




Fast 3D Registration with Accurate Optimisation and Little Learning for Learn2Reg 2021

Hanna Siebert^(✉) , Lasse Hansen , and Mattias P. Heinrich 

Institute of Medical Informatics, Universität zu Lübeck, Lübeck, Germany
{siebert,hansen,heinrich}@imi.uni-luebeck.de

Abstract. Current approaches for deformable medical image registration often struggle to fulfill all of the following criteria: versatile applicability, small computation or training times, and the being able to estimate large deformations. Furthermore, end-to-end networks for supervised training of registration often become overly complex and difficult to train. For the Learn2Reg2021 challenge, we aim to address these issues by decoupling feature learning and geometric alignment. First, we introduce a new very fast and accurate optimisation method. By using discretised displacements and a coupled convex optimisation procedure, we are able to robustly cope with large deformations. With the help of an Adam-based instance optimisation, we achieve very accurate registration performances and by using regularisation, we obtain smooth and plausible deformation fields. Second, to be versatile for different registration tasks, we extract hand-crafted features that are modality and contrast invariant and complement them with semantic features from a task-specific segmentation U-Net. With our results we were able to achieve the overall Learn2Reg2021 challenge's second place, winning Task 1 and being second and third in the other two tasks.

Keywords: image registration · convex optimisation · instance optimisation

1 Motivation

Deep-learning-based approaches for medical image registration usually involve an elaborate learning procedure and yet they often struggle with the estimation of large deformations and the versatile usability for a wide range of tasks. To address the different registration tasks of the Learn2Reg2021 challenge¹ [8], we present a fast and accurate optimisation method for image registration that requires little learning. Our method robustly captures large deformations by using discretised displacements and a coupled convex optimisation. In order to be versatile for various tasks, we include a hand-crafted feature extractor in our method that is contrast and modality invariant and still highly discriminative for local geometry.

¹ <https://learn2reg.grand-challenge.org>.

2 Methods

The main idea of our method is to perform large-deformation image registration by using a coupled convex optimisation [6] that approximates a globally optimal solution of a discretised cost function followed by an Adam-based instance optimisation to further improve the local registration accuracy. Dense correlation has already been used extensively in learning based optical flow estimation (cf. PWC-Net [16]) and end-to-end trainable 3D registration networks (cf. PDD-Net [5]), however both approaches have limitations. PWC-Net requires multiple warping steps and is difficult to extend from 2D to 3D (see [4]). PDD-Net employs a dense 3D displacements, but substantially simplifies the optimisation strategy, which may lead to some inaccuracies. ConvexAdam aims to combine the best of both worlds (learning and optimisation-based) by leveraging segmentation priors where available and relying on robust hand-crafted features and fast discrete optimisation.

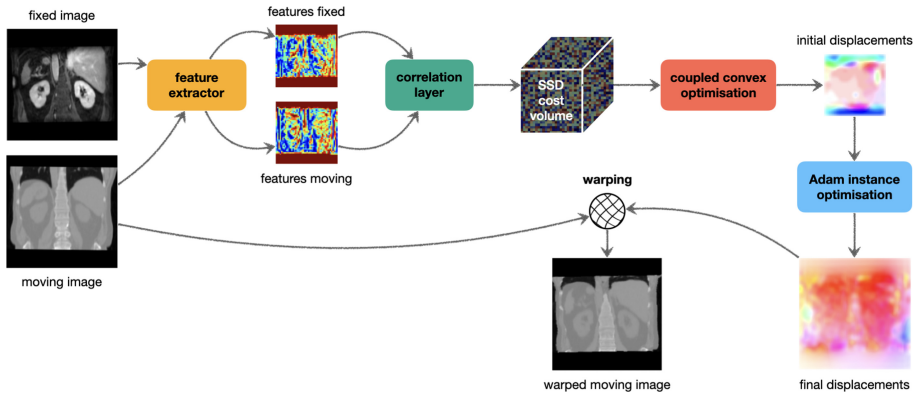


Fig. 1. The structure of our registration method. It consists of a feature extractor (MIND and/or nnUNet) and a dense correlation layer followed by a coupled convex optimisation and an Adam-based instance optimisation.

As visualised in Fig. 1, the basic structure of our registration method consists of a feature extractor, a correlation layer, a coupled convex optimisation, and an instance optimisation.

The feature extractor outputs contrast and modality invariant features from the fixed and moving input images. For this, hand-crafted MIND features [7] ensuring versatility regarding different types of registration tasks can be employed. Depending on the availability of labelled image data, automatic segmentations as provided by the nnU-Net [10] can be used instead. Different to other state-of-the-art supervised deep learning registration methods [14] we avoid using the expert labels only at the end for the warping loss, which may lead to sub-optimal results due to limited gradient backflow. We instead found that using off-the-shelf segmentation networks produce best results.

