Non-rigid image registration of brain magnetic resonance images using graph-cuts

Ronald W.K. So *, Tommy W.H. Tang, Albert C.S. Chung

Lo Kwee-Seong Medical Image Analysis Laboratory, Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong

ARTICLE INFO

ABSTRACT

We present a graph-cuts based method for non-rigid medical image registration on brain magnetic resonance images. In this paper, the non-rigid medical image registration problem is reformulated as a discrete labeling problem. Based on a voxel-to-voxel intensity similarity measure, each voxel in the source image is assigned a displacement label, which represents a displacement vector indicating which position in the floating image it is spatially corresponding to. In the proposed method, a smoothness constraint based on the first derivative is used to penalize sharp changes in the adjacent displacement labels across voxels. The image registration problem is therefore modeled by two energy terms based on intensity similarity and smoothness of the displacement field. These energy terms are submodular and can be optimized by using the graph-cuts method via α-expansions, which is a powerful combinatorial optimization tool and capable of yielding either a global minimum or a local minimum in a strong sense. Using the realistic brain phantoms obtained from the Simulated Brain Database, we compare the registration results of the proposed method with two state-of-the-art medical image registration approaches: free-form deformation based method and demons method. In addition, the registration results are also compared with that of the linear programming based image registration method. It is found that the proposed method is more robust against different challenging non-rigid registration cases with consistently higher registration accuracy than those three methods, and gives realistic recovered deformation fields.

© 2011 Elsevier Ltd. All rights reserved.

1. Introduction

Image registration has been actively and extensively studied and applied in the field of medical image analysis. By the nature of the image transformation, image registration can be classified into rigid image registration and non-rigid image registration. Rigid image transformation consists of a translation followed by a rotation in the 3D space. As such, in a rigid transformation, the degree of freedom is relatively low and the techniques for rigid image registration are becoming mature. Methods based on the maximization of mutual information [1,2] and normalized mutual information [3] are widely used in rigid medical image registration. A recent overview and more applications of the mutual-information-based medical image registration methods can be found in [4]. In contrast, non-rigid image registration is an ill-posed problem due to its supernormal high degree of freedom and inherent requirement of smoothness in the deformation field. The topic is still under active research because there are many useful applications including (1) construction of atlas across a population or a specific patient group [5], (2) recovery of deformation field during surgery with respect to a pre-treatment image [6], (3) analysis of anatomical variations across patients or during growth [7], and (4) atlas-based segmentation [8]. Overviews of the non-rigid image registration methods can be found in [9–12].

1.1. Basics for image registration problem

The task of image registration is to find spatial correspondences between two images, I and J. This is usually done by finding a transformation T such that I and $T(J)$ are spatially matched, according to an image-to-image dissimilarity measure, $C(I,T(J))$. I and J are referred as the source image and the floating image, respectively, and $T(J)$ refers to the resultant image after applying T to J. Mathematically, the registration problem can be defined as finding the optimal transformation T^* such that

$$T^* = \operatorname*{argmin}_T C(I,T(J)). \tag{1}$$

Note that a minimization is used since we assume $C(\cdot,\cdot)$ increases with the degree of dissimilarity.

For non-rigid image registration, there is a variety of methods proposed in the literature regarding how the transformation T should be modeled. Some parametric models restrict T to be of

low degree of freedom, such as affine, polyaffine [13] or control-points interpolated deformation [14] models. These models intrinsically constrain T to be smooth or elastic. They are usually capable of representing an intra-subject deformation across time because, in this case, there is a real physical underlying deformation between the images and this deformation is usually governed by physical properties such as tissue elasticity.

However, in the case of inter-subject image registration, there is no underlying physical deformation between the images and thus the registration result is just a transformation that restores the spatial correspondences of anatomical structures. Since anatomical structures can vary significantly across different subjects both geometrically and topologically, a transformation of low degree of freedom may not have enough flexibility to represent these complex structural variations. Therefore, in principal, any hard constraints on the domain of T should not be imposed. However, this will naturally create another problem. Recall that, in Eq. (1), $C(I,J(T))$ is optimized without considering the behavior of the transformation T. With the domain of T being unrestricted, T can map any points in J to any points in I and the mapping results of two adjacent points can be totally uncorrelated. One can imagine that if either sum of squared difference (**SSD**) or sum of absolute differences (**SAD**) is used as the dissimilarity function C, the optimization is equivalent to finding the transformation that brings any point p in J to a point in I, which has the nearest intensity to p's. This registration result can be less useful and unrealistic if the recovered deformation field is not smooth. Thus, T needs regularization by adding a penalizing function $S(T)$ to penalize those transformations T, which are unrealistic. By modifying Eq. (1), we obtain

$$T^* = \underset{T}{\arg\min} C(I,T(J)) + \lambda S(T), \tag{2}$$

where λ is a positive constant that controls the level of penalty for unrealistic transformation T. If T is considered as a displacement vector field, which translates any point from its original position by a vector, an unrealistic T usually means a non-smooth displacement vector field. Therefore, S is often referred as the smoothness function. In practice, the integrated magnitude of different derivatives is usually used as a criterion of smoothness. Since images I and J are usually acquired from the same imaging modality in the non-rigid image registration tasks, it is adequate to use **SSD** or **SAD** as the dissimilarity function C.

1.2. Graph-cuts method

Graph-cuts method is a powerful combinatorial optimization tool for solving energy minimization problems. Using two-dimensional multi-labels with graph-cuts method has been successfully applied into image processing and computer vision problems, such as image segmentation [15], visual correspondence for stereo, motion and image restoration [16,17], and object classification [18]. Unlike general purpose techniques such as simulated annealing, the graph-cuts method yields either a global minimum or a local minimum in a strong sense in polynomial time, under some specific conditions [16,19]. Summaries of such conditions and details of graph constructions can be found in [16,19].

In general, the graph-cuts method is used to solve discrete labeling problems by minimizing energy function E_f in the following standard form [16,19]:

$$E_f = \sum_{p \in \mathcal{P}} D_p(f_p) + \sum_{(p,q) \in \mathcal{N}} V_{p,q}(f_p,f_q). \tag{3}$$

In the above Eq. (3), $\mathcal{P}$ is the set of voxels, $\mathcal{N}$ is a neighborhood system defined on $\mathcal{P}$, $f : \mathcal{P} \rightarrow L$ is a labeling function, where L is a set of labels, and $f_i \in L$ is the label of a voxel i in f. The term $D_p(f_p)$ measures the penalty of assigning label (f_p) to voxel p and the term $V_{p,q}(f_p,f_q)$ measures the penalty of assigning labels (f_p) and

(f_q) to the neighborhood voxels p and q, respectively. The first summation is usually referred as the data term, since it is usually calculated from the observed data. The second summation is usually referred as the smoothness term, as it penalizes the discontinuity of labels across adjacent voxels.

1.3. The proposed method

In this paper, based on our prior work presented in MICCAI 2007 [20], we formulate the non-rigid medical image registration framework as a discrete three-dimensional labeling problem in which the energy system can be optimized by using the graph-cuts method. During the non-rigid registration process, each pixel in the source image is assigned a three-dimensional displacement label, which represents a displacement vector indicating its corresponding position in the floating image, according to a similarity measure. A smoothness constraint based on the first derivative is used to penalize sharp changes in the displacement labels across pixels. As such, the framework consists of two energy terms and they are based on voxel intensity similarity and smoothness of the displacement vector field. It is shown in this paper that the whole energy system is submodular and can be optimized by using the graph-cuts method via α-expansions [16]. The optimization process with the graph-cuts method is not easily trapped in local minima and the solution is guaranteed to be within a known factor of the exact minimum. This can improve the registration robustness and accuracy of the proposed approach significantly, as demonstrated though extensive experiments on magnetic resonance image volumes in this paper.

1.4. Outline

The remaining sections will be organized as follows. In Section 2, two state-of-the-art techniques and a linear programming based method for non-rigid image registration will be described. In Section 3, details about our proposed approach, including methodology and the underlying theory, will be sketched. Results of the experiments on brain magnetic resonance images obtained from the Simulated Brain Database and interpretations of the results obtained by using the two state-of-the-art methods, the linear programming based method and the proposed method will be provided in Section 4. Finally, a conclusion will be given in Section 5.

2. Related works on non-rigid image registration

This section describes two state-of-the-art and widely used non-rigid image registration methods: free-form deformations based method (**FFD**) [14], Demons based method (**DEMONS**) [21], and also a method which also formulates the non-rigid image registration problem as a discrete labeling problem: linear programming based method (**LP**) [22]. In addition, along with the closely related works, a feature based non-rigid image registration method, namely hierarchical attribute matching mechanism for elastic registration (**HAMMER**) [23], is also described in this section.

2.1. Free-form deformations based method (FFD)

Rueckert et al. [14] proposed a method for performing non-rigid registration on breast MR images. The deformation was modeled by a global affine transformation plus a local deformation. The local deformation was described by free-form deformation (**FFD**) based on B-splines. In **FFD**, a regular grid of control points on the image are first allowed to displace freely. The displacement of each voxel is then evaluated as a weighted sum of neighborhood control point displacements via the B-spline interpolation. Hard constraints are

imposed on the transformation T, so that its degree of freedom is limited by the number of control points. In this work, normalized mutual information (NMI) [3] was used as the dissimilarity function, sum of squared second-order derivatives was used as the smoothness function, and the optimization was done by using the iterated gradient descent of the transformation parameters. Since the degree of freedom of the local deformation is determined by the number of control points, it is important to choose whether to have a sparse or dense set of control points. However, both sparse and dense sets have limitations. If a sparse set of control points is used, the transformation may not allow flexible voxel movements within the grid box in order to represent complicated deformations. We will show this effect in the Experimental Results section. If a dense set of control points, e.g. each for a voxel, is used, the optimization can be both computationally inefficient and local-minima-sensitive since a tiny step size will be required for the gradient descent method.

2.2. Demons based method (**DEMONS**)

Thirion [21] proposed a diffusion-based approach to non-rigid image registration. No hard constraints were imposed on the transformation T, so that each voxel can have its own displacement. The method is based on the concepts of optical flow [24]. In each iteration, the movement of any voxel in the floating image is based on its local intensity gradient and its intensity difference with the source image at the same position. It will naturally guarantee a decrease in **SSD** or **SAD** by each iteration, if the movement steps are sufficiently small. Since all voxels can move freely, a Gaussian smoothing step is applied at the end of iteration in order to regularize the transformation. However, since the regularization is done at each iteration but not incorporated into the cost function, large displacements of voxels or sharp changes in the displacement field are not penalized but are just smoothed. Moreover, since the voxel motions are heavily depending on local intensity gradient, this method is highly sensitive to local artifacts. For example, if the boundaries of different anatomical structures are overlapping in the initial images, an incorrect voxel movement may be generated and thus the registration result may be problematic. This effect will be shown in the Experimental Results section.

2.3. Linear programming based method (**LP**)

Glocker et al. [25,22] formulated the non-rigid image registration problem as a discrete three-dimensional labeling problem using the Markov random field (MRF) model with linear programming optimization method. The MRF energy which contains the data term and smoothness term was optimized by the primal-dual schema of linear programming [26]. **LP** uses a regular control point grid to reduce the dimensionality on the variables and smooth the transformation T. The data term is the similarity measure projected on the control points, and the smoothness term is the magnitude difference of the vectors on the deformation field according to a neighborhood system on the control point grid. Similar to **FFD**, a set of control points are first displaced, then the displacement of each voxel is restricted based on the interpolation of the displacements of neighborhood control points. Therefore, the degree of freedom of **LP** is also limited by the number of control points. Hence, like **FFD**, it is facing a dilemma of choosing the number of control points and the trade-offs will be shown in the Experimental Results section. High space capacity is one of the reasons for **LP** to use control point grid in order to reduce dimensionality of the variables. In Section 3.3, we will analyze the space capacity of **LP** as well as our proposed method.

2.4. Hierarchical attribute matching mechanism for elastic registration (HAMMER)

Feature based methods use feature vectors that characterize the anatomical properties around each voxel as the signatures of that voxel. The registration process is then formulated as a feature matching and optimization problem. Shen et al. [23] proposed a feature based non-rigid registration method, namely hierarchical attribute matching mechanism for elastic registration (HAMMER). The intensity value of each voxel, the edge type and geometric invariant (GMI) features of the gray matter (GM), white matter (WM) and cerebrospinal fluid (CSF) were utilized by HAMMER to form a feature vector for each voxel. Therefore, the HAMMER algorithm requires the input images that have already been segmented into three classes of tissues (i.e., GM, WM and CSF) before registration. Then HAMMER formulates the registration process in a hierarchical manner. A set of voxels with more salient feature vectors are selected as active points from both source and floating images. These voxels are first used to drive the registration in the initial stage to estimate the transformation. During the registration process, more and more active points are added for consideration to refine the estimated transformation. At the last stage, all voxels are considered as active points.

