Novel Approaches for Multi-Modal Imaging and Fusion in Orthopaedic Research for Analysis of Bone and Joint Anatomy and Motion

by

Rupin Dalvi

B. E., University of Mumbai, 2003

M. Sc., University of Surrey, 2004

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

Master of Applied Science

in

THE FACULTY OF GRADUATE STUDIES

(Electrical & Computer Engineering)

THE UNIVERSITY OF BRITISH COLUMBIA
(VANCOUVER)

November 2009

Abstract

Faced with an increasingly aging and overweight population, our modern societies, particularly in the west, are set to witness a steep rise in various orthopaedic health problems in the coming decades, especially joint diseases such as arthritis. Better understanding of the way bones of the joints work is thus imperative for studying the nature and effects of these diseases and for finding cures. The data obtained from conventional sources such as skin markers and x-ray/fluoroscopy scans are generally useful but quite limited in terms of accuracy, quantification abilities and three-dimensional visualization potential. The continuous increase in the quality and versatility of various modern imaging modalities is enabling many new means for enhanced visualization and analysis of motion data of the joints. In this thesis we make use of ultrasound (US) and magnetic resonance (MR) imaging to facilitate robust, accurate and efficient analysis of the bones of joints in motion. We achieve this by obtaining motion data using 3D US with high temporal resolution which is then fused with a high spatial resolution, but static MRI volume of the same region (we mostly focus on the knee joint area). Our contributions include novel ways for improved segmentation and localization of the bones from image data. In particular, a highly effective method for improving bone segmentation in MRI volumes by enhancing the contrast at the bone-cartilage interface is proposed. Our contribution also focuses on robust and accurate registration of image data. To achieve this, a new method for stitching US bone volumes is proposed for generating larger fields of view. Further, we also present a novel method for US-MRI bone surface registration. The tools developed during the course of this thesis facilitate orthopaedic

research efforts aiming to improving our understanding of the workings of the joints. The tools and methodologies proposed are versatile and expected to be applicable to other applications.

Table of Contents

List of Figures

Acknowledgements

I would like to thank my supervisor Dr. Rafeef Abugharbieh for her continued support and guidance and my collaborator Dr. David Wilson for his assistance with the imaging and clinical aspects of this work. I am also very grateful to Dr. Derek Wilson, Emily McWalter, Dr. David Hunter and Ilker Hacihaliloglu for providing me with all the data that made this project possible. I would also like to thank my colleagues from the Biomedical Signal and Imaging Computing Laboratory (BiSICL), for all their help in my project and the wonderful memories over the past four years. Finally, I would like to thank my parents for their love and encouragement.

RUPIN DALVI

The University of British Columbia

November 2009

1 Introduction and Background

1.1 *Motivation*

In recent years, joint diseases have been on the rise. An ageing population has led to a huge increase in cases of arthritis (Figure 1-1) and other disorders affecting the bones and joints. In the US alone, according to 2003-2005 data from the National Health Interview Survey [1], 46.9 million (22%) of all adults have been diagnosed with arthritis. This number is set to rise to 67 million or 25% of the adult US population by 2030. Though the severity of the disease varies widely, it often takes on very painful forms that considerably restrict the mobility of the patients. In fact, 40% of the people afflicted with arthritis report that they find it "very difficult or impossible" to perform at least 1 of 9 important daily tasks including sitting and grasping. Thus arthritis is not a minor disease in either its severity or it its prevalence. It is a major disease that cripples individuals and subsequently, the society they inhabit.

Moreover, joint diseases are not limited to arthritis (as a disease), nor to the elderly (as a target group). Many joint diseases affect the young, including some like Perthes' Disease, which affects young children. Non pathological diseases, such as injuries, also add their share. Injuries to the weight bearing joints such as the hip and the knee, often requiring surgery and considerable physiotherapy, comprise some of the most common forms of sporting injuries [2].

Diseases that affect so many people are of great importance and recently, there has been an increase in the interest shown in joint disease research. One of the major areas that orthopaedic researchers are interested in is analysis of the motion of the joints that are affected. Diseases of the joints are very likely to bring about a change, both in the actual way the joint moves as well as the level of mobility, in addition, to being painful. They could also be caused by the way the joint moves.

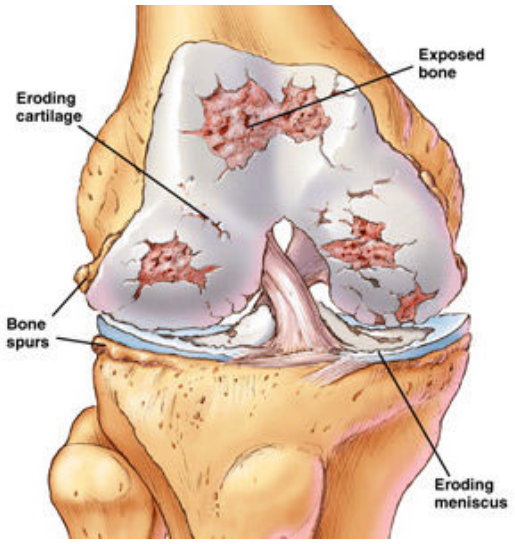

Figure 1-1 : Osteoarthritis of the knee

Study of the motion of the joints of afflicted people and their comparison to similar studies made on healthy joints could thus lead to better diagnosis and prognosis as well as better treatment planning. Moreover, it could lead to better rehabilitation assessment, following surgery, physiotherapy or pharmacological treatment, thus leading to better rehabilitation itself. It could help provide better plans to prevent the diseases or to restrict them to mild levels. The subsequent improvement in the quality of life of so many individuals would result in their leading a more active lifestyle and this could have a very positive impact on an ageing society as a whole. In the world of sports, motion analysis of the joints could lead to improvements in training that both reduce chances of injury and improve performance. Thus there is a lot to be gained, not just by the sick but also by the healthy by this research and our goal is to provide orthopaedic researchers analysis

tools for such work.

1.2 *Objectives*

The objective of the research presented is to help improve the mechanisms of joint motion analysis in people with joint diseases. Specifically, we aim to develop a system that obtains joint motion data at high speed using ultrasound (US) imaging, that is then fused with high resolution magnetic resonance imaging (MRI) volumes of the joint. This enables orthopedic researchers to visualize the motion of the joint in three dimentions (3D) and at a higher resolution. More specifically, this involves the segmentation of the tissues of interest (bone and cartilage) from the various imaging datasets and their alignment to a common frame of reference. To that end, we have developed novel techniques that, robustly and accurately segments bone and registers US and MRI image data.

The general problem facing motion analysis from an imaging perspective is that joint motion data are obtained in a variety of ways and using many modalities. By and large, the motion data themselves are collected at a low resolution. For example, in the case of US, orthopaedic scans only capture the surface of the bones with high levels of noise present at limited fields of view in order to increase the speed of acquisition. However, this means that such data are not sufficient for completing the desired motion analysis. Hence, such acquisitions often have to be merged (registered) with higher resolution data that are mostly operate as static imaging tools. Such higher resolution data (e.g., computed tomography (CT) or MRI) in turn need to be accurately segmented so that the tissue of interest (in our case, bone) is clearly localized. Once both registration and

segmentation tasks are achieved, the end result will enable proper analysis of the motion data, allowing orthopedic researchers to properly visualize and quantitatively analyze both bone morphology, pathology and motion in 3D. In summary our targeted problem comprises two distinct areas: (a) Segmentation of the high resolution MRI images and, (b) registration of the segmented MRI bones to the motion data (US volumes in our case).

1.2.1 Bone Segmentation

High resolution bone imaging data are generally provided by two modalities – CT and MRI. From an imaging perspective, each has its own set of advantages and disadvantages. Segmentation of bone from CT is relatively easy, since the bone stands out sharply against its surrounding soft tissue. Moreover, the depicted bone is cortical, which is the outer hard bone whose surface will later be imaged using US. However, CT images employs ionizing radiation, and hence poses a health risk that makes it infeasible for longitudinal studies or for research on healthy volunteers. For this reason, researchers investigated the use of MRI for acquiring high resolution static images of bone regions of interest. However, in the scenario where MRI is used, bone segmentation itself becomes a fairly difficult problem since in that case, it is the trabecular (inner) and not the cortical bone that is imaged (the cortical bone does not show up in MRI). Close to the joint, the cortical bone thins out significantly, thus the trabecular bone may come into contact with various soft tissues that have MRI intensities that are close to that of the trabecular bone. This, along with other problems such as signal fading, noise, chemical shift, image non-uniformity and partial volume effects, mean that the boundaries between the bone and the surrounding tissues become blurred making accurate bone segmentation quite difficult. In

the case of the low resolution MRI scans, the sheer lack of resolution aggravates the problem considerably. Current MRI bone segmentation methods often rely on high resolution MRI data, a condition that motion analysis data violates out of necessity due to imaging speed limitations. Using multi contrast imaging, one may be able to get higher contrast between objects of interest at lower resolutions while actually getting the images at higher speeds. The provision of such imaging approaches as well as a robust segmentation framework for motion analysis data using multi contrast MR imaging was one of this project's objectives.

1.2.2 Bone Registration

Registration of muti-modal data at differing temporal and spatial resolutions poses significant challenges. In the US-MRI case, US probes typically have a limited field of view with one probe's scan output typically not having enough information to accurately localize the corresponding region in the MRI volume, which typically has a much larger field of view. Hence, one may have to first register adjacent US volumes together stitching them into a volume with a large field of view. Following that, the second problem of registering the stitched US data to the corresponding MRI images need to be addressed.

Registration of US volumes for stitching has received some attention in the past few years. However, most of the methods in the literature have focused on the registration of deformable organs such as the liver. Given that our area of interest is the registration of bones (which are rigid) these methods end up being needlessly time consuming and in many cases, the various complexities are actually counterproductive. Hence, one of this

Thesis' objectives focused on development of a fast and accurate method for US registration for volume stitching purposes. Further, the thesis also addresses the US-MRI registration challenge where most current multi-modality alignment methods perform unsatisfactorily.

The rest of the chapter presents an overview of the various imaging modalities commonly used in orthopaedic research as well as a literature review of the methods currently being employed to deal with the problems outlined as the thesis objectives.

1.3 *Medical Imaging in Joint Disease Studies*

1.3.1 X-ray Imaging (Radiography/Fluoroscopy, CT)

Radiographs (x-ray images) and fluoroscopic data (x-ray movies) have traditionally been the standard imaging modality for studying joints, especially joint damage and bone erosion in people with problems like arthritis [74], [75] due to the relative simplicity of the equipment and the low cost of operation. However, when it comes to gauging small changes, such as those due to early bone erosion in arthritic patients, fluoroscopy/radiography is not as sensitive as other modalities [76], [77], [78]. As a result, other modalities are becoming increasingly popular. In the area of motion analysis, since fluoroscopy/radiography images only provide 2D information, they cannot be used on their own for providing full motion data. However, they can provide high temporal resolution and hence, are fused with 3D static data (such as a CT scan) to obtain 3D motion data [80], [81]. (Roentgen stereophotogrammetric analysis (RSA) [82], [83] which uses biplanar X-rays, can be used to get highly accurate motion information for

fluoroscopic images, but it requires the implantation of at least 3 fiducials in the bone, making it impractical for many studies.)

CT is a tomographic radiographic imaging method offering high-resolution and three-dimensional visualization of a large range of body tissues (Figure 1-2). CT is perhaps the most sensitive modality for studying bone information available right now, especially for assessing early bone erosions [84], [85]. However, CT has lower temporal resolution compared to fluoroscopy and moreover, the radiation dose incurred proscribes its usage for motion capture in orthopaedic settings. Thus, for motion analysis purposes, it is often used to provide high resolution static data which is fused with the temporal information provided by fluoroscopy in order to get complete motion data as stated before [80], [81].

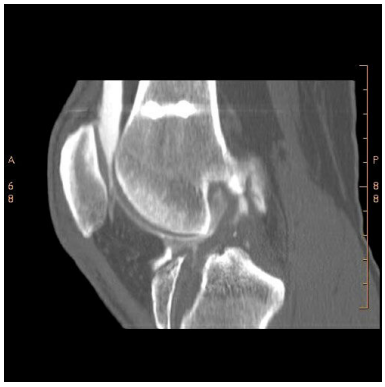

Figure 1-2: CT scan of the knee

The fluoroscopy-CT framework allows researchers to obtain motion scans of the joints with the high temporal resolution (due to the fluoroscopy) and high anatomical accuracy (due to the CT). Also, it is considerably cheaper than MRI. Hence the bulk of joint motion data obtained up to now have been using this framework. However, the use of x-rays, a form of ionizing radiation, requires that the potential risks from a procedure be carefully balanced with the benefits of the procedure to the patient. This issue makes the use of this framework impractical for researching joint motion of controls/healthy volunteers, who would be unjustifiably subjected to harmful radiation. Another problem with this framework is that soft tissues show up poorly in fluoroscopy and hence it is not very well

suited to study the changes in soft tissues due to joint disease. Nor is it a good choice for the study of soft tissues under motion. For both these reasons, MRI makes for a better option for such research.

1.3.2 Ultrasound

While the use of ultrasound is widespread in the medical community, due to its low cost, high temporal resolution and safety, it had, until recently, not found wide application in the orthopaedic field. One big reason for this is the fact that US can only image the surface of the bone and not within/beyond it (Figure 1-3).

However, in recent years, there has been growing interest in the usage of the modality in the tracking of joints during surgery as well as during joint analysis for research purposes [3], [4]. Nevertheless, the 'graininess', limited field of view (FOV) and inability to image beyond the surface restrict the mainstream usage of this modality for orthopaedic applications.

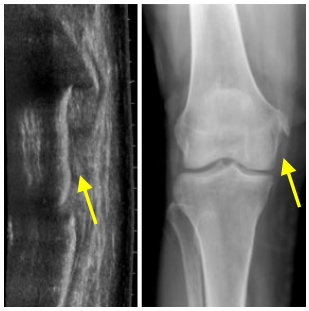

Figure 1-3: US image of the knee (left) and corresponding x-ray image (right). Note the graininess of the US image and its inability to penetrate the bone surface.

1.3.3 MRI

MRI is more sensitive than radiography/fluoroscopy when it comes to detecting subtle bone changes such as those seen during early stages of joint disease [76], [77], [78], [79] (Figure 1-4). Also, due to its good soft tissue contrast, it has been used in many cartilage studies [86], [87]. It has also been gaining traction as a modality in the joint kinematics research community [88], [91], [92], [93], [94], [93], [94].The traditional drawbacks of

MRI with regard to bone imaging are that 1) it cannot image cortical bone, and assessments have to be made based on the inner trabecular bone region and 2) it lacks the temporal resolution needed for real time motion analysis. However, with advances in MR technology, various studies of motion analysis of joints have been made possible [88], [89], [90], including some 3D motion analysis studies [91], [93], [94] using fast cine-MRI sequences. While these studies do promise improved analysis, they still have considerable limitations, especially practical ones pertaining to scanning patients with painful joint diseases. For instance, in [94], the subjects have to repeatedly flex their knees against a load for a considerable number of repetitions, thus making the method unsuitable for studying patients with arthritis. Another problem common to these studies is that the subjects have to lie down in the scanner making the knee motion artificial and unlike the normal gait that the researchers are trying to study. Nevertheless, the various benefits, including the ability to conduct longitudinal scans without exposing subjects to radiation and the ability to image soft tissues like cartilage nicely mean that MRI is becoming more commonly used as a modality for conducting orthopaedic research with.

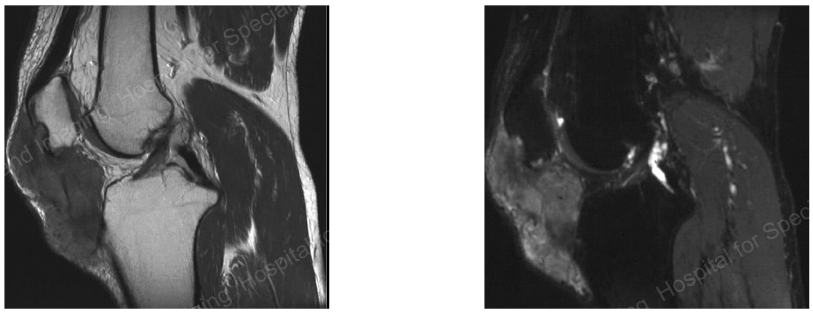

(a) T1 (b) Fat Suppressed

Figure 1-4: Different MR pulse sequences can be used to emphasize different anatomical structures . In the T1 weighted MR (a), the trabecular bone and the subcutaneous fat give higher intensity signal. In the Fat Suppressed MR (b) the same tissues are almost completely suppressed

1.3.4 Multi Spectral/Contrast MR Imaging

The usage of multi-contrast MR (T1, T2, PD, etc.) imaging for obtaining more information that may then be used for image enhancement, segmentation, etc. is made often in the area of brain imaging and angiography [95], [96]. However, it has, as yet not been used much for bone imaging and image processing. One of the reasons is probably the fact that, till very recently, bone imaging was not done using MRI. But multi contrast imaging has the potential to enhance bone images of MRI and also tremendously improve the accuracy of bone segmentation and registration.

1.4 *State-of-the-Art in US-US Volume Registration*

Due to the fact that it is cheap and safe, 3D US is becoming increasingly popular as a viable alternative imaging modality to fluoroscopy in computer-assisted orthopaedic surgery (CAOS) applications. Bone tracking using US for joint motion analysis has also been suggested [3], [4]. However, most 3D US probes have a highly limited field of view (FOV). Widening of this field of view, via the process of volume stitching, would provide great benefits to these and other clinical applications. Accordingly, 3D US volume stitching has been the subject of some recent research work. Meyer [5] used mutual information to provide affine and elastic registrations of 3D US volumes. Krücker [6] modified [5] using a sub-volume approach for faster registration. Pratikakis et al., [7] combined the minimization scheme of an automatic 3D non-rigid registration method with a multi-scale framework to register 3D US volumes. Poon et al. [8] attempted volume stitching of 3D US volumes obtained from a tracked probe. The obtained volumes were compounded using the tracking parameters and the residual errors were

corrected using two different registration methods. Wachinger et al. [9] proposed simultaneous global alignment of multiple ultrasound volumes alternative multivariate extensions based on a maximum likelihood framework. Ni et al. [10] proposed feature based alignment by registration of features obtained from the 3D US volumes using 3D SIFT. Registration of multiple freehand 3D US volumes (obtained from freehand sweeps) was attempted by Gee et al. [11]. Instead of registering the entire volume, they registered the volumes only at the dividing plane in order to improve registration speed. In a more unusual application of 3D US volume registration Pennec [12] aligned 3-D US brain volumes using a radial basis parameterization.

