Sepsis Severity Prediction Project

Status as of 4/26/22

Background

Sepsis

- What is Sepsis:
  - An extreme immune system response to infection in the bloodstream that triggers inflammation throughout the body
  - Can lead to septic shock– drop in blood pressure, organ failure, and death
- How it's diagnosed
  - Physical findings (fever, low blood pressure, etc.) and lab tests
  - Difficult to diagnose because not all elements identifying infection are always present
- General goal: predict the severity of sepsis based on images of petri dishes containing bacteria and blood cells of a patient

Dataset & Areas of Experimentation

Dataset

Images of Petri dishes:

- Red: red blood cells
- Yellow/brown: portion that is infected by the bacteria
- 492 images: low amount of data

Labels:

- Hemolytic activity: measurement of red blood cell lysis indicating blood cells being destroyed by the bacteria
- Platelet counts (Day 0 and Day 4)
- Methicillin susceptibility of bacteria on plates
- Mortality

Example Images

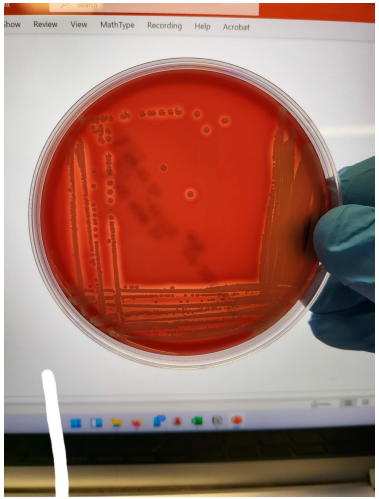

Outcome: Death
Day 4 Platelet: 101
HA: 81.52

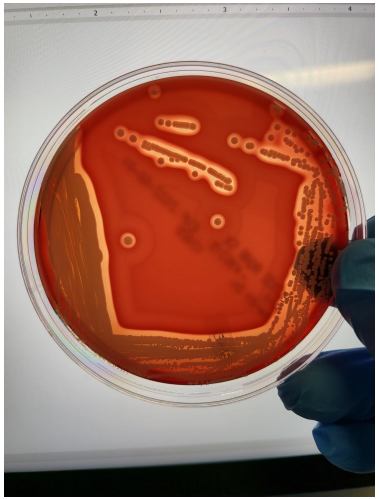

Outcome: Death
Day 4 Platelet: N/A
HA: 41.92

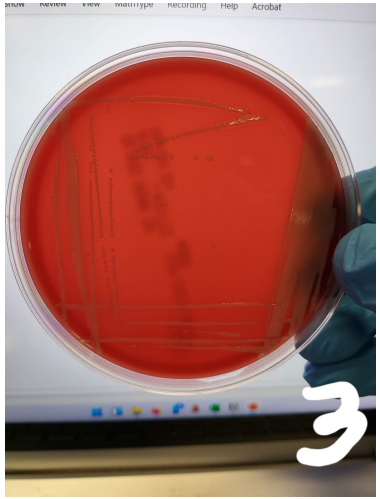

Outcome: Survive
Day 4 Platelet: 225
HA: 0

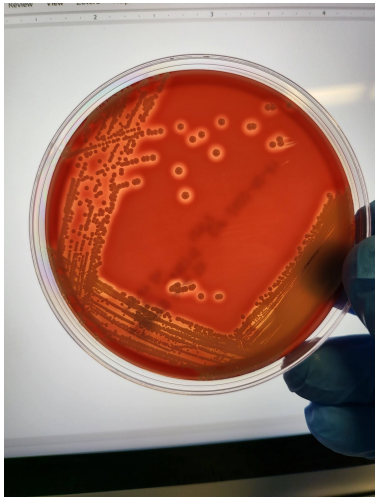

Outcome: Survive
Day 4 Platelet: 255
HA: 18.785

Areas of experimentation

- **Mortality prediction**: will a patient live or die?
- **Day 4 platelet count prediction**: will a patient's platelet count on day 4 exceed some clinically significant threshold
- **Hemolytic activity prediction**: will the hemolytic activity (a measurement of blood cell lysis) of the bacteria infecting a patient exceed some clinically significant threshold?

