Generalized Correlation Ratio for Rigid Registration of 3D Ultrasound with MR Images

A. Roche[1], X. Pennec[1], M. Rudolph[2], D. P. Auer[3], G. Malandain[1], S. Ourselin[1], L. M. Auer[2], and N. Ayache[1]

[1] INRIA, Epidaure Project, Sophia Antipolis, France
[2] Institute of Applied Sciences in Medicine, ISM, Salzburg, Austria
[3] Max Planck Institute for Psychiatry, Munich, AG-NMR, Germany

Abstract. Automatic processing of 3D ultrasound (US) is of great interest for the development of innovative and low-cost computer-assisted surgery tools. In this paper, we present a new image-based technique to rigidly register intra-operative 3D US with pre-operative Magnetic Resonance (MR) data. Automatic registration is achieved by maximization of a similarity measure that generalizes the correlation ratio (CR). This novel similarity measure has been designed to better take into account the nature of US images. A preliminary cross-validation study has been carried out using a number of phantom and clinical data. This indicates that the worst registration errors are of the order of the MR resolution.

1 Introduction

Over the last years, the development of real-time 3D ultrasound (US) imaging has revealed a number of potential applications in image-guided surgery. 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.

The registration of US with pre-operative Magnetic Resonance (MR) images will allow the surgeon to accurately localize the course of instruments in the operative field, resulting in minimally invasive procedures. At present, few papers have been published on this particular registration problem [5]. Most of the approaches that have been proposed are based on stereotactic systems. For instance, in [8] registration is achieved by tracking the US probe with a DC magnetic position sensor. Existing image-based methods match homologous features extracted from both the US and MR data. Features are user-identified in [1], while semi-automatically extracted in [2]. More recently, Ionescu *et al* [3] registered US with Computed Tomography (CT) data after automatically extracting contours from the US using watershed segmentation.

The present registration technique expands on the correlation ratio (CR) method [12]. It is an intensity-based approach as it does not rely on explicit feature extraction. In a previous work [11], we reported preliminary results of US/MR registration by maximization of CR and mutual information (MI). While results obtained using CR were more satisfactory than when using MI, the method was still lacking precision and robustness with respect to the initialisation of the transformation parameters.

In this paper, we have improved the CR method following three distinct axes: (1) using the gradient information from the MR image, (2) reducing the number of intensity parameters to be estimated, and (3) using a robust intensity distance. These extensions are presented in the following section, while section 3 proposes an original evaluation of the method accuracy using phantom and clinical data.

2 Method

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 minimizing the following cost function:

$$\min_{T,f} \quad \sum_k \left[I(x_k) - f(J(T(x_k))) \right]^2 , \tag{1}$$

which integrates over the voxel positions x_k in image I. In the following, we will use the simplified notations $i_k \equiv I(x_k)$, and $j_k^{\downarrow} \equiv J(T(x_k))$, where the arrow expresses the dependence in T. 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.

In practice, the criterion defined in eq (1) cannot be computed exactly due to the finite nature of the template image. One obvious problem is that the transformed position of a voxel will generally not match a grid point of J, such that the corresponding intensity $j_k^{\downarrow}$ is unknown. A classical approach is then to linearly interpolate $j_k^{\downarrow}$ using the eight neighbours of $T(x_k)$ in the grid of J. However, instead of interpolating the image intensity, we may directly interpolate the incremental contribution of x_k, i.e., $[i_k - f(j_k^{\downarrow})]^2$. The difference between these two approaches is illustrated in figure 1. In fact, the last method is equivalent to the so-called partial volume (PV) interpolation, originally proposed by Maes *et al* [4] in the context of joint histogram computation. We have found PV to generally outperform classical linear interpolation in terms of smoothness of the resulting registration criterion.

Another difficulty to compute eq (1) is that some points x_k may transform outside the template domain and lack eight grid neighbours. We decide not to take into account such points in the computation of the registration criterion.

