Genetic Algorithms for a Robust 3-D MR-CT Registration

Jean-Michel Rouet, Jean-José Jacq, and Christian Roux, *Senior Member, IEEE*

***Abstract*—The aim of this paper is to present an original usage of genetic algotithms as a robust search space sampler in application to 3-D medical image elastic registration.**

An overview of the standard steps of a registration algorithm is given. We focus on the genetic algorithms use and particularly on the problem of extraction of the optimal solution among the final genetic population. We provide an original encoding scheme relying on a structural approach of point matching and then point out the need for a local optimization process. We then illustrate the algorithm with a concrete registration example and assert the results with a direct multivolume rendering tool. Finally, the algorithm is applied on the vanderbilt medical image database to assert the robustness and in order to compare it with other techniques.

***Index Terms*—Genetic algorithm, global optimization, image registration, medical imaging, pattern matching.**

I. Introduction

MEDICAL IMAGING of the brain is used by surgeons and clinicians more and more to make diagnoses or plan therapies. Nevertheless, the number of modalities available and the amount of data (a set of 2-D slices constituting 3-D images) make them very difficult to use directly. Moreover, each modality gives only a limited kind of information, and very often two or more modalities of the same patient are used to obtain a good understanding of the sensed material. Here we will focus on two modalities often used for the visualization of the brain. The first one giving good structural information (i.e., good contrast for the bones) is the computed tomography (CT) scanner, whereas magnetic resonance imaging (MRI) provides good information on soft tissues (such as white matter and grey matter).

A frequent problem arises when images from different modalities (or when using only one modality but taking images at different times) need to be compared. The mental superposition is often hard to realize, particularly on 3-D images, and a more practical operation requires computer assistance while matching two 3-D images. This leads to the general problem of superimposing two different 3-D images on the same scene (possibly seen by different sensors). Registration refers to finding a geometrical transformation that corresponds any point from one 3-D image to its homologous on the other image. Furthermore, MRI sensors are known to induce slight distortions on the images and very often prevent the user from easily finding a good rigid match (a 3-D translation plus a 3-D rotation).

Manuscript received June 1, 1999; revised November 17, 1999.

The authors are with the Department of Image and Information Processing, École Nationale Supérieure des Télécommunications de Bretagne, 29285 Brest Cedex, France.

Publisher Item Identifier S 1089-7771(00)03952-2.

This paper introduces an original approach to overcome difficulties of multimodality global elastic registration. A structural description of the volumes and a combinatorial optimization process allows us to solve, in a reasonable time, a nonlinear problem presenting numerous degrees of freedom. In another way, our algorithm could be seen as a generalized stochastic Hough transform using the principles of natural selection. Thus, the great originality of our approach resides from the particular use of a genetic algorithm (from encoding to the examination of the genetic space of the solutions).

After giving some generalities about medical image registration in Section II, we will introduce in Section III the geometrical considerations of registration we used for our algorithm and focus (in Section IV) on genetic algorithms and global optimization problems. Section V will detail our registration method, and a registration example will be given in Section VI. Finally, we will discuss the validation of the algorithm on the vanderbilt RREP dataset in Section VII and then conclude with a discussion on some open issues and future directions.

II. Background

Several methods already exist to find a rigid registration between two medical images in a reasonable time, and very few are able to perform elastic registration. The major drawbacks of standard registration techniques are their sensitivity to the initial positionning of the images, the issue of dealing with multimodal images, and the prohibitive processing time of elastic registration. Since the purpose of this paper is not to give an exhaustive description of medical image registration techniques, we refer the reader to the works of Van den Elsen [1], Brown [2], Lavallée [3], and Maintz [4] who have written extensive and farily up-to-date surveys.

Nevertheless, we want to point out the main steps of a classical registration algorithms. The common ground in registration techniques lies in linear sequential data processing. It begins with an optional *preprocessing* of the raw data (mainly filtering or segmentation/extraction of interesting features on the images). Then, it requires the definition of a *fitness function* (measuring the goodness of fit between the target image and the transformed image) dealing either directly with grey-level information or with extracted or intrinsic features. Moreover, the registration requires a model of the mapping between the two image fields (*warping model*) [5]. The more common approaches make use of a global geometric model. The warping model might be chosen either rigid (translations, rotations and scaling only) or nonrigid, and either global or local. Then, one has to find the parameters of the geometric transformation (within the warping model) that maximizes

the fitness function. This step is an *optimization procedure*, which can be deterministic, iterative, or stochastic, and either global or local. Classical local optimization procedures are Hill climbings or iterative searchs, while global optimizations are often stochastic algorithms like simulated annealing, stochastic hill climbing, or genetic algorithms [6] (this last reference uses a generic genetic algorithm to address the monomodality elastic registration problem and constitutes a background to the algorithm proposed in this article). Finally, the last important step of a registration algorithm is the *fusion* and/or interpretation of the registered images. A classical approach might be to isolate the information (such as bone structures from CT scans and grey matter from MRI) of each modalities and recompose them on a hybrid volumic image. The visualization might then be achieved slice by slice or by using a 3-D rendering tool.

III. Geometric Considerations

The images we want to register are essentially CT scans and MRI images of the head (skull + brain) as shown in Section VI. Though the skull is a rigid bone structure, we believe that a rigid registration is not efficient enough to take into account, on the one hand, the slight distortions usually induced by MRI sensors and, on the other hand, the possible evolution of the soft structures (e.g., brain). As the registration must be weakly elastic and smooth, we prefer a global transformation over a local one.

A. Generic Notations Used

Given two 3-D images $F_1(p), p \in D_1 \subset I\!R^3$ and $F_2(p), p \in D_2 \subset I\!R^3$, registering F_2 on F_1 is equivalent to determining the geometric transformation $T^\star$ that matches $F_2^\star$ with F_1. T^* is globally defined as an application $T : D_1 \rightarrow I\!R^3$, and $F_2^\star$ is defined for $p \in D_1$ by

$$F_2^\star(p) = \begin{cases} F_2(T(p)), & \text{if } T(p) \in D_2 \\ 0, & \text{elsewhere.} \end{cases} \tag{1}$$

The exact definition of T will be given in Section III-B while the fitness function maximized for $T^\star$ is defined in Section III-C.

B. Global Elastic Transformation

A global elastic transformation, simple enough to be computed quickly (this is a key point while making use of a genetic algorithm where the evaluation of each individuals is the crucial point) and elastic enough to model MRI distortions, is the *trilinear transformation* which is a polynomial warping. This is the extension of a shearing along three axis, and it is defined by

$$\begin{aligned} p &= (x, y, z)^T \quad p' = (x', y', z')^T = T(p) \\ p' &= \sum_{i=0}^{1}\sum_{j=0}^{1}\sum_{k=0}^{1}(a_{i,j,k}, b_{i,j,k}, c_{i,j,k})^T x^i y^j z^k \end{aligned} \tag{2}$$

where $a_{i,j,k}$, $b_{i,j,k}$, and $c_{i,j,k}$ are the coefficients of the transformation. Pointing out that i, j, and k are either 0 or 1, the dimension of the search space $\mathcal{S}$ is then $L = 3 \times 2^3 = 24$.

With such a direct equation (2), it is not easy to express the boundaries of the search space. Nevertheless, the transformed location of eight points from image F_1 to image F_2 enables an indirect description of the transformation involved. Let $\mathbf{X_1}$ (respectively, $\mathbf{Y_1}$, $\mathbf{Z_1}$) be a vector that represents the x (respectively, y, z) coordinates of eight points from F_1 $(x_1(i), y_1(i), z_1(i))$, $i = 1, \cdots, 8$ and $\mathbf{X_2}$ (respectively, $\mathbf{Y_2}$, $\mathbf{Z_2}$) be their transformed coordinates $(x_2(i), y_2(i), z_2(i))$. Let us denote by $\mathbf{M}$ the 8×8 transformation matrix deduced from the $\mathbf{X_1}$, $\mathbf{Y_1}$, and $\mathbf{Z_1}$ components and in relation with (2). The line i of $\mathbf{M}(\mathbf{X}_1, \mathbf{Y}_1, \mathbf{Z}_1)$ is

$$\begin{aligned} &[1, x_1(i), y_1(i), z_1(i), \\ &\quad x_1(i)y_1(i), x_1(i)z_1(i), y_1(i)z_1(i), x_1(i)y_1(i)z_1(i)]. \end{aligned}$$

The transformed vectors $\mathbf{X_2}$, $\mathbf{Y_2}$, and $\mathbf{Z_2}$ are given by (3), shown at the bottom of the page.