All these works involve some form of elastic warping of the volumes to align them. Hence, many of them are time intensive and also prone to registration errors, since a volume may be warped in more than one way to closely fit another. However, most of them deal with applications where deformable registration is necessary.

In orthopaedic applications such as CAOS or joint tracking, the tissue being registered is usually bone, which removes, or at least alleviates, some of the need for deformable registration. However, the presence of reverberations, shadows, speckles, a low signal to noise ratio (SNR) and other artifacts present a challenge, even in rigid registration frameworks and may lead to poor performance of intensity-based metrics such as sum of squared differences, normalized cross-correlation and mutual information.

Traditionally, feature based registration has been discounted under such circumstances, due to the difficulty in finding salient feature point correspondences between volumes. However, Hacihaliloglu et al. [13] showed that the usage of Phase Symmetry filters

can significantly improve the signal to noise ratio in US images of bone tissue. Moreover, recently, the SIFT algorithm by Lowe [13], has been successfully used by Chen et al. [15] for rapid feature based preregistration of multi-modal images. Ni et al. [10] have used 3D SIFT features to successfully register 3D US volumes for stitching purposes. Though their algorithm takes about a minute to register the volumes, they prove that feature based registration is a valid option for 3D US alignment.

1.5 *State-of-the-Art in US-CT, US-MRI Multimodal Registration*

As noted earlier, 3D US imaging is popular in the medical community including the image guided surgery field due to its low cost, non-ionising nature and high temporal resolution. However, US images have a great deal of speckle noise and other artifacts that make visualisation as well as other image processing (segmentation, registration, etc…) cumbersome and difficult [16]. The artifacts are often misleading, since they often resemble meaningful anatomical structures but do not correspond to them. One way to use US for its benefits and yet get around its visualisation drawbacks is registration between the US images and a visualisation friendly modality, such as CT or MRI [17], [18]. This procedure promises great improvement in the safety of surgical procedures that currently use US alone – since the CT will improve visualisation. It will also improve the safety of both patient and surgeon in the cases where surgery is guided by 2D fluoroscopic images – by replacing 2D ionizing fluoroscopy with the safer and 3D US. The validity of this procedure, however, hinges on the ability to register CT (or MRI) and US images accurately, and hence this is a critical problem. This registration falls under the category of ‘multi modal registration’.

Multi-modal registration can be model based, feature based and voxel based. Voxel based methods use all the information available in the image directly to compare the source and target images using an image similarity metric. The fact that this removes the need to segment or model the data makes these methods easy to implement. Moreover the fact that all the information is being used makes the methods statistically robust as well. There are many similarity metrics used for multi-modal image registration, e.g., [19], [20], [21], [22]. Most of them have been tried on, and found to perform well on CT and MR images. Viola et al [23] and Maes et al [19] independently proposed the now industry standard mutual information (MI) metric, Studholme et al [20] suggested an improvement - normalized mutual information (NMI). A method based on partition intensity ratio (PIU) was developed by Woods [22]. Roche et al [21] proposed a method on correlation ratio (CR).

The fundamental idea underlying all these metrics is that the image intensities are assumed to be random variables and they are assumed to have identical independent distributions. The metrics then measure the correlation between the random variables. This correlation can be viewed as function (depending on the metric) of the joint probability density function (PDF) of the random variables. While this works very well for registration between modalities such as CT and MRI, the high incidence of artifacts in the US volumes that mimic anatomically valid structures means that, for US and CT volumes, registration outputs often do not correspond to the global optimum, sometimes, even a local optimum is not obtained, as shown in [24].

A solution is to detect the useful information for registration. By extracting the regions

where the images have better correspondence information, the performance of the intensity based methods can be greatly improved. Simple thresholding was used to obtain the tissues of interest by Huang [24]. However, complex anatomical structures undermine the effectiveness of this method. Roche et al [21] proposed a robust estimation of bivariate function together with a correlation ratio method to suppress the correspondence outliers between MRI and US. However, the results are highly parameter dependent and the process is time consuming. Penney et al [25] extracted a vessel probability density map from the US images and used it to register with the MRI images. However, this method needs a learning process. This in turn involves both, a large number of US images and an empirically determined threshold for MRI images. Leroy et al [26] and Wein et al [27] used noise models to detect the artifacts. This works well on artificial data, but not very well on real life volumes.

Local features can provide unique and reliable information for registering the images with less trustable information. Stewart et al [28] proposed a method to register the retinal images by using local features. The registration starts from the most accurate local feature matching and then propagates with more global feature matching. Salient regions have been used as features for registration recently because of its higher robustness. Huang et al [24] has used multiple salient regions for 2D image registrations. Recently Hacihaliloglu et al [13] showed the effectiveness of using Phase Symmetry filters to extract bone surfaces from US, opening up new possibilities for registering orthopedic US data with CT or MRI volumes.

1.6 *State-of-the-Art in MRI Cartilage and Bone Segmentation*

Segmentation, in general, is a very well researched area, and has been widely applied in many fields including computer vision, robotics and satellite imagery. However, it is typically a much more daunting task to perform segmentation for the purpose of biomedical image analysis. The segmentation algorithm has to not only successfully reproduce the results over different scans in different studies with different patients of different disease stages, but also has to realize accuracies high enough for distinguishing small changes over certain disease progression and/or populations.

In the following section, both cartilage and bone segmentations in MRI have been discussed. Many of the methods in both cases involve the finding of the bone cartilage interface – which makes these segmentation methods, to some degree, interchangeable. Hence, though this project itself does not deal with cartilage segmentation, it has been discussed.

1.6.1 Cartilage Segmentation

Cartilage segmentation has been the topic of a great deal of research. Much of the cartilage segmentation is done, at present, manually. This is time consuming and also greatly laborious and prone to human error. Most of the algorithms present in the literature are semi automated and most of the current ones use prior knowledge of the bone cartilage interface and its relationship to the shape and volume of the cartilage, or templates obtained from manually segmented volumes in order to guide the segmentation.

The watershed transform was used to segment the cartilage semi automatically in

[29] and [30]. In [29], K Nearest Neighbour classification was first used on 15 training sets in order to generate a posterior probability map which gave a gradual transition from cartilage to background. Then an empirically decided threshold (0.9) was used to threshold the cartilage. This was then region grown by growing the largest connected component. Upon this, the watershed transform contours obtained from the original image and/or the output of the KNN classifier were imposed. By clicking on the watershed-detected regions, the user could then select or deselect regions, thus fine tuning the segmentation. In [30], prior knowledge is incorporated into the equations of the watershed transform in order to improve the segmentation. It requires the input of markers to prevent over segmentation, but it tries to generate these automatically. Results were better than the standard watershed transform, but nevertheless, not very impressive (sensitivity ~90%).

Another approach based on generating a probability map was used in [31]. An initial alignment of the template to the patient scan was determined interactively using a rigid body registration for each bone. A statistical model for the intensity distribution of structures of interest was determined from a set of voxels sampled interactively. The template was converted into a set of maps by computing a distance transform of each of these structures in the registered template. This was done to indicate where the bones and cartilage are expected to be found in the patient scan The distance transform has low values where the anatomy was expected, and high values far away from where the anatomy was expected. These maps, along with the original MRI (median filtered), were used as features, and cause the classification to vary spatially, depending upon the

template anatomy. K nearest neighbour classification was then used to classify each voxel. A 3D nonlinear deformation of the template was computed with a fast elastic matching algorithm. The process of feature space generation and classification was then repeated, using this refined template.

A spline based method was used by [32], [33] and [34]. They used fat suppressed images. The spline is manually initialized by specifying the end points and, coarsely, the control points. This is used to find the cartilage midline and then two other splines, starting at the midline, travel outward in opposite directions to fit the tibial cartilage edges as specified by directional Canny edge detectors.

A more detailed description of a B-spline segmentation method was given by [35]. They used an energy minimization approach for driving the curve deformation towards cartilage boundaries. The same approach was used by [36] who used the shape learnt from manual segmentations (actually, phantom data) and the Mumford Shah model to define the energy. A similar technique was also reported more recently by Raynauld in studies regarding knee joint osteoarthritis [37], [38].

Segmentation using a probabilistic model was done, among others, by [39]. Using an atlas, a probability map of the presence of cartilage was generated considering distance and orientation with respect to the bone cartilage interface. Bones were segmented using a region growing method and then using the above information, the posterior probabilities were obtained, giving the cartilage probabilities for the given image. 0.5 was then used as a threshold to obtain the actual cartilage.

The most popular methods used for cartilage segmentation currently used are those based on Active Shape Models (ASM). These represent objects as a set of *n* labeled points (landmarks). These landmarks are extracted from a set of *s* training images either manually or automatically. A point distribution model (PDM) is then constructed to model the variation distribution of the landmarks as follows. The landmarks ($x1,y1,\ldots xn,yn$) for each training image are represented as

$$\mathrm{yi} = (\mathrm{x1}, \mathrm{y1}, \ldots \mathrm{xn}, \mathrm{yn})^{\mathrm{T}}$$

The shape vectors are aligned by scaling, rotation and translation to minimise the sum of squared distance between the landmarks. A mean shape is then calculated from the shape vectors

$$\bar{y} = \frac{1}{s}\sum_{i=1}^{s} yi \tag{1.1}$$

as well as the covariance matrix

$$S = \frac{1}{s-1}\sum_{i=1}^{s} (yi - \bar{y})(yi - \bar{y})^{T} \tag{1.2}$$

Principal component analysis (PCA) is then applied using eigenvalue decomposition of the covariance matrix. Eigenvectors corresponding to the *t* largest eigenvalues λi are retained in a matrix *P*. The number of eigenvalues to retain, *t*, is chosen such that their sum sufficiently explains the variance in the training shapes. Any shape in the training set can now be approximated by

$$y = \bar{y} + Pb \tag{1.3}$$

where b is a vector of t elements containing the model parameters, calculated by

$$b = P^T(y - \overline{y}) \tag{1.4}$$

To ensure that new shapes generated are in the allowable shape domain, the values of b are constrained to lie within the range i $\pm m\ \lambda$, where m has a value between two and three.In the literature, among many other papers, [34], [36], [40], [41], [42], [43] and [44] use variations of the ASM method described above.

1.6.2 Bone Segmentation

Though organs such as the brain, the blood vessels, the lungs, etc. have had many segmentation methods developed for them for MRI, the bones have been relatively less well addressed, and comparatively simple methods such as region growing are still used. Even those who do deal with the topic often use the bone segmentation as a coarse first step towards cartilage segmentation [39]. Kapur et al [39] proposed a two-step process where interactively-acquired seeds are followed with a texture-based region growing process resulting in a rough segmentation that is further refined using an active contour-like regularizer. Fripp et al [40] used three-dimensional (3D) active shape models initialized by affine registration to an atlas. However, they reported problems related to patella segmentation and results remained sensitive to initialization. Lorigo et al [45] used texture-based geodesic active contours [46], [47] to perform segmentation; however setting the segmentation parameters for the entire bone without leaking into the cartilage is extremely difficult. Bourgeat et al [48] used phase information of the raw MR data in addition to the magnitude (intensity) image information in order to improve segmentation of the bones of the knee. Reported results were however not very accurate, requiring

further manual refinement.

All of the previously reported segmentation methods suffer problems at weak tissue interfaces. This occurs due to a variety of factors including image blurring, noise and partial volume effects, but is often mainly attributed to poor bone-cartilage tissue contrast where different tissue intensities appear too similar and thus are not clearly distinguishable based on one single-contrast MR scan as it does not have enough discriminatory information (Fig. 2-1(a)). The influence of these complicating factors can often be reduced with more sophisticated segmentation methods such as those employing a priori information e.g. principal component analysis/active shape based methods [44], [49]. However, prior-knowledge approaches rely on assumptions such as smoothness and known shape which do not always hold in pathologically deformed bones.

An alternative approach to circumventing the problem of complicated segmentation is to improve the contrast between the bone and surrounding confounding tissues (mostly cartilage) during the imaging process itself. This could be accomplished by utilizing the MRI scanner to capture images of different contrasts i.e. using multi contrast imaging (MCI). MCI has been used in a number of application areas, most notably in brain imaging where it is commonly employed to enhance brain tissues separation [50], [51]. In bone imaging, however, application of MCI has been very limited. Ostrowitzki [52] and Carano *et al* [53] used MCI combined with clustering methods to quantify changes in bone lesions in rheumatoid arthritis using class transitional analysis. However, their interest was only in lesion changes and did not consider bone segmentation.

It is to be noted that many of the methods using ASM for cartilage segmentation also use the same methods to segment bone.

1.6.3 Objective Summary

In summary, our thesis objectives are as follows:

- MRI Segmentation:

Develop a multimodal framework for improved MRI bone segmentation (Figure 1-5)

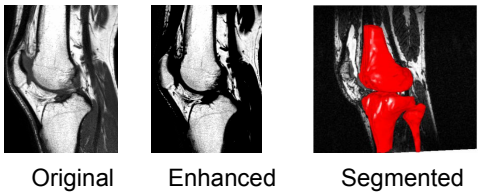

Figure 1-5: Multi-contrast MR bone segmentation

- US-US Stitching:

Develop a robust and fast stitching algorithm for increasing the field-of-view of US volumes (Figure 1-6)

Figure 1-6: US-US volume stitching

- US-MRI/CT Registration:

Develop an algorithm for US-CT/US-MRI registration (Figure 1-7)

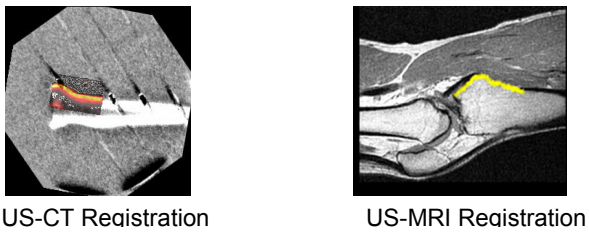

Figure 1-7: US-CT registration (left) and US-MRI registration (right)

The rest of this thesis is organized as follows:

Chapter 2 presents a method for enhanced bone tissue visualization and segmentation using multi contrast MRI. This describes the first part of our framework – robust, accurate MRI bone segmentation. The subsequent chapters deal with the processing and registration of motion capture US bone data to high resolution MRI data. In Chapter 3 a fast and efficient algorithm for the stitching of US volumes to enlarge the US field-of-view is presented. Chapter 4 deals with the multi-modal registration of US with MRI data. Chapter 5 concludes the thesis with a review of its contributions and thoughts on future improvements and projects.

2 Multi-Contrast MR Imaging for Enhanced Bone Tissue Visualization and Segmentation[1]

2.1 *Introduction*

Given the objectives of the project, the first step was to develop robust and accurate bone segmentation in MRI. In this project we used MCI for facilitating improved segmentation of bone structures in MRI based on a simple yet efficient and robust method for increasing image contrast between bone and surrounding cartilage tissue.

2.1.1 Methodology

Our proposed approach employs MR images with different tissue contrasts to enhance contrast at the bone-cartilage interface (BCI). Since trabecular bone is visible in MRI scans due to the presence of marrow fat, two scans are acquired: a T1W scan, which brightly illuminates the presence of marrow fat, and a fat suppressed (FS) scan, which suppresses and nulls the fat signal. The difference image of the two acquisitions thus results in an image where only fat remains; thus removing the surrounding cartilage and muscle tissue.

2.1.2 MR Image Acquisition

We acquired sagittal knee MR image volumes from non-arthritic volunteers (n = 9, average age = 35.6 ± 7.8 years) using a Philips Achieva 3T scanner (Philips Healthcare,

---

- [1] Versions of this work have been published.
Dalvi, R., Abugharbieh, R., Wilson, D.C., Wilson D.R.. "Highly-Automated 3D Segmentation of Femoral Bone from Hip MRI". International Society for Magnetic Resonance in Medicine Scientific Meeting (ISMRM), Berlin-Germany, 2007
Dalvi, R., Abugharbieh, R., Wilson, D.C., Wilson D.R.. "Multi-Contrast MR for Enhanced Bone Imaging and Segmentation". International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Lyon-France, 2007, pp. 5620-5623

Andover, MA, USA). A quadrature knee coil was used to obtain the T1W and FS scans for each volunteer (Fig. 2-1). The two scans were subsequently fused to increase the discriminatory power which we will later show to enable more accurate and robust tissue segmentation. The protocol parameters for the two sequences acquired were as follows:

T1W TSE (Turbo Spin Echo) scan: TR = 360 ms, TE = 10 ms, acquisition matrix 512×512, voxel size = 0.6 mm, NEX = 2, slice thickness = 2 mm, number of slices = 50, scan time = 9 min. The specific short TR was chosen to enhance bone marrow signal intensity with respect to adjacent tissues [54].

*FS scan:*TR = 11.6 ms, TE = 4.4 ms, flip angle = 15^{o}, acquisition matrix 512×512, voxel size = 0.6 mm, NEX = 2, slice thickness = 2 mm, number of slices = 50, scan time = 5 min. 3D fast gradient echo with selective water excitation (ProSet) was used.

In addition, in order to enable comparing segmentation results on lower resolution multi-contrast data to results obtained when using higher resolution single-contrast data (when the two have the same acquisition time), we also acquired sagittal images of the knee from healthy volunteers (n=2). In each case, a conventional single-contrast, high resolution T1W scan was acquired (TR = 700 ms, scan time = 652 s, dimension = 512×512×46). For the proposed multi-contrast method, two lower resolution scans (dimension = 336×336×50) were acquired for both volunteers; the first a T1W scan (TR = 180 ms, scan time = 365 s) and the second an FS scan (TR = 11.291 ms, scan time = 252 s).

