



Time-Continuous Sliding Motion Image Registration Using Stationary Velocity Fields for Respiratory Motion Interpolation

Ole Gildemeister¹ , Johannes Bostelmann¹ , Pia F. Schulz¹ ,
Andra Oltmann³ , Phillip Rostalski^{3,4} , Jan Modersitzki^{1,2} ,
and Jan Lellmann^{1,2} 

¹ Institute of Mathematics and Image Computing, University of Lübeck, Lübeck, Germany

{o.gildemeister,johannes.bostelmann,jan.lellmann}@uni-luebeck.de

² Fraunhofer Institute for Digital Medicine MEVIS, Lübeck, Germany

³ Fraunhofer IMTE, Fraunhofer Research Institution for Individualized and Cell-Based Medical Engineering, Lübeck, Germany

⁴ Institute for Electrical Engineering in Medicine, University of Lübeck, Lübeck, Germany

Abstract. We propose an image registration model for computing a dense, piecewise diffeomorphic deformation map between 3D thorax images which incorporates sliding motion as often occurring in the pleural cavity. Our approach is based on stationary velocity fields and neural implicit representations, with a particular focus on facilitating motion interpolation. This allows for generating a time-continuous model of the respiratory cycle based on end-inspiration and end-expiration images. We investigate the effect of composing deformations in motion interpolation, using a hybrid technique to enforce domain alignment. We experimentally validate our approach for registration and motion interpolation between thoracic end-inspiration and end-expiration images.

Keywords: Image Registration · Sliding Motion · Motion Interpolation · Respiratory Model

1 Introduction

Motivation. Computing dense correspondence maps between thorax images is challenging due to the complex movement of the lung, in particular along the diaphragm and pleural cavity. Reconstructing thoracic motion is an important step in various applications, such as tumor follow-up and analysis of lung ventilation. In these applications, one is typically interested in computing a single deformation to identify and quantitatively examine corresponding structures:

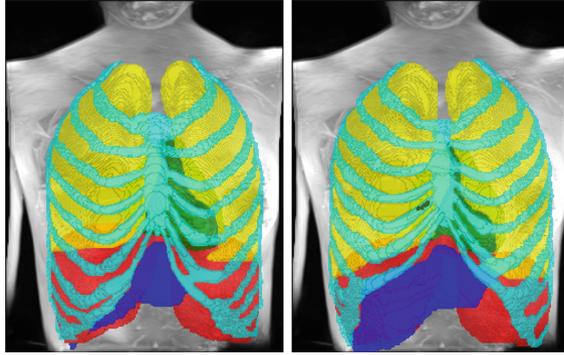


Fig. 1. Segmented organs from volumetric MRI data after expiration (**left**) and inspiration (**right**). The task of reconstructing the time-continuous motion via image registration and motion interpolation is intricate due to the complex motion of the thorax during respiration. In particular, reproducing the sliding motion in the pleural cavity between lung (**yellow**) and rib cage (**cyan**) is challenging. (Color figure online)

Given two images $I_1, I_2 : \Omega \rightarrow \mathbb{R}$ on an image domain $\Omega \subseteq \mathbb{R}^d$, the goal is to find a physically reasonable transformation $\phi : \Omega \rightarrow \mathbb{R}^d$ such that

$$I_1(x) \approx I_2(\phi(x)). \quad (1)$$

This work is motivated by the more specialized application of surface electromyography (sEMG), a prominent tool for measuring electrical muscle activities non-invasively on the skin surface. Using sEMG for the respiratory muscles has been shown to be efficient in measuring the patient’s respiratory effort and timing. It therefore is a promising approach to optimize mechanical ventilation in critically ill patients [12]. However, the signals obtained are highly dependent on the position of the surface electrodes, the muscle geometry, and the body tissue [8]. As these change throughout the respiratory cycle, a time-dependent forward model is required which incorporates the 3D tissue composition as well as muscle and electrode positions.

Obtaining time-varying 3D models via imaging is expensive and often restricted to a severely limited resolution. Therefore, we aim to interpolate the motion from two high-resolution MRI volumes at the extremes of the respiratory cycle, which can be taken under breath hold.

This requires interpolating the motion over the full respiratory cycle, i.e., finding a time-varying deformation map $\phi : \Omega \times [0, 1] \rightarrow \mathbb{R}^d$ such that $\phi(x, \cdot)$ describes the continuous trajectory of an anatomical structure located initially at point x . Notably, these trajectories need not be linear (in space as well as in time), and a simple linear interpolation of the deformation may not yield the desired result: For a rotational motion of 180 °C, for example, linear interpolation of the deformation field would warp the entire volume through the center of the rotation. Thus, even when the final deformation is diffeomorphic, this property may not be preserved at every time point in between.

