



HAL
open science

Adaptive local boundary conditions to improve deformable image registration

Eloïse Inacio, Luc Lafitte, Laurent Facq, Clair Poignard, Baudouin Denis de Senneville

► **To cite this version:**

Eloïse Inacio, Luc Lafitte, Laurent Facq, Clair Poignard, Baudouin Denis de Senneville. Adaptive local boundary conditions to improve deformable image registration. *Physics in Medicine and Biology*, 2024, 69 (16), pp.165025. <10.1088/1361-6560/ad6952>. <hal-04719267>

HAL Id: hal-04719267

<https://inria.hal.science/hal-04719267v1>

Submitted on 3 Oct 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Adaptive local boundary conditions to improve Deformable Image Registration

Eloïse Inacio^{1,2}, Luc Lafitte², Laurent Facq¹, Clair Poignard^{1,2},
Baudouin Denis de Senneville^{1,2}‡

¹ University of Bordeaux, IMB, UMR CNRS 5251, INRIA Project team Monc,
Talence, France, F-33405 Talence Cedex, France

² INRIA Bordeaux Sud Ouest, 200 Av. de la Tour, Talence, 33405, France

E-mail: eloise.inacio@inria.fr, laurent.facq@math.u-bordeaux.fr,
luc.lafitte@inria.fr, clair.poignard@inria.fr,
bdenisde@math.u-bordeaux.fr

May 2024

Abstract.

Objective: In medical imaging, it is often crucial to accurately assess and correct movement during image-guided therapy. Deformable image registration (DIR) consists in estimating the required spatial transformation to align a moving image with a fixed one. However, it is acknowledged that for DIR methods, boundary conditions applied to the solution are critical in preventing mis-registration. This poses an issue particularly when areas of interest are located near the image border. Despite the extensive research on registration techniques, relatively few have addressed the issue of boundary conditions in the context of medical DIR. Our aim is a step towards customizing boundary conditions to suit the diverse registration tasks at hand.

Approach: We analyze the behavior of two typical global boundary conditions: homogeneous Dirichlet and homogeneous Neumann. We propose a generic, locally adaptive, Robin-type condition enabling to balance between Dirichlet and Neumann boundary conditions, depending on incoming/outgoing flow fields on the image boundaries. The proposed framework is entirely automatized through the determination of a reduced set of hyperparameters optimized via energy minimization.

Main results: The proposed approach was tested on a mono-modal CT thorax registration task and an abdominal CT-to-MRI registration task. For the first task, we observed a relative improvement in terms of target registration error of up to 12% (mean 4%), compared to homogeneous Dirichlet and homogeneous Neumann. For the second task, the automatic framework provides results close to the best achievable.

Significance: This study underscores the importance of tailoring the registration problem at the image boundaries. In this research, we introduce a novel method to adapt the boundary conditions on a voxel-by-voxel basis, yielding optimized results in two distinct tasks: mono-modal CT thorax registration and abdominal CT-to-MRI

‡ Corresponding Author.

registration. The proposed framework enables optimized boundary conditions in image registration without prior assumptions regarding the images or the motion.

Keywords: Deformable Image Registration, Variational methods, Boundary conditions

1. Introduction

In medical imaging, numerous scans may be necessary on a daily basis to fully understand a patient’s condition and/or plan treatment. This is particularly true for dose calculation, for instance in the case of radiotherapy (Grégoire et al. 2020), electroporation ablation (Gallinato et al. 2019) or thermotherapy (Kok et al. 2020). As the images may not be captured simultaneously, or occasionally within the same modality, it becomes essential to align image features before drawing conclusions (Zachiu et al. 2015) (Chetty & Rosu-Bubulac 2019) (Yuen et al. 2020). Any motion event (such as breathing or peristalsis), or adjustment in the scanner’s position relative to the patient, complicate the task for the expert in accurately interpreting and comparing the images.

Deformable Image Registration (DIR) is an image processing task that consists in estimating the motion to align a moving image with a fixed image, taking into account the content of both images and sometimes prior knowledge of the motion and/or the tissue’s physical properties. It has been an active area of research in the field of image processing since the 1960s, owing to its extensive range of applications. A brief overview of existing DIR methods is provided in Appendix A. DIR can be expressed as a variational problem, offering a framework to formulate versatile DIR problems adaptable to various applications. The spatial transformation \mathbf{T} between a moving image J and a fix image I ($(I, J) \in \Omega \times \Omega$, Ω being the image domain: $\Omega \subset \mathbb{R}^3 \rightarrow \mathbb{R}$ for the three dimensional case) is estimated by minimizing the following energy functional:

$$\mathcal{E}(\mathbf{T}) = \int_{\Omega} (\mathcal{D}^{I,J}(\mathbf{T}) + \alpha \mathcal{S}(\mathbf{T})) \, \mathbf{d}\mathbf{x} \quad (1)$$

where $\mathbf{T} : \Omega \rightarrow \Omega$, \mathcal{D} is a data fidelity term which drives the motion so that relevant features are aligned between I and the registered image $K = \mathbf{T}(J)$, \mathcal{S} is a regularization term which is crucial to achieve well-posedness (Staring et al. 2007) (Larrey-Ruiz et al. 2007) (Vishnevskiy et al. 2017) and α is a parameter allowing to balance between data fidelity and regularization.

In some scenarios, only a partial field of view is available, either due to inherent limitations of the imaging sensor or the acquisition process itself, when multiple scans need to be conducted within a limited timeframe. Consequently, an additional challenge arises for DIR methods, wherein the imposition of appropriate boundary conditions is crucial for preventing mis-registration from propagating towards the image center. This becomes particularly problematic when regions of interest lie near the image boundary. Several attempts restrict the registered area using segmentation masks either automatically extracted, provided by an extra module or, in some cases, by experts (Richard & Cohen 2003) (Zhang et al. 2004). Doing so allows shearing between the zone of interest and the background as well as limits the impact of the constraint applied on the boundary, since information is still available beyond. However, most of the time, the boundary conditions are implicitly imposed either with the homogeneous Dirichlet

boundary condition (HD)–the motion is null at the border ($\mathbf{T}(\mathbf{x}) = 0, \forall \mathbf{x} \in \delta\Omega$, $\delta\Omega$ being the image boundary)– or the homogeneous Neumann boundary condition (HN) –the motion is constant along the normal border direction ($\nabla\mathbf{T} \cdot \mathbf{n} = 0$, $\mathbf{n}(\mathbf{x})$ being the vector normal to the image boundary at location $\mathbf{x} \in \delta\Omega$)– which is more adequate when objects are close to the edges (Braumann & Kuska 2005).

An important feature, when solving the minimization problem for DIR, lies in the fact that the constraints on the motion on the boundary have to be *complementing* for the volume operator \mathcal{S} so that the minimization problem is well-posed. With standard regularization terms, this condition holds for Dirichlet and Neumann conditions, but it is not the case for all regularizations, and the choice of the boundary condition is tightly linked to the operator \mathcal{S} of Eq. (1).

Our contribution encompasses the following four key aspects:

- (i) We first highlight the importance of tailoring boundary conditions, particularly in scenarios with partial fields of view where relevant information resides near the image border.
- (ii) We introduce a generic framework for voxel-wise adaptive boundary conditions in the form of a Robin-type local boundary condition tending towards a HD condition in the absence of incoming or outgoing flow fields, while relaxing the constraint under other circumstances.
- (iii) We introduce a method for detecting and incorporating incoming or outgoing flow fields within the boundary conditions, thereby accommodating the image-specific characteristics.
- (iv) We demonstrate that the proposed framework is capable of complete automation through the optimization of a reduced set of hyperparameters via DIR energy minimization.

