for the initial, three-point rigid alignment. Performance was measured using error statistics on distances between aligned pairs of seeds. Two error statistics, Root-Mean Squared Distance (RMSD) and Maximum Distance (MD) were calculated. RMSD error characterizes the success of the alignment collectively over all the seeds and MD error characterizes the worst alignment of all the seeds in a dataset. MD error is characterized because clinicians are generally concerned about characterizing the worst cases which might occur.

We represent a set of corresponding seed pairs as $U = u_i$ and $V = v_i$, where u_i corresponds to v_i. Given a set of seed pairs U and V represented in the space of an initial 3PR alignment, and given the alignment being evaluated T,

$$RMSD(U,V) = \sum_i ||u_i - T(v_i)||^2$$

and

$$MD(U,V) = \max_i(||u_i - T(v_i)||)$$

To measure the performance of the initial 3PR alignment, we calculate the RMSD and MD error as before, letting T be the identity transformation. The reason for this is that seeds have been found within a 3PR-aligned dataset. Therefore, we do not need to transform the seeds to represent the 3PR transformation.

7.3 Registration performance

The tables in Figure 7.3 and Figure 7.3 list the root-mean-squared and maximum distance errors, respectively, for five patients' images for the four registrations performed. The tables show that the aggregate quality of alignment, as measured by RMSD error, is better for MPR alignment than for 3PR and better for MPA than for MPR alignment. We also see that both IMPA for unmatched features, and MPA

for matched features, produce very similar results, indicating that IMPA converges on the MPA solution, despite ignorance of correspondences.

The tables also show that worst-case, or maximum-distance errors also decrease as we use more points or switch to an affine model. Again, as with RMSD error, IMPA worst-case error is generally similar to affine worst-case error, indicating successful convergence of the IMPA algorithm to the correct affine solution.

Alignment	A	B	C	D	E
3PR	5.6	4.2	2.8	3.1	4.9
MPR	2.5	3.3	2.4	2.4	3.4
MPA	1.9	2.2	1.9	2.2	2.5
IMPA	1.9	2.3	1.9	2.2	2.5

Figure 7-1: This figure shows RMS distance (mm) between registered pairs of seeds, using 3PR, MPR, MPA, and IMPA registration algorithms. The results are shown for five patients, A, B, C, D, and E. Registration results for each patient are listed columnwise.

Alignment	A	B	C	D	E
3PR	9.2	7.2	6.5	5.6	8.7
MPR	5.9	5.7	6.0	3.7	5.7
MPA	3.3	4.0	5.1	3.2	4.4
IMPA	3.7	4.9	5.0	3.3	4.4

Figure 7-2: This figure shows maximum distance (mm) between registered pairs of seeds, using 3PR, MPR, MPA, and IMPA registration algorithms.

We have seen that for each dataset, errors improve from 3PR to MPR and from MPR to MPA registration. We can generally characterize these trends using significance testing.

7.4 Significance tests of reduction in registration error

Statistical significance testing was used to answer the questions asked at the beginning of the chapter about (1) the effect of aligning many instead of few seeds, (2) the

effect of using a non-rigid model, and (3) the feasibility of registration on unmatched correspondences.

7.4.1 Constructing t-tests

The statistical significance of the registration error reductions was demonstrated using a t-test for significance.

RMSD errors for the each registration were collected into populations, each of whose samples are the RMSD errors for that registration on the five patients' datasets. Paired, right-sided t-tests were calculated over three pairs of error populations.

A t-test on a population X assumes that X is normally distributed, and givesan indication of the confidence with which we can say that X is centered around 0. The t-test works by characterizing X using a statistic t, which is derived using Student's (actually William Gossett's) t-distribution [25]. The confidence with which we can model X as having 0-mean is associated with the probability that t will fall within certain, predefined boundaries.

In most applications, X's non-0 mean value is considered to be significant if $p < 0.5$. This means that there is only a five percent chance that X apparently has non-0 mean, when in actuality it is 0-mean. More simply, we can be 95% sure that X has a significant non-0 mean.

