

Toward Universal Medical Image Registration via Sharpness-Aware Meta-Continual Learning

Bomin Wang^{*}, Xinzhe Luo^{*}, and Xiahai Zhuang[✉]

School of Data Science, Fudan University, Shanghai, China
bmwang21@m.fudan.edu.cn, {xzluo19, zxh}@fudan.edu.cn

Abstract. Current deep learning approaches in medical image registration usually face the challenges of distribution shift and data collection, hindering real-world deployment. In contrast, universal medical image registration aims to perform registration on a wide range of clinically relevant tasks simultaneously, thus having tremendous potential for clinical applications. In this paper, we present the first attempt to achieve the goal of universal 3D medical image registration in sequential learning scenarios by proposing a continual learning method. Specifically, we utilize meta-learning with experience replay to mitigating the problem of catastrophic forgetting. To promote the generalizability of meta-continual learning, we further propose sharpness-aware meta-continual learning (SAMCL). We validate the effectiveness of our method on four datasets in a continual learning setup, including brain MR, abdomen CT, lung CT, and abdomen MR-CT image pairs. Results have shown the potential of SAMCL in realizing universal image registration, which performs better than or on par with vanilla sequential or centralized multi-task training strategies. The source code will be available from <https://github.com/xzluo97/Continual-Reg>.

1 Introduction

Medical image registration aims to estimate the optimal spatial transformation to align the structures of interest between a pair of fixed and moving images, which is a fundamental task in medical image analysis and has been widely used in disease monitoring, surgical navigation, and image diagnostics [6, 17, 18]. Over the past decade, advancements in artificial intelligence (AI) and deep learning (DL) technologies have significantly transformed the landscape of medical image registration. DL-based methods have shown promising results in enhancing the accuracy and efficiency of image registration [7]. However, two challenges exist that hinder the real-world applications of DL-based methods.

The challenge of distribution shift. Due to the data-dependent nature of DL algorithms, their performance significantly depends on how the data are generated in specific contexts, at specific times. Therefore, most current learning-based medical image registration methods are task-specific (e.g., anatomy- and

^{*} Equal contribution. [✉] Corresponding author.

appearance-specific), and thus can be performed only on data belonging to the distribution of a single registration task [3, 8, 13]. When dealing with out-of-distribution data, significant performance decay will occur.

The challenge of centralized data collection. In highly dynamic environments like healthcare, data are typically sequentially collected over time. In addition, the issue of data privacy makes it challenging to centralize the medical data. Therefore, it is difficult to collect all data simultaneously to train a universal registration model for different tasks. Meanwhile, the problem of distribution shifts occurs when sequentially learning on different datasets using current DL methods, which can lead to catastrophic forgetting of previously learned tasks, that is the performance degrades significantly after learning new data. Currently, DL-based algorithms cannot be modified after approval for clinical use [16, 26], thereby precluding the potential for achieving universal registration models. Continual learning (CL) methods aim to endow the ability to learn continuously to DL models while reducing forgetting [15, 23, 24].

Toward Universal Medical Image Registration via Continual Learning. In stark contrast to current task-specific DL algorithms, *universal medical image registration* is a general-purpose AI for medical imaging, which performs registration on a wide range of clinically relevant tasks simultaneously while maintaining its generalizability to unseen data, and thus has huge potential to be applied to real-world clinics. To tackle the above challenges, we present the first attempt to achieve the goal of universal 3D medical image registration by proposing a CL method. The proposed method continually updates the registration network to adapt to new clinical environments and data and retain the registration performance of old tasks at the same time. After sequential training on all tasks encountered, the resulting network can perform well on all tasks. Taking four tasks as an example, we illustrate the problem considered in Fig. 1. Specifically, We utilize meta-learning with experience replay to mitigate the problem of catastrophic forgetting. To promote the generalizability of meta-continual learning, we further propose sharpness-aware meta-continual learning (SAMCL) to seek a flat loss landscape. Empirical results on four medical image registration datasets show the effectiveness of our proposed method.

