



Image matching as a diffusion process: an analogy with Maxwell's demons

Jean-Philippe Thirion

► To cite this version:

Jean-Philippe Thirion. Image matching as a diffusion process: an analogy with Maxwell's demons. Medical Image Analysis, 1998, 2 (3), pp.243–260. inria-00615088

HAL Id: inria-00615088

<https://inria.hal.science/inria-00615088>

Submitted on 17 Aug 2011

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.

Image matching as a diffusion process: an analogy with Maxwell's demons

J.-P. Thirion*

INRIA, Equipe Epidaure, 2004 Route des Lucioles BP93, 06902 Sophia-Antipolis, France

Abstract

In this paper, we present the concept of diffusing models to perform image-to-image matching. Having two images to match, the main idea is to consider the objects boundaries in one image as semi-permeable membranes and to let the other image, considered as a deformable grid model, diffuse through these interfaces, by the action of effectors situated within the membranes. We illustrate this concept by an analogy with Maxwell's demons. We show that this concept relates to more traditional ones, based on attraction, with an intermediate step being optical flow techniques. We use the concept of diffusing models to derive three different non-rigid matching algorithms, one using all the intensity levels in the static image, one using only contour points, and a last one operating on already segmented images. Finally, we present results with synthesized deformations and real medical images, with applications to heart motion tracking and three-dimensional inter-patients matching.

Keywords: deformable model, elastic matching, image sequence analysis, inter-patient registration, non-rigid matching

Received October 22, 1996; revised August 8, 1996; March 16, 1998; accepted April 13, 1998

1. INTRODUCTION

Many concepts of thermodynamics have been fruitfully applied in the field of information theory and, more specifically, to image processing. A recent example is the application of mutual entropy minimization techniques (see Viola and Wells, 1995; Maes *et al.*, 1997) to the matching of medical images acquired with different modalities.

Another instructive example is anisotropic filtering (see Perona and Malik, 1990; Catté *et al.*, 1992; Gerig *et al.*, 1992; Kimia and Siddiqi, 1994), now commonly used as a pre-processing step for medical images: it can be shown that the application of a Gaussian filter of parameter σ is equivalent to the diffusion of heat in a homogeneous material for a time duration directly related to σ . Heat propagation is a new way of looking at Gaussian filtering: variations of this concept, for example by using the object boundaries in the image to create inhomogeneities with respect to heat propagation, have

led to a large variety of new, non-linear algorithms for image filtering, even if convergence proofs have not always been established. Anisotropic filtering is not a new technique, but a new exciting way to consider a problem as ancient as image processing.

Extending the presentations in Thirion (1995, 1996), we propose an original viewpoint for image-to-image matching, also based on an analogy with thermodynamic concepts (see Figure 1). Having two images to match, the main idea is to consider the object boundaries in one image as semi-permeable membranes and to let the other image, considered as a deformable grid model, diffuse through these interfaces, by the action of effectors situated within the membranes. We illustrate this concept by an analogy with Maxwell's demons and we contrast it with more conventional viewpoints such as deformable models based on attraction. As we will see, diffusing models rely mainly on the notion of polarity (inside–outside), while attraction relies on the notion of distance, but mixed models can be devised, and we will also see how to translate the diffusing model concept into the concept of attraction, with an intermediate step corresponding to optical flow methods.

*Corresponding author

(e-mail: jean-philippe.thirion@inria.fr)

<http://www.inria.fr/epidaure/personnel/thirion/thirion.html>

Now at Focus Imaging, 449 Route des Crêtes, Sophia-Antipolis, 06560 Valbonne, France.)




Figure 1. Diffusing models: a deformed image, considered as a deformable grid, is diffusing through the contours of the objects in the static image, by the action of effectors, called demons, situated in these interfaces.

First, we recall some existing techniques concerning non-rigid matching, such as ‘snakes-based’ methods, similarity maximization techniques and optical flow. Then we present and detail our concept of diffusing models, illustrated by its application to a simplified case, when two identical discs are matched rigidly (demon 0). We present a general iterative scheme to implement diffusing models, with several possible variants (demons 1, 2 and 3) which illustrate how the concept can be applied to generate new matching algorithms. Finally, we present experimental results with several medical applications such as the tracking of deformable organs or three-dimensional (3-D) inter-patients matching.

2. NON-RIGID MATCHING TECHNIQUES

Matching is an essential task for many computer vision applications. A clear definition can be given for rigid or articulated bodies: it is to recover rigid displacements of rigid parts. The task is much more complicated for deformable objects, with plastic or elastic deformations. In this case, there is not a single definition of an ideal optimal match, but as many definitions as practical applications. Each time, one has to define precisely the set of deformations $\mathcal{T}$ which are explored (rigid, affine, spline, free-form etc), and the type of features which are used (points, curves, surfaces, intensities etc).

However, classifying all non-rigid matching techniques with a single metric is made impossible in practice by the numerous works recently performed in that domain for a large number of different goals. We only describe a few of these methods here.

2.1. The concept of attraction

A widespread (and very intuitive) way to consider matching is based on an analogy with attraction. One example is gravitation: a point P of the deformable model M is attracted by all the points P' in S which are similar. For example, let $K(P, P')$ be a similarity criterion, and $D(P, P')$ a function of the distance (not necessarily Euclidean), the force $\vec{f}$ on P ,




Figure 2. Deformable model with attraction.

induced by the attraction of all the points of S can be

$$\vec{f}(P) = \sum_{P' \in S} \frac{K(P, P')}{D(P, P')} P \vec{P}'. \quad (1)$$

M is deformed according to these forces, and also according to smoothness constraints internal to M . Such a method is computationally too expensive [$O(n^2)$]. A less expensive method is to retain in the computation only one principal attractive point of S : a point P of M is attracted by the point P' of S which is the ‘closest and most similar’ to P .

We perceive in this definition that a balance between being ‘close’ and being ‘similar’ must be determined. For example, in some methods, only contour points are used, and contour points in S are equally similar to contour points in M , making proximity the most important factor (see Figure 2). This is the case for iterative closest point (ICP) methods; see Besl and McKay (1992) and Zhang (1992) for the rigid case. This is also a basic assumption in the ‘snakes’ method introduced by Kass *et al.* (1987); see also Blake and Yuille (1992) for a review of these techniques.

Distance and similarity can be married in an elegant way. For example in Feldmar and Ayache (1994), the points have n attributes (or parameters), defined using differential geometry: these attributes are invariants associated with the points, such as principal curvatures. The coordinates x, y, z of the points are also considered as parameters, therefore each point of S is represented by a single point in an $(n + 3)$ -dimensional parametric space, where spatial coordinates and differential invariants are mixed. Then only distance is needed, like in ICP, but in the $(n + 3)$ D parametric space: it can be the Euclidean distance, or the Mahalanobis distance if uncertainty is evaluated as well (see Ayache, 1991).

Similarity can also be introduced into snakes-based techniques. For example in Benayoun and Ayache (1995), a third term is added to the energy function to be minimized, corresponding to the similarity of differential geometry properties attached to those points. This helps to reduce the well known ‘aperture problem’ (see Faugeras, 1993), which states that it is easier to recover displacements normal to the contour rather than along it.

