TrICON: Learning Diffeomorphic Lung CT Registration through Triplet Consistency

Maria Marquez-Sosa[1−2][0000−0002−6962−9881], Émile Saillard[1−3][0009−0007−1830−3885], Adrian Basarab[1−3][0000−0002−5642−7244], and Michaël Sdika[1−2][0000−0002−3129−6874]

[1] CREATIS - Centre de Recherche en Acquisition et Traitement de l'Image pour la Santé, CREATIS UMR 5220 U1294, F69100
[2] INSA-Lyon
[3] Université Claude Bernard Lyon 1

Abstract. Diffeomorphic deformable image registration is critical for the quantitative analysis of thoracic computed tomography (CT), yet recent learning-based methods promote diffeomorphisms through pairwise inverse-consistency regularization. This constraint, however, is self-referential by construction: a transformation is validated solely against its own inverse, allowing anatomically implausible correspondences to pass as readily as correct ones. We introduce `TrICON`, a population-trained framework that generalizes pairwise inverse consistency to triplet-wise consistency, jointly enforcing invertibility, closed-loop cycle consistency, and transitivity across three images via a stochastic, differentiable formulation compatible with existing registration architectures. For a given triplet, our approach requires the direct transformation between two images to agree with the transformation composed through a third, auxiliary image, yielding a correctness signal decoupled from pairwise inverse consistency. On inter-patient lung CT registration, `TrICON` outperforms representative state of the art (SOTA) methods, achieving higher DSC and lower distance based errors (HD95 and ASSD), while maintaining competitive MCD performance, and preserving near-diffeomorphic smoothness and invertibility without any explicit spatial regularization. Code is available at https://github.com/creatis-myriad/TrICON.

Keywords: Deformable image registration · Diffeomorphic registration · Transitivity consistency · Inter-patient registration · Inverse consistency · Cycle consistency

1 Introduction

Population-level studies of lung anatomy, including radiotherapy dose mapping, statistical atlas construction, and label propagation, rely on accurate *inter-patient* deformable image registration (DIR) to establish dense anatomical correspondences between CT volumes [1,2]. Downstream analyses such as voxel based analysis depend directly on the estimated deformation field, so accuracy and topology preservation are essential [1,3]. This calls for diffeomorphic transformations, which are challenging to enforce in the lung given the large respiratory motion and anatomical variability. Classical diffeomorphic registration methods, e.g., Demons [4], LDDMM [5], and ANTS SyN [6], guarantee diffeomorphism by construction but incur substantial computational cost.

Alternative approaches replace explicit smoothness regularization with algebraic consistency constraints. Early work introduced inverse consistency and identified transitivity as a fundamental property of registration operators [7,8]. Subsequent methods enforced exact symmetry and transitivity by routing transformations through a reference image, at the cost of transformation inversion and dependence on an arbitrary reference [9], and Geng et al. [10] minimized inverse-consistency and transitivity errors jointly over manifold triplets, showing that registering three manifolds at once constrains transitivity in a way that is impossible with only two. Both are instance-level optimization methods, re-solved for every new set of images and unable to amortize to unseen subjects.

Deep learning (DL) replaces this repeated optimization with a single population trained model. Existing approaches enforce inverse consistency only within image pairs [11,12], and `ICON` [13] and its extensions [14,15,16] show that approximate diffeomorphisms can emerge from stochastic inverse consistency alone, without explicit spatial regularization. Gu et al. [17] additionally introduce group wise cycle consistency, but retain explicit smoothness and topology constraints, evaluate consistency only on the voxel lattice, and consider longitudinal images from a single subject, without enforcing transitivity. Pairwise inverse consistency is inherently self-referential, as a transformation is validated only against its own inverse, allowing compensating errors to satisfy the constraint despite anatomically implausible correspondences, particularly in inter-patient registration.

