

Mono-Modalizing Extremely Heterogeneous Multi-Modal Medical Image Registration

Kyobin Choo¹, Hyunkyung Han², Jinyeong Kim³, Chanyong Yoon¹, and Seong Jae Hwang^{2*}

¹ Department of Computer Science, Yonsei University, Seoul, Republic of Korea

² Department of Artificial Intelligence, Yonsei University, Seoul, Republic of Korea

³ Yonsei University College of Medicine, Seoul, Republic of Korea
{chu, hhk, jinyeong1324, charlie2019370, seongjae}@yonsei.ac.kr

Abstract. In clinical practice, imaging modalities with functional characteristics, such as positron emission tomography (PET) and fractional anisotropy (FA), are often aligned with a structural reference (e.g., MRI, CT) for accurate interpretation or group analysis, necessitating multi-modal deformable image registration (DIR). However, due to the extreme heterogeneity of these modalities compared to standard structural scans, conventional unsupervised DIR methods struggle to learn reliable spatial mappings and often distort images. We find that the similarity metrics guiding these models fail to capture alignment between highly disparate modalities. To address this, we propose M2M-Reg (Multi-to-Mono Registration), a novel framework that trains multi-modal DIR models using only mono-modal similarity while preserving the established architectural paradigm for seamless integration into existing models. We also introduce GradCyCon, a regularizer that leverages M2M-Reg’s cyclic training scheme to promote diffeomorphism. Furthermore, our framework naturally extends to a semi-supervised setting, integrating pre-aligned and unaligned pairs only, without requiring ground-truth transformations or segmentation masks. Experiments on the Alzheimer’s Disease Neuroimaging Initiative (ADNI) dataset demonstrate that M2M-Reg achieves up to $2\times$ higher DSC than prior methods for PET-MRI and FA-MRI registration, highlighting its effectiveness in handling highly heterogeneous multi-modal DIR. Our code is available at <https://github.com/MICV-yonsei/M2M-Reg>.

Keywords: Unsupervised Deformable Medical Image Registration · Multi-Modality · Heterogeneity · Semi-Supervised Learning

1 Introduction

Medical imaging encompasses diverse modalities, each capturing distinct physiological and anatomical information. In neuroimaging, modalities such as positron emission tomography (PET) and fractional anisotropy (FA) maps exhibit significant heterogeneity from structural scans due to their specialized purposes. Thus,

* Corresponding author




Fig. 1. Discrepancy between image similarity and registration quality in extremely heterogeneous multi-modal image registration. (a) The metric values in the graph show the test set average during training. Our M2M-Reg (blue line) achieves high DSC despite low LNCC, whereas the baseline (yellow line) shows the opposite trend. This suggests that the multi-modal similarity metric does not capture registration quality in highly disparate modalities. (b) This discrepant relationship between the two metrics indicates that the baseline fails to establish meaningful PET-MRI correspondences, resulting in severely distorted warped PET images.

practical scenarios often require aligning these scans to a structural reference or spatially normalizing them to a standard MRI space for accurate interpretation and quantitative analysis [30, 5]. Such cases involve complex, spatially varying deformations, making *multi-modal deformable image registration* (DIR)—which aligns images of different modalities through non-rigid warping—indispensable.

In general, multi-modal DIR is performed in an unsupervised manner using unaligned image pairs, since obtaining ground-truth transformations that map the moving image to the fixed image is rarely feasible [4]. In this setup, the estimated transformation is iteratively refined to maximize the similarity between the fixed and warped moving images. Therefore, the performance of unsupervised DIR hinges heavily on the quality of the similarity metric for measuring image alignment. In multi-modal settings, designing an effective similarity metric is inherently challenging, as images may exhibit different intensity distributions for the same tissue type or even structural variations. To address this, various approaches have been explored, including intensity-based [17, 6], statistical [19, 25], self-similarity [12, 13], and deep representation-based [23, 27] metrics.

While these similarity metrics work well for aligning structural scans such as CT and MRI [6, 23], we find that they struggle to establish meaningful spatial correspondences when images exhibit extreme heterogeneity. Concretely, as shown in Fig. 1a, the registration network exhibits counterintuitive behavior in training dynamics: while the average similarity (LNCC) between the warped PET and MRI increases as the training progresses, the registration quality (DSC) becomes even worse than the initial state. This observation indicates that multi-modal similarity metrics do not effectively capture the complex relationship between PET and MRI, especially in regions with intricate and highly detailed anatomical structures, such as the brain (Fig. 1b).

