



ACADEMIC
PRESS

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

Computer Vision
and Image
Understanding

Computer Vision and Image Understanding 89 (2003) 272–298

www.elsevier.com/locate/cviu

Iconic feature based nonrigid registration: the PASHA algorithm

Pascal Cachier,^{a,*}¹ Eric Bardinet,^a Didier Dormont,^b
Xavier Pennec,^a and Nicholas Ayache^a

^a INRIA, Epidaure Lab, Sophia Antipolis, France

^b Neuroradiology Department and LENA UPR 640-CNRS, Salpêtrière Hospital, France

Received 17 December 2002; accepted 17 December 2002

Abstract

In this paper, we first propose a new subdivision of the image information axis used for the classification of nonrigid registration algorithms. Namely, we introduce the notion of iconic feature based (IFB) algorithms, which lie between geometrical and standard intensity based algorithms for they use both an intensity similarity measure and a geometrical distance. Then we present a new registration energy for IFB registration that generalizes some of the existing techniques. We compare our algorithm with other registration approaches, and show the advantages of this energy. Besides, we also present a fast technique for the computation of local statistics between images, which turns out to be useful on pairs of images having a complex, nonstationary relationship between their intensities, as well as an hybrid regularization scheme mixing elastic and fluid components. The potential of the algorithm is finally demonstrated on a clinical application, namely deep brain stimulation of a Parkinsonian patient. Registration of pre- and immediate postoperative MR images allow to quantify the range of the deformation due to pneumocephalus over the entire brain, thus yielding to measurement of the deformation around the preoperatively computed stereotactic targets.

© 2003 Elsevier Science (USA). All rights reserved.

Keywords: Nonrigid registration; Iconic features; Local statistics; Mixed elastic/fluid deformation models

* Corresponding author.

E-mail address: pascal.cachier@sophia.inria.fr (P. Cachier).

URL: <http://www.i-u.de>.

¹ Now with the International University in Germany.

1. Introduction

Nonrigid image registration is an important task of image processing. In medical image analysis, it is a fundamental step as soon as we want to quantify the evolution of a patient in a follow-up study, or when comparing two different patients. Consequently, it is a very creative field of research; techniques are numerous and inspired from a wide range of theories or techniques: statistics and information theory, theory of continuum mechanics or viscoelastic fluids, theory of thermodynamics, optical flow, splines, wavelets, block matching, and so on.

To get a better understanding of the different technical choices one faces when designing a registration algorithm, several classifications have been proposed [8,36,38,61]. One major axis shared by all these classifications is the *image feature axis*, i.e., the kind of information that drives the registration process.

Most, if not all, classifications split this axis into two parts: on the one hand, geometric algorithms, which use a geometric distance between segmented features in the images; on the other hand, intensity based algorithms, which use a similarity measure between the image intensities.

However, we found that the group of intensity based algorithms includes two different registration approaches that behave very differently, independently of the similarity measure or the deformation model. We formalize their difference in Section 2 by introducing the notion of iconic (i.e., image intensity related) feature based registration and show how it changes the standard classification of registration algorithms. In Section 3, we propose a new registration energy for iconic feature based registration. We show that this energy generalizes the “demons” algorithm, as well as Feldmar’s “generalized ICP,” and enables a better insight of the behavior of these algorithms. Based on this energy, we develop the PASHA² algorithm. This energy is general and may use different similarity or regularization energies: in this respect, we present Gaussian-weighted local similarity measures in Section 4, which are efficiently computed using an original convolution based technique, as well as an original mixed elastic/fluid regularization in Section 5. We compare PASHA to close algorithms in Section 6. Finally, we present in Section 7 a clinical application of PASHA for brain motion recovery during deep brain stimulation of Parkinsonian patients, and show how our algorithm recovers and propagates deformation in the brain with smooth, realistic displacement fields.

2. A new classification

Despite the large number of techniques used in registration, the main classifications found in the literature all use at least the following two major axes:

- *The deformation model*, used to regularize the registration problem. It expresses the prior knowledge we have on the shape of the transformation.

² PASHA: Pair-And-Smooth, Hybrid energy based Algorithm.

- *The image features* used by the algorithm to guide the deformation model towards (hopefully) the desired transformation.

We first briefly detail in Section 2.1 the different kinds of deformation models; we then present in Section 2.2 our original classification of the image features.

2.1. Deformation models

It is necessary to impose a deformation model to the registration algorithm, otherwise the motion of a point would be estimated independently of the motion of neighboring points, thus yielding a very discontinuous and unlikely displacement field. The deformation model should constrain the estimated transformation as much as possible using our prior knowledge on the transformation. Therefore, the deformation model greatly depends on the goal of registration: not surprisingly, one uses different models when the misalignment of the original images is due to physical deformations, to specific spatial distortions due to the image modality, or to a simple Cartesian coordinate change.

In image registration, we find three major deformation models: parametric, competitive, and incremental models, that we now shortly describe.

2.1.1. Parametric models

The parametric approach constrains the estimate T of the transformation to belong to some low dimensional transformation space \mathcal{T} . Mathematically, if $D(I, J, T)$ is some registration distance between the images I and J to be registered by a transformation T , a parametric approach solves the following minimization problem:

$$\min_{T \in \mathcal{T}} D(I, J, T).$$

Among the most popular choices of transformation space, we find rigid and affine groups, kernels such as thin plate splines [50] or tensor products of B-splines [59], and the first elements of some deformation basis such as the Fourier basis [4], wavelets [68], principal components of a training set of deformations [51], or eigenfunctions of linear operators [17].

2.1.2. Competitive regularization

Competitive models rely on the use of a regularization energy R depending on T . Whereas parametric regularization is a binary prior penalization—no transformation outside the transformation space is allowed, all the transformations inside are equiprobable—the competitive approach penalizes a transformation proportionally to its irregularity measured by the regularization energy [58]. Competitive algorithms puts in competition (hence the name) D and R by solving the following problem:

$$\min_{\forall T} D(I, J, T) + R(T) \tag{1}$$

with eventual additional boundary constraints. A popular choice of regularization energy in image registration is the energy of linear elasticity

$$R(T) = \frac{\lambda}{2} [\text{div}(T)]^2 + \mu \|dT\|^2 - \frac{\mu}{2} \|\text{rot}(T)\|$$

one can then speak of (linear) elastic registration.

2.1.3. Incremental regularization

Incremental (or “fluid”) models also rely on the use of a regularization energy R . This time, however, this energy does not depend on the transformation itself, but on its *evolution*. The simplest case occurs when the regularization energy depends on the difference between the current and the last estimate of the transformation: at iteration n , the estimate T_n of the transformation is found by minimizing

$$\min_{\forall T_n} D(I, J, T_n) + R(T_n - T_{n-1}).$$

These methods thus intrinsically depend on the initial estimate T_0 . The typical example of an incremental approach is the viscoelastic algorithm of Christensen [18]. Recently, incremental models have been generalized by Miller and Younes [42] and Trouvé [62] using a geodesic formulation of image registration, which appears to have nice properties such as topology preservation.

Let us conclude this section by observing that a deformation model may combine several of these approaches. For example, so-called biomechanical models are often a mix of parametric and competitive regularization via finite elements [27]. Techniques based on principal deformations also often use an energy to advantage the first elements of the basis.

2.2. Image features

When classified according to the image features (or “registration basis” in [38]) used to recover the transformation, registration algorithms are usually parted into two groups: geometric feature based (GFB) algorithms, and intensity based (IB) algorithm. However, IB algorithms include two groups that have little in common, except the use of an intensity similarity measure. We therefore propose to split this group into two new groups: on the one hand, standard intensity based (SIB) algorithms; on the other hand, iconic feature based (IFB) algorithms.

2.2.1. Geometric feature based registration

GFB algorithms rely on a segmentation of part or all of the images, done generally before the registration process itself. The segmented geometrical objects correspond generally to anatomical or mathematical invariants, like organ boundaries or differential invariants. Once extracted, these sets are registered by minimizing a geometrical distance between them. This distance often needs to compute explicitly a set of correspondences C between points of the two sets. Then, a transformation T that interpolates or approximates these correspondences C is found using a geometrical distance between C and T as the distance function D to be minimized, and one of the deformation models of Section 2.1.