In contrast to ICP or snakes, correlation techniques place more importance on similarity than on distance. The point P of M is attracted by P' of S which maximizes a correlation function $K(P, P')$. But a distance is also used, which is 1 if P' is within a neighbourhood of $T_i(P)$ (where T_i is the current estimated deformation) and $+\infty$ if P' is out of the correlation window.

As described, these methods are still expensive because at each iteration and for each P , a large number of possible attractive points P' in S are to be considered. In practice, there are algorithmic solutions to reduce the complexity of finding this 'closest and most similar' point. It can be the use of KD trees, like in Besl and McKay (1992) or Feldmar and Ayache (1994). It can be by reducing the set of feature points, like using the surface of organs (see Davatzikos, 1996; Thompson and Toga, 1996), or an even more compact representation such as crest lines (see Declerck *et al.*, 1995) or individual points (see Thirion, 1994; Rohr *et al.*, 1996). It can be the use of pre-computed distance maps, such as for chamfer matching techniques (see Borgefors, 1988; Malandain *et al.*, 1994).

An attractive point can also be considered to be a minimum of a potential field (an 'attractor'): by differentiating this potential field a local expression of the forces is obtained. This can be done for snakes [using, for example, the gradient, see Kichenassamy *et al.* (1995)], and also for correlation techniques. These differentiations give forces which are likely to be directed toward the 'closest most similar' point, but it is only a local approximation.

As we can see, the concept of attraction has inspired the majority of existing matching methods. Distance and similarity are central in attraction, while polarity is accessory and seldom used. We note however that there exist some works that make use of polarity—in Radeva *et al.* (1995) polarity is used to discard the forces of contours whose normals are directed away from their closest point. In the work of Chakraborty *et al.* (1996), regional information is used in the form of a third term to be minimized along with internal and external forces. This third term corresponds to the integral over the model interior of the points which are unlikely to be interior points. One can refer to the work of Ronfard (1994) for a method mostly based on polarity: we will discuss this last method in detail later on.

2.2. Optical flow methods

A special kind of method is optical flow, which is used to find small deformations in temporal sequences of images (see Horn and Schunck, 1981; Aggarwal and Nandhakumar, 1988; Barron *et al.*, 1994). At a given point P , let s be the intensity function in S and m the intensity in M (see Figure 3). The basic hypothesis of optical flow is to consider that the




Figure 3. Instantaneous velocity from image M to image S .

intensity of a moving object is constant with time, which gives, for small displacements, the optical flow equation

$$\vec{v} \cdot \vec{\nabla} s = m - s. \quad (2)$$

This constraint is not sufficient to define the velocity $\vec{v}$; see for example Simoncelli *et al.* (1991). One solution is to regularize the problem to obtain the local values of $\vec{v}$. Another solution is to consider that the end point of $\vec{v}$ is the closest point of the hypersurface m , with respect to spatial (x, y, z) translations (see Figure 3), which leads to Equation (3):

$$\vec{v} = \frac{(m - s) \vec{\nabla} s}{(\vec{\nabla} s)^2}. \quad (3)$$

This equation is unstable for small values of $\vec{\nabla} s$, leading to infinite values for $\vec{v}$. Ideally, the expression should be close to zero for small $\vec{\nabla} s$. A solution is to multiply Equation (3) by $((\vec{\nabla} s)^2 / ((\vec{\nabla} s)^2 + (m - s)^2))$, which gives Equation (4):

$$\vec{v} = \frac{(m - s) \vec{\nabla} s}{(\vec{\nabla} s)^2 + (m - s)^2} \quad (4)$$

or, $\vec{v} = \vec{0}$ if $(\vec{\nabla} s)^2 + (m - s)^2 < \epsilon$. With this expression, the optical flow can be calculated in two steps: first compute the instantaneous optical flow for every point in S , then regularize the deformation field.

In optical flow, $\vec{v}$ is considered to be a velocity because the images are two successive time frames: $\vec{v}$ is the displacement during the time interval between the two image frames. In fact, when comparing images of two different patients, there is no such temporal consideration and it is more general to consider $\vec{v}$ as being simply a displacement.

3. DIFFUSING MODELS

We now present the concept of diffusing models, with a parallel with Maxwell's demons.




Figure 4. Maxwell's demons and a mixed gas.

3.1. Maxwell's demons

The concept of demons was introduced in the 19th century by Maxwell to illustrate a paradox of thermodynamics (see Figure 4). Assume a gas composed of a mix of two types of particles a and b^a , and separated by a semi-permeable membrane. Assume also, that this membrane contains a set of 'demons', which are able to distinguish between the two types of particles, and allow particles of type a only to diffuse to side A of the membrane and particles of type b only to diffuse to the other side B. At the end, A contains only particles a , and B particles b .

This corresponds to a decrease of entropy, in contradiction with the second principle of thermodynamics. The paradox was solved because the demons generate a greater amount of entropy to recognize the particles; thus, the total entropy of the system has increased.

3.2. Demons for image processing

Let us see how to apply demons to image matching. We want to match a model image M with a scene image S : M must be deformed to resemble S as much as possible. We assume (see again Figure 1) that the contour of an object O in S is a membrane, and we scatter our demons along this contour. We assume also that we are able to determine locally, for each contour point in S , a vector perpendicular to this contour and oriented from the inside of the object to the outside (for example, the gradient of the image S can be used). We assume also that M is a deformable grid, whose vertices are particles which can be classified as 'inside' or 'outside' particles. The rigidity of the deformable grid M is determined by the relations between these particles, and different behaviours, from totally rigid to totally free form, can be obtained by varying these relations. We will see later on that many types of deformations can be used. We call M a diffusing model, and we give an informal definition for demons:

^aOriginally a and b were cool and hot particles.




Figure 5. Three iterations of a model based on attraction (top row) and a rigid diffusing model (bottom row). These examples are produced by actual implementations.

a demon is an effector situated in a point P of the boundary of an object O . It acts locally to push the model M inside O if the corresponding point of M is labelled 'inside', and outside O if it is labelled 'outside'.

Polarity (inside–outside) is central for diffusing models, and distance is optional.

3.3. Diffusing models: a simplified example

To illustrate this principle, we consider the simple case of two images displaying the same disc, and we restrict ourselves to rigid transformations. We consider how this problem is tackled having in mind the concept of attraction as well as of a diffusing model (see Figure 5).

- For attraction, it is natural to sample regularly the model disc boundary (like in snakes-based techniques), and to apply forces directed toward the closest point of the scene circle, which contains the attractive points: distance is central.
- For a diffusing model, we regularly sample the scene disc boundary, each point being a demon. The force of the demon is oriented from the inside to the outside of the disc if the corresponding model point is labelled 'outside', and the other way if the label is 'inside': polarity is central.

Both methods are iterative: at each iteration, the motion created by all the elementary forces is applied to the model. Again to simplify the comparison, the force magnitudes induced by the demons and the attractive points are constant and equal at each iteration, but decrease with each iteration, in order to allow convergence.




Figure 6. Next three iterations (top row, attractive points; bottom row, demons).




Figure 7. Example of problematic initializations: left, when the two objects to be matched do not overlap, diffusing models are inefficient. Right, with an attraction model that does not take polarity into account, and with forces decreasing with the distance, the model can get trapped in a local minimum.

