Gated CNNs for Nuclei Segmentation in H&E Breast Images

Shana Beniamin, Dimitri Androutsos, April Khademi

Electrical, Computer and Biomedical Engineering Dept., Ryerson University, Toronto, CAN.

INTRODUCTION

Convolutional Neural Networks (CNNs) outperform traditional segmentation pipelines but segmenting touching and overlapping cells is still a challenge. Gated CNNs, which process edge information in a parallel processing branch, have increased boundary segmentation accuracy on the Cityscapes dataset. We employ this technique for the purpose of improving nuclei boundary segmentation in breast cancer images.

METHODS

The GCNN architecture combines predictions learned in two parallel streams to obtain accurate segmentation maps with high quality boundaries. The regular stream can be any encoder-decoder semantic segmentation network.The edge stream is a series of residual blocks and convolutional gated layers. It uses the image gradients and high level information from the regular stream features maps to segment boundaries only. To obtain the final segmentation map, the outputs from both streams are fused together.

We employ binary cross entropy dice loss to supervise the regular stream (semantic loss) and edge stream (edge loss). The dual task loss considers the predictions from both streams to (1) prevent non-boundary pixels from dominating the loss and (2) ensure consistency between the two streams.

$$\text{Loss} = L_{semantic} + 5 * L_{edge} + L_{dualtask}$$

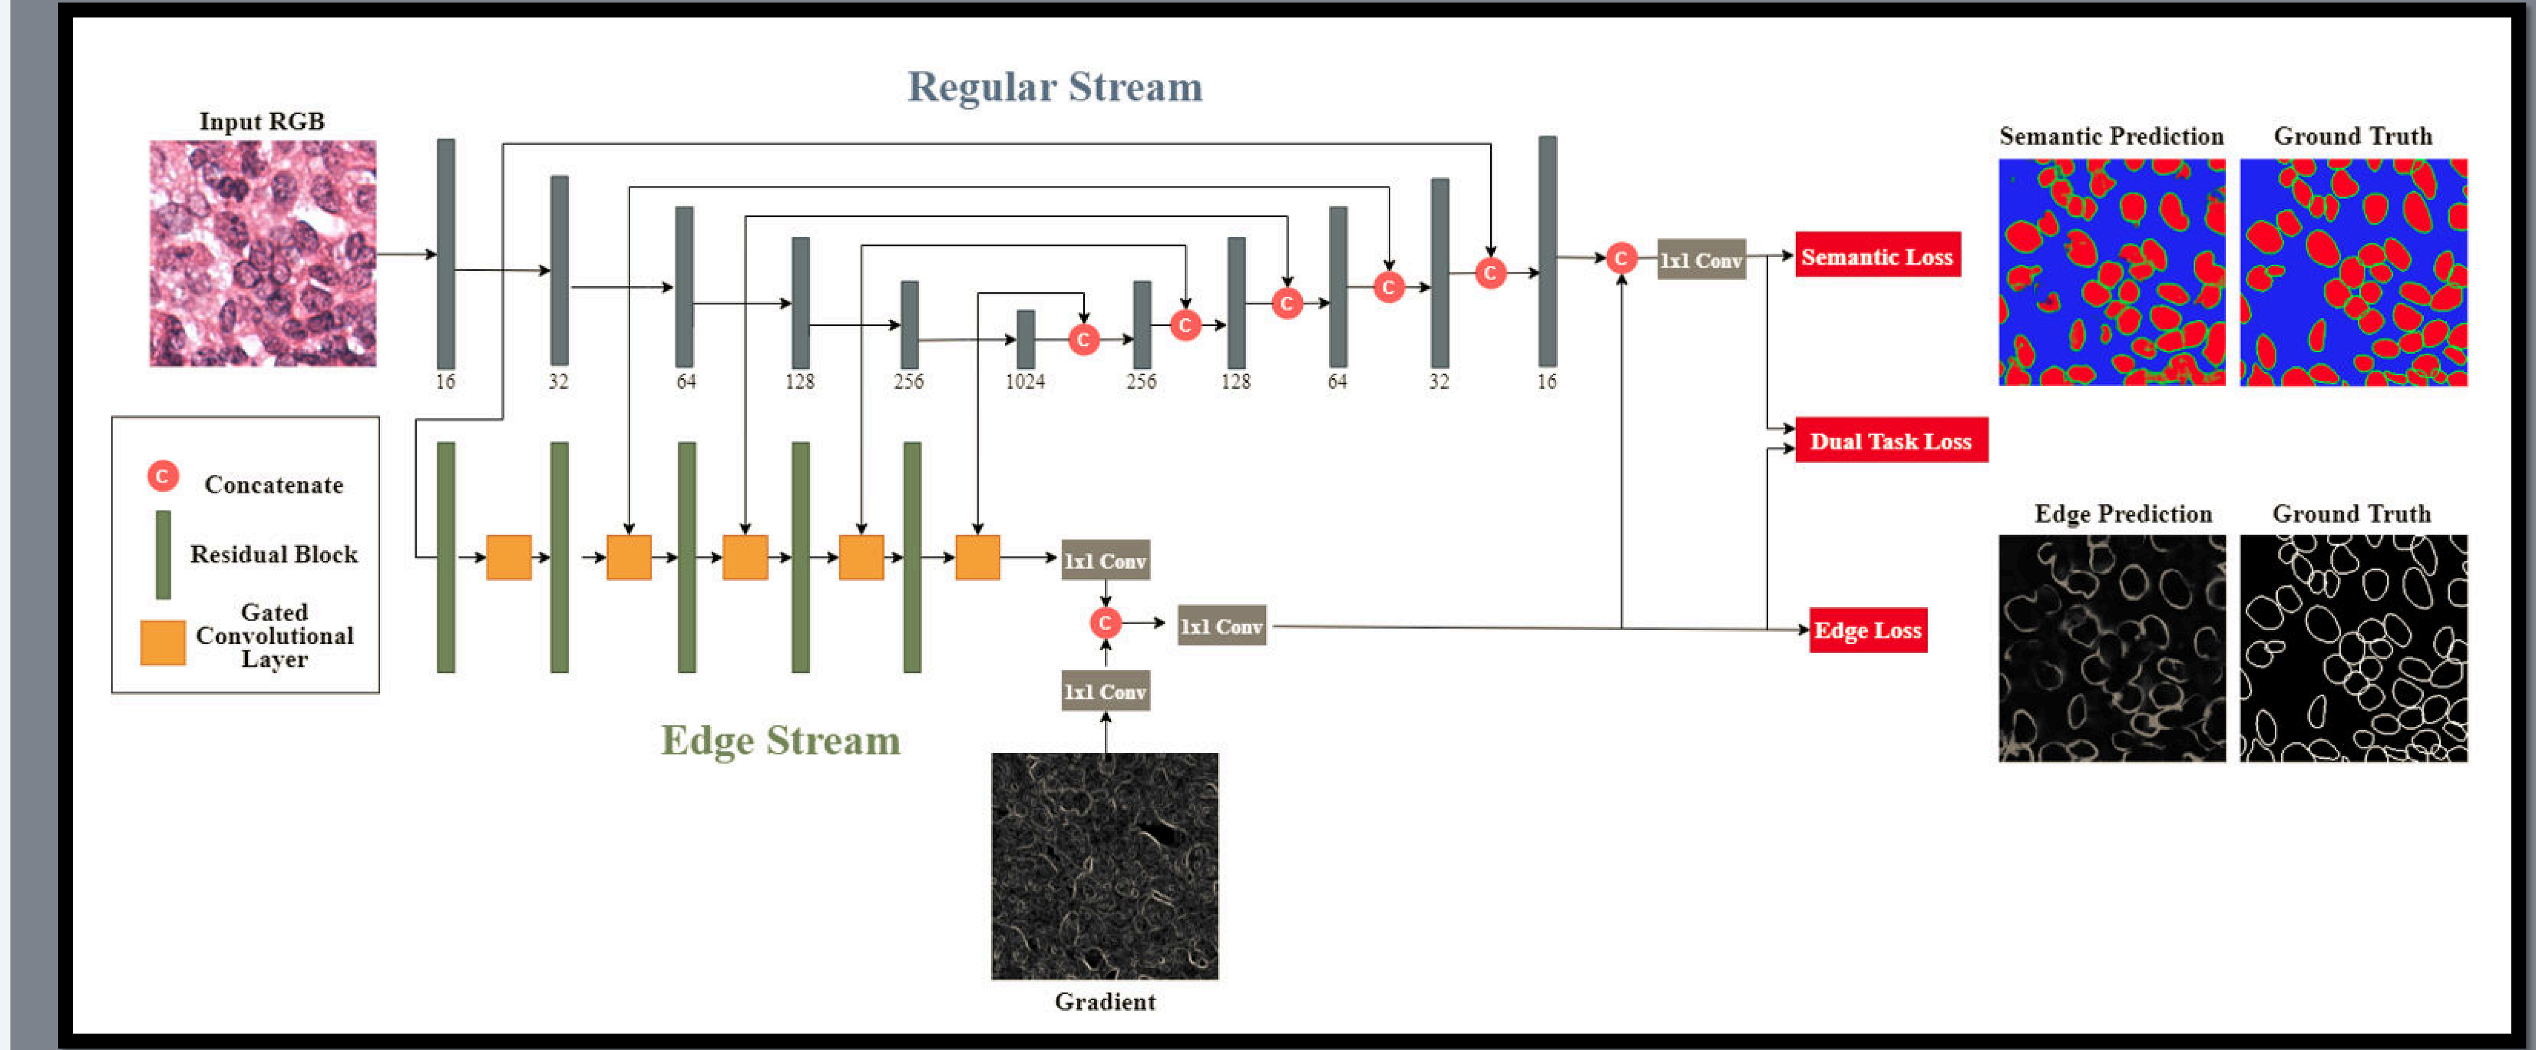

Table 1: Segmentation accuracy evaluated by average DSC.

Models	Entire Dataset (693 images) Nuclei DSC	Entire Dataset (693 images) Boundary DSC	Clustered Nuclei Dataset (317 images) Nuclei DSC	Clustered Nuclei Dataset (317 images) Boundary DSC
U-Net	0.7040 ±0.129	0.3348 ±0.090	0.7104 ±0.112	0.3392 ± 0.071
R2U-Net	0.7109 ±0.133	0.3126 ±0.089	0.7386 ±0.099	0.3199 ± 0.070
GCNN U-Net	0.7165 ±0.147	0.3509 ±0.096	0.7506 ±0.106	0.3635± 0.076
GCNN R2U-Net	0.7280 ±0.126	0.3350 ±0.089	0.7492 ±0.094	0.3394 ± 0.071

