4D NON-RIGID REGISTRATION OF RENAL DYNAMIC CONTRAST ENHANCED MRI DATA

FOO JIT SOON

NATIONAL UNIVERSITY OF SINGAPORE

2011

4D NON-RIGID REGISTRATION OF RENAL DYNAMIC CONTRAST ENHANCED MRI DATA

FOO JIT SOON
B.Eng. (Hons.), NUS

A THESIS SUBMITTED

FOR THE DEGREE OF MASTER OF ENGINEERING

DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING

NATIONAL UNIVERSITY OF SINGAPORE

2011

ACKNOWLEDGEMENTS

I am heartily thankful to my supervisor, Sun Ying, whose encouragement, guidance and support from the initial to the final stages of my project enabled me to develop a great understanding and interest of the subject.

I would also like to show my gratitude to Mahapatra Dwarikanath and Li Chao from the National University of Singapore for their valuable assistance and their insightful comments.

Lastly, I offer my regards to all of those who supported me in any respect during the completion of the project.

Foo Jit Soon

TABLE OF CONTENTS

SUMMARY

This thesis presents a near-automatic non-rigid registration algorithm requiring minimal user interaction for renal dynamic contrast enhanced (DCE) MR images. The 12 patients' dataset (24 kidney volumes) to be registered were acquired on a 1.5T scanner of size 256 x 256 x 40 (voxel resolution of 1.66mm x 1.66mm x 2.5mm) with the number of static volumes in each dataset varying from 31 to 41. A multi-level registration algorithm is proposed to first account for initial large translational errors, followed by compensating for local deformations of the kidney. A graph-cut optimization technique integrating local gradient information into an energy function solves the initial problem of 3D translational registration. A motion/noise free pseudo ground-truth dataset is then estimated from the whole time sequence of each kidney dataset obtained after translational registration. Finally, the demons algorithm is used to register each 3-D volume (as floating image) to its corresponding estimated volume (as reference image) at each time frame. Experimental results on patient data demonstrate that the proposed algorithm is able to: (1) perform initial translational registration accurately with an error of up to 5 voxels; (2) correctly estimate the pseudo ground-truth dataset, and (3) achieve non-rigid registration of 4D time-series of renal DCE MRI data.

LIST OF FIGURES

LIST OF TABLES

LIST OF SYMBOLS

X, Y, Z	Principal axes of the image volume
$E(L)$	Energy function with label L for graph-cuts
W_p	Weight term representing importance of node p
D_p	Data penalty term for node p
L_p	Label of node p
λ	Regularization constant
$V_{p,q}$	Smoothness constraint term between nodes p and q
$\mathbb{P}$	Set of all nodes
N	Set of all neighboring nodes
$\Delta x, \Delta y, \Delta z$	Translational parameters for a label
$\mathbb{Z}$	Set of all integers
F_x, F_y	Edge magnitude of the reference image in the x and y directions, respectively
G_x, G_y	Edge magnitude of the floating image in the x and y directions, respectively
$\boldsymbol{u}, \boldsymbol{u}'$	Voxel position within an image
B_p	Set of all voxels within node p
S_p	Strength of node p
S_{min}, S_{max}	Minimum and maximum strength of nodes
M_f	Magnitude of gradient in the reference image f
$d(L_p, L_q)$	Euclidean distance between displacements denoted by labels L_p and L_q
$\delta x, \delta y, \delta z$	Global translational parameters of an image volume
f, g	Reference image and floating image, respectively

$E(H(g),f')$	Energy function for pseudo ground-truth estimation of the transformed floating image $H(g)$ and the pseudo ground truth f'
H	Non-rigid transformation function
f'	Pseudo ground-truth image (reference image for non-rigid registration)
E_d	Data fidelity term
E_s	Spatial smoothness term
E_t	Temporal smoothness term
W_k	Weight matrix for the k[th] neighbor
Ds_k	First order spatial derivative operator along direction between each pixel and the k[th] neighbor
W_t	Temporal weight matrix
D_{1t}, D_{2t}	First and second order temporal derivative operators
α, β_1, β_2	Weights of different terms
ρ_{ijk}	Correlation coefficient between the intensity-time curves of the pixel at position (i,j) and its k[th] neighbor at position (i_k, j_k)
t_s	Frame number
ρ_{th}, σ	Minimal signal similarity and intensity variance within the same tissue type
L_{ij}'	Segment label at position (i,j) in the image
$E(V), E(v)$	Demons energy for the whole image and for a voxel, respectively
S	Transformation field describing the translation in the 3 principle axes of every voxel from its original position
V, v	Update of the transformation field in each iteration for the whole image and for a voxel, respectively

$\circ$	Image transformation function
σ_i	Constant for image noise
σ_x	Constant for transformation uncertainty
I	Intensity of a voxel
∇	Gradient operator
e_x, e_y, e_z	Absolute error between the translational parameters obtained from registration and the ground truth parameters in each of the 3 principle axes X, Y and Z
e	Euclidean distance of the registration error
η	Threshold error in cumulative distribution function
T	Tesla
TR	Repetition time
TE	Echo time
FOV	Field of View

CHAPTER 1

INTRODUCTION

Magnetic resonance imaging (MRI) is becoming increasingly popular due to its non-invasive nature and good reliability. However, as the data acquisition procedure is rather lengthy, the misalignment of 3-D image volumes in a given time-series occurs, induced mainly by patient movement and breathing. Image registration, the process of comparing and integrating data obtained from different or same types of measurements, is thus employed to correct these movements, with details being described by Maintz and Viergever in [1]. Of great clinical interest are automated and semi-automated image registration techniques capable of correcting patient movement and respiratory motion, as the general registration procedure is rather time consuming and labor intensive. By aligning the images across different time frames, different modalities and different scenes, post-processing algorithms such as segmentation, recognition and diagnostics can be performed with higher accuracy because minimal misalignment errors are carried forward. Outside of the medical field, image registration is also useful in other important applications such as video enhancement [2], scene representation [3] and automatic target recognition (ATR) systems [4].

1.1. Image Registration

Much research has been done in the area of medical image processing, and there exist many different algorithms for performing image registration. Registration algorithms can involve: (1) different organs/features of interest; (2) different similarity

measures; (3) different modalities; (4) different degrees of freedom. Image registration algorithms act on different objects/features of interest (organs in a body, buildings or written characters, etc.); thus the methods applied in each registration algorithm are different based on the type of image being processed, how the important details change between different images (with time) and the user's requirements. Similarity measures depict the similarity between the floating image and the reference image, with some examples being given in [5]. By using these similarity metrics to determine how similar the reference and floating images are, the registration algorithm can align images based on a pre-defined search space. Different modalities are available for imaging, and these modalities are not restricted entirely to images created by photography. In medical image processing, popular imaging modalities include MRI and computed tomography (CT). Image registration algorithms can also be classified into rigid registration, affine registration or non-rigid registration. Each of these algorithms has different degrees of freedom to define how much an image can transform. In rigid registration, translation, rotation and scaling errors within an image are recovered. Rigid registration works best for objects in which the object features do not deform over time using the same viewing angle or the same principle axes, such as buildings and vehicles. Affine registration, like rigid registration, recovers a linear transformation between images, and has a higher degree of freedom than rigid registration to account for shearing. Lastly, non-rigid registration allows the most degrees of freedom in the image transformation; the object in the floating image is transformed elastically thus ensuring a better fit to the reference image. Non-rigid registration methods are more common, due to the nature of many objects/features deforming non-rigidly across different images, such as human organs and

written alphabets/numbers. Many types of transformations exist for non-rigid registration, for example, free-form deformation (FFD) or B-splines, etc.
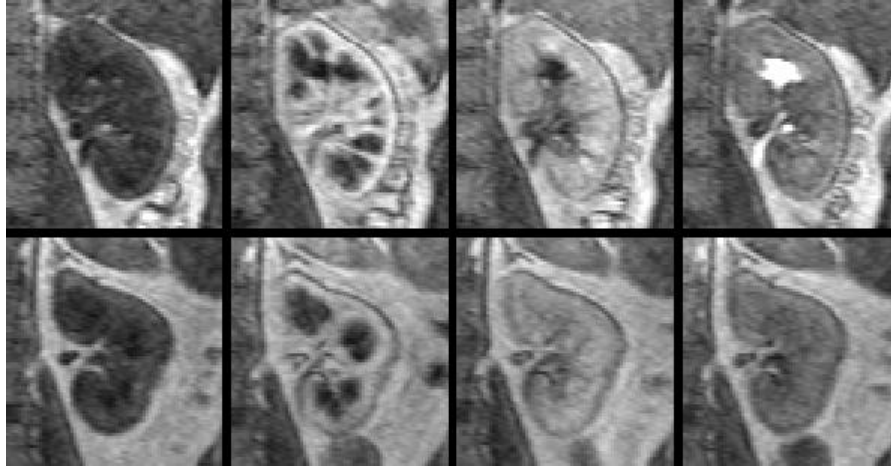
Fig.1.1. Sample kidney images from different patients (top and bottom rows) obtained at different times (different columns) with varying intensities at the cortex and medulla. The medulla (dark patches) is clearly visible in the kidney images in the second column due to the effect of contrast agent.

Image registration is complex in nature, and the difficulty of image registration is increased further when dynamic contrast enhanced-magnetic resonance imaging (DCE-MRI) is involved. Registration of DCE-MRI images is a major challenge, as the acquired images by DCE-MRI exhibit rapid intensity changes differently in different parts of the organ following the injection of a contrast agent as shown in Figure 1.1. In [6], a discussion on processing methods on DCE-MRI renal images is given.

1.2. Perfusion Magnetic Resonance Imaging

In this thesis, the patient datasets are obtained by perfusion magnetic resonance imaging (pMRI), a technique that measures the rate at which blood is delivered to the tissue. Being a subset of DCE-MRI, pMRI is a special technique for evaluating

microscopic blood flow in capillaries. With the advantages of providing images of high level of detail and ensuring that the patient is not exposed to radiation, pMRI has shown great promise for the non-invasive diagnosis of cardiovascular and renovascular diseases. As a diagnostic tool with great potential, pMRI is being used more and more extensively in both medical research and clinical practice. Some example images for pMRI are shown in Figure 1.2 and Figure 1.3. It can be observed in the MRI perfusion map of Figure 3, that "A" marks the position of a high grade brain tumor where the region demonstrates an increased capillary blood volume due to a tumor.
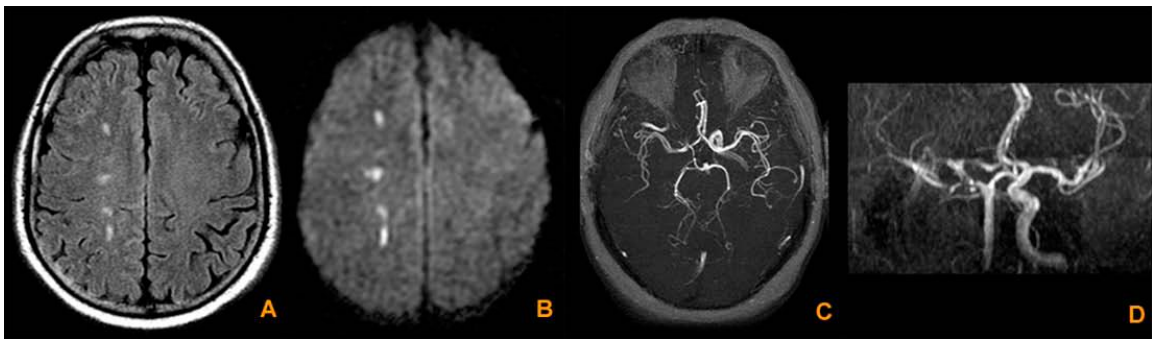
Fig.1.2. Cerebral perfusion MRI images taken from http://emedicine.medscape.com/
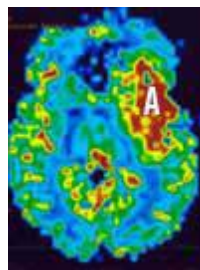
Fig.1.3. MRI perfusion map taken from http://www.rcnd.com/PerfusionMRI.html. The location marked with an 'A' indicates the position of a tumor.

The main applications of perfusion MRI include the detection of the following diseases, as listed in [7]:

- Vascular diseases: ischemic stroke (search for penumbra / mismatch corresponding to viable brain tissue), study of vasospasm in subarachnoid hemorrhage.

- Infectious or inflammatory diseases, one of whose physiopathological characteristics is hypervascularisation, can be explored in perfusion MRI.

- Tumoral diseases: perfusion MRI is used to evaluate neoangiogenesis and tumoral vascularisation, with an impact on the diagnosis or aftercare treatment of certain nervous tumors (high grade glioma, lymphoma meningioma, pilocytic astrocytoma, metastasis…).

1.3. Proposed Registration Algorithm

Due to breathing motion, non-rigid deformations exist in the kidney images over time. Such non-rigid deformations become more complex when the patient datasets contain cysts, tumors and other anomalies in the kidneys as shown in Figure 1.4; these anomalies cause the kidney to deform non-rigidly over time. Non-rigid deformations cause ambiguities over the optimal transformation parameters in a rigid registration algorithm; rigid registration algorithms are unable to account entirely for non-rigid deformations. Therefore, it is only logical to account for such non-rigid deformations within the kidney by performing a non-rigid registration which is discussed in this thesis.
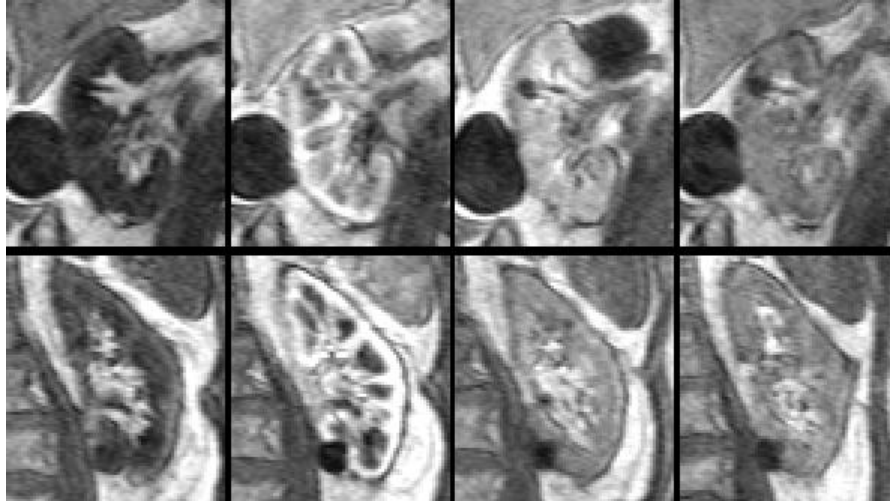
Fig.1.4. Sample kidney images from different patients (different rows) obtained at different times (different columns) with observable kidney deformations.

In order to ensure the efficiency of the registration process, a multi-level registration algorithm is proposed. Pre-processing is first performed to obtain the initial region of interest (ROI) and a 2-D kidney mask. Following the pre-processing step are two major steps in the registration framework: (1) Rigid registration to account for initial large translational errors, and (2) Non-rigid registration to determine the non-rigid local deformations of the kidney. For the initial problem of 3-D translational registration, a graph-cut optimization technique is introduced as it is efficient and robust. Robustness of the registration algorithm is ensured by combining local gradient information of the kidney into the energy function to obtain a global solution. With rigid registration performed, a smaller ROI is obtained for the non-rigid registration process increasing the overall efficiency. For the non-rigid registration of the kidney, a motion/noise free pseudo ground-truth dataset is first estimated by imposing spatiotemporal smoothness constraints on the entire time sequence of each kidney dataset obtained after translational registration. The estimation is done with reference to the method described in [8] for the registration of myocardial perfusion MRI, with certain adaptations made to fit the framework of

kidney registration. Then, the images within the pseudo-ground truth dataset will be used as reference images to register the original kidney images non-rigidly. With knowledge that the non-rigid deformations are small, an efficient demons algorithm as described in [9] is used to register each 3-D volume (as floating image) with its corresponding 3-D estimated volume (as reference image) at a particular time. The flowchart of the proposed algorithm is given in Figure 1.5.

The main contributions of the work presented in this thesis are threefold: 1) applying graph-cut to image registration using gradient information and implementing the graph by using super-nodes to link local edge information in different nodes together; 2) estimating pseudo ground-truth for the kidney images in 3D by adapting the estimation framework previously applied to 2D cardiac images; and 3) using the demons algorithm to account for small non-rigid motion in kidneys which is efficient and robust.

The rest of the thesis is organized as follows. Chapter 2 gives a literature review on the various methods proposed in rich literature to perform biomedical image registration, with the focus on registration of renal DCE-MRI image series. Chapter 3 describes the initial rigid registration algorithm to compensate for large translational motion in the kidney. Chapter 4 describes the non-rigid registration algorithm to compensate for elastic deformations of the kidney after the initial translational registration, including the estimation of the pseudo ground-truth and the demons registration algorithm. Chapter 5 presents the results qualitatively and quantitatively obtained from the multi-level registration algorithm as described in both Chapters 3 and 4. Finally, we conclude the thesis in Chapter 6.
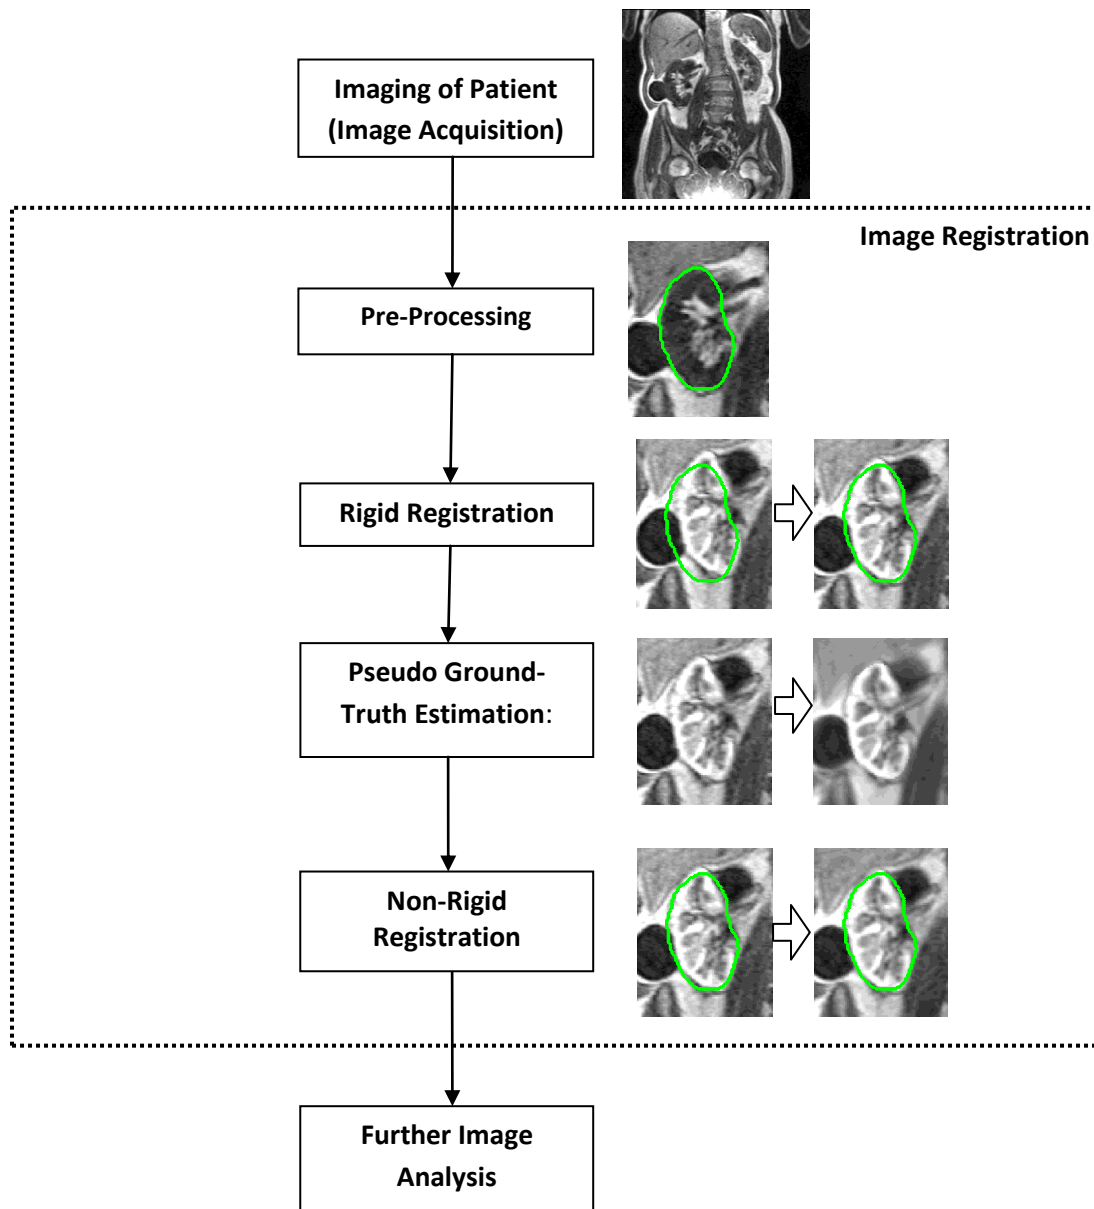
Fig.1.5. Flowchart of the proposed registration algorithm.

CHAPTER 2

LITERATURE REVIEW

This chapter provides a brief review on the various image registration methods proposed in rich literature. First, image registration methods will be classified into the different widely known categories (type of human organ, image modality, similarity measures, etc.) as briefly listed in Chapter 1. Then, a focused review will be done on the more relevant methods related to renal image registration and non-rigid image registration. The pros and cons of these methods will be analyzed. Finally, the methods will be compared in terms of the requirements of the registration algorithm in renal DCE-MRI.

As aforementioned, a multi-resolution algorithm will be proposed, consisting of both rigid and non-rigid registration methods. A rigid registration method must be able to handle large translations. Rotations are not considered as they are small and will be accounted for in the non-rigid registration step. The rigid registration method has to be robust with errors small enough to ensure that the subsequent non-rigid registration step can achieve sufficient accuracy. The non-rigid registration method, on the other hand, must be able to handle small deformations in the kidney efficiently and robustly because 3D volumes are involved.

2.1. Image Registration Algorithms

Medical image registration algorithms can be classified according to the organ of interest: brain imaging [10-15], cardiac imaging [8, 16, 17], spine imaging [18], renal

imaging [6, 16, 19-25], bladder imaging [26] and breast imaging [27]. Image registration on these organs allows a radiologist to perform an accurate diagnosis and to administer treatment effectively. These registration algorithms can be applied to different modalities, including MRI [9, 13, 14, 16, 17, 22, 24, 28], computed tomography (CT) [29, 30] and inter-modalities between the ones aforementioned [18, 31]. Each modality has its pros and cons, and displays different information about the organs that are being imaged. CT outlines the bones inside the body accurately. Even though CT is cheaper than MRI, CT uses ionizing radiation that is harmful to the human body. MRI, on the other hand, is accurate and non-ionizing, although it is more expensive and time-consuming. Moreover, MRI can produce a good contrast between tissues based on different weightings. For example, on a T_2-weighted scan, fluid-containing tissues are bright and fat-containing tissues are dark, whereas the reverse is true for T_1-weighted images. In some cases, such as when the patient has a metallic implant, MRI cannot be performed due to the high magnetic field (measured in Tesla) present during the procedure. The difficulty of image registration is high where multi-modalities are involved, as the coordinate systems used and the information (such as edges and contrast) extracted are usually different.

