Cervical Cancer Screening with Convolutional Neural Networks

Huyen Nguyen, Tucker Leavitt, Yianni Laloudakis

Stanford University – CS231N Course Project

Abstract

The type of a patient's cervix determines the type of pre-cancer treatments the patient can undergo, and the medical community would benefit from a way of efficiently classifying a patient by cervix type. Kaggle and Mobile ODT have published a collection of several thousand specular photographs of cervixes, each labeled as one of three types. We present a convolutional neural network (CNN) trained to classify the cervix images in this dataset. Our network relies on batch normalization [3] to accelerate training and dropout to reduce overfitting [4]. Our model has achieved a classification accuracy of 62%. Moving forward, we hope to improve network performance by de-noising the dataset through image segmentation and by increasing the network depth.

Motivation

When cervical cancer is caught early, it can be treated easily and effectively. The possible treatments will vary based on physiological differences in the cervix. Rural or understaffed clinics would benefit from a way of quickly and accurately classifying patients based on cervix type. Cervical cancer tends to begin in cells within the transformation zone, which could be completely ectocervical and visible (Type 1), partially endocervical but visible (Type 2), or partially endocervical and not fully visible (Type 3). Cervix types 2 and 3 may require different screening or treatment due to the placement and hidden view of precancerous lesions. Our project is to use a convolutional neural net to automate and improve this important classification process.

This project was inspired by a public Kaggle competition [1].

Figure 1: Specular photographs of cervixes from the Kaggle dataset

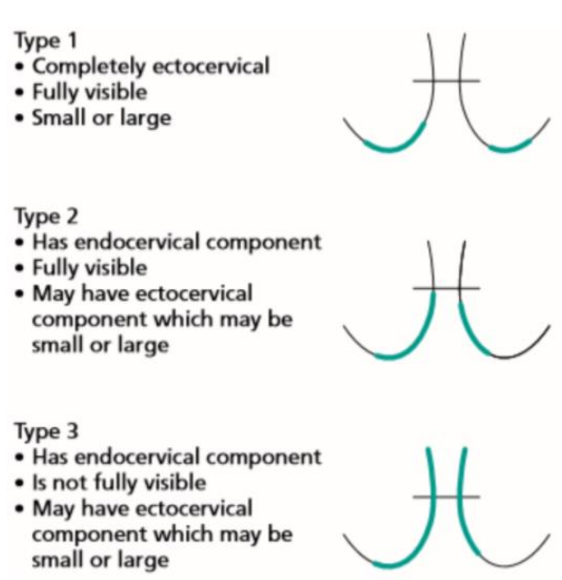

Figure 2: Characteristics of the three cervix types (taken from [1])

References

[1] MobileODT, Intel, & Kaggle Inc. (2017) Intel & MobileODT Cervical Cancer Screening. www.kaggle.com/c/intel-mobileodt-cervical-cancer-screening
[2] Wu, N., et al., TensorFlow Models, (2017), GitHub repository, https://github.com/tensorflow/models.
[3] Ioffe, S., & Szegedy, C. (2015). Batch Normalization : Accelerating Deep Network Training by Reducing Internal Covariate Shift. arXiv Preprint arXiv:1502.03167v3.
[4] Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I. & Salakhutdinov, R. Dropout: a simple way to prevent neural networks from overfitting. J. Machine Learning Res. 15, 1929–1958 (2014).

