

**Feature-Based Registration of Preoperative CT to
Intraoperative 3-D Ultrasound in Laparoscopic Partial
Nephrectomy Using A Priori CT Segmentation**

by

John Bartlett

B.Sc. Biochemistry, Queen's University, 2008

B.Cmp. (Honours) Biomedical Computing, Queen's University, 2009

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

Master of Applied Science

in

THE FACULTY OF GRADUATE STUDIES

(Electrical and Computer Engineering)

The University Of British Columbia

(Vancouver)

August 2011

© John Bartlett, 2011

Abstract

Robotic laparoscopic partial nephrectomy is a state-of-the-art procedure for the excision of renal tumours. The challenges of this surgery along with the stereoscopic interface to the surgeon make it an ideal candidate for image guidance. We propose bringing pre-operative computed tomography data to the patient's coordinate system using three-dimensional intraoperative back ultrasound. Since computed tomography and ultrasound images represent like anatomical information quite differently, we perform a manual segmentation of the computed tomography before the operation and a semi-automatic segmentation of the ultrasound intra-operatively. The segmentation of the kidney boundary facilitates a feature-based registration strategy.

Semi-automatic segmentation of kidney ultrasound images is difficult because the edges with large gradient values do not correspond to the capsule boundary seen in computed tomography. The desired edges are actually quite faint in ultrasound and poorly detected by common edge methods such as the Canny approach.

After trying a number of approaches, the best results were obtained using a novel interacting multiple-model probabilistic data association filter to select edges from ultrasound images that were filtered for phase congruency. The manual segmentation of the prior is used to guide edge detection in ultrasound. Experiments on seven pre-operative patient datasets and one intra-operative patient dataset resulted in a mean volume error ratio of 0.80 ± 0.13 from after registration to before registration. These results came after the implementation and evaluation of numerous other approaches, including radial edge filters, the covariance matrix adaptation evolution strategy, and a deformable approach using geodesic active contours.

The main contribution of this work is a method for the registration of the pre-

operative planning data from computerized tomography to the intraoperative ultrasound. For clinical use, this method requires some form of calibration with the laparoscopic camera and integration with surgical visualization tools. Through integration with emerging technologies, the approach presented here can one day augment the surgical field-of-view and guide the surgeon around important anatomical structures to the tissue that must be excised.

Table of Contents

Abstract	ii
Table of Contents	iv
List of Tables	ix
List of Figures	x
Glossary	xii
Acknowledgements	xv
1 Introduction	1
1.1 Motivation	1
1.2 Objectives	2
1.3 Contributions	3
2 Background	5
2.1 Laparoscopic Partial Nephrectomy	5
2.1.1 Partial Nephrectomy	5
2.1.2 The Laparoscopic Approach	6
2.1.3 Challenges of Laparoscopic Partial Nephrectomy	6
2.1.4 Robotic LPN	7
2.2 Ultrasound Image Segmentation	7
2.2.1 Segmentation of the Kidney in Ultrasound Images	8
2.2.2 The Interacting Multiple-Model Probabilistic Data Association Filter for Edge Detection	10

2.2.3	Level Set Methods	11
2.2.4	Discrete Dynamic Contours	12
2.2.5	Image Features from Local Phase Information	12
2.3	Image Registration	13
2.3.1	Feature-based Registration	13
2.3.2	Intensity-based Registration	14
2.3.3	Transform Model Estimation and Image Resampling	14
2.3.4	Validating Registration Results	15
2.3.5	Medical Image Registration	15
2.3.6	Optimization Strategies	17
2.4	Related Work	19
2.4.1	3-D MR to 3-D US Registration of the Brain	20
2.4.2	3-D MR to 2.5-D US Registration of the Liver	20
2.4.3	3-D CT to 2.5-D US Registration of the Kidney	21
2.4.4	Image-guided LPN with Fluoroscopy	22
3	Experiment Design	23
3.1	Data Acquisition	23
3.2	Data Preparation	25
3.3	Experimental Setup	25
3.4	Validating Registration	25
4	Methods	28
4.1	Initial Registration of the Prior CT Segmentation to the Initializa- tion Points	29
4.2	Edge Detection and Segmentation Methods	30
4.2.1	Canny Edge Detection	31
4.2.2	Edge Detection Using Phase Congruency	31
4.2.3	Radial IMMPDAF for Edge Detection in Ultrasound	32
4.2.4	Prior-guided IMMPDAF	36
4.2.5	IMMPDAF with Maximum Intensity Filter	36
4.2.6	Discarding Canny Edges Far from Prior	36
4.2.7	Level Sets and Geodesic Active Contours	37

4.3	Registration Methods	40
4.3.1	Iterative Closest Point	40
4.3.2	Unscented Kalman Filter-based Registration	40
4.3.3	CMA-based Surface Registration	40
4.3.4	Coherent Point Drift	42
4.4	Approaches	42
4.4.1	Radial IMMPDAF with ICP Registration	42
4.4.2	Radial IMMPDAF with UKF-based Registration	42
4.4.3	IMMPDAF with Canny Filter and ICP Registration	42
4.4.4	IMMPDAF with Directional Phase Filter and ICP Registration	43
4.4.5	Prior-guided IMMPDAF with ICP Registration	43
4.4.6	Prior-guided IMMPDAF on Canny Edges with ICP Registration	43
4.4.7	Prior-guided IMMPDAF on Directional Phase Edges with ICP Registration	44
4.4.8	CMA Point-Surface Registration with Radial Filter Edge Candidates	44
4.4.9	CMA Slice-wise Registration with Canny Edge Map	44
4.4.10	CMA Slice-wise Registration with Phase Edge Map	45
4.4.11	Deformable Registration with LS/GAC and Canny Filter Edges	45
5	Results	46
5.1	Initialization	46
5.2	IMMPDAF Methods	46
5.2.1	Radial IMMPDAF	46
5.2.2	Prior-guided IMMPDAF	47
5.3	CMA Methods	49
5.4	Deformable LS/GAC-based Approach	50
5.5	Intraoperative US Data	52
6	Discussion	55

6.1	Understanding the Ultrasound Images	55
6.1.1	Anatomical Significance of Visible Edges	55
6.1.2	Segmenting the Tumour	56
6.1.3	Treatment of US data	58
6.2	Initialization	58
6.3	Edge Detection Filters	59
6.3.1	Canny Edge Filter	59
6.3.2	Phase Congruency Filter	60
6.4	Radial IMMPDAF	60
6.4.1	Performance	62
6.5	Prior-guided IMMPDAF	62
6.5.1	Velocity damping	62
6.5.2	Directional edge filtering	63
6.5.3	Registration Based on Inner Edges	64
6.5.4	Canny Edge Maps	65
6.5.5	Phase Congruency Images	65
6.5.6	Performance	65
6.6	Deformable LS/GAC-based Approach	66
6.7	Registration Methods	67
6.7.1	Iterative Closest Point	67
6.7.2	UKF-based Registration	68
6.7.3	Covariance Matrix Adaptation	68
6.7.4	Other Registration Options	69
7	Conclusions	70
7.1	Future Work	71
7.1.1	Improving the Interface	71
7.1.2	Extending the IMMPDAF	72
7.1.3	Realistic Deformable Registration	72
7.1.4	Improved Error Estimation	72
7.1.5	Integration with Surgical Devices	73
	Bibliography	74

A	Vol2Strad: a Volume Conversion Tool	83
B	Setting IMMPDAF Parameters	85
	B.1 Acquisition of Phantom Data	85
	B.2 Parameter Manipulation Experiments	85
C	Patient Results for Prior-guided IMMPDAF Methods	88

List of Tables

Table 5.1	Volume error ratios for IMMPDAF methods	47
Table 5.2	Volume error ratios for CMA methods	51
Table 6.1	Volume ratios between US and CT segmentations	56
Table B.1	IMMPDAF parameter configurations	86

List of Figures

Figure 1.1	End goal: image integration in surgical field-of-view	2
Figure 1.2	CT to US registration as a part of the big picture	3
Figure 2.1	Port placement diagram for da Vinci [®] surgery	8
Figure 3.1	Approximate patient position during surgery	24
Figure 3.2	Simple GUI for US navigation and initialization point selection.	26
Figure 3.3	Volume error relationship with each DOF of rigid motion for a kidney segmentation	27
Figure 4.1	Basic workflow for all approaches	28
Figure 4.2	Overview of the initialization process	29
Figure 4.3	Prior contour and search range	30
Figure 4.4	Canny edge detection intermediates and output	32
Figure 4.5	Comparison of kidney US images before and after phase-filtering	33
Figure 4.6	Comparison between radial and prior-guided IMMPDAF . . .	35
Figure 4.7	Derivation of edge potential stopping function for LS/GAC . .	38
Figure 4.8	Volumetric edge potential stopping function for LS/GAC . . .	39
Figure 5.1	Percent volume error after initial alignment	47
Figure 5.2	Volume error for radial IMMPDAF methods	48
Figure 5.3	Volume error for prior-guided IMMPDAF methods	49
Figure 5.4	Surface centroid error for prior-guided IMMPDAF methods .	50
Figure 5.5	Volume error for CMA methods and deformable LS/GAC-based approach	51

Figure 5.6	Volume error for prior-guided IMMPDAF methods on Patient 5 intra-operative data	53
Figure 5.7	Surface centroid error for prior-guided IMMPDAF methods on Patient 5 intra-operative data	54
Figure 6.1	Kidney boundaries in US	57
Figure 6.2	Correspondence between phase edges and true kidney boundary	61
Figure 6.3	Effect of damping term on prior-guided IMMPDAF	63
Figure 6.4	Directional selectivity in kidney segmentation	64
Figure B.1	IMMPDAF output for various filter parameter combinations .	87
Figure C.1	Volume error for prior-guided IMMPDAF	89
Figure C.2	Volume error for prior-guided IMMPDAF	90
Figure C.3	Volume error for prior-guided IMMPDAF on Canny edges . .	91
Figure C.4	Volume error for prior-guided IMMPDAF on Canny edges . .	92
Figure C.5	Volume error for prior-guided IMMPDAF on phase edges . .	93
Figure C.6	Volume error for prior-guided IMMPDAF on phase edges . .	94
Figure C.7	Error map for prior-guided IMMPDAF on phase images for patients 3 to 6	95
Figure C.8	Error map for prior-guided IMMPDAF on phase images for patients 8 to 13	96

Glossary

AR	Augmented Reality
CPD	Coherent Point Drift
CPU	Central Processing Unit
CR	Correlation Ratio
CT	Computed Tomography
DDC	Discrete Dynamic Contours
DOF	Degree Of Freedom
DoG	Difference-Of-Gaussian
ES	Evolutionary Strategy
FEM	Finite Element Model
FOV	Field-Of-View
GAC	Geodesic Active Contours
GPU	Graphics Processing Unit
GUI	Graphical User Interface
ICP	Iterative Closest Point
IMMPDAF	Interacting Multiple-Model Probabilistic Data Association Filter

LPN	Laparoscopic Partial Nephrectomy
LS	Level Set
KF	Kalman Filter
MI	Mutual Information
MIS	Minimally Invasive Surgery
MSC	Mutative Strategy parameter Control
MR	Magnetic Resonance
NCC	Normalized Cross-Correlation
NSS	Nephron-Sparing Surgery
OR	Operating Room
PCA	Principle Component Analysis
PDA	Probabilistic Data Association
PET	Positron Emission Tomography
PN	Partial Nephrectomy
PRP	Percutaneous Renal Puncture
RLPN	Robotic LPN
RMS	Root Mean-Square
RN	Radical Nephrectomy
ROI	Region of Interest
SDK	Software Development Kit
SNR	Signal-to-Noise Ratio
SPECT	Single-Photon Emission Computed Tomography

TRE	Target Registration Error
TRUS	Trans-Rectal Ultrasound
UKF	Unscented Kalman Filter
US	Ultrasound

Acknowledgements

I would first like to thank my supervisors, Dr. Tim Salcudean and Dr. Purang Abolmaesumi, for guiding me through all my efforts. They saw the good side of every failed approach and always had ideas for improvements. This perseverance eventually led to success, and I am grateful to have had the opportunity to work with such accomplished researchers on this important project.

I would also like to thank Dr. Chris Ngan for his dedication to research. He brought myself and my colleagues into his OR, presented a clinical point of view at research meetings, and was available to answer clinical questions upon request.

I thank Vickie Lessoway for sharing her ultrasound expertise. Without her close involvement, it might have taken much longer for us to notice that we were targeting the wrong edges. She was also an integral member of our kidney project's data acquisition team. This team also included Steven Tang, Caitlin Schneider, Raoul Kingma and Michael Yip. Dr. Robert Rohling played a vital role on the kidney project as well.

This research would not have been possible without funding, so I would now like to recognize the financial support received from NSERC. Their investment in innovation is critical to a vast array of post-secondary research projects.

I also have to thank numerous students from the Robotics and Control Lab at UBC for their support and advice. Notably: Orcun Goksel, Guy Nir, Jing Xiang, Sara Mahdavi and Abtin Rasoulia. Caitlin Schneider contributed a number of CT segmentations as well. Troy, Hedy and Jeff: there wasn't much overlap in our work, but it's been great sharing a lab with you!

At this point I would like to thank my parents, who supported my academic pursuits all my life and certainly played a big part in getting me here. Don't worry,

I'll be coming home soon!

Last but not least, thanks to my girlfriend, Nicole, for keeping me smiling throughout the ups and downs of my research. More recently, developing tendinitis as my thesis work was drawing to a close might have been an impossible obstacle if it wasn't for her support. I don't think I can fully express my gratitude in this acknowledgement, but it's hard to imagine how I would have completed this work without her assistance and encouragement.

Chapter 1


Introduction

1.1 Motivation

There will be an estimated 5,100 new cases of kidney cancer in Canada in 2011 [17], up from 3,900 in 2001 [16], making it the cancer type with the seventh highest incidence rate in men and the eleventh highest in women [17]. Laparoscopic Partial Nephrectomies (LPNs) and Robotic LPNs (RLPNs) are increasingly being used for treatment due to decreased patient morbidity, shorter hospital stays and comparable oncologic outcomes with respect to open surgery [15, 38].

During surgery, the surgeon has the difficult task of laparoscopically locating the kidney through fat and connective tissue, all while averting injury to any vessels and nerves that are encountered. Depending on which kidney is being operated on, it may also be necessary to mobilize the liver, spleen or pancreas as well as the colon [79].

The surgeon has the benefit of a high-resolution 3-dimensional (3-D) Computed Tomography (CT) image of the patient's abdomen, but its utility is limited without knowing how this data relates to the patient's current position. Ideally, we would like the detailed CT volume to be integrated in an augmented surgical Field-Of-View (FOV) to provide real-time image guidance. CT is not available in the Operating Room (OR), but we can consider bringing it into alignment with intra-operative ultrasound (US) and the laparoscopic video using image registration techniques (see Figure 1.1).



©2010 Intuitive Surgical, Inc.

Figure 1.1: Laparoscopic video (bottom left) will ideally be augmented with intra-operative US (centre) and pre-operative CT (right) in the surgeon's FOV.

1.2 Objectives

This study proposes to aid the surgeon by bringing the pre-operative CT data into the patient's coordinate frame by registering it to an intra-operative US volume. US is commonly used intra-operatively because it is safe, inexpensive and fast. Although the usefulness of US images in image guidance is limited due to noise and artefacts, features that are commonly observed in US and CT data can be used to perform a registration of the two modalities [90]. CT to US registration makes up the first link in the chain depicted in Figure 1.2. In order to reach the end goal of integrated CT guidance, the method in this study will have to co-operate with other applications that can register the US data to the laparoscopic video or the surgical tools.

The clinical benefits will come from integration with other developing technologies, eventually leading to:

- An augmented navigation environment during LPN to improve the efficiency and safety of the operation
- Improved identification of pathologic targets

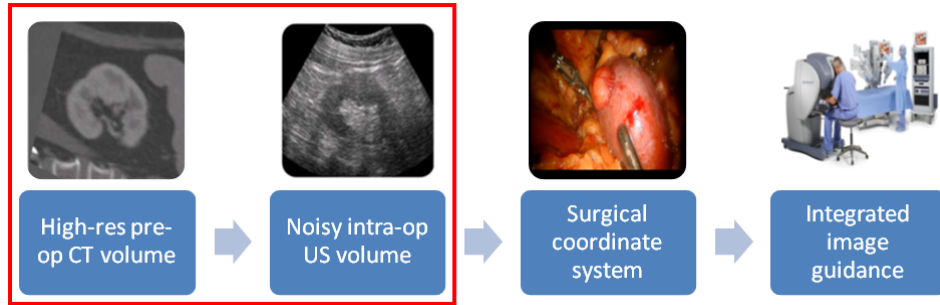


Figure 1.2: This paper will deal exclusively with the CT to US registration step (highlighted in red), but this is part of a bigger picture. Registration of the US to the surgical coordinate system and integration into an augmented display will require the use of other techniques.

- Resection guidance tools for excision of target end-organ pathology
- Surgical preplanning tools for designing individualized approaches to operative cases
- Visualization tools that will facilitate the development of new approaches to dealing with current pathologies

The proposed algorithm will take user-selected initialization points of the US kidney boundary along with a pre-operative CT segmentation of the kidney, and determine the full segmentation of the kidney in the US volume as well as the transformation required to align the US and CT images. If the US image can be registered to the surgical coordinate frame by other means, the anatomical information in the CT will be mapped to actual locations in the patient in the OR.

1.3 Contributions

- A joint segmentation and registration algorithm designed to bring a pre-operative plan into alignment with intra-operative back US during LPN.
- A rough clinical workflow, from initialization to registration, that could be incorporated into medical user interfaces and intraoperative visualization tools.

- Trials with eight patient datasets to evaluate the effectiveness of this algorithm on clinical data.
- New insight into kidney segmentation in US images that conflicts with some segmentation papers but agrees with the kidney shape and volume seen in CT.
- An extension of the Interacting Multiple-Model Probabilistic Data Association Filter (IMMPDAF) used for prostate segmentation that makes use of an a priori model along with other domain-specific information to create a specialized kidney segmentation approach.
- The application of a phase congruency filter to kidney images and an evaluation of its performance.
- An exploration of other popular edge detection and registration techniques in the context of the kidney registration problem for LPN.

Chapter 2

Background

2.1 Laparoscopic Partial Nephrectomy

2.1.1 Partial Nephrectomy

Partial Nephrectomy (PN), the excisional form of Nephron-Sparing Surgery (NSS), was originally performed only when required, that is, for tumours occurring in a solitary kidney or bilaterally. However, favourable intermediate- and long-term surgical outcomes for tumours up to 5 cm in diameter led to the inclusion of all patients with sufficiently small tumours, as it allows better post-operative renal function [3, 38, 70, 79]. Ten years after surgery, the incidence of renal insufficiency and proteinuria was shown to be higher in Radical Nephrectomy (RN) patients than PN patients. In addition, up to 10% of patients may develop a tumour in the contralateral kidney, thus increasing the dependence on the original affected kidney [3].

A 73% increase in CT and Magnetic Resonance (MR) use between 1986 and 1994 has led to more frequent detection of incidental tumours, which represented 13% of renal tumours in 1982 and 59% in 1997. Unsurprisingly, the mean size of detected renal tumours decreased by 8mm from 1988 to 2002, indicating that a greater number of patients are eligible for treatment by PN [3]. Small renal tumours are generally less aggressive, but only tumours under 3 cm in diameter in highly selected patients should be considered for observation instead of treatment [65].

2.1.2 The Laparoscopic Approach

Minimally-invasive surgical options for treatment of renal tumours are preferred, as they avoid the muscle-cutting flank incision required for open nephrectomy. Ablative surgery, which includes cryoablation and radiofrequency (RF) ablation, is a minimally-invasive option but lacks long-term studies to confirm its effectiveness [3, 43].

Clinical LPN was first reported in 1993 [86]. LPN was originally indicated only for small exophytic tumours, but has since been recommended for parenchymal tumours reaching as far as the collecting system or the renal sinus, as well as completely intrarenal tumours and even large tumours requiring heminephrectomy [43, 79].

A 3-D CT image is taken prior to the procedure, enabling the surgeon to precisely locate the boundaries of the tumour and decide whether to use a transperitoneal or retroperitoneal approach [43, 79]. The transperitoneal approach is preferable as it provides a larger working space and better suturing angles [79]. The retroperitoneal approach may be advantageous in cases where the tumour is located at one of the poles [3].

In right LPN, the liver and colon must be mobilized to expose the kidney. Mobilization of the duodenum may also be required in some cases. For LPN on the left kidney, the spleen, splenic flexure, pancreas and colon must be mobilized [79]. Laparoscopic US is used to observe any changes in the location and size of the tumour since the pre-operative CT, and measure its depth. Resection of the tumour is performed with cold scissors after the hilum is clamped [43, 79].

2.1.3 Challenges of Laparoscopic Partial Nephrectomy

In RN, the tumour remains covered by perirenal fat throughout the surgery and is usually untouched by the instruments. In contrast, PN carries a higher risk of tumour perforation and spillage due to the increased contact with the surgical tools. However, tumour seeding is avoidable when the surgeon has a good FOV and uses safe resection margins [70].

One of the challenges of the laparoscopic procedure is hemorrhage prevention by means of ischemia. By clamping the hilum, the surgeon is able to work with

an essentially bloodless FOV [79]. In open surgery, cold ischemia is attainable by applying ice. Hypothermia by this method is not feasible in laparoscopic surgery, so warm ischemia is used out of necessity. Consequently, the surgeon faces a time constraint of 20-30 minutes to perform the tumour resection once the renal artery has been clamped in order to avoid damage to the renal parenchyma [70]. Warm ischemia lasting longer than 20 minutes is associated with higher incidences of acute renal failure and chronic renal insufficiency [38].

