

Memory Efficient LDDMM for Lung CT

Thomas Polzin¹(✉), Marc Niethammer², Mattias P. Heinrich³, Heinz Handels³,
and Jan Modersitzki^{1,4}

¹ Institute of Mathematics and Image Computing,
University of Lübeck, Lübeck, Germany
polzin@mic.uni-luebeck.de

² Department of Computer Science and Biomedical Research Imaging Center,
University of North Carolina at Chapel Hill, Chapel Hill, USA

³ Institute of Medical Informatics, University of Lübeck, Lübeck, Germany

⁴ Fraunhofer MEVIS, Lübeck, Germany

Abstract. In this paper a novel Large Deformation Diffeomorphic Metric Mapping (LDDMM) scheme is presented which has significantly lower computational and memory demands than standard LDDMM but achieves the same accuracy. We exploit the smoothness of velocities and transformations by using a coarser discretization compared to the image resolution. This reduces required memory and accelerates numerical optimization as well as solution of transport equations. Accuracy is essentially unchanged as the mismatch of transformed moving and fixed image is incorporated into the model at high resolution. Reductions in memory consumption and runtime are demonstrated for registration of lung CT images. State-of-the-art accuracy is shown for the challenging DIR-Lab chronic obstructive pulmonary disease (COPD) lung CT data sets obtaining a mean landmark distance after registration of 1.03 mm and the best average results so far.

1 Introduction

COPD is the fourth leading cause of death and 24 million people are afflicted in the US alone [14]. Lung registration could support assessment of COPD phenotypes as well as disease progression [6] and improved treatment of patients and speedups in clinical workflows are expected if registration is used for follow-up inspection and/or motion estimation in treatment planning [12, 14]. Hence, registration of inhale/exhale and longitudinal data is critical, but also challenging due to the presence of large non-linear deformations, which should be diffeomorphic [12]. LDDMM can address these difficulties [2]. However, its drawbacks are large memory use and high computational costs: e.g. run times of up to three hours are reported in [17] for moderately sized data of $256 \times 192 \times 180$ voxels using 32 CPUs and 128 GB RAM. We present a new variant of LDDMM which

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-3-319-46726-9_4](https://doi.org/10.1007/978-3-319-46726-9_4)) contains supplementary material, which is available to authorized users.

- has significantly lower computational and memory demands,
- employs the well-suited distance measure Normalized Gradient Fields [11],
- and is as accurate as standard LDDMM and state-of-the-art algorithms.

Lung registration faces severe challenges. Diffeomorphic modeling of large lung motion (for inhale/exhale in the order of 50 mm [5]) is necessary [12]. Lung volume changes are large: a doubling of lung capacity is not unusual during inspiration [5]. Additionally, acquisitions with limited dose deteriorate CT quality [14].

Related Work: Sakamoto et al. [17] applied LDDMM in pulmonary computer-aided diagnosis (CAD) to analyze and detect growing or shrinking nodules in follow-up CT scans with encouraging results, but at high computational cost. A finite dimensional Lie algebra was introduced and integrated into a geodesic shooting approach by Zhang and Fletcher [19] to considerably speed up computations. Their experiments on 3D brain MRI data showed no loss of accuracy. Risser et al. [15] used appropriate kernels for LDDMM to cope with sliding motion occurring at the interface between lungs and the ribcage as this motion is not diffeomorphic. Our focus is on the registration of inner lung structures, hence we use lung segmentations as in [16] and thereby avoid sliding issues.

Contributions: We focus on the relaxation formulation of LDDMM [2], but the approach could easily be generalized to shooting formulations [1, 18]. Our scheme is based on the smoothness of transformations and velocities computed with LDDMM allowing discretizations at lower spatial resolution and consequentially substantially reduced memory requirements. The image match is computed at the original resolution thereby maintaining accurate results. Furthermore, we employ the Normalized Gradient Fields (NGF) image similarity measure [11] that aligns edges, e.g. vessels. We use NGF to cope with large volume changes influencing tissue densities and thus absorption of X-rays. The resulting gray values of the parenchyma are quite different for inhale and exhale scans. Hence, sum of squared differences might not be an appropriate similarity measure.