A more natural way to model such nonlinear motions is the *Stationary Velocity Field (SVF)* approach [1, 2, 10], in which the time-dependent deformation field ϕ is obtained as the (unique) solution of the differential equation

$$\begin{aligned} \frac{\partial \phi(x, t)}{\partial t} &= v(\phi(x, t)) \text{ for } t \in [0, 1], \\ \phi(x, 0) &= x \end{aligned} \quad (2)$$

for some velocity field $v : \mathbb{R}^d \rightarrow \mathbb{R}^d$. To compute a respiratory model, we now seek a velocity field v which maximizes the similarity between the reference (e.g., end-inspiration) image I_1 and the deformed moving (e.g., end-expiration) image $I_2 \circ \phi(\cdot, 1)$ at the end of the time interval.

While the SVF approach provides a natural way for computing nonlinear interpolations, it is not directly suited for thorax imaging, as its most prominent analytical feature is that it generates diffeomorphic transformations when $v \in C_0^1(\Omega, \mathbb{R}^d)$ [1]. This is generally a desirable property in areas of connected tissue. However, it prevents the proper capturing of the motion on sliding interfaces, along which the physically observed ground-truth deformation may exhibit discontinuities. These sliding motions occur especially at the interface between lung and rib cage (Fig. 1).

As discussed below, solutions to this problem of sliding motion have been proposed before. However, we found them unsuitable for motion interpolation. Therefore, in this work, we propose an approach combining the advantages of SVFs while allowing the generation of a consistent time-continuous model which incorporates sliding motion.

Related Work. We refer to [15, 22] for a general overview of classical image registration, to [7] for a survey of recent deep learning-based registration approaches, and to [19, Chapt. 6] for a summary on sliding boundaries. The SVF approach used in our framework can be viewed as a special case of the Large Deformation Diffeomorphic Metric Mapping (LDDMM) approach [5] with a velocity field v that is constant in time.

The incorporation of sliding motion has been widely researched in image registration. A common approach to capture these deformations is to employ spatially variant regularizers. Notable methods include those proposed in [9, 16] requiring prior information on the regions in which the sliding is expected, whereas in [14, 17, 23], regularizers are discussed which estimate the occurrence of sliding implicitly. Recently in [4], a non-differentiable kernel was introduced to adapt the LDDMM method for sliding motion, and the existence of solutions for non-smooth velocity fields was discussed.

In contrast to these approaches, we compute separate deformations for each region along whose boundary we want to encourage sliding. In [11], such an approach for modeling sliding motion is discussed in which intersecting or tearing of adjacent regions is prevented by penalizing deviations between the transformed boundaries. Instead of decoupling the deformations and employing such a soft

penalty, in our framework, we obtain the final deformation as a composition of a global exterior deformation and an interior deformation only acting on the interior regions.

An approach similar to ours was proposed in [18] in which the authors adapt an LDDMM and log-demons based framework to allow for sliding at prespecified boundaries, using a three-step process: After approximately aligning the region boundaries of both images with a global deformation, they compute region-specific deformations to encourage sliding along their respective boundaries, followed by a globally diffeomorphic refinement deformation. In comparison, in our approach, we modify the order in which the global and region-specific deformations are composed. This significantly simplifies the boundary treatment, allows for coherent motion interpolation, and supports faster integration using a scaling-and-squaring technique [2].

The use of a neural implicit representation for SVF-based image registration was discussed in [10] and recently extended to matrix groups in [6]. However, these approaches do not support sliding motion.

Contributions. We propose a piecewise diffeomorphic image registration method which facilitates motion interpolation while capturing sliding motion along the boundaries of predefined regions. Our approach is particularly suited for thorax registration and motion estimation during the respiratory cycle. To this end,

- we investigate how separate deformations acting on different parts of the domain can be composed to avoid artifacts during motion interpolation. Compared to previous works, our approach does not require an explicit consideration of the motion of region boundaries or time-varying sliding boundary conditions.
- We propose a hybrid method to ensure that interior regions remain unchanged under their corresponding interior deformation, allowing for a piecewise diffeomorphic final deformation.
- Our approach is based on stationary velocity fields parameterized by neural implicit representations.

We evaluate our method on real-world data designed to create a forward model for sEMG application.

