Non-rigid MR/US registration for tracking brain deformations

Xavier Pennec, Nicholas Ayache, Alexis Roche, Pascal Cachier
EPIDAURE, INRIA Sophia-Antipolis
BP 93, F-06902 Sophia Antipolis Cedex, France
{Xavier.Pennec, Nicholas.Ayache, Alexis.Roche, Pascal.Cachier}@sophia.inria.fr

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. 3D ultrasound (US) imaging is an innovative and low-cost modality which appears to be suited for such computer-assisted surgery tools. In this paper, we present a new image-based technique to register intra-operative 3D US with pre-operative Magnetic Resonance (MR) data. A first automatic rigid registration is achieved by the maximisation of a similarity measure that generalises the correlation ratio. Then, brain deformations are tracked in the 3D US time-sequence using a "demon's" like algorithm. Experiments show that a registration accuracy of the MR voxel size is achieved for the rigid part, and a qualitative accuracy of a few millimetres could be obtained for the complete tracking system.

1. Introduction

The use of stereotactic systems is now a quite standard procedure for neurosurgery. However, these systems assume that the brain is in fixed relation to the skull during surgery. In practice, relative motion of the brain with respect to the skull (also called brain shift) occurs, mainly due to tumour 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.

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 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 synthetize a virtual MR image that matches the current brain anatomy.

MR/US registration The idea of MR/US registration is already present in [3] 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 [6], the outlines of the 2D US image are registered to the MR surface using a Chamfer matching technique. In [10, 5, 4], the 2D US probe is optically tracked 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 [1] with the possibility of using 3D US images and a deformation computed using a spring model instead of splines. More recently, Ionescu *et al* [7] 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 esti-

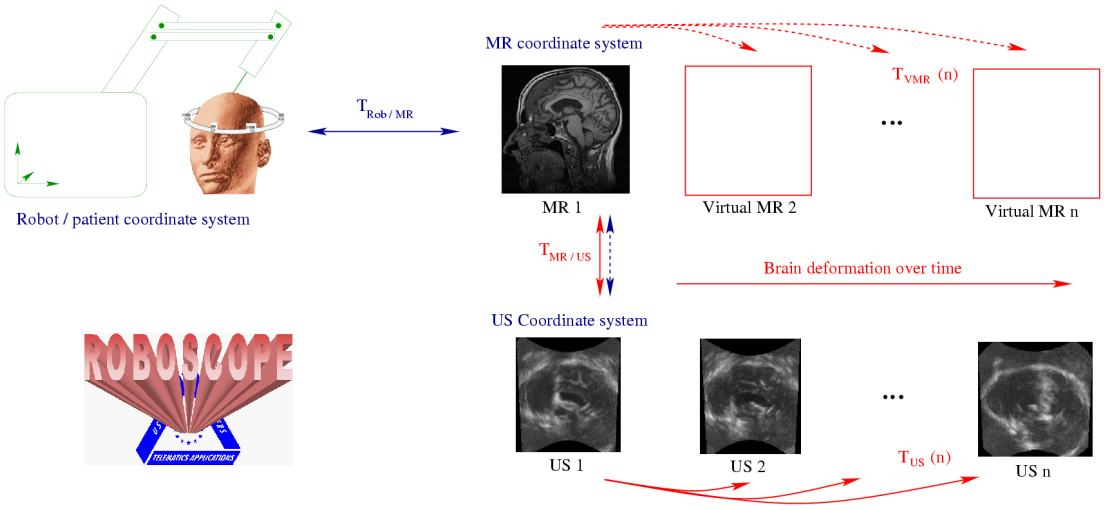

Figure 1. Overview of the image analysis part of the Roboscope project.

mation is limited to a few samples at given time-points as the user interaction is required at least to define the landmarks.

Up to our knowledge, only [8] 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.

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), then we look for the non-rigid motion within the US time-sequence.