2 Methodology

Image registration finds the spatial transformation $\phi : \Omega_F \rightarrow \Omega_M$ that registers the moving image $m : \Omega_M \rightarrow \mathbb{R}$ to the fixed image $f : \Omega_F \rightarrow \mathbb{R}$. Deep learning-based registration methods parameterize the spatial transformation by a neural network taking the image pair as the input with parameters θ , *i.e.*, $\phi_\theta = g(f, m; \theta)$. Registration is achieved by optimizing the empirical risk, *i.e.*,

$$\hat{\theta} = \arg \min_{\theta} \mathbb{E}_{(f, m) \sim \mathcal{D}} [\mathcal{L}(\phi_\theta; f, m)], \quad (1)$$

where $\mathcal{D} = \{(f_i, m_i)\}_{i=1}^N$ is a training dataset drawn *i.i.d.* from distribution $\mathcal{P}$, and $\mathcal{L}(\phi_\theta; f, m) = D(f, m \circ \phi_\theta) + R(\phi_\theta)$ is the loss function composed of a dissimilarity measure $D(\cdot, \cdot)$ and a regularization term $R(\cdot)$.




Fig. 1: Continual image registration. One can see that the image pairs differ significantly in appearance and anatomy across tasks, *i.e.*, serious distribution shift, posing a major challenge to modern learning-based registration methods.

Collecting a centralized multi-task dataset for training a universal registration network is cumbersome and unrealistic in practice. In the context of lifelong or *continual* learning (CL) [25], we train the registration network sequentially on non-stationary datasets $\{\mathcal{D}^{(t)}\}_{t=1}^T$, while achieving universal image registration. This objective can be formulated into the following optimization problem

$$\hat{\theta}^{(K)} = \arg \min_{\theta} \sum_{t=1}^K \mathbb{E}_{(f,m) \sim \mathcal{D}^{(t)}} \left[\mathcal{L}^{(t)}(\phi_{\theta}; f, m) \right], \quad \forall K \in \{1, \dots, T\} \quad (2)$$

where $\mathcal{L}^{(t)}$ is the loss function for the t -th dataset, and $\hat{\theta}^{(K)}$ are the optimal network parameters for the K tasks that have been observed. When the model learns from one individual dataset to another, it inevitably encounters the issue of generalization toward new datasets and the forgetting of knowledge on previous datasets due to distribution shifts. We propose the method SAMCL to alleviate forgetting of the registration network while promoting its generalizability.

2.1 Sharpness-Aware Meta-Continual Learning (SAMCL)

This bi-directional temporal perspective between the past and the future has motivated using meta learning as a means to achieve CL, *a.k.a.* meta-continual learning (Meta-CL). Meta-CL relies on rehearsal that replays previous task samples using a small memory buffer and encourages transfer between past and current samples. For instance, Meta Experience Replay (MER) [21] exploits the Reptile algorithm [20], which essentially optimizes the following objective:

$$\min_{\theta} \mathbb{E}_{(f_1, m_1), \dots, (f_b, m_b) \sim \mathcal{M}} \left[\sum_{i=1}^b \left(\mathcal{L}^{(t_i)}(\phi_{\theta}) - \frac{\alpha}{2} \sum_{j=1}^{i-1} \nabla_{\theta} \mathcal{L}^{(t_i)}(\phi_{\theta}) \cdot \nabla_{\theta} \mathcal{L}^{(t_j)}(\phi_{\theta}) \right) \right]$$

where $\cdot$ is the inner product, $\mathcal{M}$ is a memory buffer of size b acquired by Experience Replay (ER) [22], and $\mathcal{L}^{(t_i)}$ is the registration loss for the image pair $(f, m) \sim \mathcal{D}^{(t_i)}$. Therefore, MER tries to find a gradient direction that can simultaneously minimize losses on all datasets using samples from the memory buffer. Nevertheless, since the objective function for deformable image registration is highly *non-convex and there may exist numerous local minima* [10], leading to poor generalizability. Besides, its gradient *w.r.t.* the network parameters can be very noisy and may not represent a direction minimizing the expected loss $\mathbb{E}_{(f,m) \sim \mathcal{P}} [\mathcal{L}(\phi_\theta; f, m)]$. Inspired by the recent progress in sharpness-aware minimization (SAM) which improves the generalizability of neural networks [5], we hypothesize that integrating SAM into Meta-CL could enhance the performance on continual image registration. This new approach optimizes the objective