2.1.4 Robotic LPN

Robotic surgery is a well-established method for robotic prostatectomy, and is increasingly being used for RN and donor nephrectomy. An 8-patient study demonstrated the effectiveness of RLPN with the da Vinci[®] surgical system when dealing with complex renal tumours such as hilar tumours, endophytic tumours and multiple tumours. The authors cited 3-D stereoscopic vision, high Degree Of Freedom (DOF) instruments and motion scaling as advantages of the robotic approach, especially during tumour resection and renal reconstruction. It was shown that robotic assistance may allow minimally-invasive NSS in patients who would otherwise be recommended for open or radical nephrectomy [73].

RLPN takes place with the patient in the flank or modified flank position with pneumoperitoneum of 15 mmHg. The two robotic instrument ports and the camera port form a wide “V” with sides of approximately 8 cm. One or two assistant ports are placed near the camera port [73] (see Figure 2.1).

2.2 Ultrasound Image Segmentation

Image segmentation refers to a series of techniques for partitioning an image into non-overlapping regions that are consistent in some way but different or separated from each other. It is recognized as the first essential step of low-level vision. The segmented bodies and their boundaries effectively act as a dense, symbolic representation of significant image features [35].

US images are difficult to segment because of their low Signal-to-Noise Ratio (SNR) [62] as well as artefacts such as shadows and speckle [63]. In addition, acquisition is orientation-dependent and can be distorted by slight variations in




Figure 2.1: Port placement diagram for da Vinci[®] surgery. The robotic instrument ports (blue) and the laparoscopic camera port (green) form a wide V around the region of interest. One assistant port (yellow) is placed inferior to the camera port, and a second may be placed superiorly. The location of the patient’s kidneys is shown for reference, and the superior direction is indicated [73].

the speed of sound in different tissues. However, improvements in US imaging technology have led to increased use in both diagnostic and guidance applications, leading to a greater demand for reliable image segmentation strategies [62].

2.2.1 Segmentation of the Kidney in Ultrasound Images

Kidney segmentation in US is particularly difficult, and rarely performed in a clinical setting. Traditional segmentation methods often fail due to the lack of strong gradients or intensity clusters. Classification based on texture or speckle patterns is also difficult because of the kidney’s heterogeneous appearance in US, but some successful experiments have been reported.

Xie et al. [89] presented a texture-based approach for 2-dimensional (2-D) kidney images that uses an atlas shape model to generate an initial estimate. Textures were extracted by filtering the image with a bank of Gabor filters. After manual placement of the model as the initial segmenting curve, an energy-minimization scheme was used to find an optimal segmentation. The energy function was designed to partition the image into an inside region and an outside region, such that the inside region had high texture similarity and low texture variance, while the outside region had high texture variance. A training pattern was provided for the texture of the inside region. The output had a mean distance of 1 mm from expert manual segmentations, and it was observed that the algorithm was able to detect the kidney boundary correctly even where it did not correspond to a large intensity gradient. However, the algorithm is designed for 2-D segmentation problems and takes 40 s to run for a single US slice. Furthermore, no comparison is made with CT data and it appears that the output may be following the false inner edge boundary.

Wu and Sun [87] developed a similar method to optimize texture similarity within a segmented region, but used dynamic programming to optimize boundary selection between regions instead of a global energy function. Local gradient information and smoothness were considered as well. In contrast to the previous approach, the segmentation process was initialized with a simple ellipse. Segmentation results compared favourably with expert manual segmentations, but again the algorithm and experts appear to be following the incorrect boundary. These segmentations may be useful for US-US registration, but they are unlikely to show enough agreement with CT segmentations to allow reliable registration.

Kidney segmentation papers are less prevalent in the literature, so this section will also consider some segmentation approaches for soft-tissue US images in related applications. Echocardiography is a somewhat different problem as the shape and volume of the heart changes constantly with respect to time. Hence, a large amount of research aims to make use of 2-D+ T or 3-D+ T datasets. US images of the breast have been segmented for tumour detection, but this is quite different from segmenting an organ in the abdomen. This is also true of vascular segmentation in US, where the intensity difference across the vessel walls is greater than seen in soft-tissue images[63, 74].

Prostate segmentation, however, is a fairly close analogue as it involves soft-tissue segmentation. This field has the advantage of Trans-Rectal Ultrasound (TRUS) imaging, which enables high-quality image acquisition due to the proximity of the probe to the Region of Interest (ROI). Many techniques attack the problem by segmenting a series of 2-D TRUS slices.

The success of TRUS imaging has allowed the use of some classical segmentation techniques, including derivative edge detection, nonlinear filtering, neural network-based pixel classification methods, and texture classification methods. These methods are likely to fail with kidney images where we do not have the benefit of TRUS imaging. More recent efforts have made use of shape or speckle models to constrain the search. Successful strategies include edge detection by the IMMPDAF, deformable surfaces using level sets, and discrete dynamic contours. Local phase information has been suggested as a more robust indicator of structures of interest, but no results are reported for kidney or prostate images [63, 74].

2.2.2 The Interacting Multiple-Model Probabilistic Data Association Filter for Edge Detection

The IMMPDAF, described by Kirubarajan et al. [40], combines the IMM approach to hypothesis merging for dynamic multiple model systems as introduced by Blom and Bar-Shalom [14] with Probabilistic Data Association (PDA) [9]. In 2004, a new version of the IMMPDAF was developed for medical image segmentation by Abolmaesumi and Sirouspour using their previous work on a PDA edge filter [1, 2]. Badieli et al. used this result to develop a prostate segmentation method designed for trans-rectal US images [6–8].

In this technique, a Kalman Filter (KF) is used to detect edges while considering both process noise that allows deviation from the predicted trajectory, and measurement noise that allows deviation from the detected edges. Random signals representing process and measurement noise are characterized by Q and R , their respective covariances. The advantage of the IMMPDAF approach is that different system models with distinct Q and R values can interact to ensure that, at any given time, the detected edges are being processed by the model that is most likely to be accurate.

2.2.3 Level Set Methods

Level Set (LS) methods were originally proposed to model natural phenomena with curvature-dependent speed, such as propagating flames and crystal growth [64], but they have since been adapted for image segmentation [19] and implemented in a variety of medical segmentation and registration algorithms [66, 82–84, 92]. Instead of manipulating an explicit boundary model to obtain a segmentation, the boundary is represented by the zero level of an signed-distance function, also known as the implicit surface [22]. The implicit representation of the topology allows contour splitting and merging to be handled naturally, contrary to parametric contour models [47]. The signed-distance function varies with respect to time according to an underlying flow and the curvature of the implicit surface [64].

Caselles et al. [19] built on the work of Osher and Sethian [64] to develop an image segmentation strategy. The underlying flow was given as $g(x) = \frac{1}{1+(\nabla \hat{I})^2}$, where \hat{I} is the Gaussian-smoothed image. Thus, as the implicit surface expands according to curvature, g acts as a stopping function that diminishes the speed of the expansion around the desired edges. Malladi et al. [53] presented a similar strategy that added a reinitialization step and a narrow-band extension option. Reinitialization is done to cope with discontinuities in the velocity map arising from the non-differentiable distance function. The narrow-band option allows lazy extension of the front by only updating the signed-distance function in the neighbourhood of the zero level set. Exclusively updating a subset of the search space imposes a hard limit on the maximum step size, however. Caselles et al. [20] later presented Geodesic Active Contours (GAC) and proved that this was equivalent to solving a special case of the classical snakes method for segmentation. The authors also extended their geodesic model to handle 3-D segmentation.

LS methods have proven themselves useful in a number of segmentation problems, but they rely on distinguishable intensity groups for gradient filtering as in [19] and derivative works, or an alternative boundary detection strategy. This boundary detection step can be difficult in US images of the kidney.

2.2.4 Discrete Dynamic Contours

Discrete Dynamic Contours (DDC) are similar to LS methods in that a boundary is propagated according to some external potential energy function that depends on image features while minimizing a curvature constraint. However, the surface is explicitly represented as a discrete set of points connected by straight lines, and the forces act only on the points. This method is effective but the authors note a strong dependence on the extraction of useful image features [49].

2.2.5 Image Features from Local Phase Information

The Local Energy Model, presented by Morrone and Burr [58], proposes that line and edge features are perceived at points of maximum ‘local energy’ where an image’s Fourier components come into phase with each other. Most feature detection methods search for points of high local energy, but Kovessi [41] showed the feasibility of feature localization by direct calculation of phase congruency, by extending phase calculations to two dimensions and compensating for noise. Image phase congruency has the advantage of being a dimensionless measure of image feature significance that is invariant to illumination and magnification.

This approach was developed into a highly-localized, contrast-invariant edge detection operator that accounts for phase congruency information in multiple orientations using a covariance matrix [41, 42]. Phase information must be preserved using linear-phase filters, and the author selected log Gabor filters to allow an arbitrarily large bandwidth with a zero DC component in the even-symmetric filter. The ratio κ/ω_0 appears in the filter’s transfer function, where ω_0 represents the centre frequency of the filter. This ratio must be held constant for varying ω_0 in order to achieve constant shape over the set of filters.

Being a normalized measure, the phase filter must make some accommodation for image noise to avoid false positives from local noise patterns coming into phase by chance. A noise threshold is set at k standard deviations above the mean noise energy, where k is generally set to a higher value for noisy images.

Phase information has since been tested experimentally for feature detection in a variety of intermodal medical image registration applications including CT to fluoroscopy [25], US to MR [56, 91], and MR T1 to MR T2 [56]. Hacıhaliloglu

et al. [33] used phase information to localize bone surfaces in intraoperative US to guide orthopaedic applications without the use of ionizing radiation. It has also been used effectively to detect boundaries in echocardiography [60, 75].

2.3 Image Registration

Image registration is the process of aligning source and target images that contain complementary information in order to augment our understanding of the scene. A number of review papers on both general image registration and medical image registration are available [23, 36, 37, 52, 94]. Surface-based algorithms are covered in more detail by Audette et al. [4].

Most image registration techniques involve the following four steps: feature detection, feature matching, transform model estimation and image resampling/-transformation. In intensity-based methods, all pixels in the image are considered as features. Feature-based registration methods are well-suited to problems where most of the information is concentrated in a small area of the image [94].

2.3.1 Feature-based Registration

In feature-based registration methods, accurate registration relies on the detection of enough common features in the reference and sensed image. Thus, registration and segmentation problems can be closely intertwined in many applications [52]. Image features include lines, points, regions and surfaces. Line features are often acquired by edge detection methods such as Laplacian of Gaussian filters [54] or the Canny approach [18]. Point features include corners, regions of maximum or minimum local intensity, and local modulus maxima of the wavelet transform. The detection of region features require some form of segmentation [94]. Surfaces are commonly used in medical imaging applications as they tend to be more distinct than landmarks [36]. Some user interaction is often incorporated into feature detection in order to accurately reduce the search space, thus reducing computation time and avoiding some large mismatches [52].

Features are matched according to spatial distribution, symbolic description or correspondences in local intensity. Spatial distribution-based methods are useful for the registration of detected shapes. One of the most-used approaches for regis-

tering 3-D shapes is Iterative Closest Point (ICP) [13]. Symbolic descriptors, such as SIFT [50, 51], can be used to develop fast matching methods based on descriptor similarity [94].

2.3.2 Intensity-based Registration

Intensity-based registration methods often use correlation-like measures of image intensity to evaluate an alignment. Ergo, they are sensitive to image intensity changes and susceptible to nonspecific matching of smooth image regions. These methods are also limited by fairly flat similarity curves and high computational complexity [94]. Furthermore, they can be highly sensitive to a small set of pixels that disagree strongly in intensity, so they may have trouble with slight changes in shape between images or with the addition of unexpected artefacts [36]. Intensity-based methods are however advantageous in that they do not rely on accurate feature detection and thus do not suffer from segmentation errors [78].

It can be advantageous to apply intensity-based methods to derived images that highlight desired features. With this strategy, registration is no longer based on raw image intensities but instead on the intensities of the derived image. For example, a gradient filter can be applied to the images before registration, or the image properties in the frequency domain may be considered using Fourier-based methods. This can allow better time efficiency and robustness to noise [36, 94].

2.3.3 Transform Model Estimation and Image Resampling

After matching, the transformation between images is estimated using either a global or local model. Global models include rigid, similarity and affine transforms. Rigid transforms preserve the size and shape of the original image by allowing only rotations and translations. Similarity transforms differ only by the addition of a scaling factor, which permits a change in size. Affine transforms allow a separate scaling coefficient in each dimension, and can thus represent a wider range of mappings. More complex transformations can also be modelled by polynomial fitting. Local models, by contrast, estimate transformations for subdivisions of the image separately in order to model local deformations. In an elastic registration strategy, image regions between matched features are stretched to fit. In some

cases, the fluid model is used to smoothly handle local deformations [36, 94].

Image resampling and transformation are required to relate the images once the mapping is determined [36, 94]. A fusion step will often follow registration, in which the registered images are combined to produce a new image that integrates the information from both acquisition steps [52].

2.3.4 Validating Registration Results

Registration accuracy is best estimated by Target Registration Error (TRE), by comparing with results from another registration method, or by expert analysis. TRE is calculated as the difference between corresponding point features in the source and target images that were not used to estimate the transformation. These features should be easily localized but separate from the criteria used to determine the registration in order to avoid deceptively low results from overfitting. A gold standard method is an accepted best registration method for a given application that acts as a ground truth for comparison, but any alternate proven method will do. Expert analysis of the registration results is another alternative but is usually limited to a qualitative assessment of performance [36, 94].

2.3.5 Medical Image Registration

In medicine, image registration is often used to integrate multiview, multitemporal or multimodal data in order to produce an augmented image. Multiview analysis allows an expanded FOV or extension into another dimension. Multitemporal analysis highlights the differences between images taken at different times, possibly under different conditions. Different sensors can detect different kinds of information, and multimodal analysis aims to combine images from diverse sources. Alternatively, it can be used to register a scene to a model for localization or comparison purposes. Specific applications include combining multimodal data, observing tumour progress, verifying the effects of treatment, comparing a patient to an anatomical atlas, and surgical guidance [36, 94].

Some steps in medical registration problems can be performed by a calibrated system without an image basis. For instance, hand-held US can be calibrated for registration to an immobilized patient undergoing surgery, or CT or MR image

acquisition. Likewise, robotic surgical tools are calibrated to have a known spatial relation to the laparoscopic camera [52].

One common source of error in medical applications is patient motion during image acquisition. Intensity distortion, where similar tissue appears with different intensities in different parts of the image, can also occur in some modalities [36].

Deformable Registration

Deformable registration methods can be desirable in medical applications due to the large number of factors that can unpredictably affect the observed anatomy. Deformations can occur due to breathing, soft tissue displacement, patient position or pathologic changes [78]. For example, the kidney was observed to move approximately 17 mm with respiration [76]. Deformable methods are also useful for intersubject registration [36].

The use of Finite Element Models (FEMs) has become more prevalent as it allows the structural properties of the anatomy to be incorporated into a biomechanical model. The anatomical structure of interest is represented as a surface mesh where all point elements interact with their neighbours based on the elastic properties of the tissue [4].

Multimodal Registration

The registration of images from different modalities is often desirable in medicine as it allows information from different sources to be integrated in an augmented image. Alternatively, it can be used to bring planning data into a known coordinate frame to guide an intervention.

Feature-based methods are useful here, as reducing multimodal images to their salient features may make them more similar. A priori information should be used where possible to constrain feature detection, feature matching or the transform estimation [94]. Positron Emission Tomography (PET) and Single-Photon Emission Computed Tomography (SPECT) imaging represent a particular set of challenges for multimodal registration as they often exclusively target pathology and thus are sparse in terms of useful markers to use for registration [78].

Multimodal registration by intensity-based methods can be performed by re-

mapping the intensities of one or both images such that their intensity properties are more similar. For example, one can map high intensities to low intensities in CT images to make them resemble MR images, where bone appears dark instead of light [36].

Nowadays, the leading technique for intensity-based methods is Mutual Information (MI) [85], which comes from information theory. Information is commonly measured in terms of Shannon-Weiner entropy, H [77]:

$$H = - \sum_i^n p_i \log p_i$$

where p_i is the probability of symbol i . In MI, the validity of the alignment is determined based on joint entropy measures, or the amount of information in the combined image. If the two images are entirely unrelated, their joint entropy will be the sum of the entropies of the individual images. By contrast, similar and well-aligned images will have reduced information in their combined images as they are less independent of each other. In essence, MI evaluates how well a given mapping from source to target dataset allows prediction of local intensities in the target image from corresponding intensities in the source image. Since the analysis does not require a linear mapping of intensity values, this approach is well-suited to handle multimodal registration problems [36, 94].

2.3.6 Optimization Strategies

All registration strategies consist of a problem statement, a registration paradigm and an optimization method [52]. Optimization strategies for medical image registration problems are normally iterative methods, but some simpler problems can be analyzed for a global optimum.

One issue with iterative optimization methods is that they are sensitive to local minima. In some cases, this may be desired, as the global minimum of the cost function may represent a nonsensical alignment. To avoid stopping at an incorrect local minimum, some algorithms implement a multistart strategy, where the registration is optimized from several different starting positions in order to improve the odds of locating the correct alignment. An alternate solution is multiresolution

registration, where a coarse registration is performed first to target the general area where the finer, high-resolution registration should be performed [36].

Iterative Closest Point

ICP involves alternating matching and transformation steps. Matching consists of finding the nearest model point to each data point. Next, the data points are transformed to minimize the error between corresponding points, usually by least squares optimization. The point sets are then matched again, and this cycle continues until the alignment error reaches some threshold. Quaternions are used to compute the final transformation. The reliance on closest points makes this algorithm particularly sensitive to local minima [4, 13, 36, 93].

Evolutionary Strategies for Registration

While ICP is a valid registration tool [13, 93], it can be sensitive to local minima. An Evolutionary Strategy (ES) may be more robust, as a population of attempts undergo evaluation and selection is based only on a ranking of the population [34]. These strategies minimize an objective function through adaptation, which occurs through repeated preferential selection of random mutations of the system state [71]. In order to achieve faster convergence, modern ESs adapt the search parameters (*i.e.* step size in a given dimension) to the search space with Mutative Strategy parameter Control (MSC).

Covariance Matrix Adaptation (CMA) derandomizes MSC by analyzing the covariance matrix of the mutation distribution, essentially performing Principle Component Analysis (PCA) on previous mutation steps to determine the current mutation distribution. CMA has proven effective as it preserves the robustness of general ESs while improving performance [34]. Although CMA is a generic optimization strategy, it can manipulate registration parameters to optimize alignment if it is given an appropriate cost function.

Unscented Kalman Filter-based Registration

The Unscented Kalman Filter (UKF) is a variation of the Kalman filter used to estimate the state of a non-linear system in the presence of Gaussian noise. A

UKF-based registration algorithm was developed by Moghari and Abolmaesumi [57] that determines the variance of the registration parameters during registration. This algorithm performed better than ICP on datasets perturbed by Gaussian noise. Furthermore, data points are iteratively added to consideration, so convergence may be possible without considering the entire point set.

Coherent Point Drift

Coherent Point Drift (CPD), proposed by Myronenko and Song [61], approaches registration as a probability density estimation problem. The moving point set is represented by a set of Gaussian mixture model centroids that are moved coherently to maximize likelihood while preserving the spatial relationship between point sets. For nonrigid registration, this is accomplished by regularizing the displacement field and determining the optimal transformation with variational calculus. CPD also simultaneously determines point correspondences.

GPU-based techniques

Some modern image registration techniques make use of the highly parallelizable Graphics Processing Unit (GPU), as it generally outperforms the CPU in terms of maximum processing performance and memory bandwidth. GPUs are also equipped with special hardware for linear interpolation operations. However, since they are designed specifically for graphics operations, they are not equipped to handle random memory accesses or double-precision arithmetic. Hence, registration algorithms may require significant redesigns when they are ported to graphics hardware [30].

The optimization stage of registration is generally not a focus of GPU-based research as it is poorly parallelizable. Deformable registration methods are an exception however, as there are many parameters that can be determined in parallel [30].

2.4 Related Work

Intermodal registration techniques for US images did not develop as quickly as those for other modalities. B-mode US images typically have a lower SNR and

more artefacts than other modalities [67]. Generic surface-based registration techniques can be used if both source and target images can be segmented, as shown by Firlie et al. [28]. However, clinical viability requires US segmentation to be automatic, and reliable segmentation methods for US images of the kidney have been elusive thus far (see Section 2.2). The inherent challenges of soft-tissue US segmentation often make intensity-based registration methods a more appealing option [10].

2.4.1 3-D MR to 3-D US Registration of the Brain

Roche et al. [72] developed a registration method for MR and US images of the brain. A bivariate extension of the Correlation Ratio (CR) that considered both MR intensity and gradient magnitude was used with Powell’s optimization scheme to perform rigid registration in 3-D. Registration accuracy was on the order of MR image resolution, but this result is not examined further here as brain images are quite different than soft-tissue images of the abdomen.

2.4.2 3-D MR to 2.5-D US Registration of the Liver

Penney et al. [67] presented a method for registering 3-D pre-operative MR to sparse (2.5-D), intra-operative US slices. The objective of this work was to use information from MR intra-operatively to guide needle placement during RF ablation of liver metastases. Images representing the probability of finding a vascular structure at each voxel are derived from the MR and US images, and these vessel probability images are rigidly registered using a Normalized Cross-Correlation (NCC)-based technique. Results showed an RMS error between 2.3 and 5.5 mm.

For each patient, 15 to 28 US images were acquired at maximum exhalation using a tracked probe. Of these, 10 images were manually selected to best represent the internal structure of the liver. Vessel probability maps were calculated based on image intensity and the presence of sudden dips in intensity. In the MR images, vessel probability was calculated based on intensity alone. MR voxels that fell outside the liver were labelled as missing data and not used in the registration. The authors hypothesize that their method would work equally well with pre-operative CT images.

2.4.3 3-D CT to 2.5-D US Registration of the Kidney

Leroy et al. [44, 45] built on the work cited above to develop a method to register a pre-operative CT volume to a sparse set of tracked intraoperative US images of the kidney with the goal of guiding Percutaneous Renal Puncture (PRP). This group also performed a full guidance experiment using manual segmentation and a tracked needle to demonstrate the feasibility of their clinical goals [59]. Surface-to-surface registration with ICP resulted in average residual distances of 1 mm, 1.4 mm and 3.6° from the surface, centroid and axis, respectively.