2.2.2. Standard intensity based registration

Image intensity is another widely used information for registering images. The distance D between two images is then one of the many existing intensity similarity measure S , such as the sum of squared differences (SSD), the correlation coefficient (CC), the correlation ratio (CR), or the mutual information (MI) of the intensities of images I and J . Given a transformation T , the intensity similarity measure is computed between points lying at the same position: because of the density of the intensity information, we implicitly pair a point \mathbf{p} with $T(\mathbf{p})$.

2.2.3. Iconic feature based registration

Some algorithms relying on an intensity similarity measure do not quite work the previous way. Let us take for example block matching based rigid registration [45]: at each iteration, the intensity similarity measure is used to find correspondences between blocks only; the rigid transformation itself is found by minimizing a *geometric distance* (between paired blocks). The choice of this geometric distance gives us another degree of freedom; in that work, the geometric distance is an M-estimator, yielding excellent geometric robustness. Therefore, we need to have a third category of algorithms, which we call IFB, that use explicitly a geometrical distance in addition to the intensity similarity measure.

In this respect, iconic feature based algorithms are really intermediate between the two previous categories. On the one hand, as for GFB registration, we use a geometric distance to fit the transformation to a set of correspondences C between features. On the other hand, as for SIB registration, there is no segmentation of the images, as we use an intensity similarity measure to pair features.

In short, IFB algorithms pair points, lines, or regions, according to their intensity, and fit geometrically a transformation to these pairings. Examples of IFB algorithms are

- *Using point features*, the “demons” algorithm [60] and its extensions [6,14,30], as well as the “generalized ICP” [26].

Table 1
Our classification applied to a panel of registration algorithms

	SIB	GFB	IFB
Rigid/affine	[20,67]	[48,63]	[1,45]
Splines	[40,52]	[50]	
B-splines	[43,59]	[24,56]	
Other param.	[4,51,68]	[5]	
Competitive	[31,41]	[23,35]	[30,60]
		[12,22]	
	[65]		
Incremental	[7,17,37]	[2]	[14,46]
Geodesic	[62]	[16]	

- *Using hyperplane features*, algorithms based on the optical flow constraint [34,32], which is related to the distance of a point \mathbf{p} to the hyperplane $\nabla I(\mathbf{p}) \cdot \mathbf{x} = J(\mathbf{p}) - I(\mathbf{p})$.
- *Using volume or region features*, block matching algorithms [21,38,45].

Finally, as for deformation models, an algorithm may combine different kinds of information. Combinations of geometric features with intensity or iconic features seem particularly promising [12,22,65] (see Table 1).

3. An energy for iconic feature based registration

Looking at the existing IFB algorithms, one remarks that they generally do not minimize a global energy. IFB algorithms like the “demons” or block matching, proceed in alternating two steps. In a first step, they search for a set of correspondences C , using an intensity similarity measure. In a second step, they search for a transformation T that approximates this set of correspondences, using one of the regularization techniques presented in Section 2.1. However, even if each of these steps corresponds to an energy minimization, there is generally no global energy involved. Consequently, the analysis of these algorithms is difficult.

In this section, we propose to minimize the following energy for competitive, iconic point feature based registration of two 2- or 3-D images I and J :

$$E(C, T) = S(I, J, C) + \sigma \|C - T\|^2 + \lambda \sigma R(T). \quad (2)$$

This registration energy E depends on two variables, C and T , that are both vector fields, with one vector per pixel. C is a set of *pairings* between points: for each point of I , it gives a corresponding point in J that attracts this point. T is the estimate of the transformation: it is a smooth vector field (constrained by the regularization energy R) that is attracted by the set of correspondences C . In this energy, S is an intensity similarity energy, used to find the correspondences C , and R is a regularization energy, depending on T . Parameters σ and λ are explained later in Section 3.2.

3.1. Alternating minimization of Eq. (2): the PASHA algorithm

The energy (2) depends on two vector fields, C and T . One could minimize this energy with respect to C and T simultaneously. However, when the regularization energy R is quadratic, the alternating minimization w.r.t. C and T is appealing, because both partial minimizations are quite fast, especially the regularization part which is then solved by convolution.

PASHA is designed on that principle. It minimizes the energy (2) alternatively w.r.t. C and T . We start from $T_0 = \text{Id}$ at iteration 0 (after a possible rigid or affine initialization); then, at iteration n , we iteratively

- Find C_n by minimizing $S(I, J, C_n) + \sigma \|C_n - T_{n-1}\|^2$. This is done in PASHA using gradient descent [14].

- Find T_n by minimizing $\|C_n - T_n\|^2 + \lambda R(T_n)$. This minimization step has a closed-form solution, using a single convolution with C_n [11].

3.2. On the number of parameters for nonrigid registration

Nonrigid registration algorithms generally have only one parameter in their minimization formulation, be it the multiplicative factor of the regularization energy for competitive or incremental approaches, or the dimension of the transformation space for parametric approaches. But the choice of this regularization parameter is influenced by two distinct components: on the one hand, the noise level in the images (and more generally the confidence we have in the intensity relationship model assumed by our intensity similarity measure); on the other hand, the expected smoothness of the recovered transformation.

Our energy (2) contains *two* registration parameters σ and λ . They make it possible to separate to some extent these two different prior knowledge. λ influences the smoothness of the deformation, since it is directly related to the size of the convolution kernel used to compute T . As for σ , it is related to the noise level in the images: it penalizes the pairing of two distant points if the intensity similarity is of the order of the noise. The distinction between these two prior knowledge enables a more relevant choice of the parameters; however, the link between the optimal choice of λ and σ and the actual level of noise or the smoothness of the real transformation is not as simple as we would expect [10].

3.3. Pyramidal approach

Finally, the whole minimization process is embedded into a pyramidal framework, which consists in computing, for each input image, a sequence of smaller and smaller images, and then to register all these images, starting from the smallest one. This classical technique has two advantages. Firstly, the registration algorithm is much less sensitive to initial alignment and can go beyond local minima. Secondly, the cost of the pyramidal approach is relatively small, since the extra images are much smaller than the original one. We chose to use a Gaussian pyramid with a factor of 2 between each dimensions; actually some of the dimensions might be unchanged if the resulting image is more isotropic (a typical situation with CT scanners having a low axial resolution). The subsampling process is carried on while image dimensions are above some threshold (32 by default).

The overall registration process is thus quite fast. PASHA takes typically between 10 and 30 min to register nonrigidly two $256 \times 256 \times 128$ images on a 450 MHz Pentium.

4. Local statistic based similarity measures

In the previous section, we introduced a registration energy (2) without specifying any similarity or regularization energy. Indeed, registration algorithms seldom have

a unique set of tools: most are flexible and propose a panel of energies to better suit any registration problem [49]. This is also the case with PASHA.

In this section, we present similarity measures based on local statistics, as they are implemented in PASHA. We believe that our technique could be interesting for a wide range of applications and algorithms.

4.1. Local vs. global similarity

Standard similarity measures, such as the sum of squared differences, the cross-correlation, or the correlation ratio, are based on stationary hypothesis: they assume that the intensities of I and J are globally related by some function f , i.e., $I = f \circ J$. Even mutual information, which does not rely on a functional relational model, also makes stationary assumptions [49].

By contrast, local measures assume that the link between I and J is valid only locally. We evaluate a function f for some number of regions, whose size should be small enough for stationary assumptions to hold inside these regions, but large enough to estimate f precisely.

There may be at least two good reasons for using local measures. Firstly, stationarity may be an unrealistic assumption; for example, nonuniform biases often corrupt MR images. Secondly, with little effort we gain in robustness towards the assumption of the similarity measure; this is particularly striking when comparing the respective capabilities of the correlation and the local correlation, which can be used for multimodal registration [44,66].

Unfortunately, local measures are generally much less efficient than standard measures. For example, one has to build a local joint histogram everytime the local mutual information is computed around a voxel. However, we have developed a fast computation method based on convolution for local statistics that makes them very efficient. We will briefly present the method; a longer study can be found in [13,9].