The computation of the polynomial coefficients $\mathbf{A}, \mathbf{B}, \mathbf{C}$ is performed by the inversion of the former linear relations. It is important to notice that, given eight pairs of points (or eight correspondences between F_1 and F_2), there exists a unique trilinear transformation T managing the pairing of these points (the existence of the transformation T is ensured if every eight points and their correspondents are distinct).

Therefore, the aim of our algorithm will be to find pairs of homologous points between the two images and achieve global warping with the computed 24 coefficients.

Even if we have pointed out that eight pairs are enough to compute the transformation coefficients, a least squares extension of (3) is possible; this will imply an overdetermined system computation rather than a mere matrix inversion. These properties will be used and discussed more in detail in Section V.

C. Geometric Distance and Fitness Function

The optimality of the transformation T is defined by a criteria of distance or fitness. The distance d between two images is inversely proportional to the goodness of fit f_i of the two images. $T^\star$ is then the transformation that maximizes f_i or minimizes d for each transformation T of the search space with respect to the images F_1 and F_2. It is similar to choosing $T^\star$

$$T^\star = \arg\min_{T \in \mathcal{S}} d(T(F_1), F_2) = \arg\max_{T \in \mathcal{S}} f_i(T(F_1), F_2). \tag{4}$$

The estimation of the goodness of fit is not a trivial process. Many distances exist with different advantages and drawbacks. A grey-level correlation is not usable here because we need to calculate the fitness function many times. Thus, a fast feature based fitness function has been chosen. A common feature available on both MRI and CT-scan images is the air–skin interface.

$$\begin{cases} \mathbf{X}_2 = \mathbf{M}(\mathbf{X}_1, \mathbf{Y}_1, \mathbf{Z}_1) \cdot \mathbf{A} \\ \mathbf{Y}_2 = \mathbf{M}(\mathbf{X}_1, \mathbf{Y}_1, \mathbf{Z}_1) \cdot \mathbf{B} \\ \mathbf{Z}_2 = \mathbf{M}(\mathbf{X}_1, \mathbf{Y}_1, \mathbf{Z}_1) \cdot \mathbf{C} \end{cases} \quad \text{with} \quad \begin{cases} \mathbf{A} = [a_{0,0,0}, a_{1,0,0}, a_{0,1,0}, a_{0,0,1}, a_{1,1,0}, a_{1,0,1}, a_{0,1,1}, a_{1,1,1}]^t \\ \mathbf{B} = [b_{0,0,0}, b_{1,0,0}, b_{0,1,0}, b_{0,0,1}, b_{1,1,0}, b_{1,0,1}, b_{0,1,1}, b_{1,1,1}]^t \\ \mathbf{C} = [c_{0,0,0}, c_{1,0,0}, c_{0,1,0}, c_{0,0,1}, c_{1,1,0}, c_{1,0,1}, c_{0,1,1}, c_{1,1,1}]^t \end{cases} \tag{3}$$

A preprocessing step addresses the extraction of theses points from both modalities, providing us with two numerical surfaces. The computation of a distance between two 3-D images is then reduced to the computation of a distance between two numerical surfaces.

Let S_1 (S_2) be the surface extracted from F_1 (F_2), and $d_E(p_1, p_2)$ be the Euclidean distance between points p_1 and p_2. Then the distance[1] d_a between S_1 and S_2 with respect to the transformation T can be defined by

$$d_a(S_1, S_2 \,|\, T) = \frac{1}{\text{card } S_1} \sum_{p_1 \in S_1} \min_{p_2 \in S_2} (d_E(T(p_1), p_2)). \quad (5)$$

At this point, we notice that the creation of a distance map $\mathcal{D}_2$—a 3-D image on the domain D_2 where each point p represents $\min_{p_2 \in S_2} (d_E(p, p_2))$—greatly improves the execution time. The new expression of the distance using a distance map becomes

$$d_m(S_1, S_2 \,|\, T) = \frac{1}{\text{card } S_1} \sum_{p_1 \in S_1} \mathcal{D}_2(T(p_1)). \quad (6)$$

However, the distance d_m remains quite long to compute when the cardinal of S_1 or S_2 is large (typically within the range 30 000–60 000). Thus, we have implemented a stochastic variation of the distance d_m that we will denote d_s. A certain number ($n <$ card S_1, e.g., $n = 100$) of points are randomly chosen in the set S_1, and the distance is computed with this new point set noted S_1^n.

$$d_s(S_1, S_2 \,|\, T) = \frac{1}{n} \sum_{p_1 \in S_1^n} \mathcal{D}_2(T(p_1)). \quad (7)$$

Furthermore, we found that this distance d_s was not robust enough in the presence of outliers, i.e., first, when dealing with points whose transformed location does not belong to the distance map, and second, when a large number of points randomly chosen in S_1 does not have homologous matches in S_2 (this is a problem resulting from nonoverlapping parts of the volumes). This is why we introduced a fitness function f filtering on the large values of the distance and keeping the relevant information; in doing so, the sum is only performed on randomly chosen points p_1 whose transformed location belongs to the distance map (these are n^0 ($n^0 \leq n$) points that might be seen as part of a subset S_1^0). The corresponding fitness is given by the relationship

$$f(S_1, S_2 \,|\, T) = \frac{1}{n^0} \sum_{p_1 \in S_1^0} \exp - \frac{(\mathcal{D}_2(T(p_1)))^2}{\sigma^2}. \quad (8)$$

Fig. 1 shows the application of the fitness function [the part under the sum of (8)] on a distance map. The image in Fig. 1(a) is a slice extracted from the original 3-D distance map of a MRI volume, and the image in Fig. 1(b) features the usage of (8) on the same slice with $\sigma = 4$.

If no randomly chosen points of S_1^0 are transformed close enough (depending on the value of σ, usually $\sigma = 4$) to S_2, (8) implies that $f = 0$. If the n^0 points randomly chosen in S_1^0 have a transformed location with respect to T onto S_2, then f is maximal and equals 1.

[1]Note that the term distance used here is not a real mathematical distance, because the expression does not satisfy $d(a, b) = d(b, a)$. We only have $d(a, a) = 0$ and $d(a, b) > 0$.

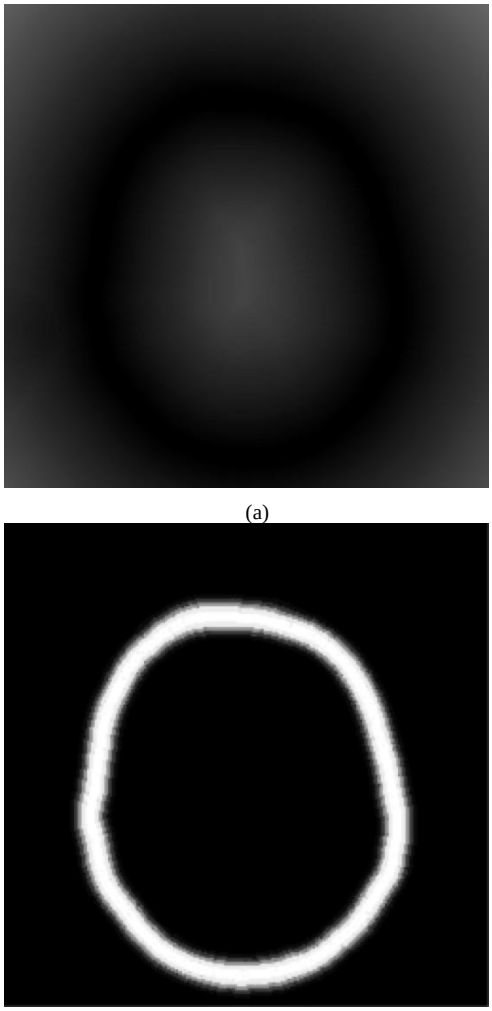

(a)

(b)

Fig. 1. (a) Distance map. (b) Fitness function mapping with $\sigma = 4$.

