Pulmonary Image Registration with elastix using a Standard Intensity-Based Algorithm 2

Marius Staring[1]*, Stefan Klein[2], Johan H.C. Reiber[1], Wiro J. Niessen[2], and Berend C. Stoel[1]

[1] Division of Image Processing (LKEB), Department of Radiology, Leiden University Medical Center, Leiden, The Netherlands
[2] Biomedical Imaging Group Rotterdam, Depts. of Radiology & Medical Informatics, Erasmus MC, Rotterdam, The Netherlands

Abstract. This document describes a modification of the algorithm used for the EMPIRE10 challenge. That algorithm obtained a 7th and 3rd position before and during the conference, but was weak in the folding category. The newly submitted modified algorithm tries to remedy this by incorporating a regularisation term in the final step of the registration.
The original algorithm, implemented in `elastix`, optimises the normalised correlation criterion, using a fast, parameter-free and robust stochastic optimisation procedure. A combination of an affine and two nonrigid B-spline transformations models the spatial relationship. The approach is embedded in a multi-resolution framework for both the image data and the transformation. No explicit regularisation was used in the original algorithm. The revised algorithm is almost identical to the original algorithm, is still fully automatic, implemented in `elastix`, but a regularisation term was added to the two nonrigid stages of the registration, thereby requiring a user-defined parameter setting.
Of the 24 submitted algorithms (2010-11-26), our revised contribution achieved the 5-th place with an average rank of 9.06 (best 6.63, worst 22.17). This compares to the 8-th place (average rank 10.19) of the unregularized original algorithm.

Key words: pulmonary image registration, evaluation, `elastix`

1 Introduction

The authors registered to the challenge under the team name "RubberBand", referring to the registration package `elastix`, developed previously by the authors [1]. Executables and source code of `elastix` are publicly available from the website `http://elastix.isi.uu.nl`, under the BSD license, which allows free academic and commercial use and permits modification of the source code. A manual for `elastix` and an example of usage can also be downloaded. In

* Corresponding author: `m.staring@lumc.nl`

addition, we created a "parameter file database", which is a collection of parameter files that proved successful, together with a short description of the clinical application for which they were used. The parameter file database can be found through the website[3] and `elastix`-users are encouraged to upload their own settings. A default parameter file can also be found here.

The goal of our contribution is to determine what a standard and generic, but fully automatic, intensity-based image registration algorithm can achieve compared to the competition. How will this relatively simple, general purpose algorithm compare to more advanced registration algorithms that are tailored to the specific application?

2 Methods

Below the precise registration framework is given. The new algorithm is a small modification of the original one, given in red.

In this paper we adopt the formulation of image registration as an optimisation problem:

$$\widehat{\boldsymbol{\mu}} = \arg\min_{\boldsymbol{\mu}} \mathcal{C}(\boldsymbol{T}_{\boldsymbol{\mu}}; I_F, I_M) = \arg\min_{\boldsymbol{\mu}} \mathcal{S}(\boldsymbol{T}_{\boldsymbol{\mu}}; I_F, I_M) + \alpha \mathcal{R}(\boldsymbol{T}_{\boldsymbol{\mu}}), \tag{1}$$

where I_F and I_M denote the fixed and moving image, respectively, $\boldsymbol{T}_{\boldsymbol{\mu}}$ is the spatial transformation relating the two and parameterised by the vector of parameters $\boldsymbol{\mu}$, $\mathcal{C}$ is the cost function that defines the quality of alignment, which is separated into a similarity measure $\mathcal{S}$ and a regularisation term $\mathcal{R}$. The parameter α weighs regularity against similarity.

Image registration is performed in three stages:

1. Affine registration using the original data, without the use of lung masks. This is done to get a coarse global alignment of the entire anatomy. Lung masks are not used, to exploit all anatomy.
2. Nonrigid registration using the processed data (see "Masking" below for an explanation), without the use of lung masks. Our experiments revealed that the use of lung masks at an early stage had a negative impact on lung boundary alignment in case of large offsets (i.e. large differences in inspiration level).
3. Nonrigid registration using the processed data, and with the use of the lung mask of the fixed image. From our experiments we learned that the match of smaller structures within the lung is improved by using a lung mask.

Several choices for the different registration components are made:

Cost function $\mathcal{C}$: Normalised Correlation Coefficient (NCC), which is suitable for mono-modal image registration, but can compensate for global intensity differences due to differences in inspiration level.

[3] `http://elastix.isi.uu.nl/wiki.php`

It is possible to include regularisation in the cost function: $\mathcal{C} = \text{NCC} + \alpha\mathcal{R}$, with a suitable choice for $\mathcal{R}$, for example the bending energy, penalising the second order derivatives [2]. For the specific application at hand, regularisation would have been beneficial for some data sets (appearance of smearing effects). However, it requires manual setting of an additional data-dependent parameter α, which is not a trivial choice, and additionally it increases the computation time. Therefore, for the sake of simplicity regularisation was omitted during the MICCAI online challenge and at the workshop, at the cost of a deduction of points in the evaluation (singularities in the deformation field). The new algorithm, however, does incorporate regularisation in the two nonrigid stages of the registration scheme. A bending energy term penalising second order derivatives [2] was chosen. A weighting factor α_k was chosen in every iteration k as:

$$\alpha_k = \alpha \left\| \frac{\partial \mathcal{S}}{\partial \boldsymbol{\mu}_k} \right\| / \left\| \frac{\partial \mathcal{R}}{\partial \boldsymbol{\mu}_k} \right\|, \tag{2}$$

where α was experimentally set to the values reported in Table 1.

