

On-the-Fly Guidance Training for Medical Image Registration

Yuelin Xin^{1*}, Yicheng Chen^{3*}, Shengxiang Ji^{4*}, Kun Han^{2*}, and Xiaohui Xie²

¹ University of Leeds, Leeds, UK

² University of California Irvine, Irvine, CA

³ Tongji University, Shanghai, China

⁴ Huazhong University of Science and Technology, Wuhan, China

Abstract. This study introduces a novel **On-the-Fly Guidance** (OFG) training framework for enhancing existing learning-based image registration models, addressing the limitations of weakly-supervised and unsupervised methods. Weakly-supervised methods struggle due to the scarcity of labeled data, and unsupervised methods directly depend on image similarity metrics for accuracy. Our method proposes a supervised fashion for training registration models, without the need for any labeled data. OFG generates pseudo-ground truth during training by refining deformation predictions with a differentiable optimizer, enabling direct supervised learning. OFG optimizes deformation predictions efficiently, improving the performance of registration models without sacrificing inference speed. Our method is tested across several benchmark datasets and leading models, it significantly enhanced performance, providing a plug-and-play solution for training learning-based registration models. Code available at: <https://github.com/cilix-ai/on-the-fly-guidance>

Keywords: image registration · on-the-fly guidance · pseudo label.

1 Introduction

Medical image registration is pivotal in medical image analysis, aiming to align two medical images by optimizing their visual similarity through a deformation field. There are two factions: traditional optimization-based methods like [3,4,21], which iteratively refine the deformation field using mathematical constraints, and modern learning-based methods [8,9], which predict deformation fields from image pairs. Both approaches are vital, with the latter gaining significant traction for its direct prediction capabilities, marking a swift evolution in the field.

Nevertheless, learning-based methods face a major hurdle: the trade-offs between weakly-supervised learning, which yields better results at the cost of extensive labeling, and unsupervised learning, which foregoes labels but directly relies on less precise image similarities for deformation field derivation. This situation prompts the critical question: *is it possible to create a training method that bypasses the need for manual labels while still benefiting from the precision of direct supervision?*

In this work, we present a novel training framework named on-the-fly guidance (OFG) that merges supervised learning with existing learning-based image registration methods to boost their performance. OFG uniquely generates pseudo-ground truth on the fly through instance-specific optimization, using these results for direct supervision. This hybrid approach combines direct prediction with iterative refinement in a two-stage process: 1) the model predicts a deformation field ϕ_{pre} , and 2) this field is optimized to produce ϕ_{opt} , which then acts as a pseudo-label. The model’s training is guided by directly comparing the initial and optimized deformation fields using $MSE(\phi_{pre}, \phi_{opt})$, ensuring a direct and efficient learning process.

OFG introduces a crucial incremental supervision method that guides models toward convergence by setting intermediate goals rather than a final objective. It optimizes the current deformation prediction in a step-by-step manner, easing the model’s learning process while balancing supervision quality with computational efficiency. This approach allows image registration models to undergo a nuanced, self-improving training process. Compared to baseline unsupervised and other pseudo-supervised methods like self-training, OFG shows consistent and significant improvements, underscoring its effectiveness.


The main contributions of our work are summarized as follows:

- Introducing OFG, a training framework that enhances existing image registration models, utilizing supervised learning without relying on labeled data.
- Using optimized pseudo-ground truth for incremental learning targets, fostering a self-enhancing cycle between the model and optimizer.
- Presenting through extensive benchmarks that our method surpasses baselines and previous state-of-the-art across different datasets and models.

2 Related Work

Weakly-Supervised & Unsupervised Training. Learning-based methods have recently overtaken traditional optimization-based approaches [16,2,11,25,7] in performance and efficiency [15,20,10,29], with supervised methods [27,23] depending on ground-truth deformation fields often derived from traditional techniques. The rise of unsupervised methods [5,26,22,28,9,8] optimize metrics like NCC to understand the dataset globally. Popular unsupervised models like VoxelMorph [5], ViT-V-Net [9], and TransMorph [8] employ a U-Net structure with CNN or ViT elements for deformation prediction and can use segmentation labels for enhanced accuracy, though with the high cost of annotations.