In Figure 5, the top three images represent three iterations using attraction, and the bottom three images are three iterations using a diffusing model. We can see that the individual forces are clearly different between both methods, and that this is not due to a change of the reference system, because the transformation is rigid. Figure 6 presents the next three iterations for attractive points (top) and diffusion (bottom). The two methods converge toward a similar result (but this is not always the case) and for small deformations, the forces also become similar.

We note that if the two discs do not overlap initially, then the demon-based method does not apply, whereas attraction still works, but there are other examples where attraction models can get stuck into local minima (see Figure 7 for an example) while diffusing models can still give the expected solution.

It is remarkable to observe that a model solely based on polarity can lead to matching.




Figure 8. Two different uses of the concept of demons. Left, in ‘anticipating snake’ for image segmentation; right, in a diffusing model, used to perform image-to-image matching.

3.4. Related works

From all the methods that we have considered, the closest work in its spirit is probably the ‘anticipating snakes’ of Ronfard (1994) used for two-dimensional (2-D) image segmentation (see Figure 8). The principle of the method is to consider a deformable contour where the external forces are replaced by ‘region-based’ forces, relying mainly on polarity, and which fit perfectly to our definition of demons. Anticipating snakes is distinct from diffusing model though, mainly because what is considered in Ronfard (1994) is a deformable contour evolving in a static image instead of a deformable image grid diffusing through a static contour. One visible consequence is that forces are applied to a deformable contour in Ronfard (1994), and to a deformable grid in our method. Along iterations, external forces are also always applied in the same location of the deformable contour in anticipating snakes while they are applied on successive locations of the deformable grid in diffusing models.

With respect to medical applications, the most related works are Bajcsy and Kovacic (1989), Gee *et al.* (1993) or Christensen *et al.* (1994b) where a global image, considered as a 2-D or 3-D grid, is also deforming. In these methods, the external forces exerted on the grid are traditional ones, based for example on the derivatives of a cross-correlation similarity measure.

It seems to us that the latter methods could be studied from the viewpoint of diffusing models, although it is still unclear to us how to do so. However, we are able to show that, to some extent, optical flow can be considered as an intermediate concept between diffusing models and attraction.

3.4.1. Optical flow behaves like diffusing models

Through each point P of a scene image S where $\vec{\nabla}s \neq 0$ is an iso-contour $s = I$, where $I = s(P)$ is constant. This iso-contour is the interface between the inside regions $s < I$ and the outside regions $s > I$. Comparing the intensities of the model image M with I also gives an automatic way to label the points of M ‘inside’ or ‘outside’.

The displacement $\vec{v}$ [Equation (4)] is comparable to the application of an elementary force $\vec{f}$ during one iteration step, whose direction is the same as $\vec{\nabla}s$ and whose orientation^b is according to $-\vec{\nabla}s$ if $m < s$, that is $m < I$ and according to $+\vec{\nabla}s$ when $m > I$. In other words, the force f pushes a point P of M toward the outside (that is, $\vec{\nabla}s$) when P is labelled outside ($m > I$), and toward the inside (that is $-\vec{\nabla}s$), when P is labelled inside ($m < I$), which is exactly our definition of a demon.

3.4.2. Optical flow also behaves like deformable models based on attraction

Consider again Figure 3: $\vec{v}$ is the shortest spatial displacement that brings the point $[P, m(P)]$ of the hypersurface corresponding to M , into the hyper-plane coming through $[P, s(P)]$, with normal $(-\vec{\nabla}s, 1)$, which is the best local approximation of the hypersurface S . If only local information are available [that is $m(P)$, $s(P)$ and $\vec{\nabla}s(P)$], it is legitimate to assume that $P' = P + \vec{v}$ is the point of S closest to P , and having the same intensity $[m(P)]$: P' is the closest point in S having the same intensity than P in M , which is by definition the behaviour of attractive points.

4. IMPLEMENTATIONS DERIVED FROM THE CONCEPT OF DIFFUSING MODELS

We propose a general scheme with several possible variants, leading to different implementations of diffusing models.

4.1. Matching as an iterative process

Similarly to deformable models based on attraction, diffusing models require an iterative scheme. The aim is to find a final transform $T \in \mathcal{T}$ (where $\mathcal{T}$ is the set of allowed deformations), between the space $\mathcal{M}$ of the model image M and the space $\mathcal{S}$ of the scene image S . T is the final evolution of a family of transforms $\{T_0, T_1, \dots, T_i, \dots\} \subset \mathcal{T}$.

At each step, the deformed version $T_i(M)$ of the model M becomes $T_{i+1}(M)$, constrained by ‘internal’ forces f_{int} created by the relations between the model points, and ‘external’ forces f_{ext} , created by the interactions between $T_i(M)$ and S . We can also change the reference system and process in $\mathcal{M}$ the interactions between M and $T_i^{-1}(S)$.

Both of the simplified examples described previously (diffusing model and attraction) illustrate an iterative scheme: in these cases, $\mathcal{T}$ is the group of rigid transforms and T_{i+1} is $\delta T_i \circ T_i$ where δT_i is the residual motion created by the set of elementary demons or attraction forces.

^bWe distinguish between direction and orientation: a straight line has a direction. An oriented line, a half line or a vector has an orientation as well as a direction, which is the way it is headed.




Figure 9. Iterative scheme in the case of diffusing models.

In the case of diffusing models, the first step is the pre-computation of the set of demons D_s , extracted from S , and the second step is an iterative estimation of the deformation T , from the model space $\mathcal{M}$ to the scene space $\mathcal{S}$ (Figure 9).

4.2. Extracting the demons from S

The demons set D_s is extracted from the scene image S . D_s can be the whole image grid (one demon per pixel or voxel), such as in Thirion (1995): in that case, the interfaces at each points are the iso-contours. D_s can be also restricted to the contour points in S , such as in the simplified example (demons 0), or to points automatically extracted by edge detection methods. The information attached to each demon can be:

- its spatial position P in S (possibly sub-pixel);
- a direction $\vec{d}$ from the inside to the outside [generally based on the gradient $\vec{\nabla}s(P)$];
- the current displacement from S to M at that point: $\vec{d} = P\vec{P}'$, where $P' = T_i^{-1}(P)$;
- information about the interface, such as the intensity at that location $s(P)$.

4.3. The iterative part

We start with an initial deformation T_0 (i.e. identity). At iteration i , we have a current estimated transform T_i , and each iteration is composed of two steps:

- For each demon $P \in D_s$, compute the associated elementary demon force $\vec{f}_i(P)$ which depends on the demon direction $\vec{d}_s$ at point P and on the polarity of M at point $T_i^{-1}(P)$.
- Compute T_{i+1} from T_i and from all the elementary demons forces $\{\vec{f}_i(P), P \in D_s\}$.

In some cases, step (ii) can be decomposed into two steps: first compute an elementary deformation δT_i from the elementary forces $\vec{f}_i$, then $T_{i+1} = \delta T_i \circ T_i$.

4.4. Possible variants

Useful variants of the scheme above can be obtained by varying:

- (i) the selection of the demon positions D_s (whole image grid, contours points etc).
- (ii) the space of deformations $\mathcal{T}$ (rigid, affine, spline, free form etc).
- (iii) the interpolation method which gives the values of M for the non-integer positions $T_i^{-1}(P)$ (linear, spline, sinc etc).
- (iv) the formula giving the force $\vec{f}$ of a demon (constant magnitude, gradient-based, optical flow like etc).