We address this limitation by extending consistency constraints to image triplets within the population trained framework of `ICON`. We propose `TrICON`, a **Tr**iplet **I**nformed **CON**sistent registration framework that remains trainable end to end without explicit spatial regularization. Building on the classical observation that triplet-level transitivity provides a consistency signal unavailable to pairwise methods [10], `TrICON` contributes the first population-trained, amortized instantiation: a composable family of stochastic, differentiable triplet consistency losses, evaluated on a perturbed identity grid [13], compatible with existing backbones and requiring no per-subject or per-dataset optimization.

2 Methods

2.1 Preliminaries

Let $\mathcal{I}$ denote the set of images defined over a common spatial domain $\Omega \subset \mathbb{R}^N$, and let F_θ denote a registration network mapping an ordered image pair (I_i, I_k) to a transformation $\Phi_{ik} : \Omega \to \Omega$. We parameterize $\Phi_{ik} = D_{ik} + \mathrm{Id}$, where $D_{ik} = \mathrm{interp}(F_\theta(I_i, I_k))$ is a continuously interpolated displacement field. Thus, Φ_{ik} is represented as the identity (I_d) plus a displacement field and is continuous

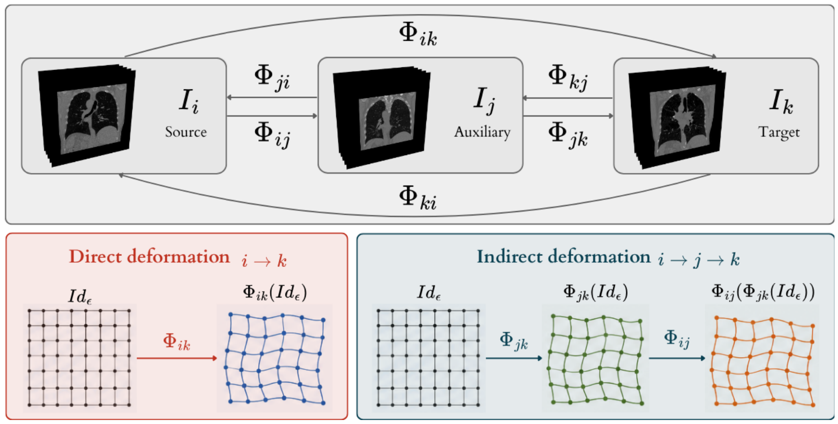

Fig. 1. Overview of the proposed framework. **Top:** Given a triplet (I_i, I_j, I_k), pairwise transformations are estimated between every image pair. **Bottom:** Transitivity is enforced by comparing the direct transformation Φ_{ik} (red) with the composed transformation $\Phi_{ij} \circ \Phi_{jk}$ (blue), both evaluated on the perturbed identity grid Id_ϵ.

and differentiable almost everywhere, with $I_i \circ \Phi_{ik} \approx I_k$. Invertibility can be obtained by enforcing the *inverse consistency* property below [9,7]:

$$\Phi_{ki} \circ \Phi_{ik} = \mathrm{Id}, \tag{1}$$

Enforcing this property can be done by construction [9] or, as in `ICON` [13], using a differentiable loss evaluated on a perturbed identity grid, $\mathrm{Id}_\epsilon = \mathrm{Id} + \epsilon$, where $\epsilon \sim \mathcal{N}(0, 1)/N$ and N the grid resolution along the last spatial axis, rather than at voxel centers. This off grid evaluation is necessary because Φ_{ij} is only piecewise linear between voxels, allowing discontinuous transformations to satisfy consistency on the lattice while remaining inconsistent elsewhere [13]. We adopt the same strategy by sampling a single perturbed grid per training iteration and sharing it across all consistency losses.

These constraints are inherently pairwise and do not enforce consistency among transformations across multiple images. We address this limitation next.

