Topology-Constrained Deep Learning for Automated Segmentation of the Left Anterior Descending Coronary Artery for Thoracic Radiotherapy

2026 JULY 19–22 VANCOUVER, BC
JOINT AAPM | COMP MEETING

Yong Yue, PhD[1], Christopher Njeh, PhD[2], Tim Lautenschlaeger, MD[2]
[1] Memorial Sloan Kettering Cancer Center, 1275 York Ave, New York, NY 10065
[2] Indiana University School of Medicine, 535 Barnhill Dr, Indianapolis, IN 46202

INTRODUCTION

Accurate LAD segmentation is important for cardiac sparing in thoracic radiotherapy. However, the LAD is thin, low-contrast, and frequently obscured on non-contrast planning CT. Standard CNN-based segmentation may produce discontinuous or truncated vessel contours. This study developed a topology-constrained 3D deep learning workflow to improve LAD continuity and anatomical completeness.

METHODS AND MATERIALS

Dataset and preprocessing. Thoracic RT planning CT scans from 46 lung cancer patients (mixed contrast and non-contrast) with expert LAD contours were used. Each scan was normalized within a heart ROI using robust median standardization, and cases were evaluated using 5-fold cross-validation.

Anatomical priors. A multiscale dual-polarity Frangi vesselness map was generated from the normalized CT to enhance tubular structures. The network input included two channels: normalized CT and vesselness.

Network and topology-constrained loss. A 3D U-Net was trained to segment the LAD. Beyond binary cross-entropy (BCE) and Dice, the composite objective incorporates: (i) soft centerline Dice (clDice) computed on differentiable skeletons of the prediction and label; (ii) a skeleton-length prior discouraging both truncation and spurious elongation; (iii) endpoint-probability enforcement with distal emphasis to preserve the thin distal segment; and (iv) surface-based regularization. Loss-term weights were ramped over a warm-up via a scheduler.

Composite loss

$$\mathcal{L} = w_{\text{bce}}\mathcal{L}_{\text{BCE}} + w_{\text{dice}}\mathcal{L}_{\text{Dice}} + w_{\text{cl}}\mathcal{L}_{\text{clDice}} + w_{\text{len}}\mathcal{L}_{\text{len}} + w_{\text{end}}\mathcal{L}_{\text{end}} + w_{\text{fus}}\mathcal{L}_{\text{surf}}$$

weights ramped $w \propto t$ over warm-up (e.g. $w_{\text{cl}}: 0.005 \rightarrow 0.10$, $w_{\text{end}}: 0.003 \rightarrow 0.05$, $w_{\text{len}}: 0.002 \rightarrow 0.04$)

Centerline Dice (clDice)

$$T_{\text{prec}} = \frac{|S_P \cap G|}{|S_P|}, \quad T_{\text{sens}} = \frac{|S_G \cap P|}{|S_G|}, \quad \text{clDice} = \frac{2\,T_{\text{prec}}\,T_{\text{sens}}}{T_{\text{prec}} + T_{\text{sens}}}$$

S_P, S_G = morphological skeletons of prediction P and ground truth G

Training. The network was optimized with Adam (learning rate 3×10^{-4}) under a cosine schedule, using foreground-biased 64^3 patch sampling (85% of crops centered on a labeled vessel voxel) with flip and rotation augmentation to concentrate gradient on the sparse LAD.

Inference and evaluation. Full-volume LAD predictions were generated using Gaussian-weighted sliding-window inference over overlapping patches. Predictions were then refined using label-free topology-aware post-processing, including vesselness- and geometry-based component selection with guarded reconnection, to enforce a single continuous LAD contour. Performance was evaluated per case and across the cohort using hard Dice, clDice, HD95, and connected-component count as a fragmentation index.

