IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received 30 May 2025, accepted 27 June 2025, date of publication 14 July 2025, date of current version 12 August 2025.

Digital Object Identifier 10.1109/ACCESS.2025.3588760

RESEARCH ARTICLE

Collision Constraints and Diffeomorphisms: Dealing With Physics in Deformable Image Registration

THOMAS ALSCHER, JENS PETERSEN, FRANÇOIS LAUZE, KENNY ERLEBEN, AND SUNE DARKNER

Department of Computer Science, University of Copenhagen, 2100 Copenhagen, Denmark

Corresponding author: Thomas Alscher (alscher@di.ku.dk)

This work was supported by the European Union's Horizon 2020 Research and Innovation Program through the Marie Skłodowska-Curie under Grant 764644.

ABSTRACT Discontinuous displacements pose a substantial challenge to image registration algorithms, which generally assume smooth deformations. Two prominent cases occur in 4DCT lung scans. The first instance exists at the lung-thorax-boundary, where due to high elasticities and expansion properties, lungs exhibit higher displacements than the surrounding thoracic cavity, leading to sliding motions along their interface. Discontinuous motion is facilitated through a piecewise registration framework, in which inter-domain boundary interactions are constrained via collision detection, and intra-domain deformation smoothness is governed by the Large Deformation Diffeomorphic Metric Mapping (LDDMM) model. The second case with discontinuous displacements is found between lung lobes along the separating lung fissures. A fissure reconstruction algorithm for sparse segmentation is presented, improving geometric lung representation. Utilizing a refined Free-Form Deformation model for lobe-on-lobe registration that permits lobar sliding and prevents self-intersection, this approach examines discontinuities around the pulmonary fissures. Both cases are evaluated on synthetic and medical data, with boundary inconsistency results and average registration metrics on the DIRLAB dataset comparable to other approaches addressing sliding motion.

INDEX TERMS Deformable image registration, LDDMM, discontinuities, collision detection.

I. INTRODUCTION

Medical imaging techniques, as a non-invasive means to visualize anatomical and functional aspects of the human body [1], have grown to become an important part of modern medicine. They contribute to medical education [2] and research [3] as well as supporting clinical professionals in reaching diagnoses and planning treatments. Since not only analysis of individual images but comparisons between two images are often needed, these images require a common domain to be studied in.

Image registration is the process of finding an alignment of two images via geometrical transformations to this common domain, while preserving and matching corresponding similarities [4]. Commonly only one, moving, image is deformed while the other, fixed, image is locked. Practical applications for image registration in the medical field incorporate atlas creation [5], surgery image guidance [6], and image fusion [7], among others. In order to structure and select appropriate registration solutions, multiple categorization systems have been described [8], [9]. Condensing these extensive works, algorithms can be classified by the region-of-interest (ROI) in the images, the registration objective, and the method employed. Being the most straightforward categorization, the ROI portrayed in the scans distinguishes between, but not confined to, lung, brain, and thorax registration [10]. The objective is usually confined by the initial domain differences between input images and can be of temporal, spatial, dimensional, or modular [11] nature. Temporal registration is usually required in intra-subject

The associate editor coordinating the review of this manuscript and approving it for publication was Orazio Gambino.

cases, where a ROI is recorded in the same patient at different time points whereas spatial registration can be found in instances where either several viewpoints in the same patient or the same ROI across different patients (inter-subject) are imaged. Lastly, multi-modal registration [12] is employed where multiple imaging modalities have to be joined for a comprehensive situation assessment. The registration method however, is usually chosen by the researcher or healthcare professional and defines the matching criteria and type of geometric transformation. Correspondences between images are generally established by either matching anatomical landmarks or image information such as intensity values. While manual landmark annotation is time consuming [13], automated approaches have been implemented to a varying degree of success. Information based methods range from directly computing intensity differences in sum-of-squared-differences (SSD) over statistical methods like normalized cross-correlation (NCC) [14] to probability based methods such as normalized mutual information (NMI) [15]. Linear geometric transformations methods, such as affine deformations, extend the classic rigid body translations and rotations [16] with shearing and scaling properties [17]. Especially in medical image registration, with complex transforms addressing non-linear behavior required, non-rigid transformations are commonly deployed. Where in affine registration each voxel is subject to the same deformation parameters, non-rigid approaches assign individual deformations to each point. This can be achieved via point-by-points deformations as introduced in Thirion et al. [18] or with parametrized basis functions, shown in Rueckert et al. [19]. Elastic regularization of non-rigid or deformable registration restricts the transformation to a physically plausible solution space, the small deformation assumptions [20] it is build on, however, struggle with large deformations [21], [22]. Introducing analogies from fluid mechanics [21], [23], notably modeling the deformation with velocity fields, allowed to overcome the restrictions imposed by linear elasticity and thereby register images of content subjected to large deformations. The mathematically rigorous extension of flow fields to flows of diffeomorphisms [24], exemplified by the large deformation diffeomorphic mapping (LDDMM), ensures two important properties of the transformation. First, it preserves the topology of the domain, preventing tearing and folding in registered organs. Second, it enables the calculation of metric distances between shapes [25], a vital measure in computational anatomy. Application of deep learning, especially in lung registration, provides fast algorithms matching the accuracy of conventional, variational approaches. Due to non-existent ground truth data and small publicly available dataset, deep learning methods lack the robustness of their variational counterparts [26].

A. RELATED WORK

Discontinuities in lung registration consists in sliding motions in two areas, the lung-thorax boundary and the lung fissures, and face a specific problem: allowing for discontinuous deformations between domains while ensuring continuity and smoothness inside the domains. In both cases, the pleura forms a boundary between sliding anatomical structures, and accurate image registration along this interface is particularly important for capturing motion in the immediate vicinity. Although the volume of this pleural boundary region is small compared to the interior of the lungs or thorax, it is clinically relevant due to the presence of nodules that may be attached to the pleura. While not all pleura-attached nodules are malignant, reported malignancy rates vary widely—from approximately 70 percent in some studies [27] to as low as 1 percent in others [28]. Regardless of the exact incidence, assessment of malignancy typically relies on shape features [29] and growth analysis over time [30], both of which require precise and reliable registration, particularly near tissue interfaces such as the pleura. Several works concerning the lung-thorax have been published [31] and can be classified either as domain split methods or local regularization.

B. LOCAL REGULARIZATION

Local regularization methods employ a single deformation field to model motions with regularizers used to identify local exceptions to continuity constraints. Papież et al. [32] utilize a bilateral filter kernel to replace the Gaussian kernel for edge preserving smoothing. While this approach is able to simultaneously identify lung boundaries its applications comes with high computational costs. An isotropic total variation regularization is introduced in Vishnevskiy et al. [33], but is limited by a linear parametrization. Gong et al. [34] balances total variation L_1 at the boundaries with L_2 regularization inside domains to present a robust method even for small segmentation errors. Sliding motion on the small scale however is still incorrectly smoothed. Local regularization methods maintain a single deformation field for the image domain and attempt to identify local exceptions to continuity constraints. With a closed mathematical solution provided under the use of segmentation masks, Schmidt-Richberg et al. [35] allow global smoothing of the deformation field in the normal direction of the boundary surface while restricting smoothing in the tangential component to the individual segments. The decoupling of the transformation into tangential and normal components in the moving domain does not represent changes encountered under deformation. Addressing this problem by first calculating the boundary motion, a diffeomorphic piecewise approach in Risser et al. [36] smooths the the velocity fields independently in each domain while matching the velocities around the boundary to the tangential direction. When using different velocity field resolutions for lung and thorax fields, the interpolation required to create dense velocity fields (DVF) introduces smoothing across the boundary.

1) SPLIT DOMAIN METHODS

Split Domain Methods model sliding motion by using independent deformation fields for lung and thorax domains. While they don't share the difficulties of modeling sliding motion as local regularization methods, their common challenge is to correctly matching interfaces. Wu et al. [37] penalizes domain overlaps by introducing unique penalty intensities. However, selection of penalty intensity values to correctly restrict the deformation is time consuming. Limiting the sliding model to two objects represented by expensive motion masks, Eiben et al. [38] calculate a term weighing the interface matching error between those masks. Berendsen et al. [39] penalize the interface matching error with a local distance measure, which allows for separation by a third domain, but is sensitive to the surface discretization. Heldmann et al. [40] and Derksen et al. [41] both calculate the interface error from the distance between deformed moving and fixed domain boundary. Insufficient parametrization can lead to a dominating matching term, restricting the registration forces. Introducing an additional third deformation field concerned with the global normal deformation relative to the interface, Delmon et al. [42] then deform the sub-domain domains with independent tangential deformation fields. Deriving the normal direction from local bases in the moving image, this approach assumes small changes in the interface. A hybrid approach utilizing a single deformation field that ensures discontinuous movement is presented by Hua et al. [43] by extending the basis functions for deformation parametrization. The b-spline functions support for an image point is restricted to locations in the same domain as the initial point. Even though deformed by one deformation field, this switch in function support can still displace voxels across the boundary creating overlapping domains. Ng and Ebrahimi [44] use CNNs to produce displacement fields preserving sliding motion by calculating the cross products of neighboring displacements. While performing fast and matching registration accuracy, local non-smooth displacements are not restricted. Even though most of these methods need segmentations, deep learning methods offer fast, robust, and accurate tools for obtaining organ segmentations [45], [46], [47]. Isotropic total variation and bilateral filtering, as the methods not requiring segmentations, struggle with boundaries between resembling structures and low contrast.

With the rise of deep learning methods, recent years have seen a growing number of learning-based models applied to image registration. While these models can still be broadly categorized into the two traditional groups discussed earlier, they warrant a dedicated discussion due to their distinct mechanisms. Learning-based approaches are typically divided into supervised, unsupervised, and GAN-based methods [26]. Fully supervised CNNs often adopt an encoder–decoder architecture to better capture large and complex deformations [48]. However, their performance is constrained by the availability of diverse and extensive training datasets. These models heavily rely on ground truth deformation fields, which are generally unavailable in real-world registration tasks. As a result, they often depend on synthetic data or deformation fields generated by classical variational methods. Weakly supervised methods attempt to overcome this limitation by aligning image pairs using annotated labels instead of full ground truth. For instance, landmark-based approaches [49] can generalize to broader datasets but remain limited by the need for accurate annotations. Unsupervised methods are conceptually closer to traditional registration techniques, as they do not require ground truth data. Instead, they optimize similarity measures between image pairs. Techniques that use edge-based normalized gradient field (NGF) metrics and second-order curvature regularization [50] provide robust performance against large, respiration-induced deformations. However, they may struggle to handle subtle displacements and sliding motion effectively. GAN-based approaches utilize a generator to produce deformation vector fields (DVFs) and a discriminator to evaluate the similarity between the warped and fixed images. While this adversarial framework enhances DVF quality, it introduces training instability and still depends on traditional similarity losses, which can limit generalization. To address this, some models replace conventional similarity metrics with perceptual loss computed from feature maps [51]. This offers a more robust image registration framework, albeit at the cost of increased computational complexity due to feature extraction. For lobe-on-lobe sliding, substantially less research has been performed. Notable first works proved the existence of sliding between lobes by Ding et al. [52], by evaluating displacement differences in lung-by-lung registrations compared to lobe-by-lobe registrations. Further works performing in-depth analyses with newly introducing measures confirmed the findings, either when lungs were deformed as a continuum [53] or on a lobe by lobe basis [54], [55]. The latest work coupled independently registered lobes and thoracic domains with finite element inspired contact constraints and agree on the importance of modeling lobar sliding [56], [57]. These contributions have provided significant work and evolved from lung wise to lobe wise registrations, yet fall short of modeling the lung domain realistically. Our contribution in this work is twofold: We improve on our previous method introduced in Alscher et al. [58] to state of the art diffeomorphic piecewise image registration regularized with collision constraints. This ensures topology preservation while allowing for discontinuities at the object boundaries. Secondly, we examine sliding between lung lobes under image registration with self-collision constraints by modeling the lung as a continuous, connected body. To the best of our knowledge, this is the first time modeling self-sliding between the lung lobes by one mesh is done. Evaluation of our approaches are performed on synthetic as well as medical data sets.

II. METHODOLOGY

According to the categorization from Sotiras et al. [8], an image registration algorithm consists of three main components: an objective function, a deformation model, and an optimization method. This section will introduce our implementation of these parts and step-by-step develop them to arrive at a final formulation.

A. IMAGE REGISTRATION ENERGY PROBLEM

We will start the formulation by developing the objective function of the image registration problem by borrowing the mathematical setting from Modersitzki [20]. Let the d-dimensional images be $I, J : \Omega \longmapsto \mathbb{R}$ with the image domain $\Omega :=]0, 1[^d \subset \mathbb{R}^d$, with I the moving image and J the fixed image. The transformation $\Phi : \Omega \longmapsto \Omega$ warps the moving image to the fixed image. A common formulation is the minimization of the energy functional $F(\Phi, I, J)$ with $x \in \Omega$ any point in I.

$$\min_{\Phi} F(\Phi, I, J) = \min_{\Phi} \int_{\Omega} M(I(x), J(\Phi(x))) + \sum_i S_i(\Phi(x))dx \tag{1}$$

The similarity measure $M : \Omega \times \Omega \longmapsto \mathbb{R}$ computes the resemblance of both images with each other and can be seen as the driving force of the registration. It can be picked from a range of measures. For intensity based image registration, commonly used are measure are sum of squared differences (SSD), cross-correlation (CC), and mutual information (MI) [59]. Deformable image registration as an ill-posed problem [60] requires additional restrictions $S : \Omega \longmapsto \Omega$ called regularizations to be imposed on the transformation Φ. The regularization term is controlling the plausibility [61] of the transformation, penalizing undesirable solutions [62]. $\sum_i S_i$ as a collection of terms with $i \in \mathbb{N}$ can either introduce prior knowledge [8] or favor certain physical models [63].