Model Architecture

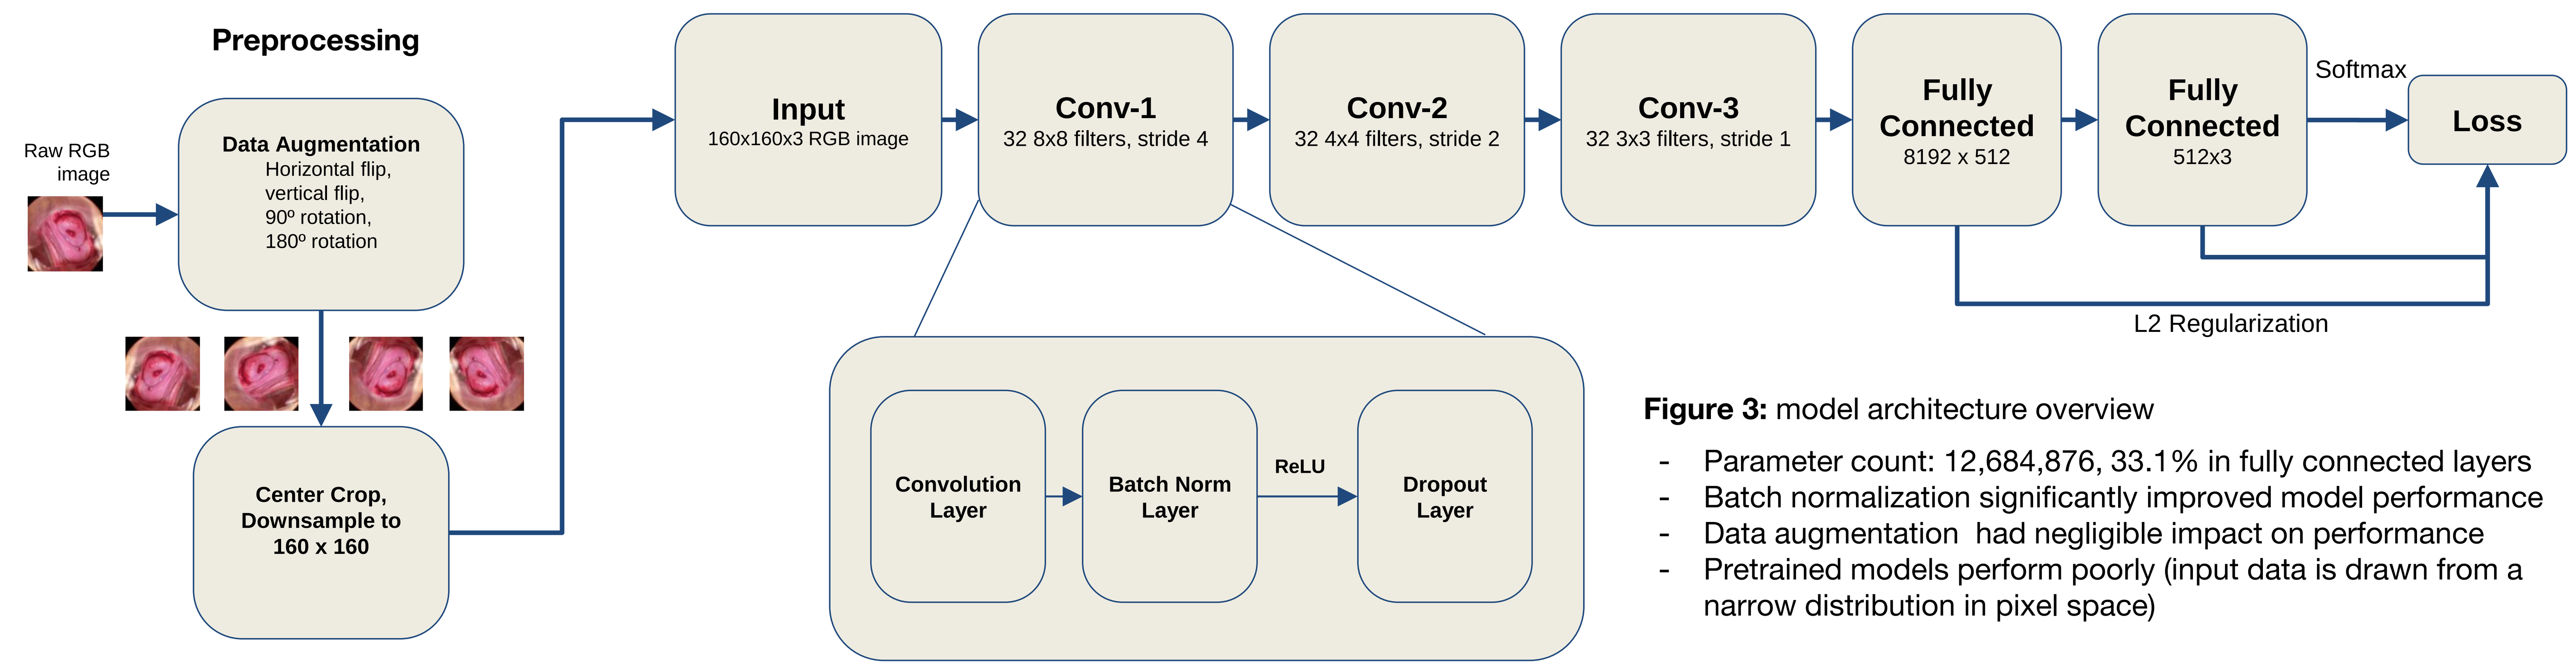

Figure 3: model architecture overview

- Parameter count: 12,684,876, 33.1% in fully connected layers
- Batch normalization significantly improved model performance
- Data augmentation had negligible impact on performance
- Pretrained models perform poorly (input data is drawn from a narrow distribution in pixel space)