As mentioned in Chapter 1, image registration algorithms can be classified into rigid registration, affine registration or non-rigid registration. For rigid registration, most algorithms make use of the different similarity/difference measures defined in [5] and their variants to align the images, for example, mutual information (MI) [18, 22, 32, 33], gradient based similarity measure [24, 31], and cross-correlation [16]. Different similarity measures make use of different image information to quantify how much the reference image and the floating image look alike. For example, gradient-based methods make use

of object edges to align the images, and MI-based methods make use of probability theory and information theory to obtain the registration result. Some other unique methods in rigid registration employ point-based algorithms [34], wavelet and Fourier transforms [23], Markov random field (MRF) optimization [35] and alignment based on mid-sagittal plane (MSP) [10]. Affine registration allows for more degrees of freedom in transforming each image, although it is more computationally intensive than its rigid counterpart. Some affine registration methods are given in [13, 15, 28]. In the case of 3D affine registration, an efficient log-Euclidean poly-affine framework is described in [36]. For non-rigid registration methods, some of which discussed in [37], most of the similarity measures mentioned above are integrated into a higher level algorithm. Many methods involve the use of affine parameters [13, 28], B-splines [21, 29] and graph-cuts [11, 12, 26, 38] to account for the local deformations of the organ due to movement and pressure. While affine registration algorithms use linear models to align the images, non-rigid registration methods attempt to align the images by non-linear models with a higher degree of freedom to represent the elasticity of the transformation. Other methods include solving Laplace's equation between closed equally-spaced contours [19], thin-plate splines [39] and the demons algorithm [9, 30]. The difference between these methods lies in their efficiency and robustness in performing the non-rigid registration task.

2.2. Methods for Registration of Renal DCE-MRI Image Series

In recent research of medical image registration, the renal DCE-MRI image registration problem has mostly been dealt with rigidly [22-25], as it is popularly

assumed that kidneys do not exhibit non-rigid movements in most healthy patients. Mahapatra [22] proposed a method using an MI-based registration model to register 3-D volumes involving rotation and translation with respect to time. A saliency model is employed to determine the importance/utility of each voxel within an image; the saliency model improves the accuracy of rigid registration of the kidney. This solution reveals an interesting aspect of image registration, by differentiating active and passive voxels within an image. Active voxels are voxels that can enhance the accuracy of the registration algorithm when a certain similarity measure is used. In contrast, passive voxels will give an inconclusive result, lowering the accuracy of the registration algorithm. For example, a voxel containing the strong edge of an object is an active voxel, beneficial to the image registration algorithm when gradient difference is used as the similarity measure; a passive voxel located in the background with low edge strength will not be able to generate reasonable results in registration and should be given lower priority when computing the cost function. Furthermore, MI is a contrast invariant measure, and is used commonly to register DCE-MRI images. However, MI is computationally intensive given a large search space because the probabilities and joint probabilities have to be recomputed for each set of transformation parameters.

Song et al. [23] applied anisotropic diffusion to pre-process the image before using wavelet and Fourier transform to detect edges and then registered 3-D volumes with time in an automatic fast Fourier processing. Anisotropic diffusion acts as a pre-processing method to remove noise while keeping the important information such as edges, lines and other details intact. The wavelet and Fourier transform then detect the edges in the frequency domain for registration. It is noted that edges are important when

aligning DCE-MRI images, as edges are the only entity that does not shift when a contrast agent washes in and out of the kidney.

Sun et al. [24] used an integrated registration technique to match images from different phases in 2-D. A multi-level solution was proposed, where gradient information was first used to align the images, followed by a level-set segmentation of the kidney to obtain the rough kidney segmentation. Registration is then refined and re-applied onto the kidney images incorporating regional homogeneity of pixel intensities to obtain the solution. This method did not consider out-of-plane motion and the kidney mask obtained is therefore in 2D. Certain methods used in this registration algorithm are notable, in particular the multi-level registration approach and the use of 2D segmentation mask with gradient information to perform registration.

Lastly, Yim et al. [25] investigated two registration methods for registering 2-D kidney slices: (1) a semi-automatic method using contours as landmarks, and (2) an adaptation of the Automated Image Registration (AIR) algorithm. While being able to register DCE-MRI images of the kidney, the proposed methods did not guarantee good results when diseased kidneys are considered. However, it is to be noted that a good kidney boundary will allow for an easier registration of DCE-MRI images, as this boundary does not change with time, unlike the medulla of the kidney that appears only during the wash-in of the contrast agent affecting the results of edge matching.

2.3. Relevant Methods for Non-Rigid Image Registration

Many solutions are available to perform non-rigid registration in medical imaging, but most of these solutions are implemented for registering other organs in the body other

than kidneys. Two similar solutions are proposed by So [11] and Tang [12]. Graph-cuts optimization with alpha expansions is carried out with MI and sum of squared difference (SSD) being integrated into the cost function. The purpose of graph-cuts serves to register 2-D images where each pixel is represented by a node in the graph and the nodes are linked in a 4-neighbourhood fashion. However, when volumetric images are considered, there is a great tendency for the processing time to increase exponentially. Graph-cuts is normally used in segmentation with the setup being shown in Figure 2.1, with a 2D organized grid having 4-neighbourhood connectivity and 2 labels representing the background and foreground. It is observed that graph-cuts is a flexible optimization tool allowing the user to define the nodes and how they are linked together, the data cost and the smoothness cost between connected nodes. Additionally, multi-labels could also be defined instead of binary-labeling which is mostly used in segmentation; multi-labels fit the context of rigid registration with a pre-defined search space.

Another solution employing graph-cuts was introduced in [38]. The proposed algorithm integrates both saliency and intensity information into a Markov Random Field (MRF) framework to register DCE-MRI kidney images. Similar to the graph-cut method as mentioned precedent using MI and SSD, edge and saliency information were both considered in the employed solution. It is observed that graph-cut is a very flexible tool which can solve MRF energy functions defined using different similarity measures and smoothness constraints.
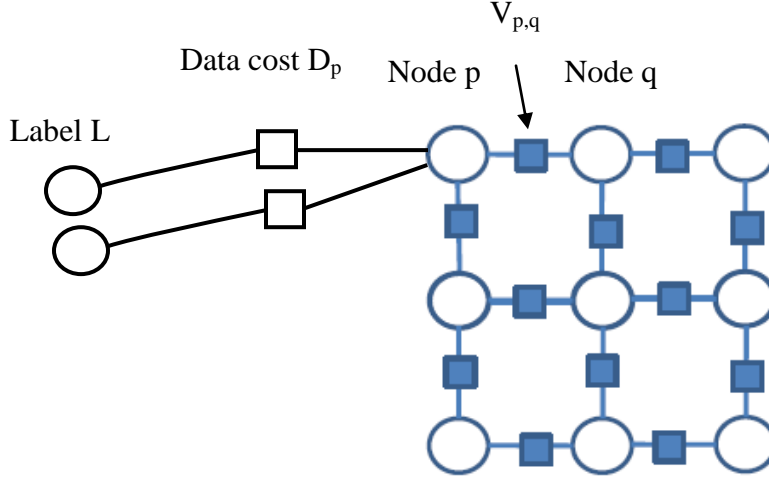
Fig.2.1. Standard grid graph.

Solutions in the form of free-form deformation (FFD) based on B-splines were introduced in [21, 29]. The transformation is defined by a deformation model based on a grid of control points described by B-splines basis functions where the grid spacing is user specified. A large grid-spacing allows for an efficient modeling of global deformations, whereas a small grid-spacing is able to detect high local deformations. For kidney registration, small local deformations are more common but this causes the FFD algorithm's complexity to greatly increase; having a small grid-spacing is undesirable due to higher computational costs. There is a tradeoff between local accuracy and computational cost, which is controlled by the grid spacing. Therefore, the grid spacing is an important parameter and should be set appropriately.

Finally, a solution based on diffusion models to perform image-to-image matching is presented – the demons model. The demons model is described in [9] with a fast variant of it being implemented in [30]. The demons algorithm performs optimization over the entire space of displacement fields, and matching is done iteratively by diffusion

based on a deformable grid model. The demons model for non-rigid registration is much faster than the other two methods described above (i.e., graph-cuts and FFD) when 3D non-rigid registration is involved. If the deformation grid is simple, where the displacement of each grid point is small, the demons algorithm is able to perform a robust non-rigid registration of the image.

2.4. Comparison of Methods

In this section, we compare the methods proposed in rich literature in terms of the requirements of the registration algorithm. First of all, we compare the two most common similarity measures, mutual information and gradient difference. As edge information is used frequently to register renal DCE-MRI images, gradient difference is the preferred choice. Besides, when rigid registration is concerned, we have a large number of labels where each label represents a certain translational transformation. In MI-based methods, when we shift the window of the kidney across the search space, the joint entropy needs to be re-computed, making the MI-based methods more computationally intensive than gradient difference based methods.

Graph-cuts optimization is employed to solve the problem of rigid registration due to the nature of graph-cuts being flexible, efficient and robust. With graph-cuts, many different parameters can be set according to application requirements: (1) it is possible to improve the emphasis on the nodes around the boundaries of the kidney as these boundaries contain useful gradient information ensuring a robust solution; (2) the use of multi-labels allows to search for the best translational parameters; (3) smoothness

between the nodes can be set in such a way that labels assigned to neighboring nodes do not represent very different transformation results; (4) sparse smoothness can be used to define the node structure and links between the nodes in any way as desired by the user. Being flexible in varying the node-structure and the size of nodes, a reasonable solution can be achieved with small translational errors, as long as the energy function follows the criteria as stated in [40]. Graph-cuts' flexibility and efficiency fit the requirements of the registration algorithm. The downsides of using graph-cuts are as follows: (1) the lost of label orders where mid-label results are not considered; (2) the computation of the data cost for each label on each of the nodes is exhaustive, even though the optimization procedure is efficient and robust.

Next, we analyze the problem of selecting a reference image to register the dataset. For rigid registration, the user selects a frame where the kidney exhibits the best boundaries surrounding it (strong edges and good gradient information). One reference frame is sufficient since the kidney boundaries are often present before, during and after contrast enhancement. For non-rigid registration, in order to minimize the intensity difference between the reference and the floating images, some existing methods [16, 41] register every two consecutive frames for the whole sequence. However, as each pair of consecutive images would still exhibit varying intensities for different tissues in the kidney, registration errors tend to accumulate and the results become undesirable towards the end of the sequence. To overcome this problem, we propose to first estimate a motion/noise free sequence and then apply non-rigid registration to corresponding pairs of images between the observed sequence and the estimated sequence. Compared to pairs of images within the observed sequence, the pairs of corresponding images have similar

intensities, and hence the registration problem is greatly eased. We call the estimated sequence 'pseudo ground truth', because it is the estimate of the image sequence that would have been acquired without being affected by motion or noise. However, to the best of our knowledge, such estimation methods do not exist for renal DCE-MRI data. Thus, an estimation method for the pseudo ground-truth will be borrowed and adapted from myocardial image registration in [8]. The pseudo ground-truth is obtained by means of intensity curve-fitting across the whole time-series volumes for each voxel, and this eliminates noise and small motion allowing for a reliable non-rigid registration. To ensure a reasonable estimation of the pseudo ground-truth, it is imperative that the translationally registered dataset does not contain large translational errors.

Finally, we compare the algorithms used to perform the non-rigid image registration. As mentioned precedent, FFD using b-splines is reliable but it is too computationally intensive when 3D image registration accounting for small local deformations is concerned. FFD using b-splines becomes unreliable if we use a large grid-spacing and thus is not a suitable candidate for non-rigid registration. The demons algorithm, on the other hand, is efficient and accounts for local deformations satisfactorily.

CHAPTER 3

TRANSLATIONAL REGISTRATION

As stated in the introduction, motion and breathing of a patient are inevitable as each patient dataset is acquired over minutes during an MRI procedure; misalignment including large translations and small deformations between kidneys in volumes acquired at different times result from such motion. To address this problem, a multi-level registration algorithm is proposed, where an initial translational alignment of kidneys is performed to eliminate large translational errors between kidney volumes. Translational registration ensures that the kidneys in each volume are kept within a small region of interest (ROI) increasing the efficiency and accuracy of the non-rigid registration step. The challenges of performing translational registration on kidneys lie in: 1) defining a metric to determine the similarity between images; 2) determining the search space; and 3) aligning kidneys robustly up to a small translational error. A good similarity metric should be sensitive to the matching of the images even in the case where a slight mismatch is present; matching kidney images will result in a high value, and this value decreases as the kidneys in the reference and floating images become further apart. The search space has to be large enough to ensure all large translational errors are accounted for, but a larger search space comes at the expense of a higher computational cost in registration. Errors are inevitable in a registration algorithm, and it is desirable if the registration errors are kept as small as possible for all the available datasets. The method to address the above challenges is discussed in this chapter.

An outline of the algorithm is given in Figure 3.1. The proposed translational registration algorithm consists of the following two steps: 1) manual selection of the ROI

and obtaining a 2D segmentation mask of the kidney in a pre-processing stage; and 2) resolving the first registration stage by using graph-cuts to register the floating image volume to the reference image volume in the 3 principle axes using a large translational search space. In the second step, we first show how the graph is set-up and the different terms (data cost, smoothness cost and weights) are defined. This is followed by performing the graph cut operation to obtain the translational parameters for each node within the graph. Finally, we perform a weighted average over the translational parameters obtained for each node within the graph to obtain a set of global translational parameters for the entire volume.

The algorithm is regarded as semi-automatic due to manual input involved in drawing a box depicting the ROI and in obtaining the seeds to be used in the grow-cut segmentation algorithm [42]. However, minimal time is spent on the manual input as compared to the main registration algorithm, and these pre-processing steps will ultimately ensure an efficient and accurate registration. To summarize, the main contributions of the translational registration step are: (1) Graph-cut is implemented with super-nodes where each of these super-nodes represents a local sub-volume of the kidney; (2) Gradient difference is incorporated into the energy function of the rigid registration algorithm.
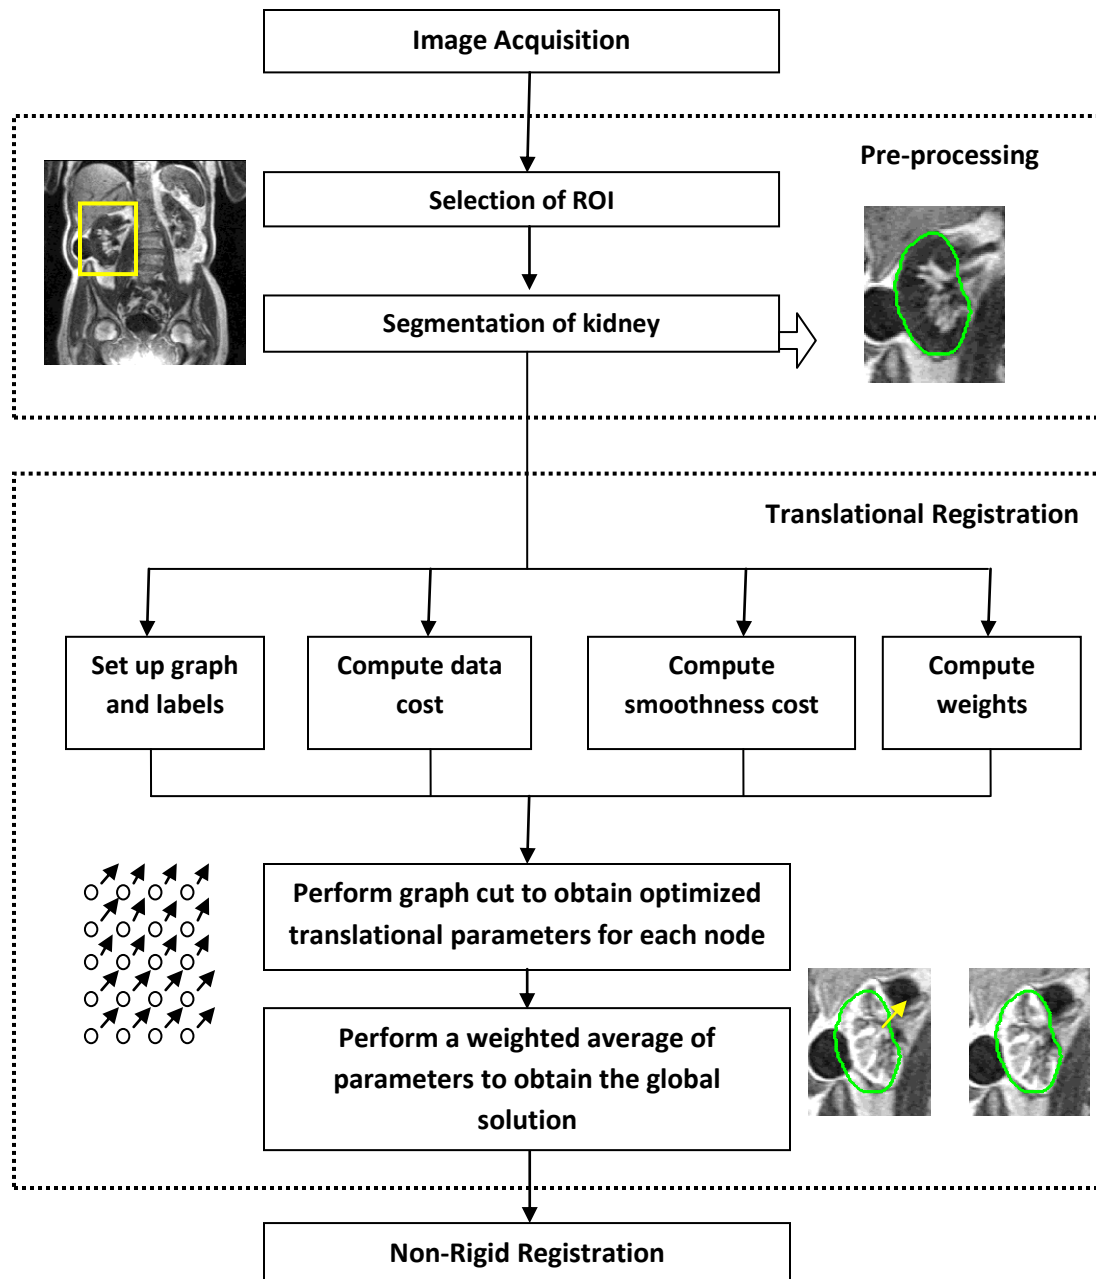
Fig.3.1. Flowchart of the translational registration algorithm

3.1. Pre-Processing

In the pre-processing step, the kidney is extracted from the reference image in the form of a rectangular box and a 2D segmentation. The rectangular box serves as a window for matching volumes at different time frames, while the 2D segmented kidney slice differentiates the kidney from its background within the ROI. For each patient dataset, two kidney time-series volumes (left and right kidneys) can be extracted and each time-series volume is registered separately. A sample time-series volume is shown in Figure 3.2. The reference frame, middle slice and the ROI are manually selected by the user, where: (1) the user determines the best time frame at which the kidney exhibits the best contrast around kidney boundaries and uses it as the reference frame, (2) the user then selects the kidney slice with the largest cross sectional area as the middle slice, and (3) the user draws a box over the kidney region to crop the ROI for registration. Five slices of the kidney, inclusive of the middle slice and the two slices immediately above and below the middle slice, are included in the volume to determine the best translational parameters efficiently in this registration stage.
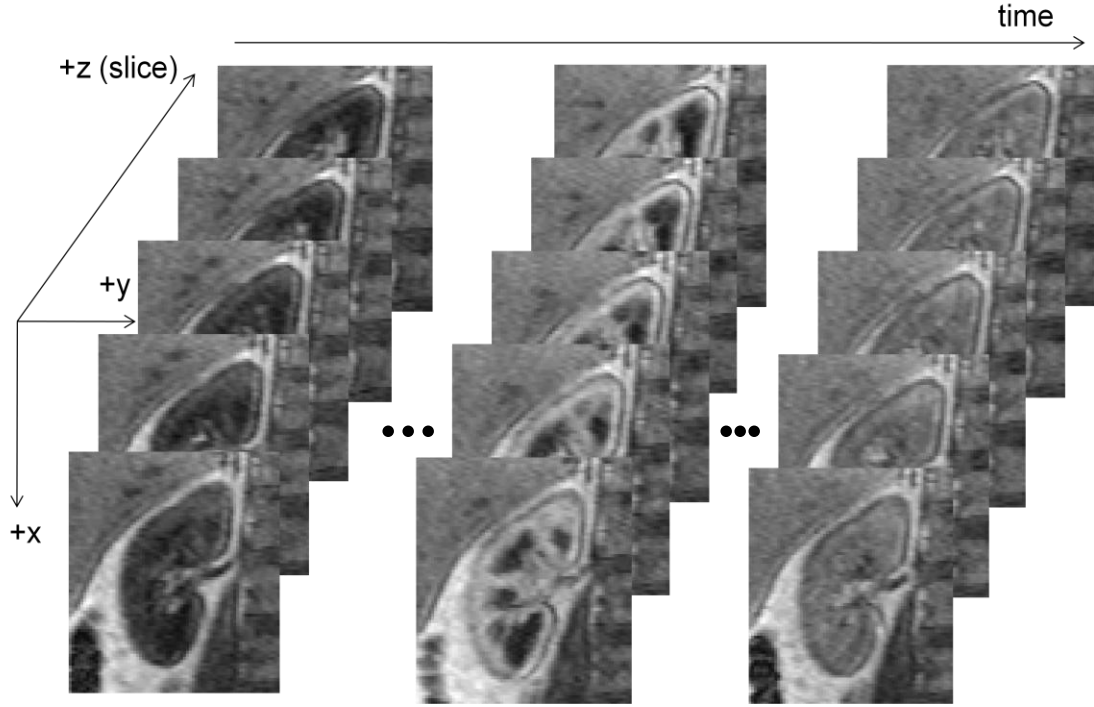
Fig.3.2.A sample time-series volume of a kidney. Each slice in the Z axis consists of a cross-section of the kidney in the XY plane. The 3-D volume of a kidney varies in intensity with time.

Upon obtaining the reference frame and the ROI, the boundary of the kidney is then obtained by 2D segmentation using a grow-cut algorithm described in [42], where the object and background seed points are determined manually by the user for only the reference image slice. The Matlab wrapper used to perform the grow-cut segmentation is obtained at [43]. Figure 3.3 shows a sample of the grow-cut segmentation over a slice of the kidney from the reference frame. The resulting segmentation gives the boundary around the kidney which is used for the subsequent registration steps. It is noted that the segmentation need not be precise, and that the user can re-input the seeds for the grow-cut segmentation if the results are undesirable.
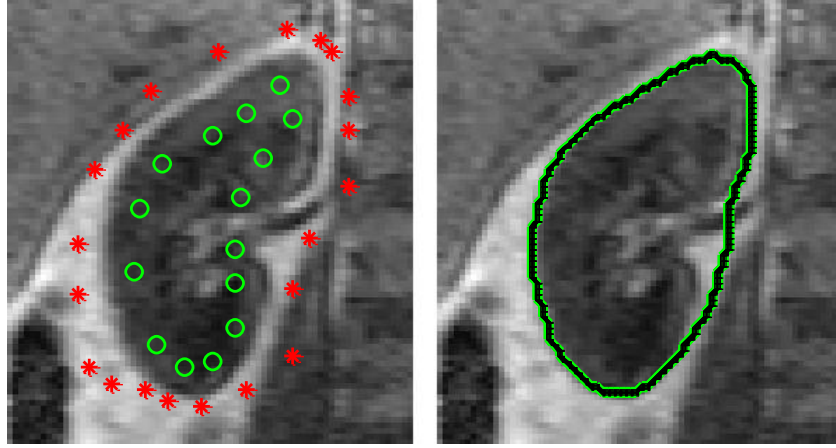
Fig.3.3.A sample slice of the kidney: (left) with the background seeds (marked with * in red) and object seeds (marked with o in green), (right) with the kidney boundary obtained by grow-cut segmentation.

