Available at www.ComputerScienceWeb.com
POWERED BY SCIENCE DIRECT®

Medical Image Analysis 7 (2003) 565–575

MEDICAL IMAGE ANALYSIS
www.elsevier.com/locate/media

A viscous fluid model for multimodal non-rigid image registration using mutual information

Emiliano D'Agostino*, Frederik Maes[1], Dirk Vandermeulen, Paul Suetens

Faculties of Medicine and Engineering, Medical Image Computing (Radiology-ESAT/PSI), Katholieke Universiteit Leuven, University Hospital Gasthuisberg, Herestraat 49, B-3000 Leuven, Belgium

Abstract

We propose a multimodal free-form registration algorithm based on maximization of mutual information. The warped image is modeled as a viscous fluid that deforms under the influence of forces derived from the gradient of the mutual information registration criterion. Parzen windowing is used to estimate the joint intensity probability of the images to be matched. The method is evaluated for non-rigid inter-subject registration of MR brain images. The accuracy of the method is verified using simulated multi-modal MR images with known ground truth deformation. The results show that the root mean square difference between the recovered and the ground truth deformation is smaller than 1 voxel. We illustrate the application of the method for atlas-based brain tissue segmentation in MR images in case of gross morphological differences between atlas and patient images.

Keywords: Non-rigid-registration; Mutual information; Viscous fluid model

1. Introduction

Combining information from multiple images, possibly acquired using different modalities, at different time points or from different subjects, requires image registration, i.e. knowledge of the geometric relationship between physically corresponding points in all images. Retrospective registration of three-dimensional (3-D) images, or the recovery of the coordinate transformation that maps points in one image volume onto their anatomically corresponding points in the other from the image content itself, is a fundamental problem in medical image analysis.

Strategies for medical image registration can be classified according to the image features used to establish geometric correspondence between both images (Maintz and Viergever, 1998). Point based or surface based registration requires localization or segmentation of corresponding anatomical landmarks or object surfaces in the images to be registered, which is non trivial in most applications, error prone and often difficult to automate. Voxel based registration approaches on the other hand, compute the registration solution by maximizing intensity similarity between both images, thereby considering all voxels in the region of overlap of the images to be registered without need for prior segmentation or preprocessing. Various voxel based registration measures have been proposed that compute intensity similarity from the intensity values directly, typically assuming the intensities of corresponding voxels to be identical (e.g. sum of squared intensity differences) or linearly related (e.g. intensity correlation), which limits their use to unimodal applications only. In contrast, maximization of mutual information (MMI) of corresponding voxel intensities assesses intensity similarity of the images to be registered from the co-occurrence of intensities in both images as reflected by their joint intensity histogram, which varies as the registration parameters are changed (Maes et al., 1997; Wells et al., 1996; Studholme et al., 1999; Pluim et al., 2000). The MMI registration criterion postulates that the statistical dependence between corresponding voxel intensities is maximal at registration, without imposing limiting constraints on the nature of this relationship (except for spatial stationarity).

The MMI criterion has been demonstrated to be highly

*Corresponding author.
E-mail address: emiliano.dagostino@uz.kuleuven.ac.be (E. D'Agostino).
[1]Frederik Maes is Postdoctoral Fellow of the Fund for Scientific Research, Flanders (Belgium).

1361-8415/03/$ – see front matter
doi:10.1016/S1361-8415(03)00039-2

successful for rigid body or affine registration of multimodal images in a variety of applications where the rigid body assumption can be assumed to be valid or local tissue distortions can be neglected (West et al., 1997). Such applications include the registration of images of the same patient or the global alignment of images of different patients or of patient and atlas images (Van Leemput et al., 1999). However, in applications where local morphological differences need to be quantified, affine registration, using only global translation, rotation, and possibly scaling and skew, is no longer sufficient and more general non-rigid registration (NRR) is required. NRR aims at recovering a dense 3-D field of 3-D displacement vectors that maps each voxel individually in one image volume onto its corresponding voxel in the other, allowing the registration to adapt to local distortions instead of being restricted to global alignment of both images only. Applications for NRR include shape analysis (to warp all shapes to a standard space for statistical comparison), atlas-based segmentation (to compensate for gross morphological differences between atlas and study images), image rectification (to correct for geometric distortion in the images) or motion analysis (to infer object motion from the deformation between consecutive frames in dynamic image sequences).

Several approaches have been proposed to extend the MMI criterion to NRR. These differ in their representation of the deformation field and in the way the variation of MI with changes in the deformation parameters is estimated. A popular representation of the deformation field is the use of smooth and differentiable basis functions with global (e.g. thin-plate splines (Meyer et al., 1997)) or local (e.g. B-splines (Rueckert et al., 1999), radial basis functions (Rohde et al., 2001)) support. While the basis functions implicitly impose local small scale smoothness on the deformation field, regularization at larger scales may require inclusion of an appropriate cost function in the registration criterion to penalize non-smooth deformations explicitly (Rueckert et al., 1999). Spline-based approaches can correct for gross shape differences, but a dense grid of control points is required to characterize the deformation at voxel level detail, implying high computational complexity unless a strategy for local adaptive grid refinement is used (Schnabel et al., 2001).

Block matching (Gaens et al., 1998) or free-form NRR approaches, using a non-parameterized expression for the deformation field, assign a local deformation vector to each voxel individually, yielding up to $3\times N$ degrees of freedom with N the number of voxels. These methods are, therefore, in general more flexible than representations using basis functions, but need appropriate constraints for spatial regularization of the resulting vector field to assure that the deformation is physically realistic and acceptable. Such constraints are typically implemented by modeling the deforming image as an elastic or viscous medium. Recently, a free-form NRR algorithm was presented by Hermosillo et al. (2001), defining the forces driving the deformation at each voxel such that mutual information is maximized and using a regularization functional derived from linear elasticity theory. However, such elastic regularizer is suitable only when displacements can be assumed to be small. In this paper we focus on the application of non-rigid image registration for inter-subject comparison of MR brain images, whereby large local deformations may have to be recovered as large morphological differences may exist in the brains of different subjects, e.g. due to the presence of enlarged ventricles in certain patients. For this particular application, a viscous fluid model is more appropriate.

In this paper, we extend the approach of Hermosillo et al. (2001) by replacing the elastic model by the viscous fluid regularization scheme of Christensen et al. (1996b) and thus generalize the method of Christensen et al. (1996b) to multimodal image registration based on MMI. The Navier–Stokes equation modeling the viscous fluid is solved by iteratively updating the deformation field and convolving it with a Gaussian filter as in (Thirion, 1998), approximating the approach of Bro-Nielsen and Gramkow (1996). The deformation field is regridded as needed during iterations as in (Christensen et al., 1996b) to assure that its Jacobian remains positive everywhere, such that the method can handle large deformations. We verified the robustness of the method by applying realistic known deformations to simulated multispectral MR brain images and evaluating the difference between the recovered and ground truth deformation fields in terms of displacement errors and of tissue classification errors when using the recovered deformation for atlas-based segmentation.

