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

Towards Improving Clinical Practice to Evaluate Rib Endpoints with Autosegmentation

AM. Haidari[1], E. Babadagli[2], P. Mehra[3], R. Zohr[1], C. Church[1]

[1]The Ottawa Hospital, Ottawa, Canada; [2]Princess Margaret Cancer Centre, Toronto, Canada; [3]Western Memorial Regional Hospital, Corner Brook, Canada

ABSTRACT

Current thoracic SBRT planning often uses a chest wall contour as a surrogate for ribs, which introduces uncertainty into rib dose assessment and fracture-risk modeling. We evaluated a treatment-planning-system-integrated, two-stage DGCNN workflow for automated rib segmentation from thoracic CT. The model was trained on 320 public cases, validated on 50 held-out cases, and clinically reviewed in 10 local patients by three clinicians. The scripted workflow generated rib contours in approximately 70 seconds per scan. Analysis of the two-stage model showed strong rib-isolation performance. Clinical review indicated that contours generally required minor edits but less effort than manual contouring. These results support TPS-integrated rib autosegmentation as a practical workflow for rib-specific dose assessment in thoracic SBRT, with future work focused on local fine-tuning and toxicity modeling.

CONTACT

Mehan Haidari
The Ottawa Hospital
mhaidari@toh.ca

INTRODUCTION

Stereotactic body radiation therapy (SBRT) delivers ablative doses for lung cancer but may expose adjacent ribs to clinically significant doses. Rib fracture remains a clinically relevant late toxicity. Recent evidence in patients treated with three-fraction lung SBRT reported a 19% cumulative rib fracture risk, with relatively low risk during the first 6 months after treatment and cumulative risk increasing rapidly between 6 and 54 months post-treatment.[1]

Rib dose-volume parameters, including near-maximum dose, V_{30}, and V_{40}, have been strongly associated with chest wall pain and rib fracture.[1–3] Current clinical practice often uses the chest wall contour as a surrogate for individual ribs, introducing geometric and dosimetric uncertainty into rib fracture risk modeling.

Although ribs are visible on computed tomography (CT), individual rib labeling remains tedious because ribs are thin, elongated, oblique, and adjacent to the vertebrae, sternum, and other osseous structures. Public labeled rib datasets, including *RibSeg* and *RibSeg v2*, have enabled point-cloud-based methods for individual rib segmentation and labeling.[4,5] Dynamic Graph CNN (DGCNN) architectures are well suited to this task because they learn local and semantic geometric relationships in point-cloud data through dynamically updated graph neighborhoods.[6]

This work evaluates the geometric performance and clinical feasibility of an in-house trained two-stage DGCNN for automated individual rib segmentation deployed through a treatment planning system (TPS) script.

METHODS AND MATERIALS

A two-stage DGCNN model for rib autosegmentation was trained on an open-source thoracic CT dataset using published point-cloud methodology.[4,5] The public cohort was split into 320 training and 50 validation cases; 10 local cases were used for clinical evaluation.

Figure 1 illustrates the Stage 1 workflow. In Stage 1, 3D CT voxels are converted to a thresholded point cloud, randomly downsampled, and processed by a shared DGCNN to isolate ribs from background anatomy. In Stage 2, the Stage 1 rib output is processed through a similar point-cloud workflow to label individual ribs. These labels informed the final composite rib structure by preserving separation between adjacent ribs.

Predictions were made within the RayStation scripting API with model-inferencing being offloaded to a local server.

Geometric accuracy was quantified using stage-specific Dice Similarity Coefficient (DSC), candidate-point accuracy, and rib-label accuracy; clinical utility was assessed by three clinicians using a 5-point Likert scale in 10 previously treated local patients.

RESULTS

Figure 2 illustrates representative final rib segmentations on local thoracic CT scans. Stage-specific validation metrics are reported to distinguish Stage 1 rib isolation from Stage 2 individual rib labeling. Metrics were calculated without additional morphological processing. The TPS-scripted workflow generated the final rib structure in approximately 70 seconds per scan.