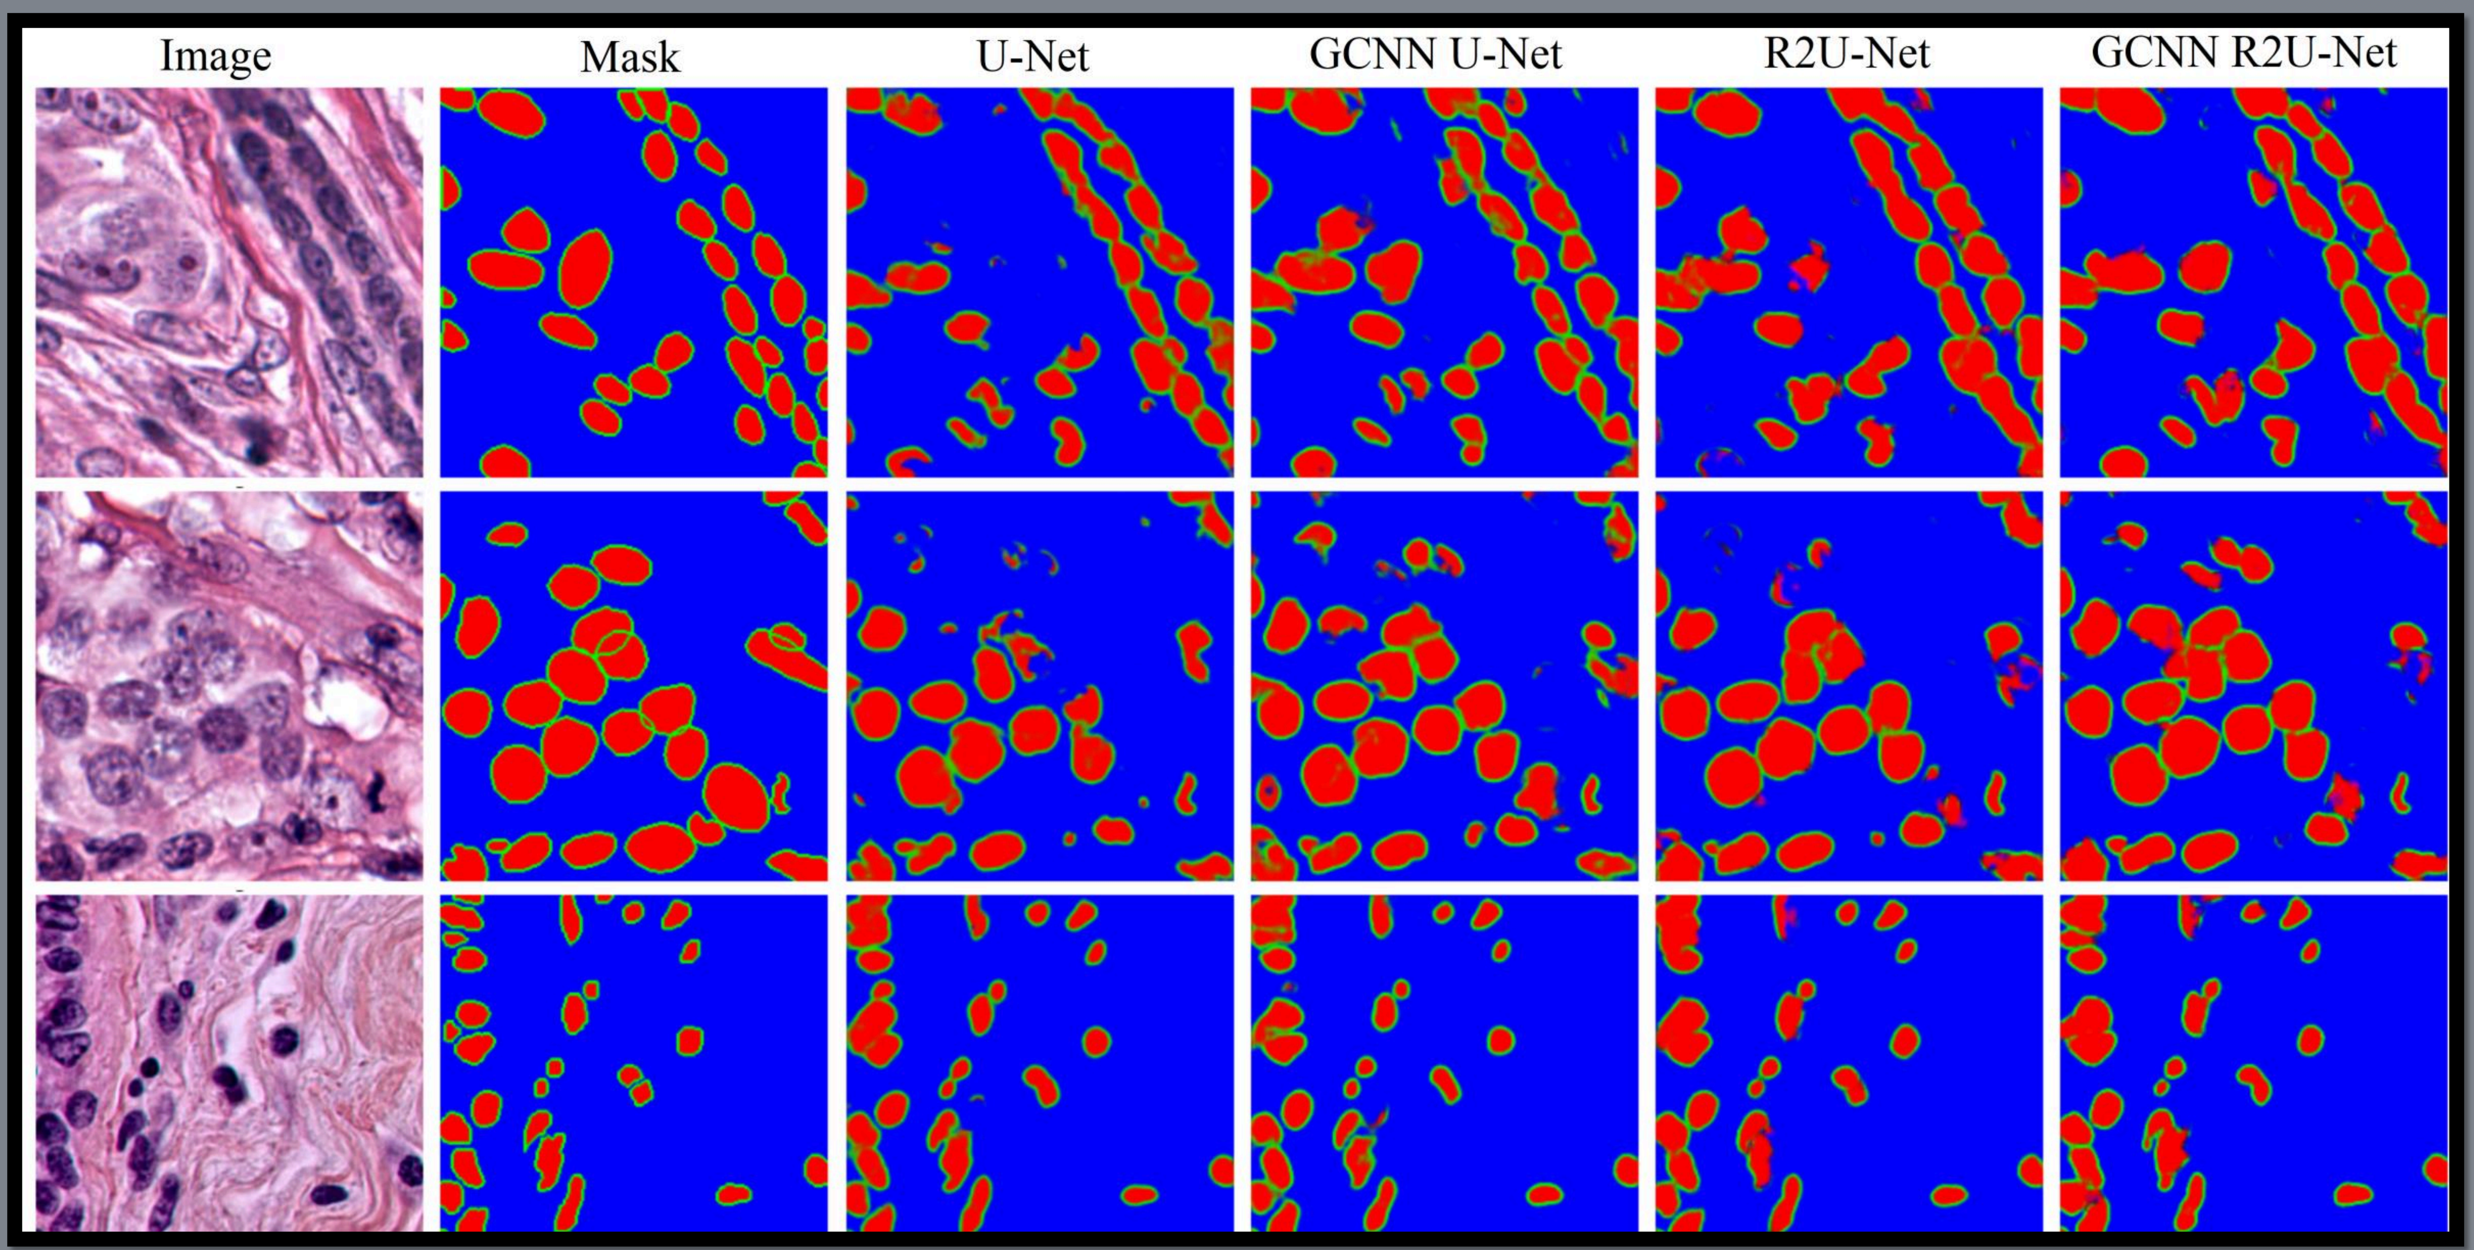

Three multicentre breast cancer datasets (TCGA, TNBC, TUPAC) were used; totaling 1126 images. Two baseline standalone models (U-Net, R2U-Net) and two GCNN models were trained using the standalone architectures as the backbone (regular stream). All four models were trained for 100 epochs, using the Adam optimizer, a learning rate of 0.0001, and a batch size of 4.