2. Method

2.1. The viscous fluid model

A template image $\mathcal{F}$ is deformed towards a target image $\mathcal{G}$ by the transformation $\vec{T}$, that is represented using an Eulerian reference frame as $\vec{T} = \vec{x} - \vec{u}(\vec{x})$, mapping fixed voxel positions $\vec{x}$ in $\mathcal{G}$ onto the corresponding points $\vec{x} - \vec{u}(\vec{x})$ in the original template $\mathcal{F}$ (Christensen et al., 1996b). The deforming template image is considered as a viscous fluid whose motion is governed by its Navier–Stokes equation of conservation of momentum. Following the argumentation in (Christensen et al., 1996b), this equation can be simplified to

$$\mu\nabla^2\vec{v} + (\mu + \lambda)\vec{\nabla}(\vec{\nabla} \cdot \vec{v}) + \vec{F}(\vec{x}, \vec{u}) = 0, \quad (1)$$

with μ and λ material parameters. We set $\mu = 1$ and $\lambda = 0$ (Wang and Staib, 2000). $\vec{v}(\vec{x}, t)$ is the deformation velocity experienced by a particle at position $\vec{x}$, that is non-linearly related to $\vec{u}$ by

$$\vec{v} = \frac{\mathrm{d}\vec{u}}{\mathrm{d}t} = \frac{\partial \vec{u}}{\partial t} + \sum_{i=1}^{3} v_i \frac{\partial \vec{u}}{\partial x_i}, \tag{2}$$

with $\vec{v} = [v_1, v_2, v_3]^{\mathrm{T}}$ and $\vec{u} = [u_1, u_2, u_3]^{\mathrm{T}}$.

$\vec{F}(\vec{x}, \vec{u})$ is the force field acting at position $\vec{x}$, that depends on the deformation $\vec{u}$ and that drives the deformation in the appropriate direction. In Section 2.2, we derive an expression for $\vec{F}$ such that the viscous fluid flow maximizes mutual information between corresponding voxel intensities of $\mathcal{G}(\vec{x})$ and $\mathcal{F}(\vec{x} - \vec{u}(\vec{x}))$. At each time instance during the deformation, the term $\vec{F}(\vec{x}, \vec{u})$ is constant, such that the modified Navier–Stokes equation can be solved iteratively as a temporal concatenation of linear equations. Solving (1) yields deformation velocities, from which the deformation itself can be computed by integration over time. In (Christensen et al., 1996b) the Navier–Stokes equation is solved by successive over relaxation (SOR), but this is a computationally expensive approach. Instead, we obtain the velocity field by simple convolution of the force field with a 3-D Gaussian kernel ϕ_s with width s (in voxels) as in (Thirion, 1998), which is an approximation of the filter kernel derived in (Bro-Nielsen and Gramkow, 1996):

$$\vec{v}^{(k)} = \phi_s \vec{F}^{(k)}, \tag{3}$$

with $\vec{F}^{(k)}$ the force field acting on $\mathcal{F}$ at iteration k. The displacement $\vec{u}^{(k+1)}$ at iteration $(k+1)$ is then given by

$$\vec{u}^{(k+1)} = \vec{u}^{(k)} + \vec{R}^{(k)} \cdot \Delta t^{(k)}, \tag{4}$$

with $\vec{R}^{(k)}$ the perturbation to the deformation field:

$$\vec{R}^{(k)} = \vec{v}^{(k)} - \sum_{i=1}^{3} v_i^{(k)} \left[\frac{\partial \vec{u}^{(k)}}{\partial x_i} \right], \tag{5}$$

and $\Delta t^{(k)}$ a time step parameter that may be adapted during iterations.

To preserve the topology of the deformed template image, the Jacobian of the deformation field should not become negative. When the Jacobian becomes anywhere smaller than some positive threshold, regridding of the deformed template image is applied as in (Christensen et al., 1996b) to generate a new template, setting the incremental displacement field to zero. The total deformation is the concatenation of the incremental deformation fields associated with each propagated template.

2.2. Force field definition

We define an expression for the force field $\vec{F}(\vec{x}, \vec{u})$ in (1) such that the viscous fluid deformation strives at maximizing mutual information $I(\vec{u})$ of corresponding voxel intensities between the deformed template image $\mathcal{F}(\vec{x} - \vec{u})$ and the target image $\mathcal{G}(\vec{x})$. We adopt here the approach of Hermosillo et al. (2001) who derived an expression for the gradient $\nabla_{\vec{u}} I$ of I with respect to the deformation field $\vec{u}$, modeling the joint intensity distribution $p_{\vec{u}}^{\mathcal{F},\mathcal{G}}(i_1, i_2)$ of deformed template and target images as a continuous function using Parzen windowing.

Mutual information I between $\mathcal{F}(\vec{x} - \vec{u})$ and $\mathcal{G}(\vec{x})$ is given by

$$I(\vec{u}) = \int\int p_{\vec{u}}^{\mathcal{F},\mathcal{G}}(i_1, i_2) \log \frac{p_{\vec{u}}^{\mathcal{F},\mathcal{G}}(i_1, i_2)}{p^{\mathcal{F}}(i_1) p_{\vec{u}}^{\mathcal{G}}(i_2)} \,\mathrm{d}i_1 \,\mathrm{d}i_2. \tag{6}$$

If the deformation field $\vec{u}$ is perturbed into $\vec{u} + \epsilon\vec{h}$, variational calculus yields the first variation of I:

$$\begin{aligned}
&\frac{\partial I(\vec{u} + \varepsilon\vec{h})}{\partial \epsilon}\bigg|_{\epsilon=0} \\
&= \int\int \frac{\partial}{\partial \epsilon} \left[p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{F},\mathcal{G}}(i_1, i_2) \log \frac{p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{F},\mathcal{G}}(i_1, i_2)}{p^{\mathcal{F}}(i_1) p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{G}}(i_2)} \right]_{\epsilon=0} \mathrm{d}i_1 \,\mathrm{d}i_2 \\
&= \int\int \left[\left(1 + \log \frac{p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{F},\mathcal{G}}(i_1, i_2)}{p^{\mathcal{F}}(i_1) p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{G}}(i_2)} \right) \frac{\partial p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{F},\mathcal{G}}(i_1, i_2)}{\partial \epsilon} \right. \\
&\quad \left. - \frac{p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{F},\mathcal{G}}(i_1, i_2)}{p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{G}}(i_2)} \frac{\partial p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{G}}(i_2)}{\partial \epsilon} \right]_{\epsilon=0} \mathrm{d}i_1 \,\mathrm{d}i_2.
\end{aligned} \tag{7}$$

Because

$$\int p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{F},\mathcal{G}}(i_1, i_2) \,\mathrm{d}i_1 = p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{G}}(i_2) \tag{8}$$

and

$$\int p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{G}}(i_2) \,\mathrm{d}i_2 = 1, \tag{9}$$

the last term of (7) reduces to

$$\begin{aligned}
&\int\int \frac{p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{F},\mathcal{G}}(i_1, i_2)}{p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{G}}(i_2)} \frac{\partial p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{G}}(i_2)}{\partial \epsilon} \,\mathrm{d}i_1 \,\mathrm{d}i_2 \\
&= \int \frac{\partial p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{G}}(i_2)}{\partial \epsilon} \frac{\int p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{F},\mathcal{G}}(i_1, i_2) \mathrm{d}i_1}{p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{G}}(i_2)} \,\mathrm{d}i_2 \\
&= \int \frac{\partial p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{G}}(i_2)}{\partial \epsilon} \,\mathrm{d}i_2 = \frac{\partial}{\partial \epsilon} \int p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{G}}(i_2) \mathrm{d}i_2 = 0,
\end{aligned} \tag{10}$$

and (7) simplifies to

$$\begin{aligned}
\frac{\partial I(\vec{u} + \epsilon\vec{h})}{\partial \epsilon}\bigg|_{\epsilon=0} = \int\int &\left(1 + \log \frac{p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{F},\mathcal{G}}(i_1, i_2)}{p^{\mathcal{F}}(i_1) p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{G}}(i_2)} \right) \\
&\times \frac{\partial p_{\vec{u}+\epsilon\vec{h}}^{\mathcal{F},\mathcal{G}}(i_1, i_2)}{\partial \epsilon}\bigg|_{\epsilon=0} \mathrm{d}i_1 \,\mathrm{d}i_2.
\end{aligned} \tag{11}$$

