RESEARCH ARTICLE

Collision-constrained deformable image registration framework for discontinuity management

Thomas Alscher ⓘ *, **Kenny Erleben** ⓘ *, **Sune Darkner** *

Department of Computer Science, University of Copenhagen, Copenhagen, Region Hovedstaden, Denmark

* alscher@di.ku.dk (TA); kenny@di.ku.dk (KE); darkner@di.ku.dk (SD)

Abstract

Topological changes like sliding motion, sources and sinks are a significant challenge in image registration. This work proposes the use of the alternating direction method of multipliers as a general framework for constraining the registration of separate objects with individual deformation fields from overlapping in image registration. This constraint is enforced by introducing a collision detection algorithm from the field of computer graphics which results in a robust divide and conquer optimization strategy using Free-Form Deformations. A series of experiments demonstrate that the proposed framework performs superior with regards to the combination of intersection prevention and image registration including synthetic examples containing complex displacement patterns. The results show compliance with the non-intersection constraints while simultaneously preventing a decrease in registration accuracy. Furthermore, the application of the proposed algorithm to the DIR-Lab data set demonstrates that the framework generalizes to real data by validating it on a lung registration problem.

Introduction

Inadequate management of discontinuities in the displacements field of image registration causes problems for widely-used algorithms utilizing smoothness regularizations [1]. One such instance is sliding motions along domain boundaries as observed in longitudinal registration of lungs. Algorithms producing a single smooth deformation field cannot account for discontinuities. A separation of the registration domain into its sliding segments is, thanks to increasingly well-performing deep learning methods for accurate individual organ segmentation [2–4], available but subsequent, independent registration of the resulting segments suffers from the possible overlap of the computed deformation fields. In order to acquire separate and congruent deformation fields without intersection, an additional constraining mechanism is required. To this end, the alternating direction method of multipliers (ADMM) is presented to enforce constraints preventing overlap of the individual components.

Extending the early development of rigid [5] registration algorithms, the implementation of diffusion inspired processes [6], called DEMONs, as well as non-rigid deformations introduced by Free-Form Deformations (FFD)[7] account for more complex image registrations.

Analogies from fluid mechanics allowed to consider large deformations [8, 9]. Topology preservation in image registration is enforced by the use of diffeomorphisms. The combination with the large deformation algorithms gave rise to another approach, called Large Deformation Diffeomorphic Metric Mapping (LDDMM)[10, 11]. One type of discontinuities, particularly present in lung and abdominal registrations, are sliding motions [12]. The sliding of the lung along the pleura and the expansion of the rib cage in the opposite direction leads to severe singularities in the deformation field and consequently invalidate the usual premise of smooth displacement fields [13].

Existing approaches aiming to deal with sliding motions can be categorized into methods computing a single deformation field for the whole image domain and those that split the domain into individual parts for each independently moving subdomain and subsequently computing as many deformation fields. Solutions computing one deformation field require additional constraints controlling smoothing across domain boundaries. One such approach of enriching the b-spline basis functions with additional information around discontinuities, i.e. in the lung case domain boundaries, is proposed in [14] and prevents information exchange over domain boundaries. Due to its built, the approach doesn't require further regularization or a specific interface smoothness. Being a completely intensity based method without explicitly penalizing overlaps, the image gradient may push voxel over the domain boundary regardless of the underlying enrichment. Further approaches utilizing the need for regularization differentiation between sub-domains can be divided into two groups [13]: direction-dependent methods and locally adaptive methods. Direction-dependent approaches decompose the deformation field into tangential and normal components. Smoothing the image globally only in the direction of the surface normals and tangentially only inside segmented domains offers a closed mathematical solution, with the draw back of stationary normal directions [15]. Using a piece wise diffeomorphic solution with a stationary velocity field implementation that smooths the velocity field domain independently in tangential direction while matching the velocity direction and magnitude on both sides of the domain interface is introduced in [16]. However, using different resolutions in different domains requires interpolation at the domain boundary effectively smoothing intensities. Locally adaptive approaches implement regularization whose weight in the loss function differs according to the location in the image domain. The first locally adaptive version investigated in this section offers to balance diffusion-based L_2 norm regularization for smooth domain interiors and L_1 norm total variation at domain boundaries, showing good results even with minor domain segmentation errors. Low magnitude updates, e.g. small scale sliding motion, can still be falsely smoothed [17]. [18] uses a classical bilateral filtering kernel to identify interface boundaries with the application of a computationally costly kernel. An alternative to bilateral filtering is isotropic total variation [1]. Both methods, bilateral filtering and total variation are challenged, however, by interfaces between similarly textured and low contrast domains. Registration algorithms splitting the image domain into independently moving sub-domains suffer more than the methods above from possible overlaps or gaps in between distinct deformation fields. [19] minimizes the overlaps by assigning unique, penalizing intensities for voxel lying outside the respective sub-domain during registration, allowing for the use of different registration parameters for each sub-domain. The definition of the right intensities can be tedious and time consuming. [20–22] drop intensity based penalization in favour if a distinct term added to the loss function penalizing overlaps and gaps. These terms can consist of the product of deformed signed distance fields [20], requiring a computationally expensive creation of motion mask and being restricted to two sub-domains, or a local distance metric to the opposing interface via the sum of sample surface points [22]. The second implementation offers a solution to sliding along curved surfaces

and surfaces that are separated by a third domain, but relies on the resolution of the sampled surface points. Another approach penalizing a set of sampled surface points is the split Gauss-Newton approach [21, 23]. Here any deviation of the point set from the deformed interface is weighted relative to the distance to the interface. If not parametrized properly, the penal term can overrule the registration forces, keeping the deformation field in its initial state. Using not n deformation fields for as many sub-domains, but $n + 1$, [24] utilizes the additional field to calculate the global normal displacements only. The normal direction is derived from local bases. The remaining deformation fields produce the sub-domain specific tangential displacements, ensuring matching normal displacements along the interface. This approach assumes smooth, small-deforming interfaces, since the bases are created in the moving image only. All aforementioned methods except [1, 18] require correct segmentations.

This work presents a registration framework, the Collision Constrained Deformable Image Registration (CC-DIR), to account for sliding motions by combining a collision constraining regularization with the registration term and solving the resulting minimization problem using the ADMM.

