Fast FEM-based Non-Rigid Registration

Karteek Popuri, Dana Cobzas, Martin Jägersand
University of Alberta
Department of Computing Science
kpopuri@ualberta.ca, dana@cs.ualberta.ca, jag@cs.ualberta.ca

Abstract

In this paper, we present a fast and accurate implementation of the diffusion-based non-rigid registration algorithm. Traditionally, finite differences are used to implement registration algorithms due to their ease of implementation. However, finite differences are sensitive to noise, and they have a narrow numerical stability range. Further, finite differences employ a uniform grid. This is often not desirable in the case of registration, as finer resolution is needed to capture the displacement field in regions that have a high number of image features, as opposed to homogeneous regions with fewer features. On the other hand, the less explored Finite Element Methods are ideal for the non-rigid registration task, as they use a non-uniform discretization of the image domain, placing points based on the local image-feature information. We present such an FEM-based implementation of a popular diffusion-based registration algorithm [8]. Originally, this algorithm was implemented using finite differences. Experimentally, we show that our implementation is much faster than the corresponding finite difference implementation, and that it achieves this computational speed without compromising the accuracy of the non-rigid registration results.

1. Introduction

Image registration is an important image processing task. It is often the first step in various computer vision algorithms. In medical imaging, images are often required to be registered to a standard reference (an atlas) for comparison, or images from two different time instants are required to be registered with each other to observe the evolution of a disease. Image registration refers to the task of finding pixel correspondences between two given images. Specifically, we want to estimate a transformation T such that a pixel $\mathbf{x}$ in image I corresponds to the pixel $T(\mathbf{x})$ in image J. In *rigid* or *affine* registration, the transformation T is a linear transformation. The affine transformation can be parametrized by a finite number of parameters and is thus constrained in the parameter space. In *non-rigid* registration, the transformation is estimated as a vector *displacement field* $\mathbf{U}(\mathbf{x})$, such that at each pixel location $\mathbf{x}$ we have $T(\mathbf{x}) = \mathbf{x} + \mathbf{U}(\mathbf{x})$. Thus, a pixel $\mathbf{x}$ with intensity $I(\mathbf{x})$ in the target image corresponds to a point $T(\mathbf{x})$ with intensity $J(\mathbf{x} + \mathbf{U}(\mathbf{x}))$. In the case of medical images, one usually estimates a non-rigid transformation to take into account the highly non-linear deformations that can occur in the underlying anatomy.

In image registration, the transformation between two images is estimated by minimizing a similarity metric (usually dependent on pixel intensities). In the case of non-rigid registration, this minimization is ill-posed as we need to search for the displacement parameters at each pixel in a large parameter space. Hence, extra regularization terms are usually added. In an early work, Broit [1] introduced the idea of elastic regularization terms to model the non-rigid deformations between images. Later, Christensen *et al.* [2] proposed the viscous fluid regularization to handle large displacements between images. In recent years, Stefenescau *et al.* [8] combined the elastic and viscous regularization ideas and proposed a novel visco-elastic regularization term. This regularization constraint was imposed by smoothing the deformation field using a diffusion process, which involves solving two Partial Differential Equations (PDEs). They used Finite Difference methods to solve the PDEs. However, Finite Difference (FD) methods are limited due to the fixed grid size. Consequently, they often fail to accurately model the highly non-uniform deformation fields that commonly arise in non-rigid registration problems. On the other hand, Finite Element Methods (FEMs), due to their ability to handle variable grid sizes, are better suited for the task of non-rigid image registration. Subsequently, there have been some efforts [3, 5] to employ FEMs for non-rigid image registration. However, these methods employ complex biomechanical models to model the elastic behaviour of the underlying anatomy in the images. These models usually are tuned to a specific type of image registration task and hence are not applicable for general non-rigid image registration.

In this paper, we explore the use of FEMs for solving the diffusion PDEs corresponding to general visco-elastic regularization constraints. Such regularization constraints can be imposed in a wider range of image registration problems. We also describe a procedure to generate non-uniform grids that adapted to the features in the image. We show the results of this FEM-based non-rigid registration methodology on 2D MRI and CT images.

2. Diffusion-based non-rigid image registration

In this section, we briefly describe the non-rigid registration algorithm proposed by Stefanescu *et al.* [8] and present our proposed Finite Element formulation of this non-rigid registration method in detail.