In the held-out validation cohort (N = 50), Stage 1 rib isolation achieved a DSC of 0.89 ± 0.02 and candidate-point accuracy of 97 ± 1%. End-to-end individual rib labeling achieved a per-rib DSC of 0.76 ± 0.19 and rib-label accuracy of 77 ± 19%. When Stage 2 was evaluated independently using ground-truth rib foreground points, performance improved to a per-rib DSC of 0.89 ± 0.22 and rib-label accuracy of 89 ± 21%.

Across 30 clinical evaluations, the median rating was 3/5 (range 3–4); 22 indicated necessary minor edits requiring less time than recontouring, while 8 indicated only nonessential stylistic edits.

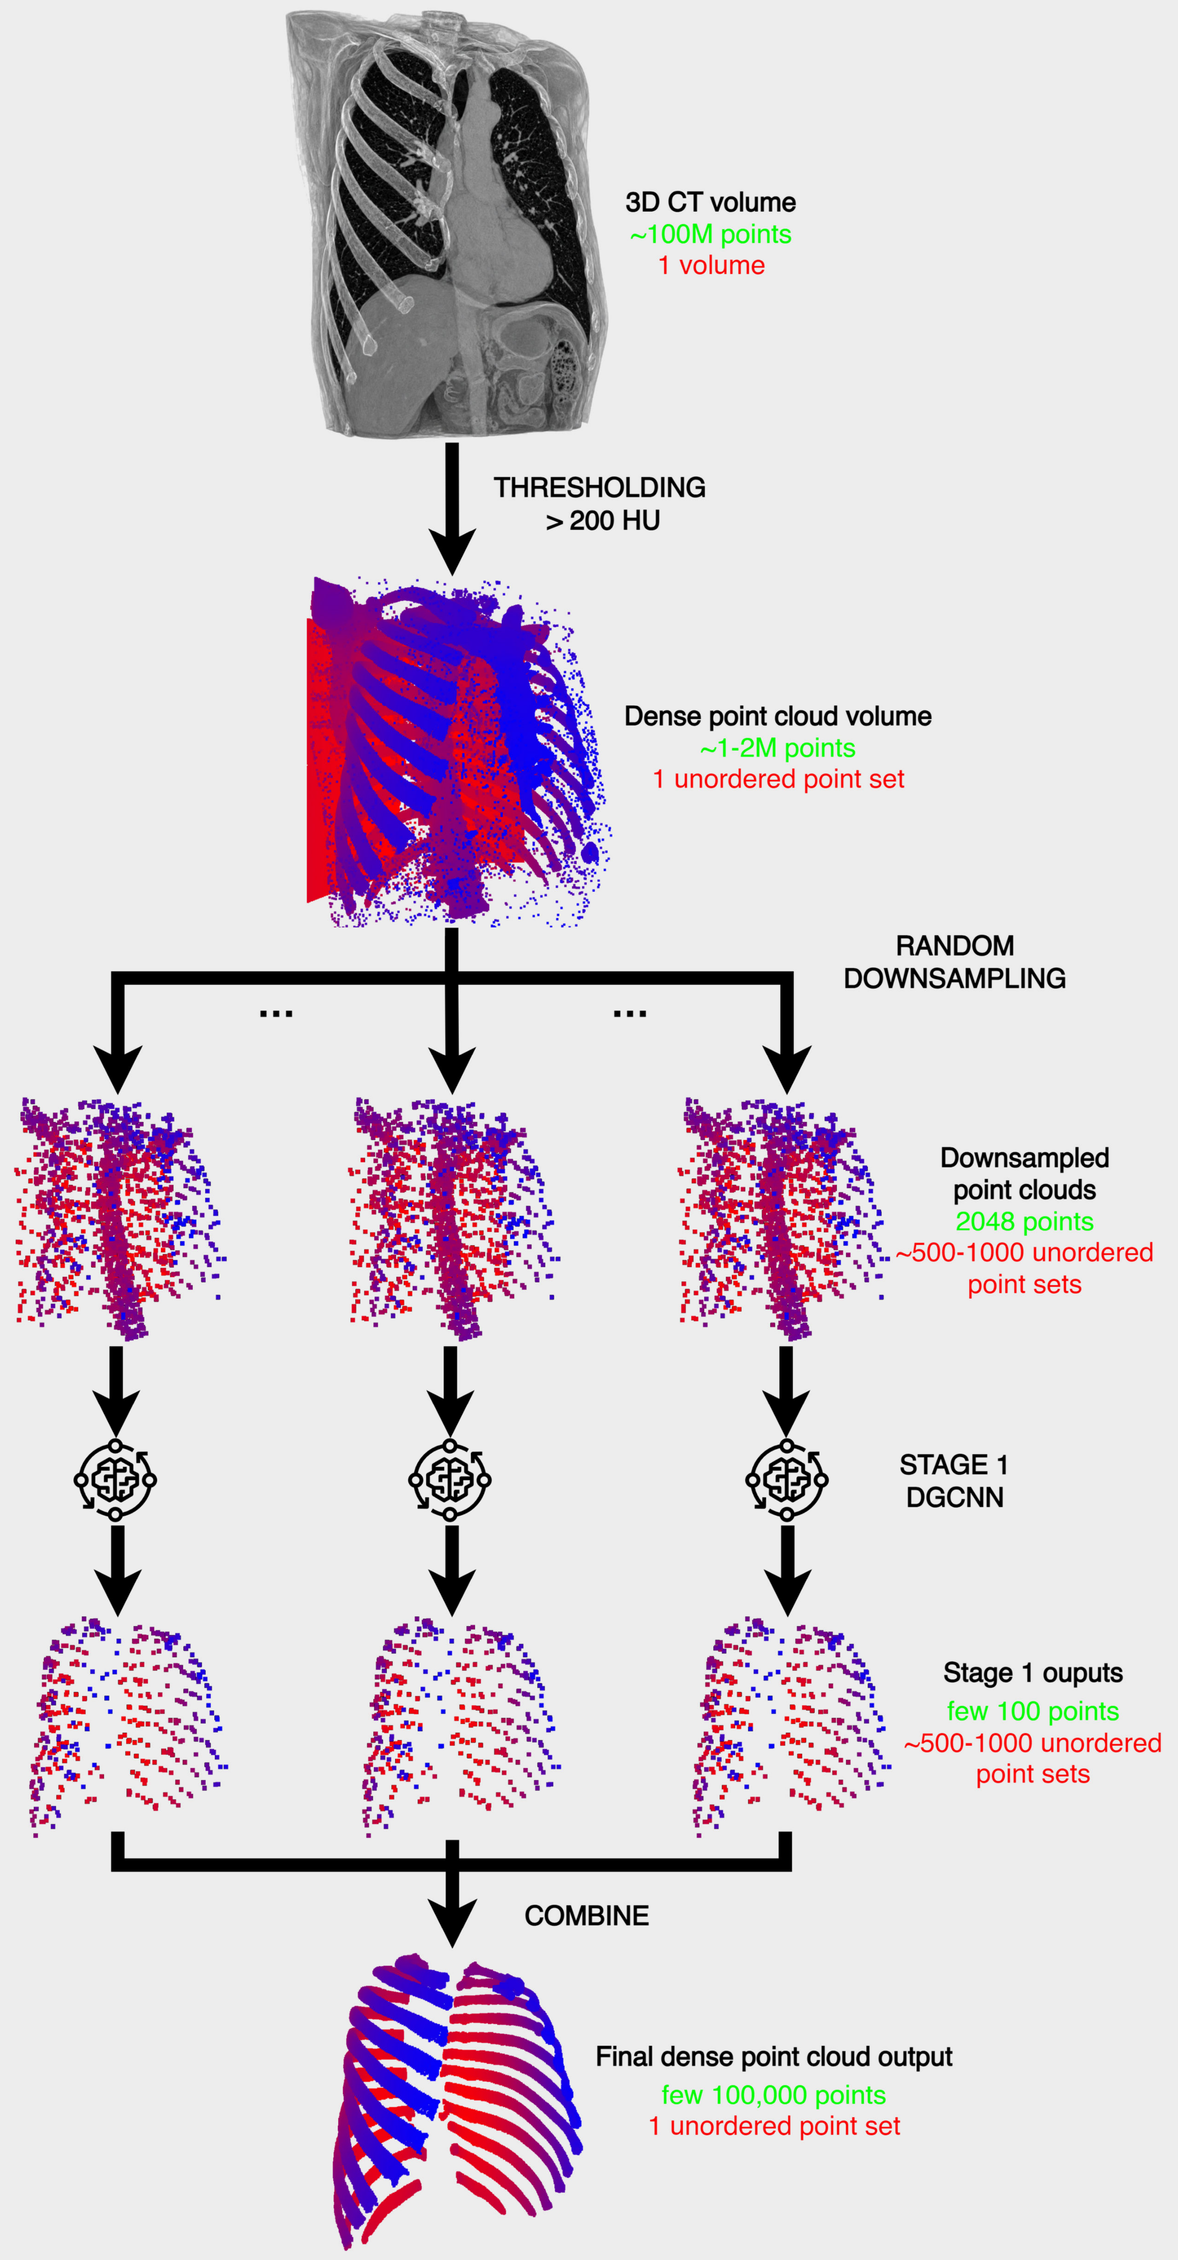

Figure 1: CT voxels are thresholded to form a point cloud, downsampled for DGCNN inference with a shared model, and aggregated to generate a final composite rib structure with adjacent ribs preserved as separate regions, following [4].