B. OPTIMIZATION MODEL

While the similarity term is aimed to solve the matching problem, the regularization term focuses on the deformation restrictions. Splitting the decision variable $\Phi(x)$ into two Φ^{Match} and Φ^{Reg} creates the dual problem and allows us to formulate the alternating direction method of multipliers (ADMM) [64] for the objective function, an optimization approach successfully used in image registration [40], [65], [66]. This method favors objective functions where each sub-problem can be solved efficiently locally [67]. Even though convergence for ADMM is only proven for convex functions, certain non-smooth and non-convex functions, like L_1 regularizations, have been shown to converged successfully [68]. This split comes at a cost by introducing a new variable u, called the dual variable. u couples the two decision variables to ensure the solution of the dual problem approximates the primal solution. Two kinds of coupling are developed to address the two intersection problems: Eulerian and Lagrangian constraints. Equation 1 with split decision variables shown in Equation 2 with two available coupling constraints. The Eulerian coupling constraints restrict differences in the deformation field Φ^{Reg} and Φ^{Match} whereas the Lagrangian constraints restrict the deformation resulting from these two fields in a set of specified coupling points $x_c \in \Omega$.

$$\begin{aligned} \min_{\Phi^{Match}, \Phi^{Reg}} & \int_{\Omega} M(I(x), J(\Phi^{Match}(x))) \\ & + \sum_i S_i(\Phi^{Reg}(x))dx \\ \text{subject to } & \Phi^{Match} - \Phi^{Reg} = 0 \text{ or} \\ & \Phi^{Match}(x_c) - \Phi^{Reg}(x_c) = 0 \end{aligned} \tag{2}$$

An exemplary overview of the iterative sequential updates for step $j+1$ with Eulerian constraints, coupling two deformation fields acting the same image domain, to the decision variables and the dual variable can be seen in Equation 3.

$$\begin{aligned} \Phi^{Match,j+1} \leftarrow & \min_{\Phi^{Match}} \int_{\Omega} M(I(x), J(\Phi^{Match}(x)))\, d\, x \\ & + \frac{\rho}{2} |\Phi^{Match} - \Phi^{Reg,j} + u^{k,j}|_2^2, \end{aligned} \tag{3a}$$

$$\begin{aligned} \Phi^{Reg} \leftarrow & \min_{\Phi^{Reg}} \int_{\Omega} \sum_i S_i(\Phi^{Reg}(x))\, d\, x \\ & + \frac{\rho}{2} |\Phi^{Match,j+1} - \Phi^{Reg} + u^{k,j}|_2^2, \end{aligned} \tag{3b}$$

$$u^{k,j+1} \leftarrow u^{k,j} + \Phi^{Match,j+1} - \Phi^{Reg,j+1}. \tag{3c}$$

C. DEFORMATION MODEL

We will develop a more specific energy functional (Equation 4) from the generalized formulation Equation 1in the next paragraphs and start by outlining the deformation model. Modeling deformation via time-varying velocity fields $V(x, t) : \Omega \times \mathbb{R} \longmapsto \mathbb{R}^d$ with $t := [0; 1]$ offers multiple advantages. Developed in 1996 by Christensen et al. [21], [69], velocity fields allow for large deformations while creating diffeomorphic transformations. Diffeomorphic behaviour can be achieved by suitable regularization of the velocity field [24]. With the right simplifications, this deformation model can easily perform as a Stationary-Velocity-Field registration (SVF) [70] or Free-Form Deformation (FFD) [19]. Lastly, the analogy to fluid mechanics allows for an easy transfer of knowledge and intuitive understanding of Lagrangian and Eulerian specifications of flow fields.

$$\inf_V \int_{\Omega} \Big(\int_0^1 ||V(x, t)||_L^2 dt + M(I(x), J(\Phi(x)))dx \Big) \tag{4}$$

The norm $||.||_L$ is responsible for both restricting the solution space to the group of diffeomorphisms and ensureing smoothness [71]. The linear operator $L = \mu\Delta + \lambda Id$, used in computational anatomy due to its corresponding order of differentiability to operators from fluid mechanics, is chosen to establish diffeomorphic transformations [72].

1) DISCRETIZATION

In computation practice, deforming images with time-dependent velocity fields over continuous domains requires spatial and time discretization as well as an appropriate integration method. First we will elaborate on the spatial discretization also often called the parametrization of the image registration approach. Apart from reducing the image registration problems dimensionality, parametrization of the deformation model can provide an improved robustness as well as faster convergence [33]. One approach is to find the voxel velocities of the underlying image by deriving them from a grid of overlaid velocity control points. The spatial resolution of the control point grid is usually more coarse than the image itself, requiring an interpolation scheme to create DVF. Additionally to their smoothing properties [73], B-Spline approximation, especially of cubic order, are supported only locally, restricting their influence to their neighborhood [19]. This makes them efficient [19] and prevents unwanted oscillations exhibited with global polynomials [74]. An example with b-splines for interpolation uses the scalar field $A : \Omega \longmapsto \mathbb{R}$ and point $p \in \Omega$ in our image domain Ω^3. To find the scalar value $A(p)$, the control point grid, defined as $\zeta \in R^{n_1 \times n_2 \times n_3}$ with an equidistant spacing at $\delta = (\delta_1, \delta_2, \delta_3)$, is evaluated as follows

$$A(p) = \sum_{l=0}^{3} \sum_{m=0}^{3} \sum_{q=0}^{3} B_l(u) B_m(v) B_q(w) \zeta_{i+l,j+m,k+q}. \tag{5}$$

The basis functions are detailed as [75]

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

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

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

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

The local coordinates are computed 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 $w = (p_3 - \delta_3 * \lfloor p_3/\delta_3 \rfloor)/\delta_3$. The control points influencing the approximation are found by $i = \lfloor p_1/\delta_1 \rfloor - \xi, j = \lfloor p_2/\delta_2 \rfloor - \xi$, $k = \lfloor p_3/\delta_3 \rfloor - \xi$, with representing the shift in support ξ along the axis and the local coordinates derive from [58]. This scheme provides a robust interpolation method for scalar fields, used for example for determining intensity values at non integer locations in images. It can also be readily extended to operate on vector fields, as demonstrated in this paper for cases such as velocity fields. The integration of any value over the spatial domain is the simple sum of values at discrete points weighted by their area or volume.

The time discretization of the velocity field is more trivial and just divides the continuous field into m timesteps. The evolution of the deformation over time is described by an ordinary differential equation (ODE) (Equation 7).

$$\frac{d}{dt}\Phi(x,t) = V(\Phi(x,t),t) \tag{7}$$

The deformation at time $t = 0$ is always the identity transform $\Phi_0(x) = \Phi(x,0) = x$. Consequently, the integration required to find the deformation Φ_t at time t, results in an initial value problem [76]. The choice of integration method can have a direct impact to the order of multiple magnitudes on the consistency error of the ODE solution [77]. Two numerical integration methods are used in this paper. A forward Euler approach in Equation 8 and a trapezoidal approach in Equation 9.

$$\Phi_{t+1}(x) = \Phi_t(x) + \Delta t V_t(\Phi_t(x)) \tag{8}$$

$$\Phi_{t+1}(x) = \Phi_t(x) + \frac{1}{2}\Delta t (V_t(\Phi_t(x)) + V_{t+1}(\Phi_{t+1}(x))) \tag{9}$$

Implementation of the explicit Euler's method is rather simple and straightforward. The trapezoidal method as an implicit approach [78] requires an evaluation of the velocity field at time steps $t+1$ and t to find the solution for the deformation at time $t+1$. A multi-step predictor-corrector method inspired by Heun's method [79] uses an initial guess from the Euler's method to converge to a solution 1.

Algorithm 1 Multi-Step Predictor-Corrector Method

```
Input: Initial condition Φ_t, time step Δt, tolerance
Output: Final state Φ_{t+1}
Φ^{pred}_{t+1} ← Φ_t + Δt · V_t(Φ_t(x))
while true do
    Φ^{cor}_{t+1} ← Φ_t + ½Δt · (V_t(Φ_t(x)) + V_{t+1}(Φ^{pred}_{t+1}(x)))
    if |Φ^{cor}_{t+1} − Φ^{pred}_{t+1}| < tolerance then
        break {Exit the loop if within tolerance}
    end if
    Φ^{pred}_{t+1} ← Φ^{cor}_{t+1} {Update the prediction}
end while
Return Φ^{cor}_{t+1} {Return the corrected value}
```

Once converged, the solution assures the properties of the integrated ODE at the evaluation times [77], also called collocation points, and justifies the trade off of the additional iterations.

2) PIECEWISE REGISTRATION AND COLLISION CORRECTION

The image domain Ω can be separated into subdomains, representing independently moving objects. These distinct objects, are regarded as closed sets $\Omega^k \subset \Omega$ with their boundaries $\delta\Omega^k$ and let us adapt Equation 4 with $D(I, J, \Phi^k) = \int_0^1 ||V^k(x,t)||_L^2 dt + M(I(x), J(\Phi(x)))$, giving

$$\inf_{V^k} \int_{\Omega^k} D(I, J, \Phi) dx. \tag{10}$$

The goal of collision correction is to correct overlap of any independently registered subdomains and consequently require an object representation with clear boundaries.

We begin by establishing a framework for handling collisions between distinct objects, referred to as inter-collision correction. The deformation of a domain is represented as $\Phi(\Omega) := \{\Phi(x)|x \in \Omega\}$. The obstacles for which collisions should be checked for Ω^k are defined by the set of domains $S_{\Omega_n} = \{\Omega_1, \ldots, \Omega_n\}$, where $n \in \mathbb{N}$, which together form Ω_{obs} as described in Equation 11.

$$\Omega_{obs} = \bigcup_{i=1}^{n} \Omega_i \quad \text{with} \quad \bigcap_{i=1}^{n} \Omega_i = \{\emptyset\} \tag{11}$$

The indicator function $\mathbf{1}$ takes the value 1 for arguments that satisfy a given condition, and 0 otherwise. This function is useful for calculating the Lebesgue integral $\int \cdot d\mu$ over the continuous set of intersecting regions and assumes that our sets are measurable. The Lebesgue integral computes a measure over the set, which, in this context, represents the area or volume of overlap. Since the indicator function is bounded and the continuous set has finite measure, the result of the integral remains finite, preventing it from diverging. The collision correction function CC is defined as $CC : \Omega \times \Omega \times \Omega \longmapsto \mathbb{N}$.

$$CC(\Phi, \Omega^k, \Omega_{obs}) = \int_{\Omega^k} \mathbf{1}(\exists x : \Phi(x) \in \Omega_{obs}) d\mu \tag{12}$$

The function for the continuous set of self-intersection, or intra-collision, is built upon this concept and is defined as follows: $SC : \Omega \times \Omega \longmapsto \mathbb{N}$.

$$SC(\Phi, \Omega^k) = \int_{\Omega^k} \mathbf{1}(\exists x : x \neq y \wedge \Phi(x) = \Phi(y) \wedge y \in \Omega^k) d\mu \tag{13}$$

These formulations are defined on continuous sets because the idealized images are also continuous. The next steps will involve discretizing these formulations and presenting our implementation. Our objects Ω^k, as closed sets, can now be viewed as a finite collection of points (e.g. pixels). This allows us to replace the Lebesgue integral with a simple summation. The inter-collision is expressed as Equation 14, and the intra-collision is expressed as Equation 15.

$$CC(\Phi, \Omega^k, \Omega_{obs}) = \sum_{x \in \Omega^k} 1(\exists x : \Phi(x) \in \Omega_{obs}) \tag{14}$$

$$SC(\Phi, \Omega^k) = \sum_{\Omega^k} 1(\exists x : x \neq y \wedge \Phi(x) = \Phi(y) \wedge y \in \Omega^k) \tag{15}$$

We now have a measure of the discretized overlapping area. However, depending on the discretization of Ω_L, a single overlapping element could indicate either a slight overlap or a significant intrusion. Larger intrusions should be penalized more heavily than smaller overlaps, so we need a depth measure to quantify the extent of the overlap. To achieve this, we aim to reformulate the approach once more. This reformulation will also help clarify the definitions of the measures used, such as the minimum set distance (MSD), which is defined as $MSD : \mathbb{R}^d \times \Omega \longmapsto \mathbb{R}$.

$$MSD(x, \Omega) = \min_x |x - y|_2 \quad \forall \quad y \in \Omega \tag{16}$$

The adapted signed distance field (ASDF), defined as $ASDF : \mathbb{R}^d \times \Omega \longmapsto \mathbb{R}$, as shown in Equation 17, now provides a measure of the depth for intersecting points, or a zero value where there is no intersection. The check of $x \in \Omega$ via winding numbers is performed by interpolating the coordinates within this adapted signed distance field (as detailed below). Note that the minimum set distance (MSD) is calculated to the boundary of the set in the case of an intersection. The loss function, in its discretized form, can now be written as Equation 18.

$$ASDF(x, \Omega) = \begin{cases} 0 & \text{if} \quad x \notin \Omega \\ -MSD(x, \delta\Omega) & \text{if} \quad x \in \Omega \end{cases} \tag{17}$$

$$CC(\Phi, \Omega^k, \Omega_{obs}) = \sum_{x \in \Omega^k} \sqrt{ASDF(\phi(x), \Omega_{obs})^2} \tag{18}$$

The inter-collision correction implementation is reproduced from Alscher et al. [58]. A combined signed distance field for all collision obstacles in the fixed image Ω_S is created and adapted to fit Equation 17. Evaluation points for inter-collision correction are derived from moving image masks for respective segment and either cover the whole segment domain or outer layers obtained from morphological operations. During image registration all evaluation points deformed from the moving image to the fixed image are evaluated according to Equation 18.

An implementation of the collision detection from the lung's perspective is presented in 2, where potential inter-collisions between the lung and thorax are identified. To perform collision detection from the thorax's perspective with respect to the lung, one can simply swap the input domains.

Algorithm 2 Inter-Collision Detection

1: **Input:** Adapted Signed Distance Field (ASDF) of Ω_{thorax}; n sampled points $p_0 = \{p_0^0, p_0^1, \ldots, p_0^n\}$ in Ω^{Lung} at time 0; iteration number i; initial deformation field Φ^{Lung}
2: **Output:** Deformation field Φ^{Lung}
3: **for** i in iterations **do**
4: loss $\leftarrow 0$
5: **for** $j = 0$ to n **do**
6: $p_1^j \leftarrow \Phi_1^{Lung}(p_0^j)$
7: loss $\leftarrow$ loss $+ \sqrt{ASDF(p_1^j, \Omega_{thorax})^2}$
8: **end for**
9: $\Phi^{Lung} \leftarrow \Phi^{Lung} - \text{lr} \cdot \nabla\Phi^{Lung}$
10: **end for**

Self-intersection is more complex to handle. While signed distance fields (SDFs) can be constructed using the concept of general winding numbers, even in the presence of self-intersecting meshes [80], we encounter some challenges. The MSD measured for the closest distance to the boundary in the spatially deformed configuration can have very small values deep inside the object, even when close to the intersecting boundary. This means that self-intersecting points, when interpolated on this signed distance field, can produce image gradients that push them further into the object. An example can be seen in Figure 1. An object consisting of an upper, blue and lower, red prong is deformed from its initial configuration (a) to a self-intersecting deformation (b). Using a SDF created from the self-intersection mesh in spatial configuration, gradients are calculated for the intersecting points from the upper, blue prong and lower, red prong. The results show that even though red points should be moved down and blue points up, both sets of points are assigned the directions increasing intersection (c). To address this, we aim to define the intersection depth by pulling the coordinates back to the non-intersecting material configuration, as it was before the deformation. To simplify the notation, let us define the set of self-intersecting points as: $SI = \{\Phi(x) : x \neq y \wedge \Phi(x) = \Phi(y) \wedge y \in \Omega^k\}$

$$SC(\Phi, \Omega^k) = \sum_{x \in SI} \sqrt{ASDF(\Phi^{-1}(x), \Omega^k)^2} \tag{19}$$

Another issue arises because the inverse of the function does not exist, as we have self-intersections, which make the function non-injective (i.e., mapping multiple points to the same value). This results in undefined behavior, and if not addressed, the gradients for intersecting points from different regions of the SDF may be assigned the same gradient. An example of this can be seen in Figure 1, where both points from the upper and lower prongs exhibit the same gradient pattern. To address this, we proceed as follows: We split our original object domain into $n \in \mathbb{N}$ subdomains. These subdomains do not overlap with each other in the same regions of the original object but, together, they fully represent it.

$$\Omega^k = \bigcup_{i=1}^{n} \Omega^{k,i} \quad \text{with} \quad \bigcap_{i=1}^{n} \Omega^{k,i} = \{\emptyset\} \tag{20}$$

With individual deformations $\Phi_k : \Omega^{k,i} \longmapsto \Omega$ mapping the subdomains, we now have, for each collection of self-intersecting points, an inverse function. $SI_i = \{\Phi_i(x) : x \in \Omega^{k,i} \wedge \Phi_i(x) = \Phi_j(y) \wedge y_j \in \Omega^{k,j} \forall j \neq i\}$

$$SC(\Phi, \Omega^k) = \sum_{i=1}^{n} \sum_{x \in SI_i} \sum_{j \neq i}^{n} \sqrt{ASDF(\Phi_j^{-1}(x), \Omega^{k,j})^2} \tag{21}$$

This self-collision detection is implemented using a point-simplex intersection check, which varies depending on the dimensionality. First, the object in its material configuration is discretized using a simplex mesh, then deformed. Afterward, all mesh vertices are checked for intersections with simplices they do not belong to, allowing for an inside-outside check. Next, the barycentric coordinates of the intersecting points within their collision partner simplices are calculated. These coordinates are then used to pull back the points to the material configuration. The SDF of the object in its undeformed configuration is interpolated at these barycentric coordinates, providing a direction for each point. This direction vector, initially in barycentric coordinates, is then warped back to the spatial configuration by computing the Cartesian coordinates within the intersecting simplex. The results, shown in our graphic, illustrate how the red points are pushed down and the blue ones are pushed up. This method allows us to assign the correct direction to points that intersect in the same area based on their initial configuration, without the computational cost of recalculating a new SDF for each intersection occurrence.