Materials and methods

Image registration

Formally, in image registration a transformation $\Phi: \Omega \to \Omega$ is sought that aims to align the source or moving image $I \in R^n$ via the transformation $\Phi \in \Omega$ to a target or fixed image $J \in R^n$. The registration problem can be written as a minimization problem of the energy term $F: \Omega \to R$

$$\min_{\Phi} F(\Phi, I, J) = \min_{\Phi} M(J, I \circ \Phi) + S(\Phi). \tag{1}$$

The image similarity $M: \Omega \to R$ measures the alignment between two images. Difference methods like Sum of Squared Differences (SSD) and Sum of Absolute Differences (SAD) are widely used as similarity measures for monomodal registrations that map the same anatomical structures, performing similarly well [25]. SSD methods penalize outliers in intensity heavier [26] and are in consequence a valid approach to properly register the relative high differences between lung tissue and the surrounding air. Image registration as an ill-posed problem has no single solution [27] and often utilizes regularization $S: \Omega \to R$. Regularization may be applied for smoothing [28] or constraining the deformation to reasonable [29] solutions. Cubic B-splines utilized in FFD act as a smoothing alternative themselves [28] and free the algorithm of additional regularization implementations.

Parametrization

Parametrizations offer a reduction in dimensionality, improved robustness [1], and decreased risk of overfitting. One example are Free-Form deformations. By deforming an overlying grid of control points, the deformation of the underlying image is calculated by interpolating the pixel- or voxelwise displacements. Widely used approximation functions for individual voxel displacements are cubic B-Splines, which have a very local influence and only affect points in their neighbourhood [7]. For an exemplary image in R^3, $\zeta \in R^{n_1 \times n_2 \times n_3}$ is the mesh control points spaced equidistantly at $\delta = (\delta_1, \delta_2, \delta_3)$. The displacement of $p \in R^3$

$$\Phi(p) = \sum_{l=0}^{3}\sum_{m=0}^{3}\sum_{q=0}^{3} B_l(u)B_m(v)B_q(w)\zeta_{i+l,j+m,k+q} \tag{2}$$

with the basis functions [30, 31]

$$B_0(u) = \frac{1}{6}(1-u)^3,$$ (3a)

$$B_1(u) = \frac{1}{6}(3u^3 - 6u^2 + 4),$$ (3b)

$$B_2(u) = \frac{1}{6}(-3u^3 + 3u^2 + 3u + 1),$$ (3c)

$$B_3(u) = \frac{1}{6}u^3.$$ (3d)

Location of the respective control points follows $i = \lfloor p_1/\delta_1 \rfloor - \xi, j = \lfloor p_2/\delta_2 \rfloor - \xi, k = \lfloor p_3/\delta_3 \rfloor - \xi$, with ξ shifting the support along the axis. The localized coordinates derive from $u = (p_1 - \delta_1{}^* \lfloor p_1/\delta_1 \rfloor)/\delta_1$, $v = (p_2 - \delta_2{}^* \lfloor p_2/\delta_2 \rfloor)/\delta_2$, and $w = (p_3 - \delta_3{}^* \lfloor p_3/\delta_3 \rfloor)/\delta_3$. By decreasing the grid resolution, cubic b-spline represent global deformations as well. Consequently, a multiresolution approach with sequentially increasing resolutions will capture both local and global deformations.

Collision detection

Collision detection algorithms can be classified according to utilized object representation [32], such as purely explicit, implicit or hybrid approaches. Explicit representations may range from simple point clouds to polygon meshes, one of the most common tools utilized [33]. Implicit representations, such as signed distance fields $SDF : \mathbb{R}^n \to \mathbb{R}$, describe geometries as a mapping function. A hybrid approach of a point cloud, representing the deformable object, with a rigid object, depicted as a SDF, has been proven to perform fast and accurately [34]. On top, this algorithm leverages existing data structures such as segmentation masks and can be directly pre-computed from segmented images by fast marching methods, allowing for fast intersection tests [33]. This sidesteps the need for object meshing, which proves hard to implement for anatomical structures [35]. Let $p_i \in \mathbb{R}^n$ be any point in the image and $\omega \subset \mathbb{R}^n$ by a closed and bounded [36] subset representing a collision object with $\delta\omega$ as its boundary. Detecting a collision of p_i with ω requires evaluating the signed distance

$$SDF_\omega(p_i) = \begin{cases} -d(p_i) & p_i \in \omega \\ 0 & p_i \in \delta\omega \\ d(p_i) & p_i \notin \omega \end{cases}$$ (4)

with $d : \mathbb{R}^n \to \mathbb{R}$

$$d(p_i) = \inf_{y \in \delta\omega} \| p_i - y \|_2$$ (5)

equating to the shortest distance to its domain boundary. The intersections are formulated as a quadratic penalty function $g : \mathbb{R}^n \to \mathbb{R}$ [37].

$$g(p_i) = \begin{cases} 0 & SDF_\omega(p_i) \geq 0 \\ \frac{\mu}{2} SDF_\omega(p_i)^2 & 0 \quad SDF_\omega(p_i) < 0 \end{cases}$$ (6)

This function weights possible collisions by the depth of intersection of colliding particles with respect to the object surface with an adjustable free parameter μ. Neglecting all influences of non-intersecting points by simple piece-wise definition g acts both as the collision detection and correction, with a value only deviating from 0 once p_i collides with ω.

Reformulating this collision constraint for the registration problem to account for sliding motion between the image domains $\omega_1 \subset I$ and its counterpart $\omega_2 \subset J$, let $X_i \in \omega_1$ be the set of all material points in the undeformed configuration. The energy formulation in the deformed configuration $x_i = X_i \circ \Phi$ reads as

$$C(\Phi, I, J) = \sum_i g(x_i) \tag{7}$$

Here, the SDF_{ω_2} is calculated in spatial space. Adding the collision constraint energy formulation to the initial registration problem changes the term to

$$\min_{\Phi} F(\Phi, I, J) + C(\Phi, I, J). \tag{8}$$

Alternating direction method of multipliers

With the constrained registration problem now consisting of $F(\Phi, I, J)$ and $C(\Phi, I, J)$, we can take advantage of this formulation by splitting the variable Φ representing the transformation into two distinct deformations Φ_R and Φ_C for the registration respectively collison problem and adding a coupling constraint.

$$\min_{\Phi_R, \Phi_C} F(\Phi_R, I, J) + C(\Phi_C, I, J) \tag{9}$$

$$\text{subject to } \Phi_R - \Phi_C = 0.$$

This formulation clearly illustrates the advantage of applying the ADMM. Utilizing a combination of dual ascent and decomposition, the ADMM splits the objective function into subproblems, favouring problems where the local optimization of theses sub-problems can be carried out efficiently [38]. The registration problems in the form of Eq 9 can be solved by iterating through the sequential updates of Φ_R, Φ_C, and a so-called dual variable u by

$$\Phi_R k + 1 \leftarrow \arg\min_{\Phi_R} F(\Phi_R, I, J) + \frac{\rho}{2} \left\| \Phi_R - \Phi_C^k + \mathbf{u}^k \right\|_2^2, \tag{10a}$$

$$\Phi_C k + 1 \leftarrow \arg\min_{\Phi_C} C(\Phi_C, I, J) + \frac{\rho}{2} \left\| \Phi_R^{k+1} - \Phi_C + \mathbf{u}^k \right\|_2^2, \tag{10b}$$

$$\mathbf{u}^{k+1} \leftarrow \mathbf{u}^k + \Phi_R^{k+1} - \Phi_C^{k+1}, \tag{10c}$$

until either convergence or a fixed number of iterations is achieved. This scaled version of the ADMM utilizes u as the accumulating sum of residuals for the coupling constraint [39] whereas the free parameter ρ allows to control its influence on the sequential updates Convergence towards an optimum is sufficiently fast for moderate accuracy, but may be slow to produce highly accurate results.