The CT volume was preprocessed with a median blur and a bi-directional Sobel gradient in order to highlight boundaries that would be significant in US. The vessel probability approach used by Penney et al. [67] was rejected here as the structure of the kidney parenchyma is less discernable than that of the liver, especially in US images.

US images were first processed with a sticks filter [24] in order to reduce speckle. Next, shadows were detected as scan lines with high correlation where the maximum acoustic interface occurs close to the probe. This information was used to generate a mask to remove shadows from further consideration.

The initial attitude was determined by point-based registration of anatomical landmarks manually chosen from the CT and US images. Next, a Powell-Brent search strategy [68] was used to optimize the CR of the overlapped images. CR measures the functional similarity between matched pixels. CR was selected over MI here because it provided a smoother cost function, lower computational complexity and better adaptability to the Powell-Brent search algorithm.

A single dataset was used, but different sets of 5 US slices out of the collected 100 were used in each trial. The ICP alignment based on manual segmentation was used as a gold standard. Registration took 80 s on a 1.7 GHz processor. Error was measured as the distance between the bronze standard CT mesh and the CT mesh after applying the detected transform in terms of average distance between surface points, surface centroids and angle of inclination of the principal axis. Out of 20 trials, all but 2 converged to between 1 mm and 6 mm error in the centroid, with an average residual distance of 5.36 mm.

2.4.4 Image-guided Laparoscopic Partial Nephrectomy with Fluoroscopy

Baumhauer et al. [11] presented a workflow for image guidance during LPN, and tested it in vitro. Their system incorporates preoperative planning with intraoperative imaging to provide the surgeon with an augmented endoscopic image that highlights cancerous tissue and risk structures.

This approach was tested on porcine kidneys in a Minimally Invasive Surgery (MIS) training unit. Agar nodules were used to represent tumours. In order to account for motion due to breathing or surgical manipulation, custom navigation aids were inserted into the phantom and tracked by a monocular endoscope. A C-arm fluoroscopic imaging device was used to acquire a 3-D cone beam image of the kidney and navigation aids. The navigation aids and kidney surface were segmented semi-automatically from the intra-operative image and registered to the pre-operative planning data using the method proposed by Benincasa et al. [12]. When 50% of the organ surface is covered by the cone beam scan, the registration achieves an accuracy less than 0.5 mm. The aligned planning data also allows virtual 3-D scenes to be generated from arbitrary perspectives.

After this initial registration, the navigation aids are tracked by the endoscopic camera to continuously update the alignment. Each registration used 5 of the 7 implanted navigation aids while the remaining 2 were used for validation. Several registration algorithms were used, with OrthIt producing the best Root Mean-Square (RMS) target visualization error at 1.36 mm.

One limitation of this approach is the requirement that 4 navigation aids must be visible to enable tracking. This could be solved by adding redundant aids or with the use of a model estimation technique, such as a Kalman or particle filter. An undesirable trait of this method is its reliance on fluoroscopic imaging, thus exposing the patient and surgical team to harmful radiation. In vivo trials are planned to further explore this approach.

Chapter 3

Experiment Design

3.1 Data Acquisition

All US data was acquired with a Sonix RP machine and a 4DC7-3/40 handheld probe from Ultrasonix Medical Corporation. The 4DC7-3/40 is a 3-7 MHz mechanical convex curvilinear probe that uses a motor to rotate the transducer array around an axis in a casing. Thus, a 3-D volume is represented as a series of approximately 80 stacked 2-D slices obtained in one sweep of the motor. Adjacent slices were separated by a motor angle, ϕ , of 0.7° .

The pre-operative US data was taken with the patient in the flank position on the CT table with the affected kidney up. An anterior-lateral approach was used as this was seen to provide the best quality images. Depending on the size of the patient, pillows or a foam wedge were placed under their waist in order to mimic the surgical position on a breaking operating table. The patient was told to hold their breath during acquisition of the US volume.

Next, the patient held their position on the table until the end of CT imaging, also with held breath. The CT image was taken approximately 5 minutes after the US volume. The technicians injected 90 cc of Optiray 350 at 2.7 cc/s prior to CT imaging, but manual segmentation confirmed that this does not affect the observed kidney shape or volume. Optiray 350 is a non-ionic, intravenous contrast agent with an iodine concentration of 350 mg/mL. Lower-resolution CT images were also take prior to contrast injection but these were not used as they were more

difficult to segment.

For the intra-operative US data, a posterior-lateral approach was used as the placement of the da Vinci[®] surgical robot made patient access difficult for the ultrasonographer (see Figure 3.1). Breath-hold was imposed on the patient by the anaesthetist to minimize movement during image acquisition. US volumes were taken prior to sterilization, prior to insufflation, and regularly afterwards until kidney mobilization introduced an air boundary that made further images useless.



Figure 3.1: Approximate patient position during surgery. The position of the robot makes US image acquisition difficult for the ultrasonographer. The patient holds a similar position in the pre-operative phase, but the break in the table must be simulated by pillows on the CT table.

In both pre-op and intra-op US imaging, the imaging depth was set between 11 and 15 cm, depending on the patient. All slices were saved as unconverted scan-lines for two or three sets of up to 10 volumes in order to maximize the likelihood of obtaining a useful image.

3.2 Data Preparation

Pre-operative CT images were anonymized and windowed for optimal soft-tissue contrast. The affected kidney was then manually segmented by slice in the 2-D scan plane using Stradwin and interpolated to obtain a 3-D mesh surface for guidance [80, 81].

For the US data, the scanlines for an entire sweep were converted to correctly spaced slices by `Vol2Strad` as described in Appendix A. Manual slicewise segmentation and interpolation was performed in Stradwin to produce a 3-D surface to act as a gold standard for validation [80, 81]. Every fourth slice of the kidney volume is used by the algorithm for edge detection, thus the effective spacing is 2.8° . The remaining slices are ignored to reduce runtime.

3.3 Experimental Setup

The scan-converted US slices and the segmented surface from the prior are loaded into Matlab¹. The operator is asked to select a set of at least twelve initialization through a point-and-click Graphical User Interface (GUI) (see Figure 3.2). These points are used to set the starting point for the registration using an ellipsoid method (see Section 4.1).

Next, edge detection is performed on the US data with guidance from the initial alignment. The detected edges are then used to improve the registration of the segmented prior to the US data. No further input is required from the operator after initialization point selection.

3.4 Validating Registration

Registration results are validated by comparing the amount of overlap between the registered CT segmentation and the gold standard segmentation from US. The success of a given registration can be measured in terms of percent volume error:

$$V.E. = 1 - \frac{V_{intersect}}{V_{union}} \quad (3.1)$$

¹Matlab[®] is a registered trademark of The MathWorks, Inc.

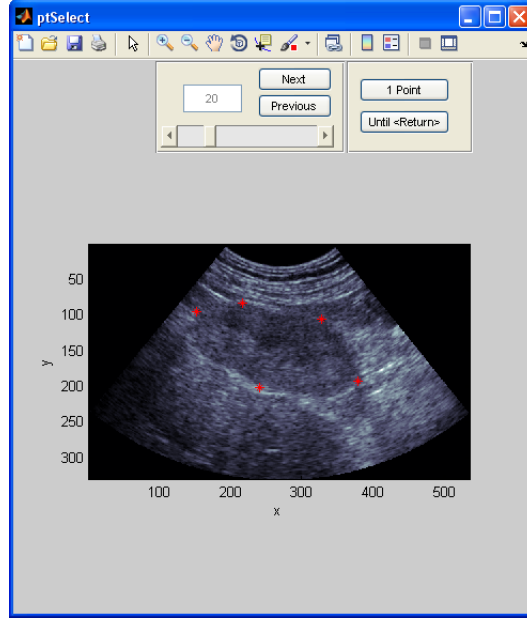


Figure 3.2: Simple GUI for US navigation and initialization point selection.

where $V_{intersect}$ is the volume of the intersection of the two segmented surfaces and V_{union} is the volume of their union. If the two surfaces are identical, their intersection will be equal to their union and the percent volume error will be zero. The relationship between volume error for a kidney segmentation and the set of rigid transformation operations is shown in Figure 3.3.

Target Registration Error (TRE) is defined as the distance between homologous non-fiducial points after registration [29]. TRE is commonly used for validating registration, but the artefacts and noise in US data make it hard to locate any anatomical point features both accurately and precisely. It might be possible to detect manually inserted fiducial points more easily, but this level of invasiveness is not permissible for pre-op procedures. Kingma et al. [39] proposed a validation method that registers CT and US images taken in the same session using an external standoff pad containing fiducial markers visible in both modalities. However, this would not be useful for US images taken during surgery, and the average TRE for kidney features was too high at 14 mm.

We can also validate the registration of two surfaces by comparing the location




Figure 3.3: Volume error relationship with each DOF of rigid motion for a kidney segmentation

of their centroids and the alignment of their principal axes. The centroid acts as a strongly localizable point feature, and the principal axis alignment can be used to estimate rotational errors. Mutual information could have been used as well but it is harder to interpret the significance of this measure.

Chapter 4

Methods

The approaches for solving this problem can usually be broken into three stages: initialization, segmentation and registration (see Figure 4.1). Initialization accepts a prior CT segmentation and a set of points from the user, and produces a rough alignment of the CT segmentation to the US volume. These points are selected manually by the user through a Matlab¹ GUI developed for this purpose (see Figure 3.2). Segmentation uses the initial alignment to guide the detection of edge features in the US images. The registration step aligns the prior segmentation to the detected edge features, returning the transformation that relates the original CT and US volumes.




Figure 4.1: Basic workflow for all approaches. The first two steps are from initialization (Section 4.1) and common to all approaches. The third and fourth steps represent the set of segmentation (Section 4.2) and registration (Section 4.3) methods, respectively.

Several different strategies were implemented for each stage of this algorithm,

¹Matlab[®] is a registered trademark of The MathWorks, Inc.

and they were combined in a modular fashion to obtain a number of registration approaches that can handle the same input. To avoid redundancy, components for each stage of the algorithm are presented separately in Sections 4.1 to 4.3, and their combination is described in Section 4.4.

4.1 Initial Registration of the Prior CT Segmentation to the Initialization Points

The initialization process brings the prior CT segmentation into rough alignment with the US data by registration with the initialization points (see Figure 4.2). A method published by Li and Griffiths [48] is used to return the polynomial representation of an ellipsoid fitted to the points. This polynomial has 10 parameters, so there is a lower bound on the number of initialization points.

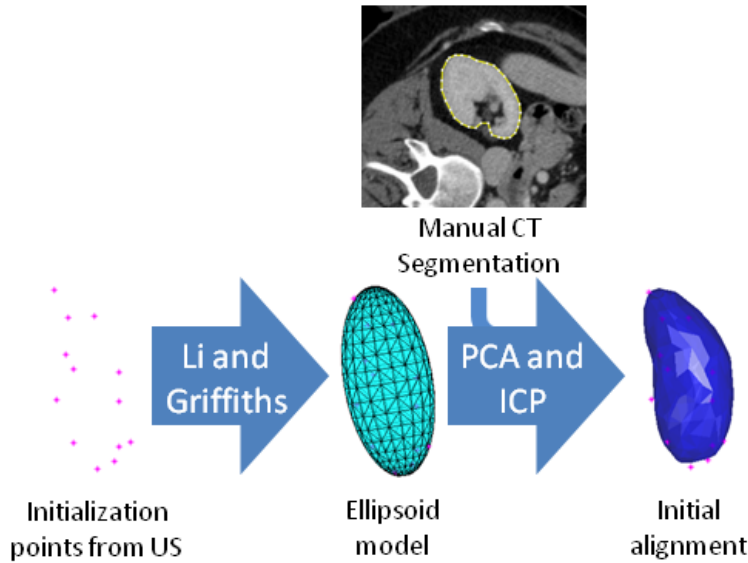


Figure 4.2: Overview of the initialization process

Major translations between the CT segmentation and the ellipsoid model of the US data are corrected by superimposing their centroids. Next, the principal axes of the two datasets are aligned by PCA. During this stage, knowledge of the approximate orientation of the kidney in both modalities is used to ensure that

corresponding faces are oriented in the same direction. The CT segmentation is then registered to the initialization points by ICP.

4.2 Edge Detection and Segmentation Methods

In all cases, the intersection of the CT-segmented prior surface in each US scan plane is used to guide edge detection. A linear interpolator generates a set of N_r equiangular guide points from the polar coordinates of the points in the intersection. The origin for the polar coordinate system is at the centroid of the intersection point set. Figure 4.3 shows the prior contour and search range for a single slice of a typical patient.

Several of the edge detection methods below take a parameter, σ , that defined the scale of features that should be detected by the filter. In early experiments, we observed that a σ -value of five pixels effectively targeted the desired features. In order to preserve compatibility across a variety of platforms, this value was later redefined as a real-world length using the image density of three pixels per millimetre as a conversion factor. Hence, $\sigma = 1.67$ mm in most cases.




Figure 4.3: A set of 120 guide points (yellow) generated from the intersection of the prior with the US scan plane. The origin for interpolation is located at the centre. The search range is contained by the cyan and magenta lines.

4.2.1 Canny Edge Detection

For Canny edge detection, all images are pre-processed with a 2-D Canny filter with $\sigma = 1.67$ mm. Here, “Canny filter” refers to the Matlab² `edge` function that returns a binary image indicating the Canny edges in the input image. The binary output is used as a mask over a gradient image generated using a 2-D gradient filter with $\sigma = 1.67$ mm. Thus, Canny edges are assigned a range of magnitudes and can be weighted appropriately in segmentation.

The steps in this edge detection approach can be seen in Figure 4.4. The Canny filter and gradient filter images are both generated from the original image, and the final output is derived by using the Canny filter’s binary output as a mask over the gradient image. The final output intensities are shown on a log scale for ease of viewing.

4.2.2 Edge Detection Using Phase Congruency

A phase filter obtained from Kovessi [41, 42] is applied to the US images (see Section 2.2.5). Filter parameters were set empirically, using work by Hacıhaliloglu [32, 33] as a reference. The images are processed in 8 different orientations with 3 wavelet scales. The smallest wavelength is set to 5 mm, and successive filters increased by a factor of 2.2. k was set to 7 to handle the inherently noisy US images, and κ/ω_0 was set to 0.4.

This filter normally detects edges by finding regions where phase congruency covariance over all orientations is at a maximum, but here we extract the phase congruency images for each orientation. We then perform a guided edge search, traversing the prior contour and selecting the appropriate orientation image in order to consider only edges that are roughly parallel to the prior segmentation and fall within a 10-mm search range. At steps where no edges are found, weak edges are placed by interpolating the path being followed around the prior contour. A sample phase-filtered image is shown in Figure 4.5.

²Matlab[®] is a registered trademark of The MathWorks, Inc.

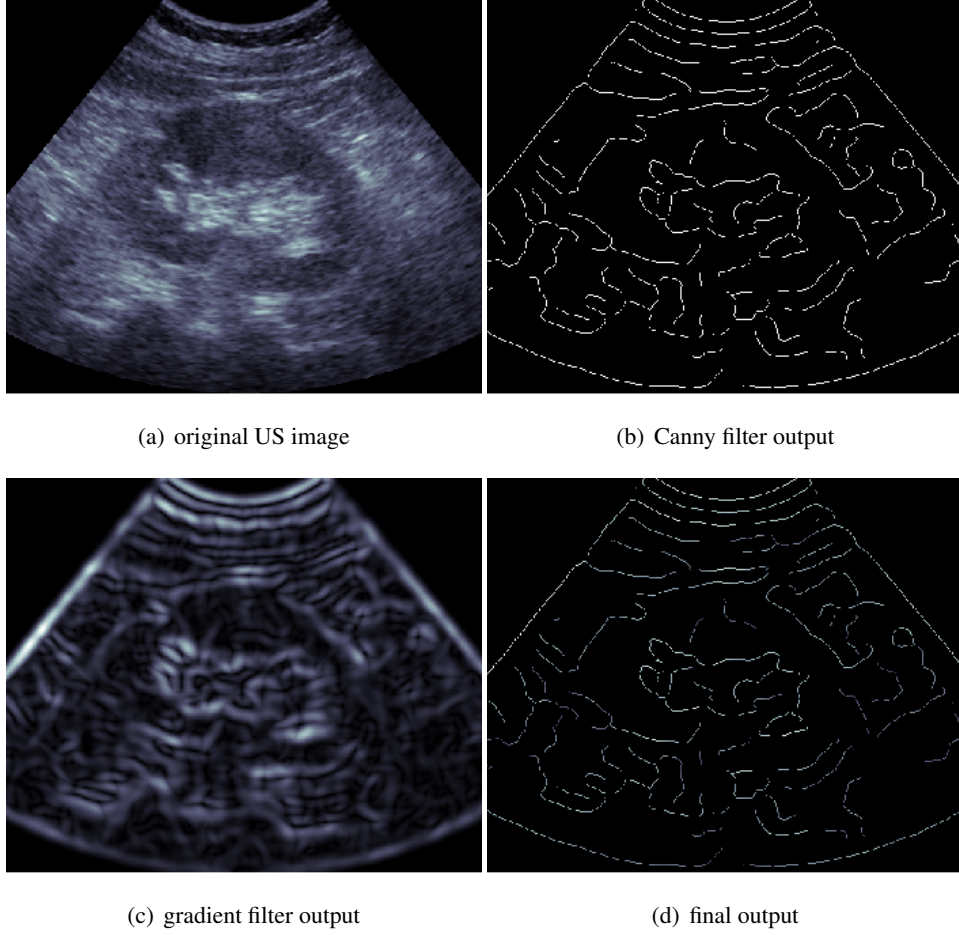


Figure 4.4: Canny edge detection intermediates and output

4.2.3 Radial IMMPDAF for Edge Detection in Ultrasound

The approach taken to kidney segmentation here is similar to the prostate segmentation method published in [1, 7, 8] and further described in [6]. The guide points predicted from the CT segmentation are used to determine a 10-mm radial search range in the neighbourhood of the prior. A star filter passes around the prior contour and selects N_c candidate edge points of maximum radial gradient for each angle θ . Each candidate for a given θ is recorded as a radius, r_c , measured from a central seed point (see Figure 4.6(a)). After applying a median filter of length 1.67




Figure 4.5: Comparison of kidney US images before and after phase-filtering

mm to the radial data, the gradient for each point p in the search range with polar coordinates (r, θ) is calculated according to the following formula:

$$gradient(r, \theta) = \sum_{i=0}^{h-1} I(r+i, \theta) - \sum_{i=1}^h I(r-i, \theta) \quad (4.1)$$

where I refers to the image intensity at the point in question and h is half of the 4-mm filter length expressed in pixels.

An adapted KF estimates x_k , the system state at step k , based on the observed edge intensities and the past output of the algorithm. In this application,

$$x_k = \begin{bmatrix} r_k \\ \dot{r}_k \end{bmatrix} \quad (4.2)$$

where r represents the radial distance of the predicted kidney boundary from the seed point at each angle θ . This segmentation strategy uses a constant velocity model, so if edges are scarce at step k , the KF estimates the new radius $r(k)$ from $r(k-1)$ assuming the same rate of change in radius \dot{r} from step $k-1$. In relation to traditional KF methods, the polar traversal of the guide contours is analogous to

the passage of time. The KF state equations for this system are as follows:

$$\begin{aligned}
\hat{x}(k|k) &= \hat{x}(k|k-1) + W(k)[z(k) - \hat{z}(k|k-1)] \\
\hat{x}(k|k-1) &= F(k-1)\hat{x}(k-1|k-1) \\
\hat{z}(k|k-1) &= H(k)\hat{x}(k|k-1) \\
F(k-1) &= \begin{bmatrix} 1 & \Delta\theta \\ 0 & 1 \end{bmatrix} \\
H(k) &= [1 \quad 0]
\end{aligned} \tag{4.3}$$

where $\hat{x}(k|k)$ is the estimated state at step k given measurements up to step k , $\hat{x}(k|k-1)$ is the estimated state at step k given measurements up to step $k-1$, $W(k)$ is the filter gain, $z(k)$ is the measurement at step k , $\hat{z}(k|k-1)$ is the predicted measurement at step k given measurements up to step $k-1$, and F and H describe our model of the system.

The measurement value z is determined at each step k using a PDA filter. $z(k)$ is given by a weighted average of the candidate radii based on the magnitude of their gradients, g , and their radial distance from $\hat{z}(k|k-1)$, the measurement prediction from the previous step:

$$z(k) = \sum_{i=1}^{N_c} \frac{p_i(k)}{\sum_{i=1}^{N_c} p_i(k)} r_i(k) \tag{4.4}$$

where $r_i(k)$ is the candidate radius i for step k , and $p_i(k)$ is a Gaussian weight function defined by:

$$p_i(k) = \frac{g_i^2(k)}{\sqrt{2\pi S(k)}} \exp\left(-\frac{(r_i(k) - \hat{z}(k|k-1))^2}{2S(k)}\right) \tag{4.5}$$

where $S(k)$ is the innovation covariance [9]. Thus, the weighting function favours points of high gradient that are close to the measurement prediction from step $k-1$.

The KF model incorporates process and measurement noise with covariances Q and R , respectively. A model with low process noise and high measurement noise is likely to default to the model prediction, resulting in smooth boundaries. In contrast, large Q and small R corresponds to an edge-following model that implies confidence in the image quality. The IMMPDAF allows two separate models with




Figure 4.6: Comparison between radial and prior-guided IMMPDAF. Here, $N_c = 5$. In (a), all angles and distances are defined with respect to the seed point and the prior contour is used only to set the search range. In (b), the prior contour defines the search path orientation, the gradient orientation, the search range, and the origin for all distance measurements.

different Q and R values to interact, enabling the algorithm to adapt its estimation strategy to handle various levels of image noise. The smooth model was designed with $Q = 10$ and $R = 20$, and the edge-following model with $Q = 10^5$ and $R = 20$. The transition probability from one model to the other was 25%.