2. Method

2.1. Generic DIR problem formulation

In the present study, we focus on the standard Tikhonov regularization:

$$\mathcal{S}(\mathbf{T}) = \|\nabla\mathbf{T}\|^2, \quad (2)$$

and the minimization of \mathcal{E} as expressed in Eq. (1) is performed thanks to the Euler-Lagrange equations, for which the following Fourier-Robin condition (see Eq. (3c)) is known to be complementing. More precisely, we focus on the following evolution problem on $\mathbf{T} = [\mathbf{T}_x, \mathbf{T}_y, \mathbf{T}_z]$ solved coordinate by coordinate, that is for $s \in \{x, y, z\}$, \mathbf{T}_s satisfies:

$$\partial_t \mathbf{T}_s - \alpha \Delta \mathbf{T}_s = -\frac{\delta \mathcal{D}^{I,J}}{\delta \mathbf{T}_s}(\mathbf{T}), \text{ in } \Omega \quad (\text{Minimization of } \mathcal{E}) \quad (3a)$$

$$\mathbf{T}_s(t=0, x) = \mathbf{T}_s^{init}(x), \quad x \in \Omega \quad (\text{Initial conditions}) \quad (3b)$$

$$\mathbf{A}_s(\nabla \mathbf{T}_s \cdot \mathbf{n}) + \mathbf{B}_s \mathbf{T}_s = \mathbf{C}_s, \text{ on } \partial\Omega \quad (\text{Boundary conditions}) \quad (3c)$$

where \mathbf{T}^{init} is a prior vector field (in this study, \mathbf{T}^{init} was identically equal to zero), $\mathbf{A}, \mathbf{B}, \mathbf{C}$ are vectors — \mathbf{A} and \mathbf{B} having positive components $\mathbf{A}_s, \mathbf{B}_s$ — defined in the subsequent section. In Eq. (3c), \mathbf{n} refers to the normal vector to the boundary of the image, or more precisely, to the boundary of the field of view of interest. In fluid mechanics, Dirichlet boundary conditions are interpreted as an imposed velocity field on the image boundary, whereas Neumann boundary conditions imply the imposition of the shear tensor along the normal direction (Abbott & Basco 1989). The Fourier-Robin boundary condition offers a way to balance between these two configurations.

2.2. Proposed adaptive local boundary condition

We proposed to drive the choice of the above vectors \mathbf{A}, \mathbf{B} and \mathbf{C} by the following considerations:

- (i) A balance between Dirichlet and Neumann conditions is achievable on a voxel-by-voxel basis.
- (ii) The source term \mathbf{C} can enforce values for both incoming and outgoing flow fields in the image boundary.

Let \mathbf{g} represent a voxelwise incoming/outgoing field map in the image boundary, we propose to choose $\mathbf{A}_s, \mathbf{B}_s, \mathbf{C}_s$ as follows:

$$\mathbf{A}_s = \beta(\mathbf{g}_s), \quad \mathbf{B}_s = 1 - \mathbf{A}_s, \quad \mathbf{C}_s = \gamma \mathbf{g}_s \quad (4)$$

where γ is a weighting factor for the source term and $\beta : \mathbb{R} \rightarrow [0, 1]$ is the 1-parameter sigmoidal function between 0 and 1 defined by:

$$\forall \lambda \in \mathbb{R}, \quad \beta(\lambda) = \begin{cases} 1 - \tanh(a - |\lambda|) & \text{if } |\lambda| \leq a \\ 1 & \text{otherwise} \end{cases} \quad (5)$$

a being the second hyperparameter that tunes the transition between Neumann ($\beta = 1$) and Dirichet ($\beta = 0$) conditions. Therefore, our proposed adaptive Robin condition (hereafter referred to as “Adaptive Robin”) relies solely on the determination of two hyperparameters: a and γ .

2.3. Estimation of incoming/outgoing flow fields on the image boundaries

At this point, we need to compute an estimation of the incoming and outgoing flow fields at the image boundary (*i.e.* \mathbf{g}), in order to implement the adaptive boundary conditions

described in Eq. (3) and (4). We recall that this mapping enables the reduction of the comprehensive parameterization of the proposed voxelwise boundary conditions to the determination of two hyperparameters: a and γ .

In this study, the proposed adaptive framework described in Eq. (4) uses an estimation of incoming/outgoing flow fields represented as a motion inverse consistency map computed on the image boundary. The motion inverse consistency relies on the hypothesis that DIR is a symmetric task, that is, switching the fix and the moving images will not impact the solution:

$$\mathbf{T}_{moving \rightarrow fix} = (\mathbf{T}_{fix \rightarrow moving})^{-1} \quad (6)$$

However, in practice, the images order may change the outcome due to asymmetry either in the problem formulation or in the numerical methods. The inverse consistency map \mathbf{g} is a vector-valued map which was computed as follows:

$$\mathbf{g}(\mathbf{T}, \mathbf{T}_{inv}) = \frac{\mathbf{T} + \mathbf{T}_{inv}(\mathbf{T}) + \mathbf{T}_{inv} + \mathbf{T}(\mathbf{T}_{inv})}{2}, \text{ on } \partial\Omega \quad (7)$$

where $\mathbf{T} = \mathbf{T}_{moving \rightarrow fix}$ and $\mathbf{T}_{inv} = \mathbf{T}_{fix \rightarrow moving}$. The calculation of \mathbf{g} with Eq. (7) was restricted to voxels located on the image boundary $\delta\Omega$.

The determination of \mathbf{g} involved the computation of forward and backward motion estimates (*i.e.* \mathbf{T} and \mathbf{T}_{inv}). This process being computationally expensive, \mathbf{g} is only estimated once per resolution step and over a small number of iterations. Given the unknown nature of the correct boundary condition at this stage, a HN boundary condition (representing the least restrictive global boundary condition) was employed.

2.4. Implemented DIR algorithm

To assess the effectiveness of the suggested adaptive boundary condition, we solved the problem outlined in Eq. (3) employing the parameter selection specified in Eq. (4). The overall framework is summarised in pseudo-code in Appendix C. The solution was obtained using the following multi-modal data fidelity term introduced in (Denis de Senneville et al. 2016):

$$\mathcal{D}^{I,J}(\mathbf{T}) = \exp(-\mathcal{C}(I, \mathbf{T}(J))), \quad (8)$$

where

$$\mathcal{C}(I, \mathbf{T}(J)) = \frac{\int_{\Gamma} |\vec{\nabla} I \cdot \vec{\nabla} \mathbf{T}(J)| \, d\mathbf{x}}{\int_{\Gamma} \|\vec{\nabla} I\|_2 \|\vec{\nabla} \mathbf{T}(J)\|_2 \, d\mathbf{x}}, \quad (9)$$

The data fidelity term $\mathcal{D}^{I,J}$ relies on the image gradient, enhancing the method's resilience to contrast reversals. Specifically, it facilitates alignment of both parallel and anti-parallel edges by considering only the absolute values of gradients.

In the aforementioned equation, Γ represents a three-dimensional patch centered around the voxel under consideration. Patch-based computation increases robustness against substantial structural alterations, encompassing transient objects among other variations.

This variational algorithm thus incorporates two hyper-parameters: (i) the patch size (small values lead to good performances but a patch size of one voxel induces instabilities in the numerical scheme ; in the following experiments, we use a patch-size of $5 \times 5 \times 5$), (ii) the regularisation factor α (it balances between the data fidelity term and the regularization in Eq.(1) ; α below 0.1 leads to an unstable numerical scheme, above that value, the algorithm is fairly robust to changes in α ; we use a value of 0.4 going forward). A complete parameter study was done in (Denis de Senneville et al. 2016) and we used those results to guide our choice of parameters.