The obtained features are fed into a correlation layer, which computes a sum-of-squared-differences (SSD) cost volume with a box filter and gives an initial best displacement for each voxel (simply taking the argmin). Therefore, we employ a search space with up to 5000 discretised displacements per voxel. The capture range can be up to at least 48 voxels in each dimension (setting for Task 2) and therefore estimate large motion accurately.

The correlation layer’s output is used to solve two coupled convex optimisation problems for efficient global regularisation: In several iterations, alternating steps are performed for similarity and smoothness optimisation, i.e. a spatially smoothed field based on the current argmin (minimal SSD costs) displacements followed by adding a penalty to the discretised SSD costs based on the discrepancy of this current globally smooth optimum.

The resulting displacements in turn are used as a starting point for an Adam-based instance optimisation in order to provide the final deformation grid used for warping of the moving input image. This step is very similar to classic optical flow estimation [15]. For this purpose, the cost function is linearised and the Adam optimiser [11] is used for gradient descent. Smoothness of the displacement field is induced by adding a B-spline deformation model and diffusion regularisation.

3 Experiments and Results

Each of the Learn2Reg2021 tasks entails certain challenges that we face with slightly varying experimental setups as outlined in the following. The complete implementation details can be found in our publicly available repository.² Table 1 presents quantitative results and Fig. 2 shows qualitative results for the individual tasks.

Task 1 Thorax-Abdomen CT-MR. The first task aims to align multimodal intra-patient data [1–3, 12]. Besides of multimodal image registration, the objectives of learning from few and noisy labels, as well as dealing with large deformations and missing correspondences are challenging. For this task, we extract hand-crafted MIND features and include an inverse-consistency constraint as introduced in [6] to enforce a minimised discrepancy between the forward and backward transformations in order to avoid implausible deformations. To further regularise the displacement field during Adam instance optimisation, we add thin plate splines yielding smooth deformation fields. As large deformations are to be expected, we chose a search space that includes discretised displacements with a capture range of 64 mm for each dimension within the scanned anatomy.

Task 2 Lung CT. The second task is to perform inspiration-expiration registration on intra-patient lung CT data [9]. In this task, there is the challenge of estimating large breathing motion for scans with only partial visibility of the lungs in the expiration scans. The displacement search range is selected in order

² <https://github.com/multimodallearning/convexAdam>.

Table 1. Results for the different Learn2Reg2021 tasks. Accuracy is measured by the Dice similarity of organ segmentations (Dice), the target registration error for anatomical landmarks (TRE), and the 95% Hausdorff distance for segmentations (HD). Robustness is measured by the 30% lowest Dice scores (Dice₃₀), Dice scores for additional segmentations (Dice_{+add}) and the 30% highest TRE values (TRE₃₀). Plausibility of the deformations is measured by the standard deviation of the logarithmic Jacobian determinant (SDlogJ). Dice similarities are reported in %, TRE and HD values are given in millimetres and inference time is given in seconds. The last table displays the challenge scores and ranks for the overall 1st, 2nd, and 3rd place.

Task 1					
	Dice	Dice ₊₉	HD	SDlogJ	time
initial	33.1	22.3	44.48	-	-
ours	75.4	73.1	20.75	0.09	1.30

Task 2				
	TRE	TRE ₃₀	SDlogJ	time
initial	10.24	16.80	-	-
ours	1.85	2.89	0.06	1.82

Task 3					
	Dice	Dice ₃₀	HD	SDlogJ	time
initial	55.9	29.7	4.07	-	-
ours	79.9	64.5	2.00	0.05	12.62

Scores and ranks			
	Task1	Task2	Task3
	score (rank)	score (rank)	score (rank)
LapIRN	0.86 (2)	0.79 (4)	0.94 (1)
convexAdam	0.88 (1)	0.83 (3)	0.82 (2)
PIMed	0.85 (4)	0.68 (6)	0.70 (5)




Fig. 2. Qualitative results of our proposed method (top row: colourmap overlay of fixed and moving image (Task 1 and 2) or segmentation (Task 3); bottom row: overlay of fixed and warped moving image or segmentation).

to capture motion with up to $42 \times 30 \times 42$ mm for the x -, y -, and z -dimension respectively. Like in the first task, MIND features of both input images are used to compute the SSD cost volume.

