DIS-CO Lung: Fast and Accurate Lung CT Alignment Using Discrete Keypoint Detection and Continuous Image Registration

Jan Rühaak[1], Thomas Polzin[2], and Mattias P. Heinrich[3]

[1] Fraunhofer MEVIS, Lübeck, Germany
[2] Institute of Mathematical Image Computing, University of Lübeck, Germany
[3] Institute of Medical Informatics, University of Lübeck, Germany
jan.ruehaak@mevis.fraunhofer.de, www.mevis.fraunhofer.de,
www.mpheinrich.de

Abstract. Despite great research effort over the previous years, accurate registration of thoracic CT scans remains challenging. In particular, large motion between full inhalation and exhalation leads to unacceptably high alignment errors. We address these challenges by combining discrete and continuous optimization and have developed a new algorithm which performs accurate sparse-to-dense alignment of thoracic CT scans. First, a discrete keypoint detection and matching using sparse graph-based correspondences is performed. Second, a continuous, deformable image registration incorporating dense intensity and the aforementioned correspondence information is employed. A volume change control mechanism is used that limits local volume change and prohibits foldings of the deformation grid. The method achieves high accuracy and very competitive run times of less than 5 minutes per registration.

Keywords: EMPIRE10, pulmonary image registration, optimization, keypoints

1 Introduction

Since its launch in 2010, the EMPIRE10 challenge [10] has grown into the most comprehensive comparison study on pulmonary image registration worldwide. So far 42 algorithms from both academia and industry have been submitted for public evaluation and comparison. The algorithms are evaluated using four different criteria: lung boundary alignment, major fissure alignment, distance of expert-annotated landmarks and deformation field singularities.

In this note, we present a novel lung image registration algorithm as submitted to the EMPIRE10 challenge on September 12th, 2016. The method consists of two phases: a graph-based matching of a large number of keypoints for the estimation of robust large-motion correspondences is followed by a continuous, deformable image registration incorporating both image intensities and keypoint information. We briefly describe our algorithm and its parameter settings used for this challenge submission in the following. For more details of the involved methods and the quantitative results, we refer to the relevant publications and the EMPIRE10 website, respectively.

2 Method

We start by describing the initial keypoint detection and keypoint matching before giving details of the deformable registration step. Let $\mathcal{F}\colon \mathbb{R}^3 \to \mathbb{R}$ denote the fixed (reference) image and $\mathcal{M}\colon \mathbb{R}^3 \to \mathbb{R}$ the moving (template) image with support in domains $\Omega_{\mathcal{F}} \subseteq \mathbb{R}^3$ and $\Omega_{\mathcal{M}} \subseteq \mathbb{R}^3$, respectively. Image registration is then commonly defined as finding a transformation $y : \Omega_{\mathcal{F}} \to \mathbb{R}^3$ encoding the spatial correspondence between the two images $\mathcal{F}$ and $\mathcal{M}$.

2.1 Preprocessing

Since the organizers of the EMPIRE10 challenge provide lung segmentations for both fixed and moving scan, we incorporate this information into the registration model. First, the lung masks are aligned by their centers of gravity, which is followed by an affine-linear registration. Second, the shape information of the segmentations is further used within a nonlinear registration of the mask images (using the same model as described in Section 2.2 and a SSD similarity metric), which yields a pre-registered moving image $\hat{\mathcal{M}}$. The segmentations are also used to mask the CT scans to the lung region, thereby in particular removing the rib cage from the images. This allows to refrain from modeling sliding motion along the ribs when focusing on the lung region only. In the following, we describe the accurate alignment of inner-lung structures (vessels, airways and fissures) using our combination of a sparse discrete correspondence search and a dense intensity-driven variational approach. Throughout these steps the same contrast-invariant similarity metric Normalized Gradient Fields (NGF) [3] is used, which is based on edge information. The underlying assumption of NGF states that corresponding structures are characterized by common change in image intensity. The choice of NGF within pulmonary image registration is motivated by the varying parenchyma densities at different breathing phases that violate, e.g., the intensity constancy assumption. We use the variant

$$\mathcal{D}(\mathcal{F}, \mathcal{M}(y)) := \int_{\Omega_{\mathcal{F}}} 1 - \left(\frac{\langle \nabla\mathcal{M}(y(x)), \nabla\mathcal{F}(x)\rangle_\eta}{\|\nabla\mathcal{M}(y(x))\|_\eta \, \|\nabla\mathcal{F}(x)\|_\eta} \right)^2 \, \mathrm{d}x \tag{1}$$

with

$$\langle f, g\rangle_\eta := \sum_{j=1}^{3} f_j g_j + \eta^2 \qquad \text{and} \qquad \|f\|_\eta := \sqrt{\langle f, f\rangle_\eta}$$

and short finite differences for image gradient computation, cf. [19]. We fix the edge parameter to $\eta = 12$ for all registrations.