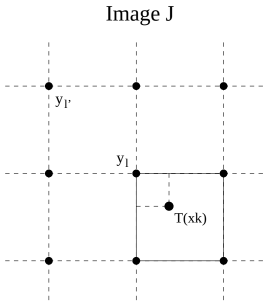

Linear interpolation weights:

$w_{kl}^{\downarrow} = \beta(T(x_k) - y_l)$, with :

$\begin{cases} \beta(z) = (1 - |z^1|)(1 - |z^2|), & \text{if } |z^1| \leq 1,\ |z^2| \leq 1 \\ \beta(z) = 0, & \text{otherwise} \end{cases}$

Intensity interpolation:

$\left[i_k - f(j_k^{\downarrow})\right]^2 \approx \left[i_k - f\left(\sum_l w_{kl}^{\downarrow} j_l\right)\right]^2$

PV interpolation:

$\left[i_k - f(j_k^{\downarrow})\right]^2 \approx \sum_l w_{kl}^{\downarrow} \left[i_k - f(j_l)\right]^2$

Fig. 1. Illustration of linear interpolation in the 2D case, and two related strategies of interpolating the registration criterion.

Doing so without particular attention, the criterion would become zero when every point x_k transforms outside J. Hence, in order to avoid an absolute minimum when the image overlap is small, we impose the additional constraint that the variance of I be large in the overlapping region. Justifications of this particular normalization strategy will be found in [12], while related normalization issues are discussed in [15, 16].

These practical considerations lead us to the following modification of eq (1):

$$C(T, f) = \frac{\sum_{k,l} w_{kl}^{\downarrow} \left[i_k - f(j_l)\right]^2}{n^{\downarrow} \mathrm{Var}(I^{\downarrow})}, \tag{2}$$

where j_l is the intensity of a voxel with coordinates y_l in the coordinate system of J. The terms depending on T are marked with an arrow: $w_{kl}^{\downarrow}$ are the linear interpolation weights, $n^{\downarrow}$ is the number of points x_k such that $T(x_k)$ has eight neighbours in the template grid, and $\mathrm{Var}(I^{\downarrow})$ is the intensity variance computed over these points.

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 [12]. 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 maximization 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 *difference* of acoustical impedance caused by successive tissue layers. Ideally, the US signal should be high at the interfaces, and zero within homogeneous tissues.

As stated above, the CR method tries to predict the intensities of the US by remapping those of the MR. Hence, uniform regions of the original MR will remain uniform in the remapped MR and, thus, this procedure is not able to account for intensity variations at the interfaces. To enable a better prediction, we propose to use the modulus of the MR gradient as an additional explanatory variable. In other terms, our template image J is now a vectorial image, $J = (M, \|\nabla M\|)$, M standing for the MR image, and we search for a function f that maps double-valued intensities to single-valued intensities. The MR gradient is practically computed by convolution with a Gaussian kernel.

At first glance, using the modulus of the MR gradient does not appear to be fully acceptable from the physics of US imaging. In fact, the US signal which is produced at an interface also depends on the tissue orientation with respect to the scan line. Thus, perhaps a more appropriate choice than $\|\nabla M\|$ would be the dot product, $\nabla M.\mathbf{u}$, where $\mathbf{u}$ is the scan direction. The main difficulty in using this last expression is that $\mathbf{u}$ is unknown before registration since it depends on the position of the US probe in the MR coordinate system.

We have not yet studied the effect of using the projected MR gradient versus the modulus. Still, we believe that there are good reasons not to take into account information from the gradient orientation, at least as a first-order approximation. Through diffraction of the ultrasound beam on interfaces, the received echo is actually less dependent on the direction propagation than would be the case with perfectly specular reflection. This, added to log-compression, tends to equalize the response values corresponding to different tissue orientations.

2.3 Parametric intensity fit

If we put no special constraint on the mapping f to be estimated, then f is described by as many parameters as there are distinct intensity values in the template image [11]. That approach makes sense as long as the number of intensity classes in J is small with respect to the number of voxels used to draw an estimate. In our case, J is a double-valued image (with, in general, floating precision encoding of the MR gradient component), and the number of parameters to be estimated becomes virtually infinite.

