PromptReg: Universal Medical Image Registration via Task Prompt Learning and Domain Knowledge Transfer*

Housheng Xie, Xiaoru Gao, and Guoyan Zheng (✉)

Institute of Medical Robotics, School of Biomedical Engineering, Shanghai Jiao Tong University, No. 800, Dongchuan Road, Shanghai 200240, China
guoyan.zheng@sjtu.edu.cn

Abstract. Most existing deep learning-based registration methods are typically constrained to dataset-specific optimization, requiring separate models for different data characteristics. In contrast, training a single model across diverse datasets presents an opportunity to create a universal registration framework capable of handling multiple domains simultaneously. However, key challenges remain in achieving effective cross-dataset adaptation while maintaining robust generalization capabilities, particularly for zero-shot registration tasks. In this work, we propose PromptReg, a universal image registration framework that incorporates prompt learning to guide the model in effectively adapting to different registration scenarios through explicit task prompts. The core of PromptReg is a Registration Prompt Generator (RPG) that generates domain-specific task prompts based on the domains of input images. Specifically, we first introduce a Static Knowledge Base (SKB) to store domain prompts and a dynamic prompt generation mechanism that projects different inputs into a shared prompt space. Then, we propose an adaptive prompt fusion strategy that combines stored domain knowledge based on the similarity between the generated dynamic prompt and the prompts in SKB, creating transferable knowledge for unseen domains. Finally, we optimize the prompt generator using domain orthogonality and task similarity losses. Our experiments show that PromptReg achieves competitive performance in universal registration and offers stronger zero-shot generalization. The code is available at https://github.com/xiehousheng/PromptReg.

Keywords: Medical image registration · Prompt learning · Zero-shot.

1 Introduction

Recent advances in deep learning have revolutionized medical image registration [2, 6, 10, 18, 20, 19], establishing learning-based approaches as superior alterna-

H. Xie and X. Gao-Contribute equally to the paper.

* This study was partially supported by the National Key R&D Program of China via project 2023YFB4706302 and by the National Natural Science Foundation of China via projects 62471293 and U20A20199.

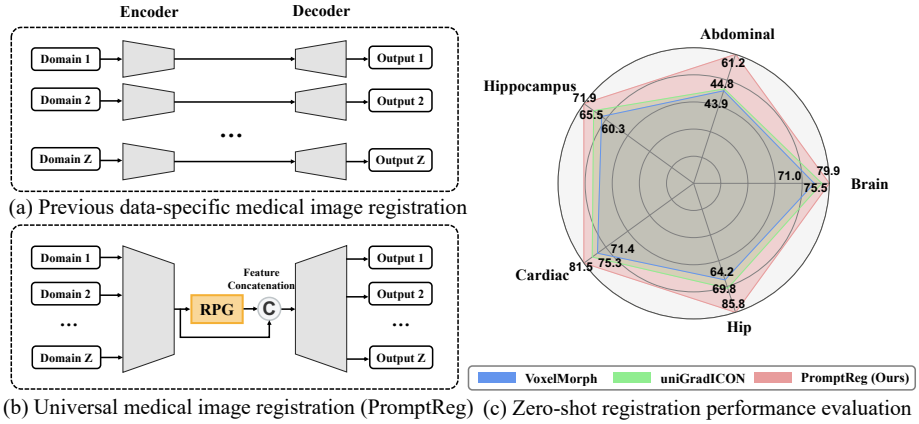

(a) Previous data-specific medical image registration

(b) Universal medical image registration (PromptReg) (c) Zero-shot registration performance evaluation

Fig. 1. (a) and (b) show differences between conventional image registration method on specific datasets and our method. (c) presents the DSC (%) performance of various methods on zero-shot task.

tives to traditional iterative optimization methods [1, 11, 13, 9]. However, their efficacy is inherently limited by the assumption of distributional consistency between training and testing datasets. Although recent technologies like uniGradICON [16] facilitate training a single model on multiple data distributions, the inherent heterogeneity of medical imaging data makes it practically impossible to achieve comprehensive coverage of all possible distributions during the training phase. This raises two critical challenges: (1) how to effectively capture and model the known domain distributions in the training data and (2) how to achieve robust zero-shot generalization on unseen domains through effective transfer of the knowledge acquired from known domains.