2 Modeling Smooth Time-Continuous Deformations Using SVF and Neural Implicit Representations

Given two images $I_1, I_2 : \Omega \rightarrow \mathbb{R}$ on an image domain $\Omega \subseteq \mathbb{R}^d$, we consider the task of finding a time-continuous physically plausible deformation $\phi : \Omega \times [0, 1] \rightarrow \mathbb{R}^d$ that warps the moving image I_2 onto the reference image I_1 in the artificial time interval $[0, 1]$ such that

$$I_1(x) \approx I_2(\phi(x, 1)) , \quad \phi(x, 0) = x \quad \forall x \in \Omega. \quad (3)$$

Disregarding sliding motion in this section, we formulate the basic SVF-based approach [10] as follows. Let S denote the *solution operator* which for a given velocity field v returns the final solution $\phi(\cdot, 1)$ of the flow equation (2). Parameterizing the velocity field v_θ by $\theta \in \mathbb{R}^n$, our goal is to minimize some loss function

$$\min_{\theta \in \mathbb{R}^n} \mathcal{L}(v_\theta; I_1, I_2) \quad (4)$$

which evaluates the similarity between I_1 and $I_2 \circ S(v_\theta)$. Specifically, we consider

$$\mathcal{L}(v; I_1, I_2, M) := -\text{NCC}(I_1, I_2 \circ S(v); M) + \alpha_1 J_1(v) + \alpha_2 J_2(S(v)), \quad (5)$$

consisting of three terms. The similarity term NCC measures the cross-correlation on some masked region $M \subset \Omega$ by

$$\text{NCC}(R, T; M) := \frac{\langle T - \bar{T}, R - \bar{R} \rangle_{L^2(M)}}{\|T - \bar{T}\|_{L^2(M)} \|R - \bar{R}\|_{L^2(M)}} \in [-1, 1], \quad (6)$$

where \bar{T}, \bar{R} denote the image means on M . This distance measure is commonly used in image registration [3, 10] and has proven to be more reliable and robust in aligning structures in our experiments, compared to the L_2 distance. We regularize the velocity field by the Frobenius norm of its Jacobian,

$$J_1(v) := \sum_{i=1}^d \int_{\Omega} |Dv_i(x)|_F^2 dx, \quad (7)$$

weighted with parameter $\alpha_1 > 0$, and include a penalty on negative or small Jacobian determinants of the final deformation,

$$J_2(\phi) := \int_{\Omega} \max\{-\det(D\phi(x)) + \epsilon, 0\} dx, \quad (8)$$

weighted with $\alpha_2 > 0$. We set ϵ to 0.1 to further discourage non-diffeomorphic self-intersections that may arise from numerical inaccuracies [10].

We generate the velocity field v_θ using an implicit neural representation in the form of a (fully-connected) multi-layer perceptron with sinusoidal activation functions [21], evaluated on an equidistant rectangular grid. The network consists of five hidden layers of 512 neurons each, with residual connections between consecutive hidden layers. This parameterization has proven effective in computing the velocity which underlies deformations; details and advantages of it are discussed in [10].

An approximation to the solution operator S can be efficiently implemented using the *scaling and squaring* approach [2]. Seeking a temporal resolution of 2^{-k} for some $k \in \mathbb{N}$, we start with a forward Euler step with step size 2^{-k} , followed by k squaring steps, exploiting the fact that

$$\phi\left(x, 2^{-(j-1)}\right) = \phi\left(\phi\left(x, 2^{-j}\right), 2^{-j}\right) \quad (9)$$

for $j = 1, \dots, k$. Recursively, we obtain $\phi(x, 1)$, with the effective temporal resolution scaling exponentially in the number of squaring steps k .

After recovering v_θ , the deformation $\phi(\cdot, t)$ can be evaluated at any time point $t \in [0, 1]$ using the same solution operator applied to the rescaled velocity field $t \cdot v_\theta$. In this manner, we can conduct motion interpolation.

3 Piecewise Diffeomorphic Deformations for Motion Interpolation

As the basic model described in Sect. 2 yields smooth deformations, it cannot incorporate sliding motion. Therefore, in the following, we discuss the specific adaptations that allow us to capture such sliding motion along the interfaces of predefined regions. For ease of notation, we omit the time dependence of ϕ throughout most of the discussion.

