Global Registration of Ultrasound to MRI Using the LC^2 Metric for Enabling Neurosurgical Guidance

Wolfgang Wein[1], Alexander Ladikos[1], Bernhard Fuerst[2], Amit Shah[2], Kanishka Sharma[2], and Nassir Navab[2]

[1] ImFusion GmbH, München, Germany
[2] Computer Aided Medical Procedures, Technische Universität München, Germany

Abstract. Automatic and robust registration of pre-operative magnetic resonance imaging (MRI) and intra-operative ultrasound (US) is essential to neurosurgery. We reformulate and extend an approach which uses a Linear Correlation of Linear Combination (LC^2)-based similarity metric, yielding a novel algorithm which allows for fully automatic US-MRI registration in the matter of seconds. It is invariant with respect to the unknown and locally varying relationship between US image intensities and both MRI intensity and its gradient. The overall method based on this both recovers global rigid alignment, as well as the parameters of a free-form-deformation (FFD) model. The algorithm is evaluated on 14 clinical neurosurgical cases with tumors, with an average landmark-based error of 2.52 mm for the rigid transformation. In addition, we systematically study the accuracy, precision, and capture range of the algorithm, as well as its sensitivity to different choices of parameters.

1 Introduction

Modern neurosurgery heavily relies on both pre-operative and interventional medical imaging, in particular MRI and US. MRI provides a good visualization of tumors, a relatively large field of view and good reproducibility. Its use as intra-operative imaging modality is possible [8], however with limited accessibility to the patient and high workflow complexity. On the other hand, US is inexpensive and easy to use, but imaging quality is reduced, fewer anatomical details are visible and, in general, US is operator-dependent and harder to interpret. Also, the field of view is limited and direction dependent. The combination of both modalities would allow to integrate high-contrast pre-operative MRI data into the interventional suite. Therefore, quick, robust and automatic alignment of MRI and US images is of high importance. In contrast to MRI, ultrasound provides real-time 2D images, which, when tracking the ultrasound transducer, can be interpreted in 3D space. This has been used in the past decades for brain examinations, for instance to localize tumors, determine their tissue and boundary properties, and detect brain shift.

Registering US and MRI images is a complex and still not satisfactorily solved problem, mostly because the nature of the represented information is completely

K. Mori et al. (Eds.): MICCAI 2013, Part I, LNCS 8149, pp. 34–41, 2013.

different for both modalities. MRI intensities are correlated to relaxation times, which in turn depend on the tissue type and hydrogen concentration, while US images are a representation of acoustic impedance transitions. Those in turn can be both reflections on large structures (hence correlating to some extent to the gradient of MRI), or reflections from the tissue inhomogeneities that cause the characteristic brightness and texture of certain tissue types (in that case correlating to the MRI intensities directly). As an additional challenge, US exhibits various direction-dependent artifacts.

Therefore, the basic and well known off-the-shelf registration approaches are known to fail, which includes registration using cost functions based on sum of squared distances, mutual information [4] or correlation ratio [12]. A method which uses a measure based on 3D gradient orientations in both US and MRI is presented in [3], however such an approach discards valuable MRI intensity information and hence requires either optimal data or close initialization. Many of the best existing approaches transform MRI and/or US intensities under application- and organ-specific considerations, in order to make them easily comparable. This is done, for example, for liver vasculature in [9], with significant effort due to learning-based pre-processing. Similarly, pseudo-US images may be generated using segmented structures from MRI [1,2,5,6]. In light of the modality-specific considerations, the most promising general strategy for robust US-MRI registration, without relying on application-specific pre-processing or segmentation, is to compare US to both the MRI intensity and its gradient, as pioneered in [12], where a global polynomial intensity relationship is fitted during registration. The alternating optimization of the rigid pose and the polynomial coefficients, as well as the fact that it is a global mapping, limit the convergence range. Higher-dimensional Mutual Information (α-MI) is theoretically suited to assess US-MRI alignment based on both intensity and gradient information (in fact, an arbitrary number of features may be used). However, current approaches are neither practical in terms of implementation effort nor computation time [11]. Powerful tools for image registration are similarity measures which are invariant to local changes, such as local normalized cross-correlation (invariant wrt. local brightness and contrast). In [13] the similarity measure Linear Correlation of Linear Combination (LC^2) is presented, which exhibits local invariance to how much two channels of information contribute to an ultrasound image. The entire method has been specially designed for US-CT registration, where a strong correlation between X-ray attenuation coefficients and acoustic impedance is known, which allows a simulation of ultrasound effects from CT. These incorporate estimates of the acoustic attenuation, multiple reflections, and shadowing, which can not directly be estimated from MRI.