As a foundation image registration model, uniGradICON exemplifies cross-dataset registration capabilities by implementing gradient inverse consistency (GradICON [17]) regularization across datasets. While showing promising capabilities in universal registration, the registration performance of the model still needs improvement. Of particular note is its limited performance in zero-shot registration tasks. These limitations arise from its architectural design, which incorporates task-specific optimization structures from existing approaches [5, 14, 4, 7, 21], rather than introducing architectural components specifically designed for universal image registration purposes.

In this work, we propose a universal image registration framework that preserves the advantages of existing popular architectures while leveraging prompt learning to guide model understanding of diverse registration scenarios through explicit task prompts. As shown in Fig. 1 (a) and (b), unlike traditional registration networks that rely on models with different weights trained for different datasets to complete various registration tasks, our method facilitates the flexible integration of explicit domain knowledge and task-specific cues by our proposed registration prompt generator (RPG), enabling robust zero-shot

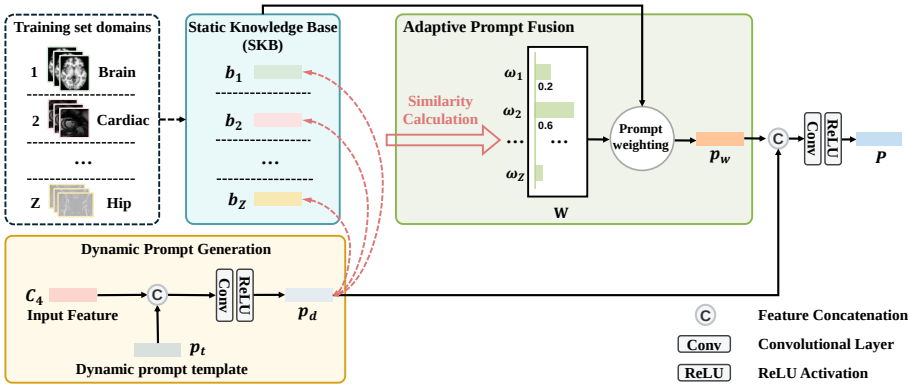

Fig. 2. Overview of the proposed Registration Prompt Generator (RPG). The input feature C_4 are utilized to generate dynamic prompts p_d through a dynamic template p_t. The p_d is compared for similarity with each prompt in the SKB and the final prompt P is derived by combining weighted SKB prompts with the dynamic prompt p_d.

generalization (shown in Fig. 1 (c)). We demonstrate that appropriate domain guidance enhances model generalization across different data distributions while maintaining architectural simplicity. Specifically, we first introduce a learnable Static Knowledge Base (SKB) that encodes domain-specific expertise from various known data sources. Subsequently, we define a dynamic prompt generation mechanism that projects different inputs into a shared prompt vector space, enabling adaptive prompt generation based on input domain characteristics. To further enhance zero-shot performance, we propose an adaptive prompt fusion strategy that rebalances the stored domain knowledge in SKB based on the similarity between the generated dynamic prompts and the prompts in SKB, generating transferable prompt knowledge for unseen domains. Finally, the optimization of the prompt generator is reinforced via domain orthogonality loss and task similarity loss, which jointly act as regularization constraints to promote effective learning.

Our contributions are twofold: (1) We introduce prompt learning to medical image registration, guiding the understanding of various registration tasks with explicit task prompts for the registration model. (2) Our proposed registration prompt generator comprises a learnable static knowledge base, a dynamic prompt generation mechanism, and an adaptive prompt fusion strategy, enabling the generation of domain-specific prompts for inputs from diverse domains.