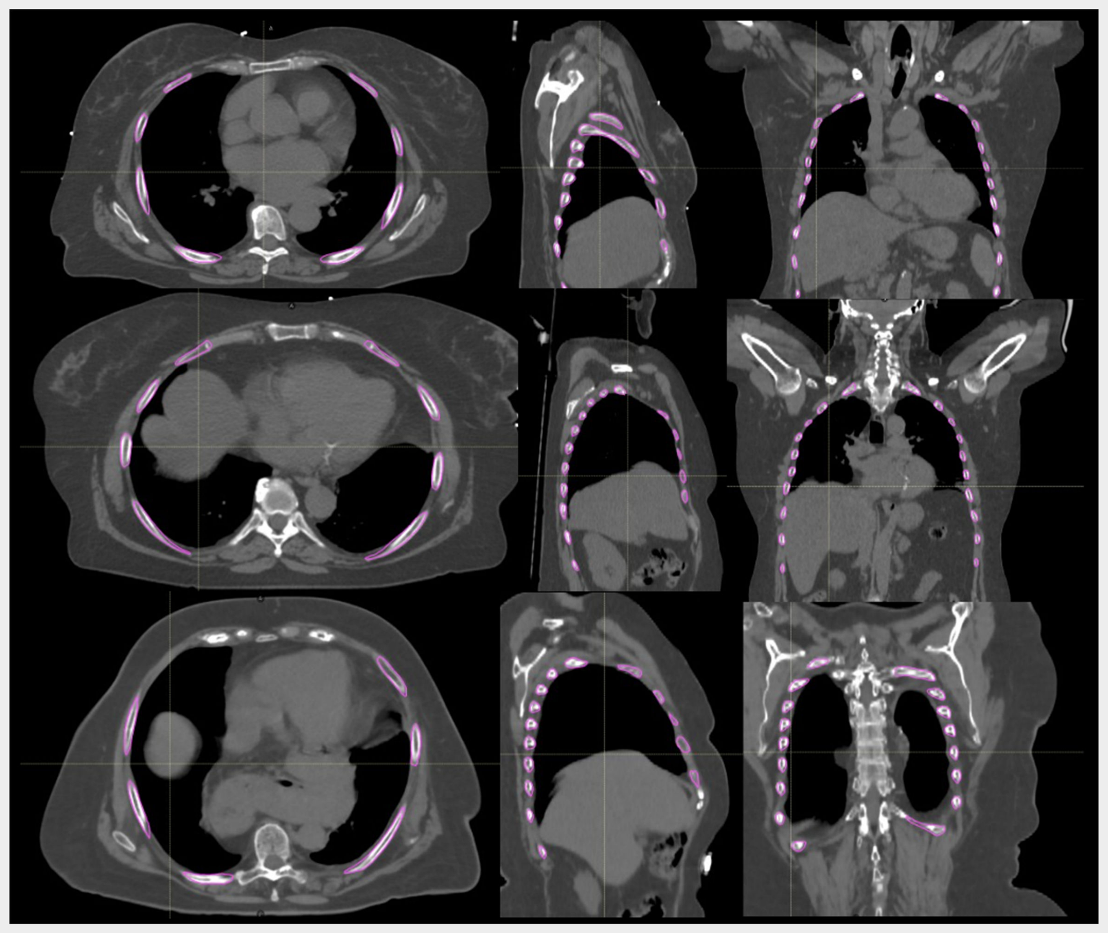

Figure 2: Final composite rib contours on local CT data are shown in magenta across three views. Each row represents a single patient.

DISCUSSION

Stage-specific results indicate that the model performs well for rib-volume identification, while individual rib labeling remains more challenging. Stage 2 was used primarily to preserve separation between adjacent ribs and produce a composite rib contour for clinical use.

Stage 1 achieved high rib-isolation accuracy; however, the lower end-to-end per-rib DSC suggests that upstream segmentation errors can propagate into Stage 2 and affect rib identity assignment. Improved Stage 2 performance using ground-truth rib foreground points indicates that both rib isolation and label classification contribute to the final error.

Clinically, the approximately 70-second runtime makes generating a composite rib structure feasible within the treatment-planning workflow, and clinician review indicated that required edits were less burdensome than manual contouring from scratch. Remaining errors were concentrated near the sternocostal and costovertebral joints that were not included in the training data.

CONCLUSIONS

