

DEFORMABLE IMAGE REGISTRATION WITH STOCHASTICALLY REGULARIZED BIOMECHANICAL EQUILIBRIUM

Pablo Alvarez ^{a,b,*} and Stéphane Cotin ^{a,b}

^a Inria, Strasbourg, France

^b ICube, University of Strasbourg, CNRS, Strasbourg, France

ABSTRACT

Numerous regularization methods for deformable image registration aim at enforcing smooth transformations, but are difficult to tune-in *a priori* and lack a clear physical basis. Physically inspired strategies have emerged, offering a sound theoretical basis, but still necessitating complex discretization and resolution schemes. This study introduces a regularization strategy that does not require discretization, making it compatible with current registration frameworks, while retaining the benefits of physically motivated regularization for medical image registration. The proposed method performs favorably in both synthetic and real datasets, exhibiting an accuracy comparable to current state-of-the-art methods.

Index Terms— Image registration, physically inspired regularization, hyperelasticity

1. INTRODUCTION

Deformable image registration (DIR) of medical images has been a subject of research for decades, given its great potential for clinical applications such as disease diagnosis [1] and interventional guidance [2], to name a few. Despite the important progress made in recent years, DIR methods still lack efficiency and robustness, which limits their transferability into clinical practice. This stems from the fact that DIR is an ill-posed problem, partly because of the characteristics of medical images (*e.g.* low contrast and noise), but also because the underlying organs may deform considerably. To address this issue, multiple regularization strategies have been proposed [3], many of which were designed specifically to strengthen transformation smoothness. Although this is a desirable property, overly-smooth transformations can reduce registration accuracy, and properly adjusting regularization strength may become an important issue. Furthermore, it is difficult to provide a physical justification for some of these commonly used regularization strategies [4].

In response, physically motivated regularization has been introduced. These methods consider the organs in medical images as hyperelastic bodies, and exploit the theory of

continuum mechanics for the design of regularization strategies [4, 5]. In addition to their strong physical basis, they also inherently strengthen desirable properties such as transformation smoothness and invertibility. However, to date, they still require complex domain discretization and resolution schemes. In this work, we extend on these ideas to propose a physically-derived regularization strategy that can be computed locally (pointwise) and does not require any discretization. As such, it facilitates integration with existing DIR frameworks and optimizers, while maintaining the desirable properties of physically motivated regularization.

1.1. The ill-posed optimization problem

Let \mathcal{F}, \mathcal{M} be the d -dimensional fixed and moving images, defined in domains Ω_0 and Ω , respectively. The registration problem consists in finding a displacement field \mathbf{u} , or equivalently, a transformation $\phi : \mathbf{x} \in \Omega_0 \mapsto \mathbf{x} + \mathbf{u}(\mathbf{x})$, that aligns structures in Ω_0 to homologous structures in Ω , *i.e.* $\mathcal{F}(\mathbf{x})$ and $\mathcal{M}(\phi(\mathbf{x}))$ are close under some similarity criterion. This is typically formulated as an optimization problem:

$$\hat{\phi} = \arg \min_{\phi} (1 - \beta) \mathcal{L}_{sim}(\phi; \mathcal{F}, \mathcal{M}) + \beta \mathcal{L}_{reg}(\phi) \quad (1)$$

where \mathcal{L}_{sim} is a metric measuring the similarity between the images, \mathcal{L}_{reg} is a regularization term and β is a normalized weighting coefficient.

In general, DIR is ill-posed, and many similarity and regularization terms have been proposed in an effort to overcome this problem [3]. In this work, we focus our attention on the regularization term \mathcal{L}_{reg} , which is of great importance for high quality and physically plausible registration results.

1.2. Regularization

In medical image registration, it is important that the estimated transformation ϕ is diffeomorphic, since it is the minimum requirement for physical plausibility. In practice, this implies that ϕ is both smooth and invertible. Various regularization strategies have been proposed to directly reinforce these properties. For instance, the diffusion [6], curvature [7] and bending energy [8] regularizers were designed to strengthen smoothness. However, in the context