Furthermore, a pseudocode representation of the implementation is presented in 3.

Algorithm 3 Inter-Lobe Collision Detection

Input: Adapted Signed Distance Field (ASDF) of Ω_{Lung}; lobe segmentations Ω^{Superior} and Ω^{Inferior} for touching lobes; tetrahedralized volume of Ω_{Lung} as n vertices $v_0 = \{v_0^0, v_0^1, \ldots, v_0^n\}$ at time t; iteration number i; initial deformation fields Φ^{Superior} and Φ^{Inferior}
Output: Updated deformation fields Φ^{Superior} and Φ^{Inferior}
$v_{0,\text{Superior}} \leftarrow$ all n_{Superior} vertices in Ω^{Superior} at time 0
$v_{0,\text{Inferior}} \leftarrow$ all n_{Inferior} vertices in Ω^{Inferior} at time 0
for i in iterations **do**
  $\text{loss}_{\text{Superior}} \leftarrow 0$
  for $j = 0$ to n_{Superior} **do**
    $v_{1,\text{Superior}}^j \leftarrow \Phi_1^{\text{Superior}}(v_{0,\text{Superior}}^j)$
    if $v_{1,\text{Superior}}^j \in \Phi^{\text{Inferior}}(\Omega^{\text{Inferior}})$ **then**
      $\text{loss}_{\text{Superior}} \leftarrow \text{loss}_{\text{Superior}} + \sqrt{\text{ASDF}(\Phi_{\text{Inferior}}^{-1}(v_{1,\text{Superior}}^j), \Omega_{\text{Inferior}})^2}$
    end if
  end for
  $\text{loss}_{\text{Inferior}} \leftarrow 0$
  for $j = 0$ to n_{Inferior} **do**
    $v_{1,\text{Inferior}}^j \leftarrow \Phi_1^{\text{Inferior}}(v_{0,\text{Inferior}}^j)$
    if $v_{1,\text{Inferior}}^j \in \Phi^{\text{Superior}}(\Omega^{\text{Superior}})$ **then**
      $\text{loss}_{\text{Inferior}} \leftarrow \text{loss}_{\text{Inferior}} + \sqrt{\text{ASDF}(\Phi_{\text{Superior}}^{-1}(v_{1,\text{Inferior}}^j), \Omega_{\text{Superior}})^2}$
    end if
  end for
  $\Phi^{\text{Superior}} \leftarrow \Phi^{\text{Superior}} - \text{lr} \cdot \nabla\Phi^{\text{Superior}}$
  $\Phi^{\text{Inferior}} \leftarrow \Phi^{\text{Inferior}} - \text{lr} \cdot \nabla\Phi^{\text{Inferior}}$
end for

The collision correction formulations are not be continuous, so we apply the introduced ADMM (Equation 3) to split these terms. We then formulate the loss functions using velocity fields, and employ a discrete version of the diffusion term D. For the inter-collision corrections, we use two velocity fields, $V^{k,Match}$ and $V^{k,CC}$, to separate the collision

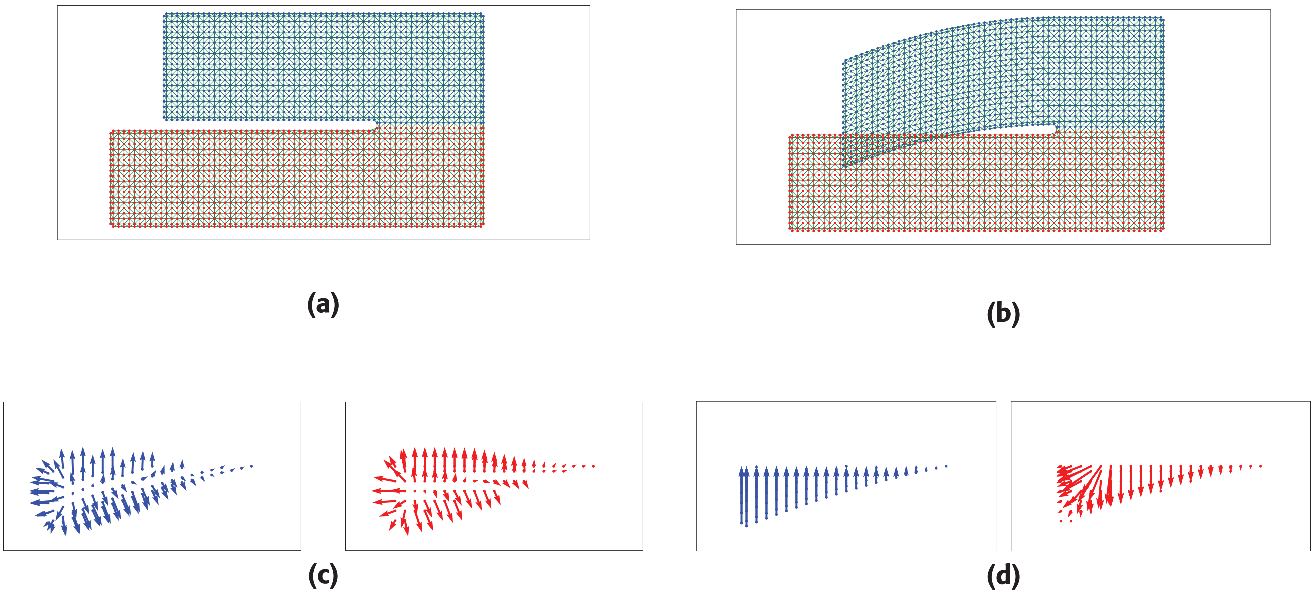

FIGURE 1. **Example of gradient in different configurations. (a) Meshed (green) prong object with colour-coded upper part (blue) and lower part (red). (b) Deformed prong object with self-intersection between upper and lower part. (c) Image gradients of signed distance field in spatial configuration. (d) Image gradients of signed distance field with pullback from material configuration.**

correction from the matching term. These fields are coupled through Eulerian constraints between them. Equation 22.

$$2V^{k,Match,j+1} \leftarrow \underset{V^{k,Match}}{\arg\min} \sum_{x\in\Omega^k} D(I, J, \Phi^{k,Match}) + \frac{\rho}{2}|V^{k,Match}(x) - V^{k,CC,j}(x) + u^{k,j}(x)|_2^2, \tag{22a}$$

$$V^{k,CC,j+1} \leftarrow \underset{V^{k,CC}}{\arg\min} \sum_{x\in\Omega^k} CC(\Phi, \Omega^k, \Omega_{obs}) + \frac{\rho}{2}|V^{k,Match,j+1}(x) - V^{k,CC}(x) + u^{k,j}(x)|_2^2, \tag{22b}$$

$$u^{k,j+1} \leftarrow u^{k,j} + V^{k,Match,j+1} - V^{k,CC,j+1}. \tag{22c}$$

For intra-collision, we also use two velocity fields, one for each subdomain of Ω^k, Ω^m, and Ω^l. These velocity fields are coupled through Lagrangian constraints at their shared boundary x_c, depending on the specific setting.

$$V^{m,j+1} \leftarrow \underset{V^m}{\arg\min} \sum_{x\in\Omega^m} D(I, J, \Phi^m) + SC(\Phi^m, \Omega^m) + \frac{\rho}{2}|x_c \circ \Phi^m - x_c \circ \Phi^{l,j} + u^j|_2^2, \tag{23a}$$

$$V^{l,j+1} \leftarrow \underset{V^l}{\arg\min} \sum_{x\in\Omega^l} D(I, J, \Phi^l) + SC(\Phi^l, \Omega^l) + \frac{\rho}{2}|x_c \circ \Phi^{m,j+1} - x_c \circ \Phi^l + u^j|_2^2, \tag{23b}$$

$$u^{j+1} \leftarrow x_c \circ \Phi^{m,j+1} - x_c \circ \Phi^{l,j+1} + u^j. \tag{23c}$$

III. EXPERIMENTS AND RESULTS

We conducted two sets of experiments to examine the effects of modeling sliding motion, one between lung and thoracic cavity wall and a second one between inferior and superior lung lobes. Both sets include explanatory 2D cases with synthetic data and 3D cases with medical data from the DIR-LAB dataset [81], [82]. In order to asses the magnitude of sliding and the performance of the registrations, we introduce several measures. For cases without annotated landmarks, the Sørensen-Dice index (DICE) is used to show the registration success. For cases with annotated landmarks, namely the experiments with DIR-LAB data, the target registration error (TRE) measures the distance between registered landmarks for the moving image to their counterparts in the fixed image. The TRE employs the snap-to-voxel approach as used in various works [33], [38], [82]. A numerical measure for sliding motion is the maximum shear stretch [55]. The eigenvalues of the decomposed deformation gradient tensor $F = \nabla\Phi$ constitute stretches in the principal directions, giving the maximum shear stretch γ_{max} as the difference between maximum and minimum components.

$$\lambda_{max} = \frac{\lambda_1 - \lambda_3}{2} \tag{24}$$

with

$$\lambda_i = \sqrt{eigenvalues\ of\ F^T F}. \tag{25}$$

Reducing areas of overlap between the domains of the piecewise approaches is a key goal of our collision corrections. In 2D cases we measure the overlap by calculating the area between polygons representing the contours of the domains. In 3D experiments, the overlap is computed by generating triangle surface meshes for each domain, registering the meshes with the respective deformations and summing the voxels in the fixed image covered by two or more meshes. Gaps are subsequently voxel covered by none of the meshes. Lastly, a measure for assessing the diffeomorphic quality of deformations is needed. Diffeomorphism are continuously differentiable bijections with a continuously differentiable inverse [83]. Local bijectivity and invertibility can be shown with positive determinants of the Jacobian of the deformation [84], requiring the resulting dense deformation fields (DDF) to have no local deformations with $det(\nabla\Phi) <= 0$. The Jacobian is computed via a central differences approach. However to evaluate the correctness or symmetry of the inverse of the deformation requires a measure called inverse consistency. Points are subjected successively to Φ and Φ^{-1} with the resulting error between the deformation and the identity $|x - \Phi^{-1}(\Phi(x))|_2$ constituting to the inverse consistency error (ICE). Evaluation of dense deforamtion fields with regards to smoothness and continuity rely heavily on the discretization and interpolation methods employed. Hence in this paper, we restrict our evaluation to visual examinations. For all experiments, SSD was used as the similarity measure. SSD is preferred for mono-modal image registration [85] and penalizes outliers more than the sum of absolute differences [62], which makes it a suitable measure in thoracic image registration with high intensity differences between air-filled lungs and dense tissue thoracic volumes. Additionally, all algorithms employ a multi-scale strategy to capture both global and local deformations as well as avoid local minima [74]. Multiple image registration passes are run sequentially from coarse to fine spatial spatial parametrizations, with the result from the former being carried over to the next scale. Gradient formulations of the objective function required by variational methods are usually obtained by computing the Euler-Lagrange equations for the functional 1. However, for complex functionals, their computation can be intricate and prone to errors [86]. To avoid this predicament we implemented all our algorithms in a differentiable programming framework capable of back-propagation and auto-differentiation, with a convenient option to run our programs on GPU. We further reduce the computational cost, by not sending all points from the moving image domain through the forward pass of our algorithm, but just selecting a subset of 10% of the complete point set. To still ensure complete coverage of the spatial domain, this subset is stochastically selected anew for every pass, yielding the same accuracy as a bigger, fixed point set.

A. SLIDING BETWEEN LUNG AND THORACIC CAVITY WALL

Our first set of experiments concerns the sliding between lung and thoracic cavity wall. We independently register lungs and thoracic cavity. For each entity, lungs and cavity, two velocity fields are coupled via Eulerian coupling constraints in the ADMM, one solving the registration problem and one for the collision correction as introduced in Equation 22. The minimization problems in the iteration updates are solved by a gradient descent approach with backtracking according to Armijo conditions [87]. To ensure LDDMM properties the velocity field is time discretized into 10 steps and numerically integrated with the trapezoidal rule.

1) SYNTHETIC INTERSECTION DATA

Our synthetic experiment highlights the differences between our LDDMM runs with and without inter-collision detection. The setup for this synthetic experiment is designed to distill the sliding image registration problem to its essential components and clearly demonstrate the effectiveness of our method. Two distinct squares, each measuring 20×20 pixels and containing internal intensity structures, are initially separated in the moving image. These squares, treated as independent objects, move freely. In the fixed image, they come into contact and then move in opposite directions, generating a sliding motion along their shared boundary—mimicking the surface interaction patterns commonly encountered in medical imaging scenarios. We employ the piecewise registration approach, meaning that each square is deformed with its own velocity field. Three successive registrations for each field are run with parametrization grids of 10, 5, and 1 pixel control point distances. The resulting registered contours are displayed in Figure 2. The collision corrected LDDMM (dCC-DIR) shows no overlap where the regular LDDMM struggles to keep the two structures separate with an overlapping area of around 5 pixel. Both algorithms show the same DICE measure of around 0.94 and piecewise diffeomorphic behaviour with 0% of $det(\nabla\Phi) <= 0$ and a mean inverse consistency of $6.5e^{-7}$ pixel.