Experiments & results

2D synthetic examples under controlled conditions, as proof of concept, as well as an application to 3D medical data to showcase its feasibility, are conducted.

Implementation

For each domain ω_r representing an independently moving object in $I \in R^n$, in this paper in both synthetic and medical experiments $r \in \{1, 2\}$, a CC-DIR is executed. The object is represented as two point clouds of i respectively j evaluation points for $\mathbf{p}_{col}, \mathbf{p}_{reg}$ in the collision respectively registration problem. Cubic B-spline parametrization (Eq 2) is used as the deformation model, while for interpolating voxel intensities at non-integer positions, linear interpolation is utilized. For updating $\Phi_C^{N+1}, \mathbf{\Phi}_R^{N+1}$ during the iterations, subsolvers are called to minimize the respective function and assign the minimizing variables $\Phi_C, \mathbf{\Phi}_R$. The subsolvers are gradient descent implementations with a backtracking according to [40], running for 2 iterations. ρ is kept at 0.5 across all experiments. In our implementation the stopping criterion for CC-DIR is a fixed set of iterations N_{fixed}.

Algorithm 1 CC-DIR in 3D

```
Inputs:
    p_col, p_reg ∈ R^{i,j×3} with p_col, p_reg ⊂ ω_r  Φ_R,Φ_C ∈ R^{3×n_1×n_2×n_3}
Initialize:
    N ← 0
    Φ_R^0 ← 0_{3×n_1×n_2×n_3}
    Φ_C^0 ← 0_{3×n_1×n_2×n_3}  u^0 ← 0_{3×n_1×n_2×n_3}
while N ≤ N_{fixed} do
    Φ_R^{N+1} ← arg min_{Φ_R} F(Φ_R, I, J) + ρ/2 ||Φ_R − Φ_C^N + u^N||_2^2
    Φ_C^{N+1} ← arg min_{Φ_C} C(Φ_C, I, J) + ρ/2 ||Φ_R^{N+1} − Φ_C + u^N||_2^2
    u^{N+1} ← u^N + Φ_R^{N+1} − Φ_C^{N+1}
end while
```

Synthetic experiments

In order to show the improvements in registration quality by the CC-DIR compared to methods without collision constraining, the synthetic experiments include a Baseline model, a modified Baseline (mod. Baseline) model for improved performance and the CC-DIR. The Baseline model consists of a B-spline FFD registration model with one deformation field and a gradient descent with backtracking as optimization method. The mod. Baseline uses the same setup, however running two FFD registrations consecutively, one for each image domain. The CC-DIR computes one deformation field for each image domain as well. However, CC-DIR uses Alg. 1 to constrain collisions between the two domains. Differences in algorithms between Baseline, mod. Baseline and CC-DIR are kept to a minimum by using the same similarity measure SSD. A small parameter study has been performed for every experiment to identify the best outcome for the synthetic setup with respect to control point grid scale, learning rate, and penalty parameter.

Data. Three simplified cases in 2D are computed for a better understanding and visualization of the proposed method, while combining different discontinuities in their respective setups. The image domain is divided by the boundary $\delta\Omega$ into Ω_L on the left side and Ω_R on the right, each containing the distinct structure I_L respectively I_R. In the first case, *Linear*, the boundary $\delta\Omega$ moves as a whole to the right-hand side while I_L and I_R are sliding in opposite directions up, respectively down. Additionally, I_L deforms under conservation of its total area. The second case, *Non − linear*, introduces a more complex, non-linear deformation of the boundary $\delta\Omega$. Both structures are subject to deformation and move along the boundary in contrary directions, again sliding along the boundary. The last case *Growth* combines deformation, translation and the sliding motion discontinuity from the second example with a growth Ω_G in the boundary region (Fig 1).
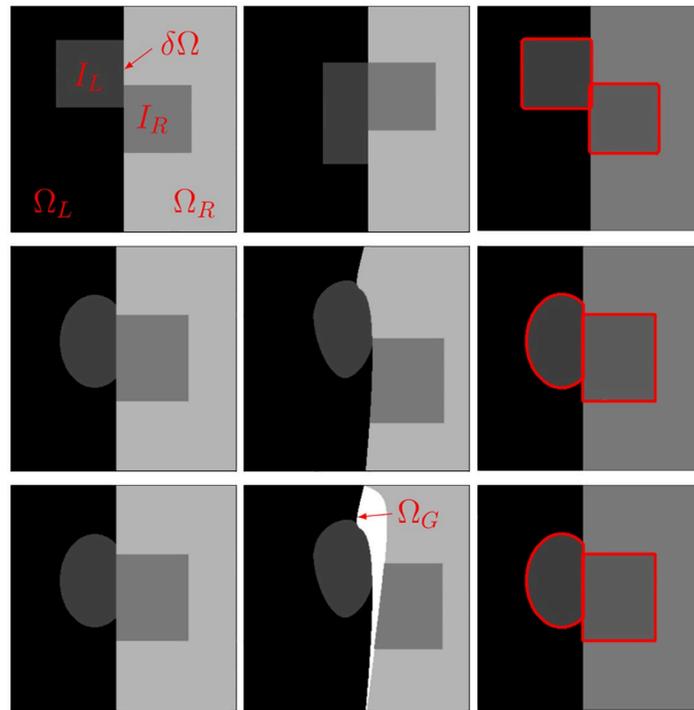
Fig 1. Synthetic experiments setup and results. From left to right: Source, target and pullback-registered CC-DIR images for Linear (Top row), Nonlinear (Center row) and Growth (Bottom row) case. The registered images show clear matching of the shapes I_L and I_R.

https://doi.org/10.1371/journal.pone.0290243.g001

Results. Two main criteria are quantified and evaluated: registration success and collision constraint violation. The registration success is easily calculated with the Sørensen-Dice index (DICE) for matching the structures I_L and I_R. For the evaluation of the collision constraint the intersection scores (IS) is introduced. The percentage of intersecting pixels with respect to the image area constitutes to IS, aiming to giving an qualitative overview of the intersection. For all three cases, the CC-DIR performs best in terms of registration as well as preventing intersection (Table 1). The DICE score indicates a high success in matching the structures, backed by a clean pullback registration visualization in Fig 1. No pixel intersects with the boundaries, so no violation of the collision constraint is detected. The mod. Baseline ranks in second, with satisfying registration results in all but the *Growth* case. Regarding intersections, this approach

Table 1. Intersection (IS) and DICE score for synthetic experiments. The purposed CC-DIR based algorithm shows no overlap in the three test cases additionally with a superior DICE score.

