



HAL
open science

Multimodal Brain Warping Using the Demons Algorithm and Adaptative Intensity Corrections

Alexandre Guimond, Alexis Roche, Nicholas Ayache, Jean Meunier

► **To cite this version:**

Alexandre Guimond, Alexis Roche, Nicholas Ayache, Jean Meunier. Multimodal Brain Warping Using the Demons Algorithm and Adaptative Intensity Corrections. RR-3796, INRIA. 1999. inria-00072863

HAL Id: inria-00072863

<https://inria.hal.science/inria-00072863>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

***Multimodal Brain Warping Using the Demons
Algorithm and Adaptative Intensity Corrections***

Alexandre Guimond — Alexis Roche — Nicholas Ayache — Jean Meunier

N° 3796

Novembre 1999

THÈME 3



*rapport
de recherche*

Multimodal Brain Warping Using the Demons Algorithm and Adaptative Intensity Corrections

Alexandre Guimond ^{*} , Alexis Roche , Nicholas Ayache , Jean Meunier [†]

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

Rapport de recherche n° 3796 — Novembre 1999 — 36 pages

Abstract: This paper presents an original method for three-dimensional elastic registration of multimodal images. We propose to make use of a scheme that iterates between correcting for intensity differences between images and performing standard monomodal registration. The core of our contribution resides in providing a method that finds the transformation that maps the intensities of one image to those of another. It makes the assumption that there are at most two functional dependences between the intensities of structures present in the images to register, and relies on robust estimation techniques to evaluate these functions. We provide results showing successful registration between several imaging modalities involving segmentations, T1 magnetic resonance (MR), T2 MR, proton density (PD) MR and computed tomography (CT). We also argue that our intensity modeling may be more appropriate than mutual information (MI) in the context of evaluating high-dimensional deformations, as it puts more constraints on the parameters to be estimated and thus permits a better search of the parameter space.

Key-words: Multimodality, Elastic registration, Intensity correction, Robust estimation, Medical imaging

This work was held in collaboration with Département d'Informatique et recherche opérationnelle, Université de Montréal, CP 6128 succ Centre-Ville, Montréal QC H3C 3J7, Canada

^{*} INRIA and Université de Montréal

[†] Université de Montréal

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

Déformation d'images cérébrales inter-modalités par algorithme des «démons» et correction d'intensité adaptative

Résumé : Cet article décrit une nouvelle méthode pour le recalage élastique tridimensionnel d'images de modalités différentes. Le schéma que nous proposons alterne entre une étape de normalisation des intensités et un recalage intra-modal standard. Notre contribution principale réside dans le calcul de la transformation reliant les intensités de deux images. La technique que nous proposons est fondée sur l'hypothèse qu'il existe une ou deux dépendances fonctionnelles entre les intensités des structures présentes dans les images à recalier. Des techniques d'estimation robuste sont mises en œuvre pour évaluer ces fonctions. Nous présentons des résultats de recalage entre différentes modalités d'images médicales : des segmentations, des images par résonance magnétique acquises selon les protocoles T1, T2, PD, et des images scanner. Nous montrons notamment que, par rapport aux méthodes de recalage fondées sur la maximisation de l'information mutuelle, celle que nous proposons contraint davantage la relation entre les intensités des images. De ce fait, elle pourrait être plus appropriée lorsqu'il s'agit d'évaluer des déformations géométriques libres.

Mots-clés : Recalage inter-modalités, recalage élastique, correction d'intensité, estimation robuste, imagerie médicale.

Contents

1	Introduction	4
2	Method	5
2.1	Modeling the Geometrical Transformation	5
2.1.1	Relation with SSD Minimization	6
2.1.2	Relation with Optical Flow	6
2.1.3	Relation with the Demons Algorithm	7
2.1.4	Relevance of the Resulting Transformation	7
2.2	Modeling the Intensity Transformation	8
2.2.1	Monofunctional Dependence Assumption	8
2.2.2	Bifunctional Dependence Assumption	10
2.3	Combining the Intensity and Geometrical Transformations	12
2.4	Mutual information	12
3	Data	13
4	Results and Discussion	13
4.1	Monofunctional Dependence	14
4.2	Bifunctional Dependence	19
4.3	Mutual Information	26
4.4	Displacement field comparison	31
5	Conclusion	32

1 Introduction

Over the last decade, automatic registration techniques of medical images of the head have been developed following two main trends: 1) registration of multimodal images using low degree transformations (rigid or affine), and 2) registration of monomodal images using high-dimensional volumetric maps (elastic or fluid deformations). The first category mainly addresses the fusion of complementary information obtained from different imaging modalities. The second category's predominant purpose is the evaluation of either the anatomical evolution process present in a particular subject or of anatomical variations between different subjects.

These two trends have evolved separately mainly because the combined problem of identifying complex intensity correspondences along with a high-dimensional geometrical transformation defines a search space arduous to traverse. Recently, three groups have imposed different constraints on the search space, enabling them to develop automatic multimodal non-affine registration techniques. All three methods make use of block matching techniques to evaluate local translations. Two of them use mutual information (MI) [30, 18] as a similarity measure and the other employs the correlation ratio [23].

The major difficulty in using MI as a similarity measure for registration is to compute the conditional probabilities of one image's intensities with respect to those of the other. To do so, Maintz *et al.* [19] proposed to use conditional probabilities after rigid matching of the images as an estimate of the real conditional probabilities after local transformations. Hence, the probabilities are evaluated only once before fluid registration. However, Gaens *et al.* [12] argued that the assumption that probabilities computed after affine registration are good approximations of the same probabilities after fluid matching is unsuitable. They also proposed a method in which local displacements are found so that the global MI increases at each iteration, permitting incremental changes of the probabilities during registration. Their method necessitates the computation of conditional probabilities over the whole image for every voxel displacement. To alleviate themselves from such computations owing to the fact that MI requires many samples to estimate probabilities, Lau *et al.* [17] have chosen a different similarity measure. Due to the robustness of the correlation ratio with regards to sparse data [23], they employed it to assess the similarity of neighboring blocks. Hence no global computation is required when moving subregions of the image.

Our method distinguishes itself by looking at the problem from a different angle. In the last years, our group has had some success with monomodal image registration using the demons method [27, 28], an optical flow variant when dealing with monomodal volumetric images. If we were able to model the imaging processes that created the images to register, and assuming these processes are invertible, one could transform one of the images so that they are both represented in the same modality. Then, we could use our monomodal registration algorithm to register them. We have thus developed a completely automatic method to transform the different structures intensities in one image so that they match the intensities of the corresponding structures in another image, and this without resorting to any segmentation method.

The rationale behind our formulation is that there is a functional relationship between the intensity of a majority of structures when imaged with different modalities. This assumption is partly justified by the fact that the Woods criterion [32] as well as the correlation ratio [23], which evaluate a functional dependence

between the images to match, have been used with success in the past, and sometimes lead to better results than MI, which assumes a more general constraint [22, 21].

The idea of estimating an intensity transformation during registration is not new in itself. For example, Feldmar *et al.* [11] as well as Barber [1] have both published methods in which intensity corrections are proposed. These methods restrict themselves to estimating affine intensity transformations in a monomodal registration context. We propose here a registration model based on one or two higher degree polynomials found using a robust regression technique to enable the registration of images from different modalities.

The remaining sections of this paper are organized in the following manner. First, we detail our multimodal elastic registration method. We then describe what kind of images were used to test our method and how they were acquired. Next, results obtained by registering different images obtained from several modalities are presented and discussed. We conclude this paper with a brief discussion on future research tracks.

2 Method

Our registration algorithm is iterative and each iteration consists of two parts. The first one transforms the intensities of anatomical structures of a source image S so that they match the corresponding structures intensities of a target image T . The second part regards the registration of S (after intensity transformation) with T using our elastic registration algorithm.

In the following, we first describe the three-dimensional geometrical transformation model and then the intensity transformation model. We believe this ordering is more convenient since it is easier to see what result must provide the intensity transformation once the geometrical transformation procedure is clarified. However, we wish to make clear to the reader that *each iteration* of our registration method proceeds first by estimating the intensity transformation and then the geometrical transformation.

2.1 Modeling the Geometrical Transformation

Many algorithms have been developed that deform one brain so its shape matches that of another [29]. The procedure used in the following work was influenced by a variety of methods, primarily the demons algorithm [27, 28]. It finds the displacement $v(x)$ for each voxel x of T so it matches the corresponding anatomical location in S . The solution is found using the following iterative scheme,

$$v_{n+1}(x) = G_\sigma \otimes \left(v_n + \frac{S \circ h_n(x) - T(x)}{\|\nabla(S \circ h_n)(x)\|^2 + [S \circ h_n(x) - T(x)]^2} \nabla(S \circ h_n)(x) \right), \quad (1)$$

where G_σ is a Gaussian filter with a variance of σ^2 , \otimes denotes the convolution, \circ denotes the composition and the transformation $h(x)$ is related to the displacement by $h(x) = x + v(x)$. As is common with registration methods, we also make use of multilevel techniques to accelerate convergence. Details about the number of levels and iterations as well as filter implementation issues are addressed in Section 4. We here show how our method can be related to other registration methods, notably the minimization of the sum of squared difference (SSD) criterion, optical flow and the demons algorithm.

