Data-Driven Tissue- and Subject-Specific Elastic Regularization for Medical Image Registration

Anna Reithmeir[1,2,3], Lina Felsner[1,3], Rickmer Braren[5,6], Julia A. Schnabel[1,2,3,4], and Veronika A. Zimmer[1,3,5]

[1] School of Computation, Information and Technology, Technical University of Munich, Germany
[2] Munich Center for Machine Learning, Germany
[3] Institute of Machine Learning in Biomedical Imaging, Helmholtz Munich, Germany
[4] School of Biomedical Engineering and Imaging Sciences, King's College London, UK
[5] School of Medicine & Health, Klinkum Rechts der Isar, Technical University of Munich, Germany
[6] German Cancer Consortium (DKTK), Partner Site Munich, Germany
`anna.reithmeir@tum.de`

Abstract. Physics-inspired regularization is desired for intra-patient image registration since it can effectively capture the biomechanical characteristics of anatomical structures. However, a major challenge lies in the reliance on physical parameters: Parameter estimations vary widely across the literature, and the physical properties themselves are inherently subject-specific. In this work, we introduce a novel data-driven method that leverages hypernetworks to learn the tissue-dependent elasticity parameters of an elastic regularizer. Notably, our approach facilitates the estimation of patient-specific parameters without the need to retrain the network. We evaluate our method on three publicly available 2D and 3D lung CT and cardiac MR datasets. We find that with our proposed subject-specific tissue-dependent regularization, a higher registration quality is achieved across all datasets compared to using a global regularizer. The code is available at https://github.com/compai-lab/2024-miccai-reithmeir.

Keywords: Spatially Adaptive Regularization · Hypernetworks.

1 Introduction

Accurate and reliable intra-patient registration of medical images is ubiquitous for image-based diagnostics and treatment, for example, in radiotherapy and disease monitoring [19]. While most state-of-the-art deep learning-based registration algorithms resort to generic and global regularization methods (regularizers) [2, 9, 14], physics-inspired regularizers, such as the linear elastic regularizer [4], are particularly well-suited for intra-patient registration. They model biomechanical principles and thus ensure the physical plausibility of the registration. First, the selection of physical parameter values is ambiguous. Prior
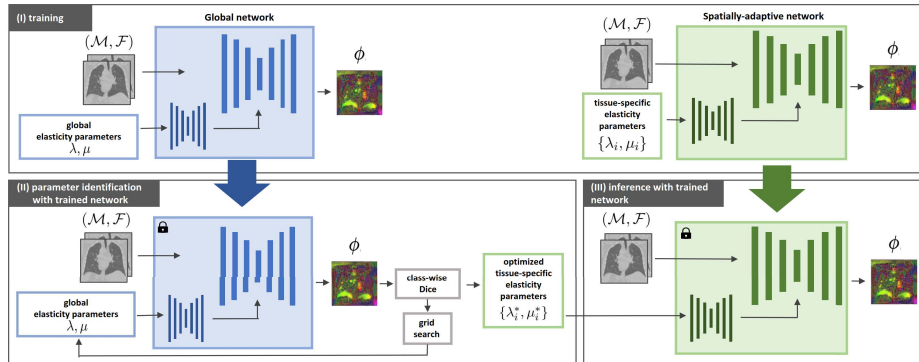
Fig. 1: Overview of our method. We train a globally and spatially adaptive network. After training, the optimal tissue-specific elasticity parameters are estimated with the global network. The spatially adaptive network then predicts the deformation field for the registration of a new image pair that follows the physical properties specified by the parameter values.

works often resort to estimations found across the literature [13,15] even though they can differ significantly not only in their absolute values but also in their ratio [8,20]. Second, different anatomical structures have different physical properties. For instance, bones are nearly rigid, while soft tissue is more elastic. A global regularizer can not accommodate such varying properties, and thus, a spatially adaptive regularizer that can vary within the image is required. However, spatial variability further complicates the choice of physical parameter values, and precise knowledge about anatomical structure locations is necessary. Third, physical properties can vary among individuals. On the one hand, age is a factor as, e.g., with increasing age, the elasticity of bones and arteries tends to decrease [5,23]. On the other hand, the presence of pathologies can impact tissue properties. For example, increased stiffness is observed in fibrotic liver and lung tissue [7,20] as well as in cardiac tissue of patients with obesity and atherosclerosis [5]. These challenges hinder the use of a one-suits-all set of physical parameter values for the physics-inspired regularization of image registration.