Tracking methods in sequences of US images There are few articles on the registration of 3D US images. [19] 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 [17], 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 [16]. 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 the shape of the cardiac ventricle in sequences of 3D US images using dedicated surface models. However, if these models could be adapted to the brain ventricles, it seems difficult to extend them to the tracking of the volumetric deformations of the whole brain.

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 [9]. We chose to focus in [11] on gradient descent techniques. Differentiating the sum of square intensity differences criterion (SSD), we showed that the demons forces proposed by Thirion in [20] were 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 [2]: the sum of Gaussian-windowed local correlation coefficients (LCC).

Overview of the article organisation The first part of this article expands on the correlation ratio (CR) method [14]. It is an intensity-based approach as it does not rely on explicit feature extraction. We have improved the method in [15] following three distinct axes: using the gradient information from the MR image, reducing the number of intensity parameters to be estimated, and using a robust intensity distance.

The second part of the article develops an automatic intensity-based non-rigid tracking algorithm suited for real-time US images sequences, based on encouraging preliminary results reported in [11, 12]. We first present the registration method for two US images and how the method is turned into a tracking algorithm.

In section 4, we present some results of the rigid MR/US registration on clinical data (a baby and a surgical case), along with the results of an original evaluation of the registration accuracy. Then, we present qualitative results of the tracking algorithm on a sequence of 3D US animal images and a qualitative evaluation of the complete tracking system on a sequence of images of an MR and US compatible phantom.

2 Rigid MR/US Registration

2.1 Correlation ratio

Given two images I and J, the basic principle of the CR method is to search for a spatial transformation T and an intensity mapping f such that, by displacing J and remapping its intensities, the resulting image $f(J \circ T)$ be as similar as possible to I. In a first approach, this could be achieved by minimising the following cost function:

$$C(T, f) = \|I - J \circ T\|^2 = \int_x [I(x) - f(J(T(x)))]^2 \quad (1)$$

This formulation is asymmetric in the sense that the cost function changes when permuting the roles of I and J. Since the positions and intensities of J actually serve to predict those of I, we will call J the "template image". In the context of US/MR registration, we always choose the MR as the template.

One problem is that we can compute this criterion only on the overlapping part of the transformed images. In order to avoid a minimum when the image overlap is small, we need to renormalise the criterion: a good choice, justified in [13], is to look for a large variance of I in the overlapping region (we are trying to register informative parts), so that the criterion becomes:

$$C(T, f) = \|I - f(J \circ T)\|^2 / \mathrm{Var}(I) \quad (2)$$

where the integrals are computed over $I \cap J \circ T$. Another important point is the discretisation scheme used to compute the criterion, leading to the choice of an interpolation scheme [13]. In this paper, we use Partial Volume.

If no constraint is imposed to the intensity mapping f, an important result is that the optimal f at fixed T enjoys an explicit form that is very fast to compute [14]. The minimization of eq (2) may then be performed by travelling through the minima of $C(T, f)$ at fixed T. This yields the correlation ratio, $\eta^2_{I|J}(T) = 1 - \min_f C(T, f)$, a measure that reaches its maximum when $C(T, f)$ is minimal. In practice, the maximisation of η^2 is performed using Powell's method.

2.2 Bivariate correlation ratio

Ultrasound images are commonly said to be "gradient images" as they enhance the interfaces between anatomical structures. The physical reason is that the amplitudes of the US echos are proportional to the squared *difference* of acoustical impedance caused by successive tissue layers. Ideally, the US signal should be high at the interfaces, and low within homogeneous tissues.

Thus, assuming that the MR intensities describe homogeneous classes of tissues amounts to consider the acoustic impedance Z as an unknown function of the MR intensities: $Z(x) = g(J(x))$. Now, when the ultrasound signal emitted from the probe encounters an interface (i.e. a high gradient of Z), the proportion of the reflected energy is $R = \|\nabla Z\|^2 / Z^2$. Adding a very simple model of the log-compression scheme used to visualise the US images, we obtain the following US image acquisition model: $I(x) = a. \log\left(\|\nabla Z\|^2 / Z^2\right) + b + \epsilon(x)$. Using $Z(x) = g(J(x))$ finally gives an unknown bivariate function: $I(x) = f\left(J(x), \|\nabla J(x)\|\right) + \epsilon(x)$. Our new correlation ratio criterion is then:

$$C(T, f) = \|I - f(J \circ T, \|\nabla J \circ T\|)\|^2 / \mathrm{Var}(I) \quad (3)$$

The MR gradient is practically computed by convolution with a Gaussian kernel.