3. Theory and methodology

3.1. Formulation of the energy function

Let I and J be, respectively, the source image and the floating image of dimension d. In this paper, we only consider I and J of the same dimension d. Without loss of generality, let $\mathbf{X} = [0,1]^d$ be the continuous domain for both images. For any spatial point $\mathbf{x} = (x_1, x_2, \ldots, x_d) \in \mathbf{X}$, $I(\mathbf{x})$ and $J(\mathbf{x})$ are the intensity values (or feature vectors in general) at $\mathbf{x}$ for both images. Consider again Eq. (2), which is the most general form of function to be optimized in the non-rigid image registration problems. $T(J)$ represents the transformed floating image. In the proposed formulation, T is represented by a displacement vector field $\mathbf{D}$ that displaces every point $\mathbf{x}$ in J away from its original position by the displacement vector $\mathbf{D}(\mathbf{x}) \in \mathbb{R}^d$ to the new point $\mathbf{x} + \mathbf{D}(\mathbf{x})$. In principle, $T(J)$ cannot be obtained, but instead we can obtain a pullback from I to J, $T^{-1}(I)$, that can be obtained by $I(\mathbf{X} + \mathbf{D})$. Therefore, in order to obtain the transformed floating image as the final product, we need to formulate $\mathbf{D}$ as the displacement vector field that displaces points in I. By modifying Eq. (2), we can get

$$\mathbf{D}^* = \operatorname*{argmin}_{\mathbf{D}} C(I(\mathbf{X}), J(\mathbf{X} + \mathbf{D})) + \lambda S(\mathbf{D}). \tag{4}$$

In general, C can be any image-to-image dissimilarity measure. Here, we tighten C as an integral of point-to-point dissimilarity measure, among the whole image domain. We further tighten this point-to-point dissimilarity measure to be an absolute difference, $|I(\mathbf{x}) - J(\mathbf{x} + \mathbf{D}(\mathbf{x}))|$, where $|\cdot|$ is the L1-norm operator. Note that the second restriction is not a strong one as $I(\mathbf{x})$ and $J(\mathbf{x} + \mathbf{D}(\mathbf{x}))$ can be image feature vectors in general. Furthermore, we use the first derivative terms as the smoothness function. It yields

$$\mathbf{D}^* = \operatorname*{argmin}_{\mathbf{D}} \int_{\mathbf{X}} |I(\mathbf{x}) - J(\mathbf{x} + \mathbf{D}(\mathbf{x}))| d\mathbf{X} + \lambda \sum_{i=1}^{d} \int_{\mathbf{X}} \|\mathbf{D}_{(x_i)}\| \, d\mathbf{X}, \tag{5}$$

where $\mathbf{D}_{(x_i)}$ is the first derivative of $\mathbf{D}$ along direction x_i and the differential element $d\mathbf{X} = dx_1 dx_2 \ldots dx_d$, and $\|\cdot\|$ is the L2-norm operator. Since everything is in the continuous domain, $\mathbf{D}$ can have infinite degree of freedom theoretically. Here, we introduce the first discretization step, by discretizing $\mathbf{X}$ into voxels. This is a natural discretizing step as the images are usually acquired in a

discretized form. For simplicity, although not limited to, we now also restrict our analysis in three-dimensional ($d = 3$) cases. Therefore, $\mathbf{X}$ becomes $\{1,2,\ldots,N_x\} \times \{1,2,\ldots,N_y\} \times \{1,2,\ldots,N_z\}$, where N_x, N_y and N_z are the spatial resolutions of the 3D images. By replacing all integrals by summations and all derivatives by finite differences, Eq. (5) now becomes

$$
\begin{aligned}
\mathbf{D}^* = \arg\min_{\mathbf{D}} &\sum_{i=1}^{N_x}\sum_{j=1}^{N_y}\sum_{k=1}^{N_z} |I(i,j,k) - J([i,j,k] + \mathbf{D}(i,j,k))| \\
&+ \lambda \sum_{i=1}^{N_x-1}\sum_{j=1}^{N_y}\sum_{k=1}^{N_z} \|\mathbf{D}(i,j,k) - \mathbf{D}(i+1,j,k)\| \\
&+ \lambda \sum_{i=1}^{N_x}\sum_{j=1}^{N_y-1}\sum_{k=1}^{N_z} \|\mathbf{D}(i,j,k) - \mathbf{D}(i,j+1,k)\| \\
&+ \lambda \sum_{i=1}^{N_x}\sum_{j=1}^{N_y}\sum_{k=1}^{N_z-1} \|\mathbf{D}(i,j,k) - \mathbf{D}(i,j,k+1)\|.
\end{aligned}
\tag{6}
$$

Note that at this stage, $\mathbf{D}(i,j,k) \in \mathbb{R}^3$ is still not discretized. Therefore, $[i,j,k] + \mathbf{D}(i,j,k)$ in Eq. (6) can be a non-integer valued vector, and $J([i,j,k] + \mathbf{D}(i,j,k))$ needs to be computed using an interpolation function. Also, when $[i,j,k] + \mathbf{D}(i,j,k)$ is outside the image domain, a pre-assigned background intensity value can be used. In Eq. (6), the first term is contributed by the differences between the image I and the transformed J. The last three terms are contributed only by the internal properties of the displacement vector field $\mathbf{D}$. Specifically, the second to fourth terms penalize the non-smoothness of $\mathbf{D}$ across the image in different spatial directions.

In principle, Eq. (6) can be optimized by any iterative optimization tools instead of using the graph-cuts method. However, in practice, the degree of freedom of $\mathbf{D}$ can still be as high as a billion since it is proportional to the number of voxels in the image volume. First, it may cost a huge amount of time for the optimization process. Second, since $\mathbf{D}$ has value in each voxel position, it is a requirement that the step size of updating D is sufficiently small in each iteration in order to ensure a smooth field. Not only adding an extra time cost, this makes the optimization process highly sensitive to local minima. Yet, Eq. (6) is still not solvable by the graph-cuts method without modifications. It is addressed in the next subsection.

3.2. Optimization via graph-cuts

Comparing with the standard form of function solvable by the graph-cuts method in Eq. (3), it is not difficult to observe that our current energy function in Eq. (6) is already in that form. If we consider a six-connected neighborhood system $\mathcal{N}$ in the 3D space, i.e., $(\mathbf{x},\mathbf{y}) \in \mathcal{N}$ iff $\mathbf{x},\mathbf{y}$ are adjacent voxels, the last three terms in Eq. (6) can be grouped together into one term, resulting the following form as in Eq. (3).

$$
\mathbf{D}^* = \arg\min_{\mathbf{D}} E_f(\mathbf{D})
$$

$$
\text{for } E_f(\mathbf{D}) = \sum_{\mathbf{x} \in \mathbf{X}} |I(\mathbf{x}) - J(\mathbf{x} + \mathbf{D}(\mathbf{x}))| + \lambda \sum_{(\mathbf{x},\mathbf{y}) \in \mathcal{N}} \|\mathbf{D}(\mathbf{x}) - \mathbf{D}(\mathbf{y})\|.
\tag{7}
$$

To convert this optimization problem to a discrete labeling problem, $\mathbf{D}(\mathbf{x}) \in \mathbb{R}^3$ should be limited into a finite set. Here, we perform the second discretization step. Also acting as a restriction of how far a voxel can be displaced, a discretized 3D window

$$
\mathcal{W} = \{0, \pm s, \pm 2s, \ldots, \pm ws\}^3
\tag{8}
$$

is chosen such that $\mathbf{D}(\mathbf{x}) \in \mathcal{W}$. Note that $\mathcal{W}$ is the discretization of the continuous 3D region $[-ws, ws]^3$ with sampling period s along all directions. Also, if $s < 1$, displacements with sub-voxel units can be considered. Now, by using $\mathcal{W}$ as the set of labels that every $\mathbf{D}(\mathbf{x})$ can be assigned, the optimization in Eq. (7) can readily be solved by

using graph-cuts via a sequence of α-expansion (α-expansion) moves [16]. This procedure can be summarized by Algorithm 1.

Given the current labeling f for the set of voxels $\mathcal{P}$ and a new label α, an α-expansion move means: For any voxel $p \in \mathcal{P}$, it is considered either keeping its current label (f_p) or changing its label to α in the next labeling f'. Obviously, an α-expansion move is a two-label problem, with label 0 meaning $f'_p = f_p$ and label 1 meaning $f'_p = \alpha$. Kolmogorov and Zabih [19] show that the graph-cuts method can find the exact minimum of a two-label problem if every energy term $V_{p,q}$ in Eq. (3) satisfies the following inequality:

$$
V_{p,q}(0,0) + V_{p,q}(1,1) \leq V_{p,q}(0,1) + V_{p,q}(1,0),
\tag{9}
$$

or equivalently, the energy terms are submodular [19,27].

We now show that any expansion move of our formulation satisfies Eq. (9). Given a current labeling f and two adjacent voxels $\mathbf{p},\mathbf{q}$ with $f_\mathbf{p} = \beta$ and $f_\mathbf{q} = \gamma$, where $\beta,\gamma \in \mathcal{W}$, an expansion move of new label $\alpha \in \mathcal{W}$ is considered.

- $V_{\mathbf{p},\mathbf{q}}(0,0)$ is the cost when both $\mathbf{p},\mathbf{q}$ choose their old labels β,γ, $\therefore V_{\mathbf{p},\mathbf{q}}(0,0) = \|\beta - \gamma\|$.
- $V_{\mathbf{p},\mathbf{q}}(1,1)$ is the cost when both $\mathbf{p},\mathbf{q}$ choose the new label α, $\therefore V_{\mathbf{p},\mathbf{q}}(1,1) = \|\alpha - \alpha\| = \mathbf{0}$.
- $V_{\mathbf{p},\mathbf{q}}(0,1)$ is the cost when $\mathbf{p}$ retains its old label β and $\mathbf{q}$ chooses the new label α, $\therefore V_{\mathbf{p},\mathbf{q}}(0,1) = \|\beta - \alpha\|$.
- $V_{\mathbf{p},\mathbf{q}}(1,0)$ is the cost when $\mathbf{p}$ chooses the new label α and $\mathbf{q}$ retains its old label γ, $\therefore V_{\mathbf{p},\mathbf{q}}(1,0) = \|\alpha - \gamma\|$.

Since $\alpha,\beta,\gamma \in \mathcal{W} \subset \mathbb{R}^3$ and $\|\cdot\|$ is the L2-norm operator, by the triangle inequality, we have $\|\beta - \gamma\| \leq \|\beta - \alpha\| + \|\alpha - \gamma\|$ for any vectors α,β,γ. Thus, the inequality in Eq. (9) is satisfied for any adjacent voxels $\mathbf{p},\mathbf{q}$ and each of our α-expansion move is globally optimal. Boykov et al. [16] have further proved that, in such a case, the α-expansion algorithm can finally converge to a local minimum, which is within a guaranteed factor of the exact minimum.

Despite the high dimensionality of the transformation model, we have shown that the proposed energy function (Eq. (7)) is submodular and can be optimized by using the graph-cuts method with α-expansion moves, by which a strong local minimum can be guaranteed. Although some of the ideas are illustrated in the three-dimensional space, the proposed method can be naturally extended into the higher dimensional spaces.

Algorithm 1. α-expansion move algorithm of graph-cuts.