2.2 TrICON: Triplet Consistency Losses

Our central contribution is to extend registration consistency from pairs to triplets, introducing an algebraic property, transitivity, and two new loss formulations to enforce it. The registration is transitive if for all (i, j, k):

$$\Phi_{ij} \circ \Phi_{jk} = \Phi_{ik}. \tag{2}$$

The direct registration from i to k equals the composed registration through an intermediate image j. When the transforms are invertible, this is equivalent to:

$$\Phi_{ij} \circ \Phi_{jk} \circ \Phi_{ki} = Id. \tag{3}$$

Transitivity constrains transformations jointly over the whole population rather than independently for each image pair.

Figure 1 summarizes the proposed framework. Given a sampled triplet of images (I_i, I_j, I_k), the registration network predicts all six pairwise directed transformations for indices $\P_{ijk} := \{(i,j),(j,i),(i,k),(k,i),(j,k),(k,j)\}$ and penalty losses are used to enforce transitivity.

Unlike [10], we formulate these constraints as differentiable, stochastic losses suitable for training a single population-level network end to end. The specific loss constructions below and their combination with a learned inverse-consistency term are, to our knowledge, new in this learning based setting.

Transitivity loss — novel. The first loss we propose to enforce transitivity is directly based on Eq. 2 evaluated on a perturbed grid:

$$\mathcal{L}_{\text{trans}} = \frac{1}{6} \sum_{(i,k)\in\mathcal{P}} \left\| \Phi_{ik}(\text{Id}_\epsilon) - \Phi_{ij}\big(\Phi_{jk}(\text{Id}_\epsilon)\big) \right\|_2^2 , \tag{4}$$

Inverse consistency loss. The transitivity loss alone does not favor invertibility: it is trivially minimized, for example, by a network that always outputs the null transform. To cope with that, we enforce invertibility separately with:

$$\mathcal{L}_{\text{inv}} = \frac{1}{6} \sum_{(i,j)\in\mathcal{P}} \left\| \Phi_{ji}\big(\Phi_{ij}(\text{Id}_\epsilon)\big) - \text{Id}_\epsilon \right\|_2^2 . \tag{5}$$

which corresponds to the inverse consistency loss of `ICON` [13] applied independently to each directed pair within the sampled triplet.

Cycle consistency loss — novel. Transitivity can also be enforced via Eq. 3:

$$\mathcal{L}_{\text{cycle}} = \frac{1}{6} \sum_{abc\in\{ijk,ikj,jki,jik,kij,kji\}} \|\varphi_{abc}(\text{Id}_\epsilon) - \text{Id}_\epsilon\|_2^2 \tag{6}$$

where $\varphi_{abc} = \Phi_{ab} \circ \Phi_{bc} \circ \Phi_{ca}$. Averaging over all six cycle orientations ensures that each pairwise transformation appears equally often in every position of the composition. This loss jointly promotes transitivity and invertibility.

Similarity loss. For completeness, the similarity loss used to drive the registration is averaged over all six ordered pairs (standard image-matching term),

$$\mathcal{L}_{\text{sim}} = \frac{1}{6} \sum_{(i,j)\in\mathcal{P}} \mathcal{S}\big(I_i \circ \Phi_{ij}, I_j\big), \tag{7}$$

where $\mathcal{S}$ is the local normalized cross correlation (LNCC) with a Gaussian kernel with $\sigma = 5$ voxels. However, any similarity metric can be used.

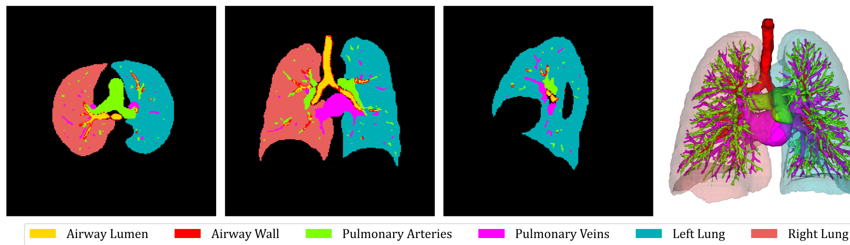

Fig. 2. The six anatomical labels of the lung segmentation: airway lumen (L_1), airway wall (L_2), pulmonary arteries (L_3), pulmonary veins (L_4), left lung (L_5), and right lung (L_6). From left to right: axial, coronal and sagittal views, and 3D rendering.