In this paper, we adapt the LC^2 formulation for the registration of interventional US to (pre-operative) MRI, and extend it to non-linear deformations. It results in a globally convergent, robust new algorithm, which we evaluate on a database of 14 patients with ground-truth information.

2 Method

Similarity Measure: Instead of correlating US intensities with two channels of simulated information from CT as in [13], we use LC^2 to correlate US with both the MRI intensity values p and its spatial gradient magnitude $g = |\nabla p|$. The local LC^2 value is computed for each pixel $\mathbf{x_i}$ in each ultrasound image, considering a neighborhood $\Omega(\mathbf{x_i})$ of m pixels. For each patch of m pixels, the contribution of MRI intensity values p and gradient magnitudes g are unknown. Therefore, we define an intensity function $f(\mathbf{x_i})$ as a function of the transformed MRI intensities $p_i = p(T(\mathbf{x_i}))$ and gradients $g_i = g(T(\mathbf{x_i})) = |\nabla p_i|$ as:

$$f(\mathbf{x_i}) = \alpha p_i + \beta g_i + \gamma, \tag{1}$$

where $y_i = \{\alpha, \beta, \gamma\}$ denotes the unknown parameters of the influence of the MRI intensities and gradients within $\Omega(\mathbf{x_i})$. They can be estimated by minimizing the difference of the intensity function and the ultrasound image intensity u_i:

$$\left\| \mathrm{M} \begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix} - \begin{pmatrix} u_1 \\ \vdots \\ u_m \end{pmatrix} \right\|^2 \quad \text{where } \mathrm{M} = \begin{pmatrix} p_1 & g_1 & 1 \\ \vdots & \vdots & \vdots \\ p_m & g_m & 1 \end{pmatrix}, \tag{2}$$

which can be solved using ordinary least squares with the pseudo-inverse of M. This results in a parameter triple y_i for each pixel $\mathbf{x_i}$, which is only depending on the neighborhood $\Omega(\mathbf{x_i})$ and therefore compensating for changing influences of tissue interfaces or organ-internal intensities. The local similarity is then:

$$S(u, M) = 1 - \frac{\sum_{\mathbf{x_i}} |u(\mathbf{x_i}) - My|^2}{\sum_{\mathbf{x_i}} Var(u(\mathbf{x_i}))} \tag{3}$$

The overall similarity is the weighted sum of eq. 3 with the local variance of the US image. This suppresses regions without structural appearance, therefore allowing to cope with ultrasonic occlusions implicitly, without the need to simulate them.

Computation: We compute LC^2 on the original 2D US image slices located in 3D-space through the tracking data, as opposed to using a 3D compounded volume. This has four main advantages:

1. An unnecessary resampling step of the ultrasound image data, which may degrade its quality, is avoided.
2. Registration can start as soon as the first US frames are available, catering to real-time applications in the operating room.
3. One may also optimize the US probe calibration parameters or further parameters expressing e.g. tracking system errors.
4. Image information within slices is inherently more consistent, because all scan-lines of an image are acquired within a short duration and void of tracking errors.