Outline: In Sect. 2 the optimization problem, including the energy with NGF distance measure, and the memory efficient implementation of LDDMM are described. We evaluate our method on 20 3D lung CT scan pairs in Sect. 3, compare the results to the state of the art and show that a coarser discretization of velocities and transformations is sufficient for an accurate registration. In Sect. 4 results and possible extensions are discussed.

2 Methods

2.1 Continuous Model and LDDMM Background

Let $I_0, I_1: \mathbb{R}^d \supset \Omega \rightarrow \mathbb{R}$ denote the moving and fixed images with domain Ω . LDDMM uses space- and time-dependent velocity ($v: \Omega \times [0, 1] \rightarrow \mathbb{R}^d$) and

transformation ($\varphi: \Omega \times [0, 1] \rightarrow \mathbb{R}^d$) fields and seeks the minimizer (v^*, φ^*) of the following energy subject to a transport constraint [2, 7]:

$$\left. \begin{aligned} E(v, \varphi) &:= \int_0^1 \langle Lv(\cdot, t), Lv(\cdot, t) \rangle dt + \frac{1}{\sigma^2} D(\varphi(\cdot, 1); I_0, I_1) \rightarrow \min, \\ \text{s. t. } \varphi_t + J_\varphi v &= \mathbf{0}, \quad \varphi(\mathbf{x}, 0) = \mathbf{x} \text{ for all } \mathbf{x} \in \Omega, t \in [0, 1] \end{aligned} \right\} \quad (1)$$

We denote the partial time derivative as $\varphi_t(\mathbf{x}, t) \in \mathbb{R}^d$ and the Jacobian with respect to the spatial coordinates as $J_\varphi(\mathbf{x}, t) \in \mathbb{R}^{d \times d}$. L is a suitable differential operator enforcing smoothness of v . In our method the Helmholtz operator $Lv := \gamma v - \alpha \Delta v$ with $\alpha, \gamma > 0$ was used, which is a standard choice for LDDMM registration, cf., e.g., [2, 7]. In all experiments we fixed $\gamma = 1$. D is a general distance measure and $\sigma > 0$ its weighting parameter. As motivated in Sect. 1 we use the Normalized Gradient Fields (NGF) distance measure [11]. NGF was successfully applied to lung CT registration [9, 16] in the following adaption of the original formulation:

$$D(\varphi_1; I_0, I_1) := \int_\Omega 1 - \left(\frac{\langle \nabla I_0(\varphi_1(\mathbf{x})), \nabla I_1(\mathbf{x}) \rangle_\eta}{\|\nabla I_0(\varphi_1(\mathbf{x}))\|_\eta \|\nabla I_1(\mathbf{x})\|_\eta} \right)^2 d\mathbf{x}, \quad (2)$$

with $\varphi_1 := \varphi(\cdot, 1)$, $\langle \mathbf{u}, \mathbf{v} \rangle_\eta := \eta^2 + \sum_{i=1}^d u_i v_i$ and $\|\mathbf{u}\|_\eta^2 := \langle \mathbf{u}, \mathbf{u} \rangle_\eta$ for $\mathbf{u}, \mathbf{v} \in \mathbb{R}^d$. The parameter $\eta > 0$ is used to decide if a gradient is considered noise or edge [11]. Throughout this paper $\eta = 100$ is used as proposed in [16].

Following the steps of [7] we add the transport equation constraint of (1) into the objective functional employing Lagrange multipliers $\lambda: \Omega \times [0, 1] \rightarrow \mathbb{R}^d$. After some straightforward calculations we obtain the necessary conditions as

$$L^\dagger Lv + J_\varphi^\top \lambda = 0, \quad (3)$$

$$\varphi_t + J_\varphi v = \mathbf{0}, \quad \varphi(\mathbf{x}, 0) = \mathbf{x}, \quad (4)$$

$$\lambda_t + J_\lambda v + \text{div}(v)\lambda = \mathbf{0}, \quad \lambda(\mathbf{x}, 1) = -(1/\sigma^2) \nabla_\varphi D(\varphi_1(\mathbf{x}); I_0, I_1) \quad (5)$$