3.2. Graph Cuts

We propose to solve the translational registration problem by means of a multi-label graph-cut optimization technique via alpha-expansions, the pros and cons of which have already been discussed in Section 2.4. As mentioned precedent, graph-cut is a flexible, efficient and robust algorithm. The loss of label orders is not a major problem as the kidney alignment will be completed with a non-rigid registration step. A graph is defined by its nodes and the links between each node and its neighbors. A graph-cut solution then looks through the entire search space denoted by the number of labels to obtain the optimal global solution (the best label to each node in the graph). The next few sub-sections will discuss in detail how the rigid image registration parameters are set up, in terms of: energy function, node topology, labels, data penalty, smoothness cost and weights. A final section will then explain how this registration step is implemented and validated.

3.2.1 Energy Function

To solve the translational registration problem, we propose to minimize the energy function, $E(L)$, based on the Bayesian labeling of first-order Markov Random Fields (MRFs) as below:

$$E(L) \; = \; \sum_{p \in P} W_p D_p(L_p) + \lambda \sum_{(p,q) \in N} V_{pq}(L_p, L_q), \tag{3.1}$$

where $L \; = \; \{L_p | p \in \mathbb{P}\}$ is a set of labels attributed to each node p within a set of nodes $\mathbb{P}$ of the image volume, W_p is the weight that controls the importance of certain nodes, D_p is the data penalty term, V_{pq} is the smoothness constraint term, λ is a constant set at 0.01 to balance the importance between the data cost and the smoothness constraint, and N is the set of all the pairs of neighboring nodes. Each parameter will be defined in the subsequent sub-sections, together with the implementation method to obtain the optimum translational values in the three axes.

3.2.2 Graph Structure

In translational registration, the image is partitioned into a number of volumes called super-nodes. The centers of these super-nodes are set in the form of a regular square grid on top of the kidney ROI. Each super-node/sub-volume contains the gradient information from the observed slice in the *XY* plane and the 2 anterior slices and 2 posterior slices in the *Z* axis; the node sizes are $25 \times 25 \times 5$ ($x \times y \times z$) for each volume partition. The size of each node in the *XY* plane is also equally sized and an overlap between each pair of neighboring nodes is implemented by using a node-spacing of $12 \times$

12 ($x \times y$) to increase the dependency between neighboring nodes. Because super-nodes are used, the number of nodes in the graph is kept small at about 20-30, thus increasing the efficiency of the algorithm. The segmented boundary, gradient map and the graph of nodes are shown in Figure 3.4.
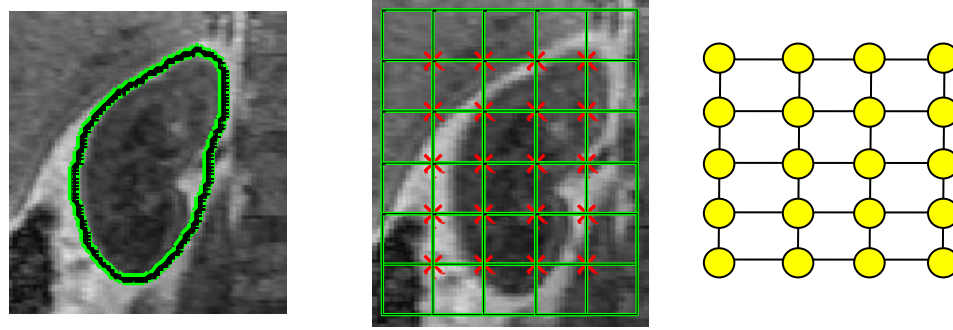
Fig.3.4. (left) ROI with segmented boundary; (middle) Node map with node-centers marked by Xs and overlapping node-windows; (right) Node distribution with links.

3.2.3 Labels

Each label, L_p, represents a certain transformation configuration for a particular node p, where $L_p \equiv \{(\Delta x, \Delta y, \Delta z) | \forall \Delta x, \Delta y, \Delta z \in \mathbb{Z}\}$. The number of labels is equal to the number of permutations of the possible values of translations in the X, Y and Z axes of the search space, which is set to be $\{-18 \leq \Delta x \leq 18, -10 \leq \Delta y \leq 10, -3 \leq \Delta z \leq 3, \Delta x, \Delta y, \Delta z \in \mathbb{Z}\}$. The search space is set with reference to the ground truth information, with an extended ± 2 voxels in each dimension of the search window. During breathing, the diaphragm contracts and causes the lungs to expand, and the kidneys in the body are pushed downwards; thus, the search space is larger in the X-axis (top to bottom of the body). The search space in the Z-axis is small as the motion of the

kidney in this axis is generally small, and that the resolution of the image is 1.66 mm $\times$ 1.66 mm $\times$ 2.5 mm where a 3-voxel shift in the X or Y-axes is equivalent in distance to a 2-voxel shift in the Z-axis.

3.2.4 Data Cost

The data cost is defined by how different the floating image is from the reference image; the greater the difference between the floating image and the reference image, the higher the data penalty. A robust measure is required to ensure that an accurate registration result is obtained for most or all of the kidney volumes. For DCE-MRI images of the kidney, edges are the most reliable features, especially those around the boundaries of the kidneys. Therefore, gradient difference is used as the main similarity measure where each node in the floating image was compared to the corresponding node in the reference image. Let $(\Delta x, \Delta y, \Delta z)$ denote the current 3D displacement, $F_x(\boldsymbol{u})$ and $F_y(\boldsymbol{u})$ denote the corresponding X-axis and Y-axis edge information of the voxel with position $(\boldsymbol{u}) \equiv \{i, j, k\}$ in the reference volume image, while $G_x(\boldsymbol{u}')$ and $G_y(\boldsymbol{u}')$ denote their counterparts in the floating volume image with $(\boldsymbol{u}') \equiv \{i + \Delta x, \ j + \Delta y, \ k + \Delta z\}$. Both F and G are obtained by performing convolution of the image feature with a Gaussian filter. The cost, D_p for a particular node p to take the label, L_p is given by:

$$D_p(L_p) = 1 - \frac{\sum_{v \in B_p} |F_x(\boldsymbol{u}).G_x(\boldsymbol{u}') + F_y(\boldsymbol{u}).G_y(\boldsymbol{u}')|}{\sum_{v \in B_p} \sqrt{(F_x(\boldsymbol{u})^2 + F_y(\boldsymbol{u})^2).(G_x(\boldsymbol{u}')^2 + G_y(\boldsymbol{u}')^2)}} \qquad (3.2)$$

where B_p is the set of all voxels within the sub-volume corresponding to node p. Each voxel, v, in the floating image is compared against its corresponding voxel in the

27

reference image in terms of edge magnitude and edge orientation. Voxels with matching

orientation angles are given a high similarity value which translates to a low data cost as

shown in (3.2).

3.2.5 Smoothness Cost

In order to ensure that neighboring nodes are not assigned to labels denoting

translations that differ a lot from each other, a smoothness constraint is introduced; the

smoothness constraint provides a force that pulls nodes together like magnets. Given a

particular 3D displacement $(\Delta x, \Delta y, \Delta z)$ for node p, and another 3D

displacement $(\Delta x', \Delta y', \Delta z')$ for neighboring node q, where $L_p \equiv (\Delta x, \Delta y, \Delta z)$

and $L_q \equiv (\Delta x', \Delta y', \Delta z')$, the smoothness cost, $V_{pq}(L_p, L_q)$, is set as follows:

$$V_{pq}(L_p, L_q) = \begin{cases} 0 & \text{if } d(L_p, L_q) = 0 \\ 0.1 & \text{if } 0 < d(L_p, L_q) \leq 1 \\ 0.2 & \text{if } 1 < d(L_p, L_q) \leq 2 \\ \infty & \text{otherwise.} \end{cases} \tag{3.3}$$

where $d(L_p, L_q)$ is the Euclidean distance between the displacements denoted by labels

L_p and L_q. It can be observed from (3.3) that if neighboring nodes are given varying

labels with Euclidean distance higher than 2, the smoothness cost will be infinity,

signifying that such an attribution of labels is impossible within a graph.

3.2.6 Weights

The weights, W_p, assigned to each node p are determined by the strength of

gradient information found within its corresponding sub-volume, and is given by

$$W_p = \begin{cases} \frac{S_p - S_{\min}}{S_{\max} - S_{\min}}, & \text{if node contains kidney boundaries} \\ 0, & \text{if node does not contain kidney boundaries} \end{cases} \quad (3.4)$$

$$S_p = \sum \sum \sum_{(i,j,k) \in B_p} M_f(i, j, k); \quad (3.5)$$

$$W'_p = W_p / \sum_{p \in P} W_p, \quad (3.6)$$

where n_p is the number of voxels within node p, S_p indicates the strength of node p, M_f represents the magnitude of the gradient of the reference image f, $S_{\max}$ and $S_{\min}$ respectively represent the minimum and maximum strengths of all the nodes. Greater emphasis is placed on active nodes with a larger overall gradient magnitude; less emphasis is placed on passive nodes with a small overall gradient magnitude. Moreover, weights are set to zero if a certain sub-volume (defined by a node) does not contain any kidney boundaries, as other edges within the kidney are not persistent throughout the whole time-sequence. In (3.6), we normalize the weights in (3.4) such that the sum of weights of the nodes will be equal to unity.

3.3 Implementation and Validation

We make use of a Matlab wrapper [44] that implements the graph-cut technique [40, 45, 46] to solve the registration problem involving only translation of each node in the *X*, *Y* and *Z* axes. By partitioning the image volumes and employing the use of graph-cuts, we are able to optimize the 3-DOF translational registration process while ensuring that dependency between neighboring nodes is observed; no two neighboring nodes have results that differ from each other by a large margin, thus obtaining the best result

globally. The resulting global translational parameters $(\delta x, \delta y, \delta z)$ are then obtained by taking the weighted average of displacement over all nodes:

$$(\delta x, \delta y, \delta z) = \sum_{p \in \mathrm{P}} W'_p \cdot (\Delta x, \Delta y, \Delta z)_p, \qquad (3.7)$$

where $(\Delta x, \Delta y, \Delta z)_p$ is the label attributed to node p after the graph-cut algorithm.

To validate the translational registration algorithm described in this chapter, many different tests are performed in a quantitative analysis. As the ground truth of the dataset denoting the optimum global 3D translation is available, it is used to obtain the registration accuracy for the different tests. The mean error, standard deviation of error, and cumulative distribution function of the errors for each volume are evaluated for each of the test results. The first test involves comparing the similarity measures, gradient difference against normalized mutual information. Three different implementations are considered: 1) gradient difference used along with graph-cuts optimization to obtain the results; 2) MI applied globally to compare the difference between the floating image and the reference image; 3) MI used along with graph-cuts optimization to register the images. The second test involves comparing different node sizes when constructing the graph. A larger node size contains more information making the algorithm more robust, but a longer time is required to compute the data cost reducing the algorithm's efficiency. Thus, the second test determines the minimal amount of information required in each super-node in order to obtain reasonably accurate results. The criterion for a result to be marked as reasonable is that the translational error obtained after registration is kept within a maximum Euclidean distance of 5 voxels. To show that the introduction of weights is able to ensure better registration accuracy, a third test is performed. Weights, as mentioned precedent, help to locate regions containing important and consistent edge

information; the regions around the kidney boundaries are used to register the images. A fourth and last test is performed on two simulated datasets, created by performing a random translational transformation on two pre-aligned datasets. The random transformation is then recovered by using the proposed registration algorithm and the errors between the obtained and applied translational parameters are analyzed visually and quantitatively. The fourth test validates the choice of using a single reference frame to register the entire dataset.

Additionally, a qualitative analysis is performed on the registration results. Apart from performing a visual analysis on the resulting registered images, a curve showing the average intensity over a few voxels of the kidney against time is constructed. If the kidney is registered correctly using the proposed registration algorithm, the intensity-time curve obtained post-registration will tend to the curve obtained by aligning the images using the ground-truth data.

CHAPTER 4

NON-RIGID REGISTRATION

The rationale of performing non-rigid registration after translational registration is to account for small rotations and also local deformations of the kidneys. Such deformations are also caused by motion and breathing of the patients. Rotation parameters or affine parameters have been used to define the deformations within a kidney with limited degrees of freedom. However, these parameters cannot account fully for deformations within a kidney, especially in 3D, and thus a model of higher order is required to fully describe such deformations. The main challenges of non-rigid registration are: 1) to obtain a high efficiency of the algorithm in 3D context, and 2) to have a good reliability of the registration algorithm performed on contrast-varying kidney images. Non-rigid registration methods as described in the literature are mostly 2D-based, thus efficiency is not a major concern for such algorithms. However, for 3D images, the computational complexity increases exponentially, thus efficient registration algorithms are required. The reliability of the non-rigid registration algorithm is dependent on two factors – the structures extracted from the image and the reference image. For DCE-MRI images, edges are the best structures which can be used to align images even as tissues display varying intensity profiles. However, a kidney shows different edge information at different time frames due to the effect of a contrast agent, so a single reference image could not be used to register all kidney images. In order to achieve optimal accuracy in non-rigid registration, a reference image should resemble that of an image from the observed sequence. Thus, a pseudo-ground truth dataset is estimated from the translationally registered dataset to serve as reference images for

registering each floating image at different time frames. A pseudo-ground truth image has all the tissues aligned with the rest of the images within a dataset, and the intensity of each voxel follows a certain intensity profile depending on the type of tissue. The pseudo-ground truth dataset, being noise and motion free, serves as the best reference images to register the entire dataset.

Figure 4.1 gives a flowchart of the non-rigid registration step. The non-rigid registration method consists of the following steps: 1) initial pseudo ground-truth dataset estimation given the translationally registered dataset; 2) resolving the non-rigid registration stage by aligning each translationally registered volume with its corresponding pseudo ground-truth volume at a particular time frame using the demons algorithm. In the first step of non-rigid registration, we compute the different terms (data fidelity, spatial smoothness and temporal smoothness) from the given dataset before performing an optimization procedure to obtain the pseudo ground-truth dataset. In the second step, we first perform a calibration on the pseudo ground-truth dataset by aligning the first static volume with its corresponding pseudo ground-truth volume. Finally, we perform a non-rigid registration on the rest of the dataset using the demons algorithm.

To summarize, the main contributions of the non-rigid registration step are: (1) Pseudo ground-truth dataset estimation previously applied to cardiac images is introduced into renal image registration; (2) Application of the demons algorithm to DCE-MRI renal images to account for non-rigid transformations.
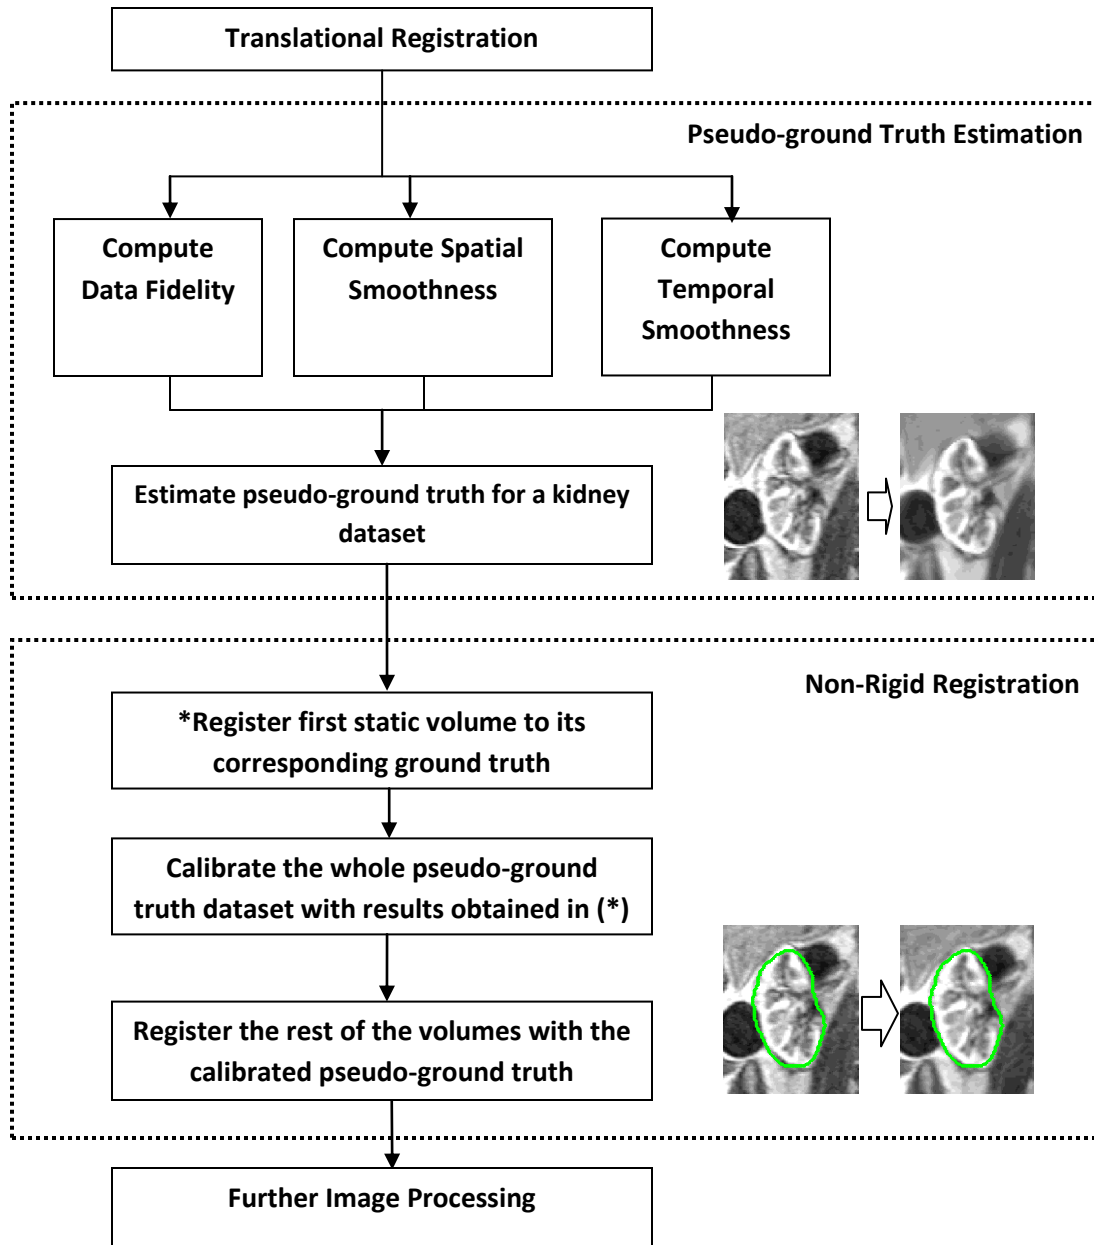
Fig.4.1. Flowchart of the non-rigid registration algorithm.

4.1. Pseudo Ground-Truth Estimation

The pseudo ground truth is an estimate of the image volume time-series that would have been acquired without being affected by noise or motion during acquisition. Figure 4.2 shows the intensity-time curve for the mean intensity across the entire kidney mask. From the curve, we can observe that the intensity change is not smooth, mainly due to non-rigid deformations and misalignments. Thus, another image volume time-series is estimated from the original dataset to produce its noise/motion free counterpart. In this step, an estimation of the pseudo ground-truth of the translational registered dataset is performed with reference to the method as given in [8]. The pseudo ground-truth dataset is then used as a reference dataset to register the translationally registered dataset. Although the algorithm has been used on cardiac images with multiple segments (background, left and right ventricles, and the myocardium), a minor adaptation is made to the algorithm to make it work for renal images with only the kidney and the background. The pseudo ground-truth estimation algorithm proposed in [8] works only on 2D images. Taking into consideration that an implementation for estimating pseudo ground-truth in 3D will result in an exponentially increasing computational complexity with similar results, it is more efficient to perform the estimation several times for each kidney volume. Thus, seven 2D slices of the kidney are extracted from each kidney time-volume and the estimation algorithm is performed seven times for each slice across the whole time sequence. In addition, since the intensities across different time-volumes affect the estimation of the pseudo ground-truth dataset, it is important that the translationally registered dataset: (1) should not contain large translational errors, and (2)

should not contain translational errors consecutively across different time-volumes. Next, we describe in detail the main components of the estimation algorithm: data fidelity, spatial smoothness, temporal smoothness and the optimization method.
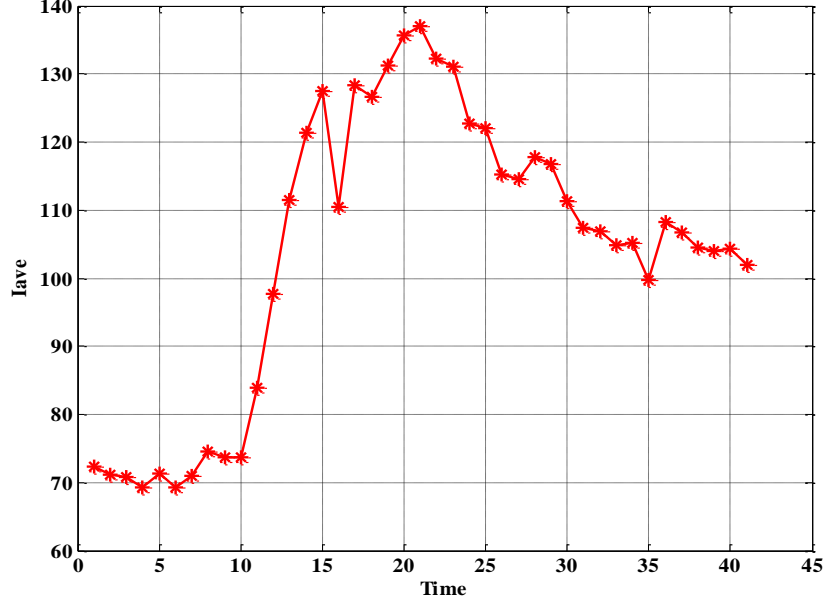
Fig.4.2. Intensity-time curve for the average of pixel intensities (I_{ave}) over the kidney mask across all frames.

4.1.1 Energy Function – Pseudo Ground Truth

The energy functional, $E(H(g), f')$, used to perform the estimation of the pseudo ground-truth dataset is given as follows:

$$E(H(g), f') = E_d(H(g), f') + E_s(f') + E_t(f'), \qquad (4.1)$$

where g represents the floating image, H represents the image transformation function, f′ represents the pseudo ground-truth image (reference image), E_d is the data fidelity term, E_s is the spatial smoothness term and E_t is the temporal smoothness term. Each of these terms within the energy function is controlled by a constant within to determine its

importance to the estimation algorithm. The data fidelity term determines the similarity of pixel intensities between the floating image and the estimated pseudo-ground truth image. The spatial smoothness term and the temporal smoothness term, on the other hand, ensure that the intensities of each pixel follow the profiles of the corresponding segmented kidney and tissue type.

4.1.2 Data Fidelity

The data fidelity term, $E_d(H(g), f')$, is defined as the sum of squared difference (SSD) between the intensities of each pixel within the reference image and the floating image, and is given by:

$$E_d(H(g), f') = (H(g) - f')^T(H(g) - f'), \qquad (4.2)$$

where f' and $H(g)$ represent the intensity vectors of both the reference and the transformed floating images, respectively. The data fidelity term ensures that the estimated pseudo-ground truth resembles the original kidney images.

