



Conditional Deformable Image Registration with Convolutional Neural Network

Tony C. W. Mok^(✉) and Albert C. S. Chung

Department of Computer Science and Engineering,
The Hong Kong University of Science and Technology, Hong Kong, China
{cwmokab, achung}@cse.ust.hk

Abstract. Recent deep learning-based methods have shown promising results and runtime advantages in deformable image registration. However, analyzing the effects of hyperparameters and searching for optimal regularization parameters prove to be too prohibitive in deep learning-based methods. This is because it involves training a substantial number of separate models with distinct hyperparameter values. In this paper, we propose a conditional image registration method and a new self-supervised learning paradigm for deep deformable image registration. By learning the conditional features that are correlated with the regularization hyperparameter, we demonstrate that optimal solutions with arbitrary hyperparameters can be captured by a single deep convolutional neural network. In addition, the smoothness of the resulting deformation field can be manipulated with arbitrary strength of smoothness regularization during inference. Extensive experiments on a large-scale brain MRI dataset show that our proposed method enables the precise control of the smoothness of the deformation field without sacrificing the runtime advantage or registration accuracy.

Keywords: Controllable regularization · Conditional image registration · Deformable image registration

1 Introduction

Deformable image registration and the subsequent quantitative assessment are crucial in a variety of medical imaging studies. Recent deep learning-based image registration (DLIR) methods [3, 5, 17, 30, 31] have achieved remarkable results and showed immense potential for time-sensitive medical imaging studies such as image-guided surgery and motion tracking. Unsupervised DLIR methods [3, 12, 23, 24] circumvent costly iterative optimization in conventional image registration approaches by re-formulating the image registration problem as a

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-87202-1_4) contains supplementary material, which is available to authorized users.

learning problem with convolutional neural networks (CNN), resulting in fast image registration. While DLIR methods have a learning formulation that differs from the conventional image registration approaches [1, 2, 28, 29], the tradeoff between registration accuracy and the smoothness of the deformation field, which is often controlled with a hyperparameter in the objective function, cannot be circumvented by DLIR methods. Typically, the optimal hyperparameter is determined using grid searching on the validation dataset [3, 23]. Ironically, despite the runtime advantage of DLIR methods, searching for the optimal hyperparameter value is notoriously time-consuming and computationally intensive in DLIR methods as the hyperparameters are fixed throughout the learning and inference phase. In DLIR methods, each grid search value requires a new DLIR model trained with the distinct hyperparameter value, and each DLIR model requires up to ~ 20 h to a few days to train from scratch [3]. As such, analyzing the effect of hyperparameters and searching for optimal regularization parameters prove to be too prohibitive in DLIR methods, leading to suboptimal registration results and limited clinical applications. Despite the computational cost of the hyperparameter searching technique, the traditional hyperparameter searching technique may not be a good solution for unsupervised DLIR methods for two thoughtful reasons. First, the optimal regularization parameter is subject to the degree of misalignment between the input images, image modality, and intensity distribution. Second, the prior knowledge of the learned model cannot be utilized in the traditional hyperparameter searching technique, resulting in a substantial computational redundancy.

In recent years, a pioneering work of Gatys et al. [9] demonstrate that CNN encodes both the content and style information of an image. Subsequent studies [4, 7, 14, 15] further illustrate that the image information can be separated by manipulating the statistics of the feature maps with feature-wise linear modulation [6] in CNN. In this paper, motivated by these studies [7, 14, 15], we propose a novel conditional image registration method and a new self-supervised learning paradigm for deformable image registration to address the inefficiency of existing hyperparameter searching technique in DLIR methods. Instead of training multiple models for searching the optimal hyperparameter, we propose utilizing a single conditional model with self-supervised learning for efficient hyperparameter tuning.

