Multi-Objective Multiple Instance Learning for Improving Prostate Cancer Detection in Ultrasounds

Mohamed Harmanani[1], Paul F. R. Wilson[1], Fahimeh Fooladgar[2], Amoon Jamzad[1], Mahdi Gilany[1], Minh Nguyen Nhat To[2], Brian Wodlinger[3], Purang Abolmaesumi[2], Parvin Mousavi[1]

[1]Queen's University, Kingston, ON, Canada [2]University of British Columbia, Vancouver, BC, Canada [3]Exact Imaging, Markham, ON, Canada

INTRODUCTION: Deep learning is typically used to detect prostate cancer (PCa) in small regions of interest (ROI) extracted from a needle trace region of the ultrasound image of a biopsy core. This approach suffers from weak labeling: ground-truth histopathology labels describe tissue properties of the entire biopsy core, and ROI labels are only an approximation of the true distribution of cancer. Multiple instance learning (MIL) approaches to PCa detection from ultrasound have recently been proposed as a solution [1] by leveraging contextual information in multiple patches. In this work, we improve the performance of MIL models by introducing a novel learning objective that takes advantage of both Core-scale and ROI-scale predictions.

Figure 1. A general pipeline for multi-objective MIL with a Transformer

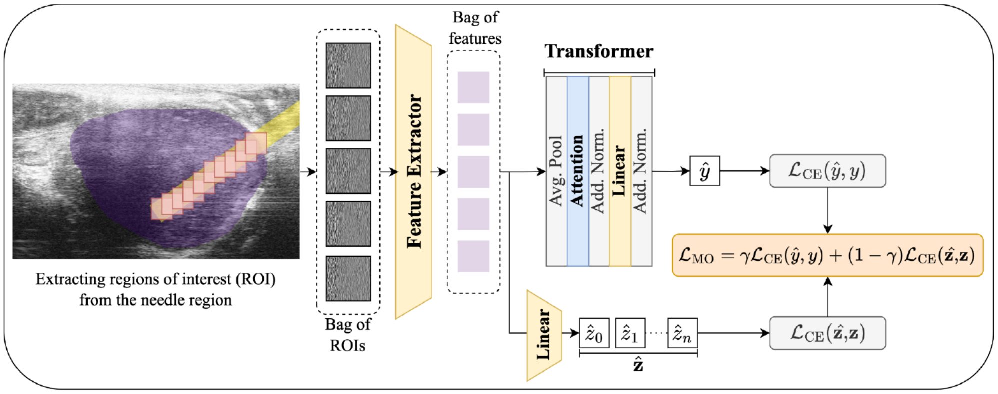

Table 1. Comparison of our method to various MIL and ROI-scale baselines

Extractor + Finetuner	Loss	AUROC	Bal. Accuracy	Sensitivity	Specificity
CCT + Linear	$\mathcal{L}_{CE}(\hat{z}_i, z_i)$	74 ± 4	69 ± 3	71 ± 7	66 ± 5
ResNet18 + Linear	$\mathcal{L}_{CE}(\hat{z}_i, z_i)$	76 ± 4	68 ± 3	65 ± 6	$\mathbf{74 \pm 4}$
CCT + MIL	$\mathcal{L}_{CE}(\hat{y}, y)$	71 ± 4	62 ± 3	62 ± 3	62 ± 3
ResNet18 + MIL	$\mathcal{L}_{CE}(\hat{y}, y)$	77 ± 2	64 ± 2	67 ± 27	57 ± 22
CCT + MO+MIL (ours)	$\mathcal{L}_{MO}$	72 ± 3	63 ± 3	53 ± 14	73 ± 9
ResNet18 + MO+MIL (ours)	$\mathcal{L}_{MO}$	$\mathbf{78 \pm 3}$	$\mathbf{71 \pm 6}$	$\mathbf{76 \pm 12}$	66 ± 20

METHODS: We use 6607 biopsy cores collected from 693 patients who underwent prostate biopsy in five centers under the guidance of Trans-rectal ultrasound (TRUS). To mitigate label imbalance, we undersample the benign cores during training in order to ensure the dataset has an equal amount of benign and cancerous cores. We compare 2 models for ROI feature extraction: the ResNet18 [2] and the Compact Convolutional Transformer (CCT) [3]. We pre-train each model using self-supervised learning [4], then finetune them on the task of ROI-scale PCa detection. We then use the models as feature extractors and train a MIL feature aggregator on top of the extractor's learned representations, using both cross-entropy (CE) loss and multi-objective (MO) loss. This workflow is shown in Figure 1. **RESULTS**: Our results are shown in Table 1. Among ROI-scale baselines with linear finetuning, the ResNet18 model achieves the highest Area Under the ROC Curve (AUROC). Both MIL models outperform the ROI-scale baselines, with the MO+MIL model obtaining the highest AUROC and Balanced Accuracy scores of 77.9 and 71.1 respectively.

CONCLUSION: Multi-objective learning combined with MIL has the potential to improve PCa detection in ultrasound data. **REFERENCES**: [1] Gilany et al., IJCARS 2023 [2] He et al., CVPR 2016 [3] Hassani et al., arXiv preprint [4] Bardes et al., ICLR 2022