The joint intensity probability is estimated from the region of overlap ν of both images (with volume V), using the 2-D Parzen windowing kernel $\psi_h(i_1, i_2)$ with width h:

$$p_{\vec{u}}^{\mathcal{F},\mathcal{G}}(i_1, i_2) = \frac{1}{V} \int_{\nu} \psi_h(i_1 - \mathcal{F}(\vec{x} - \vec{u}), i_2 - \mathcal{G}(\vec{x})) \,\mathrm{d}\vec{x}. \tag{12}$$

Inserting (12) in (11) and rearranging as in (Hermosillo et al., 2001), yields

$$\left.\frac{\partial I(\vec{u}+\epsilon\vec{h})}{\partial \epsilon}\right|_{\epsilon=0} = \frac{1}{V}\int_{\nu}\left[\frac{\partial \psi_h}{\partial i_1}L_{\vec{u}}\right](\mathcal{F}(\vec{x}-\vec{u}),\mathcal{G}(\vec{x}))\nabla F(\vec{x}-\vec{u})\vec{h}(\vec{x})\mathrm{d}\vec{x}, \tag{13}$$

with

$$L_{\vec{u}}(i_1,i_2) = 1+\log\frac{p_{\vec{u}}^{\mathcal{F},\mathcal{G}}(i_1,i_2)}{p^{\mathcal{F}}(i_1)p_{\vec{u}}^{\mathcal{G}}(i_2)}. \tag{14}$$

We, therefore, define the force field $\vec{F}$ at $\vec{x}$ to be equal to the gradient of I with respect to $\vec{u}(\vec{x})$, such that $\vec{F}$ drives the deformation to maximize I:

$$\vec{F}(\vec{x},\vec{u}) = \nabla_{\vec{u}} I = \frac{1}{V}\left[\frac{\partial \psi_h}{\partial i_1}L_{\vec{u}}\right](\mathcal{F}(\vec{x}-\vec{u}),\mathcal{G}(\vec{x}))\nabla\mathcal{F}(\vec{x}-\vec{u}). \tag{15}$$

Thus, $\vec{F}(\vec{x},\vec{u})$ is directed along the image intensity gradient of the deformed template $\mathcal{F}(\vec{x}-\vec{u})$, weighted by the impact on the mutual information of a particle in $\mathcal{F}$ at $\vec{x}-\vec{u}$ being displaced in this direction.

2.3. Joint probability estimation

The estimation of the joint image intensity probability using (12) requires a proper value for the Parzen kernel width h. This value is determined automatically using a standard leave-k-out cross-validation technique applied to the two marginal histograms (Turlach, 1993; Hermosillo, 2002). The Parzen estimator for the probability density function $f(x)$ given n samples X_i is defined by

$$\hat{f}_h(x) = \frac{1}{n}\sum_{i=1}^{n} K_h(x-X_i), \tag{16}$$

with K a symmetric kernel function such that $\int K(u)\,\mathrm{d}u = 1$ and $K_h(u) = (1/h)K(u/h)$ with h the kernel width. For our estimator, we use the Gaussian kernel:

$$K_h(u) = \frac{1}{\sqrt{2\pi h^2}}\exp\left(-\frac{u^2}{2h^2}\right). \tag{17}$$

To determine an optimal value for h, we can select h such that it maximizes the pseudo-likelihood $P(h)$ (Turlach, 1993):

$$P(h) = \prod_{i=1}^{n}\hat{f}_h(X_i). \tag{18}$$

However, since this pseudo-likelihood has a trivial maximum for $h=0$, it has been suggested to use leave-one-out cross validation, replacing $\hat{f}_h$ in (18) by

$$\hat{f}_{h,i} = \frac{1}{n-1}\sum_{j=1,j\neq i}^{n} K_h(x-X_i). \tag{19}$$

This way of choosing h minimizes the Kullback–Leibler distance between $\hat{f}_h(x)$ and $f(x)$.

The pseudo-likelihood function $P(h) = \Sigma_i \hat{f}_{h,i}$ depends on the selected samples X_i. However, as illustrated in Fig. 1, we found that the maximum of $P(h)$ is not much affected if not all image samples are accounted for, but only a subset thereof. To save computation time, we consider only 1 out of M voxels to estimate h, by simply picking the first and every M-th voxel in the image.

2.4. Implementation issues

Voxels in $\mathcal{G}$ at grid positions $\vec{x}$ with intensity i_2 are transformed into $\mathcal{F}$ and trilinear interpolation is used to determine the corresponding intensities i_1 in $\mathcal{F}$ at the transformed positions $\vec{x}-\vec{u}$. The joint histogram $H_{\vec{u}}^{\mathcal{F},\mathcal{G}}(i_1,i_2)$ of $\mathcal{F}$ and $\mathcal{G}$ within their volume of overlap $\mathcal{V}$ is constructed by binning the pairs (i_1,i_2), after appropriate linear rescaling of all values within the intensity range of either image and using 128 bins for both images.

The cross-validation scheme described in Section 2.3 is applied twice to determine $h_{\mathcal{F}}$ and $h_{\mathcal{G}}$ for the template and target image, respectively, subsampling each image by a factor $M=20$. We select $h=\max(h_{\mathcal{F}}, h_{\mathcal{G}})$ to define an isotropic 2-D Parzen Gaussian kernel $\psi_h(i_1, i_2) = K_h(i_1)K_h(i_2)$ with K_h defined in (17). The joint image intensity probability $p_{\vec{u}}^{\mathcal{F},\mathcal{G}}$ in (12) is computed by the convolution of H with a discrete approximation of ψ:

$$p_{\vec{u}}^{\mathcal{F},\mathcal{G}}(i_1,i_2) = \psi_h(i_1,i_2)H_{\vec{u}}^{\mathcal{F},\mathcal{G}}(i_1,i_2). \tag{20}$$

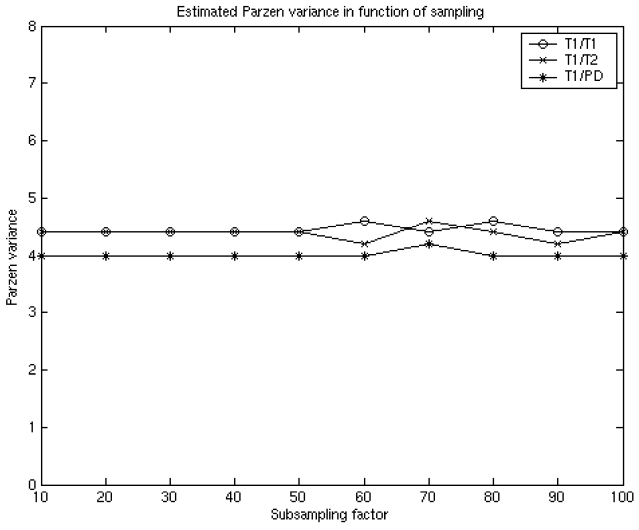

Fig. 1. Effect of image sampling on the Parzen variance estimation for three different multimodality image pairs. The maximum of the pseudo-likelihood $P(h)$ (18) varies hardly for subsampling factors M ranging from 10 to 100.