Experiment	Mode	IS [%]	DICE
Linear	Baseline	0.15	0.954
	mod. Baseline	**0.0**	0.976
	CC-DIR	**0.0**	**1.0**
Non-linear	Baseline	0.2	0.98
	mod. Baseline	0.14	0.984
	CC-DIR	**0.0**	**0.998**
Growth	Baseline	3.48	0.843
	mod. Baseline	3.44	0.876
	CC-DIR	**0.0**	**0.999**

https://doi.org/10.1371/journal.pone.0290243.t001

performs similarly well for the *Linear* case as the CC-DIR, but shows more intersections with increasing non-linearities in the boundaries. Losing registration performance quickly, the Baseline approach struggles to match the structures already in the *Linear* case. Additionally, intersections are reported for every case, hence violating the collision constraints consistently.

Medical experiments

The CC-DIR for evaluation on a medical data set is run as described in Alg. 1. Two deformation fields are computed consecutively and independently. The two independent domains of $I \in R^3$ are the lung domain $\Omega_{Lung} \subset I$ defined by a segmentation and the rib cage $\Omega_{Rib} = I \backslash \Omega_{Lung}$. The multiresolution procedure is derived from the findings in the EMPIRE challenge [41]. For the best possible results, a possible initial affine registration, followed by a multiresolution CC-DIR registration is run. The multiresolution setup consists of multiple successive CC-DIR registrations, with a coarse-to-fine control point grid adaption. The successively run CC-DIR compute updated starting coordinates for p_{col}, p_{reg}, handing these down to the next finer resolution CC-DIR. Image similarity is measured by SSD. Mirroring the synthetic experiments, the registration success and collision constraint violation are evaluated. The results are compared to a set of algorithms also tackling the sliding motion along the lung-ribcage interface [14].

Data. For the medical test setup, the DIR-LAB 4DCT data set from [42] is used with previously segmented lung masks from the Continuous Registration Challenge [43]. The registration is computed between the extreme inhale-exhale image pairs in order to ensure high deformations and sliding motions. To exclude regions with little image inforamtions, the images are cropped to exclude inmaging artifacts at the boundary regions. All CC-DIRs are run for 100 iterations.

Parameters. Even though consisting of only 10 lung scans, the DIR-Lab data set contains a variety of motion patterns, requireing individually tuned parameters. Hence, initial affine registrations were run for cases 5-10 as a gradient descent with backtracking for 100 iterations and SSD for similarity measure. The multiresolution scales in mm for the (x,y,z) axis for cases 1, 6, 7, 9, and 10 were (19.4,19.4,50), (9.7,9.7,25), and (4.85,4.85,12.5). Cases 2 had scales (23.2,23.2,50), (11.6,11.6,25), and (5.8,5.8,12.5). Case 3 had scales (23,23,50), (11.5,11.5,25), and (5.75,5.75,12.5). Case 4 used (22.6,22.6,50), (11.3,11.3,25), (5.65,5.65,12.5) as scales and case 5 (22,22,50), (11,11,25), (5.5,5.5,12.5). And finally case 8 with (38.8,38.8,100), (19.4,19.4,50), (9.7,9.7,25), and (4.85,4.85,12.5).

Target registration error. For the registration accuracy, anatomical landmarks are used. The data-set provides 300 annotated landmarks in both extreme phases. Consequently the summed position error after registration, called target registration error (TRE) over all landmarks indicates the registration outcome (Table 2). Introduced in the original data set publication [1, 20, 42], a snap-to-voxel evaluation for the TRE is used. The proposed CC-DIR algorithm's overall TRE accuracy ranks at fourth place, while matching the best benchmark in cases 1 and 2.

A coronal slice of a 3D CT lung scan in Fig 2 shows an example of the registration process. The difference image between the moving and target configuration before and after registration depict a reduction of structural differences. The deformation grid overlaid over the registered image in Fig 3 is the result of the CC-DIR. This close-up shows a clear solution to the sliding interfaces along the pleura as well as a translation in the diaphragm caused by the inhalation.

Congruent interfaces. For the evaluation of congruent, and thus physiologically plausible, rib-cage-lung interfaces, the metric introduced in [24] is used. After transforming a 3D surface

Table 2. Comparison of TREs for the DIR-LAB data set in mm. The mean TRE of the purposed CC-DIR method ranks at fourth place out of seven with an offset of 0.15 mm to the best performing algorithm.

Case	Wu (2008)	Delmon (2013)	Berendsen (2014)	Hua (2017)	Eiben (2018)	Gong (2020)	Proposed CC-DIR
1	1.1 ±0.5	1.2 ±0.6	1.0 ±0.52	1.0 ±0.51	N/A ±N/A	0.97 ±N/A	**0.92±0.97**
2	1.0 ±0.5	1.1 ±0.6	1.02 ±0.57	0.99 ±0.59	N/A ±N/A	1.02 ±N/A	**0.9±1.14**
3	1.3 ±0.7	1.6 ±0.9	1.14 ±0.89	**1.12±0.64**	N/A ±N/A	1.18 ±N/A	1.23 ±1.43
4	1.5 ±1.0	1.6 ±1.1	1.46 ±0.96	1.44 ±1.03	N/A ±N/A	**1.38±N/A**	1.49 ±1.41
5	1.9 ±1.5	2.0 ±1.6	1.61 ±1.48	**1.37±1.35**	N/A ±N/A	1.45 ±N/A	1.9 ±2.48
6	1.6 ±0.9	1.7 ±1.0	1.42 ±0.89	1.26 ±1.04	N/A ±N/A	**1.12±N/A**	1.46 ±1.41
7	1.7 ±1.1	1.9 ±1.2	1.49 ±1.06	**1.12±0.67**	N/A ±N/A	1.24 ±N/A	1.42 ±1.43
8	1.6 ±1.4	2.2 ±2.3	1.62 ±1.71	**1.18±1.22**	N/A ±N/A	1.71 ±N/A	1.26 ±1.17
9	1.4 ±0.8	1.6 ±0.9	1.3 ±0.76	**1.14±0.64**	N/A ±N/A	1.24 ±N/A	1.2 ±1.06
10	1.6 ±1.2	1.7 ±1.2	1.5 ±1.31	**1.08±0.82**	N/A ±N/A	1.24 ±N/A	1.45 ±1.72
Mean	1.47 ±0.96	1.66 ±1.14	1.36 ±0.99	**1.17±0.85**	1.21 ±N/A	1.25 ±N/A	1.32 ±1.42

mesh of the rib-cage-lung boundary $\delta\Omega$ by both displacement fields Φ_{Lung} and Φ_{Rib}, the subsequent gap and overlap between the two domains is calculated by voxelizing the respective meshes. The results in cm^3 are recorded in Table 3.

