Regularization in Image Non-Rigid Registration: I. Trade-off between Smoothness and Intensity Similarity

Pascal Cachier, Nicholas Ayache

INRIA

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Regularization in Image Non-Rigid Registration: I. Trade-off between Smoothness and Intensity Similarity

Pascal Cachier, Nicholas Ayache

N° 4188

Mai 2001

THÈME 3

ISSN 0249-6399 ISRN INRIA/RR--4188--FR+ENG

INRIA
SOPHIA ANTIPOLIS

Regularization in Image Non-Rigid Registration: I. Trade-off between Smoothness and Intensity Similarity

Pascal Cachier, Nicholas Ayache

Thème 3 — Interaction homme-machine,
images, données, connaissances
Projet Epidaure

Rapport de recherche n° 4188 — Mai 2001 — 32 pages

Abstract: In this report, we first propose a new classification of non-rigid registration algorithms into three main categories: in one hand, the geometric algorithms, and in the other hand, intensity based methods that we split here into standard intensity-based (SIB) and pair-and-smooth (P&S) algorithms.

We then focus on the subset of SIB and P&S algorithms that are *competitive*, i.e. that use a regularization energy which is minimized together with the intensity similarity energy. In SIB algorithms, these two energies are combined in a weighted sum, and thus the trade-off between them is direct. P&S algorithms alternates their respective minimization, leading to the characteristic two steps: pairing of points, and smoothing.

We theoretically compare the behavior of SIB and P&S algorithms, and more precisely, we explain why in practice the smoothness of the transforms estimated by SIB algorithms is non-uniform, thus difficult to control, while P&S algorithms estimate a motion that is more uniformly smooth. We give an example illustrating this behavior.

Very few P&S algorithms minimize a global energy. We therefore propose a new image registration energy whose minimization leads to a P&S algorithm. This energy is general, and can use any existing similarity or regularization energy. Its behavior is also compared to the previous SIB and P&S algorithms. This new energy allows uniformly smooth solutions, as for our previous P&S algorithm, while preventing registration of non-informative, noisy areas, as for SIB algorithms.

Key-words: Non-rigid registration, regularization, two-step algorithms, auxiliary variables.

Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles, B.P. 93, 06902 Sophia Antipolis Cedex (France)
Téléphone : 04 92 38 77 77 - International : +33 4 92 38 77 77 — Fax : 04 92 38 77 65 - International : +33 4 92 38 77 65

Méthodes de régularisation en recalage non rigide d'images: I. Compromis entre similarité et régularisation

Résumé : Dans ce rapport, nous proposons de classifier les algorithmes de recalage d'image en trois catégories: d'une part, les algorithmes géométriques, et d'autre part, les algorithmes iconiques, dont nous distinguons ici les algorithmes iconiques standards (*standard intensity based*, SIB) des algorithmes itérant appariement en lissage (*pair-and-smooth*, P&S).

Nous nous penchons ensuite plus précisément sur les algorithmes SIB et P&S *compétitifs*, c'est-à-dire qui utilisent une énergie de régularisation en plus d'une énergie de similarité des intensités. Les algorithmes SIB minimisent une somme ponderée de ces deux énergies : le compromis est donc direct. Les algorithmes P&S préfèrent alterner la minimisation de ces deux énergies, et procèdent donc en deux étapes caractéristiques : l'appariement des point d'une part, et le lissage de ces appariements d'autre part.

Nous discutons d'abord, d'un point de vue théorique, les comportements respectifs de ces deux types d'algorithmes. Notamment, nous tentons d'expliquer pourquoi, en pratique, la régularité des solutions fournies par les algorithmes SIB est moins uniforme, donc moins facile à choisir, que celle des solutions données par les algorithmes P&S. Nous donnons un exemple illustrant ce point.

Partant du fait que peu d'algorithmes P&S minimisent une énergie globale, nous proposons une nouvelle énergie de recalage, dont la minimisation mène à un algorithme P&S. Cette énergie reste générale et peut utiliser n'importe quelle énergie de similarité ou de régularisation déjà existante. Son comportement est également comparé aux algorithmes iconique et séquentiel précédents. Ce nouvel algorithme fourni des solutions d'une régularité uniforme, comme avec notre algorithme P&S précédent, tout en ne cherchant pas à recaler les zones bruitées non informative, comme pour les algorithmes SIB.

Mots-clés : Recalage non rigide, théorie de la régularisation, algorithmes en deux etapes, variables auxiliaires.

1 A New Classification of Non-Rigid Registration Algorithms

Intensity-based non-rigid registration of medical images has been a field of study for 20 years now since the work of (Broit, 1981), although researchers have really focused on the subject during these last 10 years since the reference paper of (Bajcsy and Kovačič, 1989).

Today, techniques of non-rigid registration are very numerous — see (Maintz and Viergever, 1998; Lester and Arridge, 1999) for recent reviews. Even if the debate is still open on a number of theoretical points (symmetry (Christensen, 1999; Ashburner et al., 2000; Cachier and Rey, 2000), registration of tensor images (Alexander and Gee, 2000; Ruiz-Alzola et al., 2000), automatic choice of the regularization strength (Chen and Suter, 1997; Meyer et al., 1998), and also less specifically to non-rigid registration, registration of three or more images (Boes and Meyer, 1999), influence of the interpolation (Pluim et al., 2000), normalization of similarity measures (Studholme et al., 1999; Holden et al., 2000), etc.), this large number of algorithms is partly due to the fact that, unlike in the rigid case, it is difficult to assess the quality of non-rigid registration, and therefore to compare the algorithms: for real experiments, there is no obvious gold standard that gives the real, dense displacement field, even if phantoms with precisely located landmarks begin to appear (King et al., 2000). On the other hand, the simulation of realistic physical deformations is very complex. Furthermore, in the case of multipatient registration, anatomies are not deduced from each other by a real motion: The result of registration is a set of pairings for which there is no ground truth to be compared to, and therefore the quality of these correspondences is very subjective.

Despite the number and the variety of the approaches of non-rigid registration algorithms, we propose to classify most of them into the following three categories:

1. **Geometric matching.** Sparse features are first extracted from the images, e.g. points selected manually, or automatically extracted crest lines or edge surfaces. Then, these features are matched together with a generally smooth non-rigid transformation. The problem of finding a closed-form formula for a transformation that interpolates or approximates the set of pairings is non trivial, even with point features and rigid transformations (Arun et al., 1987; Pennec et al., 1998), and has been the subject of many articles (Duchon, 1977; Bookstein, 1994; Davis et al., 1997; Fornefett et al., 1999; Gabrani and Tretiak, 1999; Peckar et al., 1999; Rangarajan et al., 1999; Rohr et al., 1999; Chui and Rangarajan, 2000; Suter and Chen, 2000). Alternative approaches include finite elements resolution (Kyriacou and Davatzikos, 1998; Dawant et al., 1999; Ferrant et al., 2000), multidimensional B-spline fitting (Subsol et al., 1998), diffusion of the displacement of landmarks to neighboring voxels (Andresen et al., 2000), and other extrapolation techniques (Burr, 1981; Thompson and Toga, 1996).

2. **Standard intensity-based (SIB) registration.** In these algorithms, the images are registered together using an intensity similarity measure, with a non-rigid transformation T. This transformation has two antagonist goals, namely, to be smooth and to maximize the similarity measure. The main characteristic of this class, as opposed to the following class of Pair-and-Smooth algorithms, is that the same transformation is implied both in the intensity similarity measure and in the smoothness constraint. We can divide further this category into three sub-classes:

- **Parametric.** The transformation T is constrained to minimize the intensity similarity measure in a low-dimensional space of non-rigid transformation, such as thin plate splines (Meyer et al., 1999), multidimensional B-splines (Szeliski and Coughlan, 1997; Vemuri et al., 1998; Musse et al., 1999) or wavelets (Wu et al., 2000).
- **Competitive.** The transformation T has no *a priori* shape, but minimizes a weighted sum of an intensity similarity energy and a regularization energy (Amit, 1994; Schormann et al., 1996; Benayoun and Ayache, 1998; Hata et al., 1998; Christensen, 1999; Fischer and Modersitzki, 1999; Alvarez et al., 2000; Cachier and Rey, 2000). This approach is sometimes justified by Bayesian arguments. (Miller et al., 1993; Gee, 1999).
- **Fluid.** The transformation T has no *a priori* shape, but its evolution towards the minimum of the similarity energy is constrained to be smooth (Bro-Nielsen and Gramkow, 1996; Christensen et al., 1997; Lester et al., 1999). A "fluid distance" between two transformations can actually be rigorously defined as a geodesic path: one can thus design a fluid algorithm that minimizes a weighted sum of a similarity energy and a fluid distance (Trouvé, 1998; Miller and Younes, 1999).

These three types of regularization are not exclusive, and some algorithms combine two of these techniques (Ashburner and Friston, 1999; Rueckert et al., 1999).

3. **Pair-and-Smooth (P&S) registration.** This category is intermediate between the two previous categories. Pair-and-Smooth algorithms are also intensity-based, as they use an intensity similarity energy to quantify the quality of the correspondence between the images. However, they proceed in two steps, which may alternate or not: first, they use an intensity similarity measure to find a relatively dense set of homologous points or features between both images. These correspondences are then approximated by a smooth non-rigid transformation. Their main characteristic, as opposed to SIB registration, is that they use two transformations (often implicitly): a first transformation C used to compute the similarity measure, and a second transformation T on which the smoothness constraint is hold. In this category, we find block matching algorithms, both for rigid (Althof et al., 1997; Ourselin et al., 2001) and non-rigid registration (Collins and Evans, 1997; Strintzis and Kokkinidis, 1997; Gaens et al., 2000; Maintz et al., 1998; Lau et al., 1999), registration based on the optical flow constraint (Horn and Schunk, 1981; Bricault et al., 1998; Thirion, 1998; Hellier et al., 1999; Guimond et al., 2001), ICP-like techniques (Feldmar et al., 1997; Amini et al., 1999), and other pair-and-smooth techniques (Guan et al., 1998; Hayton et al., 1999; Cachier and Pennec, 2000; Pennec et al., 2001).

Today, the use of intensity information seems to be more widespread than the use of geometric features for non-rigid registration. This fact is due to the success of intensity similarity measures in the rigid case (West et al., 1999), especially for multimodal registration, and also to the difficulty to automatically extract landmarks that remain sufficiently stable under non-rigid motion and intensity variation. Note that promising algorithms mixing both intensity and feature information are beginning to appear (Collins et al., 1998; Hellier and Barillot, 2000; Cachier et al., 2001); this point is be discussed in the context of P&S registration in section 3.7 and (Cachier, 2001).

In the following, we first compare competitive SIB and P&S methods, and show why P&S methods may have a better behavior in the context of non-rigid registration (section 2). More

specifically, we point out the fact that the smoothness of the estimated motion is much more uniform with P&S algorithms than with competitive SIB algorithms. We show this fact with a synthetic example. Then, we propose a new energy formulation for P&S registration that explains and generalizes one of our previous algorithms (Pennec et al., 2001) (section 3). We quantify the errors of all these algorithms on a set of Gaussian random transformation, and we also present results with a real experiment (section 4), showing the improvements brought by our new formulation.

2 Comparison of Competitive SIB and P&S methods

2.1 Presentation of the problem and theoretical comparison

In this section, we focus on competitive SIB and pair-and-smooth non-rigid registration algorithms. These algorithms use two energies to drive the registration of two images I and J: a similarity energy, which we note $E_{sim}(I, J, T)$, [1] and a regularization energy, which we note $E_{reg}(T)$.

Competitive SIB algorithms, as defined in section 1, balance directly the values of E_{sim} and E_{reg} by minimizing the energy

$$E(T) = E_{sim}(I, J, T) + \lambda E_{reg}(T) \tag{1}$$

λ being the regularization parameter[2]. 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 (see for example the work of (Wahba, 1990)).

However, this formulation may raise some problems in the context of registration. Unlike data fitting, the trade-off here is made between two energies that do not have the same physical dimension: E_{reg} is a geometric measure, and E_{sim} is an intensity similarity measure. The intensity similarity energy gives an "intensity distance", i.e. an idea of the *amount of intensity* necessary to change one image into the other, which is not uniformly proportional to the *amount of motion* that warps one image into the other.

We illustrate this idea in figure 1, where a 1-D image $I(x)$ is registered with a translated version of itself: $J(x) = I(x - t)$. The amount of motion necessary to register one image to the other is constant and equal to $|t|$. This misregistration is not uniformly reported on the similarity measure, and leads to small or large values depending on the local variation of the image.

This may explains why a direct trade-off between transformation smoothness and intensity similarity makes the smoothness of the transformation depend locally on the images, and more precisely on their local contrast: the higher the local variation in the difference of intensity between the images, the less smooth the transformation. In particular, the transformation is oversmoothed in low-contrast regions compared to edge regions.

The second class of algorithm is the pair-and-smooth methods, which alternate between finding correspondences C between points using the intensity similarity energy $E_{sim}(I, J, C)$

[1] In the context of non-rigid registration, the similarity energy cannot always be written $E_{sim}(I, J \circ T)$ (Cachier and Rey, 2000).

[2] Although this formulation is convenient, the regularization strength λ is sometimes preferably not chosen as a multiplicative factor but included inside the regularization energy: $E_{reg}(T, \lambda)$. This point is discussed in (Cachier, 2001).

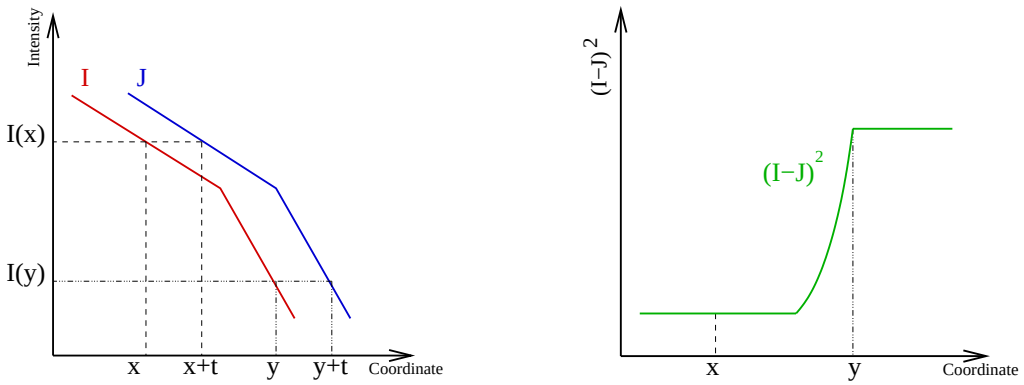

Figure 1: *A uniform error of misalignment between images I and J (left) may lead to very different intensity similarity (right): x and y have been translated by the same amount t, but $(J(x) - I(x))^2 \ll (J(y) - I(y))^2$.*

and sometimes additional energies, and finding a transformation T that approximates these correspondences, e.g. by minimizing $\int ||C - T||^2 + E_{reg}(T)$ w.r.t. T. Here, the trade off between fitting and smoothness occurs during the second step, balancing the distance between corresponding points and the regularization energy of the transformation. This is the classical trade-off of data fitting between geometrical quantities. We expect the estimated transformation to be uniformly smooth, because the distribution of the pairings is more likely to be uniform than intensity changes – and furthermore this distribution can be controlled, for example by forcing pairings to have a length below some threshold, as it is frequently done for example in block matching algorithms, or in the "demons" algorithm (see (Cachier et al., 1999) for a study of this last algorithm).

