



HAL
open science

Tracking Brain Deformations in Time-Sequences of 3D US Images

Xavier Pennec, Nicholas Ayache, Pascal Cachier

► **To cite this version:**

Xavier Pennec, Nicholas Ayache, Pascal Cachier. Tracking Brain Deformations in Time-Sequences of 3D US Images. [Research Report] RR-4091, INRIA. 2000. inria-00072541

HAL Id: inria-00072541

<https://inria.hal.science/inria-00072541>

Submitted on 24 May 2006

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

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

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

X. Pennec — P. Cachier — N. Ayache

N° 4091

Décembre 2000

THÈME 3



*Rapport
de recherche*

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

X. Pennec , P. Cachier , N. Ayache

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

Rapport de recherche n° 4091 — Décembre 2000 — 18 pages

Abstract: During a neuro-surgical 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, Augmented reality, Image guided surgery

Suivi des déformations du cerveau dans des séquences temporelles d'images ultrasonores 3D

Résumé : Lors d'une intervention neuro-chirurgicale, le cerveau se déplace et se déforme. Pour garder un positionnement correct des instruments chirurgicaux par rapport aux structures d'intérêt, on doit estimer cette déformation à partir d'images per-opératoires. Nous présentons dans cet article une étude de faisabilité d'un algorithme de suivi de déformation utilisant des séquences temporelles d'images ultrasonores 3D per-opératoires. Le traitement automatique de ce type d'images est du plus haut intérêt pour le développement de nouveaux outils chirurgicaux guidés par l'image qui soient innovants et à bas prix. La principale difficulté réside dans la nature complexe des images ultrasonores et dans la quantité de données à traiter le plus rapidement possible.

Mots-clés : Recalage non-rigide, images ultrasonores 3D, suivi temporel, réalité augmentée, chirurgie guidée par l'image

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.

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). The operation is planned on a pre-operative MRI (MR1) and 3D

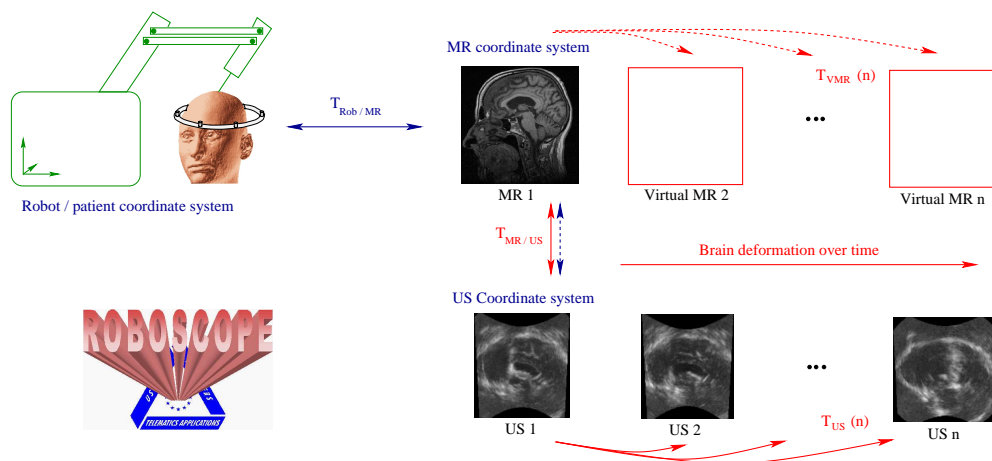


Figure 1: Overview of the image analysis part of the Roboscope project.

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 (Erbe et al., 1996) where it 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., 1998), a Chamfer matching technique is used to rigidly register the MR surface to the outlines of the 2D US image. In (Gobbi et al., 1999; Gobbi et al., 2000), the 2D US probe is optically and rigidly tracked after a proper calibration and 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. 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., 2000). 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 a Rayleigh noise and a block-matching strategy to recover the rigid motion. In (Rohling et al., 1997a), 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., 1997b). 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 find some real tracking of non-rigid motion in US images. In (Papademetris et al., 1999), the endo and epi- cardial 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 guaranty 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 linearised 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 bloc-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 char-

acteristics 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 organisation

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). We first present the registration method for two US images. We detail in section 2.1 a new parameterised deformation field. Then, we define in sections 2.2 and 2.3 the similarity and regularization energies, which are optimised 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 analysing 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 localised high intensities) is usually persistent in time and may produce reliable landmarks for successive images. 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 criterion and different optimisation strategies. Finally, we transform the registration algorithm into a tracking tool suited for time sequences.