* Corresponding author: pablo.alvarez@inria.fr

of deforming organs, their rather topological motivation is difficult to justify in physical terms. An alternative approach is to consider the underlying deforming objects as elastic bodies, and directly use their strain energy to regularize the transformation ϕ , on linearly-elastic or hyperelastic settings [4]. However, using strain energy as regularization penalizes deformation, since only rigid body motion yields zero strain energy. A more recent idea is to consider DIR as an elastostatic problem, and penalize any deviation from the solution of a hyperelastic body in equilibrium with arbitrary boundary tractions, but without body forces [5]. This equilibrium gap regularization principle, allows for large deformations, provided that they are compatible with the conservation of linear momentum of the above-mentioned elastostatic problem:

$$\nabla \cdot \left(\frac{\partial \Psi}{\partial \mathbf{F}} \right) = \nabla \cdot \mathbf{P} = \mathbf{0}, \quad (2)$$

with Ψ the strain energy density function, $\mathbf{F} = \nabla \phi$ the deformation gradient tensor and \mathbf{P} the first Piola-Kirchhoff stress tensor. As such, regularization with the equilibrium gap principle is achieved by defining

$$\mathcal{L}_{reg} = \|\nabla \cdot \mathbf{P}\|^2. \quad (3)$$

To the best of our knowledge, only Finite Element (FE) formulations of this regularization approach exist to date [5, 9]. As with the FE method in general, the accuracy depends on the quality and level of discretization, and (3) can only be enforced weakly. In this work, we propose a local (pointwise) implementation of the equilibrium gap regularization in (3) that does not require any FE discretization and can thus be enforced strongly.

2. MATERIALS AND METHODS

2.1. Analytically regularizing the equilibrium gap

Let Ψ be the strain energy density function of an isotropic, homogeneous, compressible NeoHookean material:

$$\Psi = \frac{\lambda}{4}(J^2 - 1 - 2 \ln(J)) + \frac{\mu}{2}(I_C - 3 - 2 \ln(J)), \quad (4)$$

with $I_C = \text{tr}(\mathbf{C})$ the first invariant of the right Cauchy-Green deformation tensor $\mathbf{C} = \mathbf{F}^T \mathbf{F}$, $J = \det(\mathbf{F})$, and μ and λ the so-called *Lamé parameters*.

From the constitutive equation (4), we first derive an expression for the stress of the material,

$$\mathbf{P} = \frac{\partial \Psi}{\partial \mathbf{F}} = \frac{\lambda}{2}(J^2 - 1)\mathbf{F}^{-T} + \mu(\mathbf{F} - \mathbf{F}^{-T}), \quad (5)$$

and then take its divergence to rewrite (2) as:

$$\begin{aligned} \nabla \cdot \mathbf{P} &= \frac{\lambda}{2} (2J\mathbf{F}^{-T} \nabla J + (J^2 - 1)\nabla \cdot \mathbf{F}^{-T}) \\ &\quad + \mu (\nabla \cdot \mathbf{F} - \nabla \cdot \mathbf{F}^{-T}) = \mathbf{0}. \end{aligned} \quad (6)$$




Fig. 1. Synthetic dataset generation. FE meshes in undeformed (a) and deformed (b) configurations for a hyperelastic unit square with random Dirichlet boundary conditions. The generated fixed (c) and moving (d) binary images.

It should be noted that the computation of (6) at a material point necessitates the first- and second-order derivatives of the displacement field \mathbf{u} . In a FE formulation, this is possible through the so-called shape functions of the FE spaces. Here, we suppose that analytical expressions for such derivatives are available, and therefore the computation of the equilibrium gap (3) can be performed directly. This is the case, for instance, for deformable image registration methods based on parametric transformations [10]. In this work, we exploit the recently developed Implicit Neural Representations for deformable image registration (IDIR) for our implementation.

2.2. Images and validation data

2.2.1. Synthetic hyperelasticity dataset

This dataset consists of three 2-dimensional deformation scenarios of a hyperelastic unit square constrained at its boundary with randomly generated, smooth displacements. The variational formulation of the elastostatic problem (under plane strain assumptions) described by (2) and (4) was implemented in FeniCs [11]. The displacement fields for the Dirichlet boundary conditions were generated from the sum of 12 randomly parameterized Gaussian radial basis functions with experimentally defined parameter bounds. The simulations were performed with a FE 40×40 grid mesh of quadratic quad elements, containing a total of 1600 elements and 14641 mesh nodes. We considered a highly compressible material with a unit Young's modulus and a Poisson ratio of 0.3, yielding approximately to $\mu = 0.577$ and $\lambda = 0.385$. For each deformation scenario, fixed and moving binary images were generated from the solution of the elastostatic problem as illustrated in Fig. 1. Note that similar 2D binary datasets have been used in the past for validating image registration methods, *e.g.* [4]. Our dataset, however, differs from these works in two key aspects: (i) the object is considered hyperelastic, which entails some constraints in the resulting deformation fields; and (ii) the ground truth displacement is known at each mesh node, which facilitates quantitative evaluation of registration performance.