Parallel to our work, Hoopes et al. [13] propose to learn the effects of registration hyperparameters on deformation field with Hypernetworks [10], which leverage a secondary network to generate the conditioned weights for the entire network layers. While the Hypernetworks-based method offers immense modulation potential, it adds an enormous number of parameters to the original image registration method. Alternatively, we propose a more parameter-efficient and scalable approach based on conditional instance normalization. Our method learns the effect of the regularization parameters and conditions on the feature statistics of high-dimensional layers such that the smoothness of the solution can be manipulated via arbitrary hyperparameter values during the inference phase. We further introduce a novel distributed mapping network to generate non-linear embedding with the condition variable. We present extensive exper-

iments, demonstrating that our formulation enables the precise control of the smoothness of the deformation field during the inference phase and rapid grid search of an optimal hyperparameter without sacrificing the runtime advantage or the registration accuracy of the original DLIR method.

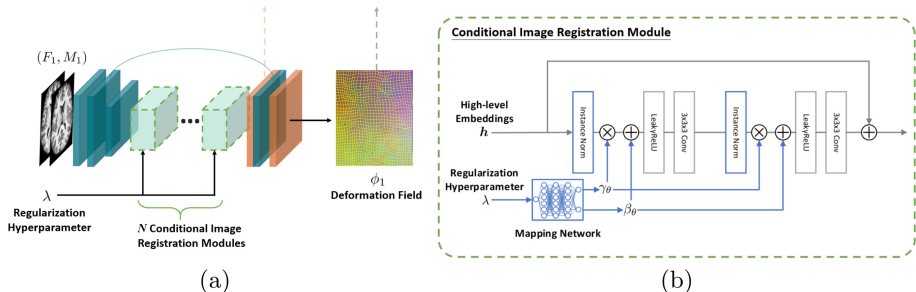


Fig. 1. Overview of the proposed (a) conditional deformable image registration method and (b) the conditional image registration module. For clarity and simplicity, we depict the first pyramid level only and illustrate the 2D formulation of our method in the figure.

2 Methods

Deformable image registration establishes a dense non-linear correspondence between a fixed image F and a moving image M , and the solution ϕ is often subject to a weighted smoothness regularization. DLIR methods often formulate the deformable image registration problem as a learning problem $\phi = f_\theta(F, M)$, in which f_θ is parameterized with CNN. Therefore, in contrast to conventional image registration approaches, the strength of the smoothness regularization is fixed throughout the training and inference phase. To address this limitation, we extend the common formulation of DLIR methods to a conditional deformable image registration setting. Instead of learning to adapt a particular weighted smoothness regularization, our proposed method learns the conditional features that correlated with arbitrary hyperparameter values. In the following sections, we describe the methodology of our proposed method.

2.1 Conditional Deformable Image Registration

Given a fixed F , a moving 3D image scan M , and a conditional variable c , we parametrize the proposed conditional image registration method as a function $f_\theta(F, M, c) = \phi$ with CNN. The proposed method works with any CNN-based DLIR methods and conditional variables. Specifically, we parametrize an example of the function f_θ with the deep Laplacian pyramid image registration network (LapIRN) and set the conditional variable to the smoothness regularization parameter λ . To condition a CNN model on a conditional variable,

a concatenation-based conditioning approach [6, 21, 32, 33] in generative models is to directly concatenate the condition variable with the input image scans. However, based on our experiments, we observed that the concatenation-based conditioning approach cannot capture a wide range of regularization parameters and bias to a limited range of hyperparameter values.

Therefore, we depart from the concatenation-based conditioning approach and extend the feature-wise linear modulation approach [4, 15] instead. We condition the hidden layers on the regularization parameter directly. In particular, the network architecture of LapIRN is comprised of L CNN-based registration networks (CRN). Each CRN consists of three major components: a feature encoder, a set of N residual blocks, and a feature decoder. We replace the N residual blocks with our proposed conditional image registration modules, as shown in Fig. 1(a). The feature encoder extracts the necessary low-level features for deformable image registration, while the feature decoder upsamples and outputs the targeted displacement fields. We only condition the hidden layers in each conditional image registration module on the hyperparameter of the smoothness regularization. We set L and N to 3 and 5 in our experiments, respectively.