The framework proposed for adapting the boundaries voxel-wise during the registration process is shown in Figure 1. The general variational formulation for DIR in Eq. (1) being valid only for small movements, a multi-resolution scheme is used. The registration begins at lower resolutions to capture global movements and uses the estimated motion as a starting point for subsequent resolution levels (Brox et al. 2009). Initial experiments indicated that an in-plane dimension of 16×16 voxels at lowest resolution level was sufficient for capturing preliminary information about larger movements. The computation of incoming/outgoing flow fields on the image boundaries (*i.e.* \mathbf{g}) underwent updates throughout the registration process, at the beginning of each resolution step, to ensure voxel-wise adaptation of boundary conditions. In the following experiments, the time step dt was set to 1, the number of iterations for the calculation of \mathbf{g} via Eq. (7) was 10, the number of iterations for the minimization of $\mathcal{E}(\mathbf{T})$ via Eq. (3) was 1000.

To properly apply the boundary conditions, the system of equations (Eq. (3a) and (3b) was solved in inner voxels (which corresponds to $\Omega \setminus \partial\Omega$), the one voxel wide perimeter (which corresponds to $\partial\Omega$) being computed separately according to the boundary conditions (Eq. (3c), see Appendix B).

2.5. Hyperparameter optimization

The two hyperparameters for the proposed adaptive Robin boundary conditions (*i.e.*, γ and a) were determined through a grid search process aimed at optimizing the minimized energy $\mathcal{E}(\mathbf{T})$ as expressed in Eq. (1). The grid search spanned values of $\gamma \in [0, 1]$ with a step size of 0.1, and $a \in \{0, 0.1, 1, 10, 100\}$.

2.6. Experimental setup

2.6.1. Mono-modal CT thorax registration: The first dataset employed for evaluating our methodology is the DIR-Lab 4DCT dataset (Castillo et al. 2009a). This dataset comprises Computed Tomography (CT) scans of lung tissue from patients undergoing treatment for esophageal cancer, thereby constituting a mono-modal registration task.

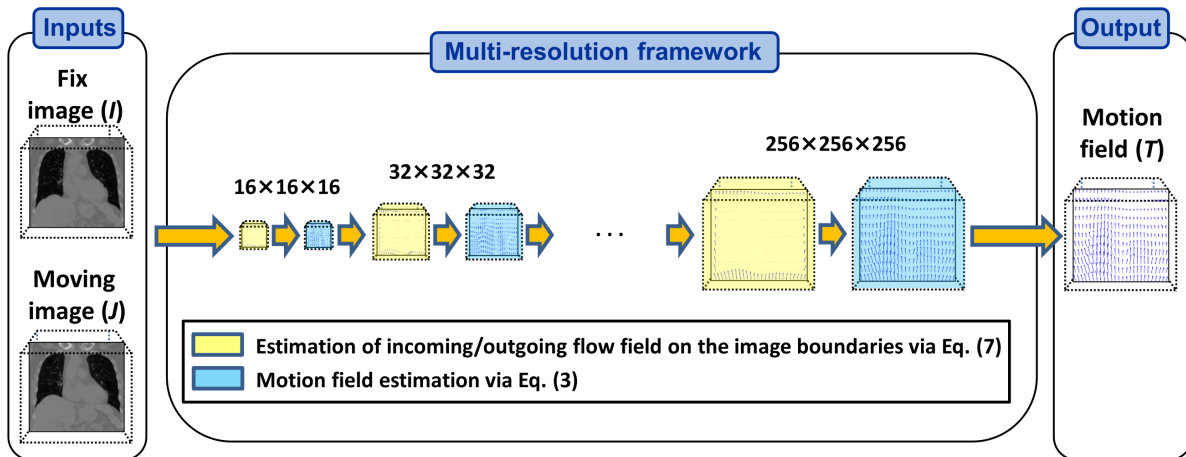


Figure 1: Generic multi-resolution framework for DIR incorporating proposed voxel-wise adaptive boundary conditions. An illustrative input image dimension of $256 \times 256 \times 256$ is displayed as a typical example. The associated pseudo-code is detailed in Appendix C. The minimization of the energy $\mathcal{E}(\mathbf{T})$ was conducted on down-sampled versions of the images, progressively refining until reaching the original image dimension (blue blocks). The computation of incoming/outgoing flow fields on the image boundaries (*i.e.* \mathbf{g}) underwent updates throughout the registration process, at the beginning of each resolution step (yellow blocks). The system of equations (Eq. (3a) and (3b)) was solved in inner voxels (which corresponds to $\Omega \setminus \partial\Omega$), the one voxel wide perimeter (which corresponds to $\partial\Omega$) being computed separately according to the boundary conditions (Eq. (3c), see Appendix B3).

From this database, we selected 10 pairs of CT scans: one acquired during inhalation and the other during exhalation. This dataset presents several challenges stemming from respiratory motion. Significantly larger displacements are evident near the diaphragm in comparison to other regions of the scan, alongside discontinuities at the interface between the lungs and the rib cage. To address this, we designate the scan acquired at extreme inhalation as the fixed image and the scan obtained at extreme exhalation as the moving image.

Each scan includes 300 manually placed landmarks within the lung regions, annotated by domain experts. These landmarks correspond to anatomical features such as blood vessels and bronchial bifurcations (for further details, please refer to (Castillo et al. 2009a)). This kind of groundtruth facilitate the computation of an error map across sparse voxels throughout the lungs, enabling localized assessment of the influence of boundary conditions on the estimation of the motion field.

2.6.2. Multi-modal CT to MRI abdomen registration: The efficacy of the proposed adaptive boundary condition was also assessed in the first task of the Learn2Reg challenge in 2021. This task involved intra-patient CT to MRI multi-modal abdominal registration. The dataset comprised a total of 122 CT and MRI scans, with only

16 of them being paired. Each scan was associated with up to four segmentation masks delineating the liver, both kidneys, and the spleen. As not all paired scans had complete ground-truth segmentation available, seven pairs were selected for the experiment, ensuring inclusion of masks for all specified organs. The fixed and moving images corresponded to MRI and CT scans, respectively.

2.7. Pre-processing:

All images used in this study were pre-processed beforehand by normalizing the intensities to a range between 0 and 1 and cropping the padding when necessary.

2.7.1. Assessment of the motion estimation process The registration outcomes derived from employing HD, HN, and the suggested adaptive Robin boundary condition (Eq. (3c)) are assessed as follows.

Mono-modal CT thorax registration: First, the target registration error (TRE) was computed using the provided 300 gold-standard landmarks in the following manner:

$$\text{TRE} = \|\vec{r}_I + (\mathbf{T}_x(\vec{r}_I), \mathbf{T}_y(\vec{r}_I), \mathbf{T}_z(\vec{r}_I)) - \vec{r}_J\|_2 \quad (10)$$

where \vec{r}_I is the coordinate of the landmark on the fix image, \vec{r}_J is the coordinate of the landmark on the moving image and \mathbf{T}_x , \mathbf{T}_y , \mathbf{T}_z are the displacements in each direction, solution to Eq. (3).

Following this, for each patient, the TRE computed for the 300 landmarks was interpolated using tri-linear interpolation onto a regular grid aligned with the image domain. Subsequently, a mean projection in the front-back direction of the resultant error map was computed to generate a 2D Mean Intensity Projection (MIP) error map.

Multi-modal CT to MRI abdomen registration: In this experiment, we made use of the provided segmentation masks by computing the average of the *Dice* coefficients obtained for the four masks available:

$$\text{Dice}(A, B) = \frac{2(A \cup B)}{|A| + |B|} \quad (11)$$

where A represents the ground-truth segmentation mask on the fixed image, and B represents the ground-truth segmentation mask on the moving image transformed with the estimated motion field. A *Dice* value approaching 1 indicates a strong overlap between the masks.