2.3 Parametric intensity fit

Since we are now looking for a bivariate intensity mapping f with floating values for the MR gradient component, one has to regularize it. We will therefore restrict our search to a polynomial function f of degree d. The number of parameters describing f then reduces to $(d+1)(d+2)/2$. In this paper, the degree was set to $d = 3$, implying that 10 coefficients were estimated. Finding the coefficient of the polynomial minimising eq (3) amounts to solve a weighted least square linear regression problem.

However, this polynomial fitting procedure adds a significant extra computational cost with respect to the unconstrained fitting and cannot be done for each transformation trial. Instead, the minimization of the criterion may be performed alternatively along T and f: (1) given a current transformation estimate T, find the best polynomial f and remap J and $\|\nabla J\|$ accordingly; (2) given a remapped image $f(J, \|\nabla J\|)$, minimise $C(T, f)$ with respect to T using Powell's method; (3), return to (1) if T or f has evolved.

2.4 Robust intensity distance

Our method is based on the assumption that the intensities of the US may be well predicted from the information available in the MR. Due to several ultrasound artefacts, we do not expect this assumption to be perfectly true. Shadowing, duplication or interference artefacts may cause large variations of the US intensity from its predicted value, even when the images are perfectly registered. To reduce the sensitivity of the registration criterion to these outliers, we propose to use a robust estimation of the intensities differences using a one-step S-estimator [18]: $\int_x [I(x) - f(J(T(x)))]^2$ is then replaced with

$$S^2(T, f) = S_0^2 / K. \int_x \rho\left([I(x) - f(J(T(x)))] / S_0\right),$$

where K is a normalisation constant that ensures consistency with the normal distribution, and S_0 is some initial guess of the scale. In our implementation, we have opted for the Geman-McClure ρ-function, $\rho(x) = \frac{1}{2}x^2/(1 + x^2/c^2)$, for its computational efficiency and good robustness properties, to which we always set a cut-off distance $c = 3.648$ corresponding to 95% Gaussian efficiency.

Initially, the intensity mapping f is estimated in a non-robust fashion. The starting value S_0 is then computed as the median of absolute intensities deviations. Due to the initial misalignment, it tends to be overestimated and may not efficiently reject outliers. For that reason, it is re-estimated after each alternated minimisation step.

3 The Tracking algorithm

The goal of this section is to estimate the brain deformations from US images time-sequences. We first detail the similarity and regularization energies minimised to find a deformation field between consecutive images of the sequence, and then how we turn this registration algorithm into a tracking tool [12].

3.1 Registering two US images

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 [11] indicated that the simplest one, the sum of square differences ($SSD(T) = \int (I - J \circ T)^2$), could be adapted. In [2], we developed a more complex similarity measure: the sum of Gaussian-windowed local correlation coefficients (LCC). Let $G \star f$ be the convolution of f by the Gaussian, $\bar{I} = (G \star I)$ be the local mean, $\sigma_I^2 = G \star (I - \bar{I})^2$ the local variance and $LC(T) = G \star \left[(I - \bar{I})(J \circ T - \overline{J \circ T})\right]$ the local correlation between image I and image $J \circ T$. Then, the global criterion to maximise is the sum of the local correlation coefficients: $LCC(T) = \int (LC(T)/\sigma_I.\sigma_{J\circ T})$.

We have shown in [11] and [2] how these criteria can be optimised using first and second order gradient descent techniques with a general free-form deformation field by computing the gradient and the Hessian of the criteria.