2.3 Combined Objective

We define the general **TrICON** objective as

$$\mathcal{L}_{\text{TrICON}} = \mathcal{L}_{\text{sim}} + \lambda \cdot \left(\mathbb{1}_{\text{inv}} \mathcal{L}_{\text{inv}} + \mathbb{1}_{\text{cycle}} \, \mathcal{L}_{\text{cycle}} + \mathbb{1}_{\text{trans}} \, \mathcal{L}_{\text{trans}} \right), \tag{8}$$

where $\mathbb{1}_{\text{inv}}, \mathbb{1}_{\text{cycle}}, \mathbb{1}_{\text{trans}} \in \{0, 1\}$ independently enable each consistency term. No explicit spatial regularization is used, with regularity arising solely from linear interpolation and the consistency constraints on perturbed grids. Section 4 evaluates all seven non-trivial combinations of active consistency terms, and compares the best performing configuration with SOTA registration methods.

3 Experiments

Data and pre-processing We trained and evaluated `TrICON` for inter-patient lung CT registration using a multi-structure annotated subset of the LUNA16 dataset [18]. The dataset contains 254 thoracic CT scans with expert annotations [19] for the six anatomical structures shown in Fig. 2. CT intensities were clipped to the Hounsfield Unit (HU) range $[-1024, 1500]$ and normalized to $[0, 1]$. All images and corresponding segmentations were resampled to isotropic 2mm spacing and a fixed spatial size of $224 \times 224 \times 224$ voxels, with image centers aligned to a common reference grid.

Data splitting and triplet sampling. Due to the large number of possible unordered triplets, $n(n-1)(n-2)/6$, we sampled 2048 unique triplets per split. Data were split at the subject level to prevent leakage. We performed five-fold cross-validation, holding out one fold for testing and splitting the remaining subjects into 90% training and 10% validation. Triplets were sampled independently within each split, ensuring that no subject appeared in multiple subsets.

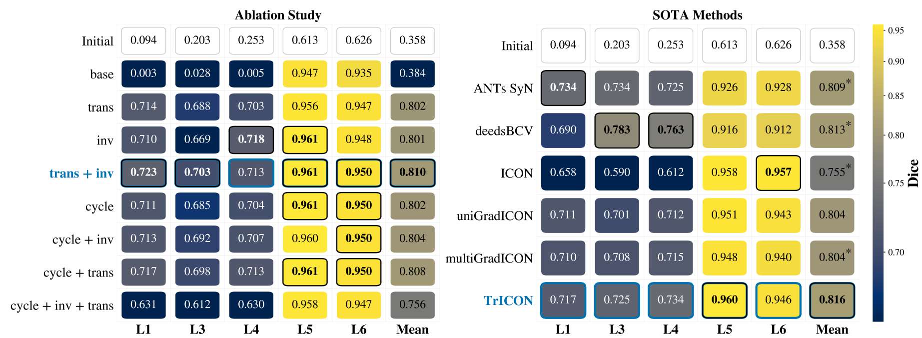

Fig. 3. Lung CT registration accuracy (DSC per label). **Left:** ablation of the proposed loss terms. **Right:** comparison with SOTA; `Initial` denotes pre-registration overlap. Cell colour indicates DSC. Best in bold, proposed method in blue; $^*p < .05$ vs. `TrICON`.