Image Processing

- Annotated and masked all images
- Reduce image sizes:
  - 3000x4000 → 600x800
  - Cropped to 350x350 to isolate the plate
- Rescale pixel values to range 0-1

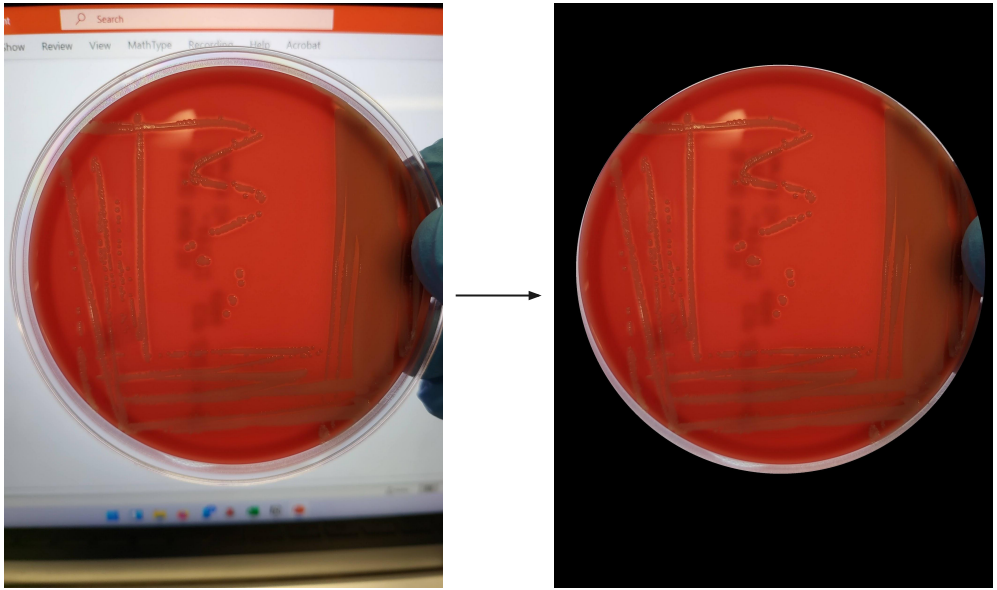

Image Processing, cont. Gabor Filters

- Gabor filters are useful for contour highlighting and texture segmentation
- May accentuate salient features in images: the halos surrounding bacterial colonies on plate

a

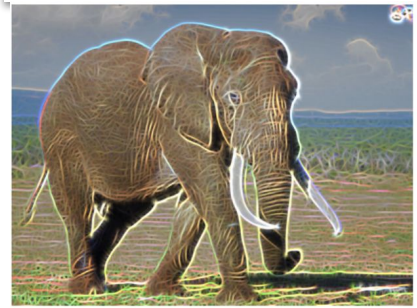

b

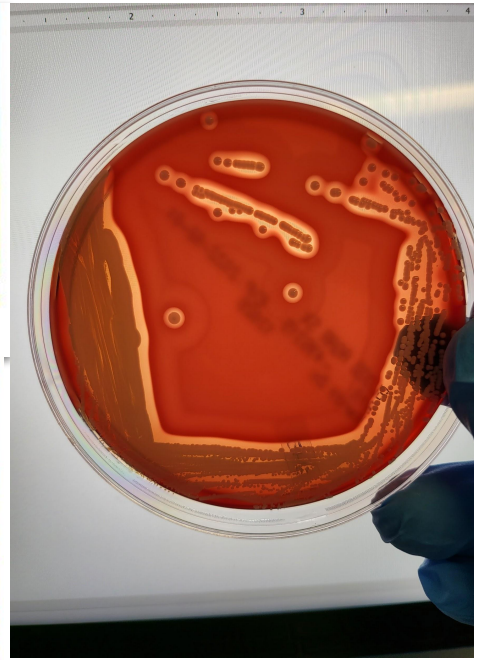

Data Augmentation

- 492 images is low for building classifiers:
  - Need to generate more training data
- Using horizontal and vertical flips of images:
  - Built into Keras ImageDataGenerator library
  - Double amount of data available for training

Model Architecture, Results, and Evaluation

Model Evaluation Metrics

- All results presented are the average across Stratified 5-fold cross validation
- Metrics presented: accuracy, F1, precision, recall, AUC

Deep Learning Models: CNN

Convolutional Neural Networks!