In light of these challenges, we propose **M2M-Reg**, an innovative framework that enables unsupervised training of multi-modal DIR models using only mono-modal similarity, thereby bypassing the limitations of multi-modal similarity metrics. Our approach trains the model for multiple cross-modal registrations, forming a cyclic source-target relationship. Specifically, during each training iteration, an additional *bridge pair* is sampled to establish a cyclic source-target mapping in which modalities alternate. By leveraging this cyclic structure, the multi-modal DIR model can be guided by mono-modal similarity as its objective, indirectly learning a more reliable multi-modal mapping. Additionally, we propose the *Gradient Cycle Consistency* (GradCyCon) regularizer, which encourages diffeomorphism on the warpings within the cycle. GradCyCon is efficiently implemented by penalizing the Jacobian of the cyclic mappings to enforce identity, as these mappings are already computed for mono-modal cycle similarity.

Remarkably, M2M-Reg further enables a new *semi-supervised* DIR paradigm with pre-aligned pairs. By simply designating available pre-aligned pairs as bridge pairs, we can incorporate prior knowledge of the true cross-modal mapping into the model, achieving performance beyond that of unsupervised methods. To the best of our knowledge, this is the first approach to jointly utilize pre-aligned and unaligned pairs within a single framework, without requiring ground-truth transformations or segmentation masks.

Contributions. Our main contributions are as follows: **(i)** We propose M2M-Reg, a novel framework that seamlessly reformulates the multi-modal DIR problem as a mono-modal one while preserving the existing architectural paradigm. **(ii)** We introduce GradCyCon, a regularization method that naturally leverages the cyclic mapping process in M2M-Reg to promote diffeomorphism. **(iii)** We further extend M2M-Reg to a pioneering semi-supervised DIR, which encompasses pre-aligned pairs alongside unaligned pairs. **(iv)** We highlight the limitations of multi-modal similarity in aligning highly heterogeneous modalities and demonstrate that M2M-Reg addresses them, achieving up to $2\times$ higher DSC in brain PET-MRI and FA-MRI registration on ADNI datasets, compared to previous state-of-the-art (SOTA) methods.

2 Methods

2.1 Preliminaries: Unsupervised Learning-Based Paradigm

Let source $S : \Omega \rightarrow \mathbb{R}$ and target $T : \Omega \rightarrow \mathbb{R}$ be grayscale 3D volumes of different modalities, where $\Omega \subset \mathbb{R}^3$ represents the set of spatial coordinates. A registration network f_θ , parameterized by θ , is trained to predict a deformation field $u_\theta^{ST} : \Omega \rightarrow \mathbb{R}^3$, which represents the displacement vectors mapping coordinates in T to their corresponding locations in S . For arbitrary coordinate $x \in \Omega$, the spatial transformation is then defined as $\Phi_\theta^{ST}(x) = x + u_\theta^{ST}(x)$, where $\Phi_\theta^{ST} : \Omega \rightarrow \mathbb{R}^3$ is the transformation function that warps S into alignment with T . The deformed moving image is obtained as $S \circ \Phi_\theta^{ST}$. Under this setting, the network parameters θ are optimized in an unsupervised manner using the following loss function:

$$\mathcal{L}_\theta = \mathcal{L}_{\text{sim}}(S \circ \Phi_\theta^{ST}, T) + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}(u_\theta^{ST}), \quad (1)$$




Fig. 2. Overview of M2M-Reg training. (a) While the conventional unsupervised multi-modal DIR approach samples a single source-target pair (S, T) , M2M-Reg samples an additional *bridge pair* (S', T') to form a cyclic structure. For semi-supervised DIR, the *bridge pair* can be designated as a pre-aligned pair. (b) The cyclic structure enables multi-modal DIR models to be guided by a mono-modal similarity objective $\mathcal{L}_{m2m}$ (Sec. 2.2). (c) For diffeomorphism, the GradCyCon regularizer ($\mathcal{L}_{gcc}$) encourages the Jacobian (∇) of the total mapping to approximate identity (Sec. 2.3).

where $\mathcal{L}_{sim}(\cdot, \cdot)$ is a multi-modal similarity metric such as local normalized cross-correlation (LNCC) [6] or mutual information (MI) [2], $\mathcal{L}_{reg}(\cdot)$ is a regularization term for enforcing diffeomorphism or promoting smoothness of the deformation field, such as Diffusion or Bending Energy [26], and λ_{reg} is a hyperparameter.