Recently, works have been introduced that learn the elastic regularization in the context of image registration [1,17]. Arratia López *et. al.* [1] adopt a hyperelastic regularizer to promote incompressibility in the registration process. While the regularizer varies between the tissue in focus and the background, it does not distinguish between different tissue types. Since this method is based on a physics-informed neural network (PINN), retraining is necessary for every image pair. Reithmeir *et. al.* [17] use a hypernetwork to learn the elasticity parameters of the linear elastic regularizer for image registration. A global regularizer is applied and tissue-specific properties are not considered. Neither of the two methods above learns patient-specific regularization parameters. The recent work of Ragoza *et. al.* [16] focuses on the elasticity parameter estimation

from medical images. They propose a PINN for spatially varying tissue elasticity reconstruction from images [16]. However, it requires a specific medical imaging modality (MR elastography) and cannot be applied to standard MR and CT images. Here, image registration is not considered.

In this work, we build upon our previous work in [17] and introduce a novel framework that learns not only tissue- but also subject-specific parameters for the elastic regularization of image registration. To this end, we employ two hypernetworks: One for the parameter estimation and one for the registration with the spatially varying elastic regularization (Fig. 1). Our main contributions are twofold:

- We propose a data-driven spatially adaptive elastic regularization that varies the elasticity properties for different anatomical structures.
- To the best of our knowledge, we are the first to introduce a registration algorithm that *learns* the *subject-specific* linear elastic regularization where the elasticity parameters of the linear elastic regularizer are estimated for each image pair individually without the need to retrain the network.

2   Methods

We propose a novel method that learns the spatially adaptive linear elastic regularization from the data. It comprises two independent neural networks that differ in their regularization scheme and which we will refer to as *global* and *spatially adaptive* network in the following (Fig. 1 (i)). The global network employs a global linear elastic regularizer and is used to extract the tissue-dependent elasticity parameters. The spatially adaptive network uses a spatially varying regularizer and is used for the regularized registration.

For the estimation of optimal tissue- and patient-specific elasticity parameters, we exploit the ability of hypernetworks [6] to adapt to hyperparameter values at test time. Hypernetworks are small networks that can learn the effect of hyperparameters on the output of a larger network. In this work, we apply two hypernetworks that learn the effect of the elasticity parameters of the linear elastic regularizer [4] on the registration. This allows us to explore the regularization parameter space for new data at test time without the need to retrain the networks.

Our method consists of three stages (see Fig. 1), which are described in detail below: (i) Training the two networks, (ii) estimation of the optimal subject- and tissue-specific elasticity parameters, and (iii) registration of new data with the identified parameters.

2.1   Training of Global and Spatially Adaptive Networks

We employ two networks to learn the globally and spatially adaptive regularized registration, respectively. Each consists of a U-Net [18] registration network and a hypernetwork [6] and predicts the deformation $\phi : \mathbb{R}^D \rightarrow \mathbb{R}^D$ between a

moving image $M : \Omega \subset \mathbb{R}^D \to \mathbb{R}$ and a fixed image $F : \Omega \subset \mathbb{R}^D \to \mathbb{R}$. We parameterize ϕ as the dense displacement field $\phi(x) = x + \mathbf{u}(x)$ for each pixel location x.

The inputs to both models are the image pair (M, F) and the Lamé parameters of the linear elastic regularizer [4] that control the elasticity properties that are modeled.

Global Model: This network is similar to [17] and based on [11]. It applies a global linear elastic regularizer, i.e., it enforces the same elasticity properties everywhere in the image domain Ω. The two global Lamé parameters $\lambda, \mu \in \mathbb{R}$ are given as inputs to the hypernetwork. The network is trained unsupervised, using the loss function in [17], with the change of constraining $\lambda, \mu \in [0, 1]$.

