SIIM CAIMI25
Conference on Artificial Intelligence in Medical Imaging
OCT 20-21 | UNIVERSITY OF CALIFORNIA, IRVINE

A Flexible nnU-Net–Based Framework for Radiotherapy Dose Prediction Across Diverse Beam Geometries and Prescriptions

Ho-hsin (Rita) Chang, University of Alabama at Birmingham; Rex Cardan, PhD; John Fiveash, MD; Samantha Simiele, PhD; Joseph Harms, PhD; Bodduluri Sandeep, PhD; Richard Popple, PhD; Carlos E. Cardenas, PhD

Introduction/Background

Radiotherapy is delivered to three out of five cancer patients, but optimizing the dose distribution for each patient is time-consuming. Deep learning offers a potential solution for automating dose prediction; however, the lack of large-scale, diverse datasets has hindered the development of a robust foundational model. To address this gap, we extended the nnDoseNet, previously published deep learning dose prediction framework, to predict 3D dose distributions across multiple disease sites, prescription doses, and beam arrangements.

Methods/Intervention

The extended nnDoseNet utilizes four input channels: simulation CT scan, target volume masks with prescriptions, organ-at-risk masks, and beam configuration maps to capture all relevant patient and treatment-specific factors. e compiled a multi-site dataset of 1,200 clinically delivered plans (the largest, to our knowledge, reported for 3D dose prediction) and split it into 1,000 cases for model training and 200 hold-out cases for independent testing. The multi-site dataset covers four tumor sites (head-and-neck, prostate, breast, and lung), encompassing a wide range of target configurations (1–5 targets per case), plan prescription doses (1.5–84 Gy), and beam geometries (3D, IMRT, and VMAT). We trained site-specific models (using 250 plans per site) as well as a single unified model on the combined dataset that pooled those same plans (250 x 4 = 1000 in total) to evaluate the benefits of multi-site learning.

Results/Outcome

Both the single-site and multi-site models demonstrated high accuracy on clinically relevant metrics when tested on a testing subset of 50 plans per site. Most predicted dose distributions had a Mean Absolute Error (MAE) within 2Gy and 5% to prescription when compared to the delivered dose (ground truth).

Conclusion

The results indicate that while there isn't a significant absolute difference between the predicted and clinical plans, a difference does emerge when comparing the predictions of the single-site and multi-site models. Notably, the multi-site model showed improved performance specifically for lung cases.

Statement of Impact

Given the convenience of a single generalizable model and its strong performance, our multi-site nnDoseNet framework could serve as a foundation for streamlining radiotherapy dose planning across cancer types.