For two populations X and Y, with corresponding values, the *paired t-test* is essentially just a t-test on their difference $X - Y$. If the paired t-test indicates that we can confidently reject a 0-mean for $X - Y$, this means that X's and Y's values are significantly different.

In our application, for any two alignments' error populations X and Y , each error value $x_i \in X$ corresponds to a value $y_i \in Y$, because both errorscome from the same patient's images. Therefore, we use a paired t-test to test the significance of X's and Y's difference.

If we are interested in testing whether X and Y are significantly different, then we use a *two-sided* t-test. If, instead, we are interested in knowing specifically whether X is significantly greather than Y, then we can use a *right-sided* t-test on $X - Y$.

3PR v. MPR	right-sided t-test for 3PR error - MPR error
MPR v. MPA	right-sided t-test for MPR error - MPA error
MPR v. IMPA	two-sided t-test

Figure 7-3: Two right-sided and one two-sided t-tests are used to answer three central questions about seed number, transformation model, and unmatched seed-based registration.

3PR v. MPR	.03	Yes
MPR v. MPA	.01	Yes
MPA v. IMPA	.37	No

Figure 7-4: The top two lines in this table show the results of paired t-tests which check if RMSD alignment error is signifcantly reduced by using MPR instead of 3PR registration, and when using MPA instead of MPR. For the first and second t-tests, the p value is less than 0.05, or 5%, indicating that MPR alignment error is significantly smaller than 3PR error, and that MPA error is significantlysmaller than MPR error. The third t-test checks the significance of the difference between the MPA and IMPA solutions. The resulting p-value is .37, indicating far exceeding the acceptible significance threshold of 0.05. Therefore, we conclude that MPA and IMPA are not significantly different.

We will use the t-test to answer our three experimental questions. We set up the t-test as shown in Figure 7-3.

7.4.2 Results of significance tests

Figure 7-4 shows the results of significance tests on our three error comparisons. Wefind that the probability that we can be 97% confident that MPR registration produces smaller RMSD alignment error than 3PR registration. We also see that we can be 99% certain that MPA produces smaller registration error than MPR. Finally, We see that MPA and IMPA are not significantly different, as the p-value for their difference is much greater than 5%. Thus, the t-tests indicate that (1) using many seeds produces better alignments than using just a few, (2) that the non-rigid affine model produces better alignments than traditional rigid models, and (3) that an iterated affine registration technique, which can align unmatched seeds, converges on the same solution as the affine solution produced with known seeds.

7.5 Visual assessment of registration quality

If we wish to further convince ourselves of the improvements afforded by registration with many seeds and affine transformations, we can view axial and sagittal slices of fused, that is, aligned and overlaid, volumes, as shown in 7.5 and 7.5.

The axial views show subtle improvements can be seen beteen MPR and 3PR, and MPA and MPR alignments. We see much greater improvement in alignment when comparing the MPR with 3PR and MPA with MPR in the sagittla views.

We can also note that MPA and IMPA alignments are almost identical, confirming that iterated affine solution with unmatched features can converge on the affine solution given known features.

7.6 Evaluation of results

We have seen from the root-mean-squared distance and maximum-distance errors, from significance testing on error reduction, and from visualization of seed alignment, that (1) many seeds produce better alignments than three seeds, (2) affine registrations are better than rigid, and (3) the iterated affine technique on unmatched correspondences successfully converges on the optimal affine registration.

7.6.1 A scheme for validation

To measure registration quality, we have calculated error over the distances between multiple aligned seed pairs. Our conclusions about registration improvements, based on these error formulations, has been shown to correspond with visual confirmations of improvements.

The technique of evaluating post-implant MR and CT registrations using multiple seeds has thus been demonstrated quantitatively and qualitatively to be useful.

