

An Efficient Locally Affine Framework for the Smooth Registration of Anatomical Structures

O. Commowick^{a,b,*} V. Arsigny^a A. Isambert^c J. Costa^a
F. Dhermain^d F. Bidault^e P.-Y. Bondiau^f N. Ayache^a
G. Malandain^a

^a*INRIA Sophia Antipolis - ASCLEPIOS Team,
2004 route des lucioles, BP93, 06902 Sophia Antipolis Cedex, FRANCE*

^b*DOSIsoft S.A., 45 avenue Carnot, 94230 Cachan, FRANCE*

^c*Service de Physique, Institut Gustave Roussy, Villejuif, FRANCE*

^d*Département de Radiothérapie, Institut Gustave Roussy, Villejuif, FRANCE*

^e*Département d'Imagerie Médicale, Institut Gustave Roussy, Villejuif, FRANCE*

^f*Département de Radiothérapie, Centre Antoine Lacassagne, Nice, FRANCE*

Abstract

Intra-subject and inter-subject non linear registration based on dense transformations requires the setting of many parameters, mainly for regularization. This task is a major issue, as the global quality of the registration will depend on it. Setting these parameters is however very hard, and they may have to be tuned for each patient when processing data acquired by different centers or using different protocols. Thus, we present in this article a method to introduce more coherence in the registration by using fewer degrees of freedom than with a dense registration. This is done by registering the images only on user-defined areas, using a set of affine transformations, which are optimized together in a very efficient manner. Our framework also ensures a smooth and coherent transformation thanks to a new regularization of the affine components. Finally, we ensure an invertible transformation thanks to the Log-Euclidean polyaffine framework. This allows us to get a more robust and very efficient registration method, while obtaining good results as explained below.

We performed a qualitative and quantitative evaluation of the obtained results on two applications: first on atlas-based brain segmentation, comparing our results with a dense registration algorithm. Then the second application for which our framework is particularly well suited concerns bone registration in the lower abdomen area. We obtain in this case a better positioning of the femoral heads than with a dense registration. For both applications, we show a significant improvement in computation time, which is crucial for clinical applications.

Key words: non linear registration, locally affine transformation, Log-Euclidean regularization, atlas-based brain segmentation

1 Introduction

Non linear registration of medical images can be used for various purposes. It allows, for example, to quantify the evolution of a disease over time [1–3]. It can also be used to compare patients between them or with respect to an anatomical atlas. In the latter case, one can use non linear registration to compute statistics with respect to a reference anatomy [4,5] or perform atlas-based segmentation [6,7] using labeled structures in the atlas to segment a patient image.

The main requirements for the use of non linear registration for atlas-based segmentation in a clinical context are divided in three points:

- (1) first, the non linear transformation must be sufficiently generic to take into account the atlas-patient variability,
- (2) however, this transformation must also be as regular as possible so that the delineated structures have meaningful contours,
- (3) finally, the registration method should also be robust, i.e. we should be able to use always the same set of parameters to delineate patients, whose images were acquired in different centers.

The first two requirements are valid for all non linear registration methods while the third one is closely related to the use in a clinical context. Clinicians are indeed usually not experts in image registration. It is therefore mandatory to be able to use a default parameter set for all patients and all acquisition protocols. In addition to this last requirement, a reasonable computational time would also be desirable.

Typically, the registration problem is solved by designing a similarity measure S (typically SSD, correlation coefficient or mutual information) that assesses the quality of the registration. In the following, we denote by R a reference image (typically an atlas image) on which we want to superimpose a floating image F using a transformation T . Then, following Brown et al. [8], the global transformation is computed by:

$$\hat{T} = \arg \min_T S(R, F \circ T). \quad (1)$$

* Corresponding author. Phone: +1 617-355-5382. Fax: +1 617-730-0635
Email address: ocommowi@gmail.com (O. Commowick).
URL: <http://olivier.commowick.org/> (O. Commowick).

Since a closed form of the solution is usually not available, this measure is optimized iteratively, i.e. a sequence of transformations $T^0, T^1, \dots, T^l, \dots, T^N$ is computed, where $T^l = T^{l-1} \circ \delta T^l$. Therefore, at each iteration l , a correction δT^l to the current transformation T^{l-1} is computed using the following equation:

$$\delta T^l = \arg \min_{\delta T} S(R, F \circ T^{l-1} \circ \delta T). \quad (2)$$

To address the first requirement, the transformation T must have a large number of degrees of freedom, for example a dense transformation made of one displacement vector per image voxel. The minimization of Eq. (2) is however likely to be trapped in a local minimum, due for example to noise in the images. To address this issue and also the second requirement, regularization terms are explicitly added to the similarity term in Eq. (2), yielding:

$$\delta T^l = \arg \min_{\delta T} S(R, F \circ T^{l-1} \circ \delta T) + \alpha E_{el}(T^{l-1} \circ \delta T) + \beta E_{fl}(\delta T) \quad (3)$$

The regularization can be done either on δT (this is the fluid regularization, popularized by [9]) or on the transformation itself T^l (this is the elastic regularization, used for example in [10]): the corresponding energy terms are respectively denoted by $E_{fl}(\delta T)$ and $E_{el}(T^{l-1} \circ \delta T)$ in Eq. (3). Other examples of elastic regularization include Bajcsy and Broit [11,12] who introduced methods where real world elastic deformations were computed. Bajcsy [13] extended later these methods to a multi-resolution scheme. Finally, one can use inhomogeneous visco-elastic regularization [14], depending on the deformability characteristics of the tissues in the images. For this last method, defining the deformability of the tissues is however a difficult problem. It has been studied for brain structures [15], but not yet for other parts of the body, where very deformable soft tissues are present.

In addition, the regularity of the transformation can be further constrained by decreasing a priori the number of degrees of freedom of its parameterization. One can use, for example, a linear combination of regularly or irregularly placed radial basis functions [16] or B-Splines [17]. However, there is a trade-off using these methods between the number of degrees of freedom and obtaining a precise deformation in a reasonable time. Other works parameterize the transformation using vortex particles [18] or geodesic interpolating splines [19]. An arbitrary number of degrees of freedom can be assigned to these transformations. We have to point out that these methods (e.g. [17,18]) do not guarantee the invertibility of the transformation, which will penalize their implementation for applications such as atlas-based brain segmentation.

These approaches reduce the dependency with respect to the regularization

parameters, but do not suppress it. A more drastic reduction of the transformation parameters may yield to the realization of the third requirement but it will also compromise the realization of the first requirement. We are however only interested in registering known structures. In this context, both the first and third requirements can be satisfied by designing an ad-hoc transformation on these regions.

Several approaches which register structures of interest in the images using few degrees of freedom have been proposed in the literature. Little et al. [20] introduced first a way to interpolate between local affine transformations, also used by Pitiot et al. [21]. Some methods have also been developed specifically for the registration of articulated structures [22,23]. They register bones but using very specific transformations, coherent with the movements in the images (rotations with respect to articulations). However, for all these methods, the affine transformations are computed independently, yielding to possible discrepancies. More recently in [24,25], more sophisticated frameworks have been introduced, allowing to compute a global invertible transformation from a set of local affine transformations and optimize it to match 2D images. However, calculations using these frameworks are costly in 2D and their generalization to 3D is again not straightforward.

In this article, we choose to use a meaningful parameterization of the transformation based on the structures present in the images. We can indeed define them easily and once and for all on a reference image or on an atlas. We therefore present a general and efficient locally affine framework for registering these fixed regions in 3D images. The registration is done by optimizing together local affine transformations and using a new Log-Euclidean regularization of these transformations, therefore ensuring coherency between them. The global transformation is then parameterized using these affine components A_i , associated to user defined areas R_i . Our framework guarantees, at the end, an invertible and anatomically consistent transformation, thanks to the use of the recently introduced Log-Euclidean polyaffine framework [26].

The remainder of the article is organized as follows: we will first present the method we adopt to register the local affine components defined by the user. Then, we will focus on our new regularization scheme, which allows us to remove discrepancies between the affine transformations of areas close to each other and to have a very smooth transformation.

Afterwards, we will demonstrate the versatility of our method by showing qualitative and quantitative results on two different applications. The first one is the segmentation of brain critical structures using atlas-to-subject registration, showing that the obtained transformations are much smoother with our method than with a dense registration method. Quantitative results using STAPLE validation method [27] also show similar or better registration

results. The second one, for which our framework is particularly well suited, is bone registration for bladder localization in the lower abdomen area. We show in this case better results on the positioning of bones (evaluated through the distance of manually pointed landmarks to those obtained automatically) than with a dense registration.

2 Method

2.1 Global Algorithm

Following the ideas presented in the introduction, we will register a floating image F on a reference image R on specific predefined regions. These regions R_i are given by the user or predefined on an anatomical atlas. The goal of our method is then to compute an affine transformation A_i for each region R_i that best matches this region in the two images. Finally, a global transformation T is interpolated from the affine components A_i .

We choose to follow an iterative framework using a multi-resolution scheme illustrated in Algorithm 1 to compute the transformation on the different specified areas. This framework is similar to the one of [28] for rigid registration, and also to the ICP framework for surface matching [29,30]. At each iteration l , we evaluate transformation corrections δA_i^l thanks to pairings obtained by Block-Matching, so that $F \circ T^l$ gets closer to R than $F \circ T^{l-1}$. To ensure a coherent transformation, we regularize the transformations corrections and the affine transformations themselves at each step of the estimation using a new regularization scheme. We detail all the steps of our algorithm in the following.