IV. Global Optimization Using a GA

This section intends to present the genetic algorithms (GA's) as a global optimization technique. After a short overview of the general functioning, we will focus on two particular GA's used (those working on real parameters and those working on indexed sets) and then we will discuss some issues regarding the GA's.

A. Basics of a Genetic Algorithm

Genetic algorithms are adaptive generate-and-test algorithms using the principles of natural population genetics and natural selection mechanisms. The main idea is to encode the parameters of a problem in a chromosome as a list of bits (see Fig. 2(a) where L parameters with a fixed length q are encoded). We then consider a population of chromosomes (potential solutions) and let them evolve with respect to three rules (we refer the reader to [7] for more details): 1) the selection/reproduction that evaluate the performance of each individual and duplicate only the good

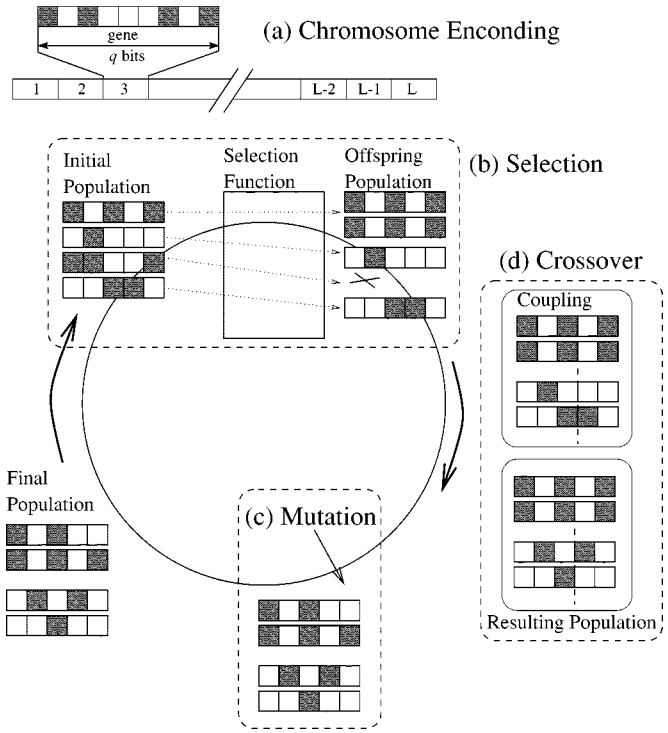

Fig. 2. A complete GA cycle from: (a) the chromosome encoding to (b) the genitors selection, (c) the mutation (random search space exploration), and (d) the crossover (hereditary conservation of parameters).

one (depending on a selection scheme) [Fig. 2(b)]; 2) the mutation which randomly modifies some bits of the chromosome and which is responsible for the search space exploration [Fig. 2(c)]; and 3) the crossover which swaps the genetic material between some elected couples of chromosomes [Fig. 2(d)]. After a few generations, and because of the nature of the exploration of the search space and the suppression of the bad individuals, the overall performance of the population increases.

It is important to notice that a postprocessing step (local optimization) is essential for the accuracy of the algorithm. Actually, even if a GA is efficient enought to avoid being trapped in local extrema, it also a slow convergence to the optimum solution. Thus, a GA might always be considered a tool giving a good approximation of the final solution and might always require the use of a local optimization procedure. Moreover, we want to point out that it is more meaningful to use the information included in all the chromosomes of the final genetic population rather than only make use of the few best elements.

B. Genetic Optimization and Image Matching

As we will see later, for our registration algorithms we used two kinds of GA. The first one, used to achieve an initial rigid registration, works directly with real parameters, and the second one works on an indexed search space.

1) GA with Real Parameters: For a rigid registration, the parameters we are looking for are real numbers representing a 3-D translation and the angles of a 3-D rotation (a total of six real numbers). We could have used a classical GA working with strings of bits that represent numbers, but, in order to avoid sampling problems, we directly encoded a chromosome as a succession of six real numbers (three for the translation and three for the rotation). Furthermore, we introduced a repartition constraint on the population over the search space using the principle of the Latin squares [8]. The latter ensures that, from the initial population elaboration to the final genetic population, a certain percentage of the population is within a subset of regions of the search space. In fact, we control that the projection of the population over each of the six axes of the search space is uniform in the sense that a minimum amount of points is present in each subdivision of the parameter space.

This new constraint entails a redefinition of the mutation and the crossover operators. The new mutation is bound to a cell of the search space, and thus it is considered to be a local search. This new mutation operator is therefore more efficient for the fine tuning of the good solutions (faster and more accurate convergence) within a cell because its scope is limited to a small area. Moreover, we notice that this scope limitation of the mutation operator is only possible because each cell of the search space is reachable via the succession of crossover operations. Concerning the crossover operator, the swapping of entire parameters between two chromosomes does not modify the distributions of the projections over the axis of the search space, thus our new constraint does not imply a modification of our crossover operator.

Moreover, we ensure a dynamic control of the homogeneity of the population (necessary because of the selection operator that might disturb the chromosome's distribution) by moving some chromosomes from over-represented regions to deserted regions.

Among all possible crossovers (one-point, cyclic, uniform), the uniform one has been chosen. For a selected pair of chromosome to undergo the crossover modification, we look iteratively through the genes (parameters) of the chromosome and decide whether or not to exchange the values of the genes (with respect to a random exchange rate). Thus, we give up the idea of a crossover point which usually is order-dependent. In addition, we have experimented with some evolutionist operators, such as the hybrid simplex method used in [9], without success. We noticed that making use of simple operators (like mutation and uniform crossover) a great number of times is more efficient than using more complicated operators (such as the evolutionist operator) which are slower.

2) GA on Indexed Sets: The second GA we used does not look for integer nor real parameters but for eight pairs of points which, given (3), is an equivalent problem. It works on an indexed set (unidimensional search space) where a potential solution (a chromosome) is a combination of eight different indices. Thus, the encoding is directly a combination of eight different indices. The mutation corresponds to changing one index by another index of the search space, while the crossover simply consists of swapping some index between two chromosomes. The only constraint added to the mutation and crossover operators is a set of coherence tests on the chromosomes. As we will see in Section V, an index represents a couple of points from the target image to the registered image. Thus, a chromosome is coherent only if the 16 points designated by the eight couples correspond to eight distinct points on the target image and eight distinct points on the registered image. So, we check that

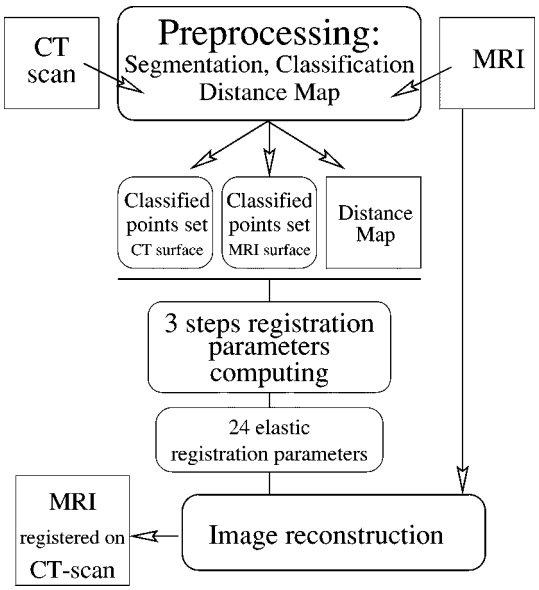

Fig. 3. General synopsis of the registration algorithm.

a mutation operator or a uniform crossover does not introduce duplicate origin or destination points.