$$\min_{\theta} \mathbb{E}_{(f_1, m_1), \dots, (f_b, m_b) \sim \mathcal{M}} \left[\sum_{i=1}^b \left(\mathcal{L}_{\text{sam}}^{(t_i)}(\phi_\theta) - \frac{\alpha}{2} \sum_{j=1}^{i-1} \nabla_{\theta} \mathcal{L}_{\text{sam}}^{(t_i)}(\phi_\theta) \cdot \nabla_{\theta} \mathcal{L}_{\text{sam}}^{(t_j)}(\phi_\theta) \right) \right] \quad (3)$$

where $\mathcal{L}_{\text{sam}}^{(t_i)}(\phi_\theta) \triangleq \max_{\epsilon: \|\epsilon\|_2 \leq \rho} \mathcal{L}^{(t_i)}(\phi_{\theta+\epsilon})$ and $\rho \geq 0$ is a hyperparameter. The motivation of SAM is to seek parameters that lie in neighborhoods having uniformly low loss value (rather than parameters that only themselves have low loss value), thus simultaneously minimizing loss value and loss sharpness. To solve the min-max optimization problem, the inner maximization problem is first approximated via a first-order Taylor expansion of $\mathcal{L}^{(t_i)}(\phi_{\theta+\epsilon})$ around $\epsilon = \mathbf{0}$. By solving a classical dual norm problem [5], we can obtain the solution

$$\begin{aligned} \hat{\epsilon}(\theta) &\triangleq \arg \max_{\epsilon: \|\epsilon\|_2 \leq \rho} \mathcal{L}^{(t_i)}(\phi_{\theta+\epsilon}) \approx \arg \max_{\epsilon: \|\epsilon\|_2 \leq \rho} \left\{ \mathcal{L}^{(t_i)}(\phi_\theta) + \epsilon^\top \nabla_{\theta} \mathcal{L}^{(t_i)}(\phi_\theta) \right\} \\ &= \arg \max_{\epsilon: \|\epsilon\|_2 \leq \rho} \epsilon^\top \nabla_{\theta} \mathcal{L}^{(t_i)}(\phi_\theta) = \rho \frac{\nabla_{\theta} \mathcal{L}^{(t_i)}(\phi_\theta)}{\|\nabla_{\theta} \mathcal{L}^{(t_i)}(\phi_\theta)\|_2}. \end{aligned} \quad (4)$$

As SAM seeks for flat local minima, there would be increasing possibilities that the gradient $\nabla_{\theta} \mathcal{L}_{\text{sam}}^{(t_i)}(\phi_\theta) \approx \nabla_{\theta} \mathcal{L}^{(t_i)}(\phi_\theta)|_{\theta+\hat{\epsilon}(\theta)}$ corresponds to the direction that minimizes the expected loss with $\hat{\epsilon}(\theta)$. Algorithm 1 summarizes the proposed algorithm for sharpness-aware meta-continual learning. We achieve universal 3D medical image registration via CL on sequentially observed datasets.

We use the Large Kernel U-Net (LK-UNet) as the backbone with 16 kernels in its first layer, which performed well on several registration tasks [12]. The objective of proposed SAMCL is to retain the registration accuracy on observed datasets after the network is sequentially trained by all the tasks, as well as improve its generalizability on unseen domains.

3 Experiments and Results

Materials and Compared Methods. The experiments were performed on four sequentially arrived datasets from the Learn2Reg challenges [9], in order

Algorithm 1: Sharpness-Aware Meta-Continual Learning (SAMCL)