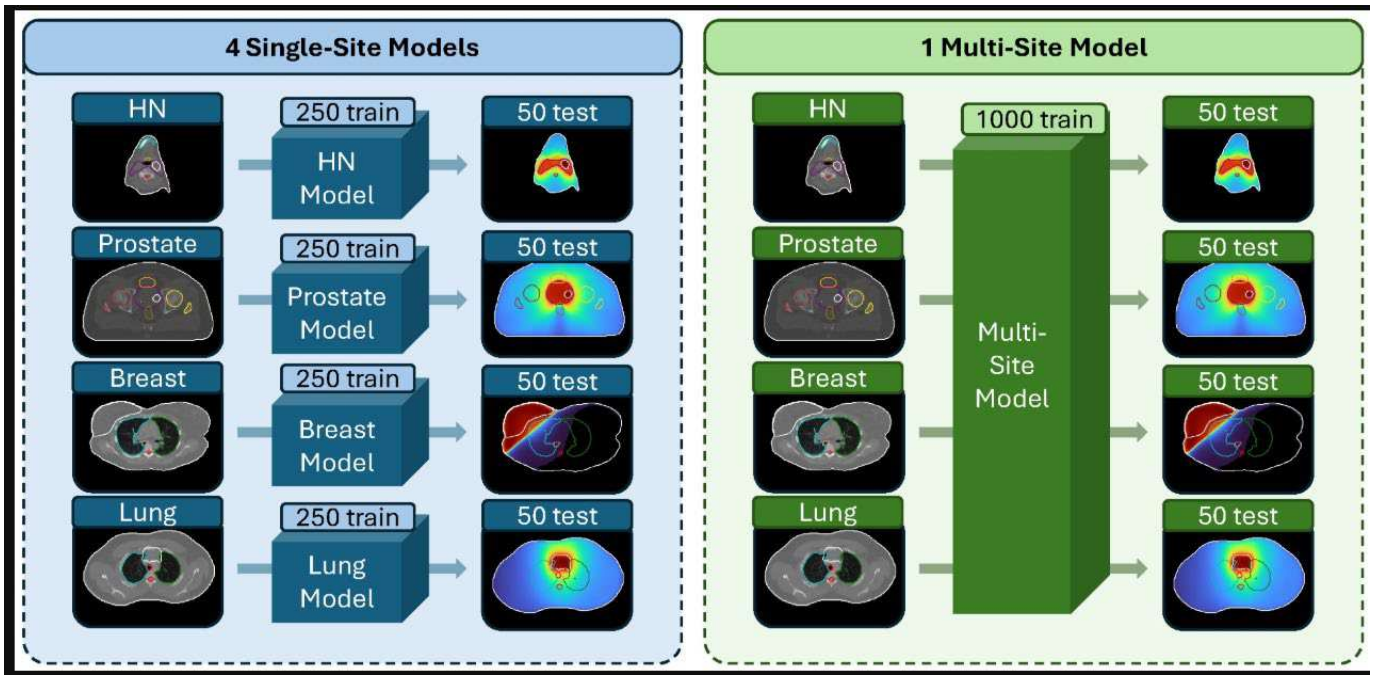

Workflow and dataset split for both single-site and multi-site models: A total of four disease sites (head and neck, prostate, breast, and lung) were included in the experiment, with a 250/50 training/testing split per site. A single-site model is trained for dose prediction specific to a particular site, while a multi-site model aims to predict doses across multiple sites.