2.1.1 Relation with SSD Minimization

In this framework, we find the transformation h that minimizes the sum of squared differences between the transformed source image and the target image. The SSD between the two images for a given transformation h applied to the source is defined as

$$SSD(h) = \frac{1}{2} \sum_{x=1}^N [S \circ h(x) - T(x)]^2. \quad (2)$$

Minimization of Equation (2) is performed using a gradient descent algorithm. Thus, differentiating the above equation we get

$$\nabla SSD(h) = -[S \circ h(x) - T(x)] \nabla(S \circ h)(x).$$

The iterative scheme is then of the form

$$h_{n+1} = h_n + \alpha [S \circ h_n(x) - T(x)] \nabla(S \circ h_n)(x),$$

where α is the step length. This last equation implies

$$v_{n+1} = v_n + \alpha [S \circ h_n(x) - T(x)] \nabla(S \circ h_n)(x). \quad (3)$$

If we set α to a constant value, this method corresponds to a first order gradient descent algorithm. By comparing Equation (3) to Equation (1), one sees that our method sets

$$\alpha = \frac{1}{\|\nabla(S \circ h_n)(x)\|^2 + [T(x) - S \circ h_n(x)]^2} \quad (4)$$

and applies a Gaussian filter to provide a smooth displacement field. Cachier *et al.* [6, 20] have shown that using Equation (4) closely relates Equation (1) with a second order gradient descent of the SSD criterion, in which each iteration n sets h_{n+1} to the minimum of the SSD quadratic approximation at h_n . We refer the reader to these articles for a more technical discussion on this subject as well as for the formula that corresponds to a true second order gradient descent.

2.1.2 Relation with Optical Flow

T and S are considered as successive time samples of an image sequence represented by $I(x, t)$, where $x = (x_1, x_2, x_3)$ is a voxel position in the image and t is time. The displacements are computed by constraining the brightness of brain structures to be constant in time so that

$$\frac{dI(x, t)}{dt} = 0. \quad (5)$$

It is well known that Equation (5) is not sufficient to provide a unique displacement for each voxel. In fact, this constraint leads to

$$f(x) = -\frac{\partial I(x, t)/\partial t}{\|\nabla_x I(x, t)\|^2} \nabla_x I(x, t), \quad (6)$$

which is the component of the displacement in the direction of the brightness gradient [15], that is

$$\nabla_x I(x, t) = \left(\frac{\partial I(x, t)}{\partial x_1}, \frac{\partial I(x, t)}{\partial x_2}, \frac{\partial I(x, t)}{\partial x_3} \right)^T.$$

Other constraints need to be added to Equation (5) to obtain the displacement components in other directions. Many methods have been proposed to fill this purpose and thus regularize the resulting vector field [2]. One that can be computed very efficiently was proposed by Thirion [28] in his description of the demons registration method using a complete grid of demons. It consists of smoothing each dimension of the vector field with a Gaussian filter G_σ . He also proposed to add $[\partial I(x, t)/\partial t]^2$ to the denominator of Equation (6) for numerical stability when $\nabla_x I(x, t)$ is close to zero, a term which serves the same purpose as α^2 in the original optical flow formulation of Horn and Schunck [15]. As is presented by Bro-Nielsen and Gramkow [5], this kind of regularization approximates a linear elasticity transformation model.

With this in mind, the displacement that maps a voxel position in T to its position in S is found using an iterative method,

$$v_{n+1}(x) = G_\sigma \otimes \left(v_n - \frac{\partial I(x, t)/\partial t}{\|\nabla_x I(x, t)\|^2 + [\partial I(x, t)/\partial t]^2} \nabla_x I(x, t) \right). \quad (7)$$

Spatial derivatives may be computed in several ways [15, 4, 26]. We have observed from practical experience that our method performs best when they are computed from the resampled source image of the current iteration. As shown in Section 2.1.1, this is in agreement with the SSD minimization. Temporal derivatives are obtained by subtracting the target images from the resampled source image of the current iteration. These considerations relate Equation (7) to Equation (1). The reader should note that the major difference between this method and other optical flow strategies is that regularization is performed *after* the calculation of the displacements in the gradient direction instead of using an explicit regularization term in a minimization framework.

2.1.3 Relation with the Demons Algorithm

Our algorithm actually is a small variation of the demons method [27, 28] using a complete grid of demons, itself closely related to optical flow as described in the previous section. The demons algorithm finds the displacements using the following formula,

$$v_{n+1}(x) = G_\sigma \otimes \left(v_n + \frac{S \circ h_n(x) - T(x)}{\|\nabla T(x)\|^2 + [S \circ h_n(x) - T(x)]^2} \nabla T(x) \right).$$

As can be seen from the last equation, the only difference from our formulation (Equation (1)) and the demons method is that derivatives are computed on the resampled source image of the current iteration. This modification was performed following the observations on the minimization of the SSD criterion.

2.1.4 Relevance of the Resulting Transformation

In the case of inter-subject non-affine registration, quantifying the accuracy of a method is difficult. One could deform an image I into I' using a known displacement field v , register I with I' and compare the result of the registration with v , but this comparison is biased by the way v is generated. For example, since our algorithm produces a smooth displacement field, if the vectors of v were to be generated randomly, the method is expected

to perform poorly. Another method is to place landmarks in the images to register and evaluate differences between landmarks after registration (See for example [31]). No such study has been performed using our algorithm.

An approach which is a generalization of the previous one is to compare manual segmentations and automatic segmentations obtained with segmentation propagation. Using this method, an expert manually segments several structures in a set of images. One image is then selected as the target and each image of the set is registered with it. The manual segmentations of each image are then deformed using the corresponding displacement fields. At this point, all segmentations should have the same shape as the ones of the target. They can then be compared, the manual result of the target serving as ground truth (See for example [13]). We believe this approach may be better suited to evaluate the quality of high dimensional transformations such as the ones obtained using the previously described algorithm (typically $3 \times 200^3 = 24 \times 10^6$ degrees of freedom). Such a study has been performed previously by Dawant *et al.* [10] for the demons algorithm and reports that contours obtained from manual and automatic segmentations are virtually indistinguishable.

Concerning the physical relevance of the transformation, it should be pointed out that our algorithm does not explicitly make sure that the resulting transformation preserves the topology of S . In our experience on high resolution data, this problem does not seem to occur when using a sigma of 1 voxel to define the Gaussian filter used for smoothing the vector field. Still, this is a characteristic we would like to do without. Better regularization methods, such as proposed by Christensen *et al.* [7] or Bro-Nielsen and Gramkow [5] should provide more satisfying transformations, at least regarding singularities.

2.2 Modeling the Intensity Transformation

Previous to each iteration of the geometrical transformation an intensity correction is performed on S so that the intensities of its structures matches those in T , a requirement for the use of Equation 1. The intensity correction process starts by defining the set C of intensity couples from corresponding voxels of T and of the current resampled source image $S \circ h$, which will be designated by S in this section for simplicity. Hence, the set C is defined as

$$C = \left\{ (S(i), T(i)); 1 \leq i \leq N \right\},$$

where N is the number of voxels in the images. $S(i)$ and $T(i)$ correspond to the intensity value of the i^{th} voxel of S and T respectively when adopting the customary convention of considering images as one-dimensional arrays. From there, we show how to perform intensity correction if one or two functional dependences can be assumed between the structures intensities.

2.2.1 Monofunctional Dependence Assumption

Our goal is to model the transformation that characterizes the mapping from voxel intensities in S to those in T , knowing that some elements of C are erroneous, i.e. that would not be present in C if S and T were perfectly matched. If we can assume a monofunctional dependence of the intensities of T with regards to the those of S

as well as additive stationary Gaussian white noise η on the intensity values of T , then we can adopt the model

$$T(i) = f(S(i)) + \eta(i), \quad (8)$$

where f is an unknown function to be estimated. This is exactly the model employed by Roche *et al.* [22, 21] which leads to the correlation ratio as the measure to be maximized for registration. In that approach, for a given transformation, the authors seek the function that best describes T in terms of S . They show that in a maximum likelihood context, the intensity function \hat{f} that best approximates f is a least squares (LS) fit of T in terms of S .