We will therefore restrict our search to a polynomial function f. Let m_l and g_l denote the intensity of the voxel with coordinates y_l, respectively in the original MR, M, and in $\|\nabla M\|$. We are searching for a mapping of the form:

$$f(m_l, g_l) = \sum_{p+q \leq d} \theta_{pq}\, m_l^p g_l^q, \qquad (3)$$

where d is the specified polynomial degree. The number of parameters describing f then reduces to $(d+1)(d+2)/2$. In all the experiments presented below, the degree was set to $d = 3$, implying that 10 coefficients were estimated. It is shown in [13] that minimizing eq (2) with respect to the polynomial coefficients brings us to a weighted least square (WLS) linear regression problem. As is standard, this is solved by the method of singular value decomposition (SVD).

This polynomial fitting procedure, however, has significant extra computational cost with respect to the unconstrained fitting. Recall that, in the basic version of the CR method, f is updated for each transformation trial. Such a strategy is no longer affordable when estimating a polynomial function. Instead, the minimization of eq (2) may be performed alternatively along T and f, resulting in the following algorithm: (1) given a current transformation estimate T, find the best polynomial f and remap J accordingly; (2) given a remapped image $f(J)$, minimize $C(T, f)$ 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, and this even when the images are perfectly registered. Such bad intensity matches are false negative.

The sensitivity of the registration criterion to false negative may be reduced by replacing the expression $(1/n^{\downarrow}) \sum_{k,l} w^{\downarrow}_{kl} \left[i_k - f(j_l)\right]^2$ in eq (2) with a robust scale estimate. A similar idea was developed in [6]. We propose here to build such an estimate from a one-step S-estimator [14]:

$$\hat{S}^2(T, f) = \frac{S_0^2}{Kn^{\downarrow}} \sum_{k,l} w^{\downarrow}_{kl}\, \rho\left(\frac{i_k - f(j_l)}{S_0}\right),$$

where ρ is the objective function corresponding to a given M-estimator, K is a normalization constant to ensure consistency with the normal distribution, and S_0 is some initial guess of the scale. The new registration criterion is then: $C(T, f) = \hat{S}^2(T, f)/\mathrm{Var}(I^{\downarrow})$.

This criterion implies few modifications of our alternate minimization strategy. As a function of T, it may still be minimized by means of Powell's method. As a function of f, the solution is found by a simple iterative WLS procedure as shown in [13], generally requiring no more than 5-6 iterations.

In our implementation, we have opted for the Geman-McClure ρ-function, $\rho(x) = \frac{1}{2}x^2/(1 + \frac{x^2}{c^2})$, 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. The normalization constant is then $K = 0.416$.

Initially, the intensity mapping f is estimated in a non-robust fashion. The starting value S_0 is then computed as the weighted median absolute deviation of the corresponding residuals, $\{|i_k - f(j_l)|\}$ (see [13] for details). Due to the initial misalignment, S_0 tends to be overestimated. Thus, it may not allow to reject efficiently bad intensity matches. For that reason, we reset S_0 at each new iteration, i.e., after completing one minimization along T and one minimization along f.

3 Experiments

The experiments related in this section were performed within the framework of the European project ROBOSCOPE[1]. The goal is to assist neuro-surgical operations using real-time 3D ultrasound images and a robotic manipulator arm. The operation is planned on a pre-operative MRI and 3D US images are acquired during surgery to track in real time the deformation of anatomical structures. In this context, the rigid registration of the pre-operative MR with the first US image (dura mater still closed) is a fundamental task to relate the position of the surgical instruments with the actual anatomical structure. This task being determinant for the global accuracy of the system, different datasets were acquired to simulate the final image quality and to perform accuracy evaluations.

It should be emphasized that all the US images provided in this project were stored in Cartesian format, which means that the actual (log-compressed) ultrasound signal is resampled on a regular cubic lattice. As a consequence, the images undergo sever interpolation artifacts (blurring) in areas which are distant from the probe. In the following, we will refer to US images as cubic images, but one has to keep in mind that this is somewhat artificial. Notably, the voxel size in Cartesian US images should not be confused with the real spatial resolution, which is in fact spatially dependent.

