

Gated CNNs for Nuclei Segmentation in H&E Breast Images

Shana Benjamin¹
 Dimitri Androutsos¹
 April Khademi¹

SBENIAMIN@RYERSON.CA
 DIMITRI@RYERSON.CA
 AKHADEMI@RYERSON.CA

¹ *Electrical, Computer and Biomedical Engineering Dept., Ryerson University, Toronto, CAN.*

Abstract

Nuclei segmentation using deep learning has been achieving high accuracy using U-Net and variants, but a remaining challenge is distinguishing touching and overlapping cells. In this work, we propose using gated CNN (GCNN) networks to obtain sharper predictions around object boundaries and improve nuclei segmentation performance. The method is evaluated in over 1000 multicentre diverse H&E breast cancer images from three databases and compared to baseline U-Net and R2U-Net.

Keywords: Nuclei Segmentation, Breast Cancer, Deep Learning, Histopathology, CNNs.

1. Introduction

Deep convolutional neural networks (CNNs) are gaining traction as they achieve higher nuclear segmentation accuracy than traditional hand-crafted pipelines. Although the U-Net architecture has been proven to perform well for this task, a major challenge for this network is distinguishing between touching and overlapping cells, which results in under-segmentation [1]. We propose the use of gated CNNs (GCNN) to improve nuclei boundary segmentation. They have been proven to increase boundary segmentation accuracy on the Cityscapes dataset [2] but the application in nuclei segmentation is novel. Three multicentre breast cancer datasets (over 1000 images) are used to validate the segmentation performance of our proposed GCNN model. The obtained results indicate strong segmentation improvement with the GCNN.

2. Methods

The GCNN is a two-stream architecture; the regular stream can be any encoder-decoder semantic segmentation network and the edge stream is a series of residual blocks and gated convolutional layers [2]. Gated convolutional layers are used to wire boundary information from the intermediate layers of the regular stream’s encoding path to the edge stream which focuses on boundary segmentation solely. The edge stream learns to predict quality edges and uses the Sobel gradient magnitudes of the input image to further highlight contours and textures. Information between the two parallel streams is fused at the end to obtain enhanced boundary predictions. The overall loss function is made up of three components:

$$Loss = \lambda_1 * L_{semantic} + \lambda_2 * L_{edge} + L_{dualtask} \quad (1)$$

where λ_1 and λ_2 are hyper-parameters, $L_{semantic}$ and L_{edge} supervise the regular and edge stream respectively by taking the Dice and Binary Cross Entropy loss, and $L_{dualtask}$ [2] compares the semantic and boundary predictions in order to ensure consistency between the two streams. The proposed work is compared to U-Net [3] and R2U-Net [4].

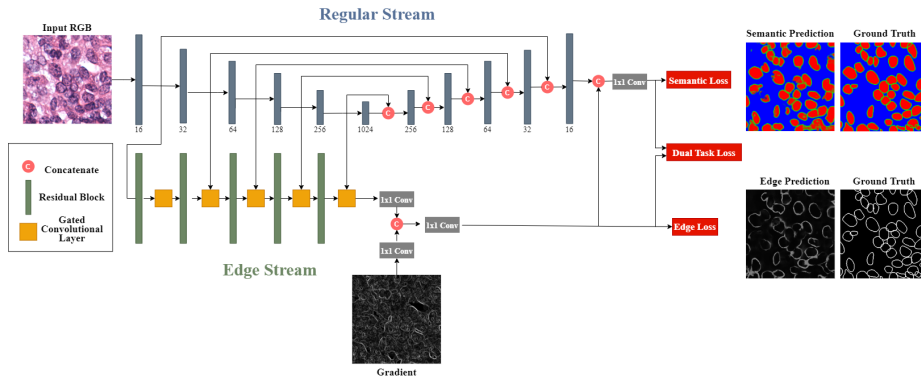


Figure 1: Proposed GCNN model architecture.

Dataset: Three multicentre datasets of invasive breast cancer disease are used in the analysis: TCGA (96 images with annotations from [5] plus an additional 212 that were annotated for this work), TNBC (200 patches from [6]), and TUPAC (618 patches from [7] that were annotated for this work). The RGB patches are 256x256, stained with H&E, cropped from whole slide images (WSI) that were all digitized at 40x magnification. To preprocess the images, Reinhard colour normalization was applied. The architectures are trained using a 60/10/30 split on TUPAC, with TCGA and TNBC being completely held out and used only for testing resulting in total of 693 test images.

Evaluation: Segmentation accuracy is validated using the DSC for both nuclei and boundary classes. To gain insight on model performance specifically on overlapping nuclei, an automated method was developed to extract a subset of the test images that have large amount of clustered nuclei: (1) a distance weight map was generated using the image ground truths [3] that highlights overlapping nuclei boundaries, (2) the maximum weight map value was found for each image and stored in an array, (3) the median of (2) was used as a threshold T to localize pixel regions in the images that have overlapping nuclei boundaries and (4) images with more than 100 pixels greater than the threshold T were retained for further analysis, resulting in 317 testing images (146 TCGA, 132 TUPAC, 39 TNBC).

3. Results

The GCNN is trained using two encoder-decoder architectures for the regular stream (U-Net and R2U-Net) which is compared to these two as baseline models. In this work $\lambda_1 = 1$ and $\lambda_2 = 5$ were found to be optimal. Each model was trained for 100 epochs, with a batch size of 4, using the Adam optimizer and a learning rate of 0.0001. Figure 2 contains nuclei segmentation results, where the GCNN predictions show the edge stream not only improves boundary segmentation (especially in clustered regions) but also promotes segmentation of faint nuclei that baseline models often miss entirely. This may be as a result of GCNN’s attention mechanism that highlights nuclei boundaries (edges); by highlighting the weak edges, the model may be more sensitive to low contrast nuclei. Table 1 reports the average DSC for the entire test set and the subset of images with lots of nuclei clusters. The GCNN offers consistent improvement in both nuclei and boundary DSC for both U-Net and R2U-Net in the regular stream. Using the entire test set, the boundary DSC was improved by