2.2 Learning Multi-Modal Registration via Mono-Modal Similarity

The conventional unsupervised multi-modal DIR approach in Sec. 2.1 [6, 21, 3] struggles to establish spatial correspondences in highly heterogeneous modalities (Fig. 1). Here, we introduce an alternative approach that effectively guides multi-modal DIR models using only mono-modal similarity comparisons.

We represent the mappings performed by the multi-modal DIR model f_θ within a single training iteration as a *graph*. As shown in Fig. 2a (left), the conventional approach samples a single source-target pair (S, T) , where f_θ predicts a directed edge (i.e., mapping) Φ_θ^{ST} between the two nodes (i.e., images). Since f_θ is inherently a multi-modal model, it only computes edges between nodes of different modalities. To incorporate mono-modal similarity into the learning process, nodes of the same modality must be connected. The most straightforward approach is to add a reverse edge Φ_θ^{TS} , allowing similarity comparison between S and $S \circ \Phi_\theta^{ST} \circ \Phi_\theta^{TS}$, as well as between T and $T \circ \Phi_\theta^{TS} \circ \Phi_\theta^{ST}$. However, this strategy risks the model converging to a trivial solution, such as an identity mapping.

Instead, we introduce a cyclic structure by independently drawing an additional *bridge pair* (S', T') from the data during each training iteration, adding

two nodes and predicting three additional cross-modal edges (Fig. 2a, right). This cyclic design provides several advantages: (i) It maximizes the number of mono-modal similarity comparisons relative to the number of computed edges, improving efficiency. (ii) The number of similarity comparisons per modality remains consistent across iterations, encouraging stable gradient updates. (iii) The cyclic paths allow composed mappings to approximate identity transformations, naturally enforcing diffeomorphic behavior (Sec. 2.3). As a result, two mono-modal paths per modality are formed, as illustrated in Fig. 2b. Leveraging these connections, we formulate an objective purely based on mono-modal similarity, guiding the multi-modal DIR process:

$$\begin{aligned} \mathcal{L}_{\text{m2m}} = & \mathcal{L}_{\text{sim}}(S \circ \Phi_{\theta}^{ST} \circ \Phi_{\theta}^{TS'}, S') + \mathcal{L}_{\text{sim}}(S' \circ \Phi_{\theta}^{S'T'} \circ \Phi_{\theta}^{T'S}, S) \\ & + \mathcal{L}_{\text{sim}}(T \circ \Phi_{\theta}^{TS'} \circ \Phi_{\theta}^{S'T'}, T') + \mathcal{L}_{\text{sim}}(T' \circ \Phi_{\theta}^{T'S} \circ \Phi_{\theta}^{ST}, T), \end{aligned} \quad (2)$$

where $\mathcal{L}_{\text{sim}}(\cdot, \cdot)$ can be any mono-modal or multi-modal similarity metric. In our case, we use LNCC, as it has been widely adopted in recent studies [6, 21, 18].

2.3 Gradient Cycle Consistency

The cyclic source-target topology in Sec. 2.2 naturally provides a condition conducive to enforcing diffeomorphic mappings, as the composition of mappings within the cycle should approximate an identity transformation (Fig. 2c). Building on this, inspired by GradICON [29], we propose the GradCyCon regularizer, which encourages the Jacobian (∇) of the total mapping to be identity ($\mathbf{I}$):

$$\mathcal{L}_{\text{gcc}} = \left\| \nabla \left[\Phi_{\theta}^{ST} \circ \Phi_{\theta}^{TS'} \circ \Phi_{\theta}^{S'T'} \circ \Phi_{\theta}^{T'S} \right] - \mathbf{I} \right\|_F^2. \quad (3)$$

We now briefly outline the theoretical intuition behind how GradCyCon, like GradICON, implicitly imposes an H^1 type regularization effect. Recall that GradICON regularizer enforces diffeomorphism by additionally computing an inverse mapping for a given source-target pair (S, T) , as follows:

$$\mathcal{L}_{\text{reg}}^{\text{GradICON}} = \left\| \nabla \left[\Phi_{\theta}^{ST} \circ \Phi_{\theta}^{TS} \right] - \mathbf{I} \right\|_F^2. \quad (4)$$