We computed all the MR/US registrations using the previously described algorithm. 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 transformation. Here, 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 [8]. After initialization, we observed that the algorithm found residual displacements in the range of 10 mm and 10 degrees. In all the experiments, the gradient norm of the MR image was computed by linear filtering using a Gaussian kernel with $\sigma = 1$ voxel.

3.1 Principle of the accuracy evaluation

To estimate the accuracy of the algorithm, one should ideally compare the result of a registration with a gold-standard. Up to our knowledge, there is no such gold-standard for MR/US registration. To get around this problem, our main idea is to use several MR and/or US images to compute registration loops and test for the residual error on test points. What we call a registration loop is a succession of transformation compositions that sould ideally lead to the identity transformation. A typical loop is a sequence of the form $US_i \rightarrow MR_i \rightarrow MR_j \rightarrow US_j \rightarrow US_i$ in the case of the Phantom data described below.

If we were given perfectly registered images within each modality, this loop would only be disturbed from the identity by errors on the two MR/US registrations. As the variances are additive, the variance of the observed error should

[1] http://www.ibmt.fhg.de/Roboscope/home.htm

roughly be: $\sigma^2_{loop} = 2\sigma^2_{MR/US}$. Unfortunately, we are not provided with a ground truth registration within each modality: we need to estimate it. This time, as we are combining one MR/MR, one US/US and two MR/US registrations, the variance of the loop error will be roughly: $\sigma^2_{loop} \simeq 2\sigma^2_{MR/US} + \sigma^2_{MR/MR} + \sigma^2_{US/US}$. The *expected* MR/US accuracy is then $\sigma_{MR/US} \simeq \sqrt{(\sigma^2_{loop} - \sigma^2_{MR/MR} - \sigma^2_{US/US})/2}$.

However, what we really measure is the maximum or *conservative* MR/US accuracy, $\sigma_{MR/US} \simeq \sigma_{loop}/\sqrt{2}$. In order to minimize the influence of intra-modality registrations errors in this figure, we need to provide very accurate MR/MR and US/US registrations. For that purpose, we designed the following algorithm.

Multiple intra-modality registration To relate n images together, we need to estimate $n-1$ rigid registrations $\bar{T}_{i,i+1}$. To obtain a very good accuracy, we chose to register all image pairs, thus obtaining $n(n-1)$ transformations $\bar{T}_{i,j}$, and estimate the transformations $\bar{T}_{i,i+1}$ that best explain our measurements in the least-square sense, i.e. that minimizes the following criteria:

$$C(\bar{T}_{1,2}, \ldots \bar{T}_{n-1,n}) = \sum_{i \neq j} \mathrm{dist}^2(\bar{T}_{i,j}, T_{i,j}),$$

where $\bar{T}_{i,j}$ is recursively defined by $\bar{T}_{i,j} = \bar{T}_{j-1,j} \circ \ldots \circ \bar{T}_{i,i+1}$ if $j > i$, and $\bar{T}_{i,j} = \bar{T}^{(-1)}_{j,i}$ if $j < i$.

We used a robust variant of the left invariant distance on rigid transformations introduced in [10]: let σ_r and σ_t be typical scales on the rotation angle and on the translation magnitude and χ^2 a threshold. If (r, t) are the rotation vector and the translation of transformation T, the distance between two transformations is

$$\mathrm{dist}^2(T_1, T_2) = \min\left(\|r_2^{(-1)} \circ r_1\|^2/\sigma_r^2 + \|t_1 - t_2\|^2/\sigma_t^2 \ ,\ \chi^2 \right).$$

The standard deviations σ_r and σ_t are manually adjusted to correspond roughly to the residual rotation and translation error after convergence. To obtain the minimum, we used a Newton gradient descent similar to the one described in [10], but on all transformations $\bar{T}_{i,i+1}$ together (see [13] for details).