2 Methods

Given paired moving ($I_m : \Omega_m \to \Re^3$) and fixed ($I_f : \Omega_f \to \Re^3$) images from an arbitrary domain $\mathcal{D}_i$, we use RDP [18] as the encoder and decoder in our framework, which extracts multi-scale hierarchical features M_l and F_l ($l = 1, 2, 3, 4$) from the input images, respectively. Subsequently, the features from stage-4 (i.e.,

M_4 and F_4) are concatenated along the channel dimension and fused through a convolution layer with ReLU activation to generate a context-rich feature representation C_4. This representation is then fed into our Registration Prompt Generator (RPG), which outputs a registration-adaptive task prompt P. Finally, the concatenated features $[C_4 \copyright P]$ are used to replace the original C_4 and input into the progressive decoder, which predicts the dense deformation field Φ. Below, we provide a detailed description of the proposed RPG as shown in Fig. 2.

2.1 Static Knowledge Base

We formally define the Static Knowledge Base (SKB) as a collection of learnable random initialization prompts $\mathbf{B} \in \mathbb{R}^{Z\times C\times H\times W\times D} = \{b_1, b_2, \ldots, b_Z\}$, where Z corresponds to the number of unique registration tasks present in the training set, C denotes the channel dimension of the context-rich feature representation C_4, and H, W, D represent the spatial dimensions of C_4. To facilitate discriminative domain-specific knowledge learning, we require each prompt in SKB to maintain its unique task characteristics by introducing an orthogonality constraint:

$$\mathcal{L}_{\text{ort}}\left(\mathbf{B}\right) = \left\| \mathbf{B}\mathbf{B}^T - \text{diag}\left(\mathbf{B}\mathbf{B}^T\right) \right\|_2, \tag{1}$$

where $\text{diag}(\cdot)$ denotes the operator that preserves only the diagonal entries, and $\|\cdot\|_2$ represents the L2-norm.

2.2 Dynamic prompt generation mechanism

Relying solely on the SKB is insufficient, due to its inherent limitation to the training distribution. In practice, during inference, unseen registration tasks lack accessible domain-specific cues from the SKB. To address this limitation, we develop a dynamic prompt generation mechanism that projects different inputs into prompt vector spaces, achieving input-adaptive feature modulation without depending on predefined domain distributions. To achieve this, we randomly initialize a learnable dynamic prompt template, denoted as p_t. Given input features C_4, we concatenate C_4 with p_t and feed $[C_4 \copyright p_t]$ into convolutional blocks for fusion, generating dynamic prompt p_d.

2.3 Adaptive Prompt Fusion Strategy

To leverage both the domain knowledge encoded in SKB and the dynamic input-specific information captured by p_d, we propose an adaptive prompt fusion strategy that seamlessly integrates these prompt representations. Considering the rich domain knowledge encoded in the SKB, we recombine this information to generate domain-specific guidance for both known and unseen domains. Specifically, for each prompt vector $b_i \in \mathbf{B}$, we compute its similarity with the dynamic prompt p_d using cosine similarity and apply softmax function to obtain adaptive weights $\mathbf{W} = \{\omega_1, \omega_2, \ldots, \omega_Z\}$:

$$\text{sim}(p_d, b_i) = \frac{\langle p_d, b_i \rangle}{\|p_d\| \cdot \|b_i\|}, \tag{2}$$

$$\omega_i = \frac{\exp(\text{sim}(p_d, b_i)/\tau)}{\sum_{k=1}^{Z} \exp(\text{sim}(p_d, b_k)/\tau)}, \tag{3}$$

where τ is the temperature parameter. Subsequently, a weighted aggregation of the prompts in SKB is computed as follows:

$$p_w = \sum_{i=1}^{Z} \omega_i \cdot b_i, \tag{4}$$

Then, we concatenate the dynamic prompt p_d with p_w to generate the concatenated feature $[p_d \copyright p_w]$. The concatenated feature is passed through a convolution block for fusion and to generate the final task prompt P.

2.4 Loss function

During the training phase, the loss function consists of four parts. The image similarity loss L_{sim} is calculated using the MIND loss [8], the smoothness loss L_{smooth} is implemented using the method from VoxelMorph [2], the domain orthogonality loss L_{ort} is introduced in Sec. 2.1, and the task similarity loss L_{ts} is defined by computing the similarity between the prompt b_i for the current task in the SKB and the weighted prompt p_w, which can be expressed as:

$$L_{ts} = \left(1 - \frac{\langle p_w, b_i \rangle}{\|p_w\| \cdot \|b_i\|}\right)^2, \tag{5}$$

The overall loss can be expressed as:

$$L = L_{sim} + \lambda_1 L_{smooth} + \lambda_2 L_{ort} + \lambda_3 L_{ts}. \tag{6}$$

where λ_1, λ_2 and λ_3 are weighting coefficients that balance the contribution of each loss term.