Extracting the MRI intensity and gradient at the presumed location in the 3D volume is performed on the GPU using its hardware tri-linear interpolation. We then compute eq. 3 using a multi-core recursive filtering strategy; however a GPU implementation is equally possible. The MRI image data is used in full resolution (typically iso-tropic voxel size of $0.5mm$). The higher-resolution US frames are down-sampled such that their pixel size is smaller than twice the voxel size (assuring that tri-linear interpolation never discards MRI voxel information within the oblique US planes). Similarly, US frames are skipped such that no overlapping planes occur, with average spacing between the image centers $<$ $1.5mm$ (mostly yielding smaller spacing in the area of interest).

Optimization of Rigid Transformation: Due to the least-squares fitting in eq. 2 which is computed for every US image pixel, an analytic derivative of LC^2 is difficult to compute. Therefore we use Bound Optimization by Quadratic Approximation (BOBYQA) [10], which internally creates own derivative approximations, resulting in fewer evaluations than most other direct search methods. While we recommend and use this optimizer throughout this paper, clinical requirements on capture range may necessitate other techniques. In particular, global optimization strategies may be chosen that perform a more thorough search within specified bounds.

Deformable Registration: The optimizer first seeks the 6 rigid transformation parameters as described above. Successively, we use a free-form deformation model with cubic splines, which is applied in the same GPU kernel which extracts MRI intensity/gradient information. More specifically, $2x2x4$ control points are placed within the bounding box of the rigidly registered ultrasound sweep. The 3-displacement vectors are then optimized for all control points using BOBYQA.

3 Experiments and Results

3.1 Clinical Data

To evaluate our method and compare the results to other publications, we used a publicly available database containing Brain Images with Tumors for Evaluation from Montreal Neurological Institute (MNI BITE) [7], with pre-operative T1-weighted MRI and pre-resection 3D freehand US from 14 patients. Initial transformations and corresponding landmarks for each US-MRI pair are included. Therefore we can provide ground truth evaluations, and denote the average Euclidean distance of the landmarks as Fiducial Registration Error (FRE).

3.2 Registration Results

Tab. 1 depicts the results of our algorithm for all 14 data sets.[1] Our rigid registration yields almost exactly the same FRE values as [3], which suggests that both methods achieve the correct optimum transformation. The errors in [11] are

[1] Computation times measured with Intel i7-3770 CPU and NVIDIA GTX 570 GPU.

Table 1. Overview of clinical data [7], previous published results [3,11], and results using our method for rigid and deformable registration including computation times

Dataset Overview and Related Methods															
Patient Number	1	2	3	4	5	6	7	8	9	10	11	12	13	14	mean
Number of Tags	37	35	40	32	31	37	19	23	21	25	25	21	23	23	-
Initial FRE (mm)	4.93	6.30	9.38	3.93	2.62	2.30	3.04	3.75	5.09	2.99	1.52	3.70	5.15	3.77	4.18±5.20
US Spacing (mm)	0.24	0.42	0.23	0.20	0.25	0.17	0.24	0.18	0.18	0.22	0.16	0.18	0.21	0.19	0.22±0.20
FRE in [3] (mm)	4.89	1.79	2.73	1.68	2.12	1.81	2.51	2.63	2.7	1.95	1.56	2.64	3.47	2.94	2.53±0.87
FRE in [11] (mm)	-	2.05	2.76	1.92	2.71	1.89	2.05	2.89	2.93	2.75	1.28	2.67	2.82	2.34	2.57±0.82
Registration Results using LC^2															
FRE Rigid (mm)	4.82	1.73	2.76	1.96	2.14	1.94	2.33	2.87	2.81	2.06	2.18	2.67	3.58	2.48	2.52 ±0.87
Precision (mm)	0.01	0.01	0.01	0.01	0.02	0.01	0.05	0.30	0.02	0.00	0.03	0.15	0.05	0.04	0.05 ±0.08
Time Rigid (sec)	5.9	8.3	11.1	5.7	7.1	8.2	18.2	8.6	6.0	23.4	17.3	25.8	8.1	7.0	11.5 ±6.8
FRE Def. (mm)	4.95	1.64	2.43	1.91	2.26	2.2	2.52	3.64	2.65	2.09	1.76	2.45	3.71	2.76	2.64 ±0.9
Time Def. (sec)	158	141	279	92	133	166	563	312	76	675	597	93	106	282	262 ±204