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 view of a real-time system, the stretch energy $E_{reg} = \int \|\nabla T\|$ (or membrane model) is particularly well suited as it is very efficiently solved by a Gaussian filtering of the transformation. Thus, the algorithm will alternatively optimize the similarity energy and smooth the transformation by Gaussian filtering.

3.2 From registration to tracking

To estimate the deformation of the brain from the first image to 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.

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, there should be small deformations and small intensity changes. 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. 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.

4 Experiments

In this section, we present quantitative results of the rigid MR/US registration algorithm on real brain images, and qualitative results of the tracking algorithm and its combination with the MR/US registration on animal and phantom sequences. The location of the US probe being linked to the pathology and its orientation being arbitrary (the rotation may be superior to 90 degrees), it was necessary to provide a rough initial estimate of the MR/US transformation. This was done using an interactive interface that allows to draw lines in the images and match them. This procedure was carried out by a non-expert, generally taking less than 2 minutes. However this user interaction could be alleviated using a calibration system such as the one described in [10]. After initialisation, we observed that the algorithm found residual displacements up to 10 mm and 10 degrees.

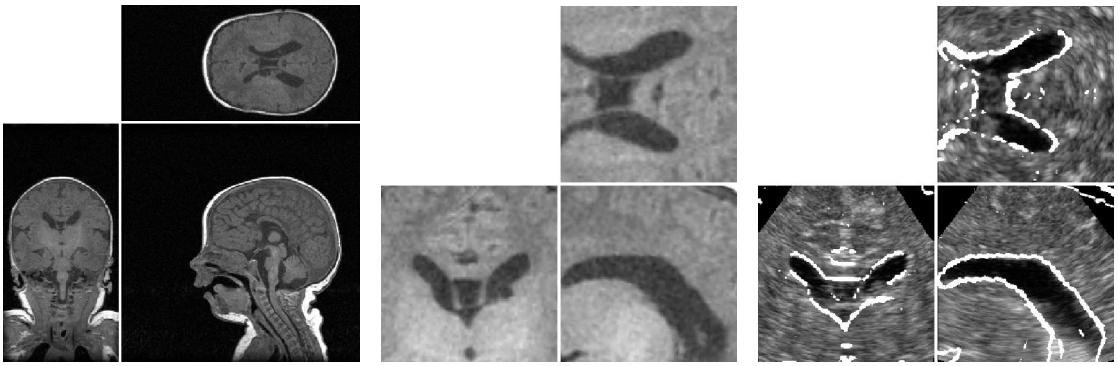

Figure 2. Example registration of MR and US images of the baby. From left to right: original MR T1 image, closeup on the ventricle area, and registered US image with MR contours superimposed.

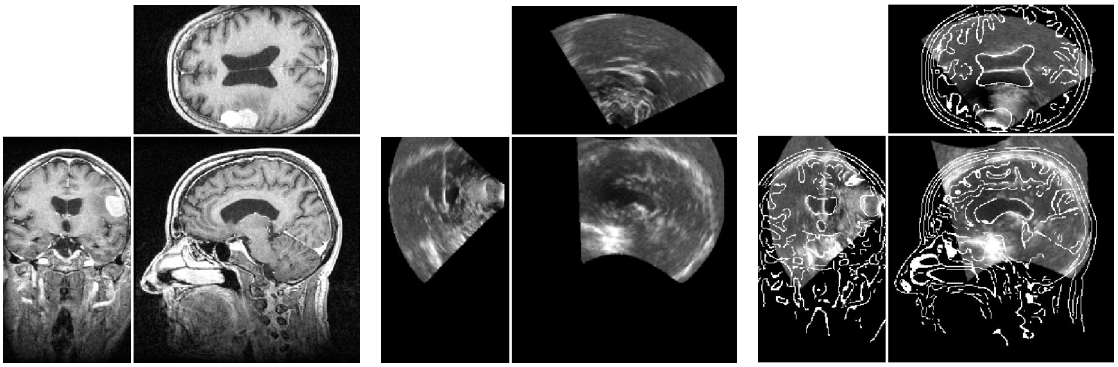

Figure 3. Example registration of MR and US images of the patient. From left to right: MR T1 image with a contrast agent, manual initialisation of the US image registration, and result of the automatic registration of the US image with the MR contours superimposed.