The major difference between our respective problems is that we seek a high-dimensional geometrical transformation. As opposed to affine registration where the transformation is governed by the majority of good matches, we have seen in Section 2.1 that using the elastic registration model, displacements are found using mainly local information (i.e. gradients, local averages, etc.). Hence, we can not expect good displacements in one structure to correct for bad ones in another; we have to make certain each voxel is moved properly during each iteration. For this, since the geometrical transformation is found using intensity similarity, the most precise intensity transformation is required. Consequently, instead of performing a standard least squares regression, we have opted for a robust linear regression estimator which will remove outlying elements of C during the estimation of the intensity transformation. To estimate f we use the least trimmed squares (LTS) method followed by a binary reweighted least squares (RLS) estimation [25]. The combination of these two methods provides a very robust regression technique with outlier detection, while ensuring that a maximum of pertinent points are used for the final estimation.

Least Trimmed Squares Computation For our particular problem, we will constrain ourselves the estimation of a polynomial function from the elements in C . We can then relate the intensity correspondences with a set of equations of the form

$$T(i) = \theta_0 + \theta_1 S(i) + \theta_2 S(i)^2 + \dots,$$

where p is the degree of the polynomial function and $\boldsymbol{\theta} = [\theta_0, \dots, \theta_p]$ needs to be estimated. A regression estimator will provide a $\hat{\boldsymbol{\theta}} = [\hat{\theta}_0, \dots, \hat{\theta}_p]$ which can be used to predict the value of $T(i)$ from $S(i)$,

$$\hat{T}(i) = \hat{\theta}_0 + \hat{\theta}_1 S(i) + \hat{\theta}_2 S(i)^2 + \dots + \hat{\theta}_p S(i)^p,$$

as well as the residual errors,

$$r(i) = T(i) - \hat{T}(i).$$

A popular method to obtain $\hat{\boldsymbol{\theta}}$ is to minimize the sum of squared residual errors,

$$\min_{\hat{\boldsymbol{\theta}}} \sum_{i=1}^N r(i)^2,$$

which leads to the standard LS solution. It is found by solving a linear system using the Singular Value Decomposition (SVD) method. See [21] for a more detailed study. This method is known to be very sensitive

to outliers and thus is expected to provide a poor estimate of the monofunctional mapping from S to T . The LTS method solves this problem by minimizing the same sum on a subset of all residual errors, thus rejecting large ones corresponding to outliers,

$$\min_{\hat{\theta}} \sum_{i=1}^h \rho(i),$$

where $\rho(i)$ is the i^{th} smallest value of the set $\{r(1)^2, \dots, r(N)^2\}$. This corresponds to a standard LS on the c values that best approximates the function we are looking for. Essentially, c/N represents the percentage of “good” points in C and must be at least $0.50 \times N$. A lesser value would allow to estimate parameters that model a minority of point which could then all be outliers. The value of c will vary according to the modalities used during registration. Assigning actual values to it is postponed to Section 4.

Our method for LTS minimization is a simple iterative technique. First, we randomly pick c/N points from C . We then iterate between calculating $\hat{\theta}$ using the standard LS technique on the selected points and choosing the h/N closest points from C . This process is carried until convergence, usually requiring less than 5 iterations and provides adequate results for our purposes. Recently, Rousseeuw and Van Driessen [24] proved that this method reduces the error on $\hat{\theta}$ at each iteration and finds a local minimum, although not guarantying the global minimum. Since we assume a good global registration previous to the estimation process, that might not make a large difference. Still, in the same paper, the authors also proposed a new efficient algorithm to find a good approximate solution of the LTS minimization. This strategy will be looked into in future implementations.

Reweighted Least Squares Computation Once $\hat{\theta}$ is obtained using LTS, we can compute the standard error σ of our points with respect to our estimate. Of course, this value will also be an estimate corresponding to

$$\hat{\sigma} = \sqrt{\frac{1}{KN} \sum_{i=1}^c \rho(i)}, \quad K = \frac{\int_{-\alpha}^{\alpha} x^2 g(x) dx}{\int_{-\alpha}^{\alpha} g(x) dx} = \frac{\int_{-\alpha}^{\alpha} x^2 g(x) dx}{c/N}, \quad (9)$$

where $g(x)$ is the Gaussian distribution $N(0, 1)$ and α is the $(0.5 + c/2N)^{\text{th}}$ quantile of $g(x)$. In Equation (9), K is a normalization factor introduced because $\sqrt{1/N \sum_{i=1}^c \rho(i)}$ is a consistent estimator of σ when the $\rho(i)$ are distributed like $N(0, \sigma^2)$. This is not the case in LTS regression since we only choose the c/N smallest $\rho(i)$.

Using this average deviation of the points from the estimates, we then perform an RLS regression by finding a new $\hat{\theta}$ that minimizes the sum of squared residual for all points within $3\hat{\sigma}$ of the previous estimate,

$$\min_{\hat{\theta}} \sum_{i=1}^N w_i \rho(i), \quad \text{where } w_i = \begin{cases} 1 & \text{if } \rho(i) \leq 3\hat{\sigma}, \\ 0 & \text{otherwise.} \end{cases}$$

This optimizes the number of points used to compute $\hat{\theta}$ by considering all the points that relate well to the LTS estimate, not only the best $c/N \times 100\%$.

2.2.2 Bifunctional Dependence Assumption

Functional dependence as expressed in Equation (8) implicitly assumes that two structures having similar intensity ranges in S should also have similar intensity ranges in T . With some combinations of multimodal

images, this is a crude approximation. For example, ventricles and bones generally give similar response values in a MR T1 weighted image while they appear with very distinct values in a CT scan. Conversely, white and black matter are well contrasted in a T1 image while corresponding to similar intensities in a CT.

To circumvent this difficulty, we have developed a strategy that enables the mapping of an intensity value in S to not only one, but two possible intensity values in T . This method is a natural extension of the previous method. Instead of computing a single function that maps the intensities of S to those of T , two functions are estimated and the mapping becomes a weighted sum of these two functions.

We start with the assumption that if a point has an intensity s in S , the corresponding point in T has an intensity t that is normally distributed around two possible values depending on s , $f_{\theta}(s)$ and $f_{\psi}(s)$. In statistical terms, this means that given s , t is drawn from a mixture of Gaussian distribution,

$$P(t|s) = \pi_1(s)N(f_{\theta}(s), \sigma^2) + \pi_2(s)N(f_{\psi}(s), \sigma^2), \quad (10)$$

where $\pi_1(s)$ and $\pi_2(s) = 1 - \pi_1(s)$ are mixing proportions that depend on the intensity in the source image, and σ^2 represents the variance of the noise in the target image. Consistently with the functional case, we will restrict ourselves to polynomial intensity functions, i.e. $f_{\theta}(s) = \theta_0 + \theta_1 s + \theta_2 s^2 + \dots + \theta_p s^p$, and $f_{\psi}(s) = \psi_0 + \psi_1 s + \psi_2 s^2 + \dots + \psi_p s^p$.

An intuitive way to interpret this modeling is to state that for any voxel, there is a binary “selector” variable $\epsilon = \{1, 2\}$ that would tell us, if it was observed, which of the two functions f_{θ} or f_{ψ} actually serves to map s to t . Without knowledge of ϵ , the best intensity correction to apply to S (in the minimum variance sense) is seen to be a weighted sum of the two functions,

$$f(s, t) = P(\epsilon = 1|s, t)f_{\theta}(s) + P(\epsilon = 2|s, t)f_{\psi}(s), \quad (11)$$

in which the weights correspond to the probability that the point be mapped according to either the first or the second function. Applying Bayes’ law, we find that for $\epsilon = \{1, 2\}$:

$$P(\epsilon|s, t) = \frac{P(\epsilon|s)P(t|\epsilon, s)}{P(t|s)},$$

and thus, using the fact that $P(\epsilon|s) = \pi_{\epsilon}(s)$ and $P(t|\epsilon, s) = G_{\sigma}(t - f_{\epsilon}(s))$, the weights are determined by

$$P(\epsilon|s, t) = \frac{\pi_{\epsilon}(s) G_{\sigma}(t - f_{\epsilon}(s))}{\pi_1(s) G_{\sigma}(t - f_{\theta}(s)) + \pi_2(s) G_{\sigma}(t - f_{\psi}(s))}, \quad (12)$$

where it should be clear from the context that $f_{\epsilon} \equiv f_{\theta}$ if $\epsilon = 1$, and $f_{\epsilon} \equiv f_{\psi}$ if $\epsilon = 2$.

In order to compute the intensity correction (Equation (11)), we now need to identify the parameters of our model, i.e. the polynomial coefficients θ and ψ , as well as the mixing proportions $\pi_1(s)$ and $\pi_2(s)$ and the variance σ^2 . For this, we employ an ad hoc strategy that proceeds as follows.