2.1 Parameterisation 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 coordinates $T(x)$ 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 regularization problems in large uniform areas (as it is the case in the phantom sequence of section 3.1).

We found that a re-parameterisation of the transformation was necessary to induce more robust transformations. 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 parameterisation tends toward the previous parameterisation. The transformation is described as being a sum of Gaussians, rather than a sum of Diracs. Therefore, the gradient descent algorithm uses the derivatives of the similarity with respect to the displacement of an entire group of voxel, which is more robust to noise.

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. Hence, it is desirable to use a similarity measure which favours 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 adapted. Let I be the reference image and $J \circ T$ the transformed image to register; the criterion to minimise 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 $LC(T) =$

$G \star [(I - \bar{I})(J \circ T - \overline{J \circ T})]$. Then, the global criterion to maximise 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 derivated 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 optimised 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 significative 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

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. In the regularization theory framework, one minimises 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 (Hellier et al., 1999; Ferrant et al., 1999).

Despite a weaker theoretical background, we chose for efficiency reasons to alternatively minimise each 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) $E_{reg} = \|\nabla T\|^2 = \int \text{Tr}(\nabla T \cdot \nabla T^T)$. Indeed, the associated Euler-Lagrange evolution equation corresponds to a Gaussian filtering of the transformation. Thus, the algorithm will alternatively optimize the similarity energy and smooth the transformation by Gaussian filtering.

2.4 Minimising 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:

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

where $\|u\|^2 = \int_x \|u(x)\|^2 \cdot dx$ is the \mathcal{L}_2 norm of the small perturbation. As, by definition, $\int_x f(x)^\top \cdot u(x) \cdot dx$ is the dot product of f and u in the space of square-integrable functions, we get by identification:

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

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

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: $\nabla_{SSD}(T+u) \simeq \nabla_{SSD}(T) + H_{SSD}(T) \cdot u$

Assuming that the Hessian matrix of the criterion is positive definite, the minimum is obtained for a null gradient, i.e.: $u = -\mathcal{H}_{SSD}^{-1}(T) \cdot \nabla_{SSD}(T)$. 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):

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

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

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

In fact, when minimising the reverse SSD criterion $\int (I \circ T^\top - J)^2$, one find that the optimal adjustment is given by (Pennec et al., 1999; Cachier et al., 1999):

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

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

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

2.5 Minimising the similarity energy for the new type of transformations

We now detail the differences induced by our new parameterisation of the free-form transformation on the SSD criterion. Using the Gaussian parameterisation 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:

$$\nabla_{SSD}(T) = 2 G_{(x_i, \sigma)} \star \left((J \circ T - I) \cdot (\nabla J \circ T) \right) \quad (5)$$

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

Thus, the Gaussian parameterisation 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 while encouraging smoother transformations. In this article, the optimisation is performed using a method similar to Levenberg-Marquardt.

2.6 From the registration to the tracking algorithm

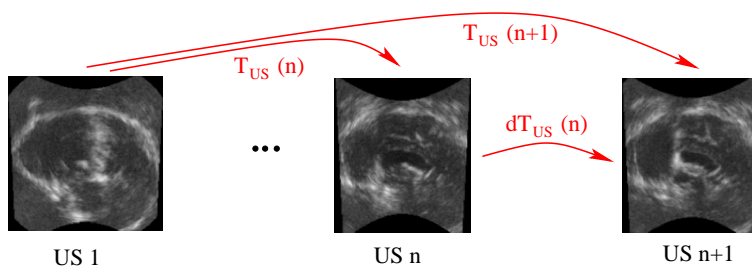
In the previous sections, we studied how to register two US images together. We now have to estimate the deformation of the brain between the first image (since the dura mater is still closed, it is assumed to correspond to the preoperative brain) and the current image of the sequence. One could think 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.