for all $\mathbf{x} \in \Omega, t \in [0, 1]$ and a differentiable D . Solving (4) means transporting the transformation maps according to the velocities whereas solving (5) is the flow of the image mismatch (given by the derivative of the distance measure) backward in time, cf. [7]. As in [7] we apply the smoothing kernel $(L^\dagger L)^{-1}$ before updating the velocities in (3) to improve numerical optimization:

$$p := v + (L^\dagger L)^{-1} (J_\varphi^\top \lambda). \quad (6)$$

2.2 Discretization and Algorithm

From now on we use $d = 3$. Velocities v and transformations φ were discretized on a nodal grid in 4D (3D + t) [11]. The number of time points was set to $n_4 = 11$. The number of points in space varied during the multi-level optimization but was chosen equal for all spatial directions, i.e. $n_1 = n_2 = n_3$. Defining $\bar{n} := n_1 n_2 n_3$ and using linear ordering in space yields arrays $\mathbf{v}, \varphi \in \mathbb{R}^{3\bar{n} \times n_4}$.

Algorithm 1. 3D multi-level LDDMM

Data: $\mathbf{I}_0^i, \mathbf{I}_1^i \in \mathbb{R}^{m_1^i \times m_2^i \times m_3^i}$, $i = 1, \dots, F$; $\alpha, \gamma, \sigma > 0$, $\mathbf{n} \in \mathbb{N}^4$
Optional Data: $\psi \in \mathbb{R}^{3\bar{n}}$
Result: $\mathbf{v}, \varphi \in \mathbb{R}^{3\bar{n} \times n_4}$

Initialize $\mathbf{v} \leftarrow \mathbf{0} \in \mathbb{R}^{3\bar{n} \times n_4}$, $\varphi_{:,1} \in \mathbb{R}^{3\bar{n}}$ as regular nodal grid on Ω or $\varphi_{:,1} \leftarrow \psi$
for $i \in \{1, 2, \dots, F\}$ **do**
 if $i > 1$ **then**
 // Interpolate solution of level $i - 1$ to level i
 $\mathbf{v} \leftarrow \text{Prolongate}(\mathbf{v}, \mathbf{n})$; $\varphi_{:,1} \leftarrow \text{Prolongate}(\varphi_{:,1}, \mathbf{n})$
 while *stopping criteria not satisfied* **do**
 for $j \in \{1, 2, \dots, n_4 - 1\}$ **do**
 Compute $\varphi_{:,j+1}$ from $\varphi_{:,j}$ and $\mathbf{v}_{:,j}$ according to (4)
 $\lambda_{:,n_4} \leftarrow -(1/\sigma^2) \nabla_{\varphi} D(\mathbf{P}\varphi_{:,n_4}; \mathbf{I}_0^i, \mathbf{I}_1^i)$
 for $k \in \{n_4 - 1, n_4 - 2, \dots, 1\}$ **do**
 Compute $\lambda_{:,k}$ from $\lambda_{:,k+1}$ and $\mathbf{P}\mathbf{v}_{:,k}$ using (5). Save only $\mathbf{P}^{\top} \lambda$.
 $\mathbf{p} \leftarrow \mathbf{v} + (\mathbf{L}_{\alpha, \gamma}^{\top} \mathbf{L}_{\alpha, \gamma})^{-1} (\mathbf{J}_{\varphi}^{\top} \mathbf{P}^{\top} \lambda)$
 $\mathbf{v} \leftarrow \mathbf{v} - \beta \mathbf{H} \mathbf{p}$ // β is computed by line search, \mathbf{H} by L-BFGS
 $n_l \leftarrow 2n_l - 1$, $l = 1, 2, 3$
 return \mathbf{v}, φ

We use central differences and Neumann boundary conditions for the discretization of div , L , L^{\dagger} and J_{φ} . NGF and its derivative were implemented according to [11]. Equations (4) and (5) were solved with a fourth order Runge-Kutta scheme. Given images $\mathbf{I}_0, \mathbf{I}_1 \in \mathbb{R}^{m_1 \times m_2 \times m_3}$ with $\bar{m} := m_1 m_2 m_3$ voxels, smoothed and downsampled versions $\mathbf{I}_0^i, \mathbf{I}_1^i \in \mathbb{R}^{m_1^i \times m_2^i \times m_3^i}$ with $m_j^i = \lfloor m_j \cdot 2^{-F+i} \rfloor$ were computed for $j = 1, 2, 3$ and $i = 1, \dots, F$ [11]. Problem (1) was then solved using a coarse-to-fine resolution multi-level strategy with $F \in \mathbb{N}$ levels.