Data: Training image pairs $\mathcal{D}^{(t)} = \{(f_i, m_i)\}_{i=1}^{N_t}$, $t = 1, \dots, T$;
Input: Learning rate α , meta-learning rate β , neighbourhood size ρ , batch size s ;
Output: Trained network parameters $\hat{\theta}^{(T)}$;

```

1 Initialization: randomly sample  $\hat{\theta}^{(0)}$ ; set memory buffer  $\mathcal{M} = \emptyset$ ;
2 for  $t = 1, \dots, T$  do
3   Initialize batches  $\mathcal{B}_t = \{(f_i, m_i)\}_{i=1}^{b_t}$ ; initialize  $\theta \leftarrow \hat{\theta}^{(t-1)}$ ;
4   for  $i = 1, \dots, b_t$  do
5      $\theta_0 \leftarrow \theta$ ;  $B_i \leftarrow (f_i, m_i)$ ;
6     for  $j = 1, \dots, s$  do
7       Compute  $\hat{\epsilon}(\theta) = \rho \nabla_{\theta} \mathcal{L}^{(t)}(\phi_{\theta}; B_i[j]) / \|\nabla_{\theta} \mathcal{L}^{(t)}(\phi_{\theta}; B_i[j])\|_2$ ;
8       Update parameters  $\theta \leftarrow \theta - \alpha \nabla_{\theta} \mathcal{L}^{(t)}(\phi_{\theta}; B_i[j])|_{\theta + \hat{\epsilon}(\theta)}$ ;
9      $\mathcal{M} \leftarrow \mathcal{M} \cup \{(f_i, m_i)\}$ ; // Reservoir sampling memory update
10     $B_i \leftarrow \text{sample}(\mathcal{M})$ ; // sample from the memory buffer
11    Compute  $\hat{\epsilon}(\theta) = \rho \nabla_{\theta} \mathcal{L}^{(t_i)}(\phi_{\theta}; B_i) / \|\nabla_{\theta} \mathcal{L}^{(t_i)}(\phi_{\theta}; B_i)\|_2$ ;
12    Update parameters  $\theta \leftarrow \theta - \alpha \nabla_{\theta} \mathcal{L}^{(t_i)}(\phi_{\theta}; B_i)|_{\theta + \hat{\epsilon}(\theta)}$ ;
13     $\theta \leftarrow \theta + \beta(\theta - \theta_0)$ ; // meta-learning update
14   $\hat{\theta}^{(t)} \leftarrow \theta$ ;
15 return  $\hat{\theta}^{(T)}$ .

```

of OASIS (brain MRI, inter-subject) [19], Abdomen CT-CT (inter-subject) [27], NLST (lung CT, intra-subject) [4], and Abdomen MR-CT (intra-subject) [4, 27]. The tasks were assigned as in Fig. 1, varying in image modalities and ROIs. All image volumes were resampled to spacing of $2 \times 2 \times 2$ mm for the OASIS and $3 \times 3 \times 3$ mm for the other datasets, and then cropped/padded to size of $112 \times 96 \times 112$. The number of training/validation/test images for each task is 290/40/84 (OASIS), 21/3/6 (Abdomen CT-CT), 294/42/84 (NLST), and 80/9/16 (Abdomen MR-CT). The intensity range of CT images was clipped to $[-200, 300]$ to cover abdominal organs. All images were linearly normalized to $[0, 1]$ as network input (for CT images this was after intensity clipping). To promote fair comparison across compared methods, no data augmentation was performed on the training set. We compared the proposed SAMCL with different baselines, including (i) independent learning on each task using different models, (ii) multi-task learning on all tasks using a single model, (iii) vanilla sequential learning over all tasks using a single model, and (iv) previous CL methods EWC [15], SI [28], GPM [24], and MER [21]. Note that independent and multi-task learning can be the upper bounds for task-specific performance and average performance of all tasks, respectively.

Evaluation Metrics and Implementation Details. The quality of registration was evaluated by anatomical labels or landmarks. For the OASIS, Abdomen

Table 1: Evaluation metrics for universal medical image registration on the OASIS/Abdomen CT-CT/NLST/Abdomen MR-CT datasets, respectively. The results for NLST in terms of TRE are marked in italics. Better performance of SAMCL than other methods is indicated by underlining. Note that the NLST data were evaluated by TREs in millimeters. Statistical significant difference indicated by a paired t -test for $p < 0.05$ was marked by asterisks.