2.1. Non-rigid registration by gradient descent

Let I be the template or the target image and let J be the source or the input image. The goal of non-rigid registration is to compute a dense deformation (displacement) field $\mathbf{U}$ between the two images. Stefanescu *et al.* [8] proposed the iterative minimization of the sum of squared distance (SSD) similarity criterion using a compositive update rule to obtain this displacement field $\mathbf{U}$. At each iteration a correction field $\mathbf{u}$ was estimated such that $SSD(I, J \circ (\mathbf{U} \circ \mathbf{u}))$ is minimized. Here $J \circ \mathbf{U}(\mathbf{x}) = J(\mathbf{x} + \mathbf{U}(\mathbf{x}))$ is the warped image and the compositive update of the displacement field $\mathbf{U}$ is given by $\mathbf{U} \circ \mathbf{u}(\mathbf{x}) = \mathbf{U}(\mathbf{x} + \mathbf{u}(\mathbf{x})) + \mathbf{u}(\mathbf{x})$, where $\mathbf{x}$ is a pixel in the image. Using a first order Taylor expansion we get:

$$\begin{aligned} &SSD(I, J \circ (\mathbf{U} \circ \mathbf{u})) \\ &\quad = \int [I(\mathbf{x}) - (J \circ (\mathbf{U} \circ \mathbf{u}))(\mathbf{x})]^2 d\mathbf{x} \\ &\quad \approx SSD(I, J \circ \mathbf{U}) \\ &\quad + \int 2[(J \circ \mathbf{U})(\mathbf{x}) - I(\mathbf{x})][\nabla(J \circ \mathbf{U})](\mathbf{x})^T \mathbf{u}(\mathbf{x}) d\mathbf{x} \end{aligned} \tag{1}$$

From the above we see that the gradient of the SSD criterion is:

$$\nabla SSD = 2[(J \circ \mathbf{U})(\mathbf{x}) - I(\mathbf{x})][\nabla(J \circ \mathbf{U})](\mathbf{x}) \tag{2}$$

Now given the current displacement field $\mathbf{U}^k$, using a gradient descent strategy we can update this displacement field by $\mathbf{u}^k = -\epsilon \, \nabla SSD$ to obtain:

$$\begin{aligned} \mathbf{U}^{k+1} &= \mathbf{U}^k \circ \mathbf{u}^k \\ &= \mathbf{U}^k \circ (-\epsilon \, 2[J \circ \mathbf{U}^k - I][\nabla(J \circ \mathbf{U}^k)]) \end{aligned} \tag{3}$$

where, ϵ (a small fraction) is a parameter of the algorithm.

2.2. Diffusion-based smoothing

In the case of non-rigid registration, the smoothing of the deformation field $\mathbf{U}$ is necessary due to two reasons: (1) to regularize the minimization problem (2) to roughly model the physical continuity of the underlying real non-rigid transformation. For this purpose, Stefanescu *et al.* [8] proposed the smoothing of the displacement field $\mathbf{U}$ and also the smoothing of the correction field $\mathbf{u}$ at every iteration.

The smoothing of the displacement field $\mathbf{U}$ is implemented on each of the components $\mathbf{U} = [U_x U_y]$ as:

$$\frac{\partial U_l}{\partial t} = \mathrm{div}(D(\mathbf{x}) \nabla U_l) \quad l \in \{x, y\} \tag{4}$$

where $D(\mathbf{x})$ is the scalar *stiffness* field. This smoothing is equivalent to the physical elastic constraint where the elasticity of various regions in the image is described by the stiffness field $D(\mathbf{x})$.

The smoothing of the correction field $\mathbf{u}$ is implemented on each of the components $\mathbf{u} = [u_x u_y]$ as:

$$\frac{\partial u_l}{\partial t} = (1 - k(\mathbf{x}))\mathrm{div}(\nabla u_l) \quad l \in \{x, y\} \tag{5}$$

where $k(\mathbf{x})$ is the *local confidence*, which measures the reliability of the correction field values obtained at each pixel p. The local confidence is estimated based on the gradient (∇J) of the source image as:

$$k(\mathbf{x}) = \exp\left(\frac{-c}{\left(\frac{\|\nabla J\|}{\lambda}\right)}\right) \tag{6}$$

From the above, we can see that $k(\mathbf{x})$ is close to 1 when the gradient is high and $k(\mathbf{x})$ is close to 0 for low gradient values. This is motivated by the fact that the edges (high gradients) in the image correspond to reliable landmarks in the image and low gradients correspond to uniform regions. This smoothing corresponds to the physical viscosity constraint.