4.2. Local statistics

Let $W : \mathbb{R}^d \rightarrow \mathbb{R}$ be a symmetric and normalized window function, and let $W_{\mathbf{p}}$ be its translation around point \mathbf{p} . Let I and J be two images. We define the local mean of I , the local correlation between I and J , the local variance of I , and the local correlation coefficient between I and J —all these statistics being centered around point \mathbf{p} —by

$$\bar{I}_{\mathbf{p}} = \int W_{\mathbf{p}}(\mathbf{x}) \cdot I(\mathbf{x}) \, d\mathbf{x},$$

$$\langle I, J \rangle_{\mathbf{p}} = \int W_{\mathbf{p}}(\mathbf{x}) \cdot (I(\mathbf{x}) - \bar{I}_{\mathbf{p}}) \cdot (J(\mathbf{x}) - \bar{J}_{\mathbf{p}}) \, d\mathbf{x},$$

$$\sigma_{\mathbf{p}}^2(I) = \langle I, I \rangle_{\mathbf{p}},$$

$$CC_{\mathbf{p}}(I, J) = \frac{\langle I, J \rangle_{\mathbf{p}}}{\sigma_{\mathbf{p}}(I)\sigma_{\mathbf{p}}(J)}.$$

We now define an intensity similarity measure by summing the local correlation coefficients for all the points of I

$$LCC(I, J) = \int CC_{\mathbf{p}}(I, J) d\mathbf{p}.$$

(Although we limit here our presentation to the local correlation coefficient criterion, different criteria can be constructed from second order statistics, such as the quadratic error to the local affine model $QELA_{\mathbf{p}} = \sigma_{\mathbf{p}}^2(J)(1 - CC_{\mathbf{p}}^2)$.)

4.3. Computation using convolutions

Fortunately, we obtain these local statistics without having to process one region at a time, as it is done for example in block matching algorithms, because we compute all of them at once using convolutions with our window function W :

$$\bar{I}_{\mathbf{p}} = W * I(\mathbf{p}),$$

$$\sigma_{\mathbf{p}}^2(I) = \bar{I}_{\mathbf{p}}^2 - \bar{I}_{\mathbf{p}}^2 = W * (I^2)(\mathbf{p}) - (W * I)^2(\mathbf{p}),$$

$$\langle I, J \rangle_{\mathbf{p}} = \bar{I} \cdot \bar{J}_{\mathbf{p}} - \bar{I}_{\mathbf{p}} \cdot \bar{J}_{\mathbf{p}} = W * (I \cdot J)(\mathbf{p}) - (W * I)(\mathbf{p}) \cdot (W * J)(\mathbf{p}),$$

where $*$ denotes the convolution operator. Therefore, only a few convolutions are necessary to compute local statistics around every voxel.

As for the derivative of $LCC(I, J \circ T)$ w.r.t. T , which are needed for gradient descent, they can also be computed using convolutions, although formulas get more complicated

$$\begin{aligned} \frac{\partial LCC}{\partial T} = & \left[I \cdot W * \frac{1}{\sigma_{\mathbf{p}}(I) \cdot \sigma_{\mathbf{p}}(J \circ T)} - J \circ T \cdot W * \frac{\langle I, J \circ T \rangle_{\mathbf{p}}}{\sigma_{\mathbf{p}}(I) \cdot \sigma_{\mathbf{p}}^3(J \circ T)} \right. \\ & \left. + W * \left(\frac{\langle I, J \circ T \rangle_{\mathbf{p}} \bar{J} \circ \bar{T}}{\sigma_{\mathbf{p}}(I) \cdot \sigma_{\mathbf{p}}^3(J \circ T)} - \frac{\bar{I}}{\sigma_{\mathbf{p}}(I) \cdot \sigma_{\mathbf{p}}(J \circ T)} \right) \right] \nabla J(T). \end{aligned}$$

By default, a linear interpolation on the images is used to get $J \circ T$ as well as $\nabla J(T)$, although we have also implemented a much slower cubic interpolation.

4.4. On the choice of the window function

The window function W should be a symmetric and normalized kernel. This leaves a large number of possibilities, and indeed, window shapes are numerous in the literature (squares, diamonds, circles, truncated Gaussians, etc.).

PASHA uses Gaussian windows, for essentially two reasons. From a theoretical point of view, the Gaussian is isotropic, therefore our similarity measure is invariant

by rotation. Also, unlike most windows used in block-matching, it is smoothly decreasing toward zero, therefore the profile of the similarity measure is much smoother, and thus easier to minimize.

From a practical point of view, since similarity measures and their derivatives are computed by convolution, we should choose a window yielding an efficient convolution. The Gaussian is particularly efficient because it is separable (it is actually the only scalar isotropic separable kernel), and 1-D Gaussian convolution can furthermore be approximated by recursive filtering, yielding a computation time independent of the size of the Gaussian kernel [25]: one can then use large Gaussian windows to compute the LCC criterion at a low cost. Although this adds yet another parameter to the algorithm, we found out that in practice a standard deviation of 1 cm would work for most brain images. Also, we showed [9] that the results are quite stable w.r.t. the window size as soon as the window is not too small.

5. Regularization

A registration algorithm can also propose different kind of regularization energies—although this is far less common than for similarity measures. A study of the different regularization energies is far beyond our scope here; one can take a look at [11] for a panel of possibilities. The important feature of PASHA is that when R is quadratic and uniform, the regularization is done using convolutions. This enables particularly simple and fast regularization techniques; we illustrate it here by presenting our mixed competitive/incremental deformation model.

5.1. Mixed competitive/incremental regularization

The registration energy (2) includes a regularization energy R depending on T , meaning that it yields a competitive algorithm (Section 2.1.1). Most of existing IFB algorithms are competitive, even if there exists parametric IFB algorithms, using for example block matching [45].

In this section we introduce a mixed competitive/incremental regularization, and show that it can be achieved at a relatively low cost when using quadratic energies. Incremental algorithms are able to recover large displacements, but they often do not preserve image topology in real applications. Mixing these models with a competitive approach drastically improves the topology of the solutions. The energy minimized by PASHA using a mixed model is now, at iteration n ,

$$S(I, J, C_n) + \sigma \|C_n - T_n\|^2 + \sigma \lambda [\omega R(T_n - T_{n-1}) + (1 - \omega)R(T_n)], \quad (3)$$

where the new parameter $\omega \in [0, 1]$ fixes the relative importance of competitive and incremental regularization ($\omega = 0$ being a pure competitive regularization and $\omega = 1$ a pure incremental regularization). The first step of the alternating minimization of this energy is unchanged (Section 3.1); during the second step, we have to minimize

$\|C_n - T_n\|^2 + \lambda[\omega R(T_n - T_{n-1}) + (1 - \omega)R(T_n)]$. When R is quadratic, T_n has a closed-form formula

$$T_n = (1 - \omega)K * C_n + \omega[T_{n-1} + K * (C_n - T_{n-1})], \quad (4)$$

where $*$ denotes convolution and K is a convolution kernel depending of R and λ (see Appendix A). The implementation of this mixed regularization scheme is thus quite fast, and enables to take the best out of both models.




Fig. 1. Two noisy images to be registered.

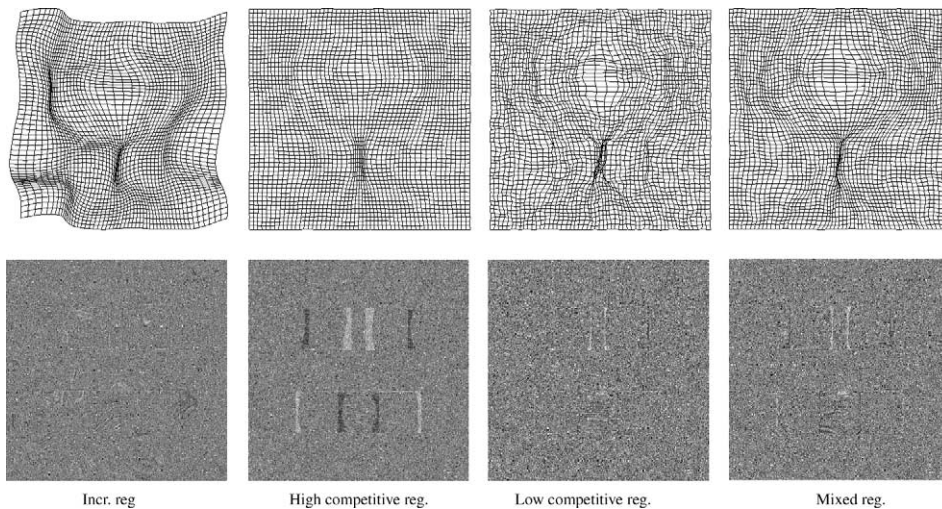


Fig. 2. Registration results on the images of Fig. 1 using PASHA and different regularization techniques. Upper row: recovered deformation field. Lower row: intensity differences between the deformed and the original images. Incremental regularization gives good image correspondences but a poor deformation field. Competitive regularization gives a good deformation field but a poor image correspondence; if we lower the competitive constraint, the deformation field becomes quite irregular. Only the mixed regularization scheme is able to match the images with a regular and plausible deformation field.

