Variational Lung Registration With Explicit Boundary Alignment

Jan Rühaak and Stefan Heldmann

Fraunhofer MEVIS, Project Group Image Registration
Maria-Goeppert-Str. 1a, 23562 Luebeck, Germany
{jan.ruehaak,stefan.heldmann}@mevis.fraunhofer.de
http://www.mevis.fraunhofer.de

Abstract. The EMPIRE10 competition [4] provides a public platform for in-depth evaluation and fair comparison of algorithms for CT lung registration. The workshop was organized as part of the MICCAI 2010 conference in Beijing, China, but the challenge continues.
We present a variational lung registration approach that is augmented by a term which drives the registration towards exact lung boundary alignment. The registration is guaranteed to be free of singularities.

Keywords: EMPIRE10, registration, variational model, evaluation, lung boundary alignment

1 Introduction

The registration of lung images has many applications in medical image processing. Prominent examples range from motion compensation in radiotherapy for better dose control to tumor monitoring in oncology. Here, accurate registration may help to robustly detect a tumor in a follow-up scan[6]. Generally speaking, lung registration always constitutes a promising approach when breathing motion is to be compensated for. The large non-linear deformations, however, make an exact registration of the lungs a challenging problem requiring advanced registration algorithms.

The EMPIRE10 challenge provides a public platform for a fair and meaningful comparison of registration algorithms. The quality of a registration is evaluated using four criteria:

1. Lung Boundary Alignment
2. Fissure Alignment
3. Landmark Distance
4. Singularities

In this note, we present a variational approach for the lung registration problem [9]. A standard non-linear registration is extended by two terms: While the first one drives the registration towards a good alignment of the lung boundaries, the second one guarantees the deformation field to be free of singularities.

2 Methods

We first briefly describe the variational registration setting. Let $\mathcal{R} : \mathbb{R}^3 \to \mathbb{R}$ denote the reference (fixed) image and $\mathcal{T} : \mathbb{R}^3 \to \mathbb{R}$ the template (i.e. moving) image with support in domains $\Omega_\mathcal{R} \subseteq \mathbb{R}^3$ and $\Omega_\mathcal{T} \subseteq \mathbb{R}^3$, respectively. Image registration tries to find a transformation $y : \Omega_\mathcal{R} \to \mathbb{R}^3$ which minimizes a suitable objective function $\mathcal{J}$ often called *joint energy functional*.

Traditionally, the objective function consists of a *distance term* $\mathcal{D}$ and a *regularization term* or regularizer $\mathcal{S}$. The distance term quantifies the similarity of images whereas the regularizer assesses the smoothness or regularity of the deformation. Hence, the joint energy functional reads

$$\mathcal{J}(y) = \mathcal{D}(\mathcal{R}, \mathcal{T}(y)) + \alpha \mathcal{S}(y) \tag{1}$$

where $\alpha > 0$ is the so-called regularization parameter controlling the balance between data fit (image alignment) and regularity of the deformation.

We use the *normalized gradient field (NGF)* distance measure

$$\mathcal{D}(\mathcal{R}, \mathcal{T}(y)) := \int_{\Omega_\mathcal{R}} 1 - \left(\frac{\langle \nabla \mathcal{T}(y(x)), \nabla \mathcal{R}(x) \rangle}{\|\nabla \mathcal{T}(y(x))\| \, \|\nabla \mathcal{R}(x)\|} \right)^2 \, \mathrm{d}x. \tag{2}$$

NGF focusses on image edges instead of intensities which make it an attractive choice especially for inhale-exhale data sets with their breathing-related intensity changes. Smaller edges caused by noise are suppressed by an edge parameter η (see e.g. [5]). We set $\eta = 100$ throughout the registration. From our experience, however, the registration result is rather insensitive to changes of this parameter, even by orders of magnitude.

The NGF distance measure is accompanied by the *curvature* regularizer

$$\mathcal{S}(y) := \frac{1}{2} \int_{\Omega_\mathcal{R}} \sum_{j=1}^{3} \|\Delta(y_j - y_j^{\mathrm{kern}})\|^2 \, \mathrm{d}x \tag{3}$$

with the kernel y^{kern} set to the result of an affine-linear preregistration in our framework. y_j denotes the j-th component function of y, y_j^{kern} analogously. The curvature regularizer penalizes second order derivatives and yields very smooth solutions.