Since the neural network f_{θ} inherently produces *independent* noise n with a small ε (i.e., $\Phi_{\theta} = \Phi + \varepsilon n$), Eq. (4) is approximated as inducing L^2 regularization on the inverse of $\nabla \Phi^{ST}$ and $\nabla \Phi^{TS}$, thereby exhibiting an H^1 -like effect [20, 10, 29]:

$$\mathbb{E}[\mathcal{L}_{\text{reg}}^{\text{GradICON}}] \approx \varepsilon^2 \left(\left\| \left[\nabla \Phi^{ST} \right]^{-1} \sqrt{\text{Det}(\nabla \Phi^{ST})} \right\|_F^2 + \left\| \left[\nabla \Phi^{TS} \right]^{-1} \right\|_F^2 \right). \quad (5)$$

In GradCyCon (Eq. (3)), grouping three consecutive mappings into one— Φ_{θ}^{ST} and $\hat{\Phi}_{\theta}^{TS} = \Phi_{\theta}^{TS'} \circ \Phi_{\theta}^{S'T'} \circ \Phi_{\theta}^{T'S}$, or $\hat{\Phi}_{\theta}^{ST'} = \Phi_{\theta}^{ST} \circ \Phi_{\theta}^{TS'} \circ \Phi_{\theta}^{S'T'}$ and $\Phi_{\theta}^{T'S}$ —yields the same formulation as Eq. (4). By applying a first-order Taylor approximation, we can express $\hat{\Phi}_{\theta}^{TS} \approx \Phi^{TS} + \varepsilon \hat{n}^{TS} + O(\varepsilon^2)$. Then, $\hat{n}^{TS}$, formed by three individual ground-truth mappings and their noise components, can be considered




Fig. 3. Qualitative comparison of PET-MRI and FA-MRI registration

independent of n^{ST} , and terms of ε^2 and higher are negligible, as assumed in the GradICON. Consequently, Eq. (3) can be derived in the same form as Eq. (5), promoting both Φ_θ^{ST} and $\hat{\Phi}_\theta^{TS}$, as well as $\hat{\Phi}_\theta^{S'T'}$ and $\Phi_\theta^{T'S}$, to be diffeomorphic.

Meanwhile, although the intermediate mappings within the cycle, $\Phi_\theta^{TS'}$ and $\Phi_\theta^{S'T'}$, do not appear to be individually regularized in Eq. (3), they naturally undergo regularization since target and bridge pairs are randomly resampled during training. From a broader training perspective, GradCyCon eventually enforces diffeomorphism for every mapping. The same holds in the semi-supervised setting (Sec. 2.4), and we empirically observe that GradCyCon performs effectively.

2.4 Semi-Supervised Learning with Pre-Aligned Pairs

We further extend our framework to incorporate scarce or multiple pre-aligned pairs, allowing the model to reference a reliable multi-modal mapping. The proposed semi-supervised approach enables learning with both unaligned and pre-aligned pairs in an unsupervised manner, offering greater flexibility than conventional supervised methods that rely on ground-truth transformations [8, 7] or segmentation masks [28, 11]. As shown at the bottom of Fig. 2a, simply designating the bridge pair (S', T') as a pre-aligned pair naturally integrates prior knowledge of cross-modal correspondence into training. In this setting, since $\Phi_\theta^{S'T'}$ can be regarded as the identity mapping, we omit $\Phi_\theta^{S'T'}$ from the formulations of our mono-modal cycle similarity (Eq. (2)) and GradCyCon regularizer (Eq. (3)), thereby simplifying optimization. As a result, one fewer mapping needs to be computed, enhancing efficiency while further improving performance.

2.5 Final Objective: Integration with Existing Models

The proposed M2M-Reg is a flexible framework that can be seamlessly applied to any unsupervised approach that follows a formulation similar to Eq. (1). It simply replaces the conventional multi-modal similarity with our mono-modal cycle similarity and incorporates GradCyCon, leading to the following final objective:

$$\mathcal{L}_\theta^{\text{final}} = \mathcal{L}_{\text{m2m}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}} + \lambda_{\text{gcc}} \mathcal{L}_{\text{gcc}}, \quad (6)$$

where $\mathcal{L}_{\text{reg}}$ is the original regularization, and λ_{reg} and λ_{gcc} are hyperparameters.