Self-training. Closely related to our research is Cyclical Self-training [6], which adopts a teacher-student approach, alternating training stages and employing pseudo labels for guidance. Our approach differs significantly in the following ways: **1)** the optimizer in [6] employs a non-learnable approach to generate deformation from the fusion of two encoded image features, which is not competitive compared with existing learning-based models, **2)** OFG generates pseudo labels incrementally for each training epoch, as opposed to updating labels between training stages which may introduce challenging shifts for the



3 Method

In this work, we present a two-stage training framework with the proposed on-the-fly guidance (OFG), using pseudo-ground truth, it embeds optimization and supervised learning in the training of registration models. 1 presents the overall structure of the proposed two-stage training framework.

$$F_{\theta}(I_f, I_m) = \phi \quad (1)$$

Optimization Stage. The optimization stage utilizes the proposed optimizer to iteratively refine the deformation field ϕ_{pre} predicted from the current training step by n steps (10 in our default setting). Subsequently, the optimized

deformation field ϕ_{opt} is used as the pseudo label to provide supervision for the current predicted deformation during the training, forming a feedback loop between the prediction model and the optimizer module.

3.2 On-the-Fly Guidance Training

Rather than relying on a fixed pseudo label derived from either a pre-trained model’s prediction or the final optimized deformation, our approach introduces on-the-fly guidance. This dynamic supervision evolves alongside the training process to control discrepancies between pseudo labels and current predictions. This approach offers two advantages: **1)** the limited number (typically 5 to 10) of optimizations incurs acceptable training overhead, **2)** the optimized deformation serves as an attainable goal for the ongoing training step, providing more direct guidance for the model. In essence, OFG delivers incremental supervision, offering step-by-step guidance for the model.

The underlying assumption behind OFG is that the model and the optimizer can form a self-improving relationship. The model can provide a reasonable prediction, and in turn, the optimizer can refine that prediction. This is validated in our experiments, as the optimizer can provide a good pseudo label even with random input parameters.

3.3 Differentiable Deformation Optimizer

Our on-the-fly guidance hinges on an effective optimizer, we explored three optimization strategies: network-based, downsampled, and our proposed optimizer (see 1 (b)) with instance optimization and high flexibility for parameter updates, the latter showing the best results (see 5). Detailed comparison in Sec. 4.3.

The proposed differentiable optimizer is simple yet powerful, taking in the deformation field generated by the prediction model as its initial parameters and optimizing it to generate the pseudo label. It features a Spatial Transformer Network (STN) [13] without extra parameters, focusing updates solely on the deformation field. During an optimization iteration, the current deformation field is applied to the moving image with the STN, yielding a warped image. An energy function will evaluate the discrepancy between the warped and the fixed image, and the distortion of the deformation field. This loss is backpropagated using Adam [14] or SGD, this essentially refines the deformation field, i.e.,

$$\phi_{opt}^{(n+1)} = \phi_{opt}^{(n)} - \eta \nabla E_{opt} \quad (2)$$

where n denotes the iteration step, η is the learning rate, and ∇E_{opt} represents the gradient of the optimization energy function, implementation detailed in 4. Notably, the optimizer’s role is limited to training, not inference, maintaining training efficiency.

3.4 Implementation Detail

Training Loss Function. The model’s training utilizes a loss function that enforces supervision from pseudo labels. The model learns by minimizing the discrepancy between the predicted deformation field ϕ_{pre} , and the optimized deformation field ϕ_{opt} , which is quantified using MSE. Implemented is as follows:

$$L_{ofg} = \frac{1}{n} \sum (\phi_{pre} - \phi_{opt})^2 \quad (3)$$

where L_{ofg} is the model’s training loss, it is essentially a MSE-based supervision. Optionally, a weight decay of 0.02 can be added to reduce overfitting for small datasets, as suggested in VoxelMorph [5].