Architecture and implementation details. Our method places no requirements on the underlying registration network and is compatible with any backbone. In this work, we use the publicly available `ICON` network [13], which predicts deformation fields using a coarse-to-fine four-stage UNet architecture with three multiresolution refinement levels and a final full-resolution refinement stage. `TrICON` was implemented in PyTorch and trained on a single NVIDIA V100 GPU using Adam (10^{-4} learning rate) and a batch size of eight triplets for up to 100 epochs.

Metrics and statistical analysis Registration accuracy was evaluated using the Dice similarity coefficient (DSC) and the 95th percentile Hausdorff distance (HD95). For thin tubular structures (airway and vascular structures, L_1, L_3, and L_4), we additionally report median closest point distance (MCD), excluding L_2, the surface of L_1. For large volumetric organs (L_5 and L_6), we report the average symmetric surface distance (ASSD). Deformation regularity was assessed by the percentage of voxels with a negative Jacobian determinants ($\%|J| < 0$). Statistical significance was tested using two sided paired Wilcoxon signed rank tests on per subject mean values, with significance set at $p < 0.05$.

4 Results and Discussion

4.1 Ablation study

We evaluated on the validation set all seven nontrivial combinations of the three consistency losses introduced in Sec. 2.2 (`inv`: invertibility, `trans`: transitivity, `cycle`: cycle consistency), together with a similarity-only baseline (`base`: $\mathcal{L}_{\text{sim}}$). For each configuration, the regularization weight was selected independently by grid search over $\lambda \in [50, 250]$, jointly considering mean DSC, mean HD95 and

Table 1. Ablation over loss combinations using $\%|J| < 0$ and surface distance metrics [mm]: HD95 (all), MCD (L_1, L_3, L_4), and ASSD (L_5, L_6). Best values in bold.

Method	**HD95** [mm] (↓)						**MCD** [mm] (↓)				**ASSD** [mm] (↓)			$\%\lvert J\rvert_{<0}$ (↓)
	L1	L3	L4	L5	L6	Mean	L1	L3	L4	Mean	L5	L6	Mean	
Initial	30.78	27.67	28.00	39.09	39.58	33.03	12.82	10.22	10.47	11.17	15.79	16.07	15.93	–
base	57.20	52.90	52.95	32.85	31.76	45.53	29.81	37.65	33.15	33.54	7.79	8.16	7.97	49.4
trans	11.96	13.33	12.36	10.17	8.94	11.35	4.59	9.37	9.21	7.72	2.83	2.89	2.86	**.341**
inv	12.05	28.17	15.50	15.38	18.71	17.96	5.02	14.46	10.89	10.12	3.26	3.93	3.59	.861
trans + inv	**11.05**	**12.79**	**12.07**	**8.65**	**7.66**	**10.44**	**4.50**	**9.28**	**9.09**	**7.63**	**2.60**	**2.68**	**2.64**	.474
cycle	11.66	14.25	12.62	9.34	8.24	11.22	4.66	9.74	9.24	7.88	2.69	2.79	2.74	.624
cycle + inv	**11.05**	13.74	12.38	9.52	8.04	10.95	4.57	9.49	9.15	7.74	2.73	2.77	2.75	.579
cycle + trans	11.23	13.29	12.21	9.02	7.74	10.70	4.55	9.64	9.14	7.78	2.64	2.70	2.67	.591
cycle + inv + trans	13.49	15.90	14.92	9.93	9.18	12.69	5.53	10.75	10.48	8.92	2.77	2.93	2.85	.590

$\%|J| < 0$. The selected values were $\lambda = 100$ for `inv` and `cycle`, $\lambda = 250$ for `trans` and `inv+trans`, and $\lambda = 150$ for the remaining configurations.