First, θ is estimated using the method described in section 2.2.1. The points not used to compute θ , in a number between 0 and $N - c$, are used to estimate ψ still using the same method. Note that if this number is less than $10 \times p$, p being the polynomial degree, functional dependence is assumed and we fall back to the method described in the previous section.

This provides a natural estimation of the “selector” variable for each voxel: the n_1 points that were used to build f_{θ} are likely to correspond to $\epsilon = 1$, while the n_2 points used to build f_{ψ} are likely to correspond to $\epsilon = 2$.

Finally, the points that are rejected while estimating ψ are considered as bad intensity matches. A natural estimator for the variance σ^2 is then

$$\hat{\sigma}^2 = \frac{n_1}{n_1 + n_2} \hat{\sigma}_1^2 + \frac{n_2}{n_1 + n_2} \hat{\sigma}_2^2,$$

where $\hat{\sigma}_1^2$ and $\hat{\sigma}_2^2$ are the variances found respectively for f_θ and f_ψ during the RLS regression (See Section 2.2.1.). Similarly, the mixing proportions are computed according to

$$\hat{\pi}_\epsilon(s) = \frac{n_\epsilon(s)}{n_1(s) + n_2(s)}, \quad \epsilon = \{1, 2\},$$

in which $n_\epsilon(s)$ is the number of voxels having an intensity s and used to build the function f_ϵ . Notice that in the case where $n_1(s) = n_2(s) = 0$ (i.e. no voxel corresponding to the intensity s has been taken into account in the computation of f_θ or f_ψ), then we arbitrarily set the mixing proportions to $\hat{\pi}_1(s) = \hat{\pi}_2(s) = 0.5$.

The intensity correction of S can now be performed by reinjecting the estimated parameters in Equations (12) and (11).

2.3 Combining the Intensity and Geometrical Transformations

As stated previously, the use of Equation 1 in Section 2.1 to model geometrical transformations makes the assumption that the intensities of structures in S match those in T . When dealing with images obtained from different modalities, this requirement is not met. To solve this problem, we presented in Section 2.2 a method to rectify intensity differences between corresponding structures in S and T using an underlying monofunctional or bifunctional model. The result of this intensity correction on S for the n^{th} iteration will be denoted by S_n^* with $S_0^* = S$. Considering this and Equation 1, the geometrical transformation is found using

$$v_{n+1}(x) = G_\sigma \otimes \left(v_n + \frac{S_n^* \circ h_n(x) - T(x)}{\|\nabla(S_n^* \circ h_n)(x)\|^2 + [S_n^* \circ h_n(x) - T(x)]^2} \nabla(S_n^* \circ h_n)(x) \right). \quad (13)$$

2.4 Mutual information

Another approach we tested to compute incremental displacements was inspired by the works of Viola *et al.* [30] and Maes *et al.* [18]. In these approaches, the rigid/affine registration between two images is formulated as a maximization of their MI. This choice is motivated by the fact that MI models the similarity between the images while resorting to assumptions that are much more general than functional dependence. Consequently, and contrary to the approach that was presented in section 2.2, they do not attempt to apply an intensity correction to one image so that it matches the other. Instead, they model the intensity dependence in a purely statistical fashion. Viola *et al.* [30] do this modeling using Parzen windowing, which enables them to differentiate the MI criterion with respect to the geometrical transformation.

As shown in Section 2.1.1, our incremental displacements closely relate with a gradient descent on the SSD criterion. Using this analogy, we implemented an alternative matching strategy where the incremental displacements given by Equation (13) are replaced with the following formula,

$$v_{n+1}(x) = G_\sigma \otimes (v_n + \alpha \nabla MI(v_n)),$$

in which α is a positive constant and $\nabla MI(v_n)$ is the gradient of the MI criterion with respect to the n^{th} displacement vector, whose mathematical expression is found in [30]. This updating corresponds to a first order gradient descent on the MI criterion, of course up to the Gaussian filtering which is used as a regularization constraint.

3 Data

Most of the data used in the following experiments were obtained from BrainWeb [3, 8, 16, 9]. This tool uses an atlas with a resolution of $1 \times 1 \times 1\text{mm}^3$ comprising nine segmented regions (See Table 1.) from which T1, T2 and PD images can be generated. Three images, one of each modality, were generated with the same resolution as the atlas, 5% noise and no intensity non-uniformity. Since they are generated from the same atlas, they represent the same underlying anatomy and are all perfectly matched.

Tissue type	Intensity	Tissue type	Intensity
Background	0	Muscle / Skin	140
CSF	28	Skin	168
Gray Matter	56	Skull	196
White Matter	84	Glial Matter	224
Fat	112	Connective	252

Table 1: Intensity values corresponding to the different tissue types present in the atlas.

We also made use of a T1 MR image and a CT image, both from different subjects and having a resolution of $1 \times 1 \times 1\text{mm}^3$. Both these images were affinely registered with the atlas using the correlation ratio method [23]. To differentiate the T1 image obtained with the atlas from the other T1 image, the latter will be referenced as SCH.

The images all respect the neurological convention, i.e. on coronal and axial slices, the patient’s left is on the left side of the image.

4 Results and Discussion

In the following section we present registration results involving images obtained from several different kinds of modalities. First, we show a typical example where monofunctional dependence can be assumed: registration of an atlas with an MR image. Then, more practical examples are shown where images from different modalities are registered and where bifunctional dependence may be assumed. We also provide registration results obtained from our implementation of MI maximization as described in Section 2.4.

The multilevel process was performed at three resolution levels, namely 4mm, 2mm and 1mm per voxel. Displacement fields at one level are initialized from the result of the previous level. The initial displacement field v_0 is set to a zero. The Gaussian filter G_σ used to smooth the displacement field has a standard deviation of 1mm. 128 iterations are performed at 4mm/voxel, 32 at 2mm/voxel and 8 at 1mm/voxel. We believe that

making use of a better stopping criterion, such as the difference of the SSD values between iterations, would probably improve the results shown below.

4.1 Monofunctional Dependence

We present here the result of registering the atlas with SCH. Since the atlas can be used to generate realistic MR images, it is safe to assume a functional dependence between the intensity of the atlas and those of SCH. Also, since SCH and the atlas are well aligned due to the affine registration, we have roughly estimated that the number of points already well matched are at least $0.80 \times N$, to which we have set the value of c . Since 10 classes are present in the atlas, the polynomial degree chosen was set to 9.

The result of registration is presented in Figure 1. The first row shows coronal, sagittal and axial slices of the atlas. In the second one are the corresponding slices of SCH. In the third and fourth rows are presented the result of registering the atlas with SCH using our algorithm. The third one shows the images resulting from the registration process. It has the same shape as the resampled atlas and intensities have been transformed using the intensity correction. The fourth one shows the result without the intensity transformation; we have simply applied to the atlas the geometrical transformation resulting from the registration procedure.

As can be seen, there is one obvious problem with this result: although the shape of the atlas seems well corrected, the CSF intensity is not (as can be seen in the ventricles and around the cortex).

This problem can also be observe by looking at the intensity transformation function presented in Figure 2 (Intensity values corresponding to the different tissue types present in the atlas are shown in Table 1.). This is due to an overlap of the CSF in the atlas and the gray and white matter in SCH, especially around the cortical area which is known to present large variations between subjects. We believe this is due to the strong smoothness constraints imposed by the Gaussian regularization which may prevent the assessment of large and uneven displacements required to match the cortex.

To verify this assumption, we registered T1 with SCH and deformed T1 using the resulting displacement field. This provided an image SCH' that closely resembles SCH (Compare images from the second rows of Figures 1 and 3) and for which we are assured that the shape variation from the atlas can be assessed by our algorithm. We then registered the atlas with SCH'. This result is presented in Figures 3. As can be observed, the CSF intensity value is now well corrected. By looking at the intensity transformation shown in Figure 4, we also notice that each atlas structure has corresponding intensity ranges in SCH that are less extended as those of Figure 2, reflecting a better match between the images. This finding puts forth that the displacement field regularization has to be able to accommodate the large uneven displacement of the cortex. To cope with large displacements, Gaussian filtering may probably be replaced with another regularization strategy such as that based on a fluid model [7] or on a non-quadratic potential energy [14].

Another difference between Figures 2 and 4 is the difference in the intensity mapping of the glial matter (Intensity 224 in the atlas.). This structure forms a small envelope around the ventricles and is represented by so few points that they are considered as outliers. This could be corrected by considering more points during the intensity transformation. In fact, c should be increased at each iteration to reflect that, during registration, S gradually aligns with T and more points in C can be considered as well matched.




Figure 1: Corresponding coronal, sagittal and axial slices of the atlas to SCH registration result. From top to bottom: Atlas; SCH; atlas with intensity correction after registration with SCH; atlas without intensity correction after registration with SCH.