Choosing $n_i = m_i$ exceeds common memory limitations and results in an extremely expensive solution of (4) and (6). As \mathbf{v} and φ are assumed to be smooth functions it is usually not necessary to use a high resolution for \mathbf{v} and φ . This motivates our choice of $n_i < m_i$. Nevertheless, we use image information at the highest resolution. Hence \mathbf{v} and φ have to be prolonged to image resolution to solve (5). This can be done gradually, i.e. for only one time point at once and reduces memory consumption substantially. Following [9], we use a prolongation matrix $\mathbf{P} \in \mathbb{R}^{3\bar{m} \times 3\bar{n}}$ that linearly interpolates \mathbf{v} and φ on the cell-centered grid points of the images. \mathbf{P} is sparse but large and does not need to be kept in memory [9]. We use matrix *notation*, but *implemented* a matrix-free operator.

After (5) is solved, the adjoint $\lambda \in \mathbb{R}^{3\bar{m} \times n_4}$ has to be brought back to grid resolution by computation of $\mathbf{P}^{\top} \lambda$ which is then used to solve (3). A memory efficient way to do this is storing $\mathbf{P}^{\top} \lambda$ instead of λ and doing computations gradually again. The numerical solution of Eqs. (3) to (6) is performed in the multi-level registration described in Algorithm 1. It is useful to start Algorithm 1 with the result of a pre-registration $\psi \in \mathbb{R}^{3\bar{n}}$. Given a (preferably diffeomorphic) ψ only the initial condition of (4) has to be adapted: $\varphi(\cdot, 0) = \psi(\cdot)$.

The discretized objective function was optimized with a L-BFGS approach [13] saving the last $M = 5$ iterate vectors for approximation of the inverse Hessian. The maximum number of iterations was set to $k_{\max} = 50$. Additional stopping criteria proposed in [11] were used to terminate the while loop in Algorithm 1. During the optimization an Armijo line search with parameters $\beta_k = 0.5^{k-1}$, $k_{\max}^{\text{LS}} = 30$ and $c_1 = 10^{-6}$ was used to guarantee a decrease of the objective function.

3 Experiments and Results

The proposed method was applied to the DIR-Lab 4DCT [4] (see Sect. 3.1) and COPD [5] (see Sect. 3.2) data sets, as they are a well-known benchmark for lung CT registration. In total 20 inhale/exhale scan pairs are available that include 300 expert annotated landmarks each. The scans were masked with a lung segmentation obtained using [10]. For all experiments a PC with 3.4 GHz Intel i7-2600 quad-core CPU was used. We implemented Algorithm 1 with Matlab and employed C++ code for computations of D , L and P .

3.1 Experiments for Varying n

Since we assume smooth velocities and transformations, we investigated the influence of a coarse grid on registration accuracy. The 4DCT data sets were affinely pre-registered to obtain ψ and Algorithm 1 was used with $\alpha = 200$, $\sigma = 0.01$ and $F = 4$ for all registrations. Initially we set $\bar{n} = 5^3$, $\bar{n} = 9^3$ and $\bar{n} = 17^3$ respectively. The final number of grid points in space for the registration on full image resolution was 33^3 , 65^3 and 129^3 accordingly. We compared the performance for different n parameters using the mean distance of expert annotated landmarks, runtime and memory consumption. See Table 1 for results. One-sided paired t-tests with significance level 0.05 were used to identify improvements in mean landmark error compared to the 65^3 method. Registrations were repeated without the finest image pyramid level to see if a reduction in image data influences accuracy. Again paired t-tests were used to compare to the methods with the same number of grid points that were registered with all levels.

Using 129^3 grid points does not provide the best results, which could be explained by the fact that the optimization is prone to local minima because more degrees of freedom are given. This is visible for case 7, where 129^3 is the fastest method but yields the worst results. Using 65^3 grid points results in significantly better accuracy compared to both 33^3 and 129^3 although improvements with respect to 33^3 are small. For memory consumption and runtime 33^3 is clearly the best choice while 129^3 requires most resources. However, memory requirements do not grow with factor 8 as there is an overhead for saving images and computing λ , therefore 65^3 offers the best compromise. Omitting the last image pyramid level results in a mean memory consumption of 0.80, 1.98 and 10.91 GB for 33^3 , 65^3 and 129^3 grid points respectively. Using the full image data is beneficial for the registration accuracy which is also confirmed by the significance tests.

