

Extension of the ICP Algorithm to non rigid Intensity-based Registration of 3D Volumes

J. Feldmar¹, J. Declerck², G. Malandain² and N. Ayache²

1: Department of Engineering Science,
Oxford University,
Parks Road, Oxford OX1 3PJ, England
E-mail: jfeldmar@robots.ox.ac.uk

2: INRIA SOPHIA, Projet EPIDAURE
2004 route des Lucioles, B.P. 93
06902 Sophia Antipolis Cedex, France.

Abstract

We present in this paper a new registration and gain correction algorithm for 3D medical images. It is intensity based. The basic idea is to represent images by 4D points (x_j, y_j, z_j, i_j) and to define a global energy function based on this representation. For minimisation, we propose a technique which does not require computing the derivatives of this criterion with respect to the parameters. It can be understood as an extension of the Iterative Closest Point algorithm [5, 56] or as an application of the formalism proposed in [13]. Two parameters enable us to develop a coarse-to-fine strategy both for resolution and for deformation. Our technique presents the advantage of minimising a well-defined global criterion, to deal with various classes of transformations (for example rigid, affine, volume spline and radial basis functions), to be simple to implement, and to be efficient in practice. Results on real brain and heart 3D images are presented to demonstrate the validity of our approach. We also explain how one can compute basic statistics on the deformation parameters to constrain the set of possible deformations by learning and to discriminate between different groups.

1 Introduction

Registration is a key problem in medical imaging. Indeed, the physician must often **compare** or **fuse** different images. The problem is as follow: *given two 3D images, find the geometric transformation that best superimposes them, with respect to some constraints*. If both images come from the same patient and from a rigid anatomical organ, then the problem is *rigid registration*. Otherwise it is *non-rigid registration*.

Registration techniques can be classified into two classes: 1) techniques using additional artificial markers and 2) techniques without such markers. Affixing markers can be very invasive or can produce unacceptable constraints. The technique proposed in this paper does not require markers.

Different methods have been proposed to try to solve the registration problem without additional markers and without user's interaction. Usually, registration is based on a representation computed from the 3D images. This representation can be high level (graphs, crest points, crest lines), intermediate level (surfaces, contours) or low level. We do not present a review here of the numerous registration methods based on high and middle level representations; complete reviews can be found in [3, 8, 23, 52].

The registration method presented in this paper is related to intensity based techniques. Most such methods were developed for 3D-3D rigid registration and try to maximise the correlation [30, 55] or the mutual information [15, 46, 53] between the two images. We explore a low level method because high or middle level representations can be difficult to compute, either be-

cause of the image acquisition modality or because it is not easy to extract features of the organs in the images.

Brain images are a good example of images that are difficult to segment. A lot of research has been done to try to solve this problem [33, 35, 54]. However, matching two MR brain images from two different patients is important in practice. Indeed, we have access to a brain image which has been manually segmented (or labelled) (courtesy of Ron Kikinis, at the Brigham and Women’s hospital, Boston) and we use it as an anatomical atlas. Hence, non rigid inter-patient registration allows us to label automatically a MR image of a new patient into anatomical regions based on knowledge of voxel-to-voxel correspondence.

Previous research has aimed to compute such a match based on crest lines, surfaces or contours [19, 23, 31, 45, 47, 49, 51] or based directly on the intensities in the images [4, 12, 16, 25, 27, 39, 50]. The algorithm presented in this paper aims to compute a global **geometric transformation** that minimises an **explicit global criterion** and uses a coarse to fine strategy to try to avoid local minima. The closest algorithm to ours is probably presented in [25]: both correct for geometric and intensity difference. However, the minimized criteria are different (see section 2) and the search strategies differ fundamentally in their nature.

Note that this idea to deform an atlas towards an image can be criticized since it is not proven that there always exists an homology between two different brains [34]. However, these deformation procedures are very useful. One of the difficulties is to segment (for point based techniques) or to correct the intensities (for intensity-based techniques) of the images that have to be

registered. This is partly because of the variable shape of the brain but also because of the gain problem in the MR images. We present in this paper a new registration and gain correction algorithm which is intensity-based. This algorithm has numerous applications. Indeed, it should enable us to perform registration when contour (or higher level feature) extraction is difficult. It is an extension of the ICP algorithm [5, 9, 11, 37, 56]. Our technique has several advantages: minimising a well-defined global criterion, to deal with various, well-defined classes of transformations (for example rigid, affine and volume spline or radial basis functions); to be simple to implement; and to be efficient in practice. Moreover, since each deformation is described by a small set of parameters, we have a highly compact information for describing a shape, which makes possible the computation of statistics to classify deformations without the problem of the curse of dimensionality.

In this paper, we first expose the representation on which the algorithm is based and the corresponding minimised criterion (section 2). Then, we describe the minimisation algorithm (section 3) and the computation of the representation (section 4). We then present initial results obtained with brain and heart images which demonstrate the validity of our approach (sections 5 and 6). We conclude by sketching future research directions.

2 A global correlation criterion

2.1 The classical criterion

The most straightforward idea for registering two images i_1 and i_2 with intensity functions $I_1(x, y, z)$ and $I_2(x, y, z)$ is to minimise the criterion:

$$C(f) = \sum_{M_i \in i_1} (I_2(f(M_i)) - I_1(M_i))^2,$$

where f is a 3D-3D geometric transformation.

If it is necessary to correct the intensity to register the two images, one can minimise the following modified criterion:

$$C'(f, g) = \sum_{M_i \in i_1} (I_2(f(M_i)) - g(I_1(M_i), M_i))^2,$$

where f is a 3D-3D geometric transformation and g is an intensity correction function. These formulations have a fundamental drawback: they imply a search for an exact superimposition of the two images even if this might not always be possible with the considered class of transformations or deformations. We want to bring closer together the points with similar intensity but we want to keep a constrained deformation.

2.2 Our criterion

In our formulation, we consider 3D images as hypersurfaces in a 4D space. Hence, an image corresponding to a function $i = I(x, y, z)$ is represented by a set of 4D points (x_j, y_j, z_j, i_j) . The first three coordinates are spatial coordinates and the fourth one is intensity coordinate (more details will be given in section 4.1).

We propose to minimise a global criterion measuring the correlation between the two images by deforming the scene image into the model image. Given our formulation, this means that we deform the 4D surface, that is, we both correct the geometry and the intensity in the image¹. The minimised energy function is:

$$E(\mathbf{f}, g) = \sum_{(\mathbf{x}_j, i_j) \in \text{SceneImage}} d((\mathbf{f}(\mathbf{x}_j), g(\mathbf{x}_j, i_j)), \mathbf{CP}_{4D}(\mathbf{f}(\mathbf{x}_j), g(\mathbf{x}_j, i_j)))^2 \quad (1)$$

where

- \mathbf{x}_j denotes the three spatial coordinates of point M_j , i.e. (x_j, y_j, z_j) and i_j denotes the intensity in the image at this point.
- \mathbf{f} is the **3D-3D** geometric transformation of the scene image. Note that it does not depend on the intensity.
- g is a **4D-1D** function which associates an intensity value to a point in the image depending on its position and its current intensity.
- \mathbf{CP}_{4D} is the function which associates with a 4D point its closest point among the points describing the model image.
- Finally, d is a “distance” function on 4D points $M = (\mathbf{x}, i)$ and $N = (\mathbf{x}', i')$,

$$d(M, N) = (\alpha_1^2(x - x')^2 + \alpha_2^2(y - y')^2 + \alpha_3^2(z - z')^2 + \alpha_4^2(i - i')^2)^{1/2}.$$

¹We choose a global function for intensity correction otherwise registration would not make sense.