Optimizer Energy Function. The energy function to be minimized in the differentiable optimizer consists of two terms: an image similarity loss term that captures the difference between the warped image $I_m \circ \phi$ and fixed image I_f and a L_2 regularization loss that imposes smoothness in ϕ :

$$E_{opt}(I_m, I_f, \phi) = NCC(I_f, I_m \circ \phi) + \sum_{p \in \Omega} \|\nabla \phi(p)\|^2 \quad (4)$$

where $\circ$ is the transform operation which warps I_m using ϕ . The similarity metric we used is the normalized cross-correlation (NCC), $\hat{I}_f(p)$ and $\hat{I}_m(p)$ represent the mean voxel value within a local window of size n^3 centered at voxel p :

$$NCC(I_f, I_m \circ \phi) = \sum_{p \in \Omega} \frac{(\sum_{p_i} (f(p_i) - \hat{f}(p))([I_m \circ \phi](p_i) - [\hat{I}_m \circ \phi](p)))^2}{(\sum_{p_i} (f(p_i) - \hat{f}(p))^2)(\sum_{p_i} ([I_m \circ \phi](p_i) - [\hat{I}_m \circ \phi](p))^2)} \quad (5)$$

4 Experiments

4.1 Experiment Conditions

Dataset and Preprocessing. We utilize three public Brain MRI datasets in our study: IXI [17], OASIS [18], and LPBA40 [24], with standard preprocessing steps including skull stripping, resampling, and affine transformation. For IXI, we use 200 volumes for training and 20 for validation; for OASIS, 200 for training and 19 for validation; and for LPBA40, 30 for training, 9 for validation. We also utilize the Abdomen CT-CT dataset [12] to evaluate the generalizability of our method on CT registration. 30 for training, and 20 for validation.

Evaluation Metrics. Our evaluation uses two primary metrics: the Dice score (DSC) [5,4] for assessing volume overlap in anatomical segmentations, indicating registration accuracy, and the Jacobian matrix to measure deformation field smoothness. The latter involves counting non-background voxels where $\%|J_\phi| < 0$, highlighting non-diffeomorphic deformation areas [1].

Baseline Methods. We validated our method based on various popular registration methods. This comparison included two traditional methods, SyN

Table 1. Evaluation results for different methods on various datasets. The OFG architecture provides significant and substantial improvement on the unsupervised learning-based methods. These results validate OFG’s effectiveness and generalizability.

Datasets	Methods	Base. DSC $\uparrow$	OFG DSC $\uparrow$	Base. $\% J_\phi < 0 \downarrow$	OFG $\% J_\phi < 0 \downarrow$
IXI [17]	SyN [3]	0.647	N/A	1.96e-6	N/A
	NiftyReg [19]	0.585	N/A	0.029	N/A
	VoxelMorph [5]	0.714	0.737 (+2.3%)	1.398	0.516 (-63.1%)
	ViT-V-Net [9]	0.716	0.738 (+2.2%)	1.543	0.545 (-64.7%)
	TransMorph [8]	0.744	0.760 (+1.6%)	1.433	0.794 (-44.6%)
	SyN [3]	0.769	N/A	1.58e-4	N/A
OASIS [18]	NiftyReg [19]	0.762	N/A	0.011	N/A
	VoxelMorph [5]	0.788	0.794 (+0.6%)	0.911	0.490 (-46.2%)
	ViT-V-Net [9]	0.794	0.809 (+1.5%)	0.887	0.487 (-45.1%)
	TransMorph [8]	0.818	0.818 (=)	0.765	0.517 (-32.4%)
	SyN [3]	0.703	N/A	1.18e-4	N/A
	NiftyReg [19]	0.691	N/A	1.13e-3	N/A
LPBA40 [24]	VoxelMorph [5]	0.658	0.666 (+0.8%)	0.288	0.023 (-92.0%)
	ViT-V-Net [9]	0.663	0.672 (+0.9%)	0.390	0.112 (-71.3%)
	TransMorph [8]	0.678	0.684 (+0.6%)	0.438	0.150 (-65.8%)




Fig. 2. Visualization of registration results on LPBA40 [24]. Demo randomly extracted from the comparison results between baseline TransMorph, VoxelMorph (row 2) and their respective model trained with OFG (row 1). OFG shows improved smoothness.