Table 1. Mean landmark distance per DIR-Lab 4DCT data set, necessary memory and runtime. Rounding to the next regular grid voxel was performed prior to computing the distance. Best values are printed bold. Significant differences to 65^3 are indicated by * and to the respective method with full image information by $^+$.

Case	Mean (mm)			Mean, omitted last image level (mm)			Memory (GB)			Runtime (h:min)		
	33^3	65^3	129^3	33^3	65^3	129^3	33^3	65^3	129^3	33^3	65^3	129^3
4DCT1	0.81	0.86	0.96	1.24	1.16	1.18	0.97	2.15	11.36	0:17	0:36	0:49
4DCT2	0.80	0.79	0.92	1.04	0.87	1.68	1.12	2.30	11.52	0:20	0:32	0:53
4DCT3	0.96	0.96	1.23	1.25	1.20	2.29	1.05	2.23	11.45	0:15	0:29	0:50
4DCT4	1.37	1.23	1.45	1.59	1.44	2.70	1.01	2.19	10.45	0:10	0:25	0:41
4DCT5	1.29	1.20	1.57	1.55	1.53	2.70	1.07	2.25	10.50	0:11	0:29	0:43
4DCT6	1.05	1.03	1.04	1.41	1.20	1.18	4.50	5.67	13.93	0:56	0:55	3:43
4DCT7	1.06	0.98	1.42	1.40	1.37	1.99	4.77	5.94	14.20	1:15	1:13	0:42
4DCT8	1.23	1.13	1.41	1.38	1.32	2.15	4.50	5.67	13.93	1:01	0:55	0:54
4DCT9	1.11	1.06	1.37	1.39	1.25	1.25	4.50	5.67	13.93	0:37	0:55	1:08
4DCT10	1.00	0.98	1.19	1.25	1.23	1.21	4.23	5.40	13.66	0:57	1:05	0:44
Avg	1.07*	1.02	1.26*	1.35 $^+$	1.26 $^+$	1.83 $^+$	2.77	3.95	12.49	0:36	0:45	1:07




Fig. 1. Overlay of a coronal slice of the fixed (blue) and (transformed) moving image (orange) of data set 10 of the DIR-Lab COPD data [5]. Aligned structures are displayed in gray or white due to addition of RGB values. Yellow circles highlight improvements. (Color figure online)

3.2 Comparison to State-of-the-Art Algorithms

In the following experiments we used a thin-plate spline (TPS) pre-registration with keypoints [8] to provide a very accurate initial estimate on the COPD data sets [5]. We ensured that ψ is diffeomorphic by removal of keypoints whose set of six nearest neighbors changes after TPS computation by more than one point, which indicates a violation of topology preservation. This procedure reduced the number of keypoints to approximately one quarter of the original number. Parameters were fixed to $\sigma = 0.1$, $\alpha = 85$ and $F = 5$. Motivated by the

Table 2. Mean and standard deviation of the distances of the 300 expert annotated landmarks per DIR-Lab COPD data set. All values are given in mm. Rounding to the next regular grid voxel was performed prior to computing the distances. Results for state-of-the-art methods are the reported ones from [3, 8, 16]. The p-values were computed performing paired t-tests with the hypothesis that the respective method is better than the proposed one. Best values are printed bold.