Therefore, the GA working on the indexed set is not used as a stand-alone optimization tool. It is introduced here as a smart search space sampler that will work in synergy with a local optimization tool. Thus, the GA will provide a valuated sampling of the search space to the final local optimization process and will be responsible for the robustness, while the local optimization improves the accuracy of the registration process.

C. Limits of the GA's

The main drawback of a GA is the risk of premature convergence to a local extremum, sometimes due to a bad initialization of the search space. This is why we use a dynamic control of the search space (with the Latin squares principle). The choice of the stopping criterion is also a main issue. How long must it take to have a good accuracy quickness ratio? Knowing these limits of a GA, we see that it is almost mandatory to connect it to a local optimization process. Thus, a GA can be seen as a parser of a search space looking for a set of good *optima* further exploited by a nonstochastic optimization algorithm. At this point, we therefore consider a GA not only as a robust optimization technique, but also as a smart search space sampler.

V. Registration Procedure

This section presents the overall approach used to solve the registration between a CT scan and an MRI image of a human head. The chart in Fig. 3 describes the different parts of the process that will be detailed in the next paragraphs. The first stage addresses the acquisition of the images and the preprocessing for the extraction of the interesting features (external interface air/skin). This leads to the extraction of two sets of points (S_1 and S_2), one for each modality. Topologically, the sets S_1 and S_2 describe numerical surfaces, and thus it is possible to compute for each element of these sets the local curvature (we use the Gaussian and maximum curvature). We define a classification function $\mathcal{C}$ returning the belonging class of a point. If p_1 and p_2, two points from S_1 and S_2, have a similar local curvature (for instance, the same Gaussian curvature sign), then $\mathcal{C}(p_1) = \mathcal{C}(p_2)$. Moreover, we compute from the second image (the MRI volume) a Euclidean distance map which is a volume where the value of each voxel represents the Euclidean distance to the numerical surface S_2 (e.g., see [10] for the computation of the distance map). For more details on the functioning of these preprocessing algorithms, we refer the reader to [11].

Then a three-step (see Section V-A) algorithm computing the parameters of the transformation is achieved. Finally, the image warping is performed in the last stage (Section V-B).

We would like to remind the reader that the underlying idea behind our robust registration process is to extract two sets of points belonging to the air–skin interface of each images and use them to perform a global rigid matching (step 1) that initializes a robust point-matching (steps 2 and 3) algorithm. The robustness of the point-matching algorithm is ensured by a good filtering of misleading point associations (the outliers) and then by the use of a GA as a smart search space sampler (step 2) further explored by a local search optimization (step 3).

A. A Three-Step Robust Registration Algorithm

1) Global Rigid Matching: A global rigid registration is used to initialize the robust point-matching algorithm. We assume that the elasticity of the transformation we are looking for is weak because the brain skull is a rigid structure and because the distortion involved by the imaging devices is not very large. Thus, a rigid matching, even if not extremely accurate, is a good initial guess of the final transformation. On the basis of such a rigid matching (along with its own confidence factor which relies on the median distance between the two numerical surfaces and with respect to the rigid transformation parameters), one may eleminate aberrant associations of nonhomologous surface points.

As stated before, the rigid transformation parameters are six real parameters, or a 3-D translation vector and three rotation angles. These parameters are determined by a GA working directly on real numbers (see Fig. 4). Note that the search space is bounded to plausible intervals (we choose a maximum translation vector and look for rotation angles between $-\pi$ and π).

2) Smart Sampling of the Point-Matching Search Space: The 24 parameters of the global trilinear transformation we are looking for are fully defined by the knowledge of any eight points from the 3-D target surface and their correspondents from the 3-D surface to register. Thus, we use a point-matching algorithm that will try to determine the correspondences between the points of the two images. Since the size of the two surfaces is too large to attempt a direct search among all the possibilities, we construct a list of possible point pairs $(\mathcal{L})$ using the best rigid registration T^r with a median error distance d_{med} and the class partitioning $\mathcal{C}$. This results in the following rule: a point p_1 from the S_1 and a point p_2 from S_2 are associable if and only if p_1 and p_2 belong to the same curvature class and satisfy (9) where d_{med} is the median error distance between S_1 and S_2 using the rigid transformation T^r (see the notation used in Section III). We set a maximum

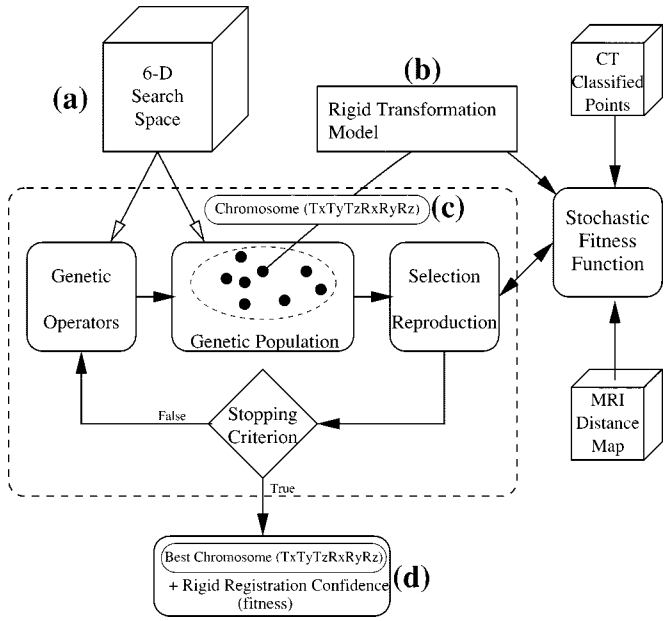

Fig. 4. Rigid registration process using a GA.

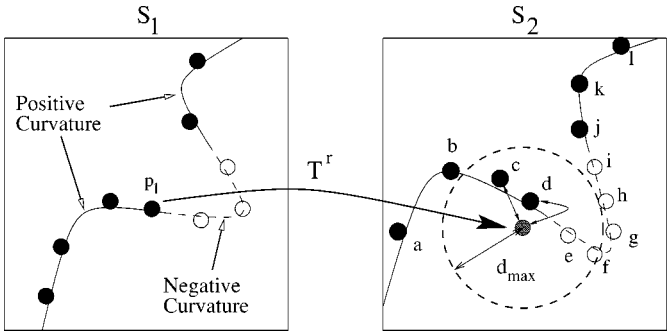

Fig. 5. The point pair selection according to the curvature class and the rigid registration mean error.

allowed distance $d_{\max}$ as the double of the median distance to eliminate the outliers. A threshold value of 2 has been chosen as a good compromise between the filtering of outliers and the search for elastic transformations

$$d_E(T^r(p_1), p_2) < 2d_{\text{med}} = d_{\max}. \tag{9}$$

Thus, the set $\mathcal{L}$ is defined by

$$\mathcal{L} = \{(p_1, p_2) \in S_1 \times S_2 / \mathcal{C}(p_1) = \mathcal{C}(p_2) \text{ and } d_E(T^r(p_1), p_2) < 2d_{\text{med}}\}. \tag{10}$$

Fig. 5 illustrates the idea of the point pair creation given two classified point sets, a rigid transformation T^r, and a median distance d_{med}. On this drawing, the only points that can be associated with p_1 from S_1 are the points noted c and d because they belong to the same curvature class (the class of positive curvature in this example), and they are in the same neighborhood (with respect to the T^r transformation).

We then use a GA algorithm to look for eight pairs through the index set ($\mathcal{L}$) resulting in good elastic transformation. The purpose of this second step is not to directly determine the best octet of index giving the elastic transformation, but rather to sample the search space (set of octet of index) by keeping elements of the search space that are present in the final genetic population. In doing so, we are using the GA as a robust stochastic sampler of a wide search space giving a new reduced search space usable by a local (deterministic) optimization process (step 3). The overview of the second step is similar to Fig. 4 where the search space shown in (a) is now $\mathcal{L}^8$, the transformation model shown in (b) is the elastic trilinear one, the chromosome encoding shown in (c) is a set of eight different indices (a point from $\mathcal{L}^8$), and finally the output of the algorithm shown in (d) is the final genetic population (set of chromosomes).

