Automated Feature Segmentation in Digital Cervigrams with a Discriminative Convolutional Neural Network

Isaac G. Freedman BPhil MPH[1,2], Thomas JS Durant MD[1,3], Harlan M. Krumholz MD SM[1,4], Wade L. Schulz MD PhD[1,3]

[1]Yale New Haven Hospital, Center for Outcomes Research and Evaluation, New Haven, CT; [2]Yale University School of Medicine, New Haven, CT; [3]Yale School of Medicine, Department of Laboratory Medicine, New Haven, CT; [4]Yale University School of Medicine, Department of Internal Medicine, New Haven, CT

Background

- ❖ Cervical cancer is the third leading cause of cancer mortality worldwide and the second most lethal cancer in developing countries; more than half of women who develop cervical cancer have not been screened appropriately.
- ❖ Visual inspection with ascetic acid (VIA) along with primary HPV testing is a cost-effective screening method in resource-limited settings.
- ❖ The first step to automated cervical cancer screening using computer vision methods is to segment the cervicographic features.

Objectives

- ❖ To train a discriminative convolutional neural network (CNN) to generate object masks, which accurately demarcate cervical regions.
- ❖ To validate the performance of the trained CNN using the standard performance metric for modern image segmentation technology.

Methods

- ❖ Used initial data sample from four datasets in the NIH cervigram database: Costa Rican Natural History Study of HPV and Cervical Neoplasia (NHS), ASCUS LSIL Triage Study (ALTS), Biopsy Study, and Costa Rica Vaccine Trial (CVT).
- ❖ Manually labeled cervical regions of interest (ROIs) from 411 cervigram images selected randomly from the four datasets.
- ❖ Trained the DeepMask/SharpMask CNN architecture to segment image priors and generate object masks.

Methods

- ❖ Validate the performance of the trained CNN by calculating the Jaccard Index, or Intersection over Union (IoU), for a set of labeled validation images.
- ❖ **<u>The SharpMask CNN architecture consists of the DeepMask feedforward CNN (left) with a bottom-up structure for image segmentation followed by refinement modules (middle and right) in a top-down structure.</u>**