[3] and NiftyReg [19] and multiple learning-based methods, VoxelMorph [5], ViT-V-Net [9], TransMorph [8]. All methods are in their default configuration.

Experiment Settings. All models were trained on RTX 4090 for 500 epochs using Adam [14], with an initial learning rate of 1e-4, batch size of 1, weight decay of 0.02. For the differentiable optimizer, we used an initial learning rate of 0.1, coupled with a default optimization step count of 10 during training.

4.2 Image Registration Results

We conducted extensive experiments on the three datasets with three baseline models to showcase the effectiveness and robustness of our proposed framework. Also comparing with the baseline of non-learning methods, see 1.

OFG on Brain MRI. We have evaluated OFG on brain MRI registration extensively using 3 popular baseline models and datasets. Our method provides a

Table 2. Abdomen CT registration. VoxelMorph uses 3-step MSE optimizer, TransMorph uses 5-step NCC optimizer. **Table 3.** Cyclical Self-training vs. OFG on LPBA40. OFG performs significantly better.

Methods	DSC	Jacob.	OFG DSC	OFG Jacob.
VoxelMorph	0.312	0.165	0.319 (+0.7%)	0.120 (-27.3%)
TransMorph	0.317	0.165	0.382 (+6.5%)	0.010 (-93.9%)

Methods	DSC	Jacob.
VoxelMorph w. CST	0.628	0.047
VoxelMorph w. OFG	0.666	0.023




Fig. 3. Visualization comparing training progress and validation DSC on LPBA40 across models. Self-training uses pre-trained network deformation fields as pseudo labels; optimized self-training enhances this with extra optimization steps. Our method achieves the best outcome, with self-training lagging due to convergence complexities.

consistent and significant margin on DSC across different models and datasets, demonstrating its effectiveness and generalizability. On the IXI dataset, OFG improved DSC by +1.6% over TransMorph. For VoxelMorph and ViT-V-Net, it increased DSC by +2.3% and +2.2%, respectively, highlighting its general applicability. While on the smaller LPBA40 dataset, OFG’s added supervision proved essential in preventing overfitting, underlining the importance of challenging supervision in sparse-data scenarios. Conversely, on the OASIS dataset, OFG showed little improvement on TransMorph, likely due to the dataset’s lesser challenge to the model, a hypothesis supported by the lowest training loss of TransMorph on OASIS over all tested cases, suggesting a reduced learning potential for this case. Importantly, for all test cases, our method significantly reduced the percentage of non-diffeomorphic voxels, preventing overly sharp deformations and improving the quality of the registration.

OFG on Abdomen CT. We briefly tested OFG on the Abdomen CT-CT dataset (see 2), with varying optimization configurations. For VoxelMorph, we used MSE as the optimizer energy function, with 3 optimization steps, resulting in a small improvement over baseline. For TransMorph, we used the default energy function, with 5 optimization steps, yielding a much greater improvement over baseline. OFG’s effectiveness on another modality showcases its robustness.

OFG vs. Self-training. We compared our result with various forms of self-training including Cyclical Self-training (CST). Our method consistently outperforms self-training methods on LPBA40, see 3. We also applied CST and OFG on VoxelMorph, and tested on LPBA40, our method provides +3.8% better DSC while halving the Jacobian, see 3. This is largely due to the OFG training strategy explained in 3.2.




Fig. 4. Ablation results on blending OFG with baseline unsupervised learning. From left to right, we evaluated optimization frequency, loss blending, and probabilistic optimization, generally showing a decrease in performance when the intensity of OFG is decreased, proving its effectiveness.

Table 4. Ablation results of the OFG Op- **Table 5.** Comparison between our optimizer on LPBA40. This evaluation focuses timer and network-based optimizer on the initial 30 epochs. **Table 5.** Comparison between our timer and network-based optimizer on the initial 30 epochs.

