Registration of bone structures in 3D ultrasound and CT data: Comparison of different optimization strategies ⋆

Susanne Winter [a] Bernhard Brendel [b] Christian Igel [a]

[a] *Institut für Neuroinformatik, Ruhr-Universität Bochum, 44780 Bochum, Germany*

[b] *Institut für Hochfrequenztechnik, Ruhr-Universität Bochum, 44780 Bochum, Germany*

Abstract

We developed a fast and robust algorithm to register intraoperative three dimensional ultrasound data of the spine with preoperative CT data. We compared different gradient based and evolutionary optimization strategies for solving the registration problem. The iRprop, a fast gradient based optimization algorithm, quickly and reliably led to higher registration rates than the two established methods BFGS and conjugate gradient descent (CG). The Covariance Matrix Adaptation evolution strategy (CMA) yielded the best results concerning registration rate and accuracy but at the cost of a slightly higher number of evaluations of the optimization criterion compared to CG and iRprop. The CMA was able to register patient data starting from a realistic misalignment in 98 % of the trials in about 15 seconds per registration.

Key words: image registration, spine, ultrasound, evolutionary optimization

1 Introduction

The interest in multimodal image registration for medical applications increased over the past years, especially concerning the registration of ultrasound images [1]. One main application of image registration is navigated surgery, where intraoperative imaging is used to register preoperative data with the

⋆ This work was an activity of the Kompetenzzentrum Medizintechnik Ruhr (KMR), Bochum. It was supported by the Bundesministerium für Bildung und Forschung (Az. 13N8079) and Siemens AG Medical Solutions, Erlangen.

Email address: `Susanne.Winter@neuroinformatik.rub.de` (Susanne Winter).

coordinate system of the operating room. After the registration the surgeon can orientate himself using the more precise preoperative data and planning information.

The usage of intraoperative CT or MR led to good registration results, but the systems are expensive and occupy much space in the operating room, data acquisition is time consuming, and CT increases the radiation exposure for the patient and the staff. In contrast, intraoperative ultrasound does not suffer from these drawbacks and the freehand data acquisition is fast and easy.

Ultrasound is very powerful in soft tissue imaging, and there exist many approaches which use volume registration algorithms to register. However, the increasing interest in intraoperative ultrasound registration of bone structures (see references in [2]) requires the development of new methods.

The disadvantage of bone imaging with ultrasound lies in the fact that only few parts of the bone surface can be visualized. Therefore, volume based registration algorithms cannot be used for bone registration. Existing approaches apply surface-surface registration methods after segmentation of bone surfaces in both ultrasound and CT images. However, the ultrasound surface segmentation is time consuming and not robust enough to be applied to real data, especially to data of the spine. The aim of our work was to develop a robust algorithm to register intraoperative three dimensional ultrasound data of the spine with preoperative CT data. Our approach [2] based on surface-volume registration overcomes the abovementioned disadvantages by avoiding the segmentation of ultrasound data. Here we concentrate on the optimization algorithm used in our registration system.

Image registration in general is an optimization problem, where parameters of a coordinate transformation are optimized to get the best mapping of one dataset to the other dataset. Various optimization strategies are used for different registration approaches [3], most of which are gradient based. There also exist some approaches with evolutionary algorithms (e.g., [4,5]) which lead to good registration rates. These methods often have the disadvantage that the results are not as exact as the results of gradient based methods. Another drawback of many optimization algorithms is the need to adjust the parameters of the optimization strategy itself manually.

In this paper we present the results of our registration algorithm in combination with different optimization strategies. Different optimization methods, where the parameters were set to default values, were tested with phantom and patient data. We show that the rate of correct registration and also the accuracy of the registration depends on the applied optimization method and come up with a rating of the optimization algorithms in our context.