2.1 Incorporating Lung Boundary Information

The EMPIRE10 data contains cross-checked lung segmentations for both reference and template scan which may be used by the contestants. Formally, these masks define corresponding sets $B_\mathcal{R} \subseteq \Omega_\mathcal{R}, B_\mathcal{T} \subseteq \Omega_\mathcal{T}$ and hence binary label functions

$$b_\mathcal{R} : \Omega_\mathcal{R} \to \{0, 1\}, \quad b_\mathcal{R}(x) = 1 \;\Leftrightarrow\; x \in B_\mathcal{R}$$

for $\mathcal{R}$ and $b_\mathcal{T}$ for $\mathcal{T}$, respectively. The registration should respect this area information, i.e. in the ideal case it holds that

$$b_\mathcal{R}(x) = b_\mathcal{T}\big(y(x)\big) \quad \forall x \in \Omega_\mathcal{R}.$$

We incorporate the segmentation information as an additional term in the objective function. For this purpose, let

$$\mathcal{B}(y) := \frac{1}{2} \int_{\Omega_{\mathcal{R}}} \Big(b_{\mathcal{T}}\big(y(x)\big) - b_{\mathcal{R}}(x) \Big)^2 \, dx.$$

The function $\mathcal{B}$ assigns a value to each deformation y which measures the agreement between the segmentation in the reference scan and the deformed template image segmentation. This leads to a new objective function

$$\mathcal{J}(y) = \mathcal{D}(\mathcal{R}, \mathcal{T}(y)) + \alpha \mathcal{S}(y) + \beta \mathcal{B}(y)$$

which is to be minimized. Note that the term $\mathcal{B}$ coincides with the sum of squared differences of the segmentation masks as binary images.

2.2 Singularities

The classical variational approach does not vouch for injectivity of the computed transformation as optimizer of the energy functional. Depending on the data to be registered and the choice of the regularization parameter α, a minimizer of (1) may exhibit severe singularities such as foldings or tearings in the corresponding deformation field and must therefore be considered physically impossible. Hence, it should be ruled out by a robust image registration method. The fourth EMPIRE10 criterion is designed to detect such implausible deformations.

From basic calculus, it is known that a continuously differentiable function y is invertible in a neighborhood of a point $p \in \mathbb{R}^d$ if and only if the Jacobian $\det \nabla y(p)$ is non-zero. Moreover, the function preserves orientation of local coordinates if the Jacobian is positive. The Jacobian also gives information how local volumes in the neighborhood of p change under transformation. A value of the Jacobian greater than 1 indicates expansion whereas values $|\det \nabla y(p)| < 1$ indicate shrinkage; the absolute value equals the factor by which the local volume is expanded or shrinked.

To this end, we introduce an additional energy similar to the approaches presented in [1, 8] that measures change of volume,

$$\mathcal{V}(y) := \int_{\Omega_{\mathcal{R}}} \psi(\det \nabla y(x)) \, dx$$

with weighting function

$$\psi(t) := \frac{(t-1)^2}{t} \quad \text{for } t > 0 \quad \text{and} \quad \psi(t) := \infty \text{ else},$$

cf. Fig. 1. The weighting function penalizes the deviation of the Jacobian from 1 and therefore measures the local change of volume. Note that ψ symmetrically penalizes volume expansion and shrinkage since $\psi(t) = \psi(1/t)$. Furthermore, ψ ensures injectivity of the deformation since $\psi(\det \nabla y) \to \infty$ as $\det \nabla y \to 0$. As a consequence $\mathcal{V}(y) = \infty$ if the Jacobian becomes negative at any point. Our

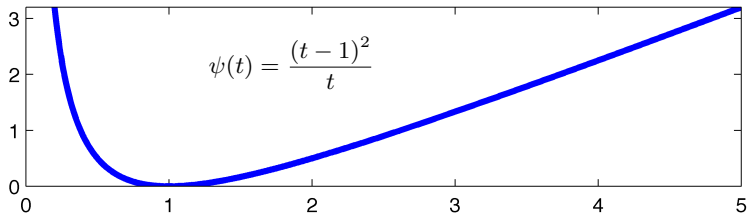

Fig. 1. Change of volume weighting function ψ

implementation is based on the idea of decomposing each voxel into tetrahedrons and then directly estimating the Jacobian of the deformation as the ratio of the volumes of a tetrahedron before and after deformation, cf. [1–3].

Summarizing, we arrive at the objective function

$$\mathcal{J}(y) = \mathcal{D}(\mathcal{R}, \mathcal{T}(y)) + \alpha\mathcal{S}(y) + \beta\mathcal{B}(y) + \gamma\mathcal{V}(y). \tag{4}$$

3 Application to Lung Registration

Let $\mathcal{R}, \mathcal{T}$ denote the reference and template scan with corresponding lung masks $B_{\mathcal{R}}$ and $B_{\mathcal{T}}$. Our lung registration approach consists of two main phases: In the first step, a rough alignment of the segmentation masks is achieved by a parametric registration. The resulting deformation is then used as starting value for the second phase in which the functional (4) is minimized.