2.2 Synthetic Sinusoidal Deformation

We illustrate the previous discussion with an example. We compare the results of two non-rigid registration algorithms, respectively competitive SIB and P&S, on a 2D synthetic experiment where we know the ground-truth transformation. These two algorithms use the same similarity and registration energies; the only difference is the way the trade-off is done between these energies.

The intensity-based algorithm is the `Asym` algorithm of (Cachier and Rey, 2000). To register two images I and J with a non-rigid transformation T, it minimizes the following energy using a gradient descent:

$$E(T) = \int (I - J \circ T)^2 + \lambda \int ||dT||^2 \tag{2}$$

The pair-and-smooth algorithm is `MAMAN`[3] (Cachier et al., 1999), an improved version of the "demons" algorithm (Thirion, 1998), with a modification in the regularization so that is uses the same regularization energy as `Asym`. Its main steps are briefly described in table 1. Here, the similarity measure used in step 2 is the SSD, and the smoothing used in step 3 is done by a convolution that minimizes

$$E(T) = \int ||T - C||^2 + \lambda \int ||dT||^2$$

[3] `MAMAN`: Matching Algorithm for Medical Acquisitions in Neurology

1. Set $n := 0$ and $T_0 := \mathrm{Id}$
2. Find T_{n+1} by minimizing $E_{sim}(I, J, T_{n+1})$ by gradient descent
3. Smooth T_{n+1}
4. Set $n := n + 1$
4. Go to step 2 until convergence

Table 1: Main steps of the MAMAN non-rigid matching algorithm.

In this experiment, we took the image 2(a) and deformed it into 2(b) with the transformation depicted in 2(c). The goal of the registration is to recover a transformation as close as possible to the original one. The advantage of using the sinusoidal transformation 2(c) is that the same pattern of deformation is repeated in the image. Therefore, we can see how this pattern is recovered depending on the local characteristics of the images.

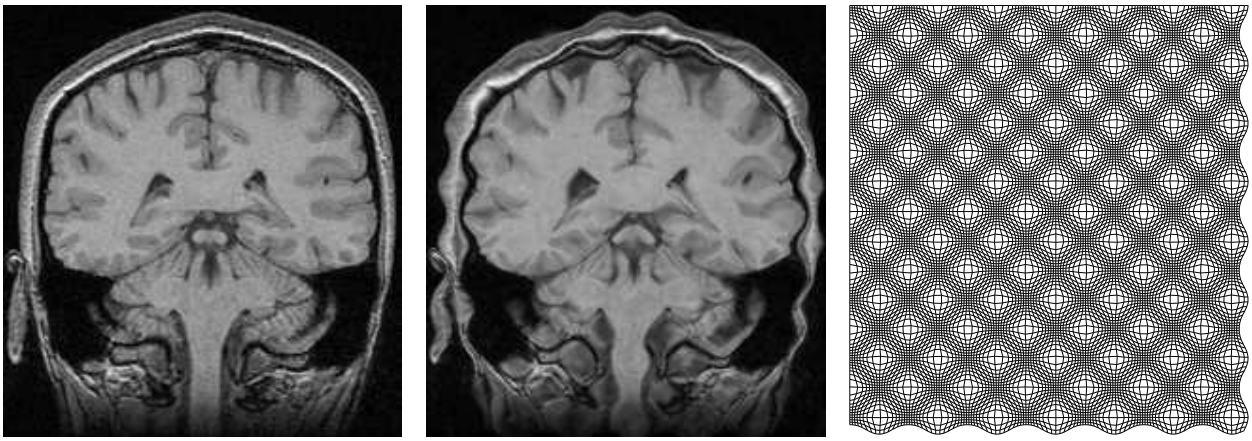

(a) Original (b) Deformed (c) Transformation

Figure 2: *A synthetic registration problem: the image on the left has been deformed into the image in the middle with the transformation depicted on the right.*

Non-rigid registration is always dependent on a regularization factor λ. As we highlighted in (Cachier et al., 1999), we cannot compare the results of two algorithms for a single regularization factor, especially here: since the regularization is not handled the same way in both algorithms, the same value of λ leads to completely different regularization strengths depending on the algorithm. The comparison is more complex than in the rigid case because we have to compare a whole set of results for a wide range of regularization factors. In this example, we have estimated that to cover approximately the same wide regularization range, λ should range from 1 to 5 for MAMAN, and 0.3 to 3000 for Asym. We have registered images 2(a) and 2(b) with Asym (figure 5 and 3) and MAMAN (figure 6 and 4) with these different regularization factors. We present both the estimated transformation (figure 3 and 4) and the deformed image (figure 5 and 6).

From these sets of results, we can draw several remarks:

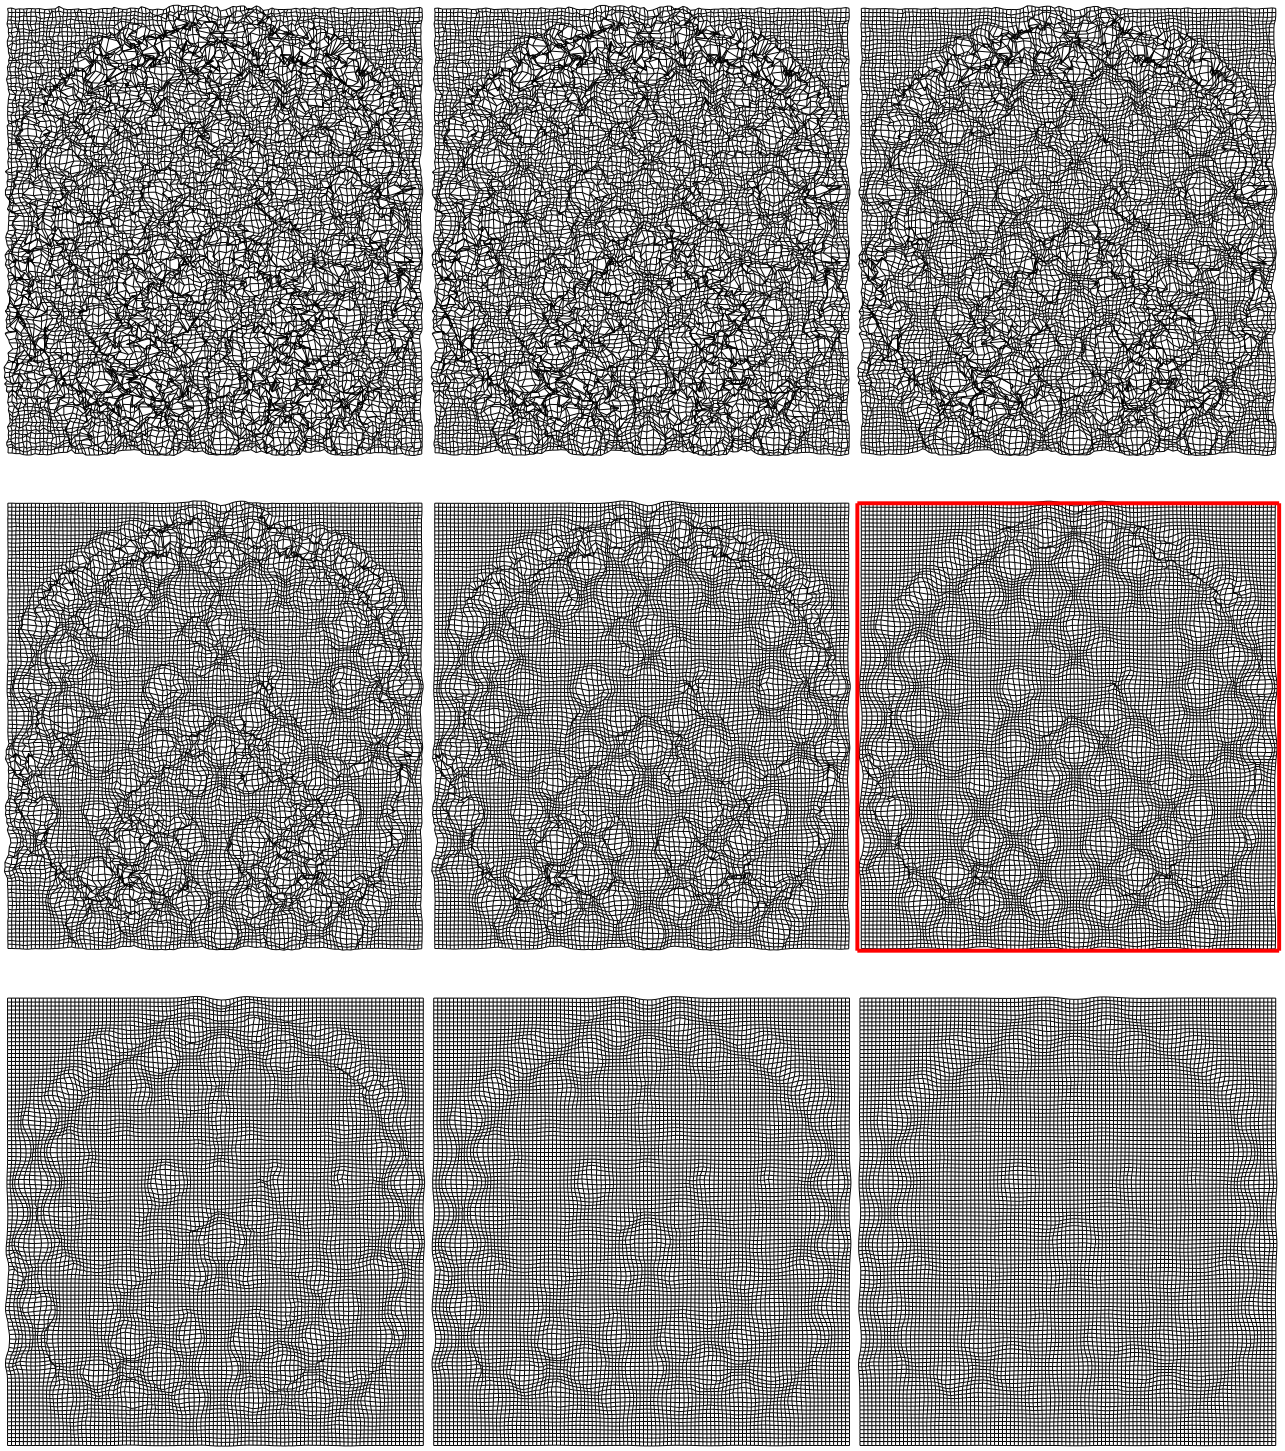

Figure 3: *Results of registration using the variational formulation (2) for regularization strengths ranging from 0.3 to 3000: as it increases, the transformation becomes quickly smooth and underestimated in plain regions, while remaining irregular near the borders of the image. The solution bordered in red is the closest to the ground truth 2(c), with an average error of 1.27 pixel.*