2.1.3 Multi-Contrast Image Fusion

All data were first preprocessed with anisotropic diffusion (edge preserving) filtering [55] applied to reduce image noise, and MRI non-uniformity correction subsequently performed using the commonly-used non-parametric non-uniform intensity normalization (N3) technique [56]. Thereafter, both images were normalized (range: 0 to 255) and their difference image was calculated by direct subtraction. (Negative values resulting from the subtraction were neglected, i.e., set to zero.) The fat/trabecular bone is the brightest tissue in the T1W image and is completely suppressed in the fat-suppressed scans. Also, post normalization the non fat/non trabecular bone tissues (muscle, cartilage, etc…) are consistently brighter in the fat-suppressed image than in the T1W image. This means that upon subtraction of the fat-suppressed image from the T1W image (and setting negative values to zero), only the fat/trabecular bone remains in the output. The resultant enhanced image thus contains all of the bone marrow as well as surrounding subcutaneous fat but excludes the majority of the surrounding cartilage and muscle (Fig. 2-1).

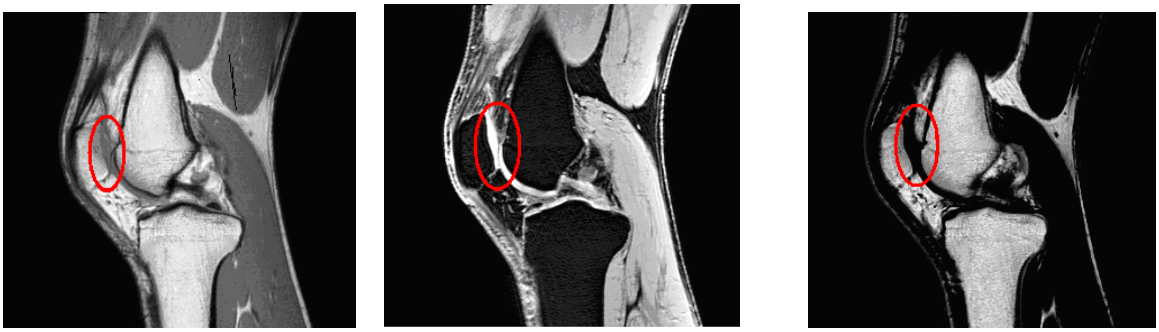

(a) (b) (c)

Figure 2-1. The proposed multi-contrast MR imaging approach. (a) Example T1W image. (b) Corresponding FS image. (c) Resultant cartilage-suppressed (difference) image. Note the poor contrast at the patella-cartilage boundary (circled) in (a) as opposed to the very clearly defined boundary in (c).

2.1.4 Bone Segmentation

To demonstrate the value of the proposed bone imaging approach, we performed segmentation using state-of-the-art methods on both the enhanced (multi-contrast based) images as well as the original T1W (single-contrast) images.

2.1.4.1 Fully Automatic Laplacian Level Set Segmentation

The first segmentation technique tested was a two step procedure incorporating a coarse initial segmentation based on region growing followed by a fine tuning step. For the *coarse* segmentation, seed points were automatically generated within the femur and the tibia. By registering the image data to be segmented to a previously labeled template image, the coordinates of the center of mass of the femur and the tibia of the registered volume were used as seed points for the segmentation process. A connectivity based region growing was then performed which is based on the mean and standard deviation of the region under consideration. First, the mean and standard deviation of all the pixel intensities currently in the region are computed. All neighboring pixels whose intensity values lie within a specified range from the mean (a user-provided factor, σ, is used to multiply the standard deviation and define a range around the mean) are then added to the region in an iterative 'region growing' manner. The values of mean and standard deviation are recalculated at each iteration and the process is repeated until a predetermined number of iterations is met.

The parameters for the coarse (confidence connected thresholding) segmentation method (viz., the number of iterations and σ) were selected as follows.

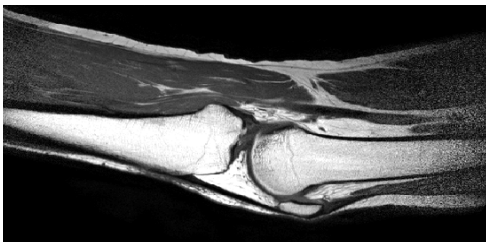

(a)

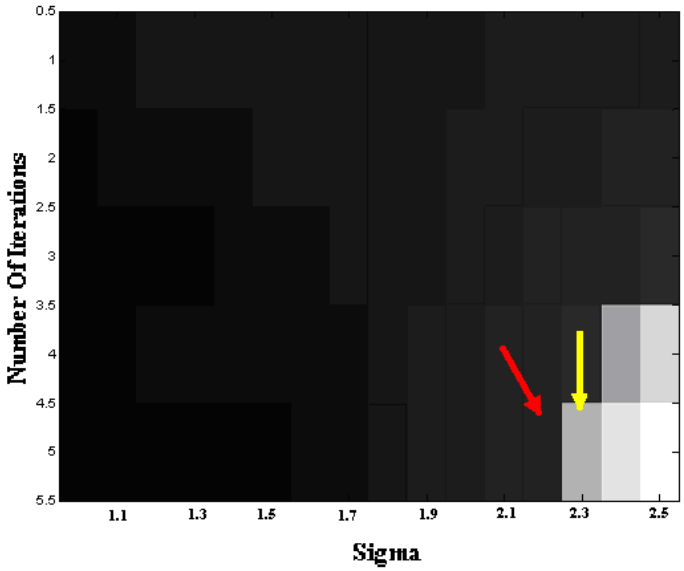

(b)

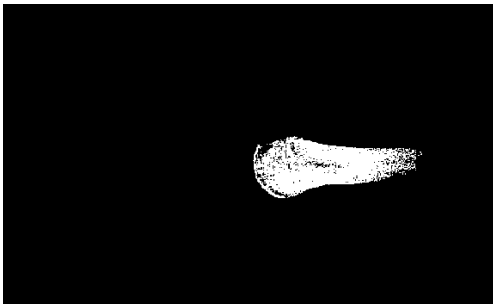

(c)

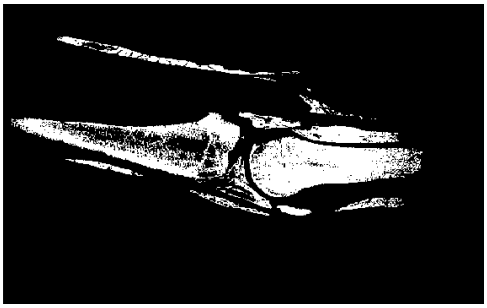

(d)

Figure 2-2: Demonstration of the process for parameter selection for the coarse segmentation (confidence connected thresholding) method. (a) original image (b) map of number of pixels in the segmented image as a function of # iterations and σ - vertical axis corresponds to number of iterations and the horizontal axis corresponds to σ. Red (oblique)arrow indicating parameter values beyond which leakage occurs (optimal values). Yellow (vertical) arrow indicating parameter values at which leakage occurs (c) segmentation of femur with optimal parameter setting (d) segmentation with parameters just beyond optimal setting.

After selecting the initial seed point as described above, the method was run over a large range of values of both number of iterations and σ and the total number of pixels in the segmented image for each set of parameters was recorded as shown in Fig. 2-2, where the vertical axis corresponds to number of iterations and the horizontal axis corresponds to σ. Given the nature of the images being dealt with, the following pattern would intuitively hold: For a given number of iterations, increasing σ increases the number of pixels in the segmented image. Once the segmentation leaks outside the bone of interest, however, there is a large surge in the number of pixels in the segmentation, since the segmentation leaks into all the surrounding fat/bone areas. By taking the horizontal gradient of the above image, one can then find the threshold points (thresholded local horizontal gradient maxima) beyond which leakage occurs, and thus, the parameters that correspond to optimal segmentation. In the event of multiple such local maxima, the optimal parameters are taken to be the ones corresponding to the highest pixel count for the segmentation. The process of parameter selection has been demonstrated in Fig. 2-2.

To *fine-tune* the coarse initial segmentation obtained, we employed an automated Laplacian level set segmentation technique [57]. Level sets are implicit representations that embed the contour in a higher dimensional function which facilitates representation of topologically complex shapes. The level-set function, denoted by $\psi(X, t)$, where X is the evolving contour at time t, is thus initialized based on the initial *coarse* contour obtained previously and then evolved iteratively. The shape of the contour at any point of

time during the curve evolution is obtained by taking the zero level set: $\Gamma((X), t) = \{\psi(X, t) = 0\}$. In essence, the Laplacian level set segmentation method attracts the evolving contour to local zero-crossings in the Laplacian (second derivative feature) image. Details of the algorithm can be found in [58]. While the algorithm can be implemented in both 3D and multi-slice 2D, we used the latter since it gave higher Dice's coefficient values.

2.1.4.2 Semi-Automated Interactive LiveWire Segmentation

The second segmentation method we used for validation employed the LiveWire approach pioneered by Barrett and Mortensen, [59], [60]. Two dimensional (2D) LiveWire was used on each slice of a given volume The method can briefly be described as follows: Let $S(q)$ be a 2D slice within an image volume, where $q = (x, y)$ is a point on the slice. Let $p=(x',y')$ be a neighboring pixel to q . A local cost map $C(p, q)$ of the original image is created using image information and path smoothness measures. Image features such as gradient magnitude cost $C_G(q)$, gradient direction cost $C_{GD}(p, q)$, Canny edge detection cost $C_C(q)$, Laplacian of Gaussian (LoG) cost $C_{LoG}(q)$, and Euclidean distance (smoothness) cost $C_d(p, q)$ (equation 2.1) were used in our implemetation.

$$C(p, q) = w_C.C_C(q) + w_{LoG}.C_{LoG}(q) + w_G.C_G(q) + w_{GD}.C_{GD}(p, q) + w_d.C_d(p, q) \qquad 2.1$$

where w_G, w_{GD}, w_C, w_{LoG}, and w_d are the weights for the respective terms. Starting at pixel q, a cost map $M(q)$ is thus created to all other pixels slicing the image using the accumulated path cost, and a graph search algorithm [59] is then used to find the globally minimal cost path. This minimal-cost path connecting the seed points selected by the user on the object boundary eventually constitute the outline of the desired segmentation. The

number of seed points needed to accurately segment the object depends on the image quality and object size and shape. In this paper, we used the LiveWire implementation of Poon et al. [61].

2.2 *Results*

For quantifying the accuracy of the segmentation, manually segmented data sets were used as the reference gold standard. In order to validate the improvement in segmentation with the proposed approach, we compared the performance of the described two segmentation methods on both the enhanced MCI and the T1W data set only (single-contrast). Dice's coefficients were used to measure segmentation accuracy (Eq.2.2):

$$Dice = 2|A \cap B| / (|A| + |B|) \qquad 2.2$$

where A and B represent the set of manually segmented and automatically segmented bone voxels, respectively. To account for the signal fade off that occurs in the areas that are not covered by the knee coils, which are near the top and bottom of the image, the accuracy was calculated only on the central parts of the image where the bone meet to form the joint.

To quantify the advantage in LiveWire segmentation using the enhanced, multi-contrast images as opposed to the original images the following evaluation was used: Seed points for the LiveWire were initialized using the manual segmentations. In the first iteration, every manual point on the boundary was used as a seed point; in the second iteration, every second point was used, then every third and so on. The spacing between the manual points used as seeds when the segmentation accuracy of the volume dipped

just below 97% (Dice's coefficient) was established. Clearly, greater spacing implies lesser number of seeds and thus reduced user input/effort and greater ease of segmentation. Table 2-1 shows the largest spacing possible between seeds for both original (single-contrast) and enhanced (multi-contrast) images for Dice's coefficient accuracy of 97%. It can be noted that our proposed multi-contrast approach leads to significant improvement in the semi automated segmentation using LiveWire with the average spacing between seeds increasing from 7 pixels for the original images and 11 pixels for the enhanced, multi-contrast images. The results shown in the paper are the ones using the default parameters for the program.

Table 2-1: Comparison results of maximum spacing in pixels between seed points for semi automated (LiveWire) segmentation of knee bones on 9 subjects from using traditional single-contrast and proposed multi-contrast approach at the same DICE accuracy of 97%.

Subject No.	Femur		Tibia		Patella	
	Single-Contrast	Multi-Contrast	Single-Contrast	Multi-Contrast	Single-Contrast	Multi-Contrast
01	6	12	7	11	7	12
02	7	11	8	11	6	11
03	8	10	8	12	8	11
04	7	12	6	11	6	12
05	8	11	9	12	7	9
06	5	9	7	11	6	10
07	6	11	7	11	8	11
08	7	10	8	11	7	11
09	8	12	5	10	9	12
Mean	6.89	10.89	7.22	11.11	7.22	11
Mean (p-values)	0.0037		0.0064		0.0059	
Std	1.05	1.05	1.20	0.60	1.20	1

Table 2-2: Comparison of quantitative results of automated Laplacian level set segmentation (using Dice's coefficients) of knee bones in 9 subjects when performed on single-contrast images and the proposed multi-contrast imaging approach.

Subject No.	Femur		Tibia		Patella	
	Single-Contrast	Multi-Contrast	Single-Contrast	Multi-Contrast	Single-Contrast	Multi-Contrast
01	94.15	96.71	95.13	96.30	88.55	89.93
02	94.35	97.16	95.24	96.32	87.39	88.80
03	94.67	96.55	96.49	97.88	88.83	91.34
04	94.17	96.55	94.48	95.51	89.59	90.07
05	95.51	97.58	95.93	96.85	88.09	89.37
06	93.23	95.95	89.68	93.94	61.52	72.87
07	95.15	96.29	95.39	96.11	90.66	91.90
08	94.87	95.99	96.18	97.49	86.30	91.38
09	94.30	96.82	95.22	96.49	89.29	89.94
Mean	94.49	**96.62**	94.86	**96.32**	85.58	**88.40**
Mean (p-values)	<<0.05		0.0174		0.0518	
Std	0.66	**0.53**	2.04	**1.14**	9.11	**5.91**

Table 2-2 shows the results of the fully automatic segmentation (in terms of Dice's coefficients) of both the single-contrast images (SCI) and the MCI using the segmentation method described above. The parameters used for the level set segmentation in our tests were empirically set to 50 for the propagation constant, 10 for the curvature scaling and 7500 for the iteration number. Quantitative results demonstrate superior segmentation of the MCI data with improvements in Dice's coefficient of 2.13%, 1.46% and 2.82% for the femur, the tibia and the patella respectively. The improvements were statistically significant for the two larger bones ($p<<0.05$ for the femur and, $p=0.0174$ for the tibia). For added clarification, the sensitivity and specificity results for the segmentation have also been documented (Tables 2-3 and 2-4).

Table 2-3: Comparison of quantitative results of automated Laplacian level set segmentation (sensitivity) of knee bones in 9 subjects when performed on single-contrast images and the proposed multi-contrast imaging approach.

Subject No.	Femur		Tibia		Patella	
	Single-Contrast	Multi-Contrast	Single-Contrast	Multi-Contrast	Single-Contrast	Multi-Contrast
01	96.0991	94.7111	94.8192	95.3769	88.3570	87.0718
02	95.7865	95.6667	95.5415	93.7591	87.5151	87.8834
03	95.0796	94.1035	94.6193	97.4533	85.5351	87.9603
04	95.7719	96.2916	95.0621	95.8312	85.8285	87.1554
05	95.1409	95.0135	93.3363	95.1940	87.4288	85.7546
06	93.8679	96.9400	94.0193	95.5620	87.5196	85.9932
07	93.6589	96.9146	95.6333	94.8358	86.4819	84.8559
08	93.8648	95.8743	95.2234	95.2316	87.3006	86.5740
09	94.6363	95.4862	94.3444	94.4038	86.8057	85.2801
Mean	94.87843	95.66683	94.7332	95.29419	86.9747	86.50319
Mean (p-values)	0.1951		0.2647		0.3740	
Std	0.923091	0.963597	0.745362	1.024821	0.899288	1.107466

Table 2-4: Comparison of quantitative results of automated Laplacian level set segmentation (specificity) of knee bones in 9 subjects when performed on single-contrast images and the proposed multi-contrast imaging approach.

Subject No.	Femur		Tibia		Patella	
	Single-Contrast	Multi-Contrast	Single-Contrast	Multi-Contrast	Single-Contrast	Multi-Contrast
01	93.0035	96.8661	93.9510	97.8755	87.2866	90.344
02	92.5395	94.979	94.6287	95.1956	87.5909	90.555
03	94.1794	94.6525	94.5130	95.1832	86.9017	89.8723
04	95.5147	96.3344	95.1162	98.4772	90.2278	91.8177
05	95.6755	97.1795	94.2952	96.2592	85.3428	89.711
06	94.1798	97.2369	91.5906	92.8602	66.9012	78.0955
07	95.8789	95.5578	94.2074	97.7569	87.6806	89.9258
08	93.971	95.1412	92.9936	95.9603	86.5032	88.5893
09	94.9832	97.6914	95.0525	96.073	88.2338	89.2585
Mean	94.436167	96.182089	94.0386	96.18234	85.1854	88.68546
Mean (p-values)	0.0048		0.0014		0.0087	
Std	1.1767443	1.1256552	1.1154	1.72076	6.983597	4.070327

A brief note on the parameter selection for the level set segmentation. Given that the level set segmentation is part of an automated method, specifically finding the optimal segmentation parameters for each bone would have been somewhat counterproductive. The results obtained in Table 2-2 were therefore obtained with a fixed set of parameters. These parameters were selected after empirically testing various parameters on one single-contrast femur volume to get a good segmentation value and then using the same parameters for all the segmentations. However, to find out whether the segmentations were robust to parameter tuning, we segmented 3 bones (2 femurs and 1 tibia) from both the single-contrast and multi-contrast data (6 cases in all) with a range of parameters and found the optimal parameters for segmenting each bone. In each instance, we segmented the other 5 cases with these parameter values and found out the mean performance. The results can be seen in Table 2-5. We found that the optimal parameters for all bones were very similar to each other and that the results were very similar to the results documented for the same bones in Table 2-2. We also found that when bones were segmented using the parameters optimized for another bone, the segmentation performance was only slightly worse. This can be seen from the fact that the mean of the optimized segmentations is only marginally better than the mean of the segmentations performed using the optimized parameters of one bone on the other bones. Thus, it would be fair to assume that the parameters used for the results obtained for Table 2-2 gave a nearly optimal segmentation performance. It should be noted that the number of iterations was set at 8000 since empirical observation showed that, unless there was leakage, the algorithm tended to converge before 8000 iterations.