We require an open *interior* subdomain $\Omega_{\text{int}} \subset \Omega$ which, in our case, represents a volume in the interior of the thorax with respect to the moving image I_2 . In particular, it covers lung, heart, diaphragm, liver, and possibly other organs in this vicinity, while explicitly excluding structures such as the rib cage and spine. Importantly, its boundary $\Gamma := \partial\Omega_{\text{int}}$ must include all surfaces along which sliding motion should be encouraged. We obtain the final deformation ϕ as the composition

$$\phi := \phi_{\text{int}} \circ \phi_{\text{ext}} \quad (10)$$

of some exterior (global) diffeomorphism $\phi_{\text{ext}} : \Omega \rightarrow \Omega$ and some interior diffeomorphism $\phi'_{\text{int}} : \Omega_{\text{int}} \rightarrow \Omega_{\text{int}}$, extended to the whole domain Ω by

$$\phi_{\text{int}}(x) := \begin{cases} \phi'_{\text{int}}(x), & \text{if } x \in \Omega_{\text{int}}, \\ x, & \text{else.} \end{cases} \quad (11)$$

This allows for a smooth global motion with a non-smooth component along the boundary Γ , while preserving the diffeomorphic property of the inner deformation field ϕ_{int} . Conveniently, this design prevents the intersecting or tearing of adjacent regions.

Both ϕ_{ext} and ϕ'_{int} are generated separately using the neural implicit representation-based SVF approach described in Sect. 2. To obtain the final deformation ϕ , we conduct a three-step optimization process.

1. We first compute the global deformation ϕ_{ext} using (4) on the original images I_1 and I_2 , with the NCC masked on the exterior domain with respect to the reference image I_1 . The regularization still acts on the entire domain Ω , yielding a globally smooth deformation.
2. We then repeat the process, restricted to the domain Ω_{int} , using $I_2 \circ S(v) \circ \phi_{\text{ext}}$ instead of $I_2 \circ S(v)$ in the NCC-term of the loss (5). This yields $\phi'_{\text{int}} = S(v)$ and thereby ϕ_{int} .
3. Finally, an end-to-end optimization step is performed for fine-tuning the network parameters using $\phi_{\text{int}} \circ \phi_{\text{ext}}$. Here, we consider the NCC over the entire domain, once on the full images and once masked to their respective exterior, and include all interior and exterior regularization terms used in steps 1 and 2.

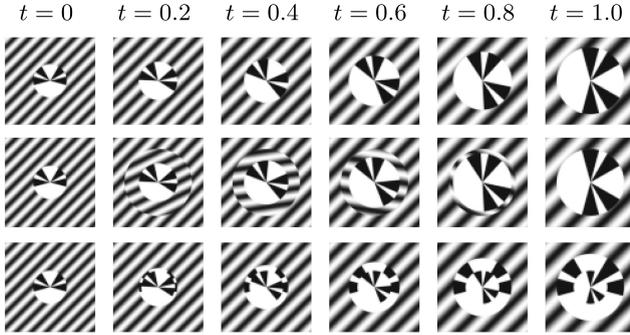


Fig. 2. For a consistent motion interpolation $I_2 \circ \phi(\cdot, t)$ with sliding boundaries, the order in which the interior and exterior deformation are composed, as well as the chosen mask, are critical. With $\tilde{\phi} = \phi_{\text{ext}} \circ \phi_{\text{int}}$, motion interpolation generates artifacts around the sliding boundary, regardless of whether ϕ_{int} is masked with respect to I_1 (**second row**) or I_2 (**third row**). With the choice $\phi = \phi_{\text{int}} \circ \phi_{\text{ext}}$ and ϕ_{int} masked with respect to I_2 , the deformation is consistent throughout the time interval (**first row**).

Deformation Order. In the following, we briefly reason the particular order chosen in the composition (10), which contrasts the approach of [18]. When instead considering the converse $\tilde{\phi} := \phi_{\text{ext}} \circ \phi_{\text{int}}$, the interior subdomain $\tilde{\Omega}_{\text{int}}$ had to be chosen with respect to the reference image I_1 in order to obtain a reasonable final deformation, yielding the respective boundary $\tilde{\Gamma}$.

In general, both approaches result in bijective deformations, but equally lose their global diffeomorphic property: Due to the discontinuity of ϕ_{int} on Γ , ϕ may not be continuous on $\phi_{\text{ext}}^{-1}(\Gamma)$, and ϕ^{-1} not on Γ . Similarly, discontinuities of $\tilde{\phi}$

may occur on $\tilde{\Gamma}$, and of $\tilde{\phi}^{-1}$ on $\phi_{\text{ext}}(\tilde{\Gamma})$. In general, this behavior is desired as sliding motion cannot be modeled in a purely diffeomorphic setting.

