

Tracking brain deformations in time-sequences of 3D US images

X. Pennec, P. Cachier and N. Ayache

*EPIDAURE, INRIA Sophia Antipolis, 2004 Rte des Lucioles, BP93, 06902 Sophia
Antipolis Cedex*

Abstract

During a neurosurgical intervention, the brain tissues shift and warp. In order to keep an accurate positioning of the surgical instruments, one has to estimate this deformation from intra-operative images. We present in this article a feasibility study of a tracking tool based on intra-operative 3D ultrasound (US) image sequences. The automatic processing of this kind of images is of great interest for the development of innovative and low-cost image guided surgery tools. The difficulty relies both in the complex nature of the ultrasound image, and in the amount of data to be treated as fast as possible.

Key words: Non rigid registration, 3D ultrasound images, Tracking

1 Introduction

The use of stereotactic systems is now a quite standard procedure for neurosurgery. However, these systems are only accurate under the assumption that the skull and the brain move together as a unique rigid body during surgery. In practice, relative motion of the brain with respect to the skull (also called brain shift) occurs, mainly due to tumor resection, cerebrospinal fluid drainage, hemorrhage or even the use of diuretics. Furthermore, this motion is likely to increase with the size of the skull opening and the duration of the operation.

Email address: {Xavier.Pennec, Pascal.Cachier, Nicholas.Ayache}@sophia.inria.fr
(X. Pennec, P. Cachier and N. Ayache).

URL: <http://www-sop.inria.fr/epidaure/> (X. Pennec, P. Cachier and N. Ayache).

Over the last years, the development of real-time 3D ultrasound (US) imaging has revealed a number of potential applications in image-guided surgery as an alternative approach to open MR and intra-interventional CT. The major advantages of 3D US over existing intra-operative imaging techniques are its comparatively low cost and simplicity of use. However, the automatic processing of US images has not gained the same degree of development as other medical imaging modalities, probably due to the low signal-to-noise ratio of US images.

1.1 Context

We present in this article a feasibility study of a tracking tool for brain deformations based on intra-operative 3D ultrasound (US) image sequences. This work was performed within the framework of the European project ROBOSCOPE, a collaboration between The Fraunhofer Institute (Germany), Fokker Control System (Netherlands), Imperial College (UK), INRIA (France), ISM-Salzburg and Kretz Technik (Austria). The goal of the whole project is to assist neuro-surgical operations using real-time 3D ultrasound images and a robotic manipulator arm (fig. 1).

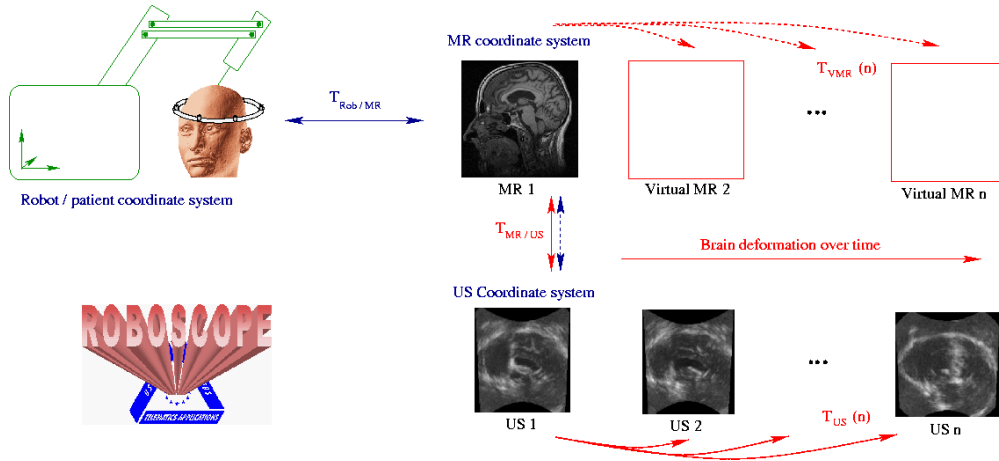


Fig. 1. Overview of the image analysis part of the Roboscope project.

The operation is planned on a pre-operative MRI (MR1) and 3D US images are acquired during surgery to track in real time the deformation of anatomical structures. The first US image (US1) is acquired with dura mater still closed and a rigid registration with the preoperative MR is performed. This allows to relate the MR and the US coordinate systems and possibly to correct for the distortions of the US acquisition device. Then, brain deformations are tracked in the time-sequence of per-operative US images. From these deformations, one can update the preoperative plan and synthesize a virtual MR image that matches the current brain anatomy.

1.2 MR/US registration

The idea of MR/US registration is already present in (Trobaugh et al., 1994a; Trobaugh et al., 1994b) where the US probe is calibrated (i.e. registered to the surgical space) and then tracked using an optical device. The MR/US registration is then obtained using the registration of the MR image to the surgical space by standard stereotactic neurosurgical procedures. (Richard et al., 1999) improved this method by designing a real-time low-cost US imaging system based on a PCI bus. In (Pagoulatos et al., 1999), the tracking of the US probe is performed with a DC magnetic position sensor. In (Erbe et al., 1996), the registration is performed by interactively delineating corresponding surfaces in all images and a visual rigid fitting of the surfaces using a 6D space-mouse. In (Hata et al., 1994), the outlines of the 2D US image is registered to the MR surface using a Chamfer matching technique. All these techniques only perform a rigid registration of the MR and the US images.

For a non rigid registration (i.e. a brain shift estimation), we have to turn to (Gobbi et al., 1999; Gobbi et al., 2000; Comeau et al., 2000), where the 2D US probe is still optically and rigidly tracked but the corresponding MR slice is displayed to the user who marks corresponding points on MR and US slices. Then, a thin plate spline warp is computed to determine the brain shift. This method is also developed in (Bucholz et al., 1997) with the possibility of using 3D US images and a deformation computed using a spring model instead of splines. More recently, Ionescu *et al* (Ionescu et al., 1999) registered US with Computed Tomography (CT) data after automatically extracting contours from the US using watershed segmentation. In these studies, there is no processing of a full time sequence of US images : the brain shift estimation is limited to a few samples at given time-points as the user interaction is required at least to define the landmarks.