Spatially adaptive Model: This network is an extension of [17] with a spatially varying linear elastic regularizer. It adapts to different elasticity properties based on the underlying tissue type at each $x \in \Omega$. This is achieved by tissue-specific Lamé parameters $\{(\lambda_i, \mu_i)| \; \forall i \in 1, ..., C\}$ with C as the number of considered tissue classes.

Two elasticity parameter maps $\mathbf{\Lambda}, \mathbf{\Gamma} : \Omega \in \mathbb{R}^D \to \mathbb{R}$ are formed with the help of the segmentation map $\mathbf{S} : \Omega \subset \mathbb{R}^D \to \mathbb{Z}$ of the fixed image F. They hold the respective tissue-specific parameters at each $x \in \Omega$, i.e., $\mathbf{\Lambda}(x) = \lambda_c$ and $\mathbf{\Gamma}(x) = \mu_c$ for $\mathbf{S}(x) = c$. The training loss function of the network is

$$\mathcal{L}(M, F, \phi, \mathbf{\Lambda}, \mathbf{\Gamma}) = (2 * \mathbb{1} - \mathbf{\Lambda} - \mathbf{\Gamma}) \; NCC(F, \phi \circ M)$$
$$+ \int_\Omega \frac{1}{4} \mathbf{\Gamma} \sum_{i,j=1}^{D} (\partial_{x_i} u_j + \partial_{x_j} u_i)^2 + \frac{1}{2} \mathbf{\Lambda} (\mathrm{div} \; \mathbf{u})^2 \; d\mathbf{x} \; , \quad (1)$$

where $\mathbb{1} : \Omega \subset \mathbb{R}^D \to \mathbb{R}$ is a constant weighting map of ones, NCC is the local normalized cross-correlation similarity metric, $\partial_x u$ are the partial derivatives of $\mathbf{u}$ and div is the divergence. To stay in line with [11, 17], we constrain the value range of the parameters $\mu_i, \lambda_i \forall i \in \Omega$ to the range $[0, 1]$ and weigh the similarity term. In the spatially-adaptive case, this weighting is voxel-based and thus dependent on the weight maps $\mathbf{\Gamma}, \mathbf{\Lambda}$.

For simplicity, we omit an additional regularization weight which controls the amount of regularization and absorb it in the elasticity parameters as in [17]. During training, the elasticity parameters are individually sampled from the uniform distribution $\mathcal{U}(0, 1)$ in each training step. The two networks are trained separately on the same training data.

2.2   Elasticity Parameter Identification

Hypernetworks allow the efficient optimization of hyperparameters at test time. We exploit this advantage and use the trained global model to estimate the optimal elasticity parameters of the regularizer. Notably, we do this not only for

each tissue type but also for each image pair individually. Thus, we obtain tissue-and subject-specific elasticity parameters, which, to the best of our knowledge, is not performed in prior works. Similar to [17], we identify the optimal parameters by a grid search over the parameter space at test time. Without increasing the computational complexity: (i) we perform a sample-based grid search to identify the optimal parameters for each subject individually, and (ii) we base the grid search on a class-wise heuristic to find the optimal parameters for each tissue type. In this work, we use the class-wise Dice score. Overall, for each test sample n, the grid search results in the parameter set $\{(\lambda_i^n, \mu_i^n)| \ \forall i \in 1, ..., C\}$. For N test samples, in total $2CN$ parameters are estimated.

2.3   Registration with Spatially Adaptive Elastic Regularization

For the registration of an unseen test sample, the optimal Lamé parameters that have been identified for this specific sample with the global model are given as input to the hypernetwork of the spatially adaptive model. A single forward pass through the trained spatially adaptive network then registers the given image pair. Since this network is trained with a spatially varying regularizer, the predicted deformation field is regularized in the different anatomical regions according to the specified local elasticity properties.

3   Experiments and Results

3.1   Implementation Details and Datasets

We implement our method with TensorFlow 2.12.1 and use the default network architecture of HyperMorph [11]. We use the Adam optimizer [12] and a learning rate of 1e-4. All networks are trained until convergence using an NVIDIA RTXA6000 GPU. The linear elastic regularizer is implemented with forward finite differences. The code will be made publicly available upon acceptance.