3) Step 3: Fine Tuning the Solution: Using directly the best element of the final genetic population is not considered as optimal. Each element of this population contains a significant piece of information. Thus, we used an algorithm managing the extraction of the information contained in each chromosome. We know that eight is the exact number of point pairs required for the computation of the 24 parameters of the global trilinear transformation. Nevertheless, with a greater number of point pairs, it is still possible to compute these 24 parameters, with the difference being that using eight pairs results in an interpolation and using more pairs result in an approximation. Thus, we see that using a greater number of couples will help to filter out the noise inherent to the segmentation of the surfaces and give the possibility of using a larger part of the genetic population. This is in great part responsible for the robustness of our registration algorithm.

The post-analysis consists in starting from the best eight pairs found by the GA, and dynamically looking for other pairs that improve the fitness when merged to the previously selected pairs until no other pair is able to significantly increase the fitness. Because the order in which the pairs are merged to form the final solution is important, it is sometimes possible to again improve the final solution by removing some of the selected pairs. Therefore, a second pass in the postanalysis algorithm consists of parsing the selected pairs and checking if a removal increases the overall fitness. Achieving this two-pass process, the information available in the final genetic population is quickly (typically a few seconds) and fully exploited.

The synergistic execution of these three steps (rigid registration, search space sampling, and local optimization) is the core of the robustness and the accuracy of our optimization algorithm.

B. Image Reconstruction and Fusion

Once the registration parameters are computed, one has to evaluate the quality of the registration. Whereas several numerical methods exists to estimate this quality, we use a visual assessment of the registration. After achieving a direct warping with respect to (2), we have the target image F_1 (as defined in Section III-A) and $F_2^\star$, which is the warping of F_2 aligned with F_1. On the one hand, the visual assessment could be done slice by slice with, for instance, a superposition of the bone structures, extracted from the CT scan, onto the soft tissue from the MRI image. On the other hand, a direct multivolume rendering (DMVR algorithm presented in [12]) can be used to evaluate the intersection and the union of the fuzzy surfaces of our volumes.

VI. Results: A Registration Example

We propose a registration example achieved on real CT scans and MRI images. After a quick presentation of the raw data, we will discuss the different registration steps, ranging from segmentation to the computation of the parameters and the visual assessment.

A. Data Sets

The data we are working on are a CT scan and an MRI volume of the head, whose physical dimensions are (256×256×91) voxels for the CT scan and (256×256×68) voxels for the MRI. It is interesting to notice that the voxels are noncubics and have a different size for both images (the scope of the CT scan is (232.82×232.82×165.81 mm^3), and (265.45×265.45×141.06 mm^3) for the MRI volume).

Our algorithms are able to deal with slightly anisotropic regular grids. There is no need to provide isotropic images which prevent a resampling stage.

Fig. 6 presents 3-D renderings of the initial CT and MR volumes. Note that the artifact volumes at the back of the CT image are part of the raw images. These volumes could have been removed in the segmentation step in order to facilitate the registration (the fitness computation). Nevertheless, we have chosen to keep these outlying structures in order to assert the robustness of our registration algorithm.

The segmentation step will not be addressed in this paper (see [11]). We will consider already segmented images, i.e., two point sets S_1 and S_2, the distance map $\mathcal{D}_2$, and the curvature classification $\mathcal{C}$ as the inputs of the registration algorithm.

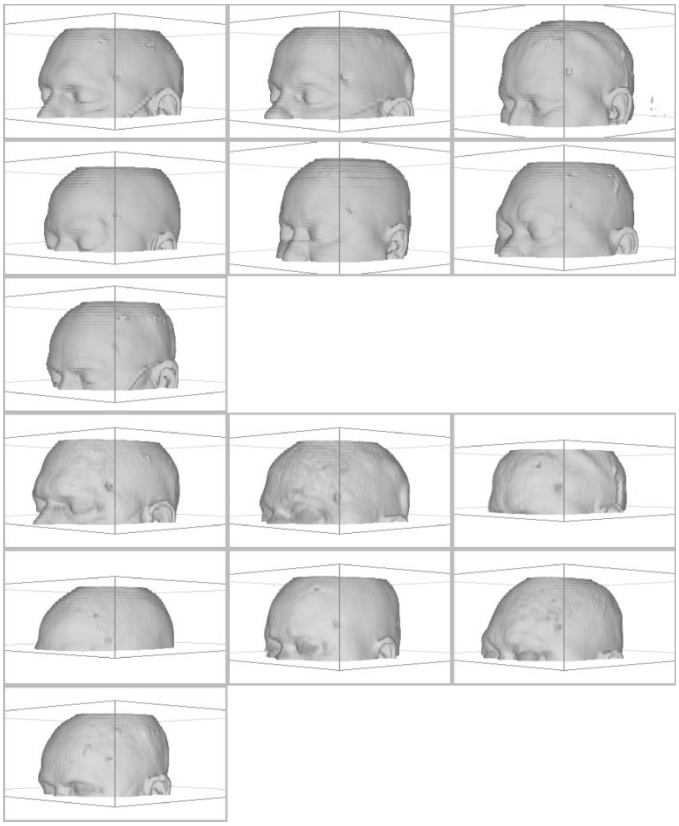

Fig. 6. Direct volume renderings of the initial CT and MR-T1 images of the Vanderbilt dataset. The seven top images (patients) correspond to the CT modality while the seven bottom images relate to the MRI modality of the respective same patients.

B. Rigid Registration

The first step of our algorithm is the establishment of a global approximation of the matching through a rigid warping. With the images used, this step is achieved in an average of less than 20 s (the PC used ran Linux 2.2.7 on a Bi-PentiumPro 233 MHz with 192 Mo EDO RAM). The quality of this registration is appreciable on Fig. 7 (top). Table I summarizes the results obtained for each step of the algorithm.

This rigid registration step requires the tuning of some parameters, mainly for the genetic algorithm. We have studied the influence of the genetic parameters and noticed that, if they were chosen in an acceptable (standard) range of values, they only interfere with the processing time. To be precise, we have chosen a population of 500 chromosomes, evolving over 100 generations. The mutation probability was set to 1% while the crossover probability was tuned at 80%. Finally, we have taken $n = 100$ stochastic points for the computation of the goodness of fit [see (7) and (8)].

The physical values of the rigid transformation found is a 3-D translation vector of $[11.30, 22.49, -15.86]$ in millimeters and the three rotation angles (along the X, Y, Z axis) $[0.135, 0.093, -0.072]$ in radians.

C. Elastic Point-Matching

The knowledge of the rigid registration fitness (74.9%) allows us to compute the confidence interval for the couple creation (with respect to the rigid transformation parameters and, whenever possible, to the differential geometry classification). Equation (9) gives a confidence interval $d_{\max} = 4.3$ mm.

With the 35 866 segmented points on the MRI image and the 63 710 points on the CT scan, more than a million couples satisfying the previous criterion are created in less than 4 min without using the curvature information. This step is the most CPU consuming part of the algorithm. We note that using curvature classes improves the rapidity of the algorithm; Table I shows the computation time and the registration quality with one curvature class (i.e., no classification), then with three curvatures classes.[2]

We then apply the second GA which will look for a minimal set of matching points. As stated above, this step is considered to be a smart search space sampling because it transforms a search space into another search space of the same kind with only the relevant points.

For this GA, we have used the same parameters as for the first GA except for the number of points for the stochastic fitness computation (here $n = 200$), and it ran for about 51 s. The best chromosome gives a stochastic fitness of 79.5%. This GA has 500 chromosomes of 8 couples each. Thus, it had 4 000 potential couples to evaluate. After the convergence of the GA, only a hundred different couples exist in the final population and

[2]Classification was made on the Gaussian curvature value of crest-line points. Four classes are set up. A class with large positive Gaussian curvature, one with large negative Gaussian curvature, a class with almost nil curvature, and a reject class for other points. The first three classes generally represent 50% of all the points of the segmented surfaces, while the reject class is not used.

