Learning iterative optimisation for deformable image registration with recurrent convolutional networks

No Author Given

No Institute Given

Abstract. We propose a fully deep learning-based approach for deformable image registration, that aims to emulate the structure of gradient-based optimisation as used in conventional registration and thus learns how to optimise. Our architecture consists of recurrent updates on a convolutional network with deep supervision and uses dynamic sampling of the cost function and hidden states to mimic gradient-based optimization, requiring fewer iterations than traditional techniques. Without pre-registration, our method achieves 2 mm TRE on the DIR-Lab COPD dataset and outperforms Adam optimisation. Our code is publicly available at: https://anonymous.4open.science/r/Learn2Optimise-BF1F/.

Keywords: Optimisation · Deep learning · Recurrent network.

1 Introduction

Deep learning-based image registration has become increasingly more prevalent in recent years and has been employed for a multitude of different medical registration tasks. One of those tasks is deformable intra-patient lung registration. However, as the results of the Learn2Reg 2021 challenge[4] have shown, methods that rely solely on deep learning do not yet achieve state-of-the-art performance regarding inter-patient lung registration. The best four submissions were either completely optimisation-based or employed instance optimisation after application of a deep learning model. Deep learning-based submissions aimed to predict displacements directly from fixed and moving image features through feedforwand networks.

We propose a novel recurrent network architecture mimicing Adam optimisation [5]. We employ iterative dynamic cost sampling to extract a minimum of required information to calculate each optimisation step and a hidden state based information flow to allow the network to calculate gradient momentum.

2 Methods

Our architecture is based on a small U-Net architecture [8], which is recurrently applied for eight iterations. A total of 45 inputs are fed into the network each

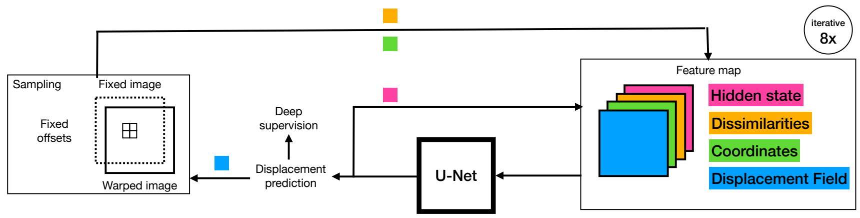

Fig. 1. Schematic depiction of our *learn to optimise* network.

iteration. They consist of the current predicted displacements, displacement coordinates for a fixed grid of subpixel offsets, dissimilarity costs for each sampled displacement and hidden states. Dissimilarity costs are calculated using MIND features. Sampled coordinates and dissimilarities can be used to approximate gradients for gradient descent. We further employ deep keypoint supervision after each iteration. A schematic representation of our method is depicted in figure 1

3 Experiments and Results

We train our network on the pulic EMPIRE10 (selected cases), DIR-Lab COPD and DIR-Lab 4DCT datasets[2,7] and evaluate the registration accuracy of our method on the COPD and 4DCT datasets. Evaluation is done in two manners. We employ our method directly on the non pre-registered image pair as well as on pairs that are pre-registered with Voxelmorph++ (VM++) [3], a variant of the Voxelmorph [1] architecture. We compare our approach with the winner of the Learn2Reg challenge LapIRN[6], evaluated with and without Adam instance optmisation. Quantitative results are shown in Table 1. Our method outperforms Adam when not using any pre-registration and achieves comparable results with pre-registration. For the 4DCT dataset we gain comparable results with LapIRN and a better TRE for the COPD dataset.

We further compared log-Jacobian determinants of displacements fields from the COPD dataset gained with our method and Adam, showing a lower mean

Table 1. Mean target registration errors in mm of our and comparison methods on the 4DCT and COPD datasets.

	Initial	LapIRN		VM++	Adam		L2O (ours)	
		w/o Adam	w. Adam		w/o Pre-Reg	w. Pre-Reg	w/o Pre-Reg	w. Pre-Reg
4DCT	6.33	2.06	1.60	4.40	2.38	1.33	2.01	1.69
COPD	11.99	3.96	3.83	5.30	7.52	2.18	4.13	2.24

standard deviation (Adam: 0.04, L2O: 0.02) and fewer cases with negative values of determinants, indicating fewer foldings for L2O (Adam: 7 out of 10 cases, L2O: 2 out of 10 cases).

When comparing the results after each interative optimisation step during inference, L2O shows faster convergence than Adam optimisation, also using less iterations in total. Using more iterations than in training does however slightly negatively impact the accuracy.

4 Discussion

We proposed a framework, that learns gradient-based optimisation steps for deformable image registration using a recurrent network.

The network uses less iterations than Adam optimisation, indicating it does not only approximate first-order gradients and imitate gradient descent, but rather learns how to optimise using all available information and overcomes this limitation of gradient descent algorithms. Furthermore it is able to use population-wide information through training data.

Even though our model does not yet always surpass Adam instance optimisation, it combines advantages of deep learning and gradient descent algorithms and yields promising results when it comes to compensating disadvantages of conventional optimisation.

References

1. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. IEEE TMI **38**(8), 1788–1800 (2019)
2. Castillo, R., Castillo, E., Guerra, R., Johnson, V.E., McPhail, T., Garg, A.K., Guerrero, T.: A framework for evaluation of deformable image registration spatial accuracy using large landmark point sets. Physics in Medicine & Biology **54**(7), 1849 (2009)
3. Heinrich, M.P., Hansen, L.: Voxelmorph++: Going beyond the cranial vault with keypoint supervision and multi-channel instance optimisation. In: WBIR (2022)
4. Hering, A., Hansen, L., Mok, T.C., Chung, A., Siebert, H., Häger, S., Lange, A., Kuckertz, S., Heldmann, S., Shao, W., et al.: Learn2Reg: comprehensive multi-task medical image registration challenge, dataset and evaluation in the era of deep learning. arXiv preprint arXiv:2112.04489 (2021)
5. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. ICLR (2015)
6. Mok, T.C., Chung, A.: Large deformation diffeomorphic image registration with laplacian pyramid networks. In: MICCAI. pp. 211–221. Springer (2020)
7. Murphy, K., Van Ginneken, B., Reinhardt, J.M., Kabus, S., Ding, K., Deng, X., Cao, K., Du, K., Christensen, G.E., Garcia, V., et al.: Evaluation of registration methods on thoracic CT: the EMPIRE10 challenge. IEEE TMI **30**(11), 1901–1920 (2011)
8. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: MICCAI. pp. 234–241. Springer (2015)