2.3. Proposed FEM-based non-rigid registration method

In general, the displacement field is computed on a uniform grid (with spacing 1 pixel $\times$ 1 pixel) defined on the image domain. Then, for a rectangular 2D image domain $I : \mathbb{R}^{L \times W} \rightarrow \mathbb{R}$ we have $\mathbf{U}^k \equiv \mathbf{U}^k(\mathbf{x}_{ij})$, where $\{\mathbf{x}_{ij}\}, i \in \{1, 2, \ldots, L\}, j \in \{1, 2, \ldots, W\}$ are the uniformly placed grid points. Such a uniform discretization of the image domain is not always desirable. This is because there might exist regions in the image where the displacement field is highly non-uniform (these are usually regions around the significant features in the image) and hence require finer discretization (spacing less than 1 pixel $\times$ 1 pixel) and on

the other hand we can have regions with a nearly constant or linear displacement field (homogeneous regions in the image) which can be well approximated using a coarse grid (spacing greater than 1 pixel $\times$ 1 pixel). For example, in the case of MRI images of the brain the regions corresponding to the ventricles and the skull exhibit relatively non-uniform deformations compared to the rest of the brain. Hence, a non-uniform discretization is often needed to estimate the displacement field accurately over the whole image domain.

Further, the diffusion PDEs (equations (4), (5)) are also solved on the uniform grid using the FD method to smooth the update and displacement fields. These smoothing steps are considerably slow even when implemented using efficient FD numerical schemes. One technique to speedup the smoothing steps (without losing accuracy of the solution) is to avoid solving the PDEs at every pixel in homogeneous regions, where the displacement field is expected to be almost constant or linear. Such a methodology can be easily implemented by using the Finite Element method instead of the Finite Differences.

In this paper, we propose to solve for the displacement field $\mathbf{U}$ on a non-uniform grid with triangular elements using the Finite Element method in order to improve the computational speed and accuracy of the non-rigid registration. The non-uniform grid is generated on the image by the Delaunay triangulation [6] of a set of input points. These points are chosen automatically so that the resulting grid is well adapted to the important features in the image (see next section for more details). Let us denote the points on the grid (henceforth called *nodes*) by $\{P_n\}$, $n = \{1, 2, \ldots, N\}$ and the triangles adjacent to the node by δ_{ni} (see Figure 1). Using the Finite Element Method, the displacement field $\mathbf{U}$ can now be approximated on the non-uniform grid by a set of piecewise-linear basis functions $\{\phi_1, \phi_2, \ldots, \phi_N\}$:

$$\mathbf{U}(\mathbf{x}) = \sum_{n=1}^{N} \mathbf{U}(P_n)\phi_n(\mathbf{x}) \qquad \forall\, \mathbf{x} \in \mathbb{R}^{L\times W} \tag{7}$$

where $\mathbf{U}(P_n)$ is the value of displacement field at the node P_n and we assign a basis function ϕ_n to each node P_n which is uniquely defined as:

$$\phi_n(\mathbf{x}) = \begin{cases} \text{is linear within each triangle } \delta_{ij} \\ 1 \quad \text{at each node } P_n \\ 0 \quad \text{at every other node } P_m \neq P_n \end{cases} \tag{8}$$

From the above we see that the displacement field needs to be computed only at the nodal locations, i.e, $\mathbf{U}(P_n)$. For this purpose we propose a novel iterative non-rigid registration strategy as follows:

Step 1 : At iteration $k = 0$, initialize $\mathbf{U}^0(P_n)$

For subsequent iterations $k > 0$

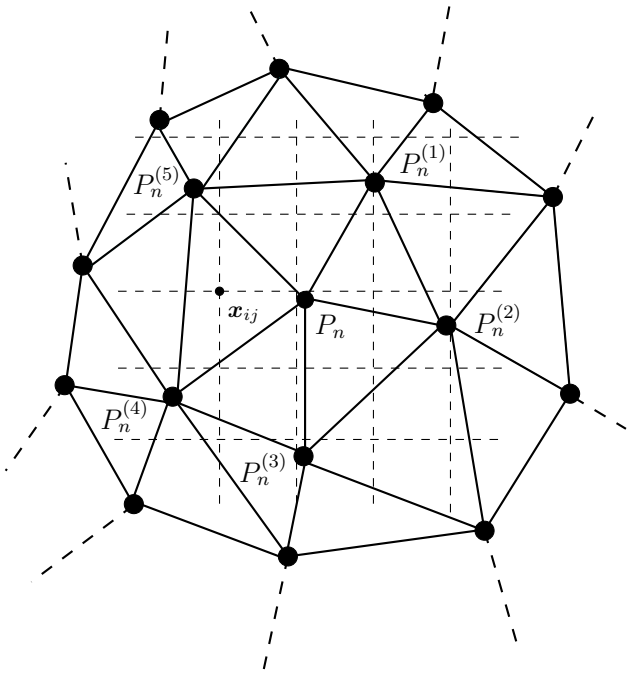

Figure 1. This figure shows a schematic of the non-uniform grid on the rectangular 2D image domain. Here, a section of the grid with a node P_n and its neighboring nodes $P_n^{(i)}$, $i \in \{1, 2, \ldots, 5\}$ is shown. We denote the triangles adjacent to P_n by $\delta_{n1} \equiv P_n^{(1)}P_nP_n^{(2)}, \delta_{n2} \equiv P_n^{(2)}P_nP_n^{(3)} \ldots \delta_{n5} \equiv P_n^{(5)}P_nP_n^{(1)}$. The corresponding uniform grid is also shown in dashed lines and a sample pixel $\mathbf{x}_{ij}$ of the uniform grid overlapping with the triangular element of the non-uniform grid is highlighted.