Algorithm 1 Overview of the registration algorithm

- 1: Creation of weight functions $\bar{w}_i(x)$ from the R_i . Section 2.2.
 - 2: Initialization of the transformation: $T^0, A_i^0 \leftarrow \text{Id}$.
 - 3: **for** $p = 1 \dots P$, iteration on pyramid levels, **do**
 - 4: **for** $l = 1 \dots L$, iterations, **do**
 - 5: Estimation of affine corrections: $\delta A_i^l \leftarrow \text{register}(R_i, F \circ T^{l-1})$. Section 2.3.
 - 6: Fluid like regularization: $\delta \tilde{A}_i^l \leftarrow \text{regul}(\delta A_i^l)$. Section 2.4.
 - 7: Composition of corrections $A_i^l = \tilde{A}_i^{l-1} \circ \delta \tilde{A}_i^l$.
 - 8: Elastic like regularization: $\tilde{A}_i^l \leftarrow \text{regul}(A_i^l)$. Section 2.4.
 - 9: Computation of T^l : $T^l \leftarrow M_2(\tilde{A}_i^l)$. Section 2.5.
 - 10: Computation of T^{final} : $T^{\text{final}} \leftarrow M_3(\tilde{A}_i^L)$. Section 2.5.
-

2.2 Preprocessing

2.2.1 User Definition of Registered Regions

To use our algorithm, we first need to define the regions R_i we want to register. We define them on the reference image. This image is indeed never deformed and the regions can thus be defined once and for all. In our case, we choose to have entire areas adopting the same affine behavior. For example, on Fig. 1, we want the transformation to be roughly affine for each one. We therefore put only one region over each eye. The areas are specified as binary sub-images.

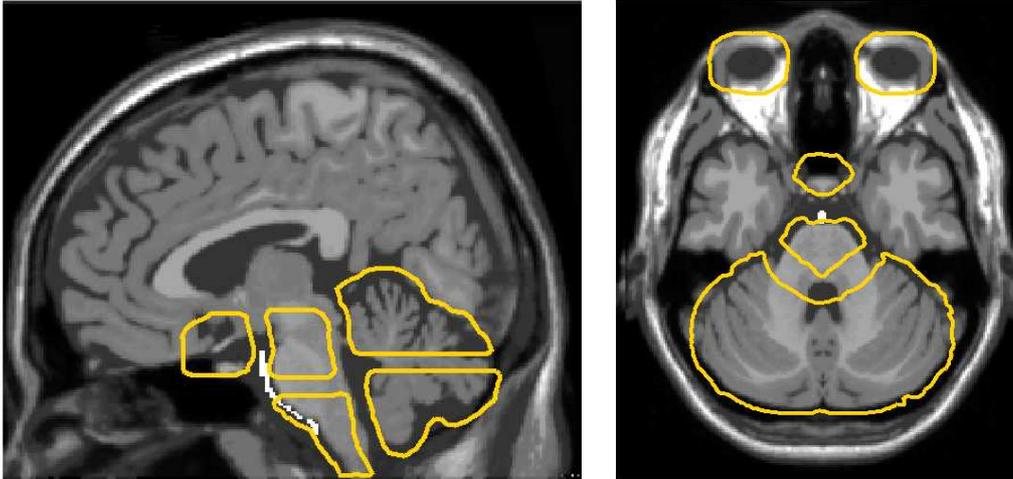


Fig. 1. **Example of predefined areas on the brain.** Each contour corresponds to a region R_i defined by the user on the atlas. The regions shown here are eroded for interpolation following the method of [21].

2.2.2 Ensuring a Minimal Distance between Regions

To ensure a smooth interpolation between the regions we have defined, we need a minimal distance between them. The binary areas can indeed overlap each other after definition. We then need to erode them as little as possible to have a minimal distance between regions while keeping them as close as possible to the original region. A way to cope with this problem has been devised in [21], illustrated on Fig. 2.

First, we superimpose all the sub-images in a single image. A series of erosions then ensures that the areas are disjoint. A distance map is computed in the background of the resulting image and a thinning algorithm is used to extract the skeleton of the background. We are then able to compute a distance to this skeleton. By removing from the binary image the voxels whose distance to the skeleton is less than a threshold ν , we ensure a minimal distance of 2ν between the regions while modifying them as little as possible.

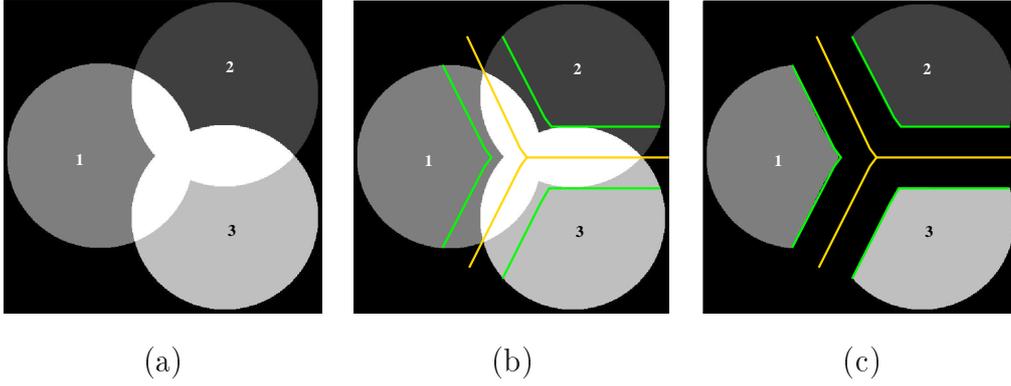


Fig. 2. **Synthetic example of the adaptive erosion of predefined structures.** These images show the process used to erode regions 1, 2 and 3 which are overlapping each other: (a) : original regions; (b) : skeleton and limits of erosion superimposed; (c) : resulting regions after erosion (see text).

2.2.3 Definition of Weighting Functions

The binary sub-images do not overlap anymore. We then associate to each region a weighting function $\bar{w}_i(x)$, defining the relative influence of the i^{th} region at point x . We have chosen to implement the weighting function for each component as a function of the minimal distance to the corresponding region R_i : $w_i(x) = 1/(1 + \alpha \text{dist}(x, R_i))$ (in our algorithm we used $\alpha = 0.5$ to ensure a smoother interpolation). Other choices are indeed possible to compute these weight functions¹.

At each point x , the weights $w_i(x)$ are then normalized so that their sum is equal to 1: $\bar{w}_i(x) = w_i(x) / \sum_{i=1}^N w_i(x)$. Because of this normalization, $\bar{w}_i(x)$ is not constant spatially inside the region R_i .

The R_i and the associated \bar{w}_i are computed on the eroded regions once and for all during the algorithm. Moreover, as we will see in section 3.1, when registering patients on an atlas, the R_i and \bar{w}_i are computed on the atlas and do not depend on the patient's image.

2.3 Updating the Transformation

At each iteration l , we update the transformation T^{l-1} by looking for local affine transformations $A_i^l = A_i^{l-1} \circ \delta A_i^l$ so that we get a better correspondence between the images over the regions R_i . We detail in the following first how we obtain pairings between the current images, and then how to combine them

¹ Convolution of the binary images with a Gaussian kernel may be another possibility to compute the weights.

to get the transformation corrections δA_i^l . For clarity, δA_i^l will be replaced by δA_i in the following.

We follow in this part the framework which was proposed in [28] to estimate a global rigid or affine transformation between two images. We now focus on the main steps of this framework and its application to our problem.

2.3.1 Pairings Estimation

We first need to choose a method to estimate correspondences between the images R and $F \circ T^{l-1}$. We have chosen to use a Block-Matching approach to estimate these correspondences. This method allows to look for big displacements while being robust to local minima in the similarity measure.

In our method, blocks are regularly spaced (typically one block every 3 voxels in each direction) inside the areas R_i defined on the reference image. For each block $B(x_v) \in R_i$, centered in x_v , we look for its best match $B(y_v)$ in $F \circ T^{l-1}$. The choice of the similarity measure to select this best match should depend of the expected relationship between the block intensities. Considering that the blocks are small (typically 5x5x5 or 7x7x7 voxels), they may contain up to two or rarely three different tissues. Assuming an affine relationship between the block intensities seems therefore reasonable [31]. Hence, we choose the local squared correlation coefficient between the two blocks as similarity measure [32], $CC^2(B(x_v), B(y)) = Cov^2(B(x_v), B(y)) / (Var(B(x_v))Var(B(y)))$. This choice also seems good as some structures (like vessels or arteries) may be present in one image and not in the other. The best pairing is therefore chosen following this equation:

$$B(y_v) = \arg \max_{B(y), y \in V(x_v)} CC^2(B(x_v), B(y)), \quad (4)$$

where $V(x_v)$ denotes the local neighborhood of x_v in image $F \circ T^{l-1}$. In the remaining, we will call CC_v the best value of the squared correlation coefficient for the block $B(x_v)$: $CC_v = CC^2(B(x_v), B(y_v))$.