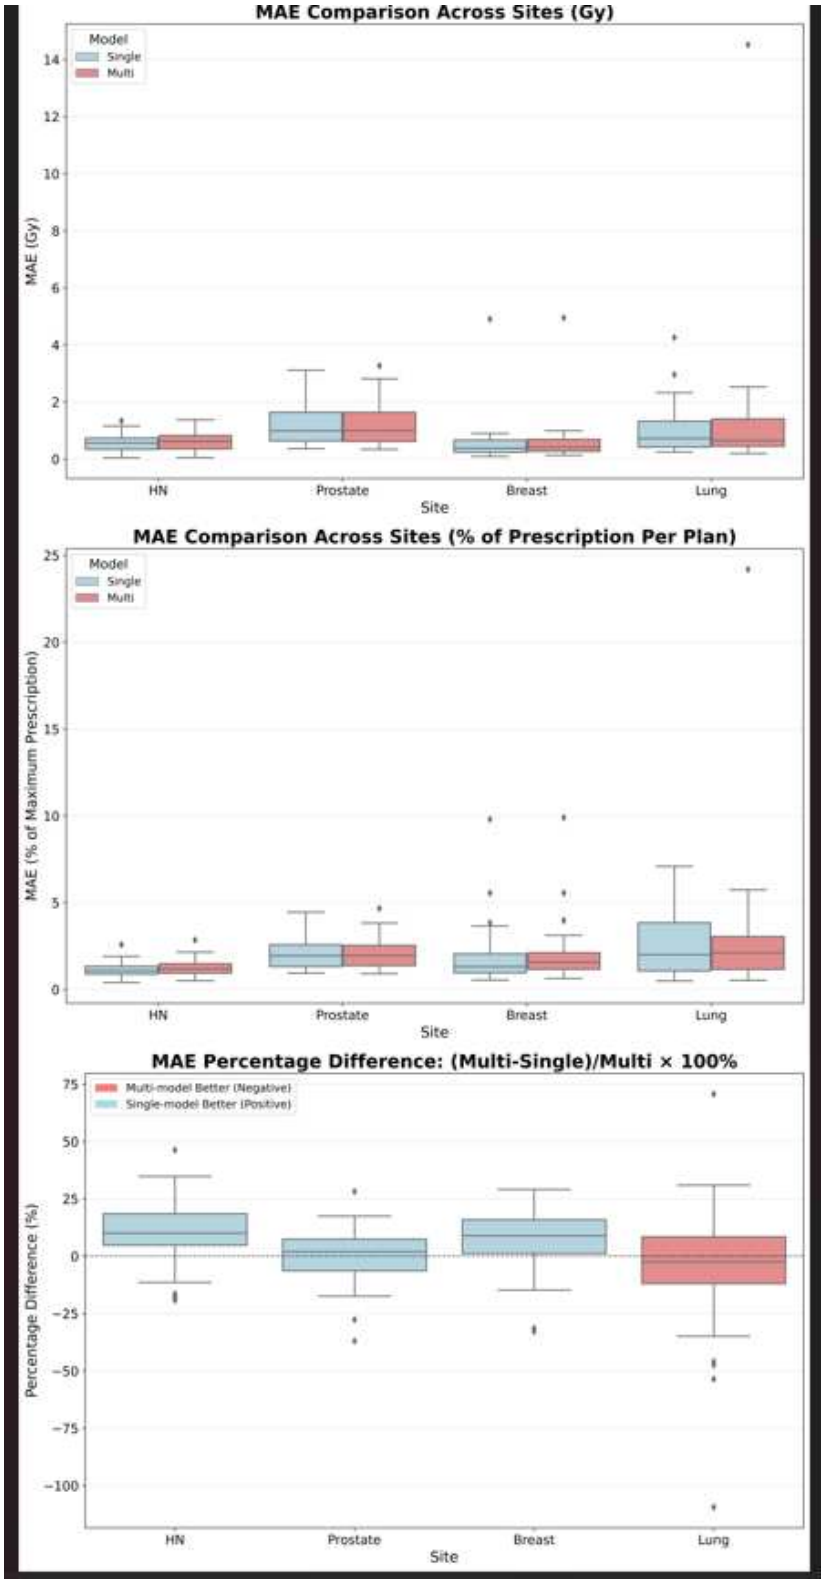

The top two rows display the Mean Absolute Error (MAE) for both single-site and multi-site trained models compared to the ground truth. These errors are presented in absolute values (Gy) and as relative values (percentage of the maximum prescription dose for each case). The bottom row shows the percentage difference, which compares the level of difference between the single-site and multi-site models.

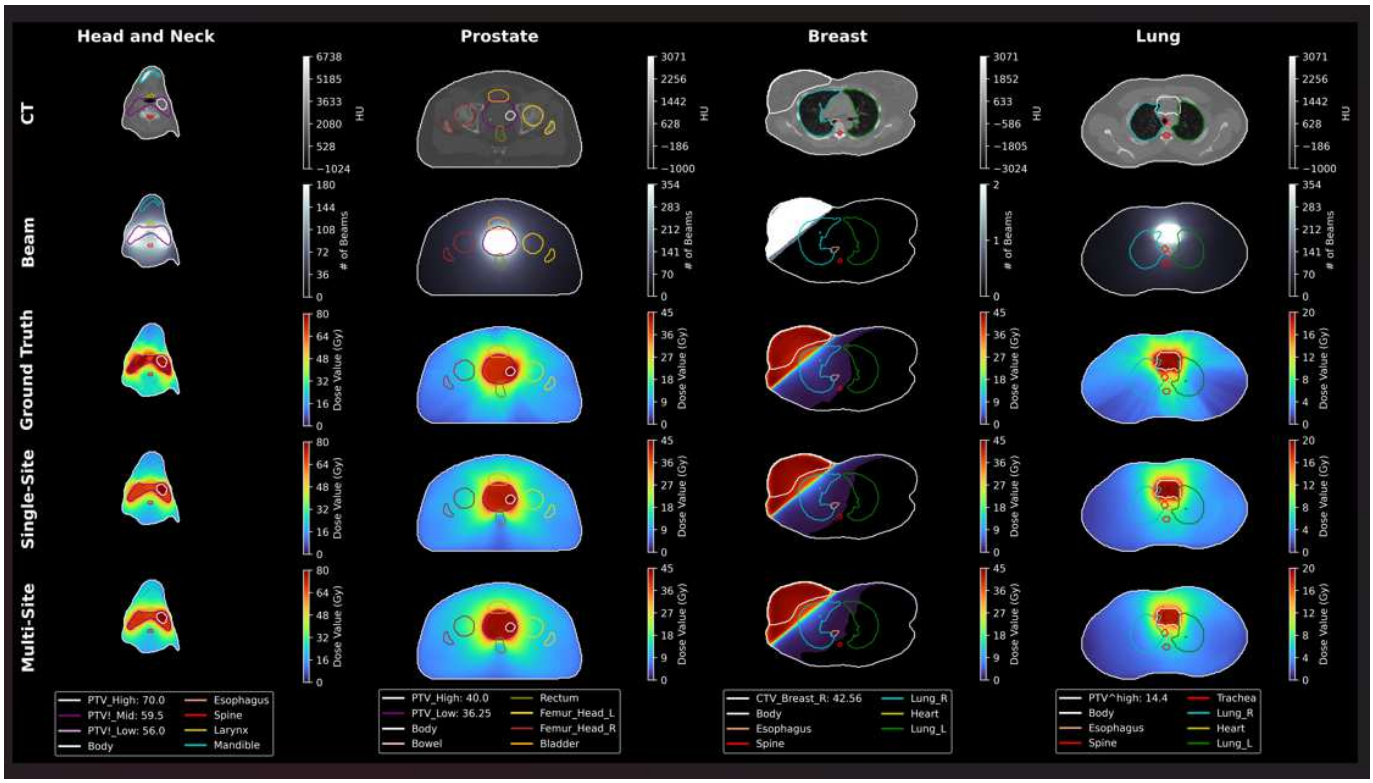

The CT, beam map, contours, and dose map of ground truth, single-site and multi-site model predictions. We include axial views for representative cases for each site among 200 independent test cases.

Keywords

Radiotherapy Treatment Planning; Dose Prediction; Deep Learning