The two-stage DGCNN enabled rapid rib autosegmentation within the treatment-planning workflow. Future work will focus on local fine-tuning, improved robustness at rib junctions, and evaluating whether rib contours improve rib dose assessment and toxicity modeling after thoracic SBRT.

REFERENCES

1. Bentsen KK, Brink C, Nielsen TB, et al. **Cumulative rib fracture risk after stereotactic body radiotherapy in patients with localized non-small cell lung cancer.** *Radiotherapy and Oncology.* 2024;200:110481.
2. Pettersson N, Nyman J, Johansson KA. **Radiation-induced rib fractures after hypofractionated stereotactic body radiation therapy of non-small cell lung cancer: a dose- and volume-response analysis.** *Radiotherapy and Oncology.* 2009;91(3):360–368.
3. Kita N, Tomita N, Takaoka T, et al. **Symptomatic radiation-induced rib fractures after stereotactic body radiotherapy for early-stage non-small cell lung cancer.** *Clinical and Translational Radiation Oncology.* 2023;43:100683.
4. Yang J, Gu S, Wei D, Pfister H, Ni B. **RibSeg Dataset and Strong Point Cloud Baselines for Rib Segmentation from CT Scans.** *MICCAI.* 2021.
5. Jin L, Gu S, Wei D, et al. **RibSeg v2: A large-scale benchmark for rib labeling and anatomical centerline extraction.** *IEEE Transactions on Medical Imaging.* 2024.
6. Wang Y, Sun Y, Liu Z, Sarma SE, Bronstein MM, Solomon JM. **Dynamic Graph CNN for Learning on Point Clouds.** *ACM Transactions on Graphics.* 2019;38(5):146.