Thus, we only use $dT_{US}(n) \circ T_{US}(n)$ as an initialisation 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.

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. 2 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.

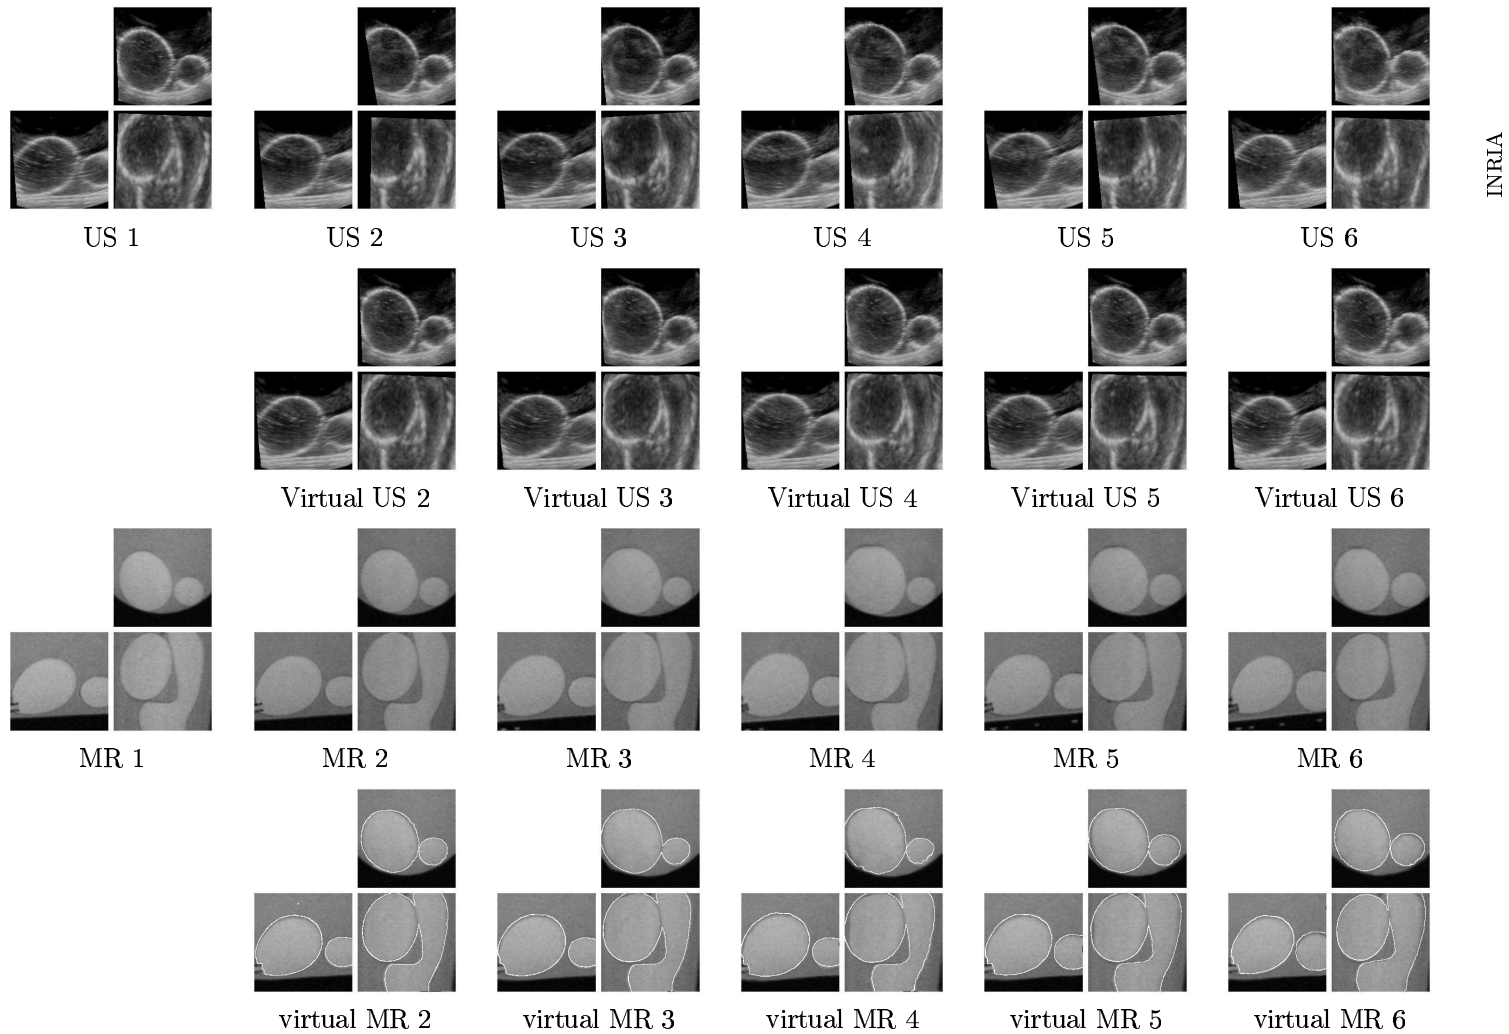


Figure 2: **Top:** 6 US images of the sequence after a rigid registration to compensate for the motion of the probe and the “virtual” US images (US 1 deformed to match the current US image) resulting from the tracking. **Bottom:** The “original” MR images (rigidly registered to the corresponding US images 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.

Results are presented in figure 2: 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 parameterisation of the transformation and especially its smoothing are designed to approximate the behaviour 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 sometimes a lack of intensity information and the deformation can only be extrapolated from the smoothing of neighbouring 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 3 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 emphasise 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