The marginal histograms $p^{\mathcal{F}}(i_1)$ and $p^{\mathcal{G}}_{\vec{u}}(i_2)$ are obtained by integrating $p^{\mathcal{F},\mathcal{G}}_{\vec{u}}$ over rows and columns, respectively. The force field $\vec{F}(\vec{x},\vec{u})$ is computed using (14) and (15), with the gradient $\nabla\mathcal{F}$ computed using a finite difference approximation.

The displacement field $\vec{u}$ is updated iteratively using (3), (5) and (4). A 3-D isotropic Gaussian kernel was used for ϕ_s in (3). The time step parameter Δt in (4) is adapted each iteration and set to

$$\Delta t^{(k)} = \max(\|\vec{R}^{(k)}\|)\cdot\ \Delta u, \tag{21}$$

with Δu (in voxels) the maximal voxel displacement that is allowed in one iteration. The impact of these parameters is explored in Section 3.1. Regridding of the deforming template image is performed when the Jacobian of $\vec{x}-\vec{u}$ becomes anywhere smaller than 0.5. Iterations are continued as long as mutual information $I(\vec{u})$ increases, with the maximum number of iterations arbitrarily set to 180. In most experiments, convergence is declared after fewer than 80 iterations when the criterion starts to oscillate.

The method was implemented in MATLAB, with image resampling, histogram computation and image gradient computation coded in C. Computation time for matching two images of size $128\times128\times80$ is about 20 min on a Linux PC with a PIII 800 MHz processor.

2.5. Example

Fig. 2 illustrates the algorithm for an MR image of the brain, for registration of the original image to an artificially deformed copy thereof. Parameter values are $h=4.5$, $\Delta u=0.6$ and $s=3$. The template, target and deformed target images are shown on the first row in Fig. 3. Fig. 2 shows the joint histogram of both images at one particular iteration during the registration process, together with the gradient weighting factor $(\partial\psi_h/\partial i_1)L_{\vec{u}}$ and the resulting X-component of the force field $\vec{F}$, velocities $\vec{v}$ and perturbation $\vec{R}$ of the deformation field. The stable and gradual increase of mutual information during iterations is shown in Fig. 2. Convergence is reached after 36 iterations. The small discontinuities that can be observed are due to regridding.

3. Results

The accuracy of the method and the impact of the implementation parameters on registration performance was investigated on simulated MR brain images, generated by the BrainWeb MR simulator (Cocosco et al., 1997) with different noise levels. The images were downsampled by a factor of 2 in each dimension to a $128\times128\times80$ grid of 2 mm^3 isotropic voxels in order to limit computer memory requirements and to increase speed performance. In all experiments the images were non-linearly deformed by known deformation fields $\vec{T}^*$ as illustrated in Fig. 3. These were generated by using our viscous fluid method (with optimal parameter values as determined in Section 3.1) to match the T1 weighted BrainWeb image to real T1 weighted images of three Periventricular Leukomalacia (PVL) patients (see Section 3.3), typically showing enlarged ventricles. The maximal displacement in $\vec{T}^*$ within the brain region is about 7 voxels. In each registration experiment, the recovered deformation field $\vec{T}$ is compared with the ground truth $\vec{T}^*$ by the root mean square (RMS) error ΔT evaluated in voxels over all brain voxels B:

$$\Delta T = \sqrt{\frac{1}{N_B}\sum_B (|\vec{T}(\vec{x}) - \vec{T}^*(\vec{x})|)^2}. \tag{22}$$

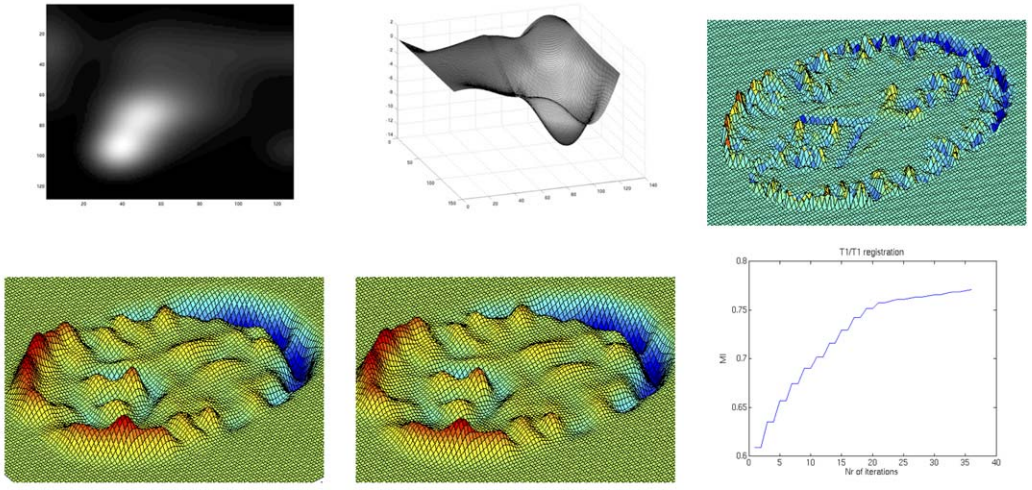

Fig. 2. Top: joint histogram (left), gradient weighting factor (middle) at a particular stage during registration and resulting X-component of force field (right). Bottom: velocities (left), perturbation to the displacement field (middle) and mutual information versus number of iterations (right). (This figure is available in colour, see the on-line version.)

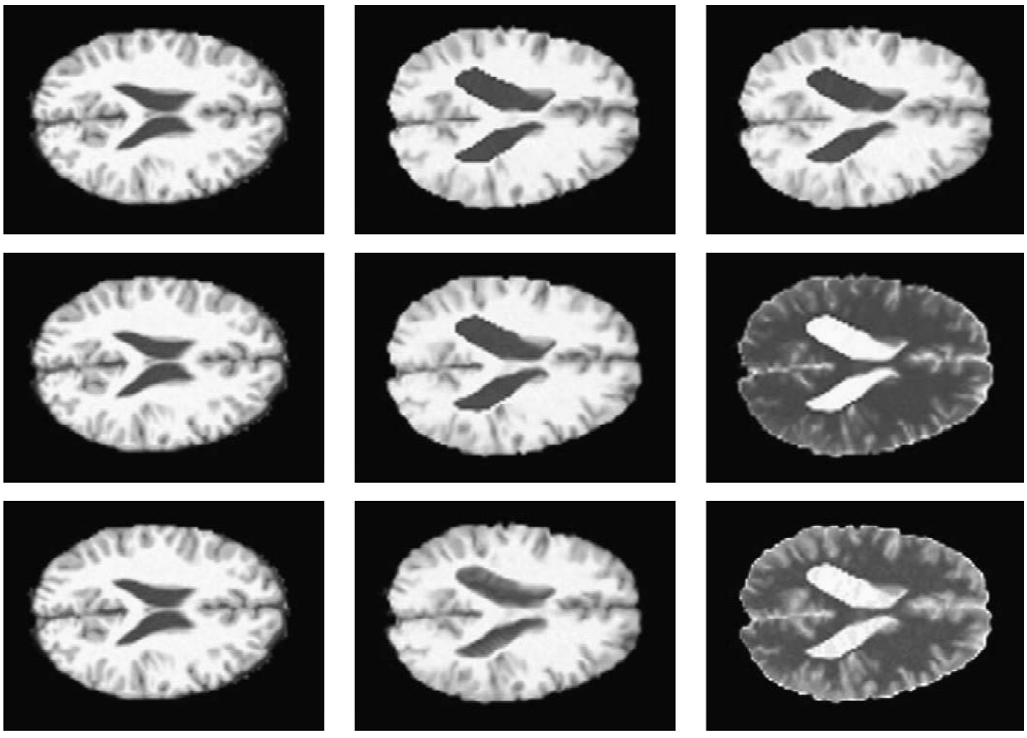

Fig. 3. Left: original BrainWeb T1 template; right: BrainWeb target image obtained by applying a known deformation; middle: template matched to target. Top: T1/T1 registration, middle: T1/T2; bottom: T1/PD.