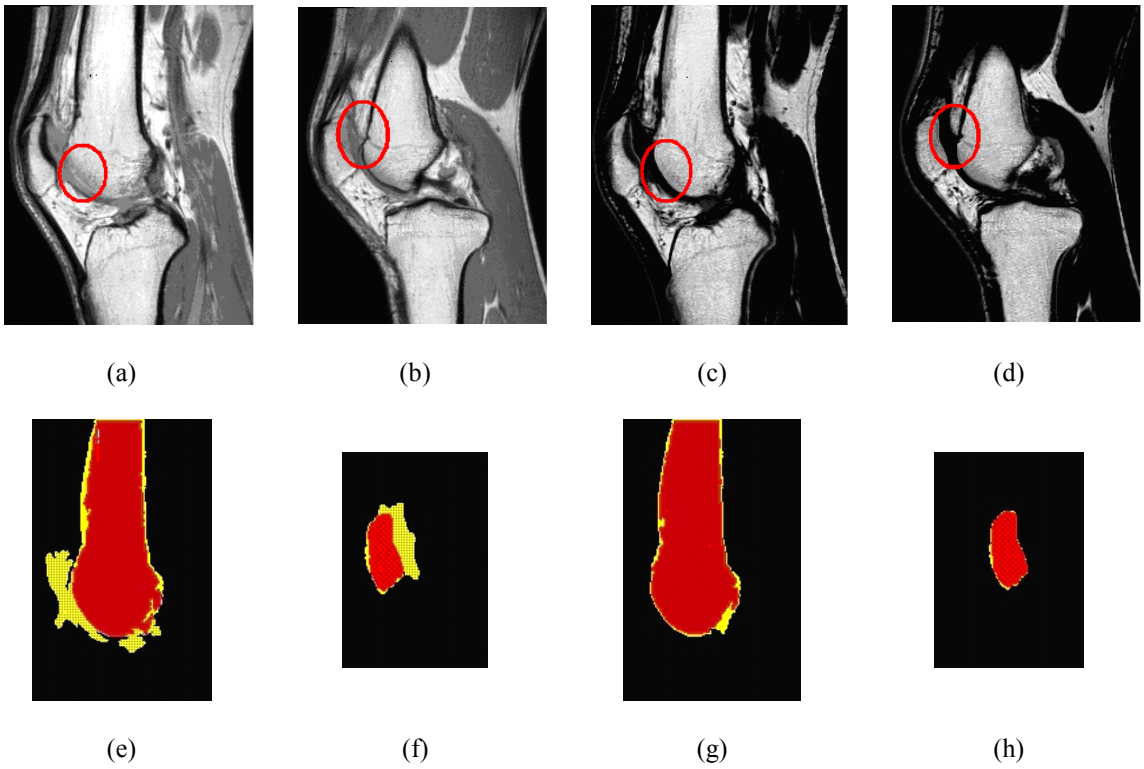

Figure 2-3: Qualitative illustration of the advantage of using the MCI enhancement. (a),(b) Original images. (e),(f) Results of segmentation using Laplacian level set segmentation on original single contrast data of femur in (a) and patella in (b). (c),(d) Contrast enhanced images. (g),(h) Corresponding segmentation using Laplacian level set segmentation on the proposed enhanced data of femur in (c and patella in (d) - note the highly superior segmentation results. Yellow indicates the areas of incorrect segmentation. Red indicates correctly segmented regions. Note the poor bone-cartilage contrast at the locations of leakage. (Red circles in (a), (b), (c) and (d) indicate points of leakage at the bone-cartilage interface.)

Figure 2-3 demonstrates the qualitative segmentation improvements achieved when using MCI compared to traditional single-contrast imaging. By sharply improving the contrast at the bone cartilage interface, our proposed approach successfully prevents leakages to which the single-contrast image data segmentation is highly prone. Figure 2-4 shows the effect of MCI enhancement on automated segmentation at the bone-subcutaneous fat interface.

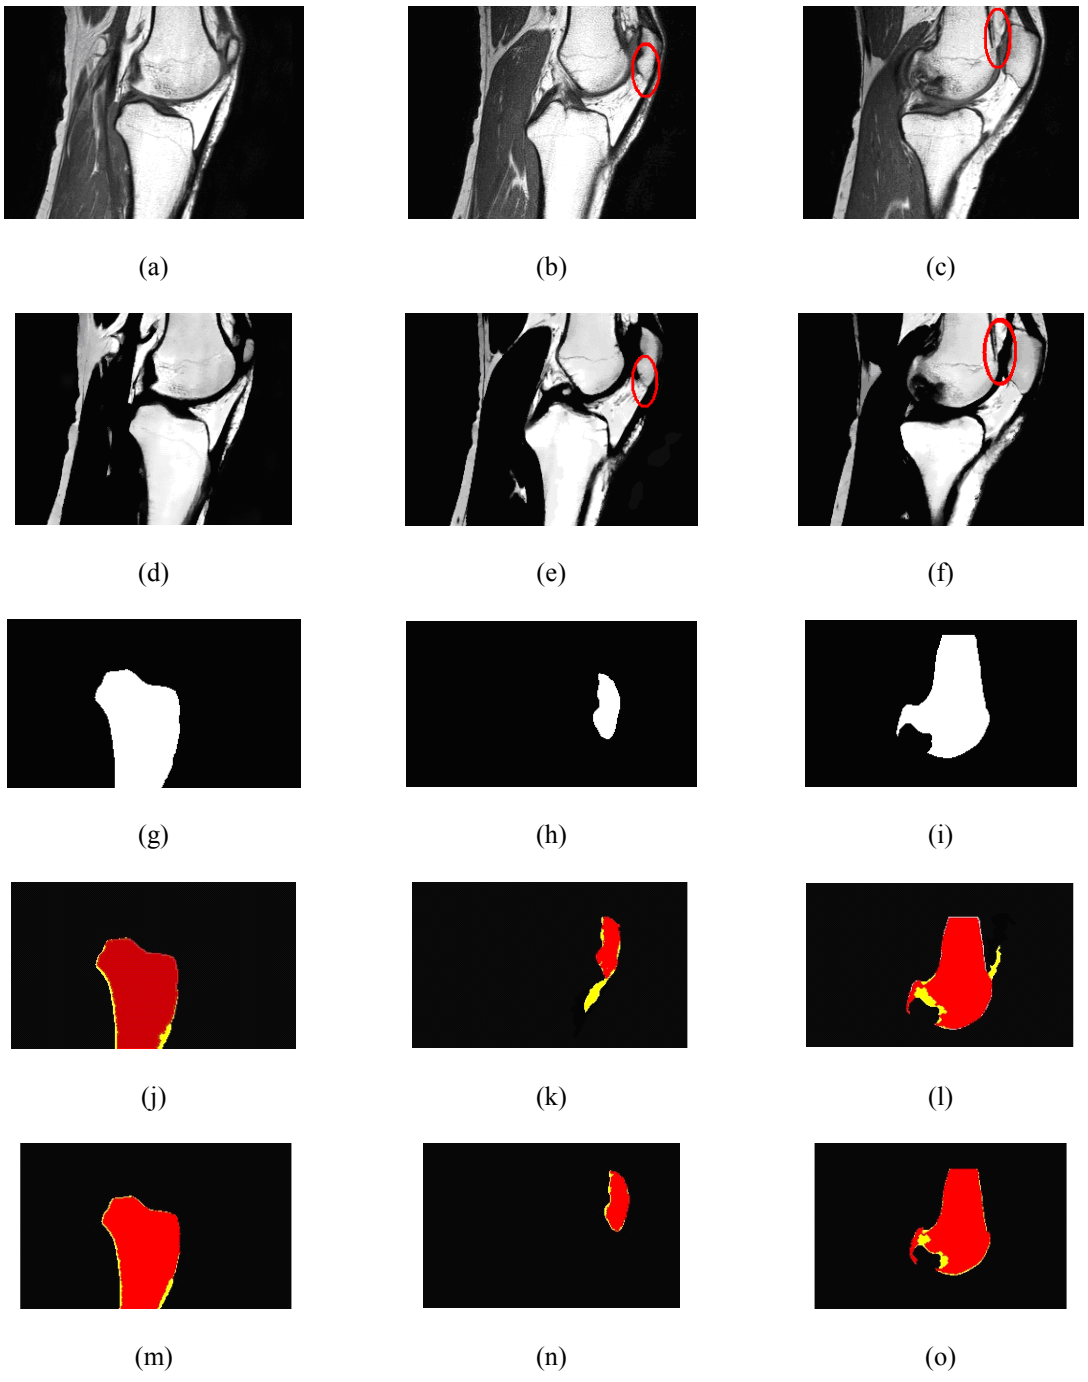

Figure 2-4: Illustration of the advantage of using the MCI enhancement at the bone-subcutaneous fat interface. Row 1: the original images. Row 2: the corresponding contrast enhanced, multi-contrast images. Row 3: manual (gold standard) segmentation. Row 4: segmentation from the original (single-contrast) images. Row 5: segmentation from the enhanced (multi-contrast) images. Yellow indicates the areas of incorrect segmentation. Red indicates correctly segmented regions. (Red circles in Rows 1 and 2 indicate points of leakage at the bone-subcutaneous fat interface.)

Table 2-5: Table demonstrating of robustness of segmentation to level set parameters. Number of iterations was set at 8000 since for most cases (where leakage did not occur) convergence was reached before then.

	Bone	Optimal Values		Dice's Coefficient	Mean Dice's Coefficient of other 5 volumes with these values
		Propagation Constant	Curvature Scaling		
Single-Contrast	Femur 3	50	10	94.8672	96.0498
	Femur 5	30	10	95.8826	95.8562
	Tibia 9	50	10	95.6593	95.9512
Multi-Contrast	Femur 3	60	8	96.8795	95.7228
	Femur 5	50	8	97.7326	95.6816
	Tibia 9	50	10	96.8429	95.8707
Mean				96.3107	95.9437

2.3 *Discussion*

In absolute terms, the accuracy of the segmentation of the MCI data was very high for both the femur and the tibia (mean Dice's coefficients of 96.6% and 96.32% respectively). It was lower for the patella (88.4%). However, this is due to the fact that the patella is a much smaller bone than the femur and the tibia and hence even small errors in segmentation (in terms of pixels over- or under-segmented) penalize the final Dice's coefficient heavily. Also, the patella is in close proximity to a bank of fat. Given that the patella is seen in the MRI scan because of the marrow fat in it, the adjoining fat pad makes segmentation very difficult even after the enhancement. Nevertheless, the segmentation on multi-contrast images outperforms segmentation on single-contrast images in the case of the patella as well. The MCI data also improved the consistency of the method as can be seen from the standard deviations in the results over the datasets.

Additional to the intended improvement at the BCI, there is also improvement at the bone fat interface. One possible reason for this is that the subtraction in the proposed method also has the side-effect of highlighting the often thin boundary between marrow and subcutaneous fat, which is an issue for the femur and especially, the patella, thus limiting a source of potential leakage which segmentation on the SCI data is prone to do. Figure 2-4 shows the effect of MCI enhancement on automated segmentation at the bone-subcutaneous fat interface.

Furthermore, from Tables 2-3 and 2-4, it can be readily seen that the main advantage of the proposed method lies in preventing leakages. This can be inferred from the fact that while the segmentation sensitivity is not statistically different for the single- and multi-contrast scenarios, the specificity - which indicates the level of leakage - shows that the segmentation of the multi-contrast volumes significantly outperforms segmentation of single-contrast volumes.

A potential argument regarding our method is that it requires two scans thus increasing the total amount of scanner time needed for imaging and that a single-contrast scan, taken at a higher resolution (with an increased scanning time) might in fact possess enough image detail to segment the data more accurately. To refute this point, we performed automated segmentation on single-contrast data taken at high resolution (512×512×50) and compared the results to those performed on corresponding multi-contrast data consisting of T1W and FS scans taken at a low resolution (336×336××50) over the *same* duration (please refer to Section 2 for data acquisition details). The results of this comparison are shown in Table 2-6, which show that the multi-contrast

enhancement holds an advantage even when the resolution is compromised to keep the SCI and MCI acquisition times the same to those of a single scan. We acknowledge that only two scans are insufficient to draw strong conclusions from. However, they do provide a strong indication of the potential advantages of the proposed methodology.

Table 2-6: Mean segmentation performance (measured as DICE coefficient) of the femur, tibia and patella in the standard single-contrast and multi-contrast approach where the single-contrast image is of a higher resolution than the multi-contrast image.

Tissue	Single-Contrast Data (high resolution)	Multi-Contrast Data (low resolution)
Femur	93.96	**94.79**
Tibia	92.55	**94.03**
Patella	88.16	**89.82**

The pulse sequences have not been optimized for the proposed method. Difficulties in collecting the data meant that the data we got was part of another group's study and was intended for other purposes and the pulse sequences had thus been optimized for them. This can be seen as a plus for the proposed method, however. Since it works well even when the images are not optimized for it, it will, intuitively, work even better if they are.

Another aspect that should be discussed is parameter selection. For the automated segmentation, we optimized the parameter selection for the coarse segmentation step in the way detailed in the Methods section. We also verified that the parameters selected for the fine tuning step (level set segmentation) were optimal or near optimal and that the segmentation of relatively robust to parameter selection within a certain range of values. While we currently do not have a tried and tested method for optimizing the weighting parameters for the LiveWire algorithm, our lab is currently working on an algorithm to do precisely that [97]. Once that algorithm is finalized and tested, we will re-evaluate our method's performance with respect to LiveWire. The results shown in the paper are the

ones using the default parameters for the program (and the ones, consequently, that most people using the algorithm are likely to use unchanged). Having said all this, it should be borne in mind that the LiveWire algorithm is a semi-automated technique and hence any comparison without the user in the loop is not very meaningful. We got two experts to segment the bones from both the original and enhanced volumes and both agreed that it was far easier, time saving and convenient to segment the bones on the enhanced volumes than on the original ones.

2.4 *Conclusions*

In this paper, we proposed a novel method for enhancing bone contrast in MR data using a simple and efficient multi-contrast acquisition approach. Our technique combines T1W and FS images, taking advantage of the fact that bone is seen in MRI scans principally due to the marrow fat within it. The advantages of our technique were demonstrated both quantitatively and qualitatively on real bone MR data where significant accuracy improvement was achieved when using multi-contrast data based segmentation compared to traditional single-contrast data. Our validation was carried out on knee data of nine volunteers. Quantitative improvements measured using DICE coefficients were demonstrated and qualitative improvements due to the contrast enhancement were also shown, manifesting in fewer visible leaks across boundaries, both at the targeted bone cartilage interface and also at the bone-subcutaneous fat interface. Superior segmentation was also achieved on lower resolution multi-contrast data compared to higher resolution single-contrast data when the two had the same acquisition time. Furthermore, our approach demonstrated increased efficiency and ease of use of semi-automated

segmentation method when used on multi-contrast images.

The proposed method promises to be a valuable tool for improving segmentation reliability and the associated reduction in analysis time and operator variability in clinical applications. This may potentially render large scale MR-based investigations of bone deformities and kinematic analyses more feasible.

The method described above provides the first part of our framework – robust, accurate MRI bone segmentation. The subsequent chapters relate the processing and registration of motion capture US bone data to high resolution structural data.

3 Fast and Accurate 3D Ultrasound Volume Stitching Using Phase Symmetry and Harris Corner Detection for Orthopaedic Applications[2]

3.1 *Introduction*

The registration phase of the project, as outlined in Chapter 1 comprises of two parts: The first part involves the stitching of multiple US volumes to generate a volume with a large field of view. The second involves registering the real time US volumes to pre-operative/pre-testing MRI volumes. This chapter deals with the first part.

Due to the fact that it is cheap and safe, 3D ultrasound (US) is becoming increasingly popular as a viable alternative imaging modality to fluoroscopy in computer-assisted orthopaedic surgery (CAOS) applications. Bone tracking using US for joint motion analysis has also been suggested [78], [79]. However, most 3D US probes have a highly limited field of view (FOV). Widening of this field of view, via the process of volume stitching, would provide great benefits to these and other clinical applications. Accordingly, 3D US volume stitching has been the subject of some recent research work. An overview of these methods was provided in Section 2.2.

Most of these works involve some form of elastic warping of the volumes to align them. Hence, many of them are time intensive and also prone to registration errors, since a

---

- [2] A version of this work has been accepted for publication. Dalvi, R., Hacihaliloglu, I., Abugharbieh, R.. "Fast and Accurate 3D Ultrasound Volume Stitching Using Phase Symmetry and Harris Corner Detection for Orthopedic Applications". SPIE Medical Imaging (MI), San Diego-USA, 2010

volume may be warped in more than one way to closely fit another. However, most of them deal with applications where deformable registration is necessary.

In orthopaedic applications such as CAOS or joint tracking, the tissue being registered is usually bone, which removes, or at least alleviates, some of the need for deformable registration. However, the presence of reverberations, shadows, speckles, a low signal to noise ratio (SNR) and other artifacts present a challenge, even in rigid registration frameworks and may lead to poor performance of intensity-based metrics such as sum of squared differences, normalized cross-correlation and mutual information.

Traditionally, feature based registration has been discounted under such circumstances, due to the difficulty in finding salient feature point correspondences between volumes. However, Hacihaliloglu et al. [13] showed that the usage of Phase Symmetry filters can significantly improve the signal to noise ratio in US images of bone tissue. Moreover, recently, the SIFT algorithm by Lowe [13], has been successfully used by Chen et al. [15] for rapid feature based preregistration of multi-modal images. Ni et al. [10] have used 3D SIFT features to successfully register 3D US volumes for stitching purposes. Though their algorithm takes about a minute to register the volumes, they prove that feature based registration is a valid option for 3D US alignment.

In this section, we use probe tracking, Phase Symmetry, Harris corner detection and sub-volume registration to obtain fast and accurate rigid registration which we will be using for bone volume stitching. Using the parameters of the tracked probe, coarse registration between the two volumes are obtained. The central slices of these volumes thus contain similar structures. These slices are then filtered using Phase Symmetry filters to

boost the signal to noise ratio and obtain a good bone surface and then Harris corner features are identified and matched. The most appropriate corresponding points are chosen based on how well they correlate with each other and also how feature rich their neighbourhoods are, and sub-volumes around these points are then registered on a volumetric basis. The parameters of this registration are used to register the two volumes.

As an additional application/verification of the algorithm, we use it to track incidental intra-surgical bone motion during CAOS, using a setup similar to that of Lavallee et al [3]. A tracked probe attached to the patient obtains volumes at periodic intervals. Each successive volume is registered to the previous one using the aforementioned method. From the tracking and registration parameters, the bone motion can be inferred. (N.B.: Here, we will be only demonstrating the bone registration, not the inference of the motion there from.)