2.2 Correspondence search

The estimation of robust correspondence fields closely follows the method presented in [5][4]. Both fixed image $\mathcal{F}$ and pre-registered moving image $\hat{\mathcal{M}}$ are resampled to 1 mm isotropic resolution. We then detect a fairly large number of keypoints $k \in K \subset \mathbb{R}^3$ (on average 3500) based on the Förstner operator [15] (and a Gaussian kernel with $\sigma = 1.4$) within the lung mask of the fixed scan. A good dispersion of the locations is achieved by omitting a fixed distinctiveness threshold and instead performing a non-maximum suppression within a window of 6^3 voxels.

A block-matching search (using patches P_k with voxels p, and a side-length of 7) is performed for each keypoint over a large number of potential displacements $d \in \{0, \pm 2, \pm 4, \ldots, \pm 32\}^3$. A discretized version of the above described NGF metric is used as similarity metric:

$$\mathcal{D}_{\mathrm{KP}}(k,d) = \frac{1}{|P_k|} \sum_{p \in P_k} 1 - \frac{\langle \nabla \hat{\mathcal{M}}(p+d), \nabla \mathcal{F}(p) \rangle_\eta^2}{\|\nabla \hat{\mathcal{M}}(p+d)\|_\eta^2 \, \|\nabla \mathcal{F}(p)\|_\eta^2}$$

Since there might still be many erroneous or ambiguous matches after this calculation, we enforce a certain regularity and find the optimal combination of displacements for all keypoints simultaneously using dynamic programming. The neighbourhood graph is approximated by a minimum-spanning tree and pairwise displacement differences are penalized with a a diffusion-like regularizer $\mathcal{R}_{\mathrm{KP}}(d_k, d_q)$ [6]:

$$\mathcal{R}_{\mathrm{KP}}(d_k, d_q) = \frac{\alpha_{\mathrm{KP}} ||d_k - d_q||^2}{\sqrt{||x_k - x_q||^2} + |\mathcal{F}(k) - \mathcal{F}(q)|/\sigma_I}$$

The robustness of the estimated correspondence field is further improved using a symmetric marginal averaging (see [5] for details). In contrast to [5], we omit a refinement step with higher keypoint density. The sparse correspondence field y_{KP} is only defined at keypoint locations and employed as keypoint penalty $\mathcal{K}$ in the subsequent dense registration.

Dense Registration Framework The intensity-driven part of our registration approach is based on [16,18]. The chosen model consists of several terms corresponding to assumptions about the nature of pulmonary image registration. These are now explained step by step.

The registration is performed within the classical derivative-based optimization framework as described in [9]. In this approach, a joint objective function typically consisting of a distance measure and a regularization term is minimized using derivative-based numerical optimization schemes such as the Gauss-Newton algorithm, cf. also [8]. The objective function of our lung registration

[4] C++ source code for the estimation of correspondence fields is publicly available at `http://mpheinrich.de/software.html`

model consists of five components. Apart from weighting parameters, the first four components are identical to the ones proposed in [18]. To keep the description self-contained, they are additionally briefly repeated here. The main component of the objective function is the NGF distance term $\mathcal{D}$, given in (1), that describes the employed concept of image similarity.

As a basic requirement for a successful lung registration, the computed deformation shall map tissue within the lung in scan $\mathcal{F}$ to tissue within the lung in scan $\mathcal{M}$, and analogously the area outside the lung in scan $\mathcal{F}$ to the area outside the lung in scan $\mathcal{M}$. To this end, lung boundary information is incorporated into the model by using an additional penalty term $\mathcal{B}$ as proposed in [17]. Let $b_{\mathcal{F}}\colon \Omega_{\mathcal{F}} \to \{0,1\}$ and $b_{\hat{\mathcal{M}}}\colon \Omega_{\mathcal{M}} \to \{0,1\}$ denote binary functions for $\mathcal{F}$ and $\mathcal{M}$ that are equal to one inside the lungs and zero otherwise. We define the penalty term $\mathcal{B}$ as

$$\mathcal{B}(y) := \frac{1}{2}\int_{\Omega_{\mathcal{F}}} \Big(b_{\mathcal{M}}\big(y(x)\big) - b_{\mathcal{F}}(x)\Big)^2 \,\mathrm{d}x.$$

Note that $\mathcal{B}$ coincides with the sum of squared differences (SSD) of the lung segmentation masks as binary images.