INPUT: Start with an arbitrary displacement vector field $\mathbf{D}$ and an energy function E_f
OUTPUT: An optimal displacement vector field $\mathbf{D}^*$
1: $\mathbf{D}^* \Leftarrow \mathbf{D}$
2: **for** $i = 1$ **to** number of α-expansion **do** // Each loop here stands for one α-expansion
3:    **for** each label $\alpha \in \mathcal{W}$ **do** // Each loop here stands for one move in an α-expansion
4:       Minimize $E_f(\mathbf{D})$ with respect to $\mathbf{D}$ and α by graph-cut (Please refer to [16] for more information).
5:       **if** $E_f(\mathbf{D}^*) > E_f(\mathbf{D})$ **then**
6:          $\mathbf{D}^* \Leftarrow \mathbf{D}$
7:       **end if**
8:    **end for**
9: **end for**
10: **return** $\mathbf{D}^*$

3.3. Space capacity analysis

For the proposed method and the linear programming based method (**LP**) [22], which model the non-rigid image registration

problem as multi-labeling problem with Markov random field based optimization, the space capacity of the optimizer is a key factor that affects the performance and generalization of the registration process. In this subsection, we analyze the space capacity of different optimizers. The optimizers of **LP** and the proposed method are the linear programming algorithm and the graph-cuts algorithm, respectively. The space capacity of the linear programming and graph-cuts algorithms in the non-rigid image registration problem is proportional to the memory required to store the optimization results. In the graph-cuts algorithm, at least $\lceil \log_2 |\mathcal{W}| \rceil$ bits are required at each voxel in the floating image for storing an assigned label, where $|\mathcal{W}|$ is

the number of labels. Note that, for each voxel, one iteration of α-expansion (i.e. one α-expansion move) just needs 1 bit to store a binary label which is indicating whether the voxel takes the move or not. Therefore, the space capacity of the graph-cuts method is $O(|\mathcal{P}| \times \lceil \log_2 |\mathcal{W}| \rceil)$, where $|\mathcal{P}|$ is the number of voxels. For the linear programming algorithm, the label assigned to a control point is indicated by a $|\mathcal{W}|$-dimension indicator vector with only one non-zero element which is equal to one. Therefore, the space capacity of **v** is $O(|\mathcal{W}|)$, and hence the space capacity of the linear programming algorithm is $O(|\mathcal{G}| \times |\mathcal{W}|)$, where $|\mathcal{G}|$ is the number of control points. If all voxels in the floating image are control points, i.e. $|\mathcal{G}| = |\mathcal{P}|$, the

Fig. 1. (a) An MR slice obtained from the simulated Brain Database, (b)–(d) show the ground truths of (b) white matter (WM) segmentation, (c) gray matter (GM) segmentation, and (d) cerebrospinal fluid (CSF) segmentation.
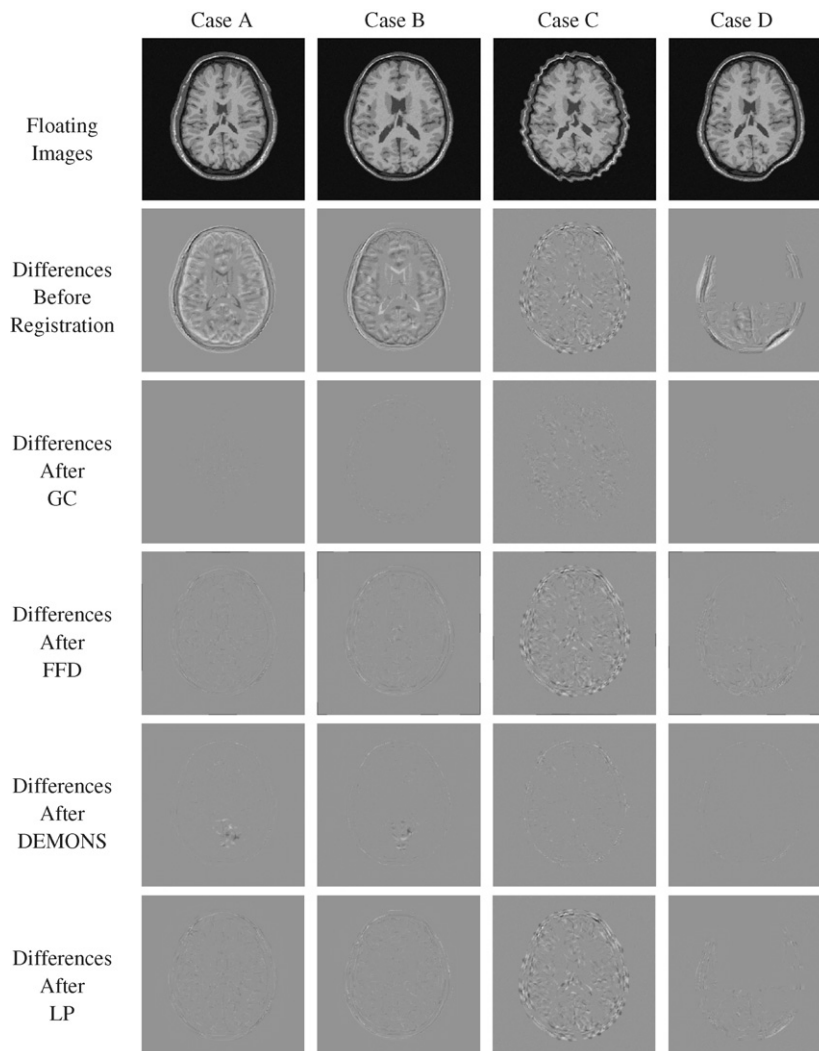

Fig. 2. The first row shows the floating images for the non-rigid registration experiments. Before registration, the differences between the source image and the floating images are shown in the second row. After registration, the differences between the source image and the deformed floating images obtained by using the proposed method (**GC**), **FFD**, **DEMONS** and **LP** are shown in the third, fourth, fifth and sixth rows, respectively.

space capacity of the linear programming algorithm is higher than that of the graph-cuts algorithm. Take a 3D non-rigid image registration as an example and assume each voxel is a control point. If the size of the floating image is $256 \times 256 \times 181$ and the number of labels is 1000, then the linear programming algorithm requires at least $(256 \times 256 \times 181 \times 1000)$ bits ≈ 1414.06 MBytes while the graph-cuts algorithm requires at least $(256 \times 256 \times 181 \times \lceil \log_2 1000 \rceil)$ bits ≈ 14.14 MBytes only. Due to the huge space capacity of the linear programming algorithm, **LP** cannot afford to have a dense set of control points or a large number of labels. This restriction causes **LP** unable to model registration cases with complicated deformation. It will be demonstrated in the Experimental Results section. In this aspect, the proposed method outperforms **LP**.

4. Experimental results

In this section, we compare the registration performance of the proposed method with the two state-of-the-art and a related image registration methods that have been introduced in Section 2, namely, the free-form deformation based method, the demons based method and the linear programming based method. Since hierarchical attribute matching mechanism for elastic registration

(HAMMER) and the proposed method are classified in different categories, the comparison between HAMMER and the proposed method is not included in this paper. For simplicity, these three methods are denoted as **FFD**, **DEMONS** and **LP**, respectively, in this section. In the meanwhile, **GC** is adopted to represent the proposed method in all figures and tables. Section 4.1 gives the details of the image datasets and values of the parameters for non-rigid registration tests. Sections 4.2–4.4 present the comparison results on registration robustness, registration accuracy, and smoothness of the recovered deformation fields. For the proposed method, the choice of window size $\mathcal{W}$ in Eq. (8) is discussed in Section 4.6. Additionally, we compare the proposed method with **LP** at three different grid sizes by seven sets of labels in Section 4.7. Finally, computation times for **FFD**, **DEMONS**, **LP** and the proposed method are given in Section 4.8.

4.1. Image datasets and parameters

In this paper, all the T1-weighted magnetic resonance (MR) image volumes and segmentation ground truths of the non-rigid registration experiments were obtained from the Simulated Brain Database [28,29]. Some example slices are shown in Fig. 1, in

Fig. 3. The first row shows the floating images for the non-rigid registration experiments. Before registration, the differences between the source image and the floating images are shown in the second row. After registration, the differences between the source image and the deformed floating images obtained by using the proposed method (**GC**), **FFD**, **DEMONS** and **LP** are shown in the third, fifth and sixth rows, respectively.

which the ground truths of the white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF) segmentations are shown. According to [28,29], all image volumes have been linearly registered using nine parameters to the same frame of reference. Therefore, in this paper, only the local part of the deformations is investigated in the experiments. Also, according to [28,29], the datasets in the Simulated Brain Database were constructed from the MR scans of different normal adults and the image volumes from the database have been intensity-normalized. All the image data used in the experiments were in unsigned-byte format so that the intensity values were within a range between 0 and 255.

For **DEMONS**, we used the implementations (version 2.8) obtained from the Insight Segmentation and Registration Toolkit (ITK) [30]. We set $\sigma = 1$, which is the standard deviation of the Gaussian kernel that is used for the deformation field smoothing after each iteration. According to the ITK implementation [30], the Gaussian kernel is approximated by a 30mm-width window with a maximum error of 0.1. All experiments were run for 500 iterations.

It is a reasonable number as the percentage change in **SSD** was less than 0.1% in the 501-st iteration in all the experiments.

For **FFD**, we used the implementations developed by Rueckert et al. [14,31], which are available at [32]. We chose 10 mm control point spacing as it was demonstrated in [14] that a control point spacing of 10 mm gave better performance than 15 or 20 mm. Smoothness penalty was dropped in the experiments since it is insignificant for a sparse set of control points. The sequence of step sizes we used were 5, 2.5, 1.25, and 0.625 mm. We run and repeated the experiments for 20 iterations for each step size. Normalized mutual information was used as the similarity measure, according to the implementation described in [14].

For **LP**, we used the implementations developed by Glocker et al. [25,22], which are available at [33]. Same as FFD, we chose 10 mm control point spacing in both 2D and 3D experiments. In order to compare with the proposed method, **SAD** was utilized as the similarity measure. Dense sampling was chosen for the "label set setting" with 30 steps in 2D experiments. Therefore, the number of

Table 1
Distributions of the absolute intensity errors within different tissue regions after performing **FFD**, **DEMONS**, **LP** and the proposed method (**GC**). Distributions are listed in the format MEAN ± SD.

Case	Tissue class	Distributions of absolute intensity errors (Mean ± SD)				
		Before registration	FFD	DEMONS	LP	GC
Case A	WM	16.66 ± 17.18	3.34 ± 3.45	3.69 ± 5.46	3.71 ± 3.87	1.07 ± 3.05
	GM	25.16 ± 25.12	3.94 ± 4.21	3.74 ± 4.58	4.31 ± 4.72	0.90 ± 3.16
	CSF	32.34 ± 24.19	5.73 ± 5.76	4.44 ± 7.35	6.09 ± 6.20	1.14 ± 3.89
	Whole image	12.35 ± 22.04	3.01 ± 4.33	2.33 ± 4.04	2.66 ± 5.00	0.35 ± 1.86
Case B	WM	13.57 ± 16.26	2.99 ± 3.49	2.56 ± 4.28	3.35 ± 3.62	0.35 ± 1.63
	GM	17.69 ± 14.66	3.84 ± 4.34	2.96 ± 4.62	4.14 ± 4.53	0.62 ± 2.50
	CSF	31.60 ± 25.13	5.91 ± 6.17	3.26 ± 6.66	5.93 ± 6.44	1.32 ± 4.89
	Whole image	9.73 ± 18.39	3.06 ± 4.95	1.87 ± 4.13	2.62 ± 5.27	0.47 ± 2.82
Case C	WM	8.09 ± 10.21	7.60 ± 9.10	3.00 ± 3.79	7.41 ± 9.18	5.87 ± 6.65
	GM	11.78 ± 13.13	10.92 ± 11.49	2.90 ± 3.82	10.59 ± 11.20	6.28 ± 7.81
	CSF	18.67 ± 19.04	16.94 ± 16.41	3.81 ± 5.09	16.87 ± 16.46	5.11 ± 8.21
	Whole image	8.87 ± 15.54	8.14 ± 13.63	2.98 ± 4.98	7.88 ± 13.36	4.04 ± 5.71
Case D	WM	3.91 ± 7.95	2.64 ± 3.38	1.28 ± 1.57	1.55 ± 2.99	0.24 ± 1.44
	GM	6.33 ± 13.41	3.43 ± 4.55	1.24 ± 1.69	2.11 ± 3.82	0.33 ± 1.99
	CSF	9.78 ± 17.18	5.24 ± 6.60	1.78 ± 2.45	3.28 ± 5.74	0.49 ± 3.05
	Whole image	6.28 ± 16.41	3.07 ± 4.71	1.57 ± 2.69	2.13 ± 4.66	0.62 ± 2.32
Case E	WM	12.25 ± 17.11	4.16 ± 4.57	2.83 ± 4.13	3.67 ± 4.17	1.29 ± 3.21
	GM	14.49 ± 19.50	5.11 ± 6.00	2.66 ± 4.44	3.97 ± 5.01	1.00 ± 3.33
	CSF	23.96 ± 24.92	7.23 ± 8.16	3.49 ± 7.18	5.58 ± 6.65	1.32 ± 4.26
	Whole image	8.84 ± 19.01	3.52 ± 5.30	1.93 ± 3.62	2.54 ± 4.68	0.46 ± 2.19
Case F	WM	9.01 ± 13.68	3.88 ± 4.13	1.42 ± 1.87	2.64 ± 3.38	0.43 ± 1.92
	GM	14.87 ± 19.99	5.09 ± 5.75	1.39 ± 1.92	3.50 ± 4.54	0.59 ± 2.46
	CSF	21.32 ± 20.33	7.05 ± 7.54	1.75 ± 2.43	5.42 ± 6.75	0.90 ± 3.82
	Whole image	12.57 ± 23.73	4.06 ± 6.00	1.68 ± 3.29	3.16 ± 5.69	0.77 ± 2.88
Case G	WM	6.28 ± 15.13	2.52 ± 3.54	2.44 ± 8.94	1.40 ± 2.83	0.36 ± 1.79
	GM	9.13 ± 19.58	3.01 ± 4.29	4.36 ± 14.29	1.89 ± 3.89	0.41 ± 1.97
	CSF	8.69 ± 18.01	3.64 ± 5.15	3.53 ± 8.00	2.25 ± 5.05	0.42 ± 2.50
	Whole image	5.76 ± 17.09	2.23 ± 4.18	3.28 ± 12.09	1.48 ± 4.17	0.50 ± 2.24
Case H	WM	14.97 ± 20.52	10.36 ± 12.70	3.28 ± 6.02	7.84 ± 11.59	4.00 ± 6.47
	GM	19.29 ± 21.94	12.87 ± 13.03	3.22 ± 5.23	9.36 ± 11.71	4.03 ± 7.45
	CSF	28.36 ± 25.21	19.57 ± 18.45	4.15 ± 5.86	14.35 ± 16.97	4.33 ± 8.60
	Whole image	14.73 ± 25.90	8.45 ± 13.26	3.86 ± 10.84	6.16 ± 11.60	2.49 ± 5.62
All cases	WM	10.59 ± 15.80	4.69 ± 6.94	2.56 ± 5.08	3.95 ± 5.20	1.70 ± 4.30
	GM	14.84 ± 19.66	6.03 ± 8.25	2.81 ± 6.35	4.98 ± 6.18	1.77 ± 4.89
	CSF	21.84 ± 23.58	8.91 ± 11.83	3.28 ± 6.05	7.47 ± 8.78	1.88 ± 5.60
	Whole image	9.89 ± 20.28	4.44 ± 8.29	2.44 ± 6.69	3.58 ± 6.80	1.21 ± 3.73
All cases excluding cases C, G & H	WM	11.08 ± 15.47	3.40 ± 3.87	2.36 ± 3.88	2.98 ± 3.60	0.68 ± 2.40
	GM	15.71 ± 19.94	4.28 ± 5.07	2.40 ± 3.82	3.61 ± 4.52	0.69 ± 2.74
	CSF	23.80 ± 24.01	6.23 ± 6.95	2.94 ± 5.78	5.26 ± 6.36	1.04 ± 4.04
	Whole image	9.95 ± 20.22	3.34 ± 5.11	1.87 ± 3.60	2.62 ± 5.06	0.53 ± 2.45