3.2 *Method*

The first step in the algorithm is to use the parameters obtained from the probe tracking to generate coarse registration. Once this is done, the central slices will contain many corresponding features. Before getting at the features, however, we must remove the noise and boost the bone signal to noise ratio in the slices.

Hacihaliloglu et al. [13] proposed the use of a ridge detector, rather than an edge detector, to identify and enhance the signal to noise ratio of the bone surface, in order to avoid the double sided response at the bone surface that edge detection produces. We have used the method they employed, i.e., Phase Symmetry – proposed by Kovesi [62] - for ridge

detection.

Phase symmetry based detection captures the major axis of symmetry of a feature at some specified spatial scale. Even signals have real (and even) Fourier transforms. Odd signals have imaginary (and odd) Fourier transforms. Generally, signals will have complex Fourier transforms, since they are neither perfectly odd nor perfectly even. The resultant phase values reflect their degree of symmetry.

Kovesi [62] gives the following analysis for 1-D phase symmetry:

For phase symmetry, a signal is convolved with a wavelet pair of band-pass quadrature filters (an odd filter and an even filter). If I is the signal and M_e^n and M_o^n denote the cosine and sine wavelets at a scale n the responses of each quadrature filter pair forms a vector,

$$[en(x); on(x)] = [\, I(x) * Men; I(x) * Mon \,], \qquad (3.1)$$

The amplitude of the transform at scale n is then given by

$$A_n(x) = (e_n(x)^2 + o_n(x)^2)^{0.5} \qquad (3.2)$$

and the phase is given by

$$\Phi_n(x) = atan2(e_n(x); o_n(x)) \qquad (3.3)$$

The absolute value of the even-symmetric filter outputs is large and that of the odd-symmetric filter outputs is small at a point of symmetry. Symmetry would thus be proportional to the difference of these absolute values. Likewise, the difference of the absolute values of the sine and cosine of the phase angle is proportional to Phase

Symmetry. Symmetry is thus given as:

$$\mathrm{Sym(x)} = \frac{\sum_{n} \lfloor A_n(\mathrm{x}) \lfloor |\cos(\phi_n(\mathrm{x}))| - |\sin(\phi_n(\mathrm{x}))| \rfloor - T \rfloor}{\sum_{n} A_n(\mathrm{x}) + \varepsilon} = \frac{\sum_{n} \lfloor \lfloor |e_n(\mathrm{x})| - |o_n(\mathrm{x})| \rfloor - T \rfloor}{\sum_{n} A_n(\mathrm{x}) + \varepsilon} \tag{3.4}$$

T is a noise threshold and is given by: $T = \mu + k\sigma$, where μ = mean of the local energy distribution, σ = the standard deviation of the same and k is a scalar factor.

For 2-D images, in addition to scales, a number of separate orientations (r) figure in the feature detection. The filters with these orientations are defined by spreading a Log-Gabor function into two dimensions. Masking a radial Log-Gabor function with an angular Gaussian tuned to ϕ_0 gives a filter tuned to a particular orientation ϕ_0. The frequency band to which the filter responds is controlled by the radial component, and the orientation to which the filter responds is controlled by the angular component. The resulting two components are then combined into a 2-D Log-Gabor function as in Equation. (3.5):

$$\mathrm{G}(\omega, \phi) = \exp[-\frac{(\log(\omega / \omega_0))^2}{2(\log(\kappa / \omega_0))^2} + \frac{(\phi - \phi_0)^2}{2\sigma_\phi}] \tag{3.5}$$

The angular bandwidth is given by:

$$\Delta\Omega = 2 \times \sigma_\phi \sqrt{2 \times \log 2} \tag{3.6}$$

where σ_ϕ = $\Delta\phi / s$ and $\Delta\Omega$ is the angular bandwidth. $\Delta\phi$ is the angular separation between neighboring orientations and is defined as $\Delta\phi = 180° / r$, where r denotes the total number of orientations used. The parameter s controls the angular overlap of the

filter's transfer function.

Over a number of scales (m) and at different orientations (r), a 2-D phase symmetry measure can then defined as in eqn (3.7):

$$PS(x,y) = \frac{\sum_r \sum_m \lfloor \lfloor |e_{rm}(x,y)| - |o_{rm}(x,y)| \rfloor - T_r \rfloor}{\sum_r \sum_m \sqrt{e_{rm^2}(x,y) + o_{rm^2}(x,y)} + \varepsilon} \tag{3.7}$$

T_r is the orientation-dependent noise threshold and is analogous to the 1-D scenario.

For efficient ridge detection for bone surfaces, we use the following parameter values, as suggested by Hacihaliloglu et al. [13]: $m = 2;\ r = 6;\ s = 1.2;\ \Delta\Omega = 25^o;\ k = 8$

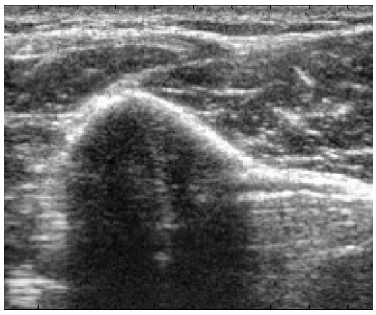

(a) Central slice of Volume1

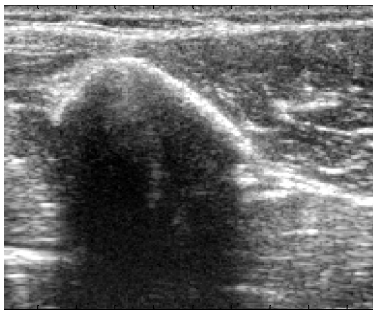

(b) Central slice of Volume2

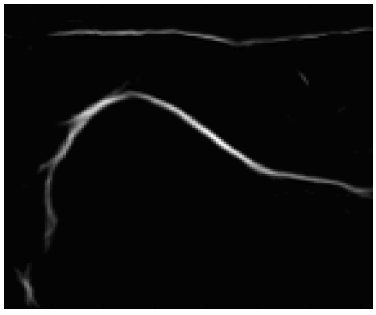

(c) Phase symmetry image of (a)

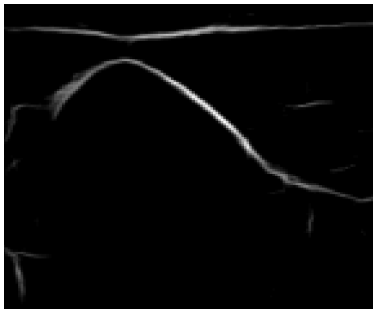

(d) Phase symmetry image of (b)

Figure 3-1: Central slices of the two volumes (in vivo radius) to be registered and their phase symmetry images

Once the image has been cleaned and the bone surface made prominent (Fig. 3-1), the Harris corner detection algorithm – as proposed by Harris and Stephens [63], is used to obtain features in both slices. The Harris corner detector works as follows: Given an image I, consider two windows, one over the area *(a,b)* and the other, shifted by (*x*,*y*). The SSD (sum of squared differences) for these two windows is:

$$SSD(x,y) = \sum_a \sum_b w(a,b)(I(a,b) - I(a+x,b+y))^2 \tag{3.5}$$

$I(a + x,b + y)$ can be approximately written as:

$$I(a+x,b+y) \approx I(a,b) + I_x(a,b)x + I_y(a,b)y \tag{3.6}$$

where I_x and I_y are partial derivatives of I

Thus, SSD is given as:

$$SSD(x,y) \approx (x \quad y) A \begin{pmatrix} x \\ y \end{pmatrix}, \tag{3.7}$$

where

$$A = \sum \sum w(a,b) \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} \tag{3.8}$$

The eigenvalues of A are calculated and the following conclusions are made: If $\lambda_1 \approx 0$ and $\lambda_2 \approx 0$, then pixel (*x*,*y*) has no features of interest. If $\lambda_1 \approx 0$ and λ_2 has some large positive value, then pixel (*x*,*y*) is on an edge. If λ_1 and λ_2 have large positive values, then pixel (*x*,*y*) is a corner.

For every corner point in one slice (Slice 1), the corresponding point in the other slice (Slice 2) is taken by correlating a window of a certain radius (in our case 29×29) around

the point with similar windows constructed around each corner point in Slice 2. Then the points which correlate maximally are considered corresponding, or matching, point pairs. For increased robustness, the procedure is performed again, but this time, for every corner point in Slice 2, a matching corner point is found in Slice 1. Only matching points that correlate maximally in both directions are considered matching point pairs. A RANSAC algorithm [64] is then used to remove outliers and find reliable correspondences between the feature points of the two slices.

From these matching points a 'salient' point pair is to be selected and the sub-volumes of the original volumes around the points in this pair are to be registered. The salient point pair is selected as follows: For each matching point pair p, generate windows of a certain size (in our case 29×29) around the points. Find the correlation C_p and the mean standard deviation $\sigma_p = (\sigma_{p1} + \sigma_{p2})/2$ of the windows. Then the point pair to be selected (Fig. 3-2) is given by

$$P = \max_p(\alpha C_p + \beta\sigma_p) \tag{3.9}$$

where α and β are weights given to each parameter.

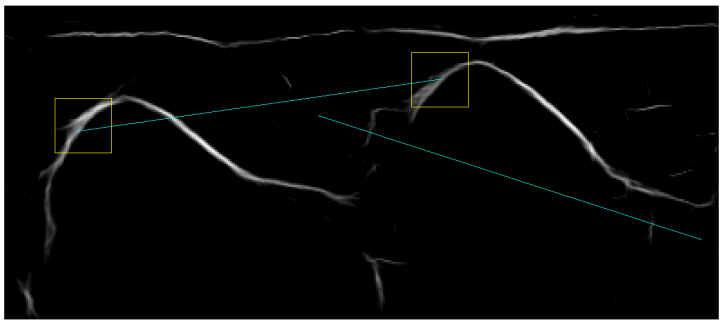

Figure 3-2: Salient point selection. The corresponding points shown in the two sub-windows above are the points around which the sub-volumes to be registered are formed.

Once the salient point pair is identified, sub-volumes (size 17×17×17) around the points are generated in both volumes and rigidly registered using volumetric means. In this case, we used normalized cross correlation and a gradient descent optimizer to register the two volumes.

Once the parameters for the registration of the two volumes are obtained, they are used to obtain the transformation parameters that are to be used on the whole volume.

3.3 *Experiments and Results*

The method was used for two different applications. The first application was registration of volumes for volume stitching. The second application was registration of volumes for tracking incidental intra-surgical bone motion during CAOS. The images were scanned using a GE Voluson 760 US scanner. The algorithm was run in MATLAB (everything apart from the sub volume registration, where ITK (C++) was used) on a P4 system with 2GB of RAM. Registration parameters were obtained, on average, in 4.5 seconds.

For the first application, the method was tested on two real (in vivo), two artificial human radial bones (Sawbone Model #1018-3, Sawbones Inc., Vashon, WA, USA), a real, in-vitro bovine femur, an artificial fetus (CIRS, Inc., Norfolk, VA, USA) and an artificial pelvis (Sawbone Model #1301-96, Sawbones Inc., Vashon, WA, USA). In the case of the artificial bones (radius and pelvis), fiducials were attached to the bones at regular intervals for quantitative assessment. The distance between the fiducials in the stitched volumes was compared to the actual distance between them on the physical bones and the accuracy of the stitching was subsequently computed (Equation 3-10).

$$Error = \left| Dist_{True} - Dist_{Stitched} \right| \quad (3\text{-}10)$$

$Dist_{True}$ = *Mean actual distance between fiducials on the physical objects*

$Dist_{Stitched}$ = *Mean distance between fiducials in the stitched volumes*

The results of the algorithm on the in vivo, real radius scans were only assessed qualitatively by visual inspection since there was no ground truth data (such as CT scans). In lieu of a tracking system (the one we had was experiencing technical difficulties which remained unresolved at the time of writing this thesis), the volumes were obtained by moving the probe along a straight line for known distances (0.5 cm) between scans and then these distances were used as tracking parameters. The quantitative results are shown in Table 3-1 and the qualitative results for the artificial radius and fetus scans are shown in Figure 3-3. The qualitative results for one of the in vivo scans are shown in Figure 3-4. Visual inspection showed that both in vivo scans had correctly registered.

For the second application, radius tracking was performed on 3 radii (3 volunteers). A probe was attached to the right forearm of each volunteer and the volunteer was told to slowly bend the arm. For each volunteer, 3 volumes were obtained at periodic intervals. Each successive volume was registered to the previous one using the aforementioned method. From the tracking and registration parameters, the bone motion could be inferred. Here, however, we only demonstrate the bone registration - which is the thrust of this method - and not the inference of the bone motion there from – partly due to the absence of an actual tracking mechanism. The image acquisition and processing hardware was the same as for the first application. Registration parameters were obtained,

on average, in 5 seconds. Again, due to an absence of ground truth for in vivo data, the registration assessment is only visual.

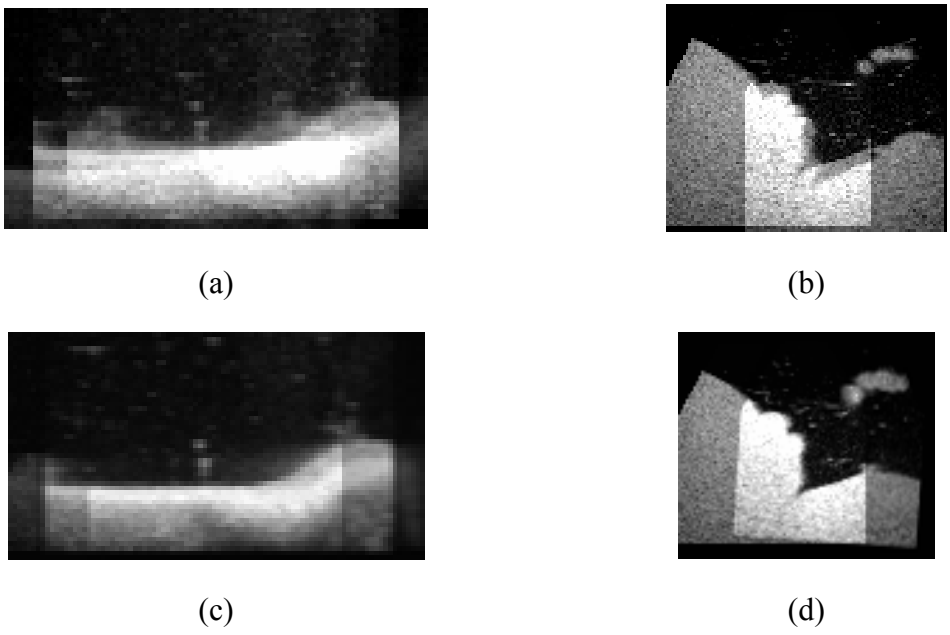

Figure 3-3: Qualitative assessment of the registration of the artificial radius (left) and the artificial fetus (right) for volume stitching. (a) and (b) indicate stitching before registration, using only the tracking. (c) and (d) show the results of stitching after tracking and registration.

Table 3-1: Mean error measurements for stitching bones

Bone	Error before registration (tracking only)	Error after tracking and registration
Artificial Radius 1	3.15mm	0.59mm
Artificial Radius 2	1.47mm	0.37mm
Bovine Femur	0.84mm	0.28mm
Artificial Pelvis	0.25mm	0.05mm

Figure 3-5 shows the qualitative result for one radius. Visual inspection showed that, for each volunteer, all 3 in vivo scans had correctly registered to each other.

(a) (b)

Figure 3-4: Qualitative assessment of the registration for the in vivo radial scans for volume stitching. Here, (a) indicates stitching before registration, using only the tracking and (b) shows the results of stitching after tracking and registration. The dark region is actually the shadow region below the bone – it was shown rather than the bone itself since the discontinuity at the overlap in (a) (circled) is more clearly highlighted in this region. Note the smooth transition in (b).

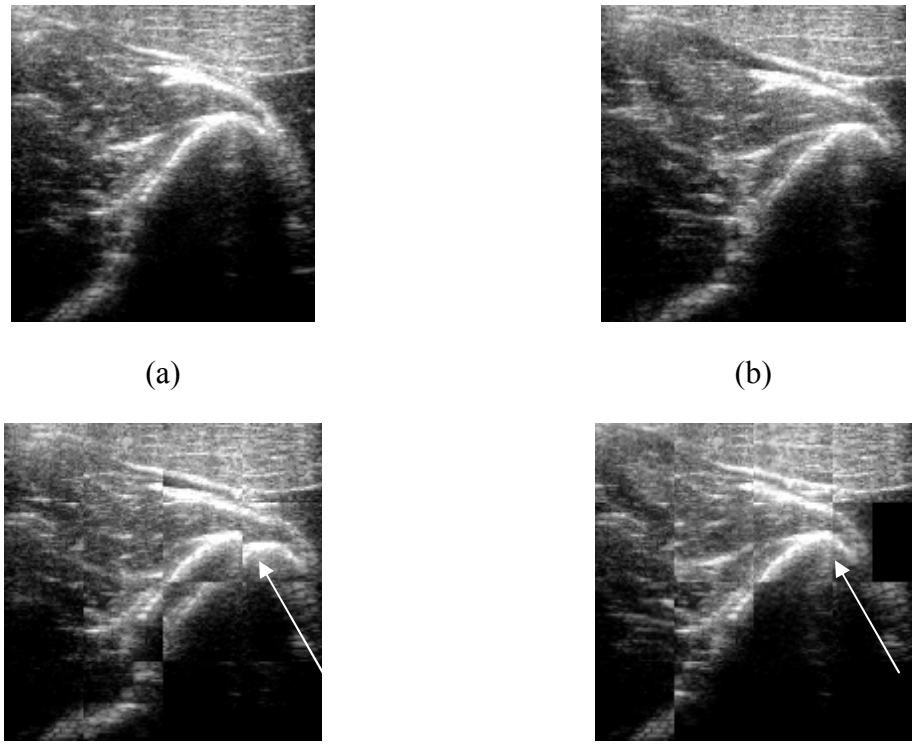

(a) (b)

(c) (d)

Figure 3-5: Qualitative assessment of the registration for the in vivo radial scans for bone tracking. Here, (a) and (b) are the central slices of the two successive volumes, (c) is a checkerboard image of the central slices of the two successive volumes before registration, using only the tracking and (d) is a checkerboard image of the central slices of the two successive volumes after tracking and registration. Note the discontinuities in (c) that have been removed in (d) (circled).