Recently, an automatic rigid registration of MR and US images was presented (Roche et al., 2001; Roche et al., 2000; Pennec et al., 2001a). This work is based on image intensities and does not rely on feature extraction. However, the estimated motion remains limited to rigid or possibly affine transformations. Up to our knowledge, only (King et al., 2000) deals with an automatic non-rigid MR/US registration: the idea is to register a surface extracted from the MR image to the 3D US image using a combination of the US intensity and the norm of its gradient in a Bayesian framework. The registration is quite fast (about 5mn), even if the compounding of the 3D US and the computation of its gradient takes about one hour. However, experiments are presented only on phantom data and our experience (see section 3) is that real US images may lead to quite different results.

1.3 Tracking methods in sequences of US images

Since non-rigid MR/US registration is a difficult problem, we chose to split it into two subproblems: first a rigid MR/US registration is performed with dura matter still closed (there is no brain shift yet), for instance using the approach of (Roche et al., 2000). Then we look for the non-rigid motion within the US time-sequence. In the literature, we found a small number of articles on the registration of 3D US images. (Strintzis and Kokkinidis, 1997) use a maximum-likelihood approach to deduce a similarity measure for ultrasound images corrupted by Rayleigh noise and a block-matching strategy to recover the rigid motion. In (Rohling et al., 1998), the correlation of the norm of the image gradient is used as the similarity measure to rigidly register two US images in replacement of the landmark-based RANSAC registration of (Rohling et al., 1997). However, these methods only deal with rigid motion and consider only two images, eluding the tracking problem.

One has to move to cardiac application to find some real tracking of non-rigid motion in US images. In (Papademetris et al., 1999), the endo- and epicardial surfaces are interactively segmented on each 2D image plane. Then, a shape-memory deformable model determines the correspondences between the points of the 3D surfaces of successive images. These correspondences are used to update an anisotropic linear elastic model (finite element mesh). The approach is appealing but relies once again on an interactive segmentation. In (Sanchez-Ortiz et al., 2000), a combination of feature point extraction (phase-based boundaries) and a multi-scale fuzzy clustering algorithm (classifying the very low intensities of intra-ventricular pixels) is used to segment the surface of the left ventricular cavity. This process is done in 2D+T and then reconstructed in 3D. Thus it exploits the whole sequence before tracking the motion itself, which is not possible for our application. These two methods are well suited for the shape of the cardiac ventricle using dedicated surface models. If they could be adapted to the brain ventricles, it seems difficult to extend them to the tracking of the volumetric deformations of the whole brain.

1.4 Intensity based non-rigid registration algorithms

Since feature or surface extraction is especially difficult in US images, we believe that an intensity-based method can more easily yield an automatic algorithm. Over recent years, several non-rigid registration techniques have been proposed. (Bajcsy and Kovačič, 1989) differentiated the linear correlation criterion and used a fixed fraction of its gradient as an external force to interact with a linear elasticity model.

(Christensen et al., 1997) show that the linear elasticity, valid for small displacements, cannot guarantee the conservation of the topology of the objects as the displacements become larger: the Jacobian of the transformation can become negative. Thus, he proposed a viscous fluid model of transformations as it can handle larger displacement. This model is also linearized in practice.

(Bro-Nielsen, 1996) started from the fluid model of Christensen and used the linearity of partial derivative equations to establish a regularization filter, several order of magnitude faster than the previous finite element method. He also justified his forces as the differential of the sum of square intensity differences criterion, but he still used a fixed fraction of this gradient, and shows that Gaussian smoothing is an approximation of the linear elastic model.

Some authors (Maintz et al., 1998) tried to apply to non-rigid registration some criteria developed for rigid or affine matching using block-matching techniques. However, these criteria require a minimal window size, thus limiting the resolution of the result. Moreover, the regularization of the displacement field is usually implicit, i.e. only due to the integration of the criterion over the window, which means that it is difficult to explicitly control the regularity of the sought transformation.

Recently, (Thirion, 1998) proposed to consider non rigid registration as a diffusion process. He introduced in the images entities (demons) that push according to local characteristics of the images in a similar way Maxwell did for solving the Gibbs paradox in thermodynamics. The forces he proposed were inspired from the optical flow equations. This algorithm is increasingly used in several teams as reported by (Dawant et al., 1999; Bricault et al., 1998; Webb et al., 1999; Prima et al., 1998). In (Pennec et al., 1999; Cachier et al., 1999), we investigated the non-rigid registration using gradient descent techniques. Differentiating the sum of square intensity differences criterion (SSD), we showed that the demons forces are an approximation of a second order gradient descent on this criterion. The same gradient descent techniques were applied to a more complex similarity measure in (Cachier and Pennec, 2000): the sum of Gaussian-windowed local correlation coefficients (LCC).

1.5 Overview of the article organization

In this article, we develop an automatic intensity-based non-rigid tracking algorithm suited for real-time US images sequences, based on encouraging preliminary results reported in (Pennec et al., 1999; Cachier et al., 1999; Pennec et al., 2001b; Pennec et al., 2001a). We first present the registration method for two US images. We detail in section 2.1 a new parameterized deformation field. Then, we define in sections 2.2 and 2.3 the similarity and regularization

energies, which are optimized in section 2.4 using a gradient descent algorithm. We show in section 2.6 how to turn the registration of two images into a tracking algorithm. In section 3, we present qualitative results of the tracking algorithm on a sequence of US images of a phantom, and quantitative results on a small sequence of US images of a dead pig brain with a simulated cyst. Our results tend to justify the choices of the similarity energy and of the model of deformations, particularly with the longer term goal of achieving a real-time tracking system.

2 The tracking algorithm

When analyzing the problem of tracking the brain deformation in 3D US time-sequences, we made the following observations. First, deformations are small between successive images in a real-time sequence, but they are possibly large deformations around the surgical tools with respect to the pre-operative image. Thus, the transformation space should allow large deformations, but only small deformations have to be retrieved between successive images. Second, there is a poor signal to noise ratio in US images and the absence of information in some areas. However, the speckle (inducing localized high intensities) is usually persistent in time and may produce reliable landmarks for successive images (Meunier and Bertrand, 1995). As a consequence, the transformation space should be able to interpolate in areas with few information while relying on high intensity voxels for successive images registration. Last but not least, the algorithm is designed in view of a real-time registration during surgery, which means that, at equal performances, one should prefer the fastest method.