labels used in the 2D experiments was $31 \times 31 = 961$. For 3D experiments, using the maximum physical memory available in a 32-bits Windows system with 3 GBytes RAM and as suggested by Glocker et al. in the implementation user guide [33], we chose sparse sampling for "label set setting" with 14 steps along each sampling direction. This can prevent the experiments from being terminated unexpectedly because of the lack of memory.

For the proposed method, we used $\lambda = 0.02 \times 255$ in Eq. (7) for 3D tests and 0.05×255 for 2D tests; and $\mathcal{W} = \{0, \pm 1, \pm 2, \ldots, \pm 7\}^3$ in Eq. (8) for 3D tests and $\{0, \pm 1, \pm 2, \ldots, \pm 15\}^2$ for 2D tests. As such, the displacement label of a voxel was chosen from a $15 \times 15 \times 15$ window (for 3D tests) and a 31×31 window (for 2D

tests) centered at that voxel. We will explain how the window size was chosen in Section 4.6. For the graph-cuts algorithm, we used the source codes provided by Kolmogorov and Zabih [19]. We repeated the α-expansion three times such that each label was examined three times in the graph-cuts algorithm.

4.2. Registration robustness

To test the robustness of **FFD**, **DEMONS**, **LP** and the proposed method, we investigate into eight different registration cases, which contain eight different two-dimensional artificial deformations. These artificial deformations can resemble different
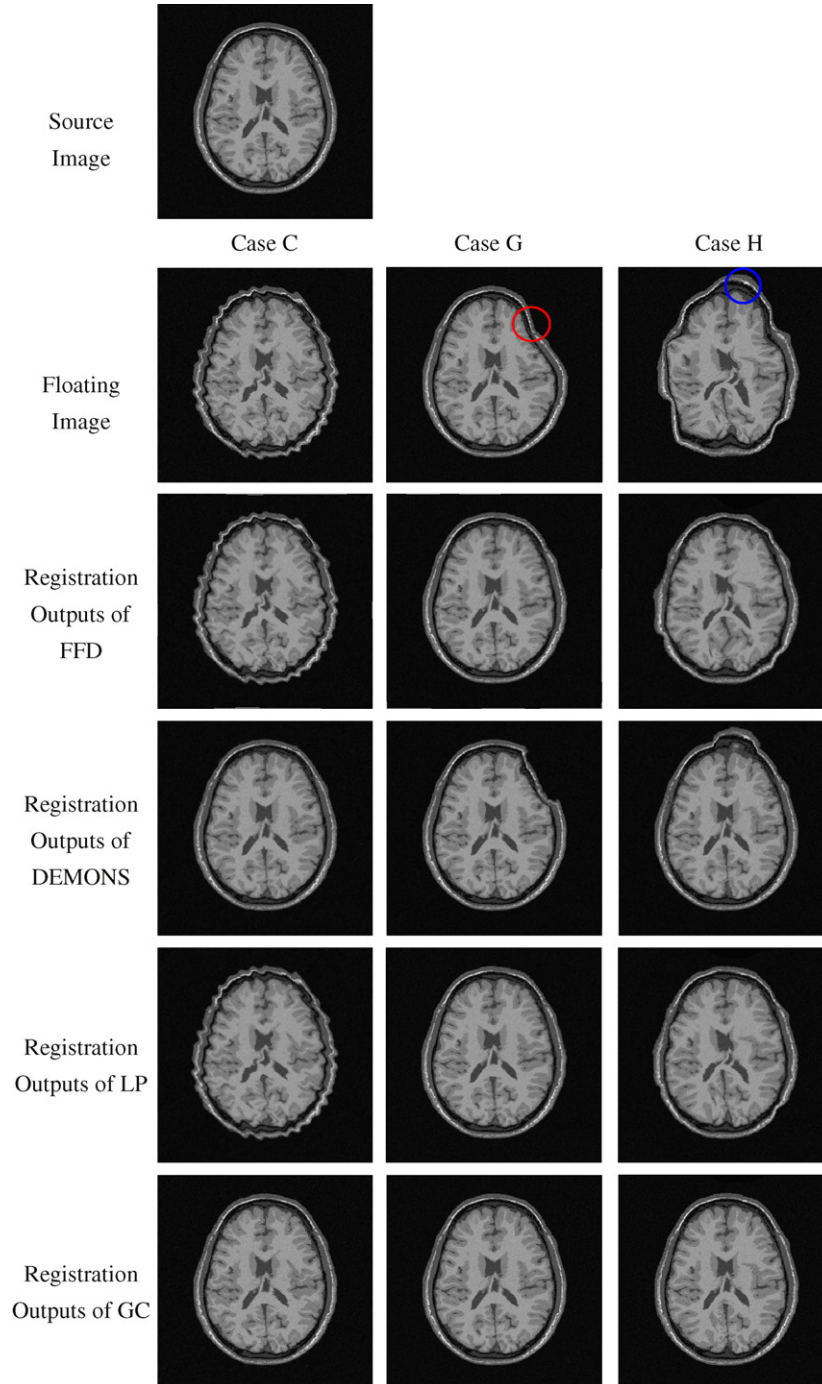

Fig. 4. *(Color images)* More details of the results obtained by using **FFD**, **DEMONS**, **LP** and the proposed method (**GC**) for Cases C, G and H. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

intra-subject and inter-subject mapping behaviors in the real world. We first used an axial slice (the 91-st (middle) slice in the image volume #04) as the source image, as shown in the left-most of Fig. 1. To the source image, eight different artificial deformations (Cases A–H) were then applied to generate eight floating images, as shown in the first rows of Figs. 2 and 3.

Cases A and B in Fig. 2 are, respectively, a squeeze-in and a squeeze-out deformations in the interior of the brain. They represent

Fig. 5. 3D registration results across two MR image volumes (#04 and #20). For each sub-figure, the middle slice across each axis of the volume is shown.

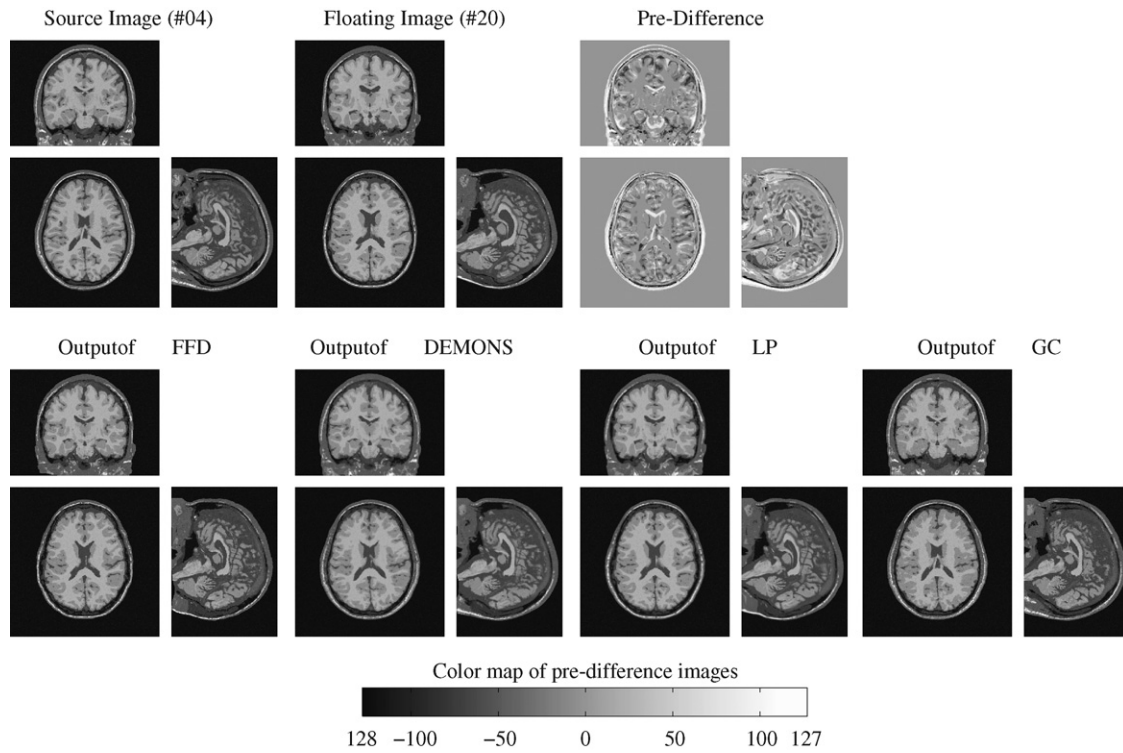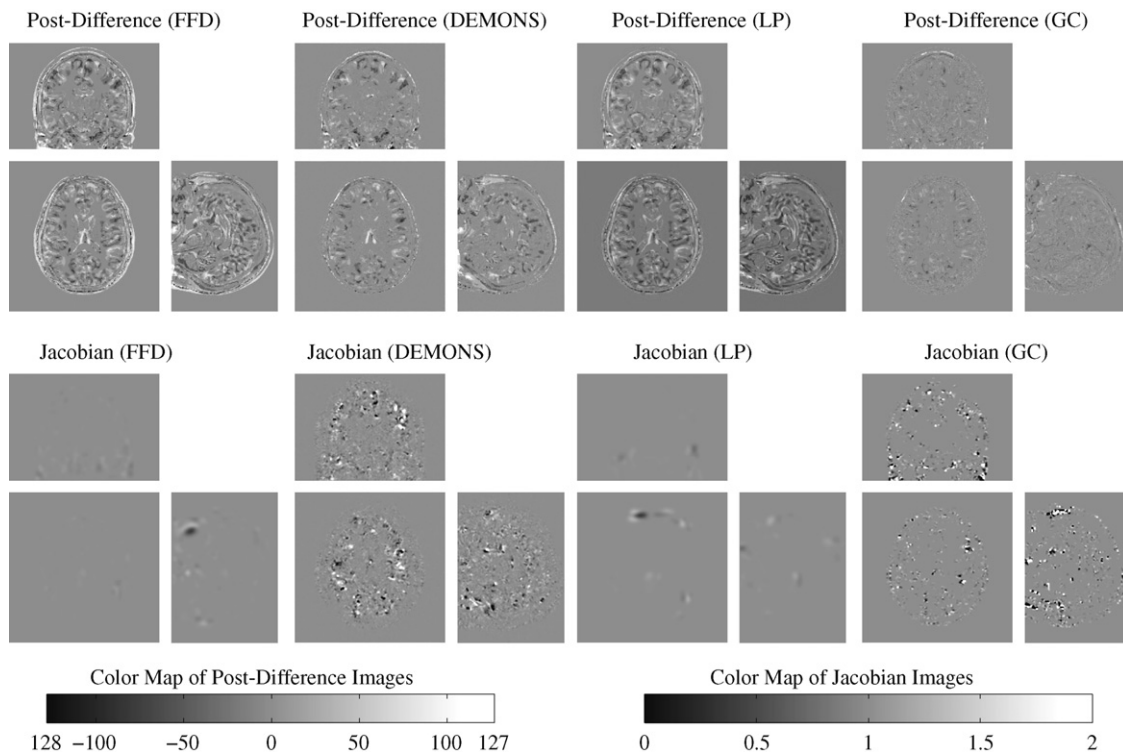**Fig. 6.** 3D registration results across two MR image volumes (#04 and #20). For each sub-figure, the middle slice across each axis of the volume is shown.

the difference in ventricle sizes across subjects or ages, as well as the structural effect of tumors on brain tissues.