The previous observation that 3PR produces a good registration, can be viewed as a validation of the 3PR result. While it may encounter some areas of problematic alignment, such as depicted in Figure 7.5, 3PR has been found to generally align most

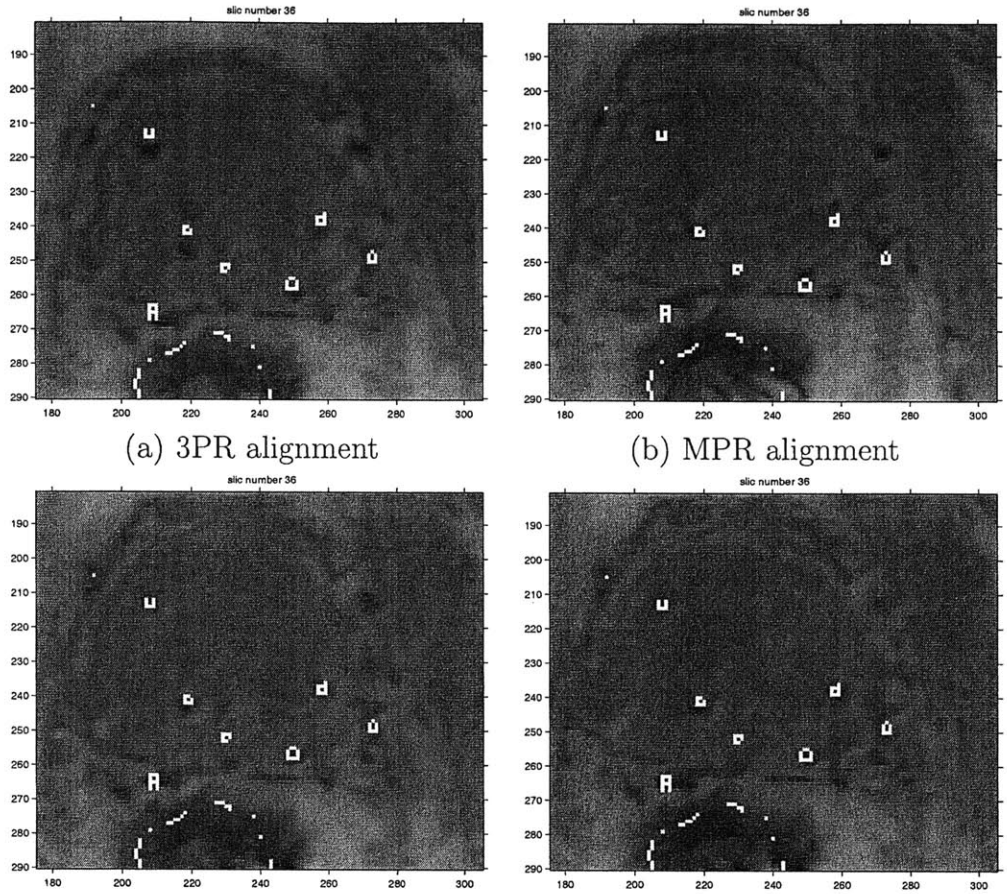

(a) 3PR alignment (b) MPR alignment

(c) MPA alignment (d) IMPA alignment

Figure 7-5: This figure shows Aligned MR images with overlaid CT edges. CT images are fixed, while MR images are reformatted according to 3PR, MPR, MPA, and IMPA registration techniques. These images are mid-volume axial slices of the volumes, zoomed in on the prostate capsule. Bright spots in the image correspond to CT seeds. Partially obscured or closely neighboring these white spots are slightly larger dark spots which generally represent seeds in MR. 3PR (a) shows some slight vertical errors in aligning the dark MR seeds with the light CT seeds. MPR (b) clearly improves on this in the leftmost seeds. MPA (c) and IMPA (d) are expectedly similar, and improve on both 3PR and MPR alignments, as can be seen over all the seeds.