3.1. Influence of parameters on registration performance

Parameters relevant for registration performance are: the width h of the Gaussian Parzen windowing kernel ψ_h used in (20) to estimate the joint intensity probability; the width s of the spatial smoothing kernel ϕ_s used in (3) to derive deformation velocities $\vec{v}$ from the force field $\vec{F}$; and the maximal allowed displacement in one iteration Δu, that determines the time step parameter Δt used in (4) by the relation (21). The sensitivity of the performance of our viscous fluid registration algorithm with respect to these parameter values was evaluated for the template and target images shown on in Fig. 3. Because the behavior was similar for all three multi-spectral combinations, we only present the results for T1/T1 registration.

Table 1 tabulates the RMS error ΔT and the CPU time until convergence for various values of the width h of the Parzen window kernel ψ_h used to estimate the joint probability density of corresponding voxel intensities, with fixed $\Delta u = 0.6$ and $s = 3$. The registration error is minimal for $h = 5$, but varies by not more than 1% within the range $4.5 \leqslant h \leqslant 5.5$. The optimal value for h estimated using the leave-one-out cross validation scheme was $h = 4.5$. For this value of h, the CPU time is maximal: for an optimal value of h, the algorithm iterates longer and reaches a more accurate registration solution.

Table 2 shows the RMS error and CPU time for various values of Δu, with fixed $s = 3$ and $h = 4.5$. For values of Δu up to 0.6, the registration error does not vary much with Δu, but for $\Delta u = 0.7$ the error sharply increases and remains equally high for values of $\Delta u > 0.7$. The CPU time gradually decreases with increasing Δu up to $\Delta u = 0.6$,

Table 1
Impact of the width h of the Parzen kernel ψ_h on RMS error ΔT (in voxels) and CPU time (in min) ($\Delta u = 0.6$, $s = 3$)

h	3	4	4.5	5	5.5	6	7	8	9
ΔT	0.53	0.46	0.403	0.399	0.404	0.41	0.46	0.53	0.60
CPU	27	26	32	29	28	26	25	23	22

Table 2
Impact of the maximal displacement Δu allowed in one iteration (in voxels) on RMS error ΔT (in voxels) and CPU time (in min) ($h = 4.5$, $s = 3$)

Δu	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
ΔT	0.43	0.40	0.41	0.42	0.41	0.42	1.12	1.12	1.12	1.12
CPU	45	33	29	27	24	19	99	91	92	90

sharply increases at $\Delta u = 0.7$ and then remains more or less constant for $\Delta u > 0.7$. This behavior can be explained by the fact that for large values of Δu, and hence large time step Δt, the computed deformation is more likely to be singular at each iteration, such that the need for regridding sharply increases. For $\Delta u \geqslant 0.7$, the algorithm keeps regridding the deformation field without converging to the registration solution and stops when the maximum number of iterations is reached. For smaller values of Δu, convergence is slower and more iterations are required. An optimal balance between speed of convergence and need for regridding is obtained for $\Delta u = 0.6$.

Table 3 shows the RMS error and CPU time for various values of the width s of the spatial smoothing kernel ϕ_s, with fixed $\Delta u = 0.6$ and $h = 4.5$. Optimal performance, both in terms of registration error and CPU time, is obtained for $s = 3$ voxels. With smaller values of s, registration evidence picked up by the force field $\vec{F}$ at one site is insufficiently propagated to neighboring sites and the algorithm fails to converge. For larger values of s, the force field $\vec{F}$ that drives the registration is overly smoothed and small scale registration features are obliterated, such that the optimum of the registration criterion to which the algorithm converges does not correspond to the true registration solution.

3.2. Validation of accuracy

The accuracy of the method is validated by matching the T1 weighted BrainWeb image to the T1, T2 or proton density (PD) weighted images that were artificially deformed by $\vec{T}^*$. The difference between the recovered deformation $\vec{T}$ and the ground-truth $\vec{T}^*$ is evaluated by the RMS error defined in (22).

We also verified the impact of possible registration errors on atlas-based segmentation. The gray matter (GM), white matter (WM) and cerebro-spinal fluid (CSF) tissue maps of the original template images are obtained using the model-based pixel classification approach of (Van Leemput et al., 1999). These probability maps are deformed onto the target image using the transformations $\vec{T}$ and $\vec{T}^*$, respectively, and subsequently hard-classified by assigning each voxel its most likely tissue class to obtain binary tissue maps M and M^*, respectively. The difference between M and M^* is measured by their overlap coefficient $O_j(M, M^*)$ for each of the three tissue types j (WM, GM, CSF):

$$O_j(M, M^*) = \frac{2V_j(M, M^*)}{V_j(M) + V_j(M^*)}, \tag{23}$$

Table 4
RMS error ΔT (in voxels) between ground truth and recovered deformation fields within the brain region for T1/T1, T1/T2 and T1/PD registration of BrainWeb simulated MR brain images at different noise levels in three cases with different ground-truth deformations

Noise level	T1/T1			T1/T2			T1/PD
	0%	3%	7%	0%	3%	7%	3%
Case 1	0.384	0.430	0.465	0.577	0.759	0.685	0.723
Case 2	0.304	0.398	0.433	0.443	0.640	0.649	0.661
Case 3	0.351	0.411	0.459	0.505	0.753	0.775	0.772

with $V_j(M, M^*)$ the volume of the voxels that are assigned to class j in both maps and $V_j(M)$ and $V_j(M^*)$ the volume of the voxels assigned to class j in each map separately. The sensitivity SE_j and specificity SP_j of the direct registration-based tissue segmentation was evaluated by the number of true positive and true negative labeled voxel fractions within each class, respectively:

$$\mathrm{SE}_j = \frac{V_j(M, M^*)}{V_j(M^*)}, \tag{24}$$

$$\mathrm{SP}_j = \frac{V_{\bar{j}}(M, M^*)}{V_{\bar{j}}(M^*)}, \tag{25}$$

with $\bar{j}$ indicating voxels not being assigned to class j.

Fig. 3 illustrates the registration of the T1 weighted image to the artificially deformed T1, T2 and PD weighted images. Table 4 shows the RMS error ΔT computed for T1 to T1, T2 and PD registration of the BrainWeb images at different noise levels (each time identical for object and target images), for 3 different ground truth deformations. All values are smaller than 1 voxel, with the most accurate results being obtained for T1/T1-matching. The overlap coefficients for WM, GM and CSF in the ground truth and recovered tissue maps are tabulated in Table 5. Specificity and sensitivity are reported for each tissue class in Table 6. The registration-based classification is illustrated in Fig. 4.