- Image classification and object recognition
- Convolutional Layer: Creates a feature map
- Pooling Layer: Reduces the sample size of a feature map
- Connected Layer: Performs Classification

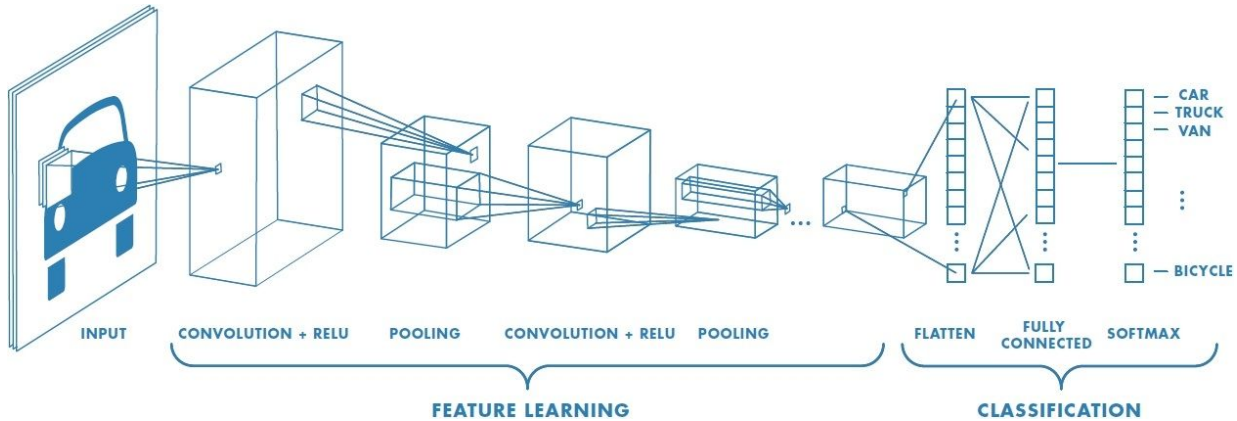

CNN: Model Architecture

- Optional Convolutional layer: 16 3x3 Gabor filters,
- Convolutional layer: 32 3x3 filters, ReLu activation, followed by 2x2 MaxPool Layer
- Convolutional layer: 32 3x3 filters, ReLu activation, followed by 2x2 MaxPool Layer
- Convolutional layer: 32 3x3 filters, ReLu activation, followed by 2x2 MaxPool Layer
- Dense layer: 512 neurons
- Optional Dropout layer (dropout = 0.5)
- Optional Dense layer: 512 neurons

Contour highlighting Basic architecture Combats overfitting

Mortality Prediction

- Predicting whether or not a patient dies
- Binary binning:
  - Class 0 (Survived): 77%
  - Class 1 (Dead): 23%
- Class imbalance remedied by oversampling minority class

Mortality Prediction Results

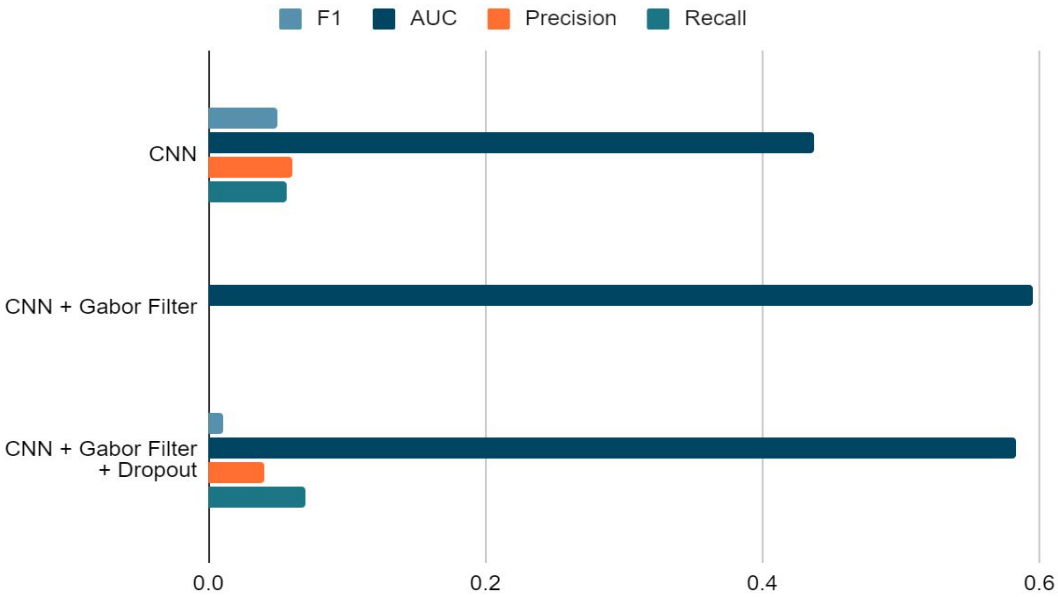

Platelet Count Prediction

- Predicting the patient's platelet range on day 4 of their infection
  - Day 4 platelet count seems to be strongly indicative of the ultimate outcome of infection
- Binary binning:
  - Class 0 (0-150): 36.4%
  - Class 1 (150+): 63.6%