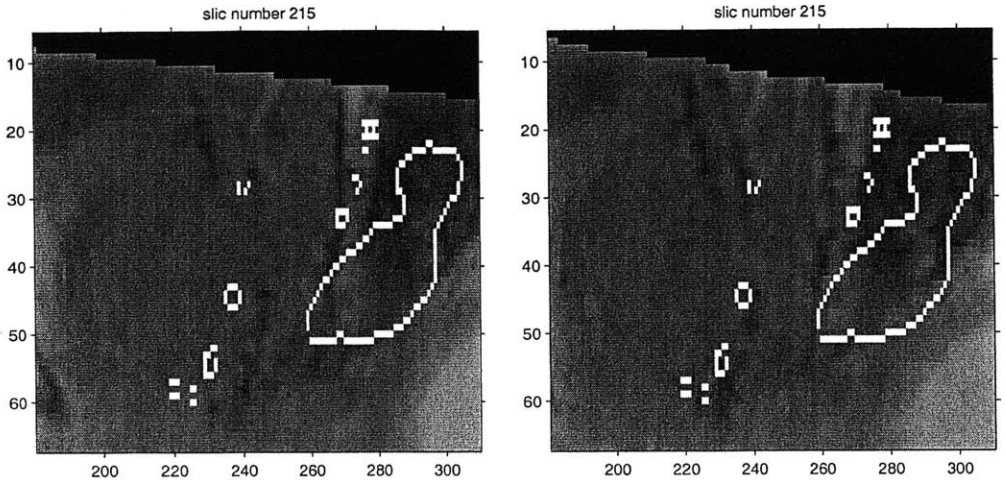

(a) 3PR alignment

(b) MPR alignment

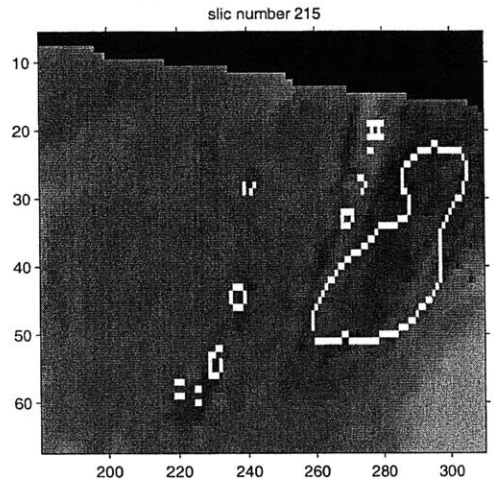

(c) MPA alignment

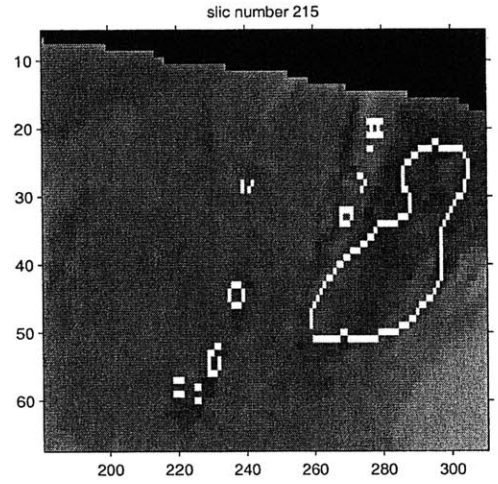

(d) IMPA alignment

Figure 7-6: These slices are sagittal views of the same alignments shown axially in Figure 7.5. Again, edges found in the CT image are overlaid on reformatted MR slices. The 3PR alignment (a) shows a consistent error (larger than a seed length) in aligning dark MR seeds with overlaid white CT seeds. MPR (b) shows improvements over 3PR. MPA (c) and IMPA (c) show improvements over both 3PR and MPR, especialy in the upper right hand corner. MPA and IMPA are very similar, as expected.

seeds within two seed lengths of each other, despite the fact that the registration is based on matches between only three seed pairs.

Manual selection of seeds and comparison of their distances in various registrations is a way to evaluate these registrations, in terms of our goal of aligning all seeds.

Chapter 8

Concluding Remarks

8.1 Summary of contributions

This thesis has addressed the challenge of accurate and precise post-implant seed localization by contributing a non-rigid model for prostate movement and an algorithm for registering unmatched, automatically found seed pairs.