The advantage of choosing ϕ over $\tilde{\phi}$ arises when viewing the deformations under time dependency, as is necessary to conduct motion interpolation. We then obtain the composed deformation at time t by

$$\phi(x, t) := \phi_{\text{int}}(\phi_{\text{ext}}(x, t), t), \quad \tilde{\phi}(x, t) := \phi_{\text{ext}}(\phi_{\text{int}}(x, t), t), \quad (12)$$

respectively, and warp the image over time with it while producing sliding motion along the discontinuities. By above argument, the coordinates of these discontinuities are only time-invariant for ϕ^{-1} and $\tilde{\phi}$, but change with t for ϕ and $\tilde{\phi}^{-1}$.

Since we evaluate the moving image I_2 in both motion interpolation and training as $I_2 \circ \phi$, this corresponds to transforming the coordinates of structures in the moving image I_2 in the inverse direction ϕ^{-1} . With the approach proposed in this work, i.e., the choice $\phi = \phi_{\text{int}} \circ \phi_{\text{ext}}$, coordinates thus first undergo the interior deformation ϕ_{int}^{-1} , and afterwards the global coordinate transformation ϕ_{ext}^{-1} . This ensures that the interior deformation acts on the correct interior information of I_2 at all times (Fig. 2, first row). This behavior is described by above statement that the coordinates of possible discontinuities in ϕ^{-1} do not change over time.

In contrast, with the converse choice $\tilde{\phi} = \phi_{\text{ext}} \circ \phi_{\text{int}}$, coordinates of structures in I_2 are first transformed globally by ϕ_{ext}^{-1} and then by the interior deformation ϕ_{int}^{-1} . This results in a time-varying interface on which the interior deformation acts, induced by the time-varying discontinuity positions of $\tilde{\phi}^{-1}$. This is unproblematic for the final deformation at $t = 1$, explaining why this topic is typically not well addressed in the literature [9, 11, 18]. In motion interpolation, however, these time-varying interfaces can cause undesirable artifacts (Fig. 2).

Enforcing Piecewise Diffeomorphisms. A key requirement of our framework is that the interior diffeomorphism in fact maps at all times from Ω_{int} to Ω_{int} , implying that the underlying velocity field must vanish along the normal direction at the boundary Γ . To ensure this, we adapt and combine two different approaches presented in the literature before: (1) We explicitly constrain the velocity field to vanish orthogonally to the boundary [18], and (2) we introduce a soft penalty against misalignments of the transformed domain arising from possible numerical issues [11]. We modify each approach to fit into our framework.

Let $\psi : \Omega \rightarrow \mathbb{R}$ denote the signed distance function with respect to Ω_{int} , so that the normalized gradient $\nabla\psi$ points towards / away from the closest point in the boundary Γ . We thus require $\langle v(x), \nabla\psi(x) \rangle = 0$ for $x \in \Gamma$. Instead of hard-clipping the velocity field \bar{v} obtained by the neural implicit representation, we use a soft-clipping approach to preserve differentiability by setting

$$\begin{aligned}
v(x) &:= \bar{v}(x) - \alpha(x)\nabla\psi(x), \text{ where} \\
\alpha(x) &:= \langle \bar{v}(x), \nabla\psi(x) \rangle - \frac{2}{\pi}\beta\psi(x) \arctan\left(\frac{\langle \bar{v}(x), \nabla\psi(x) \rangle}{\frac{2}{\pi}\beta\psi(x) + \epsilon}\right).
\end{aligned} \tag{13}$$

This soft-clip (13) constraints the component of v pointing towards the boundary to length $\beta\psi(x)$. Conveniently, ψ and $\nabla\psi$ need only be computed once.

We choose $\beta := 2^k$, motivated as follows. The scaling and squaring recursion (9) ensures that if $\phi(\cdot, 2^{-l})$ maps into Ω_{int} , then so does $\phi(\cdot, 2^{-l+1})$. Thus, it suffices to ensure that the initial Euler step, which is of the form $\phi(\cdot, 2^k) := x + 2^{-k}v(x)$, stays in Ω_{int} . If the boundary Γ is a hyperplane, the initial step towards the boundary has length $2^{-k}\langle v(x), \nabla\psi(x) \rangle$. By our choice of β , this term cannot exceed $\psi(x)$, i.e., the distance to the boundary, ensuring that the initial Euler step does not leave the specified domain.