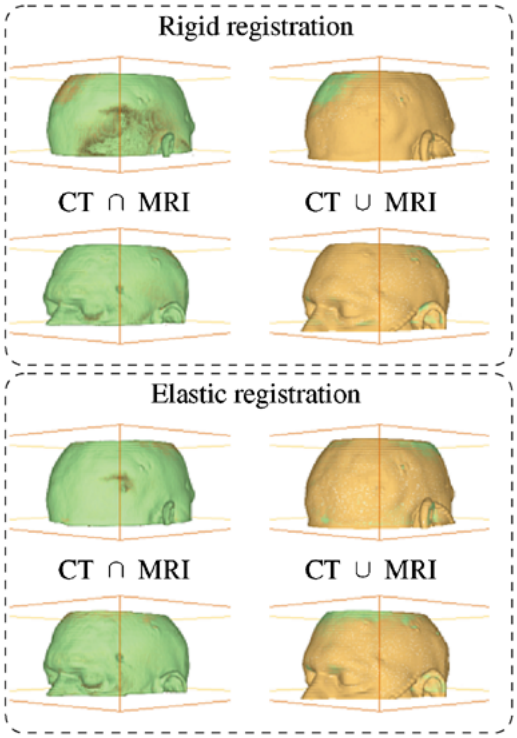

Fig. 7. Visual assessment of the median registration achieved on patient 1, from MR-T1 to CT. The fuzzy intersection renderings denotes a good overall alignment of the external surfaces, with, as predicted, a better registration in the elastic case than in the rigid approximation case.

TABLE I
GLOBAL RESULTS OF THE REGISTRATION EXAMPLE

Step	CPU Time	Fitness (%)	Figure
Segmentation	< 2 mn	-	8
Rigid Reg.	16 s	74.909 (±3.6)	9(a)
Without curvature classification			
Couples Gen.	3 mn 44 s	-	5
Elastic (Best)	51 s	79.545 (±1.191)	-
Fine tuning	1 s	80.810 (±0.872)	9(b)
Overall Proc.	4 mn 36 s	80.810	9(b)
With three curvature classes			
Couples Gen.	46 s	-	5
Elastic (Best)	48 s	79.126 (±2.526)	-
Fine tuning	1 s	81.117 (±2.600)	10
Overall Proc.	1 mn 35 s	81.117	10

constitute the new search space of the third step (local optimization).

D. Fine Tuning of the Solution

Finally, we ran the local search on the final genetic population (new search space). Given that we have a small search space (for the example used we have a hundred of couples), the fine tuning is very fast. The final fitness obtained is roughly 80% to 82%, depending whether we used or not a Gaussian curvature classification, which is not very relevant because of the stochasticity of the fitness. Better appreciations might be done on Figs. 7 (bottom) and 8.

As shown in Table I, statistics have been conducted to assert the accuracy of the algorithm. On this table we show the standard deviation of the fitness value for about 20 registrations of

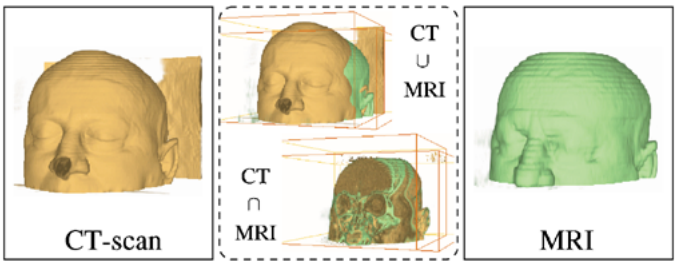

Fig. 8. Direct multi-volume renderings of the initial medical data.

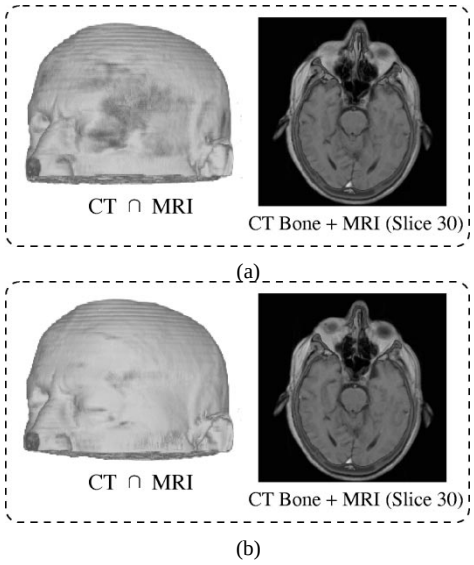

Fig. 9. (a) Rigid and (b) elastic (without curvature information) registration.

the same volumes with the same registration parameters. In addition, we have studied the influence of the genetic parameters on the behavior of the algorithm (accuracy, velocity and robustness). This study has demonstrated that the choice of the genetic parameters has a negligble influence on the accuracy, except for very degenerated values (i.e., almost no mutation and crossover, too few genetic iterations, or too small genetic population).

E. Visual Assessment

Figs. 6–8 summarize the results obtained at the different important steps of the algorithm. It is important to appreciate the information given by the multivolume renderings, especially by the intersection of two volumes. If two volumes are misregistered, it is obvious that the intersection will be noncoherent or contain some holes. Fig. 6 presents the data before registration, and Fig. 7 (top) presents the result of the rigid registration (step 1). Two different views are presented: the volume intersection on the left, and an axial slice fusion on the right. The fusion is an overlay of the bone structures (in green) extracted by thresholding on the CT image (slice 30) and the grey-level MRI information (same slice of the reconstructed image).

Fig. 7 (bottom) shows the results of the elastic registration achieved without using a curvature classification. Finally, Fig. 8 shows the optimal elastic registration achieved using three curvature classes. The presence of holes at the rear right side of the intersection volumes denotes a small misregistration of the

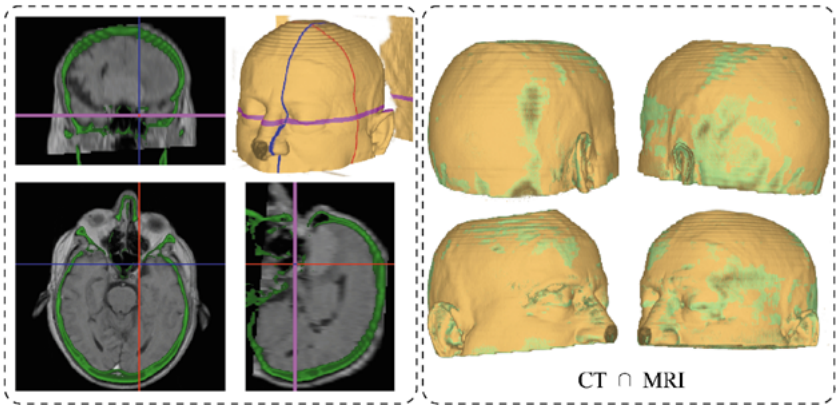

Fig. 10. Elastic registration using three curvature classes.

external surfaces, but a good alignment overall of internal structures (see slices on the left of the drawing).

The overall registration time (less than 5 min on a 233 MHz PC running Linux) is very good in comparison with the computation time obtained with other registration techniques (20 mn on a DEC ALpha Workstation for Thirion's Maxwell daemon [13], and 9 h on a MASPAR).

VII. Validation on a Medical Dataset

In order to assess the robustness of the algorithm and to compare with existing methods, we have applied the registration on the Vanderbilt RREP dataset. After a brief presentation of this dataset, we will explain the employed methodology; furthermore, we will discuss the validation of the elastic registration and then compare the rigid registration with existing methods.

A. The RREP Dataset

RREP stands for Retrospective Registration Evaluation Project (see [14] for more details). This project is designed to compare CT-MR registration techniques used by different groups. The idea is to provide a shared database of CT and MR image that each group can download and use to compute a rigid registration (PET images are also available but are not of concern for our algorithm because the air–skin segmentation of these images is particularly difficult). The latter is then compared with a gold standard registration obtained by a prospective, marker-based technique.