The experimental evidence presented in this thesis showed improvement in post-implant MR and CT registration quality by using high numbers of seeds as control points and by using a non-rigid model of prostate motion.

Automated seed-based registration has been identified as a key goal for achieving precise seed localization. Currently, the standard technique uses only three seeds to register images, but as an experimental comparison of registration using many and few seeds showed, a three-point registration may not accurately describe the relationships between all seed pairs.

Toward automated seed-based registration, an algorithm for registering unmatched seeds has been presented and experimentally demonstrated as capable of converging on the appropriate affine solution. Furthermore, a mostly automatic CT seed-finder has been presented. Together with an automatic MR seed-finder, which is a goal of future research, this work has described a scheme by which seeds can be automatically found and registered, allowing for precise alignment of all implanted seeds.

8.2 Directions for further research

8.2.1 Improving automatic seed finding

The most enticing challenge remaining in seed localization is to find an automated solution. This thesis has presented a framework for automated seed-based MR and CT registration, which leaves unsolved the problem of automatically detecting seeds in MR images. Chapter 3 discussed possible directions to pursue to achieve automatic MR seed-finding, including shape- and intensity-based analyses.

8.2.2 Other non-rigid models

The affine model is probably the simplest non-rigid model of motion, but it is limited to maintaining parallelism. The thin-plate spline [3] generalizes the affine model to describe non-parallel warpings. Thin-plate splines are described using a variable number of additional parameters, which represent control points which constrain the warping. One direction for further research is to apply a thin-plate spline model for post-implant MR and CT registration. The major challenge in this pursuit would be to decide how to assign control points to accurately represent what we understand of prostate motion. Intraoperative prostate registration has been performed using thin plate splines [4].

Bibliography

[1] *Facts about Radioactive Seed Implants (Prostate Brachytherapy).* From Kimmel Cancer Center, Thomas Jefferson University, http://www.kcc.tju.edu/RadOnc/brachy/fac.htm.

[2] A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14:239–255.

[3] Principal warps: Thin plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Mach. Intell.*, 11, 6 1989.

[4] A comparative study of warping and rigid body registration for the prostate and pelvic mr volumes. *Computerized Medical Imaging and Graphics*, 27:267–281, 2003.

[5] *Cancer Facts and Figures.* American Cancer Society, Atlanta Georgia, 2004.

[6] A. M. Agur. *Grant's Atlas of Anatomy.* Williams and Wilkins, Baltimore, MD, 9 edition, 1991.

[7] M. Albert, C. M. Tempany, D. Schultz, M.-H. Chen, R. A. Cormack, S. Kumar, M. D. Hurwitz, C. Beard, K. Tuncali, M. O'Leary, G. P. Topulos, K. Valentine, L. Lopes, A. Kanan, D. Kacher, J. Rosato, H. Kooy, F. Jolesz, D. L. Carr-Locke, J. P. Richie, and A. V. D'Amico. Late genitourinary and gastrointestinal toxicity after magnetic resonance image-guided prostate brachytherapy with or without neoadjuvant external beam radiation therapy. *Cancer*, 98(5):949–54, Sept. 2003.

[8] R. J. Amdur, D. Gladstone, K. A. Leopold, and R. D. Harris. Prostate seed implant quality assessment using mr and ct image fusion. *International Journal of Radiation Oncology, Biology, and Physics*, 43:67–72, 1999.

[9] J. N. Aronowitz. Dawn of prostate brachytherapy: 1915-1930. *International Journal of Radiation Oncology, Biology, Physics*, 54:712–8, 11 2002.

[10] R. A. Cormack. Personal communications.

[11] R. A. Cormack, H. Kooy, C. M. Tempany, and A. V. D'Amico. A clinical method for real-time dosimetric guidance of transperineal ^{125}i prostate implants using interventional magnetic resonance imaging. *International Journal of Radiation Oncology, Biology, Physics*, 46:207–214, 2000.