B. MEDICAL DATA

The dCC-DIR is evaluated on the DIR-Lab dataset, a common benchmark in 4DCT lung image registration, from exhale to inhale images. Lung and thoracic cavity are cropped to include the lung with a padding of 5 voxel and are independently registered with individual velocity fields. The domains of interest of the velocity fields are defined by segmentation mask from https://continuousregistration.grand-challenge.org/. Inter-collision correction between lungs and cavity are modeled in the fixed image domain with adapted signed distance fields for their respective counterpart. First, an affine registration is run followed by a multi-scale deformable registration at 20, 10, and 5 voxel resolution for the control point parametrization. Case 8 is an exception that it has an additional initial resolution of 40 voxel. Both Table 2 and Table 1 are extensions from Alscher et al. [58] comparing only split domain methods algorithms. To our knowledge no algorithms published since 2023 utilize split domain approaches. The registration accuracy of the non-diffeomorphic CC-DIR is retained with a small decrease of

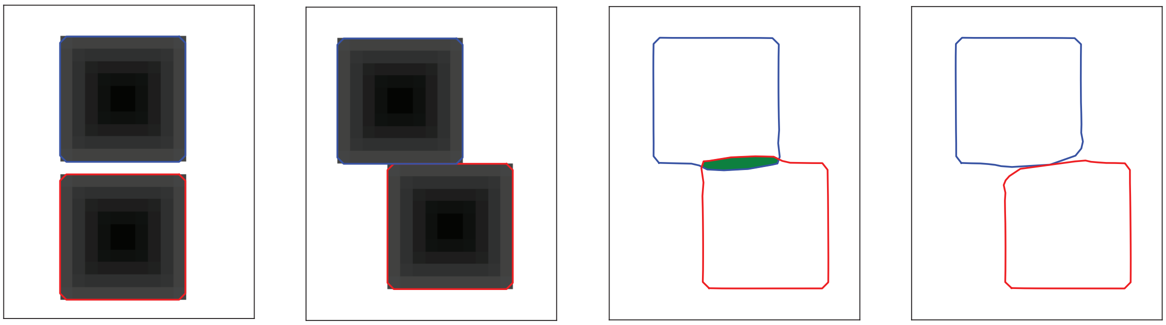

(a) (b) (c) (d)

FIGURE 2. Setup and results for synthetic intersection experiments. (a) Moving image. (b) Fixed image. (c) Registered contours with standard LDDMM resulting in overlaps. (d) Registered contours with dCC-DIR and no resulting overlaps.

TABLE 1. Extension of the gap/overlap comparison from [58] for the DIR-LAB data set by case in cm^3. The diffeomorphic dCC-DIR shows the third least overlap, but displays the highest gap value.

Case	Wu (2008)	Delmon (2013)	Berendsen (2014)	Hua (2017)	Eiben (2018)	Gong (2020)	Alscher (2023)	dCC-DIR
1	38/26	39/15	23/18	39/**2**	N/A/N/A	67/69	**22**/2	25 / 23
2	78/46	67/60	74/34	74/31	N/A/N/A	76/94	**53/6**	75 / 44
3	99/28	83/33	57/30	71/24	N/A/N/A	59/70	**53/7**	64 / 34
4	75/34	66/44	66/28	92/13	N/A/N/A	71/79	**41/7**	72 / 25
5	**11**/38	78/52	61/32	54/**6**	N/A/N/A	97/109	37/11	66 / 25
6	**10**/86	119/77	130/50	155/**11**	N/A/N/A	57/69	90/23	141 / 34
7	**10**/79	108/77	119/45	138/19	N/A/N/A	73/88	131/**14**	112 / 39
8	96/91	92/93	85/53	150/40	N/A/N/A	**80**/98	92/**39**	369 / **29**
9	61/34	**54**/44	70/51	58/14	N/A/N/A	63/72	78/**8**	**54** / 19
10	120/63	94/56	80/43	109/28	N/A/N/A	90/110	**48/19**	127 / 30
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**	113.6/29.8

TABLE 2. Extension of the TRE comparison from [58] for the DIR-LAB data set by case in *mm*. The diffeomorphic dCC-DIR retains the matching accuracy of its non-diffeomorphic counterpart, and even surpasses it in two cases. Regarding the mean TRE of the diffeomorphic CC-DIR method scores at fifth place with a 0.17 mm offset to the best performing algorithm.

Case	Wu (2008)	Delmon (2013)	Berendsen (2014)	Hua (2017)	Eiben (2018)	Gong (2020)	Alscher (2023)	dCC-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	$\mathbf{0.92 \pm 0.97}$	0.93 ± 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	$\mathbf{0.89 \pm 0.99}$
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	$\mathbf{1.05 \pm 1.06}$
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	1.48 ± 1.32
5	1.9 ± 1.5	2.0 ± 1.6	1.61 ± 1.48	$\mathbf{1.37 \pm 1.35}$	N/A ± N/A	1.45 ± N/A	1.9 ± 2.48	1.48 ± 1.74
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	1.67 ± 1.47
7	1.7 ± 1.1	1.9 ± 1.2	1.49 ± 1.06	$\mathbf{1.12 \pm 0.67}$	N/A ± N/A	1.24 ± N/A	1.42 ± 1.43	1.54 ± 1.43
8	1.6 ± 1.4	2.2 ± 2.3	1.62 ± 1.71	$\mathbf{1.18 \pm 1.22}$	N/A ± N/A	1.71 ± N/A	1.26 ± 1.17	1.46 ± 1.25
9	1.4 ± 0.8	1.6 ± 0.9	1.3 ± 0.76	$\mathbf{1.14 \pm 0.64}$	N/A ± N/A	1.24 ± N/A	1.2 ± 1.06	1.28 ± 1.06
10	1.6 ± 1.2	1.7 ± 1.2	1.5 ± 1.31	$\mathbf{1.08 \pm 0.82}$	N/A ± N/A	1.24 ± N/A	1.45 ± 1.72	1.49 ± 1.56
Mean	1.47 ± 0.96	1.66 ± 1.14	1.36 ± 0.99	$\mathbf{1.17 \pm 0.85}$	1.21 ± N/A	1.25 ± N/A	1.32 ± 1.42	1.34 ± 1.34

0.02 mm on average. This puts dCC-DIR in the middle field of those algorithms. The overlap as well as gap measure of dCC-DIR show a bigger increase to around double the overlap and gap. This still places dCC-CIR at third place out of 8 compared algorithms for the overlap measure. The average gap is outmatched by all other approaches.

An overview of the diffeomorphic measures in Table 3 shows all cases with no non positive $det(\nabla\Phi)$ and a mean inverse consistency error of 2.56*10-6 mm with a maximum of 0.00056 mm.

Two figures are supplied to show the quality of sliding motion along the lung boundary, Figure 3 displays the DDF and Figure 4 the maximum shear stretch in the coronal plane. Both images show an average of the ten DIR-Lab cases and were created by a simplified atlas registration. Case 1 was declared the atlas for our experiments and the results for Φ and γ_{max} were pulled back to the atlas domain. The deformation field in Figure 3 is changing right at the lung boundary, indicating a discontinuous, sliding movement. Figure 4 gives an overview of the shear stretch and demonstrates high maximum shear stretches along the lung boundary, matching the results from Figure 3. Quantitatively, the averaged maximum shear stretch at the lung boundary is with 1.8075 ± 2.066 around seven times higher than in the rest of the domain with 0.2409 ± 0.5185. The average time for the whole registration process with affine and multi-scale

TABLE 3. Diffeomorphic measures on the DIR-Lab cases showing % of det($\nabla\Phi$) <= 0 and inverse consistency errors. Inverse consistency is shown as the average and the maximum deviation from the identity transform after application of Φ and Φ^{-1}.

Measure	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6	Case 7	Case 8	Case 9	Case 10	Mean
$\%det(\nabla\Phi) <= 0$	0	0	0	0	0	0	0	0	0	0	0
Mean inverse consistency in mm	1.1e-06 ± 4.05e-06	1.58e-06 ± 5.44e-06	1.74e-06 ± 5.45e-06	2.3e-06 ± 6.32e-06	2.21e-06 ± 6.19e-06	3.09e-06 ± 8.20e-06	3.1-06 ± 8.24e-06	4.26e-06 ± 9.02e-06	1.79e-06 ± 5.09e-06	3.3e-06 ± 8.51e-06	2.65e-06 ± 7.27e-06
Maximum inverse consistency in mm	0.00013	0.00027	0.00019	0.00020	0.00052	0.00049	0.00039	0.00056	0.00015	0.00054	0.00056

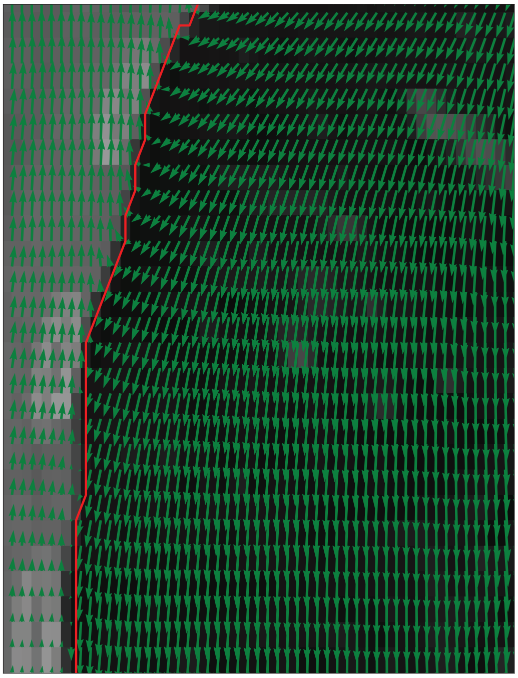

FIGURE 3. Close up of the average deformation field (green) at the lung boundary (red) in the coronal plane.

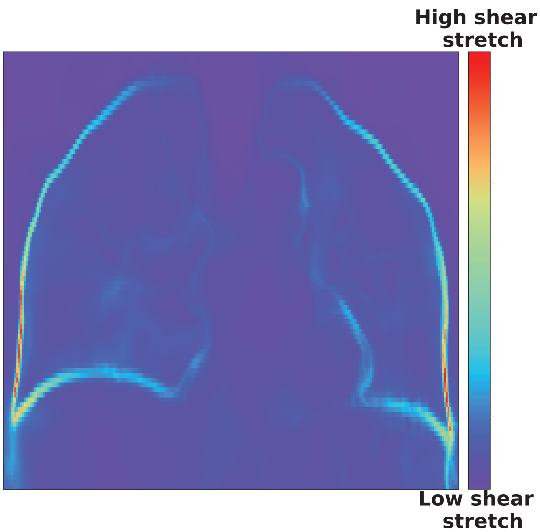

FIGURE 4. Qualitative average map of γ_{max} in the coronal plane. Red indicates high stretch, blue low shear stretch values.

deformable registration for both lung and thoracic cavity fields is 1220.805 ± 555.5444 seconds on a Xeon CPU @ 3.60GHz x12 with an GPU implementation on a GeForce RTX 3090.

C. LOBAR SLIDING

A different setup is chosen for the intra-collision correction. The object of interest is split into two parts with every part registered independently with its own velocity field. The fields are coupled via the ADMM at the shared boundary with Lagrangian coupling constraints (Equation 23). The minimization problems in the iteration updates are again solved by a gradient descent approach with backtracking. A simplex-mesh representing the whole object is created and its vertices assigned to either one of the velocity fields depending on their initial position. Deforming the vertices with different velocity fields ensures that those parts can deform freely while matching the deformation at their shared boundary. Aiming for FFD behaviour, both λ and μ in the linear operator are set to 0, the velocity field is discretized to one timesteps and integrated with the forward Euler method.

1) SYNTHETIC SELF-INTERSECTION DATA

The 2D data used for the exemplary experiments is shown in Figure 5. The setup is specifically designed to highlight the challenges posed by self-sliding contact within lung fissures. To simulate this, we constructed a pronged object of 40×60 pixels, representing the invaginated structure of a fissure. In the moving image, the prongs begin in a non-contact configuration. In the fixed domain, however, they come into contact and move in opposite directions along their shared interface, producing a sliding motion pattern similar to that observed within actual fissures. Each prong is assigned a unique internal intensity pattern, allowing them to be distinguished from one another during registration. Three experiments were run: One-Field-FFD, Piecewise-FFD, and sCC-DIR. In One-Field-FFD the whole domain was deformed by a single velocity field. For both Piecewise-FFD and sCC-DIR the image domain is virtually segmented into two velocity fields, one covering the upper prong and one covering the lower prong. The point set x_c coupling the two fields via Lagrangian constraints is the shared boundary connecting the two prongs. sCC-DIR is additionally equipped with intra-collision correction for a triangle mesh representing the two prongs like in Equation 23. This complex formulation captures the piecewise deformation of the upper and lower prongs, denoted Ω^m and Ω^l, respectively, and their individual updates to the object's representations based on the initial domains. The adapted signed distance field

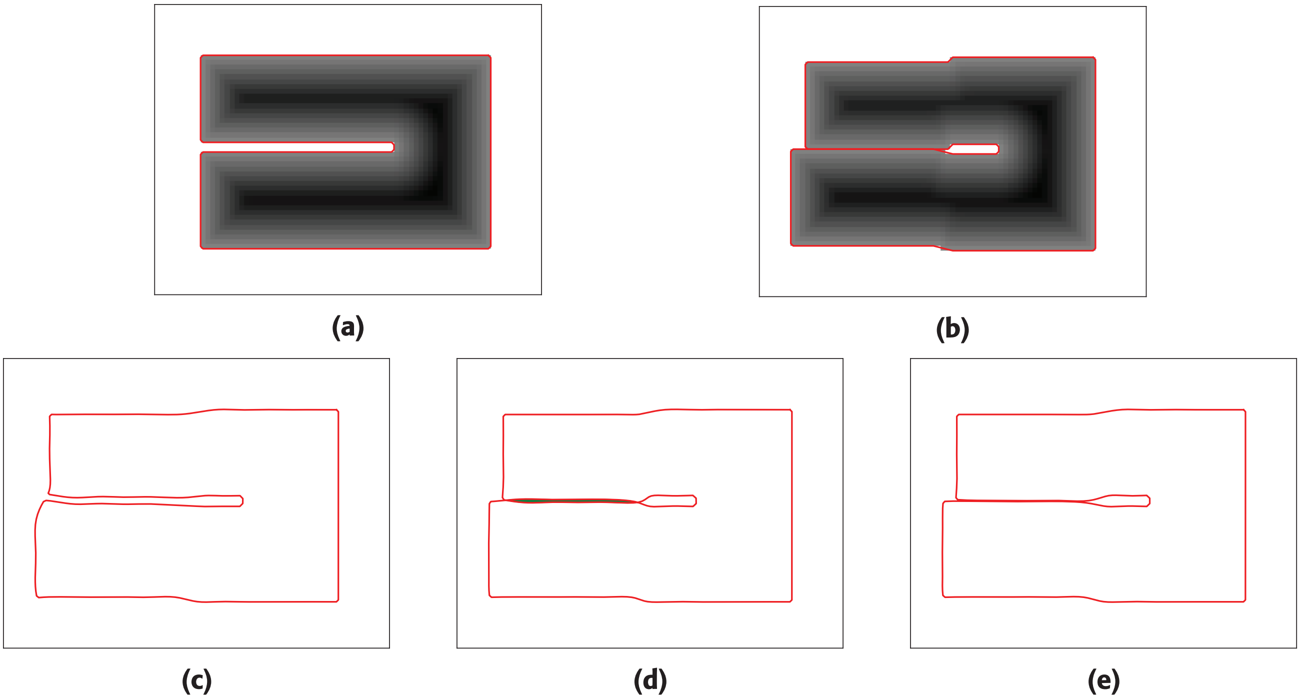

FIGURE 5. Results for synthetic self-intersection experiments. (a) Moving image with an internal gradient structure. (b) Fixed image with and internal gradient structure. (c) Registered contour for one-field-FFD. (d) Registered contour for Piecewise-FFD. (e) Registered contour for sCC-DIR.