Of course, the choice of the α_i 's is crucial. It determines how far, in terms of spatial distance and intensity, one needs to search for the matching point. In practice, $\alpha_1, \alpha_2, \alpha_3$ are the inverted voxel dimensions and α_4 is computed by solving:

$$\alpha_1^2 \Delta x^2 + \alpha_2^2 \Delta y^2 + \alpha_3^2 \Delta z^2 - \alpha_4^2 \Delta i^2 = 0,$$

where $(\Delta x, \Delta y, \Delta z)$ is the spatial displacement expressed in voxels dimensions necessary to correct for a difference in intensity Δi . In order to understand this equation, consider a point $M = (x, y, z, i)$ in the scene image and two points $M_1 = (x, y, z, i')$ and $M_2 = (x', y', z', i)$ in the model image. At stage 1, the matched 4D point is located on a segment of the hypersurface $I(x, y, z)$ which boundaries are M_1 and M_2 . If we want to match the point M with a point close to M_2 and not to M_1 , then we have to choose:

$$\alpha_4^2 \gg \frac{\alpha_1^2(x - x')^2 + \alpha_2^2(y - y')^2 + \alpha_3^2(z - z')^2}{(i - i')^2},$$

otherwise, we choose

$$\alpha_4^2 \ll \frac{\alpha_1^2(x - x')^2 + \alpha_2^2(y - y')^2 + \alpha_3^2(z - z')^2}{(i - i')^2}.$$

Of course, a key point is also to choose the definition domain of the energy function E : this constraints the functions \mathbf{f} and g . We discuss this issue in more detail in the conclusion. For example, if the two images are from the same anatomical object, and if the voxel intensities correctly represent the associated tissues, then the searched function \mathbf{f} will be a **rigid displacement** and g will be fixed to be the **identity function**.

When the two anatomical regions do not come from the same patient, or when the anatomical region is deformable, the function \mathbf{f} can be an **affine**, a **spline** or a **radial basis function** (RBF) (see appendix A). When the intensity in the images is perturbed by a distortion, one can also search g as an affine, polynomial, spline or RBF function, depending on the physical analysis of the perturbation.

2.3 Smoothing

When \mathbf{f} and g are deformation functions, it is often necessary to add a smoothing term to E . In these cases, the energy function becomes:

$$E_{smooth}(\mathbf{f}, g) = E(\mathbf{f}, g) + \lambda \sum_{(\mathbf{x}_j, i_j) \in SceneImage} Smoothness(\mathbf{f}, g, (\mathbf{x}_j, i_j)),$$

where $Smoothness(\mathbf{f}, g, (\mathbf{x}_j, i_j))$ is typically the sum of the norm of the second derivatives of \mathbf{f} and g with respect to each of their coordinates (“bending energy”) respectively at points (\mathbf{x}_j) and (\mathbf{x}_j, i_j) . The parameter λ is experimentally determined as described in the results section.

2.4 Use of the gradient information

Even if the algorithm deals directly with intensities, it can be desirable to enhance the importance of areas where the intensity varies a lot. These are the areas where the norm of the gradient is high. In such cases, we may weight each term of the energy E with the norm of the gradient at the

corresponding point:

$$E(\mathbf{f}, g) = \sum_{(\mathbf{x}_j, i_j) \in SceneImage} \|\vec{\nabla}_{SceneImage}(\mathbf{x}_j)\| \cdot d((\mathbf{f}(\mathbf{x}_j), g(\mathbf{x}_j, i_j)), \mathbf{CP}_{4D}(\mathbf{f}(\mathbf{x}_j), g(\mathbf{x}_j, i_j)))^{1/2} \quad (2)$$

This new criterion can be surprising since at boundaries, the intensity is not well defined. However, this criterion gives good results in practice.

Note that it is also possible to use information about the gradient direction. In this case, the points representing the images are no longer 4D but are 7D points: three spatial coordinates, one intensity coordinate and three gradient coordinates. Hence, the distance between two points is a compromise between the spatial distance, the difference of gradient norm and orientation and the difference of intensity. The minimisation of the corresponding energy tends to minimise this difference between the two sets of points describing the images. This way of using the gradient information is similar to the use of surface normals presented in [24] for rigid surface registration: it gave satisfactory results though a precise theoretical approach would be necessary, as exposed in [42].

3 Minimisation technique

To minimise the energy function E or E_{smooth} , we developed a technique which does not require computation of the derivatives of these functions with respect to the parameters. It can be understood as an extension of the Iterative Closest Point algorithm [5, 56] or as an application of the formalism proposed in [13].

3.1 The algorithm

The minimisation algorithm is iterative. At each iteration i , we update estimates \mathbf{f}_i and g_i of \mathbf{f} and g in two stages:

- Stage 1: we construct a set of pairs of 4D points $Match_i$ by associating with each point M_j in the scene image the point N_j such that:

$$N_j = \mathbf{CP}_{4D}([\mathbf{f}_{i-1}(\mathbf{x}_j), g_{i-1}(M_j)]),$$

where \mathbf{x}_j are the three spatial coordinates of M_j . $Match_i$ is the set of pairs (M_j, N_j) .

- Stage 2: we compute the least squares sense transformations \mathbf{f}_i and g_i corresponding to $Match_i$. These are the \mathbf{f}_i and g_i that minimise:

$$\sum_{(M_j, N_j) \in Match_i} d((\mathbf{f}(\mathbf{x}_j), g(M_j)), N_j)^2 + \sum_{(M_j, N_j) \in Match_i} Smoothness(\mathbf{f}, g, M_j).$$

For the rigid, affine, spline and radial basis transformations classes and for the smoothing terms which we use, this criterion is quadratic and the least squares estimation turns out to be the resolution of a linear system as explained in appendix A. Note from a very practical point of view that the code necessary for this estimation is available in [44] and there is almost no programming to be done.

It is straightforward to show that this algorithm minimises the defined energy and that it converges². Let us define the energy function E' :

$$E'(\mathbf{f}, g, Match) = \sum_{M_j \in SceneImage} d((\mathbf{f}(\mathbf{x}_j), g(M_j)), Match(M_j))^2 + \sum_{M_j \in SceneImage} Smoothness(\mathbf{f}, g, M_j).$$

In stage 1 of our algorithm, the variables \mathbf{f} and g are fixed and E' is minimised with respect to $Match$. Indeed, in this case, the function $Match$ minimising E' satisfies:

$$Match(M_j) = \mathbf{CP}_{4D}((\mathbf{f}(\mathbf{x}_j), g(M_j))).$$

In stage 2, the variable $Match$ is fixed and E' is minimised with respect to \mathbf{f} and g . Thus, at each stage, E' decreases. Because E' is positive, convergence is guaranteed even if it can be towards a local minimum.

This minimisation technique is efficient. Contrary to classical minimisation techniques, it is not “local” (the transformation parameters can vary a lot between two successive iterations) and it does not require either computation of the derivative of E with respect to the parameters or the tuning of parameters. On the other hand, it assumes that each point in the scene image has a correspondent in the model image, and we next address this issue.

3.2 The occlusion problem

Some points in one image do not have any correspondent in the other, for example because of occlusion, or because of the evolution of a pathology (for

²Of course, there is no guarantee that we find the global minimum. We are minimising a non-convex function and we can only prove convergence towards a local minimum.

example a tumor).

It is important to deal explicitly with this occlusion problem to get an accurate transformation. One might adopt a robust criterion (in the sense of statistics) as in [14, 29] for rigid surface registration. But then stage 2 of the algorithm would not be a linear system resolution and the algorithm would be far less efficient. We prefer the approach proposed in [56] for 3D-3D rigid surface registration. For each match (M_j, N_j) in $Match_i$, we decide if it is plausible or not. This point is important because if we accept erroneous matches, the solution will be biased and if we reject correct matches, the solution will not be accurate.

For each pair (M_j, N_j) in $Match_i$, let us note:

$$\mathbf{d}_j = (\mathbf{f}(\mathbf{x}_j), g(M_j)) - N_j.$$

One may suppose that if the registration was correct, the Mahalanobis distance

$$\delta_j = (\mathbf{d}_j - \mu)^t S^{-1} (\mathbf{d}_j - \mu),$$

where μ and \mathbf{S} are respectively the mean and the covariance of the vectors \mathbf{d}_j , would follow a χ^2 distribution. It corresponds to the assumption that the difference in each coordinate follows a gaussian distribution. Thus, we can compare the statistics of this variable δ_j with a χ^2 with 4 (or 7) degrees of freedom and decide whether a pair (M_j, N_j) is plausible by looking a χ^2 -table with an appropriate confidence value, say 95 % or 99 %. Stage 2 of the minimisation algorithm is modified so that the least squares criterion takes into account only the plausible matches in $Match_i$.

4 Representation of the images

4.1 Computing the representation

For this registration algorithm, it is fundamental to choose the 4D points representing the images. One can associate with each voxel V a point $M = (x, y, z, i)$, where x, y, z are the spatial coordinates of the voxel's centre and i its intensity. Let us call M the point representing the voxel V . One could use the representing points directly but the minimisation algorithm would then be inefficient because of the data volume.

