Improving Lung Registration by Incorporating Anatomical Knowledge: A Variational Approach

Jan Rühaak, Stefan Heldmann, and Bernd Fischer

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

Abstract. In this work, we present a novel approach for the registration of CT lung images. Therefore, we incorporate additional segmentation information yielding a significant improvement of accuracy, robustness, and reliability. The main idea of our approach is rather general and not limited to the case of lung registration. We describe a generic method for incorporating segmentation information into a variational image registration framework. Assuming that segmentation masks are available for reference and template image (e.g. masks separating lung tissue from the background), the method drives the registration process towards exact alignment of the masks. Furthermore, we extend the classical variational setting by an additional term that controls change of volumes and in particular guarantees non-singular deformation fields. Both extensions can be combined with arbitrary distance measures and regularizers, and therefore can be adapted to arbitrary registration tasks.

Keywords: registration, variational model, evaluation, lung, lobes, segmentation information

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[8]. Generally speaking, lung registration always constitutes a promising approach when breathing motion is to be compensated for. Due to the large non-linear deformations occurring in the breathing cycle, however, an exact registration of the lungs is a challenging problem and has attracted much attention in research [4].

Image registration in general, though central in many (medical) image processing applications, is inherently an ill-posed problem in the sense of Hadamard [7]. In particular, many different transformations may lead to exactly the same deformed image while exhibiting varying geometric properties – one may be injective, another one may even tear the image grid apart. Without further knowledge about the images to be registered and the process of their generation,

it is often difficult if not impossible to decide which deformation best solves a registration task.

Incorporating additionally available information therefore seems a prudent means towards a better registration. One possible way is to include landmark information, i.e. outstanding and reliably detectable points in both images, into an intensity-based registration process. If those points are known to correspond to each other beforehand, they can successfully be used to drive the registration towards a more desirable result [11, 5, 10, 13].

Unfortunately, in many situations only very few (if any) reliable landmarks are available. Very often, however, some areas are already known to correspond to each other, but the point-to-point relationship between the areas which is needed for landmark techniques is not. A typical example is the registration of a CT scan to an MR image of the same body region. It is rather trivial that e.g. the heart region in the CT must be mapped to the heart in the MR and not to any other structure, but the pointwise correspondence will usually not be known a priori. In our experience, in fact, this *area correspondence* situation is typical for many image registration problems.

In this paper, we want to show how such information can be used to help a registration algorithm. We present a method that allows for the incorporation of arbitrarily many area correspondences into a variational image registration framework. It can be combined with any distance measure and regularizer and does not require changes in the optimization algorithm. The benefit of the method is exemplified with the registration of thoracic lung CT scans using segmentation masks of the lungs for both reference and template image as additional information.

Additionally, we describe a technique that guarantees the computed transformation to be free of singularities such as foldings of the underlying image grid. Again, the method can be integrated into the variational framework in a straightforward manner and it can be combined with arbitrary distance measures and regularizers. It thus provides a very convenient extension to the variational image registration approach and ensures that the computed transformations are always physically plausible, independent of the choice of regularization parameters.

2 Methods

We start off with a brief summary of the classic variational image 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) on the one hand and regularity of the deformation on the other hand. A standard distance measure for mono-modal registration problems is the sum of squared differences (SSD)

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

whereas the *diffusive* regularizer

$$\mathcal{S}(y) := \frac{1}{2} \int_{\Omega_\mathcal{R}} \sum_{j=1}^{3} \|\nabla u_j(x)\|^2 \, dx \tag{3}$$

with the decomposition $y(x) = x + u(x)$ forms a suitable regularization term. Here, y_j denotes the j-th component function of y. The diffusive regularizer penalizes large derivatives in the deformation field and thus drives the registration towards a smooth solution.

2.1 Incorporating Segmentation Information

Assume now that additional segmentation information is available for both reference and template image, i.e. there exist 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

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

In order to enable an easy integration into existing variational registration methods, we chose to 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 or penalty 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 joint 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. The penalty approach is able to tolerate minor inaccuracies or inconsistencies in the segmentations as the registration remains image data driven. Note that the term $\mathcal{B}$ coincides with the sum of squared differences (2) of the segmentation masks as binary images.