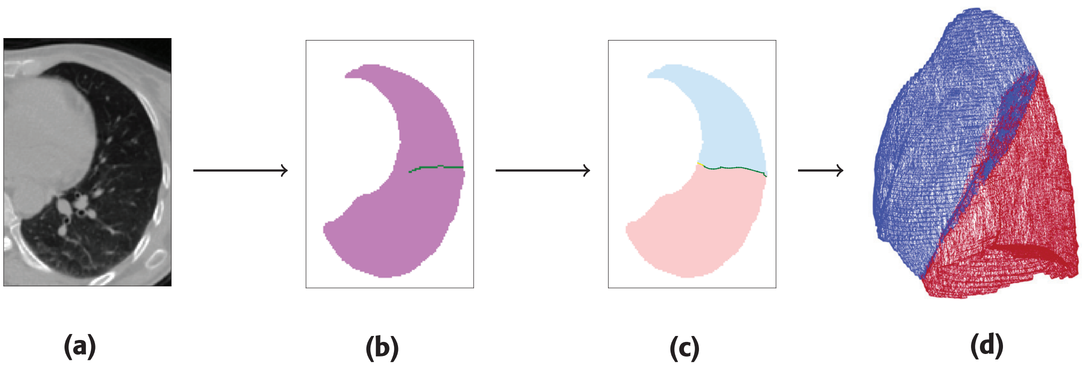

FIGURE 6. Exemplary fissure reconstruction on left lung form transverse plane. (a) Windowed original scan. (b) Original lung segmentation in purple overlaid with original sparse fissure segmentation in green. (c) Results for lobe segmentations in red (inferior) and blue (superior) with reconstructed fissure in green and connector points in yellow. (d) Wireframe rendering of 3D triangle mesh of lung partially split into inferior (red) and superior (blue) lobes.

is created from the object in the initial configuration via a marching method. The prongs are registered at two scale with 10 and 5 pixel resolutions. The resulting deformed contours in Figure 5 show a failure of One-Field-FFD to capture the self-contact between the prongs. While Piecewise-FFD and sCC-DIR create that contact, the former overshoots and generates an overlap of around 15 pixel whereas the later prevents any overlap. All three methods reach a DICE of 0.9881. The mean difference $|x_c \circ \Phi^{upper} - x_c \circ \Phi^{lower}|_2$ for the ADMM methods is 0.09 pixel.

2) FISSURE RECONSTRUCTION

To examine lobar sliding in a medical setting the lung segmentation needs to be refined. Both left and right lung are separated into lobes by lung fissures. Where the left lung is split into superior and inferior lobes by the oblique fissure, the right lung is divided into superior, middle, and inferior lobes by a oblique and a horizontal fissure. [88] Automated fissure extraction, and subsequently lobe segmentation based on fissures, has struggled for a robust method for a long time [89]. In order to combat the low intensity contrast between fissure and surrounding volume, either preprocessing fissure enhancements are employed [90], [91] or anatomical ROIs found [92]. However, deep learning approaches trained on diseased lungs have recently shown promising results [93]. We have obtained a sparse, manual fissure segmentation of the oblique fissures for all DIR-Lab cases excluding case 8. The goal of the fissure reconstruction is twofold: first, to split

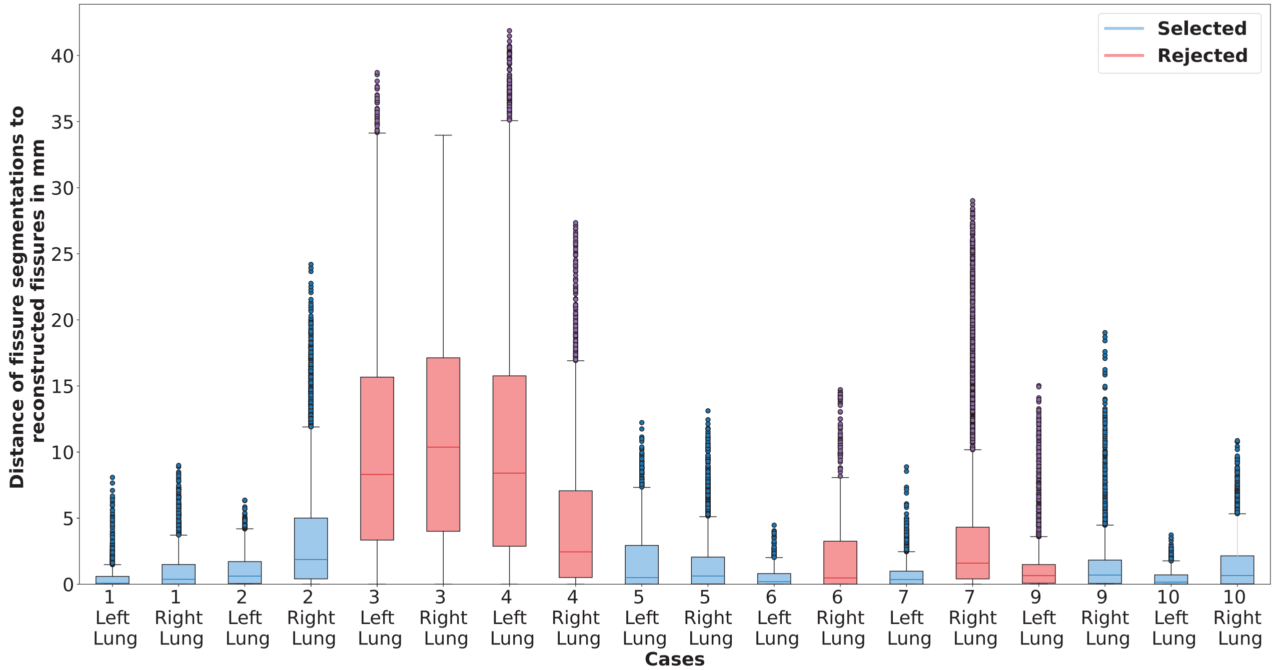

(a)

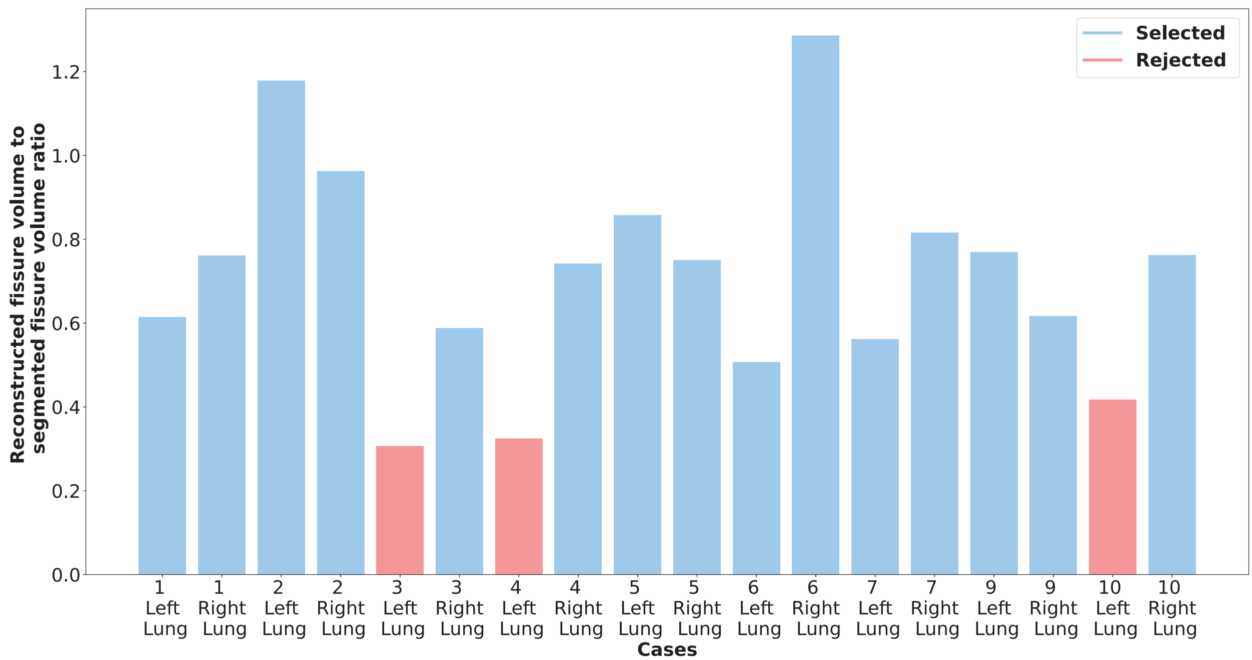

(b)

FIGURE 7. Quality assessment of the fissure reconstruction. Blue cases fit the selection criteria, red ones are rejected. (a) Boxplot for quartiles of distances between original sparse fissure segmentation and the reconstructed mesh. Left lungs for cases 3, 4, and 7, and right lungs for cases 3, 4, 6, and 9 are rejected. (b) Barplot of the volume ratio between reconstructed fissure mesh and original sparse segmentation. Left lungs for case 3 and rights lungs for cases 4 and 10 are rejected.

the lung segmentations into superior and inferior lobes along the fissure plane. Second, to generate a 3D tetrahedron mesh for self-intersection detection that represents the lung lobes split only at the segmentation points in the fissure plane while staying connected as one mesh. In order to do so, a continuous surface for each fissure is approximated and a volumetric fissure mesh generated. Our method is built on the eigenanalysis of our sparse fissure segmentation and consists of 5 consecutive steps. First, the regression plane of each fissure segmentation point cloud is found by using the two biggest eigenvectors representing the orientations with the highest variances in our dataset to span the plane. Next, the fissure points are projected onto the regression plane, filled in with sampled points and Delaunay triangulated. Afterwards, the in-plane fissure points are registered back to their original coordinates via a landmark matching based, smooth registration. The transformation is applied to all in-plane points, creating a smooth fissure surface in 3D. This smooth surface is used to separate lungs into lobes by identifying the closest triangle for each lung segmentation voxel and determining on which side of the triangle the voxel lies. For Lagrangian ADMM coupling constraints, the filler points are assigned to the set of shared tissue points x_c. Last, the set of fissure points is duplicated and both sets are extruded in opposing directions to generate a volumetric fissure mesh. The direction of this extrusion is the positive, respectively negative average of all triangle normals at the triangulation vertices corresponding with the fissure points. To generate a lung surface mesh with fissures excluded, a boolean mesh subtraction is performed. A graphic example is shown in Figure 6. Reconstructing smooth fissures that match the original segmentation is time consuming and not the focus of this work. Subsequently, instead of individually tuning parameters, only the subset of lungs with decent reconstructions from our out-of-the-box solution is selected. Two quality measures were employed to select this subset, distance error and a volume ratio. While the distance error measures the distance distribution from the sparse segmentation to the reconstructed mesh, the volume ratio displays how much the reconstructed mesh volume diverges from the original volume of the sparse segmentation. Acceptance criteria for the distance measure were formulated as below 10% of outliers, an inter-quartile range below 5 mm and a median below 2 mm. The criteria for the volume ratio is a ratio between 0.5 and 1.5. The results of the quality measures is displayed in Figure 7 Lungs fitting both criteria and consequently selected for experiments are both lungs of cases 1,2, and 5, left lungs only for cases 6 and 7, and right lung only for cases 9 and 10.

3) MEDICAL DATA

For all the selected cases One-Field-FFD, Piecewise-FFD, and sCC-DIR are run. The One-Field-FFD deformes the lung with one velocity field. For the piecewise approach and sCC-DIR, each lobe is deformed by its own velocity field with the fields being coupled via Lagrangian constraints at the set x_c described in the last paragraph. The tetrahedral mesh is obtained from the triangle surface mesh excluding the fissure volumes and its corresponding adapted signed distance field via a winding number method. Each lobe registration starts with an affine registration followed by deformable registration at 10 and 5 voxel resolution. The TRE results for all methods as well as the overlap results for the piecewise methods are given in Table 4. While the One-Field-FFD shows the highest average accuracy, the piecewise approaches follow with only 0.06 mm and 0.09 mm differences. In half of the cases the One-Field-FFD is outperformed by the piecewise methods. Regarding the overlap, Piecewise-FFD shows an average overlap of 194 mm^3 whereas sCC-DIR reduces this overlap by factor 7. Self-intersection, whether in Piecewise-FFD or sCC-DIR, occur in an average distance of 0.5336 ± 0.503 mm around their respective fissures. Figure 9 displays the mean maximum shear stretch in the oblique fissure. Piecewise-FFD and sCC-DIR results are within a 19% range of each other, while One-Field-FFD values are between 43% and 71% of the lowest piecewise results. Lobar shear modeled with sCC-DIR compared to shear in lung-thoracic cavity is in a range of 44 and 80 %. A visual inspection of sliding in the lobar fissure is shown in Figure 8. The colour coded shear map shows a left lung after registration with One-Field-FFD and Piecewise-FFD. While both cases show expectedly high shear in the lung-thorax boundary, the Piecewise-FFD additionally experiences an increase in lobar shear with the fissure location clearly distinguishable from the background. When plotting the tangential component parallel to the fissure surface of both deformations, only the blue Piecewise-FFD shows a switch in direction when crossing the boundary, indicating sliding of the lobes. The error for the filler points $|x_c \circ \Phi^{inferior} - x_c \circ \Phi^{superior}|_2$ is on average 2.0333 ± 1.8412 mm for Piecewise-FFD and 1.9877 ± 1.7679 for sCC-DIR. The average time for the registration process with affine and multi-scale deformable registration for both lung lobes is 1111.2431 ± 435.6455 seconds on the same specifications.

IV. DISCUSSION

Our approach models sliding motion between the lung and thoracic cage by employing domain-specific deformation fields that naturally allow discontinuous displacement at domain boundaries. This leads to registration accuracy, measured by Target Registration Error (TRE), that matches recent state-of-the-art benchmarks. Image registration over time can be viewed as simulating organ motion, where the image data discretizes the underlying organ shapes. Many existing domain-splitting methods handle sliding by approximating it within the image domain, separating regions, and applying smoothing at interfaces. In contrast, our method integrates collision detection as an explicit, physics-based constraint. This directly prevents tissue overlap or penetration without relying on heuristic smoothing or auxiliary loss functions that may only minimize interface inconsistencies without truly resolving collisions, resulting in more accurate and realistic sliding behavior. Moreover, collision detection is highly flexible and can handle complex, non-planar sliding interfaces, whereas domain-splitting methods often depend on simpler geometric assumptions

TABLE 4. Combined results for TREs and overlap for the self-intersection medical experiments. TREs in *mm* are shown for no registration, One-Field-FFD, Piecewise-FFD and self-intersection-correcting sCC-DIR, overlap in *mm* for Piecewise-FFD and sCC-DIR. Regarding TREs, one-field-FFD has results in the highest accuracy followed closely by the sliding-modeling methods. Overlapping volume are clearly reduced with the sCC-DIR.