2.2 Conditional Image Registration Module

Based on the assumption that the characteristics of the deformation field, i.e. smoothness, can be captured and separated by CNN, we design the conditional image registration module that takes input hidden feature maps and the regularization hyperparameter as input, and outputs hidden features with shifted feature statistics based on conditional instance normalization (CIN) [7]. Specifically, the proposed conditional image registration module adopts the pre-activation structure [11] and includes two CIN layers, each followed by a leaky rectified linear unit (LeakyReLU) activation [18] with a negative slope of 0.2 and a convolutional layer with 28 filters, as depicted in Fig. 1(b). A skip connection is added to preserve the identity of the features.

Conditional Instance Normalization. While the centralized mapping network [15] generates a conditional representation with less memory consumption and computational cost, we argue that the effective representation of the hyperparameter should be diverse and adaptable to different layers in CNN. Chen et al. [4] demonstrate that modulating layers with various depths of CNN results in inconsistent performance, which implies that hidden features of different depths hold distinct feature statistics and non-linearly correspondence to the latent code.

To maintain diverse conditional representations of the hyperparameter for each hidden level, we propose to include distributed mapping networks that learn a separate intermediate non-linear latent variable for each conditional image registration module, which is shared among all the CIN layers. Formally, given a normalized regularization hyperparameter $\lambda \in \bar{\lambda}$, the distributed mapping network $g : \bar{\lambda} \rightarrow \mathcal{Z}$ first maps λ to latent code $z \in \mathcal{Z}$. Then, the CIN layers




Fig. 2. Example axial MR slices of resulting warped images and deformation fields from the baseline method and our proposed method (CIR-DM) with $\lambda \in [0.1, 1, 4, 8]$. The standard deviation of the Jacobian determinant is shown at the upper-right corner of each resulting deformation fields.

learn a set of parameters that specialize \mathbf{z} to the regularization smoothness. The distributed mapping network is parameterized with a 4-layer multilayer perceptron (MLP). For simplicity, we set the number of perceptrons in each MLP layer and the dimensionality of the latent space to 64. The middle layers in the distributed mapping network use the LeakyReLU activation to further introduce the non-linearity into the latent code. The CIN operation for each feature map \mathbf{h}_i is defined as

$$\mathbf{h}'_i = \gamma_{\theta,i}(\mathbf{z}) \left(\frac{\mathbf{h}_i - \mu(\mathbf{h}_i)}{\sigma(\mathbf{h}_i)} \right) + \beta_{\theta,i}(\mathbf{z}), \quad (1)$$

where $\gamma_{\theta,i}, \beta_{\theta,i} \in \mathbb{R}$ are affine parameters learned from the latent code \mathbf{z} , and $\mu(\mathbf{h}_i), \sigma(\mathbf{h}_i) \in \mathbb{R}$ are the channel-wise mean and standard deviation of feature map \mathbf{h}_i in channel i . In other words, the control of smoothness regularization is learned by normalizing and shifting the feature statistics of the feature map with corresponding affine parameters $\gamma_{\theta,i}$ and $\beta_{\theta,i}$ for each channel in the hidden feature map \mathbf{h} .

2.3 Self-supervised Learning

The objective of our proposed method is to compute the optimal deformation field corresponding to the hyperparameter of smoothness regularization. Formally, this task is defined as

$$\phi^* = \arg \min_{\phi} \mathcal{L}_{sim}(F, M(\phi)) + \lambda_p \mathcal{L}_{reg}(\phi), \quad (2)$$

where ϕ^* denotes the optimal displacement field ϕ , $\mathcal{L}_{sim}(\cdot, \cdot)$ denotes the dissimilarity function, $\mathcal{L}_{reg}(\cdot)$ represents the smoothness regularization function and λ_p is uniformly sampled over a predefined range. We set the predefined range of λ_p to $[0, 10]$ empirically such that the optimal deformation field with maximum λ_p is diffeomorphic in most cases. The only difference between the objective in common unsupervised DLIR methods [3, 12, 23, 24] and our objective is that we