Thus, the simplified diffusing model example (demons 0) uses: (i) sample points of the disc contour of S for D_s ; (ii) rigid transforms for $\mathcal{T}$; (iii) no interpolation because $m(T_i^{-1}(P))$ is analytically defined and (iv) constant magnitude forces. In the following, in addition to the simplified case, we examine three other examples of diffusing models, illustrating different variants of the general scheme.

4.5. Demons 1: a complete grid of demons

The following method is especially suitable for 3-D medical image analysis (see Thirion, 1995). The general scheme is modified as follows.

- (i) All pixels of S where $\vec{\nabla}s \neq 0$ are selected to be demons ($D_s = S$).
- (ii) $\mathcal{T}$ is a free form deformation, that is, for each demon P we store the current elementary displacement $\vec{d}(P)$. To get a regular displacement field, a Gaussian filter with a given σ is applied to the whole field at each iteration. A theoretical study about how smoothing deformation fields can be used for regularization can be found in Anandan (1989).
- (iii) $m(P')$, where $P' = P + \vec{d}(P)$ is estimated using trilinear interpolation in M .
- (iv) the demon force is given by optical flow [$\vec{v}$ of Equation (4)], using only the information provided with each point: $m(P')$, $s(P)$ and $\vec{\nabla}s(P)$. More precisely, we do not define the force, but the result of the application of a force during one iteration step, which is a displacement: $\vec{d} = -\vec{v}$.

To make the algorithm faster (each voxel is considered at each iteration) and more robust with respect to initial positions, we adopt a multi-scale scheme, starting with a large number of iterations at coarse scales, down to a few iterations at the finest scale. More precisely, we use a pyramid approach where, at each scale, $\frac{1}{8}$ of the voxels and four times the number of iterations are used than that in the immediately finer scale,

which means that the computation time for the rest of the pyramid is equivalent to the time spent at the finest scale.

This gives us results qualitatively similar to other deformable grid methods such as those of Bajcsy and Kovacic (1989) or Christensen *et al.* (1994b). However, it is hard to assess quantitative differences, firstly because the definition of an 'ideal' inter-patients matching remains open to interpretation (see Thirion *et al.*, 1996).



With respect to computation time, our method is particularly time efficient, when compared to the results described by Christensen *et al.* (1996) with the use of a massively parallel computer. Besides the computational aspect of finding similar points, the method presented in Christensen (1994b, 1996) is also computationally expensive because it models exactly the physical behaviour of a viscous fluid. Bro-Nielsen and Gramkow (1996) have shown (and in 2-D) that those computations can be implemented using a convolution filter applied to the deformation field, leading to an improvement of one order of magnitude in speed. Another interesting result is that they demonstrate theoretically that the Gaussian filtering used in our method to regularize the deformation field approximates linear elasticity.

The algorithm demons 1 could have been derived from optical flow concepts too: to some extent, demons 1 can be considered as a particular kind of multi-scale, iterative version of optical flow. We derived it from diffusing models, which emphasize how optical flow can be treated as a diffusion process. This reveals possible changes in the algorithm, for example in the expression of the forces (see demons 2 and 3 below) which have less and less to do with the optical flow equation. We note however that using the optical flow formula for demons leads to particularly good results for the experiments presented in the last section.

4.6. Demons 2: demons in contours only

In matching methods, contour points are generally more important than other points in the image. Furthermore, using only the contour points of S leads to a faster algorithm. We used, for the possible variants:

- (i) the Canny-Deriche edge detector (see Deriche, 1986; Monga *et al.*, 1990) to extract D_s from S ;
- (ii) global transforms (rigid, affine) for $\mathcal{T}$, determined by least squares [like in Besl and McKay (1992)]. It could be improved by extended Kalman filtering (EKF; see Ayache, 1991) to reject outliers. We have also tested more flexible deformations using warping techniques (see Bookstein and Green, 1993; Szeliski and Lavallée, 1994; Szeliski and Coughlan, 1994; Declerck *et al.*, 1995) which extend the deformation defined at irregularly distributed points D_s to the whole image S ;

Figure 10. A first type of K functions.Figure 11. A second type of K functions.

- (iii) (tri-)linear interpolation;
- (iv) the demon behaviour is detailed below.

We now describe the demon force computation in this case. As only a subset D_s of S is considered, more sophisticated demons can be designed than in the previous method relying on the optical flow formula. Let P be a contour point in S , of intensity $s(P)$, and $\vec{n} = \vec{\nabla}(P)/\|\vec{\nabla}(P)\|$ be the oriented normal of the contour (from inside to outside). By interpolation, we compute the value $s_{out} = s(P + k\vec{n})$ and $s_{in} = s(P - k\vec{n})$ (where k is a constant, possibly 1).

P is therefore the interface between a region whose estimated intensity is s_{in} and a region whose intensity is s_{out} . The demon force is

$$\vec{f}(P) = K_{s_{in}, s_{out}}(m(P'))\vec{n} \quad (5)$$

where $P' = T_i^{-1}(P)$ and $K_{s_{in}, s_{out}}(m)$ is a function that we now detail. By definition, the demon pushes inside ($K < 0$) when P' is labelled 'inside' in M , and outside ($K > 0$) when P' is 'outside' in M .

We have considered functions K where the probability of being 'inside' is very high if $m(P') \leq s_{in}$ (then we impose $K = -k$), and the probability of being 'outside' is very high for $m(P') \geq s_{out}$ (then $K = k$). We have used the simple piecewise linear formula of Figure 10 for $K(m)$.

We have considered another type of function K , where when m is too different from s_{in} and s_{out} , P' is neither considered 'inside' nor 'outside', but simply 'irrelevant' ($K = 0$). We assume that the intensities of inside points are s_{in} , and outside points are s_{out} , but that the measure s is corrupted with Gaussian noise. We use therefore for K the sum of two Gaussian distributions $K = K^- + K^+$ (see Figure 11); $K^-(m)$ is negative and centred at s_{in} , with $K^-(s_{in}) = -k$, and $K^+(m)$ is positive, centred at s_{out} , with $K^+(s_{out}) = k$.

We still have two free parameters (σ^- and σ^+), which can be related to the level of noise in the images S and M .

Anisotropic diffusion filters mentioned in the introduction could be used efficiently to pre-process S and M and reduce the noise level, while preserving the interfaces. Both of the $K(m)$ functions described gave us good results and we kept the first one for the experiments in the present paper, but a more careful study is needed to determine when to use one or the other.

4.7. Demons 3: already segmented images

Here, the nature of the problem is different. We assume that the images S and M are already segmented, that is, for each point in image S or in image M , we have a label which refers to a given structure (for example a given organ in the case of medical images). Even though the images are segmented, we can still search for a geometric point to point correspondence between S and M . Within the general scheme, we use:

- (i) the set of interface points of S for D_s , that is, points between adjacent voxels whose labels are different (Figure 12);
- (ii) deformations similar to method 2 (rigid, affine, warping);
- (iii) no interpolation, but the label of the closest voxel in M (using linear interpolation between labels is generally meaningless);
- (iv) demons with forces of constant magnitude, the orientations are defined according to the labels (see below).