Following the encouraging results obtained in (Pennec et al., 1999; Cachier et al., 1999) for the intensity based non-rigid registration of two 3D US images, we adapt in this section the method in four different directions, according to the previous observations. We first look for more robust free-form transformations. Then we compare different image similarity criteria and different optimization strategies. Finally, we transform the registration algorithm into a tracking tool suited for time sequences.

2.1 *Parameterization of the transformation*

The brain shift is only a small deformation, but the introduction of surgical tools and the removal of tissues near the region of interest may locally introduce some large deformations. Simple transformations, like rigid or affine ones, can be represented by a small number of parameters (resp. 6 and 12 in 3D). When it comes to free-form deformations, we need to specify the coor-

ordinates $T(x)$ of each point x of the image after the transformation. Such a non-parametric transformation is usually represented by its displacement field $U(x) = T(x) - x$ (or $U = T - \text{Id}$), sampled at each voxel. This strategy proved to be successful in textured enough regions but induces convergence problems in large uniform areas (as it is the case in the phantom sequence of section 3.1) because the propagation of regularization constraints is very slow.

We found that a re-parameterization of the transformation was necessary to promote a better conditioning of the problem. We previously had a displacement t_i for each voxel position x_i . Now, t_i represents a parameter of a smooth transformation defined by:

$$T(t_1, \dots, t_n)(x) = \sum_i t_i \cdot G_\sigma(x - x_i) \quad (1)$$

Note that when σ goes to 0, the parameterization tends toward the previous parameterization. Moreover, this parameterization can still interpolate any value at each site x_i . The transformation being described as being a sum of Gaussians, rather than a sum of Diracs, the gradient descent algorithm uses the derivatives of the similarity with respect to the displacement of an entire group of voxels, which is more robust to noise, and will lead to a faster propagation of regularity constraints in uniform intensity areas.

In this article, we used a site x_i at each voxel of the image. One could think of reducing the number of sites to decrease the number of parameters and go toward smoother transformations. However, we observed that this is not equivalent to the smoothing performed in section 2.3: the Gaussian introduced here parameterizes the width of the neighborhood around voxel x_i for which the voxel intensities will have an influence on the transformation parameter t_i at x_i . Thus, it can be seen as a regularization of the similarity energy landscape, as we will see in Eq. 6, and not as a regularization of the transformation. Thus, reducing the number of sites (at a fixed Gaussian width) would only reduce the resolution of the transformation.

2.2 Similarity energy

Even if there is a poor signal to noise ratio in US images, the speckle is usually persistent in time and may produce reliable landmarks within the time-sequence (Meunier and Bertrand, 1995). Hence, it is desirable to use a similarity measure which favors the correspondence of similar high intensities for the registration of successive images in the time-sequence. First experiments presented in (Pennec et al., 1999; Cachier et al., 1999) indicated that the simplest one, the sum of square differences (SSD), could be suited. Let

I be the reference image and $J \circ T$ the transformed image to register; the criterion to minimize is:

$$SSD(T) = \int (I - J \circ T)^2$$

In (Cachier and Pennec, 2000), we developed a more complex similarity measure: the sum of Gaussian-windowed local correlation coefficients (LCC). Let $G \star f$ denote the convolution of f by the Gaussian function G . We define the local mean by $\bar{I} = (G \star I)$, the local variance by $\sigma_I^2 = G \star (I - \bar{I})^2$, and the local correlation of I and $J \circ T$ by $LC(T) = G \star [(I - \bar{I})(J \circ T - \bar{J} \circ \bar{T})]$. Then, the global criterion to maximize is the sum of the local correlation coefficients:

$$LCC(T) = \int \frac{LC(T)}{\sigma_I \cdot \sigma_{J \circ T}}$$

We showed that this criterion can be differentiated up to the second order using only recursive Gaussian convolutions which are very fast and in a time independent of the standard deviation of the Gaussian. Thus it may be optimized using a gradient descent like the SSD criterion.

We run most of the experiments presented in section 3 with the LCC criterion and we did not find significant differences with the results of the algorithm using the SSD. However, the LCC is still around 2 times slower than the SSD. Since the computation time of the US-US non-rigid registration is a key issue for real-time motion tracking, we preferred to keep the SSD criterion. We believe that this choice is justified anyway for the registration of successive images in the time sequence, but it could be reconsidered for the update of the global deformation (transformation from the first image to the current one, see section 2.6) if the sequence was to present some important intensity changes along time.

2.3 Regularization energy

In non-rigid registration, there is a trade-off to find between the similarity energy, reflected by the visual quality of the registration, and the smoothing energy, reflected by the regularity of the transformation (the term ‘‘regularity’’ should be taken in its broadest sense, since the smoothing energy may allow occasional discontinuities in the displacement field (Hellier et al., 1999)).

In the regularization theory framework, one minimizes the weighted sum of the energies: $E_{sim} + \lambda \cdot E_{reg}$. This formulation has proven to be successful for data approximation, and has been used for various approaches of non-rigid

registration algorithms (Ferrant et al., 1999). However, there is an important difference between data approximation and image registration. In data approximation, both energies measure different properties of the same object (the similarity and the smoothness of the data), while the two energies relate to different objects in image registration (the intensities of the images for the matching energy and the transformation for the regularization energy). Thus, one has to find a non linear tradeoff between the two energies.

Another widely spread method attempts to separate the image measure from the transformation measure, and could be compared with the approach of game theory. It consists in alternatively decreasing the similarity energy and the smoothing energy. This approach is chosen in many block-matching algorithms (Ourselin et al., 2000) and in some optical-flow-based techniques (Thirion, 1998). In view of a real-time system, this is particularly well suited for the stretch energy (or membrane model) $E_{reg} = \|\nabla T\|^2 = \int \text{Tr}(\nabla T \cdot \nabla T^T)$ as the associated Euler-Lagrange evolution equation corresponds to the heat propagation in a homogeneous material. Thus, one step of gradient descent corresponds to convolution of the transformation by a Gaussian with a standard deviation linked to the time step of the gradient descent (Morel and Solimini, 1995). This way, we get a simple regularization by a Gaussian smoothing of the transformation parameters t_i with a smoothing parameter (the σ_T of this Gaussian) that has a physical meaning.