2.8. Implementation and Hardware

Our experimental setup comprised a 24-core AMD Zen2 EPYC 7402 CPU paired with an Nvidia A100 GPU with 16 GB of memory. We coded our implementation utilizing Python along with the Cupy library (Okuta et al. 2017) for GPU optimization.

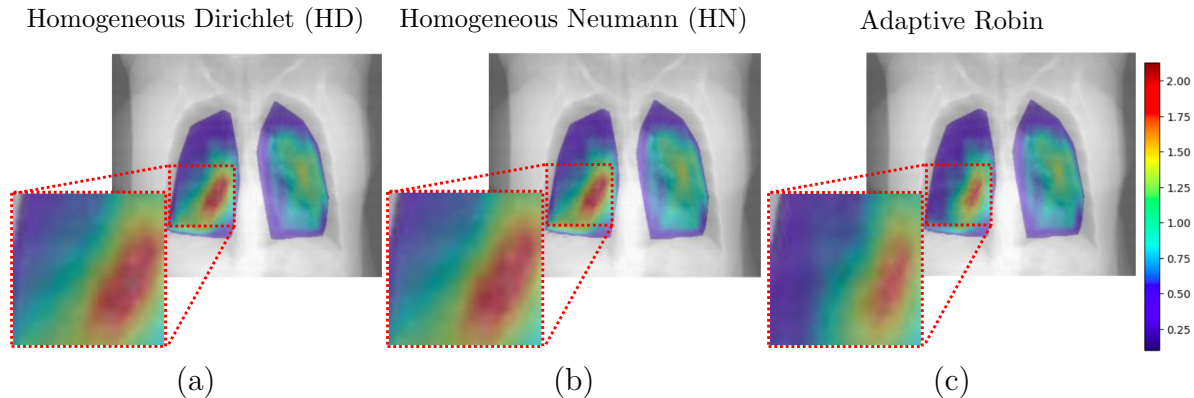


Figure 2: Example of mono-modal CT thorax registration results (case 8). Mean intensity projection of error maps are reported using HN (a), HD (b) and the proposed adaptive Robin conditions (c) (colorbar in millimeters).

3. Results

3.1. Mono-modal CT thorax registration

As depicted in Table 1, the proposed adaptive local boundary condition exhibited enhanced average TRE in contrast to HD and HN conditions ($\approx 4\%$ relative TRE decrease). Particularly notable are the differences in case 7, wherein the adaptive boundary condition yields a 12 % relative TRE enhancement.

Figure 2 provides MIP of error maps obtained using HN, HD, and the proposed adaptive Robin conditions in case 8. The selection of boundary conditions resulted in variations of up to 8 mm. As anticipated, differences are primarily observable at the image boundary, particularly near the diaphragm where substantial upward motion occurs. Notably, some of these differences extend towards the image center, within the lungs, consequently influencing the TRE at landmarks (see Fig. 2).

A typical example of a hyperparameter optimization of the proposed adaptive local boundary condition is presented in Figure 3. The DIR energy $\mathcal{E}(\mathbf{T})$ (Fig. 3a), as defined in Eq. (1), effectively aligns with the TRE metric used for image registration assessment (Fig. 3b). Please note that the DIR energy calculation was restricted to the 300 landmark locations for consistency with the computation of the TRE as defined in Equation (10). For cases 4 to 9, optimal identified hyperparameters consisted in a non-homogeneous Dirichlet boundary conditions ($a = 100$) using incoming/outgoing motion estimates as a source term (γ around 0.6). This tailored parameterization underscores the necessity for substantially greater movement near the diaphragm compared to other image edges. Consequently, a non-homogeneous Dirichlet boundary condition enabling incoming/outgoing flow near the diaphragm was deemed appropriate.

Across the entire dataset, automatic hyperparameterization based on the DIR energy led to TRE that closely approached best achievable values (see extreme right two columns in Table 1).

Table 1: Mean TRE (mm) obtained for each case of the DIR-Lab data set (mono-modal CT thorax registration) and for each tested boundary conditions, with mean and standard deviation (SD). Best scores are highlighted with bold characters.

Case	Homogeneous Neumann	Homogeneous Dirichlet	Adaptive Robin	
			Automatized	Best achievable
1	1.05	1.06	1.05	1.05
2	1.02	1.02	1.02	1.01
3	1.20	1.23	1.20	1.20
4	1.43	1.42	1.42	1.42
5	1.57	1.56	1.51	1.51
6	1.49	1.53	1.42	1.41
7	1.67	1.68	1.47	1.47
8	4.56	4.61	4.39	3.83
9	1.26	1.26	1.24	1.24
10	1.71	1.73	1.71	1.71
Mean \pm SD	1.70 \pm 1.04	1.71 \pm 1.05	1.64 \pm 0.99	1.59 \pm 0.82

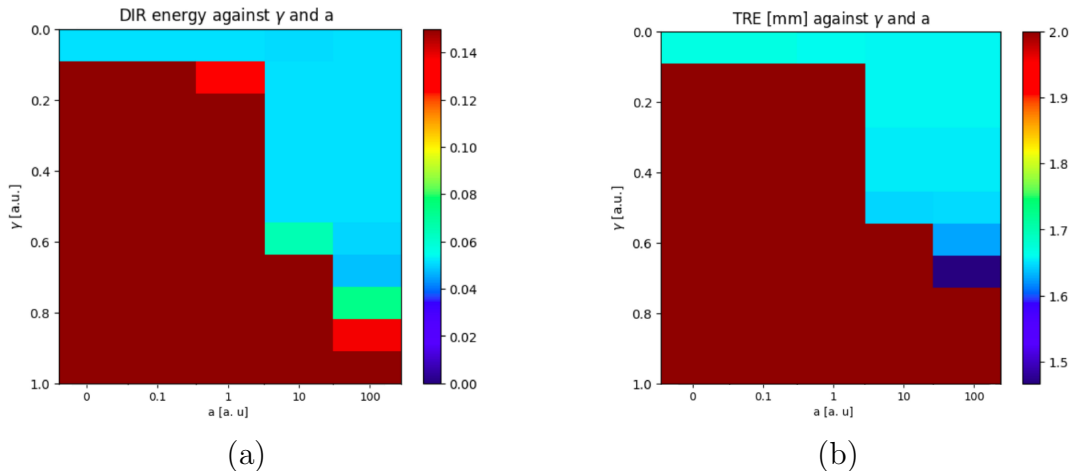


Figure 3: Typical example of hyperparameter optimization (*i.e.* γ and a) for the proposed adaptive Robin boundary conditions (mono-modal CT thorax registration, case 7). (a): grid search optimizing the minimized DIR energy $\mathcal{E}(\mathbf{T})$ as expressed in Eq. (1) (minimum energy yields near optimal hyperparameters); (b): TRE metric calculated using Eq. (10) for image registration assessment (minimum TRE yields optimal registration outcomes).

Table 2: Dice coefficients (averaged over 4 labels) obtained for each case of the Learn2Reg 2021 data set (multi-modal CT to MRI abdomen registration) and for each tested boundary conditions, with mean and standard deviation (SD). Best scores are highlighted with bold characters.

