RADIANT: A Fully Configurable Radiotherapy Dose Prediction Framework

Josiane Pafeng[1]; Adrian Celaya[2,3]; Skylar Gay[1] ; David Fuentes[1]; Mary Martel[1]; Laurence Court[1]; Tucker Netherton[1]

[1]The University of Texas MD Anderson Cancer Center , [2]Rice University, [3]Google LLC

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

PURPOSE

This work presents the Radiotherapy Dose Inference and Analysis Toolkit (RADIANT), an open-source fully configurable framework for deep learning-based dose prediction that facilitates reproducible experimentation across data preprocessing, various architectures, loss functions and training strategies.

METHODS

We applied RADIANT on cervix and prostate data created using the Radiation Planning Assistant (RPA), a web-based FDA 510(k)-cleared treatment planning system, and head and neck data from AAPM OpenKBP challenge.

- Given a training set, we applied a variety of architectures, parameters and optimization strategies to find the best dose prediction model for our data.
  - Cervix (40 models evaluated)
  - Prostate (4 models evaluated).
  - Head and neck(16 models evaluated)

RESULTS AND DISCUSSION

- **Cervix (training set)**: **(nnU-Net, MAE, Cosine)**, dose score (1.31), D_{95} error (0.66% of maximum prescription).
  - Mean dose error across all 8 models was within 1% for PTVs and 2% for OARs with respect to maximum prescription (Fig. 1).
  - Good overall performance on test set (Fig. 2).
- **Prostate (training set)**: **(nnU-Net, MAE, Cosine)**, dose score (0.88), D_{95} error (1.07% of maximum prescription).
  - Model captured details of the target and OARs on test set (Fig. 3).
- **Head and neck (training set)**: **(nnU-Net, wMAE, Polynomial)**, dose score (2.9507), DVH score (1.6801).
  - Best performance on test set (Fig. 4) ranking RADIANT 6th for dose score (2.702) and **2nd for DVH score (1.495)** on the AAPM OpenKBP leaderboard.

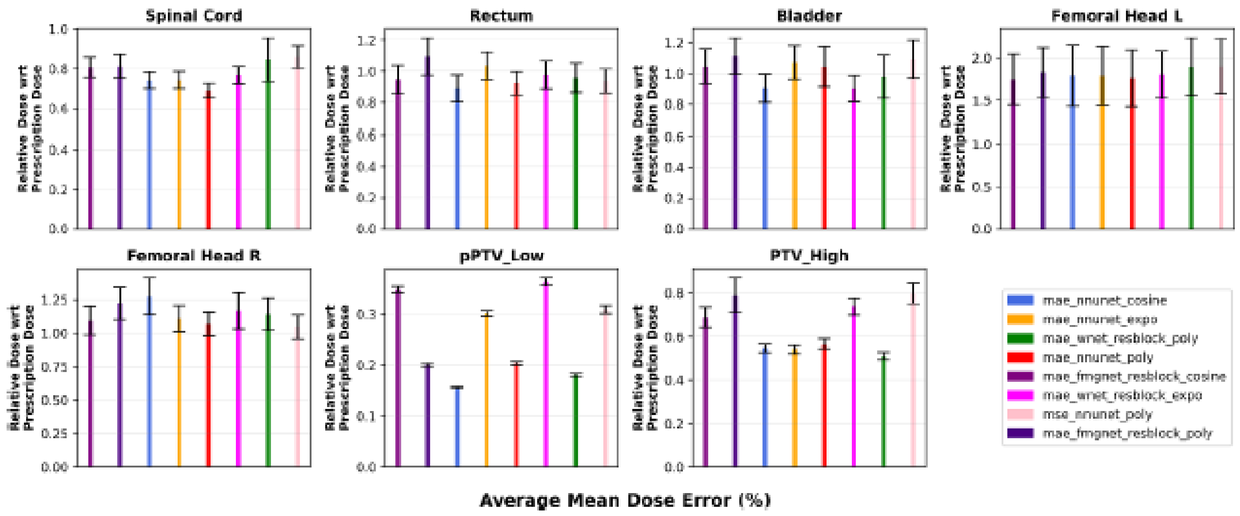

Fig. 1 Prediction of relative errors in mean dose for targets and OARs

MODELS TESTED

- **Architectures**: nnU-Net, Swin UNETR, FMG-Net, W-Net and ddU-Net.
- **Loss Functions:** MAE and MSE.
- **Learning Rates**: Cosine Annealing with warm restarts, Cosine Annealing without warm restarts, Exponential decay and Polynomial decay.
  - nnU-Net ('pocket' version) used fewer parameters than full nnU-Net (~700k vs ~30M).