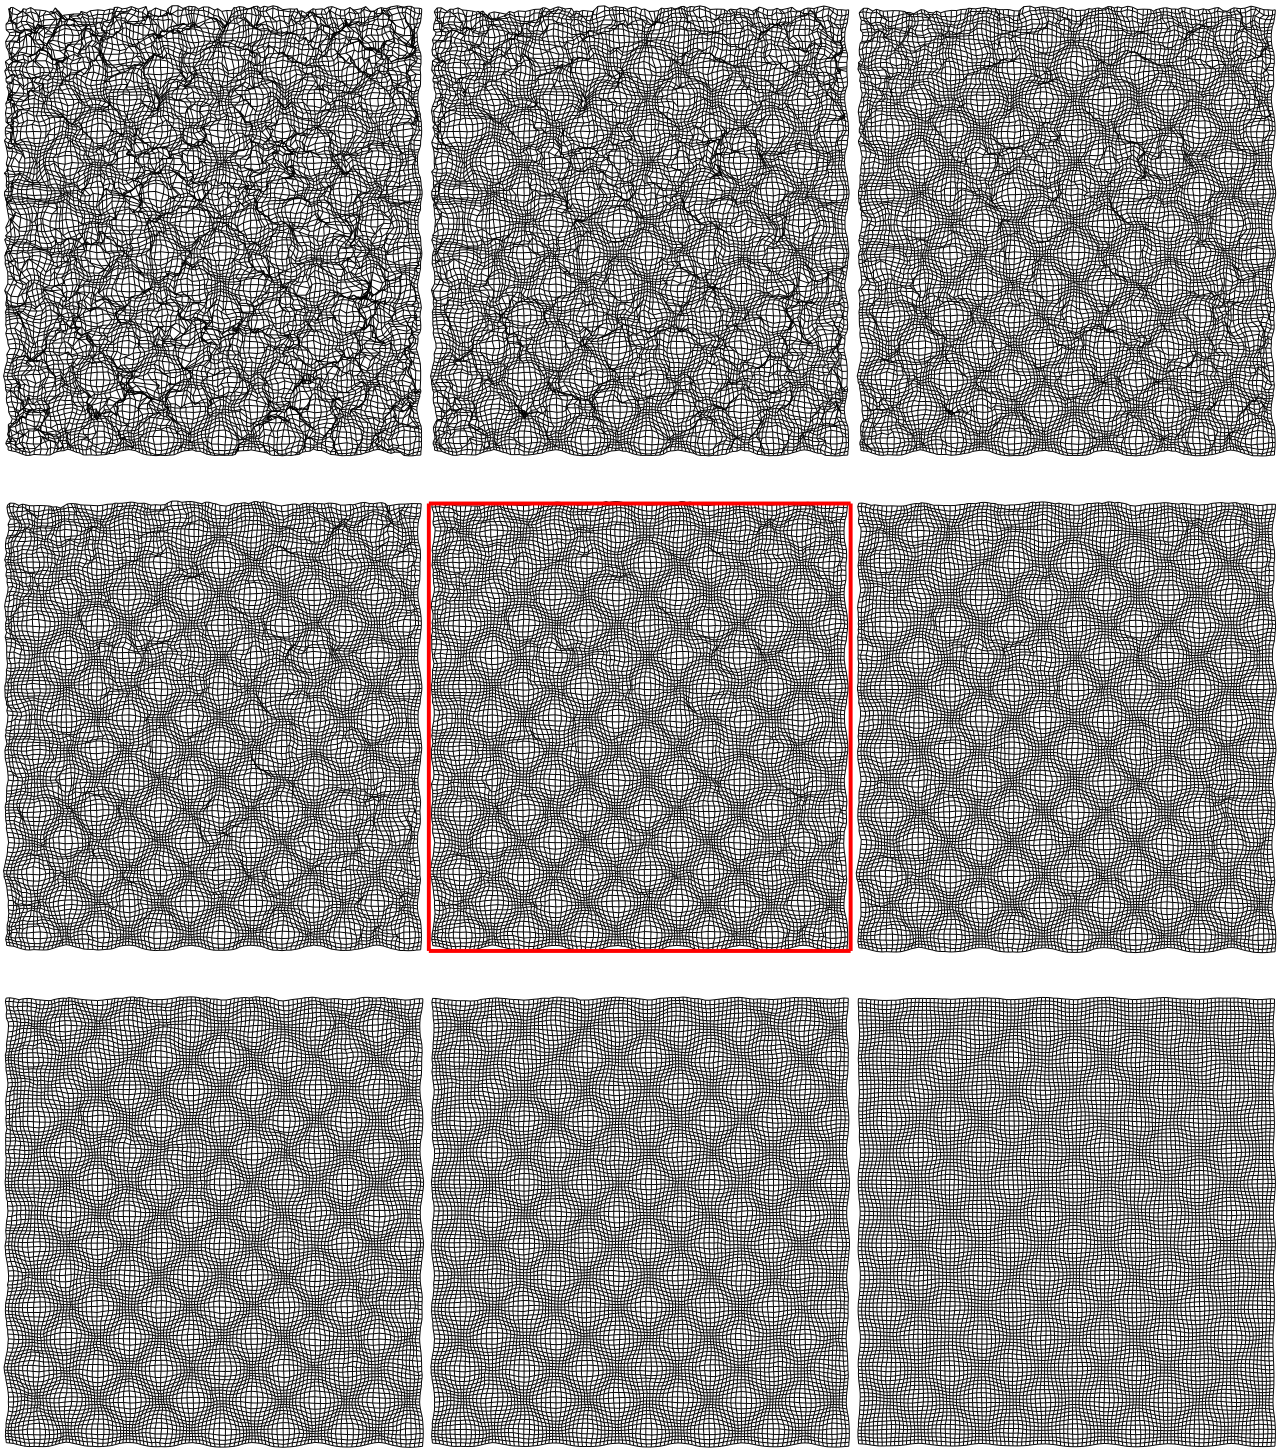

Figure 4: *Results of registration using* **MAMAN** *for regularization strengths ranging from 1.0 to 5.0. The smoothness of the transformation, as well as its closeness to the original transformation, is much more uniform across the image than for* **Asym** *(figure 3). The solution bordered in red is the closest to the ground truth 2(c), with an average error of 0.83 pixel.*

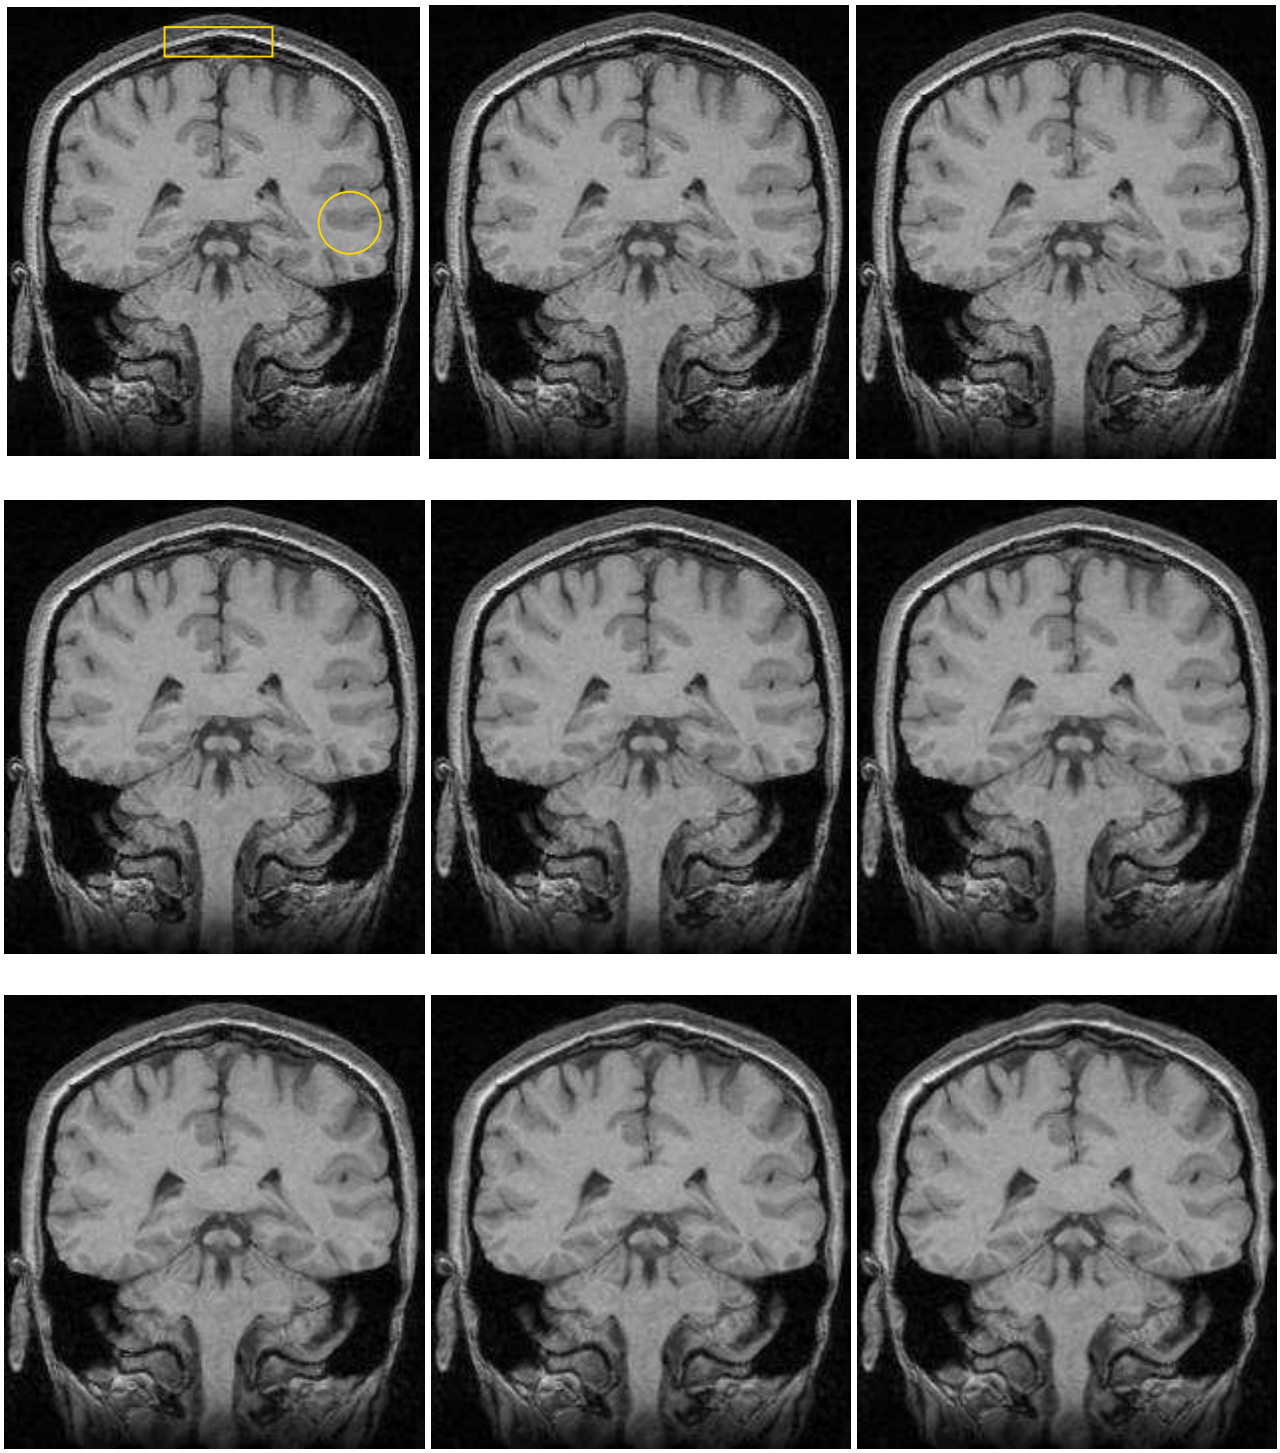

Figure 5: *Results of registration using the variational formulation (2) for increasing regularization strengths. The quality of the registration depends on the local contrast between the images, especially for moderate and high regularization strengths.*

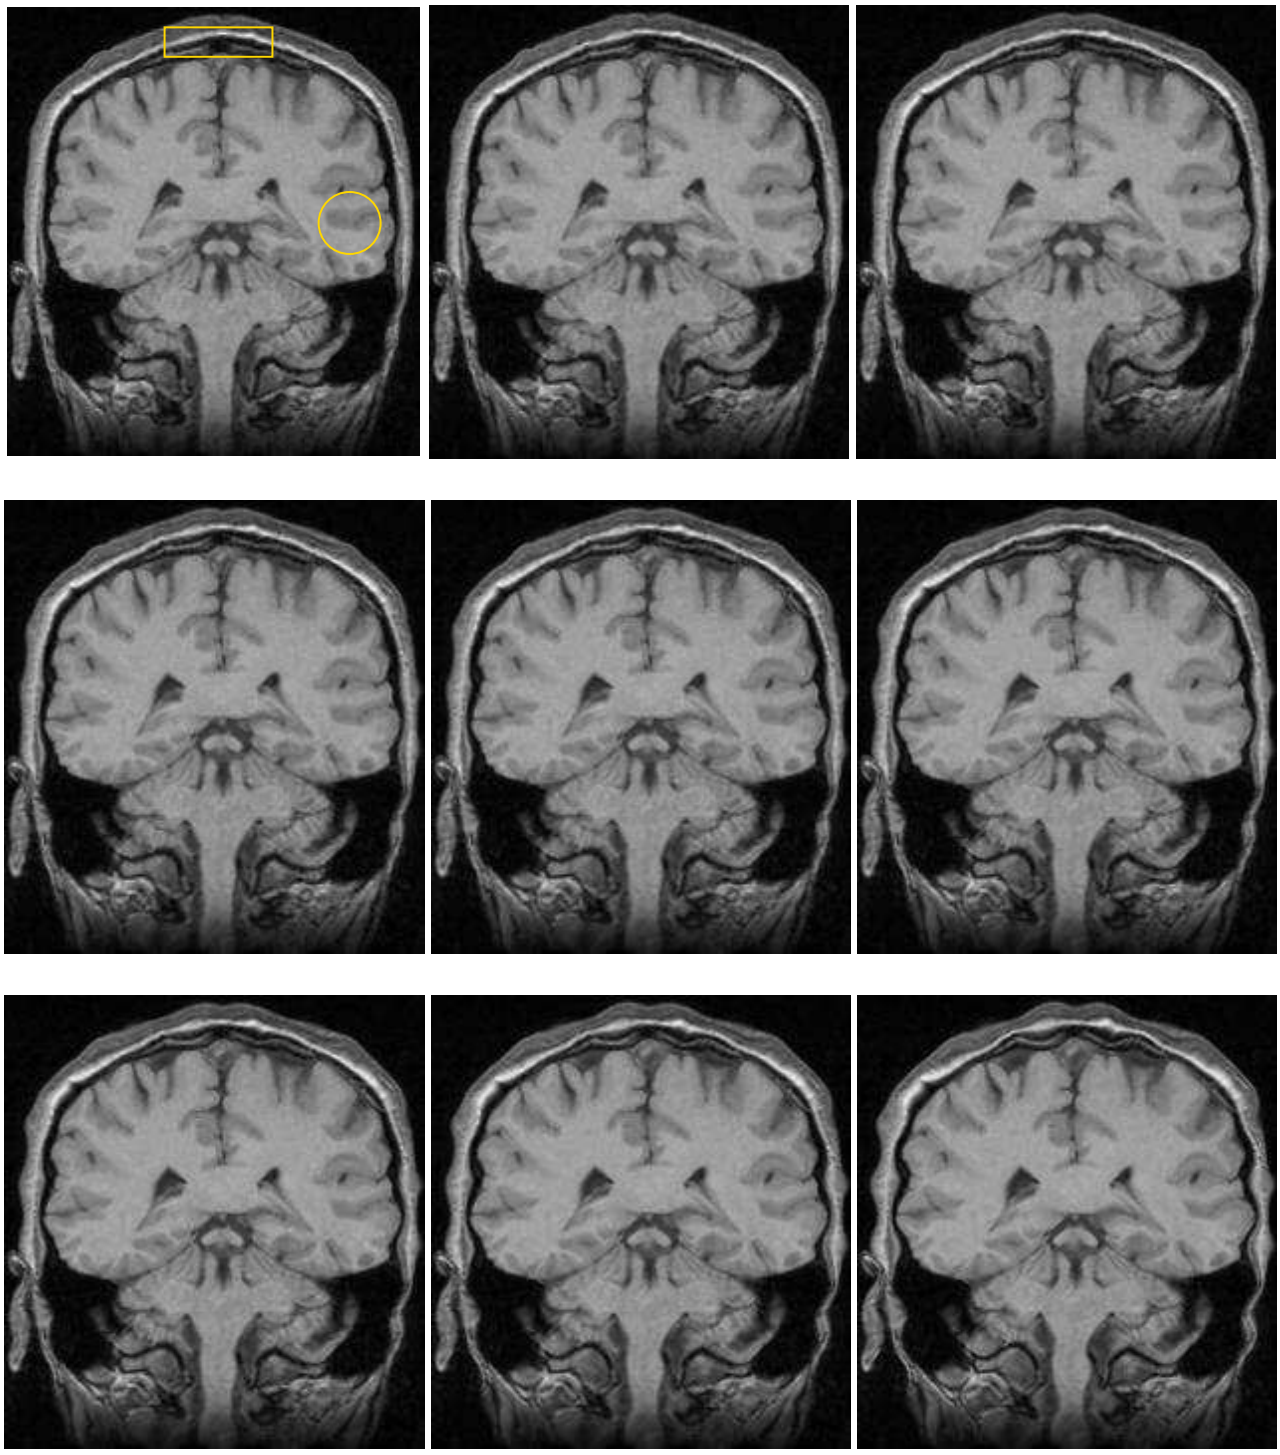

Figure 6: *Results of registration using* **MAMAN** *for increasing regularization strengths. The quality of the registration is uniformly degraded as the regularization strength increases.*

1. As expected, the higher the regularization strength, the smoother the transformation but the worst the image correspondence.