4.1.3 Spatial Smoothness Constraint

The spatial smoothness constraint term is used to ensure smoothness within different segments of an image while maintaining sharp boundaries between these segments, in this case between the kidney and the background. The kidney boundary obtained in the pre-processing stage of the translational registration is reused in this section to differentiate between the boundary pixels and the non-boundary pixels. An

estimated sequence without using spatial smoothness constraint will result in blurred boundaries due to motion, which is undesirable for non-rigid registration. By considering the intensities of non-boundary pixels in each segment less likely to be affected by motion, the ambiguities in determining the intensities of boundary pixels are resolved. The spatial smoothness term, $E_s(f')$, is therefore defined as:

$$E_s(f') = \sum_{k=1}^{K}(Ds_k f')^T W_k(Ds_k f'), \tag{4.3}$$

where K is the number of neighboring pixels considered, Ds_k is the first order spatial derivative operator along the direction between each pixel and the k^{th} operator, W_k is the weight matrix to ensure the intensity difference between the pixels from different segments (the background and the kidney) are not penalized. In this case, $K = 4$, as a 4-neighborhood system is employed.

In order for pixels belonging to the same segment to exhibit similar intensities and for their signals to maintain similar temporal dynamics, W_k, the spatial weight matrix is defined as:

$$w_k = \begin{cases} \alpha \rho_{ijk} \exp\left\{-\frac{[g(i,j,t_s)-g(i_k,j_k,t_s)]^2}{2\sigma^2}\right\}, & \text{if } \rho_{ijk} > \rho_{th} \quad \text{and } L'_{ij} = L'_{i_k j_k} \\ 0, & \text{otherwise} \end{cases} \tag{4.4}$$

$$W_k = \text{diag}(\text{vec}(w_k)), \tag{4.5}$$

where α is a positive weight term, ρ_{ijk} is the correlation coefficient between the intensity-time curves of the pixel at position (i,j) and its k^{th} neighbor at position (i_k, j_k), t_s is the frame number, L' is the label matrix for the kidney, ρ_{th} and σ defines the minimal signal similarity and the intensity variance within the same tissue type respectively. In the proposed algorithm, $\rho_{th} = 0.85$ and $\sigma = 0.3$.

4.1.4 Temporal Smoothness Constraint

The temporal smoothness constraint uses the temporal neighborhood of each frame to estimate its counterpart in the pseudo ground-truth. Two different phases are identified in the time-sequence of the kidney dataset: (1) the pre-contrast phase and (2) the perfusion phase. The pre-contrast phase is determined by locating the time frames before which the contrast agent washes into the kidney and the intensity-curve starts to rise. As the contrast agent has no effect on the background and on the kidney in the pre-contrast phase, the resulting signal at these locations should be approximately constant; the first order temporal derivative of the signals at the background or on the kidney during the pre-contrast phase should be close to zero. Then, as the contrast agent perfuses through the kidney, a piece-wise linear relationship can be observed approximately in the perfusion signals; the second order temporal derivative of the pseudo ground-truth should be penalized in this phase. The temporal smoothness term, $E_t(f')$, is defined as:

$$E_t(f') = (D_{1t}f')^T W_t (D_{1t}f') + \beta_2 (D_{2t}f')^T (D_{2t}f'), \tag{4.6}$$

where D_{1t} and D_{2t} are the first and second order temporal derivative operators, and W_t is the temporal weight matrix defined as follows:

$$w_t = \begin{cases} \beta_1, & if\ (i,j) \in \text{kidney and } t_s \in \text{precontrast phase} \\ \beta_1, & if\ L_{ij} \in \text{background} \\ 0, & \text{otherwise} \end{cases}, \tag{4.7}$$

$$W_t = \text{diag}(\text{vec}(w_t)) \tag{4.8}$$

The constants β_1 and β_2 are set as 15 and 3, respectively. A higher β_1 value will reduce the effect of noise in the pseudo-ground truth estimation and β_2 affects the amount of smoothing within the kidney.

4.1.5 Optimization Method

To obtain the pseudo ground-truth estimation f' from the translationally-aligned dataset g, we derive from (4.1) by letting $\frac{dE}{df'} = 0$ a system of linear equations as follows:

$$\left[I + \sum_{k=1}^{K}(Ds_k{}^T W_k Ds_k) + D_{1t}^T W_t D_{1t} + \beta_2 D_{2t}{}^T D_{2t}\right]f' = g \tag{4.9}$$

Let $A = \left[I + \sum_{k=1}^{K}(Ds_k{}^T W_k Ds_k) + D_{1t}^T W_t D_{1t} + \beta_2 D_{2t}{}^T D_{2t}\right]$, and the above equation becomes $Af' = g$. To solve equations of the form $Af' = $ g, the conjugate gradient method [47] is used as it is more efficient than a traditional method like Gaussian elimination [48]. Upon solving (4.9), we obtain the optimum pseudo ground-truth dataset f' which minimizes the energy function in (4.1), to be used as the reference dataset for the subsequent non-rigid registration.

4.2. Deformation Refinement with Demons Algorithm

The final step of the registration algorithm is to account for local non-rigid transformations within the kidney in 3D. In the previous step of pseudo ground-truth dataset estimation, 7 slices per volume are obtained. For 3D non-rigid transformation, 7 slices per volume is also considered, but only 5 slices of the final resulting transformed image will be evaluated because interpolation from registration causes the boundary voxels to have irregular intensities. For FFD using a small grid-spacing, the computational time becomes too long (> 1 hour for a volume time-series) and thus is undesirable. The demons algorithm is considered, because of its efficiency as compared to FFD-based methods as well as its reliability when dealing with small non-rigid

deformations. The demons algorithm is implemented following [9, 30] in which the algorithm is described in detail with the relevant equations.

The demons algorithm is an automatic non-rigid image registration algorithm, accounting for non-rigid deformations of the subject by means of a diffusion model. Demons forces are estimated using the optical flow formula. The implemented version of the demons algorithm is an improved version with a much better efficiency due to an adaptive force strength adjustment during the iterative process. Moreover, a slight adaptation of the algorithm also makes the registration edge-emphasized by using the edge strength of the image to control the optical flow in different regions. By emphasizing the edges, the smoothing/noise-removing effect caused by the dataset estimation will minimize the effects of texture distortion in the registered images. A normalization factor proposed by [49] allows the force strength to be adjusted adaptively over different iterations to determine the scale of deformation.

4.2.1 Energy Function – Demons Algorithm

The energy function (equation (3) in [50]) to minimize in the case of the demons registration algorithm, between a translationally registered image g (floating image) and a pseudo-ground truth image f' (reference image), is given below as:

$$E(\boldsymbol{V}) = (f' - g \circ (\boldsymbol{S} + \boldsymbol{V}))^2 + \frac{\sigma_i^2}{\sigma_x^2}.\boldsymbol{V^2}, \qquad (4.10)$$

where $\boldsymbol{V}$ is the update of the transformation field $\boldsymbol{S}$ describing the translation in the 3 principle axes of every voxel from its original position, $\circ$ defines the image transformation function, σ_i^2 and σ_x^2 represent the constants for the image noise due to

intensity uncertainty and transformation uncertainty, respectively. For each voxel, the energy function is described as follows:

$$E(\boldsymbol{v}) = \left(I_{f'} - I_g + \boldsymbol{v}\nabla I_g\right)^2 + \frac{\sigma_i^2}{\sigma_x^2} \cdot \boldsymbol{v^2}, \qquad (4.11)$$

where $I_{f'}$ and I_g are the intensities of the same voxel in the respective images f' and g, ∇ is the gradient operator and $\boldsymbol{v}$ is the estimated displacement (update velocity). The similarity measure used here is the squared pixel distance, with squared gradient of the transformation field being set as the smoothness regularization. It is given in [50] to replace $\sigma_i^2 = \left(I_{f'} - I_g\right)^2$ and $\sigma_x^2 = \frac{1}{\alpha^2}$, and we obtain the following:

$$E(\boldsymbol{v}) = \left(I_{f'} - I_g + \boldsymbol{v}\nabla I_g\right)^2 + \alpha^2\left(I_{f'} - I_g\right)^2 \cdot \boldsymbol{v^2}, \qquad (4.12)$$

where α is the normalization factor. It is noted that the demons registration algorithm is a local approximation, thus it is performed for a certain number of iterations until the solution converges; the displacement vector $\boldsymbol{v}$ is updated in each iteration so that any deformations in the image are accounted for progressively. In our experiments, the number of iterations is set to be 150 to ensure convergence, but if the energy function falls below a certain threshold in any iteration, the algorithm is terminated immediately since an optimized solution is found.

4.2.2 Velocity Estimation

The optical flow equation for finding small deformations is used as the basis for demons registration. The optical flow equation defines how the registration forces act on

the image to account for deformations. The estimated displacement $\boldsymbol{v}$ required for a point

in the reference image f' to match a point in the floating image g is given by:

$$\boldsymbol{v} = \frac{\left(I_{f'}-I_g\right)\nabla I_g}{|\nabla I_g|^2+\alpha^2\left(I_{f'}-I_g\right)^2} + \frac{\left(I_{f'}-I_g\right)\nabla I_{f'}}{|\nabla I_{f'}|^2+\alpha^2\left(I_{f'}-I_g\right)^2}, \qquad (4.13)$$

where α is the same normalization factor as shown in (4.12). The normalization factor is

set to be 3 in this implementation to account for small local deformations more

effectively. The first term in (4.13) is derived directly from (4.12) given that the error

gradient is zero where the error is at a minimum.

In [9], the velocity formula contained only the first term in (4.13) without the

constant α, which uses only edges in the reference image as the passive internal force.

The term $\left(I_{f'} - I_g\right)^2$ in the denominator serves to make the velocity equation more stable

in image registration. To regularize the deformation field, Gaussian smoothing is

performed on the velocity field since the estimated displacement obtained is local. The

second term in (4.13) is then introduced in [30] to improve registration convergence

efficiency and robustness.

4.3 Implementation and Validation

An adaptation of the method as described in [8] is used to estimate the pseudo

ground-truth dataset and a Matlab wrapper [51] is used to implement the demons

registration technique [9, 44] to account for the non-rigid deformations in the kidney. For

the demons registration algorithm, each translation-aligned image (floating image in g) is

registered to its corresponding estimated pseudo ground-truth image (reference image in

f'). However, calibration is performed prior to this step to remove the underlying deformation field for the reference frame. An image in f' is first registered to its corresponding reference image. The deformation field obtained is then applied to all images in f' to complete the calibration process.

In terms of results validation, many different tests were conducted to validate the proposed method. The tests investigate the registration algorithm's limits and robustness with different scales of simulated transformation applied on the dataset. A visual inspection will be performed for all the tests to verify the quality of the non-rigid registration algorithm. Graphs showing the average intensity over a few voxels of the kidney against time are also constructed, where voxels within the registered kidney should generally exhibit a smooth intensity variation across time. For quantitative analysis, the distance between kidney masks shall be evaluated mainly for the simulated datasets to identify the limits of the non-rigid registration algorithm. For the first test, the non-rigid registration's ability to overcome different translational errors in the dataset will be investigated. Second, random non-rigid transformations of different scales are applied on the image, and these deformations are recovered using the demons algorithm. The random non-rigid transformation is achieved by applying a 3D B-spline transformation defined on a regular grid with a moderate spacing between the nodes and a random displacement of each node, a sample of which is shown in Figure 4.3. The nodes of the deformation grid in Figure 4.3 are shifted randomly, but they are adjusted by the B-spline transformation, depending on the 3D parameters. The deformation in the Z-axis is restricted to 1 voxel due to a small number of slices involved. The final test includes creating a simulated dataset with both translation and non-rigid errors in the

kidney, and the registration algorithm as described in Chapters 3 and 4 will be employed
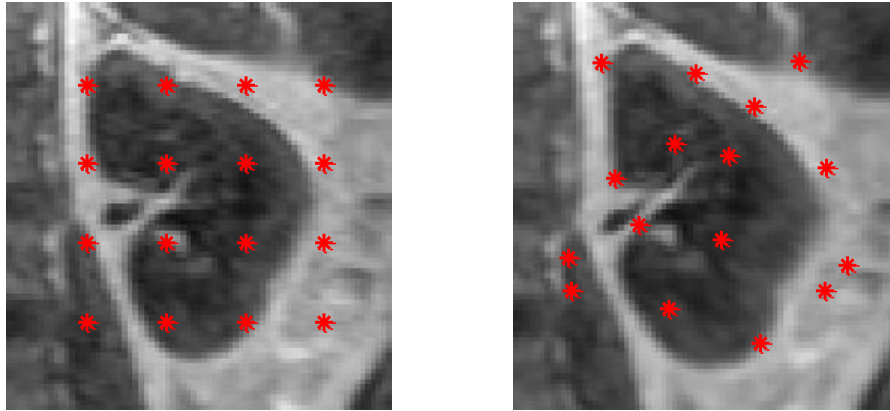
in full to register the simulated dataset.

Fig.4.3. (left) Original kidney image and (right) an example of a deformed kidney image. The nodes are marked by (*) in the images.

CHAPTER 5

RESULTS

The registration algorithm was tested on a dozen real patient datasets (24 3D+time kidney volumes). The dynamic MR images were obtained on a 1.5 T system (Avanto; Siemens, Erlangen, Germany) with a maximum slew rate of 200 T/m/s, maximum gradient strength of 45mT/m, and a torso phased-array coil. 3D T_1-weighted spoiled gradient-echo imaging was performed in the oblique coronal orientation to include the abdominal aorta and both kidneys. The following parameters were used: TR = 2.8 ms, TE = 1.1 ms, flip angle = 12^o, matrix = 256 $\times$ 256 $\times$ 20, FOV = 425 $\times$ 425$\times$ 100 mm^3, bandwidth = 650 Hz/voxel, volume acquisition time = 3 s. The 20 original 5-mm coronal partitions were interpolated to 40 2.5 mm slices so that the matrix becomes 256 $\times$ 256 $\times$ 40 (voxel resolution of 1.66mm $\times$1.66mm $\times$2.5mm).

Several unenhanced acquisitions were performed before a 4-ml bolus of Gd-DTPA (Magnevist; Berlex laboratories, Wyne, NJ, USA) was injected, followed by 20ml of saline, both at 2 ml/s. More volume acquisitions followed at regular intervals, and all the volumes (with numbers ranging from 31-41) were combined to form a dataset. The size of the kidneys ranged from 60-90 voxels and 60-80 voxels in the X and Y axes, respectively. The registration algorithms are validated by a qualitative analysis and a quantitative analysis on the main datasets and on several motion-simulated datasets. The remaining sections will discuss the efficiency and robustness of the algorithm in separate levels (rigid and non-rigid registration) and in a combined two-level registration. An Intel Core2 Duo CPU E8500 @ 3.16GHz is used to perform the entire registration algorithm on all datasets in the Matlab environment.

5.1. Analysis of Translational Registration

Various tests were conducted to validate the method qualitatively and quantitatively for translational registration as described in Chapters 3. The existing ground-truth of the dataset is used to compare the registration accuracy between different methods and the proposed method. Simulated datasets with random translational transformations in the 3 principle axes derived from pre-aligned datasets are also available to ensure the robustness of the algorithm. Four tests mentioned in sub-Section 3.2.6 are performed and discussed quantitatively with the mean and standard deviations of the error in all the 3 axes (X, Y and Z). Moreover, a cumulative distribution of the registration error displays the percentage of volumes with errors (defined by the Euclidean distance) under a certain threshold. To ensure a robust result for non-rigid registration, it is important to keep the absolute error distance for each volume within a small voxel error range ($<$ 5 voxels).

5.1.1. Qualitative Analysis of Translational Registration

Some kidney in-plane slices given in Fig. 5.1 illustrate the results of registering a simulated motion dataset with randomly applied translational motion. The first row shows the original aligned kidney images; the second row shows the respective randomly translated kidneys; and the third row displays the results of rigid registration. The registered kidney image slices generally follow the kidney boundary (contour line) of the reference frame, demonstrating good results for translational registration.
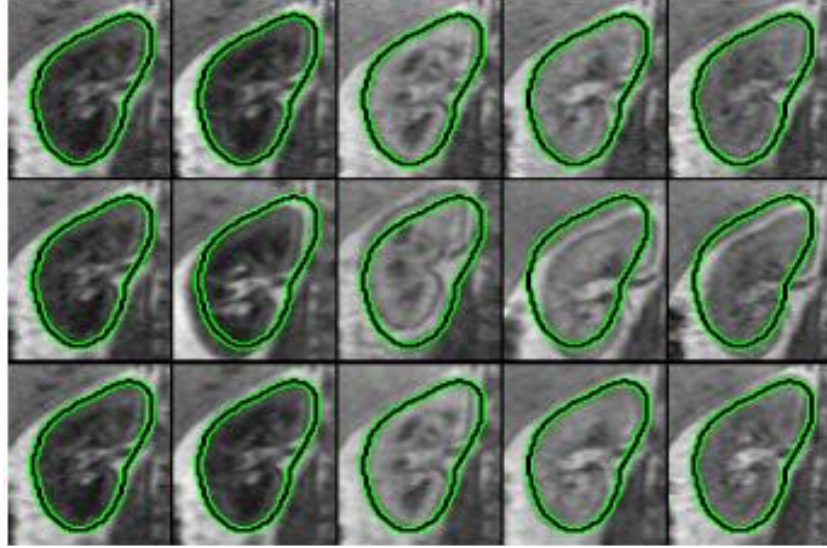
Fig.5.1. Results of rigid registration across different static time frames: (First row) some initially aligned kidney slices; (Second row) the same slices after simulated random transformations; (Third row) the registered slices. The contour line represents the boundary of the kidney in the reference slice to the left.

Moreover, three average intensity curves as shown in Fig. 5.2 displays how the average intensity of a few selected voxels of the cortex within the kidney varies with time for a dataset: 1) aligned with ground truth data; 2) registered using the proposed method; 3) pre-registration. The selected voxels of the cortex are close to the boundaries of the kidney and also the medulla, thus a slight displacement in the kidney will cause a large deviation in the intensity curve. The proposed registration method using graph-cuts has aligned the images satisfactorily as the post-registration intensity curve is close to the ground-truth aligned intensity curve. Although the intensity curve after translational registration is not smooth, it will be refined further during the non-rigid registration stage when local deformations are accounted for.
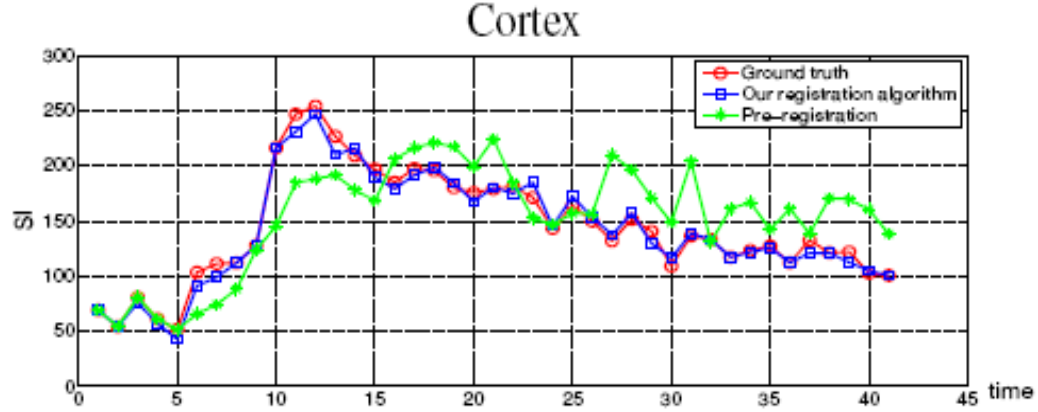
Fig.5.2. **Three average intensity curves for one of the datasets: 1) pre-registration; 2) alignment using the ground truth; 3) post-registration using our method.**

5.1.2. Quantitative analysis of Translational Registration

To quantify the accuracy of the translational registration algorithm, two types of errors are defined for each kidney volume: 1) the errors e_x, e_y and e_z representing the absolute error between the ground truth and the obtained translational value after registration in each of the axes X, Y and Z, respectively; 2) the error e denotes the Euclidean distance of the absolute errors (e_x, e_y and e_z) for each individual volume. The mean and standard deviation of the absolute errors e_x, e_y and e_z are tabulated and compared across different methods. The computational timing is also displayed for the relevant cases and analyzed whether it is feasible to use a more complex procedure to obtain a better result. The error e is used in the cumulative distribution graph where the error probability $P(e > \eta)$ signifies the percentage of volumes with registration error e larger than the threshold η. In the cumulative distribution graph of the error e, the Heaviside step function would be the most desired result.

For the first test, mutual information is compared against gradient difference in the rigid registration of the kidney images. Two forms of MI-based rigid registration methods are considered: 1) replacing gradient difference by MI in computing the data cost for the labels of each node; and 2) applying MI globally to compare the similarities between the reference and the floating image. The mean and standard deviation of error and the computational timing obtained for the MI-based rigid registration methods and the gradient-difference based registration method are shown in Table 5.1. For the MI-based registration applied globally, the error means and standard deviations obtained were much higher than the ones obtained by using gradient difference registration in the X axis. The high mean and standard deviation in the error indicate that there is a large mismatch between images for some volumes which is undesirable. In terms of computational timing, registration takes about 4 minutes to register a dataset using gradient difference, whereas 6 minutes is used to register the same dataset using MI globally. For the MI graph-cut registration, image alignment failed using MI in a graph-cut implementation with super-nodes, and the computational timing was comparatively much higher than the other two methods. The reason for the failed registration could be due to the lack of information within the super-nodes causing the graph to be unable to compute an optimal global solution. The cumulative distribution of the errors obtained for the gradient difference graph-cut registration and the global MI registration is displayed in Figure 5.3. The graphs show that the cumulative-error function for images registered by gradient difference is better than that of the images registered by mutual information. Thus, it can be concluded that gradient difference is a better choice of

similarity measure for registering DCE-MRI images due to a higher efficiency and a

more reliable result.

Table.5.1. Mean and standard deviation of error (in voxel) for rigid registration using different similarity measures: (1) Mutual information (2) Gradient Difference.