4.1 Images of a Baby

This dataset was acquired to simulate the degradation of the US images quality with respect to the number of converters used in the probe. Here, we have one MR T1 image of a baby's head and 5 transfontanel US images with different percentages of converters used. As we have no or very few deformations within the images, we can rigidly register all the US images onto our single MR.

An example result is presented in Fig.2. The visual quality of the registration is very good. In order to quantify more precisely the registration accuracy, we set up in [15] a validation scheme that uses *registration loops*. For these data, it shows that the registration accuracy is about 0.4 mm at the center of the image and 0.9 mm at the corners of the presented US image.

4.2 Patient images during tumour resection

This is an actual surgical case: two MR T1 images with and without a contrast agent were acquired before surgery. After craniotomy (dura mater still closed), a set of 3D US images was acquired to precisely locate the tumour to resect. In this experiment, we use the three US images that are large enough to contain the ventricles. Unfortunately, we could only test for the rigid MR/US registration as we have no US images during surgery.

An example of the registration results is presented in Fig.3. In this case, our validation scheme exhibit a registration accuracy of 0.6 mm at the center and 1.6 mm in the whole brain area [15]. However, when we look more carefully at the results, we find that the loop involving the smallest US image (real size $150 \times 85 \times 100$ mm, voxel size 0.63^3 mm^3) is responsible for a corner error of 2.6 mm (0.85 mm at the center) while the loops involving the two larger US images (real size $170 \times 130 \times 180$, voxels size 0.95^3 mm^3) do have a much smaller corner error of about 0.84 mm (0.4 mm at the center). We suspect that a non-rigidity in the smallest US could account for the registration inaccuracy. Another explanation could be a misestimation of the sound speed for this small US acquisition leading to a false voxel size and once again the violation of the rigidity assumption.