To avoid this problem, we compute a more compact representation of the images. The idea is to split recursively the image into “quadrilaterals” until each “quadrilateral” contains only voxels whose representative points can be approximated by a 4D hyperplane with an error smaller than ϵ , where ϵ is a parameter of the splitting algorithm³. More precisely, the algorithm is as follow:

1. Split the image I into two blocks of same size in the x (resp. y and z) direction and get the images I_{1x} and I_{2x} .
2. If x (resp. y and z) is the direction which yields the smallest error when one approximates I_{1x} and I_{2x} by a 4D hyperplane, then go to point 1 with $I = I_{1x}$ and $I = I_{2x}$ if the error is greater than ϵ .

As a result, the image is represented by a set of “quadrilaterals” with different

³The error can be the maximum of the points to plane distances or the average of these distances.

sizes containing relatively homogeneous voxels. A 4D centroid and a 4D hyperplane are attached to each “quadrilateral”.

The points M_j used to describe the scene image in the minimisation algorithm are simply the centroids attached to each "quadrilateral" obtained by recursively splitting the image. To efficiently compute the function \mathbf{CP}_{4D} , the model image is also split into “quadrilaterals”. A 4D kd-tree [43] is calculated based on the centroids B_j resulting from the recursive split. During minimisation, given a 4D point M , $\mathbf{CP}_{4D}(M)$ is computed as follows (see figure 1).

- From the kd-tree, we first find the centroid B_j the closest to M . This centroid corresponds to a “quadrilateral” Q_j .
- $\mathbf{CP}_{4D}(M)$ is the closest point to M onto the hyperplan H_j approximating the “quadrilateral” and lying in this “quadrilateral”.

4.2 A coarse-to-fine multi-resolution strategy

The parameter ϵ of the image splitting algorithm allows us to control the quality of the approximation. The smaller ϵ , the better the approximation. However, when ϵ is large, the number of points/hyperplanes describing the image is small. Thus, ϵ allows us to control **the resolution**.

It is also important to control the “quantity” of accepted deformation. We first compute rigid displacements, then affine transformations and finally spline or radial basis deformations. For the spline class of transformations, we can choose the number of control points and the parameter λ controlling

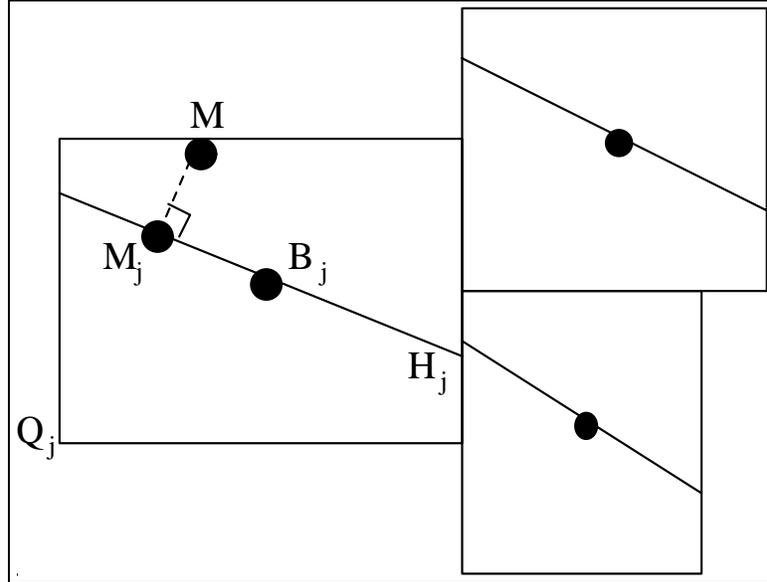


Figure 1: The computation of \mathbf{CP}_4D is done in two stages. First, we compute the centroid B_j the closest to M and then, we project M onto H_j .

the importance of the bending term in the criterion definition in order to control **the “quantity” of allowed deformation**. With radial basis functions (RBF), it is even simpler to control this quantity. Indeed, we can still choose the parameter λ , but in addition it is easy to add locally some centers where more deformation is needed. In fact, RBF do not have the topology problem encountered with splines and which makes local refinement of the deformation difficult (see [49] for such an adaptative scheme with splines). In fact, RBF are extremely convenient even if an efficient and well-justified strategy for adding centres during the iterations still has to be developed. A more complete discussion of this idea of “allowed deformation” is presented

in conclusion.

The strategy that we propose to try to avoid the local minima during minimisation uses these two properties. At the beginning, we choose a low resolution and we compute rigid displacements. The biggest structures are then first registered. Then, progressively during the minimisation process, we decrease ϵ in order to enhance the quality of the approximation of the images and we allow more and more deformation from affine transformations to spline or RBF deformations.

Of course, this description of the strategy is very qualitative. In practice, the choice of the functions controlling these resolution and deformation parameters depends on the anatomical regions to register. But they have not been difficult to find in practice.

5 Results on brain data

In this section, we present two examples of applications of our volume registration algorithm to 3D brain MR images.

5.1 Rigid registration with intensity correction

The top two images of figure 2 are two slices of same slice index of two MR images (taken at different times) of the same brain. One can observe that (1) geometric registration is necessary since the two slices do not correspond to each other (for example, the eyes are visible in one image and are not visible in the other one) (2) the left image is much brighter than the right one.

The algorithm described in this paper enables us to compute at the same time a rigid displacement to superimpose the two images and a gain coefficient to correct the intensity. The bottom image of figure 2 shows the slice corresponding to the top left image after registration and resampling of the 3D image corresponding to the top right image.

For the registration, approximately 50000 points are used to describe the images. The resolution does not vary during the iterations. The CPU time is 5 minutes. After registration, the average mean distance between matched 4D points is 0.8 mm and the average difference of intensity is 2.8 (the intensity in the images is between 0 and 255). One can observe that after registration the two slices look much more similar both from the geometric and intensity viewpoints.

Figure 3 demonstrates that one can correct at the same time for rigid displacement and non linear intensity difference. Two images A and B were simulated from the same MR image \mathcal{I} : A was created by applying a rigid displacement to \mathcal{I} and B by applying a second order polynomial multiplicative bias field. The experimental conditions are the same as in figure 2 except that the intensity correction function is a second order polynomial.

5.2 Matching with an atlas

Figures 4, 5 and 6 show an example of spline registration of two MR images of two different brains. One of the two images has been manually segmented into anatomical regions (courtesy of Ron Kikinis, at the Brigham and Women's Hospital, Boston). It can be used as an anatomical atlas. Matching