4.2.4 Prior-guided IMMPDAF

Some changes were made to the radial IMMPDAF to better deal with kidney images. The edge search was performed orthogonally to the prior contours with a range of 20 mm. Furthermore, r values were measured with respect to the prior contours instead of a central seed point (see Figure 4.6(b)). Thus, the motion of the contour with respect to k can be anticipated to some degree, and the constant velocity model will default to the shape of the prior contour instead of a circle. The seed point is not used after the guide points are generated. A damping coefficient, c , was used for velocity propagation to prevent overshooting the boundary. So instead of a truly constant velocity model, $\dot{r}(k)$ is predicted to be $c\dot{r}(k-1)$. F as defined in Equation 4.3 is thus modified as shown in Equation 4.6. A damping coefficient of 0.99 was used in this implementation.

$$F(k-1) = \begin{bmatrix} 1 & \Delta\theta \\ 0 & c \end{bmatrix} \quad (4.6)$$

The PDA calculation was altered as well to consider the direction of the orthogonal gradient. This allows the algorithm to target edges that separate inner dark patches from outer light patches, or vice-versa.

4.2.5 IMMPDAF with Maximum Intensity Filter

This version of the IMMPDAF works much the same as the techniques described in Sections 4.2.3 and 4.2.4, although it replaces the radial or orthogonal filter, respectively, with a maximum intensity search. This is not useful for unprocessed US images, but it allows edge intensities from Canny (Section 4.2.1) and phase congruency (Section 4.2.2) filters to be weighted for use in the PDA. In cases where no edges are found, a weak edge is assigned at $r(k-1)$ to support the prediction by the model.

4.2.6 Discarding Canny Edges Far from Prior

While the IMMPDAF and phase congruency filter methods automatically discard edges outside a search range, the Canny filter described in Section 4.2.1 keeps all edges in the image. To discard edges that are unlikely to be useful in registration,

we can traverse the prior contour and only keep those edges that fall within a 10-mm orthogonal search range.

4.2.7 Level Sets and Geodesic Active Contours

Canny edges are detected and weighted in each slice as in Section 4.2.1. The high-confidence initialization points are added as edge features. Next, the prior contour is traversed with an orthogonal maximum intensity search, rejecting edges that fall outside the 20-mm search range. Gaps in the edge contour are filled with weak edges, placed by interpolating the path being followed around the prior contour. Next, the edges are filtered with a 2-D Gaussian with $\sigma=1.67$ mm to create some descent. This edge strength function is then used to generate the edge potential stopping function g by applying Equation 4.7:

$$g = \frac{1}{1 + \nabla \hat{I}'} \quad (4.7)$$

where $\nabla \hat{I}'$ is the edge strength function after edge interpolation and smoothing. The derivation of this stopping function for a sample image is shown in Figure 4.7. The edge potential stopping function, shown for a single US slice plane in Figure 4.7(d), decreases in value from one to zero in the neighbourhood of detected edges in order to halt surface contraction or expansion.

The slice-based edge potential map must then be interpolated to make a 3-D grid to serve as input to GAC [20]. Once the boundaries of the grid are defined, the coordinates of each gridpoint with respect to image row, image column and angle of motor elevation are determined to allow the grid to be populated by linear interpolation of the edge potential slices. The volumetric edge potential stopping function is shown in Figure 4.8 as a series of isosurfaces. The mesh representation of the prior segmentation, obtained by interpolation of the manually-segmented CT contours (see Section 3.2), is used to create a volumetric signed-distance function ϕ_0 that defines the segmented surface implicitly as its zero-level set. ϕ is then allowed to evolve according to the GAC approach for 3-D surfaces defined by Caselles et al. [21]. The evolution was run for 30 expansion intervals of length $dt=0.3$ with the coefficient for advection in the normal direction, v , set to 0.01. Then, evolution was run for 30 contraction intervals with v equal to -0.01. At each

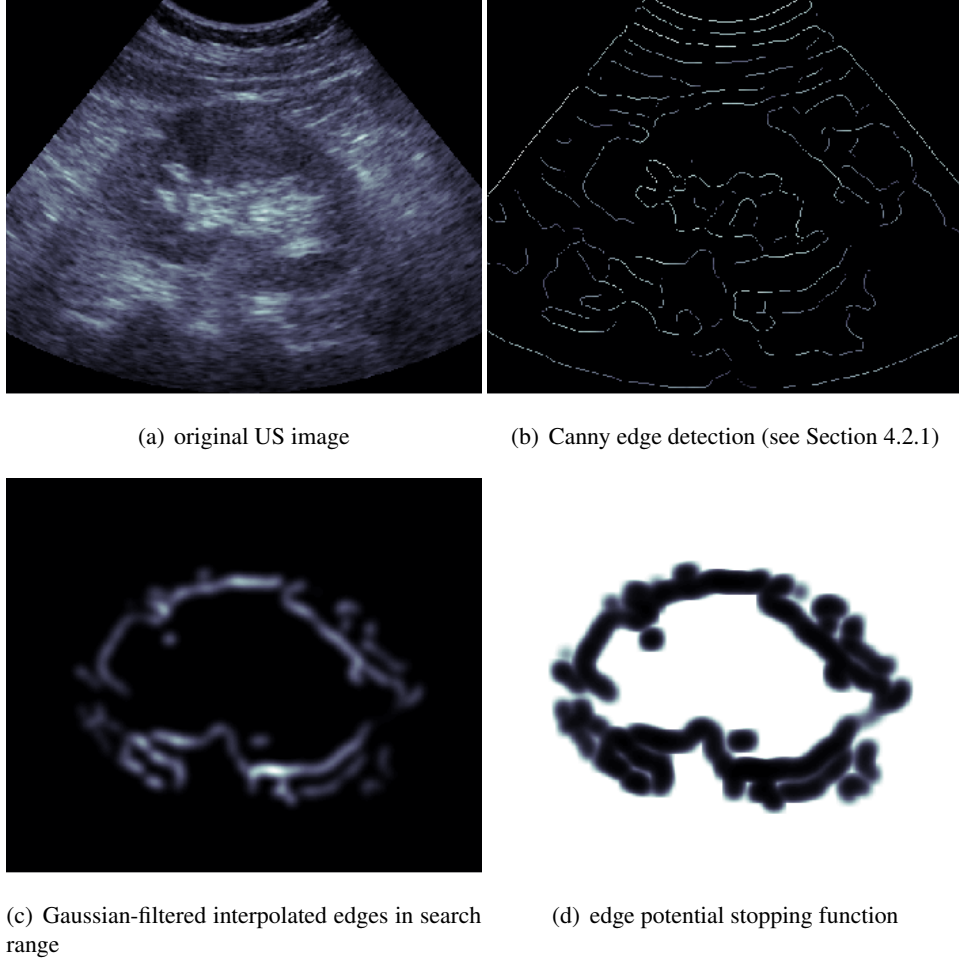


Figure 4.7: Derivation of edge potential stopping function for LS/GAC

interval, the new signed distance function was calculated using:

$$\frac{\partial \phi}{\partial t} = |\nabla \phi| \operatorname{div} \left(g(\hat{I}') \frac{\nabla \phi}{|\nabla \phi|} \right) + v g(\hat{I}') |\nabla \phi| \quad (4.8)$$

where $\nabla \phi$ is the gradient of the signed distance function and $|\nabla \phi|$ is its magnitude. The first term minimizes the curvature of the zero-level set to favour smooth boundaries and attracts the zero-level set to edge potential minima, while the second term performs expansion or contraction constrained by the edge potential stopping func-

tion.

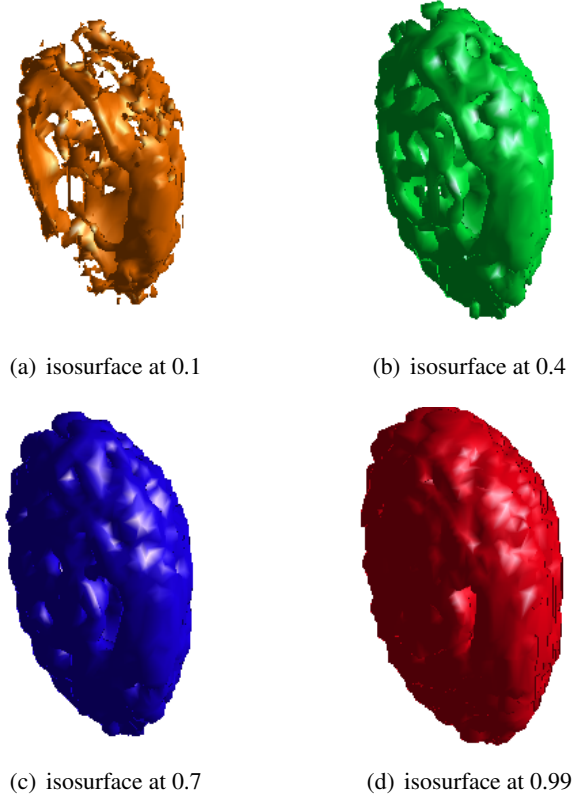


Figure 4.8: Volumetric edge potential stopping function for LS/GAC

Re-initialization is performed to ensure the signed distance function reflects the current position of the moving front. ϕ is updated according to the following equation:

$$\frac{\partial \phi}{\partial t} = \text{Sign}(\phi) |\nabla \phi| \quad (4.9)$$

where:

$$\text{Sign}(\phi) = \frac{\phi}{\sqrt{\phi^2 + |\nabla \phi|^2}}$$

The signed distance function was re-initialized every 7 intervals, as well as at the end of expansion and the end of contraction. After evolution, the registered surface is obtained from the zero-level set of the deformed ϕ function.

4.3 Registration Methods

Registration was performed using ICP, UKF, CMA and CPD strategies.

4.3.1 Iterative Closest Point

The ICP implementation used the L-BFGS [55] optimizer to find a rigid transformation between point sets. The general workings of the ICP function are covered in Section 2.3.6.

4.3.2 Unscented Kalman Filter-based Registration

A UKF-based registration method proposed by Moghari and Abolmaesumi [57] was used in some cases. The specific implementation is described in [69].

4.3.3 CMA-based Surface Registration

In order to reduce sensitivity to local minima, a CMA-based registration method was developed to register the prior segmented surface to an edge representation from US. Two approaches were taken to evaluate the quality of a given alignment, but in either case, the problem was essentially the same. The initial transformation vector was set to:

$$\begin{bmatrix} R_\alpha & R_\beta & R_\gamma & t_x & t_y & t_z \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

with initial standard deviations of 0.1 radians for the rotational terms and 2 mm for the translational terms. CMA was then called to iteratively update the position of the prior after the initial alignment (Section 4.1) with rotation about the prior centroid and translation. Two major cost functions were used to guide the registration:

A. Point-Surface Distance of Candidate Edges

Intuitively, we might consider the 3-D Euclidean distance from candidate edge points to some representation of the prior surface. However, point-point distances between candidate edges and surface vertices can be misleading if the surface is not densely populated by vertices. Thus, this cost function considers point-surface distance for each candidate edge instead.

Point-surface distance is calculated by rapidly finding the closest surface points to all candidate point using a k -D tree [26, 31], then considering the distance between each candidate point and all faces that contain its nearest neighbour [27]. In some special cases this may not be the shortest distance between the point and the surface, but the shortest distance to the set of faces that contain the nearest surface vertex should be a good approximation. The cost function to be minimized is:

$$C = \sum_{i=1}^N g_i d_i \quad (4.10)$$

where N represents the total number of candidate points over all slices, and g_i and d_i represent the gradient and distance of point i , respectively. Thus, a given pose during registration is highly penalized for strong edges that are far from the prior surface.

In order to keep the same k -D tree for all function evaluations, transformations are actually applied to the candidate edge points during iterative refinement. After convergence, the transformation is redefined to map the prior segmentation to the detected edges.

B. Slice-wise Intersection of Surface with Edge Strength Image

A slice-wise technique was later developed with the goal of making the evaluation for a given pose faster and more comprehensive. Here, an edge strength image is filtered with a Gaussian with $\sigma=1.67$ mm to scale edges with respect to distance and create descent. The intersection of the repositioned prior surface with each US scan plane is determined as described at the beginning of Section 4.2. The cost function is then redefined as:

$$C = - \sum_{j=1}^{N_s} \sum_{i=1}^{N_r} I'(x_i, y_i, j) \quad (4.11)$$

where N_s is the total number of slices, N_r is the number of points in the interpolated contour of the segmented prior, and $I'(x_i, y_i, j)$ is the value of the edge image after Gaussian smoothing at pixel i in slice j .

4.3.4 Coherent Point Drift

A CPD implementation by Myronenko and Song [61] was used for nonrigid point registration and determination of correspondences after segmentation by level sets in the deformable registration approach. This method is further described in Section 2.3.6.

4.4 Approaches

All approaches used the same initialization procedure described in Section 4.1. In this section, each approach will be outlined with reference to implementation details in Sections 4.2 and 4.3.

4.4.1 Radial IMMPDAF with ICP Registration

1. Initialization (Section 4.1)
2. Edge detection around prior contour using radial IMMPDAF (Section 4.2.3)
3. ICP rigid registration of prior segmentation to edge contours in 3-D (Section 4.3.1)
4. Perform steps 2 and 3 again to improve alignment

4.4.2 Radial IMMPDAF with UKF-based Registration

1. Initialization (Section 4.1)
2. Edge detection around prior contour using radial IMMPDAF (Section 4.2.3)
3. UKF-based registration of prior segmentation to edge contours (Section 4.3.2)
4. Perform steps 2 and 3 again to improve alignment

4.4.3 IMMPDAF with Canny Filter and ICP Registration

1. Initialization (Section 4.1)
2. Edge detection with Canny filter (Section 4.2.1)

3. Contour selection with IMMPDAF maximum intensity search (Section 4.2.5)
4. ICP rigid registration of prior segmentation to edge contours in 3-D (Section 4.3.1)
5. Perform steps 2 to 4 again to improve alignment

4.4.4 IMMPDAF with Directional Phase Filter and ICP Registration

1. Initialization (Section 4.1)
2. Edge detection with phase congruency filter (Section 4.2.2)
3. Contour selection with IMMPDAF maximum intensity search (Section 4.2.5)
4. ICP rigid registration of prior segmentation to edge contours in 3-D (Section 4.3.1)
5. Perform steps 2 and 4 again to improve alignment

4.4.5 Prior-guided IMMPDAF with ICP Registration

1. Initialization (Section 4.1)
2. Edge detection around prior contour using prior-guided IMMPDAF (Section 4.2.4)
3. ICP rigid registration of prior segmentation to edge contours in 3-D (Section 4.3.1)
4. Perform steps 2 and 3 again to improve alignment

4.4.6 Prior-guided IMMPDAF on Canny Edges with ICP Registration

1. Initialization (Section 4.1)
2. Edge detection with Canny filter (Section 4.2.1)

3. Boundary detection with prior-guided IMMPDAF (Section 4.2.4)
4. ICP rigid registration of prior segmentation to edge contours in 3-D (Section 4.3.1)
5. Perform steps 3 and 4 again to improve alignment

4.4.7 Prior-guided IMMPDAF on Directional Phase Edges with ICP Registration

1. Initialization (Section 4.1)
2. Edge detection with phase congruency filter (Section 4.2.2)
3. Boundary detection with prior-guided IMMPDAF (Section 4.2.4)
4. ICP rigid registration of prior segmentation to edge contours in 3-D (Section 4.3.1)
5. Perform steps 3 and 4 again to improve alignment

4.4.8 CMA Point-Surface Registration with Radial Filter Edge Candidates

1. Initialization (Section 4.1)
2. Detection of candidate edge points around prior contours with star filter from Section 4.2.3
3. CMA point-surface registration to optimize alignment of prior segmentation to candidate edge points in 3-D (Section 4.3.3A)

4.4.9 CMA Slice-wise Registration with Canny Edge Map

1. Initialization (Section 4.1)
2. Edge detection with Canny filter (Section 4.2.1)
3. Discard edges outside range of prior contours (Section 4.2.6)

4. CMA slice-wise registration to optimize alignment of prior segmentation to edge strength images in 3-D (Section 4.3.3B)

4.4.10 CMA Slice-wise Registration with Phase Edge Map

1. Initialization (Section 4.1)
2. Edge detection with phase congruency filter (Section 4.2.2)
3. CMA slice-wise registration to optimize alignment of prior segmentation to edge strength images in 3-D (Section 4.3.3B)

4.4.11 Deformable Registration with LS/GAC and Canny Filter Edges

1. Initialization (Section 4.1)
2. Segmentation of US volume by LS/GAC methods using segmented prior as starting point (Section 4.2.7)
 - (a) Canny edge detection
 - (b) Edge interpolation to generate edge potential stopping function in each slice
 - (c) 3-D interpolation to produce volumetric stopping function
 - (d) LS evolution using GAC to update signed distance function
 - (e) Extraction of deformed surface from zero-level set
3. Use CPD to perform nonrigid registration of prior segmentation to US segmentation (Section 4.3.4)

Chapter 5

Results

5.1 Initialization

The initialization procedure described in Section 4.1 was run 50 times for each patient with a variable number of user-selected points. With eight points or fewer, the ellipsoid was poorly constrained and the initial alignment was outside the algorithm’s capture range. Initialization was generally successful when nine points were selected, resulting in a mean volume error of $27 \pm 8\%$. The alignment improved with larger initialization point sets, but we would like to minimize the requirement for user interaction. Results are shown in Figure 5.1.

5.2 IMMPDAF Methods

5.2.1 Radial IMMPDAF

The radial IMMPDAF approaches showed little to no improvement regardless of the edge detection or registration strategy (see Table 5.1). All implementations except UKF registration showed high variability in volume error plots (see Figure 5.2). On average, the phase filter implementation worked the best out of this group, but it also resulted in worse outcomes for a significant number of trials.

The radial IMMPDAF with UKF registration was unable to converge in all cases, so the final volume error is near the starting point. This algorithm showed

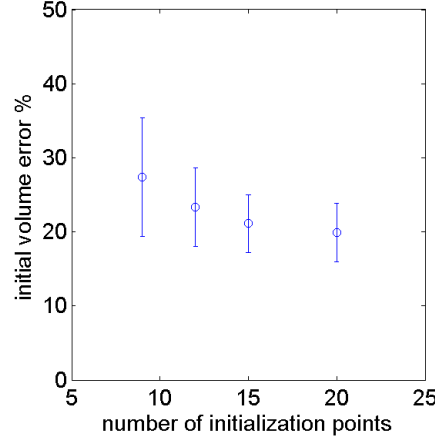


Figure 5.1: Percent volume error after initial alignment for different numbers of initialization points

Approach	Volume Error Ratio
Radial IMMPDAF	0.90 ± 0.19
Radial IMMPDAF with UKF	0.94 ± 0.18
Radial IMMPDAF with Canny Filter	0.85 ± 0.13
Radial IMMPDAF with Phase Filter	0.81 ± 0.15
Prior-guided IMMPDAF	0.80 ± 0.15
Prior-guided IMMPDAF on Canny Edges	0.84 ± 0.14
Prior-guided IMMPDAF on Phase Edges	0.80 ± 0.13

Table 5.1: Ratio of volume error after registration to volume error at initial alignment for radial and prior-guided IMMPDAF methods

modest improvements for poorly initialized volumes but essentially no improvement where initialization was good (see Figure 5.2).

5.2.2 Prior-guided IMMPDAF

While the prior-guided IMMPDAF approaches also show fairly high variability, the mean outcomes are better (see Table 5.1). The gains were least impressive when the filter was applied to Canny images, but good for phase images and direct application to US images. Prior-guided IMMPDAF on phase images performed the best overall, as volume error improved in almost all cases. Volume error plots for

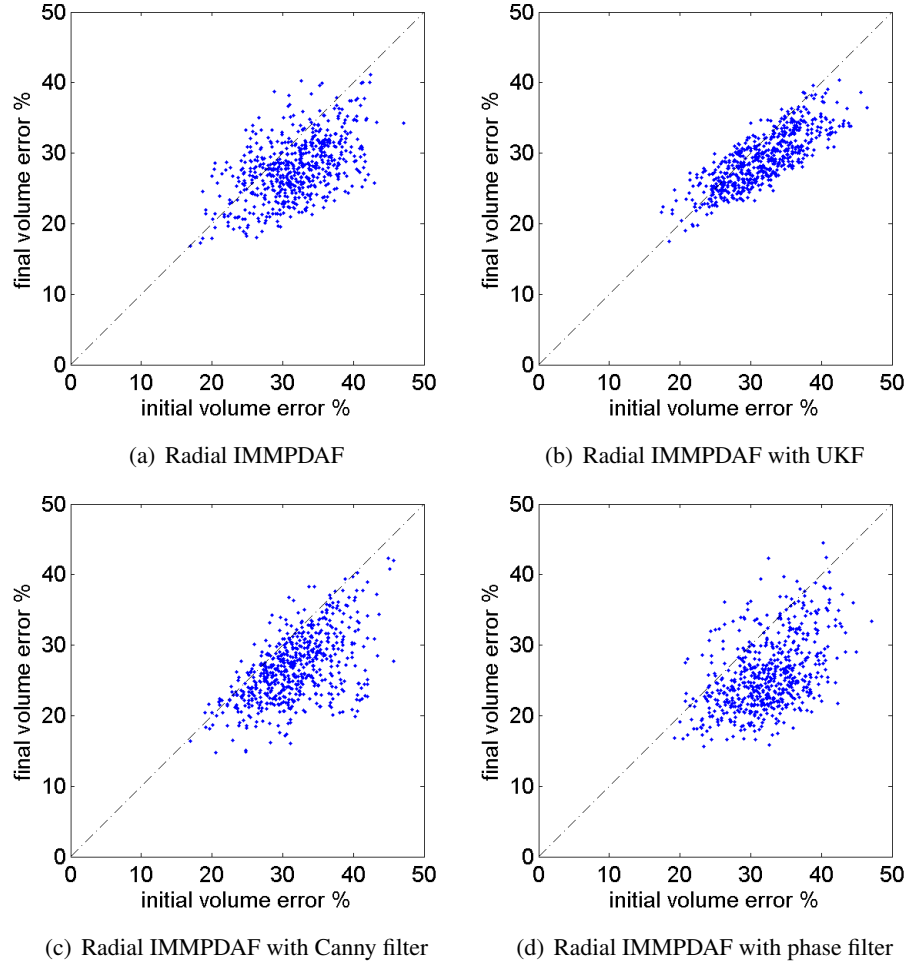


Figure 5.2: Volume error after registration vs. volume error after initial alignment for radial IMPDAF methods

these approaches are shown in Figure 5.3, and patient-specific results are available in Appendix C.

