Cancer Diagnosis using Deep Learning

SHUIJING LIU

DEPARTMENT: ECE & CS

ADVISOR: ADITYA PARAMESWARAN, JIAN PENG

Motivation and Goals

- Why cancer diagnosis is important?
  - According to medical news, cancer is the second most common cause of death in the US (22.5% of total deaths) over the past five years
  - Early diagnosis and early treatment can often prevent the spread of cancer and increase a person's chance of survival
- Why do we diagnose cancer using Deep Learning?
  - Professional pathologists are scarce resources, and pathologists may have conflict opinions on the same patient
  - Automating the diagnosis process can reduce time and costs, minimize errors, and extract a systematic rule
- Our goal is to build a binary convolutional neural network (CNN) classifier to give diagnostic predictions for biopsy images of human tissues

State of The Art

- Ciresan et al. used Deep Neural Network to detect mitosis in breast histology images by classifying each pixel in the images, which won ICPP 2012 mitosis detection competition
- Liu et al. performed tumor detection on gigapixel pathology images using multi-scale CNN, and achieved 97% AUC scores
- Few previous work studied prostate cancer, which I focused on in this project

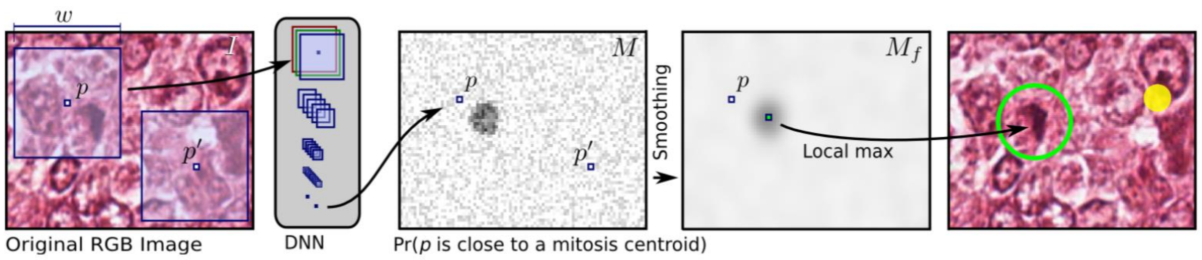

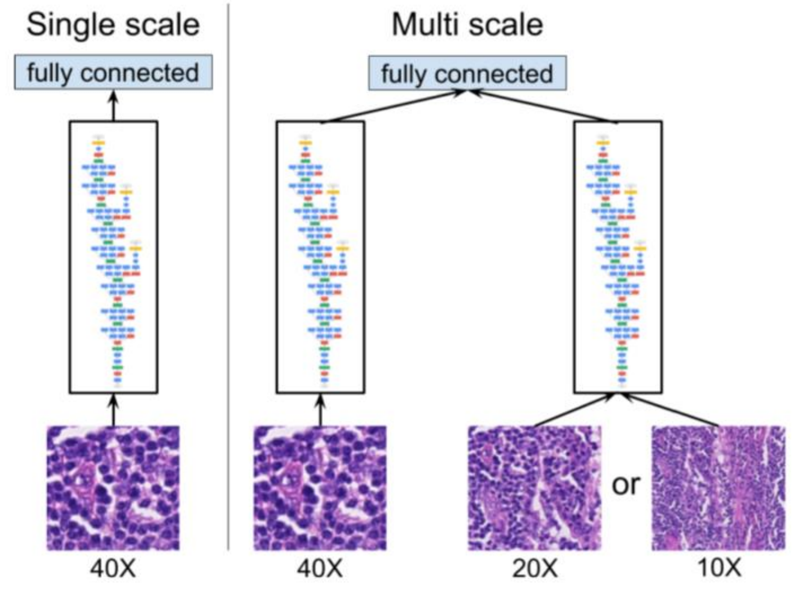

Figure 1 (top): overview of the detection approach of Ciresan et al.
Figure 2 (right): multi-scale CNN model of Liu et al.