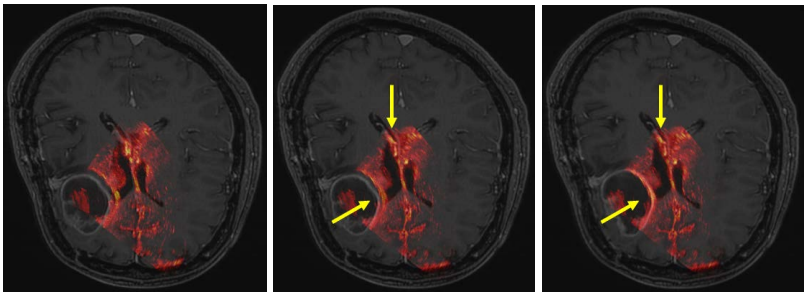

(a) Initial (b) Rigid (c) Deformable

Fig. 1. Registration result of patient 6, with US superimposed on an axial MRI slice

slightly higher (with computation times of several hours per data set), which indicates that applying their proposed deformable registration to the mostly rigid[2] data does not provide much benefit. Similarly, our deformable registration can further improve on the FRE value only in a few of the cases. We believe that the change of landmark errors induced by deformable registration lies within the range of the fiducial localization error (FLE) of the data, especially since some of the landmarks are located along boundaries, not only 3D-corner structures. Fig. 1 depicts the registration results on patient 6. It can be seen that the visual alignment is significantly improved after deformable registration.

3.3 Accuracy, Precision and Capture Range

While some initial errors are significantly away from the correct optimum (e.g. patients 2 and 3), an analysis about the suitability of our algorithm to reach the optimum under all conditions is required. Randomized trials were

[2] The pre-resection ultrasound has been acquired before opening the dura.

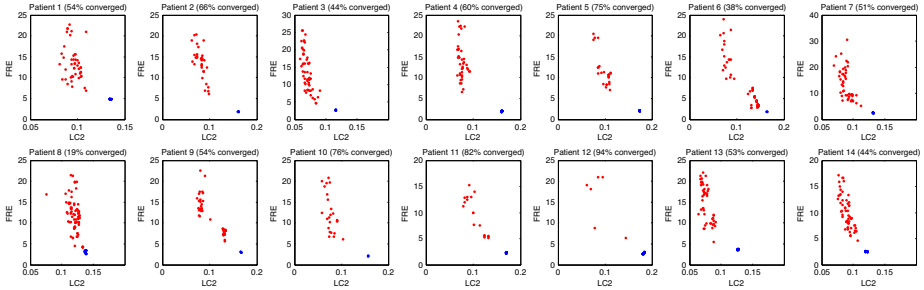

Fig. 2. Random study results of all patients. The converged results are clustered in the lower right corner (highlighted in blue, outliers red).

executed on all data sets, repeating the rigid registration for randomly displaced starting positions ($\pm 10mm/^\circ$ in all 6 parameters). Comparing the final FRE against the LC^2 metric, we discovered that the best pose is always perfectly separated by a significantly higher similarity. This proves that our method allows for global registration; the results for all patients are shown in Fig. 2, including the percentage of converged executions. The mean value of the converged results is the *accuracy*, its standard deviation the *precision*, and the range of initial FRE values from which all executions converge with a smaller than desired number of outliers is the *capture range*. The outlier behavior is also visualized in figure 4(a), where their number is plotted against the initial FRE for some patients. For an initial FRE $< 8mm$ a single local optimization with BOBYQA is sufficient.