2. For both algorithms, best results are obtained for moderate regularization strengths, while the absence of noise could suggest that the optimal value should be zero, or at least very small. An explanation of this fact is proposed in the next section.

3. In the case of competitive SIB registration (figure 3), the transformation is very irregular almost everywhere in the image for low regularization strengths. As the regularization strength increases, the transformation becomes more regular only in places where the local contrast is low. At some point, the transformation has a discontinuous derivative only on the edges with very high contrast, e.g. on the boundary of the skull. Eventually, to smooth the transformation even on these edges we have to increase further the regularization strength, but then, the transformation is correctly estimated only where the contrast is high, as it is oversmoothed everywhere else.

4. In the case of the pair-and-smooth registration (figure 4), we also start from a transformation that is irregular almost everywhere. However, as the regularization strength increases, the smoothness of the transformation is increased in the same way everywhere in the image, independently of the local contrast of the image, and the quality of the estimation is more or less uniform and decreases as the strength of the regularization increases. This is not only an aesthetic consideration: it leads to better registration results, as it could be expected from the observation of the recovered transformations (figures 3 and 4). The best registration obtained with `MAMAN` has an average error of 0.83 pixel, whereas `Asym` can only come as close as 1.27 pixel on average.

5. The results are very difficult to compare when we look at the deformed images only. Especially, for low regularization strengths, all the images look the same and match rather well the target image 2(a), even if the smoothness of the transformation, and therefore the quality of the result, varies a lot. Therefore, comparison of non-rigid registration results based on the deformed image only is not really meaningful.

6. We can nonetheless retrieve some of these differences in the deformed images. Let us compare how the gray matter in the designated circle, and the skin in the designated box, are registered as a function of the regularization strength. Using `Asym`, the gray matter is starting to be mismatched very early, for regularization strength that are relatively low. On the opposite, the piece of fat is almost perfectly registered even for the highest regularization values. This is due to the fact that the borders between the gray matter and the white matter (dark gray on light gray) have much less contrast than the borders between fat and bone (white on black) in T1 MR images. On the opposite, with `MAMAN`, the mismatch of these regions progress simultaneously with the regularization strength. See also figure 7 for a close-up on these areas.

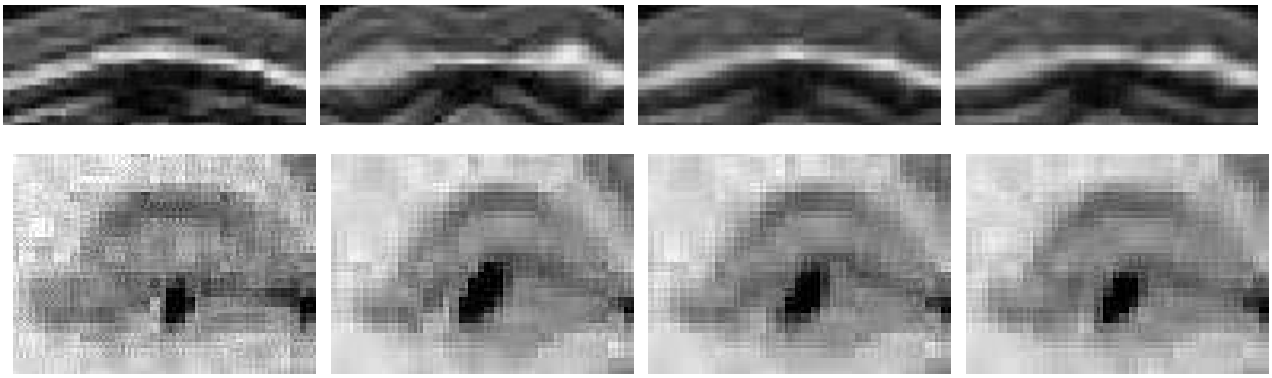

(a) Target image (b) Source image (c) Result of `Asym` (d) Result of `MAMAN`

Figure 7: *Zoom on a piece of skin (top) and gray matter (bottom, with enhanced contrast) in the image to be registered and for two results obtained respectively with* `Asym` *and* `MAMAN`*. With* `Asym`*, the same regularization strength can lead to a perfect match of the skin, while the grey matter is virtually not moved. With* `MAMAN`*, the regularization strength has a uniform impact on the registration result, and when the skin is partially matched, so is the grey matter.*

3 PASHA: A New P&S Algorithm based on a Global Hybrid Energy

3.1 One or Two Priors for Non-Rigid Registration?

Although there are a lot of different algorithms that we classified as "pair-and-smooth" (section 1), only a few of them minimize a global energy. The goal of this section is to propose a registration energy whose minimization leads naturally to a P&S algorithm.

In the `MAMAN` algorithm (Cachier et al., 1999) summarized in table 1, we are searching for correspondences between points by minimizing for instance the SSD. By doing so, small intensity changes may lead to very large displacements. This strategy is appropriate if the images are not corrupted by noise. Unfortunately, real medical images are always noisy. Thus, we do not want to pair points that are very distant from each other if their intensity difference is of the order of noise.

From this point of view, the behavior of `Asym` seems more interesting, because for low intensity differences, points are almost not registered. The parameter λ of the competitive SIB energy (1) could then be seen as being linked to the level of noise in the image.

In practice, however, this is not exactly true. In the experiment reported in figures 3 to 6, for both `Asym` and `MAMAN`, best results are obtained for moderate regularization strengths, while the absence of noise could suggest to set it to zero, or at least to a very small value. Therefore, the smoothness parameter λ seems to be more related to a prior we have on the smoothness of the transformation than on the noise level in the image. Note that in parametric SIB matching, the prior, which is the space of parametric transformation, is also linked to the smoothness of the transformation, rather than to the level of noise in the image.

Despite its good property described above, the energy (1) leads to unsatisfactory results because of the inhomogeneity of the trade-off: it leads to very non-uniform smoothness that keeps decreasing with the local contrast, as reported in section 2.2. Ideally, we want an algorithm that estimates uniformly the transformation like `MAMAN` in areas where the level of information

is significant, and that do not use the intensity information to estimate the motion in areas where the intensity differences are probably due to the noise, like for `Asym`. This means that we have to add another parameter to `MAMAN` which corresponds to some kind of threshold on the intensity similarity measure. More conceptually, we have to take into consideration *two* prior knowledges for non-rigid registration: the level of noise in the image, and the smoothness of the transformation.

3.2 A New Hybrid Energy for Non-Rigid Registration

Given a similarity energy $E_{sim}(I, J, T)$ and a regularization energy $E_{reg}(T)$, we propose to register the images I and J by introducing two non-rigid transformations C and T and minimizing

$$E(C, T) = E_{sim}(I, J, C) + \sigma \int ||C - T||^2 + \sigma\lambda E_{reg}(T) \tag{3}$$

with respect to C and T. C gives the correspondences between points of I and J, and should minimize the similarity measure while keeping close to T. T is the motion estimate, and should be both smooth and close to the set of correspondences given by C. The parameter σ is related to the level of noise in the image, and avoids correspondences C too far from T. The parameter λ is related to the prior smoothness of the transformation.

The alternate minimization of (3) w.r.t. C and T leads naturally to a P&S algorithm: in a first step, we search for pairings C between points by minimizing (3) w.r.t. C. During the second step we find the transformation estimate by minimizing (3) w.r.t. T.

One could minimize the energy (3) w.r.t. C and T simultaneously. However, when the regularization energy E_{reg} is quadratic, the alternate approach is appealing, because each step is very fast: the first step is a minimization done for each pixel separately, with only 2 or 3 degrees of freedom depending on the dimension of the image, thus very fast; then the second step is easily resolved by linear convolution. Even if the alternate minimization is known to take more iterations to converge, the overall process of minimization is relatively faster.

Based on this hybrid energy, we thus have designed a new pair-and-smooth algorithm called `PASHA`[4]. Its main steps are briefly summarized in table 2. As for `MAMAN` in table 1 we do not detail the multiresolution process, which can be found in (Cachier et al., 1999).

1. Set $n := 0$ and $T_0 := \mathrm{Id}$
2. Find C_{n+1} by minimizing $E_{sim}(I, J, C_{n+1}) + \sigma||C_{n+1} - T_n||^2$ by gradient descent
3. Find T_{n+1} by minimizing $||C_{n+1} - T_{n+1}||^2 + \lambda E_{reg}(T_{n+1})$ using convolution
4. Set $n := n + 1$
4. Go to step 2 until convergence

Table 2: Main steps of the `PASHA` non-rigid matching algorithm.

3.3 P&S Registration as a Similarity Measure Transform

The minimization of (3) can be rewritten as

$$\min_T [E^*_{sim}(I, J, T) + \sigma\lambda E_{reg}(T)] \tag{4}$$

[4] `PASHA`: Pair-And-Smooth, Hybrid-energy based Algorithm

with

$$E^*_{sim}(I, J, T) = \min_C \left[E_{sim}(I, J, C) + \sigma \int ||C - T||^2 \right] \quad (5)$$

Therefore, the minimization of the hybrid energy (3) can also be seen as the minimization of a weighted sum of two energies, a similarity energy and a regularization energy, with E^*_{sim} deduced from E_{sim} by the equation (5).

From this point of view, the concept of P&S algorithms is a mathematical transformation of standard intensity similarity measures E_{sim} into hybrid similarity measures E^*_{sim}. The regularization of these transformed similarity measures E^*_{sim} in a weighted sum, in a competitive fashion, give birth to most of P&S algorithms; but as for SIB registration, it is possible to regularize by parametric (e.g. rigid or spline) of fluid approaches. For example, in the case of `MAMAN`, fluid regularization has already been performed, and compared to competitive regularization in (Pennec et al., 1999). Block matching has also already been used for rigid registration (Ourselin et al., 2001).

3.4 Link with auxiliary variables

The formulation of the hybrid energy (3) is very similar to the one proposed in (Cohen, 1996). There is one major difference, though. In his article, Cohen aims at minimizing the energy:

$$E_{fit}(v) + E_{reg}(v) \quad (6)$$

For computational reasons, he introduces the auxiliary variable w and minimizes the energy

$$E^*_{fit}(w) + \frac{1}{2}||v - w||^2 + E_{reg}(v) \quad (7)$$

where E^*_{fit} is deduced from $E_{fit}(T)$ so that the energy (7) has the same global minimum as the energy (6). In our case, the hybrid energy (3) has certainly not the same solution as the competitive SIB energy (1). Note that most of the applications proposed in (Cohen, 1996) are purely geometric problems, and therefore the energy (6) is an homogeneous trade-off, standard in the field of data approximation.

3.5 Link with `MAMAN`

`MAMAN` as a limit case for noiseless images. `MAMAN` can be seen as being a limit case when σ tends toward 0. If $\sigma \to 0$, the weight of the similarity measure tends to infinity relatively to the closeness constraint $\int ||C - T||$; therefore, the correspondences C are found by minimizing the similarity energy $E_{sim}(I, J, C)$ alone. Then, the transformation T is found by minimizing $\int ||C-T||^2 + \lambda.E_{reg}(T)$. With this algorithm, we assume that the transformation is smooth but that there is no noise in the image, since corresponding points $\mathbf{x}$ and $C(\mathbf{x})$ have no constraint but having an intensity as similar as possible.

To illustrate this point, let us remark that in the experiment reported in figure 6 and 4, `MAMAN` recovers the motion even in areas outside of the skull, where there is no apparent information, which can seem strange. This is due to the fact that the image has been deformed without adding any noise to it, and so the assumption of noiseless images is true. Here, the algorithm relies on very subtle intensity changes to recover the motion. When there is noise in the image, this lack of robustness will prevent from an accurate recovery of the transformation, as we show in section 4.

Importance of the closeness constraint between C and T. Unfortunately, the global minimum of (3) when $\sigma \to 0$ gives poor registration results. Minimizing $E_{sim}(I, J, C)$ w.r.t. C gives a very irregular C with very large displacements, and T is a very bad estimate of the motion. Also, when $\sigma \to 0$, it is not necessary to alternate registration and regularization steps: since C does not depend on T, it can be found once and for all at the beginning of the algorithm.

The constraint that C should be close to T is essential. In most P&S algorithms, this constraint is actually more or less enforced by the minimization algorithm, which generally seeks for a local, close optimum rather than a global, remote optimum. They iterate between finding correspondences C and estimating the smooth transformation T, instead of finding the optimal correspondence C once and for all, and then deducing T. Therefore, these methods give good results, even if the global optimum of the energy is a very bad estimate of the motion.

For example, in block matching algorithms, the blocks are always translated into a few number of positions around the current estimate. In the "demons" algorithm (Thirion, 1998), the pairing is bounded by some threshold (see (Cachier et al., 1999) for a proof). In `MAMAN`, the SSD is minimized using a gradient descent with the starting guess being T, and so C is located at the local minima of the similarity energy closest to T.

However, all these implicit closeness techniques do not sufficiently constrain the solution because they are not part of the minimized energy, as we will see in the experiments of section 4.3.

3.6 Link with ICP-based image registration

(Feldmar et al., 1997) transform a n-D image $I(\mathbf{x}) : \mathbb{R}^n \to \mathbb{R}$ into an hypersurface $S : (\mathbf{x}, I(\mathbf{x}))$ of an extended space of dimension $n+1$. He proposes to see the registration of two images as a geometric matching of the two hypersurfaces deduced from the images. Roughly, his ICP-like algorithm proceeds in alternating two steps: finding corresponding points between the hypersurfaces by minimizing some distance, and finding a smooth approximation of these pairings. In this $(n+1)$-D extended space, the distance between two points lying on the surfaces generated by the image I and J is given by

$$d((\mathbf{x}_1, I(\mathbf{x}_1)), (\mathbf{x}_2, J(\mathbf{x}_2))) = [I(\mathbf{x}_1) - J(\mathbf{x}_2)]^2 + \sigma ||\mathbf{x}_1 - \mathbf{x}_2||^2 \tag{8}$$

with σ being a normalization constant, and $||.||$ the ordinary Euclidean norm.