Case C in Fig. 2 is a high frequency ripple deformation. It models some minor structural discrepancies across different brains, such as the positions and shapes of sulci, gyri and ventricles.

Case D in Fig. 2 and Cases E and F in 3 are low frequency small deformations. These cases model shape differences in skulls, white matter and gray matter across different subjects.

Case G in Fig. 3 is a low frequency but large deformation. It models the physical deformation of the brain during open-skull surgery or after severe injury.

Case H in Fig. 3 is a highly random deformation that aggregates the characteristics in the previous deformation cases.

Non-rigid registration experiments were performed using **FFD**, **DEMONS**, **LP** and the proposed method in all cases. The subtraction images of the registration outputs from the source image are shown in the last four rows in Figs. 2 and 3. Due to the space limitation within a page, registration outputs of four methods are not shown in the figures. For comparison, the pre-registration differences between the floating images and the source images are also shown in the second rows in the figures.

We list in Table 1 the distributions of the absolute intensity errors (calculated from the post-registration absolute subtraction images) within different tissue regions after applying the four registration methods. In Table 1, the values of the absolute intensity errors are listed in Mean $\pm$ SD format. A large standard deviation is a possible indicator of failure in the registration experiments since it means that there are specific regions of large intensity errors. From Table 1 and the post-registration subtraction images (Figs. 2 and 3), it is found that the four registration methods perform satisfactorily in all cases except Cases C, G and H. It is because all cases except Cases C, G and H have Mean < 5 and SD < 10 in the distributions within the whole images after applying the four registration methods. For further analysis and illustration, the registration outputs of Cases C, G and H after applying the four registration methods are shown in Fig. 4.

From the post-registration subtraction images in Fig. 2 (Case C, third row and last row) and Fig. 3 (Case H, third row and last row), and the registration outputs in Fig. 4 (Cases C and H, third row and second last row), it is found that there are noticeable discrepancies, especially around the skull regions, between the fixed source image and the registration outputs of **FFD** and **LP** in Cases C and H. These cases are situations, where transformation with low degree of freedom cannot model complicated deformations. The artificial deformations in these two cases contained some high frequency components. **FFD** and **LP** only allow control points to freely displace but restrict other pixel displacements to an interpolation of the displacements of neighborhood control points. Therefore, for the 2D cases, within a bounding box of four adjacent control points, the pixels must have highly correlated displacements. If the pixels within the bounding box show no correlation in the underlying deformation, **FFD** and **LP** do not displace the control points. As predicted, both **DEMONS** and the proposed method are capable of restoring the image in such cases since no hard constraints are imposed in the deformation models for **DEMONS** and the proposed method.

From the post-registration subtraction images in Fig. 3 (Cases G and H, second last row) and the registration outputs in Fig. 4 (Cases G and H, third last row), it is found that there are noticeable mis-registration results at some parts of the skull regions between the fixed source image and the registration outputs of **DEMONS** in Cases G and H. **DEMONS** converged in such configurations in Cases G and H because of the local minima in the optimization process. Since **DEMONS**-based non-rigid image registration involves transformation with high degree of freedom, it may easily result in local minima for some

optimization strategies. The deformations in Cases G and H were considered large as some points are displaced more than 10 pixel-units. To explain the mis-registrations in these two cases, we need to observe the initial alignments of the images. In Case G, some

Table 2
The values of the overlap measure obtained in the 15 3D inter-subject registration experiments. Simulated Brain Database image volume #04 was fixed as the source image throughout all the experiments and other 15 image volumes were used as the floating images. The last row shows the mean and the standard deviation (in format Mean $\pm$ SD) of the overlap measurements among all the 15 sets of registration experiments.

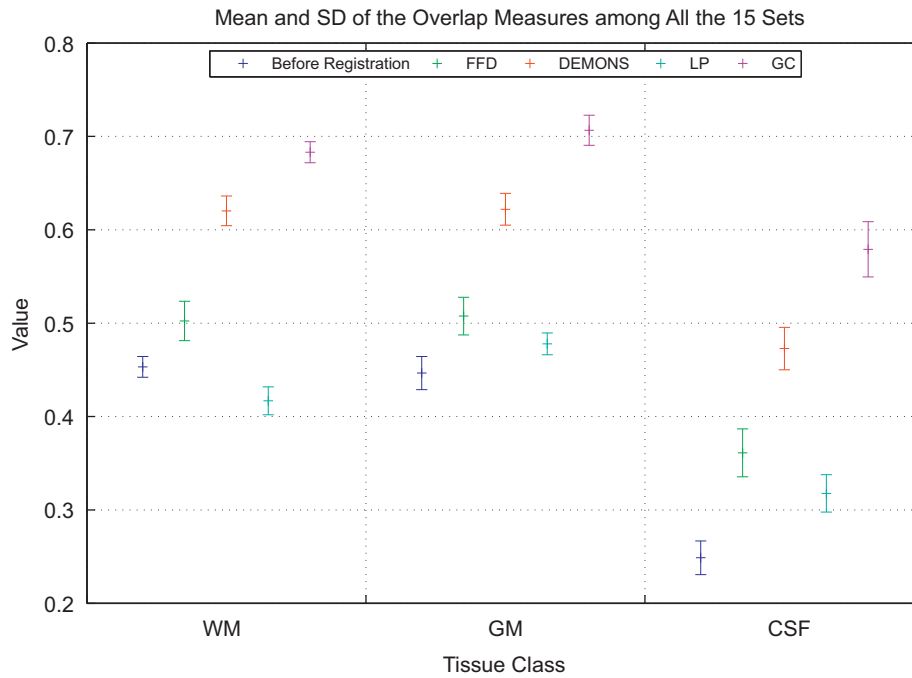| Floating image	Tissue class	The values of the overlap measure				
		Before registration	FFD	DEMONS	LP	GC
#06	WM	0.46	0.55	0.64	0.44	0.69
	GM	0.46	0.54	0.63	0.50	0.71
	CSF	0.21	0.38	0.46	0.30	0.59
#20	WM	0.45	0.50	0.62	0.41	0.69
	GM	0.46	0.52	0.64	0.48	0.72
	CSF	0.25	0.38	0.49	0.34	0.61
#38	WM	0.44	0.49	0.60	0.39	0.67
	GM	0.46	0.51	0.62	0.47	0.71
	CSF	0.26	0.35	0.46	0.31	0.57
#42	WM	0.44	0.50	0.60	0.39	0.67
	GM	0.46	0.53	0.63	0.48	0.71
	CSF	0.24	0.38	0.49	0.33	0.59
#43	WM	0.45	0.50	0.63	0.42	0.68
	GM	0.42	0.50	0.62	0.49	0.70
	CSF	0.21	0.32	0.44	0.33	0.56
#44	WM	0.43	0.46	0.60	0.41	0.68
	GM	0.43	0.49	0.62	0.47	0.70
	CSF	0.24	0.33	0.47	0.32	0.58
#45	WM	0.46	0.54	0.64	0.44	0.69
	GM	0.45	0.53	0.64	0.49	0.72
	CSF	0.25	0.40	0.51	0.35	0.63
#46	WM	0.45	0.52	0.63	0.41	0.68
	GM	0.45	0.51	0.62	0.45	0.71
	CSF	0.27	0.37	0.48	0.31	0.60
#47	WM	0.47	0.51	0.63	0.42	0.70
	GM	0.46	0.51	0.63	0.47	0.72
	CSF	0.26	0.35	0.48	0.31	0.59
#48	WM	0.44	0.50	0.60	0.43	0.66
	GM	0.43	0.50	0.58	0.48	0.69
	CSF	0.25	0.38	0.48	0.33	0.59
#49	WM	0.45	0.49	0.62	0.43	0.68
	GM	0.40	0.46	0.60	0.47	0.67
	CSF	0.24	0.35	0.46	0.32	0.54
#50	WM	0.46	0.49	0.62	0.41	0.69
	GM	0.45	0.49	0.62	0.47	0.71
	CSF	0.25	0.34	0.44	0.31	0.55
#51	WM	0.46	0.50	0.61	0.42	0.69
	GM	0.47	0.51	0.63	0.47	0.72
	CSF	0.28	0.38	0.49	0.33	0.59
#52	WM	0.47	0.51	0.64	0.41	0.70
	GM	0.46	0.52	0.65	0.48	0.72
	CSF	0.27	0.39	0.49	0.33	0.59
#53	WM	0.45	0.48	0.61	0.40	0.68
	GM	0.44	0.49	0.60	0.48	0.67
	CSF	0.25	0.33	0.43	0.26	0.51
All cases	WM	0.45 ± 0.01	0.50 ± 0.02	0.62 ± 0.02	0.42 ± 0.02	0.68 ± 0.01
	GM	0.45 ± 0.02	0.51 ± 0.02	0.62 ± 0.02	0.48 ± 0.01	0.71 ± 0.02
	CSF	0.25 ± 0.02	0.36 ± 0.03	0.47 ± 0.02	0.32 ± 0.02	0.58 ± 0.03

Fig. 7. *(Color images)* Mean and standard deviation (SD) of the overlap measurements among all the 15 sets of registration experiments. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3
The values of the overlap measure obtained in the 5 3D inter-subject registration experiments in which **FFD**, **DEMONS** and the proposed method (**GC**) were performed with a same similarity measure (sum of squared difference). Simulated Brain Database image volume #04 was fixed as the source image throughout all the experiments and other five randomly selected image volumes (#42, #43, #46, #51 and #53) were used as the floating images. The last row shows the mean and the standard deviation (in format Mean $\pm$ SD) of the overlap measurements among all the five sets of registration experiments. Note that the comparison with **LP** was skipped in this set of experiments as it had been compared with the proposed method with the same similarity measure, as listed in Table 2.

Floating image	Tissue class	The values of the overlap measure			
		Before registration	**FFD**	**DEMONS**	**GC**
#42	WM	0.44	0.48	0.60	0.63
	GM	0.46	0.47	0.63	0.70
	CSF	0.24	0.36	0.49	0.55
#43	WM	0.45	0.50	0.63	0.65
	GM	0.42	0.50	0.62	0.68
	CSF	0.21	0.33	0.44	0.48
#46	WM	0.45	0.48	0.63	0.64
	GM	0.45	0.50	0.62	0.70
	CSF	0.27	0.33	0.48	0.60
#51	WM	0.46	0.50	0.61	0.66
	GM	0.47	0.51	0.63	0.69
	CSF	0.28	0.36	0.49	0.57
#53	WM	0.45	0.47	0.61	0.64
	GM	0.44	0.49	0.60	0.64
	CSF	0.25	0.31	0.43	0.48
All cases	WM	0.45 ± 0.01	0.49 ± 0.01	0.62 ± 0.02	0.64 ± 0.01
	GM	0.45 ± 0.02	0.49 ± 0.01	0.62 ± 0.02	0.68 ± 0.03
	CSF	0.25 ± 0.02	0.34 ± 0.02	0.47 ± 0.02	0.54 ± 0.05

portions of the skull (circled red in Fig. 4) in the floating image had its whole thickness overlapping with the interior of the brain in the source image; while in Case H, some portions of the skull (circled blue in Fig. 4) in the floating image had its whole

thickness being overlapping with the background in the source image. Since **DEMONS** uses the local intensity gradient to drive the pixel movements, these initial overlapping caused some pixels move towards the incorrect directions and finally got trapped in local minima. Although the proposed transformation model also has high degree of freedom, the proposed method still survives in this situation. It is because the graph-cuts method considers the labels of all pixels in a more global manner in an α-expansion move. Once the energy barrier is overcome, a group of pixels will together pursuit a large displacement in an α-expansion step instead of moving gradually through a series of small displacements.

As shown in Figs. 2 and 3, not only the proposed method survives in all cases, it can be noted from the subtraction images that the proposed method can generally produce better registration outputs than **FFD**, **DEMONS** and **LP** in the eight cases. Quantitatively, the last row in Table 1 lists the distributions of post-registration intensity errors averaged among all cases except Cases C, G and H. The mean absolute intensity difference throughout the whole image is 0.53 after applying our proposed method, which is significantly better than those values after applying **FFD** (3.34), **DEMONS** (1.87) and **LP** (2.62). This also provides a strong evidence that using the graph-cuts method as the optimization strategy in the proposed non-rigid image registration method can yield better results than **FFD**, **DEMONS** and **LP**.