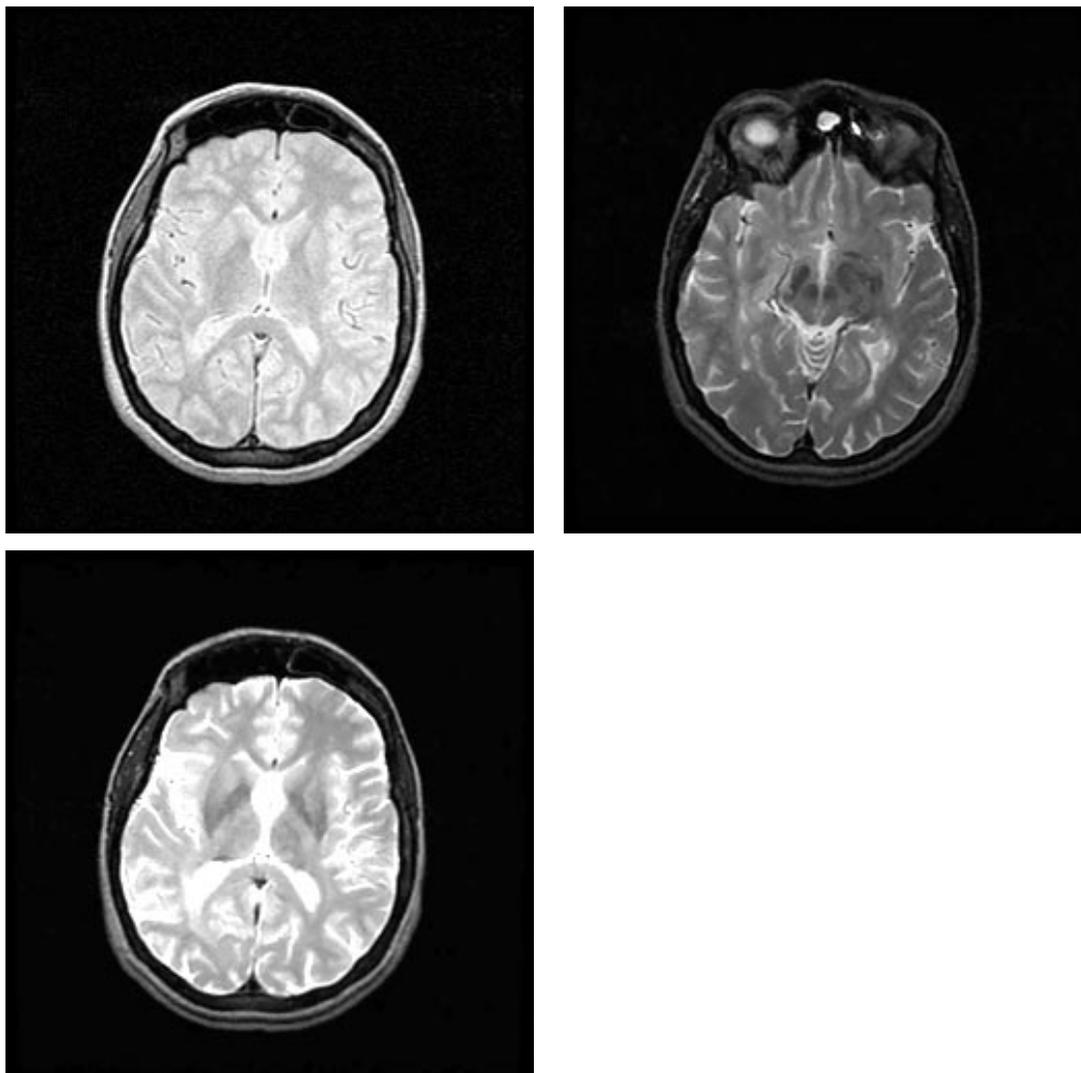


Figure 2: **Top**: two axial slices of same index of two MR images (left A, right B) of the same brain before registration. We thank Pr. Ron Kikinis, Brigham and Women's Hospital (Boston) for these images. **Bottom, left**: the slice of image B after registration, resampling and intensity correction corresponding to the slice of the top-left image. One can compare pixel by pixel the left two images.

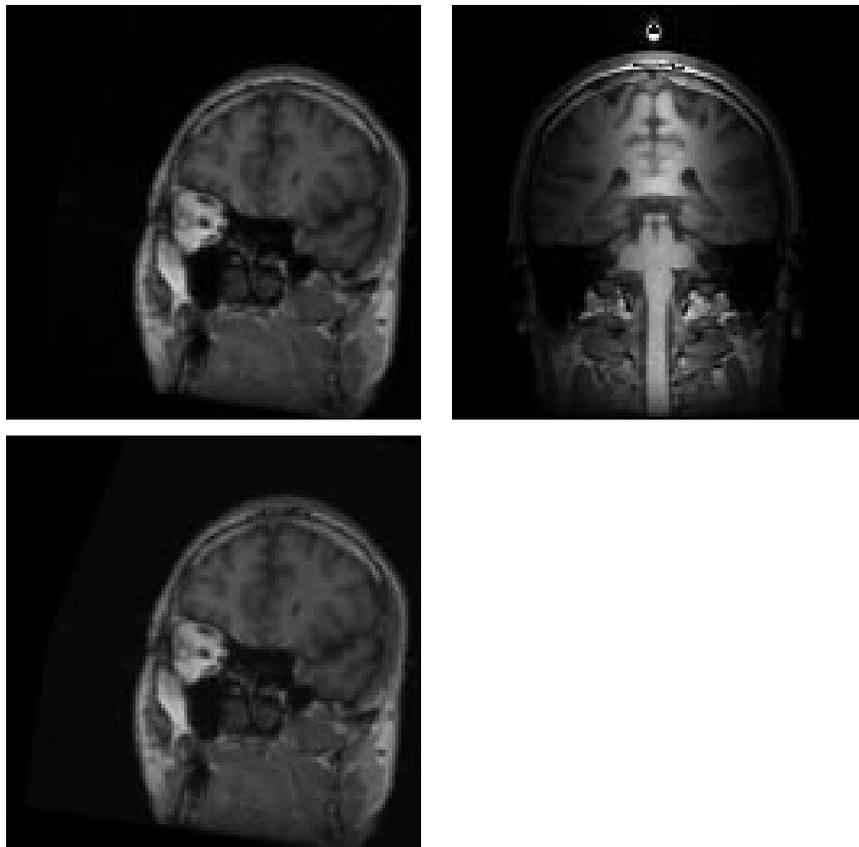


Figure 3: **Top**: two frontal slices of same index of two images simulated from the same MR image \mathcal{I} . Image A (left) is the same as \mathcal{I} but we applied a rigid displacement, image B (right) is the same as \mathcal{I} but we applied a second order polynomial multiplicative bias field. **Bottom, left**: the slice of image B after registration, resampling and intensity correction corresponding to the slice of the top-left image. One can compare pixel by pixel the left two images.

enables us to label automatically the second image from the voxel-to-voxel correspondence.

Registration uses in each image approximately 50,000 points. The computed deformation is a volume spline function (see appendix A) with $15 \times 15 \times 15$ control points. There is no intensity correction. The CPU time is 25 minutes. The resolution varies linearly from 10000 points to 50000 points during the deformation process and λ (the parameter controlling the smoothing term of the criterion) varies linearly from 5 to 2.

Registration is not perfect, though it is quite good. There are two reasons for this:

- we should use more 4D points to describe the images at the end of the process,
- the spline deformation is not sufficiently local ($15 \times 15 \times 15$ control points is the maximum that we can use because of memory limitations).

We believe that these two problems will soon be fixed as processors become more powerful.

Note that we should not have the second problem with radial basis functions since the centres can be located where needed. The difficulty then is to decide *where* they are needed. The answer may be based on discussions with an anatomist or from learning. Moreover, because the deformation is global, we can guarantee that the result of our non-rigid registration makes sense. Note also that before registration, it was quite difficult to identify in the two images the corresponding structures while it is quite easy after registration.

If a more local registration were to be necessary, the output of our algorithm could be used as the input for techniques like [12, 50] which are perhaps more sensitive to their initialization but which are more local.

6 Results on heart data

A common examination for detection of cardiac ischemia is the stress-rest comparison in myocardial perfusion studies provided by Nuclear Medicine.

Nuclear medicine imaging provides 3D density maps of blood perfusion non-invasively. In the stress-rest study, 2 perfusion maps of the heart muscle are taken: one obtained after an injection in the blood of the tracer at rest (rest image) and the other after the injection of the tracer during maximal exercise (stress image). These images are isometric, their size is 64x64x64 for a pixel size of 5 mm.

Comparing the two images provides a classification of areas of the myocardium into 3 main classes:

- The intensity distribution is *normal* in both rest and stress images.
- There is at least one region with abnormally low count rate densities in both rest and stress images. The abnormality is fixed, and this denotes a myocardial *infarction*, or in some cases a very narrow stenosis, with resting hypoperfusion and a hibernating or stunned myocardium.
- There are one or more regions of low count rate densities in the stress image, but the densities are normal in the rest image. The abnormality is said to be transient, and connotes stress *ischemia*.