Results

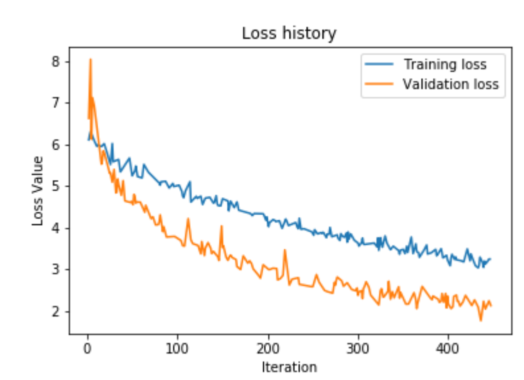

Figure 4: Total validation and training loss for an unconverged model

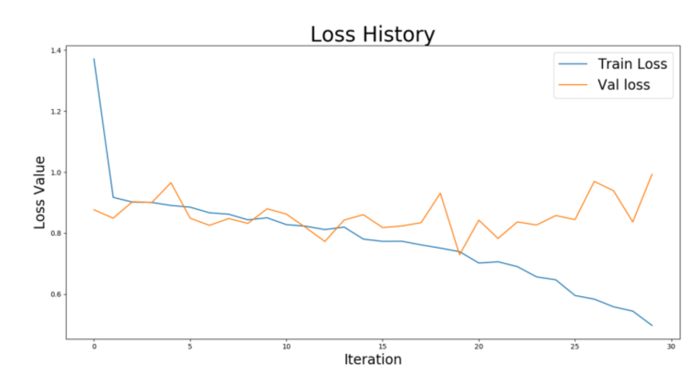

Figure 5: Average validation and training loss for a converged model

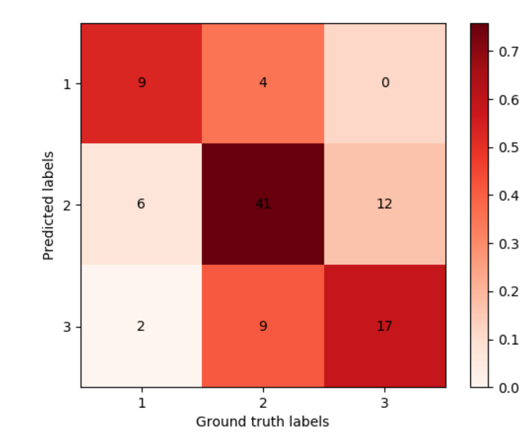

Figure 8: Confusion matrix* for 100 validation data points. Class 2 is highly over-represented in the dataset.