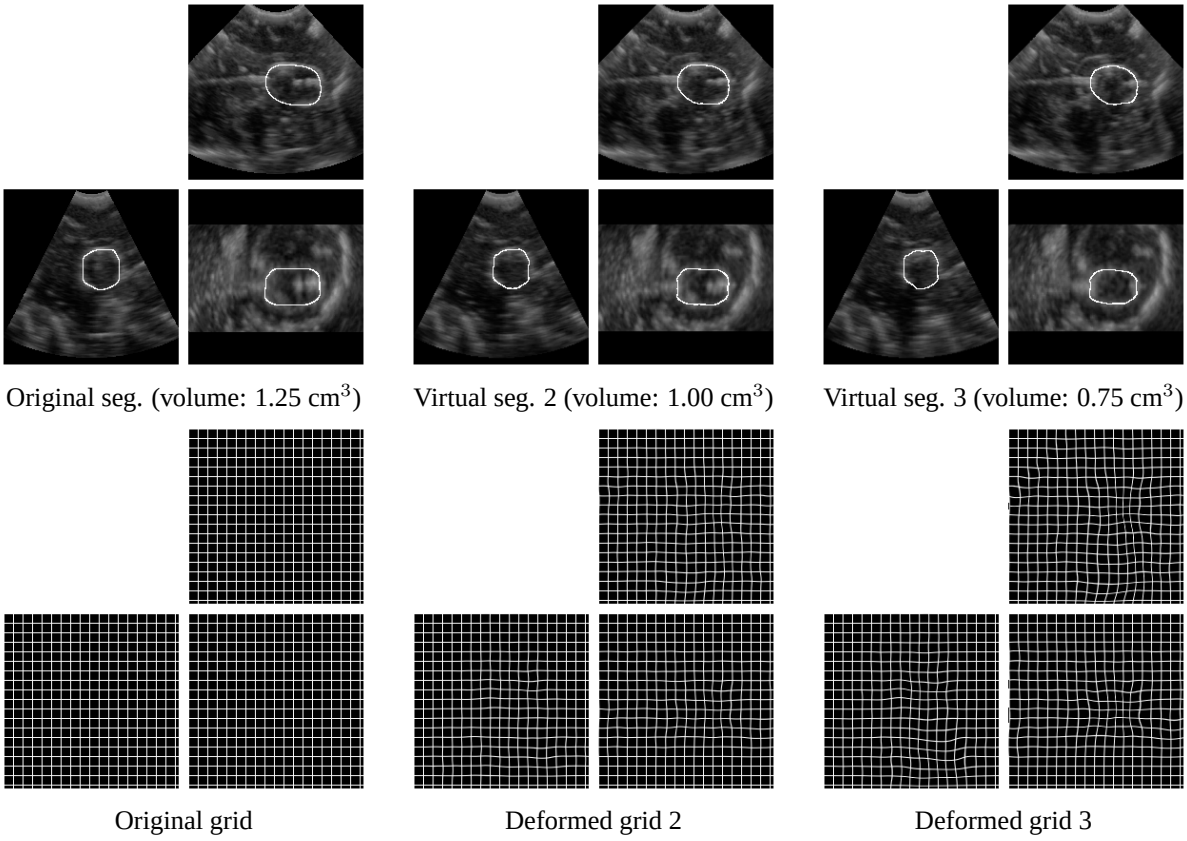

Figure 4. Top: The 3 original images of the pig brain. The segmentation of the balloon, done on the first image, is deformed according to the transformation found by the tracking algorithm and superimposed to the original US image. Bottom: deformation of a grid to visualise more precisely the location of the deformations found.

4.3 US images of an animal brain

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 synthetic image...), to emphasise the regularity of the estimated deformation, and the deformation of a segmentation of the balloon.

The correspondence between the original and the virtual (i.e. deformed US 1) images is qualitatively good. In fact, if the edges are less salient than in the phantom images (see next section), we have globally a better distribution of intensity features over the field ov 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.

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 “invent” some possibly large deformations.

4.4 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 and one 3D US image.

The first MR image is rigidly registered to the first US image. We determined that the accuracy of this registration was about 1 mm at the center of the image and 1.4 mm at