Overall, and in half of the cases specifically, the novel CC-DIR registration shows the least amount of overlap. The average resulting gap between the two domains ranks second after the interface matching algorithm [20].

Shear. In order to estimate the shear along the domain interfaces, the maximum shear stretch introduced in [44] is calculated. Using the deformation gradient tensor $F = \frac{dx}{dX}$ which maps the transformation from the undeformed configuration x to the deformed state X, the
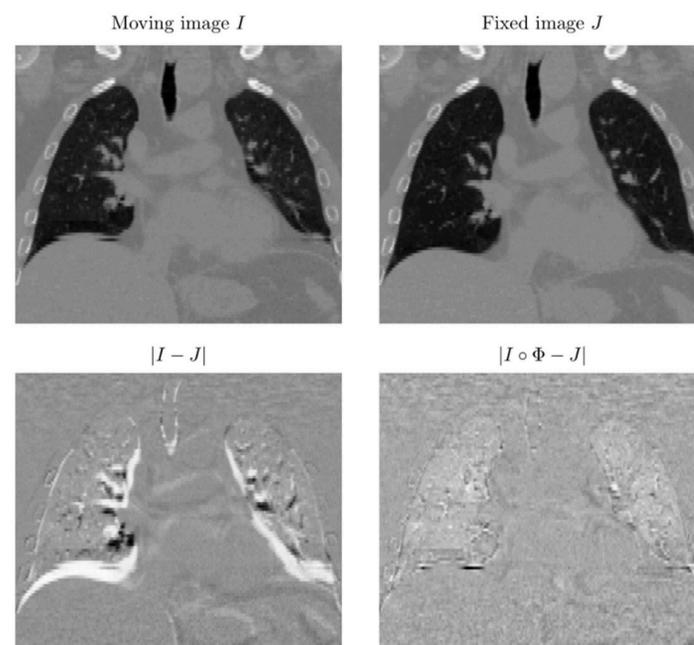

Fig 2. Overview of the reference, source and registered image. $|I - J|$ shows the intensity differences between moving and fixed image without registration, $|I \circ \Phi - J|$ after collision-based registration. A clear reduction of intensity differences and matching of anatomical structures can be observed.
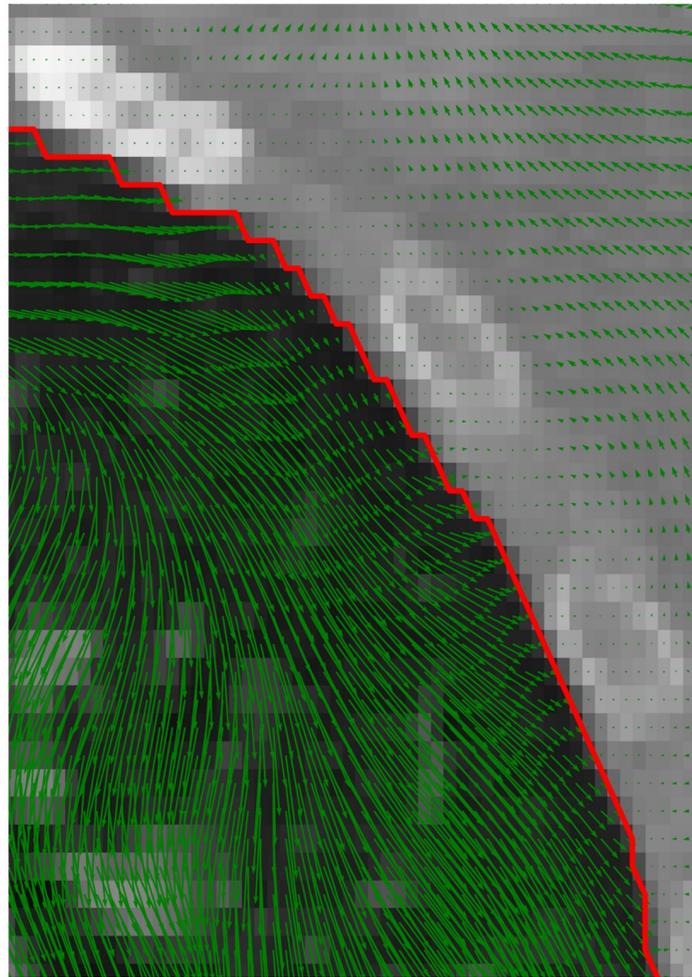
Fig 3. Close-up of the displacement field. Displacements around the interface (green arrows) show observable sliding motion along the boundary (red).

https://doi.org/10.1371/journal.pone.0290243.g003

Table 3. Comparison of gap/overlap for the DIR-LAB data set by case. The purposed CC-DIR algorithm shows the least overlap compared to the six other methods. It rates second in terms of interface gap with a mean offset of 38.1 cm^3 to the first place. All measurements are in cm^3.

Case	Wu (2008)	Delmon (2013)	Berendsen (2014)	Hua (2017)	Eiben (2018)	Gong (2020)	Proposed CC-DIR
1	38/26	39/15	23/18	39/**2**	N/A/N/A	67/69	**22**/2
2	78/46	67/60	74/34	74/31	N/A/N/A	76/94	**53/6**
3	99/28	83/33	57/30	71/24	N/A/N/A	59/70	**53/7**
4	75/34	66/44	66/28	92/13	N/A/N/A	71/79	**41/7**
5	**11**/38	78/52	61/32	54/**6**	N/A/N/A	97/109	37/11
6	**10**/86	119/77	130/50	155/**11**	N/A/N/A	57/69	90/23
7	**10**/79	108/77	119/45	138/19	N/A/N/A	73/88	131/**14**
8	96/91	92/93	85/53	150/40	N/A/N/A	**80**/98	92/**39**
9	61/34	**54**/44	70/51	58/14	N/A/N/A	63/72	78/**8**
10	120/63	94/56	80/43	109/28	N/A/N/A	90/110	**48/19**
Mean	88.2/52.5	80.0/55.1	76.5/37.4	94.0/18.8	**26.4**/34.5	73.3/85.8	64.5/**13.6**