It appears that minimizing this distance for all the image points is exactly minimizing our hybrid energy w.r.t. C (first step) when using the SSD as the similarity measure; thus, in this particular case, the formulation of Feldmar and ours are equivalent (see figure 8).

Note however that our hybrid energy (3) is more general, as it is not restricted to the SSD but can be used with more complex similarity measures, for example mutual information (Collignon et al., 1995; Wells et al., 1996), in the context of multimodal registration.

3.7 Mixing Intensity- and Feature-Based Registration

As a last evidence of the uniform smoothness given by the hybrid formulation (3), we want to give here the natural extension of this energy to the problem of mixing intensity and geometric feature in non-rigid registration of images.

The problem is now the following: in addition to the images I and J we furthermore have a discrete set of homologous points $\{(\mathbf{x}_1, \mathbf{y}_1), ..., (\mathbf{x}_p, \mathbf{y}_p)\}$, for examples anatomical landmark

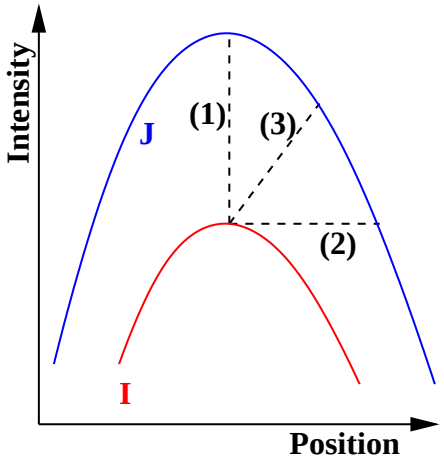

Figure 8: *Interpretation of some similarity measures in the extended space: (1) The SSD assumes image noise but no motion and uses only the distance on the intensity axis. The corresponding point used to compute the similarity has therefore the same geometric coordinates. (2)* ***MAMAN****, as well as some other P&S algorithms, assumes motion but no intensity noise, and uses only the distance on the geometric axes. The corresponding point has therefore the same intensity. (3) The ICP algorithm of (Feldmar et al., 1997), as well as our new formulation, assumes both noise and motion. The corresponding point has not necessarily the same position or intensity, but minimizes a weighted sum of intensity and geometrical distances. Competitive SIB and* ***MAMAN*** *are limit cases when* $\sigma \to 0\infty$ *and* $\sigma \to 0$, respectively.

points extracted manually in both images. We describe these pairings by a new correspondence function C_2, defined on the set $\{\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_p\}$ so that $C_2(\mathbf{x}_i) = \mathbf{y}_i$.

A natural extension of the registration energy (3) to take into these consideration these new pairings is:

$$E(C,T) = E_{sim}(I,J,C) + \sigma \int ||C-T||^2 + \sigma\alpha \sum_{i=1}^{p} ||C_2(\mathbf{x}_i) - T(\mathbf{x}_i)||^2 + \sigma\lambda E_{reg}(T) \qquad (9)$$

In this case, the second step of the alternate minimization is modified, and it is proved in (Cachier, 2001) that T has the following closed-form:

$$T(\mathbf{x}) = a.K * C(\mathbf{x}) + \sum_{i=1}^{p} \mathbf{a}_i K(\mathbf{x} - \mathbf{x}_i) \qquad (10)$$

where K is a kernel depending only on the energy E_{reg} (as previously), and the coefficients a and $\mathbf{a}_i$ are found by resolving a set of linear equations.

This formula mixes both smoothing (on the dense set of pairings C, found by the intensity similarity) and radial basis functions (centered on the points where a pairing has been explicitly given) — and furthermore, the smoothing kernel and the radial basis function are the same. The transformation thus obtained is intuitively uniformly smooth, and its smoothness depends only on the shape and size of the kernel K.

In the case where the homologous points of the $\mathbf{x}_i$ are not explicitly given, but are known to be in a certain set of points $\mathbf{y}_j$, for example after segmentation of a part of the anatomy, the correspondences C_2 have to be estimated too. The alternate minimization of (9) w.r.t. C, C_2 and T then leads to a three-step algorithm, while T remains of the form (10). An application in the case of brain matching with cortical folds segmentation is given in (Cachier et al., 2001).

4 Experiments with the PASHA Algorithm

4.1 Synthetic Sinusoidal Deformation, part II

We have run the experiment of the section 2.2 with this new PASHA algorithm to compare it to Asym and MAMAN. The energy minimized here also uses the SSD as the intensity similarity measure, and the membrane energy for regularization:

$$E(C,T) = \int (I - J \circ C)^2 + \sigma \int ||C - T||^2 + \sigma\lambda \int ||dT||^2$$

The results are reported in figures 9 and 10. The minimization of the hybrid energy (11) behaves exactly as we wanted: the estimated transformation is uniformly smooth, as with MAMAN, except in areas containing no information but noise, where the transformation is not recovered, as for Asym.

We want to emphasize the fundamental difference between PASHA and Asym. With PASHA, a transformation is not estimated if it leads to a small intensity similarity gain, and is just extrapolation; everywhere else, the transformation is estimated, and as with MAMAN is uniformly smooth. With Asym, and more generally with competitive SIB algorithms, there is not really this kind of threshold: the transformation still becomes less regular as the local contrast between the images increases.

4.2 Quantification of Errors with Gaussian Random Fields

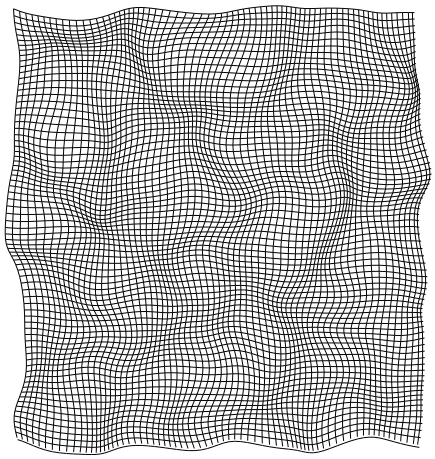

Figure 11: *An example of a Gaussian random transformation.*

We quantified the differences between PASHA and MAMAN on a set of non-rigid transformations. Following (Ruiz-Alzola et al., 2000), we chose a (non-parametric) set of isotropic Gaussian random vector fields (Adler, 1980). An example of a Gaussian random transformation is depicted in figure 11.

In this experiment, we first performed anisotropic diffusion of image 2(a) with the algorithm of (Krissian et al., 1997) in order to suppress the noise. We generate a Gaussian random transform with Gaussian noise of standard deviation $k*s$ smoothed with a Gaussian kernel of size s pixel, k being a constant, so that the "expected smoothness" of such a transformation is controlled by s and the amplitude of the deformation remains constant for all smoothness strength s.

Given a Gaussian random transformation, we deformed the image 2(a), and added white Gaussian noise of standard deviation n. We also added the same level of Gaussian noise to the original image, and matched it to the deformed image. The error of registration is then the mean distance between this estimation and the real transformation in the entire image.

We have run this experiment with MAMAN and PASHA, with three set of parameters for the deformation simulation: a reference set with standard noise and smoothness (n=10 and s=10) , a set with smoother transformations (n=10 and s=20), and a set with noisier images ($n = 15$ and $s = 10$). For each of these experiment, we used three different regularization strengths in

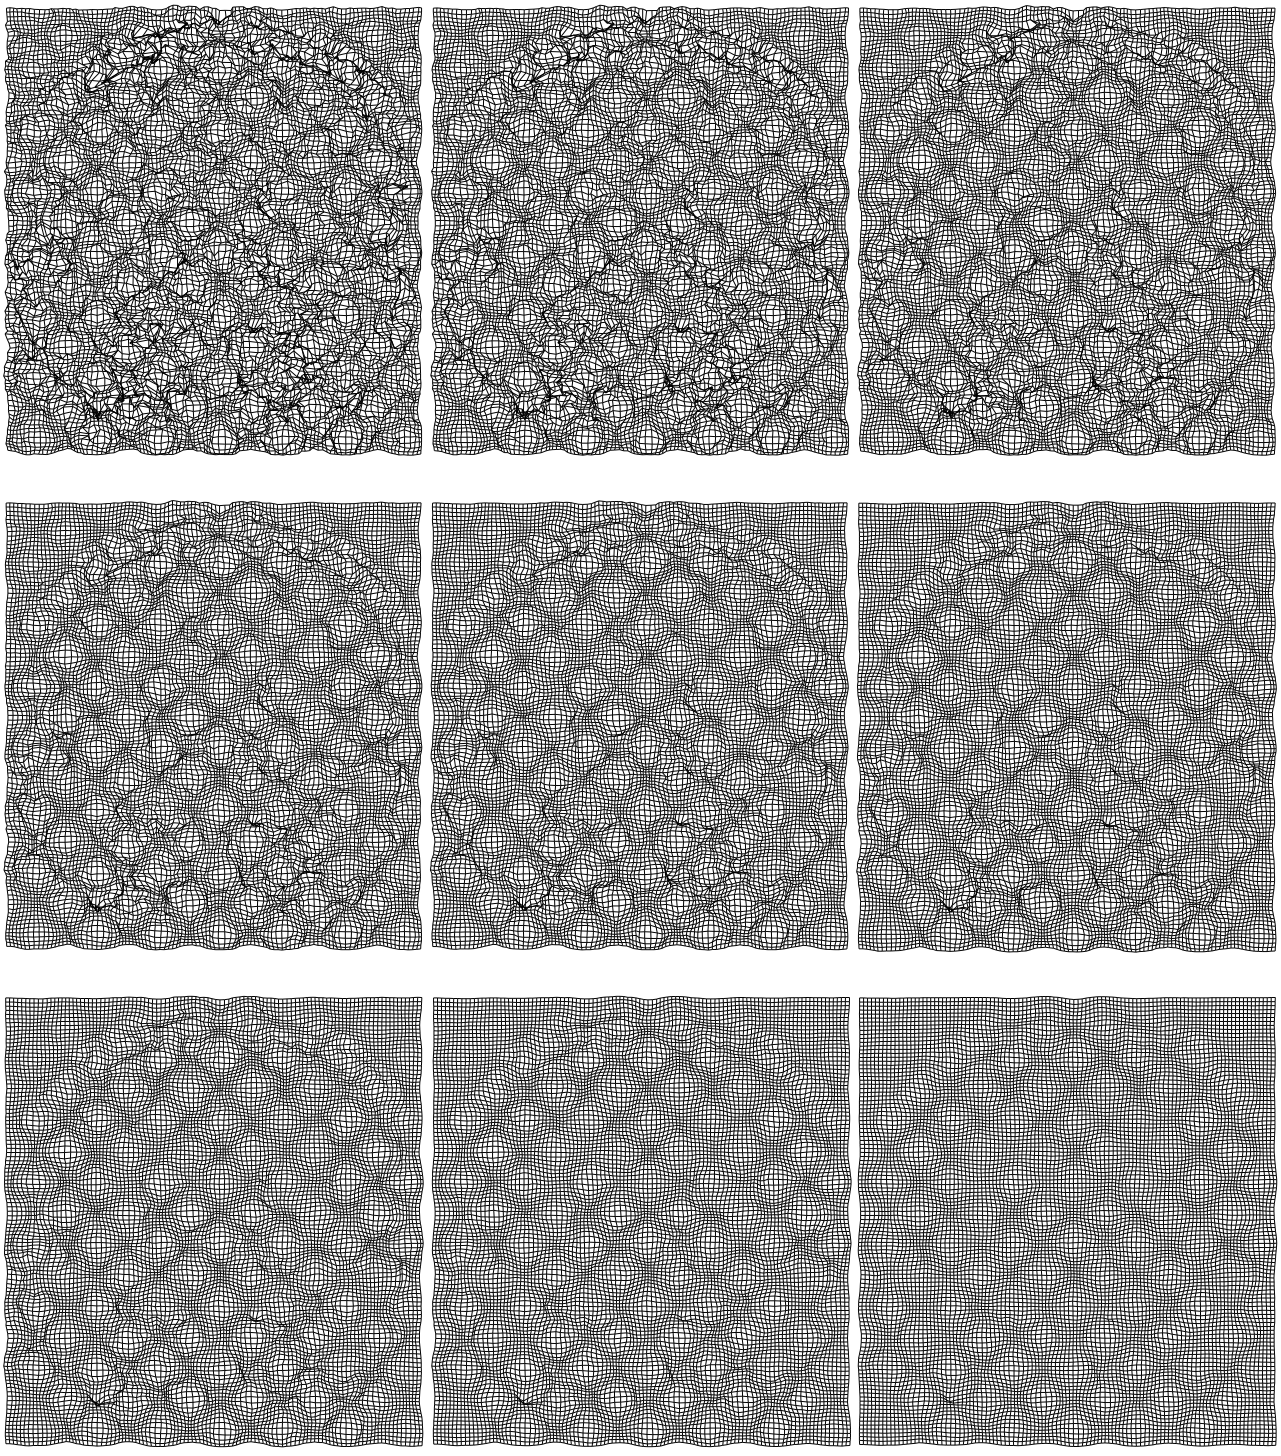

Figure 9: *Results of registration using our new hybrid energy (11) for regularization strengths ranging from 1.0 to 5.0, with* $\sigma = 5$.