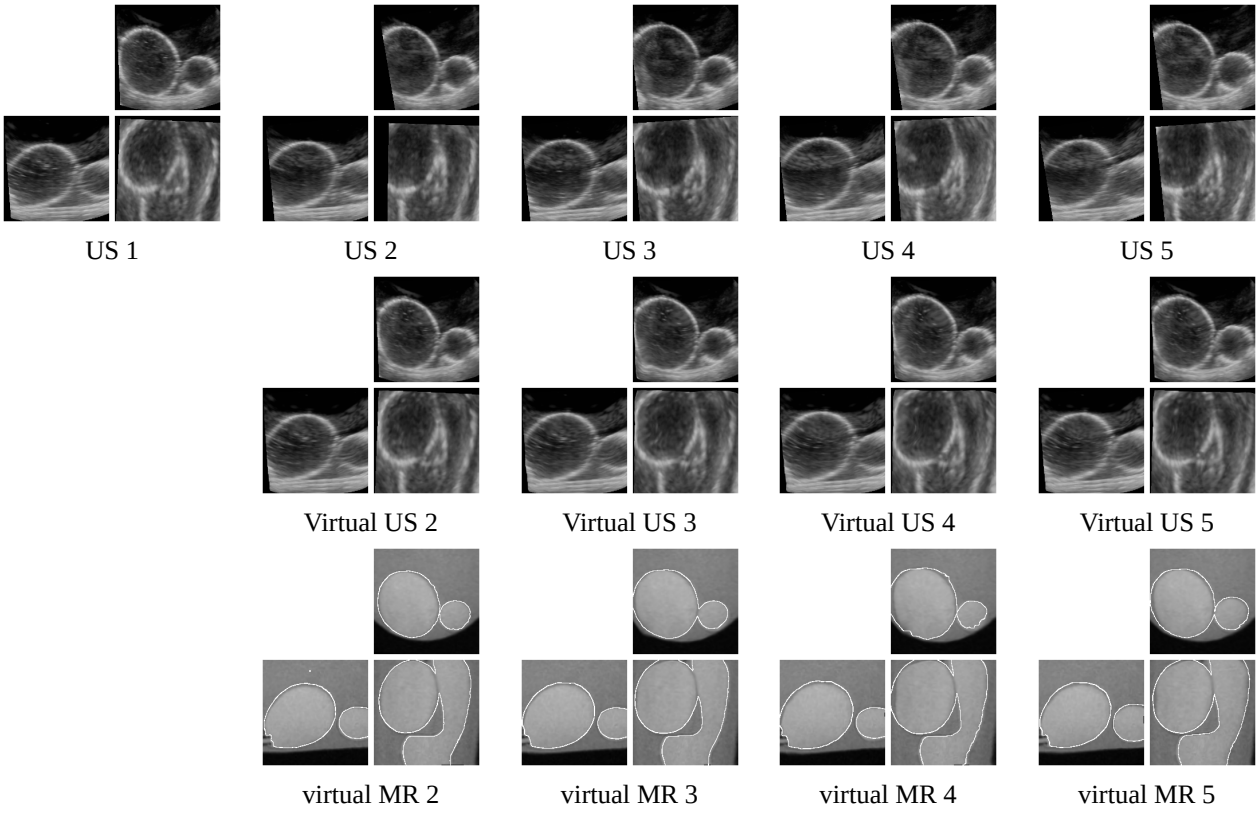

Figure 5. Beginning of the sequence of 10 images of the phantom. On top: the original US images. Middle: the "virtual" US images (US 1 deformed to match the current US image) resulting from the tracking. Bottom: the virtual MR images synthetized using the deformation field computed on the US images with the contours of the "original" MR images superimposed. 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.

the corners of the US image [15]. Then, deformations are estimated using the tracking algorithm on the US sequence, and the corresponding virtual MR image is computed. The remaining MR images can be used to assess the quality of the tracking [12].

Results are presented in Fig.5. Even if there are very few salient landmarks (all the information 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 parameterised deformation interpolates 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 often a lack of intensity information (due to the inadequate conversion from polar to Cartesian coordinates by the US machine) 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, one cannot expect this extrapolation to be very accurate.

5 Conclusion

We presented in the first part a new automated method to rigidly register 3D US with MR images. It is based on a multivariate and robust generalisation of the correlation ratio (CR) measure that allows to better take into account the nature of US images. Incidentally, we believe that the gen-

eralised CR could be considered in other registration problems where conventional similarity measures fail. Testings were performed on phantom and clinical data, and showd that the worst registration errors (errors at the Cartesian US corners) is of the order of 1.5 mm.

In the second part, we developed a tracking algorithm adapted to time sequences of US images and not only to the registration of two images. The algorithm is able to recover an important part of the deformations and issues a smooth deformation, despite the noisy nature of the US images. Experiments on animal and phantom data show that this allows to simulate virtual MR images qualitatively close to the real ones. The computation time is still far from real time but a parallelisation of the algorithms is straightforward for the computation of the image and the regularization energies.

The type of transformation is also a very sensitive choice for such a tracking algorithm. We made the assumption of a "uniform elastic" like material. This may be adequate for the brain tissues, but probably not for the ventricles and for the tracking of the surgical tools themselves. 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.

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 Dr. Ing. V. Paul at IBMT, Fraunhofer Institute for the acquisition of the pig brain images, and to Prof. Auer and his colleagues at ISM for the acquisitions of all other images.

References

[1] R. Bucholz, D. Yeh, B. Trobaugh, L. McDurmont, C. Sturm, B. C., H. J.M., L. A., and K. P. 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, 1997.