Case	Homogeneous Neumann	Homogeneous Dirichlet	Adaptive Robin	
			Automatized	Best achievable
1	0.846	0.793	0.847	0.847
2	0.695	0.608	0.695	0.695
3	0.855	0.818	0.856	0.856
4	0.866	0.760	0.852	0.866
5	0.722	0.556	0.722	0.756
6	0.810	0.529	0.810	0.810
7	0.849	0.829	0.849	0.849
Mean \pm std	0.806 \pm 0.069	0.699 \pm 0.130	0.804 \pm 0.068	0.811 \pm 0.063

3.2. Multi-modal CT to MRI abdomen registration:

Table 2 presents a summary of *Dice* scores achieved using the HD, HN, and the proposed adaptive local boundary condition across all patients in the dataset.

The HD boundary condition fails to effectively register the top section of the liver, as evidenced by the proximity of the red mask to its initial (unregistered) position (see Fig. 4c). These boundary inaccuracies propagate significantly towards the image center, particularly impacting the top left of the liver. This observation holds true for all patients within the dataset, consistently resulting in inferior outcomes with the Dirichlet boundary. Conversely, both the HN and adaptive Robin conditions yield similarly improved results.

A typical example of hyperparameter optimization for the proposed adaptive local boundary condition is illustrated in Fig. 5 (case 4). The DIR energy (Fig. 5a), as defined in Eq. (1), closely corresponds to the gold standard *Dice* values (Fig. 5b). While the non-homogeneous Dirichlet boundary condition yielded satisfactory results with an average *Dice* score around 0.80 (against 0.87 for best achievable), the optimized parameterization identified during grid search converges the adaptive local boundary condition towards a HN boundary condition ($\gamma = 0$ and $a = 0$). The results reported in Table 2 thus demonstrate that HN boundary conditions are indeed the best fit for this dataset. This customized parameterization emphasizes the necessity for significantly greater movement across all image edges.

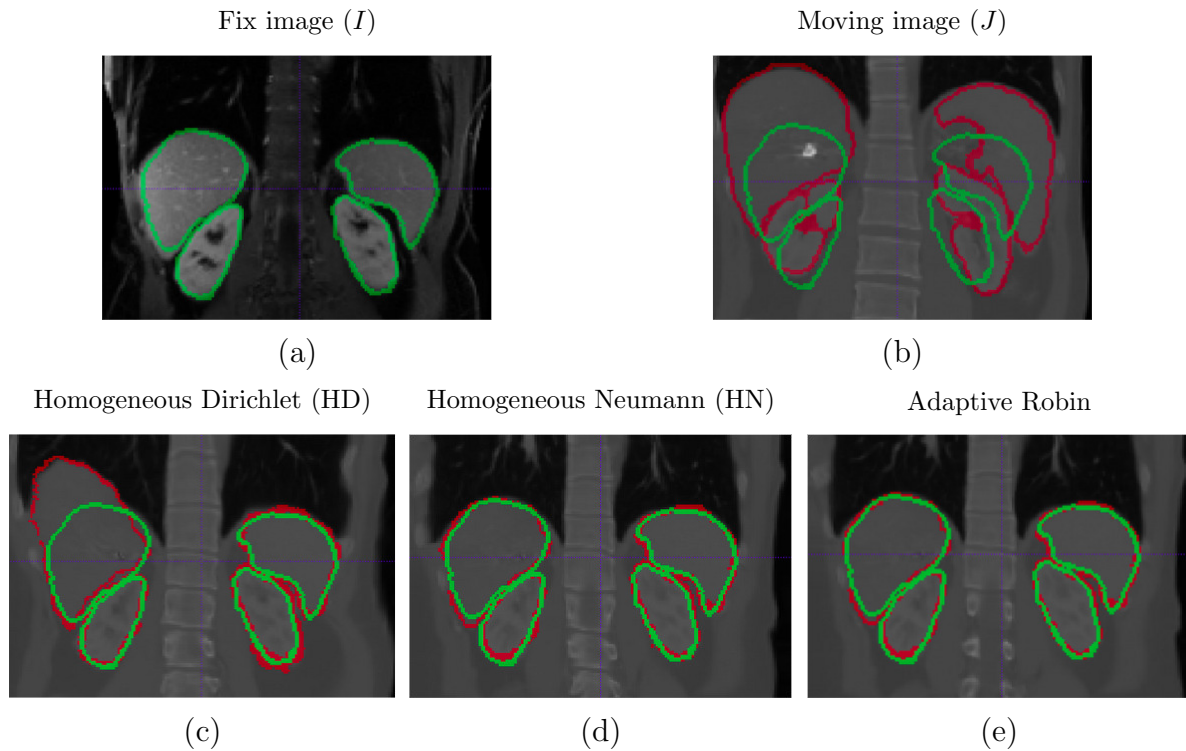


Figure 4: Example of multi-modal CT to MRI abdomen registration (case 4) showing the fix (a), moving (b) and registered images with HD (c), HN (d) and adaptive Robin (e). The segmentation contours are overlaid on all images to highlight the degree of alignment before and after registration (green: fix masks, red: registered masks). The labels include the liver, the spleen and both kidneys.

4. Discussion

To properly formulate a DIR problem applied to images acquired at distinct instants, two strategies are typically employed for handling motion in the image boundaries: either the velocity (Dirichlet) or the shear tensor along the normal (Neumann) are nullified. These global boundary conditions offer the advantage of being implementable using the discrete sine transform and the discrete cosine transform, respectively, leading to a significant reduction in computation time for CPU implementations. Our research underscores the importance of tailoring the DIR algorithm specifically at image boundaries. A generalized approach guided by a reduced set of hyperparameters is proposed to address this concern. Additionally, a straightforward and flexible numerical framework is introduced for adjusting boundary conditions during the registration process.

Presented results indicate that the suggested framework is capable of complete automation via DIR energy minimization. In both the mono-modal CT thorax registration and multi-modal CT to MRI abdomen registration tasks, the proposed adaptive boundary condition consistently produces superior or equivalent results when compared to the commonly employed global boundary conditions, namely HD or HN.




Figure 5: Typical example of hyperparameter optimization (*i.e.* γ and a) for the proposed adaptive Robin boundary conditions (multi-modal CT to MRI abdomen registration/case 4). (a): grid search optimizing the minimized energy $\mathcal{E}(\mathbf{T})$ as expressed in Eq. (1) (minimum energy yields optimal hyperparameters); (b): *Dice* metric calculated using Eq. (11) for image registration assessment (maximum *Dice* yields optimal registration outcomes).

The proposed framework customized the boundary conditions based on case-specific motion characteristics: on the one hand, within the mono-modal CT thorax dataset, significant motion is primarily concentrated near the diaphragm, with minimal motion elsewhere, hence a non-homogeneous Dirichlet boundary condition; on the other hand, in the multi-modal CT to MRI abdomen dataset, minor movements are distributed throughout the image, reflecting its reduced coverage of the abdominal region, hence a HN boundary condition. The adaptive boundary condition still enables the transition to global boundary conditions when necessary, as evidenced in the multi-modal CT to MRI abdomen registration experiments. While hyperparameters a and γ could be optimized individually for each spatial direction, this would have resulted in increased computation times for hyperparameter tuning. In the current study, identical values of hyperparameters a and γ were used across all three spatial dimensions, yet still yielding satisfactory results.

It must be noted that Deep Learning (DL) algorithms are highly valuable tools for DIR. Among its numerous advantages, DL has proven efficient in obtaining data fidelity terms by automatically learning relevant features from images. This can enhance registration accuracy in particularly challenging cases where traditional metrics struggle. The regularization term can also be effectively implemented using DL. Additionally, once trained, DL models can perform image registration rapidly, significantly reducing computational time compared to iterative optimization techniques (Fu et al. 2020). However, we emphasize that the present study does not aim to directly compete with such DL algorithms. Instead, our work focuses on addressing the specific aspect of

boundary conditions, which is relevant to both DL and non-DL algorithms. In this study, we formulate a generic variational DIR algorithm (similar to loss functions used in DL-based DIR approaches) to quantify the enhancements resulting from the proposed local boundary conditions. While highly efficient, DL approaches lack interpretability. Therefore, we decided to illustrate the proposed framework using deterministic data fidelity, regularization, and numerical schemes.