Unfortunately, precision and capture range are often not reported in the literature. Since the gradient orientation alignment (GOA) method [3] yields similar FRE values, we implemented it and re-ran the aforementioned randomized trials with it. We obtain $> 90\%$ outliers; further investigation into the cost function properties revealed that only a minor local optimum is present, see Fig. 3(b) for an example. A possible explanation is, that without further heuristics the GOA method would line up strong gradients from e.g. dura mater or skull; besides, using only gradients larger than a threshold limits the image content considered, preventing a smooth similarity increase. While we believe these to be general issues, it has to be acknowledged that better results may be obtained by changing certain implementation details such as resolution, smoothing and interpolation.

3.4 Parameter Sensitivity

To investigate the sensitivity of our method to the choice of parameters, we computed accuracy and precision for different LC^2 patch sizes, and number of US frames used (i.e. spacing in between), see Fig. 4. The registration results are similar for patch sizes 2-24, therefore our method is rather insensitive to the choice of this parameter. For all other results presented, we used a value of 9 (hence $m = (2 * 9 + 1)^2 = 361$). Overly large patches result in a global mapping of MRI intensity and gradient, removing the main advantage of LC^2 over other methods (robustness wrt. local changes of intensity-gradient relationship). Consistently good results are obtained with an average spacing $< 5\,mm$.

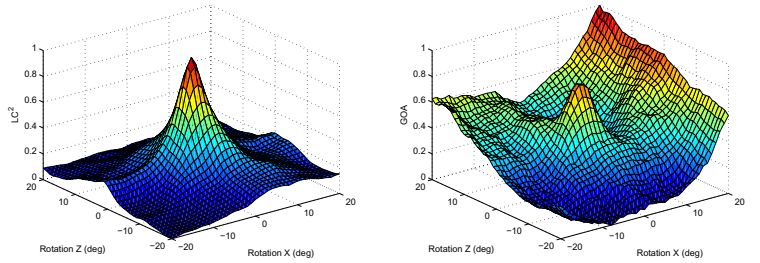

(a) LC^2 (b) GOA

Fig. 3. Plots of different cost functions for two rotation parameters on patient 2

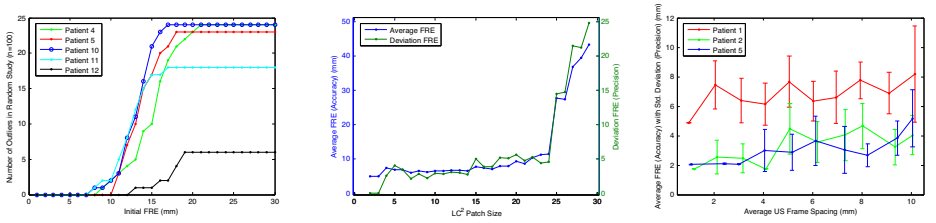

(a) Capture Range (b) LC^2 Patch Size (c) Frame Spacing

Fig. 4. Relationship between initial FRE and outliers (a); dependency of accuracy on LC^2 patch size (b) and US frame spacing (c)

For deformable registration, we chose $< 1.5\,mm$ to make sure we are not missing even smallest structures. Last but not least, we have investigated the effect of using the dot product of the MRI gradient g with the US beam direction, instead of g directly. This reduces the influence of vertical gradients, similar to the US simulation in [13]. Interestingly, this results in $10 - 25\%$ more outliers (the cost function becomes more non-linear due to the added directional dependance).

4 Conclusion

We have introduced an algorithm based on the LC^2 similarity metric, which can rigidly register US-MRI data within a few seconds, and non-linearly within a few minutes. Apart from its efficiency and global convergence, its main strength lies in its simplicity. As opposed to previous works involving the simulation of US imaging, we actually refrain from using the US beam direction and attenuation (which is more difficult with MRI as opposed to CT), and directly compare the MRI intensity and 3D gradient magnitude to US. This smoothens the topology of the cost function, while LC^2 at the same time locally picks the most suited structures. We have shown that the cost function allows for global registration, thoroughly evaluating our method on all 14 patients of an US-MRI image database. We obtain superior results both in terms of computation time and robustness with respect to previously proposed methods on the same data.