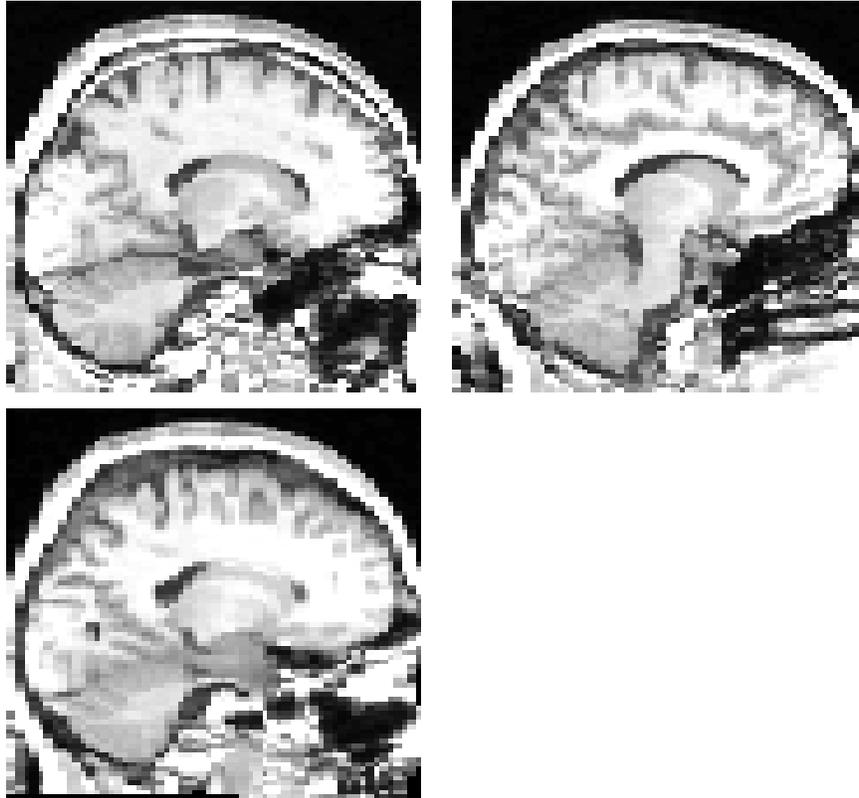


Figure 4: **Top**: two slices sagittal of same index coming from two MR images of two different brains before registration (left C, right the atlas). We thank Pr. Ron Kikinis, Brigham and Women's Hospital (Boston) for these images. **Bottom, left**: the slice of the resampled atlas corresponding to the top left image after non rigid spline registration. The two left images can be compared voxel by voxel.

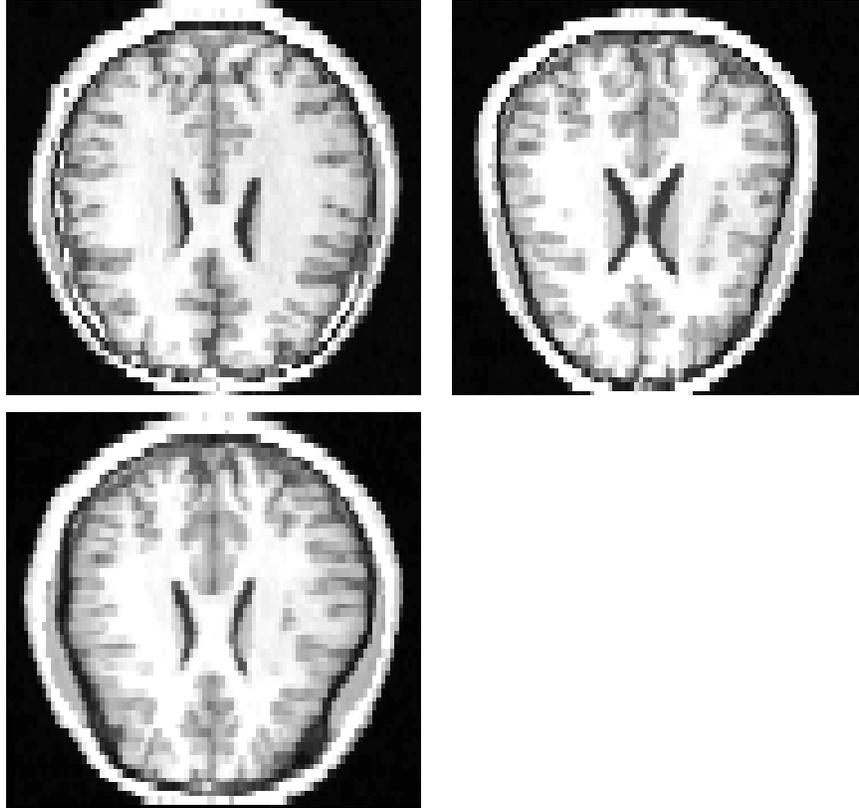


Figure 5: **Top**: two axial slices of same index coming from the same images than the ones shown figures 4 and 5 (left C, right the atlas). **Bottom, left**: the slice of the resampled atlas corresponding to the top left image after non rigid spline registration.

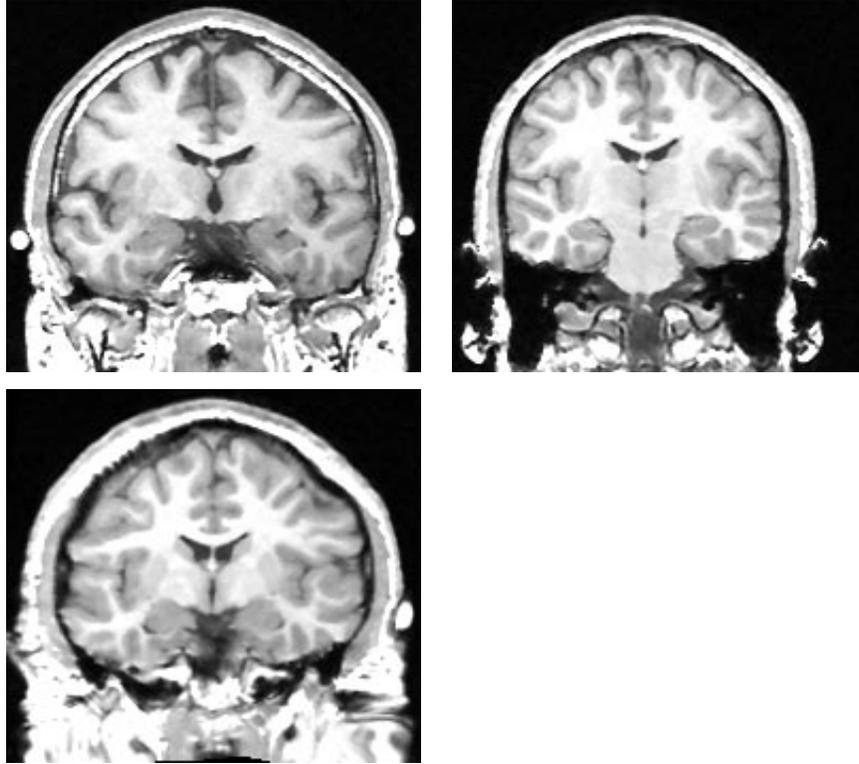


Figure 6: **Top**: two frontal slices of same index coming from the same images than the ones shown figures 4 and 5 (left C, right the atlas). **Bottom, left**: the slice of the resampled atlas corresponding to the top left image after non rigid spline registration.

6.1 Stress-rest affine registration

Because the patient is not in the same position when the stress and rest images are taken, registration is necessary. Because, for some pathologies, the heart does not have the same size at stress and at rest, we compute an affine transformation between stress and rest images. The intensity correction function is a global multiplication factor. The CPU time is 30 seconds.

The top two images of figure 7 are two slices of same index of two SPECT images of the heart: one at stress, the other one at rest. The middle image shows the result of the registration and resampling of the rest image. The bottom two images show the difference between the stress and rest images respectively before (left) and after (right) registration. One can observe that the difference image is much darker after registration. We have chosen a healthy patient for illustration because it demonstrates that a bad registration can induce wrong diagnosis. Indeed, one can see in the bottom left image bright areas which could be interpreted as ischemia. But, as demonstrated in the bottom right image (which is very dark), these bright areas actually come from bad registration. This demonstrates that registration is necessary.

We experimented with the registration algorithm presented in this paper on a database of 40 stress-rest pairs. The results are visually very good. In order to give a more quantitative quality measure, we would have to do the same experiments as we did in [28] for 3D-3D non rigid surface matching [21]. Note that it is important for this problem to deal explicitly with occlusion, because ischemia appears as occlusion in the images (it corresponds to holes present in one image and not in the other one).

6.2 Normalisation of stress-rest images

Ischemia corresponds to coronary vessels which are narrowed. One of the goals of stress-rest examination is to determine such vessels and to evaluate the importance of the occlusion. To do that, the nuclear medicine cardiologist makes use of an atlas associating with each defect, depending on its position, the defective vessel. It turns out that, if we are able to normalise the stress and rest images into a standard geometry, the diagnosis is easier [28]. Indeed, the location of the defects no longer depends on the particular shape of each heart. Thus, the association defect-vessel is easier.

