Deformable image registration with automated motion-mask extraction

Jef Vandemeulebroucke, Simon Rit, Joël Schaerer, and David Sarrut

Université de Lyon, CREATIS, CNRS UMR5220,
Léon Bérard cancer center. France.

Abstract. We participated in the empire challenge with two different implementations of deformable registration. Both use an automated extraction of a motion mask to account for sliding tissues. The first one is based on a B-spline parametrization and uses mutual information as the similarity measure. The second one is based on the Demons algorithm with an a priori lung density modification (APLDM).

1 Introduction

The empire challenge is a unique opportunity to compare in-house implementations of deformable registration to other developments in the field. We evaluated two different implementations used in previous publications of our group. These implementations include improvements proposed in these publications, mainly the B-LUT optimization [9], the motion mask extraction [12] and the a priori lung density modification (APLDM) [8].

2 Methods

First, we describe the initialization common to both registrations: the affine preregistration and the motion mask extraction. Second, we describe each deformable registration implementation.

2.1 Affine registration

The deformable registrations were initialized with the result of an automated affine transformation implemented using ITK [3]. For this registration, only the information in the lungs was retained, the rest being masked out using the provided lung masks with a constant -1200 HU density. Robustness was ensured by initializing the 3 translation parameters to align the first moment of the voxel intensities. The similarity measure was the mean of squared differences, computed in only 5% of the lungs voxels for efficiency (random selection). The optimization algorithm was the limited memory BFGS (Broyden-Fletcher-Goldfarb-Shanno) minimization with simple bounds [1], stopping when the norm of the projected gradient was below 10^{-10} mm.

2.2 Motion mask extraction

Sliding motion of the lungs with respect to the chest wall was accounted for using an automatically extracted motion mask [12]. The motion mask is a subanatomical segmentation which divides the thorax into moving (mainly the lungs, the mediastinum and the abdomen) and less-moving (mainly the outside of the patient body, the bony anatomy and surrounding tissues) regions [6]. The automated segmentation uses a level set framework to track the evolution of a moving interface, constrained by previously extracted anatomical features and regularized by a strong geometric prior. The anatomical features are the skeleton, the lungs (provided here) and the patient body. The extraction results in a smooth interface, encompassing the lungs and the mediastinum, enclosed in the rib cage and reaching the outside of the patient body in front of the abdomen. The same set of parameters was used for all reference and target images.

The motion mask was used following the procedure proposed by Wu et al. [13]. First, the intensities of the voxels in a 6 voxel border around the moving region were set to a constant -1200 HU value using morphological dilation. The rest of the voxels of the less-moving region were set to -1400 HU. The resulting pairs of images were registered using the two non-rigid registrations presented in the following. The voxels set to -1200 HU are meant to penalize in the similarity measure a potential mismatch between the borders of the moving region in the reference and the target image, but does not explicitly constrain the sliding motion. The voxels of the reference image set to -1400 HU are excluded from the similarity measure for performance enhancement. Note that the motion masks are required for both the reference and the target image. As a consequence, the segmentation should be consistent with respect to anatomical structures, as would be the case with lung masks used in the same way. The advantage of using the motion masks is that they encompass the mediastinum, thus avoiding potential inconsistencies that may be present in the lung masks near the mediastinum.

2.3 Method 1: B-splines Mutual Information with motion mask

We used an in-house implementation of non-rigid registration based on ITK [3] with multithreading, using B-splines based Free-Form Deformations (FFD) [7] and Mutual Information [4]. B-splines computation used a fast, low-memory B-LUT implementation [9], both for motion parametrization and image interpolation. A multi-resolution approach with three levels was applied to the images and the B-spline control points, the finest level having 2 mm and 32 mm isotropic spacing, respectively. Similarity was measured through Mattes implementation of the mutual information using 30 bins [5]. The optimal set of B-spline parameters was sought using the limited memory BFGS (Broyden-Fletcher-Goldfarb-Shanno) algorithm [1] starting from the affine transformation, stopping at each level when the norm of the projected gradient was below 10^{-5} mm or after 1000 iterations.