3.2 Data

MR and US compatible Phantom We developed for the European Project ROBOSCOPE an MR and US compatible phantom made of two balloons that can be inflated with known volumes in order to simulate deformations. In this experiment, we used 8 acquisitions with different balloons volumes, each acquisition consisting of one 3D MR and one 3D US image. However, we cannot directly compare the MR/US registrations as the phantom is moved between the acquisitions. Thus, the first step is to rigidly register all the MR images together and similarly for the US images.

The main problem for the multiple intra-modality registration of the phantom images is that the balloons deform inbetween acquisitions. The only rigid part is the outer part of the container. Thus, intra-modal registrations of MR images were carried out using a feature-based registration algorithm known to handle a large amount of outliers [9]. For the US images, since it is very difficult to extract meaningful features on these images, we used the robust block matching technique proposed in [7].

As we are testing the *rigid* registration, we cannot register MR and US images across acquisitions. Thus, the simplest loops we can use for accuracy estimations are the $n(n-1)$ following loops: $US_i \rightarrow MR_i \rightarrow MR_j \rightarrow US_j \rightarrow US_i$. Of course, only $n-1$ loops are independent but since the ideal value is known (the identity) there is no need to correct the estimation for the number of free parameters.

Baby dataset with degraded US images 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 and test the 30 following loops $US_i \rightarrow MR \rightarrow US_j \rightarrow US_i$ (only 5 of them being independent). For that, we still need to register the US images together. Since there are no deformations between the acquisitions (we only have a motion of the probe and a degradation of the image quality) the algorithm is much more efficient and accurate than for the Phantom.

Patient dataset during tumor resection This dataset 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 tumor to resect. In this experiment, we use the three US images that are large enough to contain the ventricles.

The two MR images were registered using the feature based method with a very high accuracy (probably overestimated as we only have two images) and we tested the loops $US_i \rightarrow MR_0 \rightarrow MR_1 \rightarrow US_i$. As the acquisition cone of the US probe is completely within the Cartesian image (see Fig. 4), the region of interest is much smaller than the images size: we took our typical points at the corners of a $80 \times 80 \times 80$ mm^3 cube centered in the image.

3.3 Results and discussion

We put in table (1) the standard deviations of the residual rotation, of the residual translation and of the displacement of the test points for the different registration involved. Since we took the origin of the images at the center, the σ_{trans} value corresponds to the mean error at the center of the image while σ_{test} corresponds to the maximum registration error within the US image (except for the patient experiment, the test points are taken at the corners of the image).

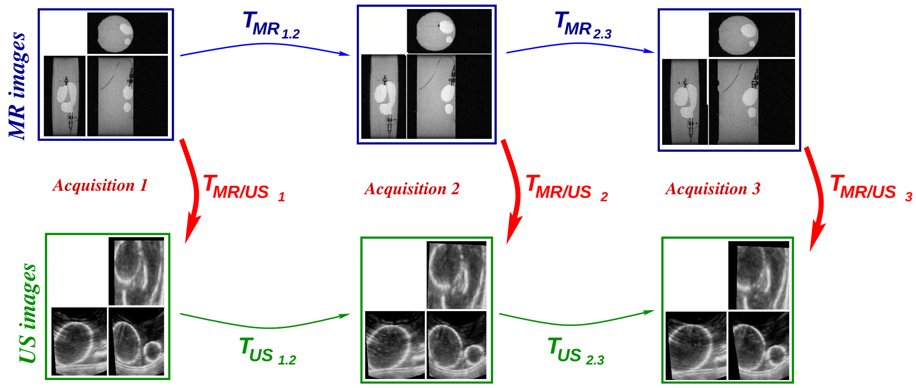

Fig. 2. MR and US images of the Phantom and the rigid registrations that are involved.