3 Experiments

3.1 Experimental Setup

Datasets. Our experiments are conducted on five datasets representing a diverse range of anatomical structures: (1) the OASIS dataset [12] (Brain), from the Learn2Reg 2021 Challenge, consists of 289 and 125 MRI scans for training and testing, respectively. (2) the Abdominal dataset (Abdominal), from the Learn2Reg 2021 Challenge, consists of 38 and 16 CT scans for training and testing, respectively. (3) the Hippocampus dataset [15] (Hippocampus), consists of 208 and 52 MRI scans for training and testing, respectively. (4) the M&Ms

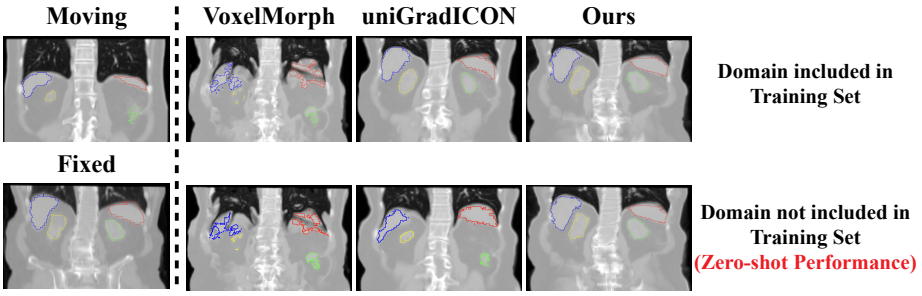

Fig. 3. Qualitative comparison of different methods on the Abdominal registration task. Each comparison method includes two settings: when the domain of the displayed image is included in the training set (top) and when the domain of the displayed image is not included in the training set (zero-shot scenario) (bottom).

dataset [3] (Cardiac), consists of 224 and 96 brain MRI scans for training and testing, respectively. (5) the in-house collected Hip dataset (Hip), consists of 45 and 16 CT scans for training and testing, respectively. These datasets contain segmentation labels for 35, 4, 2, 3, and 3 types of organs respectively, which are used to evaluate the performance of different registration methods. After pairing the images, the number of training/testing image pairs are as follows: 578/250 (Brain), 703/120 (Abdominal), 624/156 (Hippocampus), 224/96 (Cardiac), and 1980/132 (Hip).

Data Preprocessing. The CT images are clipped to Hounsfield Units (HU) in the range [-1000, 1000] and normalized to [0, 1]. For MRI data, intensity values are clipped at the 99th percentile prior to normalization. Similar to the implementation of uniGradICON, all volumes are resampled to [160, 160, 160] using trilinear interpolation for consistent dimensions across datasets. The output deformation fields are interpolated back to their original resolution for evaluation.

Implementation Details. During training, we randomly sample 100 image pairs per dataset per epoch, with a batch size of 1, training for 300 epochs in total. Set hyperparameters to τ=0.1, λ_1=0.1, λ_2=0.001, λ_3=0.001. The initial learning rate is set to 1.0E-4 and decays at each iteration following a polynomial learning rate strategy with a power of 0.9. We compare our method against two established baselines: VoxelMorph and uniGradICON. VoxelMorph is trained using the same image similarity loss, smoothness loss and hyperparameter configuration as our method. For uniGradICON, we adhere to its two-stage training scheme and loss function, executing 300 epochs for the initial stage followed by an additional 200 epochs for the subsequent refinement stage. In addition to using all datasets for training, to evaluate generalization performance, we test models under five configurations, each excluding one dataset during training to implement zero-shot registration scenarios. All quantitative experiments use the

Table 1. Quantitative comparison of DSC (%) results for different registration methods on test sets of different organ types under various training set types. The results in the zero-shot task are highlighted with light blue color.