We also considered the distance of the surface centre of gravity or centroid before and after registration. The initial error should not be large if the initialization points are selected well, but we still expect to see some improvement. Results are shown in Figure 5.4.

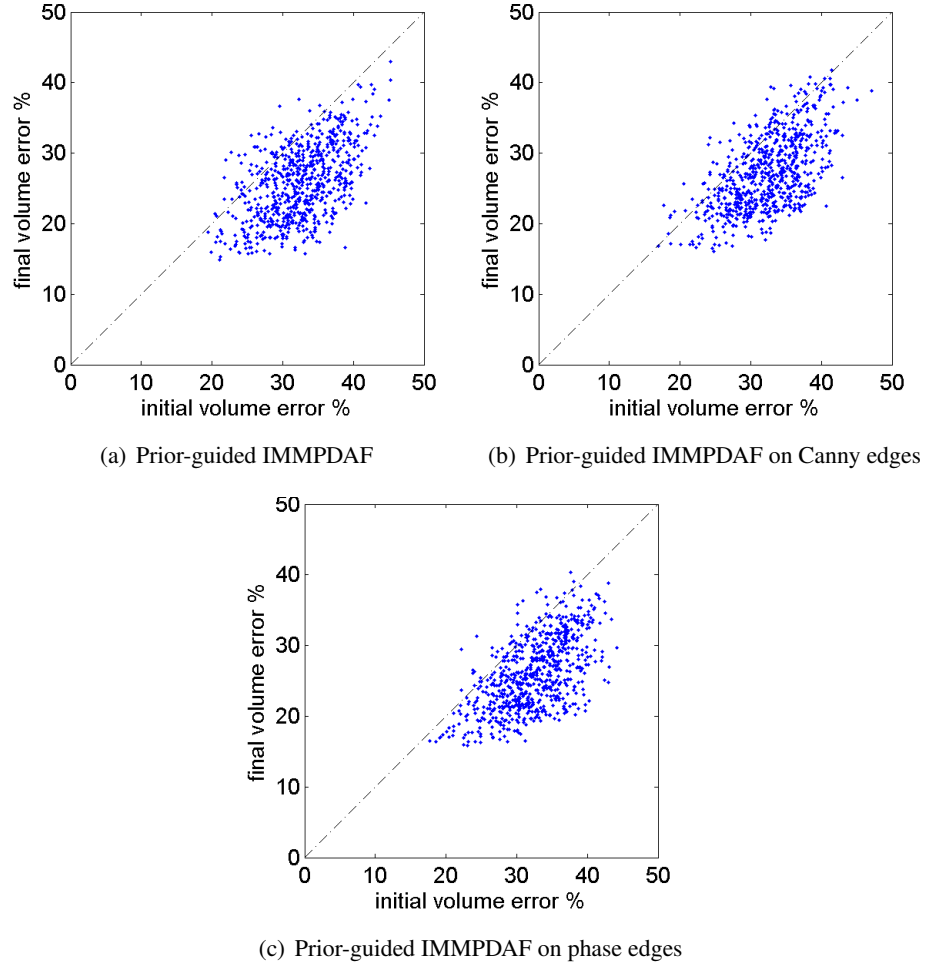


Figure 5.3: Volume error after registration vs. volume error after initial alignment for prior-guided IMMPDAF methods

5.3 CMA Methods

The point-surface CMA approach was only tested with radial filter edges due to its poor performance. Convergence took over thirty minutes and resulted in little to no improvement in the alignment. The slice-based algorithm worked faster but it was highly variable and often led to much worse outcomes. Results are shown in Table 5.2 and Figure 5.5.

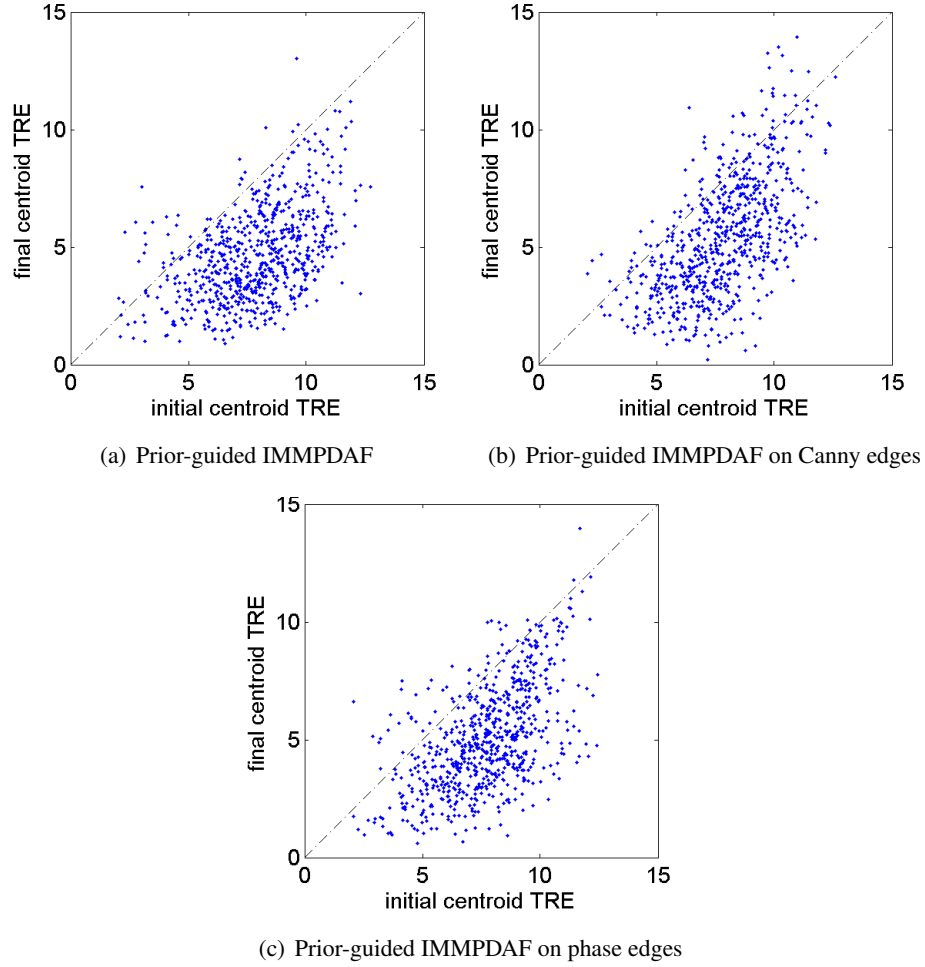


Figure 5.4: Surface centroid error after registration vs. surface centroid error after initial alignment for prior-guided IMMPDAF methods. Centroid error is measured in mm.

5.4 Deformable LS/GAC-based Approach

The deformable registration approach described in this paper did not show significant benefits, as the mean volume error ratio was 0.98 ± 0.07 . The error plot in Figure 5.5(d) shows that there was little change in all cases.

Approach	Volume Error Ratio
Point-surface CMA with Radial Filter Edges	0.89 ± 0.10
Slice-based CMA with Canny Edge Map	1.31 ± 0.45
Slice-based CMA with Phase Edge Map	1.20 ± 0.46

Table 5.2: Ratio of volume error after registration to volume error at initial alignment for point-surface and slice-based CMA methods

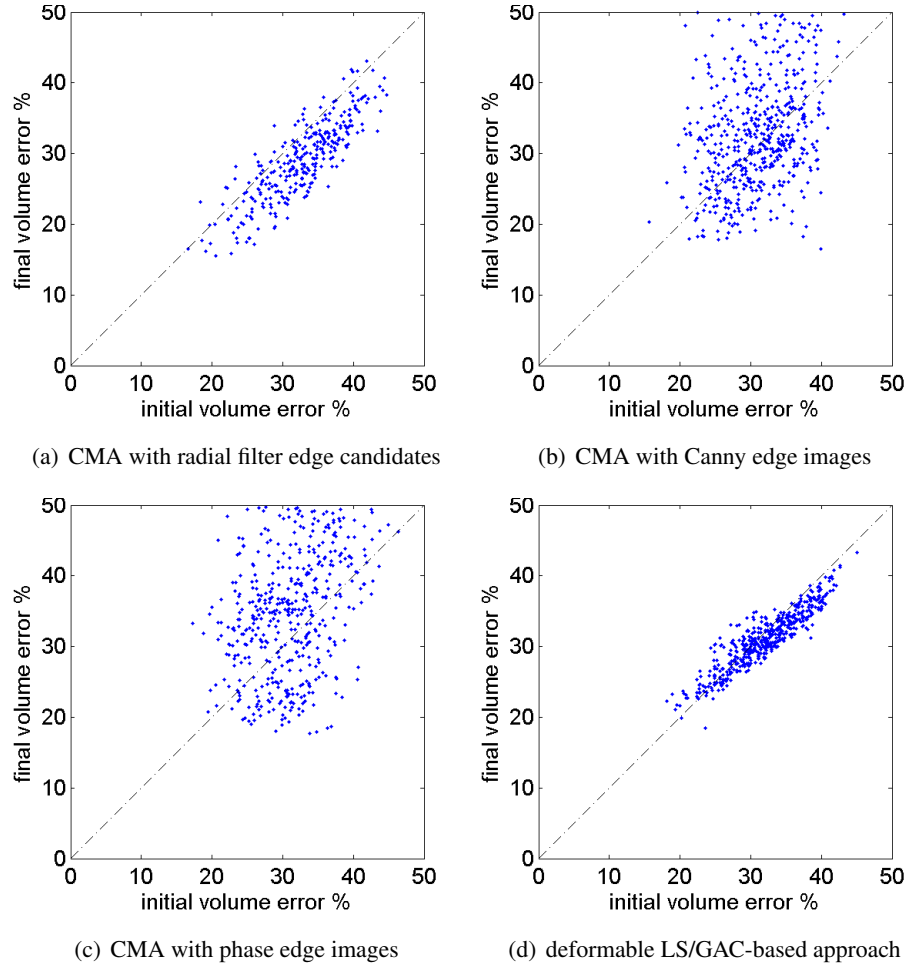


Figure 5.5: Volume error after registration vs. volume error after initial alignment for CMA methods and deformable LS/GAC-based approach

5.5 Intraoperative US Data

After the success of the prior-guided IMMPDAF methods, we decided to test the registration on intraoperative US data. In many cases, the quality of kidney images was too poor to segment due to the ultrasonographer's limited access to the patient or special surgical conditions such as pneumoperitoneum. There were also some patient datasets that were missing either pre-operative or intra-operative data due to logistical reasons. Thus, only one intra-operative patient dataset is examined here. The algorithm performed well (see volume error in Figure 5.6 and surface centroid error in Figure 5.7), but more datasets are needed to confirm this result.

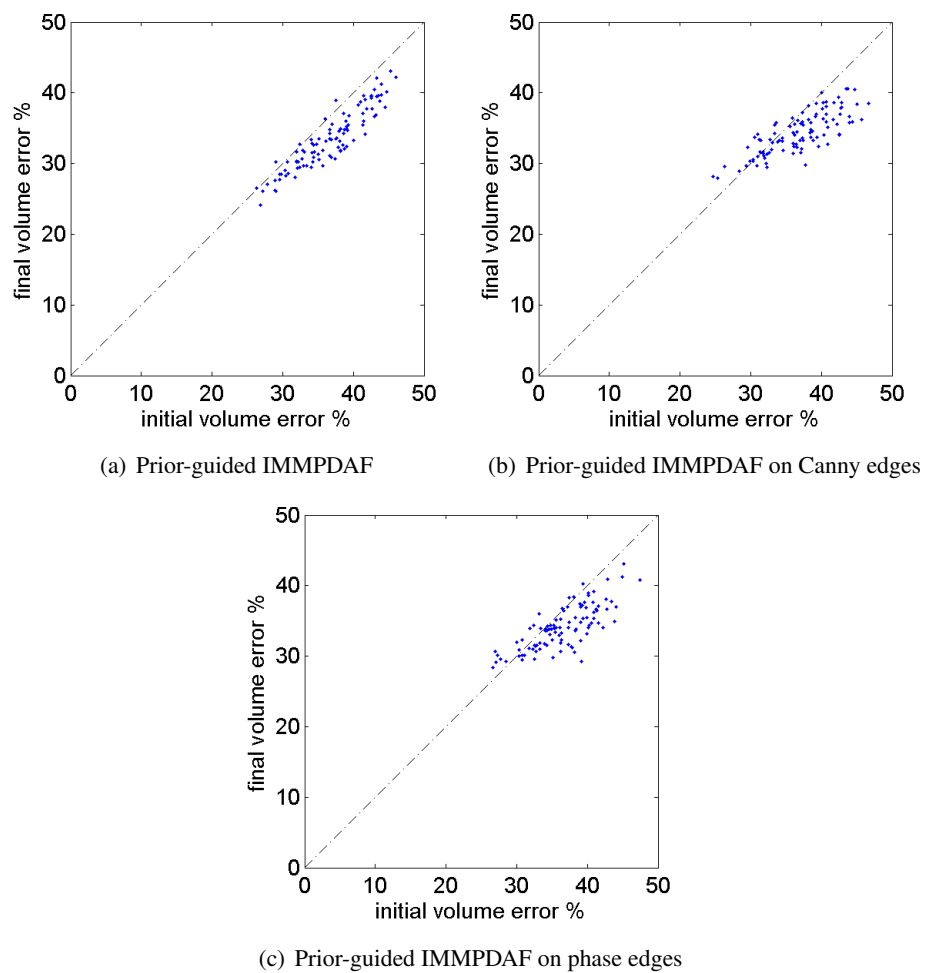


Figure 5.6: Volume error after registration vs. volume error after initial alignment for prior-guided IMMPDAF methods on Patient 5 intra-operative data

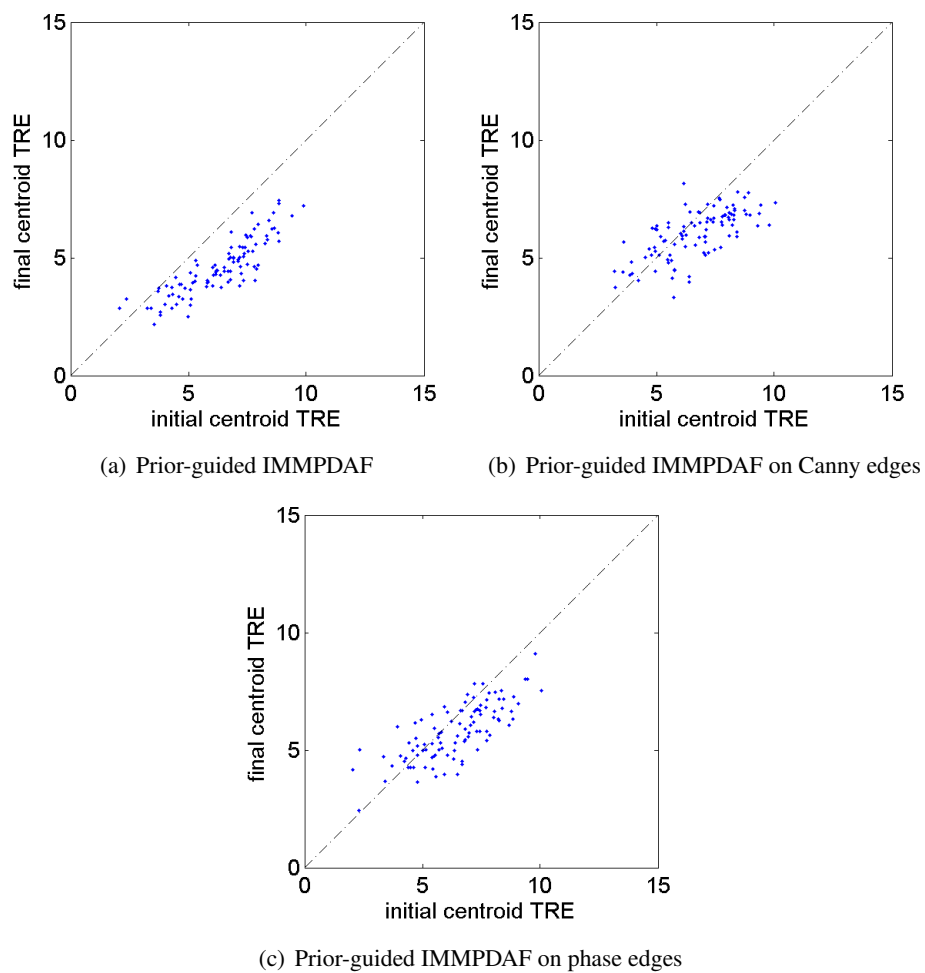


Figure 5.7: Surface centroid error after registration vs. surface centroid error after initial alignment for prior-guided IMMPDAF methods on Patient 5 intra-operative data. Centroid error is measured in mm.

Chapter 6

Discussion

6.1 Understanding the Ultrasound Images

The challenges of US segmentation have already been discussed (Section 2.2), but it should be stressed that this is not a typical image analysis problem. US images indicate the strength of echoes at different distances from the probe, where the strongest echoes occur at interfaces between different tissues. Thus, the detected image in some ways indicates anatomical gradients as opposed to anatomical structures. The aptitude for interface detection combined with the heterogeneity of individual tissues leads us to choose edge detection over pixel classification as our preferred segmentation method.

6.1.1 Anatomical Significance of Visible Edges

Large volume differences between early manual segmentations of the CT and US data suggested that different anatomical edges may be visible in the two modalities (see Table 6.1). In actuality, the strongest edges in US do not represent the true kidney boundary, and the desired edges are quite faint (see Figure 6.1). This poses a problem as traditional methods of edge detection, such as Canny and other gradient-based methods, tend to select the stronger edges when presented with these images.

The desired boundary separates the kidney capsule, which covers the renal parenchyma, from the perirenal fat. Since the nature of US images is to show tissue

Patient	Strong Edge Volume Ratio	True Edge Volume Ratio
3	0.73	1.04
4	0.75	0.96
5	0.83	0.97
6	0.87	1.02
8	0.73	1.08

Table 6.1: Volume ratios between manual segmentations of US and CT data for each patient. The second column shows that segmentations following the strong US edges typically underestimate the volume seen in CT. The third column shows that segmentations that use the true US edges agree with the CT segmentations in volume.

interfaces with high intensity, the fat-capsule and capsule-parenchyma interfaces appear together as a wide echo. Gradient-based search methods typically see the inside of this band, when it is really the leading edge of the band that indicates the kidney’s outer boundary.

On the concave side of the kidney, US beams pass through the interfaces in the reverse order. Without a definitive leading edge to rely on, gold standard segmentations may be more prone to errors in this region.

The magnitude of the gradient edge discrepancies appears to vary with the orientation of the kidney edges with respect to the direction of US beam travel. The smallest error between the strong, “expected” contour and the true boundary occurs at the top of the image, where the kidney boundary is approximately parallel to the direction of beam travel and echoes are strong. On either side of this region, the error increases as the kidney boundary becomes more parallel with the US beams.

We tried detecting the strong edges and predicting the location of the true edges from this approximate contour, but no consistent relationship between inner and outer edges was found. However, reducing the size of the CT prior by 5% allowed reasonable guidance at the strong edges detection stage.

6.1.2 Segmenting the Tumour

In all cases, the tumour boundary was even more difficult to segment in US images than the kidney itself. Even when the tumour is visible in US, it shows poor cor-




Figure 6.1: Kidney boundaries in US

respondence with the boundaries seen in CT. In addition, having extra folds in the segmentation increases the complexity of the surface and produces incorrect local minima for registration. For these reasons, the tumour was only segmented with the kidney for patients 6 and 11, where it made up an integral part of the kidney structure.

6.1.3 Treatment of US data

In early efforts, the US data was interpolated in three dimensions to create an isotropic volumetric dataset. These voxels could then be easily resliced in along the x -, y - or z -axis for segmentation. However, interpolating the volume between slices greatly reduced the image quality, as the 0.7° slice spacing was significantly larger than the axial and lateral resolution.

We considered using longitudinal slices to ensure that all slices are approximately perpendicular to the kidney surface at their intersection. We later rejected this approach when we decided that the original scan planes should be used in analysis to maximize resolution. Furthermore, it is unclear whether the boundary would be easier to detect with longitudinal slices, as the direction of travel of the US beams will be unchanged.

6.2 Initialization

Unconstrained rigid registration problems essentially require global optimization, as convexity occurs only in close proximity to the solution [10]. Optimizing the CT-US alignment over the entire search space is not computationally feasible, so we require the user to select a small set of initialization points. The bulk of the work is still done computationally, however.

We use the ellipsoid fitting method proposed by Li and Griffiths [48] to generate a model from the initialization points because it provides a closed-form solution to fitting that is robust to small errors. Although the ellipsoidal model is a simplistic approximation of the kidney shape, it appears to be a reasonable starting point. Using two overlapping ellipsoids might allow a better representation of the kidney’s asymmetry and concave region, but this does not appear to be necessary. The ellipsoid model is replaced in the next step by a prior CT segmentation.

Originally, the PCA alignment of the prior segmentation to the ellipsoid model was followed by point-based registration of the CT segmentation to the points making up the ellipsoid. Discarding the high-confidence, user-selected points in this fashion leaves the first edge detection step with two unreliable models: the US edges and the ellipsoid alignment. Instead, we register the CT segmentation to the user-selected initialization points after PCA registration. These points are used di-

rectly because they indicate areas where the user is confident in their judgement of the kidney's boundary.

UKF-based registration was considered as an alternative to ICP for the initial alignment of the manual CT segmentation to the user-selected initialization points. However, the UKF algorithm gets stuck in an infinite loop because the limited number of points cannot specify a solution unambiguously.