[12] A. V. D'Amico, R. Cormack, C. M. Tempany, S. Kumar, G. Topulos, H. M. Kooy, and C. N. Coleman. Real-time magnetic resonance image-guided interstitial brachytherapy in the treatment of select patients with clinically localized prostate cancer. *International Journal of Radiation Oncology, Biology, Physics*, 42:507–515, 1998.

[13] A. V. D'Amico, A. Davis, S. O. Vargas, A. A. Renshaw, M. Jiroutek, and J. P. Richie. Defining the implant treatment volume for patients with low risk prostate cancer: does the anterior base need to be treated? *International Journal of Radiation Oncology, Biology, Physics*, 43:587–590, 1999.

[14] A. V. D'Amico and N. J. Vogelzang. Prostate brachytherapy: Increasing demand for the procedure despite the lack of standardized quality assurance and long term outcome data. *Cancer*, 86:1632–1634, 1999.

[15] A. Dempster, N. Laird, and D. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistics Society*, 39:1–38.

[16] D. F. Dubois, J. William S. Bice, and B. R. Prestige. Ct and mri derived source localization error in a custom prostate phantom using automated image coregistration. *Medical Physics*, 28:2280–2284, 11 2001.

[17] B. K. P. Horn. *Robot Vision*, pages 437–438,450. McGraw-Hill Higher Education, Cambridge, MA, 1986.

[18] M. Van Herk, A. Bruce, A. G. Kroes, T. Shouman, A. Touw, and J. V. Lebesque. Quantification of organ motion during conformal radiotherapy of the prostate by three dimensional image registration. *International Journal of Radiation Oncology, Biology, Physics*, 33:1311–1320, 1995.

[19] S. Nag, W. Bice, K. DeWyngaert, B. Prestidge, R. Stock, and Y. Yu. The american brachytherapy society recommendations for permanent prostate brachytherapy postimplant dosimetric analysis. *International Journal of Oncology, Biology, Physics*, 46:221–230, 2000.

[20] V. Narayana, P. Roberson, R. Winfield, and P. McLaughlin. Impact of ultrasound and computed tomography prostate volume registration on evaluation of permanent prostate implants. *International Journal of Radiation Oncology, Biology, and Physics*, 39:341–346, 1997.

[21] H. Ragde, G. Grado, B. Nadir, and A. Elgamal. Modern prostate brachytherapy. *CA: A Cancer Journal for Clinicians*, 50:380–93, 2000.

[22] J. E. Sylvester, P. Grimm, J. Blasko, C. Heaney, and J. Rose. Modern prostate brachytherapy. *Oncology Issues*, 17:34–39, 2002.

[23] J. Talcott, P. Rieker, J. Clark, K. Propert, J. Weeks, C. Beard, K. Wishnow, I. Kaplan, K. Loughlin, J. Richie, and P. Kantoff. Patient-reported symptoms after primary therapy for early prostate cancer: results of a prospective cohort study. *Journal of Clinical Oncology*, 16(1):275–83, 1998.

[24] E. W. Weisstein. *Nelder-Mead Method.* From MathWorld–A Wolfram Web Resource, http://mathworld.wolfram.com/Nelder-MeadMethod.html.

[25] E. W. Weisstein. *Paired t-test.* From MathWorld–A Wolfram Web Resource, http://mathworld.wolfram.com/Pairedt-Test.html.

[26] E. W. Weisstein. *Rodrigues' Rotation Formula.* From MathWorld–A Wolfram Web Resource, http://mathworld.wolfram.com/RodriguesRotationFormula.html.

[27] E. W. Weisstein. *Rotation Matrix.* From MathWorld–A Wolfram Web Resource, http://mathworld.wolfram.com/RotationMatrix.html.

[28] W. M. Wells. Statistical object recognition, 1993.

[29] W. M. Wells. Statistical approaches to feature-based object recognition. *International Journal of Computer Vision*, 21:63–98, 1997.