Self-Supervised Deep Learning for Automated Segmentation in MR-Guided High Dose-Rate Brachytherapy

Caleb Thompson[1], Ruiyan Ni[1], Tony Tadic[2,3], Benjamin Haibe-Kains[1,3], Alexandra Rink[1,2,3]

[1]Department of Medical Biophysics, University of Toronto, Canada | [2]Department of Radiation Oncology, University of Toronto, Canada | [3]Princess Margaret Cancer Centre, University Health Network, Canada

Medical Biophysics UNIVERSITY OF TORONTO

Radiation Oncology UNIVERSITY OF TORONTO

Abstract

•MR-guided HDR-BT offers superior tumour control for cervical cancer, but is time-intensive and reliant on scarce labelled MRI data

•Self-supervised learning (SSL) via masked autoencoding (MAE) enables pretraining on unlabelled pelvic MRI, reducing dependence on large annotated datasets for organ-at-risk (OAR) segmentation

•A pretrained ResNet tested against nnU-Net demonstrated statistically significant vDSC improvements and consistent HD95 gains, under limited labelled data conditions

Objective

•Investigate SSL pretraining as a strategy to address labelled data scarcity in MR-guided HDR-BT OAR segmentation

•Pretrain and fine-tune ResNet architectures and evaluate against nnU-Net

•Determine whether SSL pretraining confers a unique advantage in low-data regimes relevant to under-resourced clinical settings

Materials & Methods

•The pretrained dataset consisted of 9,303 T2W MRI volumes from 500 patients (348 GYN, 152 GUP)

•Using 3D MAE with 75% block masking, as well as a hybrid variant combined MAE with contrastive learning

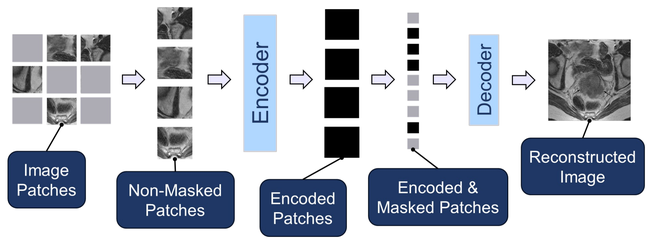

Figure 1: The masked autoencoding pretraining pipeline

•A masked autoencoding pretrained model and a hybrid pretrained model was compared to nnU-Net on each of the 3 training splits to conduct OAR segmentation on the bladder, rectum, sigmoid, and small bowel

•Model performance was evaluated on volumetric DSC (vDSC) and 95th percentile Hausdorff Distance (HD95), with paired t-tests for statistical comparisons

Split	Patients	Images
Train #1	5	6
Train #2	10	12
Train #3	15	18
Test	13	18

Table 1: Patient and image totals for each train and test split

Results

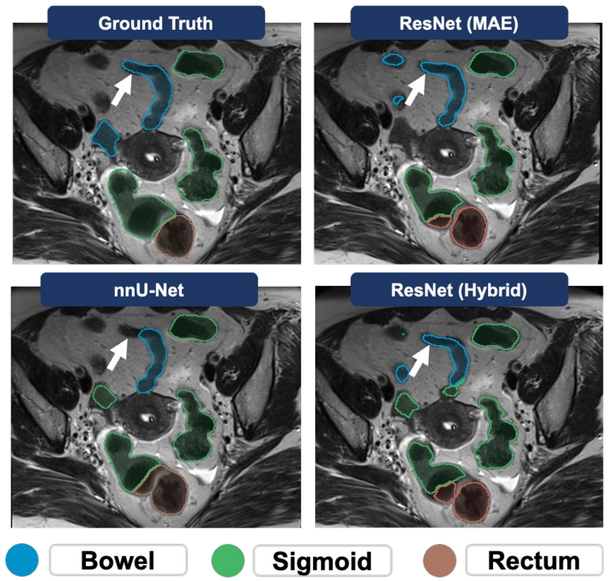

Figure 2: Comparison of baseline, nnU-Net, and pretrained ResNets for a single test case

•MAE ResNet showed consistent vDSC and HD95 gains across splits

•Largest improvements for sigmoid and small bowel

•HD95 improvements were consistent but not always statistically significant; rectum showed marginal gains

•Pretrained models showed the strongest relative advantage at smaller training sizes (5–10) scans

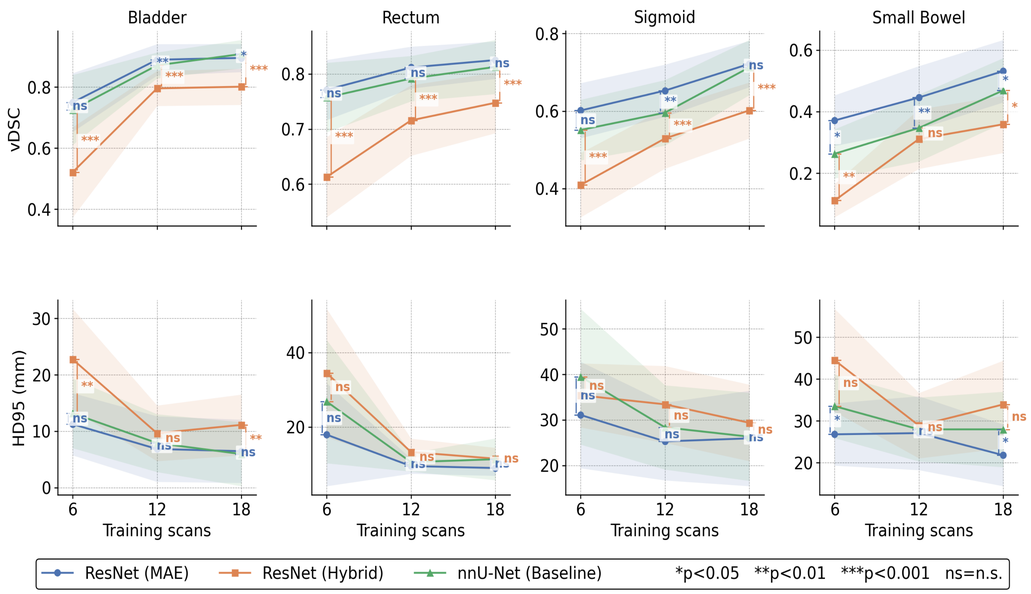

Figure 3: Baseline, MAE, and Hybrid segmentation performance

Discussion

•Hybrid SSL pretraining on unlabelled pelvic MRI yielded statistically significant vDSC gains over nnU-Net, demonstrating the value of pretraining in this domain

•SSL benefit was most pronounced at 5–10 labelled scans, supporting its use in data-scarce settings disproportionately affected by cervical cancer

•These findings support SSL as a practical tool for institutions with limited labelled data to achieve competitive segmentation performance

Conclusion

•SSL pretraining with MAE on unlabelled pelvic MRI led to statistically significant improvements of vDSC for OAR segmentation vs. nnU-Net

•Future work will expand pretraining dataset diversity, as well as comparing additional SSL strategies such as contrastive learning and self-distillation methods