Leveraging Multi-scale Cycle-consistency for Point-based Deformable Lung CT Registration

Nahyuk Lee[1] NAHYUK.LEE@VUNO.CO
Doohyun Park[1] DOOHYUN.PARK@VUNO.CO
Sungjoo Lim[1] SUNGJOO.LIM@VUNO.CO
Taemin Lee[2] KEVINLEE@KANGWON.AC.KR
[1] *Lung Vision AI Team, R&D Center, VUNO Inc., Republic of Korea*
[2] *Dept. of Electronic and AI System Engineering, Kangwon National University, Republic of Korea*

Editors: Under Review for MIDL 2025

Abstract

We propose a new point-based deformable lung CT registration network that integrates multi-scale cycle consistency into bidirectional, hierarchical deformable registration. Our method jointly estimates forward and backward flows while enforcing consistency across resolutions to learn anatomically coherent and invertible mappings. Experiments on the Lung250M benchmark demonstrate improved registration and robustness over unidirectional methods, establishing our approach as an effective and efficient registration baseline.

Keywords: Multi-scale Cycle Consistency, Point-based Deformable Lung CT Registration

1. Introduction

Deformable registration of inspiratory–expiratory lung computed tomography (CT) scans is critical for reducing respiratory motion artifact and planning radiotherapy. While voxel-based methods are widely used (Balakrishnan et al., 2019; Heinrich and Hansen, 2022), they are computationally intensive and sensitive to intensity variations. Point-based approaches offer a highly efficient alternative by directly modeling sparse anatomical structures; however, such methods remain underexplored for deformable lung CT registration.

While this point-based approach shares similar objectives with the general scene flow estimation task, existing methods (Myronenko and Song, 2010; Wu et al., 2020) rely on unidirectional flow prediction and lack mechanisms to enforce cycle consistency. These limitations restrict their ability to produce coherent and invertible deformations, thereby posing challenges for applications in anatomically sensitive medical images. Although some voxel-based methods, *e.g.* Kim et al. (2021), incorporate cycle consistency in lung deformation prediction, it is typically enforced only at a single resolution, rather than across multiple scales, *e.g.*, from coarse to fine. As a result, ensuring anatomically plausible deformations remains challenging when consistency across scales is not maintained.

In the context of deformable lung CT registration, enforcing consistency across hierarchical levels is essential for capturing both global anatomical alignment and fine-grained local deformations. Without such consistency, predictions made at different resolutions can conflict, undermining the geometric coherence of the overall transformation.

From these motivations, we propose a new point-based deformable lung CT registration framework equipped with a *multi-scale cycle consistent objective*. Building upon Point-PWC Net (Wu et al., 2020), our framework hierarchically estimates forward and backward deformations and enforces consistency across multiple resolution levels, leading to anatomically coherent and invertible mappings under an unsupervised setting. Our contributions are summarized as follows:

- We present a new ***point-based*** deformable lung CT registration framework, showing promising performance while highlighting its potential on the Lung250M benchmark.
- We introduce a ***multi-scale cycle consistent objective*** that enforces geometric consistency across multiple resolution levels throughout the hierarchy.

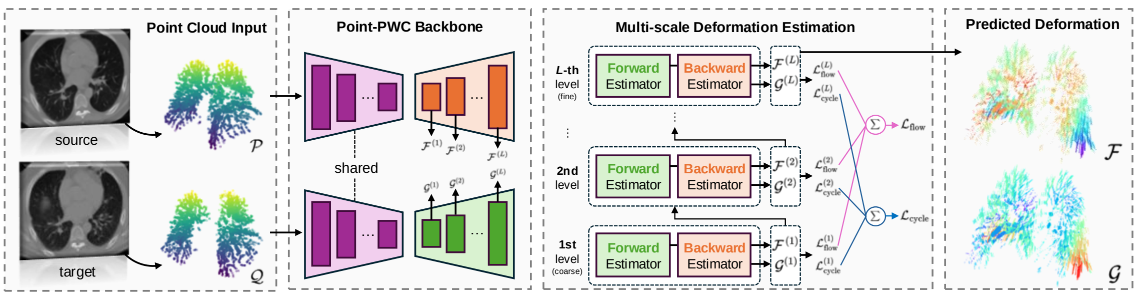

Figure 1: Bidirectional point-based network for hierarchical deformation estimation.