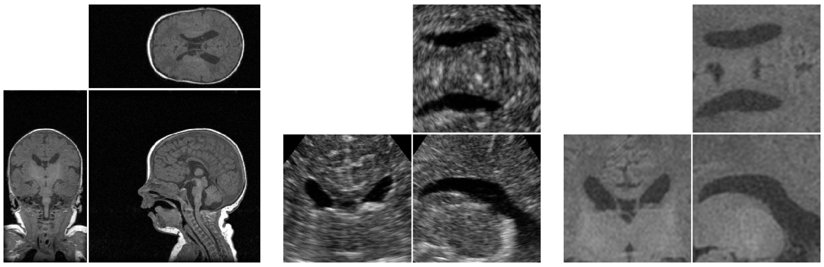

Fig. 3. Example registration of the MR and US images of the baby. From left to right: original MR T1, original US and registered MR.

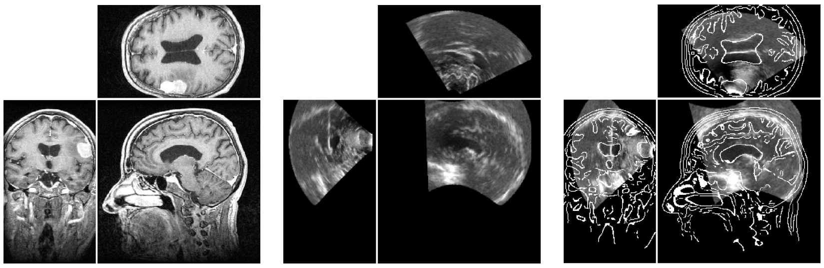

Fig. 4. Example registration of the MR and US images of the patient. From left to right: MR T1 image with a contrast agent, manual initialisation of the US image registration, and result of the automatic registration with the MR contours superimposed.

The results on the Phantom show that the MR/US registration accuracy is of the order of the MR resolution. One could probably expect a better conservative accuracy by acquiring larger US image including some rigid landmarks for multiple US/US registration. One finds the same type of results for the other

Table 1. Estimated registration errors. See text for details.

		σ_{rot} (deg)	σ_{trans} (mm)	σ_{test} (mm)
Phantom dataset	Multiple MR/MR	0.06	0.10	0.13
MR: $0.9 \times 0.9 \times 1$ mm^3	Multiple US/US	0.60	0.40	0.71
US: 0.41^3 mm^3	Conservative MR/US	1.15	1.01	1.46
	Expected MR/US	1.06	0.97	1.37
Baby dataset	Multiple US/US	0.10	0.06	0.12
MR: 0.9^3 mm^3	Conservative MR/US	1.21	0.36	0.90
US: 0.3^3 mm^3	Expected MR/US	1.21	0.36	0.89
Patient dataset	MR/MR	0.02	0.02	0.03
MR: $0.9 \times 0.9 \times 1.1$ mm^3	Conservative MR/US	1.57	0.58	1.65
US: 0.63^3 and 0.95^3 mm^3	Expected MR/US	1.57	0.58	1.65

datasets: slightly under the MR voxel size for the baby data and a bit larger for the patient data.

However, when we look more carefully at the patient 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 2.61 mm ($\sigma_{\text{trans}} = 0.84$ mm) 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 ($\sigma_{\text{trans}} = 0.39$ mm). We suspect that a non-rigidity in the smallest US could accounts for the registration difference between the two MR images. 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.

4 Conclusion

We have presented a new automated method to rigidly register 3D US with MR images. It is based on a multivariate and robust generalization of the correlation ratio (CR) measure that allows to better take into account the nature of US images. Incidentally, we believe that the generalized CR could be considered in other registration problems where conventional similarity measures fail.

Testings were performed on several phantom and clinical data, and accuracy was evaluated using an original method that does not require the knowledge of a ground truth. We estimated the worst registration errors (errors at the Cartesian US corners) to be of the order of 1 millimeter.

In our experiments, registration was tested with US images stored in Cartesian format. This obviously does not help registration owing to the fact that: 1) intensities of voxels distant from the probe are unreliable, 2) resampling on a cubic lattice artificially increases the concentration of such voxels, consequently increasing their influence on the registration criterion. The present algorithm may be straightforwardly extended to use polar US images as inputs (without interpolation of the US signal, see section 2). We believe that this could significantly improve both the accuracy and the robustness of registration.