Platelet Count Prediction Results

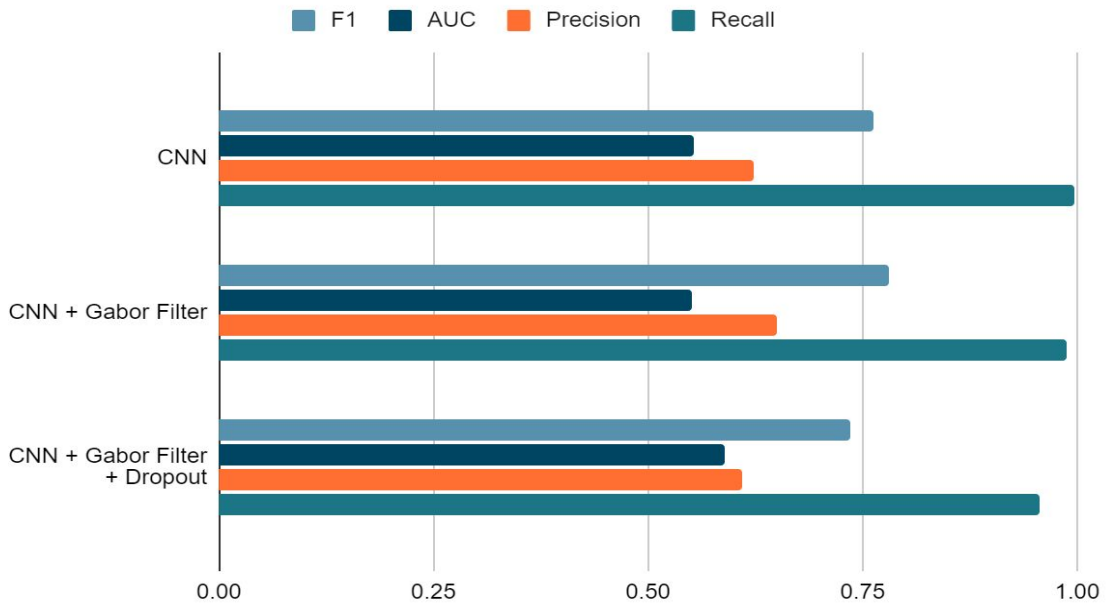

Hemolytic Activity Prediction

- Hemolytic activity is the ability of a pathogen to destroy red blood cells (hemolysis)
- Binary binning: 224 data points
  - Class 0 (0 - 65): 83.9%
  - Class 1 (66+): 16.1%