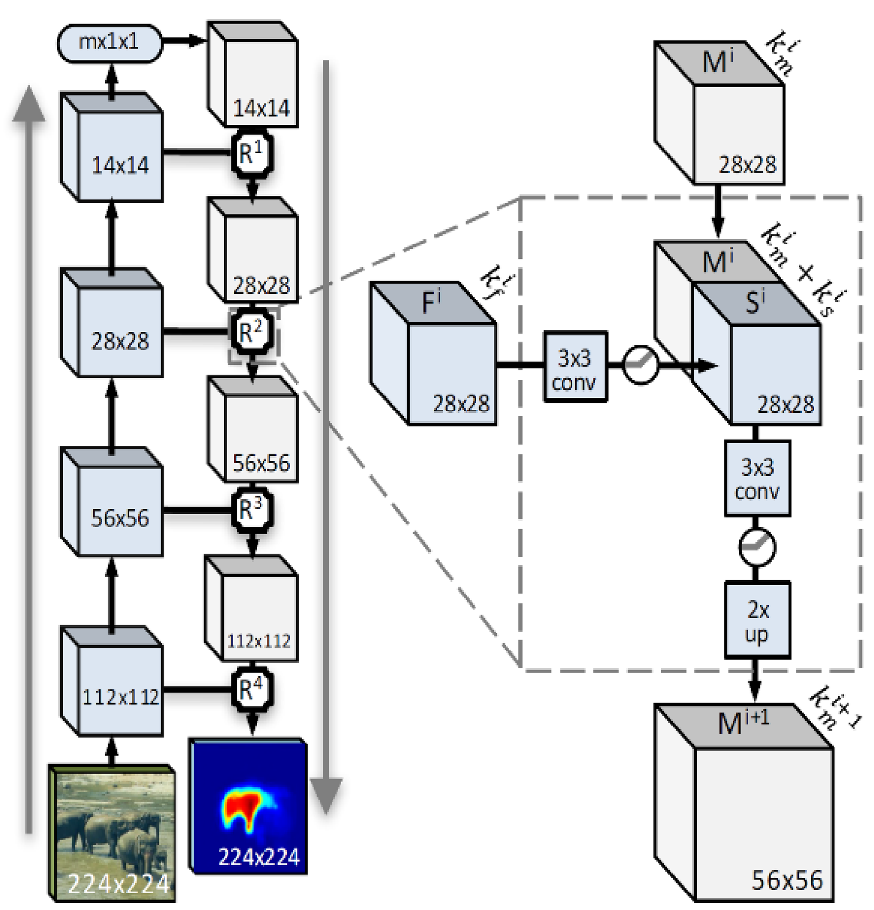

- ❖ N = 300 (72.9%) manually labeled cervigrams trained the CNN and the model was validated on N = 111 (27.1%) of the images.
- ❖ Validate the performance of the trained CNN by calculating the Jaccard Index, or Intersection over Union (IoU), for a set of labeled validation images.
- ❖ An IoU value of 1 indicates a model that completely predicts the ROI.

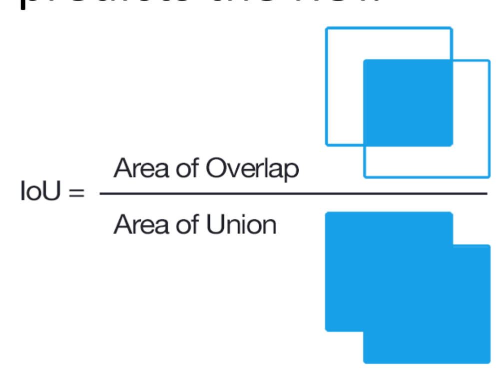

Results

- ❖ **<u>Loss function and IoU of DeepMask model during training.</u>**

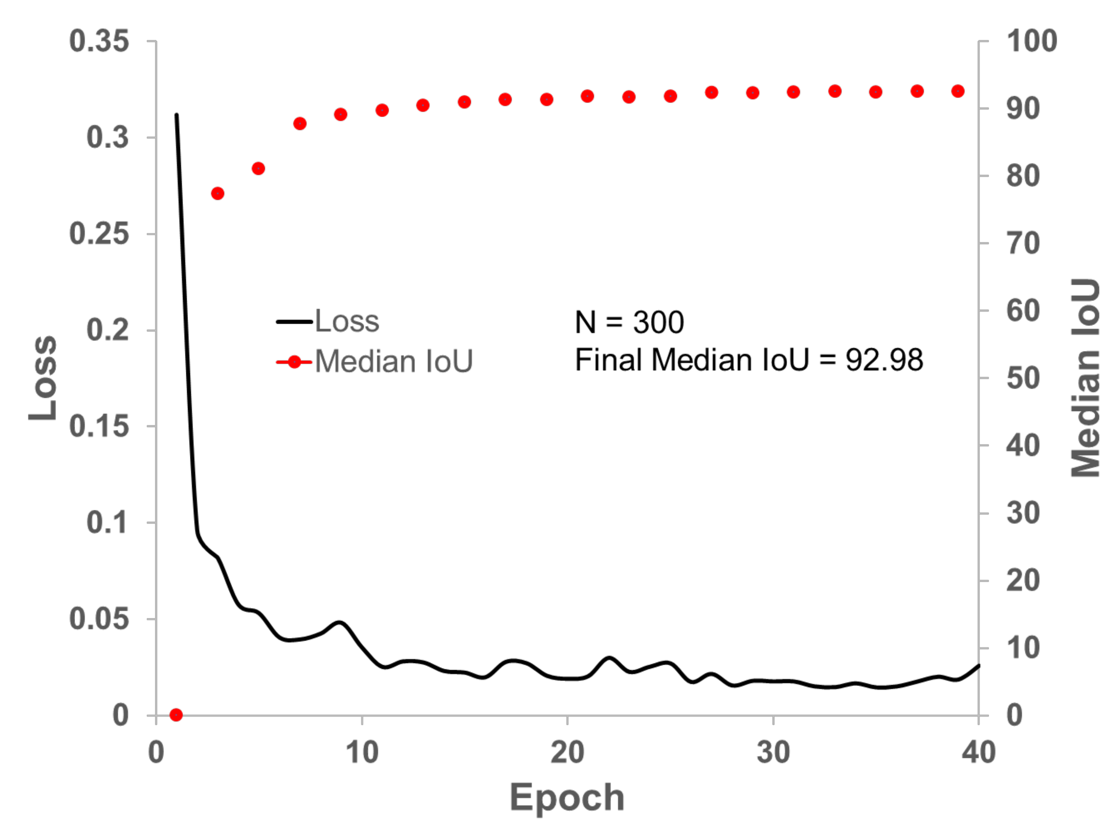

- ❖ **<u>Intersection over Union (IoU) for 300 training cervigrams and 111 validation cervigrams.</u>**

