Individual Horse Identification

Brandon Noyes, Mingyi Yang, Hady Ouyang
Faculty Advisor: Eli Shlizerman
Industry Mentor: Evan King

ELECTRICAL & COMPUTER ENGINEERING

Magic AI

PROBLEM STATEMENT

Magic AI produces a horse monitoring solution. As part of this product, they wish to be able to tell when horses have moved stables to provide more accurate analytics on the horses and update permissions for who can see the horse, if necessary. Our task was to develop a system that can accurately detect which horses are in which stalls.

DATA PREPROCESSING

To preprocess our data, we first downloaded the data from AWS, used an R-CNN to create a mask of the horse's location in the image, then saved the masked image for use in our classifier.

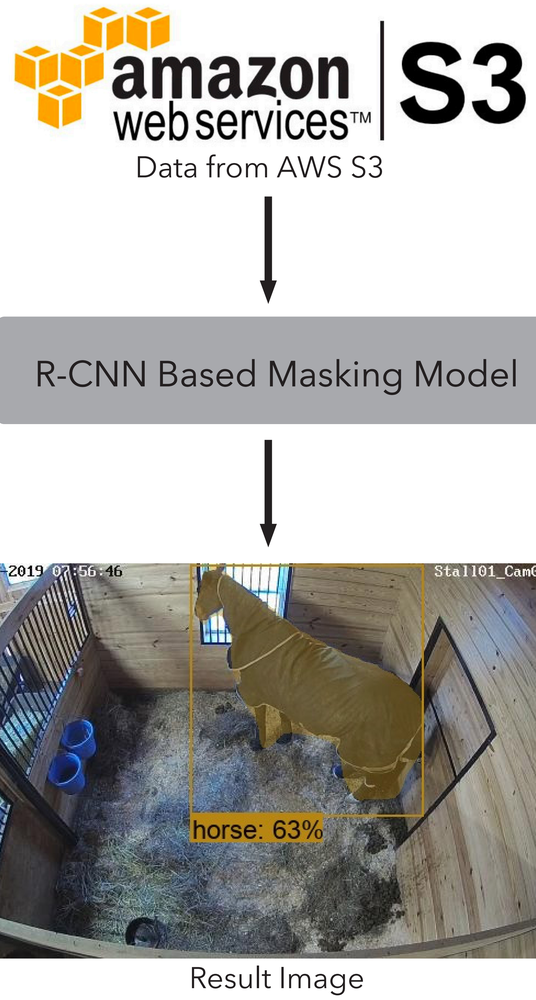

Result Image

CLASSIFICATION MODEL

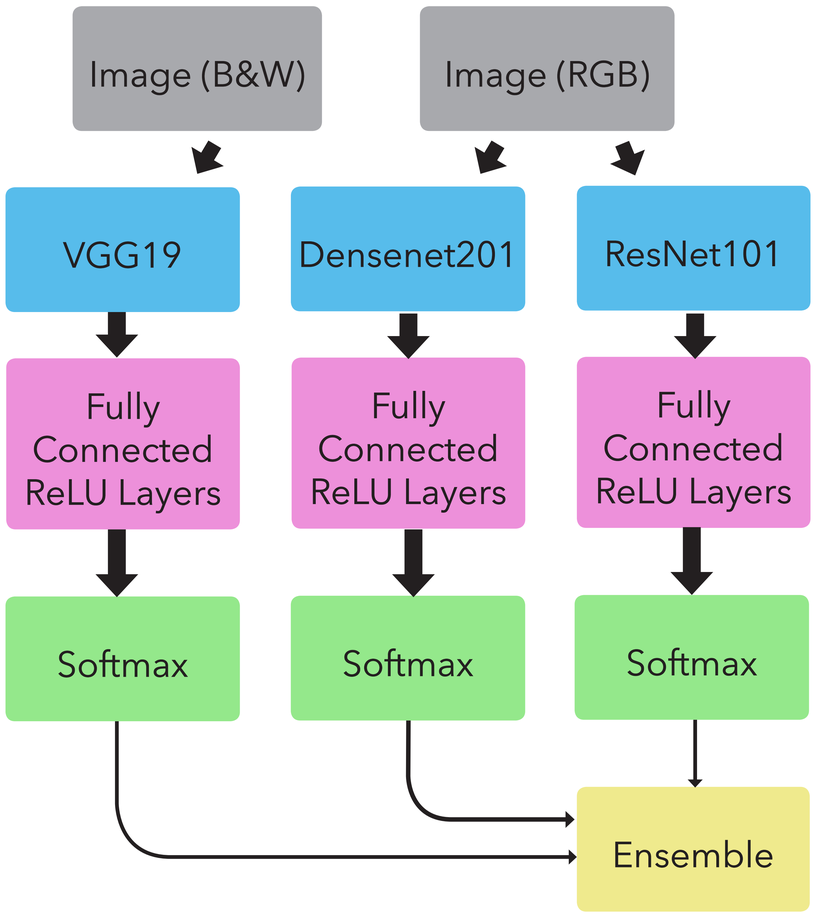

Our classifier consists of an ensemble of three models - VGG19, DenseNet201, and ResNet101. VGG19 is fed black and white images, while DenseNet201 and ResNet101 are fed color images. These networks then feed into their own set of densely connected layers - two layers of 4096 ReLU activation units for VGG19, and one layer of 2048 ReLU activation units for DenseNet201 and ResNet101. These are then fed into the final softmax layer for each model, which are then weighted and combined to form the final ensemble network.

To train these models, we used center loss as our loss function, which allows the features in the penultimate layer layer to cluster, providing some performance benefit.

RESULTS

Our results were a bit middling, as can be seen by the example confusion matrix below; while our accuracy hovered in the mid 60% range, well above the random guessing result of 10%, it is still not super accurate. However, we do not need to immediately know when horses switch locations. Thus, we can use an average of multiple frames throughout time to perform our prediction, smoothing out wrong predictions. This can be seen by the blue line in the graph below. This significantly helps our accuracy.

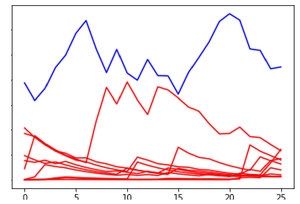

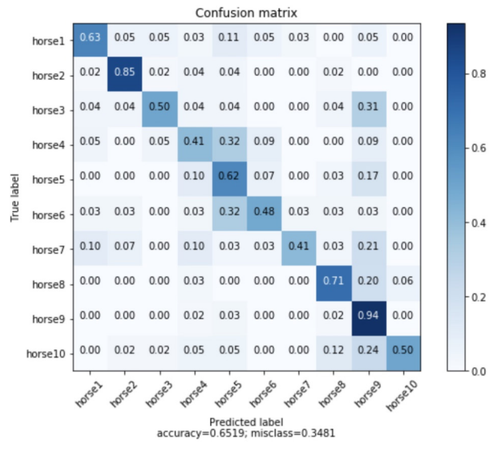

FUTURE WORK

- Improve our ensembling technique by applying a gradient boosting algorithm or secondary neural net to the second to last layers of the individual classification models, rather than using a weighted average of the softmax results.
- Further tune the models we are using to improve individual performance
- Create a more robust training set for the masking neural net so the horse is found more often, and the pixels on the mask are more accurate to the horse.
- Use knowledge of where other horses are to improve prediction results