2.3.2 Affine Transformation Estimation

At this step of the process, we have computed a collection of pairings (x_v, y_v) for each area thanks to the Block-Matching algorithm. We are then using these pairings to compute the corrections δA_i . We choose to use a separate optimization of the affine transformations: we consider each area independently. We then use only the pairings $(x_v, y_v)_{v \in R_i}$ in order to estimate δA_i . The system used to obtain the δA_i is given by the following formula, which amounts to a Weighted Least Squares (LSW) problem:

$$\delta A_i = \arg \min_{\delta A} \sum_{v: x_v \in R_i} CC_v \bar{w}_i(x_v) \|\delta A \cdot x_v - y_v\|^2. \quad (5)$$

This transformation estimation can be seen as the first similarity term of Eq. (3) presented in the introduction. In the above formulation, we choose to use $CC_v \bar{w}_i(x_v)$ as a weighting parameter for each equation of the system. We indeed consider that a block with a small $\bar{w}_i(x_v)$, i.e. located between two regions, should influence less the transformation A_i . Moreover, we want to give more importance to terms that are good pairings in order to promote their influence on the solution. CC_v varies between 0 and 1 and gives us for each block this relative goodness of the pairing. We have indeed $CC_v = 1$ when the pairing is perfect and $CC_v = 0$ when the pairing is false. We can thus use it straightforward as a weight in the LSW system.

This method is fast as we divide one large single system into several small systems. We have seen in our experiments that we obtain results similar to a coupled optimization, while running much faster. We thus choose to use this separated method as it approximates well the transformation corrections.

The energy given in Eq. (5) can be solved very efficiently as it leads to a linear over-constrained system. Given that we can have outliers due to noise or missing structures, we choose to use an Least Trimmed Squares Weighted (LTSW) estimator [33]. This minimization scheme has proved to be more robust to outliers than the classical LSW method. At a glance, instead of minimizing the total sum of the squared residuals, we will iteratively minimize the sum of the h smallest squared residuals. This method reduces drastically the influence of the outliers. Moreover, as mentioned before, we use the correlation coefficient values CC_v as weights for each equation of the system.

2.4 Regularization

At this step, we obtain at each iteration l of the algorithm a set of transformation corrections δA_i computed so that $F \circ T^l$ is closer to the reference image R than $F \circ T^{l-1}$. Without any regularization process, we can have some discrepancies in the interpolated areas between the affine components (see Fig. 3). These discrepancies are mainly due to the separate estimation of the affine corrections δA_i .

To avoid these problems, we present here a novel regularization approach, specific to locally affine transformations. The basic idea is to use the 4x4 matrix representation of 3D affine transformation in homogeneous coordinates. Interestingly, whenever the amount of rotation present in an affine transformation A is less than π radians, one can define the matrix logarithm of A , simply via

the principal logarithm of the matrix representing A . For more details on principal matrix logarithms and their practical computation, see [34]. This matrix logarithm is of the form: $\begin{pmatrix} M & v \\ 0 & 0 \end{pmatrix}$, where M is a 3x3 matrix (not necessarily invertible) and v a 3D vector. Conversely, a unique affine transformation is associated to any 4x4 matrix B of the latter form via its matrix exponential.

It has been shown in [26] that taking the logarithm of affine transformations corresponds to linearizing the (curved) affine group around the identity, while conserving excellent theoretical properties (invariance with respect to inversion in particular). This allows to perform Euclidean (i.e. vector-based) operations on affine transformations via their logarithms.

This representation of affine transformations by vectors allows the direct generalization of classical vector-based regularization techniques. For example, we can define a Log-Euclidean fluid-like energy between the affine transformation corrections (which is similar to the second regularization term of Eq. (3)):

$$E_{Reg} = \sum_{i=1}^N \sum_{j=1}^N p_{i,j} \|\log(\delta A_i) - \log(\delta A_j)\|^2, \quad (6)$$

where we have $p_{i,j} = \sum_{x \in R} \bar{w}_i(x) \cdot \bar{w}_j(x) / \sum_{x \in R} \bar{w}_i(x)$, which takes into account the spatial extensions of the components. The regularization is then done using a one fixed step gradient descent on this energy: $\delta \tilde{A}_i = \delta A_i + dt \frac{\partial E_{Reg}}{\partial \delta A_i}$. The step dt is defined to balance between the transformation corrections estimation and the regularization term: $dt = \lambda \frac{1}{N} \sum_{i=1}^N \|\log(\delta A_i)\|$. In the sequel, $\|\cdot\|$ is set to $\|M\|^2 = \text{Trace}(M \cdot M^T)$ (Frobenius norm), which has the advantage of being rotationally invariant.

Furthermore, one can define an elastic-like energy which is similar to the first regularization term of Eq. (3)) by regularizing the transformations A_i^l instead of the transformations corrections δA_i in (6). As we have two regularization terms, we define two λ parameters: $\lambda_{el} \in [0, 1]$ for the elastic regularization and $\lambda_{fl} \in [0, 1]$ for the fluid regularization.

Using this regularization scheme improves greatly the smoothness in the interpolation areas (see Fig. 3). We have used for this example a transformation composed of two areas on the cerebellum. The first locally affine registration was done without any regularization, while the second uses both elastic and fluid regularization ($\lambda_{el} = 0.3$ and $\lambda_{fl} = 0.2$). This example clearly shows the importance of the regularization. The cerebellum does not have the shape we would expect when not using regularization: there is indeed a lack of coherence between the two components yielding a result which is not consistent from an

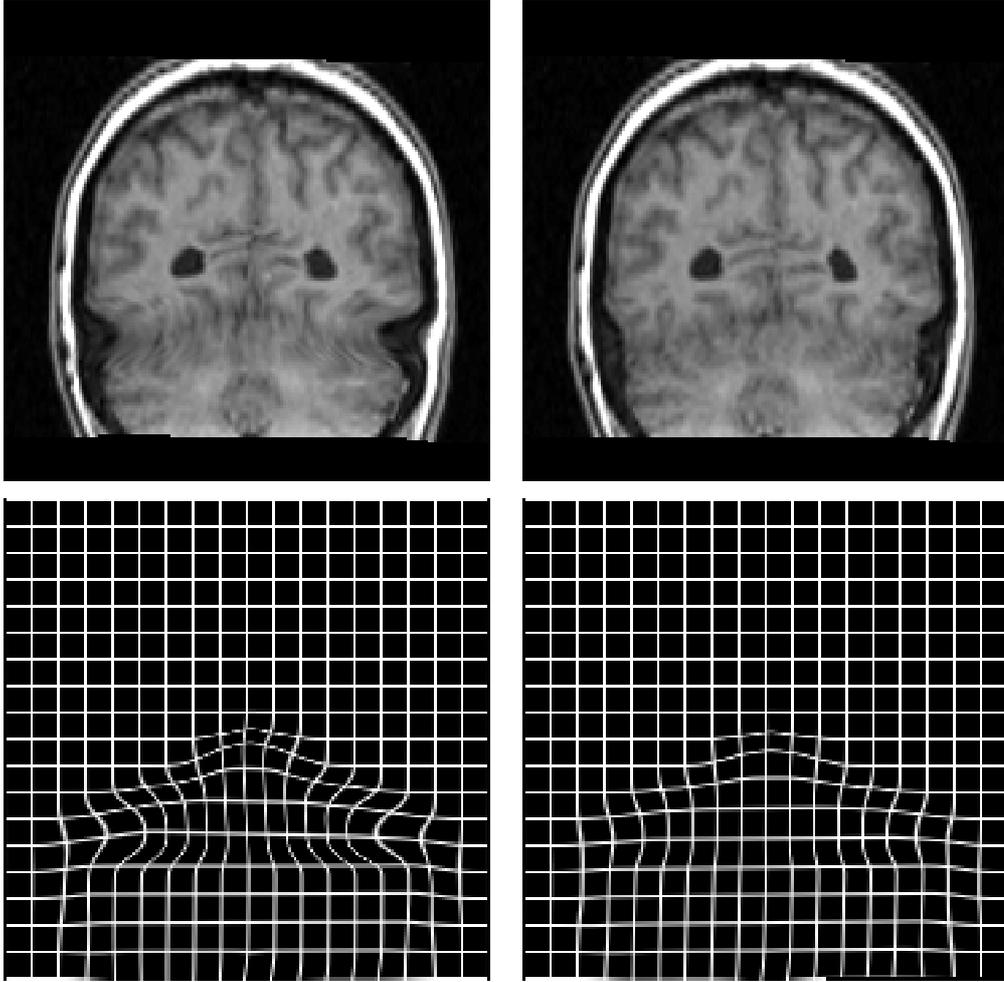


Fig. 3. **Contribution of the regularization in the registration.** From left to right: Patient image registered without regularization ($\lambda_{el} = 0$, $\lambda_{fl} = 0$) and with regularization ($\lambda_{el} = 0.3$, $\lambda_{fl} = 0.2$). Top row: coronal slices. Bottom row: corresponding deformed grids. The discrepancies in the deformation on the left image are removed thanks to our new regularization scheme. (See text for details).

anatomical point of view. The errors also propagate to the rest of the brain. Our regularization technique solves this problem and provides consistent results all over the brain. Furthermore, the regularization energies we propose here are very efficient to use in practice.

2.5 Invertible Transformation

We have detailed so far a fast method to obtain a locally affine transformation matching the image on predefined areas. This is achieved thanks to the use of separated optimization and to the use of the second method of interpolation (M_2) we describe in appendix section A.2. We chose this method of interpo-

lation, called direct averaging, as it is fast to compute and can be efficiently optimized thanks to the existence of a closed form. However, as we point in Appendix A, the transformation obtained using this method is smooth but not always invertible. This does not cause problems for the registration algorithm, as the singularities appear generally in the interpolation areas, outside the R_i . However, this can be a problem for resampling the reference image R on the floating image F , which can be of great use in some cases. We thus need to be able to invert the transformation. To ensure a smooth and invertible transformation everywhere, we therefore choose to use the third method of interpolation (Log-Euclidean polyaffine M_3 , appendix section A.3) to build the final transformation and invert it. This approach is consistent with our estimation method. It has indeed been shown in [26], that this method gives almost the same fusion as with direct averaging in regions without singularities, while removing the singularities elsewhere. However, we only use this method to compute the final transformation as it would be much more computationally expensive to use it directly for the estimation of T .