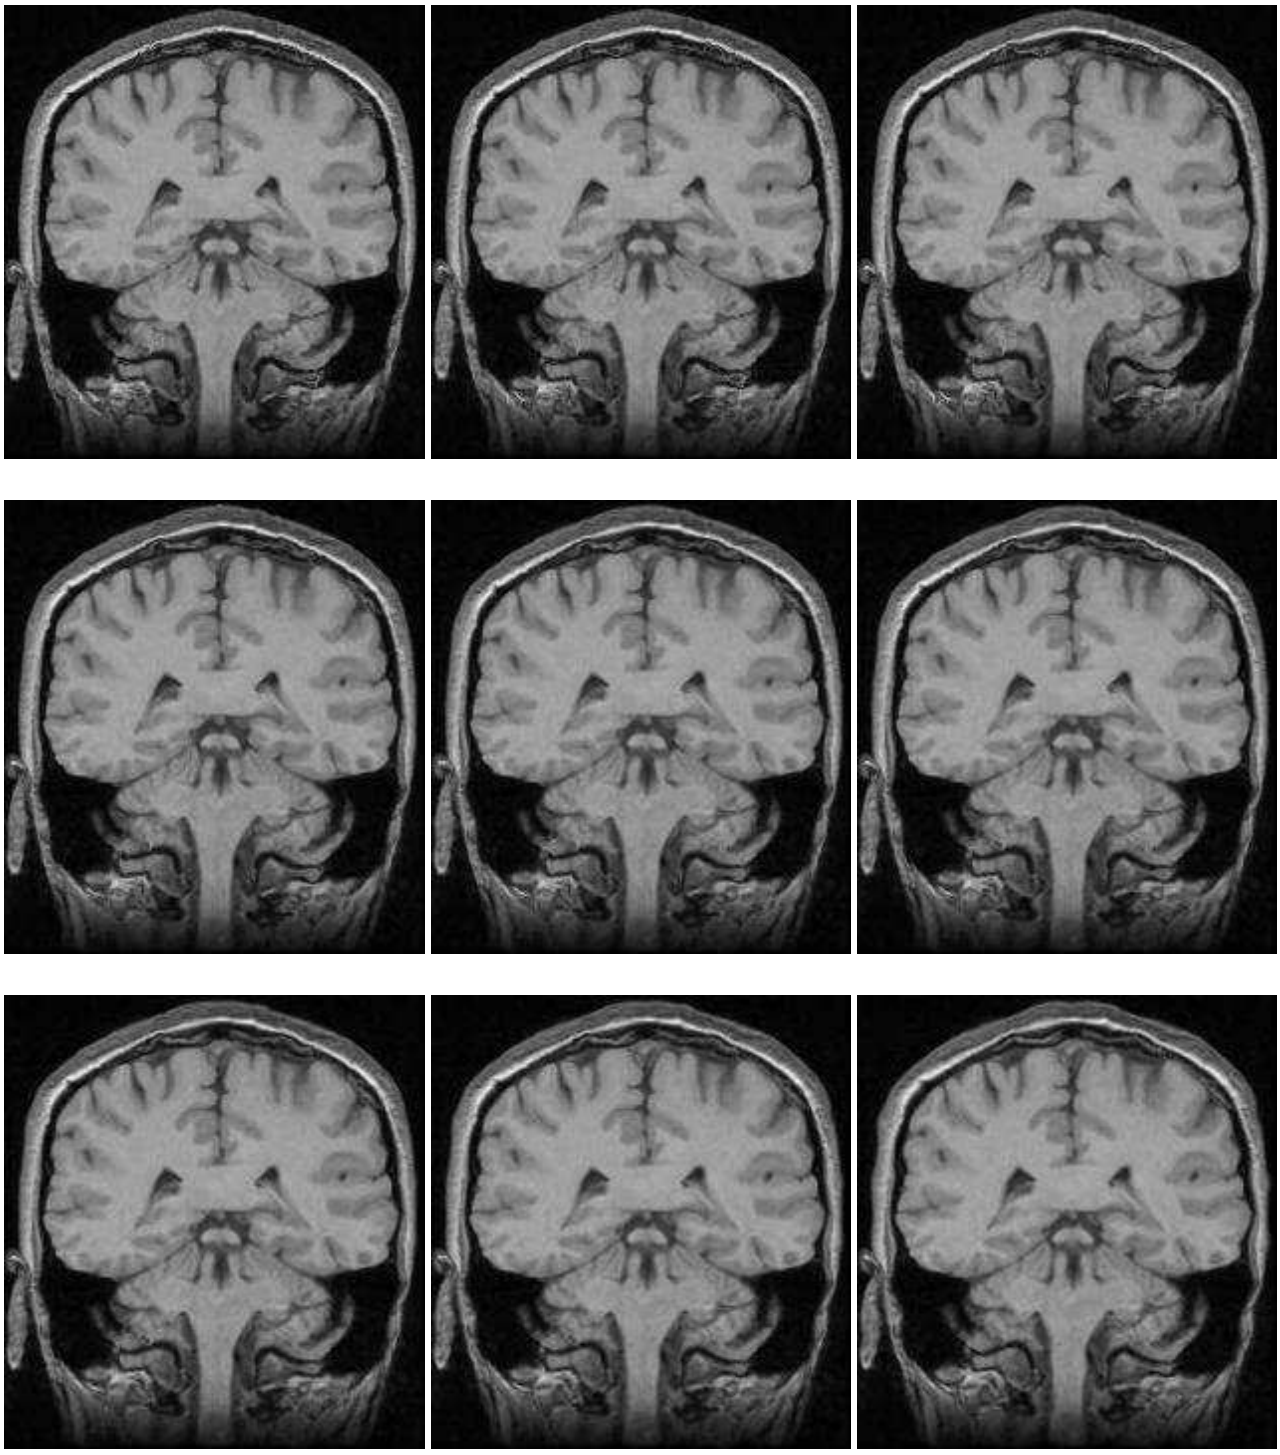

Figure 10: *Results of registration using our new hybrid energy (11) for increasing regularization strengths.*

our algorithm. Also, for PASHA, for each regularization strength, we choose different values of σ.

For all these different deformation, noise and registration parameters, we ran 20 experiments, and averaged the errors. These mean errors are plotted in figure 12.

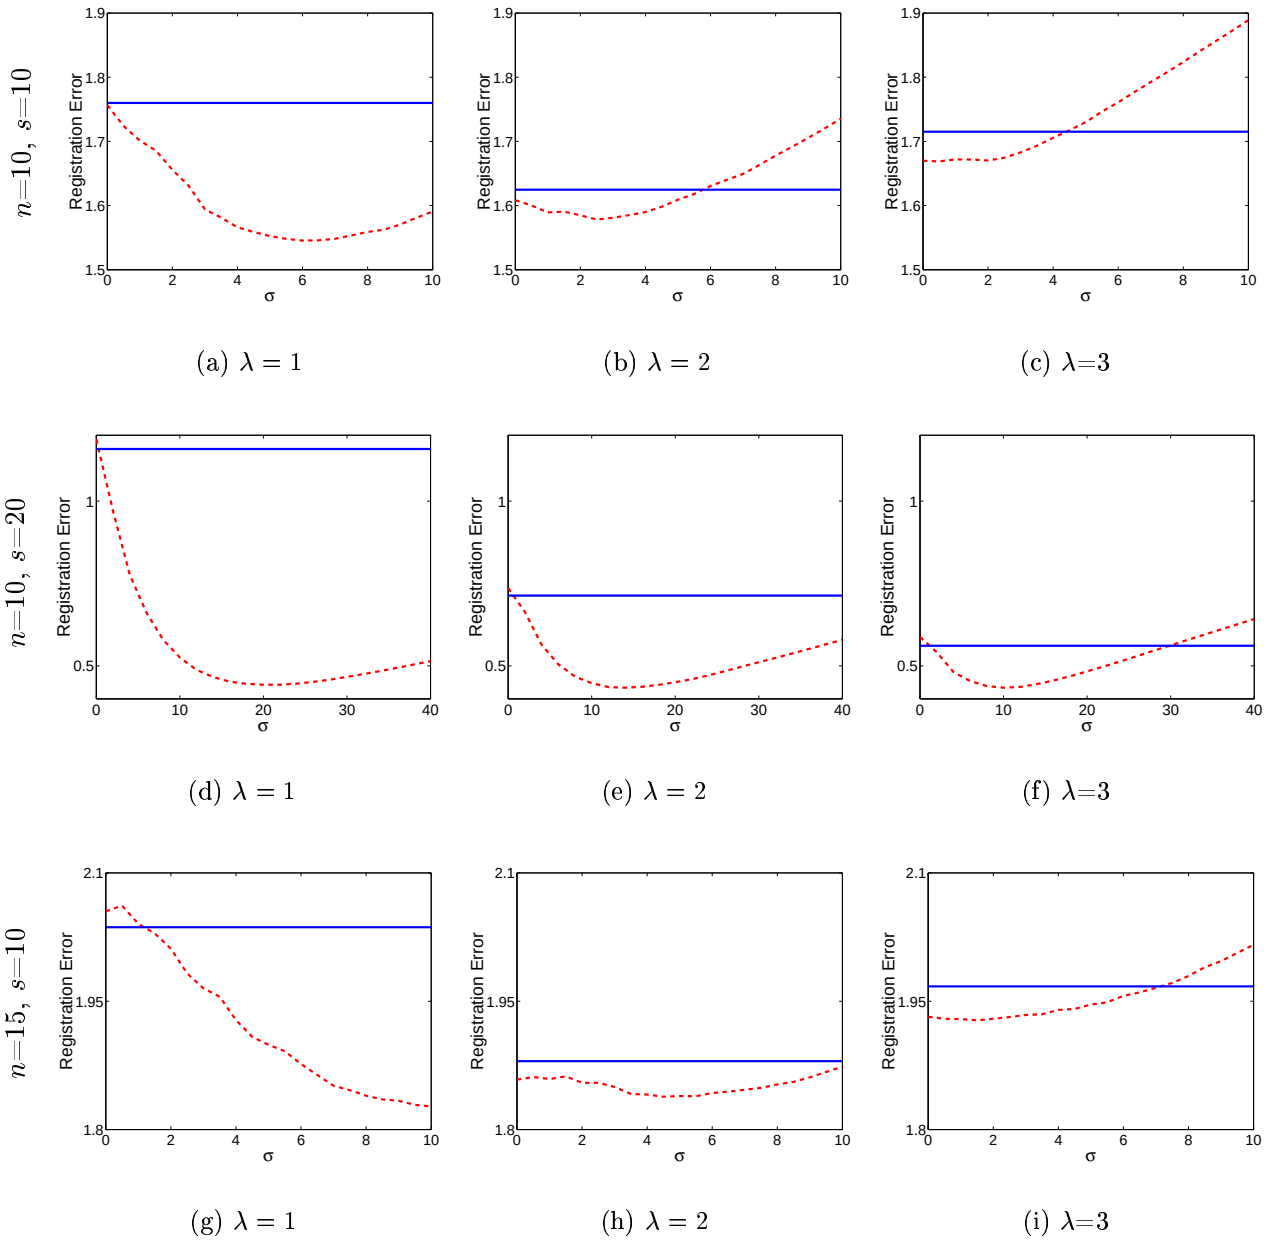

(a) $\lambda = 1$ (b) $\lambda = 2$ (c) λ=3

(d) $\lambda = 1$ (e) $\lambda = 2$ (f) λ=3

(g) $\lambda = 1$ (h) $\lambda = 2$ (i) λ=3

Figure 12: *Error of registration for different set of parameters.* ***Dashed Line:*** *Errors with* PASHA*. For each value of λ, the quality of the registration depends on the assumed level of noise σ.* ***Plain Line:*** *Error with* MAMAN*. Since there is no parameter σ in this algorithm the results do not depend on σ.* ***First Row:*** *Reference experiment: $s = 10$, $n = 10$.* ***Second Row:*** *Smoother deformations: $s = 20$, $n = 10$.* ***Third Row:*** *Noisier images: $s = 10$, $n = 15$.* ***Column, left to right:*** *Registration parameter λ= 1, 2, 3.*

According to the results, if the parameters are chosen carefully, the hybrid energy gives better results than MAMAN. This shows that relying on very subtle intensity changes can prevent us from retrieving the correct motion when dealing with noisy images.

When the transformation is smoother, for each value of λ, the optimal value of σ stay relatively stable. On the contrary, the relative score of the different regularization strength λ is changing, and high values of λ give more better results than with less smooth Gaussian transforms.

When the images are noisier, the relative score of the different value of λ do not vary very much. However, for each value of λ, the optimal value of σ is translated to higher values.

This tends to confirm that λ is much more related to the prior smoothness of the transformation, and σ to the level of noise in the image. Note however that the choice of σ and λ is less simple than expected: the best couple of registration parameters (σ, λ) for PASHA has not the same optimal regularization strength λ than MAMAN — here, the optimal λ of PASHA seems to be always lower than the one of MAMAN, as if σ still explains a part of the smoothness. Further work has to be done to understand the link between the parameters.

It is more difficult to compare Asym with MAMAN and PASHA since the regularization strength has not the same range for both class of algorithms. We have run Asym with the standard set $n = 10$ and $s = 10$, for a large number of regularization parameters. The best results were obtained for a regularization strength of $\lambda = 1000$; for this optimal value, the average error is 1.81, which is above what can achieve the two P&S algorithms (1.62 for MAMAN with $\lambda = 2$, 1.55 for PASHA with $\lambda = 1$ and $\sigma = 6$). As in section 2.2, the lack of uniform smoothness of Asym leads to greater errors in the motion estimation.