In practice, we propose to perform spline registration between stress-rest images and a template image. As a result, we get normalised stress and rest images in the geometry of the template. We have chosen a normal heart to be the template. Figure 8 shows an example of normalisation of a stress-rest pair. Intensity correction (function g) is a global multiplication factor. The CPU time is 2 minutes (Dec alpha workstation).

Again, the results on the 40 patients database look good even though a more complete validation as in [28] is necessary.

7 Conclusion

We have presented in this paper a new and efficient algorithm for registration-intensity correction of 3D images. This algorithm is an extension of the original ICP algorithm to volume registration and deals explicitly with the occlusion problem. The computed transformations are rigid, affine and spline

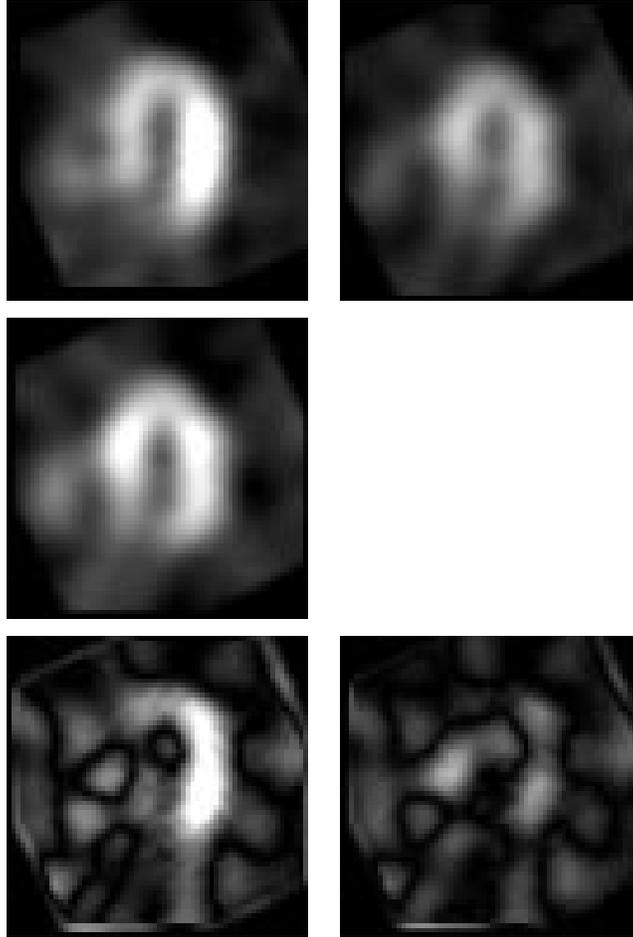


Figure 7: **Top**: two slices of same index coming from a rest (left) and a stress (right) 3D image of the same heart before registration. We thank Pr. Michael Goris, Stanford University Hospital for these images. **Middle**: the slice of same index of the stress image resampled into the geometry of the rest image after affine registration. **Bottom, left**: one slice of the 3D difference image between the stress and rest images **before** affine registration. **Bottom, right**: one slice of the 3D difference image after **after** affine registration. The ischemic areas correspond to bright areas in this 3D difference image. Note that the two difference images have been multiplied by the same global factor for visualization. Note also that bad registration (left) produces bright areas which could be interpreted as ischemia: this demonstrates that accurate registration is necessary to make a reliable diagnosis.

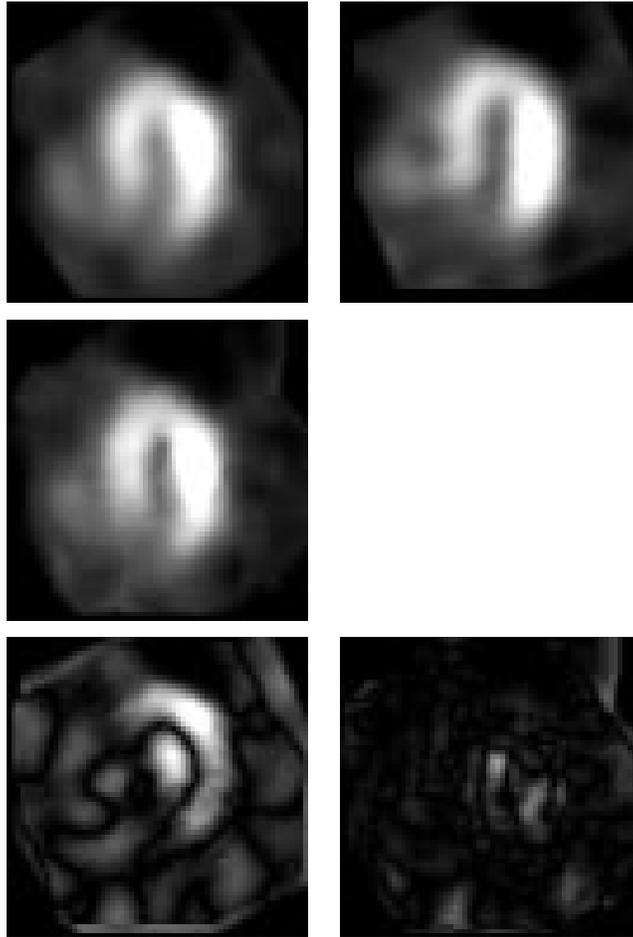


Figure 8: **Top**: two slices of same index coming from two rest images of two different hearts. The heart presented in the left image is the template. **Middle, left**: the slice of same index of the rest image (top right) resampled into the geometry of the template (top left) after non rigid spline registration. From a geometric point of view, this image and the top left image are equivalent. **Bottom**: two slices of same index of the difference image before (left) and after (right) registration. One can observe that the difference image is much darker after spline registration. Note that the two difference images have been multiplied by the same global factor for visualization.

or radial basis functions. The experiments demonstrate the validity of our approach, though a complete clinical validation is yet to be done. In future, we plan to extend this work at least in four directions.

The first extension is to use a database to determine how should be constrained the set of possible deformations for brain-atlas matching. Even if our approach enables us to minimise an explicitly defined criterion, the choice of the class of deformation is not guided by anatomy. The set of a priori possible deformations from one brain to another is not easy to determine, though the Talairach atlas is an interesting attempt. The problem is that if we choose too many parameters to describe the set of possible deformations, then the registration can fail or have no anatomical significance. On the other hand, if we don't have enough parameters then the registration won't be accurate. It is then essential both for robustness and for anatomical significance to determine the best deformation space and to describe it with as few parameters as possible.

One approach is to determine this set by learning. This has been investigated in [17, 48], [40] and [48] respectively for displacement fields, modal analysis and Fourier analysis. Similar to [17], we could compute the set of possible values of the control points on a database and then perform Principal Component Analysis on this set. At stage two of the extended ICP algorithm, we would solve the linear system under the constraint that the solution must lie in the subspace described by the first eigenvectors of PCA. It would still consists in solving a linear system and the algorithm should be faster and more robust.

The second extension would be based on basic statistics. A common

problem in medical imaging is to analyse the change in shape of an organ. For example, for the stress-rest registration problem presented in section 6, we would like to know if the shape difference between stress and rest is significant with respect to a pathology. Another problem is to compare the shape of an organ with respect to an “average” shape. For example, some psychiatrists claim that schizophrenia implies a change in asymmetry of the patients brains [18].

For both problems, one can compute two sets of deformations: one group for the controls and one group for the patients. Based on the parameters of the deformations, we could perform an Hotelling T^2 test to validate or invalidate the null hypothesis that the means of the two groups are equal [10, 32]. This would enable us to answer the question: “is there a significant difference in shape between the two groups”. If this is the case, then one should consider Fisher’s linear discriminant [10, 32]. Indeed, each coordinate would tell us how important each mode of variation is for discrimination. It would then be possible to perform 3D graphical animations to show the change in the shape of the brains between the two groups. This animation would probably be easier to interpret by the physicians than a vector of deformation. Finally, in case of significant difference between the two groups, the Fisher’s linear discriminant could allow us, given an image which is not in the database (and then for which the diagnostic is unknown), to give a measure which is related to the probability of the image of belonging to one of the two classes [6].

The third problem is the use of volume spline or RBF registration for motion tracking in sequences of 3D images (MR or SPECT, especially heart

images). In the same way as we extended the technique presented in [21] to time [20], we believe that it would be worth doing it for the algorithm presented in this paper. Indeed, time sequences of 3D images will probably be more and more common and we believe that analysis of such images presents a challenge [38, 1, 36, 41].