The ICP step originally used the prior segmentation as the moving point set and the initialization points as the fixed set, as it is more intuitive to consider the CT surface moving towards the points on the US boundary. However, ICP is faster and less prone to local minima when the moving point set is small. Thus, ICP is now used to determine the transformation from the initialization points to the rough PCA registration of the CT segmentation as the latter is over an order of magnitude larger in number than the former. The calculated transformation is then inverted before proceeding.

The results in Section 5.1 show that the initial registration step leads to an initial volume error of $27 \pm 8\%$ when nine initialization points are selected. This rough alignment should be adequate to place the CT-segmented prior into the capture range for automatic segmentation and registration.

6.3 Edge Detection Filters

6.3.1 Canny Edge Filter

The Canny filter in Matlab¹ detects a lot of noise when sigma is low enough to detect all kidney contours. Furthermore, the output condenses edge features to narrow lines, thus ignoring some of the subtle edges that are important in soft tissue US images. Thus far, the low and high edge thresholds have been determined automatically, but manual adjustment of these thresholds is unlikely to make the filter target the true, weaker edges.

Radial or prior-guided search of the filtered images was used to constrain the set of detected edges to those in the neighbourhood of the guide contours. There were cases where candidate edges could not be found at some steps. This issue

¹Matlab[®] is a registered trademark of The MathWorks, Inc.

was first dealt with by reducing the amount of Gaussian blurring, but this led to the detection of many noisy edges. This technique was later refined by accepting the measurement from the previous step in cases where no edges are found.

There were also difficulties when too many candidate edges were found in some radial or prior-guided searches. In order to retain the best edges for consideration, the binary Canny detector output was used as a mask over the Laplacian of the Gaussian of the original image, as illustrated in Figure 4.4. The resulting Canny edges thus had intensities corresponding to the magnitude of their gradients, allowing them to be sorted in a meaningful way.

Although there is no explicit measure of edge directionality in the Canny output, the KF in IMMPDAF methods should reject most edges that are perpendicular to the kidney surface. Initial testing of 2-D edge detection on kidney US images gave promising results for phantom data, but the kidney contour was more difficult to distinguish from other edges in patient data. In CMA and LS/GAC methods, the orientation of the edges is not directly considered, so these approaches may be susceptible to false minima around edges with the wrong orientation.

6.3.2 Phase Congruency Filter

The phase congruency filter was very effective in detecting salient features on the kidney boundary, but it also detected some other features arising from noise or image artefacts. Thus, approaches that were better able to differentiate the desired edges from the background performed well with phase congruency images while others did not. Figure 6.2 shows the phase congruency image in relation to the true kidney boundary from manual segmentation.

6.4 Radial IMMPDAF

The radial filter works well when the 2-D shape to be segmented is approximately circular, as is the case with prostate data. Unfortunately, this does not hold true for kidney images.

The shape of the kidney also limits the effectiveness of the ellipsoid model. Originally, the ellipsoid contours were used to guide the first segmentation. However, the ellipsoid is only useful as a rough approximation of kidney shape, so




Figure 6.2: Correspondence between phase edges and true kidney boundary: the outside edge of the bands detected by the phase filter shows good agreement with the desired contours.

contours from the prior CT segmentation are now used to guide each edge detection step. We considered using the ellipsoid as one model for IMMPDAF and CT as the other, but the CT contours provide better information in most cases.

6.4.1 Performance

The first kidney phantom results, discussed in Appendix B, showed that the edge detection filter was too biased towards smooth curves, as it was overshooting some of the sharper turns in the boundary. Some experimentation with parameter setting eventually led to somewhat better results.

However, it can be seen in Section 5.2.1 that radial IMMPDAF is not effective for edge selection in this problem. The constant velocity assumption with radial measurements is better applied to more circular targets and thus is not a good strategy for segmenting the kidney. Since we are already making use of a segmented prior, the model should be improved to integrate this information.

6.5 Prior-guided IMMPDAF

The prior-guided IMMPDAF was designed to make up for some of the shortcomings of the radial method. Primarily, it replaces the flawed model of the kidney as a circle with actual knowledge of the kidney shape from the prior CT segmentation. The IMMPDAF mechanism is essentially the same, however we consider the orthogonal distance from the prior contour instead of the radial distance from the contour centre.

6.5.1 Velocity damping

In initial results, the output contours were less smooth than ideal because the constant velocity model causes any deviation from the model to be amplified in subsequent steps. Increasing the reliance on the model by adjusting either the noise covariances or the edge weights only led to correct edges being ignored.

Rather than require the filter to stay close to the model, we can modify the model to assume the velocity $\dot{r}(k)$ will be of slightly smaller magnitude than that in the previous step. Thus, if $r(k)$ is somewhat greater than $r(k-1)$, we reject the assumption that r continues to increase at the same rate $\dot{r}(k)$ until proven otherwise. Instead, we propagate \dot{r} by multiplying by a damping coefficient. Trial and error suggests that a factor of 0.99 adequately reduces noise without preventing the filter from responding to detected edges (Figure 6.3).

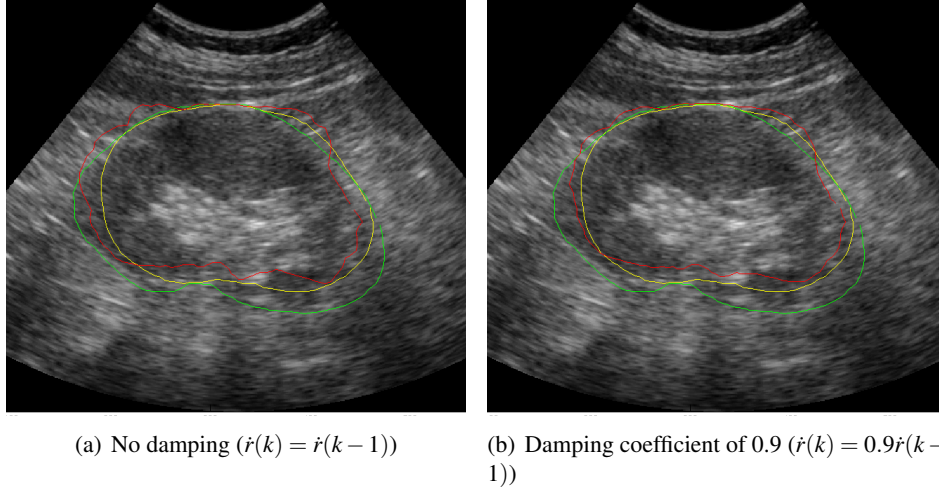


Figure 6.3: Kidney contours for patient segmentation trials performed with and without velocity damping. The yellow line represents the guide contour (segmented prior), the red line represents the output and the green line represents the gold standard (true kidney boundary).

6.5.2 Directional edge filtering

One simple benefit of the prior-guided approach over the radial IMMPDAF is that edge searching is performed orthogonally. An orthogonal edge search more accurately targets edges parallel to the kidney boundary compared to radial edge searching. Furthermore, the effective search range reflects the provided parameters more accurately and consistently as it is measured orthogonal to the prior kidney boundary.

We can also exploit the consistent orientation of step edges separating areas of different intensities. The most reliably detected edges in kidney images are located at the interface of the dark region inside the kidney and the bright boundary that surrounds it, slightly inside of the true kidney boundary. Thus, by enforcing that edge candidates are oriented with the lower-intensity region on the inside, strong edges with the opposite orientation are easily eliminated. Increasing the filter selectivity in this way produced better results with patient data and increased filter robustness when dealing with a larger search range (Figure 6.4).

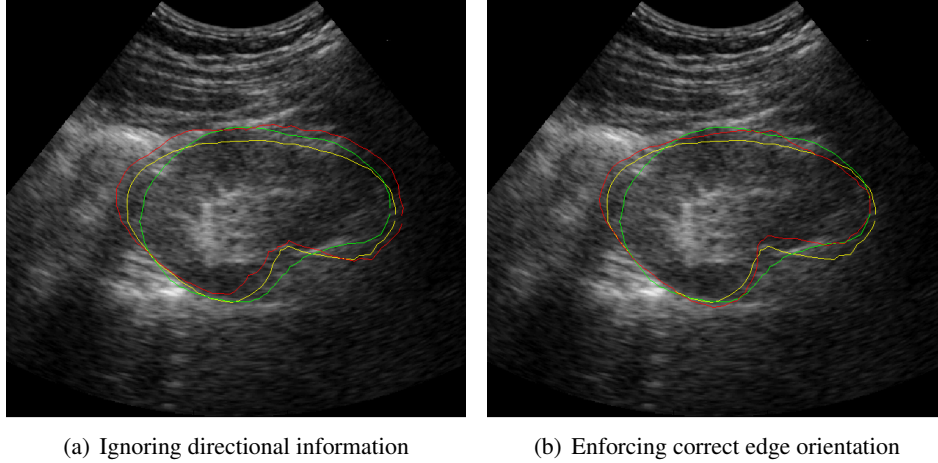


Figure 6.4: Kidney contours for trials performed on patient 3 with and without directional selectivity. The yellow line represents the guide contour (segmented prior), the red line represents the output and the green line represents the gold standard (true kidney boundary).

6.5.3 Registration Based on Inner Edges

Before knowing how to detect the true kidney boundary in US, we attempted to align the prior segmentation to the detected inner edges, relying on the strong correspondence between the manual CT and US segmentations. It may be possible to identify the true edges based on the location of the stronger inner edges, or to perform kidney registration based on the inner edges alone. For example, the prior surface can be scaled down by 5% to guide inner edge detection, and returned to its original scale after registration. This works to some degree but it is not as accurate as detecting the true kidney boundary.

The unrefined edges are noisy, so thresholding based on distance from probe or the edge intensity near the reported contour point was tried to eliminate false candidates. Removing weak edges improved the outcome only slightly, and the distance from the probe turned out to be an unreliable measure of confidence in different datasets. ICP, rigid CPD and point-surface CMA optimization all failed to improve the initial alignment with these edges. In the end it was decided that

that even weak contour points from the prior-guided IMMPDAF are useful in registration as the filter is designed to make the best possible prediction based on the available edge information and the model.

6.5.4 Canny Edge Maps

In all other methods, the quality of detected edge candidates improved with the use of a Canny filter instead of the star filter. However, the prior-guided IMMPDAF had its worst performance when applied to Canny edge maps. Segmentation results from Canny edges may be improved when the prior is used as the model, but with Canny images the filter cannot take advantage of edge direction information.

6.5.5 Phase Congruency Images

Applying the prior-guided IMMPDAF to phase congruency images worked especially well. Phase-processing produces high-intensity bands around parts of the kidney contour, and the outside of these bands roughly corresponds to the true kidney edges, as seen in Figure 6.2. Thus, IMMPDAF can pick up the desired edges with directional filtering by traversing the phase congruency image along the prior contour and searching for light-to-dark images going from inside to outside.

Since the boundary of the US image is identified as a strong feature by the phase congruency filter, it became necessary to explicitly track the image boundary and ignore images in close proximity. This was added to the prior-guided IMMPDAF code in response to some cases where the surface would migrate to the edge of the image.

6.5.6 Performance

The prior-guided IMMPDAF performed the best out of all examined methods, especially when using phase congruency images (see Figure 5.3). Phase-filtering is the slowest step in the algorithm at approximately four seconds per slice, but it may be possible to reduce this time by optimizing the phase filtering code for speed.

The centroid error was reduced in the vast majority of cases (see Figure 5.4), but since this is a segmentation-dependent measure, it does not accurately represent the confidence of the registration in a clinical sense. We can also consider the

distribution of errors over the kidney surface. It can be seen from the surface error plots in Figures C.7 and C.8 that a large amount of the alignment error occurs on the concave side of the kidney. This surface is poorly defined, so we might be underestimating the quality of the registration. Alternate error measures would be useful to determine whether this is true.

It is difficult to evaluate the effectiveness of this algorithm on intra-operative data since only one patient set was examined. Although intra-operative images are, in theory, no more difficult to process than pre-operative images, this is not the case in practise. During surgery, the ultrasonographer's access to the patient is limited so it is difficult to obtain optimal images.

Furthermore, slight changes in patient position or physiology between the pre-operative data acquisition phase and the surgical procedure can invalidate the assumption that the CT to US transformation is rigid. Nonrigidity also reduces the correctness of the prior-guidance model. This algorithm is likely to show better performance for intra-operative datasets if it is extended to find a nonrigid registration.

6.6 Deformable LS/GAC-based Approach

This image segmentation approach has been proven to work well for some problems, but its effectiveness is dependent on the edge detection method used. Here, edges were detected with the Canny method described in Section 4.2.1, which has some shortcomings for kidney US data (see Section 6.3.1).

Besides the edge detection approach, there are more sophisticated LS evolution methods available that may improve the accuracy of the surface propagation. The practice of re-initialization of the signed distance function has been criticized for its lack of theoretical basis and for causing unexpected repositioning of the zero-level set. Li et al. [46] presented a LS evolution strategy that uses an extra energy term to control deviation from the signed distance function, thus eliminating the need to re-initialize.

It is difficult to evaluate the performance of this LS/GAC implementation since we now know that it was likely misguided by the Canny filter used for edge detection. Regardless, the deformable approach to segmentation described here was not

successful.

6.7 Registration Methods

6.7.1 Iterative Closest Point

Despite its susceptibility to local minima, ICP is a common approach to 3-D point cloud registration problems. It was used as the default registration method in our research and compared with other strategies.

Shape Differences

ICP can be sensitive to differences in shape between US and CT segmentations. The kidney vasculature on the concave side makes this surface poorly defined in US and CT, thus leading to potential shape differences. In these cases, rigid ICP registration led to some misalignments that compensated for the disagreement between segmentations on the concave side.

One proposed solution was to use only the convex part of the kidney for registration, or to use a convex hull representation of the kidney surface. The problem with this approach is that a significant amount of shape information is discarded. Furthermore, the result of ICP alignment would depend heavily on the set of points that was selected to represent the convex part of the kidney. Poor correspondence between the sectors identified as convex could result in large rotational errors. The mismatches were eventually corrected by adjusting the manual segmentation strategy to assume minimum concavity.

Performance

ICP was the most reliable registration method used in this study, possibly because of its narrow search range and sensitivity to local minima. A lot of case-specific information is incorporated in the initial alignment from the original scan orientations and the initialization points, so it is appropriate to limit the search range to some degree. US artefacts tend to create false feature positives that can mislead registration algorithms that are not anchored as tightly. Thus, ICP was a good strategy to incrementally improve the alignment, provided that kidney features were detected

accurately.

6.7.2 UKF-based Registration

We expected UKF-based registration to perform better than ICP as it excludes outliers. However, in testing it was seen that the kidney’s approximate symmetry in some directions made the alignment a difficult problem. In all cases, the algorithm’s error metric fluctuated about the initial value without converging until the maximum number of iterations was reached. Thus, convergence did not occur with UKF registration using the original parameters set by Moghari and Abolmaesumi [57]. Further experimentation is required to determine the parameters for this specific application.

6.7.3 Covariance Matrix Adaptation

CMA registration strategies were considered for their robustness to local minima. The first cost function presented here is based on point-surface distances of the edge candidate points to the aligned surface. This method was rejected because it did not have a smooth descent to the desired minimum. The poor behaviour of the cost function is likely due to the fact that only the strongest candidate edges were considered in the registration, thus weaker edges were ignored. This approach was also slow as it had to iterate through point-surface distance calculations for all candidate edge points for each cost function evaluation. However, this registration algorithm produced modest improvements in alignment on average.

The slice-based version of the CMA approach was faster and more complete in its evaluation of edges, but registration was frequently unstable. Instead of aligning itself with the desired edge points, the prior segmentation would migrate to the greatest concentration of bright edges in the US volume, whether they made a kidney shape or not. Due to this instability, the slice-based CMA approaches examined in this report made the alignment worse on average.

In order to enforce the correct kidney shape in the detected edges, we considered using a modified cost function with a Difference-Of-Gaussian (DoG) filter orthogonal to the boundary of the segmented prior. This filter was used by Lowe [51] to approximate a Laplacian of Gaussian operator. The motivation for this selection

was that, when properly oriented, the negative tails would decrease the appeal of randomly-oriented bright spots, instead favouring a bright kidney boundary on a dark background. This technique might work well with a good initial alignment if the desired edges are visible, but this is seldom the case and the negative tails of the DoG filter prevent a smooth descent.

6.7.4 Other Registration Options

A volume-minimizing registration was also considered to bring the prior segmentation into alignment with the detected edges. However, the demanding computational requirements for determining the percent of the volumes of two 3-D segmentations that overlaps make this approach much slower than ICP or any of the other registration methods considered.

Chapter 7

Conclusions

LPN is a challenging surgical procedure for tumour excision that could benefit from advanced image guidance techniques. Since practical intraoperative imaging modalities are generally harmful (fluoroscopy) or prone to artefacts and noise (US), we would like to bring pre-operative CT data into an intra-operative coordinate system. Incorporating pre-operative data would also allow the integration of surgical planning data. We choose to register the CT data to intraoperative back US because it is safe, real-time and inexpensive.

This thesis presents a semi-automatic approach to this registration problem, using edge features common to US and CT images of the kidney. The proposed algorithm effectively performs a full segmentation of the kidney in US in order to determine a target surface for registration of the CT. Numerous segmentation and registration techniques were explored and tested on images from nephrectomy patients.

Over the course of this work, we learned that the true boundary of the kidney capsule is not detected well by common edge detection filters. However, phase congruency filters were able to detect the structures of interest. A specialized IMMPDAF was then used to isolate the boundaries to complete the segmentation. We conclude that applying the prior-guided IMMPDAF to phase-filtered kidney images in US produces an effective segmentation for feature-based registration. ICP showed the best performance out of the registration methods examined. The results of a small intra-operative registration experiment cannot confirm whether

the assumption of rigidity is valid for intra-operative data.

We now revisit the contributions outlined in the introduction:

- A joint segmentation and registration algorithm designed to bring a pre-operative plan into alignment with intra-operative back US during LPN.
- A rough clinical workflow, from initialization to registration, that could be incorporated into medical user interfaces and intraoperative visualization tools.
- Trials with eight patient datasets to evaluate the effectiveness of this algorithm on clinical data.
- New insight into kidney segmentation in US images that conflicts with some segmentation papers but agrees with the kidney shape and volume seen in CT.
- An extension of the IMMPDAF used for prostate segmentation that makes use of an a priori model along with other domain-specific information to create a specialized kidney segmentation approach.
- The application of a phase congruency filter to kidney images and an evaluation of its performance.
- An exploration of other popular edge detection and registration techniques in the context of the kidney registration problem for LPN.

7.1 Future Work

7.1.1 Improving the Interface

A more elaborate point-selection GUI with arbitray reslicing and contour plotting in linked 2-D and 3-D views was proposed to replace the existing interface. For example, three windows could show the US slice, the CT segmentation and the US-CT alignment at that slice. The initial alignment could even be computed on-the-fly as point selection occurs. This would indicate to the operator when a successful initial alignment has been achieved and makes initial misalignments more apparent.

After registration, some confidence measure of the registration could be provided to tell the operator how much they can trust the result. For example, we could use a confidence measure based on the distance of the initialization points from the registered surface, as these points are not used beyond the initial alignment step.

7.1.2 Extending the IMMPDAF

One of the limitations of the IMM-PDA/Kalman filter is that it only accounts for “past” information during traversal of the data, when in fact all the information is available. If the edges rankings were pre-computed for all steps, we could incorporate “future” edge information in our measurement value by including these edge candidates in the PDA step at a reduced weight. Implementing a small lookahead may help bridge the gap smoothly between strong edge candidates that are separated by several steps.

If all edges are pre-computed, we can go beyond a simple lookahead for each slice. The PDA step could include edges from neighbouring slices in its weighting of edge candidates. This may require edge processing of additional slices, but it would enforce consistency of the segmented surface in 3-D.

7.1.3 Realistic Deformable Registration

With the success of kidney segmentation in US using the IMMPDAF, an obvious extension of this work is to discard the assumption of rigidity and perform deformable registration of the prior segmentation to the detected surface. Some unpublished work from our lab examines realistic surface-to-surface registration using elasticity parameters based on the biomechanical properties of the tissue. If we interpolate the edges from each slice to generate a closed surface, nonrigid surface registration would allow this algorithm to overcome tissue deformations due to patient position, tumour growth or manipulation during surgery.

7.1.4 Improved Error Estimation

One drawback of working with kidney data is that there are no anatomical landmarks that can be precisely localized in both US and CT. This lack of reliable point targets makes it difficult to estimate TRE and provide a useful confidence measure

for the surgeon.

One solution to this issue would be to place fiducial markers that can be seen in both US and CT on the kidney. This is not practical for patient studies, but it could be performed with an animal model. This approach has the advantage of using artificial markers that can be designed for optimal visibility and localizability.

Alternatively, we could use Doppler US imaging to identify the renal artery by highlighting bloodflow. A normal renal artery has a diameter on the order of 5 mm [5], but we could likely identify the centre of this artery with better accuracy. Integrating Doppler imaging with 3-D B-mode US would be ideal, but if this is not possible we could simply use a 2-D Doppler image to identify the artery in the US volume. 3-D Doppler imaging may be useful beyond error estimation as it would provide some insight into vessel structure that could further relate the US images to a prior segmentation.

7.1.5 Integration with Surgical Devices

In order for this registration approach to be clinically viable, it must be integrated with medical interfaces and surgical visualization tools. One of the OR staff will have to make point selections to initialize the algorithm. Since the proposed algorithm registers the pre-operative plan to back US, further registration or calibration will be required to bring the planning data into the coordinate frame of the laparoscopic camera. Implementations using a GPU or other specialized hardware may be developed to reduce runtime.

Once the transformation between the preoperative CT data and the camera is known, the preoperative plan can be integrated directly in the da Vinci[®] console to provide real-time image guidance. For guidance to be effective, techniques from the domain of Augmented Reality (AR) and information visualization will be required to draw the appropriate overlays on the surgeon's FOV without creating too much clutter. The desired image guidance framework may be years away, but this thesis presents a solution to one part of this problem.