Workflow Overview

Image preprocessing:

Randomly sample many 128 x128 patches from biopsy images

Train CNN models with patches as input; predict labels for unseen testing patches

Ensemble methods:

1. Majority voting among patches to give a single label for each biopsy image;
2. Use Bootstrap Aggregating (Bagging) to improve accuracy on images

Datasets

- The first dataset (referred as dataset1) includes ~520 images of cancerous and healthy tissues

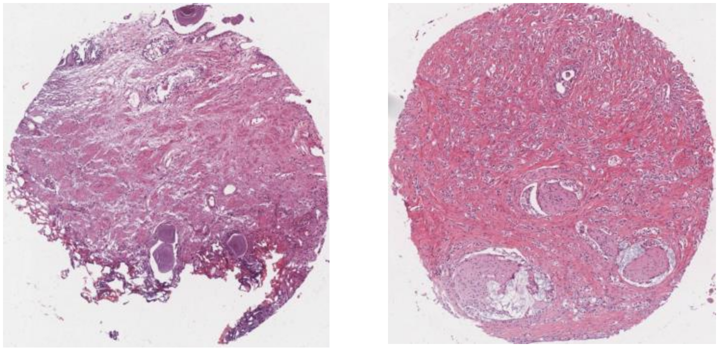

Figure 3: Sample images from dataset1 (Left: biopsy of Healthy tissues, Right: biopsy of cancerous tissues)

- The second dataset (referred as dataset2) includes ~1200 images tissues with recurred cancer and fully recovered tissues

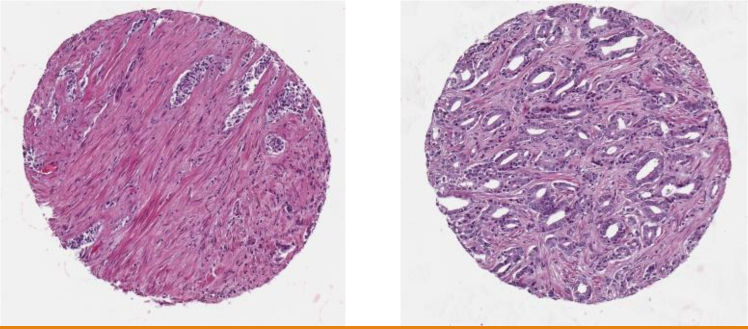

Figure 4: Sample images from dataset2 (Left: biopsy of recovered tissues, Right: biopsy of tissues with recurred cancer)

Image Preprocessing

- Given the biopsy images and their labels from pathologists, we divided the images into two classes -- positive (cancerous) and negative (healthy)
- We detected the boundary of the elliptical tissue areas; for efficient classification of CNN, we randomly sampled 100 patches with size 128 pixel x 128 pixel from red tissue areas
- The patch labels are the same as the label of the image from which they are drawn