2.4 Method 2: Demons with APLDM and motion mask

We used an in-house implementation of the Demons algorithm [11] without multithreading, using the A Priori Lung Density Modification (APLDM) preprocessing described in [8]. The method minimizes the sum of squared differences (SSD) using a second-order gradient approximation. A Gaussian regularization of the vector fields was implemented using Deriche recursive filters [2]. The APLDM preprocessing step was used to account for lung density changes due to inspiration, to which the SSD is not robust, contrarily to the mutual information. It consists in modulating the lung densities of one image according to the densities of the other in order to make them comparable. A similar multi-resolution approach with two levels was applied, the finest level having 2 mm isotropic spacing for both the images and the vector fields. The algorithm was considered to have converged if the similarity measure did not improve during 500 iterations.

3 Results

The results were visually assessed using vv (`http://vv.creatis.insa-lyon.fr/`), which is an open-source viewer developped for the particular application of lung registration [10]. Performances were assessed on an Intel Xeon bi-quad-core 2.3GHz PC with 8 Gb of RAM.

The affine registration lead to saftisfying rough alignments without manual correction. Its computation time required 6 min on average.

The motion mask registration lead to visually satisfying results for all images but one (patient 6, fixed image). The latter is a contrast-enhanced image which prevented separation of the mediastinum from the bones with the simple thresholding used to segment the skeleton. As a consequence, the lung masks instead of the motion masks were used for the pair of images of patient 6. The whole segmentation process took 5 min on average.

The quantitative performance assessment of the registrations provided by the organizers are given in tables 1 and 2. On average, method 1 gave more accurate and smoother results than method 2. The average computation times were 62 min for method 1 and 34 min for method 2.

4 Discussion

We focus here on our 2 implementations among 34 competing ones. The full comparison is written by the organizers.

Both methods are fully automated and we used only one set of parameters for each. Those parameters were tuned in previous studies on 4D CT images of the thorax, and we did not try to adjust them for the purpose of this challenge. Given the variety of the images (humans and animals, deep inspiration breathholds and free-breathing, etc...), some had characteristics very different from conventional 4D CT, including higher resolution, comparatively higher dose levels per image, the absence of motion induced artifacts and considerably larger deformation. The

	Lung Boundaries		Fissures		Landmarks		Singularities	
Scan Pair	**Score**	**Rank**	**Score**	**Rank**	**Score**	**Rank**	**Score**	**Rank**
01	0.00	14.00	0.03	6.00	3.01	15.00	0.00	11.50
02	0.00	11.00	0.00	15.00	0.57	21.00	0.00	12.50
03	0.00	23.00	0.00	12.50	0.58	20.00	0.00	12.00
04	0.00	27.00	0.00	16.50	1.35	16.00	0.00	14.00
05	0.00	13.00	0.00	16.00	0.25	24.00	0.00	13.50
06	0.00	16.00	0.00	25.00	0.43	23.00	0.00	14.00
07	0.00	9.00	0.74	10.00	2.05	11.00	0.00	20.00
08	0.00	16.00	0.08	17.00	1.33	21.00	0.00	12.50
09	0.00	17.00	0.00	6.50	0.73	22.00	0.00	13.00
10	0.01	19.00	0.00	15.00	1.63	14.00	0.00	13.50
11	0.00	9.00	0.04	13.00	1.33	16.00	0.00	11.50
12	0.00	10.00	0.00	13.50	0.32	21.00	0.00	14.50
13	0.00	19.00	0.14	25.00	1.28	25.00	0.00	13.00
14	0.11	20.00	2.68	8.00	2.94	13.00	0.00	22.00
15	0.00	18.00	0.00	25.00	0.76	21.00	0.00	12.50
16	0.00	17.00	0.08	17.00	1.34	19.00	0.00	13.50
17	0.00	6.50	0.03	3.00	1.06	19.00	0.00	14.00
18	0.00	10.00	1.36	12.00	2.96	17.00	0.00	22.00
19	0.00	14.00	0.00	12.00	0.68	27.00	0.00	14.50
20	0.00	12.00	0.86	6.00	2.36	15.00	0.00	10.50
Avg	0.00	15.02	0.30	13.70	1.35	19.00	0.00	14.22
Average Ranking Overall								15.48
Final Placement								14