3.4 *Conclusion and Discussion*

In this chapter, we have proposed and tested probe tracking, Phase Symmetry, Harris corner detection and sub-volume registration to obtain fast and accurate registration for volume stitching. Even though the registration is rigid, given the application area – orthopaedic applications – the assumption of rigidity of the tissues of interest is reasonable. Moreover, the method is very fast – taking only 4.5 seconds to obtain the registration parameters for two US volumes – and hence can be used in fields such as CAOS (computer assisted orthopaedic surgery) where speed is of importance. In non orthopaedic applications, our method could serve as an essential pre-registration step of calculating a global alignment prior to deformable. Future work will involve implementing the actual probe tracking mechanism and then validating the algorithm on various orthopaedic – and non orthopaedic data – with ground truth as well as porting the method to a faster platform such as C or C++ which could make the method fast enough to be used in real time applications.

Chapter 4 deals with the registration of the stitched US data set with the MRI volume.

4 Robust, Accurate 3D Ultrasound Volume to MRI Volume Registration Using Phase Symmetry Based Surface Separation and Registration

4.1 *Introduction*

This chapter deals with the registration of the US volumes of the bones that are obtained during the testing phase of the motion analysis to the MRI volume of the bones taken preoperatively. This is done in order to show the movement of the bones (captured with US) to the researchers using high resolution MRI data, making it much easier for the researchers to analyze the motion.

Most of the work in the literature that deals with US-CT (or US-MRI) registration has been in the field of surgical navigation procedures. Navigational procedures have become extremely important in modern day surgery. Navigation systems based on preoperatively obtained CT and/or MRI data improve the reliability and safety of minimal invasive procedures. For these systems to provide reliable performance, however, it is crucial that the preoperative datasets should be registered accurately to the volumes being obtained in the surgical room. In many surgical disciplines (orthopedics, neurosurgery, traumatology, etc.), the accurate registration of bones is of main interest.

Common methods used in the field are based on paired point registration using anatomical landmarks or fiducia1 markers. If the landmarks are anatomical, to ensure an accurate registration, the number of landmarks needed is large. Alternatively, fiducials

may be used, resulting in a smaller number of markers but this adds both to the time and invasiveness of the surgery.

Another option is using volumetric registration. In this case, complete anatomical structures can be used for registration (mostly surfaces), thus increasing the accuracy. The usage of intraoperative CT or MRI has been proposed and implemented [65], [66], but these systems have major drawbacks with respect to intraoperative applicability, costs and radiation exposure (CT).

In the light of these drawbacks, ultrasound seems to be an ideal intra-operative imaging modality. US is cheap, easy, non-ionizing and real-time, and could be used to provide intra-operative data that could be registered accurately with preoperative data.

Ultrasound, however, comes with its own set of problems. It has high noise levels, and the imaging, especially where bones are concerned, is very limited. This problem is intrinsic, and is due to physical properties of the tissues involved [67]. US waves are almost entirely reflected at from the surface of the bone, meaning that nothing beyond the bone surface gets imaged by the US. Also, given the highly specular nature of the reflections, the direction of sound propagation has to be almost or fully orthogonal to the bone surface if the image is to show up clearly.

There have been some approaches for the registration of bone structures in CT and ultrasound data sets in the literature. Most of them address long bones [68], pelvis [69], [70] and spine [71], [72], [73]. In this case, volume–volume registration methods, with traditional similarity measurements do not work well for US-CT registration. This is due

to the fact that the US shows only the bone surface and the CT shows the whole bone, resulting is two vastly different looking anatomical images. Thus, all these approaches use surface–surface registration methods. For good surface registration, however, surfaces must be extracted. Doing so in an automated manner is fairly straightforward in the case of bones in CT, but for US, this has traditionally been a challenge. Many of the papers mentioned above use artificial bones for their trials, without any overlying tissue, making segmentation of US bone surfaces easier. However, this may not be a valid way of doing things when real data is being registered.

To overcome this problem, Brendel et al [17], [18] propose a method where they do away with the need to segment the US bone. From the CT images of the bone, they obtain the bone surface which would be imaged by the US. Then they register this with the entire US volume, with the registration metric being the sum of the values of the US region overlapping the CT bone surface. They report good accuracy and high speed. However, their lumbar spine data [17] sets are artificial. Moreover, the highly distinctive shape of the vertebral bones makes it reasonable to assume that the metric described will be optimized only at the correct location. In the smoother, long bones of the knee, this may not be the case, as the other bright structures may be significant confounding factors. In [18], MRI of the knee is registered to US. However, the surface points from the MRI bone surface had to be manually determined. For both the US-CT and US-MRI registration the authors increase the contrast of the US bone surface using deep gain compensation. They indicate that the contrast increase is very important for proper registration.

We aim, in this part of the proposal, to provide a fully automated method for US-MRI volume registration for motion analysis. This involves two phases – a) registering a preoperative MRI scan with a preoperative, stitched US volume and b) registering the preoperative, registered-to- MRI US volume to "intra-operative" (taken here to mean data obtained during the motion capture) US volumes. The first phase involves quick and accurate extraction of the bone surfaces to be registered from BOTH the US and the MRI datasets and then registering them using surface registration methods. The second involves registering the two US volumes using the method used for volume stitching in Chapter 3.

(N.B.: Though the goal is to use MRI as the high resolution modality, the method has been tried out with both CT and MRI volumes as the high resolution modalities. Hence the methods section often refers to CT volumes.)

4.2 *Methods*

The first phase involves registering the stitched US volume (obtained at the end of Chapter 3) with a preoperative CT/MRI scan. [The CT/MRI would be obtained of the knee joint with a certain amount of both the femur and the tibia. The motion testing phase would involve US probes to be attached to the upper and lower leg of the subject.] Stitched-US is used to generate a greater bone area over which registration can take place. This is done to increase the accuracy of the US-CT/MRI registration. To do this, we extract the CT/MRI and US bone surfaces. To extract the relevant CT surface, we use the method used by [17], i.e., we threshold the scan to remove non bone tissue, cast rays in the direction in which the US probe will be scanning and consider the bone pels closest

to the rays in that direction. If MRI is used, the bone is segmented out using the method suggested in chapter 2 and the relevant surface is extracted by the ray casting method noted above (Fig. 4-1).

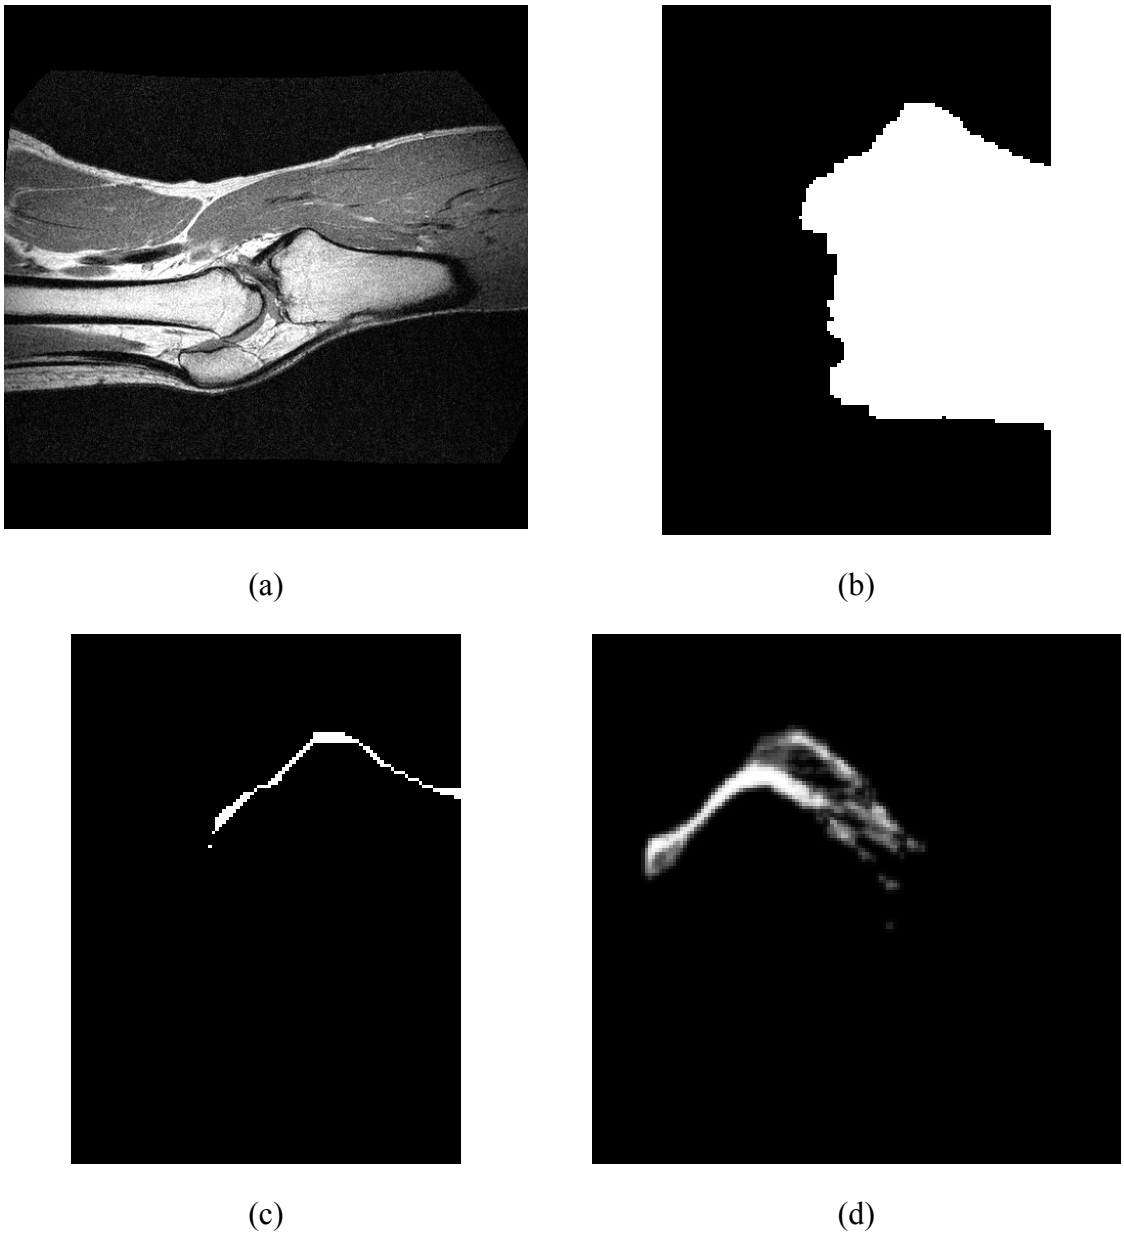

Figure 4-1: Tibial bone surface extraction from MRI scan of knee (a), segmented tibia (scaled up) (b), delineated bone surface as would be seen by US probe (c) and 3D surface (d)

To obtain the US bone surface, we perform the following steps (these steps assume that the images are taken in a transverse manner:

1. Obtain the phase symmetry images for the US scans as described in Chapter 3 (Fig. 4-2).

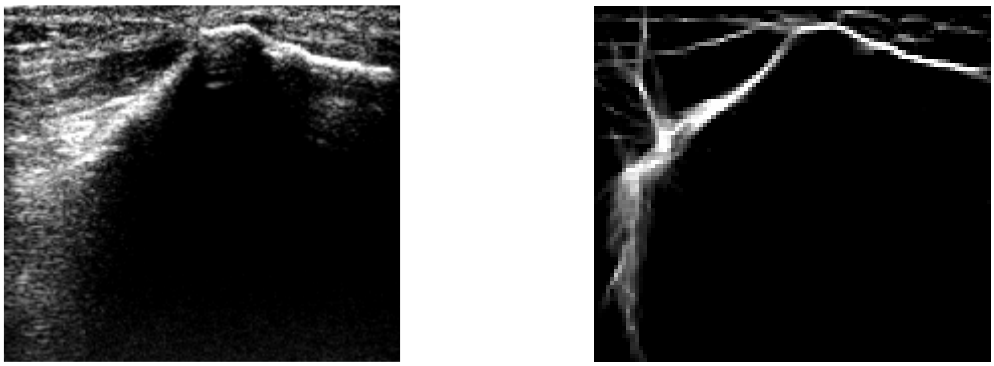

(a) (b)

Figure 4-2: Using phase symmetry to clean the US volume of the tibia: (a) US scan of tibia, (b) Phase Symmetry filtered output for (a)

2. Remove all small regions whose area is under a certain threshold.
3. We know that the US will not image any area beyond the bone surface. Hence the region under the bones will be a clean shadow, which will be made cleaner still by the phase-symmetry filtering. Hence, we cast rays along each line (row/column) from the side opposite to the probe towards the probe, and stop at the first high intensity pel. Following this, we perform morphological closing and then obtain the largest connected component in the resulting image. This will contain the bone surface – without any high intensity noise above the bone. The whole process is illustrated in Fig. 4-3.

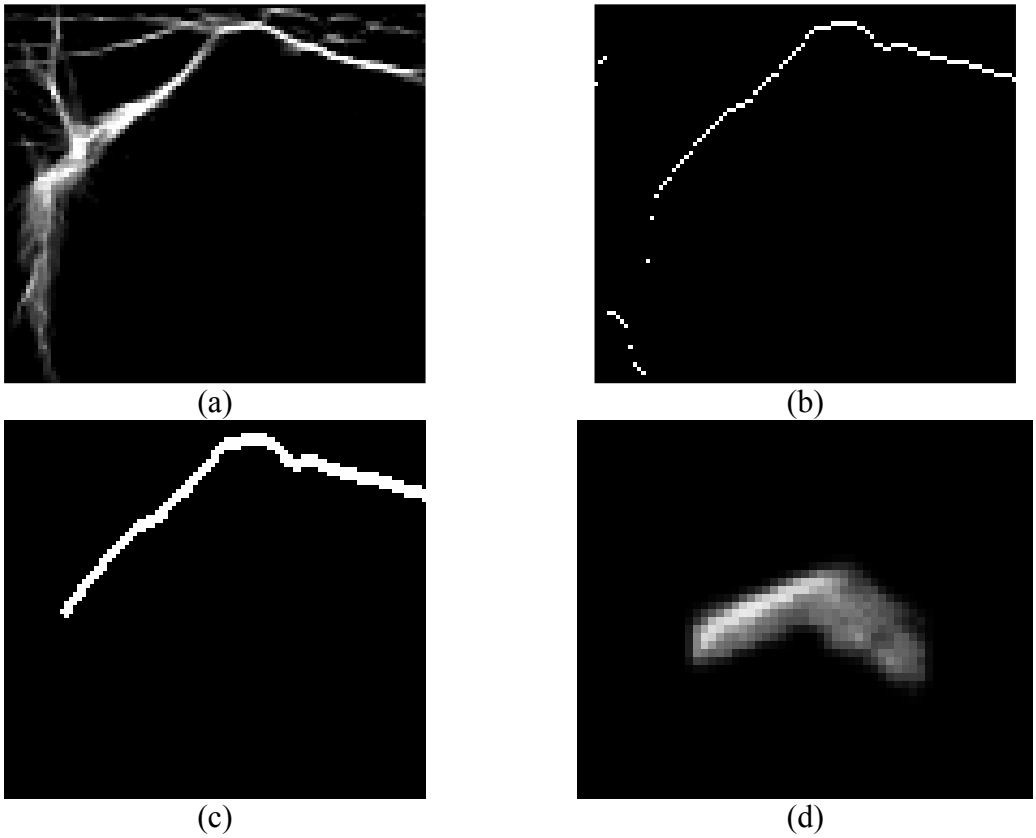

Figure 4-3: Bone surface extraction from phase symmetry volume of tibia: (a) Phase symmetry image of US scan of human tibia. (b) After ray casting from below to remove artifacts above the bone surface. (c) After closing and largest connected component selection. (d) Volume rendering of (c).

Once the US and CT/MRI bone surfaces are obtained and resampled to match resolutions, registration is performed as follows:

1. Unlike the stitching scenario, the assumption that the central slices contain almost the same features does not hold here. However, we can make the assumption that the CT/MRI and the US volumes were both captured in the same orientation (e.g. sagittal). Therefore, we can assume that the central slice (along a given orientation) of one volume will contain features that are highly similar to those in *one* of the slices of the other volume (along the same orientation). Using this we

can obtain a coarse ‘pre-registration’ as follows:

2. Hence, we take the central slice of the CT/MRI volume and register them to each slice of the US volume using feature based registration. Specifically, for each US slice, we first obtain the bone surface from the slice using the method described above and then obtain corner points from both the CT/MRI slice and the US slice using the Harris corner detector. For every corner point in the CT/MRI slice, the corresponding point in the US slice is taken by correlating a window of a certain radius (in our case 29×29) around the point with similar windows constructed around each corner point in US slice. Then the points which correlate maximally are considered corresponding, or matching, point pairs. For increased robustness, the procedure is performed again, but this time, for every corner point in US slice, a matching corner point is found in CT/MRI slice. Only matching points that correlate maximally in both directions are considered matching point pairs. A RANSAC algorithm [64] is then used to remove outliers and find reliable correspondences between the feature points of the two slices. Using these reliable point pairs, we calculate the registration parameters to align the two slices. Upon registration, the correlation between the two (registered) slices is calculated.

3. The registration parameters that register the US slice that correlates most highly with the CT/MRI slice give the in plane rotation as well as x and y translation parameters. The z translation can readily be calculated by finding the difference between the slice numbers of the two slices.

4. Transforming the US volume with these parameters gives a volume that is coarsely registered with the CT/MRI volume.

5. The coarsely registered volumes are then registered more accurately using volumetric registration, with sum of differences (SAD) as a registration metric and Powell's multidimensional search as the optimization method.

Once the registration between the stitched US volume and CT/MRI is done preoperatively, the registration between the intraoperative and preoperative US volumes can be done using the registration/stitching method detailed in Chapter 3.