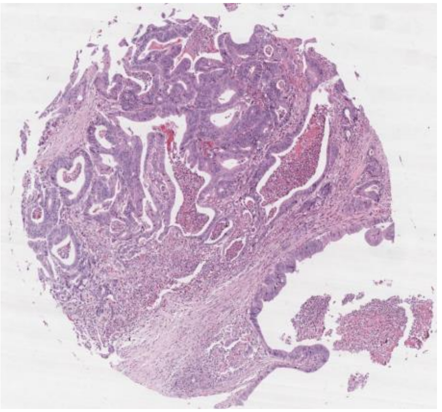

Raw biopsy image

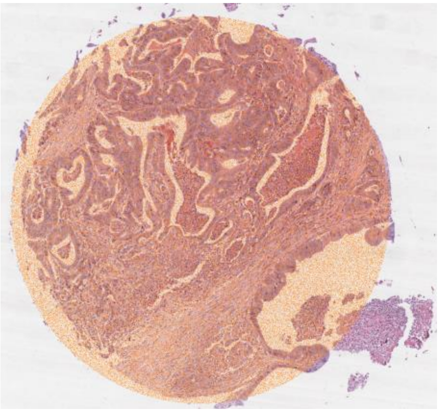

Detected elliptical boundary

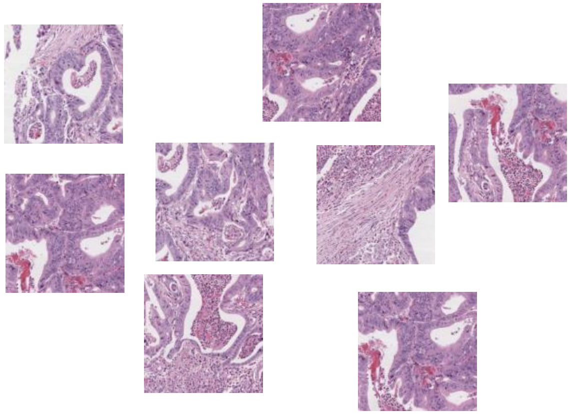

Sampled patches from the image

Deep CNN Models

- A CNN is a class of deep, feed-forward artificial neural networks that consists of multiple hidden layers
  - CNN can analyze visual imagery with minimal preprocessing
  - It is invariant of image transformations such as scale, rotation, and translation
- Previous work by Bangqi Wang used LeNet on dataset1, a simple 7-layer CNN successfully applied to handwritten digit classification
  - Accuracy is around 70%, loss is 0.6.

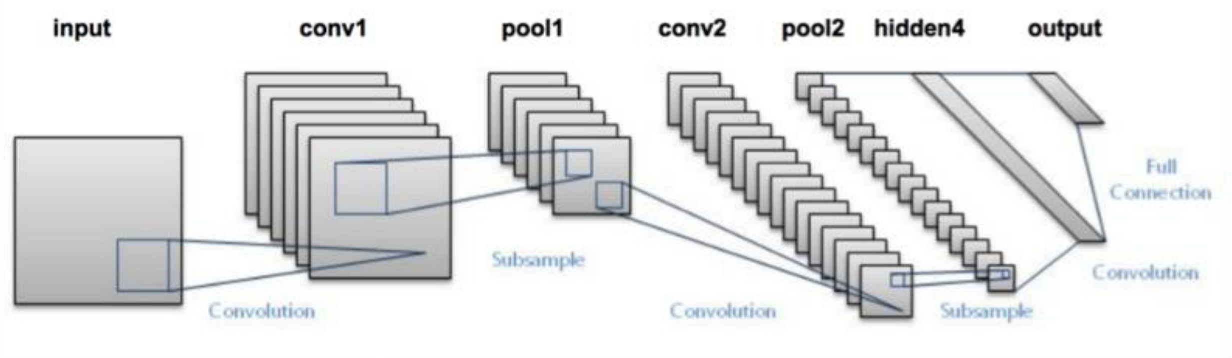

Figure 5: The LeNet architecture

Deep CNN Models

- We explored modern and deeper architectures: VGG net, Inception v3, and ResNet
  - All 3 architectures achieved very similar results: 82% accuracy on dataset1 and 63% accuracy on dataset2
  - Without sacrificing accuracy, we found that **ResNet18 with half of its convolutional layers** needs shortest training time
  - Visualization of different CNN models: http://josephpcohen.com/w/visualizing-cnn-architectures-side-by-side-with-mxnet/