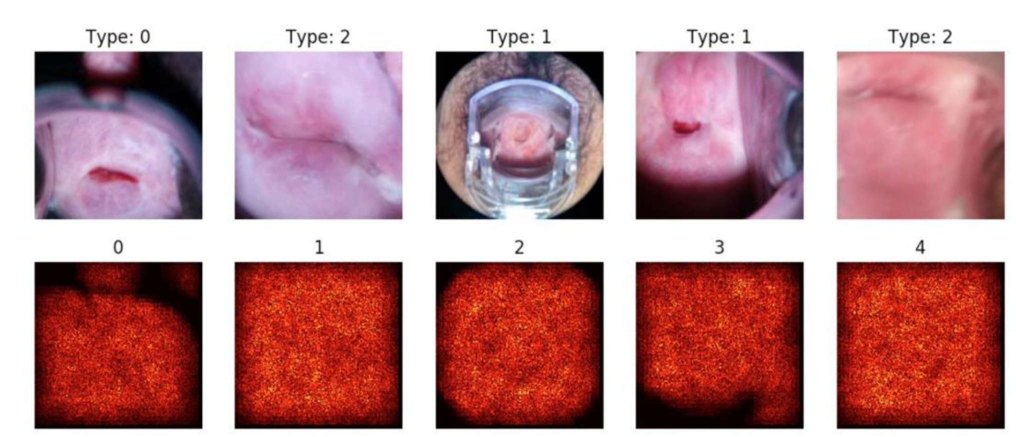

Figure 10: Saliency maps for several images. The neural net is able to distinguish cervical tissue from certain non-cervical tissue.

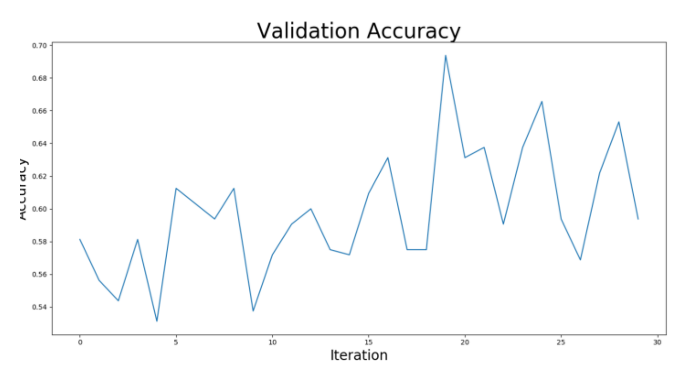

Figure 6: Validation accuracy across training

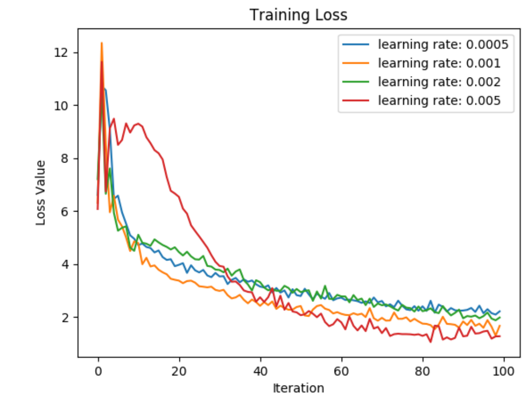

Figure 7: Training loss curves for several different learning rates.

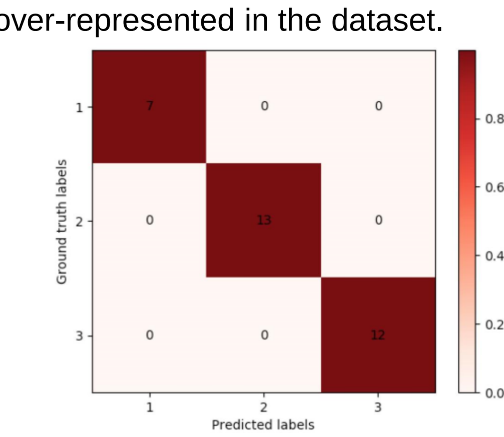

Figure 9: Confusion matrix* for 32 training data points, showing overfitting.

- Best classification results:
  - Validation accuracy: **62%**
  - Validation loss (cross entropy): **0.8048**
  - Test loss (cross entropy): **0.85474**
  - (test accuracy not reported by Kaggle)

*shading indicates the distribution density over predicted labels.

Future Work

- Manually segment the data or train a separate neural net for cervix segmentation to reduce the noisiness of the data
- Experiment with different model architectures
  - deeper models, Inception modules, ResNet architecture, ensemble of three binary classifiers (one for each cervix type)
- More meticulous parameter tuning of learning rate, regularization, and dropout
- Model Ensembling