Hemolytic Activity Prediction Results

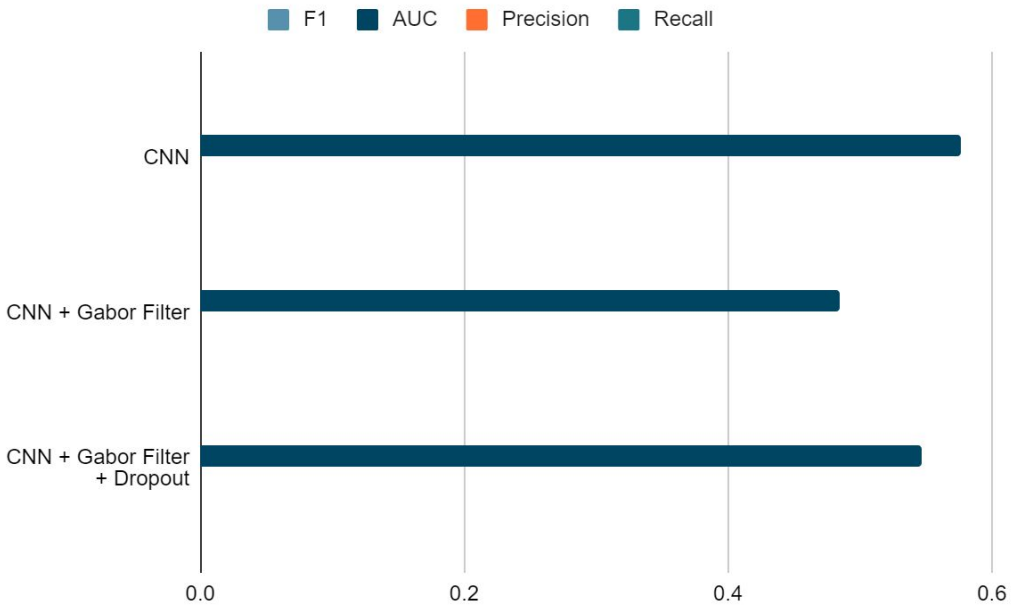

Transfer Learning

Transfer Learning

- Transfer Learning is a technique through which one can use an already pre-trained model on new data
  - Freeze convolutional layers of pre-trained model and use them for feature extraction
- We utilized transfer learning since we did not have a large dataset to train a new model
- All weights were from models trained using ImageNet