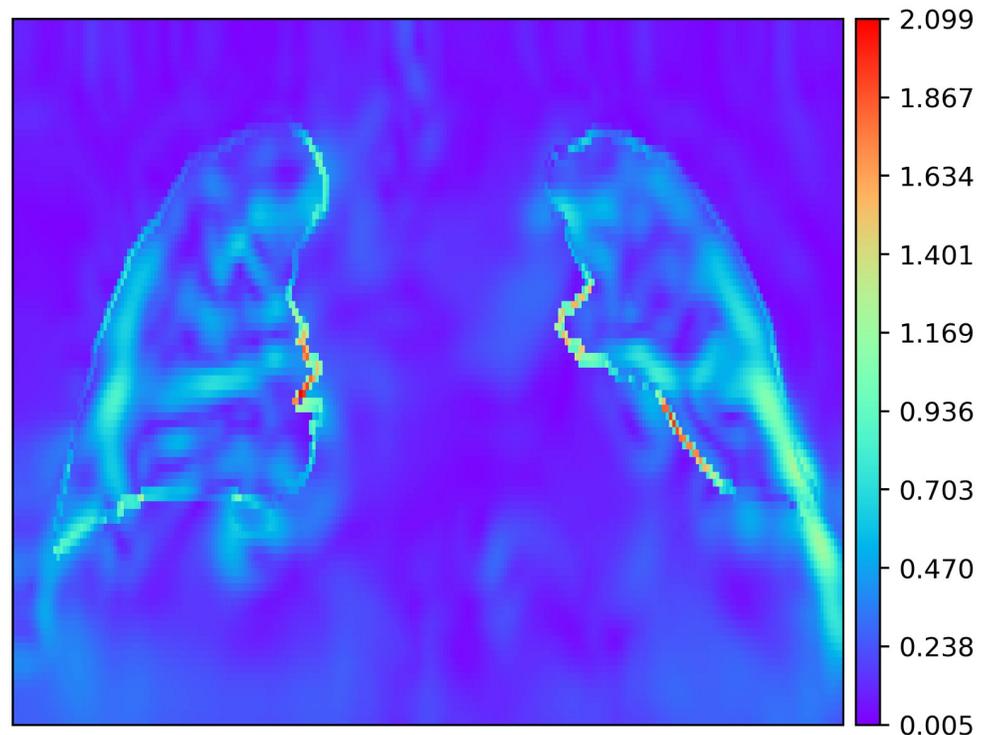https://doi.org/10.1371/journal.pone.0290243.t003

Fig 4. Shear stretch evaluation. High differences in shear stretches along lung-rib-cage interface with respect to the rest of the image are depicted.

https://doi.org/10.1371/journal.pone.0290243.g004

maximum shear stretch constitutes to

$$\lambda_{max} = \frac{\lambda_1 - \lambda_3}{2} \tag{11}$$

with

$$\lambda_i = \sqrt{eigenvalues\ of\ F^T F}. \tag{12}$$

The exemplary colour-coded representation Fig 4 clearly shows an increase in shear stretches along the interface, whereas the displacement field in Fig 3 demonstrates sliding motion. Overall the mean shear stretches amount to 1.3 ± 1.8 between the lung and rib-cage, respectively 0.2 ± 0.2 in the rest of the image. The computation time averages at 822s ± 519s on a Xeon CPU @ 3.60GHz x12 with an GPU implementation on a GeForce RTX 3090.

Discussion

Coupling deformable image registration with collision detection offers up new possibilities for supplying registration algorithms with simple physics. In order to incorporate these physics into the common registration problem formulation, an energy term (Eq 7) is needed. This setup simulates the deformation of the distinct image domains, such as organs, by calculating separate deformation fields. Discontinuities at the domain boundaries, such as sliding motion don't need to be modelled explicitly but result automatically by the collision forces minimizing the non-intersection constraint. Setting up the collision detection requires a domain representation, either as segmentation masks, surface or volumetric meshes. Acquiring these representations can be a difficult task, especially for regions with no discernible or low-contrast domain

boundaries. A problem encountered in the experiments is the lack of image information at the domain boundaries. If registration is strictly applied to the respective domain masks, the resulting gap might increase as image information outside of the sub-domain is missing. Dilating the masks to increase the tracking image forces provided an acceptable alternative. This way extensive domain overlap post-registration is prevented by the collision constraints, whereas possible gaps between the domains are closed by the image forces. Compared to single deformation field methods, CC-DIR allows to easily use different parametrizations and parameters for different sub-domains. Since no costly on the run identification of the domain boundary, as in filter based methods, has to be performed, the possibility of surface errors in low contrast regions is avoided. Not relying on normal and tangential decomposition also permits solving sliding motions along complex and non-smooth surfaces. Unless other mentioned methods registering subdomains independently using intensity based constraints, collision detection is geometric and thus modality independent, allowing the collision detection and correction to be used in different settings. Since collision detection is a fundamental area in computer graphics, it can fall back on a considerable body of research and a variety of algorithms, permitting an exchange of the collision detection method, as long as a respective energy term is formulated. An advantage of using collision detection over interface matching methods is the circumvention of redundant constraints, letting the registration energy term drive the deformation. Additionally, collision detection can be applied to scenarios where subdomains are moving freely around the image domain, with no restriction to constant contact between the subdomains, such as in computer vision or robotics. The implementation of the collision constraint in the ADMM has two advantages: First, the modularity of this optimization scheme permits not only the exchange of the collision method but also of the registration method specifications such as similarity metric or deformation model at will. Additional and existing regularization models can be freely and effectively added to the registration energy term, subsequently offering sliding motion representation at no further cost. Second, since the penalty function 6 is not bound to be differentiable, the ADMM still converges [39] and avoids performance issues of gradient-based methods [1]. However, as with many optimization strategies, convexity is of critical importance. First, the convergence of the ADMM is formally mathematically proven for convex functions only. Second, the ADMM solves the minimization for the dual problem instead of the initial primal problem. If strong duality holds, usually for convex functions [45], the dual solution is the primal solution [39]. The obtained solutions from nonlinear and non-convex problems, such as registration, might offer a lower bound and if converged, the ADMM still offers local solutions [39]. Furthermore, with additional regularization, the convexity of the objective functions might be increased and thus reduce the duality gap.

Looking at the results from the experiments in detail, more practical observations can be made. The synthetic experiment setups were chosen to analytically examine the possible improvements of the CC-DIR with collision detection compared to registration algorithms without collision detection. The results from the synthetic experiments show with stark contrast the benefits of the collision detection modeled via the CC-DIR compared to the Baseline approaches. The CC-DIR outperforms the Baseline approaches not only in terms of intersection constraints but also with respect to registration. Additionally to enforcing the non-intersection constraint, the collision detection acts as a powerful regularization term. Since the registration loss functions are likely to be non-convex, this regularization might increase convexity at least locally if not globally. Regarding the intersection constraints, the CC-DIR delivers deformation fields with no intersection at all, while the Baseline approaches have difficulties adhering to these boundaries.