We now detail the demons behaviour. Two adjacent voxels are linked by a segment of length 1, $\sqrt{2}$, $\sqrt{3}$ depending on their connectivity (6, 18, 26) (see Figure 12). Let A and B be two adjacent voxels whose labels $s_{in} = s(A)$ and $s_{out} = s(B)$ are different, and $\vec{d} = \vec{AB}$. We create a demon in $P = (A + B)/2$.

The 'force' $\vec{f}(P)$ of the demon is determined by s_{in} , s_{out} and the label $m(P')$, where $P' = T_i^{-1}(P)$, then

- if $m \neq s_{in}$ and $m \neq s_{out}$, $\vec{f} = 0$
- if $m = s_{in}$, $\vec{f} = -k\vec{d}$
- if $m = s_{out}$, $\vec{f} = k\vec{d}$.




Figure 12. Demons in a labelled image.

For a given iteration, k is a constant magnitude, which can vary at each iteration depending on the convergence strategy. For example, k can be set to k_0 voxel size and decreased linearly to zero during iterations. Accordingly, the smoothing parameter σ can be decreased from σ_0 to 0, to ensure the convergence of the algorithm. The ratio k_0/σ_0 depends on the application: it controls the rigidity of the diffusing model. The advantage of using pre-segmented images is that the interface points are labelled (s_{in} , s_{out}), and interact with the structures of M which are either s_{in} or s_{out} , but do not interact at all with other structures.

For automatic registration with a digital anatomical atlas, such a method is very suitable: a physician can spend a lot of time labelling all the voxels in the image S of a reference patient (see Figure 26), then a method using an hybrid of demons 2 and 3 can be applied to automatically match and label, at the same time, the image M of a new patient.

4.8. 'Bijectivity' of the 3-D deformation

We do not enter into the debate of what bijectivity means in the case of discrete deformation fields, but one important constraint for 3-D deformations is the possibility of inverting these deformations. With the previously described implementations, there is no guarantee that this inversion can be performed [a problem which is raised in Bro-Nielsen and Gramkow (1996)]. The bijectivity property is ensured to some extent in physically based techniques [such as those of Christensen *et al.* (1994a) or Bro-Nielsen and Gramkow (1996)] by a local control of the Jacobian (the determinant $|J|$ of the Jacobian matrix of the deformation), which is forced to be positive ($|J| > 0$). This type of bijectivity should not be confused with the 'one-to-one' correspondence of viscous fluids (see Christensen *et al.*, 1996), corresponding to incompressible matter ($|J| = 1$) which might be useful for specific types of problems.

However, the elastic fluid is probably a better model for inter-patients matching than viscous fluid^a although there

^aFirstly, the two patients are likely to have different sizes.




Figure 13. MR head slice: original (left), deformed (right).




Figure 14. Deformations corrected with demons 1 (left), with the contours of the original image superimposed (right).

are no grounds to relate these two mechanical models to anatomical differences. In the following applications we will be mainly concerned the elastic deformations.

Maintaining a constraint such that $|J| > 0$ (elasticity) is also time consuming. To simply verify that a discrete 3-D deformation field has a positive Jacobian, one can consider each cubic cell (or 8-cell) composed of 8 voxels, and verify that the deformed 8-cell is still positively oriented. One way to do this is to decompose the 8-cell into five tetrahedra, and verify that the deformed tetrahedra are still positively oriented (by computing the determinant of the matrix formed by three edges), which means about 100 floating-point operations per voxel and per iteration, which would make this measurement a bottle neck of our method, without even solving the problem (which is to impose a positive Jacobian).

We have implemented a simpler solution (see also Burr, 1981): at each iteration, we compute the direct deformation T_{12} (from I_1 to I_2 , for a regular grid in I_1), and the reverse deformation T_{21} (from I_2 to I_1 , for a regular grid in I_2), and also the residual deformation $R = T_{21} \circ T_{12}$. We then compensate equally for this residual deformation in T_{12} and T_{21} , by removing 'half' of R from each deformation. More precisely (Figure 15), a point P in image I_1 is projected into $P' = T_{12}(P)$ in I_2 , and, by interpolating T_{21} at the floating




Figure 15. The bijectivity is maintained by the computation of $R = T_{21} \circ T_{12}$ and the redistribution of the residual at each iteration of the diffusing model.

point position P' , re-projected back into $P' = T_{21}(P')$. Hence the residual vectors PP' constitute a vector field for a regular grid in I_1 . $1/2PP'$ is directly subtracted from the vector field T_{12} and the other half $1/2PP'$, is projected into I_2 by using T_{21} and added to it, to give a new T_{21} . After this operation, $T_{21} \circ T_{12}$ is very close to identity, and we have verified experimentally that this residual is less than a voxel on average, when this operation is performed at each iteration of the diffusing model.

This method requires only twice the computation of the one way transformation plus only about 20 additional floating point operations per voxel per iteration. Besides bijectivity, the advantage of this method is that it provides at the end not only the direct transformation T_{12} , but also an inverse transformation T_{21} very similar to the 'ideal' T_{12}^{-1} .

Experimentally, this bijective implementation has proved to be much more robust than the simple one for anatomically homologous organs (for example right hands with five fingers each). It tends to preserve the morphology of each specimen. For example, in the temporal lobes of the brain, the number of gyri for different patients can be different, but are still preserved: only low frequency morphometrical differences are compensated for, which allows for the comparison of local morphological differences (see Guimond *et al.*, 1997).

4.9. Partial conclusion

To some extent, diffusing models (principally demons 1) can be connected back to more classical concepts of matching. However, the concept can serve to foster the exploration of new matching algorithms by the many possible options. Diffusing models are not constrained to use forces derived from a known potential field or from a similarity measure; the concept provides a global understanding of how things are working: the deformable model diffuses through the scene interfaces in order to have more points of the model with




Figure 16. Original (left), corrected with demons 2 (right).

polarities or types similar to those at corresponding places in the scene.

5. EXPERIMENTS

We now present some experimental results of non-rigid matching obtained with implementations 1 and 2 of diffusing models, using synthesized and real data. We have used the same set of parameters for all the experiments: $\sigma = 1$ for the Gaussian filtering, four iterations at the finer scale of the pyramid, and four levels of resolution with respect to the multi-scale processing.

5.1. Experiments with simulated deformations

First, we present some results with a synthetic deformation in 2-D. Figure 13 shows a typical slice from a magnetic resonance (MR) brain image, and the same image after a deformation based on sinusoidal functions with a spatial period of 32 voxels.

We applied demons 1 to get the image of Figure 14 (CPU time, 30 s), where we have also superimposed the contours of the original image to verify visually the quality of the match, which is quite good, when compared with the large amount of deformation.

We have also successfully tested demons 2 (with warping for the deformations) on this image (see Figure 16). The quality of the match is also good, but slightly less satisfying than with demons 1. The difference, however, is very small.

The computation time is only 9 s in 2-D (three times faster, edge extraction included). This is a general result of our experiments: demons 2 is faster but less accurate than demons 1, because the density of demons is reduced in demons 2. The difference in quality is generally on the order of the voxel size. The choice of the method depends on a compromise between quality and computation time: none is better *per se*.

Lastly, we performed a test with the implementation demons 3, using warping for deformations. In this experiment only already labelled images are taken into account (the




Figure 17. Top left, original label image (I_1); top right, deformed label image (I_2). Bottom left, I_2 deformed toward I_1 using the implementation 'demos 3'; bottom right, deformed I_2 with a superimposition of I_1 contours.

grey-level images are ignored, see Figure 17). This demonstrates that polarity (or labelling) is sufficient to perform successfully the non-rigid matching of two images.