In summary, the algorithm will alternatively perform one step of gradient descent on the similarity energy E_{sim} and one step of transformation smoothing by Gaussian filtering of standard deviation σ_T .

2.4 Minimizing the similarity energy for a free-form deformation

Let T be the current estimation of the transformation and $(\nabla_J \circ T)(x)$ (resp. $(\mathcal{H}_J \circ T)(x)$) be the transformed gradient (resp. Hessian) of the image J . A perturbation by a displacement field $u(x)$ gives the following Taylor expansion:

$$(J \circ (T + u))(x) = (J \circ T)(x) + (\nabla_J \circ T)^T \cdot u(x) + \frac{1}{2} u(x)^T \cdot (\mathcal{H}_J \circ T) \cdot u(x)$$

Thus, the Taylor expansion of the criterion is:

$$\begin{aligned} SSD(T + u) = SSD(T) &+ 2 \int (J \circ T - I) \cdot (\nabla_J \circ T)^T \cdot u \\ &+ \int ((\nabla_J \circ T)^T \cdot u)^2 + \int (J \circ T - I) \cdot u^T \cdot (\mathcal{H}_J \circ T) \cdot u + O(\|u\|^2) \end{aligned}$$

where $\|u\|^2 = \int_x \|u(x)\|^2 \cdot dx$ is the \mathcal{L}_2 norm of the small perturbation. As, by definition, $\int_x f(x)^T \cdot u(x) \cdot dx$ is the dot product of f and u in the space of

square-integrable functions, we get by identification:

$$\nabla_{SSD}(T) = 2(J \circ T - I).(\nabla_J \circ T) \quad (2)$$

$$\mathcal{H}_{SSD}(T) = 2(\nabla_J \circ T).(\nabla_J \circ T)^T + 2(J \circ T - I).(\mathcal{H}_J \circ T) \quad (3)$$

Let us now approximate the criterion by its tangential quadratic form at the current transformation T . We get the following first order approximation of the criterion gradient: $\nabla_{SSD}(T + u) \simeq \nabla_{SSD}(T) + H_{SSD}(T).u$

Assuming that the Hessian matrix of the criterion is positive definite, the minimum is obtained for a null gradient, i.e.: $u = -\mathcal{H}_{SSD}^{(-1)}(T).\nabla_{SSD}(T)$. This formula require to invert the Hessian matrix $\mathcal{H}_{SSD}(T)$ at each point x of the image. To speed up the process, we approximate this matrix by the closest scalar matrix (for the \mathcal{L}_2 norm on the matrix vector space):

$$\mathcal{H}_{SSD}(T) \simeq \frac{\text{Tr}(\mathcal{H}_{SSD}(T))}{n}.\text{Id} = \frac{\|\nabla_J \circ T\|^2 + (J \circ T - I).(\Delta_J \circ T)}{3}.\text{Id}$$

where n is the space dimension (3 for us). Using this approximation, we get the following adjustment vector field:

$$u \simeq \frac{-3.(J \circ T - I).(\nabla_J \circ T)}{\|\nabla_J \circ T\|^2 + (J \circ T - I).(\Delta_J \circ T)} \quad (4)$$

In fact, when minimizing the reverse SSD criterion $f(I \circ T^{(-1)} - J)^2$, one finds that the optimal adjustment is given by (Pennec et al., 1999; Cachier et al., 1999):

$$\hat{T} = T \circ (\text{Id} + u') \quad \text{with} \quad u' = \frac{3.(I - J \circ T).\nabla_I}{\|\nabla_I\|^2 + (I - J \circ T).\Delta_I}$$

which justifies the empirical force used by Thirion's demons:

$$v = \frac{(I - J \circ T).\nabla_I}{\|\nabla_I\|^2 + \alpha.(I - J \circ T)^2}$$

In practice, we have modified the Newton optimization scheme described above into a Levenberg-Marquardt method where the adjustment vector field is given at each step by $u = -(\lambda.\text{Id} + \mathcal{H}_{SSD})^{(-1)}.\nabla_{SSD}$. Dropping the second order terms in the Hessian, we are left with:

$$u = -3.(J \circ T - I)/(\|\nabla_J \circ T\|^2 + \lambda^2).(\nabla_J \circ T) \quad (5)$$

The parameter λ performs a tradeoff between a first order gradient descent ($\lambda \gg 1$ means that we don't trust the Hessian matrix and we simply go along the gradient with a small time-step) and a second order gradient descent ($\lambda \ll 1$ means that we use our simplified Hessian matrix). At each step, λ is divided by a fixed value α (typically 5) if the similarity criterion decreased, and the criterion is re-estimated with λ multiplied by α otherwise until the criterion decreases.

2.5 Minimizing the similarity energy for the new type of transformations

We now detail the differences induced by our new parameterization of the free-form transformation on the SSD criterion. Using the Gaussian parameterization of the transformation (eq. 1), t_i is now a parameter of the transformation. Let $G_{(x_i, \sigma)} \star f$ denote the convolution by a Gaussian of variance σ centered at x_i . Deriving the SSD w.r.t. this parameter gives:

$$\nabla_{SSD}(T) = 2 G_{(x_i, \sigma)} \star \left((J \circ T - I) \cdot (\nabla J \circ T) \right) \quad (6)$$

$$\mathcal{H}_{SSD}(T) = 2 G_{(x_i, \sigma)}^2 \star \left((\nabla J \circ T) \cdot (\nabla J \circ T)^T + (J \circ T - I) \cdot (\mathcal{H}J \circ T) \right) \quad (7)$$

Thus, the Gaussian parameterization acts as a smoothing on the gradient and Hessian of the energy. Therefore, it will be more robust and may escape from previous local minima. The minimization is performed as above with a Levenberg-Marquardt method using these regularized version of the energy derivatives.

2.6 From the registration to the tracking algorithm

In the previous sections, we studied how to register two US images together. We now have to estimate the deformation of the brain between the first image (since the dura mater is still closed, it is assumed to correspond to the pre-operative brain) and the current image of the sequence. One could think of registering directly US_1 (taken at time t_1) and US_n (at time t_n) but the deformations could be quite large and the intensity changes important. To constrain the problem, we need to exploit the temporal continuity of the deformation.