Step 2 : Compute $\mathbf{U}_{ij}^k$ at each pixel $\mathbf{x}_{ij}$ on the uniform grid by interpolating $\mathbf{U}^k(P_n)$ over the triangular elements of the non-uniform grid using basis functions ϕ_n

Step 3 : Compute $\mathbf{u}_{ij}^k = -\epsilon\, 2[J_{ij} \circ \mathbf{U}_{ij}^k - I_{ij}][\nabla(J_{ij} \circ \mathbf{U}_{ij}^k)]$

Step 4 : Compute $\mathbf{u}^k(P_n) = \frac{1}{\sum_{ij}\lambda_{ij}} \sum_{ij} \lambda_{ij}\mathbf{u}_{ij}^k$. The summation is over all the pixels $\mathbf{x}_{ij}$ overlapping the triangles adjacent to the node P_n, where λ_{ij} represents the barycentric coordinate of the pixel $\mathbf{x}_{ij}$ with respect to the node P_n. (see Figure 1)

Step 5 : Update $\mathbf{U}^{k+1}(P_n) = \mathbf{U}^k(P_n) \circ \mathbf{u}^k(P_n)$

Step 6 : Smooth the displacement field $\mathbf{U}^{k+1}(P_n)$ by solving equation (4) on the non-uniform grid using the FEM method

Step 7 : Goto Step 2 and repeat until convergence

For simplicity we represented $J(\mathbf{x}_{ij}) \equiv J_{ij}$, $I(\mathbf{x}_{ij}) \equiv I_{ij}$, $\mathbf{U}(\mathbf{x}_{ij}) \equiv \mathbf{U}_{ij}$, $\mathbf{u}(\mathbf{x}_{ij}) \equiv \mathbf{u}_{ij}$.

In the above proposed non-rigid registration methodology the first novelty is in step 4, where we compute the updates at the nodes $\mathbf{u}(P_n)$ by taking a weighted average of the updates at the neighboring pixels $\mathbf{u}_{ij}$. This step is equivalent to the smoothing of the update field (see equation (5)) to model viscous fluid-like behavior. We also like to point out that this step is computationally much less expensive than the diffusion based smoothing of the update field in equation (5). The second novelty of our proposed approach is in step 6 where we perform the smoothing of the displacement field $\mathbf{U}(P_n)$ on the non-uniform grid, again which is computationally more efficient as we solve the diffusion PDE in equation (4) only at the nodal locations P_n. More details are given in the following section.

2.4. Solving the diffusion PDE using the FEM method

In this section, we present the FEM formulation to solve the diffusion PDE in equation (4). Following Weickert *et al.* [9] we can easily show that solving the diffusion PDE in equation (4), with an initial condition $w^0 = f$ until $t = \alpha$ is same as minimizing the following functional:

$$E(w) = \frac{1}{2}\int_{\Omega}\left[(w-f)^2 + \alpha(D(\mathbf{x})|\nabla w|^2)\right]d\mathbf{x} \qquad (9)$$

To determine the minimizer w of the above functional $E(w)$, we vanish the first variation of the functional $E(w)$ to get the following:

$$L(w,v) = \int_{\Omega}\left[(w-f)v + \alpha D(\mathbf{x})\nabla w.\nabla v\right]d\mathbf{x} = 0 \;\forall v \qquad (10)$$

The basic idea in the Finite Element method is to solve for the minimizer w by restricting the minimization to a finite dimensional subspace spanned by the nodal basis functions $\{\phi_1, \phi_2, \ldots, \phi_n\}$ (see equation (8)). So, to smooth the displacement field U^k (here U^k stands for both the components U_x^k, U_y^k at any iteration k) we substitute $U^k = \sum_{n=1}^{N} U^k(P_n)\phi_n$ and $v \in \{\phi_1, \phi_2, \ldots, \phi_n\}$ (from equation (7)) into the variational equation (10). Now, assuming the stiffness $D(\mathbf{x})$ to be constant over each triangular element in the non-uniform grid we obtain the following set of N linear equations (one for every node):