		TRE in mm				Overlap in mm^3	
Case	**Lung**	**No registration**	**One-Field-FFD**	**Piecewise-FFD**	**sCC-DIR**	**Piecewise-FFD**	**sCC-DIR**
1	Right	3.79 ± 2.82	$\mathbf{0.94 \pm 0.96}$	1.05 ± 1.04	1.04 ± 1.03	9.41	**0.0**
1	Left	3.92 ± 2.70	$\mathbf{0.97 \pm 0.96}$	1.07 ± 1.02	1.1 ± 1.02	4.70	**0.0**
2	Right	4.33 ± 4	0.83 ± 0.95	$\mathbf{0.8 \pm 0.98}$	0.86 ± 0.97	26.91	**3.56**
2	Left	4.34 ± 3.81	0.96 ± 1.02	$\mathbf{0.94 \pm 1.03}$	0.97 ± 1.02	57.19	**3.36**
5	Right	9.29 ± 6.02	$\mathbf{1.91 \pm 2.30}$	1.96 ± 2.46	1.92 ± 2.47	686.68	**81.68**
5	Left	5.26 ± 3.79	$\mathbf{1.34 \pm 1.33}$	1.45 ± 1.37	1.52 ± 1.41	9.08	**3.03**
6	Left	9.12 ± 5.71	$\mathbf{1.99 \pm 1.65}$	2.33 ± 2.18	2.46 ± 2.48	188.18	**21.17**
7	Left	11.79 ± 6.74	2 ± 1.55	$\mathbf{1.99 \pm 1.87}$	2.06 ± 1.75	458.69	**56.45**
9	Right	7.15 ± 3.76	$\mathbf{1.29 \pm 1.01}$	1.35 ± 1.05	$\mathbf{1.29 \pm 1.05}$	214.05	**37.64**
10	Right	7.56 ± 7.3	1.66 ± 21	$\mathbf{1.5 \pm 1.85}$	1.62 ± 1.95	289.33	**58.81**
Mean	-	6.69 ± 5.6	$\mathbf{1.39 \pm 1.53}$	1.44 ± 1.65	1.48 ± 1.7	194.42 +- 217.5	**26.55 +- 28.57**

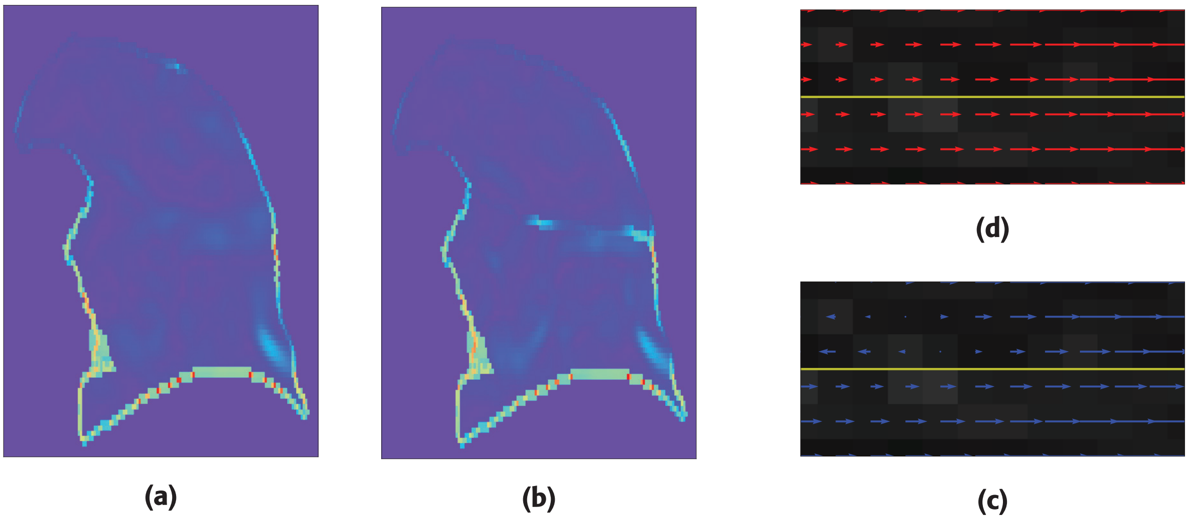

FIGURE 8. Coronal view of maximum shear stretch and sliding motion for case 7, left lung. (a) Qualitative shear map on One-Field-FFD. Brighter colour indicates higher stretch. (b) Qualitative shear map on Piecewise-FFD with more pronounced stretch along the fissure. (c) Tangential components of Φ for One-Field-FFD. Components are tangential to fissure surface (yellow). (d) Tangential components of Φ for Piecewise-FFD with direction switch over the fissure surface.

that restrict their applicability. Splitting the deformation into multiple fields also enables straightforward application of the Large Deformation Diffeomorphic Metric Mapping (LDDMM) framework within each subdomain, preserving its theoretical validity. By enforcing collision constraints between these deformation fields, our method effectively minimizes overlap. Even under diffeomorphic constraints, the amount of overlap is significantly lower compared to other split-domain algorithms. Additionally, our sampling-based collision detection, combined with signed distance fields, allows for efficient inside-outside checks via interpolation. Although the theoretical computational complexity is O(n) for n sampled points, the use of vectorized interpolation reduces this to approximately O(1) in practice, ensuring computational efficiency. This diffeomorphic nature of the deformation adds smoothness constraints to the velocity fields, influencing the local adaptability of the displacements. One result of this impacted flexibility compared to non-diffeomorphic methods is seen at the domain boundaries with an increased gap between lungs and thoracic cage.

While the TRE results of the diffeomorphic method fall well within the range of current approaches benchmarked on the DIR-Lab dataset, the performance in terms of overlap and gap may initially appear reduced compared to its non-diffeomorphic predecessor [58]. However, a closer inspection reveals that only case 8 exhibits a gap value outside the range observed in other methods. Notably, this is also the case with the highest average displacement across the dataset, reaching 15.16 mm. Such large displacements tend to induce strong penalization in the velocity field, effectively increasing the amount of regularization or smoothing. This increased smoothing can hinder the model's ability to adapt precisely to the complex interface between lung and thoracic structures, allowing the collision restoration forces to dominate. Over time, however, these penalization effects may diminish, and extending the runtime could allow the model to better adapt

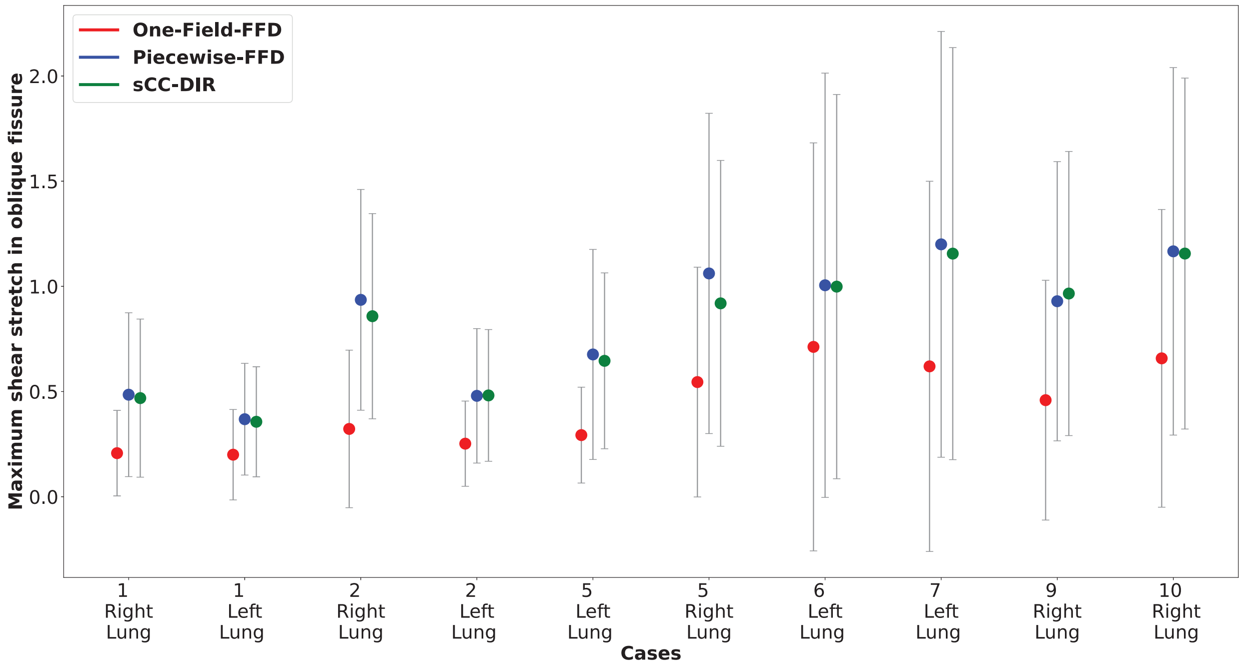

FIGURE 9. Mean maximum shear stretch in oblique fissure over registered lungs. one-field-FFD (red) shows in all cases a lower stretch while Piecewise-FFD and sCC-DIR exhibit the same behaviour.

and reduce the observed gap, potentially improving the registration quality further. A solution can lie in the adaption of the coupling constraints employed in the ADMM. When using the Eulerian formulation, the velocity field solving the registration sub-problem is heavily smoothed by the velocity field for the collision sub-problem. The collision field is only locally subject to deformation forces and prevents displacements with high magnitude or acceleration from the registration field. It can be seen as a time dependent damper or in the context of fluid mechanics the deformation fields viscosity. Two possible adaptions would be to use the Lagrangian coupling or a local relaxation of the coupling coefficient ρ.

While the diffeomorphic nature of our registration method greatly enhances anatomical plausibility, it also significantly impacts computation time. Our average runtime of 1220.8 seconds per case may initially seem high, but this must be contextualized in terms of model complexity and the substantial size of the DIR-Lab dataset. Achieving high accuracy in this setting requires dense sampling and a large number of query points, which inherently increases computational load. Among classical, numerically optimized approaches benchmarked on DIR-Lab, Wu et al. [37] and Delmon et al. [42] report runtimes of approximately 100 minutes and 86 minutes respectively—both of which are outperformed by our implementation. However, these comparisons must account for advances in computational hardware over time, which can significantly affect runtime. Learning-based methods, by contrast, dramatically reduce computation time. For example, recent approaches by Huang et al. [94], Ng and Ebrahimi [44], Zhang et al. [95], Hering et al. [50], and De Vos et al. [96] report runtimes in the range of 0.63 to 2.1 seconds. This runtimes do not take the respective training times into account. Despite their speed, these methods typically do not enforce explicit modeling of sliding motion at anatomical interfaces. In terms of target registration error (TRE), only Zhang and Hering outperform our method, while the others, although faster, do not achieve the same level of accuracy.

A. DIFFEOMORPHIC BEHAVIOR

Evaluating the diffeomorphic qualities of the transformation, proved to be less straight forward than TRE or overlap measures since none of the algorithms in Table 1 and Table 2 claimed diffeomorphic behaviour. Comparison for percentage of $det(\nabla\Phi) \leq 0$ was then performed across a spectrum of publications working on the DIR-Lab dataset. On average, Lu et al. [97] report 0.0061%, Matkovic et al. [98] 0.453 $\pm$ 0.818%, Iqbal et al. [99] 127 voxel, Risser et al. [36] 0.147%, Zhu et al. [100] values between $4.60 * 10^{-4}$% and $4.14 * 10^{-4}$% for their three algorithms, and Lui et al. [101] between 0.026% and 0.045% for another three algorithms. Only Castillo et al. [102], two out of the seven methods reported with values for $det(\nabla\Phi)$ in Ho et al. [103] and one out of the seven methods from Van der Ouderaa et al. [104] display no non positive values. With less then 20% of the methods mentioned showing this local invertibility, diffeomorphic behavior is can not be automatically assumed. Less works asses the inverse consistency. Heinrich [105] gives an overview of five methods applied to a sub-set of the DIR-Lab set, with an average inverse consistency of below 0.2 voxels. Slightly more references exist on brain image registration, where Johnson and Christensen [106] produces an inverse consistency of 0.00 to 0.009 pixel and Ashburner [107] between 0.034 and 0.032 mm. The American Association of Physicists in Medicine Task Group 132 as reported by Nenoff et al. [108] recommends a 2-3 mm inverse consistency error threshold. Being one of the few algorithms showing local invertibility and an inverse consistency below the task group recommendations, dCC-DIR can be considered state of the art.

B. LOBAR SLIDING

The existence of lobar sliding and the importance of modeling has been shown in multiple works [52], [53], [54], [56], [57]. These contributions approximate the lobar deformation by using the same lobe-by-lobe registration sCC-DIR employs. When modeling contact constraints however, the lobes are represented as individual object, whereas in reality they form a continuum with a connection at the hilum. The development of lobar sliding via self-collision detection is to the best of our knowledge the first model using a single object representation for lobe-separated lungs. In contrast to inter-collision detection, intra-collision detection is more computationally expensive. This is because each of the nn collision query points must be checked against all mm tetrahedra, resulting in a theoretical complexity of $O(m^2 \times n^2)$. While vectorization reduces this cost in practice, the point-to-tetrahedron tests remain computationally intensive. To further accelerate the process, strategies such as bounding volume hierarchies (BVH) can be employed, potentially providing significant speed-ups. The overlap in lobes without collision correction is on the mm^3 level, meaning a factor of $\frac{1}{1000}$ compared to the overlaps at the lung-thorax interface. Nevertheless, correcting the collisions produces consistent deformations while introducing no further drawbacks as the TRE accuracy on average is kept within 0.04 mm difference to Piecewise-FFD and 0.1 mm difference to dCC-DIR. A key factor for preventing overlaps is an authentic fissure reconstruction, yet the reconstruction is a difficult task, since a fine geometry has to be created from sparse segmentation. Several publicly available fissure segmentation algorithms were tested but could not produce reasonable results in a timely manner. As other published works [90], [91], our own fissure reconstruction algorithm uses approximations and manual corrections, possibly influencing the outcome with user bias. Improving the fissure reconstruction will likely improve performance and modeling right lungs with all three lobes poses an interesting challenge. While the self-intersection was developed for lung lobe interactions, other highly elastic or convex organs can benefit from our model approach. Overall, there are two major possible opportunities for future work. The first is the evident coupling of both inter-collision and intra-collision correction to achieve a fully physics-constraint lung model. This might be coupled to FEA to either provide simulations with accurate displacement mappings or use inspiration in sliding contact modeling. Adapting the variational approach to a deep learning method is the second option. The model can either be fully translated into a network loss function or a network trained purely on image registration can supply initial guesses for the deformation fields, reducing the effects of time dependent smoothing. Even though sliding motion at the lung boundaries has been substantially researched, no single standard has been established. This is apparent in the comparison of the tradeoffs of two main approaches, split domain and local regularization. While the former method ensures sliding motion by using independent deformation fields, matching of the interfaces is not guaranteed. The local regularization method suffers from the opposite dilemma.