Table 1. Quantitative comparison of PET-MRI and FA-MRI registration

Method	$\mathcal{L}_{reg}$	$\mathcal{L}_{sim}$	M2M-Reg		PET $\rightarrow$ MRI		FA $\rightarrow$ MRI	
			$\mathcal{L}_{m2m}$	$\mathcal{L}_{gcc}$	DSC $\uparrow$	% $ J_\phi _{<0} \downarrow$	DSC $\uparrow$	% $ J_\phi _{<0} \downarrow$
Initial	-	-	-	-	0.462	-	0.424	-
SyN [2]	Gaussian	MI	-	-	0.262	0.000	0.635	0.000
NiftyReg [22]	BE	LNCC	-	-	0.196	0.000	0.607	0.038
FireANTs [16]	Gaussian	MI	-	-	0.223	1.737	0.420	0.094
TransMorph [3]	Diffusion	LNCC			0.322	40.572	0.449	39.534
			$\checkmark$		0.496	8.020	0.510	10.464
			$\checkmark$ $\checkmark$		0.657	0.003	0.641	0.000
CorrMLP [21]	Diffusion	LNCC			0.237	43.149	0.064	50.106
			$\checkmark$		0.583	0.422	0.600	0.003
			$\checkmark$ $\checkmark$		0.587	0.000	0.607	0.000
GradICON [29]	GradICON	LNCC ² [6]			0.088	5.508	0.632	0.063
		LNCC		0.346	0.041	0.647	0.008	
			$\checkmark$	0.668	0.000	0.662	0.000	
			$\checkmark$ $\checkmark$	0.676	0.000	0.666	0.000	

3 Experiments

Datasets. We constructed paired PET-MRI (train: 300, test: 50) and FA-MRI (train: 230, test: 50) datasets using brain 18F-Fluorodeoxyglucose PET, FA, and T1-weighted MRI from the ADNI⁴ [24, 15] study. All data were resampled to 1.5 mm isotropic voxels with a 128 $\times$ 128 $\times$ 128 resolution, and the FA-MRI set was skull-stripped using SynthStrip [14]. Intensities were clipped at the upper 99.9th percentile and min-max normalized to [0,1]. Each pair was rigidly coregistered using SPM12 [1] for alignment. For DSC evaluation, MRI was segmented using FreeSurfer v7.1.1 [9], with the resulting labels grouped into 11 classes.

Comparison Methods. We compare optimization-based methods (SyN [2], NiftyReg [22], FireANTs [16]) and unsupervised learning models, including TransMorph [3] (a widely used transformer-based model), CorrMLP [21] (a SOTA MLP-based architecture), and GradICON [29] (a CNN-based model with a SOTA regularizer), evaluating their performance with and without M2M-Reg.

Implementation Details. M2M-Reg was implemented on the GradICON codebase and trained with an NVIDIA RTX A6000 48GB GPU using ADAM optimizer with a learning rate of 5e-5 and batch size of 8, for 50,000 iterations. The regularization weights were set to $\lambda_{reg} = 0.5$ (1.5 without M2M-Reg) and $\lambda_{gcc} = 0.1$. For TransMorph and CorrMLP, a learning rate of 5e-5, batch size of 2 were used for 50,000 iterations, with regularization weights of $\lambda_{reg} = 0.01$ (1000 without M2M-Reg) and $\lambda_{gcc} = 1$. Hyperparameters were carefully optimized across all baselines. In all methods, a single network f_θ was used to

⁴ <http://adni.loni.usc.edu>

Table 2. Performance across pre-aligned bridge pair ratios (PET $\rightarrow$ MRI)

# Pre-Aligned Pairs	0 (unsup.)	1% (3 pairs)	5%	10%	25%	50%	100%
DSC $\uparrow$	0.676	0.691	0.709	0.709	0.725	0.729	0.730
% $ J_\phi _{<0}$ $\downarrow$	0.000	0.000	0.000	0.000	0.000	0.009	0.000

perform bidirectional mappings between the two modalities. Throughout training, source-target pairs were always drawn from different subjects, except for bridge pairs in the semi-supervised setting.