The last problem concerns the validation of non-rigid registration. How is it possible to say that one algorithm is better than another ? The easiest way of doing it is maybe to ask a radiologist or an anatomist to identify features or landmarks in the images and to construct an error measure based on these points. But we believe that the answer to the question should depend on the goal of the registration. For example, we are interested in detecting blobs of activation in f-MRI data. After non-rigid registration, we use the software described in [26] to detect such blobs. The software tells us, for each blob, the probability of observing it by chance. It seems to be reasonable (even if it is debatable) to say, since registration and blobs detection are independent, that the best non-rigid registration is the one which gives the more reliable blobs. Actually we believe that for non rigid-registration used in the context of functional atlas building, the criterion should explicitly depend on the “blob’s reliability”. This is what we are going to concentrate on in the near future.

8 Acknowledgements

We wish to thank Mike Brady for encouraging us and helping us in developing the ideas presented in this paper. We thank Daniel Reissfeld for criticizing

these idea and proposing alternatives. We thank Michael Goris and Ron Kikinis for providing us with the data used for the experiments presented in this article.

A Least squares estimation of deformation

At stage two of the algorithm described in this paper, we have to estimate, given a set of matched points, the best deformation in the least squares sense corresponding to the matches. Actually, because of performance constraints but also to avoid local minima, we restrict ourself to deformations for which such an estimation leads to the resolution of a linear system. These are the deformations f of the form:

$$f(x, y, z) = (\sum_i A_i^x F_i(x, y, z), \\ \sum_j A_j^y F_j(x, y, z), \\ \sum_k A_k^z F_k(x, y, z)),$$

where the A_i^x, A_i^y, A_i^z are the parameters and the functions F_i, F_j, F_k are fixed.

From a practical point of view, the code for solving such least squares systems with Singular Values Decomposition (SVD) is available in [44]. Hence, there is almost no programming to be done for such deformations. The programmer just has to define the functions F_i, F_j, F_k .

We have chosen to implement the algorithm presented in this paper with spline and radial basis functions deformations.

A.1 3D-3D volume spline deformations

More precisely, we use the class of functions defined by tensor products of spline basis functions:

$$f(x, y, z) = \left(\begin{array}{l} \sum_{i,j,k} C_{ijk}^x B_i(x) B_j(y) B_k(z), \\ \sum_{i,j,k} C_{ijk}^y B_i(x) B_j(y) B_k(z), \\ \sum_{i,j,k} C_{ijk}^z B_i(x) B_j(y) B_k(z), \end{array} \right),$$

where the $C_{ijk} = (C_{ijk}^x, C_{ijk}^y, C_{ijk}^z)$ are the control points and the B_i are 1-D B-spline functions with regularly distributed knots. In our formulation, the class of 3D-3D spline functions is only described by moving the control points.

In the definition of the criterion, a smoothing energy is added to the least squares term on the position in order to control the regularity of the solution. This energy is expressed as a second-order Tikhonov stabilizer. For instance, for the x coordinate u of the function f :

$$SE(u) = \iiint_{\mathbb{R}^3} \left[\frac{\partial^2 u^2}{\partial x^2} + \frac{\partial^2 u^2}{\partial y^2} + \frac{\partial^2 u^2}{\partial z^2} + 2 \frac{\partial^2 u^2}{\partial x \partial y} + 2 \frac{\partial^2 u^2}{\partial x \partial z} + 2 \frac{\partial^2 u^2}{\partial y \partial z} \right]$$

The criterion is the sum of the two energies, a multiplying factor λ weights the importance of the smoothing energy with respect to the position energy.

Because the criterion is quadratic in C_{ijk}^x , C_{ijk}^y and C_{ijk}^z , the least squares minimisation in stage 2 of the algorithm presented in this paper is a linear system solution, which is quite efficient in practice.

We use 3D-3D spline functions for efficiency but also because they have interesting geometric properties:

- the 3D-3D spline functions and their derivatives are easy to compute thanks to the “de Casteljau” algorithm,
- the intrinsic rigidity properties of B-splines provide regular 3D-3D functions,
- a data point has a local influence: to evaluate a spline function at a given point, only $(K + 1)^3$ control points are necessary, where K is the spline order.

For more details about spline functions, see [7, 22, 21].

A.2 3D-3D volume radial basis deformations

RBF is the class of functions defined by linear combinations of radial functions, namely functions whose value at one point M depends only on the distance $d(C_i, M)$ between that point and the centre point C_i :

$$F_i(M) = f_i(d(C_i, M)), F_j(M) = f_j(d(C_i, M)), F_k(M) = f_k(d(C_i, M)).$$

As in the case of spline functions, one can add a smoothing term to the criterion. Radial basis functions (RBF) have been widely used in the context of probability and neural networks [6] since it can be proven that, roughly, it is possible to approximate any continuous function if the parameters are chosen correctly. Note also that thin-plate splines are special cases of RBF [7].

The choice of the centres C_i and of the functions f_i, f_j, f_k is crucial. We use functions of the form

$$e^{-\frac{d(C_i, M)^2}{\sigma_i^2}},$$

where σ_i controls the locality of the influence of each centre. We have achieved good results in practice and we believe that RBF are convenient to programme and use. We have not yet searched for an optimal automatic procedure to determine the best centres nor the σ 's. The interested reader is referred to [6]. For more details about RBF, see [2, 6, 7]

References

- [1] A. Amini, R. Curwen, A. Klein, T. Egglin, J. Pollak, F. Lee, and J. Gore. Physics based snakes, kalman snakes, and snake grids for feature localization and tracking in medical images. In *Information Processing in Medical Imaging (IPMI '95)*, Brest, June 1995.
- [2] N. Arad and D. Reisfeld. Image warping using few anchor points and radial functions. *Computer Graphics Forum*, 14(1):35–46, March 1995.
- [3] N. Ayache. Medical computer vision, virtual reality and robotics. *Image and Vision Computing*, 13(4):295–313, May 1995.
- [4] R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *CVGIP*, 46:1–21, 1989.
- [5] P. Besl and N. McKay. A method for registration of 3–D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, February 1992.
- [6] C. M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, 1995.
- [7] F. L. Bookstein. Principal Warps: Thin-Plate Splines and the Decomposition of Deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, june 1989.

- [8] L. G. Brown. A Survey of Image Registration Techniques. *ACM Computing Surveys*, 24(4):325–375, December 1992.
- [9] G. Champleboux, S. Lavallée, R. Szeliski, and L. Brunie. From accurate range imaging sensor calibration to accurate model-based 3–D object localization. In *Proceedings of the IEEE Conference on Vision and Pattern Recognition*, Urbana Champaign, June 1992.
- [10] C. Chatfield and A. Collins. *Introduction to multivariate analysis*. Chapman and Hall, 1980.
- [11] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. *Image and Vision Computing*, 10(3):145–155, 1992.
- [12] G. Christensen, R. Rabbit, M. Miller, S. Joshi, U. Grenander, T. Coogan, and D. VanEssen. Topological properties of smooth anatomic maps. In *Information Processing in Medical Imaging (IPMI '95)*, Brest, June 1995.
- [13] L. Cohen. Use of Auxiliary variables in computer vision problems. In *Proceedings of the Fifth International Conference on Computer Vision (ICCV '95)*, Boston, June 1995.
- [14] A. Colchester, J. Zhao, C. Henri, R. Evans, P. Roberts, N. Maitland, D. Hawkes, D. Hill, A. Strong, D. Thomas, M. Gleeson, and T. Cox. Craniotomy Simulation and Guidance using a Stereo Video Based tracking system (VISLAN). In *Visualization in Biomedical Computing*, Rochester, Minnesota, October 1994.
- [15] A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, and G. Marchal. Automated Multi Modality Image Registration Using Information Theory. In *Information Processing in Medical Imaging (IPMI '95)*, Brest, June 1995.