V. CONCLUSION

We have presented a method capable of modeling discontinuities such as sliding motion, either between different objects or the same object. Additionally, diffeomorphic behaviour was examined and confirmed. The registration success of our method was evaluated on synthetic and medical data, performing within the range of state of the art algorithms. This work is a contribution to improving medical image registration with physical constraints and can be the foundation to develop a more sophisticated representation of lung deformations.

ACKNOWLEDGMENT

The author's want to express their deep gratitude to Albert Alonso with his insights on integration. This article only contains the RAINBOW consortium's views and the Research Executive Agency and the Commission are not responsible for any use that may be made of the information it contains. The funders had no role in study design, data collection and analysis, decision to publish, and preparation of the manuscript.

REFERENCES

[1] H. S. Sanjay and M. Niranjanamurthy, *Medical Imaging*. Hoboken, NJ, USA: Wiley, 2023.

[2] B. Grignon, G. Oldrini, and F. Walter, ''Teaching medical anatomy: What is the role of imaging today?'' *Surgical Radiol. Anatomy*, vol. 38, no. 2, pp. 253–260, Mar. 2016.

[3] J. Franco, ''The role of medical imaging in clinical research,'' *Radiol. Technol.*, vol. 90, no. 4, pp. 353–364, Mar. 2019.

[4] J. Modersitzki, *Numerical Methods for Image Registration*. London, U.K.: Oxford Univ. Press, Dec. 2003, doi: 10.1093/acprof:oso/9780198528418.001.0001.

[5] J. Zou, B. Gao, Y. Song, and J. Qin, ''A review of deep learning-based deformable medical image registration,'' *Frontiers Oncol.*, vol. 12, Dec. 2022, Art. no. 1047215.

[6] T. De Silva, A. Uneri, M. D. Ketcha, S. Reaungamornrat, G. Kleinszig, S. Vogt, N. Aygun, S.-F. Lo, J.-P. Wolinsky, and J. H. Siewerdsen, ''3D–2D image registration for target localization in spine surgery: Investigation of similarity metrics providing robustness to content mismatch,'' *Phys. Med. Biol.*, vol. 61, no. 8, pp. 3009–3025, Apr. 2016.

[7] H. Kaur, D. Koundal, and V. Kadyan, ''Image fusion techniques: A survey,'' *Arch. Comput. Methods Eng.*, vol. 28, no. 7, pp. 4425–4447, Jan. 2021.

[8] A. Sotiras, C. Davatzikos, and N. Paragios, "Deformable medical image registration: A survey," *IEEE Trans. Med. Imag.*, vol. 32, no. 7, pp. 1153–1190, Jul. 2013.

[9] B. Zitová and J. Flusser, "Image registration methods: A survey," *Image Vis. Comput.*, vol. 21, no. 11, pp. 977–1000, Oct. 2003. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0262885603001379

[10] Y. Fu, Y. Lei, T. Wang, W. J. Curran, T. Liu, and X. Yang, "Deep learning in medical image registration: A review," *Phys. Med. Biol.*, vol. 65, no. 20, Mar. 2020, Art. no. 20TR01.

[11] H. R. Boveiri, R. Khayami, R. Javidan, and A. Mehdizadeh, "Medical image registration using deep neural networks: A comprehensive review," *Comput. Electr. Eng.*, vol. 87, Oct. 2020, Art. no. 106767.

[12] J. S. Suri, D. L. D. L. Wilson, and S. Laxminarayan, *Registration Models* (Biomedical engineering international book series). New York, NY, USA: Kluwer Academic, 2005.

[13] J. M. H. Noothout, B. D. De Vos, J. M. Wolterink, E. M. Postma, P. A. M. Smeets, R. A. P. Takx, T. Leiner, M. A. Viergever, and I. Išgum, "Deep learning-based regression and classification for automatic landmark localization in medical images," *IEEE Trans. Med. Imag.*, vol. 39, no. 12, pp. 4011–4022, Dec. 2020.

[14] J. N. Sarvaiya, S. Patnaik, and S. Bombaywala, "Image registration by template matching using normalized cross-correlation," in *Proc. Int. Conf. Adv. Comput., Control, Telecommun. Technol.*, Dec. 2009, pp. 819–822.

[15] F. Maes, D. Vandermeulen, and P. Suetens, "Medical image registration using mutual information," *Proc. IEEE*, vol. 91, no. 10, pp. 1699–1722, Oct. 2003.

[16] H. Jiang, K. S. Holton, and R. A. Robb, "Image registration of multimodality 3D medical images by chamfer matching," *Proc. SPIE*, vol. 1660, pp. 356–366, Jun. 1992.

[17] N. Chumchob and K. Chen, "A robust affine image registration method," *Int. J. Numer. Anal. Model.*, vol. 6, no. 2, pp. 1–24, 2009.

[18] J.-P. Thirion, "Image matching as a diffusion process: An analogy with Maxwell's demons," *Med. Image Anal.*, vol. 2, no. 3, pp. 243–260, Sep. 1998.

[19] D. Rueckert, L. I. Sonoda, C. Hayes, D. L. G. Hill, M. O. Leach, and D. J. Hawkes, "Nonrigid registration using free-form deformations: Application to breast MR images," *IEEE Trans. Med. Imag.*, vol. 18, no. 8, pp. 712–721, Aug. 1999.

[20] J. Modersitzki, *Numerical Methods for Image Registration*. Oxford, U.K.: OUP Oxford, 2003.

[21] G. E. Christensen, R. D. Rabbitt, and M. I. Miller, "Deformable templates using large deformation kinematics," *IEEE Trans. Image Process.*, vol. 5, no. 10, pp. 1435–1447, Oct. 1996.

[22] D. Wilson and S. Laxminarayan, *Handbook of Biomedical Image Analysis: Registration Models*, vol. 3. Cham, Switzerland: Springer, 2007.

[23] M. Bro-Nielsen and C. Gramkow, "Fast fluid registration of medical images," in *Proc. Int. Conf. Visualizat. Biomed. Comput.*, Jan. 1996, pp. 265–276.

[24] P. Dupuis, U. Grenander, and M. I. Miller, "Variational problems on flows of diffeomorphisms for image matching," *Quart. Appl. Math.*, vol. 56, no. 3, pp. 587–600, Sep. 1998.

[25] M. I. Miller, A. Trouvé, and L. Younes, "On the metrics and Euler–Lagrange equations of computational anatomy," *Annu. Rev. Biomed. Eng.*, vol. 4, no. 1, pp. 375–405, Aug. 2002.

[26] H. Xiao, X. Xue, Z. Mi, X. Jiang, Q. Xia, K. Chen, H. Li, L. Long, and K. Peng, "Deep learning-based lung image registration: A review," *Comput. Biol. Med.*, vol. 165, Sep. 2023, Art. no. 107434.

[27] L. Goodman and D. Baruah, "Lung nodules attached to the pleura: Insights from lung cancer screening," *Radiology*, vol. 310, no. 1, Jan. 2024, Art. no. 233290.

[28] Y. Zhu, R. Yip, N. You, C. I. Henschke, and D. F. Yankelevitz, "Management of nodules attached to the costal pleura at low-dose CT screening for lung cancer," *Radiol.*, vol. 297, no. 3, pp. 710–718, Oct. 2020.

[29] A. R. Larici, A. Farchione, P. Franchi, M. Ciliberto, G. Cicchetti, L. Calandriello, A. del Ciello, and L. Bonomo, "Lung nodules: Size still matters," *Eur. Respiratory Rev.*, vol. 26, no. 146, Dec. 2017, Art. no. 170025.

[30] A. El-Baz, M. Nitzken, A. Elnakib, F. Khalifa, G. Gimel'farb, R. Falk, and M. A. El-Ghar, "3D shape analysis for early diagnosis of malignant lung nodules," in *Proc. 22nd Int. Conf. Inf. Process. Med. Imag.*, Kloster Irsee, Germany. Cham, Switzerland: Springer, Jan. 2011, pp. 175–182.

[31] M. P. Heinrich and B. W. Papież, "Image registration with sliding motion," in *Handbook of Medical Image Computing and Computer Assisted Intervention*. Amsterdam, The Netherlands: Elsevier, 2020, pp. 293–318.

[32] B. W. Papiez, M. P. Heinrich, J. Fehrenbach, L. Risser, and J. A. Schnabel, "An implicit sliding-motion preserving regularisation via bilateral filtering for deformable image registration," *Med. Image Anal.*, vol. 18, no. 8, pp. 1299–1311, Dec. 2014.

[33] V. Vishnevskiy, T. Gass, G. Szekely, C. Tanner, and O. Goksel, "Isotropic total variation regularization of displacements in parametric image registration," *IEEE Trans. Med. Imag.*, vol. 36, no. 2, pp. 385–395, Feb. 2017.

[34] L. Gong, L. Duan, Y. Dai, Q. He, S. Zuo, T. Fu, X. Yang, and J. Zheng, "Locally adaptive total p-Variation regularization for non-rigid image registration with sliding motion," *IEEE Trans. Biomed. Eng.*, vol. 67, no. 9, pp. 2560–2571, Sep. 2020.

[35] A. Schmidt-Richberg, R. Werner, H. Handels, and J. Ehrhardt, "Estimation of slipping organ motion by registration with direction-dependent regularization," *Med. Image Anal.*, vol. 16, no. 1, pp. 150–159, Jan. 2012.

[36] L. Risser, M. P. Heinrich, T. Matin, and J. A. Schnabel, "Piecewise-diffeomorphic registration of 3D CT/MR pulmonary images with sliding conditions," in *Proc. 9th IEEE Int. Symp. Biomed. Imag. (ISBI)*, May 2012, pp. 1351–1354.

[37] Z. Wu, E. Rietzel, V. Boldea, D. Sarrut, and G. C. Sharp, "Evaluation of deformable registration of patient lung 4DCT with subanatomical region segmentations," *Med. Phys.*, vol. 35, no. 2, pp. 775–781, Jan. 2008.

[38] B. Eiben, H. E. Tran, M. J. Menten, U. Oelfke, D. J. Hawkes, and J. R. McClelland, "Statistical motion mask and sliding registration," in *Proc. 8th Int. Workshop Biomed. Image Registration*, Leiden, Netherlands. Cham, Switzerland: Springer, Jan. 2018, pp. 13–23.

[39] F. F. Berendsen, A. N. T. J. Kotte, M. A. Viergever, and J. P. W. Pluim, "Registration of organs with sliding interfaces and changing topologies," *Proc. SPIE*, vol. 9034, pp. 95–101, Mar. 2014.

[40] S. Heldmann, T. Polzin, A. Derksen, and B. Berkels, "An image registration framework for sliding motion with piecewise smooth deformations," in *Proc. 5th Int. Conf. Scale Space Variational Methods Comput. Vis.*, Lège-Cap Ferret, France. Cham, Switzerland: Springer, Jan. 2015, pp. 335–347.

[41] A. Derksen, S. Heldmann, T. Polzin, and B. Berkels, "Image registration with sliding motion constraints for 4D CT motion correction," in *Bildverarbeitung Für Die Medizin 2015: Algorithmen-Systeme-Anwendungen. Proceedings Des Workshops Vom 15. Bis 17. März 2015 in Lübeck*. Cham, Switzerland: Springer, Jan. 2015, pp. 335–340.

[42] V. Delmon, S. Rit, R. Pinho, and D. Sarrut, "Registration of sliding objects using direction dependent B-splines decomposition," *Phys. Med. Biol.*, vol. 58, no. 5, pp. 1303–1314, Mar. 2013.

[43] R. Hua, J. M. Pozo, Z. A. Taylor, and A. F. Frangi, "Multiresolution eXtended free-form deformations (XFFD) for non-rigid registration with discontinuous transforms," *Med. Image Anal.*, vol. 36, pp. 113–122, Feb. 2017.

[44] E. Ng and M. Ebrahimi, "An unsupervised learning approach to discontinuity-preserving image registration," in *Proc. 9th Int. Workshop Biomed. Image Registration*, Portorož, Slovenia. Cham, Switzerland: Springer, Jan. 2020, pp. 153–162.

[45] B. Pandey, D. Kumar Pandey, B. Pratap Mishra, and W. Rhmann, "A comprehensive survey of deep learning in the field of medical imaging and medical natural language processing: Challenges and research directions," *J. King Saud Univ.-Comput. Inf. Sci.*, vol. 34, no. 8, pp. 5083–5099, Sep. 2022.

[46] S. Minaee, Y. Boykov, F. Porikli, A. Plaza, N. Kehtarnavaz, and D. Terzopoulos, "Image segmentation using deep learning: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 7, pp. 3523–3542, Jul. 2022.

[47] N. Tajbakhsh, L. Jeyaseelan, Q. Li, J. N. Chiang, Z. Wu, and X. Ding, "Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation," *Med. Image Anal.*, vol. 63, Jul. 2020, Art. no. 101693.