Method	AVG Dice $\uparrow$ / <i>TRE</i> (mm) $\downarrow$	BWT Dice $\uparrow$ / <i>TRE</i> (mm) $\downarrow$
None	0.598*/0.311*/ <i>7.932*/0.353*</i>	—
Independent	0.738*/0.335*/ <i>3.259*/0.550</i>	—
Multi-task	0.743*/0.358/ <i>2.662*/0.506</i>	—
Sequential	0.348*/0.270/ <i>6.922*/0.568</i>	-0.395*/-0.129*/ <i>3.508*/0</i>
EWC [15]	0.659*/0.364*/ <i>3.836*/0.506</i>	-0.083*/-0.011*/ <i>0.386/0</i>
SI [28]	0.461*/0.303*/ <i>8.186*/0.555</i>	-0.281*/-0.084*/ <i>4.685*/0</i>
GPM [24]	0.458*/0.290*/ <i>7.095*/0.556</i>	-0.281*/-0.085*/ <i>3.574*/0</i>
MER [21]	0.698/ <i>0.331*/3.611*/0.532</i>	-0.036/-0.040/ <i>0.077*/0</i>
SAMCL	0.697/0.352/ <i>3.855/0.563</i>	-0.037/-0.039/ <i>0.306/0</i>

CT-CT and MR-CT datasets, the average Dice similarity coefficient (DSC) was computed on 35, 13, and 9 corresponding labels, respectively. For the NLST dataset, we use the target registration error (TRE) in millimeters on corresponding landmarks as the evaluation metric. Denoting $R_{i,j}$ as the average metric evaluated on task j by the model trained right after task i , we report the average performance (AVG) and backward transfer (BWT) defined as

$$\text{AVG}_i \triangleq R_{T,i}, \quad \text{BWT}_i \triangleq R_{T,i} - R_{i,i}. \quad (5)$$

AVG_i refers to the performance of the registration network on test images of task i after learning all tasks, while BWT_i signifies the forgetting of task i , i.e., the performance decay caused by subsequent tasks. Experiments were conducted using PyTorch on an NVIDIA RTXTM 3090 GPU. For each task, training was performed for 10000 iterations by the Adam optimizer [14] with a learning rate of 1×10^{-4} . The batch size was set to 4. The memory buffer size of SAMCL was set to 200 while the meta-learning rate was chosen as 0.25. The ρ parameter of SAM was set to 0.05 as it produced the best performance. Besides, since the images from different tasks are annotated by different labels, the loss function for each task also varies. Specifically, for the OASIS, Abdomen CT-CT and Abdomen MR-CT tasks the loss function consists of a negative local normalized cross-correlation (LNCC) with window size of 3 [2] and a soft dice loss [11], while for NLST images the loss function comprises a negative LNCC and a landmark distance loss based on TRE. The registration is performed symmetrically between the fixed and moving images using a diffeomorphic transformation parameterized by stationary velocity fields. A regularization term composed of membrane and




Fig. 2: Registration results on an exemplar OASIS and Abdomen CT image pair using different training strategies. The top row visualizes the difference map between the fixed image and the registered moving image. Note that CT intensities were clipped to the range $[-200, 300]$.

bending energy is also included to enforce deformation smoothness [1], whose weights were set to 1 during training.

Comparison Study. Table 1 presents the results of the comparison study. In terms of average registration accuracy (AVG Dice/TRE), SAMCL outperformed previous CL methods like EWC, SI and GPM on at least one of the four tasks. Note that for the Abdomen CT-CT task, both independent and multi-task learning can only achieve minor improvement due to the scarcity of training data and large inter-subject deformations. In terms of backward transfer (BWT Dice/TRE) which indicates the stability and robustness to knowledge forgetting of a model, SAMCL achieved significant improvement over sequential training and performed better than or on par with previous CL methods, which demonstrates its advantage in a continual learning setup.

Furthermore, SAMCL performed comparable to independent training and centralized multi-task learning, indicating its potential for universal image registration. Fig. 2 visualizes an exemplar registered image pair from the OASIS and Abdomen CT data using different training strategies, where SAMCL outperformed previous CL methods and achieved comparable performance to centralized multi-task training. Please refer to the supplementary material for more visualization of the results.