Further developments also include non-rigid registration, in order to correct for false distance artefacts in US, as well as for tissue deformations due to brain shift and operative manipulations.

Acknowledgements This work was partially supported by la Région PACA (France), and 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).

References

1. H. Erbe, A. Kriete, A. Jödicke, W. Deinsberger, and D.-K. Böker. 3D-Ultrasonography and Image Matching For Detection of Brain Shift During Intracranial Surgery. In *Appl. of Comp. Vision in Med. Image Processing*, volume 1124 of *Excerpta Medica – Int. Congress Series*, pages 225–230. Elsevier, 1996.
2. 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 *Visualization in Biomedical Computing*, volume 2359 of *Proc. SPIE*, pages 252–259, Rochester, MN, USA, October 1994. SPIE Press.
3. G. Ionescu, S. Lavallée, and J. Demongeot. Automated Registration of Ultrasound with CT Images: Application to Computer Assisted Prostate Radiotherapy and Orthopedics. In *Proc. MICCAI'99*, volume 1679 of *LNCS*, pages 768–777, Cambridge, UK, 1999. Springer Verlag.
4. F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality Image Registration by Maximization of Mutual Information. *IEEE Transactions on Medical Imaging*, 16(2):187–198, 1997.
5. J. B. A. Maintz and M. A. Viergever. A survey of medical image registration. *Medical Image Analysis*, 2(1):1–36, 1998.
6. C. Nikou, F. Heitz, J.-P. Armspach, and I.-J. Namer. Single and multimodal subvoxel registration of dissimilar medical images using robust similarity measures. In *Conf. on Med. Imaging*, volume 3338 of *SPIE*, pages 167–178. SPIE Press, 1998.
7. S. Ourselin, A. Roche, S. Prima, and N. Ayache. Block Matching: a General Framework to Improve Robustness of Rigid Registration of Medical Images. In *MICCAI'2000*, Pittsburgh, USA, October 2000.
8. N. Pagoulatos, W.S. Edwards, D.R. Haynor, and Y. Kim. Interactive 3-D Registration of Ultrasound and Magnetic Resonance Images Based on a Magnetic Position Sensor. *IEEE Trans. on Information Technology in Biomed.*, 3(4):278–288, 1999.
9. X. Pennec, N. Ayache, and J.P. Thirion. Chap. 31: Landmark-based registration using features identified through differential geometry. In *Handbook of Medical Imaging*. Academic Press, 2000. In press.
10. X. Pennec, C.R.G. Guttmann, and J.P. Thirion. Feature-based registration of medical images: Estimation and validation of the pose accuracy. In *Proc. MICCAI'98*, number 1496 in LNCS, pages 1107–1114, Cambridge, USA, 1998. Springer Verlag.
11. A. Roche, G. Malandain, and N. Ayache. Unifying Maximum Likelihood Approaches in Medical Image Registration. *Int. J. of Imaging Systems and Technology*, 11:71–80, 2000.
12. A. Roche, G. Malandain, X. Pennec, and N. Ayache. The Correlation Ratio as a New Similarity Measure for Multimodal Image Registration. In *Proc. MICCAI'98*, volume 1496 of *LNCS*, pages 1115–1124, Cambridge, USA, 1998. Springer Verlag.
13. 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. Technical report, INRIA, 2000. In press.
14. Peter J. Rousseeuw and Annick M. Leroy. *Robust Regression and Outlier Detection*. Wiley Series In Probability And Mathematical Statistics, first edition, 1987.
15. C. Studholme, D. L. G. Hill, and D. J. Hawkes. An overlap invariant entropy measure of 3D medical image alignment. *Pattern Recognition*, 1(32):71–86, 1998.
16. P. Viola and W. M. Wells. Alignment by Maximization of Mutual Information. *Intern. J. of Comp. Vision*, 24(2):137–154, 1997.