4.3. Registration accuracy

To evaluate the registration accuracy, we performed 15 sets of 3D inter-subject registration experiments by using **FFD**, **DEMONS**, **LP** and the proposed method. One image volume (#04) was fixed as the source image throughout all the non-rigid registration experiments and other 15 image volumes were used as the floating images. The spatial resolutions of all image volumes are $256 \times 256 \times 181$ voxels with 1 mm isotropic voxel spacings. Figs. 5 and 6 show one of the 15 registration test results obtained by using the four registration methods. In Figs. 5 and 6, from the pre-difference and post-difference images, it is shown that the registration result

of the proposed method is better than those of **FFD**, **DEMONS** and **LP**. For all the 3D inter-subject registration tests, Table 2 lists the pre-registration and post-registration overlap measure values of the three tissue classes, white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF), after applying the four registration methods. We adopted the overlap measure $\#(A \cap B)/\#(A \cup B)$, as proposed by Crum et al. [34], where A and B denote the regions of the two images that belong to a specific tissue class. In addition, Fig. 7 shows the values of the overlap measure averaged among all 15 sets of registration experiments. From Table 2 and Fig. 7, it is found that the proposed method can consistently achieve higher registration accuracy than both **FFD**, **DEMONS** and **LP** in all the 15 sets of 3D registration experiments.

In order to demonstrate that the high accuracy of the proposed method is achieved by using the MRF model as well as the optimization strategy, rather than the similarity measure, we ran the proposed method and **FFD** on five image pairs with the sum of squared difference (**SSD**) as the similarity measure. The registration results were then compared with that of **DEMONS** with **SSD** as the similarity measure. Note that the comparison with **LP** was skipped in this set of experiments as it had been compared with the proposed method with the same similarity measure, as listed in Table 2. For the five image pairs, the randomly selected image volumes (#42, #43, #46, #51 and #53) were, respectively, used as the floating images, and image volume (#04) was fixed as the source image. Table 3 shows the overlap measure values of the three tissue classes after applying **FFD**, **DEMONS** and the proposed method. It is shown in Table 3 that the proposed method outperforms **FFD** and **DEMONS** in terms of accuracy under the same similarity measure.

4.4. Smoothness of the recovered deformation fields

Another crucial element to consider in registration result evaluation is the property of the recovered deformation fields
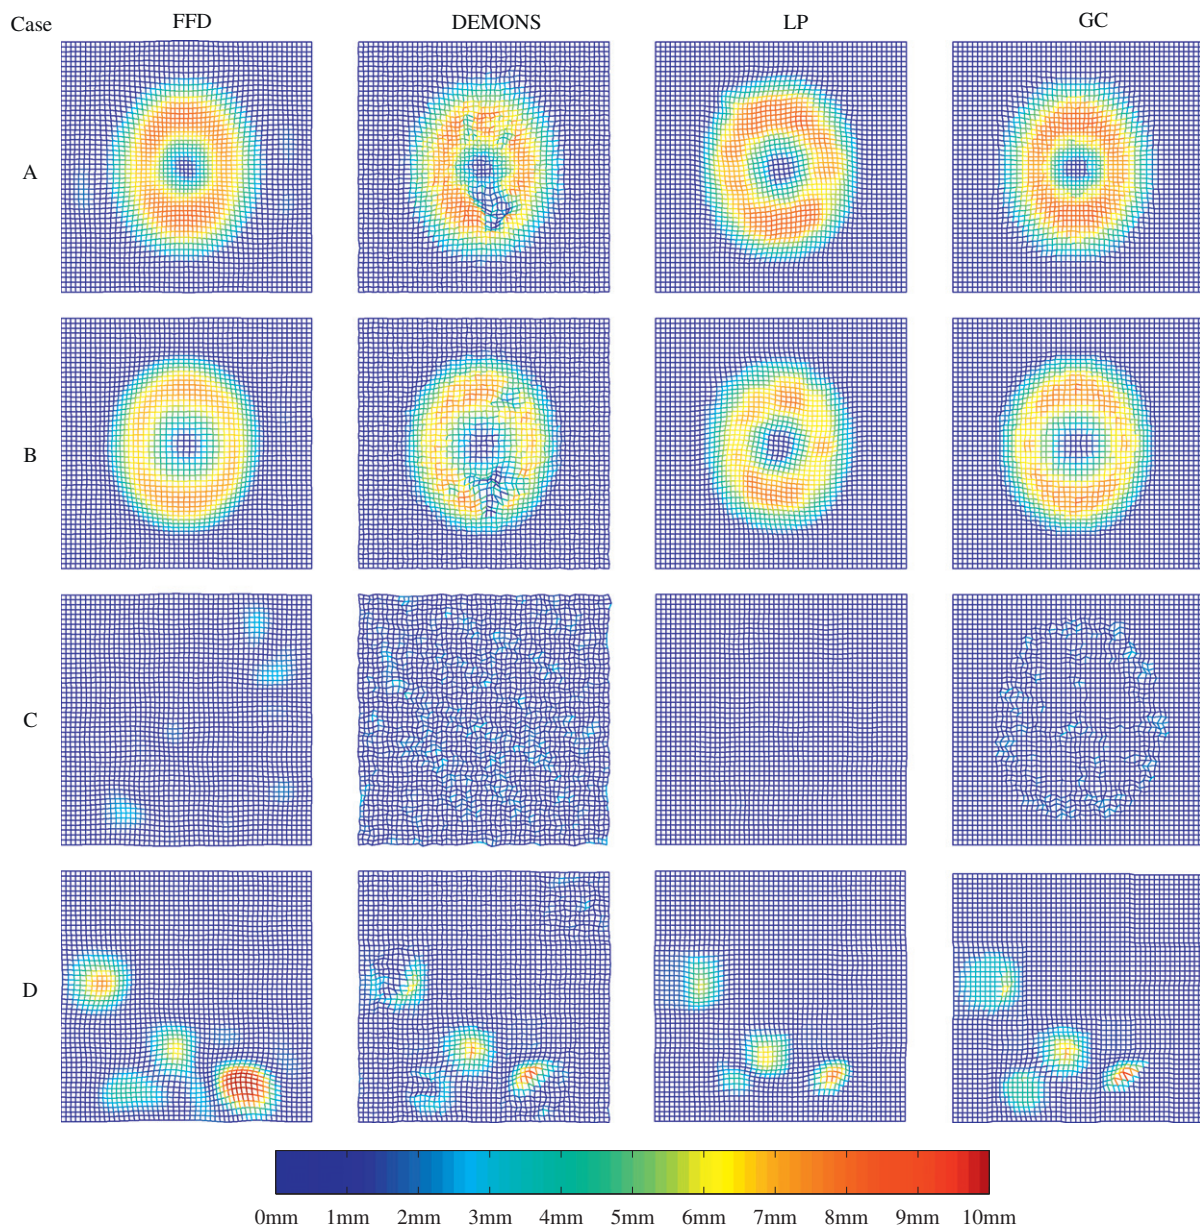

Fig. 8. (Color images) Recovered deformation fields obtained by using **FFD**, **DEMONS**, **LP** and the proposed method (**GC**) in Cases A–D. Red color represents large displacements while blue color represents small displacements. The color coding scheme for the displacement magnitudes is given in the last row. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

after registration, especially for those registration methods where no hard constraint is imposed on the deformation fields. It is because even if the registration output image is similar to the fixed source image, the recovered deformation field can be unrealistic.

To illustrate and compare the smoothness of the recovered deformation fields from the four registration methods, we plot the fields for Cases A–H (refer to Section 4.2 for the details about each case) in Figs. 8 and 9. Cases A and B correspond to a squeeze-in and a squeeze-out deformations. As predicted and shown in Fig. 8 (the first two rows), the fields recovered by **FFD** and **LP** are more smoother than **DEMONS** and the proposed method. It is because **FFD** and **LP** internally constrain the fields by the B-spline based transforms, and both **DEMONS** and the proposed method do not impose hard constrains on the deformation fields. Comparing with **DEMONS**, it can be observed that the field recovered by the proposed method is smoother than **DEMONS**. It can be the

effect induced by several factors however the real underlying cause is still being investigated. One of the possible reasons is that **DEMONS** only smooths the deformation fields after each iteration whereas the proposed method embeds the smoothness term in the energy function during the optimization process.

Another example is Case C, which is corresponding to ripple distortions. As shown in Fig. 8 (third row), it is observed from the recovered deformation fields that **FFD** and **LP** are not able to recover the ripple deformations while **DEMONS** and the proposed method are able to recover the deformations. From the same figure for this case, it can also be observed that the proposed method does not transform pixels in the background of the image while **DEMONS** transforms almost all background pixels. It is because the energy function, as stated in Eq. (6), that the proposed method optimizes contains the similarity measure as well as the smoothness term. In the background regions, the smoothness penalty due to the movements of pixels outweighs

Fig. 9. *(Color images)* Recovered deformation fields obtained by using **FFD**, **DEMONS**, **LP** and the proposed method (**GC**) in Cases E–H. Red color represents large displacements while blue color represents small displacements. The color coding scheme for the displacement magnitudes is given in the last row. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
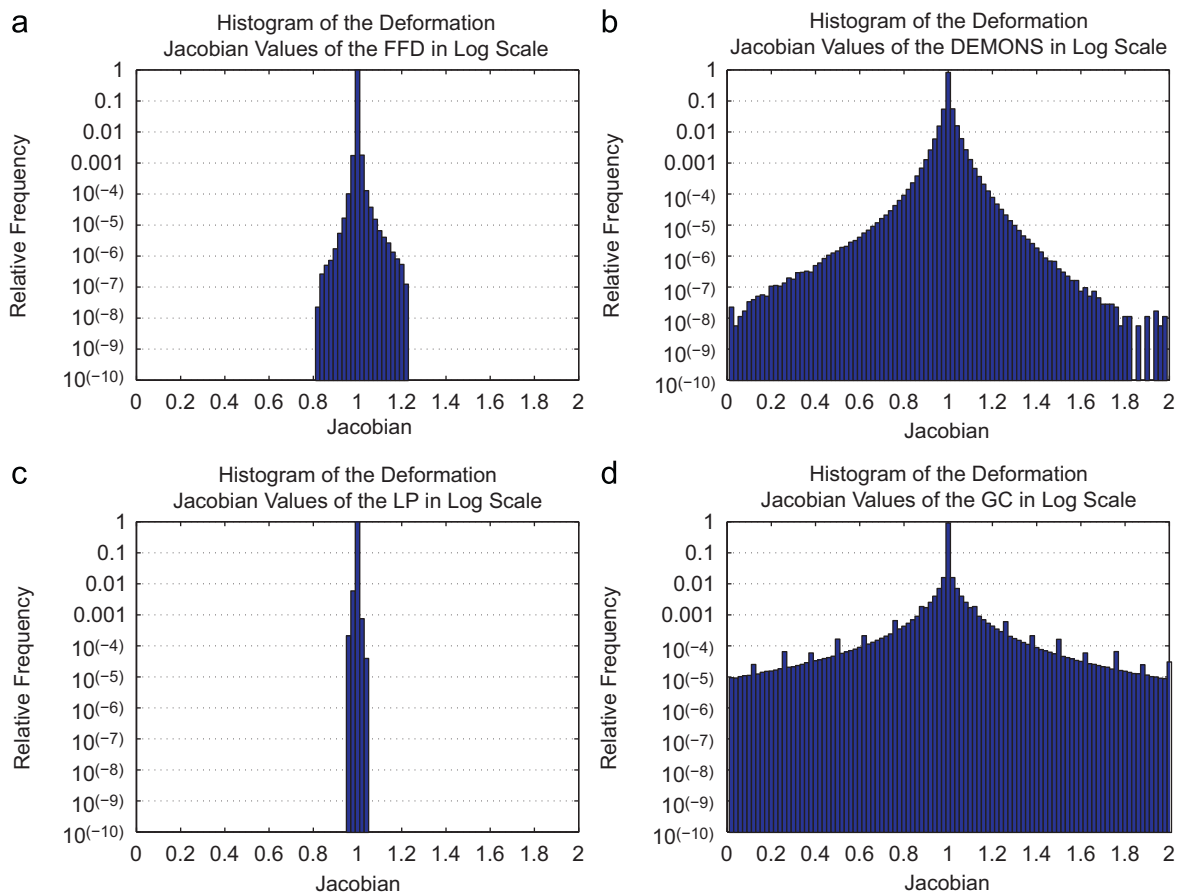
Fig. 10. The log-scaled histograms of the Jacobian values of the deformation fields recovered by (a) **FFD**, (b) **DEMONS**, (c) **LP** and (d) the proposed method (**GC**).

the decrease in **SAD** contributed by those movements and it causes the proposed method not to pursuit those movements.

For other cases, in general, it is observed that the deformation fields recovered by the proposed method in Cases D–H are realistic.