Let us illustrate the advantage of this mixed regularization with an example: in Fig. 1, we have two noisy images presenting some large deformations. Incremental regularization is able to register the images, but being sensitive to noise, the recovered displacement field is aberrant (among others, images are symmetric, but the deformation field is not). Competitive regularization is much more robust to noise, but it is more difficult to recover large deformations: with a high regularization strength, the deformation looks plausible but the images are not well registered; with a low regularization strength, the deformation becomes very irregular, while the images are still not matched. With mixed regularization, we combine the advantages of both approaches, and are able to match the images while having a smooth and likely deformation field. All these results are reported in Fig. 2.

6. Comparison with other approaches

6.1. Comparison with the “demons” algorithm and its extensions

The “demons” algorithm [60] and its extensions [6,14,30] are a limit case of PASHA when the parameter σ of Eq. (2) tends to zero (S being then the SSD, and R the energy yielding Gaussian filtering). Indeed, in that case, the closeness constraint between T and C disappears during the first step, and we end up minimizing the SSD alone: this is exactly how these algorithms work. Following Section 3.2, this means that these algorithms assume that the images are noiseless, or more generally, that the intensity correspondence assumed by the similarity measure is strictly enforced.

This closeness constraint between C and T during the first step is yet essential to obtain good results. As a matter of fact, almost all existing IFB algorithms implicitly enforce this constraint, not within their energy like PASHA, but via their minimization scheme: for example, block matching algorithms search iteratively for a position of the window close to the previous one rather than everywhere in the image; the “demons” algorithm uses a bounded correspondence field [14]; our previous algorithm [14] uses a gradient descent starting from $C = T$, therefore C is the local minima of S closest to T . These weak constraints avoid totally aberrant results we would get otherwise, but are still not sufficient: we illustrate this point on a real example.

6.1.1. Illustration on a real 3-D experiment

In the experiment of Fig. 3, we have two $256 \times 256 \times 128$ T1 MR images of the brain of two different subjects that we register both with PASHA and with our own version of the “demons” algorithm [14], using the SSD similarity measure and the same Gaussian kernel for regularization in both cases. Besides the natural noise of MRIs, the intensity conservation is slightly violated, among others, in the background of the images: the background intensity in a rectangle near the skull is slightly higher (around 2 units for an intensity range of 256) in the target than in the source image.

Because of the difference in their closeness constraints between C and T , these algorithms give different results. Using the “demons,” the background is dilated

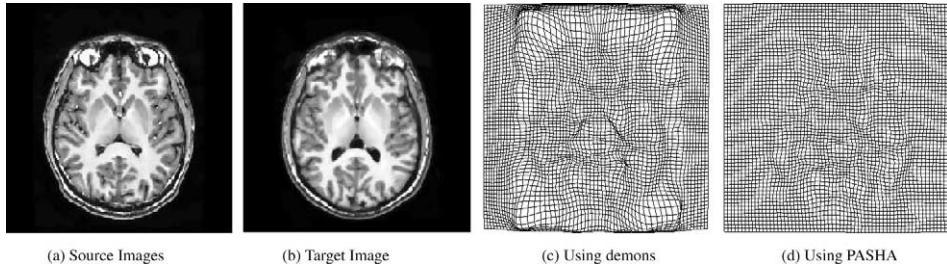


Fig. 3. A real 3-D registration experiment: original images and deformation field found by the “demons” and PASHA.

around the skull. Indeed, if the intensity conservation hypothesis was really true, there would really have been some important deformations: in this case the “demons” would behave correctly. However this is not the case; with PASHA, this unnatural deformation is not recovered when we set a positive value for σ . This helps not to pair remote points that have only a small intensity difference. Similar but visually less illustrative problems also occurs in the white matter, which is far too deformed by the “demons” (one can see some topology problems occurring in the center of the deformation field) and correctly deformed by PASHA (see [10] for further illustrations). Note that because these differences in the deformation occur precisely in low contrast regions, the deformed images are virtually identical. This illustrates the hazard of looking at deformed images only.

6.2. Comparison with Feldmar’s “generalized ICP”

In [26], Feldmar proposed to register two images using their intensity profiles. The intensity profile of an n -D image $I : [0, 1]^n \rightarrow \mathbb{R}$ is the hypersurface $S(I)$ in $(n + 1)$ -D made of all the points $(\mathbf{x}, I(\mathbf{x})) \in [0, 1]^n \times \mathbb{R}, \forall \mathbf{x} \in [0, 1]^n$. To register two images I and J , he first computes the surfaces $S(I)$ and $S(J)$; then, his algorithm roughly consists in alternating the following two steps: finding the closest point on $S(J)$ of each point of $S(I)$, and approximating these correspondences by an affine or local affine transformation. During the first step, the distance between two points located resp. on $S(I)$ and $S(J)$ is given by

$$d^2((\mathbf{x}_1, I(\mathbf{x}_1)), (\mathbf{x}_2, J(\mathbf{x}_2))) = [I(\mathbf{x}_1) - J(\mathbf{x}_2)]^2 + \sigma \|\mathbf{x}_1 - \mathbf{x}_2\|^2$$

σ being a normalization constant and $\|\cdot\|$ the standard Euclidean norm. It appears that this step is strictly equivalent to our first step when S is the SSD similarity measure. Despite the fact that Feldmar’s algorithm is parametric whereas PASHA is competitive, PASHA is conceptually a generalization of the hypersurface matching approach, which is satisfying because Feldmar’s is one of the few IFB algorithms that minimizes an energy. The advantage of PASHA is that its energy (2) is not restricted to the SSD; virtually any similarity measure can be used. With the original intensity profile formulation of Feldmar, the extension to multimodal registration is less clear.

6.3. Comparison with “auxiliary variables”

Auxiliary variables have been introduced for minimizing complex, nonquadratic competitive energies of the form

$$D(v) + R(v). \quad (5)$$

Basically, the idea is to split a coupled, nonlinear multidimensional minimization into two minimizations, one of which has coupled parameters but is quadratic, and the other being nonlinear but having decoupled parameters. Geman [29] first developed this technique for nonlinear regularization, and introduced the auxiliary variable w to transform Eq. (5) into

$$D(v) + \|M \cdot v - w\|^2 + R^\dagger(w), \quad (6)$$

where R^\dagger is the Legendre transform (or conjugate function) of R , and M is a matrix depending on R , so that the minimum of Eq. (6) is obtained for the same value of v as for Eq. (5). Cohen [19] applied Gemans’ framework to the similarity term, minimizing

$$D^\dagger(w) + \|v - w\|^2 + R(v),$$

where D^\dagger is the Legendre transform of D , and proposed to explain some “two-steps” algorithm by this transformation. While auxiliary variables give energies that look like Eq. (2), the fundamental difference is that they are used to simplify the minimization of Eq. (5), and can be applied to geometric as well as to intensity based registration. Our formulation, by contrast, concerns intensity based registration only; and above all, the transformation obtained after minimizing (2) is quite different from the minimizer of $S(I, J, T) + \lambda R(T)$. This is the topic of the next section.

6.4. Comparison with competitive SIB algorithms

To register two images I and J , a competitive SIB algorithm minimizes the energy $S(I, J, T) + \lambda R(T)$ w.r.t. T , where S is an intensity similarity measure, R a regularization energy, and λ a regularization parameter. This seems to be a natural way to combine these two energies, and indeed, this formulation has proven to be successful in the field of data fitting and approximation [64]. Furthermore, it has been justified for registration using Bayesian arguments [28,41].

However, in the context of intensity based registration, this formulation may raise some problems. Typically, competitive SIB algorithms tend to give very important and sharp displacements at the edges of the object, and almost constant displacement in less contrasted regions. We will observe this behavior in the following experiment; but it can also be observed in other papers where the deformation has been presented, e.g. [3,28]. We explain this by the fact that S and R have different physical dimensions: the trade-off is inhomogeneous. The intensity similarity is an intensity distance, i.e., an idea of the *amount of intensity* necessary to fade one image into the other, which is not uniformly proportional to the *amount of motion* that warps

one image into the other. Consequently, the smoothness of transformations found by competitive SIB algorithms is very variable across the image, and the smoothness parameter is therefore tricky to choose.