2.6 Evaluation Methodology

2.6.1 Existing Registration Methods

We have presented so far a complete framework to register images on predefined regions using a locally affine transformation. In the following, we will present experiments on two applications: atlas-based brain critical structures segmentation and lower abdomen registration. In order to compare our algorithm with existing dense registration methods, we used for each application an existing dense registration algorithm.

For brain registration (section 3.1), we used a method which was specifically designed for this task [14], named Runa. This method is very similar to the "Demons" method [35] and to Pasha [10]. The main improvement of Runa with respect to these methods is the use of an inhomogeneous visco-elastic regularization. It minimizes in an iterative process the following energy:

$$\begin{aligned}
 E = & S(R, F \circ T) + \beta \sum_{\alpha \in \{x_0, x_1, x_2\}} \int [1 - k(x)] \left\| \nabla \frac{\partial T_\alpha}{\partial t} \right\|^2 \\
 & + \gamma \sum_{\alpha \in \{x_0, x_1, x_2\}} \int D(x) \|\nabla T_\alpha\|^2
 \end{aligned} \tag{7}$$

where S is a similarity term, typically the SSD. The fluid regularization (second term) is weighted by a confidence term $k(x)$ giving more importance to displacements where the image gradient is high. The elastic regularization

(third term) is also weighted by a term $D(x)$, which depends of the type of tissue in the reference image. The white matter will thus have an elastic behavior, while cerebro-spinal fluid will have a fluid behavior). This allows them to obtain very precise deformations and to be more independent of regularization parameters. More details on this algorithm can be found in [14].

We could not use Runa directly in the case of bone registration. Runa is indeed designed specifically for brain registration and its use for lower-abdomen images would have needed further studies on the inhomogeneous regularization to use. Cachier et al. [10] instead parameterizes the transformation as a dense deformation field and uses homogeneous elastic regularization. This method derives from the Demons method [35] and minimizes iteratively the following energy:

$$E_{Pasha} = S(R, F \circ C) + \sigma \|C - T\|^2 + \sigma \lambda E_{el}(T). \quad (8)$$

where C and T correspond to two transformations estimated alternatively and S is again the similarity term of the equation, typically a SSD. The last term is here an homogeneous regularization over the whole image. The second term ensures that the regularized transformation T and the estimated transformation C do not go too far from each other during the registration process. This method will be named DT in section 3.2

2.6.2 Quantitative Validation Method: STAPLE

In order to compare quantitatively the brain segmentations obtained by each registration method, we have used a dataset of six images where manual delineations of the brainstem have been done by seven experts separately. To evaluate our results and compare them to all these manual delineations, we have used a widely used method in the literature: STAPLE [27]. This method allows to compute from several segmentations both the underlying ground truth and the performance parameters of the experts.

This is done using an Expectation Maximization algorithm, where the underlying ground truth is the hidden parameter. First, the probability for each voxel to belong to the ground truth is computed, knowing the current estimates of the performance parameters and the manual delineations. Then, the performance parameters are evaluated, knowing these new probabilities. Warfield et al. use sensitivity and specificity as parameters for each expert.

We have chosen in our experiments to use STAPLE to get the performance parameters of the manual delineations as well as the underlying ground truth. We then used this ground truth to compute a posteriori the performance parameters of each automatic segmentation method. All those computations

were done on a dilated mask of the union of the manual image segmentations and all the automatic segmentations. This therefore ensures a larger variation range for the specificity than when using all the background voxels as true negatives, and the same number of true negatives is used for the evaluation of each automatic segmentation method.

3 Experiments and Results

We present the evaluation of our new registration framework on two different applications. First, we show some qualitative and quantitative results on brain atlas registration for atlas-based segmentation. Then, we present results on the registration of bones in the lower abdomen area.

3.1 *Brain Critical Structures Segmentation for Radiotherapy Planning*

The first application considered here takes place in the frame of conformal brain radiotherapy. The planning of this treatment requires the accurate delineation of the tumor and of the critical structures, in order to define precisely the irradiation beams during the treatment. This task is very tedious to do manually, and also not reproducible.

The use of a brain atlas as the one proposed in [7] can then be very useful to automatically delineate the critical structures in the brain. This atlas is composed of two images: a simulated MRI from the BrainWEB² [36–39] and its segmentation, which was done manually by an expert.

The method we follow here consists of bringing the patient image on the simulated MRI and then applying the inverse transformation to the segmentation image to get the critical structures. We first globally position the patient on the atlas using an affine transformation. The second step is then to refine the result locally by using a non linear registration algorithm.

We will show examples of segmentation results and their evaluation on two image databases. The first database, coming from the Centre Antoine Lacasagne (CAL) in France, consisted of about 50 T1 patients acquired with 2 mm slice thickness. Among these patients, six had the brainstem delineated manually by seven different experts. This database was then used to perform a first qualitative evaluation (section 3.1.2) followed by a quantitative evaluation on the brainstem (section 3.1.3).

² See web site: <http://www.bic.mni.mcgill.ca/brainweb/>

A second database of images was then used to validate our method in clinical conditions at the Institut Gustave Roussy (IGR) in France. This database consisted of 22 T1 images injected with gadolinium and had a larger slice thickness (3 mm). As the atlas was made from a T1 image without injection, this may introduce errors in the segmentation. This database was used to perform a semi-quantitative study [40]. We will report the results obtained on this database in section 3.1.4.

3.1.1 Registration Parameters

The first task to use our algorithm is to define the areas to register. This is easy in this case and done once and for all, as we can use our atlas as the reference image for all the patients. In the following example, we register twelve critical structures: the cerebellum, the eyes, the optic chiasma, the gray nuclei (pallidum, putamen, caudate and thalami) and the brainstem.

The atlas provides us with an efficient way to select the local regions to register. For small structures such as the eyes, the optic chiasma or the gray nuclei, an affine behavior can reasonably be assumed between the images. We thus simply dilated each related label in the atlas and used it as a region. For bigger structures, like the cerebellum or the brainstem, the deformations are more complex. One affine transformation per structure is not sufficient. We therefore choose to split arbitrarily each of these structures into two areas. We thus register these structures using two affine transformations.

This gives a total of twelve registered regions (two on the brainstem, two on the cerebellum, one on the chiasma, on each eye and on each gray nucleus) predefined on the atlas. Some of these predefined regions are shown in section 2.2 in Fig. 1.

For all the experiments presented in this section, we have used the same parameters for the registration algorithms. One of our goals for this application is indeed to get a method that can be used on images from different centers and acquired using different parameters, without needing the optimization of the parameters for each patient. The best set of parameters for each method was determined qualitatively on the 50 patients by looking for parameters that gave good results everywhere on all patients.

We have thus chosen for the dense registration method (Runa [14]) parameters ($\alpha = 1$ and $\beta = 0.1$ in (3)) that gave qualitatively good results for all images. The same was done for the locally affine framework. All these experiments use a single set of parameters (block size: 7x7x7, one block every 3 voxels in each direction, $\lambda_{el} = 0.3$, $\lambda_{fl} = 0.2$ and a LTS cut at 70 %).

3.1.2 Qualitative Comparison with Runa [14]

Our first experiment was to compare the registration results (images and deformed grids) obtained using a globally affine registration, a dense registration algorithm (Runa), and our algorithm.

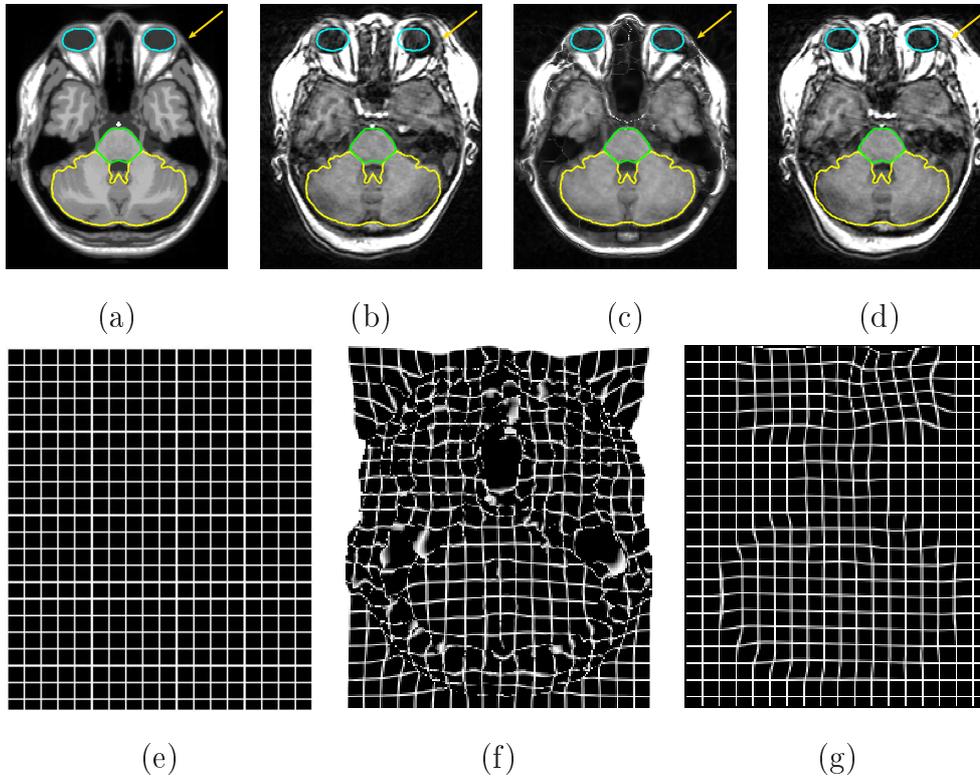


Fig. 4. **Qualitative comparison of the registration results on one patient.** **Top row:** (a): Atlas reference image. Patient resampled on atlas using (b): a global affine registration, (c): a specific dense algorithm (Runa) [14], (d): our algorithm. **Bottom row:** Deformation grids of the non linear part of the registration using (e): a global affine registration, (f): a specific dense algorithm (Runa) [14], (g): our algorithm. We can see that we obtain a much smoother and realistic deformation with our method than with a dense registration.