Designs	DSC	Opt. Time	VRAM
None	0.635	N/A	12.07 GB
Network-based	0.642	13.604	14.35 GB
DownSample	0.610	0.097	12.44 GB
<i>Ours</i>	0.654	3.228	13.07 GB

Designs	DSC	$\% J_\phi < 0$	Opt. Time
VoxelMorph-1	0.691	0.456	1.174
VoxelMorph-2	0.690	0.383	2.314
VoxelMorph-5	0.693	0.385	5.856
<i>Ours</i>	0.731	0.503	3.018

4.3 Ablation Study

OFG Intensity Ablation. To show OFG’s deciding factors and how they influence performance, we blended OFG with baseline unsupervised loss in 3 different forms: **1)** Optimization frequency: only use optimizer every n epochs, i.e., decreased optimization frequency. **2)** Loss weight composition: adding NCC loss into the loss function, i.e., $L = \alpha L_{ofg} + \beta L_{NCC}$. **3)** Probabilistic optimization: only randomly optimizes a portion of the image instances during training. As shown in Fig 4, we observed a decrease in performance when the intensity of OFG decreases, in all three forms. Notably, a low optimization frequency resembles the training strategy used in Cyclical Self-training. This result suggests a higher optimization frequency (intensity) provides improved performance.

Optimizer Design. We also evaluated 2 other optimizer designs for our framework, including: **1)** Network-based optimizer: using a network capable of fitting a general transformation to optimize deformation fields, in our case, we used n cascaded VoxelMorph. **2)** Downsampled optimizer: to improve the computational overhead, a downsampled optimizer reduces all dimensions in half, with only 1/8 of the updatable parameters. Table 4 and 5 show the proposed design achieves the best performance.

Optimization Steps. We assessed the impact of optimization steps ranging from 1 to 15 on training outcomes to balance computational efficiency with optimization quality. Findings indicate that 5 to 10 steps offer optimal balance, enhancing optimization quality without significantly lengthening training time (only with a 10 to 18% increase), with no notable benefits from exceeding this range. See Fig. 1. in supplementary material for detailed results.

Self-improving Relationship. OFG is based on the concept that the model and optimizer can enhance each other, with the optimizer’s robustness being key to generating high-quality pseudo labels in various scenarios. Our findings indicate that the optimizer can effectively refine initial deformations, even those generated randomly or from models with random initialization, leading to significant improvements. See Fig. 2. in supplementary material for detailed results.

5 Conclusion

This work introduces On-the-Fly Guidance, a training framework that successfully applies supervised-style training to learning-based registration models. Demonstrating significant improvements on benchmark datasets, especially with deformation smoothness, OFG has proven its effectiveness and generalizability. OFG only comes with limited training overhead and no inference overhead. The flexibility of our method allows future work to focus on aspects such as improving the efficiency of the optimizer, using dynamic optimization steps, altering the optimizer design, and so on.

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* p. 95–113 (Oct 2007)
2. AV, D., A, B., NS, R., P, G.: Patch-based discrete registration of clinical brain images. *Patch Based Tech Med Imaging* (2016) (2016)
3. Avants, B., Epstein, C., Grossman, M., Gee, J.: Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Medical Image Analysis* **12**(1), 26–41 (2008), special Issue on The Third International Workshop on Biomedical Image Registration – WBIR 2006
4. Bajcsy, R., Kovačič, S.: Multiresolution elastic matching. *Computer Vision, Graphics, and Image Processing* **46**(1), 1–21 (1989)
5. 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 (aug 2019)
6. Bigalke, A., Hansen, L., Mok, T.C.W., Heinrich, M.P.: Unsupervised 3d registration through optimization-guided cyclical self-training. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 677–687. Springer Nature Switzerland, Cham (2023)
7. C, C., L, W., LD, S., JG, C., MI, M., JT, R.: Large deformation diffeomorphic metric mapping registration of reconstructed 3d histological section images and in vivo mr images. *Front Hum Neurosci* (2010)