Fig. 2. Registration accuracy on synthetic dataset for bending energy regularization (Bending) and the proposed analytical equilibrium gap regularization (Physics) with varying β , across deformation scenarios. Mean (solid line), first quantile (lower bound) and third quantile (upper bound) of registration error as percentage of unit length.

2.2.2. DIRLab 4DCT dataset

The DIRLab 4DCT dataset [12] contains ten cases of lung CT images taken at various phases of a respiratory cycle, with varying image resolutions ranging from 256×256 to 512×512 for the in-plane resolution and 256 or 512 for the number of slices. As in previous works, we evaluated registration performance only using the end-inhalation (fixed) and end-exhalation (moving) images, for which a set of 300 paired anatomical landmarks is available. In addition, we extracted lung masks from each image using an automatic method [13] in order to define regions of interest for transformation optimization, as is common practice in the literature [14, 15].

2.3. Implementation within the IDIR framework

Recently, implicit neural representations for deformable image registration (IDIR) have been proposed [15, 16]. These learning techniques implicitly represent the transformation ϕ through neural networks, using as input spatial coordinates \mathbf{x} , and as output the corresponding displacements \mathbf{u} . When combined with periodic activation functions, these networks allow the computation of high-order spatial derivatives of the displacement field \mathbf{u} , therefore providing a framework for the implementation of complex regularization schemes as the proposed analytical equilibrium gap regularizer.

We implemented the IDIR framework in PyTorch following [15]. Given the invariance of IDIR to image resolution, we used the same architecture for all registration problems: 3 hidden layers of 256 sine-activated units and d -dimensional input/output layers, accordingly. The network weights were initialized following [17]. For each registration task, the optimization problem (1) was solved stochastically using the Adam optimizer and 10000 randomly sampled points, pooled from the whole fixed domain for the synthetic dataset, and from the lung masks for the DIRLab 4DCT dataset.

The mean squared error and normalized cross-correlation




Fig. 3. Qualitative registration results on synthetic dataset for bending energy regularization (solid blue) and the proposed analytical equilibrium gap regularization (dashed orange) with varying β , on top of the ground truth deformed geometry (solid black).

were used as \mathcal{L}_{sim} for the synthetic and DIRLab 4DCT datasets, respectively. For the regularization \mathcal{L}_{reg} , the bending energy [8] and our proposed analytical equilibrium gap regularizer were implemented. For the latter, we chose a unit Young’s modulus and zero Poisson’s ratio, yielding $\mu = \frac{1}{2}, \lambda = 0$. This choice reduces the registration bias for the synthetic dataset (same hyperelastic model but different parameters), and also avoids over-penalization of volume changes that are expected in the DIRLab 4DCT dataset.

3. RESULTS

3.1. Synthetic hyperelasticity dataset

Considering that the functional of the optimization problem (1) is a convex combination of the similarity \mathcal{L}_{sim} and regularization \mathcal{L}_{reg} terms, the regularization weight β affects both. Therefore, we chose values spanning the unit interval to evaluate the influence of the regularization weight in registration accuracy, namely $\beta \in \{0.0, 0.001, 0.01, 0.1, 0.5, 0.9\}$. In Fig. 2, the distribution of percent error across all deformation scenarios is presented; and, to further illustrate such influence, Fig. 3 depicts the result of forward-warping a uniform grid with the estimated transformations for one of the three deformation scenarios.

3.2. DIRLab 4DCT dataset

We evaluated registration accuracy in the DIRLab dataset using snap-to-voxel target registration errors (TRE) for the 300 landmarks available per case, for both bending energy and our proposed analytical equilibrium gap regularization strategies. The results of this quantitative evaluation are presented in Table 1, along with those for the current state-of-the-art method




Fig. 4. Cumulative distribution of TRE with bending energy regularization (Bending) and the proposed analytical equilibrium gap regularization (Physics) with varying β for the DIRLab 4DCT dataset.

in the DIRLab dataset [18] for comparison. The regularization weight for this evaluation was set to $\beta = 0.001$, which was found to provide optimal results across cases.

Table 1. Mean (standard deviation) snap-to-voxel TRE (mm) with bending energy (Bending) and the proposed analytical equilibrium gap (Physics) regularizers, along with the state-of-the-art method for the 4DCT DIRLab dataset.