2 Methods

We obtained 3D-ultrasound and CT data from a phantom of three synthetic vertebrae in a water bath. We also acquired 3D-ultrasound data of a patient

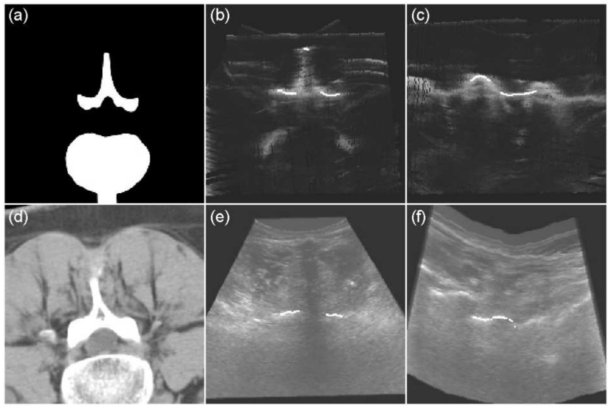

Fig. 1. Data slices; (a-c) phantom data, (d-f) patients data; (a) and (d) axial CT; (b) and (e) axial ultrasound slices; (c) and (f) sagittal ultrasound slices. The white contours in the ultrasound images show the defined optimal position of the bone surface points of one vertebra.

who had a spiral-CT for diagnostic purposes of the lumbar spine. In these datasets two vertebrae were imaged. All datasets were resampled to get a resolution of $0.5 \times 0.5 \times 0.5\,\mathrm{mm}^3$.

In the preprocessing of the CT data we extracted those bone surface points that can be visualized with ultrasound. The content of ultrasound images depends highly on the direction of the ultrasound propagation. For the ultrasound data acquisition of the lumbar spine we can presume that the transducer is moved from dorsal over the spinous processes. So the imaging plane of the ultrasound correlates with an axial CT slice and the direction of the sound propagation correlates to the posterior-anterior axis of the volume dataset.

All bone surface points visible in the ultrasound image were extracted from the CT data by thresholding (200 Hounsfield units for patient data). Additionally, we verified the angle between the surface and the sound propagation in order to include only parts of the surface which are nearly orthogonal to the sound propagation. The surface points of different neighboring vertebrae were separated manually.

We use a rigid surface-volume registration algorithm, where six transformation parameters, three for rotation and three for translation, have to be optimized. Since the tissue-bone interface is imaged as a bright area in the ultrasound data, the optimization criterion was chosen as the sum of the ultrasound gray values covered by the transformed surface points. The objective is to maximize this gray value sum.

We implemented different optimization methods: the Broyden-Fletcher-Goldfarb-Shanno algorithm (BFGS) and the method of conjugate gradient descent (CG), which are gradient based algorithms; the improved Resilient Backpropagation (iRprop^+) [6,7], which is also gradient based and was originally developed for the training of artificial neural networks, and the Covariance Matrix Adaptation (CMA) evolution strategy [8,9], an elaborated evolutionary algorithm for real-valued optimization. All these methods were used with standard