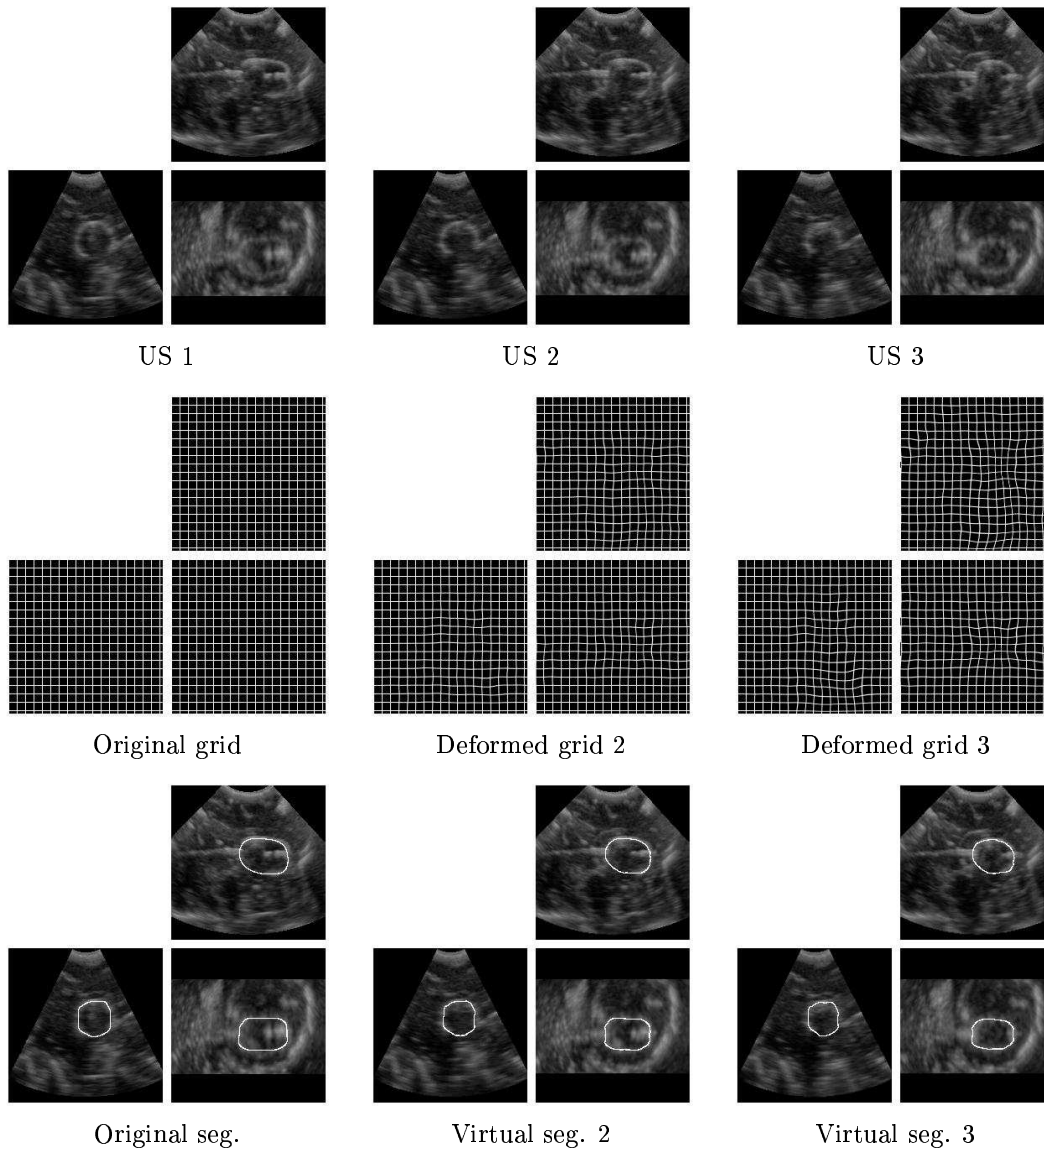


Figure 3: **Top:** The 3 images of the pig brain. **Middle:** deformation of a grid to visualise more precisely the location of the deformations found. **Bottom:** We segmented the balloon on the first image. Then, this segmentation is deformed using the transformation found and superimposed to the original US image.

(see bottom line of Fig. 3). 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
Original balloon volume (cm ³)	1.25	1.00	0.75
Relative volume ratio		0.8	0.6
Measured balloon volume	1.29	1.12	0.92
Measured volume ratio		0.87	0.82

The measurements indicates that we are overestimating the volume (underestimating the deformation) by 8% for image 2 and 27% for image 3. 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. The estimated deformation at the occlusion being computed thanks to the regularization of the deformation field from neighbouring 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. 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 neighbouring 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 parallelised. 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 parameterisation of deformations and to validate the accuracy of the estimation.

Acknowledgements

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 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.
- 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. (2000). Ultrasound/mri overlay with image warping for neurosurgery. In *Proc of MICCAI'00*, LNCS 1935, pages 106–114.
- Gobbi, D., R.M., C., 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.
- Hata, N., Suzuki, M., Dohi, T., Iseki, H., Takakura, K., and Hashimoto, D. (1998). Registration of Ultrasound Echography for Intraoperative Use: A Newly Developed Multi-property Method. *SPIE*, 2359.