Scan	Bending	Physics	isoPTV [18]
4DCT 01	0.77 (0.90)	0.79 (0.93)	0.76 (0.90)
4DCT 02	0.69 (0.88)	0.78 (0.91)	0.76 (0.89)
4DCT 03	0.91 (1.03)	0.90 (1.05)	0.90 (1.05)
4DCT 04	1.35 (1.22)	1.33 (1.27)	1.24 (1.29)
4DCT 05	1.27 (1.49)	1.22 (1.50)	1.12 (1.44)
4DCT 06	1.12 (1.06)	1.13 (1.05)	0.85 (0.89)
4DCT 07	1.10 (0.98)	1.05 (0.98)	0.80 (1.28)
4DCT 08	1.29 (1.36)	1.24 (1.28)	1.34 (1.93)
4DCT 09	1.17 (0.99)	1.17 (1.03)	0.92 (0.94)
4DCT 10	1.13 (1.05)	1.12 (1.06)	0.82 (0.89)
Average	1.19	1.18	0.95

The effect of regularization weight β was also evaluated for the DIRLab dataset, but in terms of cumulative TRE, as in previous works [14, 15]. These results are depicted in Fig. 4.

4. DISCUSSION AND CONCLUSION

The proposed analytical equilibrium gap regularization strategy resulted in comparable or even improved registration accuracy than bending energy regularization, for our experiments with synthetic and DIRLab 4DCT datasets. Both regularization strategies performed well in the evaluated datasets, but the proposed regularization was more robust to changes

in regularization weight β . This is clearly illustrated in Figs. 2 and 4, where the registration accuracy was still acceptable with the proposed regularization strategy even for very high values of regularization weight ($\beta \in \{0.5, 0.9\}$), while being comparable with the state-of-the-art in the DIRLab 4DCT dataset for an optimal weight ($\beta = 0.001$, see Table 1). In practice, it is impossible to choose *a priori* a value for β yielding optimal accuracy, and therefore, the lesser its influence in registration accuracy, the better. We want to emphasize that the bending energy regularization, as well as other non-physically motivated regularization strategies, may, by construction, over-regularize the transformation ϕ in a large deformation setting. This can be seen in Fig. 3, where a decrease in deformation with increasing β is observed for bending energy regularization, effect that is significantly lower for our proposed physically motivated regularization. In fact, with high regularization weight, the estimated transformation approaches an affine transformation, which is not penalized by this type of regularization.

There are, however, some limitations that should be addressed in future work. A first limitation concerns our current implementation, which depends on the IDIR framework. Although we tried to limit the impact of the network hyperparameters by reutilizing those proposed in a previous study [15], these may affect registration accuracy and a sensibility study is therefore welcomed. We recall, however, that the IDIR framework was used herein only as a proxy for studying the performance of our proposed regularization strategy, and other registration approaches could be used, provided that the first and second derivatives of the displacement field are available. For instance, we plan an implementation in the parametric image registration framework Elastix in the future. A second limitation is the comparison to bending energy regularization only. Nonetheless, as mentioned above, other non-physically inspired regularization strategies will most likely fail in large deformation settings (with high regularization weight) as they inherently penalize deformation. A more informative comparison would be to the equilibrated warping approach [5, 9], since it uses the same mechanical quantity for regularization as in this work. The difference is that we propose a local (pointwise) form in (6) optimized stochastically, as opposed to a variational form weakly optimized globally in [5]. Although we do not require any discretization, it is unclear whether both strategies would yield the same solution, with the same efficiency and precision. We plan to investigate these considerations in future work.

Finally, extensions to the proposed regularization strategy are possible. For instance, other hyperelastic constitutive laws such as the Money-Rivlin or Yeoh models may be implemented similarly. Also, the tissue parameters in (6) can be tuned in for specific applications, to enforce quasi-incompressibility, for example. We are confident that physically motivated regularization can lead to more meaningful registration methods for the solution of a variety of problems.

5. COMPLIANCE WITH ETHICAL STANDARDS

This research study uses human data from the 10-years old DIRLab 4DCT dataset made available in open access [12]. No ethical statement is provided with the data.