Fig. 3 (left) reports DSC results. The `base` configuration performed well on the lungs but poorly on the airway and vascular structures, confirming that image similarity alone is insufficient for thin, high-curvature anatomy. Among all configurations, `inv+trans` achieved the highest mean DSC and the best overlap on the airway and vascular labels. Combining all three consistency losses `cycle+inv+trans` yielded the lowest performance among the regularized configurations. The surface distance metrics in Table 1 followed the same trend: `inv+trans` achieved the lowest mean HD95, MCD, and ASSD, whereas `base` remains the weakest configuration overall. Introducing any consistency loss reduced the fraction of folding voxels relative to `base` by two orders of magnitude while simultaneously improving DSC across all labels, confirming that consistency regularization alone is sufficient to recover near-diffeomorphic transformations in the absence of explicit spatial smoothing.

These results answer the question from Sec. 2.3: whether invertibility, cycle consistency, and transitivity impose complementary constraints, or whether some arise implicitly from others. The consistent advantage of `inv+trans` across all metrics indicates that invertibility and transitivity provide complementary regularization rather than redundant supervision. Cycle consistency alone improved substantially over `base` and performed on par with the other individual consistency terms, showing a useful constraint in isolation. However, adding it on top of the `inv+trans` provided no further benefit and instead reduced accuracy below every other regularized configuration. We hypothesize that this stems from $\mathcal{L}_{\text{cycle}}$ back-propagating gradients through a three-transformation composition rather than two; we did not verify this directly, and confirming it through a training-dynamics analysis (e.g., gradient norms across configurations) is left to future work. We therefore adopt `inv+trans` as the final `TrICON` model.

4.2 Comparison with SOTA

We compared `TrICON` against two classical registration methods, `ANTS SyN` [6] and `deedsBCV` [20]; the `ICON` model [13] retrained on our dataset; and the foun-

Table 2. SOTA comparison on inter-patient lung CT registration using $\%|J| < 0$ and surface distance metrics (mm): HD95 (all), MCD (L_1, L_3, L_4), and ASSD (L_5, L_6). Best values in bold; p-values vs. `TrICON` (* $p < .05$).

Method	**HD95** [mm] (↓)						**MCD** [mm] (↓)				**ASSD** [mm] (↓)			$\%\lvert J\rvert_{<0}$ (↓)
	L1	L3	L4	L5	L6	Mean	L1	L3	L4	Mean	L5	L6	Mean	
`Initial`	30.78	27.67	28.00	39.09	39.58	33.03*	12.82	10.22	10.47	11.17*	15.79	16.07	15.93*	–
`ANTS SyN`	15.30	13.93	14.23	11.59	11.16	13.24*	6.81	7.29	7.63	7.24	3.50	3.84	3.67*	**.000**
`deedsBCV`	12.90	12.30	12.68	12.27	11.64	12.36*	5.71	6.51	6.87	6.36	3.89	4.08	3.98*	.121
`ICON`	15.99	35.23	32.95	23.61	26.34	26.82*	5.95	18.93	18.20	14.36*	4.51	5.09	4.80*	.117
`uniGradICON`	14.19	**12.51**	12.82	9.63	8.43	11.52*	6.37	**6.83**	**7.22**	**6.81**	2.74	2.93	2.84*	.113
`multiGradICON`	14.47	12.86	13.15	9.70	8.89	11.81*	6.42	6.89	7.25	6.85	2.91	3.08	3.00*	.033
`TrICON`	**11.19**	12.99	**12.61**	**8.59**	**7.52**	**10.58**	**4.60**	9.53	9.21	7.78	**2.65**	**2.79**	**2.72**	.425

dation models `uniGradICON` [15] and `multiGradICON` [16], which were trained on substantially larger and more diverse datasets.

DSC results are reported in Fig. 3 (right) and surface distance metrics in Table 2. `TrICON` achieved the highest mean DSC and the lowest mean HD95 among all evaluated methods, together with the lowest mean ASSD, while remaining competitive on MCD. Compared with `ICON` baseline, it improves every surface metric, with the largest gains on HD95 for the pulmonary vasculature. While `TrICON` produces more folding voxels than the comparison methods, the fraction stays under 0.5%, keeping transformations close to diffeomorphic despite the improved surface-distance and overlap accuracy.