First, assuming that we already have the deformation $T_{US}(n)$ from image US_1 to US_n , we register US_n with the current image US_{n+1} , obtaining the transformation $dT_{US}(n)$. If the time step between two images is short with respect to the deformation rate (which should be the case in real-time sequences at a rate ranging from 1 to 5 images per second), this registration should be easy.

Moreover, the intensity changes should be small. For this step, we believe that the SSD criterion is well adapted.

Then, composing with the previous deformation, we obtain a first estimation of $T_{US}(n+1) \simeq dT_{US}(n) \circ T_{US}(n)$. However, the composition of deformation fields involves interpolations and just keeping this estimation would finally lead to a disastrous cumulation of interpolation errors:

$$T_{US}(n+1) = dT_{US}(n) \circ dT_{US}(n-1) \dots dT_{US}(2) \circ dT_{US}(1)$$

Moreover, a small systematic error in the computation of $dT_{US}(n)$ leads to a huge drift in $T_{US}(n)$ as we go along the sequence.




Fig. 2. The deformations computed in the tracking algorithm.

Thus, we only use $dT_{US}(n) \circ T_{US}(n)$ as an initialization for the registration of US_1 to US_n . Starting from this position, the residual deformation should be small (it corresponds to the correction of interpolation and systematic error effects) but the difference between homologous point intensities might remain important. In this case, the LCC criterion might be better than the SSD one despite its worse computational efficiency.

One of the main consequences is that the first US image will have to be of very high quality since it will be the only reference for tracking deformations along the whole sequence. One possibility consists in acquiring several images of the still brain in order to compute a mean image of better quality. Another possibility consists in performing some anisotropic diffusion on US_1 to improve its quality.

3 Experiments

In this section, we present qualitative results of the tracking algorithm on a sequence of US images of a phantom, and quantitative results on a small sequence of US images of a dead pig brain with a simulated cyst. Experiments were performed using the SSD and the LCC criterion without significant differences in the results. Since the LCC is around 2 times slower than the SSD,

we present here results and computation times for the SSD criterion. All 3D-US images were acquired using a commercial 3D-US volume scanner Voluson 530 D from Kretz Technology (4-9 MHz, 90 degrees aperture).

3.1 A Phantom study

Within the ROBOSCOPE project, an MR and US compatible phantom was developed by Prof. Auer and his colleagues at ISM (Austria) to simulate brain deformations. It is made of two balloons, one ellipsoid and one ellipsoid with a “nose”, that can be inflated with known volumes. Each acquisition consists in one 3D MR image and one 3D US image (see Fig. 3 for an example). The goal is to use the US sequence to track the deformations and compute the corresponding virtual MR images from the first MR image. Then, the original MR images can be used to assess the quality of the tracking.

Since the US probe cannot enter the MR machine, it was removed for the MR acquisitions. Thus, we had to compensate for the apparent motion of the probe by first computing a rigid registration of all the US images together. Then we run the deformation tracking algorithm. The registration of each image of the sequence takes between 10 and 15 minutes on a standard PC running linux.

Results are presented in figure 3: on the first line, we show the original US images after the rigid registration. The second line represents the first US image deformed to match the above US image. On the last two lines, we show the MR image registered to the original US (our “ground truth”) and the virtual MR produced by the tracking algorithm. To assess the quality of the tracking, we superimposed on the virtual MR images the contours of the balloons extracted from the “original” MR images.

Even if there are very few salient landmarks (all the information in the US images is located in the thick and smooth balloons boundaries, and thus the tracking problem is loosely constrained), results are globally good all along the sequence. This shows that the SSD criterion correctly captures the information at edges and that our regularized free-form deformation field is able to interpolate reasonably well in uniform areas.

When looking at the virtual MR in more details, one can however find some places where the motion is less accurately recovered: the contact between the balloons and borders of the US images. Indeed, the parameterization of the transformation and especially its smoothing are designed to approximate the behavior of a uniform elastic like body. If this assumption can be justified for the shift of brain tissues, it is less obvious for our phantom where balloons are placed into a viscous fluid. In particular, the fluid motions between the two balloons cannot be recovered. On the borders of the US images, there is