   We evaluate our method on three intra-patient datasets in 2D and 3D; (I) *NLST*: The lung CT data of the National Lung Screening Trial dataset provided by the Learn2Reg challenge [9, 21]. We use the middle coronal slice of each of the 209 annotated images and a train/val/test split of 169/10/30. We clip the intensities to (-980, 1518) HU. (II) *ACDC*: The 2D cardiac MRI dataset is taken from the Automatic Cardiac Diagnosis Challenge (ACDC) [3]. We use the middle slices of the end-systolic and end-diastolic frames, a split of 50/20/30 and we crop the images to size (128,128). All images are resampled to isotropic pixel size, and their intensities are normalized. (III) *L2R-Lung*: The 3D lung CT dataset is taken from the Learn2Reg challenge [10] and contains 30 inhale/exhale image pairs. The fixed images are only partly visible. We use a split of 20/4/4, clip the intensities to (-980, 600) HU and padded the images to size (192,128,192).

   In addition to the already provided lung segmentation masks in the NLST and L2R-Lung datasets, we generate liver and bone (vertebrae, ribs, and sternum) masks with the TotalSegmentator tool [22]. The ACDC dataset provides segmentations for the right ventricle, myocardium, and left ventricle blood pool. Thus, for all three datasets $C = 3$.
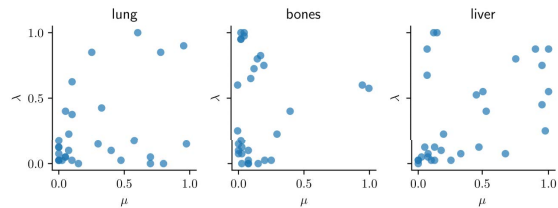
Fig. 2: The subject-specific tissue-wise regularization parameters λ, μ that have been estimated with the global network for the NLST dataset. Each blue dot corresponds to one test sample. The results show that the optimal parameters vary between tissue types and subjects.

3.2   Analysis of Subject-Specific Elasticity Parameters

In the first experiment, we analyze the sample-specific elasticity parameters that are estimated with the trained global network as described in Subsection 2.2. We use the class-wise Dice score as the grid search optimization heuristic and extract the optimal Lamé parameters for each tissue type and subject individually. The results for the NLST dataset are shown in Fig. 2. It can be seen that the values for both parameters (λ, μ) vary between the subjects. Within each tissue class, a tendency towards low values for μ and λ can be observed. However, for each tissue type, multiple samples can be found for which high values are optimal. Furthermore, the value distribution differs between the tissue types. For example, the bone-specific μ generally favors values that are smaller than 0.5, while the liver-specific values stretch out over the full parameter space. Similar results are obtained for the other two datasets (see Appendix A).

3.3   Comparison of Regularization Methods

In the second experiment, we evaluate our proposed method. To compare the different regularization schemes to each other, we use them within the same registration framework, namely [11]. we compare our subject-specific spatially varying regularization (spatially adaptive models) to the global regularization (global models) of [11,17]. The global models are evaluated with the estimated optimal global regularization parameters according to the mean Dice score. The spatially adaptive models are evaluated with the estimated optimal tissue-specific regularization parameters according to the class-wise Dice scores. Furthermore, for better comparison, we estimate tissue- and subject-specific regularization weights in the same manner for the diffusion regularizer of the HyperMorph model [11]. Here, the regularization weights specify the amount of smoothness per tissue type. As evaluation metrics, we use the mean and class-wise Dice scores (DSC) and the 95th percentile Hausdorff distance (HD95) between the fixed and predicted warped segmentations, the target registration error (TRE, Euclidean distance) between the fixed and predicted warped keypoints and the fraction of negative Jacobian determinants of the displacement field ($\%negJ$).