As non-continuous sliding motion along the rib cage does not have to be recovered when using lung segmentations, the expected remaining motion is very smooth. Based on this observation, the curvature regularizer [2]

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

is selected, where y^{kern} is a given transformation and y_j, y_j^{kern} denote the j-th component of y and y^{kern}, respectively. The curvature regularizer penalizes second order derivatives of the deviation of y from y^{kern} which is typically set to the result of a pre-registration.

Although the curvature regularizer generates smooth transformations, it cannot safeguard against physically implausible deformations exhibiting large local volume change or even foldings. To this end, we add an additional term similar to the approaches presented in [1,14] that directly incorporates change of volume,

$$\mathcal{V}(y) := \int_{\Omega_{\mathcal{F}}} \psi(\det \nabla y(x)) \,\mathrm{d}x$$

with weighting function

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

We call $\mathcal{V}$ *volume change control*. As $\psi(t) = \psi(1/t)$, deviations of the Jacobian from 1, i.e. local volume expansion or shrinkage, are symmetrically penalized. In addition, ψ ensures (local) injectivity of the deformation since $\psi(\det \nabla y) \to \infty$ as $\det \nabla y \to 0$. Hence, $\mathcal{V}(y) = \infty$ if the Jacobian becomes negative at any point. The volume change control therefore fulfills two important functions: It prevents foldings and ensures that large changes in volume are penalized.

Integration of Keypoint Information Keypoint information has been successfully used to support intensity-based algorithms for pulmonary image registration, cf., e.g., [4,12,13]. We follow the approach of Han [4] and add the keypoints as soft constraints to the model, thereby also limiting the influence of possible erroneous keypoint matches. We wish to minimize the sum of all squared Euclidean distances for all keypoint locations between our final transformation y and the sparse keypoint transformation y_{KP}:

$$\mathcal{K}(y) := \sum_{i=1}^{|K|} \|y(x_i) - y_{\mathrm{KP}}(x_i)\|_2^2$$

With weighting parameters $\alpha, \beta, \gamma, \delta > 0$, the full model is given by

$$\mathcal{J}(y) = \mathcal{D}(\mathcal{F}, \mathcal{M}(y)) + \alpha\mathcal{S}(y) + \beta\mathcal{B}(y) + \gamma\mathcal{V}(y) + \delta\mathcal{K}(y). \tag{2}$$

Optimization The minimization of the joint objective function (2) is performed with the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) quasi-Newton optimization method [11]. We employ matrix-free derivative calculation schemes for all objective function components except the keypoint distance term to reduce runtime and memory footprint and allow for parallel computation, cf. [7]. The optimization is performed with a multi-level strategy ranging from coarse to fine. The multi-level pyramid is generated using Gaussian smoothing and subsequent downsampling by a factor of two in each dimension. Five levels are used, with the original image as the finest level.

Additionally, the deformation grid resolution is also embedded into the multi-level strategy. We employ a grid of $128 \times 128 \times 128$ cells on the finest level. For each coarser level, the deformation grid size is reduced by a factor of two in each dimension, leaving a deformation resolution of $8 \times 8 \times 8$ at the coarsest resolution.

For all registrations, we used $\alpha_{\mathrm{KP}} = \frac{1}{45}, \alpha = 2, \beta = 1, \gamma = 0.001$ as parameters. The weight δ of the keypoint matching term $\mathcal{K}$ was adaptively determined such that $\mathcal{D}(\mathcal{F}, \mathcal{M}(y)) = \delta\mathcal{K}(y)$ at the beginning of the optimization on the coarsest level of the multilevel scheme. Similarly, we chose β for the lung mask term adaptively such that $\mathcal{D}(\mathcal{F}, \mathcal{M}(y)) = 10\beta\mathcal{B}(y)$. We reduced the weight by the factor 10 due to the fact, that the lung masks are aligned quite well after the preprocessing.