Decoupling the similarity matching and the regularization as it is done in PASHA seems to help to tackle this problem. The regularization, i.e., the fitting of T to C is now completely geometric and thus homogeneous, as for classical data fitting. The inhomogeneous trade-off is transferred in the estimation of C , which seems to affect less the regularity of the solution. We are thus able to obtain deformations with uniform smoothness, much like those obtained with parametric algorithms (e.g., B-splines), but without the folding problems often occurring with pure parametric algorithms [43,53] thanks to the regularization energy R .

6.4.1. The sinusoidal experiment

We illustrate the previous point by comparing the results given by PASHA and *Asym* [15], which is a competitive SIB algorithm that minimizes

$$\int (I - J \circ T)^2 + \lambda \int \|dT\|^2 \quad (7)$$

by gradient descent, where dT is the Jacobian matrix of T . In this experiment, PASHA uses exactly the same similarity and regularization energies: the energy minimized is

$$\int (I - J \circ C)^2 + \sigma \|C - T\|^2 + \sigma \lambda \int \|dT\|^2. \quad (8)$$

These algorithms thus differ only in the way the trade-off is made. In SIB algorithms, the regularization energy is in direct competition with the similarity measure, and these variations are not uniformly proportional to the deformation because of image contrast differences. Our next experience exacerbates this point by using a couple of MR images having information both on very contrasted, visible areas (all the anatomical borders of the head) and also on very low contrast regions, like inside the background or inside the white matter, where subtle but stable intensity variations make it possible to recover the deformation in these areas. The original MR image




Fig. 4. A synthetic 2-D experiment. From left to right: the original MR image, the deformed image, and the synthetic transformation used to deform the image.

has been deformed using a sinusoidal transformation (Fig. 4), and without adding any noise to it afterwards; therefore, even though the original noise gives texture to the plain areas (which makes it possible to recover the deformation on these parts),

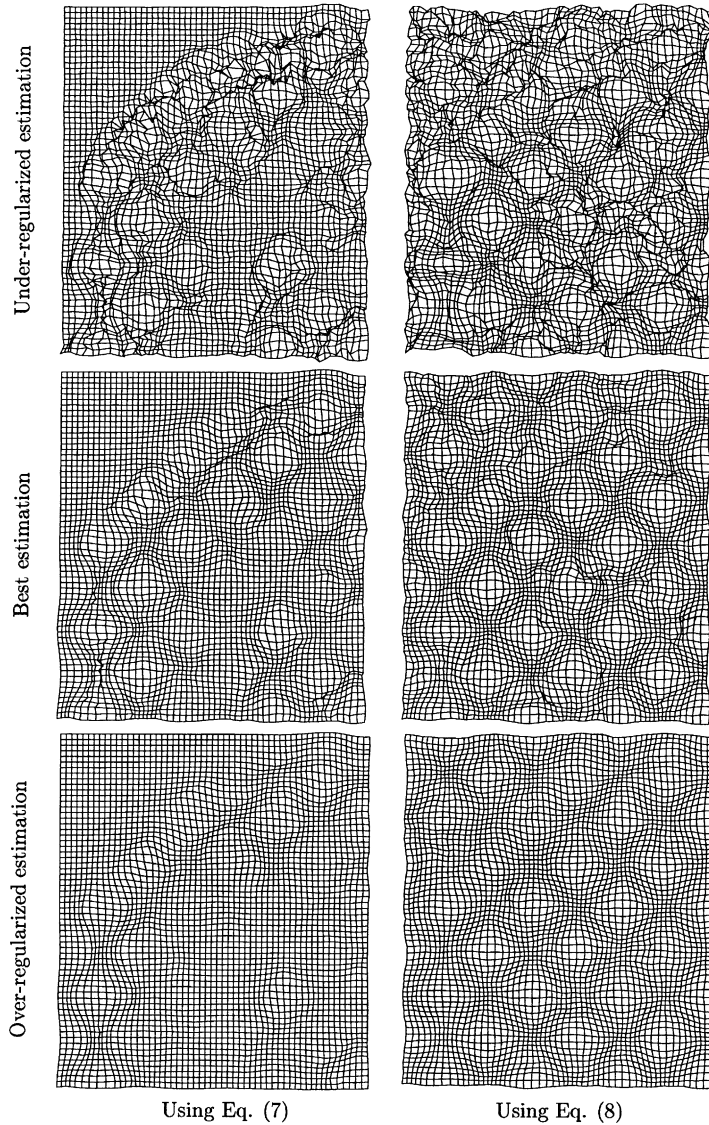


Fig. 5. Details of results obtained with Eq. (7) and with PASHA, Eq. (8) corresponding to the upper left quarter of the images of Fig. 4. Results obtained by Eq. (7) have a very contrasted smoothness depending on the local intensity variation, for any value of λ . Consequently, the deformation is not well recovered. Using Eq. (8), the estimated transformation has a much more uniform smoothness, and is thus closer to the real transformation.

the couple of images can be considered almost noiseless (although a small uncertainty remains due to the linear interpolation used to compute the deformed image). The point of using a sinusoidal transformation is that the same deformation pattern is repeated several times in the image; therefore, we can study how this pattern is recovered depending on the local characteristics of the image. This registration problem looks easy; yet this experiment shows that competitive SIB algorithms have a hard time to recover this simple transformation, contrary to PASHA.

Since the regularization technique of these algorithms are different, we cannot compare their results for the same numerical value of λ . Therefore, we have registered these images with Eq. (7) and with Eq. (8) for a high number of parameters; among all these deformations, we find the best result (i.e., closest to the real transformation) for both algorithms. We report these optimal transformations, as well as sub- and over-constrained transformations, in Fig. 5. As previously discussed, the results obtained with Eq. (7) have a very nonuniform smoothness: it is almost constant and largely underestimated on a large part of the image, except on points lying on contrasted boundaries (like between the fat and the cerebrospinal fluid) where it can be discontinuous. It is impossible to recover the deformation near the skull and in the white matter at the same time, even though there is sufficient information in both areas to do so. At best, the recovered deformation has a mean error of 1.23 pixel. With PASHA, the smoothness of the transformation is much more uniform, and we are thus able to recover accurately the sinusoidal deformation pattern everywhere in the image. The error drops by 33% to 0.83 pixel in mean.

If noise is added to the images, differences are less illustrative, although still significant; PASHA can of course not recover the deformation in low contrast areas anymore, but the uniform smoothness of PASHA yields a better approximation of the transformation in areas where it can actually be estimated, just like for this experiment.

6.5. Comparison with incremental and geodesic algorithms

Incremental (“fluid”) algorithms were introduced in nonrigid registration by Christensen et al. [18] to tackle large deformations. Although these algorithms are indeed able to register very different images, they are also known to give unrealistic deformations on real images, when the assumptions made by the similarity measure are violated.

Recently, however, a new class of techniques based on geodesic formulations has generalized the incremental approach [42,62]. These techniques are able to tackle large deformations while preserving image topology. Although the interest of these new methods is obvious, we want to show that on most of the examples proposed in these papers, a classical incremental regularization could work as well. We present results on the famous “C” experiment in Fig. 6 (which seems to be a standard for such algorithms) obtained with PASHA using a pure incremental Gaussian regularization ($\omega = 1$ in Eq. (3)). For a correct set of registration parameters, the recovered transformation matches perfectly the images while keeping an excellent topology (Fig. 7).

Results obtained on synthetic, simple images are impressive; but the incremental approach often fails on complex or noisy images, as illustrated in Section 5. The geo-




Fig. 6. The “C” experiment. Left: Template image. Middle: target image. Right: deformed image obtained by PASHA.




Fig. 7. Same images as Fig. 6 with a pattern added to the images in order to visualize the transformation. Left: template image. Right: deformed template image.

desic approach may prove its superiority over the classical incremental approach on this kind of images.

7. A clinical application to neurosurgery

Previous sections presented theoretical aspects of PASHA, as well as some comparisons with other approaches. Despite the excellent results obtained previously, it is important to show that the algorithm works for real studies as well. Actually, our algorithm has been used already in several studies, mainly for brain tracking in ultrasound images [47], and for multipatient MRI registration [12]—not to mention all the studies based on the “demons” algorithm (e.g., [23,30]), of which PASHA is a generalization. In this section we propose to expose a novel application of PASHA to preoperative pneumocephalus.