4.3 *Results*

Since real life CT/MRI data is difficult to obtain, especially matched with US data, we have so far been able to try the method out on only one US-MRI knee dataset pair and one US-3D Fluoroscopy pair. Qualitative results of the US-MRI registration are shown in Figure 4-4. The result of the US-3D Fluoro registration is shown in Figure 4-5. It should be noted that attempting to directly register the CT/MRI and US bone surfaces (without the pre-registration described above) using volume–volume registration methods, with mutual information and cross correlation, resulted in very poor registrations (not shown).

While only data from one knee (real) and one radius (artificial) were available for testing the method (qualitatively), the initial results seem promising. The US volumes seem to be correctly registered to the corresponding region in the MRI/3D Fluoro volumes. Further testing is necessary to validate the actual usefulness and effectiveness of the proposed

method.

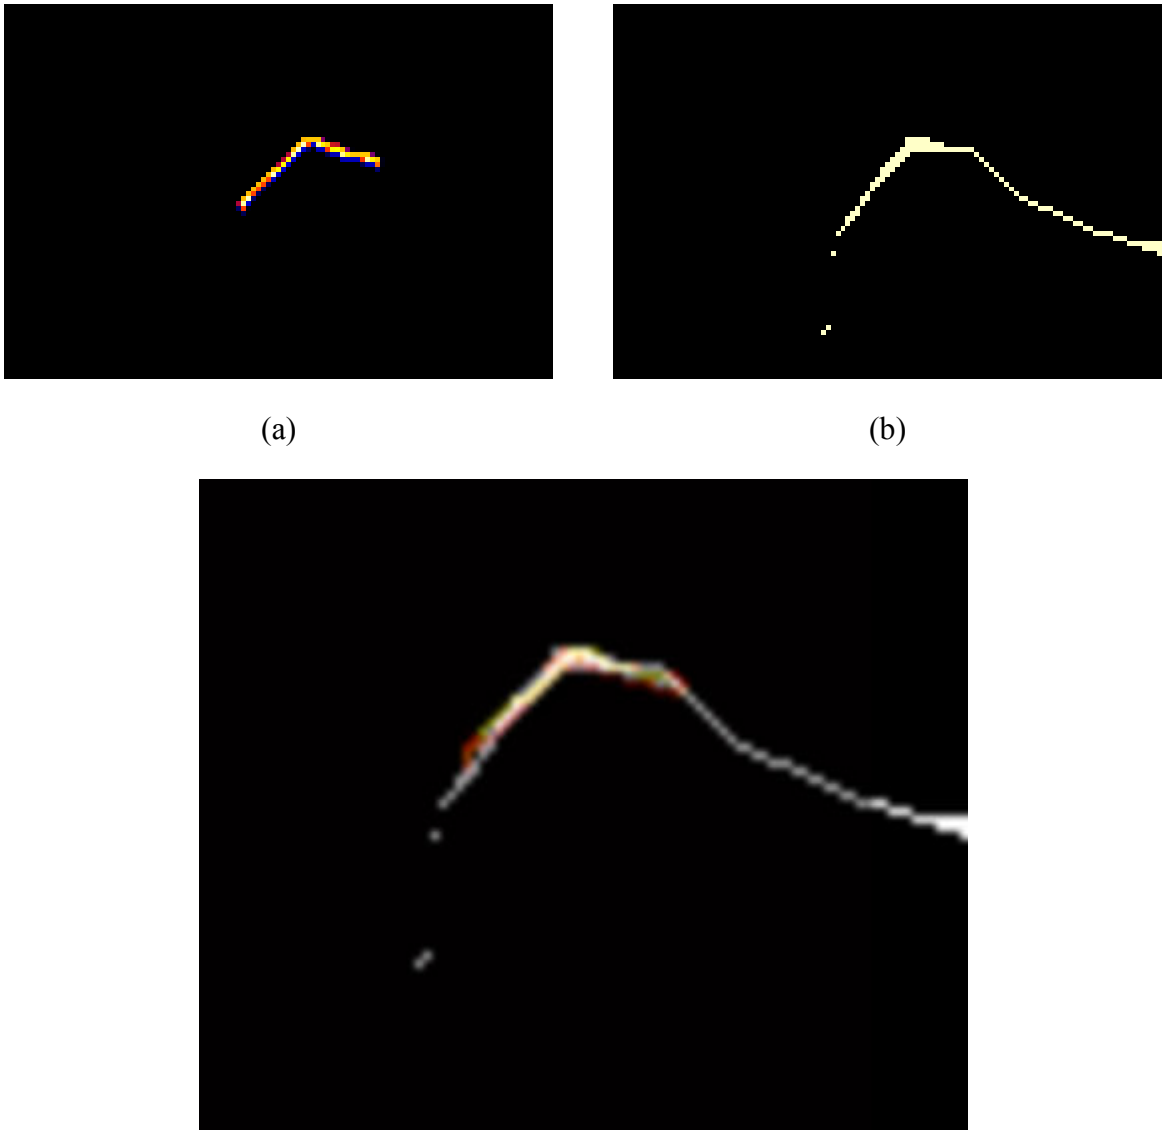

Figure 4-4: Qualitative result of US-MRI registration: (a) cross section of bone surface extracted from US dataset, (b) cross section of bone surface extracted from MRI dataset, (c) cross sectional overlay of registered US/MRI bone surfaces

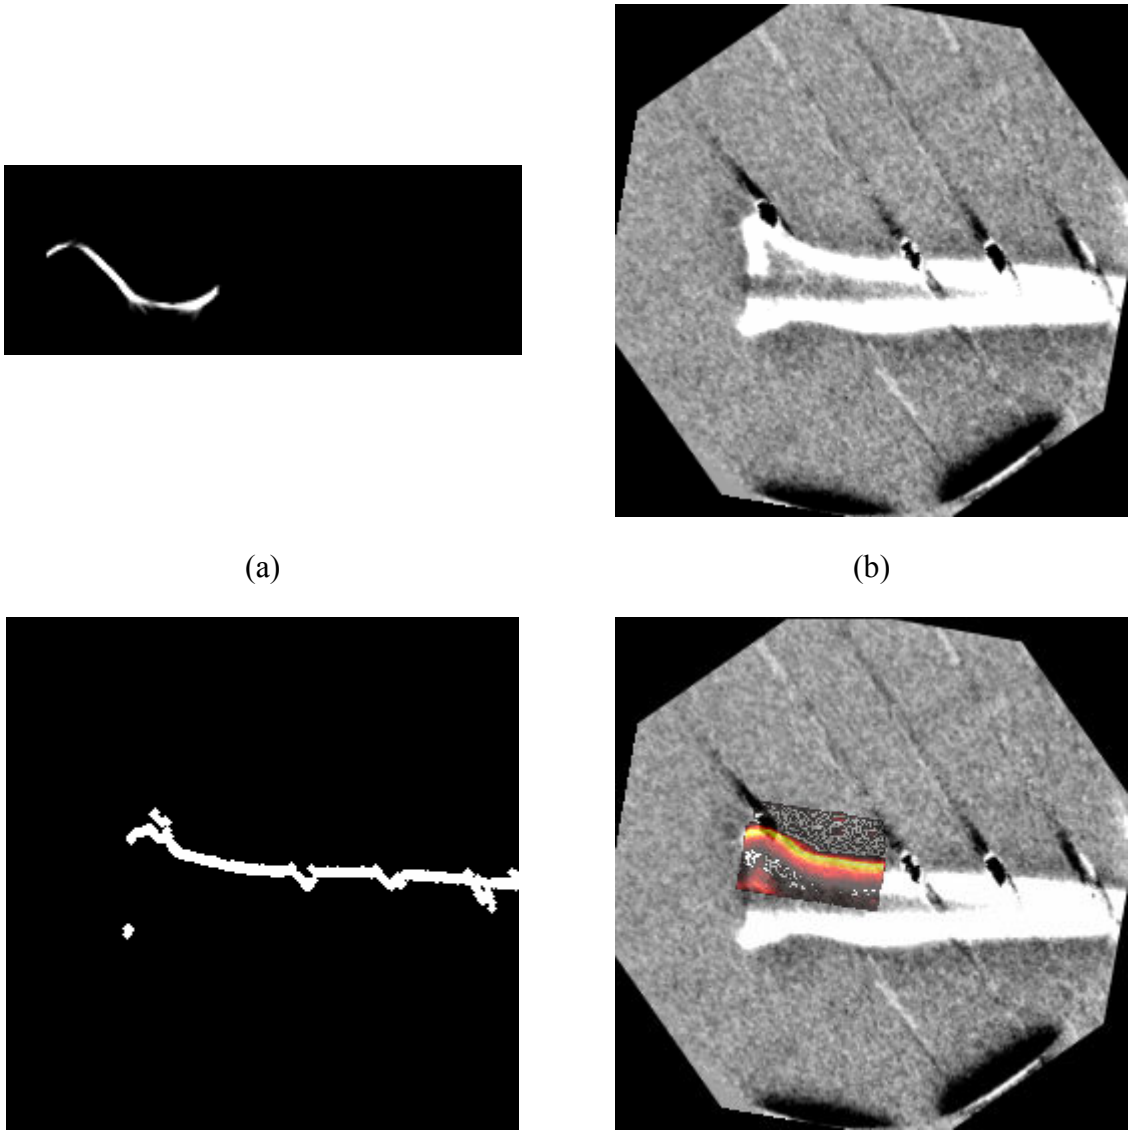

Figure 4-5: Qualitative result of US-3D fluoroscopy registration: (a) cross section of bone surface extracted from US dataset, (b) cross section of bone from 3D- fluoroscopy dataset, (c) cross section of bone extracted from 3D- fluoroscopy dataset, (d) cross sectional overlay of registered US/3D- fluoroscopy bone surfaces

4.4 *Conclusions*

We have proposed, in this Chapter, a method for the registration of US-CT/MRI data. This allows us to register preoperatively a US volume to a CT/MRI volume. Intra-operative US scans can then be quickly registered to the preoperative US volume, thus registering them to the CT/MRI volume. Thus, we can obtain a feasible motion analysis mechanism based on preoperative US and CT/MRI volumes and intraoperative US volumes.

5 Discussions and Conclusions

5.1 *Introduction*

This chapter is intended to serve as a summary of the contribution of the thesis. It is also used to point out aspects of the work done that could be further investigated, improved or applied in different directions.

5.1.1 Thesis Contributions

The primary goal of this project was to provide a robust and accurate motion analysis framework in order to attain the following main objective: to help improve the analysis of the motion of the joints of people with joint diseases by orthopaedic researchers.

This overall objective translated into the following sub goals:

1. To develop a robust segmentation framework for bone MRI using multi contrast MR imaging:

In this project we used multi contrast imaging (MCI) for facilitating improved segmentation of bone structures in MRI based on a simple yet efficient and robust method for increasing image contrast between bone and surrounding cartilage tissue.

We obtained MR imaging at different tissue contrasts that were then combined to enhance the contrast at bone-cartilage interface (BCI). The method was subsequently shown to produce sustained and significant improvement in bone segmentation when both automated and semi-automated segmentation methods were used.

2. Register multiple shifted US volumes together and stitch them to generate a volume with a large field of view:

We proposed a phase symmetry based method that would quickly and accurately register two US bone volumes together. We showed that the method produced very fast results that would, upon porting to a faster language and code optimization, lead to near real time or real time registration.

3. Register US bone volume to MRI volume:

Finally, we developed a method for registration of data from US and MRI together. We used phase symmetry and post processing to accurately obtain the organ surface in US and register it to the corresponding surface in MRI. This method would require considerable additional testing and validation, but the initial test makes it seem promising.

5.1.2 Future Work

The work done so far involves two distinctly separate parts – the enhancement of the bone images in MRI and the subsequent improvement in their segmentation, and the registration of US volumes (used to capture subject motion) to static, high resolution CT/MRI data in order to present the motion of the joint in high resolution 3D. Future work would initially involve combining these two parts. While we have developed procedures for combining the motion data from the US and the high resolution, static CT/MRI data, considerably more testing would be needed to validate the methods. Registration of the US and MRI would immediately be carried out near the joint surfaces, where the cortical bone thins out and hence, the surfaces obtained from both

the MRI and the US data correspond. However, this could be inconvenient for many applications. Hence later on, registration of bone data taken from the central parts of the bone will have to be done. This will involve two approaches:

1. It will involve estimating the cortical bone in the MRI scans (since what is seen in the scans and segmented is the inner bone) and then registering that to the US data (which sees the cortical bone surface).
2. Another approach would be to use new, ultra short echo time sequences being developed in MRI that allow for the visualization of cortical bone. These scans would allow us to visualize the cortical bones. Then, using the methodology of Chapter 2, we could perform multi contrast difference imaging using these scans and conventional scans (say T2 weighted scans) to separate out just the cortical bone. If that could be done, the registration step would be the same would not require estimation of the cortical surface in MRI and this would make registration a lot more robust.

Future work would also involve applying the methods developed in the course of the project to other areas of the body. For instance, there is a lot of research being conducted on surgical intervention guidance for kidney and prostate biopsies and surgeries. The methods developed above could be adapted to those fields as well.

Finally, once validations on all the proper data have been performed, porting the code to C++ and optimizing it for speed should be done.

Bibliography

[1] http://www.cdc.gov/nchs/fastats/arthrits.htm
[2] http://www.arthroscopy.com/sp13008.htm
[3] Lavallee, S., Plaskos, C., Schers, J., Troccaz, J.: Ultrasonic Bone Motion Tracking System, http://www.faqs.org/patents/app/20090018445.
[4] Chen, T. K., Abolmaesumi, P., Pichora, D., Ellis, R.: A system for ultrasound-guided computer-assisted orthopaedic surgery. Computer Aided Surgery. 10(5):281-292, (2005).
[5] Meyer, C., Boes, J., Kim, B., Bland, P., LeCarpentier, G., Fowlkes, J., Roubidoux, M., Carson, P.: Semiautomatic registration of volumetric ultrasound scans. Ultrasound Med. Biol., vol. 25, no.3, pp. 339–347, (1999).
[6] Krücker, J., LeCarpentier, G., Fowlkes, J., Carson, P.: Rapid elastic image registration for 3-D ultrasound. IEEE Trans. Med. Imag., vol. 21, pp. 1384–1394, (2002).
[7] Pratikakis, I., Barillot, C., Hellier, P., Mémin, E.: Robust multiscale deformable registration of 3-D ultrasound images. Int. J. Image Graphics, to be published.
[8] Poon, T., Rohling, R.: Three-dimensional extended field-of-view ultrasound. Ultrasound in Medicine and Biology 32(3), 357–369 (2003).
[9] Wachinger, C., Wein, W., Navab, N.: Three-dimensional ultrasound mosaicing. In: MICCAI 2007, Part II. LNCS, vol. 4792, pp. 327–335. Springer, Heidelberg (2007).
[10] Dong Ni, Yingge Qu, Xuan Yang, Yim-Pan Chui, Tien-Tsin Wong, Simon S. M. Ho, Pheng-Ann Heng: Volumetric Ultrasound Panorama Based on 3D SIFT. MICCAI 2008, pp 52-60, (2008).
[11] Gee, A., Treece, G., Prager, R., Cash, C., Berman, L.: Rapid registration for wide field of view freehand three-dimensional ultrasound. IEEE Trans. Med. Imaging 22(11), 1344–1357 (2003).
[12] Pennec, X., Cachier, P., Ayache, N.: Tracking brain deformations in time-sequences of 3-D us images. Pattern Recogn. Lett., vol. 24, no. 4–5, pp. 801–813, (2003).
[13] Hacihaliloglu I., Abugharbieh R., Hodgson A., Rohling R.. "Bone Surface Localization in Ultrasound Using Image Phase Based Features". Ultrasound in Medicine and Biology (UMB), 2009, volume 35, number 9, pp. 1475-1487
[14] Lowe, D.: Distinctive image features from scale-invariant keypoints. International Journal of Computer Vision 60(2), 91-110, (2004).
[15] Chen, J., & Tian, J.: Rapid Multi-modality preRegistration based on SIFT descriptor. Engineering in Medicine and Biology Society, EMBS '06. 28th Annual International Conference of the IEEE, (2006).
[16] Angelsen B., Ultrasound Imaging. , Emantec, Trondheim, Norway (2000).
[17] Brendel B., Winter S., Rick A., Stockheim M., Ermert H, Registration of 3D CT and ultrasound datasets of the spine using bone structures, Computer Aided Surgery 7 (3) (2002) 146– 155.

[18] Dekomien C.; Hold S. et al.: Registration of intraoperative 3D ultrasound with preoperative MRI data for navigated surgery – first results at the knee. In Proc CAOS, 2007

[19] Maes F., Collignon A., and Vandermeulen D.. Multimodality image registration by maximization of mutual information. IEEE Trans. on Med. Imag., 16(2):181–187, 1997.

[20] Studholme C., Hawkes D.J., and Hill D.L.G.. An overlap invariant entropy measure of 3d medical image alignment. Pattern Recognition, 32(11):71–86, 1999.

[21] Roche A., Malandain G., Ayache N., and Pennec X.. Multimodal image registration by maximization of the correlation ratio. INRIA Research Report, (3378), 1998.

[22] Woods R., Mazziotta J., and Cherry S.. MRI-PET registration with automated algorithm. Journal of Computer Assistant Tomography, 17(4):536–546, 1993.

[23] Viola P. and Wells W. III, Alignment by maximization of mutual information, (1995), pp. 16–23. IEEE, 1995.

[24] Huang X.S., Hill N.A., Ren J., and Peters T.M.. Dynamic 3d ultrasound and MR image registration of the beating heart. In Medical Image Computing and Computer Assisted Intervention , MICCAI 2005, volume 3750 of Lecture Notes in Computer Science, pages 171–178, 2005.

[25] Penney G., Blackall J., Hamady M., Sabharwal T., Adam A., and Hawkes D.J.. Registration of freehand 3d ultrasound and magnetic resonance liver images. Medical Image Analysis, 8(1):81–91, 2004.

[26] Leroy A., Mozer P., Payan Y., and Troccaz J.. Rigid registration of freehand 3d ultrasound and CT-scan kidney images. In Medical Image Computing and Computer Assisted Intervention , MICCAI 2005, volume 3216 of Lecture Notes in Computer Science, pages 837–844. Springer, Sep 2004.

[27] Wein W., Barbara R., and Nassir N.. Automatic registration and fusion of ultrasound with CT for radiotherapy. In Medical Image Computing and Computer Assisted Intervention , MICCAI 2005, Lecture Notes in Computer Science, pages 303–311. Springer, Oct 2005.