We show on Fig. 4 one example of a floating (a patient selected from the CAL database) image resampled to match the reference (atlas) image, with the contours of the atlas superimposed. This example was selected as it illustrates well the registration problems that can arise with dense registration when using only one set of parameters for all patients. We can first see that the two non linear methods perform well from a qualitative point of view on the brainstem and the cerebellum. The transformed structures are indeed well delineated by the reference contours. The eyes seem not quite as good for the dense registration method (arrows on Fig. 4). We can also see that, by constraining the transformation using the anatomical a priori available in the atlas, we have been able to remove strong local deformations (see the deformed

grids on Fig. 4), which may result in irregular contours when applying the transformation to the atlas structures for segmentation.

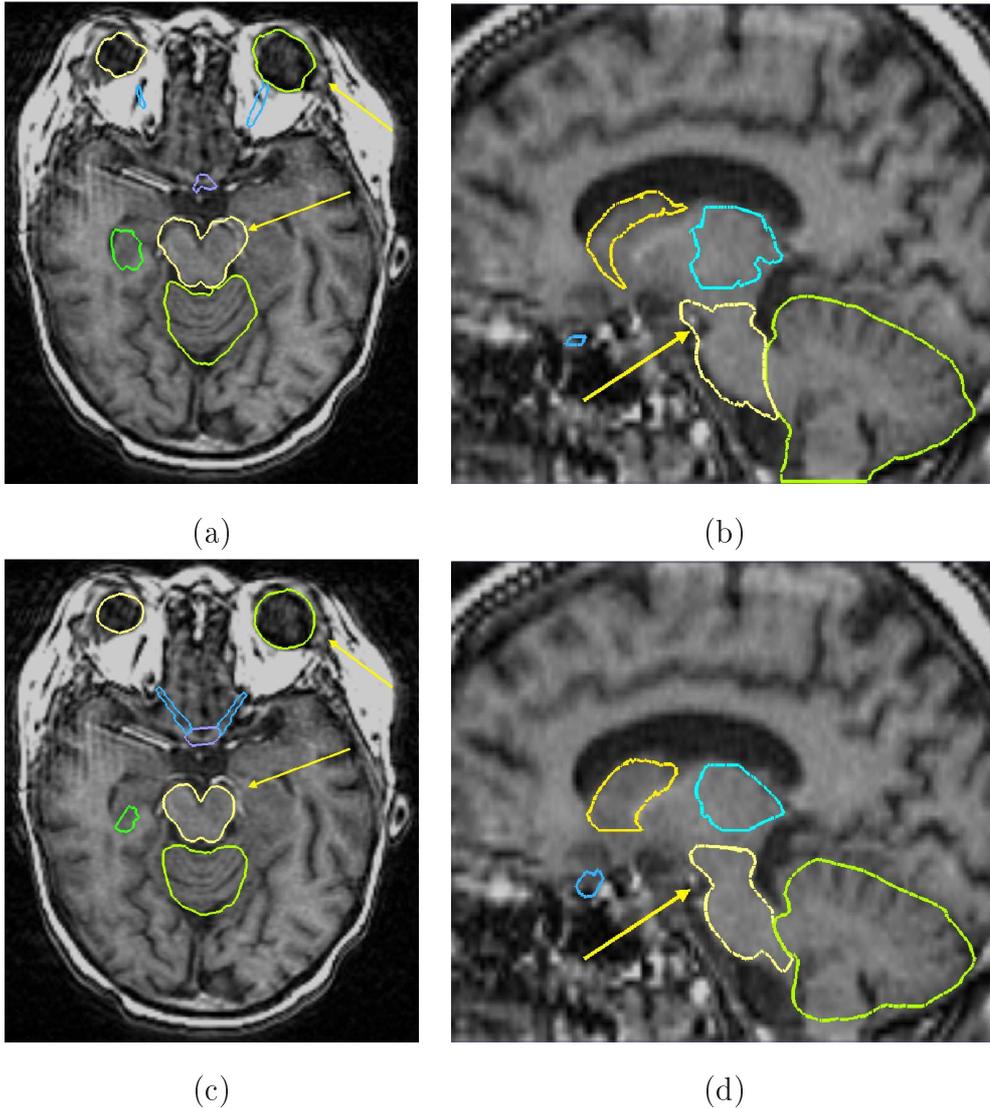


Fig. 5. **Qualitative comparison of the segmentation results on one patient from CAL.** Contours obtained using a dense transformation (top row) and using our new framework (bottom row). The contours obtained with our method are more regular and do not include the arteries in front of the brainstem. The eyes and gray nuclei contours are also less noisy.

We then use the obtained transformations for each algorithm (dense registration: Runa, locally affine registration: LAF) to resample the segmentation image and get the structure contours on the patient image. We show the qualitative results in Fig. 5. Again, the contours are much smoother, with our method than with a dense registration, for example on the eyes or on the gray nuclei. They are also more precise, mostly on the brainstem and the eyes. This fact was not obvious on the first registration results. The behavior of the obtained dense transformation is due to the fact that the dense registration

algorithm does not take into account any prior on the structures present in the images. This can lead to irregular contours when some structures are missing in the atlas (e.g. the artery in front of the brainstem). Our framework is better adapted as it is able to take into account priors on the structures to register. The transformation is also, by construction, less sensitive to local minima of the similarity measure. Finally, the computation time is faster (10 minutes as opposed to 40 minutes on a 3 GHz computer) than with the dense registration.

We also present here qualitative results obtained on the images from the IGR database. These images are much more difficult to delineate than the images from the CAL because of their larger slice thickness and of the injection of gadolinium that is not present in our atlas. These results, shown in Fig. 6, were obtained using the same set of parameters as for the patients from the CAL.

Again, better delineations are obtained using the locally affine method when compared to Runa. This is particularly true on small structures. The optic chiasma for example is only seen on one or two slices in the patient and is therefore very difficult to segment, even more when using one parameter set for all patients. On this structure, the results obtained by the locally affine are visually better thanks to the a priori constraint on the transformation.

3.1.3 CAL Database: Qualitative and Quantitative Evaluation

The 50 images of the CAL database were used for a first qualitative validation. These segmentation results have been presented to an expert (a radiotherapist) who visually inspected them and considered them as very satisfactory. In addition to this qualitative evaluation, we also perform a quantitative study on the six patients, for whom we have the brainstem delineated by seven experts. These patients were registered on the atlas using the two non linear registration methods. We have then used the atlas to get the binary segmentation of the brainstem for each image. For each one, we have compared this automatic segmentation with the seven manual segmentations performed by the different experts using the STAPLE algorithm [27], described in section 2.6. We show the associated sensitivity and specificity results in Table 1. We also report, for each sensitivity/specificity couple, the distance to the best achievable result ($Sens. = 1, Spec. = 1$) to give a quick overview of the quality of the result.

These figures confirm that the results obtained using our method are similar to the ones obtained using the dense transformation algorithm. We indeed obtain only one distance which is a not quite as good with our algorithm while the others are better than with dense registration.