[48] T. Sentker, F. Madesta, and R. Werner, "GDL-FIRE^{4D}: Deep learning-based fast 4D CT image registration," in *Proc. Int. Conf. Med. image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, 2018, pp. 765–773.
[49] Y. Hu, M. Modat, E. Gibson, N. Ghavami, E. Bonmati, C. M. Moore, M. Emberton, J. A. Noble, D. C. Barratt, and T. Vercauteren, "Label-driven weakly-supervised learning for multimodal deformable image registration," in *Proc. IEEE 15th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2018, pp. 1070–1074.
[50] A. Hering, S. Häger, J. Moltz, N. Lessmann, S. Heldmann, and B. van Ginneken, "CNN-based lung CT registration with multiple anatomical constraints," *Med. Image Anal.*, vol. 72, Aug. 2021, Art. no. 102139.
[51] L. Duan, X. Ni, Q. Liu, L. Gong, G. Yuan, M. Li, X. Yang, T. Fu, and J. Zheng, "Unsupervised learning for deformable registration of thoracic CT and cone-beam CT based on multiscale features matching with spatially adaptive weighting," *Med. Phys.*, vol. 47, no. 11, pp. 5632–5647, Nov. 2020.
[52] K. Ding, Y. Yin, K. Cao, G. E. Christensen, C. Lin, E. A. Hoffman, and J. M. Reinhardt, "Evaluation of lobar biomechanics during respiration using image registration," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Cham, Switzerland: Springer, Jan. 2009, pp. 739–746.
[53] R. Amelon, K. Cao, K. Ding, G. E. Christensen, J. M. Reinhardt, and M. L. Raghavan, "Three-dimensional characterization of regional lung deformation," *J. Biomechanics*, vol. 44, no. 13, pp. 2489–2495, Sep. 2011.
[54] R. Amelon, K. Cao, J. M. Reinhardt, G. E. Christensen, and M. L. Raghavan, "Estimation of lung lobar sliding using image registration," *Proc. SPIE*, vol. 8317, pp. 416–423, Mar. 2012.
[55] R. E. Amelon, K. Cao, J. M. Reinhardt, G. E. Christensen, and M. L. Raghavan, "A measure for characterizing sliding on lung boundaries," *Ann. Biomed. Eng.*, vol. 42, no. 3, pp. 642–650, Mar. 2014.
[56] A. E. Galloy, R. E. Amelon, J. M. Reinhardt, and M. L. Raghavan, "Contact mechanics model of lung lobar sliding," *Appl. Eng. Sci.*, vol. 10, Jun. 2022, Art. no. 100098.
[57] A. E. Galloy, J. M. Reinhardt, and M. L. Raghavan, "Role of lung lobar sliding on parenchymal distortion during breathing," *J. Appl. Physiol.*, vol. 135, no. 3, pp. 534–541, Sep. 2023.
[58] T. Alscher, K. Erleben, and S. Darkner, "Collision-constrained deformable image registration framework for discontinuity management," *PLoS ONE*, vol. 18, no. 8, Aug. 2023, Art. no. e0290243.
[59] F. P. M. Oliveira and J. M. R. S. Tavares, "Medical image registration: A review," *Comput. Methods Biomechanics Biomed. Eng.*, vol. 17, no. 2, pp. 73–93, Mar. 2012.
[60] E. Angelini, Y. Jin, and A. Laine, *Handbook of Biomedical Image Analysis-registration Models*, 2005.
[61] M. Burger, J. Modersitzki, and L. Ruthotto, "A hyperelastic regularization energy for image registration," *SIAM J. Sci. Comput.*, vol. 35, no. 1, pp. B132–B148, Jan. 2013.
[62] D. Hill, P. G. Batchelor, M. P. Holden, and D. J. Hawkes, "Medical image registration," *Phys. Med. Biol.*, vol. 46, no. 3, pp. R1–R45, Feb. 2001.
[63] S. Darkner, M. S. Hansen, R. Larsen, and M. F. Hansen, "Efficient hyperelastic regularization for registration," in *Proc. 17th Scandin. Conf. Image Anal.*, Ystad, Sweden. Cham, Switzerland: Springer, Jan. 2011, pp. 295–305.
[64] S. Boyd, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, 2010.
[65] H. Ji, Z. Zhang, P. Xue, M. Ren, and E. Dong, "Image registration method based on distributed alternating direction multipliers," *J. Med. Biol. Eng.*, vol. 44, no. 4, pp. 582–595, Aug. 2024.
[66] R. Sanyal, Sk. M. Ahmed, M. Jaiswal, and K. N. Chaudhury, "A scalable ADMM algorithm for rigid registration," *IEEE Signal Process. Lett.*, vol. 24, no. 10, pp. 1453–1457, Oct. 2017.
[67] E. Wei and A. Ozdaglar, "Distributed alternating direction method of multipliers," in *Proc. IEEE 51st IEEE Conf. Decis. Control (CDC)*, Dec. 2012, pp. 5445–5450.
[68] Y. Wang, W. Yin, and J. Zeng, "Global convergence of ADMM in nonconvex nonsmooth optimization," *J. Sci. Comput.*, vol. 78, no. 1, pp. 29–63, Jun. 2018.
[69] M. F. Beg, M. I. Miller, A. Trouvé, and L. Younes, "Computing large deformation metric mappings via geodesic flows of diffeomorphisms," *Int. J. Comput. Vis.*, vol. 61, no. 2, pp. 139–157, Feb. 2005.
[70] V. Arsigny, O. Commowick, X. Pennec, and N. Ayache, "A log-Euclidean framework for statistics on diffeomorphisms," in *Proc. 9th Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, Copenhagen, Denmark. Cham, Switzerland: Springer, Jan. 2006, pp. 924–931.
[71] B. B. Avants, P. T. Schoenemann, and J. C. Gee, "Lagrangian frame diffeomorphic image registration: Morphometric comparison of human and chimpanzee cortex," *Med. Image Anal.*, vol. 10, no. 3, pp. 397–412, Jun. 2006.
[72] U. Grenander and M. I. Miller, "Computational anatomy: An emerging discipline," *Quart. Appl. Math.*, vol. 56, no. 4, pp. 617–694, Dec. 1998.
[73] N. J. Tustison and B. B. Avants, "Explicit B-spline regularization in diffeomorphic image registration," *Frontiers Neuroinform.*, vol. 7, p. 39, Dec. 2013.
[74] M. L. Kessler, "Image registration and data fusion in radiation therapy," *Brit. J. Radiol.*, vol. 79, no. 1, pp. S99–S108, Sep. 2006.
[75] S. Lee, G. Wolberg, and S. Y. Shin, "Scattered data interpolation with multilevel B-splines," *IEEE Trans. Vis. Comput. Graphics*, vol. 3, no. 3, pp. 228–244, Aug. 1997.
[76] S. Smale and M. W. Hirsch, *Differential Equations, Dynamical Systems, and Linear Algebra*, vol. 60. Amsterdam, The Netherlands: Elsevier, 1974.
[77] S. Darkner, A. Pai, M. G. Liptrot, and J. Sporring, "Collocation for diffeomorphic deformations in medical image registration," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 7, pp. 1570–1583, Jul. 2018.
[78] D. F. Griffiths and D. J. Higham, *Numerical Methods for Ordinary Differential Equations: Initial Value Problems*, vol. 5. Cham, Switzerland: Springer, 2010.
[79] S. Rosloniec, *Fundamental Numerical Methods for Electrical Engineering*, vol. 18. Cham, Switzerland: Springer, 2008.
[80] A. Jacobson, L. Kavan, and O. Sorkine-Hornung, "Robust inside-outside segmentation using generalized winding numbers," *ACM Trans. Graph.*, vol. 32, no. 4, pp. 1–12, Jul. 2013.
[81] E. Castillo, R. Castillo, J. Martinez, M. Shenoy, and T. Guerrero, "Four-dimensional deformable image registration using trajectory modeling," *Phys. Med. Biol.*, vol. 55, no. 1, pp. 305–327, Jan. 2010.
[82] R. Castillo, E. Castillo, R. Guerra, V. E. Johnson, T. McPhail, A. K. Garg, and T. Guerrero, "A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets," *Phys. Med. Biol.*, vol. 54, no. 7, pp. 1849–1870, Apr. 2009.
[83] L. Younes, *Shapes and Diffeomorphisms*, vol. 171. Cham, Switzerland: Springer, 2010.
[84] G. E. Christensen and H. J. Johnson, "Consistent image registration," *IEEE Trans. Med. Imag.*, vol. 20, no. 7, pp. 568–582, Jul. 2001.
[85] J. Ulysses and A. Conci, "Measuring similarity in medical registration," in *Proc. 17th Int. Conf. Syst., Signals Image Process.*, 2010, pp. 1–4.
[86] M. Grabner, T. Pock, T. Gross, and B. Kainz, "Automatic differentiation for gpu-accelerated 2D/3D registration," in *Advances in Automatic Differentiation*. Cham, Switzerland: Springer, 2008, pp. 259–269.
[87] L. Armijo, "Minimization of functions having Lipschitz continuous first partial derivatives," *Pacific J. Math.*, vol. 16, no. 1, pp. 1–3, Jan. 1966.
[88] F. H. W. Martini, M. J. Timmons, and R. B. Tallitsch, *Human Anatomy* (Pearson Custom Library), 7th ed., Harlow, U.K.: Pearson, 2014.
[89] T. Doel, D. J. Gavaghan, and V. Grau, "Review of automatic pulmonary lobe segmentation methods from CT," *Computerized Med. Imag. Graph.*, vol. 40, pp. 13–29, Mar. 2015.
[90] J. Wang, M. Betke, and J. P. Ko, "Pulmonary fissure segmentation on CT," *Med. Image Anal.*, vol. 10, no. 4, pp. 530–547, Aug. 2006.
[91] B. Lassen, J.-M. Kuhnigk, O. Friman, S. Krass, and H.-O. Peitgen, "Automatic segmentation of lung lobes in CT images based on fissures, vessels, and bronchi," in *Proc. IEEE Int. Symp. Biomed. Imag., Nano Macro*, Apr. 2010, pp. 560–563.
[92] S. Ukil and J. M. Reinhardt, "Anatomy-guided lung lobe segmentation in X-ray CT images," *IEEE Trans. Med. Imag.*, vol. 28, no. 2, pp. 202–214, Feb. 2009.
[93] D. Carmo, J. Ribeiro, S. Dertkigil, S. Appenzeller, R. Lotufo, and L. Rittner, "A systematic review of automated segmentation methods and public datasets for the lung and its lobes and findings on computed tomography images," *Yearbook Med. Informat.*, vol. 31, no. 1, pp. 277–295, Aug. 2022.
[94] Q. Huang, L. Ren, T. Quan, M. Yang, H. Yuan, and K. Cao, "MA-VoxelMorph: Multi-scale attention-based VoxelMorph for nonrigid registration of thoracoabdominal CT images," *J. Innov. Opt. Health Sci.*, vol. 18, no. 1, pp. 1–18, Jan. 2025.

[95] Y. Zhang, X. Wu, H. M. Gach, H. Li, and D. Yang, "GroupRegNet: A groupwise one-shot deep learning-based 4D image registration method," *Phys. Med. Biol.*, vol. 66, no. 4, Feb. 2021, Art. no. 045030.

[96] B. D. de Vos, F. F. Berendsen, M. A. Viergever, H. Sokooti, M. Staring, and I. Isgum, "A deep learning framework for unsupervised affine and deformable image registration," *Med. Image Anal.*, vol. 52, pp. 128–143, Feb. 2019.

[97] J. Lu, R. Jin, and E. Song, "Pyramid convolutional recurrent network for serial medical image registration with adaptive motion regularizations," *IEEE Trans. Radiat. Plasma Med. Sci.*, vol. 8, no. 7, pp. 800–813, Sep. 2024.

[98] L. A. Matkovic, Y. Lei, Y. Fu, T. Wang, A. H. Kesarwala, M. Axente, J. Roper, K. Higgins, J. D. Bradley, T. Liu, and X. Yang, "Deformable lung 4DCT image registration via landmark-driven cycle network," *Med. Phys.*, vol. 51, no. 3, pp. 1974–1984, Sep. 2023.

[99] M. Z. Iqbal, I. Razzak, A. Qayyum, T. T. Nguyen, M. Tanveer, and A. Sowmya, "Hybrid unsupervised paradigm based deformable image fusion for 4D CT lung image modality," *Inf. Fusion*, vol. 102, Feb. 2024, Art. no. 102061.

[100] Z. Zhu, Q. Li, Y. Wei, and R. Song, "Hierarchical multi-level dynamic hyperparameter deformable image registration with convolutional neural network," *Phys. Med. Biol.*, vol. 69, no. 17, Sep. 2024, Art. no. 175007.

[101] W. Liu, Z. Ren, and X. Li, "Weakly supervised medical image registration with multi-information guidance," *Meas. Sci. Technol.*, vol. 35, no. 4, Apr. 2024, Art. no. 046202.

[102] E. Castillo, "Quadratic penalty method for intensity-based deformable image registration and 4DCT lung motion recovery," *Med. Phys.*, vol. 46, no. 5, pp. 2194–2203, May 2019.

[103] T. T. Ho, W. J. Kim, C. H. Lee, G. Y. Jin, K. J. Chae, and S. Choi, "An unsupervised image registration method employing chest computed tomography images and deep neural networks," *Comput. Biol. Med.*, vol. 154, Jan. 2023, Art. no. 106612.

[104] T. F. A. V. D. Ouderaa, I. Isgum, W. B. Veldhuis, and B. D. D. Vos, "Deep group-wise variational diffeomorphic image registration," in *Proc. 2nd Int. Workshop Thoracic Image Anal.*, Lima, Peru. Cham, Switzerland: Springer, Jan. 2020, pp. 155–164.

[105] M. Heinrich, "Deformable lung registration for pulmonary image analysis of MRI and CT scans," Ph.D. dissertation, Oxford Univ., Oxford, U.K., 2013.

[106] H. J. Johnson and G. E. Christensen, "Consistent landmark and intensity-based image registration," *IEEE Trans. Med. Imag.*, vol. 21, no. 5, pp. 450–461, May 2002.

[107] J. Ashburner, "A fast diffeomorphic image registration algorithm," *NeuroImage*, vol. 38, no. 1, pp. 95–113, Oct. 2007.

[108] L. Nenoff, F. Amstutz, M. Murr, B. Archibald-Heeren, M. Fusella, M. Hussein, W. Lechner, Y. Zhang, G. Sharp, and E. Vasquez Osorio, "Review and recommendations on deformable image registration uncertainties for radiotherapy applications," *Phys. Med. Biol.*, vol. 68, no. 24, Dec. 2023, Art. no. 24TR01.

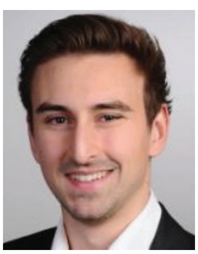

THOMAS ALSCHER received the M.Sc. degree in medical engineering from OTH Regensburg, Germany, in 2019. He is currently pursuing the Ph.D. degree with the Department of Computer Science, University of Copenhagen, Denmark. His research interests include biomedical image processing, computer vision, and simulations.

JENS PETERSEN is currently a Lecturer with the Department of Computer Science, University of Copenhagen.

FRANÇOIS LAUZE received the degree in mathematics in France, the Ph.D. degree in algebraic geometry from the University of Nice Sophia Antipolis (now Université de la Côte d'Azur), in 1994, and the Ph.D. degree from the IT University of Copenhagen, Denmark, in 2004. He spent some years in Burkina Faso, West Africa, where he taught mathematics at the University of Ouagadougou. He worked on variational methods in image analysis, for motion-compensated inpainting and motion recovery. In 2008, he moved to the University of Copenhagen, where he is currently an Associate Professor. He has since mainly worked with variational methods, and geometry for image analysis (Riemannian, but also some metric geometry). More recently, he worked on segmentation problems, using variational methods, and deep learning, and has spent some time with inverse problems in photometric stereo, tomographic imaging, some geometric deep learning, and 3D vision with refraction, and some geometric deep learning problems.

KENNY ERLEBEN received the Ph.D. degree in computer science, in 2005. He is currently a Full Professor of computer graphics and simulation with the University of Copenhagen. His research is focused on how to build simulators and use data-driven approaches together with simulation. Hence, his work ranges from inventing new numerical methods for simulation to using simulation to generate data for machine learning or learning physical models from real-world data. In March 2024, he founded the company DARERL, a provider of digital twin technology for computerized testing and design of medical devices.

SUNE DARKNER received the master's degree in applied mathematics, in 2004, and the Industrial Ph.D. degree in shape and deformation analysis of the human ear canal from the Department of Informatics and Mathematical Modelling, Technical University of Denmark (DTU), in 2009. After holding a position at an energy company as a Data Analyst, he rejoined the Department of Informatics and Mathematical Modelling, DTU, in 2009, as a Postdoctoral Researcher. He currently holds a position as a Professor of image analysis with the Department of Computer Science, University of Copenhagen. He founded a software company building databases for the telecommunication industry, and in collaboration with Oticon A/S, a large hearing aid manufacturer. His research interests include image registration and classification, optimization and regularization, and computational physics.