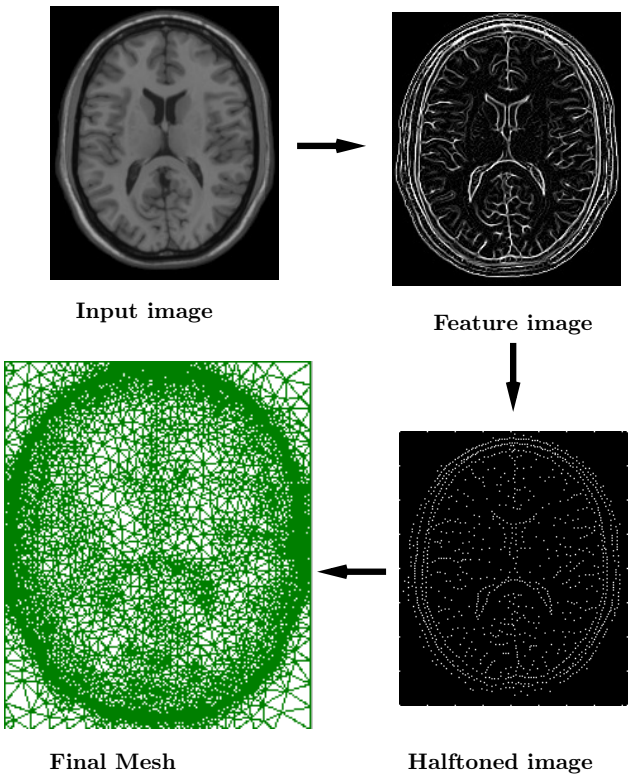

Figure 2. Illustration of the image adaptive non-uniform grid generation. Here the grid is generated on an input MRI image. Notice the greater density of triangular elements in the skull and the ventricle regions.

$$L(U^k, \phi_m) = \sum_{n=1}^{N}(U^k(P_n) - U^{k*}(P_n))\int_{\Omega}\phi_n\phi_m d\mathbf{x} + \alpha\sum_{n=1}^{N}\int_{\Omega} D(\mathbf{x})\nabla\phi_n.\nabla\phi_m d\mathbf{x} = 0 \qquad (11)$$
$$\Omega \equiv \mathbb{R}^{L\times W} \quad m = \{1, 2, \ldots, N\}$$

where U^{k*} is the non-smooth displacement field estimated at the iteration k. We would like to bring to the readers's attention that the above system of linear equations is very sparse as the nodal basis functions vanish almost everywhere except at the respective nodes and their adjacent triangles (see equation (8)). Further we point out that the nodal integrals ($\int_{\Omega}\phi_n\phi_m$, $\int_{\Omega}\nabla\phi_n\nabla\phi_m$) are precomputed (analytically) at the start of the non-rigid registration as they are independent of the displacement field estimates at each iteration. This greatly improves the computational efficiency of our FEM-based non-rigid registration method.

3. Generation of the non-uniform grid

The generation of a non-uniform grid that is well adapted to the features in the image is very important for a good performance of our non-rigid registration methodology. For this purpose we implemented the image-adaptive grid generation strategy proposed by Yang *et.al* [10]. The central idea of this method is to place grid nodes in the image domain so that their spatial density varies according to the local image content. The subsequent Delaunay triangulation and the refinement steps automatically ensure that the fine triangular elements are placed in image regions with high-frequency features while the coarse elements are used to represent the smooth areas. The steps to generate the non-uniform grid are summarized below:

- Given an input 2D image $f(x, y)$ compute the feature map $\sigma(x, y) = \left(\frac{G(x,y)}{A}\right)$ where $G(x, y) = \max |f''_\theta(x, y)| \quad \theta \in [0, 2\pi]$ and A is a normalizing constant
- Halftone the feature image $\sigma(x, y)$ to obtain a binary image
- Input the locations of the white pixels in the binary image as initial grid nodes to a Delaunay grid generation algorithm
- Refine the grid generated from the above step to obtain the final image adapted non-uniform grid

See Figure 2 for an illustration of the above procedure.

4. Implementation details

We implemented our non-rigid registation method primarily in MATLAB. We made use of the MEX interface in MATLAB to implement the smoothing PDEs in both the FD and FEM cases. Further, the computation of the nodal updates in the FEM case was also performed using MEX code. The TRIANGLE package [7] was used for the Delaunay triangulation and refinement step of our grid generation procedure. We ran our registration method on a Intel Quad-Core 2.66GHz CPU with 4GB RAM.