Mortality Prediction Results: Transfer Learning

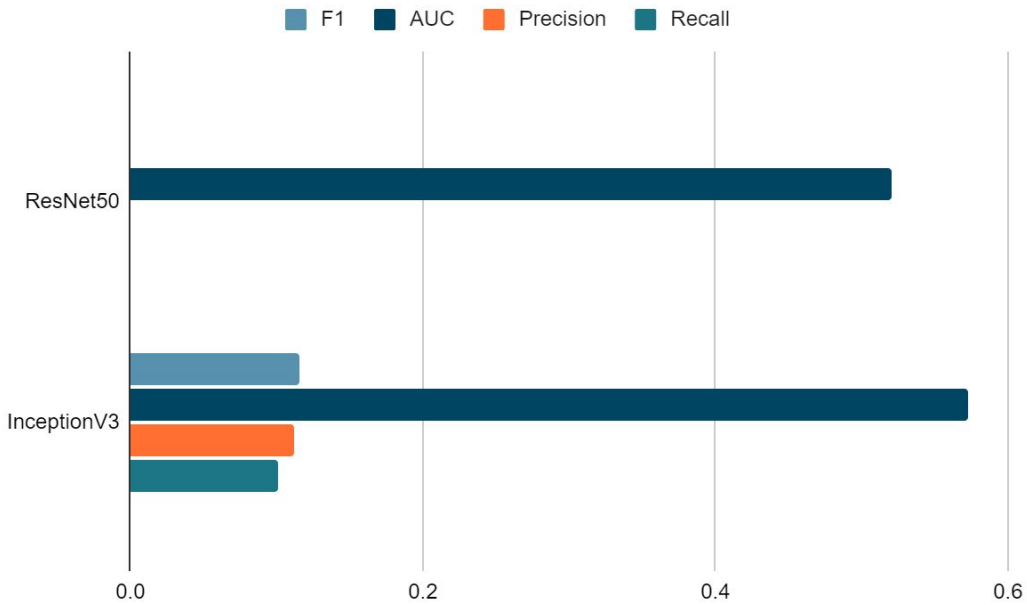

Hemolytic Activity Prediction Results: Transfer Learning

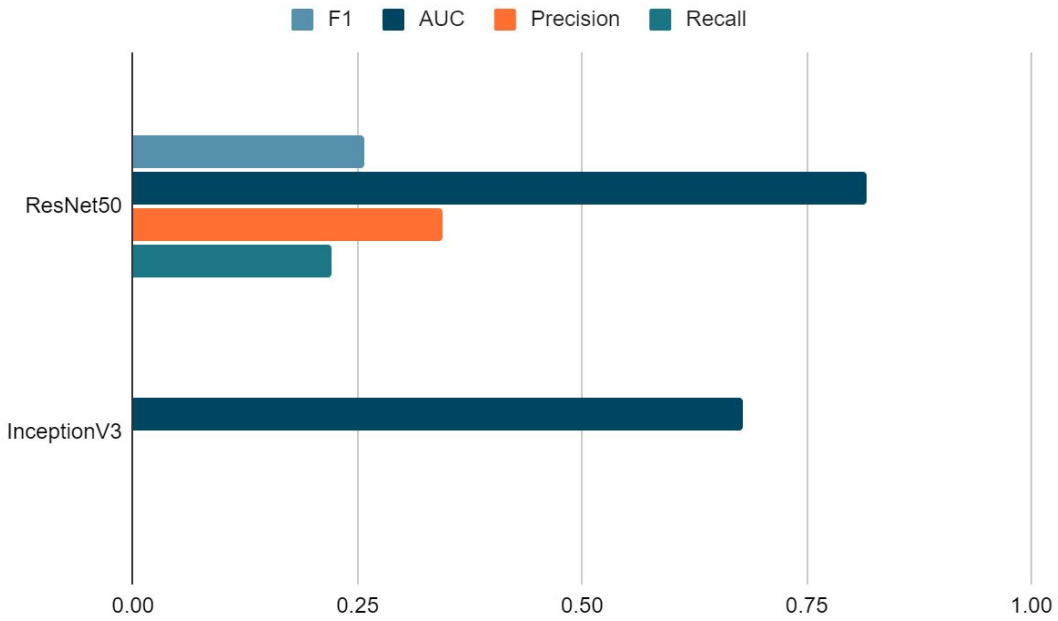

Modeling on the “extremes”

- The following results are from experiments conducted using only the upper and lower 25% of the data (when sorted by hemolytic activity)
  - This reduces the dataset size to 112 images (56 each from upper and lower 25%)
  - Data from the lower 25% is labeled as 0, and data from the upper 25% is labeled as 1

HA Prediction Results: Transfer Learning w/ InceptionV3 + Outlier Data (Top/Bottom 25%)

Metric	5-fold CV Mean	5-fold CV Standard Dev
Accuracy	0.732	0.138
F1	0.734	0.173
Precision	0.722	0.178
Recall	0.822	0.212
AUC	0.881	0.153

Classical Modeling

Binary Classification & Local Ternary Pattern

Binary Classification: Platelet count prediction

- Input Features:
  - Percentage of the image occupied by red blood cell
  - Methicillin Susceptibility
- Platelet count Classes:
  - Class 0: 0 <= x < 150
  - Class 1: x >= 150
- Classifiers:
  - Decision Tree
  - KNN
  - Random Forest
  - Logistic Regression
  - SVC

Feature Extraction: Image Processing

- Used OpenCV and HSV colorspace to detect portion of red blood cells that were infected by the bacteria

- Calculated percentage of the image occupied by the survived red blood cells

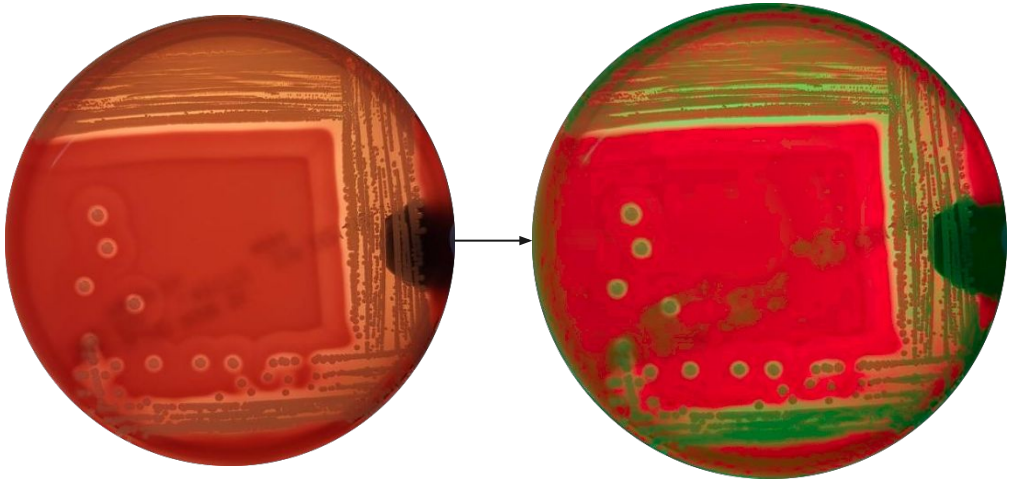

Platelet Prediction Results: Classical (Bacteria Ratio + Methicillin Susceptibility*)

- Logistic Regression (AUC)
  - Stratified 5-fold cross validation
  - Mean AUC: 0.531
  - AUC standard deviation: 0.100