3.1 Pre-Registration

In the pre-registration step, first the centers of gravity of reference and template mask are computed and aligned. In most cases, this step is redundant, but sometimes the lung masks do not overlap at all in world coordinates which leaves the optimization algorithm trapped in the local minimum of zero deformation.

Subsequently, the lung masks are registered using a standard 3D affine-linear parametric transformation model (see e.g. [5]) with the SSD distance measure. For this purpose, the masks are downsampled by a factor of five in each dimension preceded by a Gaussian smoothing. The downsampled masks are again smoothed in order to spread the boundary information over a larger area.

The optimization is carried out by a Gauss-Newton algorithm. The whole pre-registration phase takes less than ten seconds on a Mac Pro 2x2.4 GHz Quad-Core with 32 GB RAM. Our C++ implementation is not heavily optimized for performance, yet. However, some parts as interpolation and volume computations have been parallelized on the CPU using OpenMP.

3.2 Non-linear Registration

Having roughly aligned the mask shape, the main non-linear registration phase begins. We use a multi-level strategy ranging from course to fine on the CT scans which are masked using the provided segmentations. The multi-level pyramid is generated using Gaussian smoothing and subsequent downsampling by a factor of two. We use four levels for the registration and stop at the second finest level for computational reasons. As discussed in section 2, we use normalized gradient fields (2) as distance measure together with the curvature regularizer (3).

Additionally, the deformation grid is also embedded into the multi-level setting. We do not assign a grid point to every voxel, but employ a grid of $64 \times 64 \times 64$ cells on the finest level used in the registration. In between, the deformation is linearly interpolated which can be regarded as a small additional regularization. For each coarser level, the deformation grid size is reduced by a factor of two in each dimension.

We solve the optimization of (4) by a *first-discretize-then-optimize* approach. This means, the terms of the objective function are discretized first, yielding a finite dimensional optimization problem. The minimization is performed by the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) quasi-Newton optimization method [7]. The runtime of our algorithm ranges from 1:17 to 5:56 minutes on our system depending on the resolution of the CT scans (see table 1 for details). For all registrations, $\alpha = 5, \beta = 1$ and $\gamma = 1$ were used as parameters.

4 Results and Discussion

We have presented an extended variational approach for CT lung registration that allows for the incorporation of lung segmentation information. It guarantees invertibility of the computed deformation.

References

1. M. Burger, J. Modersitzki, and L. Ruthotto, "A hyperelastic regularization energy for image registration," *SIAM Journal on Scientific Computing (submitted)*, 2011.
2. E. Haber and J. Modersitzki, "Image registration with guaranteed displacement regularity," *International Journal of Computer Vision*, vol. V71, no. 3, pp. 361–372, 2007.
3. E. Haber and J. Modersitzki, "Numerical methods for volume preserving image registration," *Inverse Problems*, vol. 20, no. 5, pp. 1621–1638, 2004.
4. K. Murphy, B. van Ginneken, J. Reinhardt, S. Kabus, K. Ding, X. Deng, and J.P.W. Pluim, "Evaluations of Methods for Pulmonary Image Registration: The EMPIRE10 Study", Medical Image Analysis for the Clinic: A Grand Challenge, Workshop Proceedings MICCAI 2010, `http://empire10.isi.uu.nl`.
5. J. Modersitzki, "FAIR: Flexible Algorithms for Image Registration," SIAM, Philadelphia, 2009.

Scan	Runtime (min:s)
1	3:22
2	4:38
3	1:58
4	4:17
5	2:01
6	1:32
7	5:38
8	3:19
9	2:35
10	3:04
11	4:20
12	3:12
13	1:27
14	5:56
15	3:30
16	1:51
17	1:30
18	4:25
19	3:31
20	4:42
21	4:45
22	3:54
23	1:17
24	2:39
25	2:54
26	1:56
27	3:04
28	3:06
29	3:31
30	1:21

Table 1. Running time of the non-linear registrations

6. J.H. Moltz, M. Schwier, and H.-O. Peitgen, "A General Framework for Automatic Detection of Matching Lesions in Follow-Up CT", Proceedings of the 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro, Boston, MA, USA, June 28 - July 1, 2009, pp. 843-846.
7. J. Nocedal and S.J. Wright, "Numerical optimization", Springer, 2006.
8. T. Rohlfing, C. R. Maurer, D. A. Bluemke, and M. A. Jacobs, "Volume-preserving nonrigid registration of mr breast images using free-form deformation with an incompressibility constraint," *IEEE Transactions on Medical Imaging*, vol. 22, pp. 730–741, 2003.
9. J. Rühaak, S. Heldmann, B. Fischer: "Improving Lung Registration by Incorporating Anatomical Knowledge: A Variational Approach", Fourth International Workshop on Pulmonary Image Analysis, MICCAI 2011, Toronto, Canada.