5. Experiments

In this section, we describe the experiments conducted to evaluate the performance of our non-rigid registration method. We considered registration between the following three pairs of images: synthetic MRI source to synthetic MRI target, real patient MRI to atlas MRI, and CT scans of a single patient taken at two different times. In all three cases, we compared the computational speed and accuracy of our FEM-based method with the FD-based non-rigid registration. In the case of FD-based registration, the smoothing steps (for both the displacement and update fields) were one to two orders of magnitude slower than the rest of the steps. Similarly, in the case of FEM-based registration, the computation of the nodal updates and the smoothing of the displacement field took one to two orders of magnitude longer than the rest of the registration steps. Hence, below, we only report the computational time taken for the smoothing steps in the FD case and compare it with the time taken to perform the nodal updates along with the smoothing of the displacement field in the FEM case (see Table 1).

Exp.	Run time FEM	Run time FD
Synth. MRI	13.07 sec	57.06 sec
Real MRI	21.31 sec	112.66 sec
CAT scan	65.59 sec	303.72 sec

Table 1. This table compares the computational speed of the FD-based and the FEM-based non-rigid registration methods on three different image registration tasks. The run times reported above correspond to the time taken for the smoothing steps in the FD case and time taken to perform the nodal updates along with the smoothing of the displacement field in the FEM case. We observe that the FEM-based registration method is $\sim 4-5$ times faster than the FD-based registration method.

In the first experiment, we warped a synthetic brain MRI image (217×181) using a sinusoidal displacement field. This introduced large "wave-like" deformations in the skull and in the ventricle regions. We registered the original synthetic MRI image to the warped synthetic MRI image using both the FEM and FD non-rigid registration methods. We set the stiffness field to $D(\mathbf{x}) = 0.1$ in the ventricle region and to $D(\mathbf{x}) = 0.5$ in the rest of the brain. The registration results are shown in Figure 3. We can clearly observe that the FEM-based registration performs better, especially in the skull region. This is due to the fact that the FEM-based method uses a finer discretization in the skull region, and therefore it is able to better approximate the relatively non-uniform displacements there. On the other hand, the FD-based method fails to capture this non-uniform displacement field as it is limited by the uniform grid size. Further, the FEM-based method was considerably faster as it only took 13.07 sec, compared to the 57.06 sec taken by the FD based method.

In the second experiment, we performed registration between the MRI image of a patient's brain and an atlas using

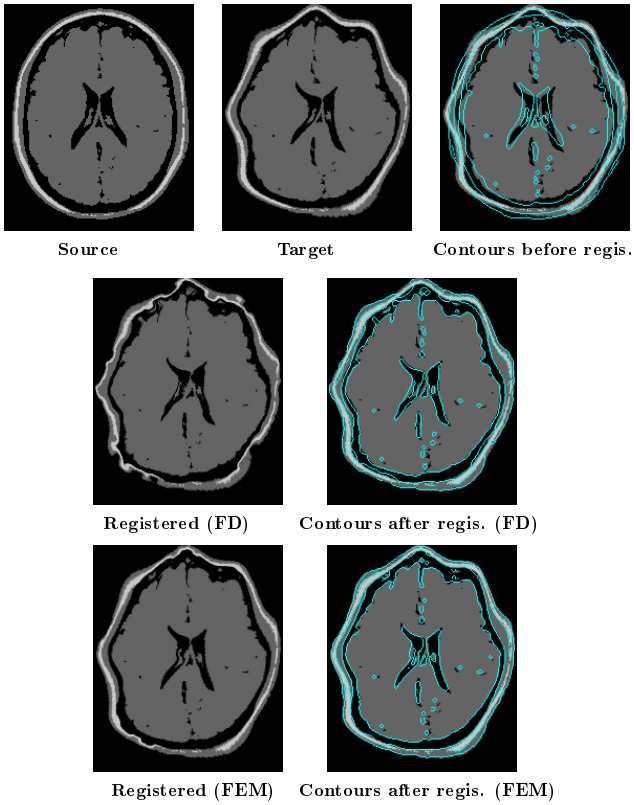

Figure 3. This figure shows the registration results between a synthetic brain MRI image and the same image warped with a sinusoidal displacement field. The stiffness field was set to $D(\mathbf{x}) = 0.1$ in the ventricle region and $D(\mathbf{x}) = 0.5$ in the rest of the brain. Here, we observe that the FEM-registration result better matches with the target image in the skull region, when compared to the FD-registration result.