Table 1: Quantitative results for the linear elastic (*elas*) and diffusion (*diff*) regularized models. The models with global *(g)* and spatially adaptive and subject-specific *(sas)* regularization parameters are compared. The results are shown as mean(standard deviation). For the class DSC the classes correspond to [lung, bone, liver] for the NLST/L2R-Lung and [right ventricle, myocardium, left ventricle blood pool] for the ACDC dataset. The results show that with the tissue- and subject-specific regularization, higher registration quality is achieved. For metric descriptions see Sec. 3.3

2D	Model	Class DSC↑	Mean DSC↑	HD95↓	%negJ↓
NLST	initial	[0.914, 0.531, 0.775]	0.749(0.064)	11.755(10.328)	-
	diff-g [11]	[**0.972**, 0.799, **0.855**]	0.885(0.067)	8.227(10.293)	**0.022**(0.007)
	diff-sas (ours)	[**0.972**, **0.801**, **0.855**]	**0.886**(0.065)	**8.157**(9.998)	0.038(0.028)
	elas-g [17]	[0.969, 0.784, **0.855**]	0.879(0.068)	9.211(10.636)	**0.009**(0.004)
	elas-sas (ours)	[**0.970**, **0.796**, 0.852]	**0.883**(0.067)	**8.598**(10.621)	0.040(0.025)
ACDC	initial	[0.748, 0.520, 0.667]	0.645(0.119)	7.714(3.162)	-
	diff-g [11]	[0.897, 0.757, **0.888**]	0.847(0.052)	2.696(1.776)	**0.006**(0.005)
	diff-sas (ours)	[**0.903**, **0.771**, **0.888**]	**0.854**(0.056)	**2.659**(1.967)	0.011(0.008)
	elas-g [17]	[0.886, 0.735, 0.867]	0.829(0.066)	**3.139**(1.682)	**0.002**(0.002)
	elas-sas (ours)	[**0.888**, **0.743**, **0.869**]	**0.833**(0.068)	3.175(2.010)	0.009(0.004)

3D	Model	Class DSC↑	Mean DSC↑	TRE↓/HD95↓	%negJ↓
L2R-Lung	initial	[0.911, 0.494, 0.758]	0.721(0.061)	3.227(0.734)/16.092(7.328)	-
	diff-g [11]	[0.959, 0.693, 0.829]	0.827(0.023)	3.041(0.551)/12.865(7.452)	0.030(0.014)
	diff-sas (ours)	[**0.960**, **0.695**, **0.839**]	**0.832**(0.022)	**2.568**(0.532)/**12.481**(7.579)	**0.027**(0.010)
	elas-g [17]	[**0.961**, 0.678, 0.822]	0.820(0.022)	2.645(0.484)/13.330(7.942)	0.013(0.005)
	elas-sas (ours)	[0.959,**0.684**,**0.842**]	**0.828**(0.024)	**2.638**(0.512)/**12.499**(6.786)	**0.006**(0.002)