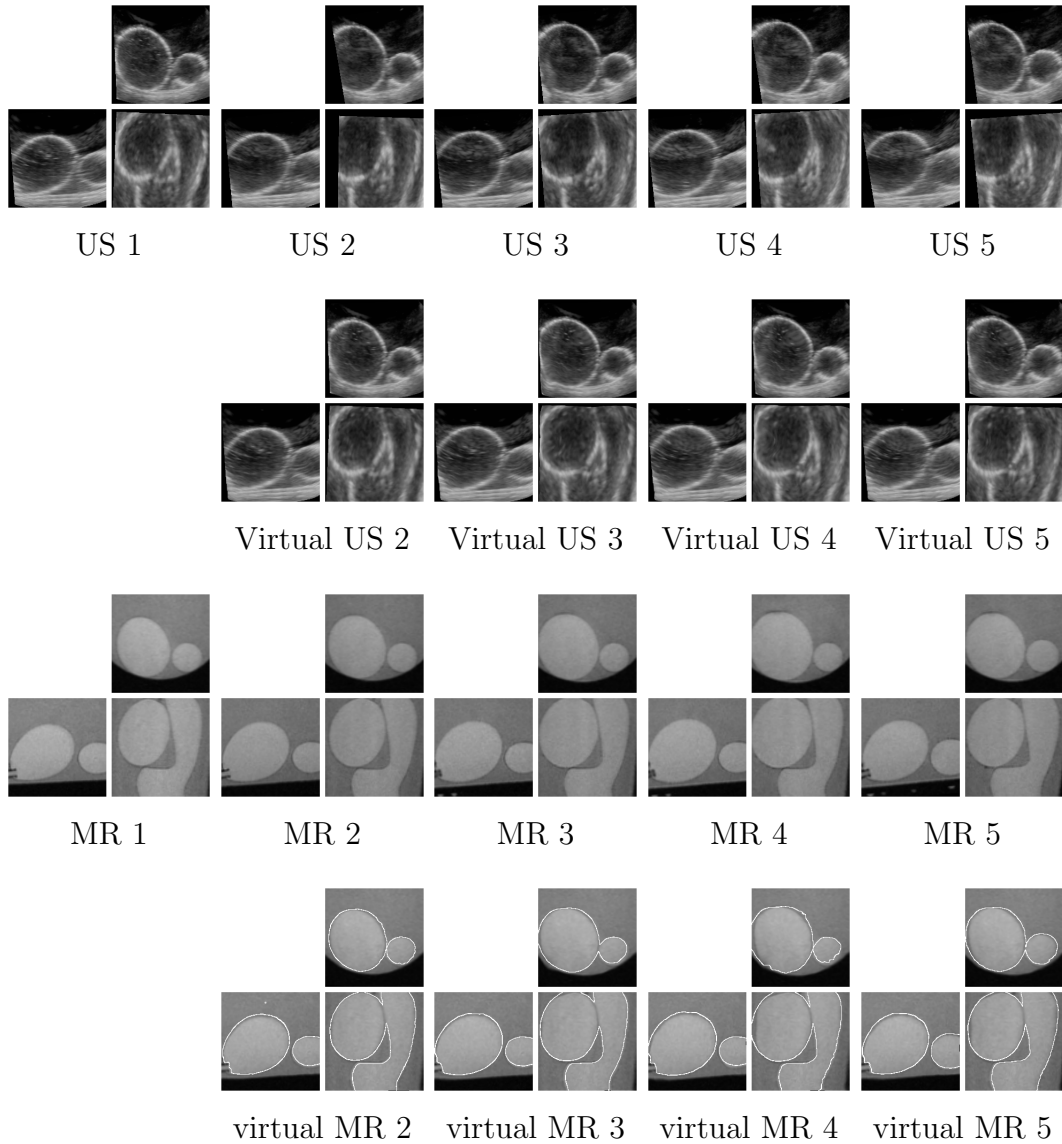


Fig. 3. **Tracking deformations on a phantom.** In this figure, each triplet of 2D images represents 3 orthogonal views resliced from the 3D image. **Top:** The first 5 images of the sequence of 10 images after a rigid registration to compensate for the motion of the probe and the “virtual” US images (US 1 deformed to match the current US image) resulting from the tracking. **Bottom:** The “original” MR images (rigidly registered to the corresponding US images to correct for the probe motion and the phantom motion between MR acquisitions) and the virtual MR image synthesized using the deformation field computed on the US images. To assess the quality of the tracking, we superimposed the contours of the “original” MR images. The volume of the balloons ranges from 60 to 90 ml for the ellipsoid one and 40 to 60 ml for the more complex one.

sometimes a lack of intensity information and the deformation can only be extrapolated from the smoothing of neighboring displacements. Since we are not using a precise geometrical and physical model of the observed structures like in (Skrinjar and Duncan, 1999), one cannot expect this extrapolation to be very accurate.

As a conclusion from this experiment, one can say that elastic-like deformations are qualitatively well tracked in the sequence if there are some salient intensity landmarks surrounding the area of interest.

3.2 Real (pig) brain images

This dataset was obtained by Dr. Ing. V. Paul at IBMT, Fraunhofer Institute (Germany) from a pig brain at a post-lethal status. A cyst drainage has been simulated by deflating a balloon catheter with a complete volume scan at three steps. We present in figure 4 the results of the tracking. Since we have no corresponding MR image, we present on the two last lines the deformation of a grid (a virtual MR image...), to emphasize the regularity of the estimated deformation, and the deformation of a segmentation of the balloon. The registration of each image of the sequence takes between 10 and 15 minutes on a standard PC running linux.

The correspondence between the original and the virtual (i.e. deformed US 1) images is qualitatively very good. In fact, if the edges are less salient than in the phantom images, we have globally a better distribution of intensity features over the field of view due to the speckle in these real brain images. One should also note on the deformed grid images that the deformation found is very smooth.

To obtain a quantitative measurement of the transformation, we segmented the first image and we deformed this segmentation according to the estimated transformation field (see bottom line of Fig. 4). We can now compare the volume of the deformed balloon with its theoretical value. In fact, since the segmentation originally overestimates the balloon volume, we have to compare the ratio between the deformed volume and the original one.

Image number	1	2	3	4
Original balloon volume (cm ³)	1.25	1.00	0.75	0.5
Relative volume ratio		0.8	0.6	0.4
Measured balloon volume	1.28	1.10	0.80	0.67
Measured volume ratio		0.86	0.62	0.53




Fig. 4. **Tracking deformations on a pig brain.** In this figure, each triplet of 2D images represents 3 orthogonal views resliced from the 3D image. **Top:** The 4 images of the pig brain with a deflating balloon simulating a cyst drainage. **Middle:** deformation of a grid to visualize more precisely the location of the deformations found. These images correspond to the deformation of an image of a 3D grid (a “virtual MR” image) with strips orthogonal to each 2D resliced plane: they allow to visualize the in-plane deformation for each 2D slice. **Bottom:** We segmented the balloon on the first image. Then, this segmentation is deformed using the transformation found and superimposed to the corresponding original US image.

The measurements indicates that we are overestimating the volume (underestimating the deformation) by 7.5% for image 2, by 3.3% for image 3, and by 30% for image 4. However, one should note that volume measurements are very sensitive as they relate to the cube of the balloon dimension: this corresponds to an error of less than one millimeter on the balloon diameter. This could be explained by an occlusion of the lower part of the balloon probably due to an air bubble trapped inside the balloon during the experience: on US 4, almost the entire lower half of the balloon is shadowed by the air bubble. In these conditions, one cannot expect a perfect retrieval. The estimated deformation at the occlusion being computed thanks to the regularization of the

deformation field from neighboring structures, it is expected to be less than the real deformations (maximal at the balloon boundaries).

Reducing the smoothing of the transformation could allow the algorithm to find a closer fit. However, this could allow some unwanted high frequency deformations due to the noise in the US images. We believe that it is better to recover the most important deformations and miss some smaller parts than trying to match exactly the images and have the possibility to create some possibly large deformations.

