desert conservation
PROGRAM

Final Project Report

Executive Summary

Vehicle-induced fatalities pose a considerable risk to desert tortoise (*Gopherus agassizii*) populations near Nevada highways. Water culverts under highways offer an alternative crossing method, though the extent of their use by tortoises is unknown. For this project, we used nearly one million wildlife camera images from 18 different culverts to identify desert tortoise occurrences under Highways 93 and 95 near Las Vegas, Nevada. With a custom annotation software, we manually labeled 231 tortoise images from the project and an additional 393 desert tortoise images from the iNaturalist and GBIF databases. We fine-tuned Google's Inception v3 image classification model to detect tortoises and enhanced the training with data augmentations to simulate different travel directions and lighting conditions. We also developed a workflow to manage image metadata and added camera-reported temperatures. The fine-tuned model classified 564 images containing desert tortoises, and all results were combined with metadata for all images in a PostgreSQL database. This approach saved a significant amount of time and resources over manual inspection of the image set and the resulting detections can be used to analyze the biological importance of tortoise occurrences.

Introduction

Roads are a mainstay on the landscape throughout America and in southern Nevada; however, roads can have an extreme negative effect on tortoise populations. Increased mortality on, and population depletion adjacent to, roads and highways contributes to reduced connectivity and gene flow from one side of the road to the other. One solution to that problem has been to install fences to keep tortoises off roadways and allow for that habitat to be utilized. This solves the mortality issues but still leaves the issues of connectivity and gene flow unaddressed. Culverts are in place to allow water to flow from one side of the road to the other but very little is known about how tortoises may use these culverts for connectivity. This project will examine photos from trail cameras placed inside culverts to identify what species use the culverts for crossing and how often it occurs.

The objective of this project is to process trail camera images from culverts along Highways 93 and 95. SWCA trained a convolutional neural network (CNN) to recognize desert tortoises in trail camera images. This model was then used to detect tortoise occurrences as they cross under the highways. The deliverable will be a spreadsheet of tortoise detections, with image file names, dates, times, and temperatures as reported by the cameras. The total number of images to be processed is 1,000,000.

Methods and Materials

We used over 984,000 wildlife camera images collected between 2018–2021 and from 20 different culverts to identify desert tortoise occurrences under Highways 93 and 95 near Las Vegas, Nevada. With a custom annotation software, we manually labeled 231 tortoise images from the project and an additional 393 desert tortoise images from the iNaturalist and GBIF databases. We fine-tuned Google's Inception v3 image classification model to two classes: tortoises versus not_tortoise and enhanced the training data with random augmentations to simulate different travel directions, lighting conditions, and image qualities with varying levels of rotation, zoom, pixel noise, and contrast. Eighty percent of the training data was used for training, while 20% was used for validation after each iteration of the network; this CNN was ultimately constructed to classify each image as one of the two classes and output a probability for the classification (Figure 1). We optimized the model to recover true positives, at the cost of some false positives that needed manual validation. The model was developed in a Jupyter notebook with Tensorflow and Keras in the Python language. Training, validation, and inference were all performed on an Amazon Web Services cloud computing instance. We also developed a workflow to manage image metadata and added camera-reported temperatures with automated optical character recognition, using R packages 'exifr', 'stringr', and 'tesseract'.

Tortoise Classifier Architecture

Type	Patch size/stride or remarks	Output size
Input	NaN	299x299x3
Conv	3x3/2	149x149x32
Conv	3x3/1	147x147x32
Conv padded	3x3/1	147x147x64
pool	3x3/2	73x73x64
Conv	3x3/1	71x71x80
Conv	3x3/2	35x25x192
Conv	3x3/1	35x35x288
3x Inception	Block 1	17x17x768
5x Inception	Block 2	8x8x1280
2x Inception	Block 3	8x8x2048
Pool	8x8	2048
Linear	Logits	2048
Dropout	NaN	2048
Linear	Dense	1024
Linear	Classifier	1

Figure 1. Convolutional Neural Network architecture. This diagram shows the different layers used in the deep learning process. Images were reshaped to 299x299 pixels, and the final output was a classification to either 0 or 1, with a probability between 0 and 1.

Results and Evidence of the Results

The final model recovered 99.8% of the tortoise images and 86.3% of the not_tortoise images in the validation data set (Figure 2). False positives were typically a pair of checkered shoes with rounded edges or dark objects with a rounded shadow from the culvert above. This required a bit more manual validation of predicted tortoises, which was a fair tradeoff to ensure that the model was tuned to find all the tortoises in the data set.