Acknowledgments. The authors affiliated with Technische Universität München are partially supported by the EU 7th Framework Program (FP7/2007-2013 and FP7/ICT-2009-6) under Grant Agreements No. 256984 (EndoTOFPET-US) and No. 270460 (ACTIVE) and by a Marie Curie Early Initial Training Network Fellowship under contract number (PITN-GA-2011-289355-PicoSEC-MCNet).

References

1. Comeau, R.M., Sadikot, A.F., Fenster, A., Peters, T.M.: Intraoperative Ultrasound for Guidance and Tissue Shift Correction in Image-Guided Neurosurgery. Medical Physics 27, 787 (2000)
2. Coupé, P., Hellier, P., Morandi, X., Barillot, C.: 3D Rigid Registration of Intraoperative Ultrasound and Preoperative MR Brain Images Based on Hyperechogenic Structures. Journal of Biomedical Imaging 2012, 1 (2012)
3. De Nigris, D., Collins, D.L., Arbel, T.: Fast and Robust Registration Based on Gradient Orientations: Case Study Matching Intra-Operative Ultrasound to Pre-Operative MRI in Neurosurgery. In: Abolmaesumi, P., Joskowicz, L., Navab, N., Jannin, P. (eds.) IPCAI 2012. LNCS, vol. 7330, pp. 125–134. Springer, Heidelberg (2012)
4. Huang, X., Hill, N., Ren, J., Guiraudon, G., Boughner, D., Peters, T.: Dynamic 3D Ultrasound and MR Image Registration of the Beating Heart. In: Medical Image Computing and Computer-Assisted Intervention 2005, pp. 171–178 (2005)
5. King, A., Rhode, K., Ma, Y., Yao, C., Jansen, C., Razavi, R., Penney, G.: Registering Preprocedure Volumetric Images With Intraprocedure 3-D Ultrasound Using an Ultrasound Imaging Model. IEEE Trans. Med. Imag. 29(3), 924–937 (2010)
6. Kuklisova-Murgasova, M., Cifor, A., Napolitano, R., Papageorghiou, A., Quaghebeur, G., Noble, J.A., Schnabel, J.A.: Registration of 3D Fetal Brain US and MRI. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) MICCAI 2012, Part II. LNCS, vol. 7511, pp. 667–674. Springer, Heidelberg (2012)
7. Mercier, L., Del Maestro, R., Petrecca, K., Araujo, D., Haegelen, C., Collins, D.: Online Database of Clinical MR and Ultrasound Images of Brain Tumors. Medical Physics 39, 3253 (2012)
8. Moriarty, T., Kikinis, R., Jolesz, F., Black, P., Alexander, E.: Magnetic resonance imaging therapy. Intraoperative MR imaging. Neurosurg. Clin. N Am. 7, 323–331 (1996)
9. Penney, G., Blackall, J., Hamady, M., Sabharwal, T., Adam, A., Hawkes, D., et al.: Registration of Freehand 3D Ultrasound and Magnetic Resonance Liver Images. Medical Image Analysis 8(1), 81–91 (2004)
10. Powell, M.J.: The BOBYQA Algorithm for Bound Constrained Optimization without Derivatives. Cambridge Report NA2009/06, University of Cambridge (2009)
11. Rivaz, H., Collins, D.L.: Self-similarity Weighted Mutual Information: A New Non-rigid Image Registration Metric. In: Ayache, N., Delingette, H., Golland, P., Mori, K. (eds.) MICCAI 2012, Part III. LNCS, vol. 7512, pp. 91–98. Springer, Heidelberg (2012)
12. Roche, A., Pennec, X., Malandain, G., Ayache, N.: Rigid Registration of 3-D Ultrasound with MR Images: A New Approach Combining Intensity and Gradient Information. IEEE Transactions on Medical Imaging 20(10), 1038–1049 (2001)
13. Wein, W., Brunke, S., Khamene, A., Callstrom, M., Navab, N.: Automatic CT-Ultrasound Registration for Diagnostic Imaging and Image-Guided Intervention. Medical Image Analysis 12(5), 577 (2008)