Looking at the registration success for the medical data set measured by the TREs, this algorithm ranks in fourth out of seven, with an average distance on 0.15 mm to the best

ranking algorithm. However, with the finest voxel resolution of 0.97mm in the DIR-Lab data set, these differences can be regarded as minimal and may be caused by the image discretization. Furthermore, due to its advantageous framework build, replacing the registration methods is not only simple, but any improvement to the registration methods will also likely increase the TRE accuracy. With the least overlap and the second lowest gap between the domains compared to other algorithms, the CC-DIR produces separate yet congruent deformation fields that minimize physiologically not feasible transformations. Even though the CC-DIR doesn't prevent overlapping completely, two adaptions may offer an improvement: individually tuned hyper-parameters and an alternative constraint formulation. The implementation as quadratic penalty functions has the possible disadvantage of producing inexact solutions as the iterates may be drawn to points that violate the equality constraints but satisfy optimality conditions [37]. Exchanging the quadratic penalty function with an exact penalty function might yield the desired solution, however the resulting non-smoothness should be taken into account.

Conclusion

The CC-DIR successfully introduces a collision constrained method into the field of image registration. With its simple framework build, it allows to combine state-of-the-art registration methods with collision detection for discontinuity preservation. No major drawback in terms of registration quality have been observed, thus providing an effective method to account for sliding motion while simultaneously ensuring congruent interfaces. Inspired by the results in this paper, a future development of this framework is the exchange of the registration method with diffeomorphic counterparts. Looking at image registrations of the liver or prostate, which frequently utilize the Finite-Element-Method to simulate deformations [46], an image driven pathway to simulations may open up new possible developments. The advantageous movement parametrization based on basis functions can be used to easily derive FE models. A reformulation with tetrahedral interpolation and movement parametrization coupled with collision detection between tetrahedral meshes and deformable signed distance fields can even boost this process to fuse image registration with bio-mechanical analysis.

Author Contributions

Conceptualization: Thomas Alscher, Kenny Erleben, Sune Darkner.

Formal analysis: Thomas Alscher, Kenny Erleben, Sune Darkner.

Investigation: Thomas Alscher.

Methodology: Thomas Alscher, Sune Darkner.

Project administration: Thomas Alscher.

Software: Thomas Alscher, Sune Darkner.

Supervision: Kenny Erleben, Sune Darkner.

Validation: Thomas Alscher.

Visualization: Thomas Alscher.

Writing – original draft: Thomas Alscher.

Writing – review & editing: Kenny Erleben, Sune Darkner.

References

1. Vishnevskiy V, Gass T, Szekely G, Tanner C, Goksel O. Isotropic total variation regularization of displacements in parametric image registration. IEEE transactions on medical imaging. 2016; 36(2):385–395. https://doi.org/10.1109/TMI.2016.2610583 PMID: 27654322

2. Pandey B, Pandey DK, Mishra BP, Rhmann W. A comprehensive survey of deep learning in the field of medical imaging and medical natural language processing: Challenges and research directions. Journal of King Saud University-Computer and Information Sciences. 2021;.

3. Minaee S, Boykov YY, Porikli F, Plaza AJ, Kehtarnavaz N, Terzopoulos D. Image segmentation using deep learning: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence. 2021;.

4. Tajbakhsh N, Jeyaseelan L, Li Q, Chiang JN, Wu Z, Ding X. Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation. Medical Image Analysis. 2020; 63:101693. https://doi.org/10.1016/j.media.2020.101693 PMID: 32289663

5. Jiang H, Holton KS, Robb RA. Image registration of multimodality 3D medical images by chamfer matching. In: Acharya RS, Cogswell CJ, Goldgof DB, editors. Biomedical Image Processing and Three-Dimensional Microscopy. vol. 1660. International Society for Optics and Photonics. SPIE; 1992. p. 356–366.

6. Thirion JP. Image matching as a diffusion process: an analogy with Maxwell's demons. Medical Image Analysis. 1998; 2(3):243–260. https://doi.org/10.1016/S1361-8415(98)80022-4 PMID: 9873902

7. Rueckert D, Sonoda LI, Hayes C, Hill DL, Leach MO, Hawkes DJ. Nonrigid registration using free-form deformations: application to breast MR images. IEEE transactions on medical imaging. 1999; 18(8):712–721. https://doi.org/10.1109/42.796284 PMID: 10534053

8. Christensen GE, Rabbitt RD, Miller MI. Deformable templates using large deformation kinematics. IEEE transactions on image processing. 1996; 5(10):1435–1447. https://doi.org/10.1109/83.536892 PMID: 18290061

9. Bro-Nielsen M, Gramkow C. Fast fluid registration of medical images. In: International Conference on Visualization in Biomedical Computing. Springer; 1996. p. 265–276.

10. Beg MF, Miller MI, Trouvé A, Younes L. Computing large deformation metric mappings via geodesic flows of diffeomorphisms. International journal of computer vision. 2005; 61(2):139–157. https://doi.org/10.1023/B:VISI.0000043755.93987.aa

11. Avants BB, Schoenemann PT, Gee JC. Lagrangian frame diffeomorphic image registration: Morphometric comparison of human and chimpanzee cortex. Medical Image Analysis. 2006; 10(3):397–412. https://doi.org/10.1016/j.media.2005.03.005 PMID: 15948659

12. Kabus S, Franz A, Fischer B. Variational image registration with local properties. In: International Workshop on Biomedical Image Registration. Springer; 2006. p. 92–100.

13. Heinrich MP, Papież BW. Chapter 13—Image registration with sliding motion. In: Zhou SK, Rueckert D, Fichtinger G, editors. Handbook of Medical Image Computing and Computer Assisted Intervention. The Elsevier and MICCAI Society Book Series. Academic Press; 2020. p. 293–318.