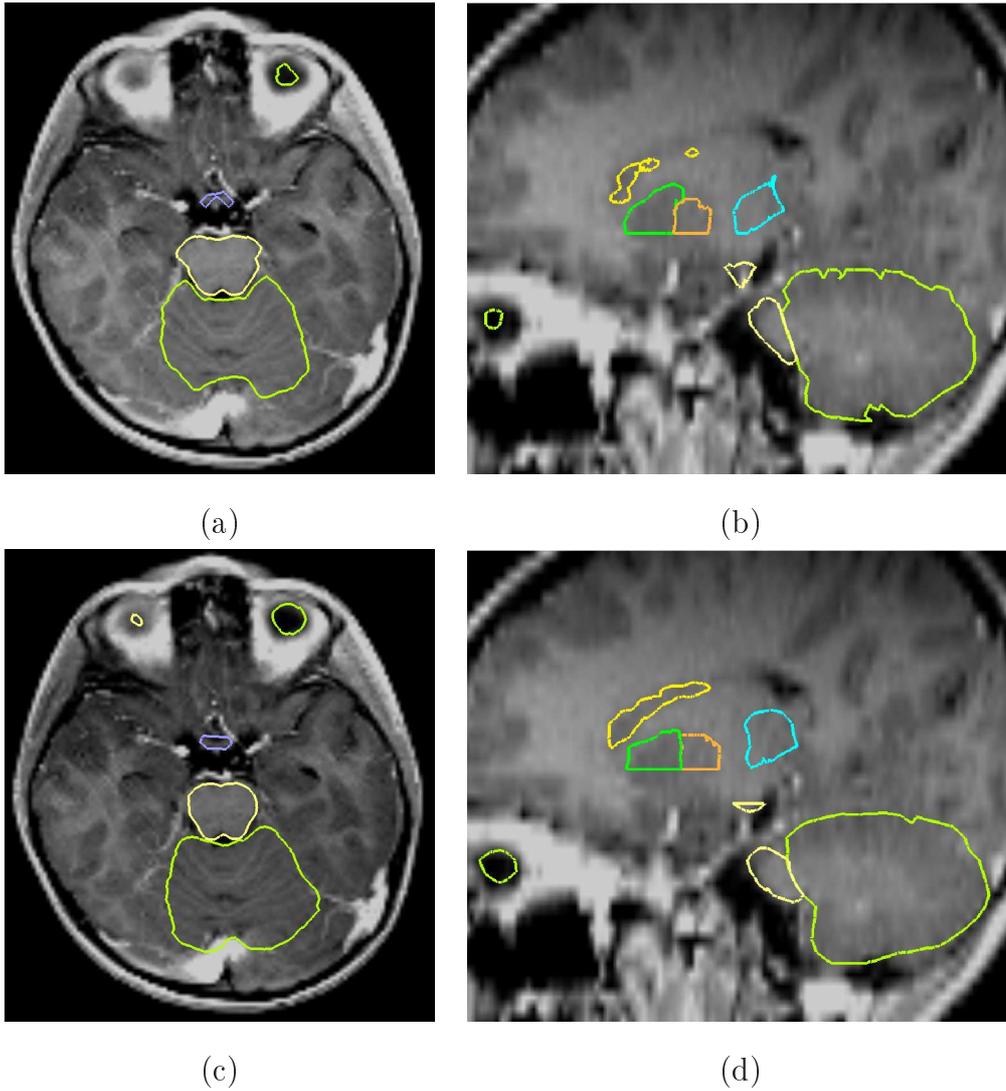


Fig. 6. **Qualitative Segmentation Results on a patient from IGR.** Comparison of the segmentations obtained using Runa ((a), (b)) and our locally affine framework ((c), (d)) on injected T1 images with 3mm slice thickness. Left column: axial slices ; right column: sagittal slices.

3.1.4 Semi-quantitative Validation in Clinical Conditions

Finally, this work on atlas-based brain critical structures segmentation is also currently being validated in clinical routine [40]. A semi-quantitative validation has been performed on the IGR database described in section 3.1. An expert rated the automatic segmentations by assigning for each structure a qualitative indice of conformation. This index varies between 0 and 5:

- 0 corresponds to a total disagreement of the expert with the automatic delineation
- 1 or 2 corresponds to a partial agreement ($< 50\%$ of the volume)
- 3 or 4 corresponds to a better agreement ($> 50\%$ of the volume)

Patient #	1	2	3	4	5	6
Sensitivity (Runa)	0.88	0.86	0.84	0.79	0.85	0.78
Specificity (Runa)	0.86	0.91	0.79	0.94	0.89	0.91
Distance to (1,1) (Runa)	0.19	0.17	0.27	0.22	0.19	0.24
Sensitivity (LAF)	0.86	0.85	0.82	0.84	0.84	0.76
Specificity (LAF)	0.89	0.91	0.83	0.92	0.90	0.93
Distance to (1,1) (LAF)	0.18	0.17	0.25	0.18	0.19	0.25

Table 1

Quantitative registration results on the brainstem with a dense method (Runa) and our method (LAF). Sensitivity and specificity obtained using STAPLE algorithm based on seven experts segmentations for six patients. The results we obtain are at least similar to the results obtained with a dense registration.

- 5 corresponds to a perfect agreement.

The agreement was visually evaluated on the basis of the shape of each structure, on the contours regularity and on their distance to the expected contours. This index is called ASC index in Fig. 7. The average results of this evaluation, done in clinical conditions on the 22 patients, are presented in Fig. 7 on the following structures: the optic chiasma, the eyes, the pituitary gland, the cerebellum, the optic nerves and the brainstem. Again these results confirm that the results are better using the locally affine framework when compared to Runa, particularly on small structures such as the pituitary gland and the optic chiasma. More details on this validation can be found in [40].

This study is intended in a near future to include more patients and use quantitative measures of quality. However, it already confirms that, thanks to its robustness and efficiency, the locally affine method is particularly adapted to the delineation of brain structures, even in difficult cases such as the images from IGR.

3.2 Bone Registration in the Lower Abdomen Area

The second application also takes place in the frame of conformal radiotherapy planning. The aim is to develop an automatic method to localize the bladder and the prostate in the lower abdomen area. This will be based on CT images, as it is the classical image acquired for this area. However, the organs we want to delineate have a very low contrast on CT images. They also exhibit a high variability in shape, size and contrast, depending on its emptiness and on the presence or not of a contrast agent. Using an anatomical atlas as for the brain is thus not adequate in this case and it is very difficult to automatically

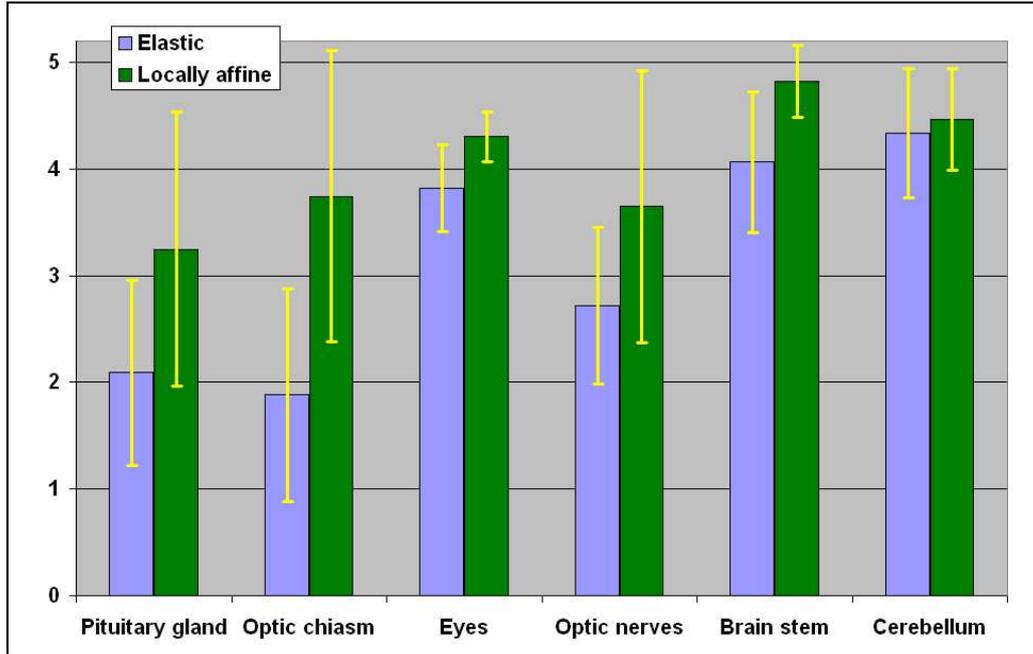


Fig. 7. **Semi-quantitative comparison of Segmentation Results.** Comparison using the ASC index performed at IGR between Runa method (called elastic) and the locally affine framework (called multi-affine). The error bars superimposed show the standard deviation of the results.

determine their position with accuracy.

One possible method to achieve this goal is the use of deformable models. These methods however require a precise initialization of the model. This can be done using an atlas where a reference image is associated to an initial position of deformable models. We thus present in this application a first step to estimate automatically the initial position of the bladder and prostate in the lower abdomen area. For that purpose, we must initially bring the patient image into the space of the reference image. As the variability of the organs we are looking for is high, we would prefer to bring the patient in this common space while not deforming too much the surrounding tissues; thus avoiding misregistrations that may appear when comparing a bladder fully contrasted to a bladder half contrasted for example.

We therefore established a set of salient points in the pelvic and leg bones and aim at registering these stable surrounding structures. We present here a feasibility study on this step, using five areas (the two femoral heads, the sacrum and the two lower ischiatic tuberosities). The regions around these points are used as affine component localizations in our algorithm. The evaluation will then be performed on these reference points to assess this registration step.