We have also examined the deformation fields obtained in the registration experiments on the 15 image volumes (refer to Section 4.3 for the details of the experiments). We calculated the values of the Jacobian of the deformation fields recovered by **FFD**, **DEMONS**, **LP** and the proposed method at all voxels in the 15 image volumes. Jacobian is the determinant of the Jacobian matrix formed by all first-order partial derivatives of the transformation vector field. It characterizes the factor of volume changes in the field for nearly any point p. A value greater than one indicates a volume expansion near p while less than one indicates a volume shrinking near p. We adopted the procedure for calculating the Jacobian as suggested in [35]. When determining the Jacobian matrices, the first-order partial derivatives of the deformation fields were obtained by applying convolution with a Gaussian derivative filter. The standard deviation of the Gaussian derivative filter was set to 2, which was twice the spatial resolution of the voxels [35]. Then, the values of Jacobian can be evaluated by finding the determinant of the corresponding Jacobian matrix. The last row in Fig. 6 shows the Jacobian images of the deformation fields recovered by the four methods in one of the 3D registration tests. The Jacobian image of the proposed method is comparable to **DEMONS** in general, except that in some regions the proposed method produces more extreme values than **DEMONS**. However, it can be observed from the registration image pair (the top row in Fig. 5) that there are significant topological changes in those regions, indicating that it is essential

to have large volume changes in order to register the image pair accurately and the accuracy obtained by the proposed method is listed in the second row of Table 2. Moreover, the overall histograms of the deformation Jacobian values are plotted in Figs. 10(a)–(d), for **FFD**, **DEMONS**, **LP** and the proposed method, respectively. It is observed from the histograms that the majority of the Jacobian have values near one. Although the deformation fields obtained by the proposed method may not be diffeomorphic, the experimental results have demonstrated that those deformation fields are still realistic and comparable to the deformation fields recovered by **DEMONS**. Therefore, the proposed method can be used in different application including clinical studies, atlas-based segmentation, etc.

4.5. Accuracy under different dissimilarity measures, intensity inhomogeneity and noise

We performed additional 24 3D experiments to study the effects on registration accuracy of the proposed method under different dissimilarity measures, intensity inhomogeneity (**RF**) and noise. The tested dissimilarity measures are sum of absolute difference (**SAD**) and sum of squared difference (**SSD**). Three **RF** levels (0%, 20% and 40%) and four noise levels (0%, 3%, 5% and 7%) were added to image volume (#20) to generate 12 image volumes as the floating images, and image volume (#4) was fixed as the source image throughout all the experiments. Therefore, there were 12 3D experiments for each dissimilarity measure. Note that the intensity inhomogeneity was formed by the inhomogeneity field A from the Simulated Brain Database.

Table 4 lists the experimental results. From the values of the overlap measure, there are four observations: (1) **SAD**

Table 4
The values of the overlap measure of the proposed method with two dissimilarity functions, sum of absolute difference (**SAD**) and sum of squared difference (**SSD**), obtained in the 24 3D inter-subject registration experiments with three intensity inhomogeneity (**RF**) levels and four noise levels. Simulated Brain Database image volume #04 was fixed as the source image throughout all the experiments and other 12 image volumes, generated by adding different levels of **RF** and noise to the image volume #20 of simulated Brain Database, were used as the floating images. The last row of each dissimilarity measure and the last column show the means and standard deviations (in format Mean ± SD) of the overlap measurements among a given **RF** level and a given noise level, respectively. Note that the values of overlap measure of white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF) before registration are, respectively, 0.4542, 0.4566 and 0.2489.

Measure	Floating image	Tissue class	The values of the overlap measure			Average over **RF**
			0% **RF**	20% **RF**	40% **RF**	
SAD	0% Noise	WM	0.6850	0.6503	0.5260	**0.6205 ± 0.0836**
		GM	0.7200	0.7132	0.6661	**0.6998 ± 0.0293**
		CSF	0.6119	0.6099	0.5914	**0.6044 ± 0.0113**
	3% Noise	WM	0.6770	0.6348	0.4784	**0.5967 ± 0.1046**
		GM	0.7206	0.7090	0.6485	**0.6927 ± 0.0387**
		CSF	0.6105	0.6031	0.5767	**0.5967 ± 0.0178**
	5% Noise	WM	0.6709	0.6209	0.4507	**0.5808 ± 0.1154**
		GM	0.7189	0.7032	0.6363	**0.6861 ± 0.0439**
		CSF	0.6061	0.5967	0.5651	**0.5893 ± 0.0215**
	7% Noise	WM	0.6638	0.6066	0.4309	**0.5671 ± 0.1214**
		GM	0.7168	0.6965	0.6262	**0.6798 ± 0.0475**
		CSF	0.6022	0.5887	0.5546	**0.5818 ± 0.0246**
	Average over noise	WM	**0.6742 ± 0.0090**	**0.6282 ± 0.0187**	**0.4715 ± 0.0413**	
		GM	**0.7191 ± 0.0017**	**0.7054 ± 0.0073**	**0.6443 ± 0.0172**	
		CSF	**0.6077 ± 0.0044**	**0.5996 ± 0.0091**	**0.5719 ± 0.0158**	
SSD	0% Noise	WM	0.6140	0.6125	0.5685	**0.5984 ± 0.0258**
		GM	0.6669	0.6795	0.6673	**0.6712 ± 0.0072**
		CSF	0.5956	0.5979	0.5901	**0.5945 ± 0.0040**
	3% Noise	WM	0.6135	0.6120	0.5357	**0.5871 ± 0.0445**
		GM	0.6717	0.6815	0.6554	**0.6695 ± 0.0132**
		CSF	0.5951	0.5939	0.5816	**0.5902 ± 0.0075**
	5% Noise	WM	0.6134	0.6102	0.5131	**0.5789 ± 0.0570**
		GM	0.6730	0.6808	0.6457	**0.6665 ± 0.0184**
		CSF	0.5914	0.5877	0.5724	**0.5838 ± 0.0101**
	7% Noise	WM	0.6130	0.6044	0.4934	**0.5703 ± 0.0667**
		GM	0.6734	0.6777	0.6354	**0.6622 ± 0.0233**
		CSF	0.5850	0.5774	0.5588	**0.5737 ± 0.0135**
	Average over noise	WM	**0.6135 ± 0.0004**	**0.6098 ± 0.0037**	**0.5277 ± 0.0322**	
		GM	**0.6712 ± 0.0030**	**0.6799 ± 0.0017**	**0.6510 ± 0.0136**	
		CSF	**0.5918 ± 0.0049**	**0.5892 ± 0.0089**	**0.5757 ± 0.0134**	

outperforms **SSD** in 0% **RF** and 20% **RF**. This can be observed by simply looking at the values of the overlap measure. (2) Both **SAD** and **SSD** are robust in noisy environment. It is noticed that the standard deviations (SDs) shown in the last row of each dissimilarity measure under 0% **RF** are relatively small (< 0.01). This means that the variation of the values of the overlap measure in different noise level under 0% **RF** is small, and hence proves that both dissimilarity measures are robust in noisy environment. (3) **SSD** is more robust than **SAD** under intensity inhomogeneity. The SDs shown in the last column can verify this conclusion since all the SDs of **SSD** in the last column are smaller than the corresponding SD of **SAD**. (4) 40% **RF** is quite challenging for the both dissimilarity measures. It is observed that there is a drop in the values of the overlap measure, especially for the white matter, when the **RF** was changed from 20% to 40%. It is because that the intensity range of white matter in the source image (i.e. #4) is close to the intensity range of gray matter in the floating images in which 40% **RF** was added. Therefore, some gray matter voxels of the floating images were mismatched to the white matter voxels of the source image. In short, the quantitative experimental results demonstrate that the proposed method can achieve high accuracy and robustness under different reasonable levels of noise and intensity inhomogeneity.

4.6. Choice of window $\mathcal{W}$

The window $\mathcal{W}$ in Eq. (8) controls how far a point in the image can move. This is essential for transforming a conventional continuous-domain non-rigid registration problem to a discrete labeling optimization problem that can be solved by using the proposed graph-cuts based image registration method. Thus, $\mathcal{W}$ needs to be chosen carefully such that not only it is capable of capturing most real-life deformation cases, but also time complexity is considered, since the number of α-expansion moves in each iteration of the optimization process is equal to $|\mathcal{W}|$. Fig. 11 gives the distribution of the displacement magnitudes of the recovered deformation fields estimated by using **DEMONS** in the 15 3D inter-subject experiments, as described in the Section 4.3. **DEMONS** is chosen since no hard constraints are imposed in the deformation model, which can give more related and useful statistics as the reference for the proposed method. Since the images have been already linearly pre-registered before the experiments, Fig. 11 can be considered as the general population statistics of the local deformations across different subjects. From the figure, it is observed that the displacement magnitudes of the recovered deformation fields were 99.6% bounded within seven voxel-units. Therefore, $\mathcal{W} = \{0, \pm 1, \pm 2, \ldots, \pm 7\}^3$ was chosen for the 3D inter-subject experiments, as

presented in Section 4.3, so that it can cover the displacements of almost all voxels and make the running time of the proposed time
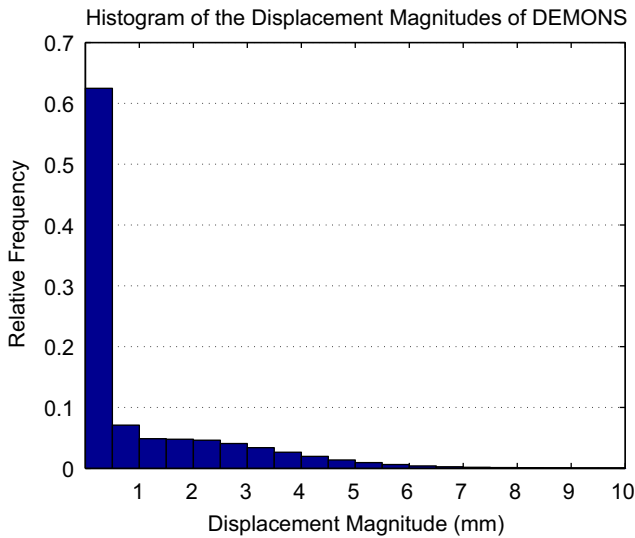

Fig. 11. Distribution of the displacement magnitudes in the deformation fields recovered by using **DEMONS** in the 15 3D inter-subject experiments, as described in Section 4.3.

comparable to that of **FFD**. Running times of the four registration methods are given in Section 4.8.

4.7. Comparing with dense sampling LP at different grid sizes

In Section 4.3, it reveals that **LP** with sparse sampling performs no better than the other three methods in 3D non-rigid image registration. To evaluate the performance of **LP** with dense sampling and the physical memory requirement of **LP** (refer to Section 3.3 for the details of the space capacity analysis), extra 56 3D experiments were performed on two image pairs. Image volume #04 from the Simulated Brain Database was the source image of these two image pairs, and image volumes #20 and #45 were the floating images, respectively. We performed the proposed method, and **LP** at three different grid sizes (60, 30, and 15 mm control point spacing) on each image pair with seven sets of labels $|L|$. Tables 5 and 6 show the pre-registration and post-registration overlap measure values of the three tissue classes, white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF), after applying the registration methods. The failed cases of **LP** in the tables mean that the implementation of **LP** was terminated unexpectedly before the registration processes were finished. Since **LP** failed more cases when the grid size was increased, we found that these failed cases were caused by the huge physical memory requirement of **LP**. For all other cases, it is shown that the proposed method outperforms **LP** under different grid sizes and numbers of

Table 5
The values of the overlap measure and computation time (**Time**) of the proposed method (**GC**) and **LP** (in three different grid sizes) obtained in the 28 3D inter-subject registration experiments with seven sets of labels $|L|$. Simulated Brain Database image volumes #04 and #20 were fixed as the source image and the floating image, respectively, throughout all the experiments.

$|L|$	Tissue class	The values of the overlap measure				
		Before registration	GC	LP @ $5 \times 5 \times 4$	LP @ $10 \times 10 \times 7$	LP@ $18 \times 18 \times 13$
27	WM	0.45	0.59	0.43	0.41	0.45
	GM	0.46	0.61	0.46	0.48	0.51
	CSF	0.25	0.43	0.26	0.32	0.36
	Time	–	943 s	856 s	1881 s	2370 s
125	WM	0.45	0.66	0.39	0.41	Failed
	GM	0.46	0.68	0.45	0.47	
	CSF	0.25	0.54	0.26	0.29	
	Time	–	4351 s	2775 s	6845 s	
343	WM	0.45	0.68	0.39	0.41	Failed
	GM	0.46	0.71	0.45	0.47	
	CSF	0.25	0.58	0.26	0.29	
	Time	–	11 485 s	7044 s	17 205 s	
729	WM	0.45	0.68	0.39	Failed	Failed
	GM	0.46	0.72	0.45		
	CSF	0.25	0.60	0.26		
	Time	–	24 541 s	16 289 s		
1331	WM	0.45	0.69	0.39	Failed	Failed
	GM	0.46	0.72	0.45		
	CSF	0.25	0.61	0.26		
	Time	–	43 215 s	29 144 s		
2197	WM	0.45	0.69	Failed	Failed	Failed
	GM	0.46	0.72			
	CSF	0.25	0.61			
	Time	–	71 746 s			
3375	WM	0.45	0.69	Failed	Failed	Failed
	GM	0.46	0.72			
	CSF	0.25	0.61			
	Time	–	86 689 s			

Table 6
The values of the overlap measure and computation time (**Time**) of the proposed method (**GC**) and **LP** (in three different grid sizes) obtained in the 28 3D inter-subject registration experiments with seven sets of labels $|L|$. Simulated Brain Database image volumes #04 and #45 were fixed as the source image and the floating image, respectively, throughout all the experiments.