It was observed that the HD, HN, and non-homogeneous Dirichlet boundary conditions yielded the best results in the experiments. Conversely, the non-homogeneous Neumann conditions never produced optimal results. One possibility is that the source term \mathbf{C} is not well-suited for non-homogeneous Neumann conditions. However, since both HN and non-homogeneous Dirichlet conditions produced the best results, it appears that our choice of \mathbf{C} is generally satisfactory. Another potential explanation for the suboptimal performance of non-homogeneous Neumann boundary conditions could be the absence of applied shear stress within the field of view between the fixed and moving images in both tested datasets. In minimally invasive interventional procedures, such as radiofrequency (Zhang et al. 2016) or irreversible electroporation procedures (Gallinato et al. 2019), needles are introduced, generating local displacement due to shear stress in the normal direction. It can be anticipated that non-homogeneous Neumann boundary conditions may be more appropriate in such scenarios. Future studies on percutaneous ablations will further confirm or refute this assertion.

It is important to underline that the datasets evaluated in this study primarily focus on assessing respiratory motion. To our knowledge, there is currently no database specifically designed for evaluating registration in slow physiological motions such as peristalsis. Nevertheless, we anticipate that our method will effectively generalize to accommodate these types of movements as well. Moreover, it must be noticed that the datasets used in this study do not account for scenarios where parts of the organs are not fully covered within the image field of view. In cases where the field of view vary within the same dataset, we anticipate that our method will need individualized hyperparameter optimization to accurately capture motion. However, situations where the field of view differs significantly between the fixed and moving images may require careful investigation (Sadowsky et al. 2007) (Greve & Fischl 2009), as our DIR formulation assumes that most information present in one image is also visible in the other.

In this paper, we estimated incoming and outgoing flow fields utilizing the motion inverse consistency metric. Our underlying hypothesis is that discrepancies between the forward and backward motion estimates are interpreted as material flow exiting or entering the field of view. The utilization of motion inverse consistency was driven by its practical advantage of relying solely on motion estimates. This characteristic makes this metric compatible with both mono- and multi-modal registration tasks, as it does not rely on image intensity. Furthermore, it does not require any additional information such as landmarks or image segmentations. Nevertheless, this method may result in erroneous estimation of incoming/outgoing flow fields in several scenarios. For instance,

when there is limited or no tissue contrast present at the image boundary, false negative flows may be identified, as the registration process is predominantly influenced by the regularization term. Moreover, false positive flows may be detected in the presence of transient objects due to fluctuations in the data fidelity term. For these reasons, the motion inverse consistency metric may not be suitable in image regions involving needle insertions, as encountered in percutaneous minimally invasive interventional procedures. Application specific solution may be investigated thanks to the highly modular aspect of the proposed framework expressed in Eq. (4).

In the current study, we conducted an exhaustive search for hyperparameter determination to thoroughly document the behaviour of the registration algorithm with respect to the boundary conditions. However, the grid search strategy on our experimental setup using our implementation proved to be time-consuming. The computation time for hyperparameter determination is inherently influenced by the processing time required for image registration. For example, it took approximately 4 minutes to register a $192 \times 149 \times 142$ image using our experimental setup. Consequently, the entire grid search process took about four hours, rendering it impractical for clinical use. However, our results indicate that various strategies can be employed in practice to significantly reduce the computation time required for hyperparameter determination. First, under similar image acquisition conditions, hyperparameter optimization may be performed only once, as it will result in similar parameter settings. For instance, in the mono-modal CT thorax (DIR lab) dataset, significant motion is primarily concentrated near the diaphragm, with minimal motion elsewhere. Hence, the parameterization of the proposed boundary condition mainly corresponds to a non-homogeneous Dirichlet condition with $\gamma = 0.6$. In the multi-modal CT to MRI abdomen dataset, minor movements are distributed throughout the image, reflecting its reduced coverage of the abdominal region, which mainly favors the homogeneous Neumann boundary condition. For better results, the optimisation can be refined for each image pair. The processing time for hyperparameter determination can be reduced by restricting the search area. For example, in the multi-modal CT to MRI abdomen (Learn2Reg) dataset, γ can be set to zero, allowing the optimization to focus on the parameter a , which defines the trade-off between Dirichlet and Neumann conditions. **Notably, multi-step and multi-resolution search strategies may also be advantageous by providing coarse estimates of hyperparameters, which can then be locally refined. For example, considering that the inhomogeneous Neumann conditions could be discarded for the two datasets involved in this study (consistent with the absence of imposed external inhomogeneous displacement, such as percutaneous insertion), the following three extremities of the hyperparameter grid can be initially considered: NH ($a=0, \gamma=0$), DH ($a=100, \gamma=0$), and inhomogeneous Dirichlet ($a=100, \gamma=1$). Subsequently, the two best pairs can be refined using a bisection method. After three levels of successive subdivisions, the complete optimization process requires the calculation of six DIR, which could be accomplished in 24 minutes on our test platform. It is also noteworthy that performing the optimization on 2-fold downsampled images further reduces the computation time**

by a factor of 4 without degrading the optimum hyperparameter calculation, thereby reducing the computation time for case-specific hyperparameter tuning to 6 minutes on our test platform. We anticipate that the hyperparameter tuning procedure can also be accelerated using a deep learning approach, utilizing the DIR functional as the loss function for unsupervised learning. This will be addressed in a forthcoming study. Ultimately, the computation time for hyperparameters can be significantly reduced by decreasing the processing time required for DIR. For example, learning-based approaches can conveniently circumvent the requirement to repeat the entire intensive DIR optimization scheme for each new image pair (Fu et al. 2020).

5. Conclusion

A DIR toolbox is essential in the field of medical imaging due to the growing need to align images of the same scene acquired across time and/or modalities. This study emphasizes the significance of customizing the DIR algorithm at image boundaries. A generalized formulation for this issue is presented, along with a direct and adaptable framework for tailoring boundary conditions throughout the registration process. Its adaptability stems from its capacity to shift between homogeneous Neumann, homogeneous Dirichlet, homogeneous Robin, and their non-homogeneous counterparts, conditioned by a reduced number of hyper-parameters. Additionally, we demonstrated the feasibility of fully automatizing such boundary conditions.

The suggested local adaptive Robin boundary condition yields results closely approaching the best achievable outcomes for both examined mono-modal CT thorax registration and multi-modal CT to MRI abdomen registration tasks.

The impact of boundary conditions on more complex motion regularization term and different DIR formulations will need to be addressed in future works. Moreover, future studies will challenge the proposed method for other non-rigid multi-sensor registration scenarios (*i.e.* CT to CBCT, MRI to CBCT, MRI to echography, etc). A crucial and forthcoming perspective will involve investigating the impact of dose delivery, particularly in procedures such as radiotherapy (Grégoire et al. 2020), thermotherapy (Zhang et al. 2016) and irreversible electroporation (Gallinato et al. 2019).

Acknowledgements

This study has been partly granted by the Plan Cancer Projet MECI, n°21CM119-00. For the purpose of Open Access, a CC-BY public copyright licence has been applied by the authors to the present document and will be applied to all subsequent versions up to the Author Accepted Manuscript arising from this submission.

Experiments presented in this paper were carried out using the PlaFRIM experimental testbed, supported by Inria, CNRS (LABRI and IMB), Université de Bordeaux, Bordeaux INP and Conseil Régional d'Aquitaine (see <https://www.plafrim.fr/>).