7.1. Introduction

MR image guided brain surgery is an actively developing field. In most of the cases, the techniques are based on the use of volumetric preoperative MR acquisitions.

These techniques implicitly assume that a preoperative acquisition gives a faithful and precise representation of the brain anatomy during the intervention.

A major limit of these techniques is the brain deformation during the surgical intervention, hence leading to anatomical differences, which can be significant, with the preoperative MR images.

To overcome this limit, there has been recent interest in quantifying brain deformation during neurosurgery [33,54,55].

One interesting example of this type of procedures is functional neurosurgery for Parkinson's disease. This intervention is based on the stereotactic introduction of electrodes in a small, deeply located, nucleus of the brain, called the subthalamic nucleus.

This nucleus is targeted on preoperative stereotactic MR acquisitions. During the intervention, which is performed in the operating room, outside the MR unit, an electrophysiological and clinical study is performed with the electrodes to check the preoperatively determined target position.

This exploration is time-consuming and can lead to the development of a pneumocephalus (presence of air in the intracranial cavity) because of CSF leak. This pneumocephalus leads to a brain shift which can yield a significant deformation of the entire brain and thus can cause potential errors in the preoperatively determined position of the stereotactic targets.

Therefore, computing accurately the deformation induced by the pneumocephalus over the entire brain appears to be a key issue, as it will allow to quantify the deformation occurred around the stereotactic targets and check if correction of the preoperative localization would have been needed.

We propose a method to quantify brain deformation from pre- and immediate postoperative MR acquisitions, based on nonrigid registration. Our method is related to the approach developed in [55] for estimating tissue deformation induced by intracranial electrode implantation in patients with epilepsy; however, deformations provided by Studholme et al. [55] are limited to coarse B-splines, while ours are vector fields. Therefore, the method presented below allows a finer analysis of the deformation. It is also substantially faster.

7.2. Material and methods

One patient with bilateral subthalamic lateral implantation was studied. The subthalamic targets were determined one day before the intervention using 3-D stereotactic IR-FSPGR MR acquisition. Then, the patient had bilateral implantation of depth electrodes at the level of the subthalamic nuclei. MR control using the same acquisition as preoperative was performed the day after the implantation. Voxel size of both acquisitions was $0.9375 \times 0.9375 \times 1.3$ mm, and image dimension was $256 \times 256 \times 124$.

On the control acquisition, presence of air collection was clearly observed at the anterior part of the cranial cavity.

The methodology to quantify brain deformation from these pre- and immediate postoperative MR acquisitions consisted in three steps: robust rigid registration of

pre- and postoperative MR acquisitions using [49]; segmentation of the cortex on the registered images; nonrigid registration of the resulting images.

One can observe some changes at the scalp level, which are not caused by the pneumocephalus. Searching for a nonrigid deformation accounting for these changes would be rather complex, error-prone, and not fundamentally useful for our task. Therefore, we segmented the brain before computing the nonrigid residual deformation. This was done automatically using successively thresholding with the CSF value, morphological opening, erosion, maximal connected component extraction, and dilation.

From the segmented and rigidly registered images, a nonrigid registration was performed, using PASHA. We used the Gaussian-weighted local correlation coefficient, presented in Section 4, as the similarity measure between the images, which is particularly well suited for MRIs because of frequent nonuniform bias. We also used the mixed regularization of Section 5, with $\omega = 0.6$, which enabled a total recovery of the deformation, and that, with a very smooth transformation.

Once the deformation field was computed, it was straightforward to identify regions in the brain that underwent deformations higher than a given value, using the norm of the deformation field, hence allowing to estimate the deformation due to pneumocephalus around the preoperatively localized targets.




Fig. 8. Rigid and nonrigid registration of pre- and postoperative MR images of the same patient. Axial (first row) and sagittal (second row) slices show subtraction images after rigid (left column) and nonrigid (right column) registration.

7.3. Results

Nonrigid registration using PASHA was performed on the couple of $256 \times 256 \times 124$ segmented brain MRI (registration is fully 3-D here). It took 8 min on a 450 MHz Pentium III running Linux. Fig. 8 shows axial and sagittal slices of the subtraction images after rigid and nonrigid registration of the preoperative and postoperative volumes, in order to visually assess the quality of the registration. Notice in particular the consistent matching obtained around the prefrontal lobe, where deformation caused by the pneumocephalus was largest.

Once the deformation was computed (the maximal value of the deformation norm was 5.79 mm), it was possible to examine its spatial distribution, and especially to




Fig. 9. Nonrigid registration of pre- and postoperative MR images of the same patient, starting from the rigidly registered images, after brain extraction. Columns correspond to three different axial slices. First row: deformation field; second row: norm of the deformation; third row: postoperative image. On the last two rows, isolines of the deformation norm (1, 2, and 3 mm) are superimposed on the images.

look at regions of brain that suffered deformations higher than a given value. Fig. 9 shows, on the first row, axial slices through the deformation field itself, the displacement of each voxel being represented by an arrow. Then, on the next two rows, isolines of the deformation norm are superimposed on the same axial slices, of the deformation norm, and of the postoperative volume (after brain extraction). The iso-values were set to 1, 2, and 3 mm (3 mm corresponding to the smallest region). As expected, the maximal values of the deformation coincided with the prefrontal lobe where the pneumocephalus was observed. Then, as can be seen, the deformation smoothly decreased, reaching values under 1 mm at the level of the basal ganglia. Also, deformation around the targets (identified by the dark holes located between the putamen and the thalamus on the postoperative images) was significantly lower than 1 mm (around 0.7 mm for the anterior part of the targets, and around 0.4 mm for the posterior part). Note that the localized deformation area (greater than 1 mm) at the level of the target on slice 134 was due to the electrode MR signal.

7.4. Discussion

We have proposed a method for quantifying brain deformation induced by pneumocephalus that occurs during deep brain stimulation on Parkinsonian patients. This method is based on nonrigid registration of pre- and immediate postoperative MR acquisitions.

On the patient under study, results confirmed surgeon intuition, i.e., pneumocephalus does not affect target preoperative localization. Indeed, the deformation was significantly inferior than 1 mm around the targets, knowing that the voxel size of the MR images was around 1 mm^3 .

To confirm the result obtained on this patient, a clinical study is presently conducted on a wide range of Parkinsonian patients.

The methodology presented here could obviously be used in other contexts, for example for dystonia or epilepsy.

8. Conclusion

In this paper, we have first highlighted a fundamental difference that exists between intensity based registration algorithms. On the one hand, standard intensity based (SIB) algorithms use an intensity similarity measure to quantify the quality of the registration. On the other hand, iconic feature based (IFB) algorithms use a geometric distance between homologous geometric features, whose pairing is based on intensities. This last category includes the “demons” algorithm, the “generalized ICP,” block matching and optical flow constraint based methods.

We have then proposed an energy for competitive IFB registration. The alternating minimization of this energy leads to a two-step algorithm we named PASHA, which we compared to some other existing registration techniques. In particular, it appears that PASHA generalizes the “demons” as well as the “generalized ICP” methods, adding new features (more robustness than the “demons” by adding an

image noise assumption, more flexibility than the generalized ICP by using similarity measures other than the SSD).

Concerning intensity similarity, we have proposed an efficient way to compute Gaussian-weighted local statistics, which are very robust towards nonstationary assumptions. As for the regularization energy, a closed form formula exists using convolutions when this energy is quadratic, even for mixed incremental/competitive regularization that we have presented here. We refer to [11] for a deeper discussion on isotropic vectorial filters, including generalized Gaussian filtering with cross-effects between components.

Among the different existing applications of PASHA [12,47], we have chosen to show the performance of our approach on a new clinical application, namely the study of preoperative pneumocephalus during deep brain stimulation of a Parkinsonian patient. The nonrigid registration of pre- and immediate postoperative MR images allowed us to quantify, in a very short time, the deformation induced by the pneumocephalus around the preoperatively determined target position and to assert that this position had not been biased by the pneumocephalus in this case.