Table 1. Results of method 1 for each scan pair, per category and overall. Rankings and final placement are from a total of 34 competing algorithms.

	Lung Boundaries		Fissures		Landmarks		Singularities	
Scan Pair	**Score**	**Rank**	**Score**	**Rank**	**Score**	**Rank**	**Score**	**Rank**
01	0.14	23.00	0.01	5.00	2.57	14.00	0.02	28.00
02	0.00	25.00	0.00	15.00	0.79	30.00	0.00	26.00
03	0.00	26.00	0.00	12.50	0.69	25.00	0.00	29.00
04	0.02	31.00	0.00	16.50	3.49	29.00	0.02	30.00
05	0.00	26.00	0.00	16.00	0.62	28.00	0.00	13.50
06	0.00	16.00	0.00	1.00	0.79	31.00	0.00	31.00
07	0.31	26.00	0.68	9.00	3.38	18.00	0.18	28.00
08	0.04	24.00	0.00	3.50	0.98	14.00	0.00	28.00
09	0.00	28.00	0.00	6.50	0.77	24.00	0.02	30.00
10	0.05	29.00	0.05	32.00	4.86	27.00	0.35	30.00
11	0.26	27.00	0.07	15.00	2.06	22.00	0.13	29.00
12	0.07	30.00	0.00	13.50	0.50	24.00	0.00	14.50
13	0.00	24.00	0.07	6.00	1.02	16.00	0.00	27.00
14	0.28	23.00	6.64	21.00	7.95	21.00	1.56	30.00
15	0.00	24.00	0.00	7.00	0.83	24.00	0.00	28.00
16	0.01	32.00	0.00	6.50	1.19	16.00	0.03	29.00
17	0.00	21.00	0.06	28.00	1.00	17.00	0.00	14.00
18	0.45	27.00	1.65	14.00	2.93	15.00	0.47	30.00
19	0.00	29.00	0.00	12.00	0.64	24.00	0.00	14.50
20	0.27	28.00	2.82	16.00	1.75	11.00	0.04	27.00
Avg	0.09	25.95	0.60	12.80	1.94	21.50	0.14	25.82
Average Ranking Overall								21.51
Final Placement								27

Table 2. Results of method 2 for each scan pair, per category and overall. Rankings and final placement are from a total of 34 competing algorithms.

challenge results show both the relative robustness of the chosen parameters and room for improvement.

The results of the registrations were as expected for method 1. We achieved average accuracy on boundaries and fissures. The motion field is overall very smooth and shows little folding. The B-spline implementation did not include a regularization other than the intrinsic B-spline smoothness. This is satisfactory as long as lungs contain high-contrast structures homogeneously distributed in space, and the B-spline control points are sufficiently spaced (32 mm here). However, further refinement of the registration would require a regularization and we intend to add it.

The results obtained for method 2 were surprisingly worse. In particular the matching of the lung borders performs poorly, especially when considering the good performance on the fissures. The way the motion masks were used in the registration process could be responsible for this, and should be investigated further. The method also obtained poor accuracies for landmarks. One reason could be that the Demons are not meant for large deformations such as the ones observed in some pairs of images, e.g. patient 1. The APLDM enhances the conventional Demons for this difficulty but does not fully tackle it. Note that more similar results were observed for the two methods on small deformations. The Demons showed more singularities than most algorithms (average rank: 26). This is a known downside of this method, since there were no specific mechanisms preventing folding as in the diffeomorphic version. The result indicates however that regularization (Gaussian smoothing of the motion field using $\sigma = 2\ mm$ at the finest level) might have to be increased in the future.