3.3. Application: atlas-based segmentation in case of gross morphological differences

PVL is softening of the brain tissue near the ventricles, starting in premature babies. This softening occurs because brain tissue in that area has died due to ischemia. In patients affected by PVL, the brain presents gross morphological differences compared with a normal brain. These are most apparent in the ventricles that are strongly

Table 3
Impact of the width s (in voxels) of the spatial smoothing kernel ϕ_s on RMS error ΔT (in voxels) and CPU time (in min) ($h = 4.5$, $s = 3$)

s	1	2	2.5	3	4	5	6	7	8	9
ΔT	1.12	1.12	1.12	0.42	0.46	0.64	0.70	0.79	0.88	0.98
CPU	84	86	87	19	20	24	20	22	20	22

Table 5
Overlap coefficient for different tissue classes of tissue maps obtained with ground truth and recovered deformation fields for T1/T1, T1/T2 and T1/PD registration of BrainWeb simulated MR brain images in three cases with different ground-truth deformations

	T1/T1			T1/T2			T1/PD		
	WM	GM	CSF	WM	GM	CSF	WM	GM	CSF
Case 1	0.928	0.918	0.870	0.860	0.832	0.765	0.860	0.845	0.784
Case 2	0.925	0.928	0.897	0.860	0.846	0.784	0.856	0.837	0.782
Case 3	0.928	0.926	0.880	0.846	0.827	0.758	0.855	0.803	0.7413

Noise level was 3% in each case.

enlarged. Brain tissue segmentation of such images using automated atlas-guided intensity-based pixel classification as described in (Van Leemput et al., 1999) may fail, if the gross morphological differences between atlas and patient images are not corrected for. Indeed, the initial classification, derived from a digital brain atlas which was constructed from normal brain images and which is globally aligned with the patient image by affine registration, labels large portions of the enlarged ventricles as WM. The initial estimate of the tissue intensity parameters (i.e. mean and spread) is thus not reliable and it is, therefore, unlikely that the iterative segmentation process, which alternates between pixel classification and parameter estimation, converges to the correct segmentation solution. Viscous fluid based registration of atlas and patient images using the method presented above is illustrated in Fig. 5. The original 3-D T1-weighted MR patient brain image was downsampled to $128\times128\times80$ 2 mm^3 voxels. Registration with the BrainWeb T1, T2 or PD-weighted atlas image was performed with optimal values of $\Delta u = 0.6$ and $s = 3$ as derived in Section 3.1. Registration converged after typically 40 iterations. The result of CSF segmentation obtained with the method of (Van Leemput et al., 1999) with affine and with viscous fluid atlas-to-patient matching is shown in Fig. 6. Note how the segmentation of the enlarged ventricles is much improved by using viscous fluid based atlas warping. The use of the viscous fluid regularizer was indeed necessary to get good results.

Table 6
Sensitivity (SE) and specificity (SP) of the proposed method in recovering WM GM and CSF, for T1/T1, T1/T2 and T1/PD registration of BrainWeb simulated MR brain images in three cases with different ground-truth deformations

	T1/T1		T1/T2		T1/PD	
	SE	SP	SE	SP	SE	SP
Case 1						
WM	0.932	0.973	0.884	0.922	0.925	0.895
GM	0.962	0.893	0.874	0.818	0.848	0.854
CSF	0.957	0.985	0.928	0.935	0.922	0.930
Case 2						
WM	0.928	0.984	0.894	0.936	0.928	0.905
GM	0.969	0.900	0.867	0.837	0.861	0.826
CSF	0.967	0.981	0.938	0.913	0.891	0.939
Case 3						
WM	0.935	0.977	0.891	0.920	0.945	0.888
GM	0.951	0.909	0.832	0.828	0.837	0.789
CSF	0.971	0.969	0.935	0.886	0.848	0.936

Noise level was 3% in each case.

4. Discussion

We present an algorithm for non-rigid multimodal image registration using a viscous fluid model by defining a force field that drives the deformation such that mutual information of corresponding voxel intensities is maximized. Our method is in fact the merger of the mutual information based registration functional presented in (Hermosillo et al., 2001) with the viscous fluid regularization scheme of (Christensen et al., 1996b). Hence, our method is an extension of the approach of (Christensen et al., 1996b) to multi-modal data and of the approach of (Hermosillo et al., 2001) to cases with large deformations.

The joint intensity probability of the images to be matched is estimated using Parzen windowing and is differentiable with respect to the deformation field. The size of the Parzen windowing kernel needs to be properly chosen such that the criterion is a more or less smooth function of the deformation field. This choice is related to the image noise. The extension of the Parzen estimator is automatically computed using a leave-*k*-out cross validation technique maximizing an empirical likelihood of the marginal densities (Hermosillo, 2002; Turlach, 1993).

Another relevant implementation parameter is the time step Δt or the maximal displacement Δu allowed at each iteration that is specified to update the displacements after solving the Navier–Stokes equation. Selecting a larger value for Δt will result in larger displacement steps and a more frequent regridding of the template as the Jacobian of the transformation is more likely to become non-positive. A smaller value of Δt on the other hands implies a larger number of iterations for convergence. An appropriate choice for the parameter Δt is related to the width s of the spatial smoothing kernel used to compute the velocities, given the driving forces. Optimal values for $\Delta u = 0.6$ and $s = 3$ were experimentally determined for the BrainWeb simulated MR images ($128\times128\times80$ grid, 2 mm^3 voxels). Experiments with actual patient images as illustrated in Section 3.3 showed that these values where also generally