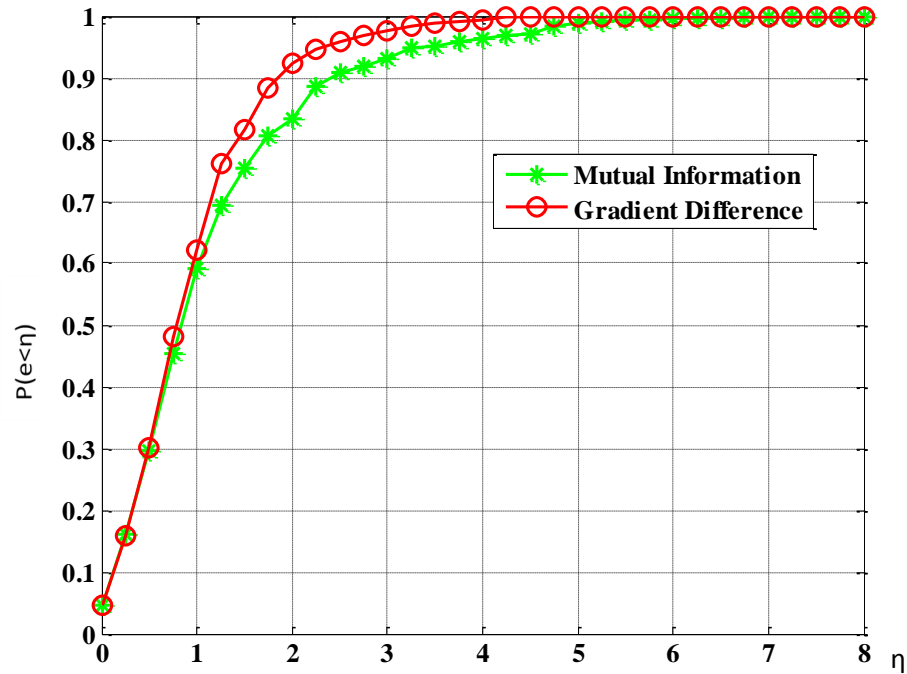| Similarity Measure	Mean Error			Standard Deviation of Error			Computational Time (per dataset)
	e_x	e_y	e_z	e_x	e_y	e_z	
Mutual Information (global)	0.8497	0.4394	0.5314	1.0542	0.4930	0.5049	~6 minutes
Gradient Difference (graph-cut)	0.5717	0.4593	0.5205	0.5409	0.5594	0.4895	~4 minutes
Mutual Information (graph-cut)	8.9881	5.2369	1.3819	6.2728	3.8905	0.9659	~10 minutes

Fig.5.3. Cumulative distribution function for the error *e* in registration for similarity measures MI (*) and gradient difference (o).

For the second test, the registration accuracy and computational timing using different node sizes in a graph are compared. Different node sizes between 11×11 and 25×25 are tested and analyzed, with the spacing between nodes set as half the node size in both the X and Y axes. The mean and standard deviation of errors obtained for all the different registration algorithms using graph-cuts of different node sizes are given in Table 5.2.

Table.5.2. Mean and standard deviation of error (in voxel) for rigid registration using different node sizes.

Node Size $(X \times Y)$	Mean Error			Standard Deviation of Error			Computational Time in Seconds (per volume)
	e_x	e_y	e_z	e_x	e_y	e_z	
25×25	0.5717	0.4593	0.5205	0.5409	0.5594	0.4895	5.4
21×21	0.5761	0.4693	0.5221	0.5561	0.5677	0.4959	4.3
17×17	0.6441	0.5136	0.5241	0.7501	0.7390	0.5038	4.0
13×13	0.7459	0.5830	0.5219	0.9763	0.8579	0.4912	3.7
11×11	0.9236	0.6431	0.5380	1.3635	0.9996	0.5253	3.5

The mean and standard deviation of the error becomes smaller with a larger node size, as seen from the results in Table 5.2. When the node size is large at 21×21 or 25×25, the results of translational registration are similar. The lower mean and standard deviation of the error comes at a higher computational cost, affecting the efficiency of the registration algorithm. Figure 5.4 shows the cumulative distribution function of the error distances given in all the different implementations using single sized nodes. It is

observed that the results improved with a large node-size, with a small difference between the results obtained with nodes sizes of 21 $\times$ 21 and 25 $\times$ 25. Since the most important criterion for a good translational registration is for the maximum value of the error distance, e, to be less than 5 voxels, the node size is chosen to be sufficiently large so that the maximum error distance is small. It can be seen in Figure 5.4(B) that only two cases (of node sizes of 21 $\times$ 21 and 25 $\times$ 25) are able to obtain a maximum error distance of less than 5 voxels. A node size of 25 $\times$ 25 will be used in the subsequent tests.

For the third test, the registration accuracy is compared for two implementations using the same node-structure (regular spaced grid), the same node-size (25 $\times$ 25) and the same similarity measure (gradient difference); only the weights will be varied in this test. For the first method, the weights allocated to each node are determined by the ratio of the amount of edge information present within the node to the amount of edge information throughout the whole kidney. Weights are excluded in the second registration method. The mean and standard deviation of errors obtained for the both methods are tabulated in Table 5.3. It is noted that the computational complexity is approximately the same for both methods and will not be shown in the table.
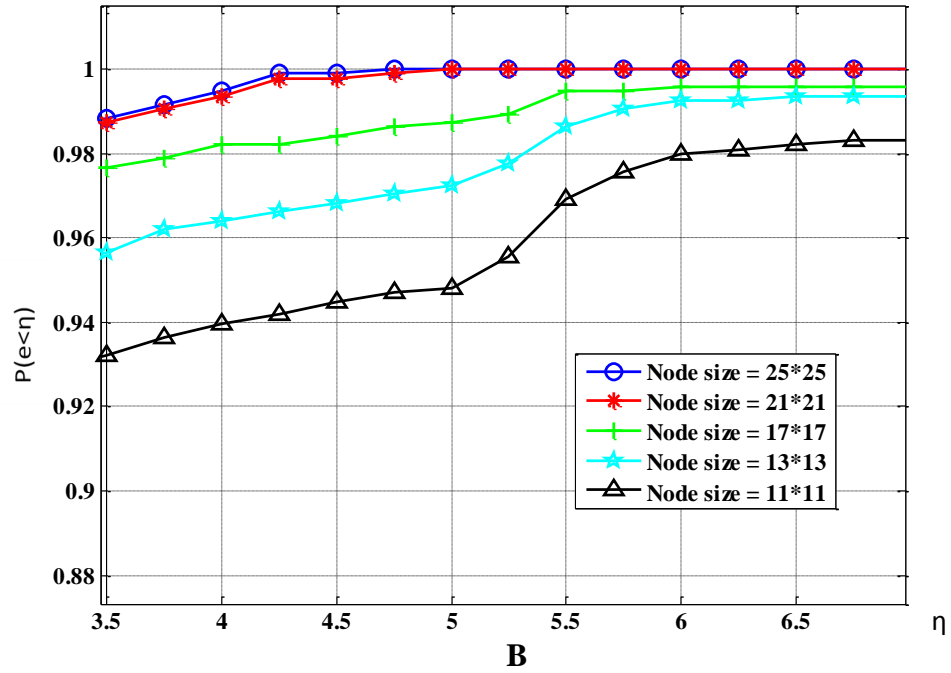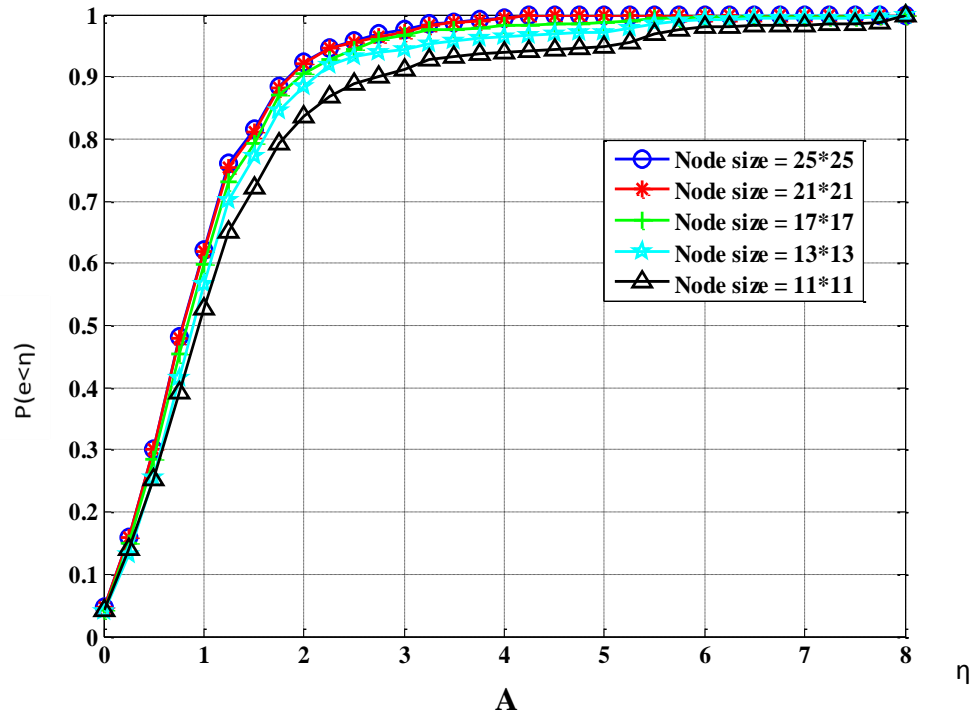
Fig.5.4. (A) Cumulative distribution curves for the registration error using different node sizes (B) Zoomed in result of the graph in A.

Table.5.3. Mean and standard deviation of error (in voxel) for rigid registration: (1) using weights and (2) without using weights.

Registration Method	Mean Error			Standard Deviation of Error		
	e_x	e_y	e_z	e_x	e_y	e_z
w/ weights	0.5717	0.4593	0.5205	0.5409	0.5594	0.4895
w/o weights	0.5964	0.4639	0.5251	0.5769	0.5560	0.4755

The mean and standard deviation of error for rigid registration are similar for both methods with and without using weights; although the mean error for the registration algorithm without weights is slightly higher in each of the 3 principle axes. As previously mentioned, weights allow the registration algorithm to place more emphasis on active nodes containing more kidney edge information and less emphasis on passive nodes containing less edge information, thus obtaining a more robust solution. Figure 5.5 shows the cumulative distribution function of the errors given in both the methods, and it can be observed that the first method using weights is better by only a small margin. Table 5.4 gives the maximum error distance for each of the patient datasets (24 kidney datasets). From Table 5.4, it is observed that the difference between the two methods is not significant for almost all of the patients except for patient 10, where the maximum error varied beyond the criteria of 5 voxels. The maximum error distance is 6.1985 and 4.6205 for the registration algorithm without using weights and for the registration algorithm using weights, respectively; the registration algorithm using weights fits the criterion of having a maximum error distance of less than 5 voxels. Table 5.4, however, does not

show that the overall registration results improved with the introduction of weights as some datasets registered better whereas some did not. Taking into account that the maximum and mean error distance reduced, weights are included in the formulation of a translational registration algorithm.

In the fourth and last test, the translational registration algorithm is tested on two simulated datasets that are created by performing random translations on two pre-aligned datasets. The mean and standard deviation of the error (in voxel) are negligible, with only a 1 voxel error in the X axis for one of the volumes in the simulated datasets. This concludes that the translational registration algorithm is robust in the presence of translational errors with a single reference frame.
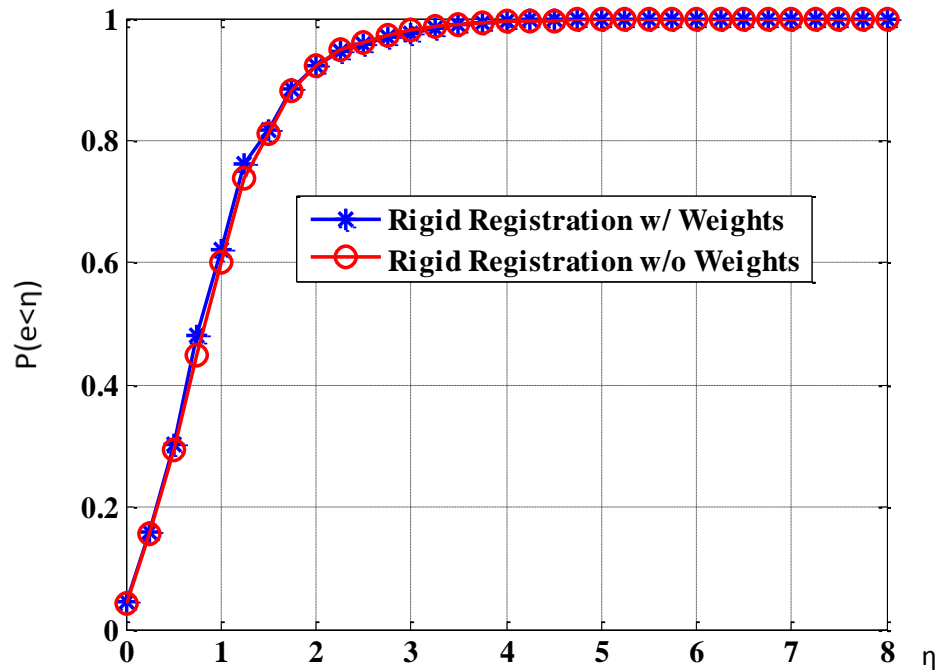

Fig.5.5. Graph on cumulative distribution function for the registration error in graph-cut: (1) incorporating weights based on amount of edge information (*); (2) without incorporating weights (o)

Table.5.4. Maximum error distance (in voxel) for each kidney dataset. The patient with the worst results is highlighted.

Patient Number	Left Kidney		Right Kidney	
	w/ weights	w/o weights	w/ weights	w/o weights
1	2.7216	2.0980	0.9865	1.3701
2	1.0714	1.0714	1.3147	1.8442
3	1.4660	1.6427	2.8522	2.8522
4	1.0875	1.0875	2.1697	2.1697
5	2.0645	2.0497	2.2513	1.9682
6	1.2015	1.3210	2.0243	2.0243
7	1.2745	1.2745	1.9010	1.7275
8	2.0689	2.0689	1.7222	1.7593
9	1.3519	1.3519	1.6430	1.6430
10	4.5656	6.1985	3.7474	4.6205
11	1.5641	1.5641	1.5207	1.5207
12	3.2418	2.6824	2.7913	2.4736

5.2. Analysis of Estimated Pseudo Ground-Truth Dataset

A visual inspection of the resulting estimated dataset and the intensity-time curves is done to interpret the results of pseudo ground-truth dataset estimation. Figure 5.6 shows some instances of the pseudo ground-truth being estimated from two datasets. The pseudo ground-truth estimation is able to eliminate noise and minor motion from the original dataset by imposing spatiotemporal smoothness constraints on each voxel in a given time-series volume. The main edges around the boundary of the kidney and the medulla are preserved even though smoothing has occurred in the pseudo ground-truth

dataset. Figure 5.7 shows the results of curve fitting in three graphs by using: (1) average intensity over the kidney mask; (2) intensity of a chosen voxel of the cortex; and (3) intensity of a chosen voxel of the medulla. It is observed that the proposed method as described in Chapter 4.1 is able to produce a good estimation of the pseudo ground-truth, even in the presence of noise and slight motion, as long as errors are not consistent.

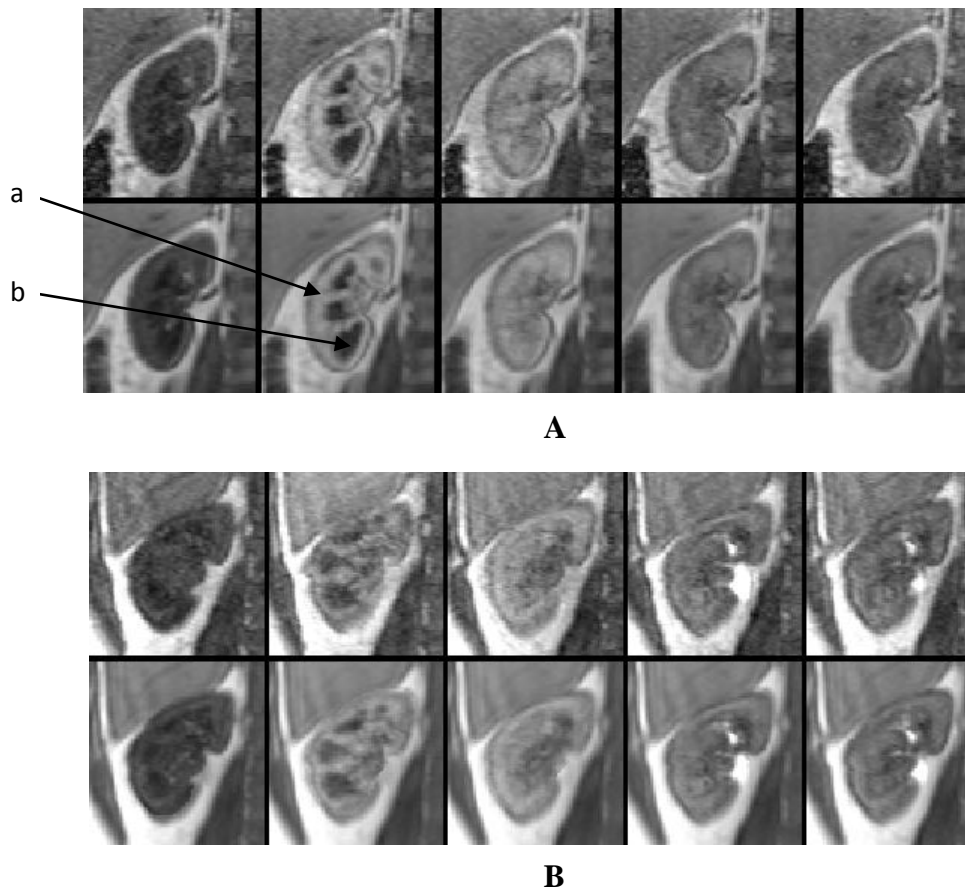**A**

B

Fig.5.6. A and B gives different datasets showing: (top row) sample kidney slices from the resulting volume after translational registration; (bottom row) estimated pseudo ground-truth of the same kidney slices. Point marked 'a' is a voxel of the cortex and point marked 'b' is a voxel of the medulla within the kidney, the intensity-time curve of which are shown in Figure 5.7.
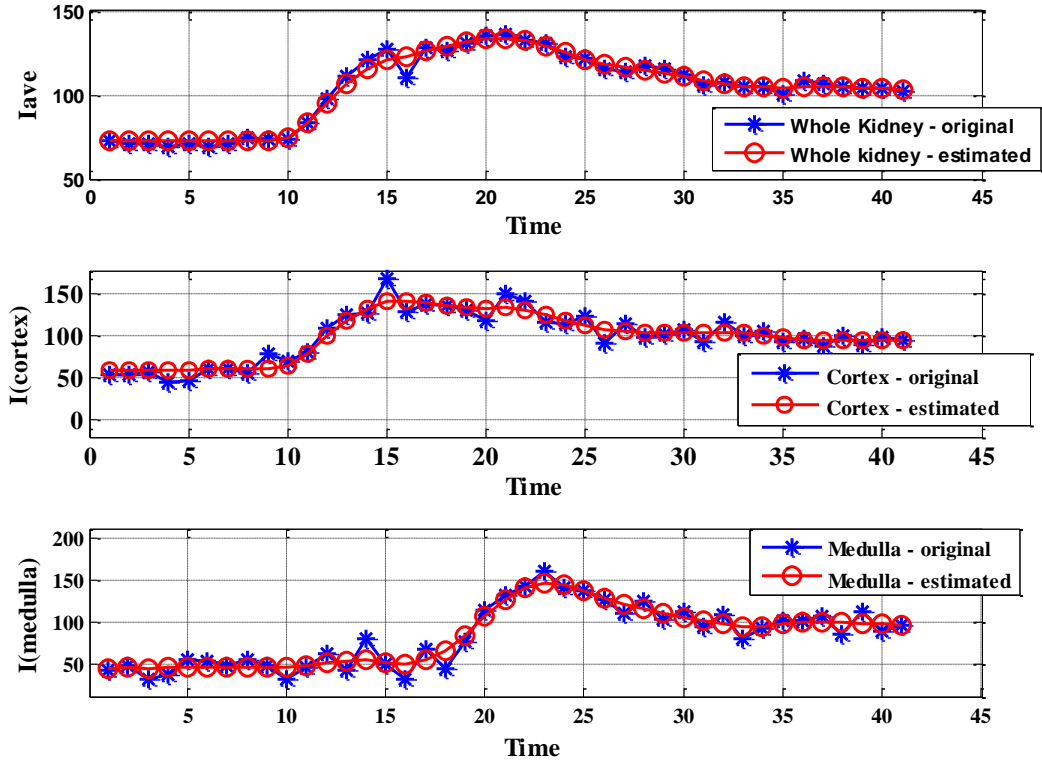
Fig.5.7. Graphs show the intensity-time curves between the translationally registered dataset (*) and the estimated pseudo ground-truth dataset (o) for: (top) the average intensities of the kidney across the mask; (middle) the intensity of a cortex voxel marked 'a' in Fig 5.6; (bottom) the intensity of a medulla voxel marked 'b' in Fig 5.6.

5.3. Analysis of Non-Rigid Registration

In this section, a visual assessment/qualitative analysis on the general results of non-rigid registration will first be given. Then, to quantify the results of non-rigid registration, the distance between kidney masks is computed for the simulated datasets to verify the limits of registration. As mentioned previously, another set of tests was conducted to validate the non-rigid registration algorithm as described in Chapter 4. The first test will establish the tolerance of the demons algorithm to register volumes with translational errors of up to 5 voxels. The second test will then determine the tolerance of

the demons algorithm to register volumes with pre-applied free-form deformations over three different levels with the maximum error of a node in the grid to be 6 voxels. Using the demons algorithm, non-rigid registration takes roughly 1 minute to register a 3D volume of 7 slices non-rigidly (~ 40minutes a dataset), compared to more than an hour per 3D volume using FFD. Thus, the demons algorithm is much more efficient.

5.3.1. Qualitative analysis of Non-Rigid Registration

Non-rigid registration transforms a floating image (the translationally registered dataset) with reference to a reference image (the estimated pseudo ground-truth dataset) to account for deformations that exist in an image. A sample dataset in Figure 5.8 gives the general results where small local deformations are accounted for in non-rigid registration by the demons algorithm. It is observed that the non-rigid registration results are good; small local deformations are accounted for in all slices. Apart from having kidney boundaries that match between the registered image and the pseudo ground-truth dataset, the original texture of the translational registered dataset is also preserved with minor smoothing effect.

Not to neglect are some examples where the deformation of the kidney is larger. In the next sample dataset as shown in Figure 5.9, it is observed in the translationally registered dataset that the kidneys shown did not fit into the kidney mask (represented by a contour line) exactly. The demons algorithm is also able to account for such errors and correctly register non-rigidly these images.