Case	Mean \pm Standard deviation					
	Initial	MILO [3]	MRF [8]	NLR [16]	Prereg.	Proposed
COPD1	26.34 \pm 11.43	0.93 \pm 0.92	1.00 \pm 0.93	1.33 \pm 1.55	1.15 \pm 1.00	0.90 \pm 0.93
COPD2	21.79 \pm 6.47	1.77 \pm 1.92	1.62 \pm 1.78	2.34 \pm 2.88	2.18 \pm 2.10	1.56 \pm 1.67
COPD3	12.64 \pm 6.39	0.99 \pm 0.91	1.00 \pm 1.06	1.12 \pm 1.07	1.19 \pm 1.03	1.03 \pm 0.99
COPD4	29.58 \pm 12.95	1.14 \pm 1.04	1.08 \pm 1.05	1.54 \pm 1.61	1.32 \pm 1.12	0.94 \pm 0.98
COPD5	30.08 \pm 13.36	1.02 \pm 1.23	0.96 \pm 1.13	1.39 \pm 1.38	1.18 \pm 1.21	0.85 \pm 0.90
COPD6	28.46 \pm 9.17	0.99 \pm 1.08	1.01 \pm 1.25	2.08 \pm 3.01	1.27 \pm 1.45	0.94 \pm 1.12
COPD7	21.60 \pm 7.74	1.03 \pm 1.08	1.05 \pm 1.07	1.10 \pm 1.28	1.32 \pm 1.45	0.94 \pm 1.25
COPD8	26.46 \pm 13.24	1.31 \pm 1.76	1.08 \pm 1.24	1.57 \pm 2.08	1.47 \pm 1.94	1.12 \pm 1.56
COPD9	14.86 \pm 9.82	0.86 \pm 1.06	0.79 \pm 0.80	0.99 \pm 1.29	1.02 \pm 1.10	0.88 \pm 0.98
COPD10	21.81 \pm 10.51	1.23 \pm 1.27	1.18 \pm 1.31	1.42 \pm 1.44	1.51 \pm 1.39	1.17 \pm 1.28
Average	23.36 \pm 10.11	1.13 \pm 1.23	1.08 \pm 1.16	1.49 \pm 1.76	1.36 \pm 1.38	1.03 \pm 1.16
p-value	$4.8 \cdot 10^{-7}$	$5.5 \cdot 10^{-3}$	0.054	$9.5 \cdot 10^{-4}$	$1.5 \cdot 10^{-5}$	–

results of Sect. 3.1 we used $\mathbf{n} = (65, 65, 65, 11)$ on the finest level. Resulting mean and standard deviations of landmark distances and p-values of one-sided paired t-tests are given in Table 2. The filtering of the keypoints results in a diffeomorphic pre-registration with worse mean landmark distances as visible by comparing the columns MRF and Prereg. respectively.

We compared the proposed method to the MILO [3] and MRF [8] algorithms, which are the two best ranked methods for registration of DIR-Lab COPD data sets. Results of the NLR [16] algorithm are reported because it also uses the NGF distance measure. The proposed registration achieves in most cases the best result and performed significantly better than the pre-registration, MILO [3] and NLR [16]. Mostly it is also better than MRF, which may have singularities while our method is diffeomorphic. The registrations took on average 46 min and used at most 5.9 GB of RAM. Figure 1 shows an exemplary central coronal slice of data set 10 without registration, after TPS pre-registration and after the proposed LDDMM registration. The TPS pre-registration captures the large motion and aligns most of the vessels. The subsequent LDDMM registration improves the alignment of lung boundaries and nearby vessels as well as small vessels as highlighted by the yellow circles.

4 Discussion

In this paper a memory efficient LDDMM scheme was introduced that employs the well-suited NGF distance measure for the registration of pulmonary CT data. Although the method was tested only for lung CT images it is applicable for other anatomical sites or modalities. We accomplish at least a 25-fold reduction of memory requirements (cf. supplementary material) compared to full resolution LDDMM and obtain diffeomorphic mappings without impairment of registration quality. The proposed method achieves the currently lowest average landmark error (1.03 mm) and advances the state of the art on challenging lung CT data [5]. A trustworthy registration is vital for clinical applications such as COPD classification or CAD of lung nodule evolution [6, 14]. A possible extension of our method would be a lung mask similarity term as described in [16] to improve the lung boundary alignment. Another option is the integration of sliding motion into the model as in [15]. We also plan to submit results to the EMPIRE10 challenge [12] for additional evaluation of our method.

Acknowledgments. This work was supported by a fellowship of the German Academic Exchange Service (DAAD) and NSF grant ECCS-1148870.