Transformation: An affine registration is performed prior to nonrigid registration to accommodate for global offset and differences in inspiration level. Subsequent nonrigid transformations are modelled by B-splines [2], embedded in a multi-grid setting.

Optimisation: To optimise (1) we opt for an iterative procedure, called adaptive stochastic gradient descent (ASGD) [3].

Sampling strategy: A relatively small number of samples (2000) are drawn randomly each iteration from the fixed image (off the voxel grid), to compute an approximate gradient of the cost function.

Interpolation: During registration a linear interpolator is used to compute the spatial derivative of the moving image.

Hierarchical strategy: For the image data Gaussian pyramids are used with sub-sampling, to increase robustness. For the B-spline transform a multi-grid approach is used, starting with a coarse control point grid in the first resolution, only capable of modelling coarse deformations. In subsequent

Table 1. Parameter settings for stages 1-3, resolution levels R1-R5.

Stage	Iterations		R1	R2	R3	R4	R5
1. Affine	1000						
2. Nonrigid without mask	1000	Grid spacing (mm)	80	80	40	20	10
		Downsample factor	16	8	4	2	1
		α			0.05		
3. Nonrigid with mask	2000	Grid spacing (mm)	80	40	20	10	5
		Downsample factor	4	3	2	1	1
		α			0.05		

resolutions the B-spline grid is gradually refined, thereby introducing the capability to match smaller structures.

Masking: In the last step of the registration procedure we have used lung masks, which were created automatically by the EMPIRE10 organisers.

The above image registration algorithm is fully automatic. Exact `elastix` settings that were used in the experiments have been made available via the parameter file database, see `http://elastix.isi.uu.nl/wiki.php`, see `par0011`. The registration settings for each experiment can be inspected in detail, and the parameter files can be downloaded for reproducing our results or for use in other applications.

3 Experiments and Results

Data sets and the scoring methodology are described in [4].

3.1 Runtime

All registration were performed on an Intel Xeon E5620 @ 2.40GHz, 24GB RAM, Ubuntu Linux 64 bit. The mean run time of `elastix` for each stage is given in Table 2. On average the registration took 106 minutes, of which 7 minutes were spent to automatically compute the optimisation parameters by the ASGD optimiser. The computation time in stage 3 is longer than that of stage 2 due to a doubling in the number of iterations and a more involved computation of $\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x})$, since at stage 3 $\boldsymbol{T}_{\boldsymbol{\mu}}(\boldsymbol{x})$ is a composition of three transforms.

3.2 Results

Visual inspection of the affine registration showed that all scans were successfully matched globally. Automatic scoring was performed on the final result after stage 3, by the EMPIRE10 organisers. The results are divided in four categories: lung boundary match, fissure match, landmark precision, and the presence of singularities in the deformation field. A comparison to other participants can be found at `http://empire10.isi.uu.nl/mainResults.php`. At the date of submission

Table 2. Average runtime in minutes for each stage of the registration.

	registration			ASGD		
stage	mean	min	max	mean	min	max
1. Affine	1	0	1	0	0	0
2. Nonrigid without mask	24	23	25	3	2	4
3. Nonrigid with mask	81	76	98	4	2	12
total	106			7		

(2010-11-26) out of the 24 submitted algorithms, the revised algorithm achieved the 5-th place with an average rank of 9.06 (best 6.63, worst 22.17).

Compared to the original submission, somewhat better scores were obtained in the fissure and landmarks categories, and notably in the folding category. Lung boundary alignment decreases somewhat, probably due to over-smoothing the deformation field in some cases.

4 Discussion and Conclusion

In conclusion, a standard, fully automatic, intensity-based image registration algorithm achieved a ranking of 5 out of 24. The revised algorithm improved on the original contribution. The implementation is publicly available from `http://elastix.isi.uu.nl`.

References

1. Klein, S., Staring, M., Murphy, K., Viergever, M., Pluim, J.: elastix: a toolbox for intensity-based medical image registration. IEEE Transactions on Medical Imaging **29**(1) (January 2010) 196 – 205
2. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L.G., Leach, M.O., Hawkes, D.J.: Non-rigid registration using free-form deformations: Application to breast MR images. IEEE Transactions on Medical Imaging **18**(8) (1999) 712 – 721
3. Klein, S., Pluim, J.P.W., Staring, M., Viergever, M.A.: Adaptive stochastic gradient descent optimisation for image registration. International Journal of Computer Vision **81**(3) (2009) 227–239
4. Murphy, K., van Ginneken, B., Reinhardt, J., Kabus, S., Ding, K., Deng, X., Pluim, J.: Evaluation of methods for pulmonary image registration: The EMPIRE10 study. (2010)