- [16] D. Collins, A. Evans, C. Holmes, and T. Peters. Automated 3D Segmentation of Neuro-Anatomical Structures. In *Information Processing in Medical Imaging (IPMI '95)*, Brest, June 1995.
- [17] T. Cootes, A. Hill, and C. Taylor. Rapid and accurate Medical Images Interpretation Using Active Shape Models. In *Information Processing in Medical Imaging (IPMI '95)*, Brest, June 1995.
- [18] T. J. Crow. Temporal Lobe asymmetry as the key to the aetiology of schizophrenia. *Schizophrenia Bulletin*, 16:433—443, 1990.
- [19] C. Davatzikos. Nonlinear registration of brain images using deformable models. In *Proceedings of MMBIA*, San Francisco, June 1996.
- [20] J. Declerck, J. Feldmar, and N. Ayache. Definition of a 4D continuous polar transformation for the tracking and the analysis of the LV motion. In *CVRMed-MRCAS'97*, Grenoble, March 1997. (Also INRIA Research Report # 3039).
- [21] J. Declerck, J. Feldmar, M. Goris, and N. Ayache. Automatic registration and alignment on a template of cardiac stress and rest SPECT images. In *MMBIA '96*, San Francisco, June 1996.
- [22] G. Farin. *Curves and surfaces for computer aided geometric design, a practical guide*. Academic Press, 1988.
- [23] J. Feldmar and N. Ayache. Rigid, Affine and Locally Affine Registration of Free-Form Surfaces. *The International Journal of Computer Vision*, 18(2), May 1996. Published in two parts in ECCV'94 (rigid and affine) and CVPR'94 (locally affine). Also INRIA Research Report # 2220.
- [24] J. Feldmar, N. Ayache, and F. Betting. 3D-2D projective registration of free-form curves and surfaces. *The Journal of Computer Vision and Image Understanding*, 1997. Published in three parts: CVRMed'95, ICCV'95 and IPMI'95. Also INRIA Research Report 2434.

- [25] K. Friston et al. Spatial Registration and Normalization of Images. *Human Brain Mapping*, 2:165–189, 1995.
- [26] K. Friston et al. Statistical Parametric Maps in Functional Imaging: A General Linear Approach. *Human Brain Mapping*, 2, February 1995.
- [27] J. Gee, L. Lebricquer, C. Barrillot, and D. Haynor. Probabilistic matching of Brain Images. In *Information Processing in Medical Imaging (IPMI '95)*, Brest, June 1995.
- [28] M. Goris, J. Declerck, J. Feldmar, M. Goris, and N. Ayache. Elastic Transformations in Referential Quantification of Scintigraphic Images. In *CAR'96*, Paris, June 1996.
- [29] W. Grimson, T. Lozano-Perez, W. W. III, G. Ettinger, S. White, and R. Kikinis. An Automatic Registration Method for Frameless Stereotaxy, Image Guided Surgery, and Enhanced Reality Visualization. In *IEEE Proceedings of Computer Vision and Pattern Recognition 1994 (CVPR'94)*, Seattle, USA, June 1994.
- [30] F. Hemler, P. V. D. Elsen, T. Sumanaweera, S. Napel, J. Drace, and J. Adler. A quantitative Comparison of Residual Errors for three different Multimodality registration techniques. In *Information Processing in Medical Imaging (IPMI '95)*, Brest, June 1995.
- [31] S. Joshi et al. Hierarchical brain mapping via a generalized Dirichlet solution for mapping brain manifolds. In *Proceedings of the SPIE Conf on Geom. Methods in Applied Imaging*, July 1995.
- [32] W. J. Krzanowski. *Principles of Multivariate Analysis*. Oxford University Press, 1988.
- [33] G. Malandain. *Filtrage, topologie et mise en correspondance d'images médicales multidimensionnelles*. PhD thesis, Ecole Centrale de Paris, Septembre 1992.

- [34] J.-F. Mangin, V. Frouin, I. Bloch, B. Bendriem, and J. Lopez-Krahe. Fast Nonsupervised 3D Registration of PET and MR Images of the Brain. *Journal of Cerebral Blood Flow and Metabolism*, 14(5):749–762, 1994.
- [35] J.-F. Mangin, F. Tupin, V. Frouin, I. Bloch, R. Rougetet, J. Regis, and J. Lopez-Krahe. Deformable Topological Models for Segmentation of 3D Medical Images. In *Information Processing in Medical Imaging (IPMI '95)*, Brest, June 1995.
- [36] E. McVeigh. MRI of Myocardial Function: Motion Tracking Techniques. *Magnetic Resonance Imaging*, 14:137–150, 1996.
- [37] C.-H. Menq, H.-T. Yau, and G.-Y. Lai. Automated precision measurement of surface profile in CAD-directed inspection. *IEEE Trans. RA*, 8(2):268–278, 1992.
- [38] F. Meyer, R. Constable, A. Sinusas, and J. Duncan. Tracking Myocardial deformation using spatially-constrained velocities. In *Information Processing in Medical Imaging (IPMI '95)*, Brest, June 1995.
- [39] M. Miller and et.al. Mathematical Textbook of Deformable Neuranatomies. In *Proc. of the National Academy of Sciences*, volume 90, pages 11944–11948, 1996.
- [40] C. Nastar, B. Moghaddam, and A. Pentland. Generalized Image Matching: statistical Learning of Physically-Based Deformations. In *Proceedings of the Fourth European Conference on Computer Vision (ECCV '96)*, Cambridge, May 1996.
- [41] J. Park, D. Metaxas, and L. Axel. Volumetric Deformable Models with Parameters Functions: a new approach to the 3D Motion Analysis of the LV from MRI-SPAMM. In *Proceedings of the Fifth International Conference on Computer Vision (ICCV '95)*, Boston, June 1995.

- [42] X. Pennec and N. Ayache. Randomness and Geometric Features in Computer Vision. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR'96)*, San Francisco, Cal, June 1996.
- [43] F. P. Preparata and M. I. Shamos. *Computational Geometry, an Introduction*. Springer Verlag, 1985.
- [44] W. H. Press et al. *Numerical Recipes in C*. Cambridge University Press, 1987.
- [45] S. Sandor and R. Leahy. Towards Automated Labelling of the Cerebral Cortex using a Deformable Atlas. In *Information Processing in Medical Imaging (IPMI '95)*, Brest, June 1995.
- [46] C. Studholme, D. Hill, and D. Hawkes. Multi-Resolution Voxel Similarity Measures for MR-Pet Registration. In *Information Processing in Medical Imaging (IPMI '95)*, Brest, June 1995.
- [47] G. Subsol. *Construction automatique d'atlas anatomiques à partir d'images médicales tridimensionnelles*. PhD thesis, Ecole Centrale, 1995.
- [48] G. Szekely, A. Kelemen, C. Brechbüler, and G. Gerig. Segmentation of 3D Object from MRI Volume Data Using Constrained Elastic Deformations of Flexible Fourier Surface Models. In *First International Conference on Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed '95)*, Nice, April 1995.
- [49] R. Szeliski and S. Lavallée. Matching 3-D Anatomical surfaces with non-rigid volumetric deformations. In *Proceedings of the IEEE Workshop on Biomedical Images Analysis (WBIA'94)*, Seattle, Washington, June 1994. Also in AAAI 1994 Spring Symposium Series. Application of Computer Vision in Medical Image Processing, Stanford University, 1994.

- [50] J.-P. Thirion. Fast Non-Rigid Matching of 3D Medical Images. In *Medical Robotics and Computer Aided Surgery (MRCAS'95)*, pages 47–54, Baltimore, November 1995.
- [51] P. Thompson and A. Toga. A surface-based technique for warping three-dimensional images of the brain. *IEEE Trans. on Medical Imaging*, 15(4):402–417, august 1996.
- [52] P. van den Elsen, E. Pol, and M. Viergever. Medical Image Matching. A Review with classification. *IEEE Engineering in Medicine and Biology*, 12(4):26–39, 1993.
- [53] P. Viola and W. M. Wells III. Alignment by Maximisation of Mutual Information. In *Proceedings of the Fifth International Conference on Computer Vision (ICCV '95)*, Boston, June 1995.
- [54] W. Wells III, E. Grimson, R. Kikinis, and F. Jolesz. Adaptive Segmentation of MRI Data. In *First International Conference on Computer Vision, Virtual Reality and Robotics in Medicine (CVRMed '95)*, Nice, April 1995.
- [55] R. Woods, S. Cherry, and J. Mazziotta. Rapid Automated Algorithm for Aligning and Reslicing PET images. *Journal of Computer Assisted Tomography*, 16(1):1–14, 1992.
- [56] Z. Zhang. Iterative Point Matching for Registration of Free-Form Curves and Surfaces. *the International Journal of Computer Vision*, 13(2):119–152, 1994. Also Research Report No.1658, INRIA Sophia-Antipolis, 1992.