Deep CNN Models

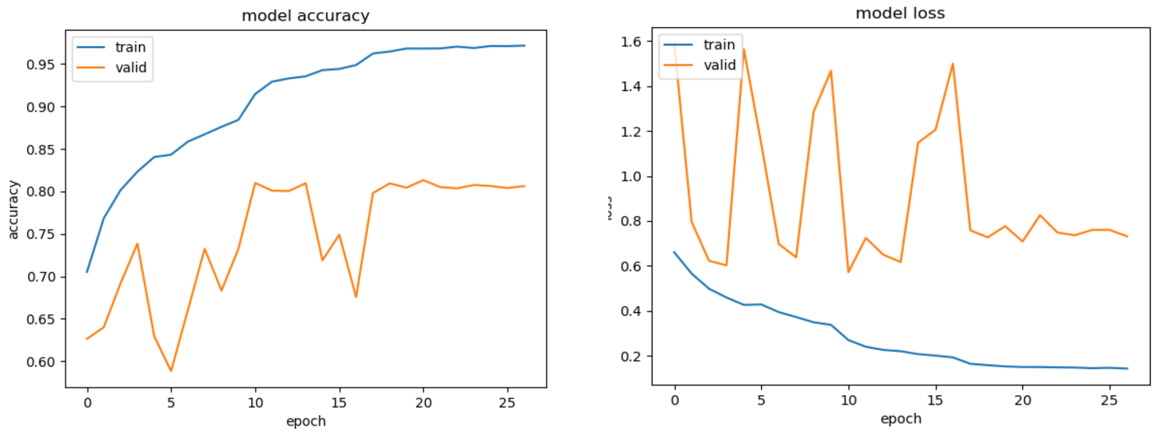

Figure 6: Learning curves of ResNet on dataset1
(Left: accuracy v.s. epoch; right: loss v.s. epoch)

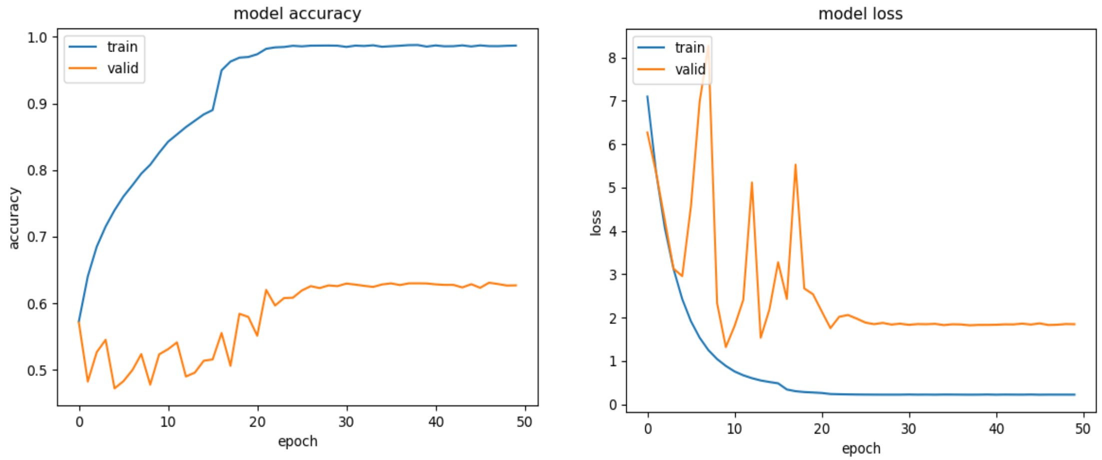

Figure 7: Learning curves of ResNet on dataset2
(Left: accuracy v.s. epoch; right: loss v.s. epoch)

Reduce Overfitting on Dataset2

- As the figures shown, 100% accuracy on training data and 63% accuracy on testing data indicates severe overfitting problem on dataset2
- To solve this, we 1. added Dropout layers, 2. reduced the size of our network, and 3. increased the number of patches for training; unfortunately, none of these methods improved accuracy