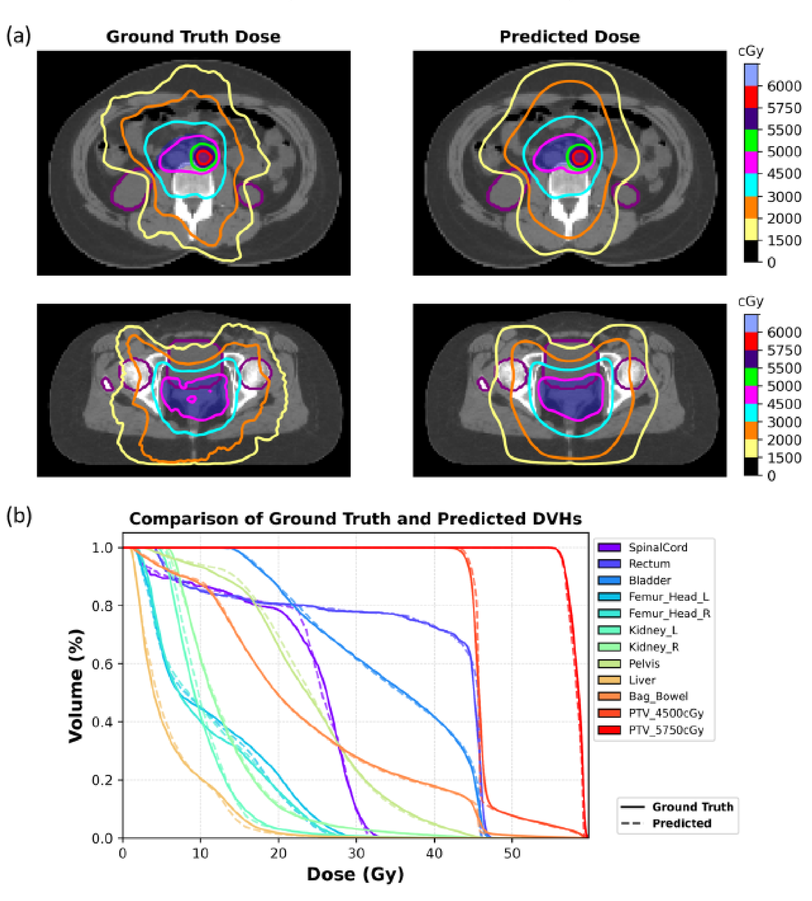

Fig. 2 Comparison between ground truth RPA data and predicted dose for a cervix test patient.