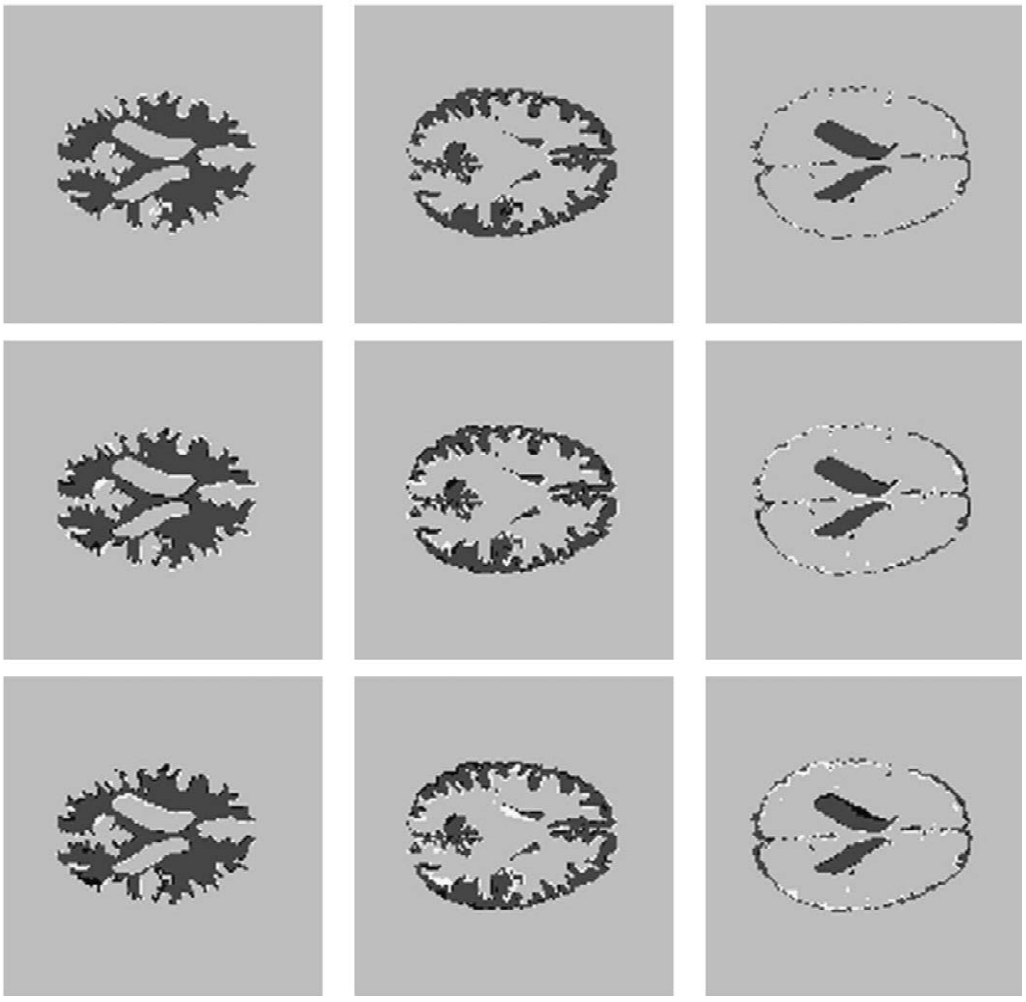

Fig. 4. Mismatched WM (left), GM (middle) and CSF (right) voxels of atlas-based segmentation using the recovered versus the ground truth deformation for the images in Fig. 3. Top: T1/T1 registration; middle: T1/T2; bottom: T1/PD. Voxels are identified using different colors as true negatives (TN, light gray), true positives (TP, dark gray), false positives (FP, white) and false negatives (FN, black). For each class c, these are defined, respectively, as: TP = c & c^*, TN = $\bar{c}$ & $\bar{c}^*$ FP = c & $\bar{c}^*$ and FN = $\bar{c}$ & c^*, where * indicates the ground truth and $\bar{c}$ is the complement to class c.

appropriate for real T1, T2 and PD brain images with similar dimensions and voxel sizes, yielding quick and stable convergence to a visually accurate registration result.

We validated our algorithm using simulated T1, T2 and PD images from BrainWeb with different noise levels and different realistic ground truth deformations generated by registration of the simulated image with real patient images. Although the RMS error was found to be subvoxel small in all cases, T1/T1 registration gave more accurate results than T1/T2 or T1/PD registration. The contrast between GM and WM especially is much better in T1 than

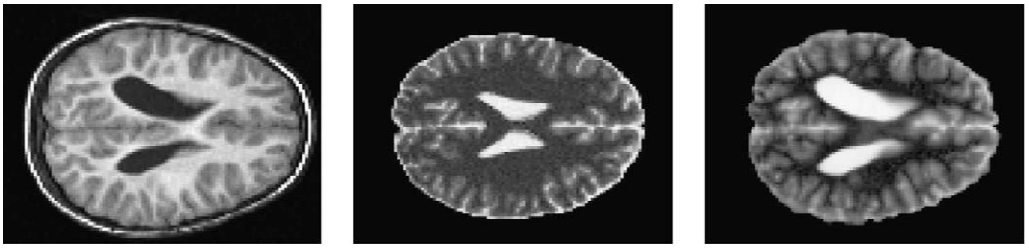

Fig. 5. Left: T1-weighted PVL patient brain MR image, showing enlarged ventricles (target image); middle: original T2-weighted BrainWeb atlas image (template image); right: atlas image warped to patient image (iteratively deformed template image). The blurring of the deformed template image is caused by repeated trilinear interpolation in subsequent iterations of the algorithm.

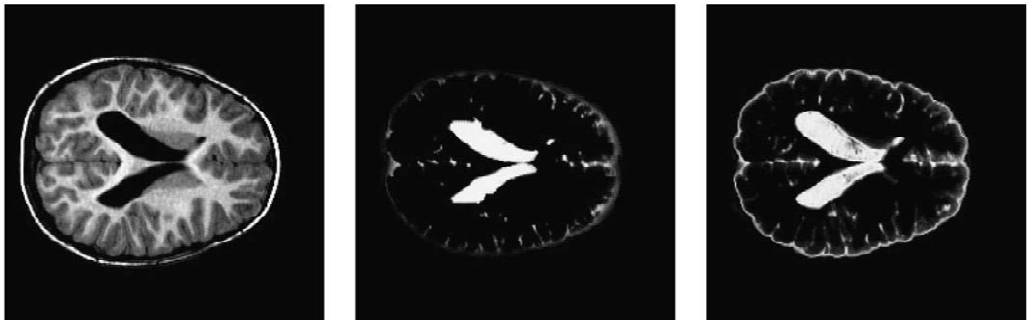

Fig. 6. Left: T1 MPRAGE patient image; middle: CSF segmented by atlas-guided intensity-based tissue classification using the method of (Van Leemput et al., 1999) with affine atlas-to-patient registration; right: CSF segmented after viscous fluid matching of the atlas.

in T2 or PD and the algorithm succeeds better at recovering the interface between both tissues in T1 than in T2 or PD. We also compared T1-with-T2 versus T2-with-T1 registration and found that somewhat better results are obtained using T1 as the template image. This can be explained by the fact that the forces driving the registration depend on the gradient of the template image, which is better defined in T1 than in T2 at the interface between white and GM.

In this paper, we focussed on non-rigid inter-subject MR brain image registration for which the viscous fluid regularization scheme is well suited because of its ability to recover large deformations. Nevertheless, the same mutual information force field may be combined with other regularizers, such as elastic or curvature-based kernels (Hermosillo et al., 2001; D'Agostino et al., 2003). Fig. 7 shows the result for warping the T1-weighted BrainWeb images shown on the top row of Fig. 3 using both the viscous fluid scheme presented here and the linear elastic model of (Christensen et al., 1996a). The elastic model is obtained by replacing the deformation velocity $\vec{v}$ in (1) by the displacement $\vec{u}$ and smoothing $\vec{u}$ instead of $\vec{v}$ at each iteration using a similar Gaussian kernel. As shown in Fig. 7, the deformation obtained with the elastic regularizer is clearly much more smooth than that obtained with the viscous fluid scheme, especially around the ventricles and the brain surface. A comparative evaluation of the performance of various regularization schemes in different applications is outside the scope of this paper.

5. Conclusions

We have presented a multimodal free-from registration algorithm based on MMI that models the images as a viscous fluid. The forces deforming the images are defined