2.2 Singularities

A serious weakness of the classical variational approach is that is 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.

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, 12] 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

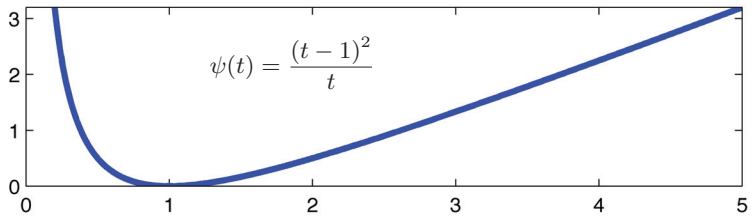

Fig. 1. Change of volume weighting function ψ

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

We now describe the application of our method to the registration of lungs in thoracic CT scans. For all our tests we used the images taken from the preceding MICCAI Grand Challenge EMPIRE10 [4] (image courtesy Bram van Ginneken, University Nijmegen Medical Centre, The Netherlands).

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 start 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. [7]) with the SSD distance measure (2). For this purpose, the masks are downsampled by a factor of five in each dimension preceded by a Gaussian smoothing. The result is again smoothed in order to spread the boundary information of the masks over a larger area.

There is a second reason to use downsampled masks besides the obvious gain in runtime and memory consumption: As we aim for a robust removal of the largest parts of the deformation between the images, we are not interested in small details of the mask shape. Moreover, these details can cause the optimization algorithm to get stuck in local minima. Hence, the reduced complexity and extended smoothness of the resampled masks is desirable for an effective pre-registration.

The optimization is achieved by a standard 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.

Fig. 2 shows the lung masks of a breathhold inspiration and expiration scan pair before registration, Fig. 3 shows the parametric registration result. A major part of the comparably large deformation is already removed in this step.

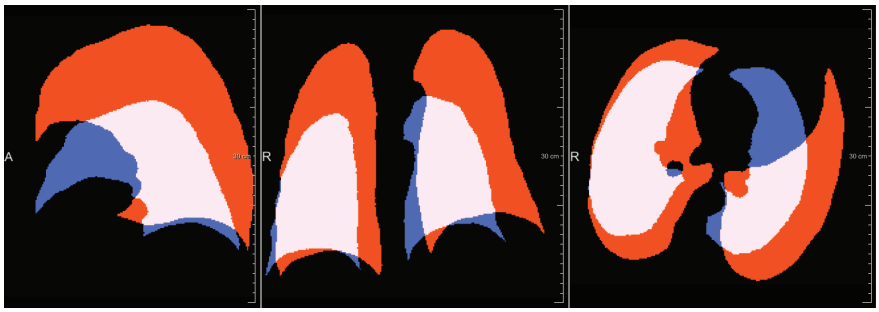

Fig. 2. Lung masks before registration in sagittal, coronal and axial view (from left to right). Note the large volume difference between inhale (orange mask) and exhale state (blue mask).

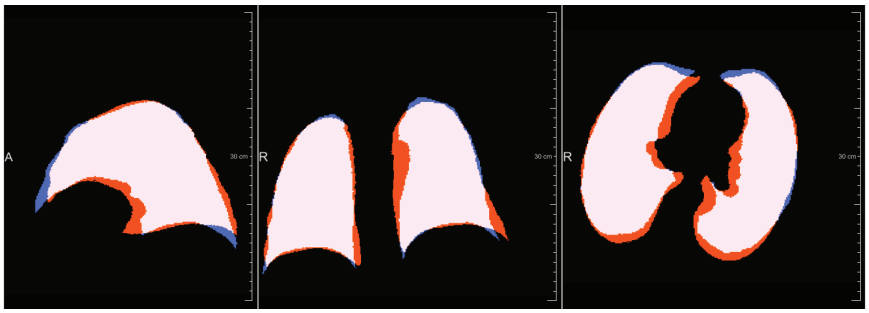

Fig. 3. Result of the affine-linear pre-registration of the lung masks in sagittal, coronal and axial view (from left to right).

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. The reasons for this are similar as for the mask registration in the first phase: Omitting small details on coarser levels in general leads to good alignment of larger structures and gives a better start value on the next, richer level because the main deformations should already have been corrected for.