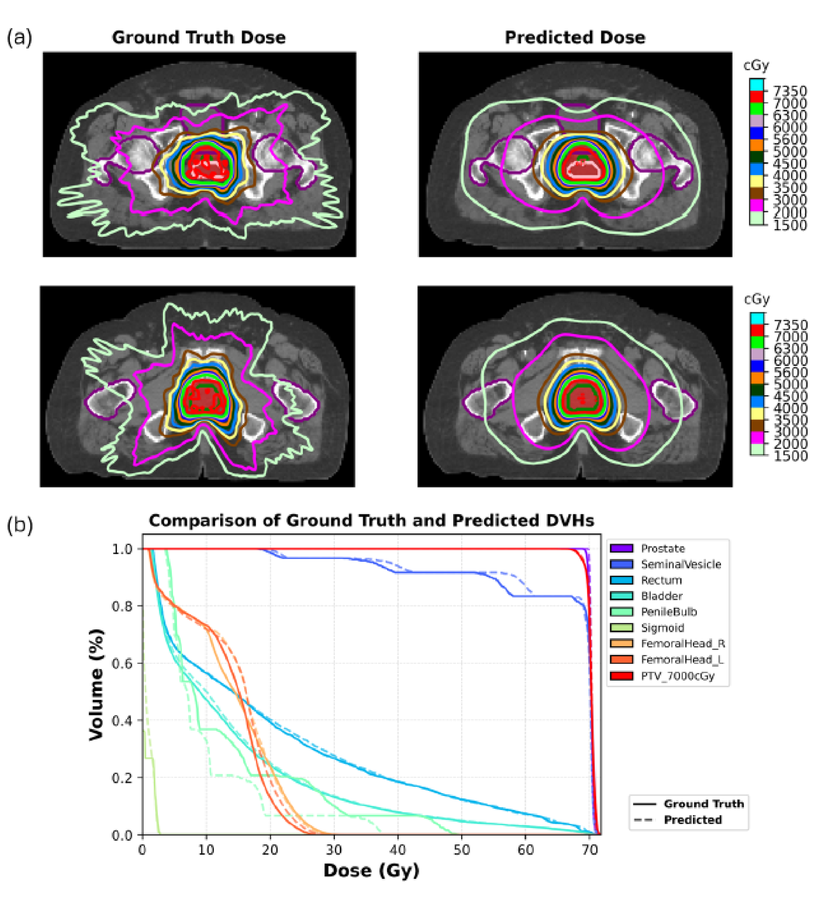

Fig. 3 Comparison between ground truth RPA data and predicted dose for a prostate test patient.

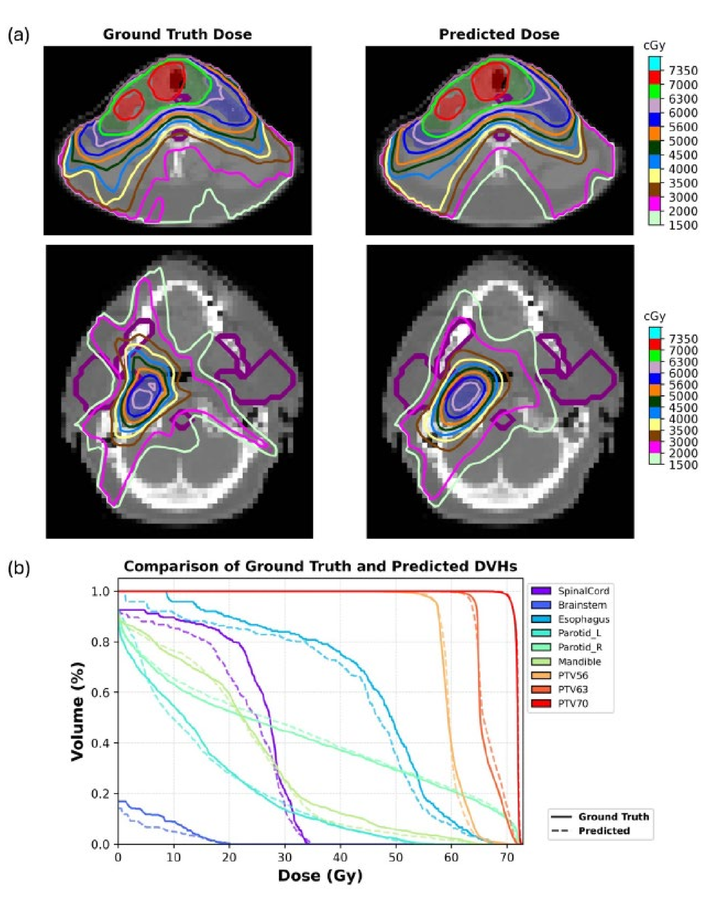

Fig. 4 Comparison between ground truth RPA data and predicted dose for a head and neck test patient.

Josiane Pafeng (jlpafeng@mdanderson.org)
QR code of article and GitHub repo

CONCLUSIONS

RADIANT, a fully configurable, open-source platform for radiotherapy dose prediction enables reproducible experimentation across diverse architectures, loss functions, and optimization strategies.
It is flexible, uses fewer parameters compared to most architectures, while achieving accurate performance on cervical, prostate and head and neck datasets, including top-tier DVH accuracy on the OpenKBP benchmark dataset.
RADIANT addresses a critical need for scalable tools in radiotherapy AI research.