2. Proposed Approach

Our method formulates point-based deformable lung CT registration as a bi-directional dense flow estimation problem. Given source and target point clouds $\mathcal{P} = \{\mathbf{p}_i\}_{i=1}^{N}$ and $\mathcal{Q} = \{\mathbf{q}_j\}_{j=1}^{M}$, we aim to learn a pair of forward and backward deformation fields $\mathcal{F} = \{f_i\}_{i=1}^{N}$ and $\mathcal{G} = \{g_j\}_{j=1}^{M}$ that minimize the alignment error over ground-truth correspondences $\mathcal{C}$:

$$\mathcal{F}^* = \arg\min_{\mathcal{F}} \sum_{(i,j)\in\mathcal{C}} \|f_i(\mathbf{p}_i) - \mathbf{q}_j\|_2^2, \qquad \mathcal{G}^* = \arg\min_{\mathcal{G}} \sum_{(i,j)\in\mathcal{C}} \|\mathbf{p}_i - g_j(\mathbf{q}_j)\|_2^2. \tag{1}$$

For joint optimization of bi-directional deformations, we extend the Point-PWC Net (Wu et al., 2020) by introducing *symmetric decoders* for each prediction[1], as shown in Figure 1. For each resolution level $l \in \{1, \ldots, L\}$, both decoders are supervised with a flow loss $\mathcal{L}_{\text{flow}}$:

$$\mathcal{L}_{\text{flow}}^{(l)} = \sum_{\mathbf{p}_i\in\mathcal{P}^{(l)}} \|\hat{f}_i^{(l)}(\mathbf{p}_i) - f_i^{(l)}(\mathbf{p}_i)\|_2 + \sum_{\mathbf{q}_j\in\mathcal{Q}^{(l)}} \|\hat{g}_j^{(l)}(\mathbf{q}_j) - g_j^{(l)}(\mathbf{q}_j)\|_2, \tag{2}$$

where $\hat{f}_i^{(l)}$ and $\hat{g}_j^{(l)}$ denote the predicted deformation fields, and $f_i^{(l)}$, $g_j^{(l)}$ are the ground-truth fields at level l, respectively. To further regularize the model and promote geometric reversibility, we introduce a *multi-scale cycle consistency loss* at each level by applying sequential forward and backward flows:

$$\mathcal{L}_{\text{cycle}}^{(l)} = \text{CD}\left(\mathbf{T}_{\hat{\mathcal{G}}}^{(l)} \circ \mathbf{T}_{\hat{\mathcal{F}}}^{(l)}(\mathcal{P}^{(l)}), \mathcal{P}^{(l)}\right) + \text{CD}\left(\mathbf{T}_{\hat{\mathcal{F}}}^{(l)} \circ \mathbf{T}_{\hat{\mathcal{G}}}^{(l)}(\mathcal{Q}^{(l)}), \mathcal{Q}^{(l)}\right), \tag{3}$$

1. We refer the readers to the work of Wu et al. (2020) for further network details.

Table 1: Quantitative comparison.

Method	TRE	25%	50%	75%
initial	16.25	10.14	15.94	21.76
CPD (Myronenko and Song, 2010)	3.13	1.51	2.28	3.58
Point-PWC (Wu et al., 2020)	2.85	1.52	2.33	3.54
Ours	**2.24**	**1.16**	**1.79**	**2.77**

Table 2: Ablation studies.

Ref.	Bi-dir. Estimation	Cycle Consistency	TRE	25%	50%	75%
(a)	✗	✗	2.54	1.29	2.10	3.12
(b)	✓	✗	2.32	1.25	2.01	2.98
Ours	✓	✓	**2.24**	**1.16**	**1.79**	**2.77**

where $\mathbf{T}^{(l)}_{\hat{\mathcal{F}}}(\cdot)$ and $\mathbf{T}^{(l)}_{\hat{\mathcal{G}}}(\cdot)$ denote transformation functions that deform a point cloud by applying the corresponding deformation fields; for instance, $\mathbf{T}^{(l)}_{\hat{\mathcal{F}}}(\mathcal{P}^{(l)}) = \sum_i \hat{f}_i(\mathbf{p}_i^{(l)})$. $\mathrm{CD}(\cdot,\cdot)$ denotes the Chamfer Distance. The final objective aggregates the flow and cycle consistency losses over all levels: $\mathcal{L} = \sum_{l=1}^{L}\left(\mathcal{L}^{(l)}_{\text{flow}} + \lambda \cdot \mathcal{L}^{(l)}_{\text{cycle}}\right)$, where λ set to 0.1.