Two-sided paired Wilcoxon signed-rank tests on per-subject mean values showed that `TrICON` differed significantly from `ANTS SyN`, `multiGradICON` and `deedsBCV` on DSC, HD95 and ASSD but not on MCD. Against `uniGradICON`, differences were significant on HD95 and ASSD only.

Inference cost is a decisive factor for clinical deployment. The classical methods are markedly slower, with `ANTS SyN` and `deedsBCV` both requiring more than 10 minutes per image pair on CPU, whereas all learning based methods register a pair in under 10 seconds on an NVIDIA RTX 2000 GPU. `TrICON` achieves an average inference time of 6 seconds, a speed-up of two orders of magnitude over the classical baselines.

These results show that `TrICON` reaches performance competitive with foundation models trained on far larger and more heterogeneous data, despite being trained on a single dataset. Its improvement over the retrained `ICON` baseline further indicates that the observed gains arise from the proposed triplet consistency losses rather than from triplet sampling alone, and suggests that triplet consistency and large-scale pretraining may offer complementary benefits.

5 Conclusions

We introduced `TrICON`, a triplet-based extension of pairwise inverse-consistent registration that introduces transitivity as a population-level consistency constraint. By formulating and systematically evaluating inverse consistency, cycle consistency, and transitivity as stochastic, differentiable losses and evaluating all their combinations, we showed that inverse consistency and transitivity are complementary and jointly optimal among the tested configurations, whereas adding cycle consistency on top of both yields no further benefit and, under a shared regularization weight, degrades performance. The resulting model outperforms both classical and learning-based baselines, at an inference cost of 6 seconds. These findings indicate that triplet relationships provide supervision beyond what pairwise registration can capture, and that not all consistency constraints are equally beneficial in practice. Future work will explore gradient-based formulations of the proposed losses and evaluate `TrICON` across additional registration backbones, anatomical domains, and larger, more diverse image populations.

References

1. Jiang, J., Choi, C.M.S., Thor, M., Deasy, J.O., Veeraraghavan, H.: Tumor aware recurrent inter-patient deformable image registration of computed tomography scans with lung cancer. Medical Physics **52**, 938–950 (2024). https://doi.org/10.1002/mp.17536
2. Palma, G., Monti, S., Cella, L.: Voxel-based analysis in radiation oncology: A methodological cookbook. Physica Medica **69**, 192–204 (2020). https://doi.org/10.1016/j.ejmp.2019.12.013
3. Choi, C.M.S., Jiang, J., Mankuzhy, N.P., Nadkarni, N., Madhavan, S., Wu, A.J., Deasy, J.O., Thor, M., Rimner, A., Veeraraghavan, H.: Tumor-conditioned inter-patient registration using planning computed tomography for voxel-based analysis to predict radiation pneumonitis in lung cancer patients. Physics and Imaging in Radiation Oncology **37** (2026). https://doi.org/10.1016/j.phro.2026.100907
4. Vercauteren, T., Pennec, X., Perchant, A., Ayache, N.: Diffeomorphic demons: Efficient non-parametric image registration. NeuroImage **45**(1), S61–S72 (2009). https://doi.org/10.1016/j.neuroimage.2008.10.040, mathematics in Brain Imaging
5. Beg, M.F., Miller, M.I., Trouvé, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. International Journal of Computer Vision **61**(2), 139–157 (2005). https://doi.org/10.1023/B:VISI.0000043755.93987.aa
6. 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). https://doi.org/10.1016/j.media.2007.06.004
7. Christensen, G.E., Johnson, H.J.: Consistent image registration. IEEE Transactions on Medical Imaging **20**(7), 568–582 (2001). https://doi.org/10.1109/42.932742
8. Christensen, G.E., Johnson, H.J.: Invertibility and transitivity analysis for nonrigid image registration. Journal of Electronic Imaging **12**(1), 106–117 (2003). https://doi.org/10.1117/1.1526494
9. Škrinjar, O., Bistoquet, A., Tagare, H.: Symmetric and transitive registration of image sequences. International Journal of Biomedical Imaging (2008). https://doi.org/10.1155/2008/686875
10. Geng, X., Kumar, D., Christensen, G.E.: Transitive inverse-consistent manifold registration. In: Christensen, G.E., Sonka, M. (eds.) Information Processing in Medical Imaging. pp. 468–479. Springer Berlin Heidelberg, Berlin, Heidelberg (2005). https://doi.org/10.1007/11505730_39