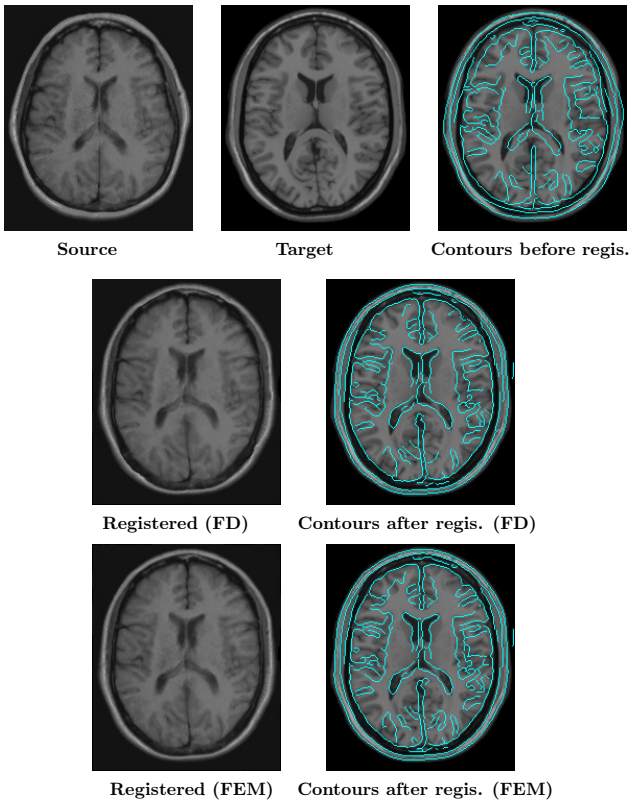

Figure 4. This figure shows the registration results between a patient's brain MRI image and an atlas. Here, the stiffness field was set to $D(\mathbf{x}) = 0.01$ in the ventricle region and to $D(\mathbf{x}) = 0.1$ in the rest of the brain.

both the FD and FEM-based registration methods. Both the atlas and the patient images were of dimensions 217×181. We aligned the patient and atlas images using affine registration before performing the non-rigid registration. Here, we set the stiffness field at $D(\mathbf{x}) = 0.01$ in the ventricle region and $D(\mathbf{x}) = 0.1$ in the rest of the brain. The results of the non-rigid registration can be seen in Figure 4. In this case, both the FD and FEM methods performed well with effectively identical results. When superimposed on the atlas image, the contours in the registered images show a good matching between the features in the registered and atlas images. However, again the FEM based method took much less computational time, at 21.31 sec compared to the FD-based method which took 112.66 sec.

In the final experiment, we considered CT scans of a patient's liver and spleen taken at two different points in time. The CT scans were of size 445×391, almost 4 times the size of the MRI images. We performed both the FEM and FD registrations at this high resolution. Here, the difference between the computational speeds of the FEM- and FD-based methods is more evident. The FEM method took about a minute (65.59 sec) whereas the FD-based method took over 5 minutes (303.72 sec). Looking at the patient's two CT scans suggests a relatively small magnitude of the deformation field between the two images. Both the FEM- and FD-based registration methods are successful in estimating this deformation field, as can be seen in Figure 5. Again, we note that the CT scans were affinely registered before being input to the non-rigid registration methods.

6. Conclusion

In this paper we presented a fast Finite Element Method based non-rigid registration method. Our method employed a grid with variable resolution, placing finer elements in image regions with numerous features and coarser elements in homogenous regions. This improved the computational speed and accuracy of our registration method when compared to the traditional Finite Difference based implementations of non-rigid registration. However, we only evaluated our method on 2D images in this paper. Our method can be easily extended to 3D images and, we are considering it in our future work. Further, in our work we chose the stiffness field by trial and error. In the future we intend to explore the possibility of learning the stiffness field from a set of training images.

Acknowledgements

We are extremely thankful to Dr. Vickie Baracos from the Department of Oncology, Cross Cancer Institute at the University of Alberta for providing us with the CT scan of a patient's liver to perform the experiments reported in this paper. We obtained the T1 MRI brain scan of a patient from the "The whole brain atlas " online resource [4]. We are thankful for the clinical advice provided by Dr. Albert Murtha from the Department of Oncology, Cross Cancer Institute at the University of Alberta, during the course of this work. We would also like to acknowledge Neil Birkbeck, David Lovi from the Department of Computing Science, University of Alberta, for their useful inputs.

References

[1] C. Broit. Optimal registration of deformed images. *Dissertations available from ProQuest*, 1981.

[2] G. Christensen, S. Joshi, M. Miller, et al. Volumetric transformation of brain anatomy. *IEEE Transactions on Medical Imaging*, 16(6):864–877, 1997.

[3] M. Ferrant, A. Nabavi, B. Macq, F. Jolesz, R. Kikinis, and S. Warfield. Registration of 3D intraoperative MR images of the brain using a finite element biomechanical model. *IEEE Transactions on Medical Imaging*, 20(12):1384–1397, 2001.