[28] Stewart C.V., Tsai C.L., and Roysam B.. The dual-bootstrap iterative closest point algorithm with application to retinal image registration. IEEE Transaction on Medical Imaging, 22(11):1379–1394, 2003.

[29] Dam E., Folkesson J., Pettersen P., and Christiansen C., "Semi-automatic knee cartilage segmentation", Proc. SPIE Int. Soc. Opt. Eng. 6144, 614441, 2006.

[30] Grau V., et al. 'Improved Watershed Transform for Medical Image Segmentation Using Prior Information'. IEEE Trans. Med. Imaging 23(4): 447-458 , 2004.

[31] Warfield S., Winalski C., Jolesz F. and Kikinis R., "Automatic Segmentation of MRI of the Knee", http://splweb.bwh.harvard.edu:8000/pages/papers/simonw/index.html

[32] Lynch J., Zaim S., Zhao J., Stork A., Peterfy C.G., and Genant H.K.. Cartilage segmentation of 3D MRI scans of the osteoarthritic knee combining user knowledge and active contours. SPIE, pages 925–935, 2000.

[33] Lynch J., Zaim S., Zhao J., Peterfy C.G., and Genant H.K.. Automating measurement of subtle changes in articular cartilage from MRI of the knee by combining 3D image registration and segmentation. In Proceedings SPIE '01, pages

431–439, 2001.
[34] Dam E., Folkesson J., Pettersen P., and Christiansen C., “Semi-automatic knee cartilage segmentation”, Proc. SPIE Int. Soc. Opt. Eng. 6144, 614441, 2006.
[35] Stammberger T., Eckstein F., Michaelis M., Englmeier K.H., and Reiser M.. Interobserver reproducibility of quantitative cartilage measurements: Comparison of B-spline snakes and manual segmentation. Magnetic Resonance Imaging, 17(7):1033–1042, 1999.
[36] Tejos, C.; Hall, L.D.; Cardenas-Blanco, A. , “Segmentation of articular cartilage using active contours and prior knowledge”, 26th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2004. IEMBS '04. Vol.1, Iss., 1-5 Sept. 2004 Pages: 1648- 1651 Vol.3
[37] Raynauld J.-P., Kauffmann C., Beaudoin G., Berthiaume M.-J., de Guise J. A., A. Bloch D., Camacho F., Godbout B., Altman R. D., Hochberg M., Meyer J. M., Cline G., Pelletier J.-P., and Martel-Pelletier J.. Reliability of a quantification imaging system using magnetic resonance images to measure cartilage thickness and volume in human normal and osteoarthritic knees. Osteoarthritis and Cartilage, 11(5):351–360, May 2003.
[38] Raynauld J.-P., Martel-Pelletier J., Berthiaume M.-J., Labont ´e F., Beaudoin G., de Guise J.A., Bloch D.A., Choquette D., Haraoui B., Altman R.D., Hochberg M.C., Meyer J.M., Cline G.A., and Pelletier J.-P.. Quantitative magnetic resonance imaging evaluation of knee osteoarthritis progression over two years and correlation with clinical symptoms and radiologic changes. Arthritis & Rheumatism, 50(2):476–487, February 2004.
[39] Kapur T., Beardsley P., Gibson S., Grimson W., and Wells W., “Model-based segmentation of clinical knee MRI.”, In IEEE International Workshop on Model-Based 3D Image Analysis ’98, pages 97–106, 1998.
[40] Fripp, J.; Crozier, S.; Warfield, S.; Ourselin, S. , “Automatic Initialization of 3D Deformable Models for Cartilage Segmentation”, Digital Image Computing: Techniques and Applications, 2005. DICTA '05. Proceedings 2005, Vol., Iss., 06-08 Dec. 2005
[41] Cootes T.F., Hill A., Taylor C.J., and Haslam J.. Use of active shape models for locating structures in medical images. Image and Vision Computing, 12(6):355–365, July-August 1994.
[42] Cootes T.F., Taylor C.J., Cooper D.H., and Graham J.. Active shape models - their training and application. Computer Vision and Image Understanding, 61(1):38–59, 1995.
[43] Fripp J., et al. ‘Automatic segmentation of the bone and extraction of the bone-cartilage interface from magnetic resonance images of the knee’, Physics in Medicine and Biology, Volume: 52 Start Page: 1617, 2007.
[44] Fripp J., Crozier S., Warfield S.K., Ourselin S.: Automatic Initialization of 3D Deformable Models for Cartilage Segmentation. DICTA 2005: 74, 2005.
[45] Lorigo L., et al. ‘Segmentation of Bone in Clinical Knee MRI Using Based Geodesic Active Contours’. In MICCAI, October 1998.
[46] Caselles V., Kimmel R., Sapiro G.: Geodesic Active Contours. ICCV 1995: 694-699
[47] Caselles V., Kimmel R., Sapiro G.. International Journal of Computer Vision

22(1), 61–79, Kluwer Academic Publishers, 1999.
[48] Bourgeat, P et al. ‘T Based Segmentation of the Knee Bones in MRI using Phase Information’. 3rd IEEE International Symposium on Biomedical Imaging: Macro to Nano,. 742- 745, 2006.
[49] Schmid J, and Magnenat-Thalmann N. MRI bone segmentation using deformable models and shape priors. in Proc. of Medical Image Computing and Computer-Assisted Intervention. New York City, NY, USA. 2008; pp. 44-51
[50] Alfano B, Amatoo U, Antoniadis A, and Larobina M. Segmentation of MR brain images through discriminant analysis. Journal of Neuroscience Methods. 30 December 2003; 131(1-2):65-74.
[51] Andersen AH, Zhang Z, Avison MJ, Gash DM. Automated segmentation of multispectral brain MR images. Journal of Neuroscience Methods. 31 December 2002; 122(1):13-23.
[52] Ostrowitzki S, Rédei J, Lynch JA, Carano RA, Zaim S, Miaux Y, Genant HK. Use of multispectral magnetic resonance imaging analysis to quantify erosive changes in the hands of patients with rheumatoid arthritis: short-term and long-term longitudinal studies. Arthritis & Rheumatism. 2004; 50:716–724.
[53] Carano RA, Lynch JA. Quantification of bone lesions in rheumatoid arthritis by MRI multispectral tissue class transition analysis. Seminars in Musculoskeletal Radiology. Dec 2004; 8(4):301-11.
[54] Gold GE, Han E, Stainsby J, Wright G, Brittain J, Beaulieu C. Musculoskeletal MRI at 3.0 T: relaxation times and image contrast. American Journal of Roentgenology. August 2004; 183(2):343-351.
[55] Perona P, Malik J. Scale-space and edge detection using anisotropic diffusion. IEEE Transactions on Pattern Analysis and Machine Intelligence. July 1990; 12(7):629-639.
[56] Sled JG, Zijdenbos AP, Evans AC. A nonparametric method for automatic correction of intensity nonuniformity in MRI data. IEEE Transactions on Medical Imaging. February 1998; 17(1):87-97.
[57] Sethian J. Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry. Cambridge Press. 1998.
[58] Yoo T, editor. Insight into Images 'Principles and Practice for Segmentation, Registration and Image Analysis. A K Peters Ltd.; 2004.
[59] Barrett, W, and Mortensen, E. Interactive live-wire boundary extraction. Medical Image Analysis, 1997; 1(4): 331–341.
[60] Barrett, W., Mortensen, E. Fast, accurate, and reproducible live-wire boundary extraction. in Proc. of Visualization in Biomedical Computing, Hamburg, Germany, 1996; pp 183-192.
[61] Poon M., Hamarneh G., Abugharbieh R. Segmentation of complex objects with non-spherical topologies from 3D medical images using 3D livewire. in Proc. of SPIE Medical Imaging, San Diego-USA, 2007; pp. 651231 1-10.
[62] Kovesi, P.: Symmetry and Asymmetry from Local Phase, AI 1997. In: Tenth Australian Joint Conference on Artificial Intellegence, pp. 185–190 (1997)
[63] Harris, C., Stephens, M.: A combined corner and edge detector. Proceedings of the 4th Alvey Vision Conference: pp 147—151, (1988).
[64] Fischler, M., Bolles, R.: Random Sample Consensus: A Paradigm for Model

Fitting with Applications to Image Analysis and Automated Cartography. Comm. of the ACM, Vol 24, pp 381-395, (1981).
[65] Gering D., Nabavi A., Kikinis R., Grimson W.E.L., Hata N., Everett P., Jolesz F. and Wells W., An integrated visualization system for surgical planning and guidance using image fusion and interventional imaging. In: Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI) (1999), pp. 809–819.
[66] Zonneveld F.W., Intra-operative CT: implementation of the Tomoscan M. Medicamundi 42 (1998), pp. 6–11.
[67] Angelsen B.A.J., Ultrasound Imaging. , Emantec, Trondheim, Norway (2000).
[68] Ault T. and Siegel M.W., Frameless patient registration using ultrasonic imaging. In: Proc. Int. Symposium on Medical Robotics and Computer Assisted Surgery (1994), pp. 74–82 1 (2) .
[69] Carrat L., Tonetti J., Merloz P. and Troccaz J., Percutaneous computer assisted iliosacral screwing: clinical validation. In: Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI) (2000), pp. 1229–1237.
[70] Amin D.V., Kanade T., DiGioia A., Jaramaz B., LaBarca R. and Nikou C., Non-invasive ultrasound registration of the pelvic bone surface for THR. In: Presented at the 47th Annual Meeting of the Orthopaedic Research Society, San Francisco (CA), USA (2001).
[71] Muratore D.M., Dawant B.M. and Galloway R.L., Vertebral surface extraction from ultrasound images for technology-guided therapy. In: Proc. SPIEMedical Imaging vol. 3661 (1999), pp. 1499–1510.
[72] Weber P.K., Peter L., Voss G., Schlegel J.C. and Harland U., A system for ultrasound-based intraoperative navigation in spine surgery. In: Proc. IEEE Int. Ultrasonics Symp. (2001), pp. 1361–1364.
[73] Ioppolo J., Kowal J. and Nolte L.P., Ultrasonic registration techniques. In: Presented at the 2nd Annual Meeting of the International Society for Computer Assisted Orthopaedic Surgery (CAOS International), Santa Fe (NM), USA (2002).
[74] American College of Rheumatology Subcommittee on Rheumatoid Arthritis. Guidelines for the management of rheumatoid arthritis: 2002 Update. Arthritis Rheum.;46:328–346. doi: 10.1002/art.10148; 2002.
[75] Boers M, Felson DT. Clinical measures in rheumatoid arthritis: which are most useful in assessing patients? J Rheumatol.;21:1773–1774; 1994.
[76] McQueen FM, Stewart N, Crabbe J, Robinson E, Yeoman S, Tan PL, McLean L. Magnetic resonance imaging of the wrist in early rheumatoid arthritis reveals a high prevalence of erosions at four months after symptom onset. Ann Rheum Dis. 1998;57:350–356.
[77] Backhaus M, Kamradt T, Sandrock D, Loreck D, Fritz J, Wolf KJ, Raber H, Hamm B, Burmester GR, Bollow M. Arthritis of the finger joints: a comprehensive approach comparing conventional radiography, scintigraphy, ultrasound, and contrast-enhanced magnetic resonance imaging. Arthritis Rheum. 1999;42:1232–1245. doi: 10.1002/1529-0131(199906)42:6<1232::AID-ANR21>3.0.CO;2-3. [PubMed]
[78] Klarlund M, Østergaard M, Jensen KE, Madsen JL, Skjødt H, Lorenzen I. Magnetic resonance imaging, radiography, and scintigraphy of the finger joints: one year follow up of patients with early arthritis. The TIRA Group. Ann Rheum Dis.

2000;59:521–528. doi: 10.1136/ard.59.7.521. [PubMed]
[79] Lindegaard HM, Vallø J, Hørslev-Petersen K, Junker P, Østergaard M. Low-cost, low-field dedicated extremity magnetic resonance imaging in early rheumatoid arthritis: a 1-year follow-up study. Ann Rheum Dis. 2006;65:1208–1212. doi: 10.1136/ard.2005.049213. [PubMed]
[80] D. Tomazevic, B. Likar and F. Pernus, Gold standard 2D/3D registration of X-ray to CT and MR images. Medical Image Computing and Computer-Assisted Intervention (MICCAI 02) (2002).
[81] Yamazaki T. , Watanabe T., Tomita T., Sugamoto K., Ogasawara M., Sato Y., Yoshikawa H. and Tamura S. 3D kinematics of normal knee using X-ray fluoroscopy and CT images. World Congress on Medical Physics and Biomedical Engineering, South Korea, 2006.
[82] Yuan, X.H. and Ryd, L., 2000. Accuracy analysis for RSA: a computer simulation study on 3D marker reconstruction. Journal of Biomechanics 33, pp. 493–498.
[83] Selvik, G., 1989. Roentgen stereophotogrammetry. A method for the study of the kinematics of the skeletal system. Acta Orthopaedica Scandinavica 60 Suppl 232, pp. 1–51.
[84] Goldbach-Mansky R, Woodburn J, Yao L, Lipsky PE. Magnetic resonance imaging in the evaluation of bone damage in rheumatoid arthritis: a more precise image or just a more expensive one? Arthritis Rheum. 2003;48:585–589. doi: 10.1002/art.10819. [PubMed]
[85] Møller Døhn U, Ejbjerg BJ, Court-Payen M, Hasselquist M, Narvestad E, Szkudlarek M, Møller JM, Thomsen H, Østergaard M. Are bone erosions detected by magnetic resonance imaging and ultrasonography true erosions? A comparison with computed tomography in rheumatoid arthritis metacarpophalangeal joints. Arthritis Res Therapy. 2006;8:R110. doi: 10.1186/ar1995.
[86] Eckstein F., Cicuttini F., Raynauld J., Waterton J., Peterfy C. Magnetic resonance imaging (MRI) of articular cartilage in knee osteoarthritis (OA): morphological assessment Osteoarthritis and Cartilage, Volume 14, Issue null, Pages 46-75, 2006.
[87] Bolbos R, Benoit-Cattin H, Langlois JB, Chomel A, Chereul E, Odet C, Pastoureau P, Janier M, Beuf O. Knee cartilage thickness measurements using MRI: a 4(1/2)-month longitudinal study in the meniscectomized guinea pig model of OA., Osteoarthritis Cartilage. Jun;15(6):656-65. Epub 2007 Feb 14, 2007.
[88] Muhle C, Brossmann J, Heller M. Kinematic CT and MR imaging of the patellofemoral joint. Eur Radiol 1999;9(3):508–18.
[89] Anderst WJ, Tashman S. A method to estimate in vivo dynamic articular surface interaction. J Biomech 2003;36(9):1291–9.
[90] Powers CM, Ward SR, Fredericson M, et al. Patellofemoral kinematics during weightbearing and non-weight-bearing knee extension in persons with lateral subluxation of the patella: a preliminary study. J Orthop Sports Phys Ther 2003;33(11):677–85.
[91] Sheehan FT, Drace JE. QuantitativeMR measures of three-dimensional patellar kinematics as a research and diagnostic tool. Med Sci Sports Exerc 1999;31(10):1399–405.
[92] Wilson DR., McWalter EJ., Johnston JD.,The Measurement of Joint Mechanics and their Role in Osteoarthritis Genesis and Progression, Rheum Dis Clin N Am 34 605–

622, 2008.
[93] Sheehan FT, Zajac FE, Drace JE. In vivo tracking of the human patella using cine phase contrast magnetic resonance imaging. J Biomech Eng 1999;121(6):650–6.
[94] Rebmann AJ, Sheehan FT. Precise 3D skeletal kinematics using fast phase contrast magnetic resonance imaging. J Magn Reson Imaging 2003;17(2):206–13.
[95] Nandhagopal R., McKeown M., Stoessl J., Functional. Imaging in Parkinson's Disease. Neurology 2008;70:1478-1488
[96] Korosec F., Frayne R., Grist T., Mistretta C., Time-resolved contrast-enhanced 3D MR angiography. Magnetic Resonance in Medicine, Volume 36 Issue 3, Pages 345 – 351, 2005.
[97] Rao J., Hamarneh G., and Abugharbieh R.. Adaptive Contextual Energy Parameterization for Automated Image Segmentation. in Lecture Notes in Computer Science, International Symposium on Visual Computing: Special Track on Optimization for Vision, Graphics and Medical Imaging: Theory and Applications (ISVC OVGMI), volume 5875-I, pages 1089-1100, 2009.

Appendix A - Publications

Journals

R. Dalvi, R. Abugharbieh, D. R. Wilson. "Multi-Contrast MR Imaging for Enhanced Bone Tissue Visualization and Segmentation". Submitted to Magnetic Resonance in Medicine (MRM), under revision

Conferences

R. Dalvi, R. Abugharbieh, "Efficient Feature Based Multi Slice to Volume Registration Using Phase Congruency". International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Vancouver-Canada, 2008, pp. 5390-5393

R. Dalvi, R. Abugharbieh, D. C. Wilson, D. R. Wilson. "Multi-Contrast MR for Enhanced Bone Imaging and Segmentation". International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), Lyon-France, 2007, pp. 5620-5623

R. Dalvi, R. Abugharbieh, M. Pickering, P. Smith. "Registration of 2D to 3D Joint Images Using Phase-Based Mutual Information". SPIE Medical Imaging (MI), San Diego-USA, 2007, pp. 651209 1-9

R. Dalvi, R. Abugharbieh, D. C. Wilson, D. R. Wilson. "Highly-Automated 3D Segmentation of Femoral Bone from Hip MRI". International Society for Magnetic Resonance in Medicine Scientific Meeting (ISMRM), Berlin-Germany, 2007.

R. Dalvi, I. Hacihaliloglu, R. Abugharbieh. "Fast and Accurate 3D Ultrasound Volume Stitching Using Phase Symmetry and Harris Corner Detection for Orthopedic Applications". SPIE Medical Imaging (MI), San Diego-USA, 2010 (accepted)

Appendix B – Statement of Imaging Ethics

The live MRI human data were taken at Boston University by Dr. David Hunter and Dr. Derek Wilson (CREB #: H04-70225) at the University of British Columbia (UBC) Hospital and the US scans were taken by Mr. Ilker Hacihaliloglu at UBC. All of them have cleared the required ethical concerns. No live human test data were collected by me (Rupin Dalvi) personally, or under my name.