$|L|$	Tissue class	The values of the overlap measure				
		Before registration	GC	LP @ $5 \times 5 \times 4$	LP @ $10 \times 10 \times 7$	LP @ $18 \times 18 \times 13$
27	WM	0.46	0.60	0.40	0.42	0.47
	GM	0.45	0.61	0.45	0.48	0.51
	CSF	0.25	0.43	0.26	0.31	0.37
	Time	–	928 s	1001 s	1985 s	2018 s
125	WM	0.46	0.67	0.40	0.43	Failed
	GM	0.45	0.69	0.45	0.48	
	CSF	0.25	0.55	0.26	0.31	
	Time	–	4230 s	2829 s	7289 s	
343	WM	0.46	0.69	0.40	0.43	Failed
	GM	0.45	0.71	0.45	0.48	
	CSF	0.25	0.60	0.27	0.31	
	Time	–	11 383 s	7135 s	18 446 s	
729	WM	0.46	0.70	0.40	Failed	Failed
	GM	0.45	0.72	0.45		
	CSF	0.25	0.62	0.26		
	Time	–	24 335 s	14 466 s		
1331	WM	0.46	0.70	0.40	Failed	Failed
	GM	0.45	0.73	0.45		
	CSF	0.25	0.63	0.26		
	Time	–	42 734 s	25 798 s		
2197	WM	0.46	0.70	Failed	Failed	Failed
	GM	0.45	0.73			
	CSF	0.25	0.63			
	Time	–	66 421 s			
3375	WM	0.46	0.69	Failed	Failed	Failed
	GM	0.45	0.72			
	CSF	0.25	0.63			
	Time	–	90 046 s			

labels. Note that the machine used in the experiments was equipped with 3 GBytes of physical memory and it is the maximum physical memory supported by 32-bits Windows systems.

4.8. Computation times

Using the experiment configurations as given in Section 4.1, the computation times for performing one set of 2D registration experiment using **FFD**, **DEMONS**, **LP** and the proposed method were all less than one minute. The computation times for performing one set of 3D registration experiment using **FFD**, **DEMONS**, **LP** and the proposed method were approximately 24 h, 6 h, 10 min and 24 h, respectively. Note that, regarding the configurations given in Section 4.1, sparse sampling was used in **LP** to prevent the registration experiments from being terminated unexpectedly because of the lack of memory. Therefore, the computation time of **LP** under this configuration is relatively faster than the other compared methods. The computation times of **LP** and the proposed method in different configurations can be found in Tables 5 and 6. It is shown that the computation time of the proposed method is comparable to that of **LP**, especially when dense control point grid was used in **LP**, while the accuracy of the proposed method is higher than that of **LP**. All experiments were conducted in a machine equipped with an Intel Core 2 Duo (2.13 GHz) CPU, 3 GBytes of physical memory and Windows XP operating system. All programs were implemented in C++ programming language.

5. Conclusion and future work

In this paper, we propose a method for the non-rigid medical image registration problems. First, a flexible non-parametric deformation model is adopted in the proposed formulation. It allows every voxel to displace freely and this is crucial for recovering the complex deformation fields in the medical image registration problems. The model also considers the smoothness of the adjacent displacements. Second, we suggest an energy function associated with the proposed deformation model for recovering the displacement vector field **D** during the non-rigid image registration process. This function considers the intensity similarity between the image volumes together with the smoothness requirement of the deformation field simultaneously, given the current transformation. Finally, despite the supernormal high degree of freedom in **D** as well as its smoothness requirement, we have shown that our energy function is submodular and can be optimized by using the graph-cuts method, which can provide a solution within a guaranteed factor of the exact minimum.

Experimental results on artificially distorted brain magnetic resonance images have demonstrated that the proposed method shows robustness against different challenging registration cases, e.g. large deformation, ripple distortion. It can be explained by the flexibility of the proposed deformation model, as well as the power of the graph-cuts method to perform optimization in a relatively global manner. Moreover, according to the registration results on the realistic brain phantoms obtained from the Simulated Brain Database, the proposed method can consistently achieve higher registration accuracy than the two state-of-the-art methods: free-form deformations based method [14] and demons based method [21], and the linear programming based method [22]. It is also shown that the recovered deformation fields are realistic.

In the future, we will address the quality of the transformation in terms of bijectivity and inverse consistency such that a diffeomorphic deformation field can be recovered. Moreover, since our formulation is simple and general, it deserves different extensions, such as incorporating feature point constraints.

Finally, although all the experiments presented in this paper were carried out on the brain magnetic resonance images, the proposed method can also be applied to other imaging modalities in theory and will be tested on different types of medical images and application scenarios.

References

[1] W. Wells, P. Viola, H. Atsumi, S. Nakajima, R. Kikinis, Multi-modal volume registration by maximization of mutual information, Medical Image Analysis 1 (1) (1996) 35–51.
[2] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, P. Suetens, Multimodality image registration by maximization of mutual information, IEEE Transactions on Medical Imaging 16 (2) (1997) 187–198.
[3] C. Studholme, D. Hill, D. Hawkes, An overlap invariant entropy measure of 3D medical image alignment, Pattern Recognition 32 (1) (1999) 71–86.
[4] J. Pluim, J.A. Maintz, M. Viergever, Mutual information based registration of medical images: a survey, IEEE Transactions on Medical Imaging 22 (8) (2003) 986–1004.
[5] A. Rao, G. Sanchez-Ortiz, R. Chandrashekara, M. Valdés, R. Mohiaddin, D. Rueckert, Construction of a cardiac motion atlas from MR using non-rigid registration, in: Functional Imaging and Modeling of the Heart FIMH, 2003, pp. 141–150.
[6] S. Warfield, M. Ferrant, X. Gallez, A. Nabavi, F. Jolesz, R. Kikinis, Real-time biomechanical simulation of volumetric brain deformation for image guided neurosurgery, in: Supercomputing, 2000, p. 23.
[7] D. Rueckert, A. Frangi, J. Schnabel, Automatic construction of 3D statistical deformation models using non-rigid registration, in: Medical Image Computing and Computer-Assisted Intervention (MICCAI), Lecture Notes in Computer Science, vol. 2208, 2001, pp. 77–84.
[8] B. Vemuri, J. Ye, Y. Chen, C. Leonard, Image registration via level-set motion: applications to atlas-based segmentation, Medical Image Analysis 7 (1) (2003) 1–20.
[9] A. Toga, J. Mazziotta, Brain Mapping: The Methods, second ed., Academic Press, 2002.
[10] B. Zitová, J. Flusser, Image registration methods: a survey, Image and Vision Computing 21 (2003) 977–1000.
[11] J. Modersitzki, Numerical Methods for Image Registration, Oxford University Press, 2004.
[12] A. Gholipour, N. Kehtarnavaz, R. Briggs, M. Devous, K. Gopinath, Brain functional localization: a survey of image registration techniques, IEEE Transactions on Medical Imaging 26 (4) (2007) 427–451.
[13] V. Arsigny, X. Pennec, N. Ayache, Polyrigid and polyaffine transformations: a novel geometrical tool to deal with non-rigid deformations—application to the registration of histological slices, Medical Image Analysis 9 (6) (2005) 507–523.
[14] D. Rueckert, L.I. Sonoda, C. Hayes, D. Hill, M. Leach, D. Hawkes, Non-rigid registration using free-form deformations: application to breast MR images, IEEE Transactions on Medical Imaging 18 (8) (1999) 712–721.
[15] Y. Boykov, V. Kolmogorov, Computing geodesics and minimal surfaces via graph cuts, in: International Conference on Computer Vision (ICCV), vol. 1, 2003, pp. 26–33.
[16] Y. Boykov, O. Veksler, R. Zabih, Fast approximate energy minimization via graph cuts, IEEE Transactions on Pattern Analysis and Machine Intelligence 23 (11) (2001) 1222–1239.
[17] S. Roy, Stereo without epipolar lines: a maximum-flow formulation, International Journal of Computer Vision 34 (2–3) (1999) 147–161.
[18] J. Winn, N. Jojic, Locus: learning object classes with unsupervised segmentation, in: International Conference on Computer Vision (ICCV), 2005, pp. 756–763.
[19] V. Kolmogorov, R. Zabih, What energy functions can be minimized via graph cuts?, IEEE Transactions on Pattern Analysis and Machine Intelligence 26 (2) (2004) 147–159.
[20] T.W.H. Tang, A.C.S. Chung, Non-rigid image registration using graph-cuts, in: Medical Image Computing and Computer-Assisted Intervention (MICCAI), Lecture Notes in Computer Science, vol. 4791, 2007, pp. 916–924.
[21] J.-P. Thirion, Image matching as a diffusion process: an analogy with maxwell's demons, Medical Image Analysis 2 (3) (1998) 243–260.
[22] B. Glocker, N. Komodakis, G. Tziritasc, N. Navabb, N. Paragios, Dense image registration through mrfs and efficient linear programming, Medical Image Analysis 12 (6) (2008) 731–741.
[23] D. Shen, C. Davatzikos, HAMMER: hierarchical attribute matching mechanism for elastic registration, Transactions on Medical Imaging 21 (11) (2002) 1421–1439.
[24] B. Horn, B. Schunck, Determining optical flow, Technical Report AIM-572, MIT, 1980.

[25] B. Glocker, N. Komodakis, N. Paragios, G. Tziritas, N. Navab, Inter and intra-modal deformable registration: continuous deformations meet efficient optimal linear programming, in: Information Processing in Medical Imaging (IPMI), Lecture Notes in Computer Science, vol. 4584, 2007, pp. 408–420.

[26] N. Komodakis, G. Tziritas, N. Paragios, Fast, approximately optimal solutions for single and dynamic mrfs, in: IEEE Computer Vision and Pattern Recognition Conference (CVPR), 2007, pp. 1–8.

[27] V. Kolmogorov, C. Rother, Minimizing nonsubmodular functions with graph cuts—a review, IEEE Transactions on Pattern Analysis and Machine Intelligence 29 (7) (2007) 1274–1279.

[28] B. Aubert-Broche, A. Evans, D. Collins, A new improved version of the realistic digital brain phantom, NeuroImage 32 (1) (2006) 138–145.

[29] B. Aubert-Broche, M. Griffin, G. Pike, A. Evans, D. Collins, Twenty new digital brain phantoms for creation of validation image data bases, IEEE Transactions on Medical Imaging 25 (11) (2006) 1410–1416.

[30] The national library of medicine insight segmentation and registration toolkit, ⟨http://www.itk.org/⟩.

[31] J. Schnabel, D. Rueckert, M. Quist, J. Blackall, A.C. Smith, T. Hartkens, G. Penney, W. Hall, H. Liu, C. Truwit, F. Gerritsen, D. Hill, D. Hawkes, A generic framework for non-rigid registration based on non-uniform multi-level free-form deformations, in: Medical Image Computing and Computer-Assisted Intervention (MICCAI), Lecture Notes in Computer Science, vol. 2208, 2001, pp. 573–581.

[32] Image registration toolkit, ⟨http://www.homes.doc.ic.ac.uk/dr/software/⟩.

[33] Drop—deformable registration using discrete optimization, ⟨http://www.mrf-registration.net/deformable/index.html⟩.

[34] W. Crum, D. Rueckert, M. Jenkinson, D. Kennedy, S. Smith, A framework for detailed objective comparison of non-rigid registration algorithms in neuroimaging, in: Medical Image Computing and Computer-Assisted Intervention (MICCAI), 2004, pp. 679–686.

[35] P. Hellier, C. Barillot, et al., Retrospective evaluation of intersubject brain registration, IEEE Transactions on Medical Imaging 22 (9) (2003) 1120–1130.

R.W.K. So received the B.Eng. degree (first class Honors) in computer engineering and the M.Phil. degree in computer science and engineering in 2007 and 2009, respectively, from The Hong Kong University of Science and Technology. He is currently pursuing the Ph.D. degree in computer science and engineering at the Hong Kong University of Science and Technology.

His research interests include medical image processing and analysis, vascular segmentation, lumen center extraction and medical image registration.

T.W.H. Tang received the B.Eng. degree (first class Honors) in computer science and the M.Phil. degree in computer science in 2005 and 2007, respectively, from The Hong Kong University of Science and Technology.

His research interests include medical image processing and analysis, and medical image registration.

A.C.S. Chung received the B.Eng. degree (first class Honors) in computer engineering from The University of Hong Kong in 1995 and the M.Phil. degree in computer science from The Hong Kong University of Science and Technology in 1998.

He joined the Medical Vision Laboratory, University of Oxford, Oxford, U.K., as a doctoral research student with a Croucher Foundation scholarship and graduated in 2001. He was a Visiting Scientist at the Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, in 2001. He is currently an Associate Professor with the Department of Computer Science and Engineering, The Hong Kong University of Science and Technology. His research interests include medical image analysis, image processing, and computer vision.

Dr. Chung won the 2002 British Machine Vision Association Sullivan Thesis Award for the best doctoral thesis submitted to a U.K. University in the field of computer or natural vision. Authorized licensed.