Bibliography

- [1] P. Abolmaesumi and M. Sirouspour. An interacting multiple model probabilistic data association filter for cavity boundary extraction from ultrasound images. *IEEE Transactions on Medical Imaging*, 23(6):772–784, 2004. → pages 10, 32
- [2] P. Abolmaesumi, M. Sirouspour, and S. Salcudean. Real-time extraction of carotid artery contours from ultrasound images. In *IEEE Symposium on Computer-Based Medical Systems*, pages 181–186. IEEE, 2000. → pages 10
- [3] M. Aron, G. Haber, and I. Gill. Laparoscopic partial nephrectomy. *British Journal of Urology International*, 99(5b):1258–1263, 2007. → pages 5, 6
- [4] M. Audette, F. Ferrie, and T. Peters. An algorithmic overview of surface registration techniques for medical imaging. *Medical Image Analysis*, 4(3): 201–217, 2000. → pages 13, 16, 18
- [5] S. Aytac, H. Yigit, T. Sancak, and H. Ozcan. Correlation between the diameter of the main renal artery and the presence of an accessory renal artery. *Journal of Ultrasound in Medicine*, 22(5):433–439, 2003. → pages 73
- [6] S. Badiei. Prostate segmentation in ultrasound images using image warping and ellipsoid fitting. Master’s thesis, The University of British Columbia, 2007. → pages 10, 32
- [7] S. Badiei, S. Salcudean, J. Varah, and W. Morris. Prostate segmentation in 2D ultrasound images using image warping and ellipse fitting. *Medical Image Computing and Computer-Assisted Intervention*, pages 17–24, 2006. → pages 10, 32
- [8] S. Badiei, S. Salcudean, J. Varah, N. Usmani, N. Chng, A. Kruk, I. Spadinger, and W. Morris. 3D prostate segmentation in TRUS images

using image warping and ellipsoid fitting. *Brachytherapy*, 6(2):95, 2007. → pages 10, 32

- [9] Y. Bar-Shalom and E. Tse. Tracking in a cluttered environment with probabilistic data association. *Automatica*, 11(5):451–460, 1975. → pages 10, 34
- [10] M. Baumann, P. Mozer, V. Daanen, and J. Troccaz. Towards 3D ultrasound image based soft tissue tracking: a transrectal ultrasound prostate image alignment system. In *Medical Image Computing and Computer-Assisted Intervention*, pages 26–33. Springer-Verlag, 2007. → pages 20, 58
- [11] M. Baumhauer, T. Simpfendorfer, B. Müller-Stich, D. Teber, C. Gutt, J. Rassweiler, H. Meinzer, and I. Wolf. Soft tissue navigation for laparoscopic partial nephrectomy. *International Journal of Computer Assisted Radiology and Surgery*, 3(3):307–314, 2008. → pages 22
- [12] A. Benincasa, L. Clements, S. Herrell, and R. Galloway. Feasibility study for image-guided kidney surgery: Assessment of required intraoperative surface for accurate physical to image space registrations. *Medical Physics*, 35:4251, 2008. → pages 22
- [13] P. Besl and N. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 239–256, 1992. ISSN 0162-8828. → pages 14, 18
- [14] H. Blom and Y. Bar-Shalom. The interacting multiple model algorithm for systems with Markovian switching coefficients. *IEEE Transactions on Automatic Control*, 33(8):780–783, 1988. → pages 10
- [15] A. Breda, S. Stepanian, J. Liao, J. Lam, G. Guazzoni, M. Stifelman, K. Perry, A. Celia, G. Breda, P. Fornara, et al. Positive margins in laparoscopic partial nephrectomy in 855 cases: a multi-institutional survey from the United States and Europe. *The Journal of Urology*, 178(1):47–50, 2007. → pages 1
- [16] Canadian Cancer Society’s Steering Committee on Cancer Statistics. *Canadian Cancer Statistics*, 2001. → pages 1
- [17] Canadian Cancer Society’s Steering Committee on Cancer Statistics. *Canadian Cancer Statistics*, 2011. → pages 1
- [18] J. Canny. A computational approach to edge detection. *Readings in Computer Vision: Issues, Problems, Principles, and Paradigms*, 184, 1987. → pages 13

- [19] V. Caselles, F. Catté, T. Coll, and F. Dibos. A geometric model for active contours in image processing. *Numerische Mathematik*, 66(1):1–31, 1993. → pages 11
- [20] V. Caselles, R. Kimmel, and G. Sapiro. Geodesic active contours. *International Journal of Computer Vision*, 22(1):61–79, 1997. ISSN 0920-5691. → pages 11, 37
- [21] V. Caselles, R. Kimmel, G. Sapiro, and C. Sbert. Minimal surfaces: A geometric three dimensional segmentation approach. *Numerische Mathematik*, 77(4):423–451, 1997. ISSN 0029-599X. → pages 37
- [22] D. Cremers, M. Rousson, and R. Deriche. A review of statistical approaches to level set segmentation: integrating color, texture, motion and shape. *International Journal of Computer Vision*, 72(2):195–215, 2007. → pages 11
- [23] W. Crum, T. Hartkens, and D. Hill. Non-rigid image registration: theory and practice. *British Journal of Radiology*, 77(Special Issue 2):S140, 2004. → pages 13
- [24] R. Czerwinski, D. Jones, and W. O’Brien Jr. Ultrasound speckle reduction by directional median filtering. In *Proceedings of the International Conference on Image Processing*, volume 1, pages 358–361. IEEE, 1995. → pages 21
- [25] R. Dalvia, R. Abugharbieha, M. Pickeringb, J. Scarvellic, and P. Smithd. Registration of 2D to 3D joint images using phase-based mutual information. In *Proceedings of SPIE*, volume 6512, pages 1–9. Citeseer, 2007. → pages 12
- [26] M. de Berg, O. Cheong, and M. van Kreveld. *Computational Geometry: Algorithms and Applications*. Springer, 3 edition, 2008. → pages 41
- [27] D. Eberly. Distance between point and triangle in 3D. Technical report, Geometric Tools, LLC, <http://www.geometrictools.com/>, 2008. → pages 41
- [28] E. Firle, S. Wesarg, G. Karangelis, and C. Dold. Validation of 3D ultrasound: CT registration of prostate images. In *Proceedings of SPIE*, volume 5032, page 354, 2003. → pages 20
- [29] J. Fitzpatrick, J. West, and C. Maurer Jr. Predicting error in rigid-body point-based registration. *IEEE Transactions on Medical Imaging*, 17(5): 694–702, 1998. → pages 26

- [30] O. Fluck, C. Vetter, W. Wein, A. Kamen, B. Preim, and R. Westermann. A survey of medical image registration on graphics hardware. *Computer Methods and Programs in Biomedicine*, pages 1–13, 2010. doi:10.1016/j.cmpb.2010.10.009. → pages 19
- [31] J. Friedman, J. Bentley, and R. Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Transactions on Mathematical Software (TOMS)*, 3(3):209–226, 1977. → pages 41
- [32] I. Hacihaliloglu. *Towards A Novel Minimally Invasive 3D Ultrasound Imaging Based Computer Assisted Orthopaedic Surgery System for Bone Fracture Reduction*. PhD thesis, The University of British Columbia, 2010. → pages 31
- [33] I. Hacihaliloglu, R. Abugharbieh, A. Hodgson, and R. Rohling. Bone surface localization in ultrasound using image phase-based features. *Ultrasound in Medicine & Biology*, 35(9):1475–1487, 2009. → pages 13, 31
- [34] N. Hansen and A. Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation*, 9(2):159–195, 2001. → pages 18
- [35] R. Haralick and L. Shapiro. Image segmentation techniques. *Computer Vision, Graphics, and Image Processing*, 29(1):100–132, 1985. → pages 7
- [36] D. Hill, P. Batchelor, M. Holden, and D. Hawkes. Medical image registration. *Physics in Medicine and Biology*, 46:R1–R45, 2001. → pages 13, 14, 15, 16, 17, 18
- [37] B. Hutton and M. Braun. Software for image registration: algorithms, accuracy, efficacy. In *Seminars in Nuclear Medicine*, volume 33, pages 180–192. Elsevier, 2003. → pages 13
- [38] G. Janetschek. Laparoscopic partial nephrectomy: how far have we gone? *Current Opinion in Urology*, 17(5):316–321, 2007. ISSN 0963-0643. → pages 1, 5, 7
- [39] R. Kingma, R. Rohling, and C. Ngan. Registration of CT to 3D ultrasound using near-field fiducial localization: A feasibility study. *Computer Aided Surgery*, 16(2):54–70, 2011. → pages 26
- [40] T. Kirubarajan, Y. Bar-Shalom, W. Blair, and G. Watson. IMMPDAF for radar management and tracking benchmark with ECM. *IEEE Transactions on Aerospace and Electronic Systems*, 34(4):1115–1134, 1998. → pages 10

- [41] P. Kovesi. Image features from phase congruency. *Videre: Journal of Computer Vision Research*, 1(3):1–26, 1999. → pages 12, 31
- [42] P. Kovesi. Phase congruency detects corners and edges. In *The Australian Pattern Recognition Society Conference*, pages 309–318. Citeseer, 2003. → pages 12, 31
- [43] B. Lane and A. Novick. Nephron-sparing surgery. *British Journal of Urology International*, 99(5b):1245–1250, 2007. → pages 6
- [44] A. Leroy, P. Mozer, Y. Payan, and J. Troccaz. Rigid registration of freehand 3D ultrasound and CT-scan kidney images. *Medical Image Computing and Computer-Assisted Intervention*, pages 837–844, 2004. → pages 21
- [45] A. Leroy, P. Mozer, Y. Payan, and J. Troccaz. Intensity-based registration of freehand 3D ultrasound and CT-scan images of the kidney. *International Journal of Computer Assisted Radiology and Surgery*, 2:31–41, 2007. ISSN 1861-6410. URL <http://dx.doi.org/10.1007/s11548-007-0077-5>. 10.1007/s11548-007-0077-5. → pages 21
- [46] C. Li, C. Xu, C. Gui, and M. Fox. Level set evolution without re-initialization: a new variational formulation. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 430–436. IEEE. → pages 66
- [47] C. Li, C. Xu, C. Gui, and M. Fox. Distance regularized level set evolution and its application to image segmentation. *IEEE Transactions on Image Processing*, 19(12):3243–3254, 2010. → pages 11
- [48] Q. Li and J. Griffiths. Least squares ellipsoid specific fitting. In *Proceedings of the Geometric Modeling and Processing*, pages 335–340. IEEE, 2004. → pages 29, 58
- [49] S. Lobregt and M. Viergever. A discrete dynamic contour model. *IEEE Transactions on Medical Imaging*, 14(1):12–24, 1995. → pages 12
- [50] D. Lowe. Object recognition from local scale-invariant features. In *IEEE International Conference on Computer Vision*, volume 2, pages 1150–1157, 1999. → pages 14
- [51] D. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. → pages 14, 68

- [52] J. Maintz and M. Viergever. A survey of medical image registration. *Medical image analysis*, 2(1):1–36, 1998. → pages 13, 15, 16, 17
- [53] R. Malladi, J. Sethian, and B. Vemuri. Shape modeling with front propagation: A level set approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(2):158–175, 1995. → pages 11
- [54] D. Marr and E. Hildreth. Theory of edge detection. *Proceedings of the Royal Society of London, Series B: Biological Sciences*, 207(1167):187, 1980. → pages 13
- [55] H. Matthies and G. Strang. The solution of nonlinear finite element equations. *International Journal for Numerical Methods in Engineering*, 14(11):1613–1626, 1979. → pages 40
- [56] M. Mellor and M. Brady. Phase mutual information as a similarity measure for registration. *Medical Image Analysis*, 9(4):330–343, 2005. → pages 12
- [57] M. Moghari and P. Abolmaesumi. Point-based rigid-body registration using an unscented Kalman filter. *IEEE Transactions on Medical Imaging*, 26(12):1708–1728, 2007. → pages 19, 40, 68
- [58] M. Morrone and D. Burr. Feature detection in human vision: A phase-dependent energy model. *Proceedings of the Royal Society of London, Series B: Biological Sciences*, 235(1280):221–245, 1988. → pages 12
- [59] P. Mozer, A. Leroy, Y. Payan, J. Troccaz, E. Chartier-Kastler, and F. Richard. Computer-assisted access to the kidney. *The International Journal of Medical Robotics and Computer Assisted Surgery*, 1(4):58–66, 2005. → pages 21
- [60] M. Mulet-Parada and J. Noble. 2D+T acoustic boundary detection in echocardiography. *Medical Image Analysis*, 4(1):21–30, 2000. → pages 13
- [61] A. Myronenko and X. Song. Point set registration: Coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12):2262–2275, 2010. → pages 19, 42
- [62] J. Noble. Ultrasound image segmentation and tissue characterization. *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine*, 224(2):307–316, 2010. → pages 7, 8
- [63] J. Noble and D. Boukerroui. Ultrasound image segmentation: A survey. *IEEE Transactions on Medical Imaging*, 25(8):987–1010, 2006. → pages 7, 9, 10

- [64] S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed: algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, 1988. → pages 11
- [65] S. Pahernik, S. Ziegler, F. Roos, S. Melchior, and J. Thüroff. Small renal tumors: correlation of clinical and pathological features with tumor size. *The Journal of Urology*, 178(2):414–417, 2007. → pages 5
- [66] N. Paragios, M. Rousson, and V. Ramesh. Knowledge-based registration & segmentation of the left ventricle: A level set approach. In *IEEE Workshop on Applications of Computer Vision*, pages 37–42. IEEE, 2002. → pages 11
- [67] G. Penney, J. Blackall, M. Hamady, T. Sabharwal, A. Adam, and D. Hawkes. Registration of freehand 3D ultrasound and magnetic resonance liver images. *Medical Image Analysis*, 8(1):81–91, 2004. → pages 20, 21
- [68] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. Numerical recipes in C. 1992. → pages 21
- [69] A. Rasouljan. Group-wise CT to ultrasound registration of lumbar spine. Master’s thesis, Der Technischen Universität München, 2009. → pages 40
- [70] J. Rassweiler, C. Abbou, G. Janetschek, and K. Jeschke. Laparoscopic partial nephrectomy: The european experience. *Urologic Clinics of North America*, 27(4):721–736, 2000. → pages 5, 6, 7
- [71] I. Rechenberg. *Evolutionsstrategie: Optimierung technischer Systeme nach Prinzipien der biologischen Evolution*, volume 15. Frommann-Holzboog, 1973. → pages 18
- [72] A. Roche, X. Pennec, G. Malandain, and N. Ayache. Rigid registration of 3-D ultrasound with MR images: a new approach combining intensity and gradient information. *IEEE Transactions on Medical Imaging*, 20(10):1038–1049, 2001. → pages 20
- [73] C. Rogers, A. Singh, A. Blatt, W. Linehan, and P. Pinto. Robotic partial nephrectomy for complex renal tumors: surgical technique. *European Urology*, 53(3):514–523, 2008. → pages 7, 8
- [74] K. Saini, M. Dewal, and M. Rohit. Ultrasound imaging and image segmentation in the area of ultrasound: A review. *International Journal of Advanced Science and Technology*, 24:41–59, 2010. → pages 9, 10

- [75] G. Sanchez-Ortiz, J. Declerck, M. Mulet-Parada, and J. Noble. Automating 3D echocardiographic image analysis. In *Medical Image Computing and Computer-Assisted Intervention*, pages 687–696. Springer, 2000. → pages 13
- [76] L. Schwartz, J. Richaud, L. Buffat, E. Touboul, and M. Schlienger. Kidney mobility during respiration. *Radiotherapy and Oncology*, 32(1):84–86, 1994. → pages 16
- [77] C. Shannon. Communication in the presence of noise. *Proceedings of the Institute of Radio Engineers*, 37(1):10–21, 1949. → pages 17
- [78] P. Slomka and R. Baum. Multimodality image registration with software: state-of-the-art. *European Journal of Nuclear Medicine and Molecular Imaging*, 36:44–55, 2009. → pages 14, 16
- [79] M. Spaliviero and I. Gill. Laparoscopic partial nephrectomy. *British Journal of Urology International*, 99(5b):1313–1328, 2007. ISSN 1464-410X. → pages 1, 5, 6, 7
- [80] G. Treece, R. Prager, and A. Gee. Regularised marching tetrahedra: improved iso-surface extraction. *Computers and Graphics*, 23(4):583–598, 1999. → pages 25, 83
- [81] G. Treece, R. Prager, A. Gee, and L. Berman. Surface interpolation from sparse cross-sections using region correspondence. *IEEE Transactions on Medical Imaging*, 19(11):1106–1114, 2000. → pages 25, 83
- [82] A. Tsai, A. Yezzi Jr, W. Wells, C. Tempany, D. Tucker, A. Fan, W. Grimson, and A. Willsky. A shape-based approach to the segmentation of medical imagery using level sets. *IEEE Transactions on Medical Imaging*, 22(2):137–154, 2003. → pages 11
- [83] B. Vemuri and Y. Chen. Joint image registration and segmentation. *Geometric Level Set Methods in Imaging, Vision, and Graphics*, pages 251–269, 2003. → pages
- [84] B. Vemuri, J. Ye, Y. Chen, and C. Leonard. Image registration via level-set motion: Applications to atlas-based segmentation. *Medical Image Analysis*, 7(1):1–20, 2003. → pages 11
- [85] P. Viola and W. Wells III. Alignment by maximization of mutual information. *International Journal of Computer Vision*, 24(2):137–154, 1997. → pages 17

- [86] H. Winfield, J. Donovan, A. Godet, and R. Clayman. Laparoscopic partial nephrectomy: initial case report for benign disease. *Journal of Endourology*, 7(6):521–526, 1993. → pages 6
- [87] C. Wu and Y. Sun. Segmentation of kidney from ultrasound B-mode images with texture-based classification. *Computer Methods and Programs in Biomedicine*, 84(2-3):114–123, 2006. → pages 9
- [88] J. Xiang. Registration of 3D ultrasound to computed tomography images of the kidney. Master’s thesis, The University of British Columbia, 2010. → pages 85
- [89] J. Xie, Y. Jiang, and H. Tsui. Segmentation of kidney from ultrasound images based on texture and shape priors. *IEEE Transactions on Medical Imaging*, 24(1):45–57, 2005. → pages 9
- [90] W. Zhang, R. Rohling, and D. Pai. Surface extraction with a three-dimensional freehand ultrasound system. *Ultrasound in Medicine & Biology*, 30(11):1461–1473, 2004. ISSN 0301-5629. → pages 2
- [91] W. Zhang, J. Noble, and J. Brady. Real time 3-D ultrasound to MR cardiovascular image registration using a phase-based approach. In *IEEE International Symposium on Biomedical Imaging*, pages 666–669. IEEE, 2006. → pages 12
- [92] Y. Zhang, B. Matuszewski, L. Shark, and C. Moore. Medical image segmentation using new hybrid level-set method. In *BioMedical Visualization*, pages 71–76. IEEE, 2008. → pages 11
- [93] Z. Zhang. *Iterative point matching for registration of free-form curves*. Citeseer, 1992. → pages 18
- [94] B. Zitova and J. Flusser. Image registration methods: a survey. *Image and Vision Computing*, 21(11):977–1000, 2003. → pages 13, 14, 15, 16, 17

Appendix A

Vol2Strad: a Volume Conversion Tool

A specialized volume conversion tool, `Vol2Strad`, was developed to convert raw US data into correctly spaced slices. The program takes an Ultrasonix `.vol` file as an argument, and prompts the user for the desired sweep. It then reads the data file's header info and calculates the relative position of all slices using the known probe parameters.

The scanlines for each slice were interpolated in 2-D using the Pando SDK and probe parameters from Ultrasonix to produce an image in Cartesian coordinates. Slices were then placed at the correct position in 3-D space and written to files that can be read by Stradwin for viewing and segmentation [80, 81].

Through this work, it was discovered that the imaging depth indicated by Ultrasonix's `Propello` interface at the time of acquisition is only a rough estimate. For proper scan conversion, the imaging depth is calculated as:

$$depth = (samples/line) \left(\frac{c}{2f_{samp}} \right) \quad (A.1)$$

where c is the speed of sound in the tissue, approximated at 1540 m/s, and f_{samp} is the sampling frequency of the transducer.

The converter was validated using a Model 049 QA phantom made by Computerized Imaging Reference Systems. A known 20-mm sphere was segmented

from the phantom images with negligible error.

Appendix B

Setting IMMPDAF Parameters

B.1 Acquisition of Phantom Data

The porcine kidney phantom images used for filter calibration were acquired from Xiang [88]. Kidneys were excised from live pigs, flushed with saline, and injected with a gelatin contrast agent. They were then placed in a container with an agar solution for US and CT imaging. Images from kidney 8 were used for the filter experiments.

B.2 Parameter Manipulation Experiments

The standard IMMPDAF function for edge detection in US was parameterized to allow fine control over edge detection. Parameters include the radial search range (SR), the number of candidate edge points per radius (N_r), the number of iterations (N_{iter}), the median filter length (l_{med}), the gradient filter length (l_f) and the covariances of the process noise (Q) and measurement noise (R).

An experiment was designed to determine the optimal settings for edge detection in kidney images. The IMMPDAF algorithm was tried on pig phantom data with a variety of parameter combinations, five of which are presented in Table B.1. The first run serves as a control, the second uses two equivalent models that favour edge-following, the third uses two equivalent models that favour smoothness (model-following), the fourth features a longer gradient filter for edge

Trial	Q	R	l_f
control	[10 10 ⁵]	[20 20]	6 pixels
edge	[10 ⁵ 10 ⁵]	[1 1]	6 pixels
smooth	[5000 5000]	[20 20]	6 pixels
long	[10 10 ⁵]	[20 20]	10 pixels
best	[10 ⁴ 10 ⁶]	[10 1]	8 pixels

Table B.1: Selected parameter configurations for IMMPDAF parameter refinement trials

detection and the fifth is the combination that was eventually selected.