5.2. 3-D image sequences displaying cardiac motions

The two major drawbacks of non-rigid matching are reduced for image sequence analysis: the acquisition device being the same, the intensities of the scanned object do not vary much from frame to frame, and the position of the object in one frame is a good initialization for the next one.

Many deformable surface techniques have been proposed to study cardiac sequences, such as in Bardin *et al.* (1994), where a parametric surface is fit to the data and the deformation is analysed using modal decomposition. A very interesting study about the relevance of surface deformable models to heart motion analysis can be found in Shi *et al.* (1995), based on clinical experiments with dogs. However, these techniques require a segmentation of the object surface and the problem of extending the deformation to the whole volume remains to be solved.

In the following experiment, we present an example of the analysis of a CT scan of the heartbeat of a dog, acquired with the Mayo Clinic DSR (courtesy of Dr Richard Robb). To consider extreme conditions, we match directly




Figure 18. Corresponding diastolic and systolic slice before matching (dog).




Figure 19. Corresponding diastolic and systolic slices after 3-D matching and re-sampling.

the diastolic and systolic 3-D images. In Figure 18, we present two corresponding slices before elastic matching, and in Figure 19, the same slices after 3-D matching and re-sampling. With 3-D tools, we verified visually the quality of the match in all the other parts of the volumes. To get an idea of the smoothness of the non-rigid transform, we artificially added a regular grid tag to the image before deformation (see Figure 20). The computation time, for the matching of the two 100^3 voxels images was about 5 min on a DEC alpha workstation.

Figures 21 and 22 present similar experiments for nuclear medicine 64^3 SPECT images (a human heart), with a CPU time of about 1 min. Figure 23 is a 3-D visualization of the same results.

We have also analysed whole cardiac gated sequences (eight 3-D images for a whole cardiac cycle). According to the physics of the SPECT acquisition, we have added two constraints: the cyclicity of the sequence, which is imposed by another set of iteration taking the whole sequence into account and performing temporal filtering, and the conservation of the radioactive matter. This last constraint




Figure 20. Artificially tagged image (2-D grid of planes) to show the 3-D deformation.




Figure 21. Corresponding diastolic and systolic slice before matching (SPECT, human).




Figure 22. Corresponding diastolic and systolic slice after 3-D matching and re-sampling (SPECT, human).

states that, when the myocardium contracts, the density of radioactive material is increased, and thus the image intensity is correspondingly enhanced. This can be easily modelled within the design of the demons, by multiplying $m(P_i)$ by the Jacobian of the transform, which represents the local variation of the density. The analysis of one whole time series takes only about 10 min.




Figure 23. 3-D visualization of the left ventricle of the heart in SPECT images. Left is the diastole, middle is the systole, right is the surface of the diastole deformed toward the systole by the 3-D vector field obtained with demons 1.




Figure 24. Two slices of two different patients ($256 \times 256 \times 128$ voxels).

To conclude, we have an automatic, fast algorithm to estimate a deformation field between frames in 3-D image sequences, that does not necessitate a segmentation of the object, which is advantageous for SPECT images. A good point is that the physics of nuclear medicine acquisition can be taken into account in our scheme. The challenge now is to extract from this 4-D flow field the parameters that precisely characterize a pathology. Of course, these examples are given only to illustrate potential applications of diffusing models, but much harder work remains to be performed to validate the results from a medical viewpoint and to lead to useful clinical applications.

5.3. Inter-patients matching: the 3-D case

Matching the images of two different patients is very important for medical image applications. We present in Figure 24 slices from two different patients, extracted from their 3-D $256 \times 256 \times 128$ MR images. We applied demons 1 and demons 2, but in 3-D, to the two blocks of data. The computation time, with the bijective computation, is only about 30 min on a Dec Alpha workstation, which is reasonably fast, with respect to the size of the data under




Figure 25. The two different patients after automatic matching.




Figure 26. The deformation field is used to adapt the 3-D anatomical atlas presented in Figure 27 to the new patient. Top row: one slice of the MR image of the originally labelled subject along with the corresponding labelling. Bottom row: left, a slice of a new subject and right, the inferred labelling obtained by the application of the deformation field to the labelled reference subject. The inference is visually satisfying, up to the local morphological differences.

consideration. Figure 25 presents the result of the matching: differences between the deformed patient and the reference patient is now hard to perceive. Only anatomical details which are topologically different between the two patients are still different. Again, demons 1 performs slightly better than demons 2, but at the expense of more computation time.

Because we have a 3-D completely labelled version of the reference image (courtesy of Dr Kikinis, Brigham and




Figure 27. A 3-D display of the segmented image of the brain, courtesy of Dr Ron Kikinis, and rendered by Gerard Subsol.

Women's Hospital, Boston, see Figure 27), we are able to apply the computed deformation field to it, and superimpose result on the image of the new patient (Figure 26), providing an automatic labelling of the patient. We note, however, that the inferred labels are mostly correct only where there are no local morphological differences between the two subjects. We hope to improve the automatic labelling with a more local, final pass of segmentation, taking the tissue types into account, such as in the segmentation presented by Wells III *et al.* (1995). Dr Benoît Dawant, at the University of Vanderbilt, TN, has used our method to compare the automatic labelling of images obtained by inference thanks to the deformation field that we can produce with several independent manual segmentation (see Dawant *et al.*, 1998).

5.4. A database of brain images

Finally, we present the results of the automatic non-rigid matching performed on a database of brain images. In order to perform this experiment, we have modified our method to estimate automatically a global bias and gain between the two images to be matched. We do this simply by incorporating into the iterative scheme an estimation of the two parameters based on the fitting of a line (with outliers rejection) to the transfer function between the two images, which gives, for each intensity in one image, the average intensity in the other image, based on the correspondence T_i




Figure 28. One slice (out of a 128) of the original images of nine different patients. Shapes and intensities are very different.

estimated at iteration i . We can see in Figure 28 a zoom on the central part of the brain in nine different patients, which shows both the differences in intensities and in shapes. Figure 29 shows the result of the automatic matching: after re-sampling, morphometrical differences between patients are hard to perceive, while local morphological differences still exist.

Figure 30 shows several slices of the reference image, and the average of the 10 images before and after matching with this reference image. The whole process is entirely automatic, and it takes about 5 min CPU time to match two 128^3 3-D images. It shows the quality of the matching, and also that the central parts of the brains are indeed more stable anatomically than the peripheral parts (which appear fuzzier).

5.5. Other medical applications

Our non-rigid matching method has been embedded into a more general system to automatically explore large databases of 3-D medical images in order to find similarities between patients, and to compute average representative specimens (see Guimond *et al.*, 1997). It is also being used for an ongoing large clinical study about Schizophrenia (E.C. project Biomorph) to measure and compare brain asymmetry between individuals. It has also been used for the computation of average activation maps in SPECT images (see Migneco *et al.*, 1997): similar ongoing studies are under way




Figure 29. The same nine patients after non-rigid matching, re-sampling and intensity correction. The computation is performed entirely in 3-D. Note that the morphometrical differences are compensated for, but not the local morphological differences. Our technique has been already applied, with qualitatively identical results, to more than a hundred cases.

on a larger scale for fMRI and for PET activation images using competing non-rigid matching techniques (see Collins *et al.*, 1994). Finally, we have used our method to study and quantify the evolution of multiple sclerosis plaques in time series of 3-D brain MRI (see Thirion and Calmon, 1997).