Ablation and Generalization Study. Fig. 3(a) performs the ablation study on using different memory buffer sizes for our proposed SAMCL. One can observe that for the first three tasks (OASIS, Abdomen CT-CT and MR-CT), increasing buffer size has a positive effect on the registration accuracy. This is because




Fig. 3: Ablation and generalization study of the proposed SAMCL on (a) using various memory buffer sizes for experience replay, and (b) using sharpness-aware minimization (SAM) for in- and out-of-domain generalization tasks.

using a larger memory size improves the preservation of knowledge on previous tasks. We also investigate the effectiveness of the sharpness-aware minimization (SAM) in terms of its performance gains for in- and out-of-domain generalization, where a model trained by early tasks is tested on unseen data from later tasks. Fig. 3(b) presents the bar plots for different generalization tasks with or without SAM. The result shows that for both in-domain (Abdomen CT-CT/MR-CT) and out-of-domain (Abdomen CT-CT to MR-CT, Abdomen CT-CT to NLST, OASIS to NLST) generalization tasks, our model equipped with SAM outperforms the vanilla Meta-CL without SAM, highlighting the advantage of SAMCL as a training strategy for simultaneously improving model generalizability.

4 Conclusion and Discussion

In this paper, we introduce the first effort toward accomplishing the goal of universal 3D medical image registration through a continual learning approach. We utilize meta-learning with experience replay to tackle the catastrophic forgetting problem. We propose to use sharpness-aware training to enhance the generalization performance. Empirical results show that the proposed method can achieve promising registration performance, comparable to centralized multitask and independent training strategies. The sequence in which tasks are presented to the model may affect the universal registration performance. In future work, we will evaluate the task order robustness of SAMCL with more datasets from different modalities and anatomies.

Acknowledgments. This work was funded by the National Natural Science Foundation of China (grant No. 62372115, 61971142 and 62111530195).

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

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. *Medical Image Analysis* **12**(1), 26–41 (2008)
3. Che, T., Wang, X., Zhao, K., Zhao, Y., Zeng, D., Li, Q., Zheng, Y., Yang, N., Wang, J., Li, S.: Amnet: Adaptive multi-level network for deformable registration of 3d brain mr images. *Medical Image Analysis* **85**, 102740 (2023)
4. Clark, K.W., Vendt, B.A., Smith, K.E., Freymann, J.B., Kirby, J.S., Koppel, P., Moore, S.M., Phillips, S.R., Maffitt, D.R., Pringle, M., Tarbox, L., Prior, F.W.: The cancer imaging archive (tcia): Maintaining and operating a public information repository. *Journal of Digital Imaging* **26**, 1045–1057 (2013)
5. Foret, P., Kleiner, A., Mobahi, H., Neyshabur, B.: Sharpness-aware minimization for efficiently improving generalization. In: *International Conference on Learning Representations* (2021)
6. Gupta, S., Gupta, P., Verma, V.S.: Study on anatomical and functional medical image registration methods. *Neurocomputing* **452**, 534–548 (2021)
7. Haskins, G., Kruger, U., Yan, P.: Deep learning in medical image registration: a survey. *Machine Vision and Applications* **31**, 1–18 (2020)
8. Hering, A., Häger, S., Moltz, J., Lessmann, N., Heldmann, S., van Ginneken, B.: Cnn-based lung ct registration with multiple anatomical constraints. *Medical Image Analysis* **72**, 102139 (2021)
9. Hering, A., Hansen, L., Mok, T.C.W., Chung, A.C.S., Siebert, H., Hager, S., Lange, A., Kuckertz, S., Heldmann, S., Shao, W., Vesal, S., Rusu, M., Sonn, G.A., Estienne, T., Vakalopoulou, M., Han, L., Huang, Y., Brudfors, M., Balbastre, Y., Joutard, S., Modat, M., Lifshitz, G., Raviv, D., Lv, J., Li, Q., Jaouen, V., Visvikis, D., Fourcade, C., Rubeaux, M., Pan, W.Q., Xu, Z., Jian, B., Benetti, F.D., Wodzinski, M., Gunnarsson, N., Qiu, H., Li, Z., Grossbrohmer, C., Hoopes, A., Reinertsen, I., Xiao, Y., Landman, B.A., Huo, Y., Murphy, K., van Ginneken, B., Dalca, A.V., Heinrich, M.P.: Learn2reg: Comprehensive multi-task medical image registration challenge, dataset and evaluation in the era of deep learning. *IEEE Transactions on Medical Imaging* **42**, 697–712 (2021)
10. Hill, D.L.G., Batchelor, P.G., Holden, M., Hawkes, D.J.: Medical image registration. *Physics in Medicine & Biology* **46**, R1 – R45 (2001)
11. Hu, Y., Modat, M., Gibson, E., Li, W., Ghavami, N., Bonmati, E., Wang, G., Bandula, S., Moore, C.M., Emberton, M., et al.: Weakly-supervised convolutional neural networks for multimodal image registration. *Medical Image Analysis* **49**, 1–13 (2018)
12. Jia, X., Bartlett, J., Zhang, T., Lu, W., Qiu, Z., Duan, J.: U-net vs transformer: Is u-net outdated in medical image registration? In: *International Workshop on Machine Learning in Medical Imaging*. pp. 151–160. Springer (2022)
13. Kang, M., Hu, X., Huang, W., Scott, M.R., Reyes, M.: Dual-stream pyramid registration network. *Medical Image Analysis* **78**, 102379 (2022)
14. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *International Conference on Learning Representations* (2015)
15. Kirkpatrick, J., Pascanu, R., Rabinowitz, N.C., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D.,