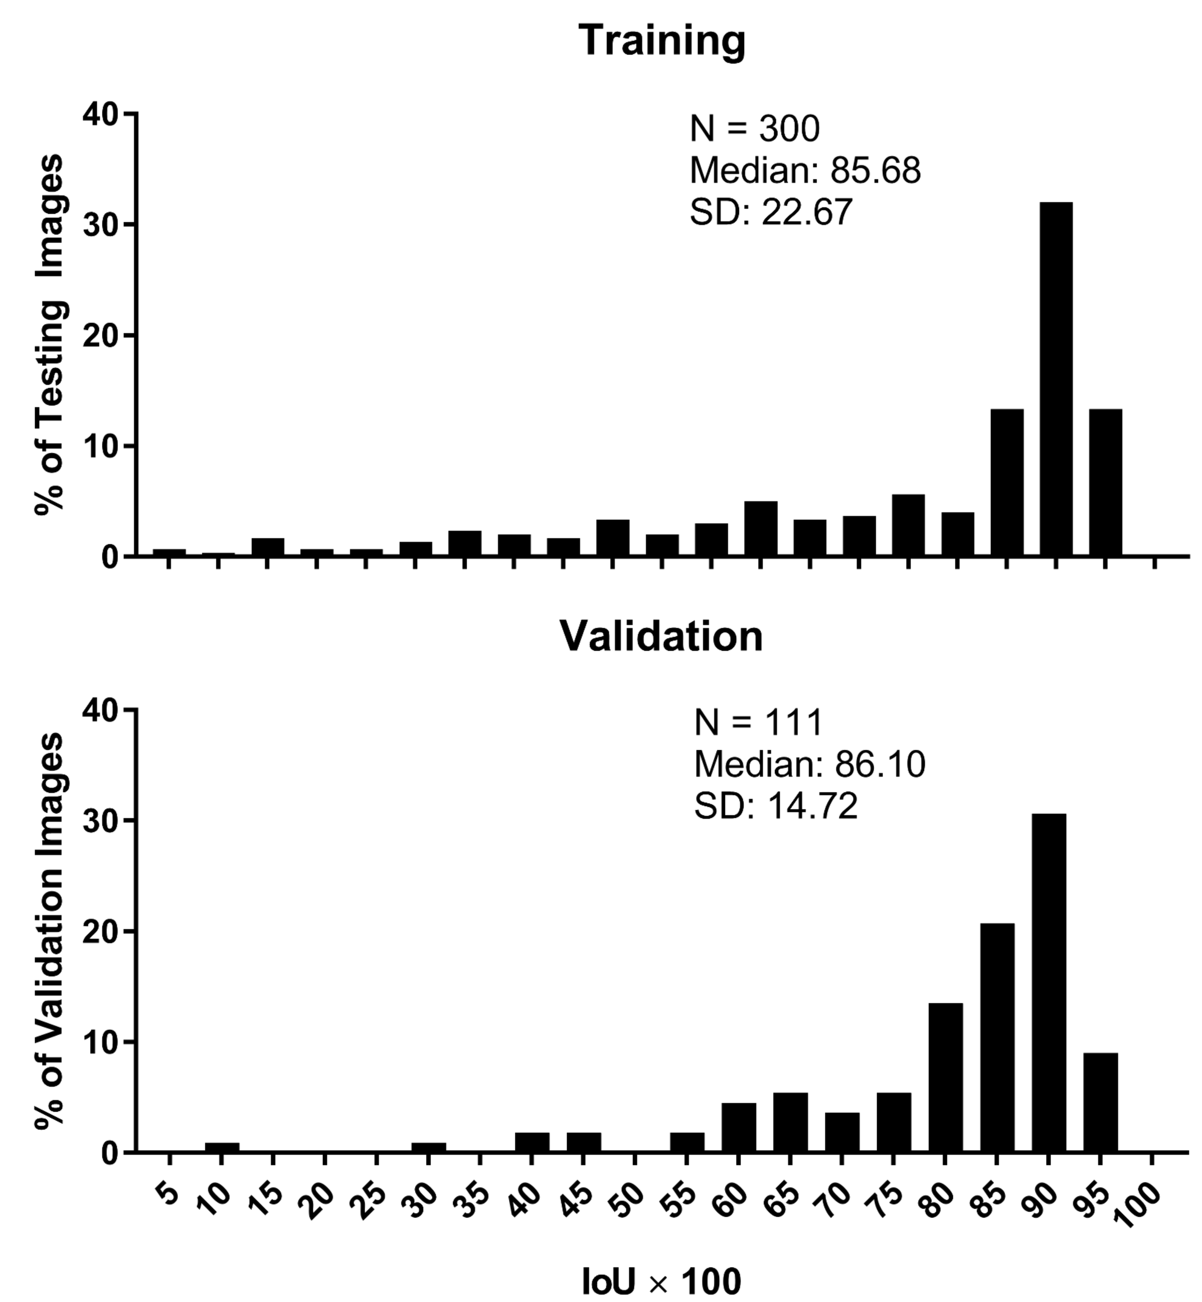

Results

- ❖ **<u>Representative cervigram (a) with segmentation mask (b).</u>**
- ❖ Red line indicates manual label. Image is automated segmentation with IoU = 92%.

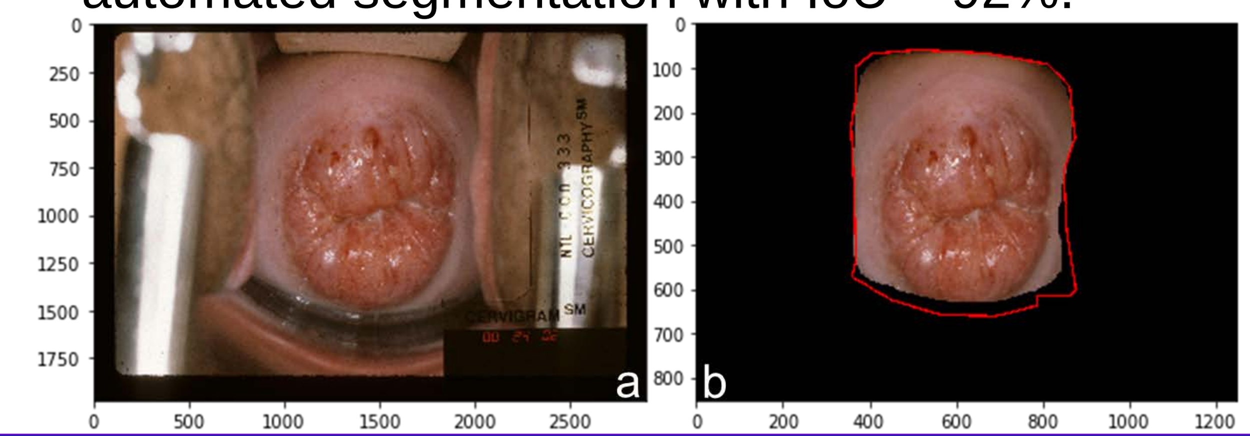

Conclusions

- ❖ Discriminative CNN architecture yields state of the art image segmentation of cervigrams.
- ❖ Model trained on a small fraction of the pilot dataset (14%). Training the model on a larger number of images will likely yield higher segmentation accuracy (IoU).
- ❖ Automatically segmented cervigrams from our model trained on the complete dataset will next be used to train a classification CNN to predict malignancy.

References

- ❖ Ponka D, et al. CMAJ 2014; 186(18):1394
- ❖ Gordi SJ, et al. NEJM 2004; 353(20):2158-68
- ❖ Pinheiro PO, et al. CVPR 2016; ArXiv:1603.08695

Acknowledgements

We would like to acknowledge the NCI CISP User Committee for consultation and access to the cervigram dataset.