[4] K. Johnson and J. Becker. The whole brain atlas. `www.med.harvard.edu/AANLIB/home.html`.

[5] D. Loeckx, L. Roose, F. Maes, D. Vandermeulen, and P. Suetens. An Elasticity Penalty: Mixing FEM and Non-rigid Registration. *4th European Conference of the International Federation for Medical and Biological Engineering*, 22, 2008.

[6] F. Preparata and M. Shamos. *Computational geometry: an introduction*. Springer, 1985.

[7] J. Shewchuk. Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator. *Lecture Notes in Computer Science*, 1148:203–222, 1996.

[8] R. Stefanescu, X. Pennec, and N. Ayache. Grid powered nonlinear image registration with locally adaptive regularization. *Medical Image Analysis*, 8(3):325–342, 2004.

[9] J. Weickert, J. Heers, C. Schnorr, K. Zuiderveld, O. Scherzer, and H. Siegfried Stiehl. Fast parallel algorithms for a broad class of nonlinear variational diffusion approaches. *Real-Time Imaging*, 7(1):31–45, 2001.

[10] Y. Yang, N. Wernick, and G. Brankov. A fast approach for accurate context-adaptive mesh generation. *IEEE transactions on image processing*, 12(8):866–881, 2003.

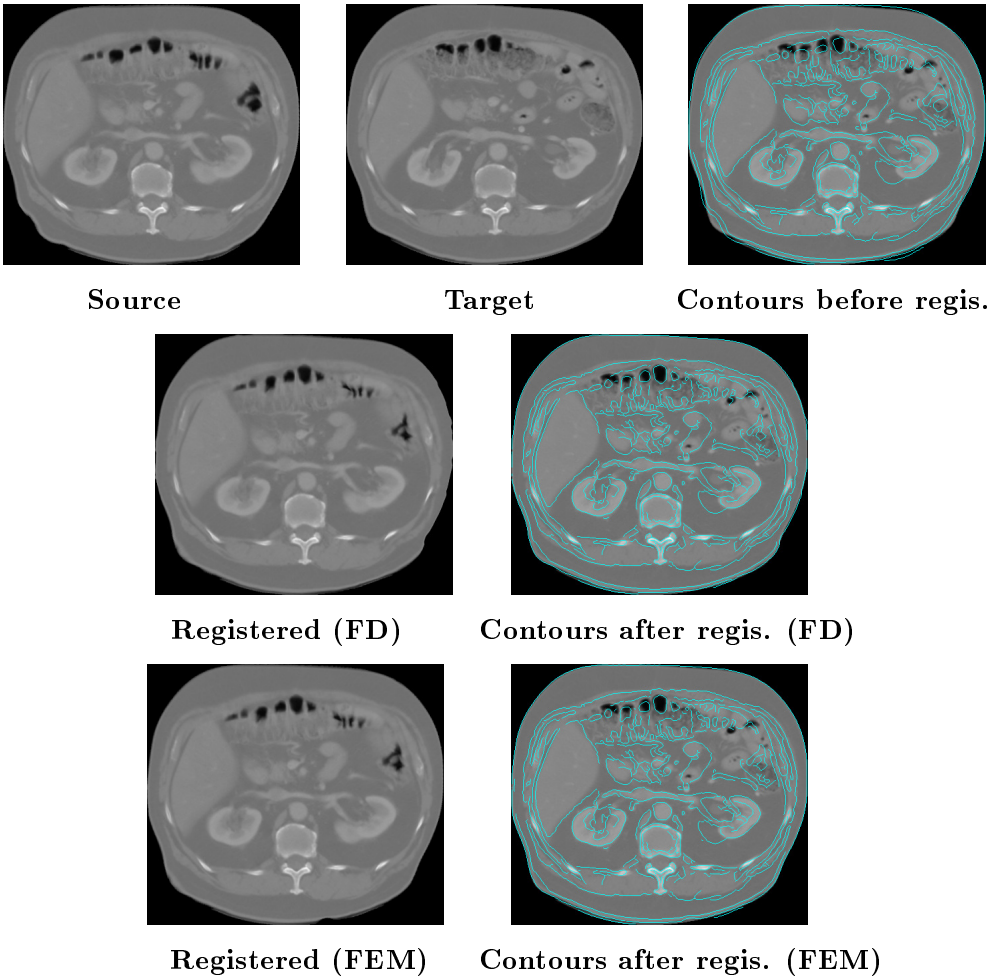

Figure 5. This figure shows the registration results between two CT scans of a patient taken at two different time points. We performed the registration on these high-resolution (445×391) CT scans at their original resolution.