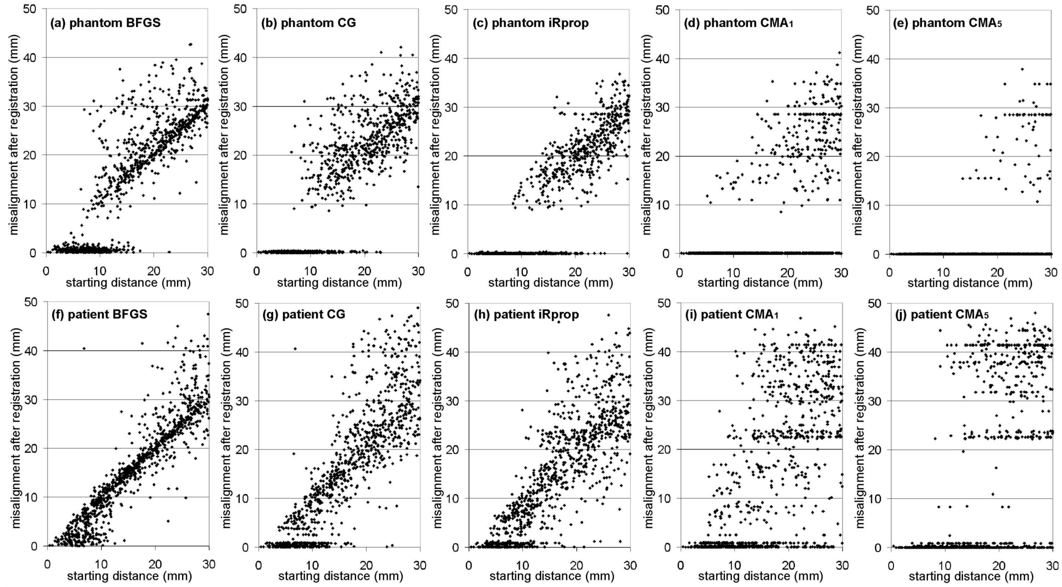

Fig. 2. Misalignment after 1000 registration trials with different optimization algorithms; (a-e) one phantom vertebra; (f-j) one patient vertebra.

parameters as described in the literature. To get a reference for the registration results we had to define the optimal position of the surface points in the ultrasound dataset. This was done by a manual preregistration, followed by multiple local optimizations and a visual control of the position with the best optimization criterion. Some slices of the data are shown in figure 1. In the ultrasound images the bone surface points of one vertebra are shown in their defined optimal position.

3 Results

Two sets of 1000 starting positions were generated for the evaluation. These positions are deviations from the defined optimal position. The direction and the length of the translation vector were drawn from uniform distributions. The same was done for the three rotation parameters. The first set of starting positions varied between 0 and 30 mm for the translation and 0 and 0.35 rad for the rotation. These are relative large deviations regarding the spinal structures. The second set varied between 0 and 15 mm for the translation and 0 and 0.2 rad for the rotation. This is a realistic misalignment for an intraoperative starting position. To quantify the misalignment of a registration and to measure the starting distance of a trial, we measured the average distance of the surface points to their optimal position. Because of the combination of the rotational and translational deviation we had starting positions for the vertebra between 0 and 32 mm for the first set and between 0 and 16 mm for the second. Results which have a distance of less than 1 mm after registration were regarded as correctly registered. The registration rates of the phantom vertebrae for the first set of starting positions were as follows, the numbers in

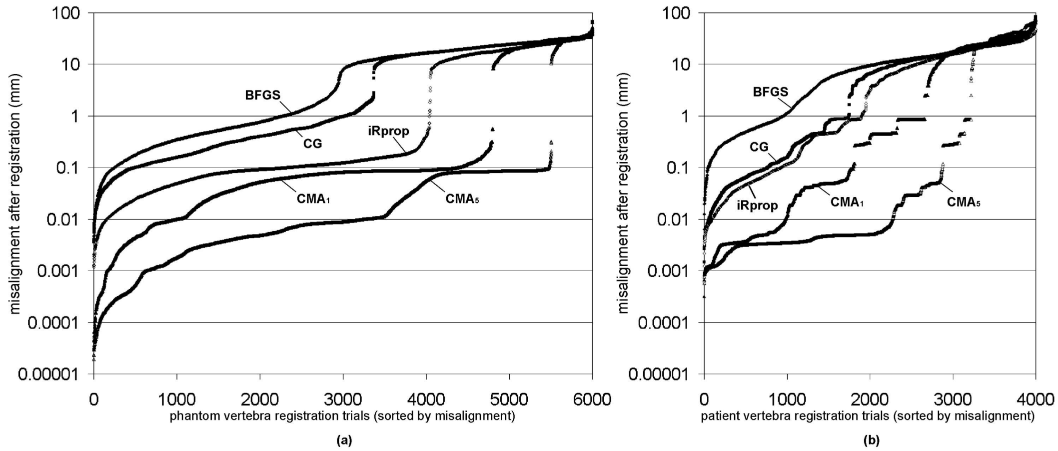

Fig. 3. Resulting misalignments for different optimization strategies; (a) all 6000 registration trials for three phantom vertebrae; (b) all 4000 registration trials for two patient vertebrae.

brackets are the results for the patient vertebrae: BFGS 25.9 % (16.2 %); CG 35.5 % (28.6 %); iRprop 48.8 (30.2 %); CMA 62.6 % (43.5 %). The results for the second set of starting positions were: BFGS 49.5 % (31.7 %); CG 65.6 % (58.7 %); iRprop 85.8 % (65.2 %); CMA 97.5 % (89.5 %). The results with CMA could be improved by considering five independent populations. Then the registration rates for the first set of starting positions were 83.5 % (62.7 %) and for the second set 100 % (98.4 %).