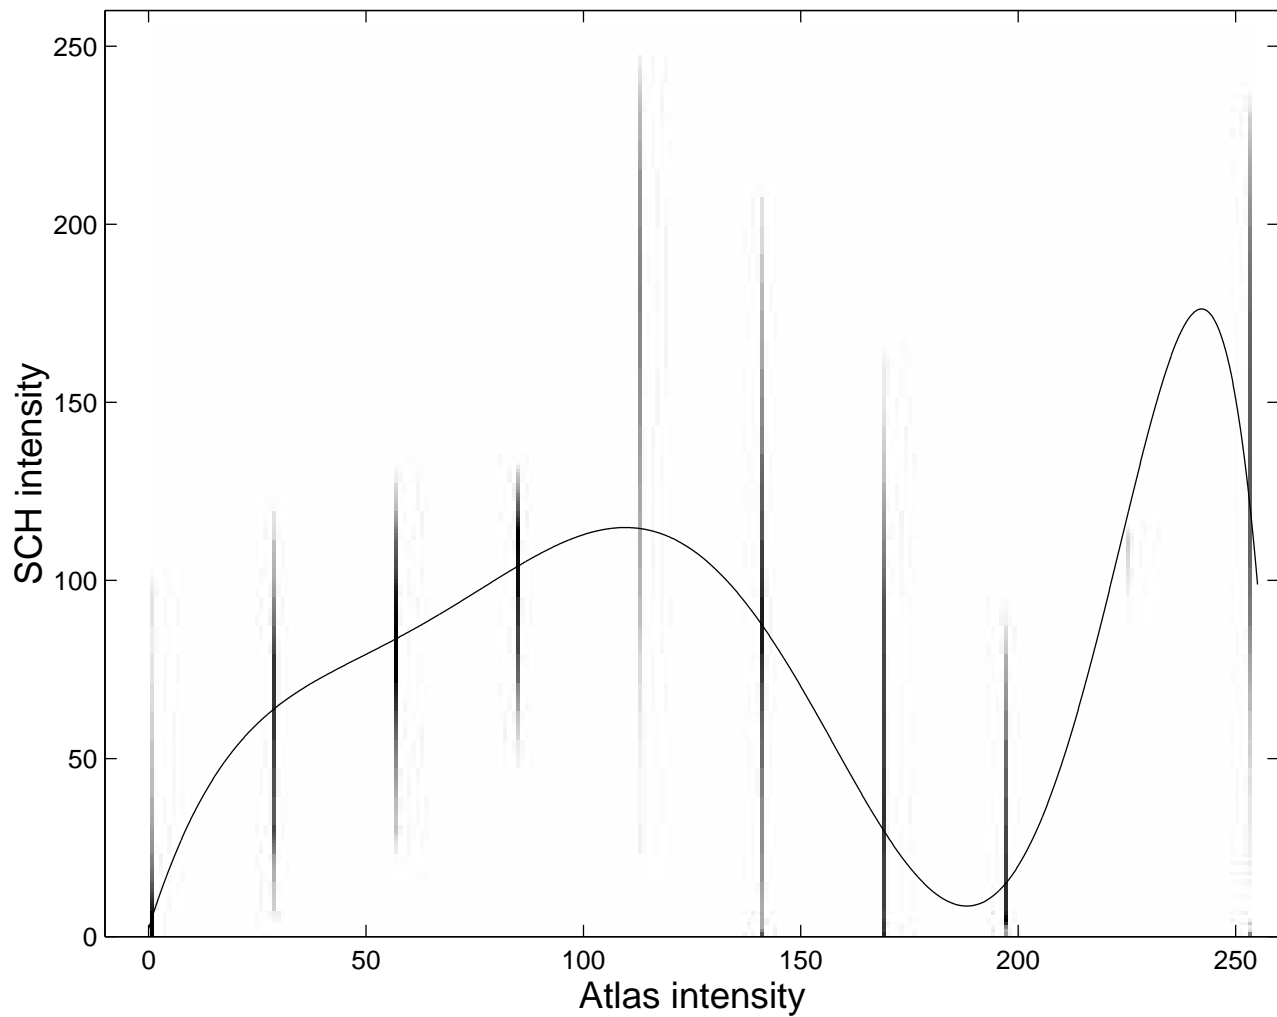


Figure 2: Intensity transformation found by registering the atlas with SCH and assuming monofunctional dependence. The function is overlaid on the joint histogram of the two images after registration. The tissue types corresponding to the different atlas intensities are presented in Table 1.




Figure 3: Corresponding coronal, sagittal and axial slices of the atlas to SCH' registration result. From top to bottom: Atlas; SCH'; atlas with intensity correction after registration with SCH'; atlas without intensity correction after registration with SCH'.




Figure 4: Intensity transformation found by registering the atlas with SCH' and assuming monofunctional dependence.

4.2 Bifunctional Dependence

When registering images from different modalities, monofunctional dependence may not necessarily be assumed. Here, we applied the method described in Section 2.2.2 where two polynomial functions of degree 12 are estimated. This number was set arbitrarily to a relatively high value to enable important intensity transformations.

Figure 5 presents the result of registering T1 with CT. Using these last two modalities, most intensities should be mapped to gray and only the skull, representing a small portion of the image data, should be mapped to white. After affine registration almost all voxels are well matched. Hence, in this particular case, we have chosen a high value for c set to $0.90 \times N$.

As we can see in Figure 5, the skull, shown in black in the MR image and in white in the CT scan, is well registered and the intensity transformation adequate. The top graph of Figure 6 presents the joint histogram of the two images after registration. This histogram is color-coded and ranges from red representing high point densities to blue depicting low point densities. The bottom graph of this figure presents the functions f_{θ} and f_{ψ} found during the registration process. The red line corresponds to f_{θ} and the blue one to f_{ψ} . The line width for a given intensity s is proportional to the value of the corresponding $\pi_c(s)$. The gray values represent the joint histogram after registration.

As can be observed in Figure 6, the polynomials found fit well with the high density clusters of the joint histogram. Still, some points need to be addressed.

We can observe that due to the restricted polynomial degree, f_{θ} , (shown in red) oscillates around the CT gray value instead of fitting a straight line. This is reflected in the intensity corrected image, shown in the third row of Figure 5, where the underlying anatomy can still be observed by small intensity variations inside the skull. This artifact has insubstantial consequences during the registration process since the difference between most of the voxels of T and S is 0, resulting in null displacements. The displacements driving the deformation will be those of the skull and the skin contours, and will be propagated in the rest of the image with the Gaussian filtering of the displacement field.

We also notice that f_{ψ} (shown in blue), which is mainly responsible for the mapping of the skull, does not properly model the cluster it represents for intensities smaller than 5. The mapping for these intensities is slightly underestimated. This may have two causes. First, as in the previous case, it might be due to the restricted polynomial degree. Second, we can notice that some of the background values in T1 that have an intensity close to 0 are mapped to gray values in the CT which correspond to soft tissues. This means that some of the background in the T1 is matched with the skin in the CT. This has the effect of “pulling” f_{ψ} closer to the small cluster positioned around (2,65). If the underestimation of f_{ψ} arises because of the second reason, letting the algorithm iterate longer might provide a better result.

In Figures 7 and 9, we present the result of registering T2 and PD respectively with SCH. Figures 8 and 10 shows the corresponding intensity transformations. For these experiments, c was set to $0.60 \times N$, a value we have found to be effective for these types of modalities after affine registration.

One observation that can be made by looking at the joint histograms of Figures 8 and 10 is that there seems to be a functional dependence between the images intensities, i.e. there is a function that can go through the




Figure 5: Corresponding coronal, sagittal and axial slices of T1 to CT registration result. From top to bottom: T1; CT; T1 with intensity correction after registration with CT; T1 without intensity correction after registration with CT.