11. Kim, B., Kim, D.H., Park, S.H., Kim, J., Lee, J.G., Ye, J.C.: CycleMorph: Cycle consistent unsupervised deformable image registration. Medical Image Analysis **71** (2021). https://doi.org/10.1016/j.media.2021.102036
12. Chatterjee, S., Bajaj, H., Siddiquee, I.H., Subbarayappa, N.B., Simon, S., Shashidhar, S.B., Speck, O., Nürnberger, A.: MICDIR: Multi-scale inverse-consistent deformable image registration using UNetMSS with self-constructing graph latent. Computerized Medical Imaging and Graphics (2023). https://doi.org/10.1016/j.compmedimag.2023.102267
13. Greer, H., Kwitt, R., Vialard, F.X., Niethammer, M.: ICON: Learning regular maps through inverse consistency. In: 2021 IEEE/CVF International Conference on Computer Vision (ICCV). pp. 3396–3405 (2021). https://doi.org/10.1109/ICCV48922.2021.00338
14. 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: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023). https://doi.org/10.1109/CVPR52729.2023.01734
15. 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: Linguraru, M.G., Dou, Q., Feragen, A., Giannarou, S., Glocker, B., Lekadir, K., Schnabel, J.A. (eds.) Medical Image Computing and Computer Assisted Intervention (MICCAI 2024). pp. 749–760. Springer (2024). https://doi.org/10.1007/978-3-031-72069-7_70
16. Demir, B., Tian, L., Greer, H., Kwitt, R., Vialard, F.X., Estépar, R.S.J., Bouix, S., Rushmore, R.J., Ebrahim, E., Niethammer, M.: MultiGradICON: A foundation model for multimodal medical image registration. In: International Workshop on Biomedical Image Registration (WBIR). pp. 3–18. Springer (2024). https://doi.org/10.1007/978-3-031-73480-9_1
17. Gu, D., Cao, X., Ma, S., Chen, L., Liu, G., Shen, D., Xue, Z.: Pair-wise and group-wise deformation consistency in deep registration network. In: Martel, A.L., Abolmaesumi, P., Stoyanov, D., Mateus, D., Zuluaga, M.A., Zhou, S.K., Racoceanu, D., Joskowicz, L. (eds.) International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI). pp. 171–180. Lecture Notes in Computer Science, Springer (2020). https://doi.org/10.1007/978-3-030-59716-0_17
18. Colin, J., Arnaud Arindra, A.S., Traverso, A., Van Ginneken, B.: LUNA16: Lung Nodule Analysis 2016. https://luna16.grand-challenge.org/Data/ (2016)
19. Liu, J., Zhang, Z., Niu, B., Kang, S., Ren, J., Wang, L., Xu, K.: A Custom Annotated Dataset for Segmentation of Pulmonary Veins, Arteries, and Airways. Scientific Data **12**, 1806 (2025). https://doi.org/10.1038/s41597-025-06074-6
20. Heinrich, M.P., Jenkinson, M., Brady, M., Schnabel, J.A.: MRF-based deformable registration and ventilation estimation of lung ct. IEEE Transactions on Medical Imaging **32**(7), 1239–1248 (2013). https://doi.org/10.1109/TMI.2013.2246577