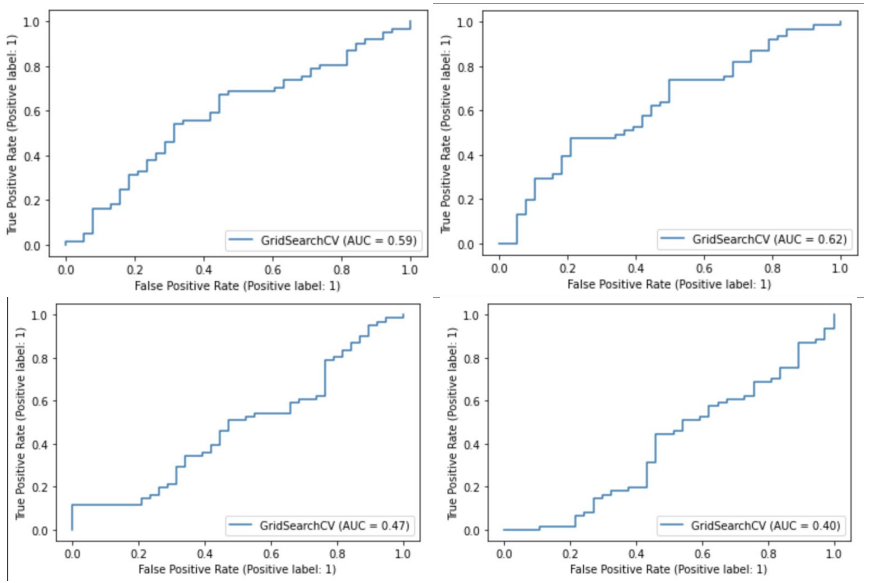

Platelet Prediction Results: Classical (Bacteria Ratio + Methicillin Susceptibility*)

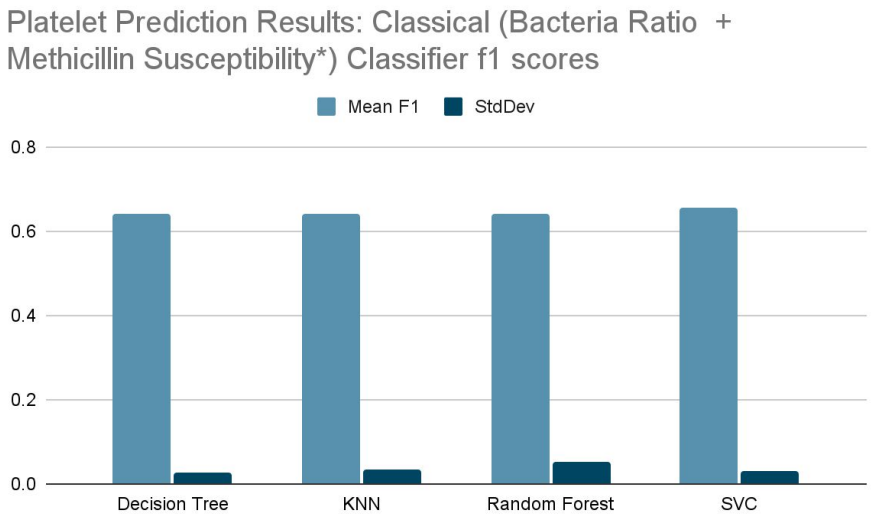

- Other different classifiers: Decision Tree, KNN, Logistic Regression, Random Forest, SVC

Model	Mean F1	StdDev F1
Decision Tree	0.640	0.027
KNN	0.641	0.034
Logistic Regression	0.751	0.026
Random Forest	0.640	0.053
SVC	0.656	0.032

Local Ternary Pattern

Classical Modeling, cont. Local Ternary Pattern

- Technique for segmenting textures in an image
- More noise robust than Local Binary Patterns
- Can yield feature vector of size 3^8 when using 3x3 kernels
    - Amenable to techniques like SVM, tree-based methods, etc