We present here results of inter-patient registration on five patients using these anatomical landmarks. All the patients' images are registered with respect to

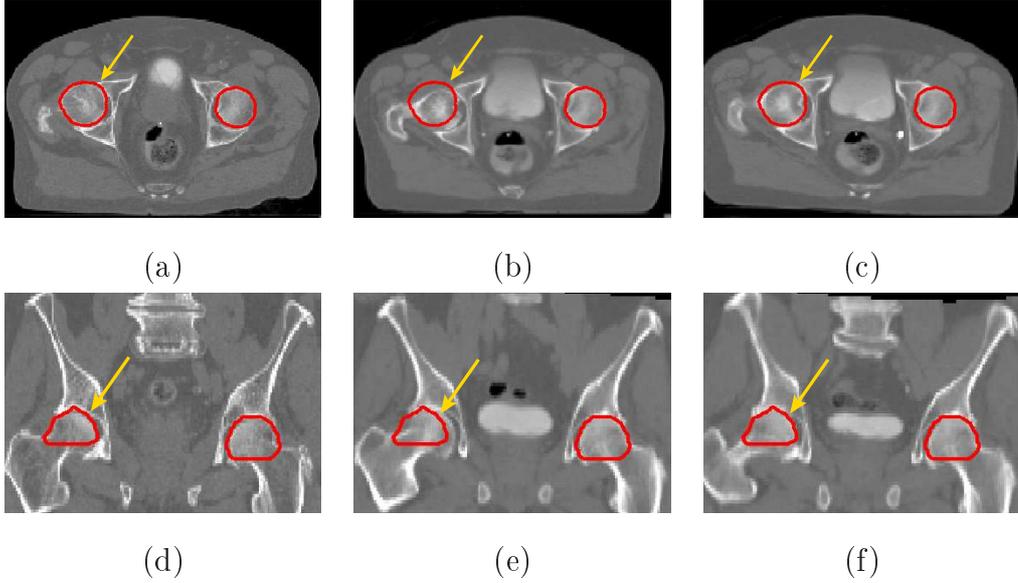


Fig. 8. **Registration result on the pelvis with femoral heads contours superimposed.** First row: axial slices. Second row: coronal slices. (a), (d): Reference image, (b), (e): floating image registered on atlas using a global affine transformation, (c), (f): floating image registered using our framework. Our method achieves a good positioning of the structures.

a sixth image which is taken as the reference image. The process consists of two stages: a global affine registration is performed using a block-matching algorithm [28], and then our algorithm is applied. We compare qualitatively the results of one of those registrations in Fig. 8. We can see a significant qualitative improvement of the registration result with respect to a global affine registration. The right femoral head is indeed misplaced with the global transformation and is well registered with our method (arrows on Fig. 8). Moreover, the information contained in the images outside the regions used in the registration remains consistent from an anatomical point of view.

We also compare the results we obtain using our method with results obtained using a dense registration algorithm [10], described rapidly in section 2.6. We will call this method DT in our experiments.

Qualitatively, we obtain very similar results. We thus performed a quantitative evaluation of the results on the position of the femoral heads in the images. To do that, we apply for each method the obtained transformation to the femoral heads centers of the reference image and compare them with the ones in the floating images. We show in Table 2 the norm of the Euclidean distance between the transformed landmarks and the corresponding landmarks in the floating images.

We can see in this table that the results obtained using our method are at least as good as the results obtained through dense elastic registration. We obtain

Patient #	1	2	3	4	5
Left head (AFFI)	6.22	11.40	2.02	4.41	13.68
Left head (DT)	4.12	3.10	1.91	2.32	3.55
Left head (LAF)	3.00	3.22	0.55	1.82	2.63
Right head (AFFI)	3.72	0.87	7.75	4.59	7.00
Right head (DT)	2.55	1.52	2.16	1.28	3.95
Right head (LAF)	1.34	1.38	1.14	0.79	2.93

Table 2

Registration results on femoral head centers. Distances in millimeters between the expected femoral head centers and those obtained from the registration (affine: AFFI ; locally affine: LAF ; dense transformation: DT). These figures show the ability of our method to cope with large deformations similarly or even better than with a dense registration.

also a stronger difference between the two non linear methods. This is due to large deformations in soft tissues surrounding the bones. These deformations indeed propagate when using a dense transformation in the elastic registration, leading to poorer results. Our algorithm computation time is also much lower (3 minutes as opposed to 10 minutes). Moreover, the goal was to place all the patients in a common space while deforming the soft tissues as little as possible. Our method performs the registration based on specific areas and ensures consistent results all over the image. It is then much more adapted to this type of application than a dense transformation solution, which tries to match the entire floating image.

4 Conclusion and Future Work

We have introduced in this article a novel framework for locally affine registration of anatomical structures. This allows us to adapt the transformation complexity to the specific registration task to perform. This method is also more robust, thanks to a more constrained transformation, and can therefore deal with images coming from different centers, as demonstrated in [40]. Thanks to our novel Log-Euclidean regularization framework, we ensure a smooth and coherent transformation. Finally, the use of the Log-Euclidean polyaffine framework at the end of the algorithm ensures a smooth and invertible transformation everywhere on the image.

We have used this framework so far on two applications, for which this framework is well adapted. On those applications, we have shown a significant qualitative improvement with respect to a dense registration method. The contours

are indeed much more regular, while delineating well the structures. Our algorithm is also much less dependent of the parameters, as we have always used the same parameters for all the experiments on the brain. Finally, we have shown quantitative results using sensitivity and specificity, showing results at least similar to the ones obtained using a dense registration algorithm.

In our experiments, we have chosen to use for each method one parameter set for all patients. This set was determined qualitatively on several patients, in order to get a good result everywhere. An interesting study would be to determine these parameter sets by using quantitative measures on many structures everywhere in the brain. Moreover, in order to have a better idea of the dependency of the algorithms to their parameters, it would be interesting to study the variations of the quantitative measures when the parameters are varying. This would add a robustness criterion to the evaluation of the methods.

The quantitative results on the brain have shown the limits of using only voxel-based overlap quantitative measures. We have indeed some segmentations that look visually better than the dense registration, because of the less irregular contours and that give nearly the same quantitative results as with the dense registration. Thus, one future work will be to introduce a measure of smoothness in the quality measures. This could be of great interest, as this is one of the important points for medical doctors. We will also extend our validation, mainly on brain critical structures segmentation. The goal will be to add more patients and also to extend this validation to more structures.

An other field of experiments will be to estimate the influence of the areas on the registration result. They are indeed predefined and not refined during the registration algorithm. It would be interesting to see first the influence of small changes in the predefined areas. Then, we could study the influence of the number and shape of the areas (for example by defining several areas using different methods on the cerebellum). An other point to look at is the refinement of the areas during or after the registration. We could use the clustering of the obtained deformation field over our regions as in [21]. We could also define a measure of the goodness of the registration inside an area and use it for refining the predefined regions.

Finally, this algorithm can be used as a link between global registration and dense non linear registration. This can be done by defining areas all over the image, register the images and splitting the regions that are misregistered. By looping over these two last steps, we would have an algorithm getting closer to the dense transformation while keeping smooth transformations. Another way could be to register the images using three steps instead of two: global registration followed by our algorithm and, to get the small residual deformations, a dense registration. As the residual deformations will be smaller, we will then be able to regularize substantially more and get smoother deformations.

Acknowledgments

This work was partially founded by ECIP project MAESTRO (IP CE503564) and ANRT.

References

- [1] D. Rey, G. Subsol, H. Delingette, N. Ayache, Automatic detection and segmentation of evolving processes in 3D medical images: Application to multiple sclerosis, *Medical Image Analysis* 6 (2) (2002) 163–179.
- [2] W. R. Crum, R. I. Scahill, N. C. Fox, Automated hippocampal segmentation by regional fluid registration of serial MRI: validation and application in Alzheimer’s disease, *Neuroimage* 13 (5) (2001) 847–55.
- [3] R. I. Scahill, J. M. Schott, J. M. Stevens, M. N. Rossor, N. C. Fox, Mapping the evolution of regional atrophy in Alzheimer’s disease: unbiased analysis of fluid-registered serial MRI, *Proc Natl Acad Sci U S A* 99 (7) (2002) 4703–7.
- [4] P. Thompson, M. Mega, K. Narr, E. Sowell, R. Blanton, A. Toga, Brain image analysis and atlas construction, in: M. Fitzpatrick, M. Sonka (Eds.), *Handbook of Medical Image Proc. and Analysis*, SPIE, 2000, Ch. 17.
- [5] A. Toga, P. Thompson, Mapping brain asymmetry, *Nature Reviews Neuroscience* 4 (1) (2003) 37–48.
- [6] B. M. Dawant, S. L. Hartmann, J.-P. Thirion, F. Maes, D. Vandermeulen, P. Demaerel, Automatic 3-D segmentation of internal structures of the head in MR images using a combination of similarity and free-form transformations: Part I, methodology and validation on normal subjects, *IEEE Transactions on Medical Imaging* 18 (10) (1999) 909–916.
- [7] P.-Y. Bondiau, G. Malandain, S. Chanalet, P.-Y. Marcy, J.-L. Habrand, F. Fauchon, P. Paquis, A. Courdi, O. Commowick, I. Rutten, N. Ayache, Atlas-based automatic segmentation of MR images: validation study on the brainstem in radiotherapy context, *Int J Radiat Oncol Biol Phys* 61 (1) (2005) 289–98.
- [8] L. Brown, A survey of image registration techniques, *ACM Computing Surveys* 24 (4) (1992) 325–376.
- [9] G. E. Christensen, R. D. Rabbitt, M. I. Miller, Deformable templates using large deformation kinematics, *IEEE Transactions on Image Processing* 5 (10) (1996) 1435–1447.
- [10] P. Cachier, E. Bardinet, D. Dormont, X. Pennec, N. Ayache, Iconic Feature Based Nonrigid Registration: The PASHA Algorithm, *CVIU — Special Issue on Nonrigid Registration* 89 (2-3) (2003) 272–298.