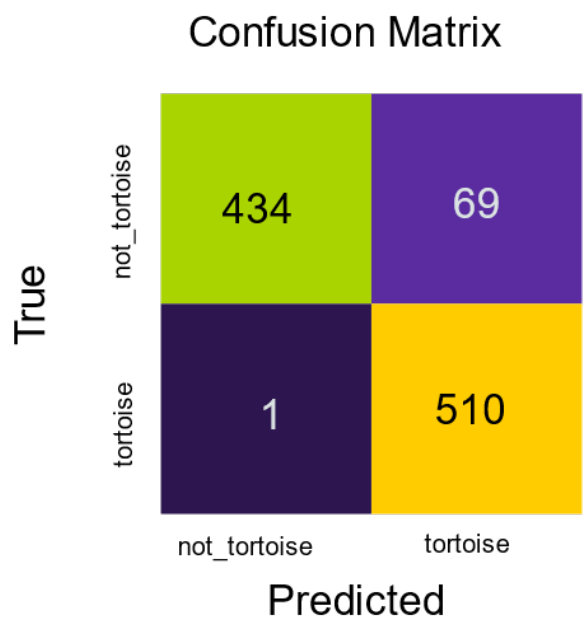

Figure 2. Confusion matrix depicting false positives and false negatives in the 20% validation data set used to test the final model. The model recovered all but one of the true positive tortoise images.

The model was able to detect a total of 564 tortoise images with a high average confidence (93.5%) in those classifications (Figure 3). Some of the lower probability classifications were tortoises at the far edge of a picture frame, with only the outer edge of the shell showing; these were detected during manual validation of the results and added to the final detection set.

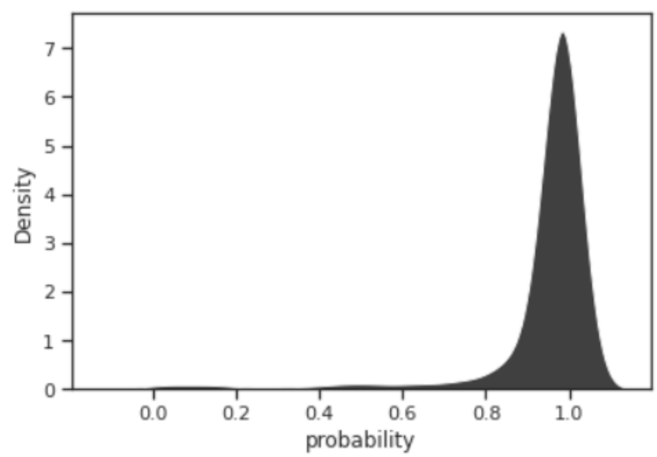

Figure 3. Distribution of probabilities in the 564 tortoise classifications, as a metric of the model's confidence in those determinations.

Tortoise detections were from 12 different culverts: five under highway 93 and 7 under highway 95; culvert 1 had the most activity at highway 93, whereas culvert 7 had the most activity on highway 95 (Figure 4). Detections were from all three seasons, but mostly in Spring (Spring: n = 415, Summer: n = 115, and Fall: n = 34).

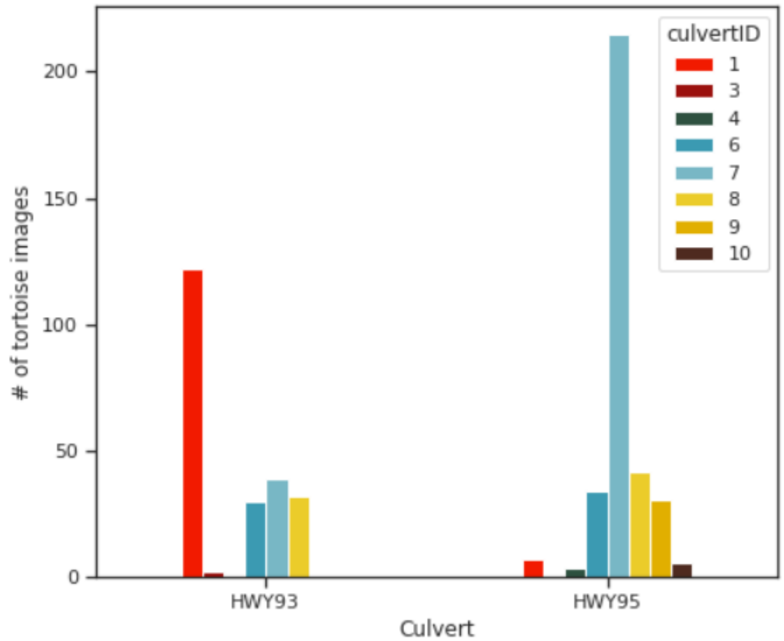

Figure 4. Bar chart displaying the number of images classified as tortoise at each culvert with detections.