- Hellier, P., Barillot, C., Mémin, E., and Pérez, P. (1999). Medical Image Registration with Robust Multigrid Techniques. In *Proc. of MICCAI'99*, LNCS 1679, pages 680 – 687, Cambridge, UK. Springer.
- 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*.
- 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.
- 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., Cachier, P., and Ayache, N. (1999). Understanding the “demon’s algorithm”: 3D non-rigid registration by gradient descent. In *Proc. of MICCAI'99*, LNCS 1679, pages 597–605, Cambridge, UK.
- 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.
- 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. (1997a). Automatic Registration of 3D Ultrasound Images. Technical Report TR 290, Cambridge University Engineering Dep.
- Rohling, R. N., Gee, A. H., and Berman, L. (1997b). Three-Dimensional Spatial Compounding of Ultrasound Images. *Medical Image Analysis*, 1(3):177–193.
- 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).
- Webb, J., Guimond, A., Roberts, N., an D. Ch adwick, P. E., Meunier, J., and Thirion, J.-P. (1999). Automatic Detection of Hippocampal Atrophy on Magnetic Resonance Images. *Magnetic Resonance Imaging*, 17. To appear.



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

Unité de recherche INRIA Lorraine : Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - B.P. 101 - 54602 Villers lès Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38330 Montbonnot St Martin (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - B.P. 105 - 78153 Le Chesnay Cedex (France)

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