Task 3 Whole Brain MR. The third task deals with the registration of inter-patient T1-weighted brain MRI [13]. Here, the main challenge is to precisely align small structures of variable shape. For this reason, we chose a displacement capture range of 16 mm for each dimension within the scanned brain structures. As this task comprises a large amount of labelled image data, nnU-Net predictions for segmentation guidance are employed. We use the nnU-Net predictions in the form of inverse class-weighted one-hot encodings as features for our method's optimisation steps.

4 Conclusion

Our contribution to the Learn2Reg2021 challenge showed that image registration can be performed fast and accurately using an optimisation strategy with little learning. It is highly parallelisable on a GPU and robust by using a large search space of discretised displacements. Smoothness of the deformation fields could be induced by a global convex regularisation, diffusion regularisation, and B-spline interpolation. By using an efficient Adam-based instance optimisation, our method yields very precise results and by integrating a modality-invariant feature extractor, we achieve a wide versatility. We were able to achieve the overall Learn2Reg2021 challenge's second place, winning Task 1, being second in Task 3, and being third in Task 2.

References

1. Akin, O., et al.: Radiology data from the cancer genome atlas kidney renal clear cell carcinoma [TCGA-KIRC] collection. Cancer Imaging Arch. (2016). <https://doi.org/10.7937/K9/TCIA.2016.V6PBVTDR>
2. Clark, K., et al.: The Cancer Imaging Archive (TCIA): maintaining and operating a public information repository. J. Digit. Imaging **26**(6), 1045–1057 (2013)
3. Erickson, B., et al.: Radiology data from the cancer genome atlas liver hepatocellular carcinoma [TCGA-LIHC] collection. Cancer Imaging Arch. (2016). <https://doi.org/10.7937/K9/TCIA.2016.IMMQW8UQ>
4. Gunnarsson, N., Sjölund, J., Schön, T.B.: Learning a deformable registration pyramid. In: Shusharina, N., Heinrich, M.P., Huang, R. (eds.) MICCAI 2020. LNCS, vol. 12587, pp. 80–86. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-71827-5_10
5. Heinrich, M.P.: Closing the gap between deep and conventional image registration using probabilistic dense displacement networks. In: Shen, D., et al. (eds.) MICCAI 2019. LNCS, vol. 11769, pp. 50–58. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32226-7_6

6. Heinrich, M.P., Papież, B.W., Schnabel, J.A., Handels, H.: Non-parametric discrete registration with convex optimisation. In: Ourselin, S., Modat, M. (eds.) WBIR 2014. LNCS, vol. 8545, pp. 51–61. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-08554-8_6
7. Heinrich, M.P., Jenkinson, M., Papież, B.W., Brady, S.M., Schnabel, J.A.: Towards realtime multimodal fusion for image-guided interventions using self-similarities. In: Mori, K., Sakuma, I., Sato, Y., Barillot, C., Navab, N. (eds.) MICCAI 2013. LNCS, vol. 8149, pp. 187–194. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40811-3_24
8. Hering, A., et al.: Learn2Reg: comprehensive multi-task medical image registration challenge, dataset and evaluation in the era of deep learning (2021)
9. Hering, A., Murphy, K., van Ginneken, B.: Learn2Reg challenge: CT lung registration - training data, May 2020. <https://doi.org/10.5281/zenodo.3835682>
10. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nat. Meth.* **18**(2), 203–211 (2021)
11. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
12. Linehan, M., et al.: Radiology data from the cancer genome atlas cervical kidney renal papillary cell carcinoma [KIRP] collection. *Cancer Imaging Arch.* (2016). <https://doi.org/10.7937/K9/TCIA.2016.ACWOGBEF>
13. Marcus, D.S., Wang, T.H., Parker, J., Csernansky, J.G., Morris, J.C., Buckner, R.L.: Open Access Series of Imaging Studies (OASIS): cross-sectional MRI data in young, middle aged, nondemented, and demented older adults. *J. Cogn. Neurosci.* **19**(9), 1498–1507 (2007)
14. Mok, T.C.W., Chung, A.C.S.: Large deformation image registration with anatomy-aware Laplacian pyramid networks. In: Shusharina, N., Heinrich, M.P., Huang, R. (eds.) MICCAI 2020. LNCS, vol. 12587, pp. 61–67. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-71827-5_7
15. Papenberg, N., Bruhn, A., Brox, T., Didas, S., Weickert, J.: Highly accurate optic flow computation with theoretically justified warping. *Int. J. Comput. Vis.* **67**(2), 141–158 (2006)
16. Sun, D., Yang, X., Liu, M.Y., Kautz, J.: PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 8934–8943 (2018)