- Clopath, C., Kumaran, D., Hadsell, R.: Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences* **114**, 3521 – 3526 (2016)
16. Lee, C.S., Lee, A.Y.: Clinical applications of continual learning machine learning. *The Lancet Digital Health* **2**(6), e279–e281 (2020)
 17. Maes, F., Collignon, A., Vandermeulen, D., Marchal, G., Suetens, P.: Multimodality image registration by maximization of mutual information. *IEEE Transactions on Medical Imaging* **16**(2), 187–198 (1997)
 18. Maintz, J.A., Viergever, M.A.: A survey of medical image registration. *Medical Image Analysis* **2**(1), 1–36 (1998)
 19. 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. *Journal of Cognitive Neuroscience* **19**, 1498–1507 (2007)
 20. Nichol, A., Achiam, J., Schulman, J.: On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999* (2018)
 21. Riemer, M., Cases, I., Ajemian, R., Liu, M., Rish, I., Tu, Y., Tesauro, G.: Learning to learn without forgetting by maximizing transfer and minimizing interference. In: *International Conference on Learning Representations* (2019)
 22. Rolnick, D., Ahuja, A., Schwarz, J., Lillicrap, T., Wayne, G.: Experience replay for continual learning. *Advances in Neural Information Processing Systems* **32** (2019)
 23. Rolnick, D., Ahuja, A., Schwarz, J., Lillicrap, T.P., Wayne, G.: Experience replay for continual learning. In: *Neural Information Processing Systems* (2018)
 24. Saha, G., Garg, I., Roy, K.: Gradient projection memory for continual learning. In: *International Conference on Learning Representations* (2021)
 25. Thrun, S.: Is learning the n-th thing any easier than learning the first? In: *Advances in Neural Information Processing Systems* (1995)
 26. Vokinger, K.N., Feuerriegel, S., Kesselheim, A.S.: Continual learning in medical devices: Fda’s action plan and beyond. *The Lancet Digital Health* **3**(6), e337–e338 (2021)
 27. Xu, Z., Lee, C.P., Heinrich, M.P., Modat, M., Rueckert, D., Ourselin, S., Abramson, R.G., Landman, B.A.: Evaluation of six registration methods for the human abdomen on clinically acquired ct. *IEEE Transactions on Biomedical Engineering* **63**, 1563–1572 (2016)
 28. Zenke, F., Poole, B., Ganguli, S.: Continual learning through synaptic intelligence. In: *International Conference on Machine Learning*. pp. 3987–3995. PMLR (2017)