The multi-level pyramid is generated using Gaussian smoothing and subsequent downsampling by a factor of two. We use three levels for the registration and stop at the second finest level for computational reasons. The sum of squared differences (2) is used as distance measure together with the diffusive 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 $64\times64\times64$ 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 carried out using the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) quasi-Newton optimization method [9]. The runtime of our algorithm ranges from 4 to 15 minutes on our system depending on the resolution of the CT scans.

Figs. 4 and 5 show the registration results. The reference scan is overlayed with the lung boundary of the deformed template image. For all registrations, $\alpha = 10^3, \beta = 10^5$ and $\gamma = 10^4$ were used as parameters. For comparison, a registration with the same parameters but with $\beta = 0$, i.e. without the area correspondence term, is displayed. Although both registrations show a good match of the lung boundaries, adding the area correspondence term leads to an improvement especially in the lung tips. The deformation fields show no singularities as expected by the construction of the energy functional.

In Fig. 6, an additional overlay of the main lung vessels is shown. Both results show a satisfactory alignment of the vessel structure indicating that the standard variational approach already provides a good starting point. It can, however, be further improved by adding specifically tailored terms to the objective function as done here.

In order to further analyze the improvement achieved by our method, we performed a quantitative evaluation of eight breathhold inspiration and expiration scan pairs. We registered the scan pairs with the additional area correspondence term and compared the results to the registration without the extension. For evaluation, the dice coefficient and the volume of the difference image of reference and deformed template mask were computed.

Table 1 shows the results of the evaluation. Our approach leads to an average reduction of the difference volume by a factor of 1.69. The average dice coeffi-

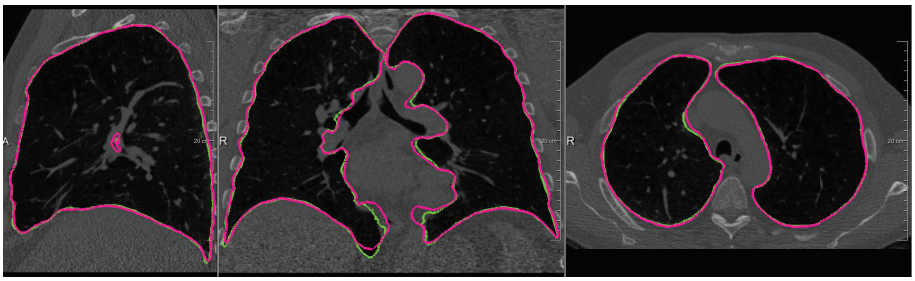

Fig. 4. Reference scan overlayed with lung contours of deformed template image in sagittal, coronal and axial view. The green contour shows the registration result with $\beta = 10^4$, the red contour represents the registration result with $\beta = 0$, i.e. without the additional area correspondence term.

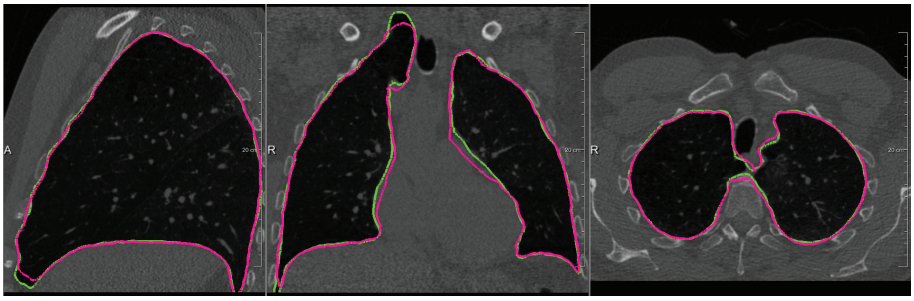

Fig. 5. Example of a breathhold inspiration and expiration scan pair. The reference scan is overlayed with lung contours of deformed template image with $\beta = 0$ (red) and $\beta = 10^4$ (green). Note the improved alignment especially in the lung tips.