Despite its structural simplicity, this extension offers a key advantage: it enables explicit enforcement and observation of cycle consistency at multiple scales. As a result, the model learns anatomically coherent, invertible deformations with improved stability and accuracy across hierarchical resolutions.

3. Experimental Results

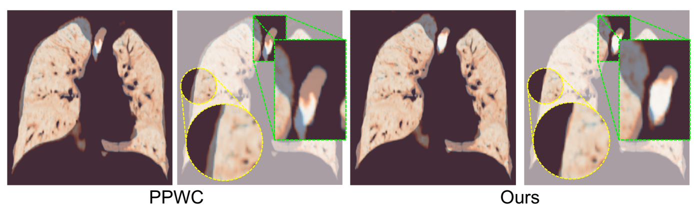

Figure 2: Qualitative comparison.

Table 1 present the quantitative comparisions with existing point-based methods. Following evaluation protocol of Falta et al. (2023), we evaluate our model[2] with Target Registration Error (TRE) as the primary metric. Mean and $25, 50, 75$ percentile TREs are reported to capture registration error across varying anatomical regions and difficulty levels. Our model consistently outperforms both optimization-based method, *e.g.* CPD (Myronenko and Song, 2010), and learning-based method, *e.g.* Point-PWC (PPWC) (Wu et al., 2020), across all metrics with achieving a mean TRE of 2.24 mm. These results highlight the effectiveness of our bi-directional architecture and multi-scale cycle consistency regularization. Figure 2 presents the qualitative comparison with PPWC, highlighting improved registration while preserving anatomical structures.

Further, to assess the contribution of each component, we conduct ablation studies as shown in Table 2. Introducing bidirectional estimation alone (Tab. 2b) yields performance gains over the unidirectional baseline (Tab. 2a), while further incorporating multi-scale cycle consistency leads to the best registration result.

4. Conclusion

We propose a ***point-based framework*** for deformable lung CT registration that incorporates ***multi-scale cycle consistency*** with bidirectional dual objectives. By extending Point-PWC Net with symmetric decoders and enforcing consistency across hierarchical levels, our approach produces anatomically coherent and invertible mappings. Experiments on Lung250M confirm its efficacy, offering a strong baseline for future lung CT registration.

2. Our model is trained for 1,500 epochs using Adam (Kingma and Ba, 2014) with a learning rate of 1e-3 and batch size 4. Ground-truth correspondences are obtained from corrField (Heinrich et al., 2015).

References

Guha Balakrishnan, Amy Zhao, Mert R Sabuncu, John Guttag, and Adrian V Dalca. Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging*, 38(8):1788–1800, 2019.

Fenja Falta, Christoph Großbröhmer, Alessa Hering, Alexander Bigalke, and Mattias Heinrich. Lung250m-4b: a combined 3d dataset for ct-and point cloud-based intra-patient lung registration. *Advances in Neural Information Processing Systems*, 36:54819–54832, 2023.

Mattias P Heinrich and Lasse Hansen. Voxelmorph++ going beyond the cranial vault with keypoint supervision and multi-channel instance optimisation. In *International workshop on biomedical image registration*, pages 85–95. Springer, 2022.

Mattias P Heinrich, Heinz Handels, and Ivor JA Simpson. Estimating large lung motion in copd patients by symmetric regularised correspondence fields. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part II 18*, pages 338–345. Springer, 2015.

Boah Kim, Dong Hwan Kim, Seong Ho Park, Jieun Kim, June-Goo Lee, and Jong Chul Ye. Cyclemorph: cycle consistent unsupervised deformable image registration. *Medical image analysis*, 71:102036, 2021.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Andriy Myronenko and Xubo Song. Point set registration: Coherent point drift. *IEEE transactions on pattern analysis and machine intelligence*, 32(12):2262–2275, 2010.

Wenxuan Wu, Zhi Yuan Wang, Zhuwen Li, Wei Liu, and Li Fuxin. Pointpwc-net: Cost volume on point clouds for (self-) supervised scene flow estimation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 88–107. Springer, 2020.