4.3 Real 3D Experiment

We registered two $256 \times 256 \times 128$ MR images of two different patients, with MAMAN and PASHA, using the SSD as the similarity measure and a Gaussian kernel for regularization (figure 13). The registration took around 16 minutes for both algorithms, on a Pentium II 500MHz. Prior

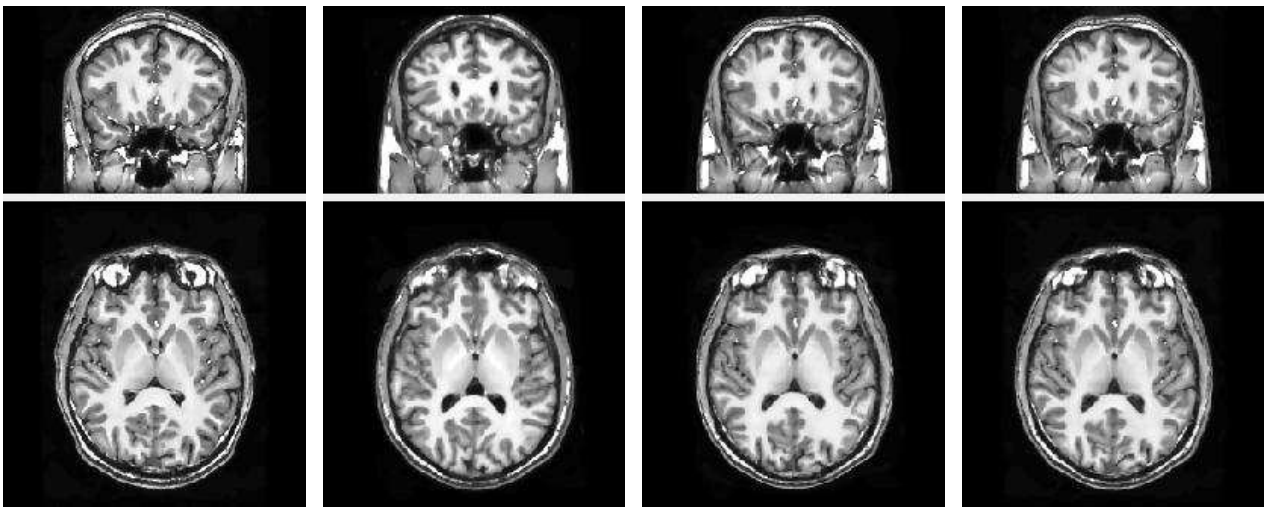

(a) Source Image (b) Target Image (c) MAMAN result (d) PASHA result

Figure 13: *A real 3D registration experiment: Original and deformed images.*

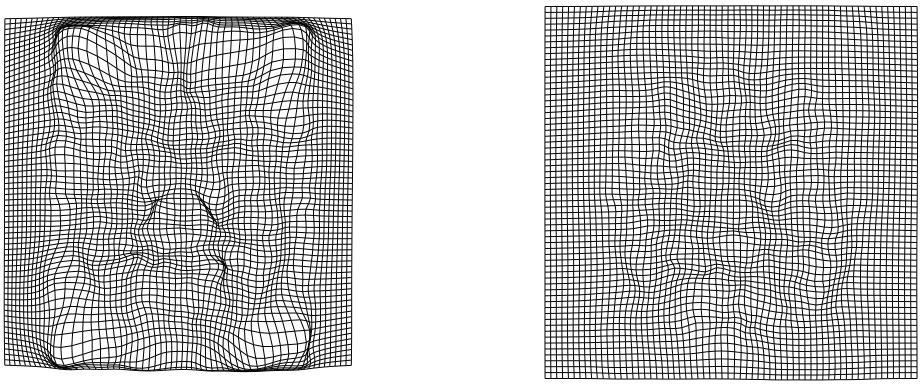

(a) Result of MAMAN (b) Result of PASHA

Figure 14: *Transformations found by MAMAN and PASHA. The differences between the two algorithms are made obvious here, and are mainly due to slightly different levels of intensity in the white matter (see also figure 15) and in the background between the source and target images. These strong differences in the recovered transformation leads nonetheless to very small intensity differences, as can be seen in figure 16. This illustrates that looking at the similarity of voxel intensities only can be misleading.*

to non-rigid registration, these images have been anisotropically diffused (Krissian et al., 1997), their intrinsic as well as relative bias have been removed (Mangin, 2000), and they have been linearly registered with the algorithm of (Roche et al., 2000). The level of noise in these images is thus negligible.

However, there still remains some small violation of the intensity conservation assumption in the images. The first violation is that the average intensity level of the background is slightly higher in the target image than in the source image. The second violation is that the average intensity level of the white matter is slightly lower in the target image than in the source image.

These slight intensity differences lead to some problems with MAMAN, as it can be seen in figure 14. The background in the target image is stretched to the border of the source image to minimize the SSD. Also, at some places, the white matter is artificially dilated (figure 15). The PASHA algorithm is much more robust towards these slight violations of the intensity conservation assumption: the points in the background are not moved despite the intensity differences, and the inside of the white matter is not inflated. Note that these differences typically appear for points with very slight intensity differences; therefore, their impact on the deformed image and the final intensity difference with the target image is very small, as it can be seen in figure 16.

5 Conclusion

We have proposed a new classification of image non-rigid registration algorithms. The geometric algorithms use sparse features extracted from the images. The standard intensity based (SIB) algorithms use a single transformation that register the images with an intensity similarity measure while constrained to remain smooth. The pair-and-smooth (P&S) algorithms proceed in two step: the first step pairs homologous points according to the intensity similarity measure, and the second step approximates these pairings with a smooth transformation.

We have compared the competitive SIB and P&S approaches, which both use two energies to register the images: a similarity energy, and a regularization energy. We have shown that the regularization parameter of the competitive SIB algorithms is not directly related to the smoothness of the transformation, which tends to be very non-uniform: the higher the local contrast between the images, the more irregular the transformation. With P&S algorithms, the regularization strength has a uniform influence on the smoothness of the estimated transformation. We have shown on a synthetic experiment this fundamental difference between both approaches, with two instances of competitive SIB and P&S algorithms called respectively

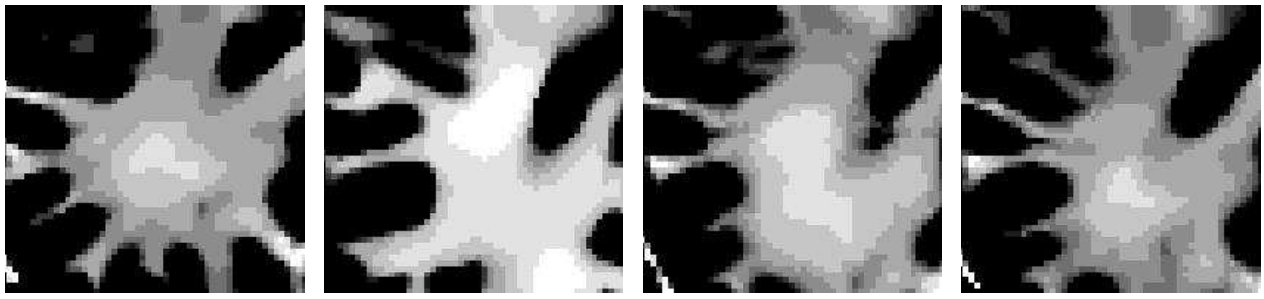

(a) Original (b) Target (c) Result of `MAMAN` (d) Result of `PASHA`

Figure 15: *Close-up on a part of the white matter, with intensities saturated in the range [45;53]. The white matter intensity is slightly higher in the target image than in the original image. Therefore, while the boundary of the deformed white matter is almost the same in images 15(c) and 15(d),* `MAMAN` *attempts to dilate the areas of the white matter with the higher intensities to minimize the intensity differences, visualized here by the expansion of the light gray region in 15(c) compared to its original position in 15(a). With* `PASHA`*, the closeness constraints forbids such deformations.*

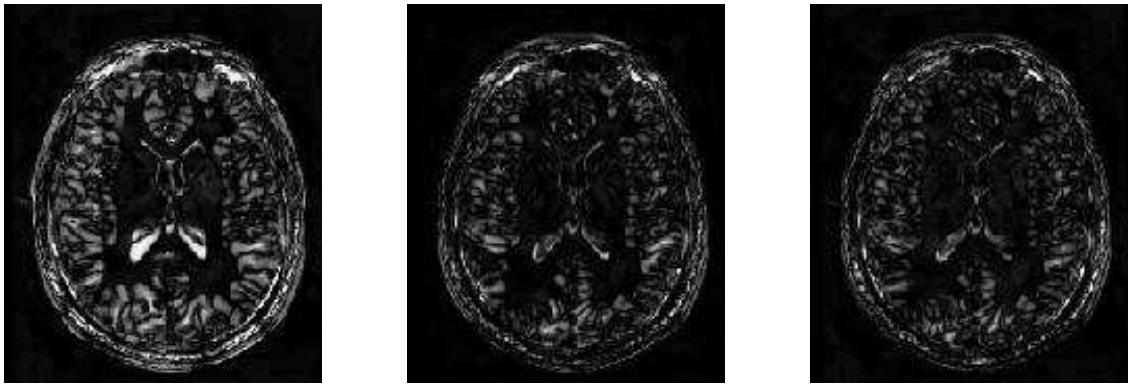

(a) `Before matching` (b) `MAMAN` (c) `PASHA`

Figure 16: *Intensity difference between target image and deformed original image, with* `MAMAN` *and* `PASHA`*. From this point of view, both methods give very similar results, although* `PASHA` *seems slightly better. Perfect registration is not possible in that case, as it is multipatient registration.*

`Asym` and `MAMAN`. We have also shown that these consideration is not only aesthetic but leads in practice to less accurate results.

We have introduced a hybrid energy for image registration, whose minimization can lead naturally to a P&S algorithm. This energy does not use a single parameter, but two, which are linked to two prior knowledges: the supposed levels of image noise and transformation smoothness. This energy allows to better understand `MAMAN` (and other P&S algorithms as well) as an algorithm that does not assume any image noise. Our new algorithm `PASHA`, based on this energy, can therefore be seen as a generalization of `MAMAN`. Like competitive SIB algorithms, the motion with `PASHA` is not recovered if it leads to a small intensity similarity gain. However, the fundamental difference is that for high intensity similarity gain, the smoothness is uniform for `PASHA`, while the transformation still becomes less regular as the local contrast increases with SIB algorithms.

We have shown on a set of noisy images deformed by Gaussian random transformations the advantages of `PASHA` compared to our two other algorithms. These experiments also showed that the parameters of the hybrid energy are not as easy to pick as expected, as they do not totally independently correspond to the level of image noise and transformation smoothness.

On a real 3D experiment, we have shown that the hybrid energy significantly increases the robustness of our P&S algorithm towards the hypothesis made by the similarity measure (e.g. intensity conservation for the SSD criterion). These differences are easy to see on the recovered transformation. However, since they typically appear on areas with small intensity differences, the deformed image looks almost identical.

Our new energy formulation of image registration depends on an intensity similarity energy and a regularization energy. While intensity similarities has been studied and compared in numerous papers, e.g. (Roche et al., 2000), regularization energies have been relatively less focused on. We propose to continue our study in (Cachier, 2001) and focus on regularization energies, and to deduce their associated filters and basis functions, particularly in the context of P&S registration.

Acknowledgments

Many thanks to Hervé Delingette and Xavier Pennec for their proofreading and very helpful remarks.

Contents

References

Adler, R. J. (1980). *The Geometry of Random Fields*. John Wiley & Sons.

Alexander, D. C. and Gee, J. C. (2000). Elastic Matching of Diffusion Tensor Images. *Comp. Vision and Image Understanding*, 77(2):233 – 250.

Althof, R. J., Wind, M. G. J. K., and Dobbins, J. T. (1997). A Rapid and Automatic Image Registration Algorithm with Subpixel Accuracy. *IEEE Trans. on Medical Imaging*, 16(3):308 – 316.

Alvarez, L., Weickert, J., and Sánchez, J. (2000). Reliable Estimation of Dense Optical Flow Fields with Large Displacements. *Int. J. of Comp. Vision*, 39(1):41 – 56.

Amini, A. A., Chen, Y., and Abendschein, D. (1999). Comparison of Landmark-Based and Curve-Based Thin-Plate Warps for Analysis of Left-Ventricular Motion from Tagged MRI. In *Proc. of MICCAI'99*, volume 1679 of *LNCS*, pages 498 – 507, Cambridge, UK. Springer.

Amit, Y. (1994). A Nonlinear Variational Problem for Image Matching. *SIAM J. on Scientific Comp.*, 15(1):207 – 224.

Andresen, P. R., Bookstein, F. L., Conradsen, K., Ersbøll, B. K., Marsh, J. L., and Kreiborg, S. (2000). Surface-Bounded Growth Modeling Applied to Human Mandibles. *IEEE Trans. on Medical Imaging*, 19(11):1053 – 1063.

Arun, K. S., Huang, T. S., and Blostein, S. D. (1987). Least-Squares Fitting of Two 3-D Point Sets. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 9(5):698 – 700.

Ashburner, J., Andersson, J. L. R., and Friston, K. J. (2000). Image Registration Using a Symmetric Prior – in Three Dimensions. *Human Brain Mapping*, 9(4):212 – 225.

Ashburner, J. and Friston, K. J. (1999). Nonlinear Spatial Normalization Using Basis Functions. *Human Brain Mapping*, 7:254 – 266.

Bajcsy, R. and Kovačič, S. (1989). Multiresolution Elastic Matching. *Comp. Vision, Graphics and Image Processing*, 46:1–21.

Benayoun, S. and Ayache, N. (1998). Dense Non-Rigid Motion Estimation in Sequences of Medical Images Using Differential Constraints. *Int. J. of Comp. Vision*, 26(1):25 – 40.

Boes, J. L. and Meyer, C. R. (1999). Multi-Variate Mutual Information for Registration. In *Proc. of MICCAI'99*, volume 1679 of *LNCS*, pages 606 – 612, Cambridge, UK. Springer.

Bookstein, F. L. (1994). *Functional Neuroimaging: Technical Foundations*, chapter 10: Landmarks, Edges, Morphometrics, and the Brain Atlas Problem, pages 107–119. Academic Press.

Bricault, I., Ferretti, G., and Cinquin, P. (1998). Registration of Real and CT-Derived Virtual Bronchoscopic Imag es to Assist Transbronchial Biopsy. *IEEE Trans. on Medical Imaging*, 17(5):703–714.

Bro-Nielsen, M. and Gramkow, C. (1996). Fast Fluid Registration of Medical Images. In *Proc. of VBC'96*, volume 1131 of *LNCS*, pages 267–276, Hamburg, Germany. Springer.

Broit, C. (1981). *Optimal Registration of Deformed Images*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania.

Burr, D. J. (1981). A Dynamic Model for Image Registration. *Comp. Graphics and Image Proc.*, 15(2):102–112.

Cachier, P. (2001). Regularization in Image Non-Rigid Registration: II. Orthotropic Energies, Filters and Splines. Technical report, INRIA. to appear.

Cachier, P., Mangin, J.-F., Pennec, X., Rivière, D., Papadopoulos-Orfanos, D., Régis, J., and Ayache, N. (2001). Multisubject Non-Rigid Registration of Brain MRI using Intensity and Geometric Features. In *Proc. of MICCAI'01*. in press.

Cachier, P. and Pennec, X. (2000). 3D Non-Rigid Registration by Gradient Descent on a Gaussian-Windowed Similarity Measure using Convolutions. In *Proc. of MMBIA'00*, pages 182 – 189, Hilton Head Island, USA.
`http://www-sop.inria.fr/epidaure/personnel/Pascal.Cachier/publi.html`.

Cachier, P., Pennec, X., and Ayache, N. (1999). Fast Non-Rigid Matching by Gradient Descent: Study and Improvements of the "Demons" Algorithm. Technical Report RR-3706, INRIA.

Cachier, P. and Rey, D. (2000). Symmetrization of the Non-Rigid Registration Problem using Inversion-Invariant Energies: Application to Multiple Sclerosis. In *Proc. of MICCAI 2000*, pages 472 – 481, Pittsburgh, USA.
`http://www-sop.inria.fr/epidaure/personnel/Pascal.Cachier/publi.html`.

Chen, F. and Suter, D. (1997). Elastic Spline Models for Human Cardiac Motion Estimation. In *IEEE Nonrigid and Articulated Motion Workshop*, pages 120 – 127, Puerto Rico. IEEE Computer Society.