Fig.5.8. Sample dataset showing different 2D slices (different rows) of a kidney after: (left) translational registration; (middle) pseudo-ground truth estimation; (right) non-rigid registration. Points A, B, C and D shows some regions where there were local deformations, which are resolved after non-rigid registration as shown in points A*, B*, C* and D*.
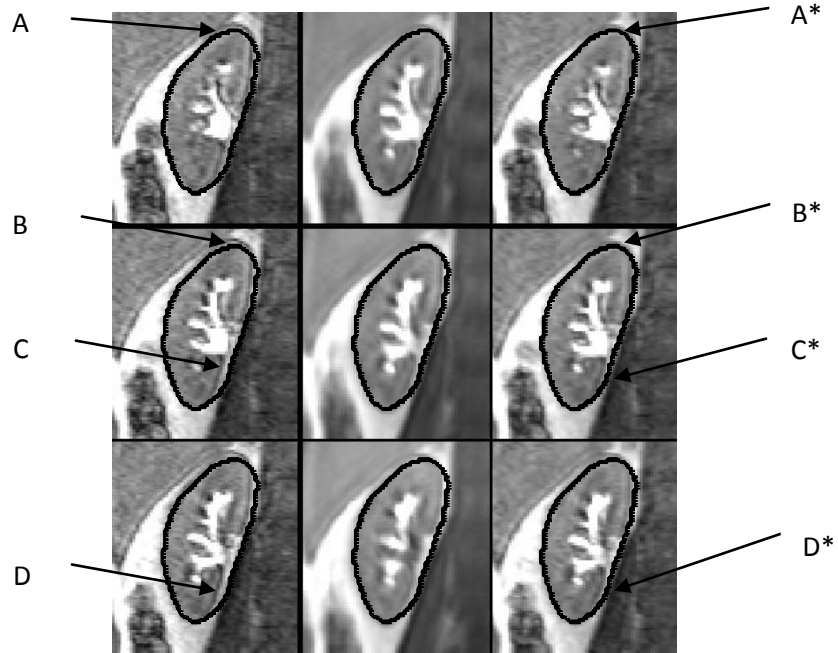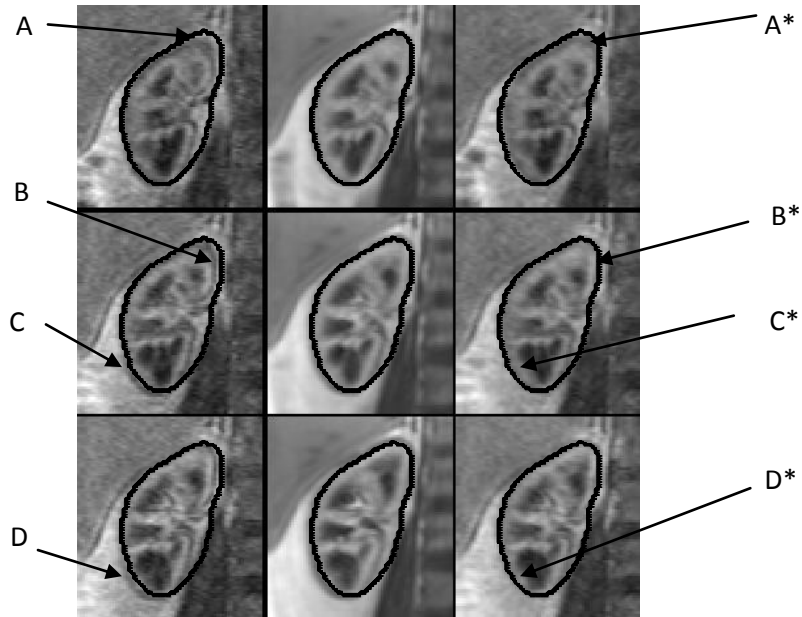

Fig.5.9. Sample dataset showing different 2D slices (different rows) of a more deformed kidney after: (left) translational registration; (middle) pseudo-ground truth estimation; (right) non-rigid registration. Points A, B, C and D shows some regions where there were larger local deformations, which are resolved after non-rigid registration as shown in points A*, B*, C* and D*.

61

Figure 5.10 shows the same sample dataset as in Figure 5.9, but with an additional 4th column showing the difference of images before and after non-rigid registration. It can be observed in Figure 5.10 that the edges of the kidney have shifted after non-rigid registration, showing how non-rigid deformation in the kidney is accounted for. Some results for the simulated dataset are provided in a later sub-section to demonstrate that non-rigid registration is indeed accurate and robust. Furthermore, note that some non-rigid transformation within the kidney exists, marked by white patches within the kidney in the 4th column of images. This transformation is caused by: (1) aligning the medulla within the kidney and (2) the smoothening of edges within the kidney during non-rigid registration.
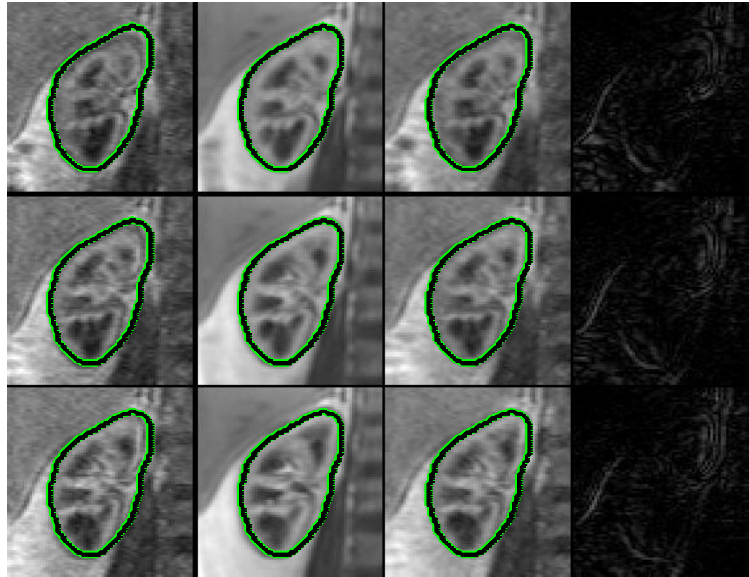
Fig.5.10. Same sample dataset as shown in Fig 5.9 showing different 2D slices (different rows) of a more deformed kidney after: (1st column) translational registration; (2nd column) pseudo-ground truth estimation; (3rd column) non-rigid registration; (4th column) Difference between images in 1st column and 3rd column.

In the first test, different scales of translational transformations (1 to 5 voxels) on a single axis are applied to certain time-volumes, and the demons registration algorithm is used to recover the pre-applied transformations. As the maximum error in the rigid

registration case is less than 5 voxels, the scales being considered are reasonable. Figure 5.11 shows the results of performing demons registration in three graphs: (1) average intensity over the kidney mask; (2) intensity of a chosen voxel of the cortex; and (3) intensity of a chosen voxel of the medulla. We observe that the intensity curve post-registration is generally smoother and closer to the pseudo ground-truth intensity curve than the intensity curve pre-registration, signifying that non-rigid registration is successful.
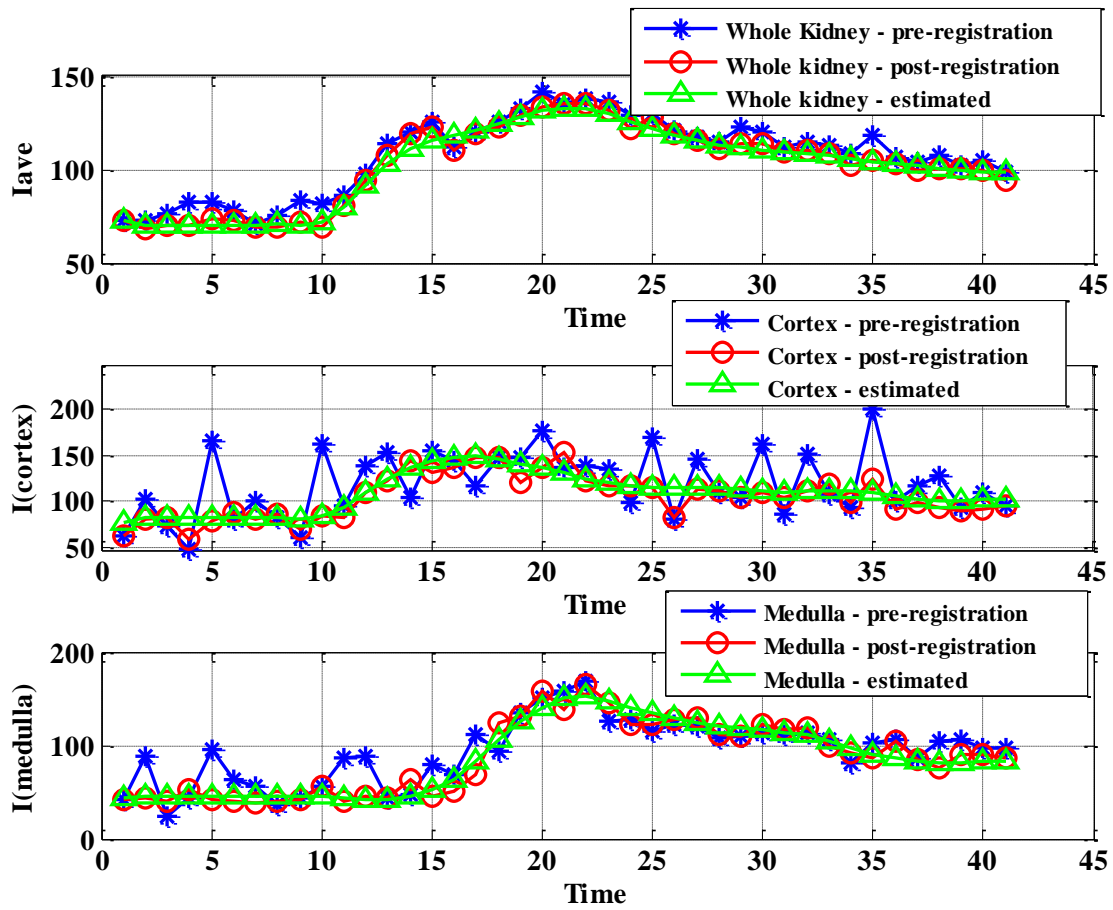

Fig.5.11. Graphs show the intensity-time curves between the pre-registered dataset (*), the post-registered dataset (o) the and the estimated pseudo ground-truth dataset (Δ) for: (top) the average intensities of the kidney across the mask; (middle) the intensity of a cortex voxel marked 'a' in Fig 5.13; (bottom) the intensity of a medulla voxel marked 'b' in Fig 5.13.

Figures 5.12 and 5.13 show the results of non-rigid registration applied to two different datasets with different translational errors applied to each and every time frame except for the reference frame. In both figures, the 5^{th} column shows the absolute difference of kidney images between the pre-aligned kidney and the deformed kidney and the 6^{th} column shows the absolute difference of kidney images between the pre-aligned kidney and the post-registered kidney. It is observed in Figure 5.12 that the differences due to simulated translational error are accounted for by non-rigid registration as the absolute difference between images was eliminated in most cases. In Figure 5.13, due to a lower contrast between the kidney and the background tissues, the registration result did not account for all of the translational errors. For most of the volumes, the reasonable results obtained reinforce the criterion of having a maximum residual translational error after rigid registration to be less than 5 voxels.
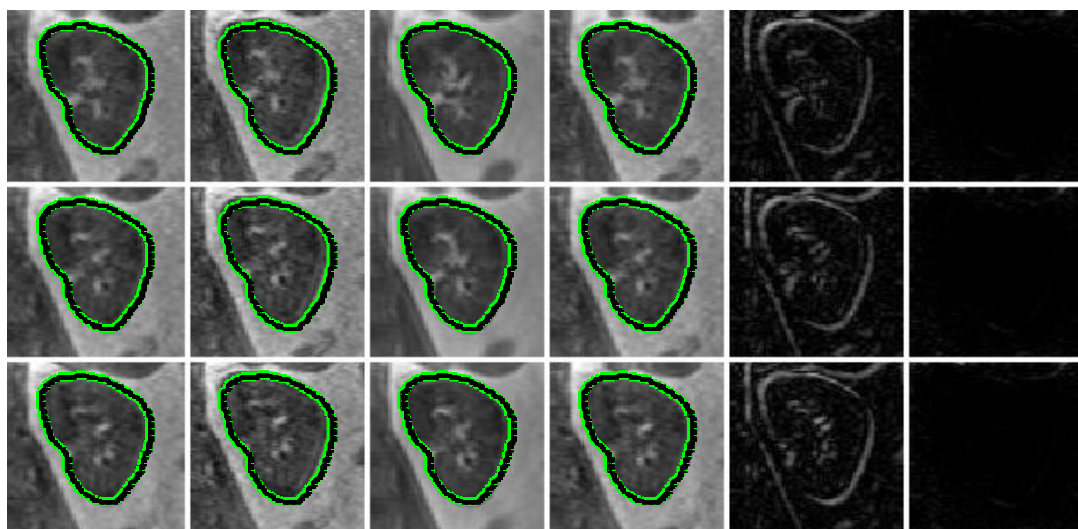
A
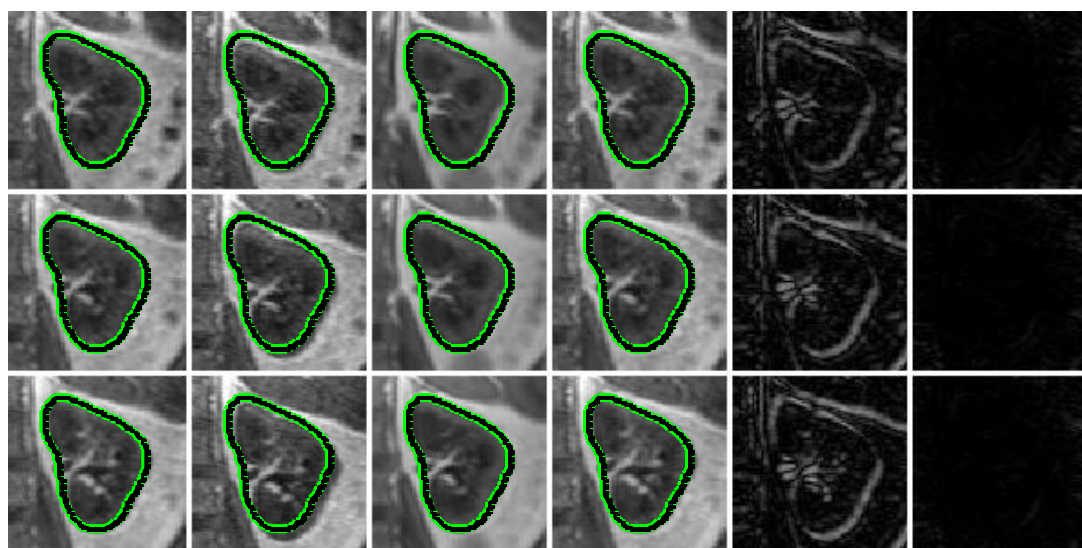
B

Fig.5.12. 2 different datasets A and B displaying different 2D slices (different rows) of: (1ˢᵗ column) an aligned kidney; (2ⁿᵈ column) kidney after simulated translation; (3ʳᵈ column) estimated pseudo-ground truth kidney; (4ᵗʰ column) registered kidney; (5ᵗʰ column) absolute difference of kidney between columns 1 and 2; (6ᵗʰ column) difference of kidney between columns 1 and 4.
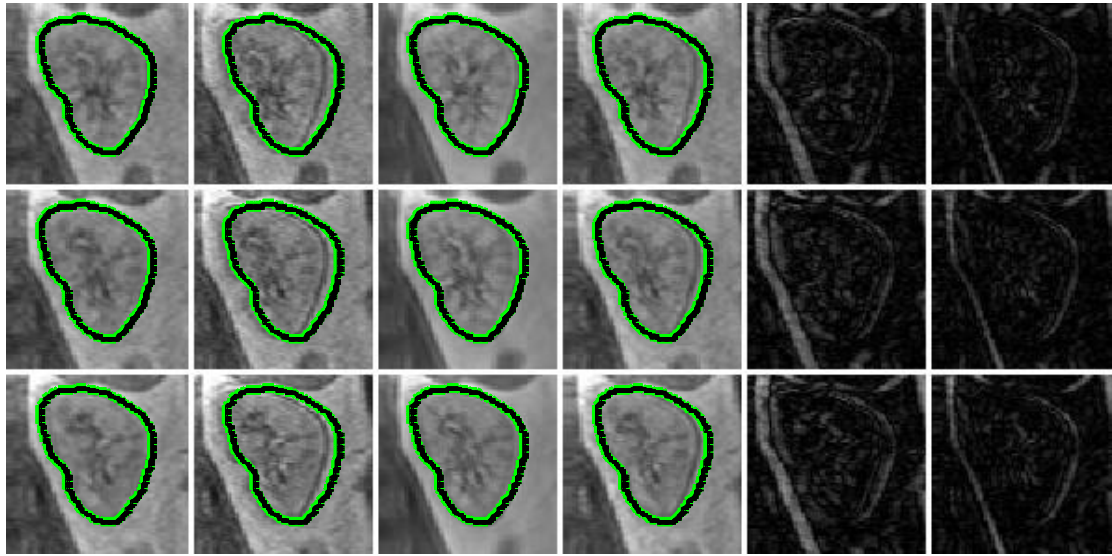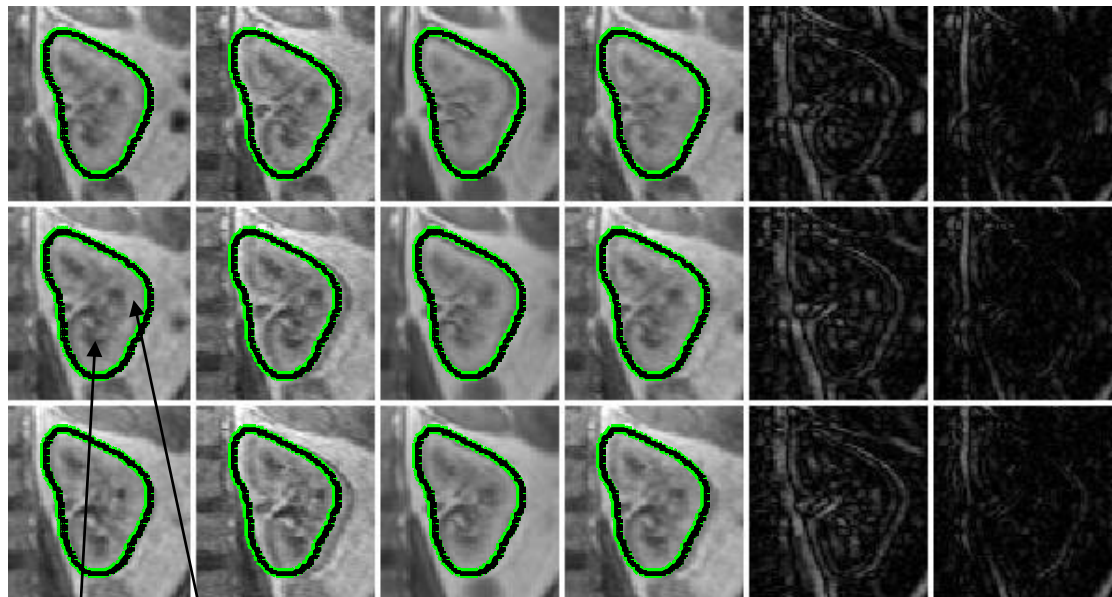
Fig.5.13. Two datasets, A and B, at another time frame showing different 2D slices (different rows) in low contrast of: (1st column) an aligned kidney; (2nd column) kidney after large simulated translation; (3rd column) estimated pseudo-ground truth kidney; (4th column) registered kidney; (5th column) absolute difference of kidney between columns 1 and 2; (6th column) difference of kidney between columns 1 and 4. Point marked 'a' is a voxel of the cortex and point marked 'b' is a voxel of the medulla within the kidney, the intensity-time curve of which are shown in Figure 5.11.

For the second test, random freeform-deformation fields are applied to the kidney images and registration is performed to recover these pre-applied non-rigid transformations. We obtained Figure 5.14 which shows the results of performing demons registration in three graphs: (1) average intensity over the kidney mask; (2) intensity of a chosen voxel of the cortex; and (3) intensity of a chosen voxel of the medulla. Similar to the first test, the intensity curve post-registration is smoother and closer to the pseudo ground-truth intensity curve than the intensity curve pre-registration, signifying that non-rigid registration is also successful in this scenario.

In Figures 5.15 and 5.16 we use different random transformations to different datasets, and we compare the results of registration visually by comparing: (1) the absolute difference of kidney between the pre-aligned kidney and the deformed kidney in the 5^{th} column and; (2) the absolute difference of kidney between the pre-aligned kidney and the post-registered kidney in the 6^{th} column. It can be observed in Figure 5.15 that the differences due to simulated deformation of the kidney are accounted for by non-rigid registration as the absolute difference has reduced. Some minor differences still remain partly due to the effect of smoothening by a non-rigid transformation. In Figure 5.16(A), a larger deformation field is applied, and the results show that the demons algorithm is still able to register the images to a reasonable degree, even though interpolation caused more smoothing to occur within the image volume. In Figure 5.16(B), the absolute error appears to have not been reduced, as the applied deformation is too large, and contrast between the kidney and the background is low due to the kidney exhibiting higher intensities. The non-rigid registration thus could not accurately register the kidney. This shows that the demons registration algorithm is sensitive to the contrast between the

kidney and the background tissues; when the contrast is small between the kidney and the