5.6. Validation with real images

The major problem of inter-patients matching is how to validate the different methods. In Thirion *et al.* (1996), we carefully compared three different inter-patients matching techniques with a database of CT scans of dry skulls: one matching method was based on the feature points defined by an anatomist, one was based on deformable 3-D crest lines, and the last technique was the demons-based method. The main results of this study were:

- the three matching methods gave mutually consistent results, with an average point to point distance of about 3 voxels;
- this average distance was reduced to about 2 voxels when used to compute average patients' models;
- the average influence of the choice of a reference patient on the resulting average model computations was less than 1 voxel.




Figure 30. Left, five consecutive slices of the reference image; right, the average of 10 patients without matching; centre is the average of the 10 patients after matching and re-sampling. The fuzziness of the edges of the brain in the middle images express the remaining local morphological differences between the subjects, which have been reduced.

As mentioned, a study being performed at the University of Vanderbilt with our tool provides another independent validation of the application to automatic segmentation.

VIDEO DEMONSTRATIONS

Two video sequences illustrating the 3-D deformation field measured between two patients can be obtained from: http://www.inria.fr/epidaure/Gallery/atlas_movie.html and [atlas_movie2.html](http://www.inria.fr/epidaure/Gallery/atlas_movie2.html).

6. CONCLUSION

We have proposed the concept of a diffusing model for image-to-image matching, and the related concept of demons to emphasize the role of polarity information in matching methods. A diffusing model corresponds to the ‘diffusion’ of a deformable model grid into the image of a scene, where the boundaries of the scanned objects are considered to be semi-permeable membranes, filtering the points of the model according to their polarity (inside or outside). We have also shown that optical flow can be considered as an intermediate step between diffusing models and more classical matching methods based on attraction.

Diffusing models and demons can be tools for examining matching methods from an original viewpoint, to generate ideas leading to new methods, or to improve existing ones by taking into account polarity, whose importance is generally underestimated. Considering matching as a diffusion process might also be a way to introduce new ideas from thermodynamics into image matching.

With respect to applications, we have shown different ways to apply these ideas to the medical field, such as organ tracking or 3-D inter-patients matching. Developing and testing such 3-D non-rigid matching tools is now a very important challenge for diagnosis, for surgery planning and control, and more generally, for a better understanding of the human anatomy and its variability.

ACKNOWLEDGEMENTS

I wish to thank Nicholas Ayache, Michael Brady, Morten Bro Nielsen, Jérôme Declerck, Gérard Subsol, Xavier Pennec, Hervé Delingette, Sylvain Prima, Alexandre Guimond and the Epidaure people in general for stimulating discussions and specific comments about the ideas presented in this paper. Special thanks for Gregoire Malandain who implemented the Canny-Deriche-Monga edge detection, and to Allen Sanderson for his careful proof-reading of early versions of this paper.

REFERENCES

- Aggarwal, J. K. and Nandhakumar, N. (1988) On the computation of motion from sequences of images—a review. *Proc. IEEE*, 76, 917–935.
- Anandan, P. (1989) A computational framework and an algorithm for the measurement of visual motion. *Int. J. Comput. Vision*, 2, 283–310.
- Ayache, N. (1991) *Artificial Vision for Mobile Robots: Stereo Vision and Multisensory Perception*. MIT Press, Cambridge, MA.
- Bajcsy, R and Kovacic, S. (1989) Multiresolution elastic matching. *Comp. Vision, Graphics Image Processing*, 46, 1–21.
- Bardinet, E., Cohen, L. D. and Ayache, N. (1994) Fitting of iso-surfaces using superquadrics and free-form deformations. In *Proc. IEEE Workshop on Biomedical Images Analysis (WBIA'94)*, Seattle, WA.
- Barron, J. L., Fleet, D. J. and Beauchemin, S. S. (1994) System and experiment: performance of optical flow techniques. *Int. J. Comp. Vision*, 12, 43–77.
- Benayoun, S. and Ayache, N. (1995) Dense non-rigid motion estimation in sequences of 3D images using differential constraints. *Int. J. Comput. Vision*, 26, No 1.
- Besl, P. J. and McKay, N. D. (1992) A method for registration of 3-d shapes. *IEEE PAMI*, 14, 239–255.
- Blake, A. and Yuille, A. (1992) *Active Vision*. MIT Press, Cambridge, MA.
- Bookstein, F. L. and Green, W. D. K. (1993) A feature space for derivatives of deformations. In Barrett, H. H. and Gmitro, A. F. (eds), *Information Processing in Medical Imaging, (IPMI'93), Lecture Notes in Computer Science*, Vol. 687, pp. 1–16. Springer-Verlag, Flagstaff, AZ.
- Borgefors, G. (1988) Hierarchical chamfer matching: a parametric edge matching algorithm. *IEEE Trans. PAMI*, 10, 849–865.
- Bro-Nielsen, M. and Gramkow, C. (1996) Fast fluid registration of medical images. In Höhne, K. H. and Kikinis, R. (eds), *Visualization in Biomedical Computing VBC'96, Lecture Notes in Computer Science*, Vol. 1131, pp. 267–276. Springer-Verlag, Hamburg.
- Burr, D. J. (1981) A dynamic model for image registration. *Comp. Graphics Image Processing*, 15, 102–112.
- Catté, F., Lions, P.-L., Morel, J.-M. and Coll, T. (1992) Image selective smoothing and edge detection by nonlinear diffusion. *SIAM J. Numer. Anal.*, 29, 182–193.
- Chakraborty, A., Staib, L. H. and Duncan, J. S (1996) Deformable boundary finding in medical images by integrating gradient and region information. *IEEE Trans. Med. Imag.*, 15, 859–870.
- Christensen, G. E., Miller, M. I. and Vannier, M. (1994a). A 3D deformable magnetic resonance textbook based on elasticity. In *AAAI Spring Symp. Series: Applications of Computer Vision in Medical Image Processing*, pp. 153–156. Stanford University, Stanford, CA.
- Christensen, G. E., Rabbitt, R. D. and Miller, M. I. (1994b). 3D brain mapping using a deformable neuroanatomy. *Phys. Med. Biol.*, 39, 609–618.