In general, however, the boundary of Ω_{int} is not a hyperplane, so above argument does not suffice everywhere. Therefore, when optimizing for the interior deformation, we include an additional soft penalty J_3 in the loss (5), weighted with parameter $\alpha_3 > 0$, to encourage that it maps Ω_{int} inversely consistent onto itself:

$$J_3(v_\theta; M) := \frac{1}{2}(\text{DS}(M \circ S(v_\theta), M) + \text{DS}(M \circ S(-v_\theta), M)). \tag{14}$$

Here, $\text{DS}(X, Y) = 2|X \cap Y|/(|X| + |Y|)$ denotes the Dice score, a measure of the overlap of two regions, and the mask M is given by Ω_{int} . Note that for the SVF approach it holds $S(-v_\theta) = S(v_\theta)^{-1}$ up to discretization.

4 Numerical Results

Implementation Details. We implemented our framework in PyTorch 2.3.0. All experiments were performed on a 24-core AMD EPYC 74F3 system with 256GB of RAM, 3x NVIDIA A100, and CUDA 12.0. The results presented here were obtained with weight parameters $\alpha_1 = 0.001$, $\alpha_2 = 10$ and $\alpha_3 = 1$, a SIREN hyper-parameter ω_0 of 4.6 [21], Adam [13] as optimizer with 180 iterations, and a fixed the learning of 0.004. In the end-to-end-optimization, we used $\alpha_1 = 0.01$.

Dataset. We evaluate our method on a dataset consisting of 3D MRI sequences of the torso recorded in the end-expiratory and end-inspiratory phase of the same patient. For both phases, four differently weighted images of the same scene are given, allowing to distinguish more anatomical structures in the thorax. The data of both phases is complemented by segmentations of six regions: Diaphragm, Lung, Heart, Liver, Rib Cage, and Thorax. The dataset originates from an ongoing project on surface electromyography, is currently not publicly available and comes without a known ground-truth deformation. It is visualized in Fig. 1.

Experimental Results. Figure 3 illustrates the necessity of explicitly modeling sliding motion. It shows the result after each registration phase of our three-step framework in terms of the difference images between the reference image I_1 and the warped moving image $I_2 \circ \phi$, warped by the respective final deformation $\phi(\cdot, 1)$, in one coronal 2D slice. Including the separate interior deformation allows to correctly capture the sliding motion occurring in the thorax.

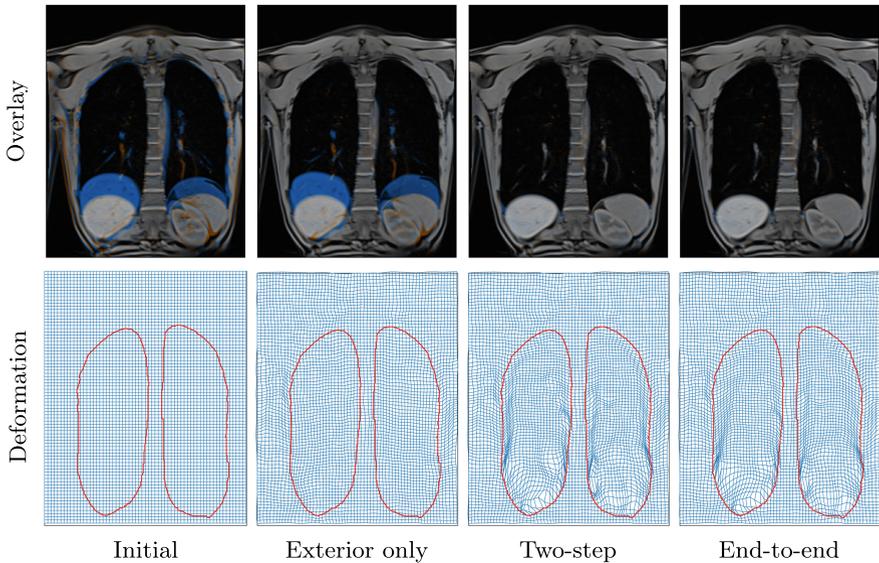


Fig. 3. The exterior deformation ϕ_{ext} aligns exterior structures and the lung boundaries along the rib cage (**second from left**). Adding the interior deformation as $\phi_{\text{int}} \circ \phi_{\text{ext}}$ (**second from right**), also inner structures are aligned, featuring sliding motion along the interfaces (**red**). An end-to-end refinement step further improves the results (**right**). The warped moving (end-expiration) image is colored blue, while the reference (end-inspiration) image is colored orange; gray values indicate good alignment. The warped grids depict the corresponding deformations, the mask boundary is drawn in red. (Color figure online)

To quantitatively evaluate the registration accuracy, we computed the Dice scores for aligning the segmented regions using different registration approaches (Fig. 4). Compared to the results obtained with the classical SVF approach without sliding boundaries, the SyN-toolbox [3], and a recent ResNET-LDDMM approach [20], our method provides a more accurate final deformation.