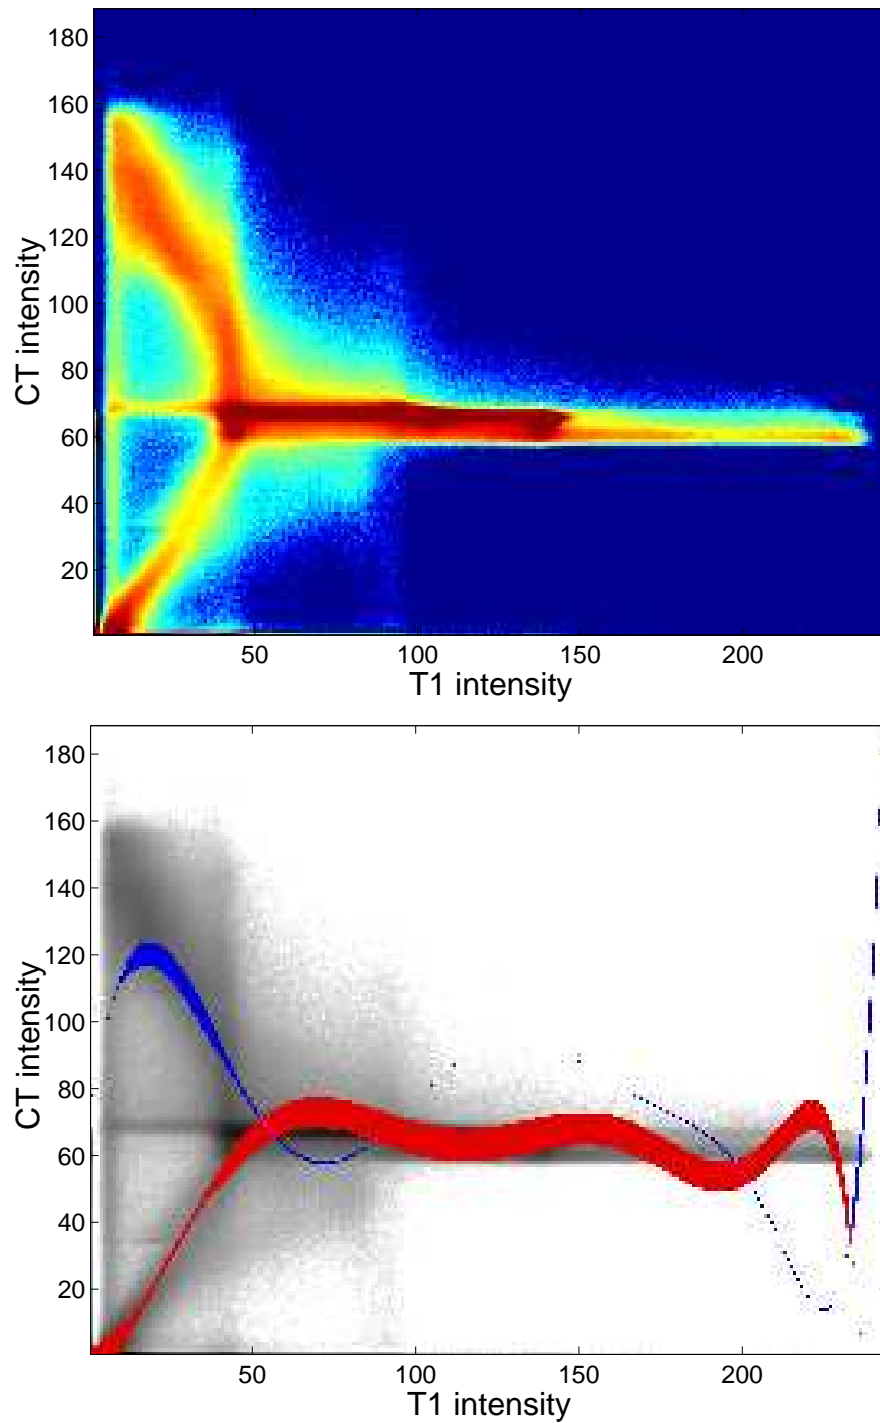


Figure 6: Intensity transformation found when registering T1 with CT and assuming bifunctional dependence. The top graph presents the joint histogram of the two images after registration. This histogram is color-coded and ranges from red representing high point densities to blue depicting low point densities. In the bottom graph, the red line corresponds to f_{θ} and the blue one to f_{ψ} . The line width for a given intensity value s in T1 corresponds to the value of the corresponding $\pi_{\epsilon}(s)$. The gray values represent the joint histogram after registration.




Figure 7: Corresponding coronal, sagittal and axial slices of T2 to SCH registration result. From top to bottom: T2; SCH; T2 with intensity correction after registration with SCH; T2 without intensity correction after registration with SCH.




Figure 8: Intensity transformation found when registering T2 with SCH and assuming bifunctional dependence.




Figure 9: Corresponding coronal, sagittal and axial slices of PD to SCH registration result. From top to bottom: PD; SCH; PD with intensity correction after registration with SCH; PD without intensity correction after registration with SCH.

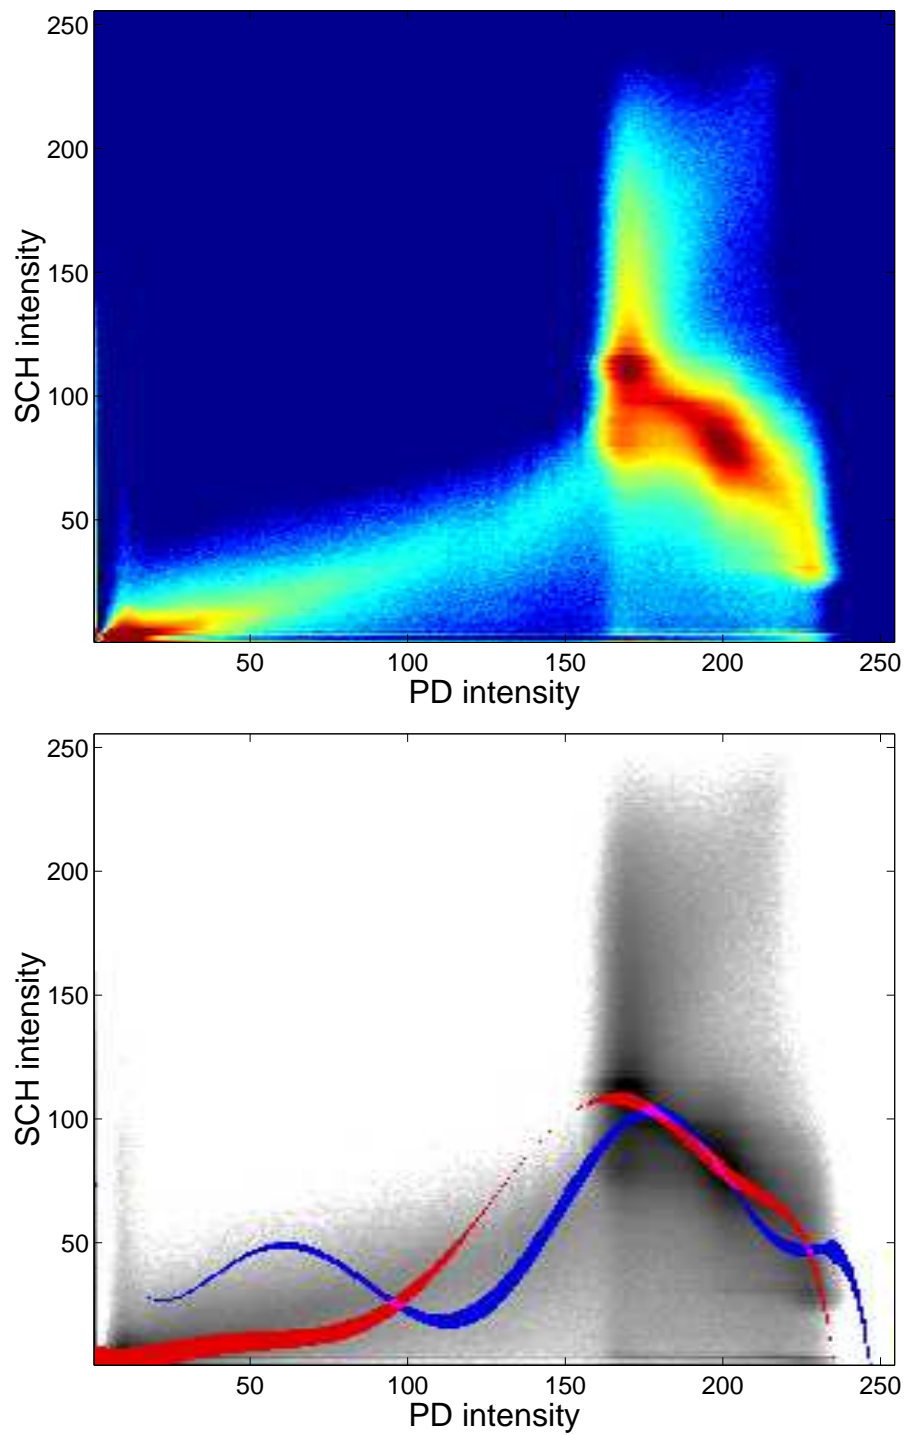


Figure 10: Intensity transformation found when registering PD with SCH and assuming bifunctional dependence.

major clusters of the joint histograms. This is also reflected by the closely similar shapes of the corresponding f_{θ} and f_{ψ} .