Appendix

A. Brief overview of DIR methods

Deformable image registration is an image processing task that consists in estimating the motion to align a moving image with a fixed image, taking into account the content of both images and sometimes prior knowledge of the motion and/or the tissue’s physical properties. It has been an active area of research in the field of image processing since the 1960s, owing to its extensive range of applications. In pursuit of this goal, a variety of methods are available including the followings.

In parametric motion estimation, the motion field is parameterized using either generic, simpler motions such as translation, rotation, scaling, or more comprehensive parameterizations like B-splines (Klein et al. 2007) (Vandemeulebroucke et al. 2011) (Vishnevskiy et al. 2017).

In variational methods, the problem is formulated as a functional optimization. The energy functional to be minimized consists of a data fidelity term, which drives the registered image closer to the fixed image, and a regularization term, which addresses the ill-posed nature of the problem and potentially enhances the anatomical plausibility of the solution (Castillo et al. 2009b) (Muenzing et al. 2010) (Denis de Senneville et al. 2016).

Machine learning, increasingly used in image processing, can also assist in solving registration problems in several ways. Such approach has proven efficient in obtaining task-specific data fidelity terms by automatically learning relevant features from images and solving the energy functional more efficiently (Cao et al. 2018) (Fan et al. 2019) (Balakrishnan et al. 2019) (Fechter & Baltas 2020) (Fu et al. 2020).

B. Numerical implementation of the boundary conditions

B1. Homogeneous Dirichlet

For homogeneous Dirichlet, we simply set the value 0 in the one voxel wide image perimeter (which corresponds to $\partial\Omega$).

B2. Homogeneous Neumann

To apply a homogeneous Neumann condition, $\nabla\mathbf{T}_s \cdot \mathbf{n}$ for $s \in \{x, y, z\}$ was computed using the following second-order approximation with a spatial step size ds :

$$\nabla\mathbf{T}_s^0 \cdot \mathbf{n} = -\frac{-3\mathbf{T}_s^0 + 4\mathbf{T}_s^1 - \mathbf{T}_s^2}{2ds} \quad (\text{B1})$$

On the other side of the boundary, one have:

$$\nabla\mathbf{T}_s^N \cdot \mathbf{n} = \frac{\mathbf{T}_s^{N-2} - 4\mathbf{T}_s^{N-1} + 3\mathbf{T}_s^N}{2ds} \quad (\text{B2})$$

where the superscript gives the number of voxels along the direction considered.

Substituting into $\nabla \mathbf{T}_s^{0,N} \cdot \mathbf{n} = 0$ and solving for \mathbf{T}_s^0 and \mathbf{T}_s^N respectively, we obtain:

$$\mathbf{T}_s^0 = \frac{4\mathbf{T}_s^1 - \mathbf{T}_s^2}{3} \quad (\text{B3})$$

$$\mathbf{T}_s^N = \frac{4\mathbf{T}_s^{N-1} - \mathbf{T}_s^{N-2}}{3} \quad (\text{B4})$$

These conditions were imposed in the one voxel wide image perimeter (which corresponds to $\partial\Omega$).

B3. Adaptive Robin boundary condition

The adaptive Robin boundary condition takes the form:

$$\beta(\mathbf{g}_s)(\nabla \mathbf{T}_s \cdot \mathbf{n}) + (1 - \beta(\mathbf{g}_s))\mathbf{T}_s = \gamma \mathbf{g}_s \quad (\text{B5})$$

where \mathbf{n} is the normal to the boundary $\delta\Omega$, $\beta : \mathbb{R} \rightarrow [0, 1]$ is computed using Eq. (5), and γ is one of the two hyper-parameters for the adaptive Robin boundary condition.

Making the same approximation as in Appendix B2, substituting in Eq. (B5) and solving for \mathbf{T}_s^0 and for \mathbf{T}_s^N with $ds = 1$:

$$\mathbf{T}_s^0 = \frac{\beta_s^0(4\mathbf{T}_s^1 - \mathbf{T}_s^2) + 2\gamma \mathbf{g}_s^0}{\beta_s^0 + 2} \quad (\text{B6})$$

$$\mathbf{T}_s^N = \frac{\beta_s^N(4\mathbf{T}_s^{N-1} - \mathbf{T}_s^{N-2}) + 2\gamma \mathbf{g}_s^N}{\beta_s^N + 2} \quad (\text{B7})$$

These conditions were imposed in the one voxel wide image perimeter (which corresponds to $\partial\Omega$).

C. Pseudo-code

Algorithm 1 corresponds to the primary framework for registration for typical $256 \times 256 \times 256$ images, encompassing the calculation of the voxelwise incoming/outgoing boundary flow designed to guide boundary conditions, the initialization of the motion field for the forthcoming estimation process at the current resolution, the registration algorithm, and the multi-resolution scheme (see Fig. 1).

Algorithm 2 computes the incoming/outgoing flow field \mathbf{g} via Eq. (7). It corresponds to the yellow blocks delineated in Figure 1. Note that i is a local variable determining the iteration number (max=10).

Algorithm 3 estimates the motion field between fixed and moving images through the implemented DIR algorithm, which incorporates the suggested local Robin boundary conditions, via Eq. (3). It corresponds to the blue blocks delineated in Figure 1. Note that i is a local variable determining the iteration number (max=1000).

Algorithm 1: Primary framework for registration

Input : $(I, J, \mathbf{T}^{init}) = (\text{fix image, moving image, initial condition})$
Output: $\mathbf{T} = \text{motion field}$
 $\mathbf{T} \leftarrow \text{Resample } \mathbf{T}^{init} \text{ to } (16 \times 16 \times 16)$
for $res \in \{(16 \times 16 \times 16), (32 \times 32 \times 32), \dots, (256 \times 256 \times 256)\}$ **do**
 $I' \leftarrow \text{Resample } I \text{ to } res$
 $J' \leftarrow \text{Resample } J \text{ to } res$
 $\mathbf{T} \leftarrow \text{Resample } \mathbf{T} \text{ to } res$
 $\mathbf{g}, \mathbf{T} \leftarrow \text{Algorithm 2}(I', J', \mathbf{T})$
 $\mathbf{T} \leftarrow \text{Algorithm 3}(I', J', \mathbf{T}, \mathbf{g})$
end

Algorithm 2: Calculation of incoming/outgoing flow field in the boundary

Input : $(I, J, \mathbf{T}) = (\text{fix image, moving image, motion field})$
Output: $\mathbf{g} = \text{Voxelwise in/out flow field in the boundary}$
 $\mathbf{T} = \text{motion field}$
 $D \leftarrow \text{data fidelity term calculated from } I, J \text{ and } \mathbf{T} \text{ using Eq. (8)}$
 $\mathbf{T}_{inv} \leftarrow \mathbf{T}$
 $J_{cur} \leftarrow J \text{ transformed with } \mathbf{T}$
 $i \leftarrow 0$
while $i < 10$ **do**
 Update \mathbf{T} boundaries according to homogeneous Neumann boundary conditions (Appendix B2)
 $\mathbf{T} \leftarrow \mathcal{L}(I, J_{cur}, D, \mathbf{T})$
 $i \leftarrow i + 1$
end
 $i \leftarrow 0$
while $i < 10$ **do**
 Update \mathbf{T}_{inv} boundaries according to homogeneous Neumann boundary conditions (Appendix B2)
 $\mathbf{T}_{inv} \leftarrow \mathcal{L}(J_{cur}, I, D, \mathbf{T}_{inv})$
 $i \leftarrow i + 1$
end
 $\mathbf{g} \leftarrow \text{incoming/outgoing field map calculated from } \mathbf{T} \text{ and } \mathbf{T}_{inv} \text{ using Eq. (7)}$