Training Set Types	Methods	DSC (%) in different datasets				
		Brain	Abdominal	Hippocampus	Cardiac	Hip
w/o Brain	VoxelMorph	71.0 ± 0.2	45.3 ± 2.5	62.2 ± 1.9	73.3 ± 0.8	67.8 ± 1.3
	uniGradICON	75.5 ± 0.1	55.2 ± 1.4	72.6 ± 1.1	81.5 ± 0.5	85.8 ± 0.8
	Ours	79.9 ± 0.1	66.0 ± 2.2	75.7 ± 0.8	84.6 ± 0.3	92.6 ± 0.3
w/o Abdominal	VoxelMorph	73.5 ± 0.2	43.9 ± 2.4	62.0 ± 2.0	73.5 ± 0.7	66.2 ± 1.4
	uniGradICON	78.1 ± 0.1	44.8 ± 1.9	72.1 ± 1.3	80.9 ± 0.5	81.1 ± 1.1
	Ours	80.5 ± 0.1	61.2 ± 2.2	75.1 ± 0.8	84.2 ± 0.3	92.5 ± 0.2
w/o Hippocampus	VoxelMorph	73.0 ± 0.2	43.5 ± 2.3	60.3 ± 2.0	73.5 ± 0.8	65.0 ± 1.4
	uniGradICON	77.7 ± 0.1	53.7 ± 1.3	65.5 ± 2.1	81.0 ± 0.5	83.7 ± 0.8
	Ours	81.0 ± 0.1	70.1 ± 1.6	71.9 ± 1.2	84.8 ± 0.3	93.3 ± 0.2
w/o Cardiac	VoxelMorph	72.0 ± 0.2	44.6 ± 2.4	60.9 ± 1.9	71.4 ± 0.9	67.9 ± 1.3
	uniGradICON	77.5 ± 0.1	51.8 ± 1.4	72.2 ± 1.3	75.3 ± 0.8	82.4 ± 1.1
	Ours	81.2 ± 0.1	66.3 ± 2.0	75.1 ± 0.9	81.5 ± 0.5	91.7 ± 0.4
w/o Hip	VoxelMorph	72.9 ± 0.2	43.6 ± 2.3	61.3 ± 2.0	73.0 ± 0.7	64.2 ± 1.4
	uniGradICON	78.1 ± 0.1	51.4 ± 1.3	72.0 ± 1.3	80.7 ± 0.5	69.8 ± 1.7
	Ours	81.4 ± 0.1	67.3 ± 2.0	75.4 ± 1.0	84.9 ± 0.3	85.8 ± 1.1
All domains	VoxelMorph	73.1 ± 0.2	43.6 ± 2.3	60.5 ± 2.0	73.8 ± 0.7	65.6 ± 1.4
	uniGradICON	77.8 ± 0.1	55.3 ± 1.4	72.1 ± 1.4	82.0 ± 0.4	85.7 ± 0.7
	Ours	81.3 ± 0.1	68.7 ± 2.0	74.8 ± 1.0	84.7 ± 0.3	93.4 ± 0.2

average Dice Similarity Coefficient (DSC) metric to quantify the volume overlap between the warped segmentation mask and the fixed image segmentation mask.

3.2 Experimental Results and Discussion

Comparison Study. Table. 1 shows the DSC evaluation results obtained from various registration methods across six different experimental configurations. These configurations include scenarios where specific datasets are excluded from the training process (for example, w/o Brain indicates training without the Brain training dataset). The results show that while VoxelMorph can be trained in a multi-dataset environment, its accuracy is limited due to its simplistic architecture, which struggles with complex deformations. In contrast, uniGradICON uses a hierarchical network that captures complex nonlinear deformations, improving performance but still falling short of ideal results. Our method outperforms others by incorporating prior knowledge through prompts. The zero-shot generalization ability of universal registration models is crucial for cross-task registration. Compared to other methods, our method achieves the best zero-shot generalization performance, underscoring the efficacy of our task prompt in improving model adaptability across diverse cross-task settings. Fig. 3 presents a qualitative comparison of different methods on the Abdominal dataset, showing warped labels overlaid on the warped images. When the testing domain is included in the training set, the training configuration follows the 'All domains' settings in Table. 1. Meanwhile, for the domain not included in the training

Table 2. Ablation study results of DSC (%) on test sets for different organ types. The results in the zero-shot task are highlighted with light blue color.