The quantitative results are shown in Tab. 1. The results for the diffusion regularizer, in general, outperform the results for the linear elastic regularizer. Similar observations are made in [17]. In terms of the class-wise DSC, all spatially adaptive models outperform the global models in at least two out of three tissue classes. For the ACDC dataset and the elastic regularizer, all three class-wise DSCs show higher performance with the spatially adaptive regularization, even though this is the only case where a lower HD95 is obtained. Furthermore, the mean DSC is consistently higher for the spatially adaptive models. Interestingly, for the 2D datasets, the amount of folding, i.e., the fraction of negative Jacobian determinants, increases when spatially adaptive regularization is applied, while it decreases in the 3D case. Visual inspection of the results reveals that for some test samples, a high amount of folding is observed inside the lung and background (compare second-to-top row of Fig.3). Note, however, that nearly no folding is observed here in the liver and bone structures. The qualitative results in Fig. 3 furthermore confirm the superior registration quality for the proposed method: The breathing motion of the lower lung lobes (NLST dataset, second-to-top row) and the omnidirectional expansion of the two heart ventricles (ACDC dataset, bottom row) is better reflected in the predicted deformations as can be seen in the deformation plots and difference images. Further qualitative examples can be found in Appendix B.
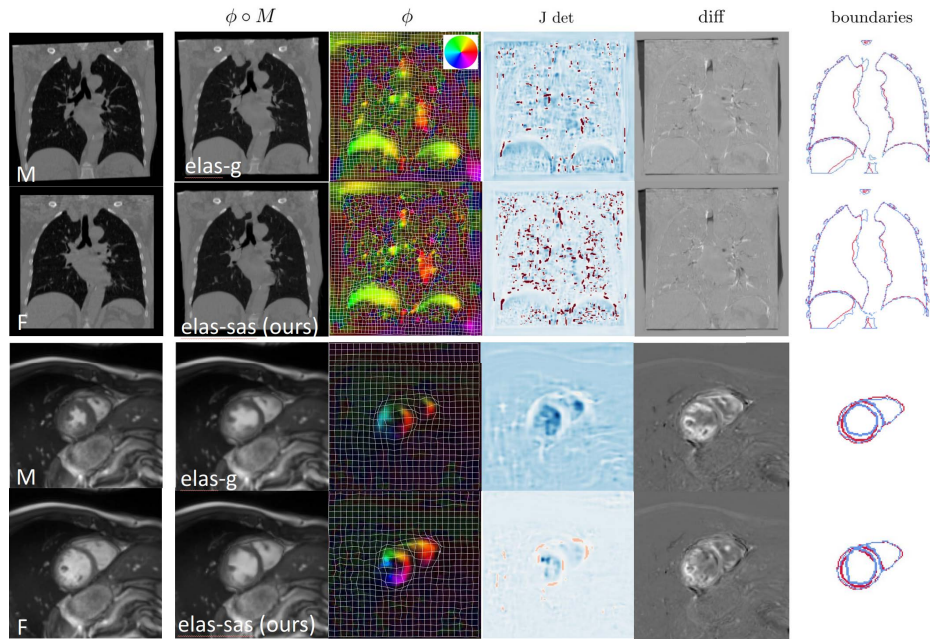
Fig. 3: Qualitative results for the NLST (top) and ACDC (bottom) datasets. The left column shows the moving (M) and fixed (F) images. From left to right: Warped moving image, predicted deformation, Jacobian determinant (blue: positive values, red: negative values), difference image between fixed and warped image, segmentation map boundaries. The results show that the subject-specific spatially adaptive regularization (elas-sas) leads to more plausible deformation than the global regularization (elas-g).

4   Discussion and Conclusion

In this work, we proposed a novel method that estimates tissue- and subject-specific parameters for the physics-inspired regularization of data-driven image registration. We focused on the parameters of the linear elastic regularizer and evaluated the method on three diverse datasets.

Our analysis of the subject-specific elasticity parameters that have been estimated with the global model in Sec.3.2 confirms that the parameters can differ not only between tissue types but also between individual subjects. This highlights the importance of the adaptation of physics-inspired regularizers to the local anatomical structures and to the individual patients. The comparison of the different regularization methods in Sec.3.3 shows that physiologically more meaningful deformations are obtained with spatially adaptive, subject-specific linear elastic regularization. While our method results in more folding, we observe consistently increased class Dice scores for the bones in Tab.1, which indicates

that small anatomical structures are better registered with our method than with a globally regularized model.

A current limitation of this work is that the estimated elasticity parameters cannot be interpreted in a physically meaningful manner. The parameter μ is the shear modulus and serves as a measure of stiffness: Rigid materials have higher values for μ than more elastic materials. This, however, is not reflected in our estimations since, e.g., the bone-specific μ do not favor larger values compared to the liver- or lung-specific μ. A possible reason for this could be the uniform sampling space of the parameters in our method, and we plan to investigate this further in the future. Also, we aim to extend the evaluation of our method to additional datasets and more tissue types, as well as to explore gradient-based elasticity parameter optimization approaches.

In conclusion, we proposed a novel image registration method that learns tissue- and sample-specific elasticity parameters for physics-inspired regularization. We believe that the patient-specific adaptation of the regularization is fundamental for successful intra-patient registration and that it can further improve data-driven registration methods in the future.