3.1 Evaluations

Quantitative Results. We evaluated registration quality using the Dice similarity coefficient (DSC) and deformation quality via the percentage of negative Jacobian determinants (% $|J_\phi|_{<0}$). In Table 1, most baselines exhibited severe % $|J_\phi|_{<0}$ and showed little to no DSC improvement, even performing worse than the initial state in PET $\rightarrow$ MRI. Applying M2M-Reg to baselines (highlighted in blue) significantly improved DSC (up to $2\times$) while nearly eliminating % $|J_\phi|_{<0}$, demonstrating our framework’s effectiveness in handling high heterogeneity. Even using only mono-modal cycle similarity ($\mathcal{L}_{m2m}$) led to substantial improvements, confirming its role in guiding the model to learn meaningful mappings. GradCyCon ($\mathcal{L}_{gcc}$) further enhanced both metrics by preventing unrealistic deformations. Notably, its effect was less pronounced in GradICON, likely because GradCyCon performed a similar implicit H^1 -type regularization.

Qualitative Results. In Fig. 3, baselines struggled to establish spatial mapping between disparate modalities. In PET-MRI registration, they either forced PET into MRI-like shapes (TransMorph, CorrMLP) or severely collapsed its structures (GradICON). In FA-MRI registration, they artificially expanded FA’s white matter to resemble cortical regions in MRI. In both cases, SyN failed to deform regions like the ventricles sufficiently. In contrast, M2M-Reg accurately aligned to the target anatomy while preserving the source modality’s structures.

Semi-Supervised Setting. We evaluated semi-supervised DIR by gradually increasing the proportion of pre-aligned bridge pairs in M2M-Reg within GradICON. In Table 2, performance improved monotonically, with DSC increasing by up to 5.4%p, proving its significant impact. Notably, even using just 3 pairs improved DSC by 1.5%p, and employing 25% of pairs achieved performance comparable to 100%, highlighting the effectiveness of the proposed framework.

4 Conclusion

Recognizing that multi-modal similarity is inadequate for registering highly heterogeneous modalities, we introduced M2M-Reg, a paradigm-shifting framework

that reformulates multi-modal deformable image registration (DIR) as a mono-modal problem. By leveraging only mono-modal similarity within a novel cyclic learning scheme, M2M-Reg circumvented the limitations of conventional unsupervised DIR, which relies on multi-modal similarity, while preserving architectural flexibility. Coupled with GradCyCon, which approximates diffeomorphic transformations through cycle consistency, our framework effectively captures complex heterogeneous mappings in an unsupervised manner. Empirical results demonstrated that M2M-Reg consistently improved performance across diverse baselines and datasets, successfully handling highly disparate modalities. Beyond its strong unsupervised performance, our semi-supervised approach further maximizes the utility of pre-aligned pairs. This adaptability positions M2M-Reg as a practical, scalable solution for advancing multi-modal medical image analysis.

Acknowledgments. This work was supported in part by the IITP RS-2024-00457882 (AI Research Hub Project), IITP 2020-II201361, NRF RS-2024-00345806, NRF RS-2023-00219019, and NRF RS-2023-002620.

Disclosure of Interests. The authors have no competing interests.

References

1. Statistical parametric mapping (spm). Scholarpedia **3**(4) (2008)
2. Avants, B.B., Tustison, N., Song, G., et al.: Advanced normalization tools (ants). *Insight j* **2**(365), 1–35 (2009)
3. 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)
4. Chen, J., Liu, Y., Wei, S., Bian, Z., Subramanian, S., Carass, A., Prince, J.L., Du, Y.: A survey on deep learning in medical image registration: New technologies, uncertainty, evaluation metrics, and beyond. *Medical Image Analysis* p. 103385 (2024)
5. Choo, K., Joo, J., Lee, S., Kim, D., Lim, H., Kim, D., Kang, S., Hwang, S.J., Yun, M.: Deep learning-based precontrast ct parcellation for mri-free brain amyloid pet quantification. *Clinical Nuclear Medicine* pp. 10–1097 (2022)
6. Demir, B., Tian, L., Greer, H., Kwitt, R., Vialard, F.X., Estépar, R.S.J., Bouix, S., Rushmore, R., Ebrahim, E., Niethammer, M.: Multigradicon: A foundation model for multimodal medical image registration. In: *International Workshop on Biomedical Image Registration*. pp. 3–18. Springer (2024)
7. Eppenhof, K.A., Lafarge, M.W., Moeskops, P., Veta, M., Pluim, J.P.: Deformable image registration using convolutional neural networks. In: *Medical Imaging 2018: Image Processing*. vol. 10574, pp. 192–197. SPIE (2018)
8. Eppenhof, K.A., Pluim, J.P.: Pulmonary ct registration through supervised learning with convolutional neural networks. *IEEE transactions on medical imaging* **38**(5), 1097–1105 (2018)
9. Fischl, B.: Freesurfer. *Neuroimage* **62**(2), 774–781 (2012)
10. Greer, H., Kwitt, R., Vialard, F.X., Niethammer, M.: Icon: Learning regular maps through inverse consistency. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3396–3405 (2021)