learn to optimize the objective function over a predefined range of hyperparameter instead of a fixed hyperparameter value. To exemplify our proposed learning paradigm, we follow [24] and instantiate the objective function with a similarity pyramid and a diffusion regularizer on the spatial gradients of displacement fields. We also adopt a progressive training scheme to train the network in a coarse-to-fine manner. Mathematically, the objective function for each pyramid level $l \in L$ is defined as

$$\mathcal{L}_l(F, M(\phi), \phi, \lambda_p) = \sum_{i \in [1..l]} -\frac{1}{2^{(l-i)}} NCC_w(F_i, M_i(\phi)) + \lambda_p \|\nabla \phi\|_2^2, \quad (3)$$

where λ_p is sampled uniformly in $[0, 10]$ for each iteration and $NCC_w(\cdot, \cdot)$ denotes the local normalized cross-correlation (NCC) with window size w , in which w is set to $1 + 2i$. It is worth noting that our proposed learning paradigm does not introduce extra computational cost to the original objective function and can be easily transferred to various DLIR applications with minimum efforts.

3 Experiments

Data and Pre-processing. We evaluate our method on brain atlas registration tasks. We use 425 T1-weighted brain MR scans from the OASIS [19, 20] dataset and 40 brain MR scans from the LPBA40 [26, 27] dataset. The OASIS dataset contains subjects aged from 18 to 96, and 100 of the included subjects were diagnosed with very mild to moderate Alzheimer’s disease. We follow [24] and perform standard pre-processing, including skull stripping, affine spatial normalization, intensity normalization, and subcortical structures segmentation, for each MR scan using FreeSurfer [8]. For the OASIS dataset, subcortical segmentation maps of 26 anatomical structures serve as the ground truth for the evaluation of our method. For the LPBA40 dataset, the brain MR scans in atlas space and its subcortical segmentation map of 56 anatomical structures, which are delineated by experts, are used in our experiments. We resample all MR scans with isotropic voxel sizes of 1^3 mm and center-cropped all the pre-processed image scans to $144 \times 192 \times 160$. We randomly split the OASIS dataset into 255, 20, and 150 volumes and split the LPBA40 dataset into 28, 2, and 10 volumes for training, validation, and test sets, respectively. We randomly select 3 and 2 MR scans from the test sets as atlases in OASIS and LPBA40, respectively. Finally, we register each subject to the chosen atlas using the baseline method and different conditional deformable image registration methods. In summary, there are 441 and 16 combinations of test scans from OASIS and LPBA40, respectively, included in the evaluation.

Implementation. Our proposed method and the other baseline methods are implemented with PyTorch 1.7 [25] and deployed on the same machine, equipped with an Nvidia Titan RTX GPU and an Intel Core (i7-4790) CPU. We build our

method on top of the official implementation of LapIRN available in [22]. We adopt Adam optimizer [16] with a fixed learning rate 0.0001. We normalize $\bar{\lambda}$ to [0,1]. We train all the methods from scratch (60000 iterations in OASIS and 40000 iterations in LPBA40). The source code will be published online.

Baseline Methods. We compare our method with the original LapIRN [24] with a fixed hyperparameter (denoted as baseline). Specifically, we train seven distinct LapIRNs with different regularization hyperparameters $\lambda \in [0.1, 0.5, 1, 2, 4, 8, 10]$. For each hyperparameter value λ , we select the top-3 models with the highest Dice score on the validation set for evaluation to alleviate the model variation. We further compare it with a concatenation-based conditioning approach (denoted as the traditional method) [6, 21, 33], which simply concatenates the regularization hyperparameter with the input scans in LapIRN to achieve conditional image registration. An ablation study of the variant of our proposed method is performed using either the 8-layer MLP centralized mapping network [15] with latent space 256 (denoted as CIR-CM) and the proposed distributed mapping network (denoted as CIR-DM). For each condition deformable image registration method, we adopt the same training scheme and select the top-3 models with the highest Dice score ($\lambda = 0.1$) on the validation set for evaluation.