4 Discussion and conclusion

The algorithm presented here partly fills the goals of the ROBOSCOPE project: it is able to recover an important part of the deformations along the sequence and issues a smooth deformation, despite the noisy nature of the US images. Experiments show that this allows to simulate virtual MR images qualitatively very close to the real ones. Quantitative measurements remains to be done, but it seems that an accuracy of 1 to 2 mm is achievable in the areas where there is an elastic deformation. This is encouraging since the accuracy of the clinicians without per-operative imaging is estimated to be around 3 to 5 mm. However, some improvements of the algorithm will likely be needed to cope with non-elastic deformations in the CSF, the skull, and with the introduction of the surgical tools.

We observed that the SSD criterion is well adapted for the registration of successive images in the time-sequence and performs well on our examples for the update of the global transformation. However, it is possible that other types of sequences with intensity changes may require a more complex criterion like the LCC.

The type of transformation is a very sensitive choice for such a tracking algorithm. We made the assumption of a “uniform elastic” material. This may be adequate for the brain tissues, but probably not for the ventricles and for the tracking of the surgical tools themselves. Indeed, they will penetrate into the brain without any elastic constraint with the neighboring tissues. A specific adaptation of the algorithm around the tools will likely be necessary. Another possibility for errors is the occlusion of a part of a structure visible in the US, for instance the shadowing by the endoscope.

The computation time is still far from real time for a continuous tracking of deformations during surgery but the implementation was focused on generic components in order to test different criteria and gradient descent methods. A dedicated re-implementation of the method may gain a factor 4 to 8, leading to

a clinically useful tool for brain shift estimation (one estimation every minute or 2). To be further accelerated and reach real-time video-rate for instance, the algorithm must be parallelized. This would impose stronger hardware requirements but it is rather straightforward for both the computation of the image similarity and the regularization energies.

There are different parameters to tune in the algorithm but we believe that most of them could be adjusted for specific types of US images sequences. More sequences are anyway necessary to validate the estimation of the deformation.

In conclusion, we developed a tracking algorithm adapted to time sequences of US images and not only the registration of two images. Experiments on a phantom and on a real (pig) brain sequence show that the main part of the deformation is retrieved with a smooth deformation field. The image similarity criterion being independent from the type of transformation used, it could be changed in the future to better fit the assumptions on the US images depending on the application considered. We have shown here that the SSD criterion performs reasonably well in view of real-time considerations, even if a specific parallel version has to be designed in order to meet all the time requirements. However, more experiments will be needed to choose the best adapted parameterization of deformations and to validate the accuracy of the estimation.

Acknowledgments

This work was partially supported by the EC-funded ROBOSCOPE project HC 4018, a collaboration between The Fraunhofer Institute (Germany), Fokker Control System (Netherlands), Imperial College (UK), INRIA (France), ISM-Salzburg and Kretz Technik (Austria). The authors address special thanks to Prof. Auer and his colleagues at ISM (Austria) for the acquisition of the phantom sequence, and to Dr. Ing. V. Paul at IBMT, Fraunhofer Institute (Germany) for the acquisition of the pig brain images.

References

- Bajcsy, R. and Kovačič, S. (1989). Multiresolution Elastic Matching. *Computer Vision, Graphics and Image Processing*, 46:1–21.
- Bricault, I., Ferretti, G., and Cinquin, P. (1998). Registration of Real and CT-Derived Virtual Bronchoscopic Images to Assist Transbronchial Biopsy. *Transactions in Medical Imaging*, 17(5):703–714.
- Bro-Nielsen, M. (1996). *Medical image registration and surgery simulation*. PhD thesis, IMM-DTU.
- Bucholz, R., Yeh, D., Trobaugh, B., McDurmont, L., Sturm, C., C., B., J.M.,

- H., A., L., and P., K. (1997). The correction of stereotactic inaccuracy caused by brain shift using an intraoperative ultrasound device. In *Proc of CVRMed-MRCAS'97*, LNCS 1205, pages 459–466.
- Cachier, P. and Pennec, X. (2000). 3D non-rigid registration by gradient descent on a gaussian-windowed similarity measure using convolutions. In *Proc. of IEEE Workshop on Mathematical Methods in Biomedical Image Analysis (MMBIA'00)*, pages 182–189, Hilton Head Island, South Carolina, USA. IEEE Computer society.
- Cachier, P., Pennec, X., and Ayache, N. (1999). Fast non-rigid matching by gradient descent: Study and improvements of the "demons" algorithm. Research Report 3706, INRIA.
- Christensen, G. E., Joshi, S. C., and Miller, M. I. (1997). Volumetric Transformation of Brain Anatomy. *IEEE Trans. on Medical Imaging*, 16(6):864–877.
- Comeau, R., Sadikot, A., Fenster, A., and Peters, T. (2000). Intraoperative ultrasound for guidance and tissue shift correction in image-guided neurosurgery. *Med. Phys.*, 27(4):787–800.
- Dawant, B., Hartmann, S., and S., G. (1999). Brain atlas deformation in the presence of large space-occupying tumors. In *Proc. of MICCAI'99*, LNCS 1679, pages 589–596, Cambridge, UK.
- Erbe, H., Kriete, A., Jödicke, A., Deinsberger, W., and Böker, D.-K. (1996). 3D-Ultrasonography and Image Matching for Detection of Brain Shift During Intracranial Surgery. *Computer Assisted Radiology*, pages 225–230.
- Ferrant, M., Warfield, S. K., Guttman, C. R. G., Mulkern, R. V., Jolesz, F. A., and Kikinis, R. (1999). 3D Image Matching using a Finite Element Based Elastic Deformation Model. In *Proc. of MICCAI'99*, LNCS 1679, pages 202 – 209, Cambridge, UK.
- Gobbi, D., Comeau, R., and Peters, T. (1999). Ultrasound probe tracking for real-time ultrasound/MRI overlay and visualization of brain shift. In *Proc of MICCAI'99*, LNCS 1679, pages 920–927.
- Gobbi, D., Comeau, R., and Peters, T. (2000). Ultrasound/MRI overlay with image warping for neurosurgery. In *Proc of MICCAI'00*, LNCS 1935, pages 106–114.
- Hata, N., Suzuki, M., Dohi, T., Iseki, H., Takakura, K., and Hashimoto, D. (1994). Registration of Ultrasound echography for Intraoperative Use: A Newly Developed Multiproperty Method. In *Proc. of Visualization in Biomedical Computing (VBC'94)*, volume 2359 of *SPIE Press*, pages 251–259, Rochester, MN, USA.
- Hellier, P., Barillot, C., Mmin, E., and Prez, P. (1999). Medical Image Registration with Robust Multigrid Techniques. In *Proc. of MICCAI'99*, LNCS 1679, pages 680 – 687, Cambridge, UK. Springer.
- Ionescu, G., Lavallée, S., and Demongeot, J. (1999). Automated Registration of Ultrasound with CT Images: Application to Computer Assisted Prostate Radiotherapy and Orthopedics. In *Proc. MICCAI'99*, volume