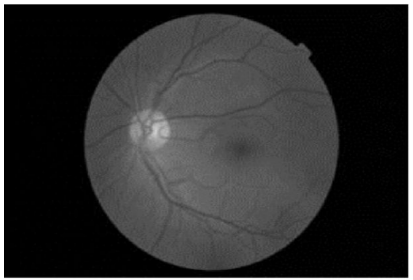

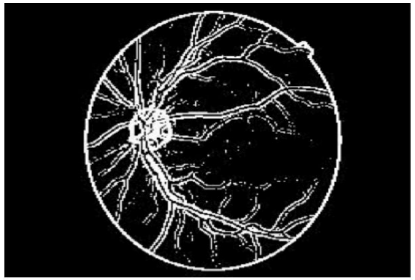

Source for idea

LTP cont'd

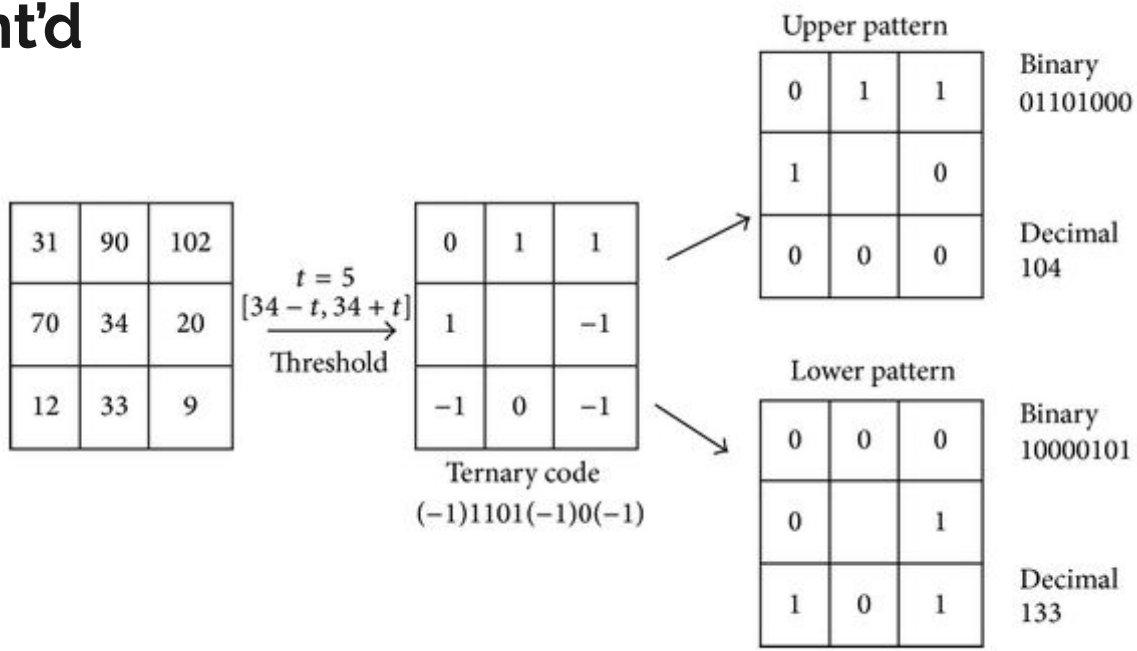
31 90 102
70 34 20
12 33 9
t = 5
[34 − t, 34 + t]
Threshold
0 1 1
1 −1
−1 0 −1
Ternary code
(−1)1101(−1)0(−1)
Upper pattern
0 1 1
1 0
0 0 0
Binary
01101000
Decimal
104
Lower pattern
0 0 0
0 1
1 0 1
Binary
10000101
Decimal
133

Platelet Count Prediction Results: (LTP)

Model	Mean F1	StdDev F1	Mean Precision	StdDev Precision	Mean Recall	StdDev Recall	Mean AUC	StdDev AUC
Decision Tree	0.601	0.036	0.412	0.010	0.396	0.041	0.508	0.019
KNN	0.389	0.011	0.441	0.425	0.383	0.191	0.592	0.041
Random Forest	0.614	0.042	.606	0.025	0.633	0.114	0.486	0.047
SVC	0.388	0.090	0.679	0.198	0.309	0.109	0.592	0.046

Conclusion

- Methods used:
  - Deep Learning: CNN
  - Transfer Learning: ResNet50 & InceptionV3
  - Classical Modeling: Binary Classification & Local Ternary
- Most modeling techniques were not successful
  - Discrimination between outliers of the dataset seems feasible
- Speculation of the underlying reason:
  - Lack of data, noise due to inclusion of samples with zero hemolytic activity

Future Work

- Model without using data points with 0 hemolytic activity
  - May be introducing additional noise
- Model using only methicillin resistant samples
  - Discrimination between classes might be more feasible with this subgroup
- Move towards multiclass prediction
  - I.e. less risk, unsure, high risk