Daily timing of detections was bimodal with tortoises mainly moving between 6 AM – 11 AM or 4 PM – 7 PM (Figure 5) and temperature during movement was typically between 75–95 °F (Figure 6).

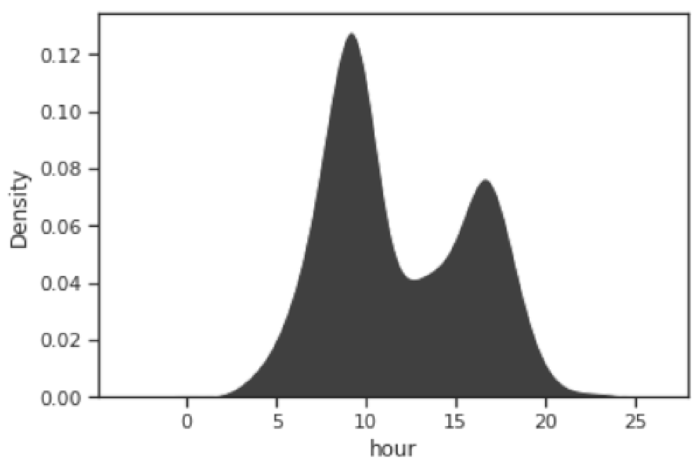

Figure 5. Distribution of movement times when tortoises when detected. Zero is midnight, 10 is 10 AM, and 20 is 8 PM.

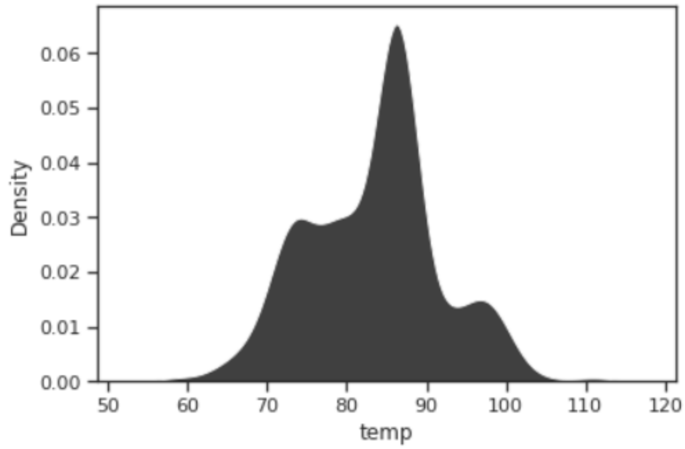

Figure 6. Distribution of temperature when tortoises were detected, in Fahrenheit.

Evaluation/Discussion of Results

Taking the approach of leveraging existing models allowed us to fine-tune the neural network to focus on the new class, Mojave Desert Tortoise. Google's Inceptionv3 model was trained on 1000 different classes, many of which were animals. However, no tortoises were included in the base network. Because the base network was already able to distinguish 1000 classes, we were able to retrain the last few layers to focus on tortoises and classify everything else as not_tortoise.

Using available tools, we were able to save months of time pulling metadata from image embeddings and from the surface of the images to add to the final results file. Results from this study can now be combined with GPS tag studies for biological inference about tortoise movements.

Conclusion

During the process of building this model and validation of the results along the way, we observed many other species, including snakes, small mammals & birds, badgers & foxes, bats, cats, lizards, roadrunners, and even several humans (excluding workers changing SD cards). Similar approaches to monitor wildlife with trail cameras and programmatically process the images should be feasible using the workflow developed in this project.

Recommendations

We observed many images that were triggered by vegetation or trash directly under cameras. If possible, instruct future workers to ensure any loose material be cleared out of the area near cameras to avoid sequential false triggers. Alternately, some newer trail cameras contain temperature (PIR) sensors that may help; however, they may not work well in hot climates when there might not be much of difference between animal and ambient temperatures. The most common other wildlife observed were snakes, lizards, ground squirrels, rabbits, and small birds; similar models could be developed for those species. Another approach would be to develop a multi-class model that could detect broad classes of animals: small mammals, medium-sized mammals, birds, and reptiles. Such a model could be used to first filter out the large number of false triggered images and sort the true animal images into different folders for either manual inspection or further development in a species-specific model. There are still 442,125 images from 2017 that could be processed with the tortoise classification model to add more detections to the 564 found thus far.