Figure 5 visualizes the final result. Our combined approach allows for smooth motion interpolation of the thorax during the respiratory cycle. It captures the sliding motion along the interface between thoracic interior and rib cage at all time steps without producing significant artifacts. Unlike the classical approach without sliding motion, it does not generate unnaturally deformed ribs.

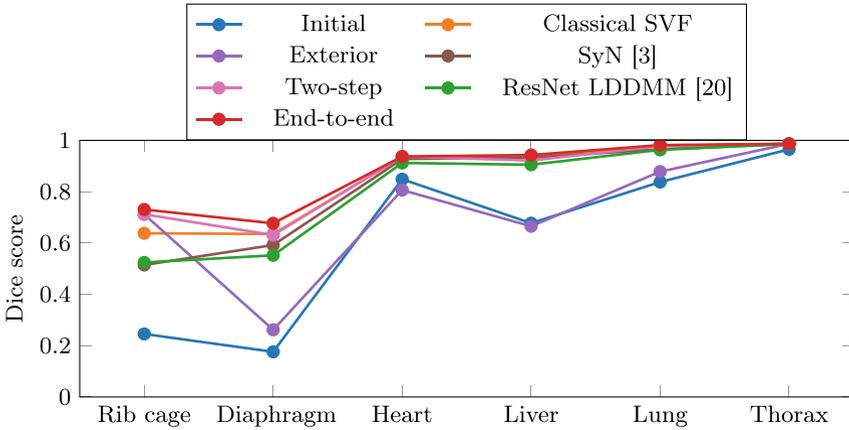


Fig. 4. Comparison of different image registration approaches in terms of Dice scores. Values close to 1 indicate a good alignment of the segments after the deformation. Without incorporating sliding interfaces, the SVF approach aligns especially the rib cage and diaphragm quite poorly (**orange**). Our sliding motion framework captures its deformation significantly better (**pink**); additional end-to-end optimization further improves the result (**red**). Overall, the approach outperforms the comparison methods SyN [3] and ResNet LDDMM [20] significantly, especially in terms of the rib cage alignment. (Color figure online)

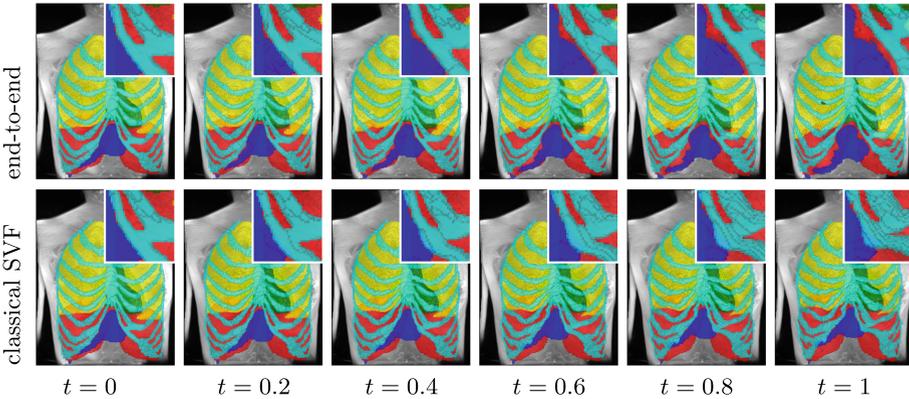


Fig. 5. Full motion interpolation of the respiratory cycle applied to a segmented MR image. Without sliding interfaces (**bottom**), the rib cage is deformed unnaturally due to a combination of an interior downward motion and the expansion of the rib cage. Our sliding motion approach (**top**) results in more natural deformations throughout the process.

5 Discussion and Future Work

In this work, we have found that accurately modeling sliding interfaces for motion interpolation requires a more careful analysis than estimating only a single deformation field. For this setting, our two-step SVF approach combined with an implicit neural field parameterization allows to considerably reduce artifacts, as it is able to capture sliding motion in a time-continuous framework.

A limitation of our approach is that it requires predefined regions along whose boundaries sliding is encouraged. The integration of automatic segmentation into the framework might therefore be a promising modification, also in view of recent developments in deep-learning. In the context of respiratory motion interpolation, the question whether it is sufficiently expressive to assume the velocity field to be constant in time, remains subject to further research.

In the future, an interesting and mathematically challenging direction would be to incorporate sliding motion between the pulmonary lobes, which is more subtle and intricate to detect. Finally, while our approach requires solving an optimization problem for each new image pair, there is no inherent barrier to generate the velocity fields using another trainable network, allowing the extension of our approach to a fully data-driven method.