PASHA is able to recover a smooth and precise deformation in a fairly small amount of time. Its main drawback, shared by many registration algorithms, is that a number of parameters still have to be chosen manually, especially the regularization strength. Therefore, the registration quality depends ultimately on the experience of the user. Work still has to be done on the extension of classical regularization parameter estimation techniques to the particular case of image registration, as it has already begun to some extent, e.g., in [39,57]. We believe that the difficulty encountered in this topic is related to the discussion of Section 3.2, namely that the regularization does not depend only on the image noise but should also take into account a prior knowledge on the deformation smoothness. Until an effective solution is found, low computation time is a key factor for user interaction, which is one of the strength of PASHA.

Appendix A. On the closed form formula (4)

Let

$$E(T) = \|C_n - T_n\|^2 + \lambda[\omega R(T_n - T_{n-1}) + (1 - \omega)R(T_n)]$$

be the energy to minimize. Let us note \hat{V} the Fourier transform of a vector field V . We suppose in the following that R is a quadratic regularization energy that can be expressed in the Fourier domain as

$$R(T) = \int \|L \cdot \hat{T}\|^2,$$

L being a square matrix function. Let us then rewrite the energy E in the Fourier domain

$$E(T) = \|\hat{C}_n - \hat{T}_n\|^2 + \lambda \int \left[\omega \|L \cdot (\hat{T}_n - \hat{T}_{n-1})\|^2 + (1 - \omega) \|L \cdot \hat{T}_n\|^2 \right].$$

The formal derivative of this energy w.r.t. \hat{T}_n is simply

$$2(\hat{T}_n - \hat{C}_n) + 2\lambda \left[\omega L \cdot (\hat{T}_n - \hat{T}_{n-1}) + (1 - \omega)L\hat{T}_n \right].$$

Setting this derivative to zero, and solving for \hat{T}_n , we get

$$\hat{T}_n = (1 - \omega)\hat{K}\hat{C}_n + \omega \left[\hat{K} \cdot (\hat{C}_n - \hat{T}_{n-1}) + \hat{T}_{n-1} \right],$$

where $\hat{K} = (\text{Id} + \lambda \cdot L)^{-1}$. Getting back to the spatial domain, we obtain Eq. (4).

References

- [1] R.J. Althof, M.G.J.K. Wind, J.T. Dobbins, A rapid and automatic image registration algorithm with subpixel accuracy, *IEEE Trans. Med. Imag.* 16 (3) (1997) 308–316.
- [2] P.R. Andresen, F.L. Bookstein, K. Conradsen, B.K. Ersbøll, J.L. Marsh, S. Kreiborg, Surface-bounded growth modeling applied to human mandibles, *IEEE Trans. Med. Imag.* 19 (11) (2000) 1053–1063.
- [3] J. Ashburner, J.L.R. Andersson, K.J. Friston, Image registration using a symmetric prior—in three dimensions, *Human Brain Mapp.* 9 (4) (2000) 212–225.
- [4] J. Ashburner, K.J. Friston, Nonlinear spatial normalization using basis functions, *Human Brain Mapp.* 7 (1999) 254–266.
- [5] F.L. Bookstein, Principal warps: thin-plate splines and the decomposition of deformations, *IEEE Trans. Pattern Anal. Mach. Intell.* 11 (6) (1989) 567–585.
- [6] I. Bricault, G. Ferretti, P. Cinquin, Registration of real and CT-derived virtual bronchoscopic images to assist transbronchial biopsy, *IEEE Trans. Med. Imag.* 17 (5) (1998) 703–714.
- [7] M. Bro-Nielsen, C. Gramkow, Fast fluid registration of medical images, in: *Proc. VBC'96*, vol. 1131 of LNCS, Hamburg, Germany, Springer, Berlin, 1996, pp. 267–276.
- [8] L.G. Brown, A survey of image registration techniques, *ACM Comput. Surveys* 24 (4) (1992) 276–325, See also p. 12.
- [9] P. Cachier, Recalage non rigide d'images médicales volumiques: contributions aux approches iconiques et géométriques, PhD thesis, École Centrale Paris, January 2002.
- [10] P. Cachier, N. Ayache, Regularization in image non-rigid registration: I. Trade-off between smoothness and similarity, Technical Report RR-4188, INRIA, 2001.
- [11] P. Cachier, N. Ayache, Isotropic energies, filters and splines for vector field regularization, *J. Math. Imag. Vision* (to appear).
- [12] P. Cachier, J.-F. Mangin, X. Pennec, D. Rivière, D. Papadopoulos-Orfanos, J. Régis, N. Ayache, Multisubject non-rigid registration of brain MRI using intensity and geometric features, in: *Proc. MICCAI'01*, vol. 2208 of LNCS, 2001, pp. 734–742.
- [13] P. Cachier, X. Pennec, 3D non-rigid registration by gradient descent on a Gaussian-windowed similarity measure using convolutions, in: *Proc. MMBIA'00*, vol. 1935 of LNCS, 2000, pp. 182–189.
- [14] P. Cachier, X. Pennec, N. Ayache, Fast non-rigid matching by gradient descent: study and improvements of the “Demons” algorithm, Technical Report RR-3706, INRIA, June 1999.
- [15] P. Cachier, D. Rey, Symmetrization of the non-rigid registration problem using inversion-invariant energies: application to multiple sclerosis, in: *Proc. MICCAI 2000*, 2000, pp. 472–481.
- [16] V. Camion, L. Younes, Geodesic interpolating splines, in: *Proc. EMMCVPR'01*, vol. 2134 of LNCS, 2001, pp. 513–527.
- [17] G.E. Christensen, S.C. Joshi, M.I. Miller, Volumetric transformation of brain anatomy, *IEEE Trans. Med. Imag.* 16 (6) (1997) 864–877.
- [18] G.E. Christensen, R.D. Rabitt, M.I. Miller, Deformable templates using large deformation kinematics, *IEEE Trans. Image Process.* 5 (10) (1996) 1435–1447.

- [19] L.D. Cohen, Auxiliary variables and two-step iterative algorithms in computer vision problems, *J. Math. Imag. Vision* 6 (1) (1996) 59–83.
- [20] A. Collignon, F. Maes, D. Delaere, D. Vandermeulen, P. Suetens, G. Marchal, Automated multi-modality image registration based on information theory, in: *Information Processing in Medical Imaging*, Kluwer Academic Publishers, Dordrecht, 1995, pp. 263–274.
- [21] D.L. Collins, A.C. Evans, ANIMAL validation and applications of nonlinear registration based segmentation, *Int. J. Pattern Recog. Artif. Intell.* 11 (8) (1997) 1271–1294.
- [22] D.L. Collins, G. Le Goualher, A.C. Evans, Non-linear cerebral registration with sulcal constraints, in: *Proc. MICCAI'98*, vol. 1496 of LNCS, 1998, pp. 974–984.
- [23] B.M. Dawant, S.L. Hartmann, S. Gadamsetty, Brain atlas deformation in the presence of large space-occupying tumors, in: *Proc. MICCAI'99*, vol. 1679 of LNCS, 1999, pp. 589–596.
- [24] J. Declerck, N. Ayache, E.R. McVeigh, Use of a 4D planispheric transformation for the tracking and the analysis of LV motion with tagged MR images, Technical Report RR-3535, INRIA, October 1998.
- [25] R. Deriche, Recursively implementing the Gaussian and its derivatives, in: *Proc. 2nd Int. Conf. on Image Processing*, 1992, pp. 263–267.
- [26] J. Feldmar, J. Declerck, G. Malandain, N. Ayache, Extension of the ICP algorithm to nonrigid intensitybased registration of 3D volumes, *Comp. Vision Image Understanding* 66 (2) (1997) 193–206.
- [27] M. Ferrant, S.K. Warfield, A. Nabavi, F.A. Jolesz, R. Kikinis, Registration of 3D intraoperative MRI of the brain using a finite element biomechanical model, in: *Proc. MICCAI'00*, vol. 1935 of LNCS, 2000, pp. 19–28.
- [28] J.C. Gee, On matching brain volumes, *Pattern Recog.* 32 (1999) 99–111.
- [29] D. Geman, C. Yang, Nonlinear image recovery with half-quadratic regularization, *IEEE Trans. Image Process.* 4 (7) (1995) 932–946.
- [30] A. Guimond, A. Roche, N. Ayache, J. Meunier, Three-dimensional multimodal brain warping using the demons algorithm and adaptive intensity corrections, *IEEE Trans. Med. Imag.* 20 (1) (2001) 58–69.
- [31] N. Hata, T. Dohi, S. Warfield, W. Wells, R. Kikinis, F.A. Jolesz, Multimodality deformable registration of pre- and intraoperative images for MRI-guided brain surgery, in: *Proc. MICCAI'98*, vol. 1496 in LNCS, 1998, pp. 1067–1074.
- [32] P. Hellier, C. Barillot, E. Mémin, P. Pérez, Hierarchical estimation of a dense deformation field for 3D robust registration, *IEEE Trans. Med. Imag.* 20 (5) (2001) 388–402.
- [33] D.L.G. Hill, C.R. Maurer, A.J. Martin, S. Sabanathan, W.A. Hall, D.J. Hawkes, D. Rueckert, C.L. Truweit, Assessment of intraoperative brain deformation using interventional MR imaging, in: *Proc. MICCAI'99*, vol. 1679 of LNCS, 1999, pp. 910–919.
- [34] B.K.P. Horn, B.G. Schunk, Determining optical flow, *Artif. Intell.* 17 (1981) 181–203.
- [35] S.K. Kyriacou, C. Davatzikos, S.J. Zinreich, R.N. Bryan, Nonlinear elastic registration of brain images with tumor pathology using a biomechanical model, *IEEE Trans. Med. Imag.* 18 (7) (1999) 580–592.
- [36] H. Lester, S.R. Arridge, A survey of hierarchical non-linear medical image registration, *Pattern Recog.* 32 (1999) 129–149.
- [37] H. Lester, S.R. Arridge, K.M. Jansons, L. Lemieux, J.V. Hajnal, A. Oatridge, Non-linear registration with the variable viscosity fluid algorithm, in: *Proc. IPMI'99*, vol. 1613 of LNCS, June/July, 1999, pp. 238–251.
- [38] J.B.A. Maintz, M.A. Viergever, A survey of medical image registration, *Med. Image Anal.* 2 (1) (1998) 1–36.
- [39] C. Meyer, J. Boes, B. Kim, P. Bland, Evaluation of control point selection in automatic, mutual information driven 3D warping, in: *Proc. MICCAI'98*, vol. 1496 of LNCS, October, 1998, pp. 944–951.
- [40] C.R. Meyer, J.L. Boes, B. Kim, P.H. Bland, Probabilistic brain atlas construction: thin plate spline warping via maximization of mutual information, in: *Proc. MICCAI'99*, vol. 1679 of LNCS, September, 1999, pp. 631–637.