Training Set Types	Methods	DSC (%) in different datasets				
		Brain	Abdominal	Hippocampus	Cardiac	Hip
w/o Abdominal	w/o RPG	79.7 ± 0.1	56.2 ± 3.0	74.6 ± 1.0	83.8 ± 0.3	92.2 ± 0.3
	w/o R. Cons.	80.6 ± 0.1	57.9 ± 2.7	75.1 ± 1.0	84.3 ± 0.3	92.2 ± 0.3
	Ours	80.5 ± 0.1	61.2 ± 2.2	75.1 ± 0.8	84.2 ± 0.3	92.5 ± 0.2
w/o Hippocampus	w/o RPG	78.9 ± 0.1	63.6 ± 1.8	69.8 ± 1.3	84.1 ± 0.3	90.0 ± 0.5
	w/o R. Cons.	81.1 ± 0.1	68.7 ± 1.6	71.4 ± 1.3	84.7 ± 0.3	93.0 ± 0.2
	Ours	81.0 ± 0.1	70.1 ± 1.6	71.9 ± 1.2	84.8 ± 0.3	93.3 ± 0.2
All domains	w/o RPG	78.5 ± 0.1	63.4 ± 2.3	73.6 ± 1.1	83.8 ± 0.3	90.9 ± 0.4
	w/o R. Cons.	80.1 ± 0.1	64.7 ± 2.1	74.4 ± 1.0	84.4 ± 0.3	92.5 ± 0.2
	Ours	81.3 ± 0.1	68.7 ± 2.0	74.8 ± 1.0	84.7 ± 0.3	93.4 ± 0.2

set (shown at the bottom of Fig. 3), the training configuration follows the 'w/o Abdominal' settings.

Ablation Study. Table. 2 presents the results of ablation studies on three selected training sets to evaluate the contribution of different components. The results reveal that removing the prompt generator (w/o RPG), leaving only the naive registration network (vanilla RDP), leads to substantial performance degradation across all datasets in all training set types. This is due to the absence of task-specific prompts, which are essential for capturing underlying distributions across multiple datasets. Moreover, this variant has limited zero-shot generalization. When we only remove the regularization constraints (w/o R. Cons.) consisting of domain orthogonality loss L_{ort} and task similarity loss L_{ts}, the model still underperforms compared to the complete framework. This is because without the guidance of regularization constraints, the model must adaptively adjust the prompt information in SKB and generate task-specific prompts to enhance registration accuracy across different datasets. These results validate the effectiveness of our regularization design in guiding the network to better differentiate between tasks.

4 Conclusion

In this work, we propose PromptReg, a universal medical image registration framework based on prompt learning, aimed at guiding the model to adapt to different registration scenarios through explicit task prompts. At the core of our approach is the Registration Prompt Generator (RPG), which stores prompt information for known domains in a Static Knowledge Base (SKB) and projects different inputs into a shared prompt vector space through a dynamic prompt generation mechanism. Additionally, we introduce an adaptive prompt fusion strategy, effectively leveraging stored domain knowledge to generate transferable prompt information for unseen domains. Furthermore, we incorporate domain

orthogonality loss and task similarity loss into the optimization process. Extensive experimental results demonstrate that PromptReg achieves competitive performance in cross-dataset registration tasks and exhibits robust capabilities in zero-shot generalization, showcasing promising universal image registration performance.

Disclosure of Interests. The authors declare no conflict of interest.