[2] P. Cachier and X. Pennec. 3D non-rigid registration by gradient descent on a gaussian-windowed similarity measure using convolutions. In *Proc. of MMBIA'00*, pages 182–189, Hilton Head Island, USA, June 2000. IEEE Comput. society.

[3] H. Erbe, A. Kriete, A. Jodicke, W. Deinsberger, and D.-K. Boker. 3D-Ultrasonography and Image Matching for Detection of Brain Shift During Intracranial Surgery. *Computer Assisted Radiology*, pages 225–230, 1996.

[4] D. Gobbi, R. Comeau, and T. Peters. Ultrasound/MRI Overlay with Image Warping for Neurosurgery. In *Proc of MICCAI'00*, LNCS 1935, pages 106–114, Pittsburgh, Oct. 2000.

[5] D. Gobbi, C. R.M., and T. Peters. Ultrasound Probe Tracking for Real-Time Ultrasound/MRI Overlay and Visualization of Brain Shift. In *Proc of MICCAI'99*, LNCS 1679, pages 920–927, Cambridge, UK, Sept. 1999.

[6] N. Hata, M. Suzuki, T. Dohi, H. Iseki, K. Takakura, and D. Hashimoto. Registration of Ultrasound echography for Intraoperative Use: A Newly Developed Multiproperty Method. In *Proc. of VBC'94*, SPIE 2359, pages 251–259, Rochester, MN, USA, Oct. 1994.

[7] G. Ionescu, S. Lavallee, and J. Demongeot. Automated Registration of Ultrasound with CT Images: Application to Computer Assisted Prostate Radiotherapy and Orthopedics. In *Proc. MICCAI'99*, LNCS 1679, pages 768–777, Cambridge (UK), Oct. 1999.

[8] A. King, J. Blackall, G. Penney, P. Edwards, D. Hill, and D. Hawkes. Baysian Estimation of Intra-operative Deformation for Image-Guided Surgery Using 3-D Ultrasound. In *Proc of MICCAI'00*, LNCS 1935, pages 588–597, Oct. 2000.

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

[10] N. Pagoulatos, W. Edwards, D. Haynor, and Y. Kim. 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, Dec. 1999.

[11] X. Pennec, P. Cachier, and N. Ayache. Understanding the "demon's algorithm": 3D non-rigid registration by gradient descent. In *Proc. of MICCAI'99*, LNCS 1679, pages 597–605, Cambridge, UK, Sept. 1999.

[12] X. Pennec, P. Cachier, and P. Ayache. Tracking brain deformations in time sequences of 3D US images. In *Proc. of IPMI'01*, LNCS, Davis, CA, USA, June 2001. To appear.

[13] A. Roche, G. Malandain, and N. Ayache. Unifying Maximum Likelihood Approaches in Medical Image Registration. *International Journal of Imaging Systems and Technology: Special Issue on 3D Imaging*, 11:71–80, 2000.

[14] A. Roche, G. Malandain, X. Pennec, and N. Ayache. The correlation ratio as a new similarity measure for multimodal image registration. In *Proc. of MICCAI'98*, LNCS 1496, pages 1115–1124, Cambridge, USA, Oct. 1998.

[15] A. Roche, X. Pennec, M. Rudolph, D. P. Auer, G. Malandain, S. Ourselin, L. M. Auer, and N. Ayache. Generalized Correlation Ratio for Rigid Registration of 3D Ultrasound with MR Images. In *Proc. of MICCAI'00*, LNCS 1935, pages 567–577, Oct. 2000. Submitted to IEEE TMI.

[16] R. N. Rohling, A. H. Gee, and L. Berman. Three-Dimensional Spatial Compounding of Ultrasound Images. *Medical Image Analysis*, 1(3):177–193, 1997.

[17] R. N. Rohling, A. H. Gee, and L. Berman. Automatic registration of 3-D ultrasound images. *Medicine and Biology*, 24(6):841–854, July 1998.

[18] P. J. Rousseeuw and A. M. Leroy. *Robust Regression and Outlier Detection*. Wiley, 1987.

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

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