Figure 2 shows the misalignment after 1000 registration trials for each of the different optimization algorithms. Roughly speaking, for each starting distance each point in a plot corresponds to a local optimum. Dots not on the zero line indicate premature convergence of the optimization algorithm. Obviously, the gradient based methods – in particular BFGS and CG – are more prone to getting stuck in local optima than the CMA. In case of the phantom data, for all optimization strategies there was a clear distinction between results close to the optimum and trials with incorrect registrations. The distribution of the results obtained with BFGS was much larger compared to the other optimization strategies. In case of the patient data, such an obvious separation between correctly and incorrectly registered trials could not be found when using BFGS. The CG and the iRprop showed a very similar distribution of the results. The CMA evolution strategy stopped in fewer local optima, especially when using five independent populations (CMA_5). Distinct local maxima could be identified at different misalignments, for example in figure 2 in the phantom data at a misalignment of 28 mm or in the patient data at 23 mm. In figure 3 all resulting misalignments for the registration trials applied to the three phantom and the two patient vertebrae were sorted by the distance to the optimum. In case of the phantom data, iRprop and CMA rarely led to misalignments between 1 and 10 mm. In case of the patient data, this was only true for the CMA algorithm. Here, the average number of computations of the optimization criterion required to find a solution was 1689 for BFGS, 5506 for CG, 694 for iRprop, and 1486 for the CMA. On an AMD Athlon 1600+, 1000 computations correspond approximately to 2.2 seconds for the registration of one vertebra.

4 Conclusion

The iRprop is a very fast gradient based optimization algorithm which quickly and reliably led to higher registration rates than the two established methods BFGS and CG. The CMA evolution strategy gave the best results concerning registration rate and accuracy but at the cost of a slightly higher number of evaluations of the optimization criterion compared to CG and iRprop. We could even improve the results of CMA by using five independent populations. By doing so, it was possible to correctly register the phantom data for starting distances up to 16 mm in 100 %. Although there were more local maxima in the patient data, it was possible to reach registration results of 98 % for realistic starting positions. We conclude that state-of-the-art evolution strategies are particularly well suited for solving registration problems.

References

[1] J. Pluim, J. Fitzpatrick, Image registration, IEEE Trans Med Imaging 8 (22) (2003) 1341–1343.

[2] B. Brendel, S. Winter, A. Rick, M. Stockheim, K. Schmieder, H. Ermert, Registration of 3D CT- and ultrasound-datasets of the spine using bone structures, Comput Aided Surg 7 (2002) 146–155.

[3] J. Maintz, M. Viergever, A survey of medical image registration, Medical Image Analysis 2 (1) (1998) 1–36.

[4] J.-M. Rouet, J.-J. Jacq, C. Roux, Genetic algorithms for a robust 3-D MR-CT registration, IEEE Trans Inf Technol Biomed 4 (2) (2000) 126–136.

[5] G. Kagadis, K. Delibasis, G. Matsopoulos, N. Mouravliansky, P. Asvestas, G. Nikiforidis, A comparative study of surface- and volume-based techniques for the automatic registration between CT and SPECT brain images, Medical Physics 29 (2) (2002) 201–13.

[6] M. Riedmiller, Advanced supervised learning in multi-layer perceptrons - from backpropagation to adaptive learning algorithms, Computer Standards and Interfaces (1994) 265–278.

[7] C. Igel, M. Hüsken, Empirical evaluation of the improved Rprop learning algorithm, Neurocomputing 50 (C) (2003) 105–123.

[8] N. Hansen, A. Ostermeier, Completely derandomized self-adaptation in evolution strategies, Evol Comput 9 (2) (2001) 159–195.

[9] N. Hansen, S. Kern, Evaluating the CMA evolution strategy on multimodal test functions, in: X. Yao, et al. (Eds.), Parallel Problem Solving from Nature (PPSN VIII), Vol. 3242 of LNCS, Springer, 2004, pp. 282–291.