Acknowledgments. Andra Oltmann and Philipp Rostalski were supported by European Union - European Regional Development Fund, the Federal Government and Land Schleswig-Holstein, Project: “Diagnose- und Therapieverfahren für die Individualisierte Medizintechnik (IMTE)”, Project No. 12 420 002.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Arguillere, S., Trélat, E., Trouvé, A., Younes, L.: Shape deformation analysis from the optimal control viewpoint. *J. Math. Pures. Appl.* **104**(1), 139–178 (2015)
2. Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-euclidean framework for statistics on diffeomorphisms. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 924–931. Springer (2006)
3. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image Anal.* **12**(1), 26–41 (2008)
4. Bao, L., Lu, J., Ying, S., Sommer, S.: Sliding at first-order: higher-order momentum distributions for discontinuous image registration. *SIAM J. Imaging Sci.* **17**(2), 861–887 (2024)
5. Beg, M.F., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *Int. J. Comput. Vis.* **61**, 139–157 (2005)
6. Bostelmann, J., Gildemeister, O., Lellmann, J.: Stationary velocity fields on matrix groups for deformable image registration. arXiv preprint [arXiv:2410.10997](https://arxiv.org/abs/2410.10997) (2024)

7. Chen, J., et al.: A survey on deep learning in medical image registration: new technologies, uncertainty, evaluation metrics, and beyond. *Med. Image Anal.* 103385 (2024)
8. Farina, D., Cescon, C., Merletti, R.: Influence of anatomical, physical, and detection-system parameters on surface EMG. *Biol. Cybern.* **86**(6), 445–456 (2002)
9. Fu, Y., Liu, S., Li, H.H., Li, H., Yang, D.: An adaptive motion regularization technique to support sliding motion in deformable image registration. *Med. Phys.* **45**(2), 735–747 (2018)
10. Han, K., et al.: Diffeomorphic image registration with neural velocity field. In: *IEEE Winter Conference on Applications of Computer Vision*, pp. 1869–1879 (2023)
11. Heldmann, S., Polzin, T., Derksen, A., Berkels, B.: An image registration framework for sliding motion with piecewise smooth deformations. In: *Aujol, J.-F., Nikolova, M., Papadakis, N. (eds.) SSVM 2015. LNCS*, vol. 9087, pp. 335–347. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-18461-6_27
12. Jonkman, A.H., Warnaar, R.S.P., Baccinelli, W.a.: Analysis and applications of respiratory surface emg: report of a round table meeting. *Crit Care* **28**(1) (2024)
13. Kingma, D.P.: Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
14. Mang, A., Biros, G.: Constrained H^1 -regularization schemes for diffeomorphic image registration. *SIAM J. Imaging Sci.* **9**(3), 1154–1194 (2016)
15. Modersitzki, J.: *Numerical methods for image registration*. OUP Oxford (2003)
16. Pace, D.F., Aylward, S.R., Niethammer, M.: A locally adaptive regularization based on anisotropic diffusion for deformable image registration of sliding organs. *IEEE Trans. Med. Imaging* **32**(11), 2114–2126 (2013)
17. Papi ez, B.W., Heinrich, M.P., Fehrenbach, J., Risser, L., Schnabel, J.A.: An implicit sliding-motion preserving regularisation via bilateral filtering for deformable image registration. *Med. Image Anal.* **18**(8), 1299–1311 (2014)
18. Risser, L., Vialard, F.X., Baluwala, H.Y., Schnabel, J.A.: Piecewise-diffeomorphic image registration: application to the motion estimation between 3D CT lung images with sliding conditions. *Med. Image Anal.* **17**(2), 182–193 (2013)
19. Schmidt-Richberg, A.: *Registration methods for pulmonary image analysis: integration of morphological and physiological knowledge*. Springer (2014)
20. Schulz, P.F., et al.: Image registration for a dynamic breathing model. In: *BVM Workshop*, accepted. Springer (2025)
21. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. *Adv. Neural. Inf. Process. Syst.* **33**, 7462–7473 (2020)
22. Viergever, M.A., Maintz, J.A., Klein, S., Murphy, K., Staring, M., Pluim, J.P.: A survey of medical image registration - under review. *Med. Image Anal.* **33**, 140–144 (2016)
23. Vishnevskiy, V., Gass, T., Szekely, G., Tanner, C., Goksel, O.: Isotropic total variation regularization of displacements in parametric image registration. *IEEE Trans. Med. Imaging* **36**(2), 385–395 (2016)