Amongst the images of the database, CT and MR couples from seven different patients are available. For almost each patient there exist six different MR images (T1-weighted, T2-weighted, Positron Density, and three other obtained with a geometrical correction of the three first ones), and one CT-scan. For more information, we invite the reader to visit the Vanderbilt home page.[3] Nevertheless, notice that the VOI covered by each volume of the database varies greatly. Some volumes only correspond to the top of the brain whereas some other maps show from the nose to the top of the head. This feature is particularly important to asses the robustness of the algorithm. Moreover, note that we only used fpur kinds of MR images for our different registrations. We have not used the T2-weighted MR images because the contrast around the air–skin region was not sufficient to allow a good segmentation. Fig. 9 shows direct volume renderings of the seven CT scan volumes and the seven T1-weighted MR volumes of the dataset.

The registrations of the RREP images addressed in this section were achieved without using any local curvature information because the slice spacing was not accurate enough to enable use of a differential geometry.

B. Methodology and Results

1) Elastic Validation: On the one hand, for each of the seven patients of the database, we have computed 11 elastic registrations, and then we have analyzed the stochastic performance of each registration from steps 1 to 3. For each registration, we considered the CT image as the target image because less distortions occur during the acquisition process of these images and thus are more reliable with respect to the true topology of the head. In summary, 4×11 registrations are achieved for each of the seven patients. The main numerical results are reported in Table II. This table first points out the fastness of the algorithm along the different steps (the CPU times are presented in seconds). Then, the stochastic performances denote an improvement between the rigid registrations and the final results which mean that elastic registration was needed. The standard deviation (based on 11 registrations) is rather small on all cases which means that all the registrations are quantitatively comparable. Moreover, we want to point out that we have realized a visual assessment of the median results (the fifth out of the 11 results). These visual assessments made on the same basis as those of Section VI have shown a good overall alignment of the images of each patient (see also Fig. 10 presenting median MR-T1 to CT rigid and elastic registration achieved on the first patient of the database).

2) Rigid Results Comparison: On the other hand, the main purpose of the Vanderbilt project was to compare rigid registration methods. Given that our method is not designed to produce effective rigid registration, we have tried to use the bare results of the first step of our algorithm (genetic rigid registration) to compare with those obtained by other techniques and, at the same time, assess the robustness of this rigid registration step. The overall comparison is presented in Table III. The rather average results shown in this table might be explained in two

[3][Online]. Available: http://www.vuse.vanderbilt.edu/~jayw/

TABLE II
MAIN RESULTS OBTAINED WITH THE APPLICATION OF OUR ALGORITHM ON THE VANDERBILT DATASET. THIS TABLE SHOWS THE MEAN STOCHASTIC PERFORMANCE (AND THE STANDARD DEVIATION) OBTAINED ON 11 DIFFERENT EXECUTIONS OF THE ALGORITHM

#	modality	stochastic performance (step 1) in percent	CPU time (sec) step 1	stochastic performance (step 2) in percent	stochastic performance (step 3) in percent	CPU time (sec) step 2+3
1	MR T1	91.92 (1.00)	15.7	94.89 (0.49)	96.69 (0.30)	99.2
	MR T1R	94.31 (0.51)	18.5	95.77 (0.46)	97.08 (0.21)	96.6
	MR PD	89.54 (1.30)	15.6	91.95 (1.20)	93.97 (1.04)	92.1
	MR PDR	95.45 (0.76)	15.5	95.81 (0.42)	97.23 (0.39)	94.2
2	MR T1	87.47 (2.69)	15.1	88.33 (1.14)	91.52 (1.46)	75.6
	MR T1R	86.39 (2.57)	15.0	85.92 (2.05)	89.81 (1.69)	74.9
	MR PD	87.27 (1.99)	15.1	87.71 (1.39)	90.36 (1.30)	73.9
	MR PDR	88.75 (2.37)	15.2	88.13 (0.78)	90.56 (1.07)	65.8
3	MR T1	77.42 (3.20)	16.0	77.21 (3.05)	81.84 (3.60)	62.8
	MR T1R	76.18 (3.96)	14.2	77.21 (1.58)	81.68 (2.56)	63.8
	MR PD	73.90 (5.11)	16.5	74.92 (1.76)	80.15 (1.92)	63.9
	MR PDR	78.84 (4.15)	16.6	74.26 (2.93)	78.87 (3.09)	61.4
4	MR T1	82.69 (3.01)	14.0	85.33 (1.51)	89.05 (1.55)	61.1
	MR T1R	83.07 (2.12)	14.2	84.88 (1.51)	88.80 (1.74)	63.1
	MR PD	83.78 (3.32)	13.7	84.92 (1.91)	88.10 (1.42)	61.6
	MR PDR	84.49 (1.94)	13.9	82.94 (1.37)	87.42 (1.66)	63.1
5	MR T1	85.40 (2.13)	18.3	88.99 (1.14)	91.75 (1.12)	85.5
	MR T1R	86.78 (2.73)	18.5	89.19 (1.03)	91.84 (1.12)	91.4
	MR PD	83.44 (0.00)	18.0	88.48 (0.96)	91.84 (0.82)	92.6
	MR PDR	85.58 (2.38)	18.7	88.01 (1.60)	90.43 (1.57)	71.9
6	MR T1	84.45 (3.50)	17.6	84.95 (1.90)	88.36 (1.99)	66.2
	MR PD	83.39 (1.91)	18.0	82.60 (2.06)	86.56 (2.19)	67.3
	MR PDR	88.16 (2.58)	17.5	85.13 (1.33)	88.09 (1.41)	65.5
7	MR T1	87.16 (1.34)	16.3	91.58 (1.41)	93.92 (1.07)	87.1
	MR T1R	85.70 (0.91)	16.2	88.68 (1.10)	91.79 (1.65)	84.1
	MR PD	89.00 (2.14)	16.0	91.81 (0.75)	94.00 (0.69)	89.7
	MR PDR	91.36 (1.30)	16.0	93.33 (0.61)	95.32 (0.52)	99.0

TABLE III
COMPARISON OF RIGID REGISTRATION ERRORS OBTAINED BY DIFFERENT GROUPS AND COMPUTED ON THE BASIS OF A GOLD STANDARD TRANSFORMATION

Method	CT to MR DP			CT to MR DPR			CT to MR T1			CT to MR T1R		
author	mean	med.	max.	mean	med.	max.	mean	med.	max.	mean	med.	max.
Barillot	2.38	1.92	6.93	2.28	1.71	5.95	2.13	1.62	6.35	1.91	1.41	5.86
Collignon	2.04	2.09	3.83	0.89	0.81	2.50	1.90	1.53	6.69	1.03	0.72	3.81
Elsen	2.54	2.01	6.55	1.69	1.11	5.32	2.12	1.63	6.05	1.22	0.93	2.61
Harkness	10.86	3.12	49.60	9.99	3.06	45.86	10.46	3.39	51.81	11.68	3.38	48.26
Hemler	3.14	2.37	10.45	1.78	1.66	3.69	2.68	1.37	10.97	1.08	1.00	2.12
Hill	2.00	1.94	4.05	0.89	0.73	2.36	1.36	1.17	2.78	0.87	0.71	2.35
Hsu	1.86	1.67	5.07	1.47	1.46	2.72	2.73	2.51	7.05	2.43	2.38	5.78
Huang	2.16	2.01	5.03	1.13	1.01	2.93	1.81	1.64	4.87	1.66	1.52	3.26
Luo	1.76	1.71	3.56	1.08	0.97	2.66	1.22	1.10	2.99	1.15	1.03	2.81
Maintz	5.41	4.15	18.97	3.78	2.97	10.15	5.68	5.05	12.85	5.05	4.94	14.33
Malandain	10.41	4.00	59.00	10.22	4.04	62.66	10.08	4.32	61.43	11.43	5.42	60.64
Nikou 1	3.06	2.60	5.80	3.00	2.95	5.35	2.72	2.56	6.43	2.45	2.75	4.59
Nikou 2	2.67	2.31	6.18	2.01	1.86	5.07	1.93	1.50	4.36	1.75	1.43	4.54
Noz	6.89	7.80	13.86	5.93	4.61	11.57	4.58	3.32	10.39	4.71	3.40	9.61
Pelizzari	1.96	1.93	4.30	2.16	2.04	4.65	2.79	2.74	7.27	2.34	2.23	5.95
Robb	7.06	5.46	22.19	6.74	5.48	22.14	6.73	5.24	21.77	7.67	5.90	22.24
Rouet	4.36	3.88	15.25	4.27	4.18	9.97	3.39	2.75	12.48	5.60	4.52	20.34
Thevenaz 1	2.01	2.04	4.56	0.85	0.83	1.71	1.69	1.40	4.67	1.04	0.95	4.02
Thevenaz 2	1.94	1.69	5.19	1.15	1.07	4.62	1.72	1.56	5.98	1.07	0.84	4.15
Thevenaz 3	2.15	1.93	4.18	0.99	1.01	1.64	1.61	1.53	4.17	1.05	0.89	3.11