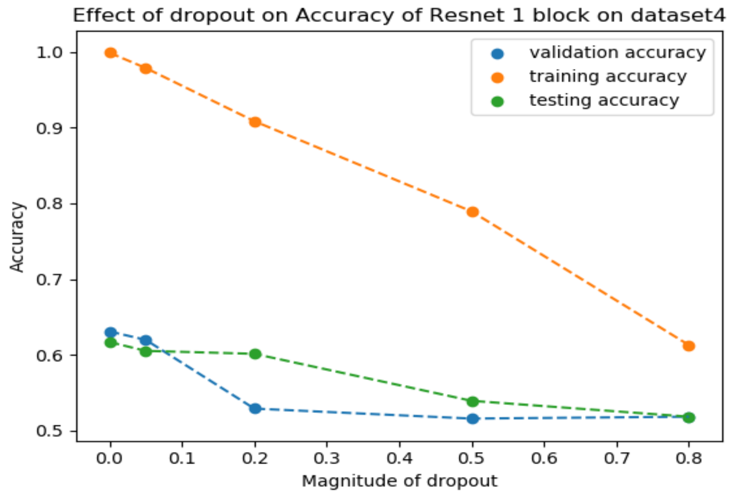

Figure 8: Accuracy v.s. Dropout magnitude for 1 block of ResNet

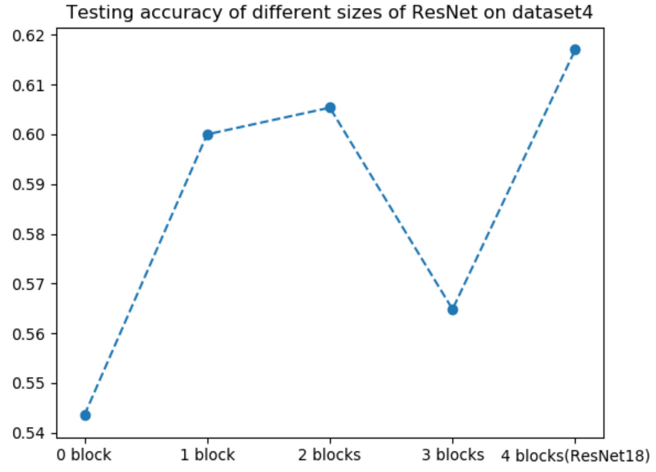

Figure 9: Accuracy v.s. number of layers for ResNet

Reduce Overfitting on Dataset2

- To debug our CNN model, we visualized the predictions returned by CNN models, but no clear visual pattern was found