background tissues, only a small amount of deformation could be recovered.
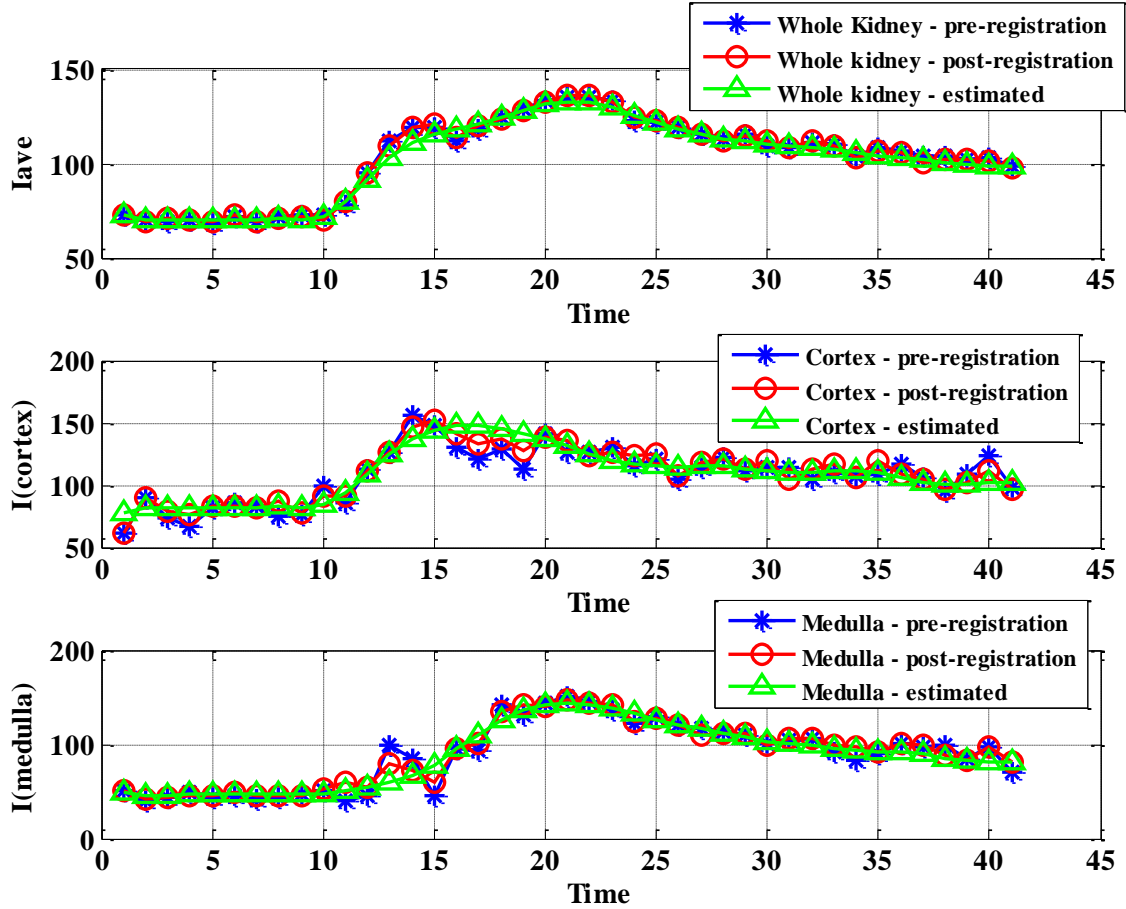
Fig.5.14. Graphs show the intensity-time curves between the pre-registered dataset (*), the post-registered dataset (o) and the estimated pseudo ground-truth dataset (Δ) for: (top) the average intensities of the kidney across the mask; (middle) the intensity of a cortex voxel marked 'a' in Fig 5.16; (bottom) the intensity of a medulla voxel marked 'b' in Fig 5.16.
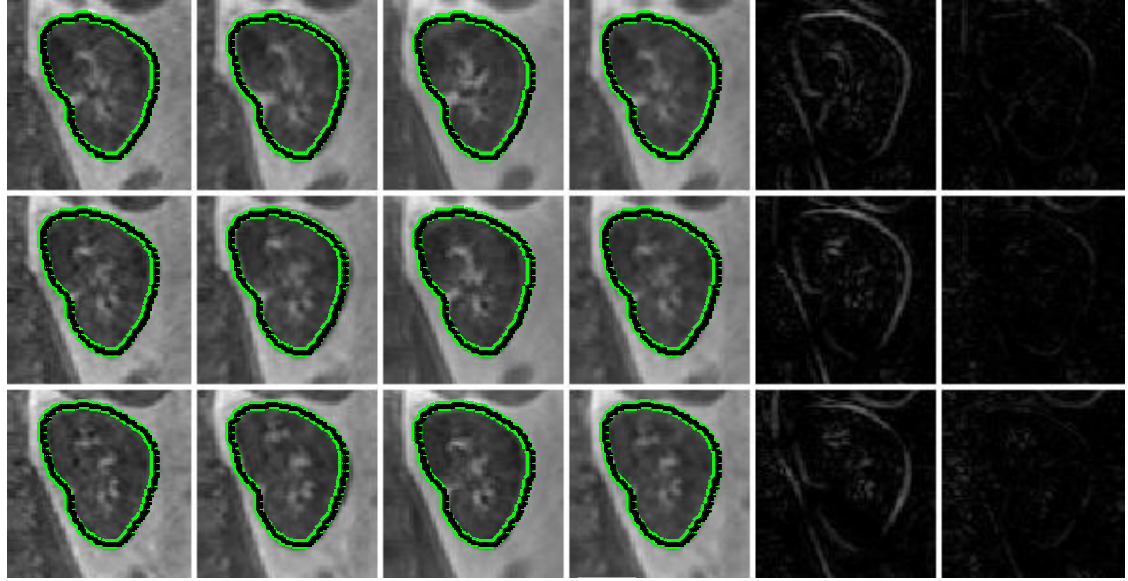
A
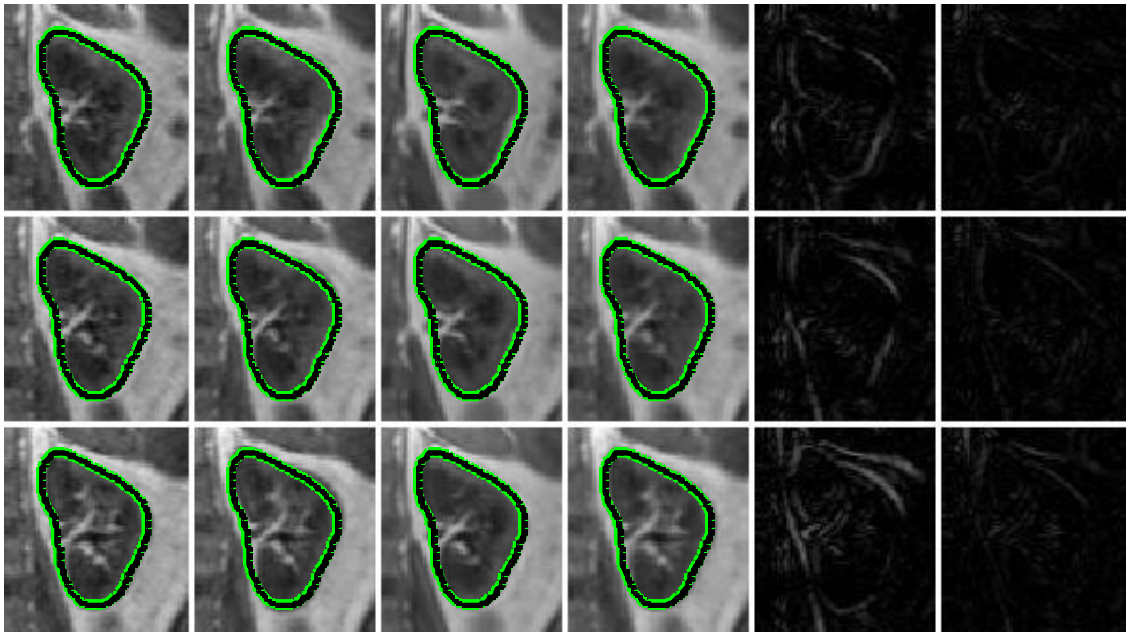
B

Fig.5.15. Sample datasets A and B showing different 2D slices (different rows) of:
(1st column) an aligned kidney; (2nd column) kidney after a different simulated deformation; (3rd column) estimated pseudo-ground truth kidney; (4th column) registered kidney; (5th column) absolute difference of kidney between columns 1 and 2; (6th column) difference of kidney between columns 1 and 4.
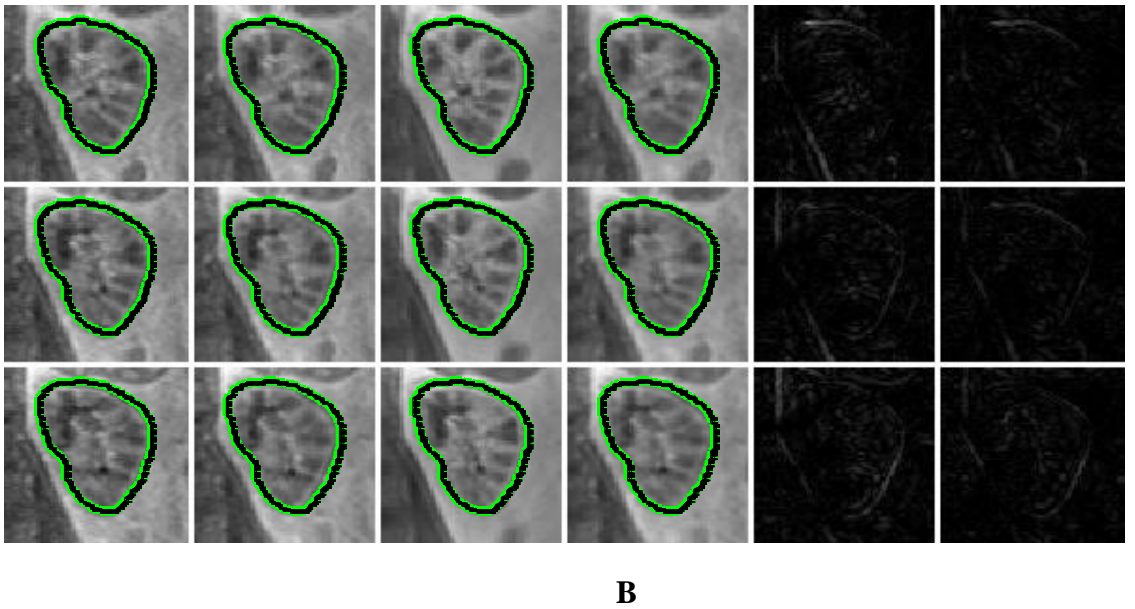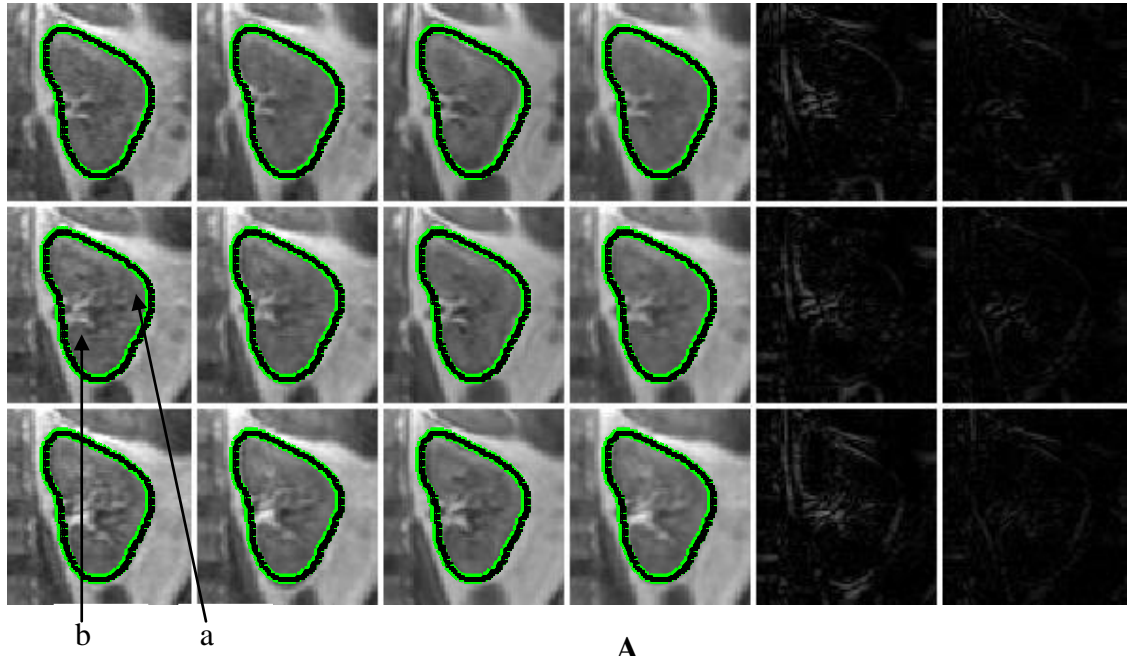
Fig.5.16. Two datasets, A and B, at another time frame showing different 2D slices (different rows) in low contrast of: (1st column) an aligned kidney; (2nd column) kidney after another simulated deformation; (3rd column) estimated pseudo-ground truth kidney; (4th column) registered kidney; (5th column) absolute difference of kidney between columns 1 and 2; (6th column) difference of kidney between columns 1 and 4. Point marked 'a' is a voxel of the cortex and point marked 'b' is a voxel of the medulla within the kidney, the intensity-time curve of which are shown in Figure 5.14.

5.3.2. Quantitative analysis of Non-Rigid Registration

We attempt to perform a quantitative analysis on the results of non-rigid registration by means of comparing the kidney mask segment before and after registration for the simulated datasets where the dataset is pre-aligned and the simulated transformation applied is known. In this way, we are able to determine the strength of the registration algorithm during the different stages of contrast within the kidney.

Three images are present: (1) the pre-aligned image; (2) the transformed image (simulated image) and (3) the registered image. First, the kidney mask is obtained for each of the three images. To quantify the registration results for the translationally simulated dataset, we compare the Euclidean distance for each volume within the dataset between: (1) the centroid of the mask before registration to the mask obtained from the pseudo ground-truth; and (2) the centroid of the mask after registration to the mask obtained from the pseudo ground-truth. For the deformed simulated dataset, two distance measures are obtained by computing the average absolute distance between each pixel in the mask boundary of the pre-aligned image to: (1) the mask of the transformed image and (2) the mask of the registered image. For each test, the distance measures for all the simulated datasets are tabulated, and the variation of the two distance measures across different time frames in a particular dataset is displayed in the form of a graph.

Table.5.5. Translation simulation: mean and standard deviation of the distance (in voxels) between the kidney masks (1) pre-registration and (2) post-registration.

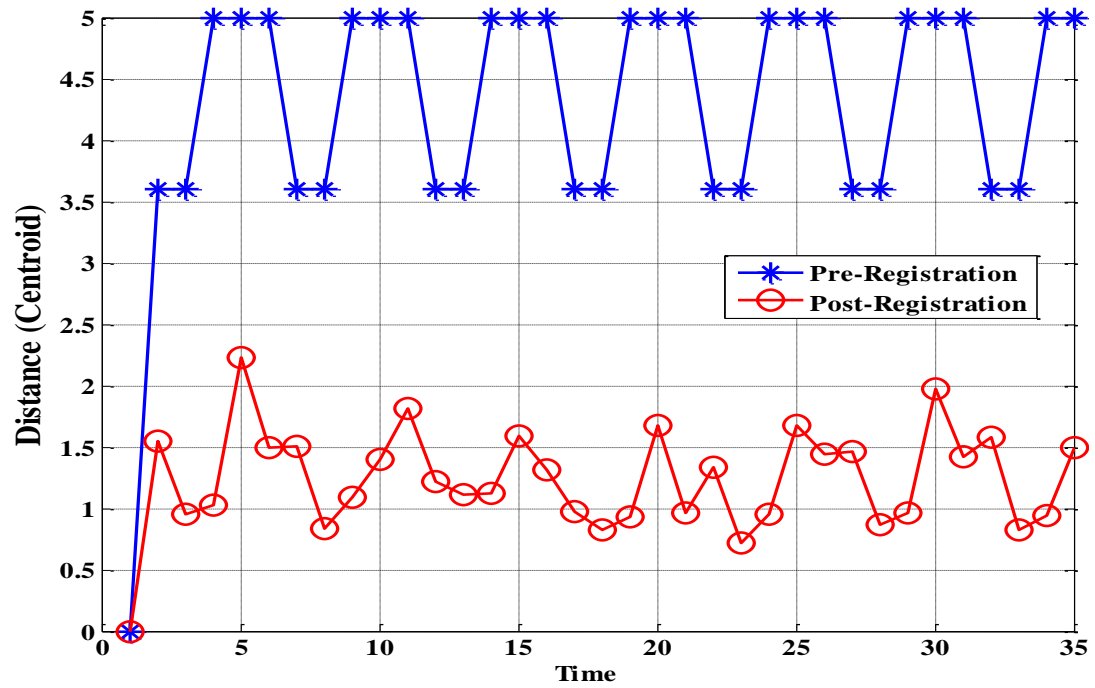| Dataset	Translation level	Mean		Standard Deviation	
		Pre-Registration	Post-Registration	Pre-Registration	Post-Registration
1	Low	4.5119	1.2312	0.6736	0.3497
	Medium	2.7326	1.0963	0.3690	0.2661
	High	1	0.4955	0	0.0661
2	Low	4.4258	1.2724	0.6966	0.3685
	Medium	2.6854	1.1543	0.3816	0.1907
	High	1	0.5108	0	0.0860

Fig.5.17. Translation simulation: curves showing the centroid distance between the kidney masks to the pseudo ground-truth kidney mask over time: (1) pre-registration and (2) post-registration.

For the first test involving the datasets with simulated translation, the results are shown in Table 5.5 and Figure 5.17. Table 5.5 gives the mean and standard deviation of the distance between each centroid of the kidney mask to the centroid of the pseudo ground-truth kidney mask before and after registration for all simulated datasets. The low, medium and high translation levels represent simulated translation distances of 1, 3 and 5 voxels for each volume respectively. The variation of the same distance measure across all time frames of a particular dataset is shown in Figure 5.17.

From Table 5.5, it is observed that the centroid distance between the kidney masks increases with a larger simulated translation. The centroid distance between the kidney masks also improved post-registration for all cases where only simulated translation is involved. The registration results for the low level of simulated translation are good, even though the results obtained for a medium and high levels of simulated translation are still considered reasonable.

When simulated deformations are considered in the second test, the results obtained vary from the first and are shown in Table 5.6 and Figure 5.18. Similarly, Table 5.6 gives the mean and standard deviation of the average distances between the kidney mask boundaries before and after registration for all simulated datasets, and Figure 5.18 shows the variation of the distance measure across all time frames of a simulated dataset. The low, medium and high deformation levels define deformations of up to 2, 4 and 6 voxels of each node in the grid, respectively.

Table.5.6. Deformed simulation: mean and standard deviation of the distance (in voxel) between the kidney masks (1) pre-registration and (2) post-registration.

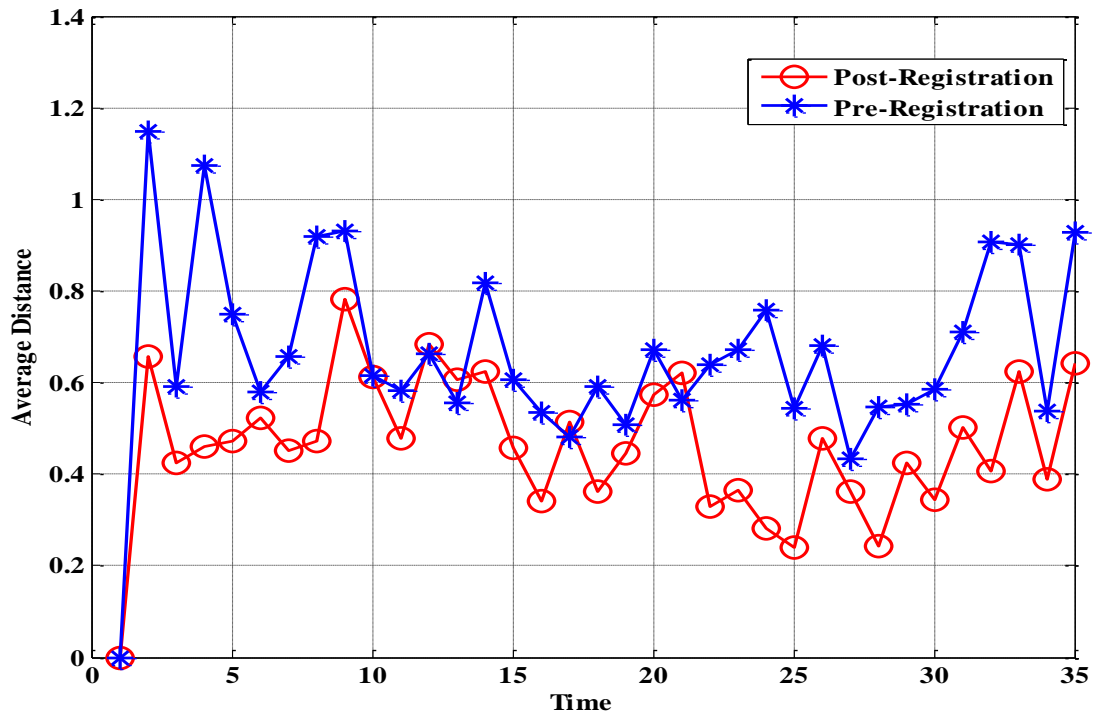| Dataset	Deformation level	Mean		Standard Deviation	
		Pre-Registration	Post-Registration	Pre-Registration	Post-Registration
1	Low	0.4191	0.3842	0.1055	0.1101
	Medium	0.5326	0.4192	0.1010	0.1064
	High	0.6601	0.4845	0.1420	0.1418
2	Low	0.3906	0.3591	0.0953	0.1314
	Medium	0.5448	0.4073	0.1263	0.1229
	High	0.6833	0.4759	0.1739	0.1318

Fig.5.18. Deformed simulation: curves showing the average distance between the kidney masks over time: (1) pre-registration and (2) post-registration.

From Table 5.6, it is observed that the mean distance between the kidney masks increases with a larger simulated deformation. The mean distance between the kidney masks also improved post-registration for all cases where only simulated deformations is involved. It is noted, however, that moderate and large simulated deformations become increasingly difficult to recover using the non-rigid registration algorithm. Some cases as shown in Figure 5.18 have a slightly larger mask distance post-registration, but this is not entirely attributed to misalignment, as the blurred boundaries due to transforming an image affect the grow-cut segmentation accuracy. The registration algorithm works better when a small simulated deformation is present; small simulated deformation fields can be recovered more easily using the proposed registration algorithm.

5.4. Registration Results on Simulated Datasets

For the final test, we have applied both a translation and a free-form deformation on all the kidney volumes except for a reference volume to determine the robustness of the multi-level algorithm as a whole. The values used are all randomized for each and every volume. The translational values considered in the X, Y and Z axis are as follows: $\{-5 \leq \Delta x \leq 5, -5 \leq \Delta y \leq 5, -1 \leq \Delta z \leq 1, \Delta x, \Delta y, \Delta z \in \mathbb{Z}\}$, and the non-rigid transformation grid is set with nodes being shifted by random values with 3 different levels of deformation (up to 2, 4 and 6 voxels). The purpose of this simulated dataset test is to ensure that the proposed multi-level registration algorithm is able to properly register the kidney volumes given a reasonable translation error and different levels of deformation of the kidney. The results will be evaluated separately for the rigid

registration and the non-rigid registration steps to obtain a more objective analysis of the entire registration algorithm. A sample of the simulated datasets is shown in Figure 5.19.

For rigid registration, we attempt to recuperate the translational error that was applied to the dataset while simulating the dataset using the rigid registration algorithm described precedent. The mean and standard deviation of the error are computed and given in Table 5.7.
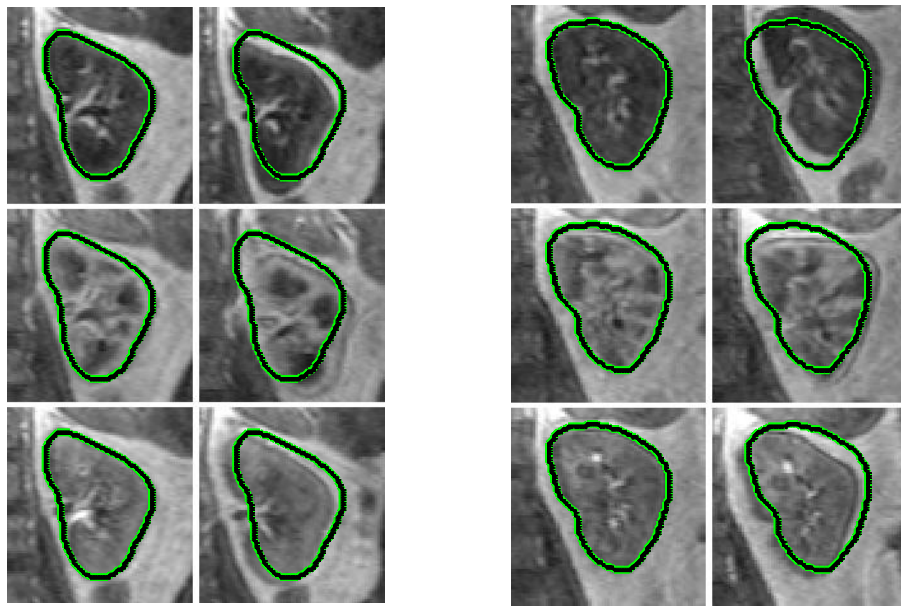
Fig.5.19. Two sample datasets (left and right) showing different 2D slices (different rows) of the kidney: (1st column) pre-aligned; (2nd column) after simulated transformation.

From Table 5.7, it is observed that when the non-rigid deformation level is low (up to 2 voxels shift per node in the deformation grid), the mean and standard deviation of the error is small. When the deformation level becomes higher (up to 6 voxels shift per node in the deformation grid), it becomes increasingly difficult for employing gradient difference because the local edges are deformed causing the orientation to become distorted. Thus, the mean and standard deviation of the error becomes higher for both of

the simulated datasets. The errors are still reasonable as the maximum error distance is still under 5 voxels for all simulated cases.

Table.5.7. Mean and standard deviation of error (in voxel) for rigid registration of different datasets with different simulated deformation levels.