- Christensen, G. E., Miller, M. I. and Vannier, M. W. (1996) Individualizing neuro-anatomical atlases using a massively parallel computer. *IEEE Computer*, 32–38.
- Collins, D. L., Neelin, P., Peters, T. M. and Evans, A. C. (1994) Automatic 3D intersubject registration of MR volumetric data in standardized Talairach space. *J. Comp. Assis. Tomogr.*, 18, 192–205.
- Davatzikos, C. (1996) Nonlinear registration of brain images using deformable models. In *Workshop on Mathematical Methods in Biomedical Image Analysis*, pp. 94–103, San Francisco, CA. IEEE.
- Dawant, B. M., Thirion, J.-P., Maes, F., Vandermeulen, D. and Demaerel, P. (1998) Automatic 3d segmentation of internal structures of the head in MR images using a combination of similarity and free form transformations. In *SPIE Medical Imaging 1998*, Vol. 3338, San Diego, CA.
- Declerck, J., Subsol, G., Thirion, J.-P. and Ayache, N. (1995) Automatic retrieval of anatomical structures in 3D medical images. In *Computer Vision, Virtual Reality and Robotics in Medicine, Lecture Notes in Computer Science*, Vol. 905, pp. 153–162. Springer-Verlag, Nice.
- Deriche, R. (1986) Using Canny's criteria to derive a recursively implemented optimal edge detector. *Int. J. Comp. Vision*, 6, 167–187.
- Faugeras, O. (1993) *Three-Dimensional Computer Vision: a Geometric Viewpoint*. The MIT press, Cambridge, MA.
- Feldmar, J. and Ayache, N. (1994) Locally affine registration of free-form surfaces. In *IEEE Conf. on Computer Vision and Image Processing (CVPR'94)*, pp. 496–500, Seattle, WA.
- Gee, J. C., Reivich, M. and Bajcsy, R. (1993) Elastically deforming 3D atlas to match anatomical brain images. *J. Comp. Assis. Tomogr.*, 17, 225–236.
- Gerig, G., Kübler, O., Kikinis, R. and Jolesz, F. A. (1992) Nonlinear anisotropic filtering of MRI data. *IEEE Trans. Med. Imag.*, 11, 221–232.
- Guimond, A., Subsol, G. and Thirion, J.-P. (1997) Automatic MRI database exploration and applications. *Int. J. Patt. Recogn. Artif. Intell.*, 11, 1345–1365.
- Horn, B. P. K. and Schunck, B. G. (1981) Determining optical flow. *Artific. Intell.*, 17, 185–203.
- Kass, M., Witkin, A. and Terzopoulos, D. (1987) Snakes: active contour models. *Int. J. Comp. Vision*, 1, 312–331.
- Kichenassamy, S., Kumar, A., Olver, P., Tannenbaum, A. and Yezzi, A. (1995) Gradient flows and geometric active contour models. In *5th Int. Conf. on Computer Vision (ICCV'95)*, pp. 810–815, Cambridge, MA. IEEE.
- Kimia, B. and Siddiqi, K. (1994) Geometric heat equation and nonlinear diffusion of shapes and images. In *Conf. on Computer Vision and Pattern Recognition (CVPR'94)*, pp. 113–120, Seattle, WA, USA.
- Maes, F., Collignon, A., Vandermeulen, D., Marchal, G. and Suetens, P. (1997) Multimodality image registration by maximization of mutual information. *IEEE Trans. Med. Imag.*, 16, 187–198.
- Malandain, G., Fernández-Vidal, S. and Rocchisani, J. M. (1994) Improving registration of 3-D medical images using a mechanical based method. In *Proc. 3rd Eur. Conf. on Computer Vision (ECCV '94)*, pp. 131–136, Stockholm, Sweden. *Lecture Notes in Computer Science*, Vol. 801.
- Migneco, O., Thirion, J.-P., Benoit, M., Malandain, G., Robert, P., Ayache, N. and Darcourt, J. (1997) Spect memory activation studies thanks to non-rigid automated 3D image registration. In *The Joint CVRMed-MRCAS Conf.*, pp. 487–490, Grenoble.
- Monga, O., Deriche, D., Malandain, G. and Cocquerez, J. P. (1990) 3D edge detection by recursive filtering and edge tracking. In *Int. Conf. on Pattern Recognition (ICPR)*, Atlantic City, USA.
- Perona, P. and Malik, J. (1990) Scale-space and edge detection using anisotropic diffusion. *IEEE PAMI*, PAMI-12, 629–639.
- Radeva, P., Serrat, J. and Martí, E. (1995) A snake for model-based segmentation. In *5th Int. Conf. on Computer Vision (ICCV'95)*, pp. 816–821, Cambridge, MA. IEEE.
- Rohr, K., Stiehl, H. S., Sprengel, R., Beil, W., Buzug, T. M., Weese, J. and Kuhn, M. H. (1996) Point-based elastic registration of medical image data using approximating thin-plate splines. In Höhne, K. H. and Kikinis, R. (eds), *Visualization in Biomedical Computing VBC'96, Lec. Notes in Computer Science*, Vol. 1131, pp. 297–306. Springer-Verlag, Hamburg.
- Ronfard, R. (1994) Region-based strategies for active contour models. *Int. J. Comp. Vision*, 13, 229–251.
- Shi, P., Robinson, G., Chakraborty, A., Staib, L., Constable, R., Sinusas, A. and Duncan, J. (1995) A unified framework to assess myocardial function from 4d images. In *Computer Vision, Virtual Reality and Robotics in Medicine, Lecture Notes in Computer Science*, Vol. 905, pp. 327–337. Springer-Verlag, Nice.
- Simoncelli, E. P., Adelson, E. H. and Heeger, D. J. (1991) Probability distributions of optical flow. In *Computer Vision and Pattern Recognition, CVPR'91*, pp. 310–315, Lahaina, Maui, Hawaii. IEEE.
- Szeliski, R. and Coughlan, J. (1994) Hierarchical spline-based image registration. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR'94)*, pp. 194–201, Seattle, WA.
- Szeliski, R. and Lavallée, S. (1994) Matching 3-d anatomical surfaces with non-rigid volumetric deformations. In *Proc. IEEE Workshop on Biomedical Images Analysis (WBIA'94)*, Seattle, WA. Also in AAAI 1994 Spring Symposium Series. Application of Computer Vision in Medical Image Processing, Stanford University, 1994.
- Thirion, J.-P. (1994) Extremal points: definition and application to 3D image registration. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pp. 587–592, Seattle, WA.
- Thirion, J.-P. (1995) Fast non-rigid matching of 3D medical images. In *Medical Robotics and Computer Aided Surgery (MRCAS'95)*, pp. 47–54, Baltimore.
- Thirion, J.-P. (1996) Non-rigid matching using demons. In *Computer Vision and Pattern Recognition, CVPR'96*, San Francisco, CA.

- Thirion, J.-P. and Calmon, G. (1997) Measuring lesion growth from 3D medical images. In *Nonrigid and Articulated Motion Workshop (NAM'97)*, Puerto Rico. IEEE.
- Thirion, J.-P., Subsol, G. and Dean, D. (1996) Cross validation of three inter-patients matching methods. In *Visualization in Biomedical Computing, VBC'96, Lecture Notes in Computer Science*, Vol. 1131, pp. 327–336, Hamburg.
- Thompson, P. and Toga, A. W. (1996) A surface-based technique for warping three-dimensional images of the brain. *IEEE Trans. Med. Imag.*, 15, 402–417.
- Viola, P. and Wells, W. M. III (1995) Alignment by maximization of mutual information. In *5th Int. Conf. on Computer Vision, ICCV'95*, pp. 16–23, Cambridge, MA. IEEE.
- Wells III, W. M., Grimson, W. E.L., Kikinis, R. and Jolesz, F. A. (1995) Adaptive segmentation of mri data. In *Computer Vision, Virtual Reality and Robotics in Medicine, Lecture Notes in Computer Science*, Vol. 905, pp. 59–69. Springer-Verlag, Nice.
- Zhang, Z. (1992) *Iterative Point Matching for Registration of Free-form Curves*. Technical Report 1658, INRIA.