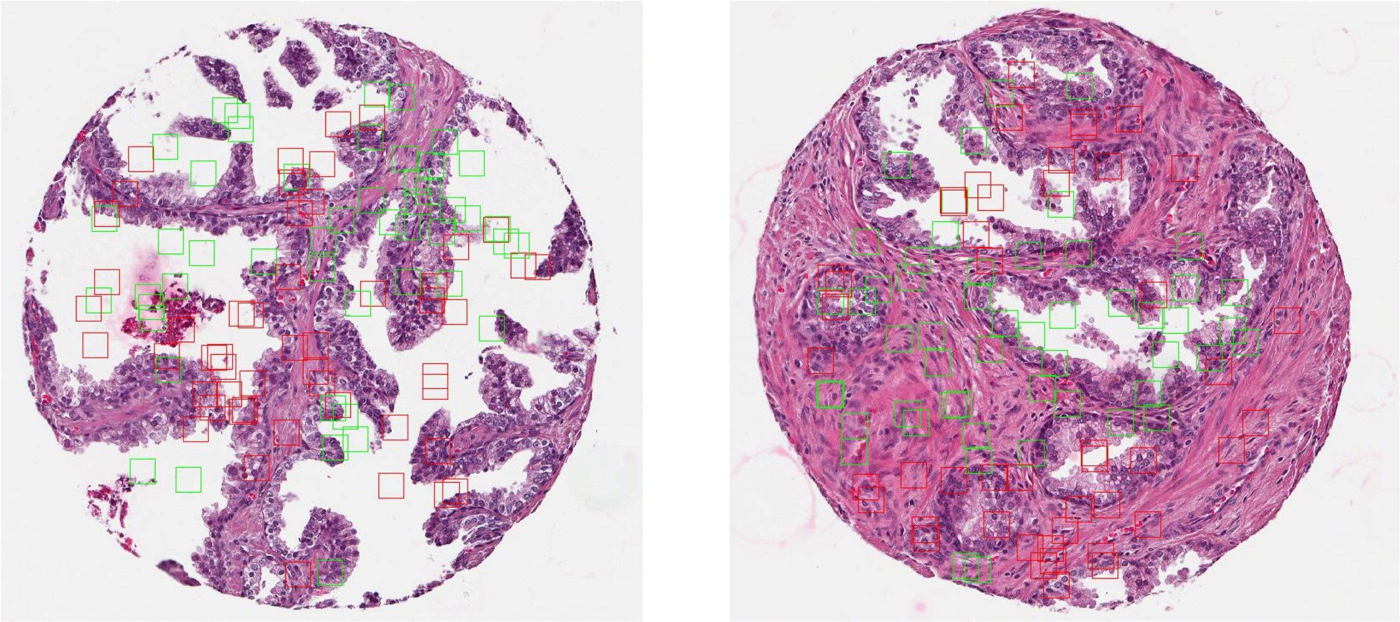

Figure 10: Visualization of predictions on test images (Green: locations of patches with correct predictions, red: locations of patches with wrong predictions

Ensemble Method (step 1): Majority Voting

- We used the patch labels returned by CNN models to predict a label for the image they came from by majority voting
- Accuracy increased for 5% ~ 10% on dataset1, but remained the same on dataset2

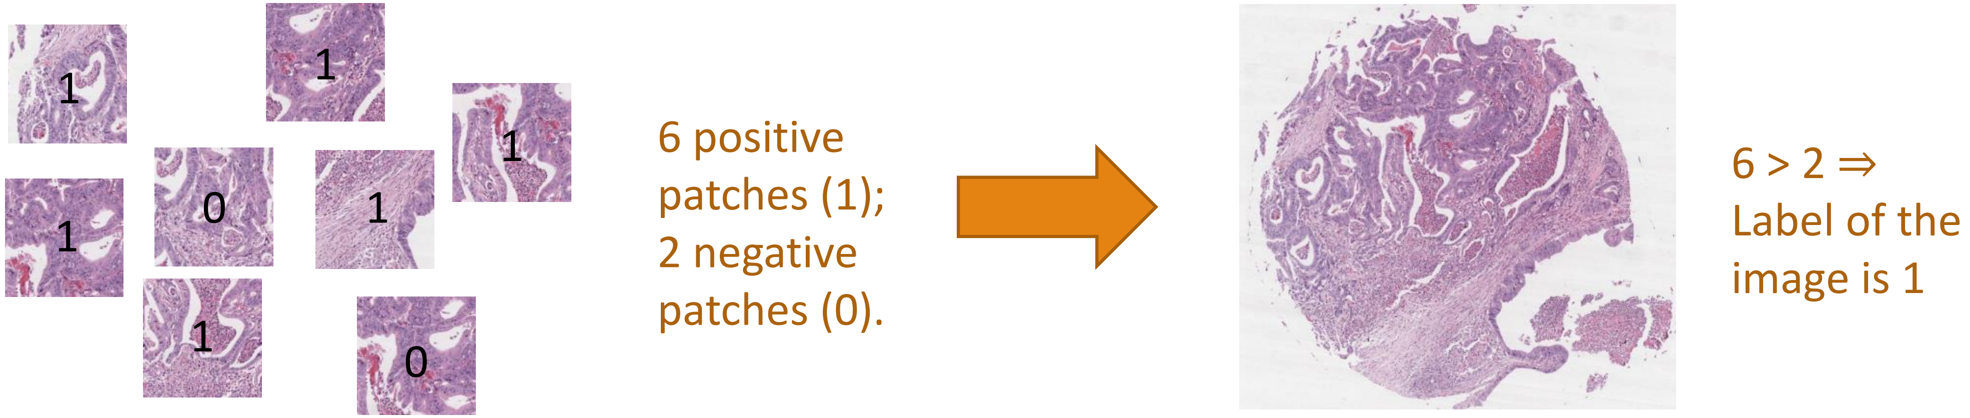

Figure 11: Majority voting process

Ensemble Method (step 2): Bagging

- After we gave labels for images by majority voting, we used Bootstrap Aggregating, or Bagging, which combines many models by averaging the output or voting
- Accuracy increased to nearly 100% on dataset1, but remained the same on dataset2