Table 1. Quantitative results of the mean DSC and mean $\text{std}(|J_\phi|)$ over seven hyperparameter values on the OASIS and LPBA40 datasets. Initial: spatial normalization.

Method	OASIS						LPBA40					
	DSC	%DSC	$\text{std}(J_\phi)$	$\% \text{std}(J_\phi)$	T_{train}	T_{test}	DSC	%DSC	$\text{std}(J_\phi)$	$\% \text{std}(J_\phi)$	T_{train}	T_{test}
Initial	0.552	-	-	-	-	-	0.560	-	-	-	-	-
Baseline	0.770	-	1.157	-	200.3h	0.204s	0.729	-	0.697	-	143.2h	0.206s
Traditional	0.780	+1.41%	0.970	+4.62%	28.4h	0.212s	0.722	-1.01%	0.440	-12.35%	20.3h	0.210s
CIR-CM	0.767	-0.57%	0.900	-5.23%	28.8h	0.227s	0.721	-1.14%	0.473	-5.73%	20.5h	0.225s
CIR-DM	0.770	-0.19%	0.963	-3.78%	28.5h	0.216s	0.728	-0.17%	0.552	-3.89%	20.4h	0.218s

Measurement. We register each scan in the test set to an atlas, propagate the anatomical segmentation map of the moving image using the resulting deformation field with the nearest-neighbour interpolation, and measure the overlap of the segmentation maps using Dice similarity coefficient (DSC). We also measure the standard deviation of the Jacobian determinant on the deformation fields ($\text{std}(|J_\phi|)$), representing the smoothness and local orientation consistency of the deformation field. Moreover, we compare each individual solution from all conditional methods to the solution of the corresponding test case generated from the baseline method, and measure the average difference (in percentage) of the mean Dice score (%DSC) and the standard deviation of the Jacobian determinant on the deformation fields ($\% \text{std}(|J_\phi|)$) over the total number of test cases. Finally,

we measure the total training time in hours (T_{train}) and the average inference time per case in seconds (T_{test}) for each method. We repeat the experiment with seven distinct hyperparameter values λ . An ideal conditional image registration algorithm should achieve comparably registration accuracy and quality with the baseline method.

Results and Discussions. Table 1 presents a comprehensive summary of the results of each method in the OASIS and LPBA40 datasets. Figure 2 illustrates qualitative results compare to the baseline method and Fig. 3 shows detail results of each method over seven distinct hyperparameter values in the OASIS dataset. We demonstrate that not only does our method achieves highly consistent results with the baseline method, our method significantly reduces the total training time needed to generate solutions with diverse complexities.