Thus, we have registered PD with SCH using the monofunctional model and $c = 0.8$. This result is presented in Figure 11. As can be seen by comparing the images of the third rows of Figures 9 and 11, the intensity transformation does not correct as well the CSF intensities and the distinction between the different structures is less contrasted.

This may be explained by a closer look at our bifunctional intensity modeling. Equation 10 reflects the assumption that if an anatomical point has an intensity s in S , the corresponding point has an intensity t in T that is distributed normally around two possible values depending on s . But it makes no assumption about how the intensities in S are distributed. This models the intensities of S without noise, which may not necessarily be well justified, but enables the use of linear regression to estimate the intensity transformation.

The effect of noise in S is reflected in the joint histograms by enlarging clusters along the x axis. This, added to bad matches and partial volume effect, creates many outliers in C and makes the assessment of the true intensity transformation more difficult and more resistant to our robust regression technique. Preprocessing of S using for example anisotropic diffusion may narrow the clusters and provide better results [22].

Adding the estimation of a second function in the bifunctional model helps counter the effect of noise on S . For example, the CSF in the PD image has intensity values ranging from about 200 to 240 and gray matter from about 175 to 210. In SCH, these ranges are about 30 to 70 and 55 to 80 respectively. As can be seen in Figure 10, f_{θ} models well the gray matter cluster but fails to reflect the CSF transformation. This is also well depicted in Figure 11 in which the CSF and gray matter intensity transformation is modeled using a single polynomial. In this case, the CSF is often mapped as gray matter. Estimating the second polynomial f_{ψ} solves this problem by considering the CSF cluster.

Note that for the registration of the atlas to SCH and of the T1 to the CT, we have always deformed the image with the most information, i.e. the one with the higher number of perceivable structures. This is simply because our algorithm permits many structures of the deformed image to be mapped to a single intensity, as is the case when transforming a T1 image into a CT image (See Figure 5.). But a single intensity in the deformed image can be mapped to at most 2 intensities in the target image. For example, if we used the CT image as the image to be deformed, the dominant gray intensity value in this image would have to be mapped to gray matter, white matter, CSF, etc. This would require more than two functions to be estimated and complicates the algorithm. Hence, it is always better to use the image with the most number of structures visible as the source image.

4.3 Mutual Information

We present in Figure 14 the result of registering the T1 image with the CT image using the MI method described in Section 2.4. A typical difference between using our bifunctional method instead of our MI implementation can be appreciated by comparing the images from the third rows of Figures 5 and Figure 14. As can be seen, the contours of corresponding structures do not match after registration using MI. The head contours seem to be attracted by the image borders, which means that the driving forces have misleading directions in this




Figure 11: Corresponding coronal, sagittal and axial slices of PD to SCH registration result using the mono-functional model. From top to bottom: PD; SCH; PD with intensity correction after registration with SCH; PD without intensity correction after registration with SCH.




Figure 12: Results of registration from different modalities to SCH. Left column, an axial view of the SCH and corresponding contours extracted using a Canny-Deriché edge detector. On the middle and right columns, the contours are overlaid with the results of affine and elastic registration, respectively. From top to bottom: T1; ATLAS; T2; PD.




Figure 13: Results of registration from different modalities to SCH. Left column, a coronal view of the SCH and corresponding contours extracted using a Canny-Deriche edge detector. On the middle and right columns, the contours are overlaid with the results of affine and elastic registration, respectively. From top to bottom: T1; ATLAS; T2; PD.

region. This outcome might be due to the fact that Parzen windowing provides too few information on how to match the intensities of the images. As a consequence, the direction of the MI gradient, from which the local driving force are derived, might be unreliable. Many registrations were performed using the MI criterion with varying values for α and several Parzen window sizes. The results we obtained using this strategy were much less convincing than the ones we obtained using our bifunctional method; the source image deformed very little unless the α parameter was set to a high value, in which case the displacements looked random.




Figure 14: Corresponding coronal, sagittal and axial slices T1 to CT registration result using MI. From top to bottom: T1; CT; T1 after registration with CT.

The poor results obtained with our implementation of MI might sound surprising as MI performs generally very well in rigid/affine registration. Although it would be irrelevant to conclude that MI is not well-suited for elastic registration, a partial explanation may be given. What basically does MI is to model the dependence between the images by estimating their intensity joint distribution among the set of every possible distribu-

tions [22]. This set has a dimension $k_s \times k_t - 1$, where k_s and k_t are the numbers of intensity levels respectively in the source and target images. In the context of rigid/affine registration, where no more than 12 geometrical parameters need to be estimated, such an enormous search space is usually affordable. However, with elastic registration, the geometrical transformation is granted many degrees of freedom and maximizing MI might then become an under-constrained problem.

Our method distinguishes itself from MI by imposing strong constraints on the intensity space. In the case where monofunctional dependence is assumed, only $p + 1$ parameters are estimated to model the intensity dependence, p being the polynomial degree. When assuming a bifunctional relationship, this number becomes $2(p + 1) + 1 + 2k_s$. To give a numerical comparison in the typical case of one-byte encoded images ($k_s = k_t = 256$), and with a polynomial degree $p = 25$ (an upper bound), our method estimates 26 parameters in the monofunctional version and 565 parameters in the bifunctional version, whereas MI requires 65535 parameters to be estimated.




Figure 15: T1 to CT registration results. First row, left: an axial view of the CT image. On the same row, the other pictures show the CT contours overlaid with the corresponding view of the T1 image, respectively after affine registration (second from the left), elastic registration using the bifunctional model (third), and elastic registration using MI (fourth). The second row presents the same overlay for a coronal view.

4.4 Displacement field comparison

Since the atlas, the T1, the T2 and the PD images have all been registered with SCH, it is relevant to compare some statistics of the resulting displacement fields to assess if our algorithm provides consistent results across modalities.

We computed statistics regarding the difference between any two of these displacement fields. The length of the vectors of the resulting difference fields were calculated. Each cell of Table 2 presents, for each combination of displacement fields, the median length, the average length with the corresponding standard deviation and the maximum length of the difference field.

The two largest average error are 1.76mm and 1.58mm and were found when registering the atlas with T1 and PD respectively. This may be explained by the intensity correction bias for the CSF that would tend to attenuate displacements and produce larger errors, a problem invoked in Section 4.1. Aside from these, the average error length varies between 0.97mm and 1.40mm and the median error is between 0.85mm and 1.32mm. These are values in the range of the image resolution of 1.0mm. Note also that all the standard deviations are below this value.

We point out that these are global measures that are presented to provide an idea of the differences between the displacement fields. They do not strictly provide a validation of the method, but do show a certain coherence between the different results we obtained.

5 Conclusion

In this paper, we introduced an original method to perform non-rigid registration of multimodal images. This iterative algorithm is composed of two sections: the geometrical transformation and the intensity transformation. We have related the geometrical transformation model to several popular registration concepts: SSD, optical flow and the demons method. Two intensity transformation models were described which assume either monofunctional or bifunctional dependence between the images to match. Both of these models are built using robust estimators to enable precise and accurate transformation solutions. Result of registration were presented and showed that the algorithm performs very well for several kinds of modalities including T1 MR, T2 MR, PD MR, CT and segmentations, and provides consistent results across modalities. Our algorithm was compared with the maximization of the MI criterion and seems to be more apt at evaluating high-dimensional deformations, as it puts more constraints on the parameters to be estimated and thus permits a better search of the parameter space.

Acknowledgments

Alexandre Guimond is funded in part by Le Fonds FCAR, Québec, Canada. Alexis Roche is funded in part by la région Provence-Alpes-Côte d’Azur, France.

Many thanks to Grégoire Malandain for providing us with a (marvelous) Canny-Deriché edge detector as well as various image processing tools.

		Atlas-SCH	Atlas-SCH'	T1-SCH	T2-SCH	PD-SCH
Atlas-SCH	median		0.85	1.46	1.13	1.67
	average		0.97	1.58	1.23	1.76
	std. dev.		0.60	0.84	0.63	0.79
	maximum		4.57	6.99	5.14	7.10
Atlas-SCH'	median	0.85		0.89	0.95	1.15
	average	0.97		1.00	1.07	1.25
	std. dev.	0.60		0.55	0.60	0.64
	maximum	4.57		4.48	5.46	7.43
T1-SCH	median	1.46	0.89		1.00	1.01
	average	1.58	1.00		1.18	1.16
	std. dev.	0.84	0.55		0.78	0.71
	maximum	6.99	4.48		7.17	8.08
T2-SCH	median	1.13	0.95	1.00		1.32
	average	1.23	1.07	1.18		1.40
	std. dev.	0.63	0.60	0.78		0.68
	maximum	5.14	5.46	7.17		6.86
PD-SCH	median	1.67	1.15	1.01	1.32	
	average	1.76	1.25	1.16	1.40	
	std. dev.	0.79	0.64	0.71	0.68	
	maximum	7.10	7.43	8.08	6.86	

Table 2: Statistics regarding the displacements difference between each type of registration. Each cell present the median length, the average length with the corresponding standard deviation and the maximum length. All measures are in millimeters.