- [41] M.I. Miller, G.E. Christensen, Y. Amit, U. Grenander, Mathematical textbook of deformable neuroanatomies, *Proc. Natl. Acad. Sci. USA* 90 (24) (1993) 11944–11948.
- [42] M.I. Miller, L. Younes, Group actions, homeomorphisms, and matching: a general framework, *Int. J. Comput. Vision* 41 (1/2) (2001) 61–84.
- [43] O. Musse, F. Heitz, J.-P. Armspach, Topology preserving deformable image matching using constrained hierarchical parametric models, *IEEE Trans. Image Process.* 10 (7) (2001) 1081–1093.
- [44] S. Ourselin, Recalage d’images médicales par appariement de régions, Application à la construction d’atlas histologiques, PhD thesis, Université de Nice Sophia-Antipolis, Janvier 2002.
- [45] S. Ourselin, A. Roche, G. Subsol, X. Pennec, N. Ayache, Reconstructing a 3D structure from serial histological sections, *Image Vision Comp.* 19 (1–2) (2001) 25–31.
- [46] X. Pennec, P. Cachier, N. Ayache, Understanding the “Demons” algorithm: 3D non-rigid registration by gradient descent, in: *Proc. MICCAI’99*, vol. 1679 of LNCS, September, 1999, pp. 597–605.
- [47] X. Pennec, P. Cachier, N. Ayache, Tracking brain deformations in time-sequences of 3D us images, *Pattern Recognition Letters, Special Issue on Ultrasonic Image Processing and Analysis* 24 (4–5) (2003) 801–813.
- [48] A. Rangarajan, H. Chui, J.S. Duncan, Rigid point feature registration using mutual information, *Med. Image Anal.* 3 (4) (1999) 425–440.
- [49] A. Roche, G. Malandain, N. Ayache, Unifying maximum likelihood approaches in medical image registration, *Int. J. Imag. Systems Technol.* 11 (2000) 71–80.
- [50] K. Rohr, M. Fornefett, H.S. Stiehl, Approximating thin-plate splines for elastic registration: integration of landmark errors and orientation attributes, in: *Proc. IPMI’99*, vol. 1613 of LNCS, June/July, Springer, Berlin, 1999, pp. 252–265.
- [51] D. Rueckert, A.F. Frangi, J.A. Schnabel, Automatic construction of 3D statistical deformation models using non-rigid registration, in: *Proc. MICCAI’01*, vol. 2208 of LNCS, October, 2001, pp. 77–84.
- [52] D. Rueckert, C. Hayes, C. Studholme, P. Summers, M. Leach, D.J. Hawkes, Non-rigid registration of breast MR images using mutual information, in: *Proc. MICCAI’98*, October, 1998, pp. 1144–1152.
- [53] J.A. Schnabel, C. Tanner, A.D. Castellano Smith, M.O. Leach, C. Hayes, A. Degenhard, R. Hose, D.L.G. Hill, D.J. Hawkes, Validation of non-rigid registration using finite element methods, in: *Proc. IPMI’01*, vol. 2082 of LNCS, Davis, USA, 2001, pp. 344–357.
- [54] O.M. Skrinjar, J.S. Duncan, Real time 3D brain shift compensation, in: *Proc. of IPMI’99*, vol. 1613 of LNCS, June, 1999, pp. 42–55.
- [55] C. Studholme, E. Novotny, I.G. Zubal, J.S. Duncan, Estimating tissue deformation between functional images induced by intracranial electrode implantation using anatomical MRI, *NeuroImage* 13 (2001) 561–576.
- [56] G. Subsol, J.Ph. Thirion, N. Ayache, A general scheme for automatically building 3D morphometric anatomical atlases: application to a skull atlas, *Med. Image Anal.* 2 (1) (1998) 37–60.
- [57] D. Suter, F. Chen, Left ventricular motion reconstruction based on elastic vector splines, *IEEE Trans. Med. Imag.* 19 (4) (2000) 295–305.
- [58] R. Szeliski, Bayesian modeling of uncertainty in low-level vision, *Int. J. Comput. Vision* 5 (3) (1990) 271–301.
- [59] R. Szeliski, J. Coughlan, Spline-based image registration, *Int. J. Comput. Vision* 22 (3) (1997) 199–218.
- [60] J.-P. Thirion, Image matching as a diffusion process: an analogy with Maxwell’s demons, *Med. Image Anal.* 2 (3) (1998) 243–260.
- [61] A.W. Toga, P. Thompson, The role of image registration in brain mapping, *Image Vision Comput.* 19 (2001) 3–24.
- [62] A. Trouvé, Diffeomorphisms groups and pattern matching in image analysis, *Int. J. Comput. Vision* 28 (3) (1998) 213–221.
- [63] M. Van Herk, K.G.A. Gilhuijs, J. de Munck, A. Touw, Effect of image artifacts, organ motion, and poor segmentation on the reliability and accuracy of 3D chamfer matching, *Comput. Aided Surg.* 2 (1997) 346–355.

- [64] G. Wahba, Spline Models for Observational Data, vol. 59 of CRMS-NSE Regional Conf. Series in Applied Mathematics, Society for Industrial and Applied Mathematics, 1990.
- [65] Y. Wang, L.H. Staib, Elastic model based non-rigid registration incorporating statistical shape information, in: Proc. MICCAI'98, October, 1998, pp. 1162–1173.
- [66] J. Weese, P. Rösch, T. Netsch, T. Blaffert, M. Quist, Gray-value based registration of CT and MR images by maximization of local correlation, in: Proc. MICCAI'99, vol. 1679 of LNCS, September, 1999, pp. 656–663.
- [67] W.M. Wells, P. Viola, H. Atsumi, S. Nakajima, R. Kikinis, Multi-modal volume registration by maximization of mutual information, *Med. Image Anal.* 1 (1) (1996) 35–51.
- [68] Y.-T. Wu, T. Kanade, C.-C. Li, J. Cohn, Image registration using wavelet-based motion model, *Int. J. Comput. Vision* 38 (2) (2000) 129–152.