References

1. 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)
2. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. IEEE transactions on medical imaging **38**(8), 1788–1800 (2019)
3. Campello, V.M., Gkontra, P., Izquierdo, C., Martin-Isla, C., Sojoudi, A., Full, P.M., Maier-Hein, K., Zhang, Y., He, Z., Ma, J., et al.: Multi-centre, multi-vendor and multi-disease cardiac segmentation: the m&ms challenge. IEEE Transactions on Medical Imaging **40**(12), 3543–3554 (2021)
4. Cao, Y., Cao, W., Wang, Z., Yuan, G., Li, Z., Ni, X., Zheng, J.: A light-weight rectangular decomposition large kernel convolution network for deformable medical image registration. Biomedical Signal Processing and Control **95**, 106476 (2024)
5. Chang, Y., Li, Z., Yang, N.: Corrmorph: Unsupervised deformable brain mri registration based on correlation mining. IEEE Journal of Biomedical and Health Informatics (2024)
6. Chen, J., Frey, E.C., He, Y., Segars, W.P., Li, Y., Du, Y.: Transmorph: Transformer for unsupervised medical image registration. Medical image analysis **82**, 102615 (2022)
7. Hasan, M.K., Luo, Y., Yang, G., Yap, C.H.: Feedback attention to enhance unsupervised deep learning image registration in 3d echocardiography. IEEE Transactions on Medical Imaging (2025)
8. Heinrich, M.P., Jenkinson, M., Bhushan, M., Matin, T., Gleeson, F.V., Brady, M., Schnabel, J.A.: Mind: Modality independent neighbourhood descriptor for multimodal deformable registration. Medical image analysis **16**(7), 1423–1435 (2012)
9. Heinrich, M.P., Jenkinson, M., Brady, S.M., Schnabel, J.A.: Globally optimal deformable registration on a minimum spanning tree using dense displacement sampling. In: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2012: 15th International Conference, Nice, France, October 1-5, 2012, Proceedings, Part III 15. pp. 115–122. Springer (2012)
10. Khor, H.G., Ning, G., Sun, Y., Lu, X., Zhang, X., Liao, H.: Anatomically constrained and attention-guided deep feature fusion for joint segmentation and deformable medical image registration. Medical Image Analysis **88**, 102811 (2023)
11. Klein, S., Staring, M., Murphy, K., Viergever, M.A., Pluim, J.P.: Elastix: a toolbox for intensity-based medical image registration. IEEE transactions on medical imaging **29**(1), 196–205 (2009)
12. 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**(9), 1498–1507 (2007)

13. Modat, M., Ridgway, G.R., Taylor, Z.A., Lehmann, M., Barnes, J., Hawkes, D.J., Fox, N.C., Ourselin, S.: Fast free-form deformation using graphics processing units. Computer methods and programs in biomedicine **98**(3), 278–284 (2010)
14. Pham, X.L., Luu, M.H., van Walsum, T., Mai, H.S., Klein, S., Le, N.H., Chu, D.T.: Cman: Cascaded multi-scale spatial channel attention-guided network for large 3d deformable registration of liver ct images. Medical Image Analysis **96**, 103212 (2024)
15. Simpson, A.L., Antonelli, M., Bakas, S., Bilello, M., Farahani, K., Van Ginneken, B., Kopp-Schneider, A., Landman, B.A., Litjens, G., Menze, B., et al.: A large annotated medical image dataset for the development and evaluation of segmentation algorithms. arXiv preprint arXiv:1902.09063 (2019)
16. Tian, L., Greer, H., Kwitt, R., Vialard, F.X., San José Estépar, R., Bouix, S., Rushmore, R., Niethammer, M.: unigradicon: A foundation model for medical image registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 749–760. Springer (2024)
17. Tian, L., Greer, H., Vialard, F.X., Kwitt, R., Estépar, R.S.J., Rushmore, R.J., Makris, N., Bouix, S., Niethammer, M.: Gradicon: Approximate diffeomorphisms via gradient inverse consistency. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18084–18094 (2023)
18. Wang, H., Ni, D., Wang, Y.: Recursive deformable pyramid network for unsupervised medical image registration. IEEE Transactions on Medical Imaging (2024)
19. Wu, N., Xing, J., Zhang, M.: Tlrn: Temporal latent residual networks for large deformation image registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 728–738. Springer (2024)
20. Xin, Y., Chen, Y., Ji, S., Han, K., Xie, X.: On-the-fly guidance training for medical image registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 694–705. Springer (2024)
21. Zhu, Z., Song, R., Wei, Y.: Symmetric deformable image pairwise registration by optimizing neural fields network architectures for guaranteeing transformations consistency. Biomedical Signal Processing and Control **104**, 107453 (2025)