RESULTS

GCNN consistently provides improvement for both nuclei and boundary segmentation as the GCNN models outperform the corresponding baselines every time. Notably, DSC scores in the subset data, which are representative of GCNN's edge-based contribution, show GCNN with U-Net experiences a 4% and 2.4% increase in mean nuclei and boundary DSC compared to baseline systems.

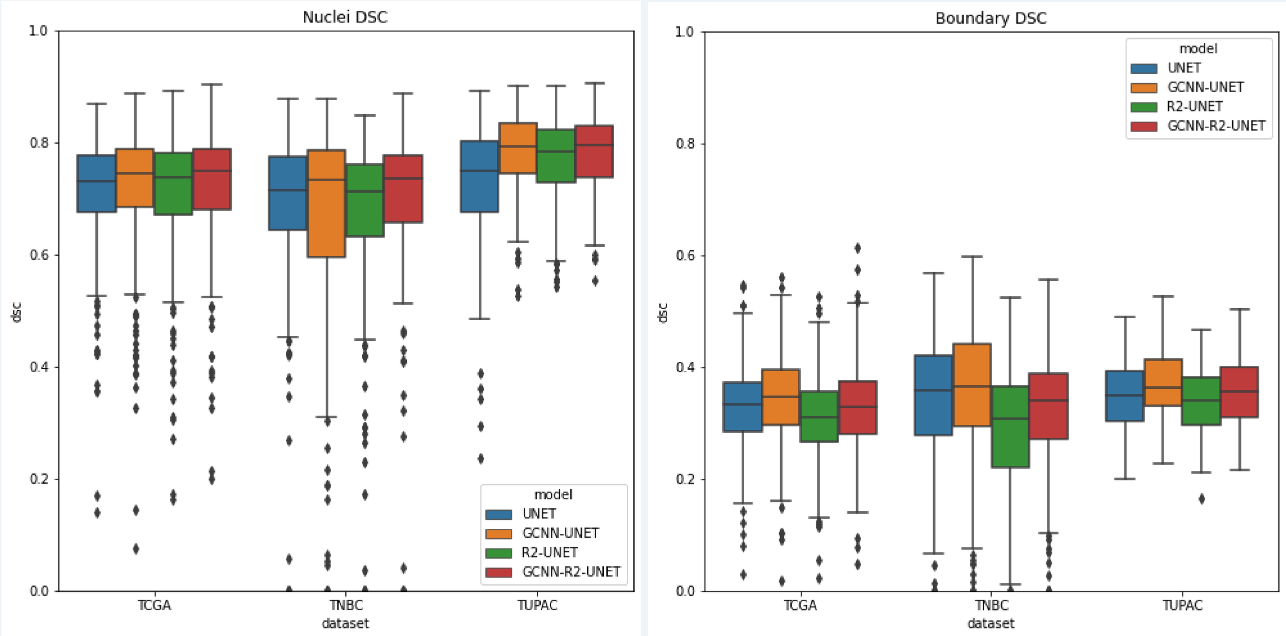

DISCUSSION & CONCLUSION

By implementing GCNNs, our models are better able to segment boundaries, separate touching nuclei, and detect low contrast nuclei that are often undetected by the baseline models. Furthermore, since segmentation performance improved for all three test datasets, our models are able to generalize well to unseen data.