- 1679 of *Lecture Notes in Computer Science*, pages 768–777, Cambridge (UK).
- King, A., Blackall, J., Penney, G., Edwards, P., Hill, D., and Hawkes, D. (2000). Bayesian estimation of intra-operative deformation for image-guided surgery using 3-d ultrasound. In *Proc of MICCAI'00*, LNCS 1935, pages 588–597.
- Maintz, J. B. A., Meijering, E. H. W., and Viergever, M. A. (1998). General Multimodal Elastic Registration based on Mutual Information. *Image Processing*.
- Meunier, J. and Bertrand, M. (1995). Ultrasonic Texture Motion Analysis: Theory and Simulation. *IEEE Transactions on Medical Imaging*, 14(2).
- Morel, J.-M. and Solimini, S. (1995). *Variational Methods in Image Segmentation*. Progress in Nonlinear Differential Equations and Their Applications. Birkhuser.
- Ourselin, S., Roche, A., Prima, S., and Ayache, A. (2000). Block matching: a general framework to improve robustness of rigid registration of medical images. In *Proc of MICCAI'00*, LNCS 1935, pages 557–566.
- Pagoulatos, N., Edwards, W., Haynor, D., and Kim, Y. (1999). Interactive 3-D Registration of Ultrasound and Magnetic Resonance Images Based on a Magnetic Position Sensor. *IEEE Transactions on Information Technology In Biomedicine*, 3(4):278–288.
- Papademetris, X., A.J., S., D.P., D., and Duncan, J. (1999). 3D cardiac deformation from ultrasound images. In *Proc. of MICCAI'99*, LNCS 1679, pages 421–429, Cambridge, UK.
- Pennec, X., Ayache, N., Roche, A., and Cachier, P. (2001a). Non-rigid MR/US registration for tracking brain deformations. In Press, I. C. S., editor, *Proc of Int. Workshop on Medical Imaging and Augmented Reality (MIAR 2001), 10-12 June 2001, Shatin, Hong Kong*, pages 79–86.
- Pennec, X., Cachier, P., and Ayache, N. (1999). Understanding the “demon’s algorithm”: 3D non-rigid registration by gradient descent. In Taylor, C. and Colchester, A., editors, *Proc. of 2nd Int. Conf. on Medical Image Computing and Computer-Assisted Intervention (MICCAI'99)*, volume 1679 of *LNCS*, pages 597–605, Cambridge, UK. Springer Verlag.
- Pennec, X., Cachier, P., and Ayache, P. (2001b). Tracking brain deformations in time sequences of 3D US images. In Insana, M. and Leahy, R., editors, *Proc. of IPMI'01*, volume 2082 of *LNCS*, pages 169–175, Davis, CA, USA. Springer Verlag.
- Prima, S., Thirion, J.-P., Subsol, G., and Roberts, N. (1998). Automatic Analysis of Normal Brain Dissymmetry of Males and Females in MR Images. In *Proc. of MICCAI'98*, volume 1496 of *Lecture Notes in Computer Science*, pages 770–779.
- Richard, W., Zar, D., LaPresto, E., and Steiner, C. (1999). A low-cost PCI-bus-based ultrasound system for use in image-guided neurosurgery. *Computerized Medical Imaging and Graphics*, 23(5):267–276.
- Roche, A., Pennec, X., Malandain, G., and Ayache, N. (2001). Rigid regis-

- tration of 3D ultrasound with MR images: a new approach combining intensity and gradient information. *IEEE Transactions on Medical Imaging*, 20(10):1038–1049.
- Roche, A., Pennec, X., Rudolph, M., Auer, D. P., Malandain, G., Ourselin, S., Auer, L. M., and Ayache, N. (2000). Generalized Correlation Ratio for Rigid Registration of 3D Ultrasound with MR Images. In *Proc. of MICCAI'00*, LNCS 1935, pages 567–577, Pittsburgh, USA.
- Rohling, R. N., Gee, A. H., and Berman, L. (1997). Three-Dimensional Spatial Compounding of Ultrasound Images. *Medical Image Analysis*, 1(3):177–193.
- Rohling, R. N., Gee, A. H., and Berman, L. (1998). Automatic registration of 3-D ultrasound images. *Ultrasound in Medicine and Biology*, 24(6):841–854.
- Sanchez-Ortiz, G., Declerck, J., Mulet-Parada, M., and Noble, J. (2000). Automatic 3D echocardiographic image analysis. In *Proc. of MICCAI'00*, LNCS 1935, pages 687–696, Pittsburgh, USA.
- Skrinjar, O. and Duncan, J. (1999). Real time 3D brain shift compensation. In *Proc of IPMI'99*, pages 42–55, Visegrad, Hungary.
- Strintzis, M. G. and Kokkinidis, I. (1997). Maximum Likelihood Motion Estimation in Ultrasound Image Sequences. *IEEE Signal Processing Letters*, 4(6).
- Thirion, J.-P. (1998). Image matching as a diffusion process: an analogy with Maxwell's demons. *Medical Image Analysis*, 2(3).
- Trobaugh, J., Richard, W., Smith, K., and Bucholz, R. (1994a). A low-cost PCI-bus-based ultrasound system for use in image-guided neurosurgery. *Computerized Medical Imaging and Graphics*, 18(4):235–246.
- Trobaugh, J., Trobaugh, D., and Richard, W. (1994b). Three-dimensional imaging with stereotactic ultrasonography. *Computerized Medical Imaging and Graphics*, 18(5):315–323.
- Webb, J., Guimond, A., Roberts, N., and D. Chadwick, P. E., Meunier, J., and Thirion, J.-P. (1999). Automatic Detection of Hippocampal Atrophy on Magnetic Resonance Images. *Magnetic Resonance Imaging*, 17(8):1149–1161.