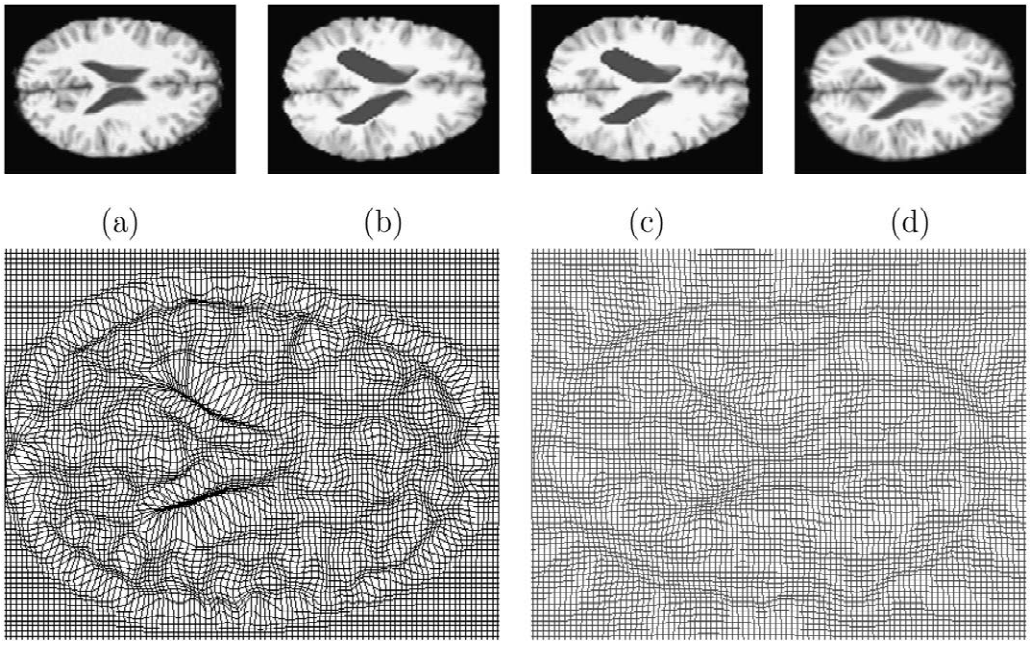

Fig. 7. Top: original template image (a) deformed to target image (b) using the viscous fluid regularizer (c) presented here and the elastic regularization scheme of (Christensen et al., 1996a) (d). Bottom, deformation fields obtained with viscous fluid (left) and elastic (right) regularizers.

as the gradient of mutual information with respect to the deformation field, using Parzen windowing to estimate the joint intensity probability. We have validated our method for matching simulated T1/T1, T1/T2 and T1/PD MR brain images, showing that the method performs quite well in both mono and multi-modal conditions. We illustrated the benefit of viscous fluid inter-subject MR brain image registration for atlas-based brain tissue segmentation in case of gross morphological differences between patient and atlas images. Future work includes the introduction of more spatial information and more specific intensity models into the similarity criterion in order to make the registration more robust.

References

Bro-Nielsen, M., Gramkow, C., 1996. Fast fluid registration of medical images. In: Kikinis, R., Hoehne, K. (Eds.). Visualization in Biomedical Computing (VBC'96). Lecture Notes in Computer Science, vol. 1131. Springer, Hamburg, Germany, pp. 267–276.

Christensen, G., Miller, M., Grenander, U., Vannier, M., 1996a. Individualizing neuroanatomical atlases using a massively parallel computer. IEEE Computer 29 (1), 32–38.

Christensen, G., Rabbitt, R., Miller, M., 1996b. Deformable templates using large deformation kinetics. IEEE Transactions on Image Processing 5 (10), 1435–1447.

Cocosco, C., Kollokian, V., Kwan, R.-S., Evans, A., 1997. Brainweb: online interface to a 3d MRI simulated brain database. NeuroImage 5 (4, part 2/4), S425, Proceedings of Third International Conference on Functional Mapping of the Human Brain. Copenhagen, May 1997.

D'Agostino, E., Modersitzki, J., Maes, F., Vandermeulen, D., Fischer, B., Suetens, P., June 2003. Free-form registration using mutual information and curvature regularization. In: Workshop on Biomedical Image Registration (WBIR'03). Philadelphia, PE, USA, in press.

Gaens, T., Maes, F., Vandermeulen, D., Suetens, P., 1998. Non-rigid multimodal image registration using mutual information. In: Wells, W., Colchester, A., Delp, S. (Eds.). Medical Image Computing and Computer-Assisted Intervention (MICCAI'98). Lecture Notes in Computer Science, vol. 1496. Springer, Berlin, Cambridge, MA, pp. 1099–1106.

Hermosillo, G., May 2002. Variational methods for multimodal image matching. Ph.D. thesis. Université de Nice (INRIA-ROBOTVIS), Sophia Antipolis, France.

Hermosillo, G., Chef d'Hotel, C., Faugeras, O., February 2001. A variational approach to multi-modal image matching. Tech. Rep. 4117, INRIA-ROBOTVIS, Sophia Antipolis, France.

Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P., 1997. Multi-modality image registration by maximization of mutual information. IEEE Transactions on Medical Imaging 16 (2), 187–198.

Maintz, J., Viergever, M., 1998. A survey of medical image registration. Medical Image Analysis 2 (1), 1–36.

Meyer, C., Boes, J., Kim, B., Bland, P., Wahl, R., Zasadny, K., Kison, P., Koral, K., Frey, K., 1997. Demonstration of accuracy and clinical versatility of mutual information for automatic multimodality image fusion using affine and thin plate spline warped geometric deformations. Medical Image Analysis 1 (3), 195–206.

Pluim, J., Maintz, J., Viergever, M., 2000. Image registration by maximization of combined mutual information and gradient information. IEEE Transactions on Medical Imaging 19 (8), 809–814.

Rohde, G., Aldroubi, A., Dawant, B., 2001. Adaptive free-form deformation for interpatient medical image registration. In: Sonka, M., Hanson, K. (Eds.). Medical Imaging 2001: Image Processing. Proceedings of the SPIE, vol. 4322. SPIE Press, Bellingham, WA; San Diego, CA, USA, pp. 1578–1587.

Rueckert, D., Sonoda, L., Hayes, C., Hill, D., Leach, M., Hawkes, D., 1999. Nonrigid registration using free-form deformations: application to breast MR images. IEEE Transactions on Medical Imaging 18 (8), 712–721.

Schnabel, J., Rueckert, D., Quist, M., Blackall, J., Castellano-Smith, A., Hartkens, T., Penney, G., Hall, W., Liu, H., Truwit, C., Gerritsen, F., Hill, D., Hawkes, D., 2001. A generic framework for non-rigid registration based on non-uniform multi-level free-form deformations. In: Niessen, W., Viergever, M. (Eds.). Medical Image Computing and Computer-Assisted Intervention (MICCAI '01). Lecture Notes in Computer Science, vol. 2208. Springer, Berlin, Utrecht, Netherlands, pp. 573–581.

Studholme, C., Hill, D., Hawkes, D., 1999. An overlap invariant entropy measure of 3D medical image alignment. Pattern Recognition 32 (1), 71–86.

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

Turlach, B., 1993. Bandwidth selection in kernel density estimation: a review. Discussion paper 9317, Institut de Statistique, UCL, Louvain la Neuve, Belgium.

Van Leemput, K., Maes, F., Vandermeulen, D., Suetens, P., 1999. Automated model-based tissue classification of MR images of the brain. IEEE Transactions on Medical Imaging 18 (10), 897–908.

Wang, Y., Staib, L., 2000. Physical model-based non-rigid registration incorporating statistical shape information. Medical Image Analysis 4 (1), 7–20.

Wells, III W., Viola, P., Atsumi, H., Nakajima, S., Kikinis, R., 1996. Multi-modal volume registration by maximization of mutual information. Medical Image Analysis 1 (1), 35–51.

West, J., Fitzpatrick, J., Wang, M., Dawant, B., Maurer, Jr C., Kessler, R., Maciunas, R., Barillot, C., Lemoine, D., Collignon, A., Maes, F., Suetens, P., Vandermeulen, D., van den Elsen, P., Napel, S., Sumanaweera, T., Harkness, B., Hemler, P., Hill, D., Hawkes, D., Studholme, C., Maintz, J., Viergever, M., Malandain, G., Pennec, X., Noz, M., Maguire, Jr G., Pollack, M., Pellizari, C., Robb, R., Hanson, D., Woods, R., 1997. Comparison and evaluation of retrospective intermodality brain image registration techniques. Journal of Computer Assisted Tomography 21, 554–566.