3 Results

The deformable registration, which was performed on a quad-core processor, achieves very fast runtimes of on average 2:33 minutes. The keypoint search in its current implementation (on a single core) requires a similar amount of time, but could easily be speeded up by exploiting its inherent parallelism. For detailed quantitative results, the reader is referred to the EMPIRE10 website `http://empire10.isi.uu.nl/res_lbeckdisco.php`. To summarize, we obtain very low target registration errors, as well as very good fissure alignments. The deformation fields contain no singularities (thus no folding) and the lung masks are perfectly aligned for most cases. Overall, our algorithm ranks among the best approaches of the ongoing challenge at the time of submission. Since our method is computationally much faster than most competing methods, it provides an excellent choice for general purpose pulmonary image registration.

References

1. Burger, M., Modersitzki, J., Ruthotto, L.: A hyperelastic regularization energy for image registration. SIAM Journal on Scientific Computing 35(1), B132–B148 (2013)
2. Fischer, B., Modersitzki, J.: Curvature based image registration. Journal of Mathematical Imaging and Vision 18(1), 81–85 (2003)
3. Haber, E., Modersitzki, J.: Intensity gradient based registration and fusion of multi-modal images. Methods of Information in Medicine 46, 292–9 (2007)
4. Han, X.: Feature-constrained Nonlinear Registration of Lung CT Images. In: Medical Image Analysis for the Clinic: A Grand Challenge. pp. 63–72 (2010)
5. Heinrich, M.P., Handels, H., Simpson, I.J.: Estimating Large Lung Motion in COPD Patients by Symmetric Regularised Correspondence Fields. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015, pp. 338–345. Springer (2015)
6. Heinrich, M., Jenkinson, M., Papież, B., Glesson, F., Brady, M., Schnabel, J.: Edge- and detail-preserving sparse image representations for deformable registration of chest MRI and CT volumes. In: Information Processing in Medical Imaging – IPMI 2013. pp. 463–474. Springer (2013)
7. König, L., Rühaak, J.: A Fast and Accurate Parallel Algorithm for Non-Linear Image Registration using Normalized Gradient Fields. In: Biomedical Imaging (ISBI), 2014 IEEE 11th International Symposium on. pp. 580–583 (2014)
8. Modersitzki, J.: Numerical Methods for Image Registration. Oxford University Press (2004)
9. Modersitzki, J.: FAIR: Flexible Algorithms for Image Registration. SIAM (2009)
10. Murphy, K., Van Ginneken, B., Reinhardt, J.M., Kabus, S., Ding, K., Deng, X., Cao, K., Du, K., Christensen, G.E., Garcia, V., et al.: Evaluation of registration methods on thoracic CT: the EMPIRE10 challenge. Medical Imaging, IEEE Transactions on 30(11), 1901–1920 (2011)
11. Nocedal, J., Wright, S.: Numerical optimization. Springer (1999)
12. Polzin, T., Rühaak, J., Werner, R., Handels, H., Modersitzki, J.: Lung registration using automatically detected landmarks. Methods of Information in Medicine 53, 250–256 (2014)
13. Polzin, T., Rühaak, J., Werner, R., Strehlow, J., Heldmann, S., Handels, H., Modersitzki, J.: Combining Automatic Landmark Detection and Variational Methods for Lung CT Registration. In: Fifth International MICCAI Workshop on Pulmonary Image Analysis. pp. 85–96. Nagoya, Japan (September 2013)
14. Rohlfing, T., Maurer Jr, C.R., Bluemke, D.A., Jacobs, M.A.: Volume-preserving nonrigid registration of MR breast images using free-form deformation with an incompressibility constraint. Medical Imaging, IEEE Transactions on 22(6), 730–741 (2003)
15. Rohr, K.: Landmark-Based Image Analysis Using Geometric and Intensity Models. Springer (2001)

16. Rühaak, J., Heldmann, S.: Variational Lung Registration With Explicit Boundary Alignment. Grand Challenges in Medical Image Analysis (2012)
17. Rühaak, J., Heldmann, S., Fischer, B.: Improving Lung Registration by Incorporating Anatomical Knowledge: A Variational Approach. In: 4th International Workshop on Pulmonary Image Analysis. pp. 147–156 (2011)
18. Rühaak, J., Heldmann, S., Kipshagen, T., Fischer, B.: Highly Accurate Fast Lung CT Registration. In: SPIE Medical Imaging, Image Processing. pp. 86690Y–86690Y–9. International Society for Optics and Photonics (2013)
19. Rühaak, J., König, L., Hallmann, M., Papenberg, N., Heldmann, S., Schumacher, H., Fischer, B.: A Fully Parallel Algorithm for Multimodal Image Registration Using Normalized Gradient Fields. In: Biomedical Imaging (ISBI), 2013 IEEE 10th International Symposium on. pp. 572–575 (2013)