points. First, surface based methods are now known to be less accurate for rigid registration than methods using voxel-similarity measures. This point could be solved using a segmentation of the inner cortical bone in both images and adding these new points into the stochastic point-matching algorithm. Secondly, unlike other rigid registration techniques reported in this table, our test report addresses the MRI to CT registration case, and not CT to MRI (i.e., the more difficult as it needs robust algorithms); this implies that our matching process is to valuate transformation from some outlier CT skin points that have no counterpart in the MRI list of points (this biases the distance measurements). In order to compare our technique with others on the same basis, we have to submit the CT to MRI rigid registration results, which is underway.

Finally, we want to point out that, as noted by West [14], using numerical-only techniques to assert and validate registration techniques is not enough. One should always use a visual assessment and require an expert point of view.

VIII. Conclusions

This paper has presented a robust 3-D elastic registration method which uses genetic algorithms. Its originality comes from the combination of a structural approach (segmentation and classification of structures) with a global and robust search space sampler (two genetic algorithms) and a local optimization process for accuracy.

A registration example using this algorithm has been described. The overall processing time of the elastic registration on a standard 233 MHz PC is less than 5 min which is a very good result in comparison with other elastic registration techniques. Moreover, 3-D multivolume renderings asserting the accuracy of the registration have been presented, and multiple registrations asserting the robustness have also been achieved.

Moreover, the overall registration process has undergone an extensive test work on the RREP (Retrospective Registration Evaluation Project [14]) database of Vanderbilt University (USA). We carried out a series of tests over seven different patients (head) while matching MRI-T1, MRI-T1R MRI-PD, and MRI-PDR images on the corresponding CT dataset. More than 300 tests of the proposed procedure have been achieved and allowed us to assert the robustness of the algorithm.

Future work will be related to a more comprehensive comparison of the rigid and elastic registration results with those obtained by means of other techniques. Investigations regarding the possible segmentation of cortical bone by making use of the rigid registration are also planned and will probably help us to totally alleviate the skin elasticity problem.

Acknowledgment

The authors would like to thank S. Lavallée for providing the MRI and CT-scans of the first head example used in this paper. Other images and their standard transformation(s) were provided as part of the project, "Evaluation of Retrospective Image Registration", National Institutes of Health, 1 R01 NS33926-02, Principal Investigator, J. Michael Fitzpatrick, Vanderbilt University, Nashville, TN.

References

[1] P. A. Van den Elsen, "Medical image matching—A review with classification," *IEEE EMBS Mag.*, pp. 26–39, Mar. 1993.

[2] L. Brown, "A survey of image registration techniques," *ACM Computing Surveys*, vol. 24, no. 4, pp. 325–376, 1992.

[3] S. Lavallée, "Computer integrated surgery and therapy: State of the art," in *Three Dimensionnal Biomedical Imaging*, C. Roux and J. L. Coatrieux, Eds: IOS Press, Nov. 1996.

[4] J. B. A. Maintz and M. A. Viergever, "A survey of medical image registration," *Medical Image Analysis*, vol. 2, no. 1, pp. 1–36, 1998.

[5] G. Wolberg, *Digital Image Warping*. Los Alamos, CA: IEEE Computer Society Press, 1990.

[6] J.-J. Jacq and C. Roux, "Registration of 3-D images by genetic optimization," *Patt. Recognit. Lett.*, vol. 16, no. 8, pp. 823–841, Aug. 1995.

[7] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, MA: Addison-Wesley, 1989.

[8] H. J. Ryser, *Mathématiques Combinatoires*: Monograph. Dunod, 1969.

[9] J. Yen, J. C. Liao, B. Lee, and D. Randolph, "A hybrid approach to modeling metabolic systems using genetic algorithm and simplex method," presented at the SMC'95, 1995.

[10] T. Saïto and J. I. Toriwaki, "New algorithms for euclidean distance transformation of a n-dimensional digitized picture with applications," *Patt. Recognit.*, vol. 27, no. 7, pp. 66–88, July 1994.

[11] J.-J. Jacq and C. Roux, "Automatic detection of articular surfaces in 3-D image through minimal subset random sampling," in *Computer Vision, Virtual Reality and Robotics in Medicine. CVRMed'97*, N. Ayache, Ed, Grenoble, France: Springer-Verlag, Mar. 1997, vol. 1205, pp. 73–82.

[12] ——, "A direct multi-volume rendering method aiming at comparisons of 3-D images and models," *IEEE Trans. Inform. Technol. Biomed.*, vol. 1, pp. 30–43, Mar. 1997.

[13] J.-P. Thirion, "Fast non-rigid matching of 3-D medical images," INRIA Sophia Antipolis, Tech. Rep. 2547, May 1995.

[14] J. West, J. M. Fitzpatrick, M. Wang, B. Dawant, C. R. Maurer, R. M. Kessler, and R. J. Maciunas, "Retrospective intermodality registration techniques: Surface-based versus volume-based," *Computer Vision, Virtual Reality and Robotics in Medicine. CVRMed'97*, vol. 1205, pp. 151–160, Mar. 1997.

Jean-Michel Rouet received the degree from the École Nationale Supérieure des Télécommunications de Bretagne, Brest, France in 1996. He is currently working toward the Ph.D. degree at the Department of Image and Information Processing at the same school and works in collaboration with the Laboratoire de Traitement de l'Information Medicale-LATIM-of Brest (a joint laboratory between the School of Medicine, the University Hospital of Brest and ENSTBr).

His main research interests are medical imaging, genetic algorithms, 3-D renderings, and computer programming.

Jean-José Jacq received the Ph.D. degree in electronics from the University of Bretagne Occidentale, France, in 1986.

He is a researcher in the Image and Information Processing Department of Télécom Bretagne, France,a French Grande École dedicated to Telecommunications. His research interests are in medical imaging, image registration, genetic algorithm, pattern recognition, and geometric modeling. More generally, his research efforts aim at promoting use of voxel-based approaches in the various application fields of image visualization and image analysis. His current research topic addresses object partitioning and object morphometry from manifold representation of their surfaces.

Christian Roux (M'84–SM'95) graduated from the École Normale Supérieure de Cachan, France, in 1978 and received the Ph.D. degree from the Institut National Polytechnique de Grenoble, France, in 1980.

He is a Professor at the École Nationale Supérieure des Télécommunications de Bretagne, Brest. He has been a Visiting Professor at the Medical Image Processing Group, Department of Radiology, University of Pennsylvania, Philadelphia. His research interests are related to advanced information processing in medicine and biology. He has published more than 70 papers in this domain. He is an Associate Editor of the IEEE Transactions on Medical Imaging, and Vice President of the IEEE Engineering in Medicine and Biology Society. He is one of the founders of the IEEE International Summer School on Biomedical Imaging in 1994.