11. Heinrich, M.P., Hansen, L.: Voxelmorph++ going beyond the cranial vault with keypoint supervision and multi-channel instance optimisation. In: International workshop on biomedical image registration. pp. 85–95. Springer (2022)
12. Heinrich, M.P., Jenkinson, M., Bhushan, M., Matin, T., Gleeson, F.V., Brady, M., Schnabel, J.A.: Mind: Modality independent neighbourhood descriptor for multi-modal deformable registration. *Medical image analysis* **16**(7), 1423–1435 (2012)
13. 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: Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013: 16th International Conference, Nagoya, Japan, September 22–26, 2013, Proceedings, Part I 16. pp. 187–194. Springer (2013)
14. Hoopes, A., Mora, J.S., Dalca, A.V., Fischl, B., Hoffmann, M.: Synthstrip: Skull-stripping for any brain image. *NeuroImage* **260**, 119474 (2022)
15. Jack Jr, C.R., Bernstein, M.A., Fox, N.C., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P.J., L. Whitwell, J., Ward, C., et al.: The alzheimer’s disease neuroimaging initiative (adni): Mri methods. *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine* **27**(4), 685–691 (2008)
16. Jena, R., Chaudhari, P., Gee, J.C.: Fireants: Adaptive riemannian optimization for multi-scale diffeomorphic registration. arXiv preprint arXiv:2404.01249 (2024)
17. Lange, F.J., Arthofer, C., Bartsch, A., Douaud, G., McCarthy, P., Smith, S.M., Andersson, J.L.: Mmorf—fsl’s multimodal registration framework. *Imaging Neuroscience* **2**, 1–30 (2024)
18. Ma, T., Zhang, S., Li, J., Wen, Y.: Iirp-net: Iterative inference residual pyramid network for enhanced image registration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11546–11555 (2024)
19. 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)
20. Mang, A., Biros, G.: Constrained h^1 -regularization schemes for diffeomorphic image registration. *SIAM journal on imaging sciences* **9**(3), 1154–1194 (2016)
21. Meng, M., Feng, D., Bi, L., Kim, J.: Correlation-aware coarse-to-fine mlps for deformable medical image registration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9645–9654 (2024)
22. 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)
23. Mok, T.C., Li, Z., Bai, Y., Zhang, J., Liu, W., Zhou, Y.J., Yan, K., Jin, D., Shi, Y., Yin, X., et al.: Modality-agnostic structural image representation learning for deformable multi-modality medical image registration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11215–11225 (2024)
24. Petersen, R.C., Aisen, P.S., Beckett, L.A., Donohue, M.C., Gamst, A.C., Harvey, D.J., Jack Jr, C., Jagust, W.J., Shaw, L.M., Toga, A.W., et al.: Alzheimer’s disease neuroimaging initiative (adni) clinical characterization. *Neurology* **74**(3), 201–209 (2010)
25. Qiu, H., Qin, C., Schuh, A., Hammernik, K., Rueckert, D.: Learning diffeomorphic and modality-invariant registration using b-splines. In: Medical imaging with deep learning (2021)

26. Rueckert, D., Sonoda, L.I., Hayes, C., Hill, D.L., Leach, M.O., Hawkes, D.J.: Non-rigid registration using free-form deformations: application to breast mr images. *IEEE transactions on medical imaging* **18**(8), 712–721 (1999)
27. Song, X., Xu, X., Yan, P.: Dino-reg: General purpose image encoder for training-free multi-modal deformable medical image registration. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 608–617. Springer (2024)
28. Tan, Z., Zhang, H., Tian, F., Zhang, L., Sun, W., Lu, H.: Progressively coupling network for brain mri registration in few-shot situation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 623–633. Springer (2023)
29. 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)
30. Wang, Y., Shen, Y., Liu, D., Li, G., Guo, Z., Fan, Y., Niu, Y.: Evaluations of diffusion tensor image registration based on fiber tractography. *Biomedical engineering online* **16**, 1–20 (2017)