8. 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 (nov 2022)
9. Chen, J., He, Y., Frey, E.C., Li, Y., Du, Y.: Vit-v-net: Vision transformer for unsupervised volumetric medical image registration (2021)
10. Chen, Q., Li, Z., Lui, L.M.: A learning framework for diffeomorphic image registration based on quasi-conformal geometry. *CoRR* **abs/2110.10580** (2021)
11. Glocker, B., Komodakis, N., Tziritas, G., Navab, N., Paragios, N.: Dense image registration through mrfs and efficient linear programming. *Medical Image Analysis* **12**(6), 731–741 (2008), special issue on information processing in medical imaging 2007
12. Hering, A., Hansen, L., Mok, T.C.W., Chung, A.C.S., Siebert, H., Häger, S., Lange, A., Kuckertz, S., Heldmann, S., Shao, W., Vesal, S., Rusu, M., Sonn, G., Estienne, T., Vakalopoulou, M., Han, L., Huang, Y., Yap, P.T., 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., Xu, Z., Jian, B., De Benetti, F., Wodzinski, M., Gunnarsson, N., Sjölund, J., Grzech, D., Qiu, H., Li, Z., Thorley, A., Duan, J., Großbröhmer, C., Hoopes, A., Reinertsen, I., Xiao, Y., Landman, B., Huo, Y., Murphy, K., Lessmann, N., 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**(3), 697–712 (2023). <https://doi.org/10.1109/TMI.2022.3213983>
13. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial transformer networks (2016)
14. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization (2017)
15. Liu, R., Li, Z., Fan, X., Zhao, C., Huang, H., Luo, Z.: Learning deformable image registration from optimization: Perspective, modules, bilevel training and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(11), 7688–7704 (2022)
16. Loeckx, D., Maes, F., Vandermeulen, D., Suetens, P.: Nonrigid image registration using free-form deformations with a local rigidity constraint. In: Barillot, C., Haynor, D.R., Hellier, P. (eds.) *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2004*. pp. 639–646. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)
17. London, I.C.: Information extraction from images (2023), <https://brain-development.org/ixi-dataset/>
18. 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 (09 2007)
19. Centre for Medical Image Computing, University College London, U.: Niftyreg (2023), <http://cmictig.cs.ucl.ac.uk/wiki/index.php/NiftyReg>
20. Nazib, A., Fookes, C., Perrin, D.: A comparative analysis of registration tools: Traditional vs deep learning approach on high resolution tissue cleared data (2018)
21. Shen, D., Davatzikos, C.: Hammer: hierarchical attribute matching mechanism for elastic registration. *IEEE Transactions on Medical Imaging* **21**(11), 1421–1439 (2002)
22. Shen, Z., Han, X., Xu, Z., Niethammer, M.: Networks for Joint Affine and Non-parametric Image Registration. *arXiv e-prints arXiv:1903.08811* (Mar 2019)

23. Sokooti, H., de Vos, B., Berendsen, F., Lelieveldt, B.P.F., Išgum, I., Staring, M.: Nonrigid image registration using multi-scale 3d convolutional neural networks. In: Descoteaux, M., Maier-Hein, L., Franz, A., Jannin, P., Collins, D.L., Duchesne, S. (eds.) *Medical Image Computing and Computer Assisted Intervention - MICCAI 2017*. pp. 232–239. Springer International Publishing, Cham (2017)
24. University of Southern California, L.o.N.I.: Loni probabilistic brain atlas (lpba40) (2023), <https://loni.usc.edu/research/atlas>
25. Thirion, J.P.: Image matching as a diffusion process: an analogy with maxwell’s demons. *Medical Image Analysis* **2**(3), 243–260 (1998)
26. de Vos, B.D., Berendsen, F.F., Viergever, M.A., Sokooti, H., Staring, M., Išgum, I.: A deep learning framework for unsupervised affine and deformable image registration. *Medical Image Analysis* **52**, 128–143 (feb 2019)
27. Yang, X., Kwitt, R., Styner, M., Niethammer, M.: Quicksilver: Fast predictive image registration - a deep learning approach (2017)
28. Zhang, Y., Pei, Y., Zha, H.: Learning dual transformer network for diffeomorphic registration. In: de Bruijne, M., Cattin, P.C., Cotin, S., Padoy, N., Speidel, S., Zheng, Y., Essert, C. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*. pp. 129–138. Springer International Publishing, Cham (2021)
29. Zou, J., Gao, B., Song, Y., Qin, J.: A review of deep learning-based deformable medical image registration. *Frontiers in Oncology* **12** (2022)