Algorithm 3: Calculation of the motion field estimate

Input : $(I, J, \mathbf{T}, \mathbf{g}) =$ (fix image, moving image, motion field,
incoming/outgoing flow)

Output: \mathbf{T} = motion field

$i \leftarrow 0$

while $i < 1000$ **do**

$D \leftarrow$ data fidelity term calculated from I, J and \mathbf{T} using Eq. (8)

Update \mathbf{T} boundaries according to \mathbf{g} (Appendix B3)

$J_{cur} \leftarrow J$ transformed with \mathbf{T}

$\mathbf{T} \leftarrow \mathcal{L}(I, J_{cur}, D, \mathbf{T})$

Update \mathbf{T} boundaries according to \mathbf{g} (Appendix B3)

$i \leftarrow i + 1$

end

References

- Abbott, M. & Basco, D. (1989). *Computational Fluid Dynamics: An Introduction for Engineers*, pp. 1–425.
- Balakrishnan, G., Zhao, A., Sabuncu, M. R., Guttag, J. & Dalca, A. V. (2019). Voxelmorph: A learning framework for deformable medical image registration, *IEEE Transactions on Medical Imaging* **38**(8): 1788–1800.
- Braumann, U. D. & Kuska, J. P. (2005). Influence of the boundary conditions on the result of non-linear image registration, *IEEE International Conference on Image Processing 2005*, Vol. 1, pp. 1129–1132.
- Brox, T., Rosenhahn, B., Gall, J. & Cremers, D. (2009). Combined region and motion-based 3d tracking of rigid and articulated objects, *IEEE transactions on pattern analysis and machine intelligence* **32**(3): 402–415.
- Cao, X., Yang, J., Zhang, J., Wang, Q., Yap, P. T. & Shen, D. (2018). Deformable image registration using a cue-aware deep regression network, *IEEE Transactions on Biomedical Engineering* **65**(9): 1900–1911.
- Castillo, E., Castillo, R., Guerra, R., Johnson, V., McPhail, T., Garg, A. & Guerrero, T. (2009a). A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets, *Physics in Medicine and Biology* **54**: 1849–70.
- Castillo, E., Castillo, Zhang, Y. & Guerrero, T. (2009b). Compressible image registration for thoracic computed tomography images, *Journal of Medical and Biological Engineering* **29**(5): 222–233.
- Chetty, I. J. & Rosu-Bubulac, M. (2019). Deformable registration for dose accumulation, *Seminars in Radiation Oncology* **29**(3): 198–208. Adaptive Radiotherapy and Automation.
- Denis de Senneville, B., Zachiu, C., Ries, M. & Moonen, C. T. W. (2016). Evolution: an edge-based variational method for non-rigid multi-modal image registration, *Physics in Medicine and Biology* **61**(20): 7377.
- Fan, J., X., C., Wang, Q., Yap, P. T. & Shen, D. (2019). Adversarial learning for mono- or multi-modal registration, *Medical Image Analysis* **58**: 101545.
- Fechter, T. & Baltas, D. (2020). One-shot learning for deformable medical image registration and periodic motion tracking, *IEEE Transactions on Medical Imaging* **39**(7): 2506–2517.
- Fu, Y., Lei, Y., Wang, T., Curran, W. J., Liu, T. & Yang, X. (2020). Deep learning in medical image registration: a review, *Physics in Medicine & Biology* **65**(20): 20TR01.
- Gallinato, O., Denis de Senneville, B., Seror, O. & Poignard, C. (2019). Numerical workflow

- of irreversible electroporation for deep-seated tumor, *Physics in Medicine and Biology* **64**(5): 055016.
- Greve, D. N. & Fischl, B. (2009). Accurate and robust brain image alignment using boundary-based registration, *Neuroimage* **48**(1): 63–72.
- Grégoire, V., Guckenberger, M., Haustermans, K., Lagendijk, J. J. W., Ménard, C., Pötter, R., Slotman, B. J., Tanderup, K., Thorwarth, D., van Herk, M. & Zips, D. (2020). Image guidance in radiation therapy for better cure of cancer, *Mol Oncol* **14**(7): 1470–1491.
- Klein, S., Staring, M. & Pluim, J. P. W. (2007). Evaluation of optimization methods for nonrigid medical image registration using mutual information and b-splines, *IEEE Transactions on Image Processing* **16**(12): 2879–2890.
- Kok, H. P., Cressman, E. N. K., Ceelen, W., Brace, C. L., Ivkov, R., Grüll, H., Ter Haar, G., Wust, P. & Crezee, J. (2020). Heating technology for malignant tumors: a review, *Int J Hyperthermia* **37**(1): 711–741.
- Larrey-Ruiz, J., Morales-Sánchez, J. & Verdú-Monedero, R. (2007). Generalized regularization term for non-parametric multimodal image registration, *Signal Processing* **87**(11): 2837–2842.
- Muenzing, S., van Ginneken, B. & Pluim, J. (2010). Knowledge driven regularization of the deformation field for pde based non-rigid registration algorithms, *Medical Image Analysis for the Clinic - A Grand Challenge* pp. 127–136.
- Okuta, R., Unno, Y., Nishino, D., Hido, S. & Crissman (2017). Cupy : A numpy-compatible library for NVIDIA GPU calculations, *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)* .
URL: <https://api.semanticscholar.org/CorpusID:41278748>
- Richard, F. J. P. & Cohen, L. D. (2003). A new image registration technique with free boundary constraints: application to mammography, *Computer Vision and Image Understanding* **89**(2): 166–196. Nonrigid Image Registration.
- Sadowsky, O., Chintalapani, G. & Taylor, R. H. (2007). Deformable 2d-3d registration of the pelvis with a limited field of view, using shape statistics, *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2007: 10th International Conference, Brisbane, Australia, October 29–November 2, 2007, Proceedings, Part II 10*, Springer, pp. 519–526.
- Staring, M., Klein, S. & Pluim, J. P. W. (2007). A rigidity penalty term for nonrigid registration, *Medical Physics* **34**(11): 4098–4108.
- Vandemeulebroucke, J., Rit, S., Kybic, J., Clarysse, P. & Sarrut, D. (2011). Spatio-temporal motion estimation for respiratory-correlated imaging of the lungs, *Medical physics* **38**: 166–78.
- Vishnevskiy, V., Gass, T., Szekely, G., Tanner, C. & Goksel, O. (2017). Isotropic total variation regularization of displacements in parametric image registration, *IEEE Transactions on Medical Imaging* **36**(2): 385–395.
- Yuen, J., Barber, J., Ralston, A., Gray, A., Walker, A., Hardcastle, N., Schmidt, L., Harrison, K., Pöder, J., Sykes, J. R. & Jameson, M. G. (2020). An international survey on the clinical use of rigid and deformable image registration in radiotherapy, *Journal of Applied Clinical Medical Physics* **21**(10): 10–24.
- Zachiu, C., Denis de Senneville, B., Moonen, C. T. W. & Ries, M. (2015). A framework for the correction of slow physiological drifts during mr-guided hifu therapies: Proof of concept, *Med Phys* **42**(7): 4137–48.
- Zhang, B., Moser, M. A., Zhang, E. M., Luo, Y., Liu, C. & Zhang, W. (2016). A review of radiofrequency ablation: Large target tissue necrosis and mathematical modelling, *Physica Medica* **32**(8): 961–971.
- Zhang, T., Orton, N. P., Mackie, T. R. & Paliwal, B. R. (2004). Technical note: A novel boundary condition using contact elements for finite element based deformable image registration, *Medical physics* **31**(9): 2412–5.