- [11] R. Bajcsy, C. Broit, Matching of deformed images, in: Proc. of the 6th International Conference on Pattern Recognition, 1982.
- [12] C. Broit, Optimal registration of deformed images, Ph.D. thesis, University of Pennsylvania (1981).
- [13] R. Bajcsy, S. Kovačič, Multiresolution elastic matching, *Computer Vision, Graphics, and Image Processing* 46 (1989) 1–21.
- [14] R. Stefanescu, X. Pennec, N. Ayache, Grid powered nonlinear image registration with locally adaptive regularization, *Medical Image Analysis* 8 (3) (2004) 325–342, MICCAI 2003 Special Issue.
- [15] O. Commowick, R. Stefanescu, P. Fillard, V. Arsigny, N. Ayache, X. Pennec, G. Malandain, Incorporating statistical measures of anatomical variability in atlas-to-subject registration for conformal brain radiotherapy, in: J. Duncan, G. Gerig (Eds.), MICCAI 2005, Part II, Vol. 3750 of LNCS, Springer Verlag, Palm Springs, CA, USA, October 26-29, 2005, pp. 927–934.
- [16] G. Rohde, A. Aldroubi, B. Dawant, The adaptive bases algorithm for intensity based nonrigid image registration, *IEEE Transactions on Medical Imaging. Special issue on image registration.* 22 (2003) 1470–1479.
- [17] D. Rueckert, L. L. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, D. J. Hawkes, Nonrigid registration using free-form deformations: Application to breast MR images, *IEEE Transactions on Medical Imaging* 18 (8) (1999) 712–721.
- [18] A. Cuzol, P. Hellier, E. Mémin, A novel parametric method for non-rigid image registration, in: G. Christensen, M. Sonka (Eds.), 19th International Conference on Information Processing in Medical Imaging, IPMI’2005, no. 3565 in Lecture Notes in Computer Science, Glenwood Springs, USA, 2005, pp. 456–467.
- [19] V. Camion, L. Younes, Geodesic interpolating splines, in: *Energy Minimization Methods in Computer Vision and Pattern Recognition*, 2001, pp. 513–527.
- [20] J. Little, D. Hill, D. Hawkes, Deformations incorporating rigid structures, *Computer Vision and Image Understanding* 66 (2) (1997) 223–232.
- [21] A. Pitiot, E. Bardinet, P. M. Thompson, G. Malandain, Piecewise affine registration of biological images for volume reconstruction, *Medical Image Analysis* 10 (3) (2006) 465–483.
- [22] X. Papademetris, D. P. Dione, L. W. Dobrucki, L. H. Staib, A. J. Sinusas, Articulated rigid registration for serial lower-limb mouse imaging., in: J. Duncan, G. Gerig (Eds.), Proceedings of the 8th Int. Conf. on Medical Image Computing and Computer-Assisted Intervention - MICCAI 2005, Part II, Vol. 3749 of LNCS, Springer Verlag, 2005, pp. 919–926.
- [23] A. du Bois d’Aische, M. de Craene, B. Macq, S. Warfield, An articulated registration method, in: *IEEE International Conference on Image Processing*, Genova, Italy, 2005.

- [24] R. Narayanan, J. A. Fessler, H. Park, C. R. Meyer, Diffeomorphic nonlinear transformations: A local parametric approach for image registration., in: IPMI, 2005, pp. 174–185.
- [25] 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.
- [26] V. Arsigny, O. Commowick, X. Pennec, N. Ayache, A fast and Log-Euclidean polyaffine framework for locally affine registration, Research Report RR-5865, INRIA (March 2006).
- [27] S. K. Warfield, K. H. Zou, W. M. Wells, Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation, *IEEE Transactions on Medical Imaging* 23 (7) (2004) 903–921.
- [28] S. Ourselin, A. Roche, S. Prima, N. Ayache, Block matching: A general framework to improve robustness of rigid registration of medical images, in: A. DiGioia, S. Delp (Eds.), *Third International Conference on Medical Robotics, Imaging And Computer Assisted Surgery (MICCAI 2000)*, Vol. 1935 of LNCS, Springer, Pittsburgh, Pennsylvania USA, 2000, pp. 557–566.
- [29] P. Besl, N. McKay, A method for registration of 3D shapes, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14 (2) (1992) 239–256.
- [30] Z. Zhang, Iterative point matching for registration of free-form surfaces, *International Journal of Computer Vision* 13 (2) (1994) 119–152.
- [31] G. Malandain, E. Bardinet, K. Nelissen, W. Vanduffel, Fusion of autoradiographs with an MR volume using 2-D and 3-D linear transformations, *NeuroImage* 23 (1) (2004) 111–127.
- [32] A. Roche, G. Malandain, N. Ayache, Unifying maximum likelihood approaches in medical image registration, *International Journal of Imaging Systems and Technology: Special Issue on 3D Imaging* 11 (1) (2000) 71–80.
- [33] P. Rousseeuw, Least median of squares regression, *Journal of the American Statistical Association* 12 (1984) 871–880.
- [34] S. H. Cheng, N. Higham, C. Kenney, A. Laub, Approximating the logarithm of a matrix to specified accuracy, *SIAM J. Matrix Anal. Appl.* 22 (4) (2001) 1112–1125.
- [35] J.-P. Thirion, Image matching as a diffusion process: an analogy with maxwell’s demons, *Medical Image Analysis* 2 (3) (1998) 243–260.
- [36] C. Cocosco, V. Kollokian, R.-S. Kwan, A. Evans, BrainWeb: Online interface to a 3D MRI simulated brain database, *NeuroImage* 5 (4).
- [37] R.-S. Kwan, A. Evans, G. Pike, MRI simulation-based evaluation of image-processing and classification methods, *IEEE Transactions on Medical Imaging* 18 (11) (1999) 1085–97.

- [38] D. Collins, A. Zijdenbos, V. Kollokian, J. Sled, N. Kabani, C. Holmes, A. Evans, Design and construction of a realistic digital brain phantom, *IEEE Transactions on Medical Imaging* 17 (3) (1998) 463–468.
- [39] R.-S. Kwan, A. Evans, G. Pike, An extensible MRI simulator for post-processing evaluation, in: *Visualization in Biomedical Computing (VBC'96)*, Vol. 1131 of LNCS, 1996, pp. 135–140.
- [40] A. Isambert, F. Dhermain, A. Beaudre, G. Malandain, O. Commowick, J. Diaz, F. Bidault, P. Bondiau, J. Bourhis, M. Ricard, D. Lefkopoulos, Requirements for the use of an atlas-based automatic segmentation for delineation of organs at risk (OAR) in conformal radiotherapy (CRT): quality assurance (QA) and preliminary results for 22 adult patients with primary brain tumors, in: *Acts of the 24th European Society for Therapeutic Radiology and Oncology (ESTRO 2005)*, 2005.
- [41] D. Sheppard, A two-dimensional interpolation function for irregularly spaced data, in: *23rd National Conference of the ACM*, 1968, pp. 517–524.

A Locally Affine Transformation

We detail in this appendix how to create a global locally affine transformation from a set of N affine transformations A_i defined on regions R_i . The main idea is that we want to have a global transformation keeping the properties of the local affine transformations on their areas of definition. We detail here three methods to achieve this goal.

These three methods require for each region R_i an affine transformation A_i and a non-negative weight function $\bar{w}_i(x)$. This function \bar{w}_i can be of any form like a Gaussian or a distance function. Furthermore, we assume that the weights are normalized, i.e. for all x , $\sum_{i=1}^N \bar{w}_i(x) = 1$.

A.1 Transformation Incorporating Rigid Structures

This method has first been proposed by [20] and is also used in [21]. Weight functions $w_i(x)$ are basically defined as an inverse distance to the region R_i and renormalized to give the \bar{w}_i . In this approach, pairs of rigid structures are selected in the input images, along with linear transformations. A number of pairs of outer landmarks further constrain the interpolation scheme, which uses Hardy multi-quadric basis functions to interpolate in between the areas. This results in applying the affine transforms to the user-defined structures while ensuring a smooth interpolation in between them.

This scheme has the advantage to keep an exact affine transformation inside the R_i . However, it does not ensure always the invertibility of the final transformation, which is desirable when using the registration for the delineation of structures. Furthermore, it has been used so far in 2D and seems complex to use and to optimize in full 3D.

A.2 Direct Averaging Method

In this case, having defined the weight functions, the way to fuse the affine components has been given in [41]. The transformation $T = M_2(A_i)$ is simply computed as the weighted average of the displacements generated from each A_i with respect to the weights $\bar{w}_i(x)$:

$$T(x) = \sum_{i=1}^N \bar{w}_i(x) A_i \cdot x. \quad (\text{A.1})$$

The transformation obtained using (A.1) is smooth. Using an inverse distance function as in section A.1 allows us to get roughly affine transformations inside the regions R_i . However, as pointed in [25], the resulting transformation obtained is not invertible in general. Nevertheless, this method is fast and simple to implement. Moreover, optimizing the affine transformations in this framework is very easy and efficient.

A.3 Log-Euclidean Polyaffine Framework

To remedy the problem of invertibility in the preceding method and ensure an invertible transformation all over the image, [25] proposed a polyaffine framework, which was improved and rendered faster thanks to the Log-Euclidean framework recently in [26].

This method, noted $T = M_3(A_i)$, consists in averaging infinitesimal displacements generated from the A_i according to the weights $\bar{w}_i(x)$. The value of the transformation at point x is then obtained by integrating the trajectory of point x between time 0: $x(0) = x$ and time 1: $x(1) = T(x)$. This results in an Ordinary Differential Equation (ODE):

$$\dot{x} = \sum_i \bar{w}_i(x) \log(A_i).x. \quad (\text{A.2})$$

The Log-Euclidean polyaffine transformations are, by construction, always invertible, and their inverse can be very efficiently computed. With this method, the estimation of the affine transformations is however computationally expensive and much more complicated than with the direct averaging. Moreover, the results are very close to the preceding method except in interpolating areas where non invertible parts of the transformation were present.