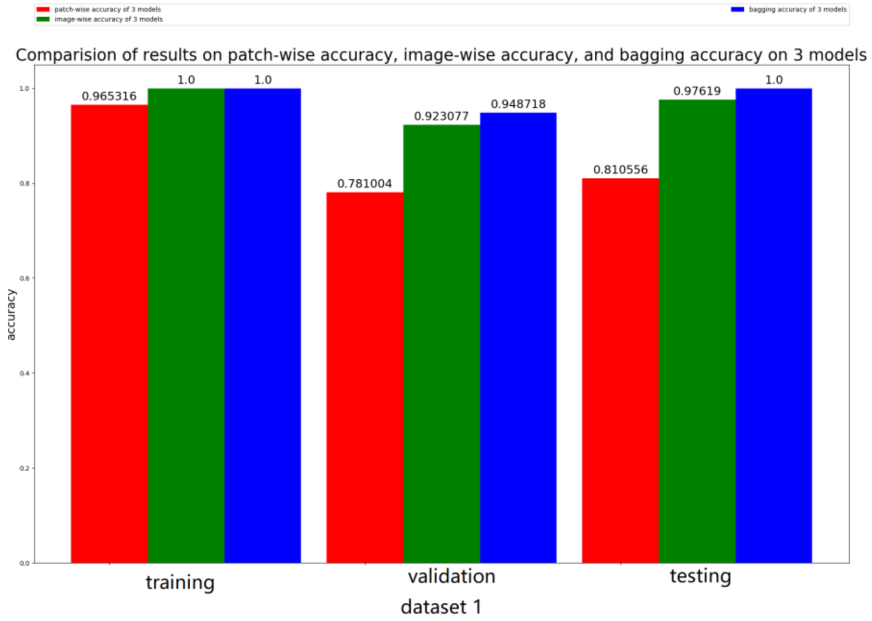

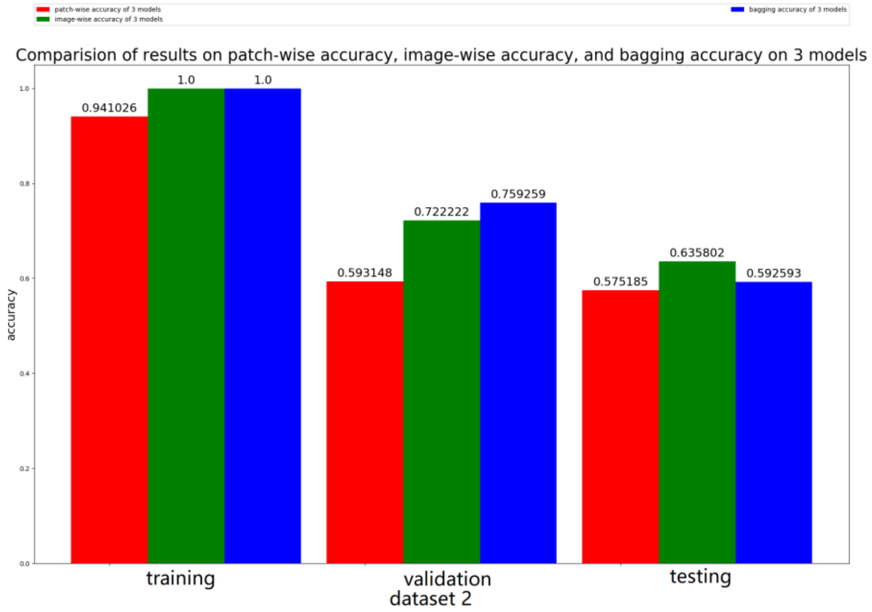

Figure 12: Bar graph of patch-wise accuracy, image-wise accuracy, and Bagging accuracy (3 models) on training, validation, and testing data

Conclusions

- For dataset1, the nearly perfect accuracy obtained by ensembled CNN models is encouraging
  - Deep CNN models achieved strong performance on our data
  - Ensemble methods reduced the variance of predictions and improved the results significantly
- For dataset2, overfitting problem is not solved
  - Need to explore new classification models
  - On the other hand, the image data might be inseparable due to the lack of distinctive features between two classes

Future Work

- Classification on other biology datasets
- Use object detection and segmentation to locate cancerous tissues