As expected, the edge-following configuration produces a noisy boundary because the filter is unable to tell the edges from the background (Figure B.1(b)). The smooth configuration, shown in Figure B.1(c), illustrates the drawbacks of disregarding too many edges, as well as the limitations of the constant velocity radial distance model. It can be seen that the filter “overshoots” the contour as it passes the poles in its clockwise sweep. Then, as it corrects itself, it develops an inward velocity that takes it inside the true contour until it is corrected by edge measurements. These difficulties are not present in the prior-guided IMMPDAF, as the sharp corners at the poles are predicted by the model, and the damping coefficient reduces overshooting.

Using a longer radial gradient filter, as in Figure B.1(d), can help to reduce the number of less significant intensity changes that attract the filter. In the end, the best results were obtained using an intermediate filter length (Figure B.1(e)). The process noise was increased relative to the measurement noise in order to allow more deviation from the model around high-curvature areas. These parameters were mostly kept constant with the patient data, although there were some algorithmic changes.




Figure B.1: IMMPDAF output for the filter parameter combinations described in Table B.1. The yellow dotted line shows the contours from the prior segmentation, the red shows the IMMPDAF output, and the green shows the gold standard segmentation.

Appendix C

Patient Results for Prior-guided IMMPDAF Methods

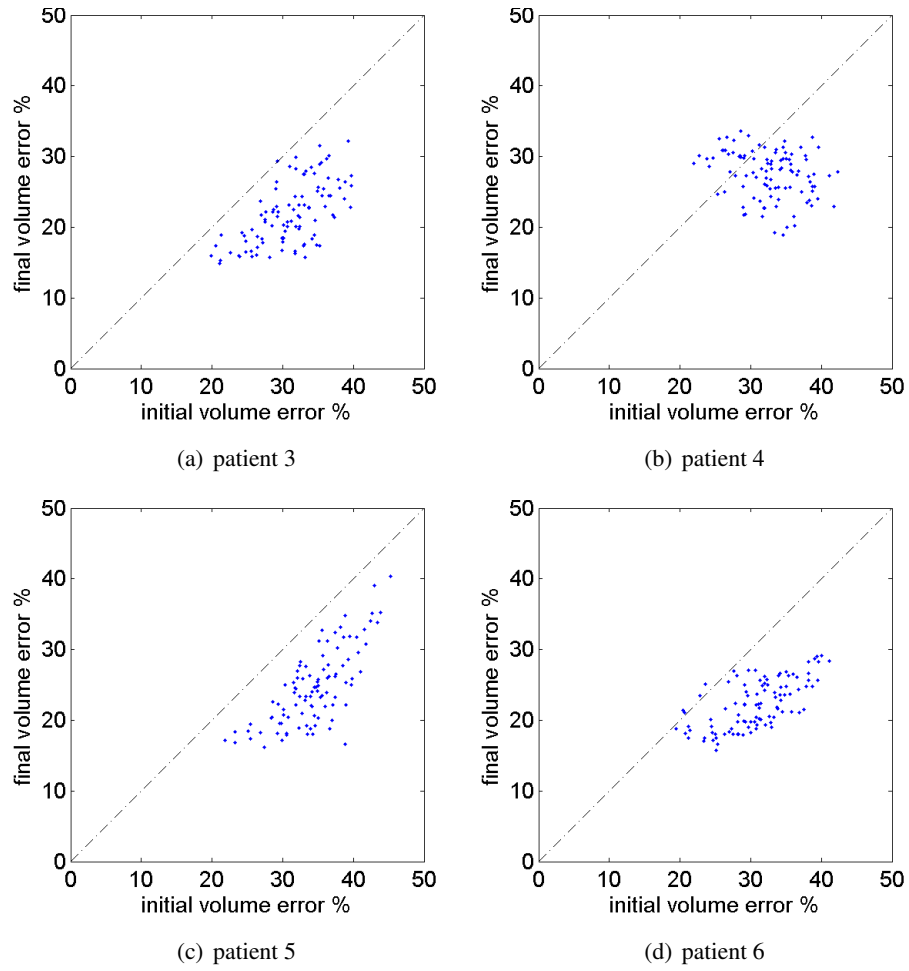


Figure C.1: Volume error after registration vs. volume error after initial alignment for prior-guided IMMPDAF

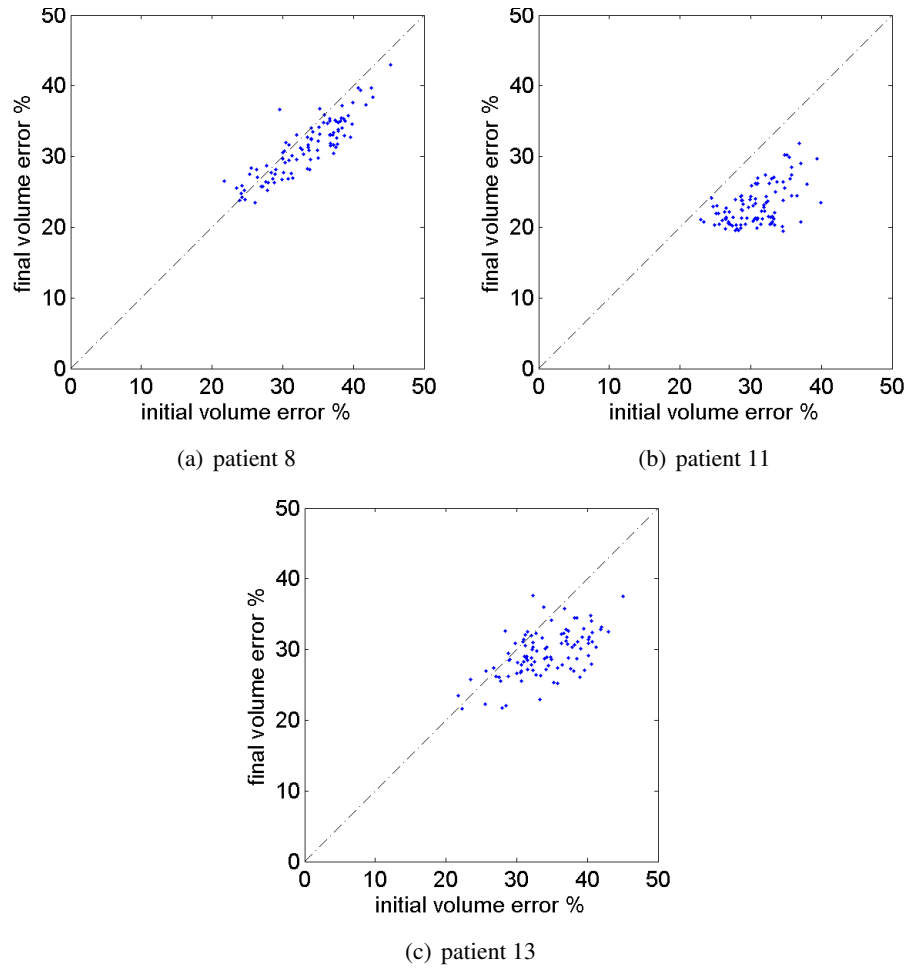


Figure C.2: Volume error after registration vs. volume error after initial alignment for prior-guided IMMPDAF




Figure C.3: Volume error after registration vs. volume error after initial alignment for prior-guided IMMPDAF on Canny edges

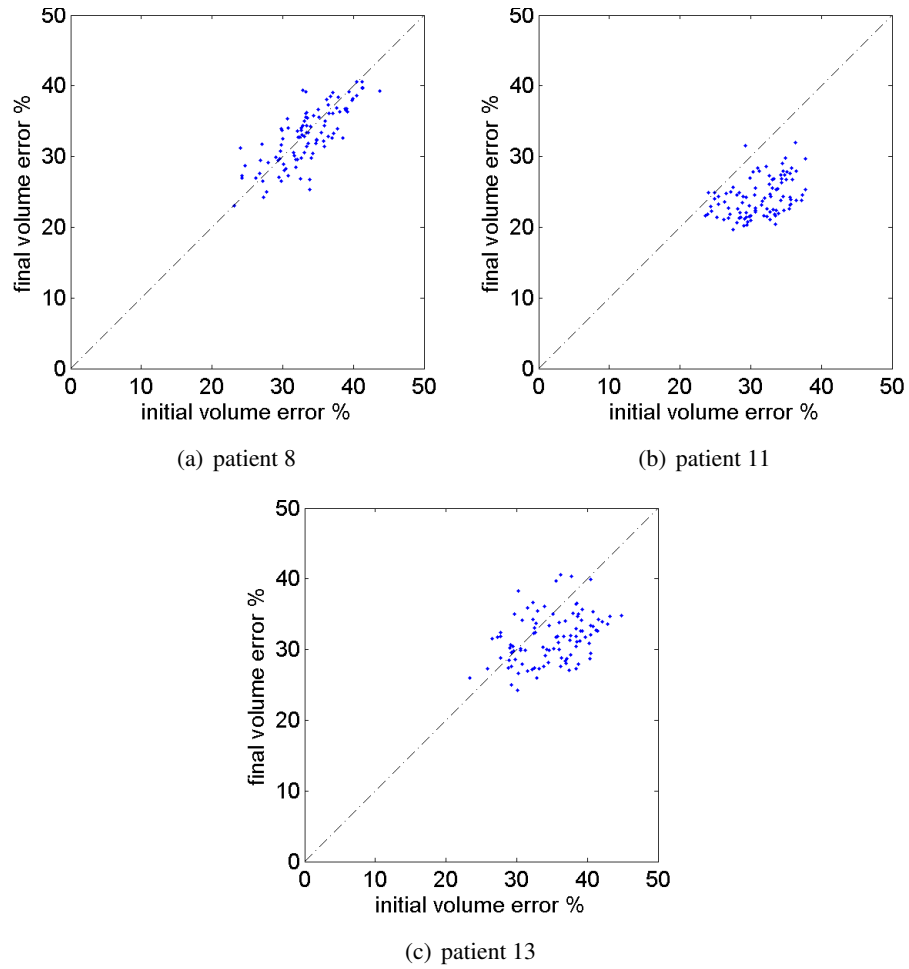


Figure C.4: Volume error after registration vs. volume error after initial alignment for prior-guided IMMPDAF on Canny edges

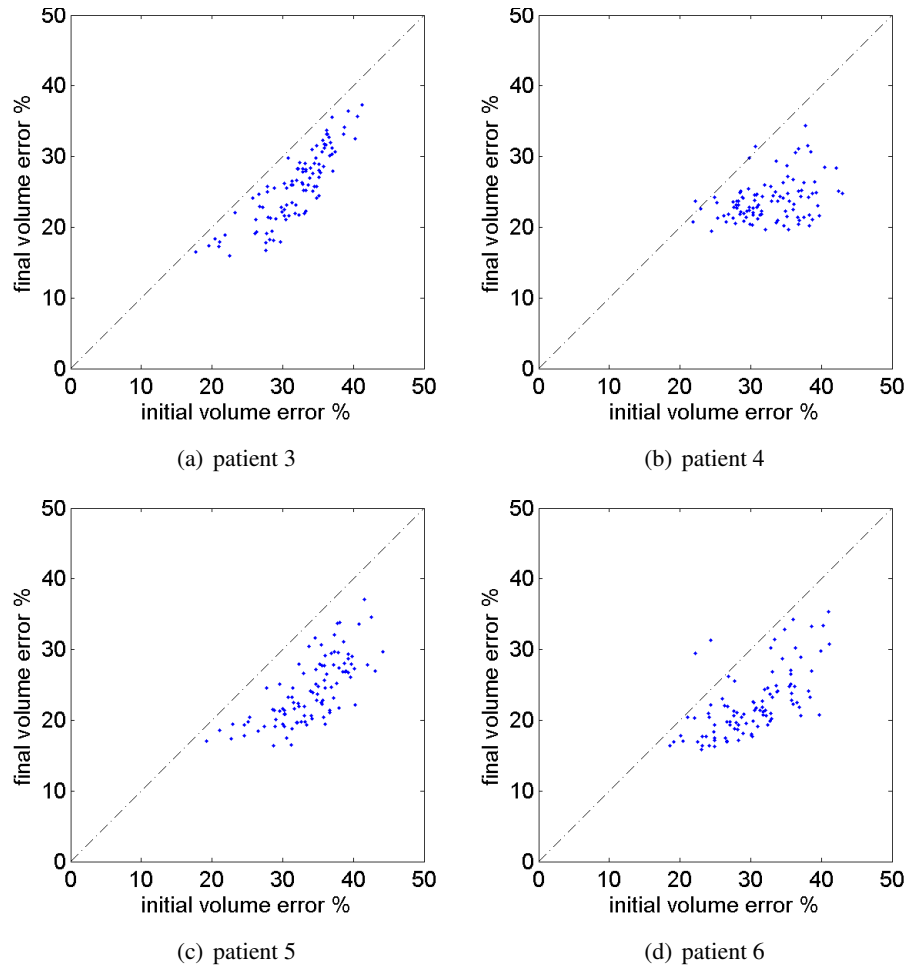


Figure C.5: Volume error after registration vs. volume error after initial alignment for prior-guided IMMPDAF on phase edges

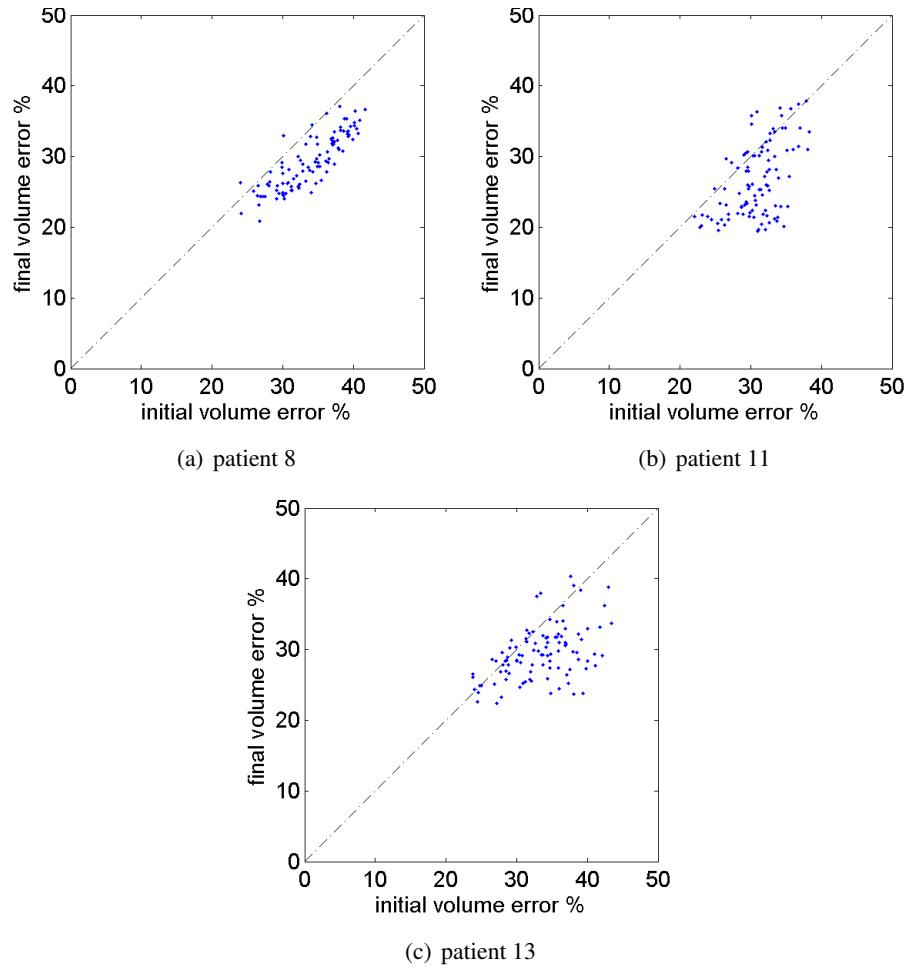


Figure C.6: Volume error after registration vs. volume error after initial alignment for prior-guided IMMPDAF on phase edges

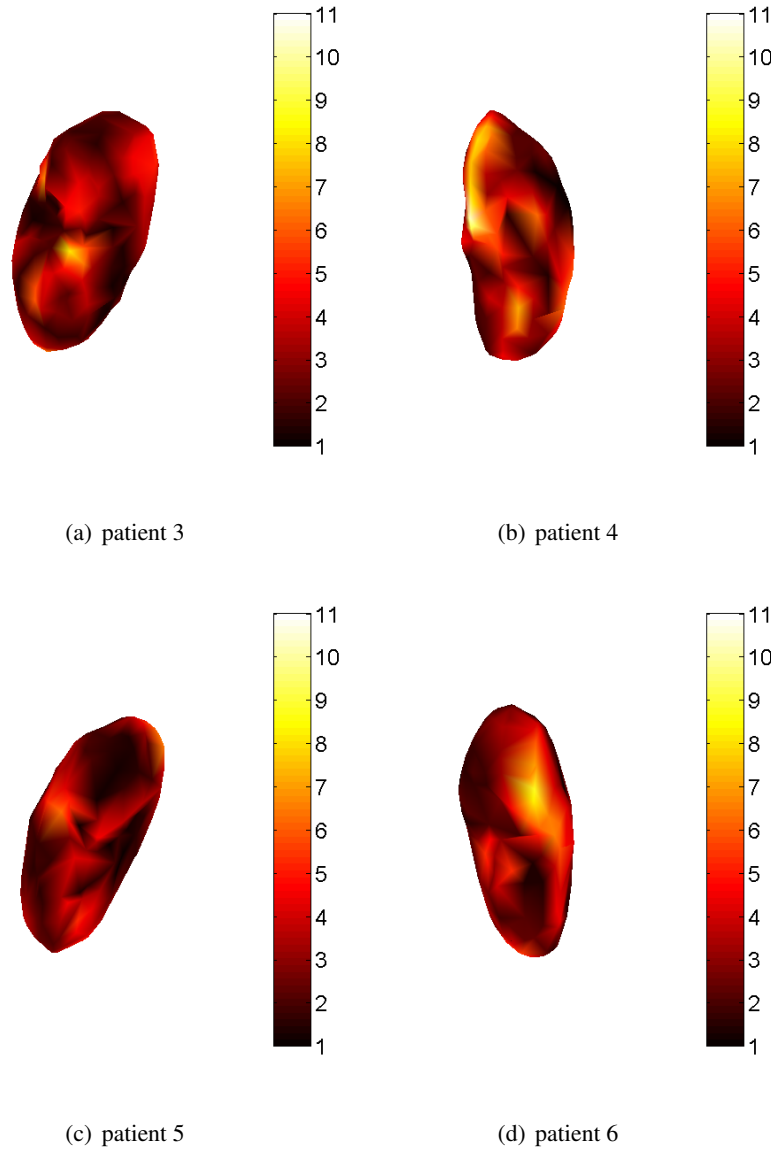


Figure C.7: Colour map showing final alignment error after prior-guided IMMPDAF on phase images for patients 3 to 6. Error is shown in mm for all locations on surface of gold standard.

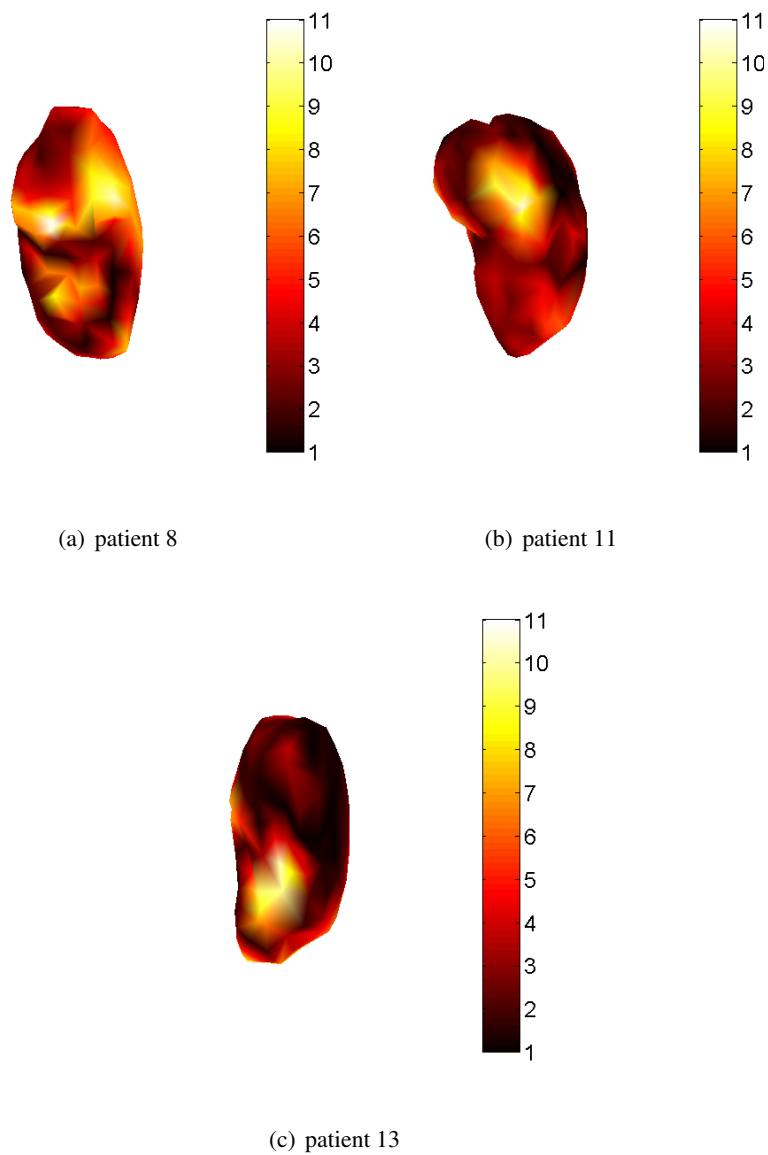


Figure C.8: Colour map showing final alignment error after prior-guided IMMPDAF on phase images for patients 8 to 13. Error is shown in mm for all locations on surface of gold standard.