Christensen, G. E. (1999). Consistent Linear-Elastic Transformations for Image Matching. In *Proc. of IPMI'99*, volume 1613 of *LNCS*, pages 224 – 237, Visegrád, Hungary. Springer.

Christensen, G. E., Joshi, S. C., and Miller, M. I. (1997). Volumetric Transformation of Brain Anatomy. *IEEE Trans. on Medical Imaging*, 16(6):864–877.

Chui, H. and Rangarajan, A. (2000). A Feature Registration Framework using Mixture Models. In *Proc. of MMBIA'00*, pages 190 – 197, Hilton Head Island, USA.

Cohen, L. D. (1996). Auxiliary variables and two-step iterative algorithms in computer vision problems. *J. of Mathematical Imaging and Vision*, 6(1):59 – 83.

Collignon, A., Maes, F., Delaere, D., Vandermeulen, D., Suetens, P., and Marchal, G. (1995). *Information Processing in Medical Imaging*, chapter Automated Multi-Modality Image Registration based on Information Theory, pages 263 – 274. Kluwer Academic Publishers.

Collins, D. L. and Evans, A. C. (1997). ANIMAL Validation and Applications of Nonlinear Registration Based Segmentation. *Int. J. of Pattern Recognition and Artificial Intelligence*, 11(8):1271 – 1294.

Collins, D. L., Le Goualher, G., and Evans, A. C. (1998). Non-Linear Cerebral Registration with Sulcal Constraints. In *Proc. of MICCAI'98*, volume 1496 of *LNCS*, pages 974 – 984, Cambridge, USA. Springer.

Davis, M. H., Khotanzad, A., Flamig, D. P., and Harms, S. E. (1997). A Physics-Based Coordinate Transformation for 3D Image Matching. *IEEE Trans. on Medical Imaging*, 16(3):317 – 328.

Dawant, B. M., Hartmann, S. L., and Gadamsetty, S. (1999). Brain Atlas Deformation in the Presence of Large Space-Occupying Tumors. In *Proc. of MICCAI'99*, volume 1679 of *LNCS*, pages 589 – 596, Cambridge, UK. Springer.

Duchon, J. (1977). *Constructive Theory of Functions of Several Variables*, chapter Spline minimizing rotation-invariant semi-norms in Sobolev spaces, pages 85 – 100. Springer-Verlag.

Feldmar, J., Declerck, J., Malandain, G., and Ayache, N. (1997). Extension of the ICP Algorithm to Nonrigid Intensity-Based Registration of 3D Volumes. *Comp. Vision and Image Understanding*, 66(2):193–206.

Ferrant, M., Warfield, S. K., Nabavi, A., Jolesz, F. A., and Kikinis, R. (2000). Registration of 3D Intra-operative MRI of the Brain using a Finite Element Biomechanical Model. In *Proc. of MICCAI'00*, volume 1935 of *LNCS*, pages 19 – 28, Pittsburgh, USA. Springer.

Fischer, B. and Modersitzki, J. (1999). Fast Inversion of Matrices arising in Image Processing. *Numerical Algorithms*, 22(1):1–11.

Fornefett, M., Rohr, K., and Stiehl, H. S. (1999). Elastic Registration of Medical Images Using Radial Basis Functions with Compact Support. In *Proc. of CVPR'99*, pages 402 – 407, Fort Collins, USA.

Gabrani, M. and Tretiak, O. J. (1999). Surface-Based Matching using Elastic Transformations. *Pattern Recognition*, 32(1):87 – 97.

Gaens, T., Maes, F., Vandermeulen, D., and Suetens, P. (2000). Non-Rigid Multimodal Image Registration using Mutual Information. In *Proc. of MICCAI'98*, volume 1496 of *LNCS*, pages 1099 – 1106, Cambridge, USA. Springer.

Gee, J. C. (1999). On Matching Brain Volumes. *Pattern Recognition*, 32:99 – 111.

Guan, W.-G., Xie, L., and Ma, S.-D. (1998). Deformable Registration of Digital Images. *J. of Comp. Science and Technology*, 13(3):246 – 260.

Guimond, A., Roche, A., Ayache, N., and Meunier, J. (2001). Three-Dimensional Multimodal Brain Warping Using the Demons Algorithm and Adaptive Intensity Corrections. *IEEE Trans. on Medical Imaging*, 20(1):58–69.

Hata, N., Dohi, T., Warfield, S., Wells, W., Kikinis, R., and Jolesz, F. A. (1998). Multimodality Deformable Registration of Pre- and Intraoperative Images for MRI-Guided Brain Surgery. In *Proc. of MICCAI'98*, number 1496 in LNCS, pages 1067 – 1074, Cambridge, USA. Springer.

Hayton, P. M., Brady, M., Smith, S. M., and Moore, N. (1999). A Non-Rigid Registration Algorithm for Dynamic Breast MR Images. *Artificial Intelligence*, 114:124–156.

Hellier, P. and Barillot, C. (2000). Coupling Dense and Landmark-Based approaches for Non Rigid Registration. Technical Report 1368, IRISA.

Hellier, P., Barillot, C., Mémin, E., and Pérez, P. (1999). Medical Image Registration with Robust Multigrid Techniques. In *Proc. of MICCAI'99*, volume 1679 of *LNCS*, pages 680 – 687, Cambridge, UK. Springer.

Holden, M., Hill, D. L. G., Denton, E. R. E., Jarosz, J. M., Cox, T. C. S., Rohlfing, T., Goodey, J., and Hawkes, D. J. (2000). Voxel Similarity Measures for 3-D Serial MR Brain Image Registration. *IEEE Trans. on Medical Imaging*, 19(2):94 – 102.

Horn, B. K. P. and Schunk, B. G. (1981). Determining Optical Flow. *Artificial Intelligence*, 17:181–203.

King, A. P., Blackall, J. M., Penney, G. P., Edwards, P. J., Hill, D. L. G., and Hawkes, D. J. (2000). Bayesian Estimation of Intra-operative Deformation for Image-Guided Surgery using 3D Ultrasound. In *Proc. of MICCAI'00*, volume 1935 of *LNCS*, pages 588 – 597, Pittsburgh, USA. Springer.

Krissian, K., Malandain, G., and Ayache, N. (1997). Directional Anisotropic Diffusion Applied to Segmentation of Vessels in 3D Images. In *Scale-Space theory in Computer Vision*, pages 345–348, Utrecht, The Netherlands.

Kyriacou, S. K. and Davatzikos, C. (1998). A Biomechanical Model of Soft Tissue Deformation, with Applications to Non-Rigid Registration of Brain Images with Tumor Pathology. In *Proc. of MICCAI'98*, pages 531 – 538, Cambridge, USA. Springer.

Lau, Y. H., Braun, M., and Hutton, B. F. (1999). Non-Rigid 3D Image Registration Using Regionally Constrained Matching and the Correlation Ratio. In *Proc. of WBIR'99*, pages 137 – 148, Bled, Slovenia.

Lester, H. and Arridge, S. R. (1999). A Survey of Hierarchical Non-Linear Medical Image Registration. *Pattern Recognition*, 32:129 – 149.

Lester, H., Arridge, S. R., Jansons, K. M., Lemieux, L., Hajnal, J. V., and Oatridge, A. (1999). Non-Linear Registration with the Variable Viscosity Fluid Algorithm. In *Proc. of IPMI'99*, volume 1613 of *LNCS*, pages 238 – 251, Visegrád, Hungary. Springer.

Maintz, J. B. A., Meijering, E. H. W., and Viergever, M. A. (1998). General Multimodal Elastic Registration based on Mutual Information. *Image Processing*.

Maintz, J. B. A. and Viergever, M. A. (1998). A Survey of Medical Image Registration. *Medical Image Analysis*, 2(1):1–36.

Mangin, J.-F. (2000). Entropy minimization for automatic correction of intensity nonuniformity. In *Proc. of MMBIA'00*, pages 162 – 169.

Meyer, C., Boes, J., Kim, B., and Bland, P. (1998). Evaluation of Control Point Selection in Automatic, Mutual Information Driven 3D Warping. In *Proc. of MICCAI'98*, number 1496 in LNCS, pages 944 – 951, Cambridge, USA. Springer.

Meyer, C. R., Boes, J. L., Kim, B., and Bland, P. H. (1999). Probabilistic Brain Atlas Construction: Thin Plate Spline Warping via Maximization of Mutual Information. In *Proc. of MICCAI'99*, volume 1679 of *LNCS*, pages 631 – 637, Cambridge, UK. Springer.

Miller, M. and Younes, L. (1999). Group actions, homeomorphisms, and matching: a general framework. In *Proc. of IEEE Workshop. on Statistical and Comp. Theories of Vision*, Fort Collins, USA.

Miller, M. I., Christensen, G. E., Amit, Y., and Grenander, U. (1993). Mathematical Textbook of Deformable Neuroanatomies. *Proc. of the National Academy of Science*, 90(24):11944–11948.

Musse, O., Heitz, F., and Armspach, J.-P. (1999). 3D Deformable Image Matching using Multiscale Minimization of Global Energy Functions. In *Proc. of CVPR'99*, Fort Collins, USA.

Ourselin, S., Roche, A., Subsol, G., Pennec, X., and Ayache, N. (2001). Reconstructing a 3D Structure from Serial Histological Sections. *Image and Vision Comp.*, 19(1-2):25 – 31.

Peckar, W., Schnörr, C., Rohr, K., and Stiehl, H. S. (1999). Parameter-Free Elastic Deformation Approach for 2D and 3D Registration Using Prescribed Displacements. *J. of Mathematical Imaging and Vision*, 10:143 – 162.

Pennec, X., Cachier, P., and Ayache, N. (1999). Understanding the "Demons" Algorithm: 3D Non-Rigid Registration by Gradient Descent. In *Proc. of MICCAI'99*, volume 1679 of *LNCS*, pages 597 – 605, Cambridge, UK. Springer.

Pennec, X., Cachier, P., and Ayache, N. (2001). Tracking brain deformations in time-sequences of 3d us images. *Pattern Recognition Letters - Special Issue on Ultrasonic Image Processing and Analysis*. in press.

Pennec, X., Guttman, C. R. G., and Thirion, J.-P. (1998). Feature-Based Registration of Medical Images: Estimation and Validation of the Pose Accuracy. In *Proc. of MICCAI'98*, number 1496 in LNCS, pages 1107 – 1114, Cambridge, USA. Springer.

Pluim, J. P. W., Maintz, J. B. A., and Viergever, M. A. (2000). Interpolation Artefacts in Mutual Information-Based Image Registration. *Comp. Vision and Image Understanding*, 77(2):211 – 232.

Rangarajan, A., Chui, H., and Duncan, J. S. (1999). Rigid Point Feature Registration using Mutual Information. *Medical Image Analysis*, 3(4).

Roche, A., Malandain, G., and Ayache, N. (2000). Unifying Maximum Likelihood Approaches in Medical Image Registration. *Int. J. of Imaging Systems and Technology*, 11:71–80.

Rohr, K., Fornefett, M., and Stiehl, H. S. (1999). Approximating Thin-Plate Splines for Elastic Registration: Integration of Landmark Errors and Orientation Attributes. In *Proc. of IPMI'99*, volume 1613 of *LNCS*, pages 252 – 265, Visegrád, Hungary. Springer.

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

Ruiz-Alzola, J., Westin, C.-F., Warfield, S. K., Nabavi, A., and Kikinis, R. (2000). Nonrigid Registration of 3D Scalar, Vector and Tensor Medical Data. In *Proc. of MICCAI'00*, volume 1935 of *LNCS*, Pittsburgh, USA. Springer. to appear in Medical Image Analysis.

Schormann, T., Henn, S., and Zilles, K. (1996). A New Approach to Fast Elastic Alignment with Applications to Human Brains. In *Proc. of VBC'96*, volume 1131 of *LNCS*, pages 337 – 342, Hamburg, Germany. Springer.

Strintzis, M. G. and Kokkinidis, I. (1997). Maximum Likelihood Motion Estimation in Ultrasound Image Sequences. *IEEE Signal Processing Letters*, 4(6).

Studholme, C., Hill, D. L. G., and Hawkes, D. J. (1999). An Overlap Invariant Entropy Measure of 3D Medical Image Alignment. *Pattern Recognition*, 32(1):71 – 86.

Subsol, G., Thirion, J. P., and Ayache, N. (1998). A General Scheme for Automatically Building 3D Morphometric Anatomical Atlases: application to a Skull Atlas. *Medical Image Analysis*, 2(1):37–60.

Suter, D. and Chen, F. (2000). Left Ventricular Motion Reconstruction Based on Elastic Vector Splines. *IEEE Trans. on Medical Imaging*, 19(4):295 – 305.

Szeliski, R. and Coughlan, J. (1997). Spline-Based Image Registration. *Int. J. of Comp. Vision*, 22(3):199–218.

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

Thompson, P. and Toga, A. W. (1996). A Surface-Based Technique for Warping 3-Dimensional Images of the Brain. *IEEE Trans. on Medical Imaging*, 15(4):402–417.

Trouvé, A. (1998). Diffeomorphisms Groups and Pattern Matching in Image Analysis. *Int. J. of Comp. Vision*, 28(3):213–221.

Vemuri, B. C., Huang, S., Sahni, S., Leonard, C. M., Mohr, C., Gilmore, R., and Fitzsimmons, J. (1998). An Efficient Motion Estimator with Application to Medical Image Registration. *Medical Image Analysis*, 2(1):79–98.

Wahba, G. (1990). *Spline Models for Observational Data*, volume 59 of *CRMS-NSE Regional Conf. Series in Applied Mathematics*. Society for Industrial and Applied Mathematics.

Wells, W. M., Viola, P., Atsumi, H., Nakajima, S., and Kikinis, R. (1996). Multi-Modal Volume Registration by Maximization of Mutal Information. *Medical Image Analysis*, 1(1):35–51.

West, J., Fitzpatrick, J., Wang, M., Dawant, B., Maurer, C., Kessler, R., and Maciunas, R. (1999). Retrospective Intermodality Registration Techniques for Images of the Head: Surface-based versus Volume-based. *IEEE Trans. on Medical Imaging*, 18(2):144 – 150.

Wu, Y.-T., Kanade, T., Li, C.-C., and Cohn, J. (2000). Image Registration using Wavelet-Based Motion Model. *Int. J. of Comp. Vision*, 38(2):129 – 152.

Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - B.P. 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Lorraine : Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - B.P. 101 - 54602 Villers lès Nancy Cedex (France)
Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)
Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38330 Montbonnot St Martin (France)
Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - B.P. 105 - 78153 Le Chesnay Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, B.P. 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399