6. REFERENCES

- [1] Azira Khalil, Siew-Cheok Ng, Yih Miin Liew, and Khin Wee Lai, “An overview on image registration techniques for cardiac diagnosis and treatment,” *Cardiology Research and Practice*, vol. 2018, pp. 1–15, Aug. 2018.
- [2] Pablo Alvarez, Simon Rouz e, Michael I. Miga, Yohan Payan, Jean-Louis Dillenseger, and Matthieu Chabanas, “A hybrid, image-based and biomechanics-based registration approach to markerless intraoperative nodule localization during video-assisted thoracoscopic surgery,” *Medical Image Analysis*, vol. 69, pp. 101983, Apr. 2021.
- [3] A. Sotiras, C. Davatzikos, and N. Paragios, “Deformable medical image registration: A survey,” *IEEE Transactions on Medical Imaging*, vol. 32, no. 7, pp. 1153–1190, July 2013.
- [4] Martin Burger, Jan Modersitzki, and Lars Ruthotto, “A hyperelastic regularization energy for image registration,” *SIAM Journal on Scientific Computing*, vol. 35, no. 1, pp. B132–B148, Jan. 2013.
- [5] M. Genet, C. T. Stoeck, C. von Deuster, L. C. Lee, and S. Kozerke, “Equilibrated warping: Finite element image registration with finite strain equilibrium gap regularization,” *Medical Image Analysis*, vol. 50, pp. 1–22, Dec. 2018.
- [6] Bernd Fischer and Jan Modersitzki, “Fast diffusion registration,” in *Inverse problems, image analysis, and medical imaging*, pp. 117–127. American Mathematical Society, 2002.
- [7] Bernd Fischer and Jan Modersitzki, “Curvature based image registration,” *Journal of Mathematical Imaging and Vision*, vol. 18, no. 1, pp. 81–85, 2003.
- [8] 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 Transactions on Medical Imaging*, vol. 18, no. 8, pp. 712–721, 1999.
- [9] Lik Chuan Lee and Martin Genet, “Validation of equilibrated warping—image registration with mechanical regularization—on 3d ultrasound images,” in *Functional Imaging and Modeling of the Heart*, pp. 334–341. Springer International Publishing, 2019.
- [10] Daniel Rueckert, Paul Aljabar, Rolf A. Heckemann, Joseph V. Hajnal, and Alexander Hammers, “Diffeomorphic registration using b-splines,” in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2006*, pp. 702–709. Springer Berlin Heidelberg, 2006.
- [11] Anders Logg, Kent-Andre Mardal, and Garth Wells, *Automated Solution of Differential Equations by the Finite Element Method*, Springer Berlin Heidelberg, 2012.
- [12] Richard Castillo, Edward Castillo, David Fuentes, Moiz Ahmad, Abbie M. Wood, Michelle S. Ludwig, and Thomas Guerrero, “A reference dataset for deformable image registration spatial accuracy evaluation using the COPDgene study archive,” *Phys. Med. Biol.*, vol. 58, no. 9, pp. 2861–2877, May 2013.
- [13] Johannes Hofmanninger, Florian Prayer, Jeanny Pan, Sebastian R ohrich, Helmut Prosch, and Georg Langs, “Automatic lung segmentation in routine imaging is primarily a data diversity problem, not a methodology problem,” *European Radiology Experimental*, vol. 4, no. 1, aug 2020.
- [14] Jan Ruhaak, Thomas Polzin, Stefan Heldmann, Ivor J. A. Simpson, Heinz Handels, Jan Modersitzki, and Mattias P. Heinrich, “Estimation of Large Motion in Lung CT by Integrating Regularized Keypoint Correspondences into Dense Deformable Registration,” *IEEE Trans. Med. Imaging*, vol. 36, no. 8, pp. 1746–1757, Aug. 2017.
- [15] Jelmer M. Wolterink, Jesse C. Zwienenberg, and Christoph Brune, “Implicit neural representations for deformable image registration,” in *Medical Imaging with Deep Learning*, 2022.
- [16] Louis D. Van Harten, Jaap Stoker, and Ivana Iřgum, “Robust deformable image registration using cycle-consistent implicit representations,” *IEEE Transactions on Medical Imaging*, pp. 1–1, 2023.
- [17] Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein, “Implicit neural representations with periodic activation functions,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, Eds. 2020, vol. 33, pp. 7462–7473, Curran Associates, Inc.
- [18] Valery Vishnevskiy, Tobias Gass, Gabor Szekely, Christine Tanner, and Orcun Goksel, “Isotropic total variation regularization of displacements in parametric image registration,” *IEEE Transactions on Medical Imaging*, vol. 36, no. 2, pp. 385–395, Feb. 2017.