WORKFLOW

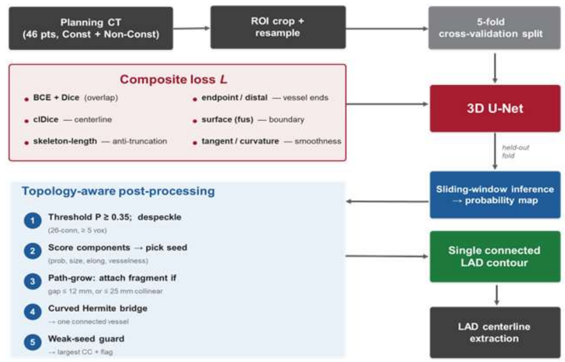

Figure 1. LAD Segmentation workflow

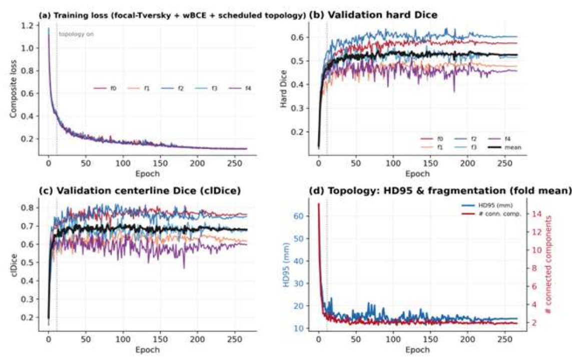

Figure 2. Five-fold training dynamics: (a) composite loss; (b) validation hard Dice with cross-fold mean; (c) centerline Dice; (d) fold-mean HD95 and connected-component count. The dotted line marks where the scheduled topology weight reaches full strength.

Table 1. Held-out test performance by cross-validation fold.

Fold	N	Hard	clDice	HD95 (mm)	nCC
f0	10	0.570	0.798	8.9	1.40
f1	9	0.485	0.653	20.8	1.44
f2	9	0.628	0.829	6.4	1.11
f3	9	0.524	0.745	12.6	1.11
f4	9	0.481	0.658	22.5	1.33
Overall	46	0.538	0.738	14.1	1.28

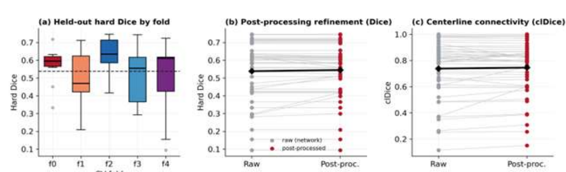

Figure 3. Held-out test results (N=46): (a) hard Dice by fold; (b) paired raw-to-post-processed hard Dice and (c) paired raw-to-post-processed clDice.

RESULTS

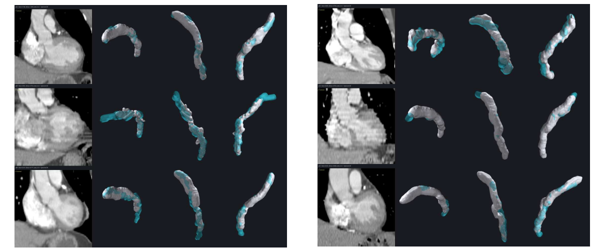

LAD in contrast CT, white = prediction, blue = GT

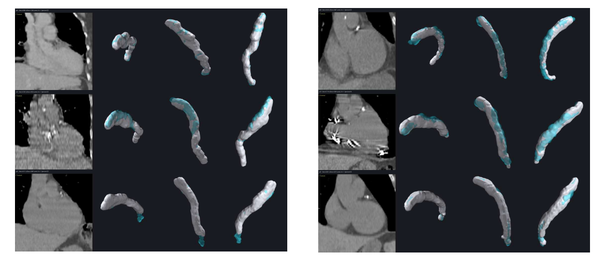

LAD in non contrast CT, white = prediction, blue = GT

Figure 4. Representative cases of LAD segmentations for contrast and non-contrast CT

The topology-constrained framework improved LAD segmentation accuracy and anatomical completeness. In 5-fold cross-validation (N = 46), sliding-window inference achieved a hard Dice of 0.538 ± 0.164, clDice of 0.738 ± 0.224, HD95 of 14.1 ± 5.6 mm, and a mean of 1.28 connected components per case, indicating that the network output was already largely continuous. Performance was similar between contrast and non-contrast CT cases, with hard Dice values of 0.541 and 0.530, respectively. Given that the network predictions were mostly connected, Post-processing served mainly as a topological guarantee, enforcing a single connected component in every case and providing modest accuracy refinement.

CONCLUSIONS

A topology-constrained 3D deep learning framework improved automated LAD segmentation on thoracic RT planning CT and produced anatomically continuous predictions. This approach provides a practical foundation for routine LAD contour generation and subsequent LAD dosimetry extraction, supporting future investigations of LAD dose response relationships and associations with post-RT cardiac events.

REFERENCES

1. Ronneberger O, et al. U-Net: convolutional networks for biomedical image segmentation. MICCAI 2015.
2. Cicek O, et al. 3D U-Net: learning dense volumetric segmentation from sparse annotation. MICCAI 2016.
3. Shit S, et al. clDice: a topology-preserving loss for tubular structure segmentation. CVPR 2021.