References

1. Ashburner, J., Friston, K.J.: Diffeomorphic registration using geodesic shooting and Gauss-Newton optimisation. *NeuroImage* **55**(3), 954–967 (2011)
2. Beg, M.F., Miller, M.I., Trounev, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *IJCV* **61**(2), 139–157 (2005)
3. Castillo, E., Castillo, R., Fuentes, D., Guerrero, T.: Computing global minimizers to a constrained B-spline image registration problem from optimal l1 perturbations to block match data. *Med. Phys.* **41**(4), 041904 (2014)
4. Castillo, E., Castillo, R., Martinez, J., Shenoy, M., Guerrero, T.: Four-dimensional deformable image registration using trajectory modeling. *Phys. Med. Biol.* **55**(1), 305–327 (2010)
5. Castillo, R., Castillo, E., Fuentes, D., Ahmad, M., Wood, A.M., et al.: A reference dataset for deformable image registration spatial accuracy evaluation using the COPDgene study archive. *Phys. Med. Biol.* **58**(9), 2861–2877 (2013)
6. Galbán, C.J., Han, M.K., Boes, J.L., Chughtai, K., Meyer, C.R., et al.: Computed tomography-based biomarker provides unique signature for diagnosis of COPD phenotypes and disease progression. *Nat. Med.* **18**, 1711–1715 (2012)
7. Hart, G.L., Zach, C., Niethammer, M.: An optimal control approach for deformable registration. In: *IEEE CVPR Workshops*, pp. 9–16 (2009)
8. Heinrich, M.P., Handels, H., Simpson, I.J.A.: Estimating large lung motion in COPD patients by symmetric regularised correspondence fields. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) *MICCAI 2015*. LNCS, vol. 9350, pp. 338–345. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-24571-3_41](https://doi.org/10.1007/978-3-319-24571-3_41)
9. König, L., Rühaak, J.: A fast and accurate parallel algorithm for non-linear image registration using normalized gradient fields. In: *IEEE ISBI*, pp. 580–583 (2014)

10. Lassen, B., Kuhnigk, J.M., Schmidt, M., Krass, S., Peitgen, H.O.: Lung and lung lobe segmentation methods at Fraunhofer MEVIS. In: Proceedings of the Fourth International Workshop on Pulmonary Image Analysis, pp. 185–199 (2011)
11. Modersitzki, J.: FAIR: Flexible Algorithms for Image Registration. SIAM (2009)
12. Murphy, K., van Ginneken, B., Reinhardt, J.M., Kabus, S., Ding, K., et al.: Evaluation of registration methods on thoracic CT: the EMPIRE10 challenge. *IEEE Trans. Med. Imaging* **30**(11), 1901–1920 (2011)
13. Nocedal, J., Wright, S.: Numerical Optimization. Springer, New York (2006)
14. Regan, E.A., Hokanson, J.E., Murphy, J.R., Lynch, D.A., Beaty, T.H., et al.: Genetic Epidemiology of COPD (COPDGene) study design. *COPD* **7**, 32–43 (2011)
15. Risser, L., Vialard, F.X., Baluwala, H.Y., Schnabel, J.A.: Piecewise-diffeomorphic image registration: application to the motion estimation between 3D CT lung images with sliding conditions. *Med. Image Anal.* **17**(2), 182–193 (2013)
16. Rühaak, J., Heldmann, S., Kipshagen, T., Fischer, B.: Highly accurate fast lung CT registration. In: SPIE 2013, Medical Imaging, p. 86690Y-1-9 (2013)
17. Sakamoto, R., Mori, S., Miller, M.I., Okada, T., Togashi, K.: Detection of time-varying structures by large deformation diffeomorphic metric mapping to aid reading of high-resolution CT images of the lung. *PLoS ONE* **9**(1), 1–11 (2014)
18. Vialard, F.X., Risser, L., Rueckert, D., Cotter, C.J.: Diffeomorphic 3D image registration via geodesic shooting using an efficient adjoint calculation. *IJCV* **97**(2), 229–241 (2012)
19. Zhang, M., Fletcher, P.T.: Finite-dimensional Lie Algebras for fast diffeomorphic image registration. In: Ourselin, S., Alexander, D.C., Westin, C.-F., Cardoso, M.J. (eds.) IPMI 2015. LNCS, vol. 9123, pp. 249–260. Springer, Heidelberg (2015). doi:[10.1007/978-3-319-19992-4_19](https://doi.org/10.1007/978-3-319-19992-4_19)