Large Animal Hazard Detection

Using computer vision to prevent collisions with livestock

Kipp Sherer; Stanford University, Stanford, California 94305

The Problem

Every year, vehicle collisions with large livestock cause untold damage to life and property.[1] In Arizona, like in other open-range states, many drivers don't realize the real hazard and frequency of livestock on the road until it is too late. Using video to detect the hazard and notify drivers of impending hazard with alert activated signage could be a cost effective method to prevent accidents.

Main challenges for this project are the availability of a realistic dataset and the variation of images that could be inputs. For instance, identifying cows at night, or mixed in with vegetation.

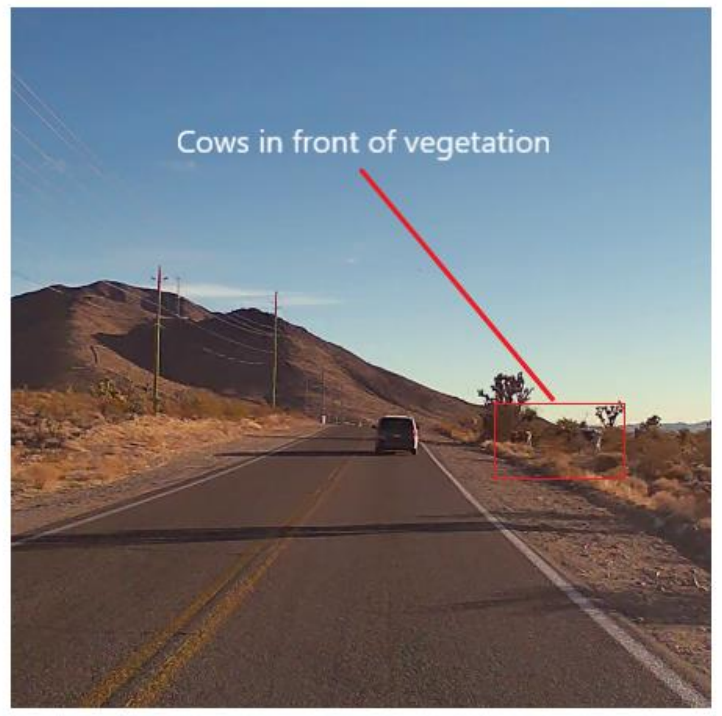

Background

Studies have shown that motorist alert signage that is activated specifically at the time and place of the hazard have higher effectiveness in preventing collisions with wildlife than signs alone.[2] Current radar based systems activate whenever there is a car even if there is no current animal hazard, essentially a false positive from the point of view of the motorist.

Using video to detect a hazard and notify drivers of impending hazard with alert activated signage could be an effective method to prevent accidents.

Contact Information:
Kipp Sherer ksherer@stanford.edu

Methods

Datasets

The initial plan for a dataset was to use a derivative of the public image datasets, that approach ended up having several critical problems.

The images were very different than what is expected to be encountered. For instance, the bottom right example is an image collected.

Creation of datasets were achieved by extracting video frames from a vehicle dash cam and applying data augmentation techniques.

Random shifting, random rotation (up to +/- 12 degrees), and image reflection. For additional analysis, images were labeled with whether they were day or night as well.

Datasets	Total Images	Positive	Negative
Training	54,274	14,962	39,312
Dev	3000	1,000	2,000
Test (Complete)	3000	1,000	2,000
Test (Dark)	1,568	224	1,344

Models

The approach used was to try three different architectures. Base models were derivatives from either MobileNetV2[3], ResNet101[4], or custom CNNs.

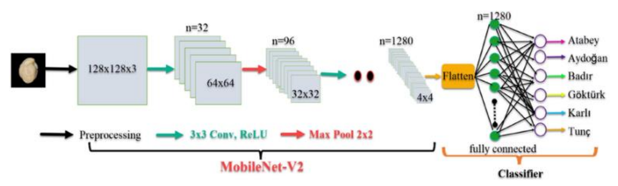

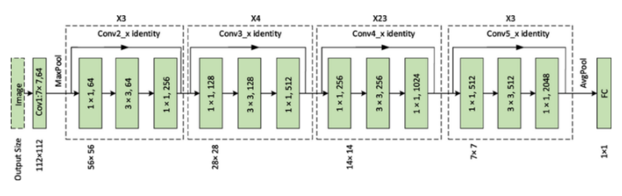

Each model takes a 224x224x3 image for input and outputs a binary classification of whether or not a livestock hazard is present. Models were trained using binary cross entropy for the loss function, Adam optimizer, and adjusting amount of training and learning rate for performance.