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

References

1. Arratia López, P., Mella, H., Uribe, S., Hurtado, D.E., Sahli Costabal, F.: Warppinn: Cine-mr image registration with physics-informed neural networks. Medical Image Analysis **89**, 102925 (2023)
2. Balakrishnan, G., Zhao, A., Sabuncu, M., Guttag, J., Dalca, A.: Voxelmorph: A learning framework for deformable medical image registration. IEEE Transactions on Medical Imaging **38**, 1788–1800 (2019)
3. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Ballester, M.A.G., et al.: Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: is the problem solved? IEEE transactions on medical imaging **37**(11), 2514–2525 (2018)
4. Broit, C.: Optimal Registration of Deformed Images. Graduate School of Arts and Sciences, University of Pennsylvania (1981)
5. Cocciolone, A.J., Hawes, J.Z., Staiculescu, M.C., Johnson, E.O., Murshed, M., Wagenseil, J.E.: Elastin, arterial mechanics, and cardiovascular disease. American Journal of Physiology-Heart and Circulatory Physiology **315**(2), H189–H205 (2018)
6. Ha, D., Dai, A.M., Le, Q.V.: Hypernetworks. arXiv preprint arXiv:1609.09106 (2017)
7. Haak, A.J., Tan, Q., Tschumperlin, D.J.: Matrix biomechanics and dynamics in pulmonary fibrosis. Matrix Biology **73**, 64–76 (2018)

8. Hagemann, A., Rohr, K., Stiehl, H., Spetzger, U., Gilsbach, J.: Biomechanical modeling of the human head for physically based, nonrigid image registration. IEEE Transactions on Medical Imaging **18**(10), 875–884 (1999)

9. Hering, A., Hansen, L., Mok, T.C.W., et al.: Learn2reg: Comprehensive multi-task medical image registration challenge, dataset and evaluation in the era of deep learning. IEEE Transactions on Medical Imaging **42**(3), 697–712 (2023)

10. Hering, A., Murphy, K., van Ginneken, B.: Learn2reg challenge: Ct lung registration - training data [data set]. zenodo. http://doi.org/10.5281/zenodo.3835682 (2020)

11. Hoopes, A., Hoffmann, M., Fischl, B., Guttag, J., Dalca, A.V.: Hypermorph: Amortized hyperparameter learning for image registration. Information Processing in Medical Imaging (IPMI), vol. 12729, p. 3–17 (2021)

12. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)

13. Min, Z., Baum, Z.M.C., Saeed, S.U., Emberton, M., Barratt, D.C., Taylor, Z.A., Hu, Y.: Non-rigid medical image registration using physics-informed neural networks. In: Information Processing in Medical Imaging (IPMI) (2023)

14. Mok, T.C.W., Chung, A.C.S.: Large deformation diffeomorphic image registration with laplacian pyramid networks. Medical Image Computing and Computer Assisted Intervention (MICCAI), vol. 12263, pp. 211–221 (2020)

15. Qin, C., Wang, S., Chen, C., Bai, W., Rueckert, D.: Generative myocardial motion tracking via latent space exploration with biomechanics-informed prior. Medical Image Analysis **83**, 102682 (2023)

16. Ragoza, M., Batmanghelich, K.: Physics-informed neural networks for tissue elasticity reconstruction in magnetic resonance elastography. In: Medical Image Computing and Computer Assisted Intervention (MICCAI). pp. 333–343 (2023)

17. Reithmeir, A., Schnabel, J.A., Zimmer, V.A.: Learning physics-inspired regularization for medical image registration with hypernetworks. In: Medical Imaging 2024: Image Processing. International Society for Optics and Photonics, SPIE (2024 (accepted))

18. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 234–241 (2015)

19. Rueckert, D., Schnabel, J.A.: Medical Image Registration, pp. 131–154. Springer, Berlin, Heidelberg (2011)

20. Talwalkar, J.A.: Elastography for detecting hepatic fibrosis: options and considerations. Gastroenterology **135**(1), 299–302 (2008)

21. Team, N.L.S.T.R.: Data from the national lung screening trial (nlst) [data set]. the cancer imaging archive. (2013)

22. Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., Heye, T., Boll, D.T., Cyriac, J., Yang, S., Bach, M., Segeroth, M.: Totalsegmentator: Robust segmentation of 104 anatomic structures in ct images. Radiology: Artificial Intelligence **5**(5) (2023)

23. Zioupos, P.: Ageing human bone: Factors affecting its biomechanical properties and the role of collagen. Journal of Biomaterials Applications **15**(3), 187–229 (2001)