Dataset	Deformation level	Mean Error			Standard Deviation of Error		
		e_x	e_y	e_z	e_x	e_y	e_z
1	Low	0.0857	0.1714	0.1143	0.2840	0.3824	0.3228
	Medium	0.3143	0.2286	0.1143	0.4710	0.4260	0.3228
	High	0.4000	0.4000	0.2286	0.4971	0.4971	0.4902
2	Low	0.1951	0.0732	0.0488	0.4012	0.2637	0.2181
	Medium	0.4634	0.1707	0.0976	0.5049	0.3809	0.3004
	High	0.6829	0.4146	0.1220	0.5674	0.4988	0.3313

Next, we attempt to estimate the pseudo ground-truth of the dataset and to register the images non-rigidly using the demons algorithm. Figure 5.20 shows the sample registration results. It is observed that the estimated pseudo ground-truth dataset in the 3rd column has suffered from greater blurring effects due to the entire simulated dataset undergoing too much consecutive deformations. But the main edges of the kidney remain in the estimated pseudo ground-truth and that is why the resulting non-rigid registration volume still resembles the original image volume. In reality, the kidney volumes do not exhibit large random deformations as used in the simulated datasets. Therefore, the

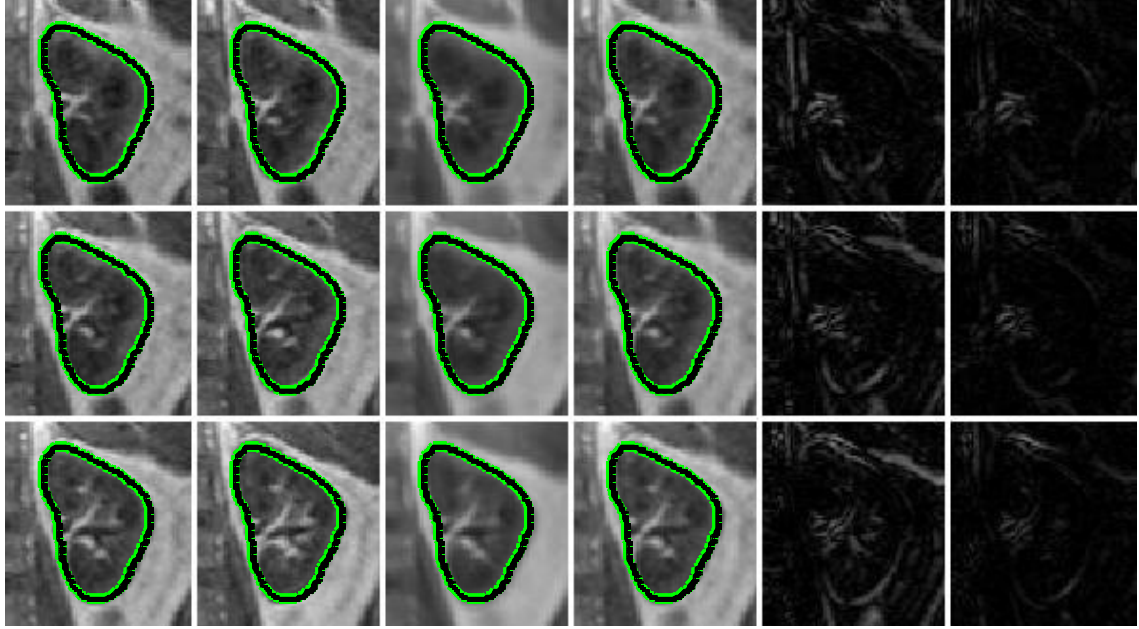pseudo ground-truth estimation is more accurate and a better registered dataset is obtained.

A
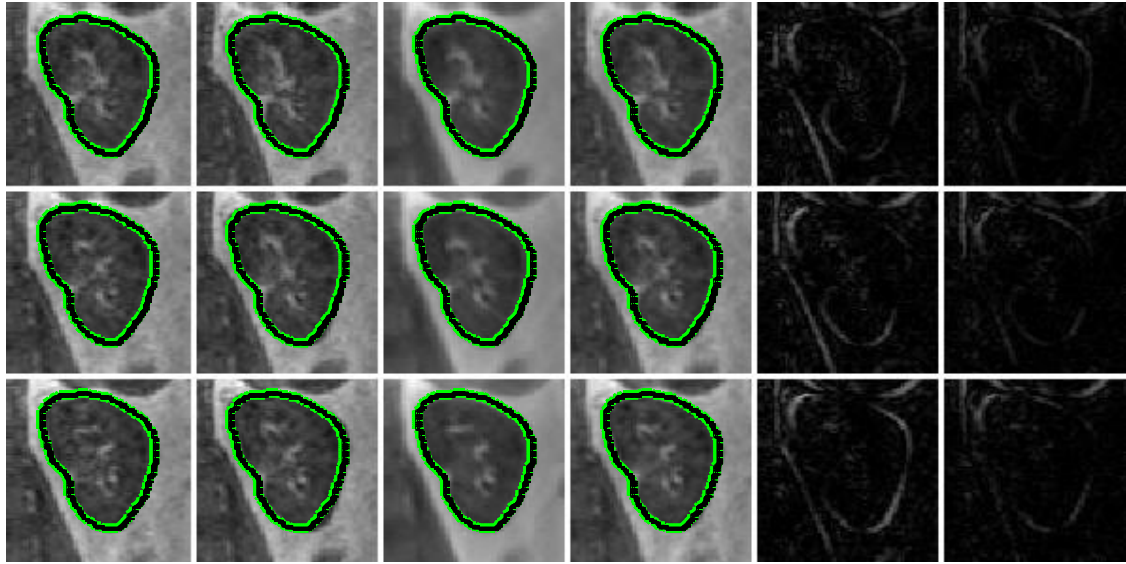
B

Fig.5.20. Sample datasets A and B showing different 2D slices (different rows) of: (1st column) an aligned kidney; (2nd column) kidney after simulated deformation; (3rd column) estimated pseudo-ground truth kidney; (4th column) registered kidney; (5th column) absolute difference of kidney between columns 1 and 2; (6th column) difference of kidney between columns 1 and 4.

To quantify the results obtained for the non-rigid registration step objectively, we include the translational results obtained from the first rigid registration step in the computation of the kidney mask which is compared to the kidney mask post-registration. The non-rigid registration results are shown in Table 5.8 and Figure 5.21. Table 5.8 gives the mean and standard deviation of the distance measures before and after registration for all simulated datasets. The variation of the distance measure across all time frames of a simulated dataset is shown in Figure 5.21. The low, medium and high deformation levels define deformations of up to 2, 4 and 6 voxels of each node in the grid, respectively.

From Table 5.8, it is observed that the mean distance between the kidney masks increases with a larger simulated deformation. The mean distance between the kidney masks also improved post-registration for all cases where only simulated deformations is involved. It is noted, however, that moderate and large simulated deformations become increasingly difficult to recover using the non-rigid registration algorithm. The registration algorithm works better when a small simulated deformation is present; small simulated deformation fields can be recovered more easily using the proposed registration algorithm. In Figure 5.21, some volumes have an increased distance between masks post-registration which signifies that the kidneys in these volumes are unable to recover their pre-applied deformations by performing non-rigid registration.

Table.5.8. Translation + deformed simulation: mean and standard deviation of the distance (in voxel) between the kidney masks (1) pre-registration and (2) post-registration.

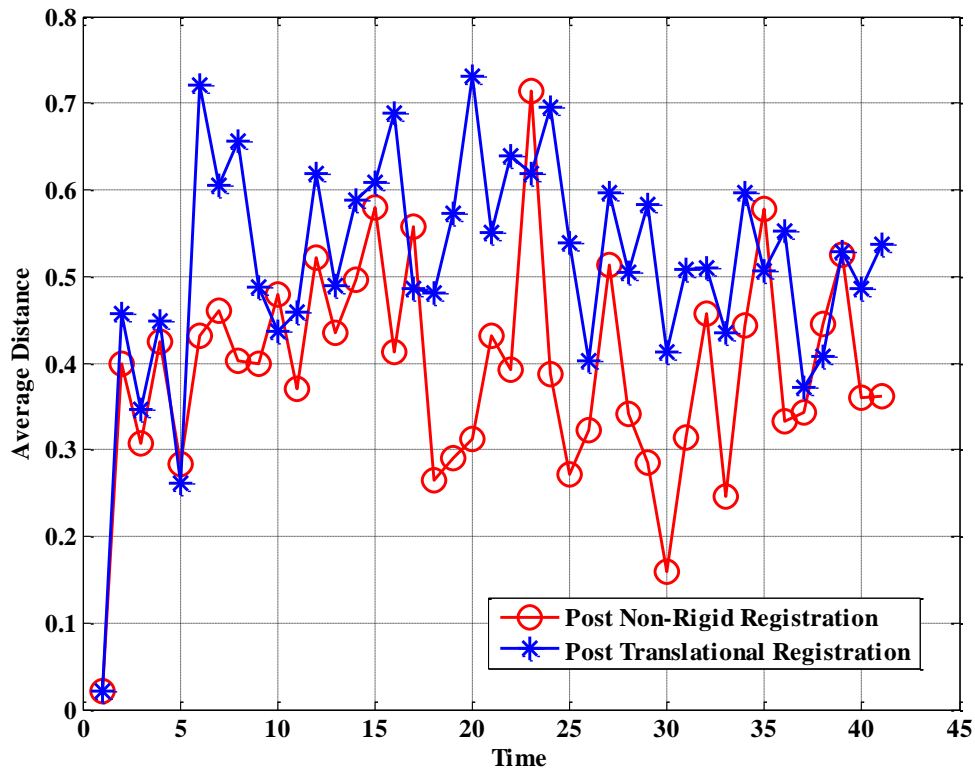| Dataset	Deformation level	Mean		Standard Deviation	
		Pre-Registration	Post-Registration	Pre-Registration	Post-Registration
1	Low	0.3775	0.2846	0.1431	0.1114
	Medium	0.5281	0.4013	0.1043	0.1088
	High	0.6885	0.4858	0.1801	0.1350
2	Low	0.3596	0.2591	0.0989	0.0880
	Medium	0.4766	0.3753	0.1003	0.0933
	High	0.6270	0.5070	0.1563	0.1194

Fig.5.21. Translation + deformed simulation: curves showing the average distance between the boundaries of the kidney masks over time: (1) post translational registration and (2) post non-rigid registration.

CHAPTER 6

CONCLUSION

In this thesis, a semi-automatic non-rigid registration algorithm for renal images is investigated, detailed and analyzed. A multi-level approach was proposed, where a rigid registration step accounts for large initial translational errors before a non-rigid registration step accounts for local deformations of the kidney. To the best of our knowledge, all but one of the registration methods for renal images found in rich literature is rigid-based, as it is widely assumed that kidneys do not exhibit non-rigid motion. But this assumption is not true, even for healthy kidneys. Transformations (translation and non-rigid deformations) are caused by mainly patient's motion and breathing, and in the case of a diseased kidney, cysts, tumors and other anomalies will contribute to the misalignments.

For rigid registration, a graph-cut method was proposed. A graph-cut solution offers efficiency in computational timing and flexibility in defining the node structure, node size and node links. Super-nodes are considered where each super-node contains a certain volume of the kidney, which is unique because most imaging methods employing graph cuts represents each node by a pixel/voxel. Coupled with gradient difference as the main similarity measure, graph-cuts is able to make use of local gradient information to obtain a reasonable global solution. Other graph structures are also possible with graph-cuts, but the other structures tested do not give better results than the regular grid structure as proposed in this thesis.

It is concluded with several tests that in our graph-cuts implementation, a larger node size increases the robustness of the algorithm, but at the expense of higher

computational cost. Thus, it is best to select a suitable node size that allows the registration algorithm to achieve reasonable errors. Moreover, the use of weights which signify the amount of gradient information contained within a super-node allows for a more robust solution, as the registration results became better in terms of mean error and also in terms of the maximum error distance. The weights ensure that only nodes around the boundary of the kidney are considered, as the kidney boundary is consistent over time unlike the outline of the medulla which appears only during the contrast phase. Lastly, mutual information is not used as the main similarity measure because of its poor computational timing and unreliable results for certain volumes.

For the non-rigid registration, the reference images are formed by means of estimating the pseudo ground-truth for each dataset. An estimation method was borrowed from myocardial image registration and adapted to fit the renal image registration context. There exist a few limitations to this stage of the non-rigid registration; it is imperative that the initial stage of translational registration must be fairly accurate in order for the dataset estimation to be good. In addition, the kidney must not be deformed heavily in random directions in order for the estimation to be more accurate with lesser blurring effects. This limitation was examined in the previous chapter where the registration algorithm was tested on the various simulated datasets with various pre-defined transformations (translation, B-spline deformation or both).

The demons algorithm is then used to register the volumes non-rigidly. The parameters are set in such a way that only small deformations are accounted for. Several simulated datasets with different known random transformations are applied on the pre-aligned images and the registration algorithm is used to recover the transformations with

different degrees of freedom: (1) only translations; (2) only non-rigid deformations with different levels (small/moderate/large); and (3) translations and non-rigid deformations with different levels of non-rigid motion. It is observed that the demons algorithm is able to recover translations up to a maximum of 5 voxels satisfactorily. Thus, a criterion for evaluating rigid registration in the first step is that the maximum error distance must be less than 5 voxels for each volume in the datasets. The demons algorithm is able to recover small and moderate deformations applied to the pre-aligned images sufficiently. But, when larger deformations are concerned, there exists several volumes where the registration results are not reasonable, even though the fitting of the kidney masks has improved. It is noticed that the volumes where the demons algorithm is unable to register accurately have a low contrast between the kidney and the background tissues (during which the contrast agent washes into the kidney). The demons algorithm is edge-emphasized; thus, when the edge becomes weaker due to a low contrast, the registration algorithm will not be able to obtain an optimal solution. Moreover, for the simulated dataset where both translation and non-rigid deformations are applied, the translational error after rigid registration increases with the amount of deformation pre-applied. This is due to the use of gradient difference as the similarity measure, where the edge orientation could not be determined properly. But, given the three different levels of deformation, the translational registration is still able to account for translational errors with a small error distance (<5 voxels).

As a continuation to the research work as presented in this thesis, several other considerations could be made in: (1) the graph cut algorithm to incorporate other graph structures, but how the nodes are linked within each graph still needs to be investigated,

since we can only have one realization of the smoothness cost for all the neighborhood links present; (2) a self-validation method to validate the results of the rigid registration without the use of the ground truth provided and re-run the registration algorithm using another setup (different reference frame/different node sizes/different grid structures) if needed; (3) the demons registration algorithm where a multi-level solution could be introduced to allow non-rigid registration to account for global errors first before accounting for more local deformations by adjusting the parameters for each different level; and (4) a more robust quantitative measure to determine the robustness of the non-rigid registration algorithm.

BIBLIOGRAPHY

1.   Maintz, J.B.A. and Viergever, M.A., *A Survey of Medical Image Registration.* Medical Image Analysis, 1998. **2**(1): p. 1-36

2.   Irani, M. and Peleg, S., *Motion Analysis for Image Enhancement: Resolution, Occlusion, and Transparency.* Journal of visual communication & image representation, 1993. **4**(4): p. 324-335

3.   Sun, Z. and Tekalp, A.M., *Image Registration Using A 3-D Scene Representation.* in *ICIP*. 1998. **1**: p. 828-832

4.   Stevens, M.R. and Beveridge, J.R., *Precise Matching of 3-D Target Models to Multisensor Data.* in *IEEE Transactions on Image Processing.* 1997. **6**(1): p. 126-142

5.   Penney, G.P., et al., *A Comparison of Similarity Measures for Use in 2-D–3-D Medical Image Registration.* in *IEEE Transactions on Medical Imaging.* 1998. **17**(4): p. 586-595

6.   Lv, D., et al., *Dynamic Contrast-Enhanced Magnetic Resonance Images of the Kidney*, in *IEEE Engineering in Medicine and Biology Magazine.* 2008. p. 36-41.

7.   E-MRI, *Main applications of perfusion MRI.* Campus Medica 2007; Available from: http://www.imaios.com/en/e-Courses/e-MRI/Cerebral-perfusion-imaging/applications.

8.   Li, C., Sun, Y., and Chai, P., *Pseudo ground truth based nonrigid registration of myocardial perfusion MRI.* Medical Image Analysis, 2011: p. 449-459

9.   Thirion, J.-P., *Image matching as a diffusion process: an analogy with Maxwell's demons.* Medical Image Analysis, 1998. **2**(3): p. 243-260

10.   Favretto, F.O., Bergo, F.P.G., and Falcao, A.X., *A Fast and Automatic Method for 3D Rigid Registration of MR Images of the Human Brain.* in *Computer Graphics and Image Processing.* 2008: p. 121-128

11.   So, R.W.K. and Chung, A.C.S., *Non-Rigid Image Registration by Using Graph-Cuts with Mutual Information.* in *ICIP.* 2010: p. 4429-4432

12.   Tang, T.W.H. and Chung, A.C.S., *Non-rigid Image Registration Using Graph-cuts.* MICCAI, 2007: p. 916-924

13.   Jenkinson, M. and Smith, S., *A Global Optimisation Method for Robust Affine Registration of Brain Images.* Medical Image Analysis, 2001. **5**: p. 143-156

14. Nyúl, L.G., Udupa, J.K., and Saha, P.K., *Incorporating a Measure of Local Scale in Voxel-Based 3-D Image Registration*. in *IEEE Transactions on Medical Imaging*. 2003. **22**(2): p. 228-237

15. Commowick, O., et al., *An Efficient Locally Affine Framework for the Registration of Anatomical Structures*. ISBI, 2006: p. 478-481

16. Gupta, S.N., et al., *Fast Method for Correcting Image Misregistration Due to Organ Motion in Time-Series MRI Data*. Magnetic Resonance in Medicine, 2003. **49**: p. 506-514

17. Shen, D., et al., *Consistent Estimation of Cardiac Motions by 4D Image Registration*. in *MICCAI*. 2005: p. 902-910

18. Okker, B.H., et al., *Accurate and Fully Automatic 3D Registration of Spinal Images Using Normalized Mutual Information*. in *IEEE International Workshop on Biomedical Circuits & Systems*. 2004: p. S3/1-5

19. Khalifa, F., et al., *Non-invasive Image-Based Approach for Early Detection of Acute Renal Rejection*. in *MICCAI*. 2010: p. 10-18

20. Zöllnera, F.G., et al., *Assessment of 3D DCE-MRI of the kidneys using non-rigid image registration and segmentation of voxel time courses*. Computerized Medical Imaging and Graphics, 2009. **33**: p. 171-181

21. R. Sance, et al., *Alignment of 3D DCE-MRI Abdominal Series for Optimal Quantification of Kidney Function*. 5th International Symposium on image Signal Processing and Analysis, 2007: p. 413-417

22. Mahapatra, D. and Sun, Y., *Rigid Registration of Renal Perfusion Images Using a Neurobiology-Based Visual Saliency Model*. EURASIP Journal on Image and Video Processing, 2010: p. 1-16

23. Song, T., et al., *An automated three-dimensional plus time registration framework for dynamic MR renography*. Journal of visual communication & image representation, 2010. **21**(1): p. 1-8

24. Sun, Y., Jolly, M.-P., and Moura, J.M.F., *Integrated registration of dynamic renal perfusion MR images*. in *Int. Conf. Image Processing*. 2004: p. 1923-1926

25. Yim, P.J., et al., *Registration of Time-Series Contrast Enhanced Magnetic Resonance Images for Renography*. in *CBMS*. 2001: p. 516-520

26. Weibel, T., et al., *Endoscopic Bladder Image Registration Using Sparse Graph-Cuts*. in *ICIP*. 2010: p. 157-160

27. Guo, Y., Suri, J., and Sivaramakrishna, R., *Image Registration for Breast Imaging: A Review*. in *EMBC*. 2005. Shanghai, China: p. 3379-3382

28. Buerger, C., Schaeffter, T., and King, A.P., *Hierarchical Adaptive Local Affine Registration for Respiratory Motion Estimation from 3-D MRI*. in *ISBI*. 2010: p. 1237-1240

29. Li, D., et al., *Multiscale Registration for Noisy Medical Images*. in *ICCDA*. 2010. **1**: p. 469-473

30. Wang, H., et al., *Validation of an accelerated 'demons' algorithm for deformable image registration in radiation therapy.* Phys. Med. Biol., 2005. **50**: p. 2887–2905

31. Livyatan, H., Yaniv, Z., and Joskowicz, L., *Gradient-Based 2-D/3-D Rigid Registration of Fluoroscopic X-Ray to CT*. in *IEEE Transactions on Medical Imaging* 2003. **22**(11): p. 1395-1406

32. Pluim, J.P.W., Maintz, J.B.A., and Viergever, M.A., *Mutual-Information-Based Registration of Medical Images: A Survey*. in *IEEE Transactions on Medical Imaging*. 2003. **22**(8): p. 986-1004

33. Ma, Y. and Tian, J., *The Algorithm of Rapid Medical Image Registration by Using Mutual Information*. in *ICBBE*. 2010

34. Fookes, C., Williams, J., and Bennamoun, M., *Global 3D Rigid Registration of Medical Images.* ICIP, 2000. **2**: p. 447-450

35. Glocker, B., et al., *Linear Image Registration through MRF Optimization*. in *ISBI*. 2009: p. 422-425

36. Arsigny, V., et al., *A Fast and Log-Euclidean Polyaffine Framework for Locally Linear Registration.* Journal of Mathematical Imaging and Vision, 2009. **33**(2): p. 222-238

37. Crum, W.R., Hartkens, T., and Hill, D.L.G., *Non-rigid image registration: theory and practice.* The British Journal of Radiology, 2004. **77**: p. 140-153

38. Mahapatra, D. and Sun, Y., *Nonrigid Registration of Dynamic Renal MR Images Using a Saliency Based MRF Model*. in *MICCAI*. 2008: p. 771-779

39. Qiu, Z., Tang, H., and Tian, D., *Non-rigid medical image registration based on the thin-plate spline algorithm*. in *IEMBS*. 2009. **1**: p. 522-527

40. Kolmogorov, V. and Zabih, R., *What Energy Functions can be Minimized via Graph Cuts?* IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI). **26**(2): p. 147-159.February 2004.

41. Hennemuth, A., et al., *A Comprehensive Approach to the Analysis of Contrast Enhanced Cardiac MR Images*. in *IEEE Trans. Med. Imag.* 2008. **27**: p. 1592-1610

42.   Vezhnevets, V. and Konouchine, V., *"GrowCut" - Interactive Multi-Label N-D Image Segmentation by Cellular Autonoma.* 2005

43.   Lankton, S., *GrowCut Segmentation in Matlab*. Available from: http://www.shawnlankton.com/2008/03/growcut-segmentation-in-matlab/. 2008.

44.   Bagon, S., *Matlab Wrapper for Graph Cut*. Available from: www.wisdom.weizmann.ac.il/~bagon. 2006.

45.   Boykov, Y. and Kolmogorov, V., *An Experimental Comparison of Min-Cut/Max-Flow Algorithms for Energy Minimization in Vision.* IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI). **26**(9): p. 1124-1137.September 2004.

46.   Boykov, Y., Veksler, O., and Zabih, R., *Fast Approximate Energy Minimization via Graph Cuts.* IEEE transactions on Pattern Analysis and Machine Intelligence (PAMI). **23**(11): p. 1222-1239.2001.

47.   Hestenes, M.R. and Stiefe, E., *Methods of Conjugate Gradients for Solving Linear Systems.* Journal of Research of the National Bureau of Standards, 1952. **49**(6): p. 409-436

48.   Poschel, J., *Inverse Spectral Theory - Gaussian Elimination*. Vol. 130. 1987. 169-176

49.   Pennec, X., Cachier, P., and Ayache, N., *Understanding the "Demon's Algorithm": 3D Non-rigid Registration by Gradient Descent*. in *MICCAI*. 1999. **LNCS 1679**: p. 597-606

50.   Kroon, D.J. and H.Slump, C., *MRI Modalitiy Transformation in Demon Registration*. in *ISBI*. 2009: p. 963-966

51.   Kroon, D.J., *Multimodality Non-Rigid Demon Algorithm Image Registration*. 2008; Available from: http://www.mathworks.com/matlabcentral/fileexchange/21451?product=IP.