References

- [1] D. C. Barber. Registration of low resolution medical images. *Physics in Medicine and Biology*, 37(7):1485–1498, 1992.
- [2] J. L. Barron, D. J. Fleet, and S. S. Beauchemin. Performance of optical flow techniques. *International Journal of Computer Vision*, 12(1):43–77, January 1994.
- [3] Simulated brain database. <http://www.bic.mni.mcgill.ca/brainweb/>.
- [4] J. W. Brandt. Improved accuracy in gradient-based optical flow estimation. *International Journal of Computer Vision*, 25(1):5–22, 1997.
- [5] M. Bro-Nielsen and C. Gramkow. Fast fluid registration of medical images. In K. H. Höhne and R. Kikinis, editors, *Proceedings of the 4th International Conference Visualisation in Biomedical Computing (VBC'96, held in Hamburg, Germany, September 22-25, 1996)*, volume 1131 of *Lecture Notes in Computer Science*, pages 267–276. Springer-Verlag, 1996.
- [6] P. Cachier, X. Pennec, and N. Ayache. Fast non rigid matching by gradient descent: Study and improvements of the “demons” algorithm. Technical Report 3706, Institut National de Recherche en Informatique et en Automatique, June 1999. <http://www.inria.fr/RRRT/RR-3706.html>.
- [7] G. E. Christensen, R. D. Rabbitt, and M. I. Miller. Deformable templates using large deformation kinematics. *IEEE Transactions in Medical Imaging*, 5(10):1435–1447, October 1996.
- [8] C. A. Cocosco, V. Kollokian, R. K.-S. Kwan, and A. C. Evans. Brainweb: Online interface to a 3D MRI simulated brain database. *NeuroImage, Proceedings of the Third International Conference on Functional Mapping of the Human Brain (HBM'97, Conference held in Copenhagen, Denmark, May 19–23 1997)*, 5(4):S425, May 1997.
- [9] D. L. Collins, A. P. Zijdenbos, V. Kollokian, J. G. Sled, N. J. Kabani, C. J. Holmes, and A. C. Evans. Design and construction of a realistic digital brain phantom. *IEEE Transactions in Medical Imaging*, 17(3):463–468, June 1998.
- [10] B. M. Dawant, J.-P. Thirion, F. Maes, D. Vandermeulen, and P. Demaerel. Automatic 3D segmentation of internal structures of the head in MR images using a combination of similarity and free form transformations. In K. M. Hanson, editor, *Medical Imaging 1998: Image Processing (MI'98, Conference held in San Diego (CA), USA, February 23–26 1998)*, volume 3338 of *SPIE Proceedings*, pages 545–554, San Diego (CA), USA, February 1998. International Society for Optical Engineering.
- [11] J. Feldmar, J. Declerck, G. Malandain, and N. Ayache. Extension of the ICP algorithm to non-rigid intensity-based registration of 3D volumes. *Computer Vision and Image Understanding*, 66(2):193–206, May 1997.

- [12] T. Gaens, F. Maes, D. Vandermeulen, and P. Suetens. Non-rigid multimodal image registration using mutual information. In W. M. Wells, A. Colchester, and S. Delp, editors, *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'98, Conference held in Cambridge (MA), USA, October 11–13, 1998)*, volume 1496 of *Lecture Notes in Computer Science*, pages 1099–1106. Springer-Verlag, 1998.
- [13] J. W. Haller, A. Banerjee, G. E. Christensen, M. Gado, S. C. Joshi, , M. I. Miller, Y. Sherline, M. W. Vannier, and J. G. Csernansky. Three-dimensional hippocampal MR morphometry with high-dimensional transformation of a neuroanatomic atlas. *Radiology*, 202:504–510, 1997.
- [14] P. Hellier, C. Barillot, E. Mémin, and P. Pérez. Medical image registration with robust multigrid techniques. In *Proceedings of the Second International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'99, Conference held in Cambridge, England, September 19–22, 1999)*, volume 1679 of *Lecture Notes in Computer Science*, pages 680–687, Cambridge, England, October 1999.
- [15] B. K. P. Horn and B. G. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, August 1981.
- [16] R. K.-S. Kwan, A. C. Evans, and G. B. Pike. An extensible MRI simulator for post-processing evaluation. In K. H. Höhne and R. Kikinis, editors, *Proceedings of the 4th International Conference Visualisation in Biomedical Computing (VBC'96, held in Hamburg, Germany, September 22-25, 1996)*, volume 1131 of *Lecture Notes in Computer Science*, pages 135–140. Springer-Verlag, 1996.
- [17] Y. H. Lau, M. Braun, and B. F. Hutton. Non-rigid 3d image registration using regionally constrained matching and the correlation ratio. In *International Workshop on Biomedical Image Registration (WBIR'99, Workshop held in Bled, Slovenia, August 30–31, 1999)*, 1999. In Print.
- [18] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality image registration by maximization of mutual information. *IEEE Transactions in Medical Imaging*, 16(2):187–198, 1997.
- [19] J. B. A. Maintz, E. H. W. Meijering, and M. A. Viergever. General multimodal elastic registration based on mutual information. In K. M. Hanson, editor, *Medical Imaging 1998: Image Processing (MI'98, Conference held in San Diego (CA), USA, February 23–26 1998)*, volume 3338 of *SPIE Proceedings*, pages 144–154, Bellingham (WA), USA, April 1998. International Society for Optical Engineering.
- [20] X. Pennec, P. Cachier, and N. Ayache. Understanding the “demon’s algorithm”: 3d non-rigid registration by gradient descent. In *Proceedings of the Second International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'99, Conference held in Cambridge, England, September 19–22, 1999)*, volume 1679 of *Lecture Notes in Computer Science*, pages 597–605. Springer-Verlag, 1999.
- [21] A. Roche, G. Malandain, and N. Ayache. Unifying maximum likelihood approaches in medical image registration. *International Journal of Imaging Systems and Technology*, 1999. Submitted. <http://www.inria.fr/RRRT/RR-3741.html>.

- [22] A. Roche, G. Malandain, N. Ayache, and S. Prima. Towards a better comprehension of similarity measures used in medical image registration. In *Proceedings of the Second International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'99, Conference held in Cambridge, England, September 19–22, 1999)*, volume 1679, pages 555–566. Springer-Verlag, October 1999.
- [23] A. Roche, G. Malandain, X. Pennec, and N. Ayache. The correlation ratio as a new similarity measure for multimodal image registration. In W. M. Wells, A. Colchester, and S. Delp, editors, *Proceedings of the First International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI'98, Conference held in Cambridge (MA), USA, October 11–13, 1998)*, volume 1496 of *Lecture Notes in Computer Science*, pages 1115–1124. Springer-Verlag, 1998. <http://www.inria.fr/RRRT/RR-3378.html>.
- [24] P. J. Rousseeuw and K. Van Driessen. Computing LTS Regression for Large Data Sets. Technical report, Statistics Group, University of Antwerp, 1999. <http://win-www.uia.ac.be/u/statis/publicat/fastlts/readme.html>.
- [25] Peter J. Rousseeuw and Annick M. Leroy. *Robust Regression and Outlier Detection*. Wiley series in probability and mathematical statistics. John Wiley & Sons, 1987.
- [26] E. P. Simoncelli. Design of multi-dimensional derivative filters. In *International Conference on Image Processing*, Austin, USA, November 1994. IEEE.
- [27] J.-P. Thirion. Fast non-rigid matching of 3D medical images. Technical Report 2547, Institut National de Recherche en Informatique et en Automatique, Sophia-Antipolis, 1995. <http://www.inria.fr/RRRT/RR-2547.html>.
- [28] J.-P. Thirion. Image matching as a diffusion process: an analogy with Maxwell's demons. *Medical Image Analysis*, 2(3):243–260, 1998.
- [29] Arthur W. Toga. *Brain Warping*. Academic Press, 1999.
- [30] P. Viola and W. M. Wells. Alignment by maximization of mutual information. *International Journal of Computer Vision*, 24(2):137–154, 1997.
- [31] R. P. Woods, S. T. Grafton, J. D. G. Watson, N. L. Sicotte, and J. C. Mazziotta. Automated image registration: II. intersubject validation of linear and nonlinear models. *Journal of Computer Assisted Tomography*, 22(1):153–165, 1998.
- [32] R. P. Woods, J. C. Mazziotta, and S. R. Cherry. MRI-PET registration with automated algorithm. *Journal of Computer Assisted Tomography*, 17(4):536–546, 1993.



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