14. Hua R, Pozo JM, Taylor ZA, Frangi AF. Multiresolution eXtended Free-Form Deformations (XFFD) for non-rigid registration with discontinuous transforms. Medical Image Analysis. 2017; 36:113–122. https://doi.org/10.1016/j.media.2016.10.008 PMID: 27894001

15. Schmidt-Richberg A, Werner R, Handels H, Ehrhardt J. Estimation of slipping organ motion by registration with direction-dependent regularization. Medical Image Analysis. 2012; 16(1):150–159. https://doi.org/10.1016/j.media.2011.06.007 PMID: 21764627

16. Risser L, Vialard FX, Baluwala HY, Schnabel JA. Piecewise-diffeomorphic image registration: Application to the motion estimation between 3D CT lung images with sliding conditions. Medical Image Analysis. 2013; 17(2):182–193. https://doi.org/10.1016/j.media.2012.10.001 PMID: 23177000

17. Gong L, Duan L, Dai Y, He Q, Zuo S, Fu T, et al. Locally adaptive total p-variation regularization for non-rigid image registration with sliding motion. IEEE Transactions on Biomedical Engineering. 2020; 67(9):2560–2571. https://doi.org/10.1109/TBME.2020.2964695 PMID: 31940514

18. Papież BW, Heinrich MP, Fehrenbach J, Risser L, Schnabel JA. An implicit sliding-motion preserving regularisation via bilateral filtering for deformable image registration. Medical image analysis. 2014; 18(8):1299–1311. https://doi.org/10.1016/j.media.2014.05.005 PMID: 24968741

19. Wu Z, Rietzel E, Boldea V, Sarrut D, Sharp GC. Evaluation of deformable registration of patient lung 4DCT with subanatomical region segmentations. Medical physics. 2008; 35(2):775–781. https://doi.org/10.1118/1.2828378 PMID: 18383700

20. Eiben B, Tran EH, Menten MJ, Oelfke U, Hawkes DJ, McClelland JR. Statistical motion mask and sliding registration. In: Biomedical Image Registration: 8th International Workshop, WBIR 2018, Leiden, The Netherlands, June 28-29, 2018, Proceedings 8. Springer; 2018. p. 13–23.

21. Heldmann S, Polzin T, Derksen A, Berkels B. An image registration framework for sliding motion with piecewise smooth deformations. In: International Conference on Scale Space and Variational Methods in Computer Vision. Springer; 2015. p. 335–347.

22. Berendsen FF, Kotte AN, Viergever MA, Pluim JP. Registration of organs with sliding interfaces and changing topologies. In: Medical Imaging 2014: Image Processing. vol. 9034. SPIE; 2014. p. 95–101.

23. Derksen A, Heldmann S, Polzin T, Berkels B. Image registration with sliding motion constraints for 4D CT motion correction. In: Bildverarbeitung für die Medizin 2015. Springer; 2015. p. 335–340.

24. Delmon V, Rit S, Pinho R, Sarrut D. Registration of sliding objects using direction dependent B-splines decomposition. Physics in Medicine & Biology. 2013; 58(5):1303. https://doi.org/10.1088/0031-9155/58/5/1303 PMID: 23388109

25. Ulysses J, Conci A. Measuring similarity in medical registration. In: IWSSIP 17th Int. Conf. on Systems, Signals and Image Processing; 2010.

26. Hajnal JV, Hill DL. Medical image registration. CRC press; 2001.

27. Sotiras A, Davatzikos C, Paragios N. Deformable Medical Image Registration: A Survey. IEEE Transactions on Medical Imaging. 2013; 32(7):1153–1190. https://doi.org/10.1109/TMI.2013.2265603 PMID: 23739795

28. Tustison N, Avants B. Explicit B-spline regularization in diffeomorphic image registration. Frontiers in Neuroinformatics. 2013; 7:39. https://doi.org/10.3389/fninf.2013.00039 PMID: 24409140

29. Burger M, Modersitzki J, Ruthotto L. A hyperelastic regularization energy for image registration. SIAM Journal on Scientific Computing. 2013; 35(1):B132–B148. https://doi.org/10.1137/110835955

30. Bankman I. Handbook of medical image processing and analysis. Elsevier; 2008.

31. Lee S, Wolberg G, Shin SY. Scattered data interpolation with multilevel B-splines. IEEE transactions on visualization and computer graphics. 1997; 3(3):228–244. https://doi.org/10.1109/2945.620490

32. Lin M, Gottschalk S. Collision detection between geometric models: A survey. In: Proc. of IMA conference on mathematics of surfaces. vol. 1. Citeseer; 1998. p. 602–608.

33. Ericson C. Real-time collision detection. Crc Press; 2004.

34. Fuhrmann A, Sobotka G, Groß C. Distance fields for rapid collision detection in physically based modeling. In: Proceedings of GraphiCon. vol. 2003. Citeseer; 2003. p. 58–65.

35. Teo JCM, Chui CK, Wang ZL, Ong SH, Yan CH, Wang SC, et al. Heterogeneous meshing and biomechanical modeling of human spine. Medical Engineering & Physics. 2007; 29(2):277–290. https://doi.org/10.1016/j.medengphy.2006.02.012 PMID: 16679044

36. Chan T, Zhu W. Level set based shape prior segmentation. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). vol. 2. IEEE; 2005. p. 1164–1170.

37. Wright S, Nocedal J, et al. Numerical optimization. Springer Science. 1999; 35(67-68):7.

38. Wei E, Ozdaglar A. Distributed alternating direction method of multipliers. In: 2012 IEEE 51st IEEE Conference on Decision and Control (CDC). IEEE; 2012. p. 5445–5450.

39. Boyd S, Parikh N, Chu E. Distributed optimization and statistical learning via the alternating direction method of multipliers. Now Publishers Inc; 2011.

40. Armijo L. Minimization of functions having Lipschitz continuous first partial derivatives. Pacific Journal of mathematics. 1966; 16(1):1–3. https://doi.org/10.2140/pjm.1966.16.1

41. Murphy K, van Ginneken B, Reinhardt JM, Kabus. Evaluation of Registration Methods on Thoracic CT: The EMPIRE10 Challenge. IEEE Transactions on Medical Imaging. 2011; 30(11):1901–1920. https://doi.org/10.1109/TMI.2011.2158349 PMID: 21632295

42. Castillo R, Castillo E, Guerra R, Johnson VE, McPhail T, Garg AK, et al. A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets. Physics in Medicine and Biology. 2009; 54(7):1849–1870. https://doi.org/10.1088/0031-9155/54/7/001 PMID: 19265208

43. Niessen W, Mendrik A, Murphy K, Modat M, Jef V;. https://continuousregistration.grand-challenge.org/.

44. Amelon RE, Cao K, Reinhardt JM, Christensen GE, Raghavan ML. A measure for characterizing sliding on lung boundaries. Annals of biomedical engineering. 2014; 42(3):642–650. https://doi.org/10.1007/s10439-013-0920-5 PMID: 24114112

45. Boyd S, Boyd SP, Vandenberghe L. Convex optimization. Cambridge university press; 2004.

46. Goksel O, Vishnevsky V, Carrillo AG, Tanner C. Imaging of sliding visceral interfaces during breathing. In: 2016 IEEE 13th International Symposium on Biomedical Imaging (ISBI); 2016. p. 298–301.