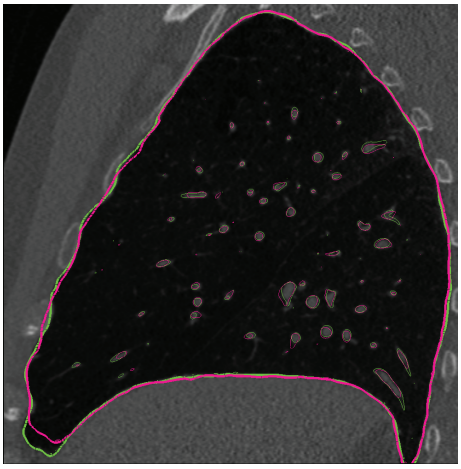

Fig. 6. Sagittal view with contours of main lung vessels shown as overlay (red for $\beta = 0$, green for $\beta = 10^4$). Both registrations show a satisfactory vessel alignment.

cient of 0.968 without the area correspondence term already indicates a good alignment of the lung boundaries. Using the area term, a further improvement to 0.982 could be achieved.

Scan	$\text{vol}_{\text{diff}}, \beta = 0$	$\text{vol}_{\text{diff}}, \beta = 10^5$	Dilated Mask	Dice, $\beta = 0$	Dice, $\beta = 10^5$
1	176.52	123.15	239.10	0.973	0.981
2	272.94	224.86	351.93	0.977	0.981
3	149.74	137.77	330.60	0.986	0.987
4	407.57	172.95	311.61	0.962	0.983
5	280.57	209.78	277.03	0.971	0.978
6	279.78	142.74	267.82	0.969	0.984
7	419.67	118.41	202.51	0.919	0.976
8	145.17	135.11	295.02	0.985	0.986

Table 1. Quantitative evaluation results. For comparison, the volume of the difference image of reference mask and dilated reference mask is given. A kernel of 3x3x3 voxels was used for dilation.

4 Summary and Discussion

In this work, we have presented an extended generic variational approach for CT lung registration that allows for the incorporation of a-priori known area correspondences. It prevents extreme non-physical volume changes and guarantees invertibility of the computed deformation. We have shown how this extension improves the alignment of lung boundaries in the registration of lungs from CT images.

In the public EMPIRE10 challenge [4], our approach achieved a 10^{th} place out of 27 contributions with no singularites and rank three in the lung boundary alignment criterion. A detailed analysis of the comparably unsatisfactory test cases shows that our optimization scheme did not perform well on all pyramid levels. Hence, we expect further improvements by modifications in our optimization method.

In the future, we would like to extend our approach by using a label image of the lung lobes instead of a lung segmentation. In contrast to mere separating the image into lung parenchyma and background, the differentiation into five lung tissue types is expected to further steer the registration towards plausible deformations that respect the lung and lobe boundaries. Moreover, as the fissures separating the lung lobes are often hardly visible and offer low contrast, we expect the registration to benefit even more from this additional information. Unfortunately, this low contrast also makes lung lobe segmentation a rather demanding and difficult problem.

Therefore, we look very much forward to the results of the LOLA11 [6] lung and lobe segmentation challenge on MICCAI 2011 in Toronto, Canada.

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. T. Lange, N. Papenberg, J. Olesch, B. Fischer, and P.M. Schlag, "Landmark Constrained Non-rigid Image Registration with Anisotropic Tolerances", World Congress on Medical Physics and Biomedical Engineering, Springer, IFMBE Proceedings, vol. 25/IV, pp. 2238-2241, 2009.
6. The Lobe and Lung Analysis Challenge 2011, `http://www.lola11.com/`
7. J. Modersitzki, "FAIR: Flexible Algorithms for Image Registration," SIAM, Philadelphia, 2009.
8. 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.
9. J. Nocedal and S.J. Wright, "Numerical optimization", Springer, 2006.
10. J. Olesch, N. Papenberg, T. Lange, M. Conrad, and B. Fischer, "Matching CT and Ultrasound Data of the Liver by Landmark Constrained Image Registration", SPIE, vol. 7261,72610G, 2009.
11. N. Papenberg, J. Olesch, T. Lange, P.M. Schlag, and B. Fischer, "Landmark Constrained Non-parametric Image Registration with Isotropic Tolerances", Bildverarbeitung für die Medizin, Springer, Informatik aktuell, pp. 122-126, 2009.
12. 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.
13. K. Rohr, "Landmark-Based Image Analysis", Computational Imaging and Vision, vol. 21, Springer 2001.