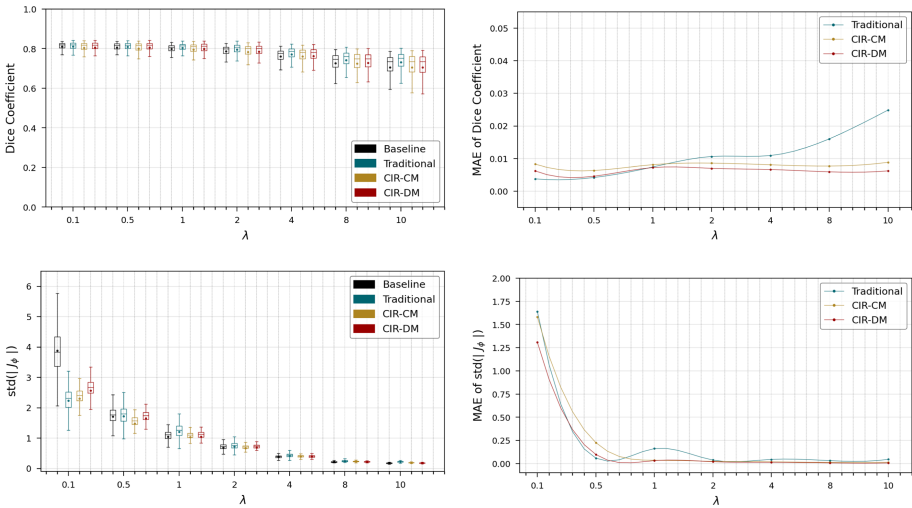


Fig. 3. Quantitative results over seven distinct hyperparameter values on the OASIS dataset. First row: the boxplot of Dice scores and the mean absolute error (MAE) of DSC compared to the baseline method. Second row: the boxplot of $\text{std}(|J_\phi|)$ and the MAE of $\text{std}(|J_\phi|)$ compared to the baseline method. The MAE of DSC (and the $\text{std}(|J_\phi|)$) is computed by averaging the absolute difference of individual solutions between the targeting methods and baseline method over the total number of test cases.

Specifically, all methods under our proposed conditional framework only required one trained model to generate solutions with seven distinct hyperparameter values of the smoothness regularization λ , resulting in $\sim 7x$ faster total training time than the baseline method. Interestingly, we find that the complexity of the resulting deformation fields ($\text{std}(|J_\phi|)$) at $\lambda = 0.1$ declines significantly (-32% to -41%) while maintains comparable Dice scores with the baseline

method, indicating that our methods produce even more desirable (smoother) solutions than the baseline method. In contrast to methods based on conditional instance normalization, the traditional method achieves a consistently higher average Dice score and standard deviation of the Jacobian determinant than the baseline method on the OASIS dataset when $\lambda \geq 2$ as shown in Fig. 3, indicating the traditional method tends to bias to a limited range of λ . Compare to CIR-CM, our distributed mapping network design is in every way superior to the centralized mapping network in the context of conditional deformable image registration, as shown in Fig. 3. Importantly, our method achieves only -0.19% (-0.17% on LPBA40) difference of mean Dice score compared to the results of baseline method on OASIS, and the average inference time of CIR-DM is ~ 0.21 seconds, highlighting the fact that CIR-DM is the only method that enables precise control of the deformation field regarding diverse λ without sacrificing the registration accuracy or the runtime advantage of DLIR methods.

4 Conclusion

In summary, we have presented a novel conditional deformable image registration framework and self-supervised learning paradigm for deep learning-based deformable image registration. Our method learns the conditional features that are correlated with the regularization hyperparameter by shifting the feature statistics. It is demonstrated that our method enables precise control of the smoothness regularization in the inference phase without sacrificing the runtime advantage or the registration accuracy of the original DLIR method. Extensive experiments on brain atlas registration have been carried out, demonstrating that the results of our method consistently align with the results of the original DLIR method, and our method is superior to the common conditional approaches with diverse hyperparameter values. In principle, the proposed conditional image registration framework can be easily transferred to arbitrary CNN-based image registration approaches for controllable regularization of the deformation field and rapid hyperparameter tuning.

References

1. Ashburner, J.: A fast diffeomorphic image registration algorithm. *Neuroimage* **38**(1), 95–113 (2007)
2. Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image Anal.* **12**(1), 26–41 (2008)
3. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Gutttag, J., Dalca, A.V.: An unsupervised learning model for deformable medical image registration. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 9252–9260 (2018)
4. Chen, T., Lucic, M., Houlsby, N., Gelly, S.: On self modulation for generative adversarial networks. In: *International Conference on Learning Representations* (2019)

5. Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning for fast probabilistic diffeomorphic registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 729–738. Springer (2018). https://doi.org/10.1007/978-3-030-00928-1_82
6. Dumoulin, V., et al.: Feature-wise transformations. *Distill* **3**(7), e11 (2018)
7. Dumoulin, V., Shlens, J., Kudlur, M.: A learned representation for artistic style. International Conference on Learning Representations (2017)
8. Fischl, B.: Freesurfer. *Neuroimage* **62**(2), 774–781 (2012)
9. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2414–2423 (2016)
10. Ha, D., Dai, A., Le, Q.V.: Hypernetworks. [arXiv:1609.09106](https://arxiv.org/abs/1609.09106) (2016)
11. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: European Conference on Computer Vision. pp. 630–645. Springer (2016). https://doi.org/10.1007/978-3-319-46493-0_38
12. Hering, A., van Ginneken, B., Heldmann, S.: mlvirnet: multilevel variational image registration network. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 257–265. Springer (2019). https://doi.org/10.1007/978-3-030-32226-7_29
13. Hoopes, A., Hoffmann, M., Fischl, B., Guttag, J., Dalca, A.V.: Hypermorph: amortized hyperparameter learning for image registration. In: International Conference on Information Processing in Medical Imaging (2021)
14. Huang, X., Belongie, S.: Arbitrary style transfer in real-time with adaptive instance normalization. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 1501–1510 (2017)
15. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4401–4410 (2019)
16. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
17. Krebs, J., et al.: Robust non-rigid registration through agent-based action learning. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 344–352. Springer (2017). https://doi.org/10.1007/978-3-319-66182-7_40
18. Maas, A.L., Hannun, A.Y., Ng, A.Y.: Rectifier nonlinearities improve neural network acoustic models. In: International Conference on Machine Learning (ICML), vol. 30, p. 3 (2013)
19. Marcus, D.S., Wang, T.H., Parker, J., Csernansky, J.G., Morris, J.C., Buckner, R.L.: Oasis brains - open access series of imaging studies. <https://www.oasis-brains.org/>. Accessed on 01 March 2021
20. 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)
21. Mirza, M., Osindero, S.: Conditional generative adversarial nets. [arXiv:1411.1784](https://arxiv.org/abs/1411.1784) (2014)
22. Mok, T.C., Chung, A.: Official implementation of laplacian pyramid image registration network. <https://github.com/cwmok/LapIRN>. Accessed on 01 March 2021
23. Mok, T.C., Chung, A.: Fast symmetric diffeomorphic image registration with convolutional neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4644–4653 (2020)

24. Mok, T.C., Chung, A.C.: Large deformation diffeomorphic image registration with laplacian pyramid networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention, pp. 211–221. Springer (2020). https://doi.org/10.1007/978-3-030-59716-0_21
25. Paszke, A., Gross, S., Chintala, S., et al.: Automatic differentiation in pytorch. In: NIPS-W (2017)
26. Shattuck, D.W., et al.: Lpba40 atlases download. <https://resource.loni.usc.edu/resources/atlasses-downloads/>. Accessed 0n 01 March 2021
27. Shattuck, D.W., et al.: Construction of a 3d probabilistic atlas of human cortical structures. *Neuroimage* **39**(3), 1064–1080 (2008)
28. Thirion, J.P.: Image matching as a diffusion process: an analogy with maxwell’s demons. *Med. Image Anal.* **2**(3), 243–260 (1998)
29. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration. *NeuroImage* **45**(1), S61–S72 (2009)
30. de Vos, B.D., Berendsen, F.F., Viergever, M.A., Staring, M., Išgum, I.: End-to-end unsupervised deformable image registration with a convolutional neural network. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support, pp. 204–212. Springer (2017). https://doi.org/10.1007/978-3-319-67558-9_24
31. Yang, X., Kwitt, R., Styner, M., Niethammer, M.: Quicksilver: fast predictive image registration—a deep learning approach. *NeuroImage* **158**, 378–396 (2017)
32. Zhang, H., et al.: Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 5907–5915 (2017)
33. Zhang, H., et al.: Stackgan++: Realistic image synthesis with stacked generative adversarial networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **41**(8), 1947–1962 (2018)