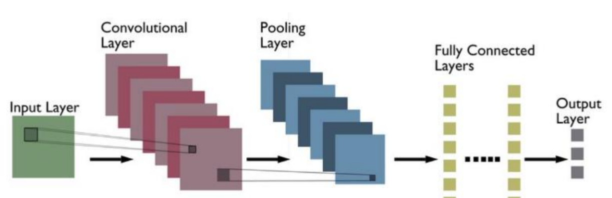

Top Experiments

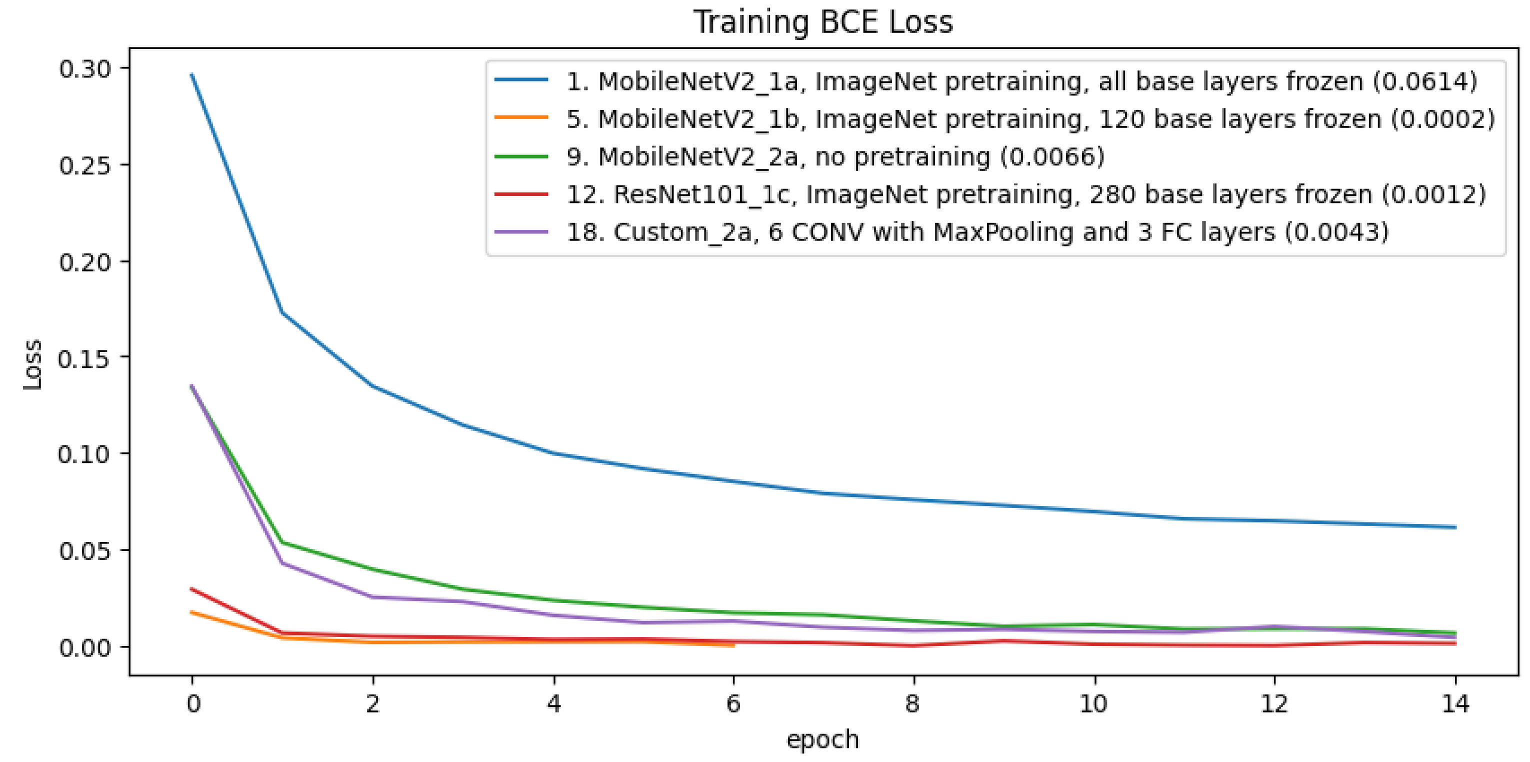

Analysis

The results of the experiments far exceeded the target goal of an accuracy above 85% and an optimal recall. A factor in why these results were achieved could be that the image data collected did not have as much variation as would have been preferred. If the images in the test set were very similar to the training data, even if rotated or shifted, they may be easy for the models to detect.

To further differentiate the performance of the models, they were also evaluated against the subset of the test data consisting of night images. In most cases, there was only a slight decrease in performance. Much less than expected. Two notable exceptions were the MobileNetV2_1a_nuky model which had a 0.04 drop in recall and the MobileNetV2_2a_etti model, which had a 0.15 drop in recall.

Conclusions

1. **The model architecture is not as significant in the performance as far as this problem and evaluation were defined. Derivatives of each very different base model were able to achieve comparable results for this task.**
2. **Using models pretrained with imagenet did offer some improvement in the end result over the same model, trained from scratch.**
3. **Adjusting the image standardization made a significant improvement in the custom CNN models. This could be in part because the dark images would normalize to near zero using a straight factor method.**

References

1. Transportation Systems Management and Operations. ARIZONA MOTOR VEHICLE CRASH FACTS 2023. https://azdot.gov/sites/default/files/2024- 07/2023- Crash- Facts.pdf. July 2024.
2. Todd Williams et al. Arizona Statewide Wildlife-Vehicle Conflict Study. https://azdot.gov/sites/default/files/media/2021/08/WVC_Final_Report_July30_2021.pdf. July 2021.
3. Mark Sandler et al. MobileNetV2: Inverted Residuals and Linear Bottlenecks. 2019. arXiv: 1801.04381 [cs.CV].
4. Kaiming He et al. "Deep residual learning for image recognition". In: Proceedings of the IEEE conference on computer vision and pattern recognition. 2016, pp. 770–778.