A common remark to both registration methods concerns the image resolution of 2 mm at which the registrations were performed. Improved accuracy, particularly for landmarks and fissures are expected when using the full resolution images. This refinement was found very costly for some image pairs, and therefore not performed in this study. We currently only have developed CPU implementations of our registration software. The advantage is that they can be run on any PC. Of course, it comes at a cost compared to GPU implementations: although we improved the performances with various strategies (B-LUT, multithreading, etc...), the registrations took about an hour for method 1 at 2 mm^3 resolution. This resolution is probably sufficient for radiotherapy applications since we typically use 3 mm slice spacing for treatment planning. Future work includes performance improvement to perform registrations at full resolution for fine resolutions such as some of the challenge images.

5 Conclusion

This challenge allowed us to evaluate two implementations of deformable registration. It increased our confidence in those implementations but also showed potential areas for improvement.

References

1. Richard H. Byrd, Peihuang Lu, Jorge Nocedal, and Ciyou Zhu. A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM Journal on Scientific Computing*, 16(5):1190, 1995.
2. R. Deriche. Recursively implementing the gaussian and its derivatives. Technical report, INRIA, 1993.
3. Luis Ibanez, W.J. Schroeder, Lydia Ng, Josh Cates, and InsightConsortium. The ITK Software Guide, August 2005.
4. Frederik Maes, André Collignon, Dirk Vandermeulen, Guy Marchal, and Paul Suetens. Multimodality image registration by maximization of mutual information. *IEEE transactions on Medical Imaging*, 16(2):187–198, 1997.
5. David Mattes, David R. Haynor, Hubert Vesselle, Thomas K. Lewellen, and William Eubank. PET-CT image registration in the chest using free-form deformations. *IEEE Transactions on Medical Imaging*, 22(1):120–128, 2003.
6. Eike Rietzel and George T. Y. Chen. Deformable registration of 4D computed tomography data. *Medical Physics*, 33(11):4423, 2006.
7. D Rueckert, LI Sonoda, C Hayes, DLG Hill, and MO Leach. Nonrigid registration using free-form deformations: application to breast MR images. *IEEE Transactions on Medical Imaging*, 18(8):712–721, 1999.
8. David Sarrut, Vlad Boldea, Serge Miguet, and Chantal Ginestet. Simulation of four-dimensional CT images from deformable registration between inhale and exhale breath-hold CT scans. *Medical Physics*, 33(3):605, 2006.
9. David Sarrut and Jef Vandemeulebroucke. B-LUT: Fast and low memory B-spline image interpolation. *Computer Methods and Programs in Biomedicine*, 99:172–178, December 2010.
10. Pierre Seroul and David Sarrut. VV: a viewer for the evaluation of 4D image registration. *MIDAS Journal*, (Medical Image Computing and Computer-Assisted Intervention MICCAI'2008, Workshop - Systems and Architectures for Computer Assisted Interventions):1–8, 2008.
11. J.P. Thirion. Image matching as a diffusion process: an analogy with Maxwell's demons. *Medical Image Analysis*, 2(3):243–260, 1998.
12. Jef Vandemeulebroucke, Olivier Bernard, Jan Kybic, Patrick Clarysse, and David Sarrut. Automatic motion mask extraction for deformable registration of the lungs. In *XVIth International Conference on the Use of Computers in Radiation Therapy*, Amsterdam, June 2010.
13. Ziji Wu, Eike Rietzel, Vlad Boldea, David Sarrut, and Gregory C. Sharp. Evaluation of deformable registration of patient lung 4DCT with subanatomical region segmentations. *Medical physics*, 35:775, 2008.