



Using Pre-trained Convolutional Neural Networks to Classify Wildlife Animals

Jimmie Hagle, Gurman Gill

Department of Computer Science, Sonoma State University, Rohnert Park

Introduction:

In our research we implement a machine learning algorithm using pre-trained convolutional neural networks to classify wildlife animals. The term "trained" refers to the models algorithm weights that are adjusted to better classify a given set of training images. We compare classification accuracy using models trained on individual trail cam images versus images of all cameras. The hypothesis here is that individual cameras will perform better even though their datasets are smaller due to similarity in the data.

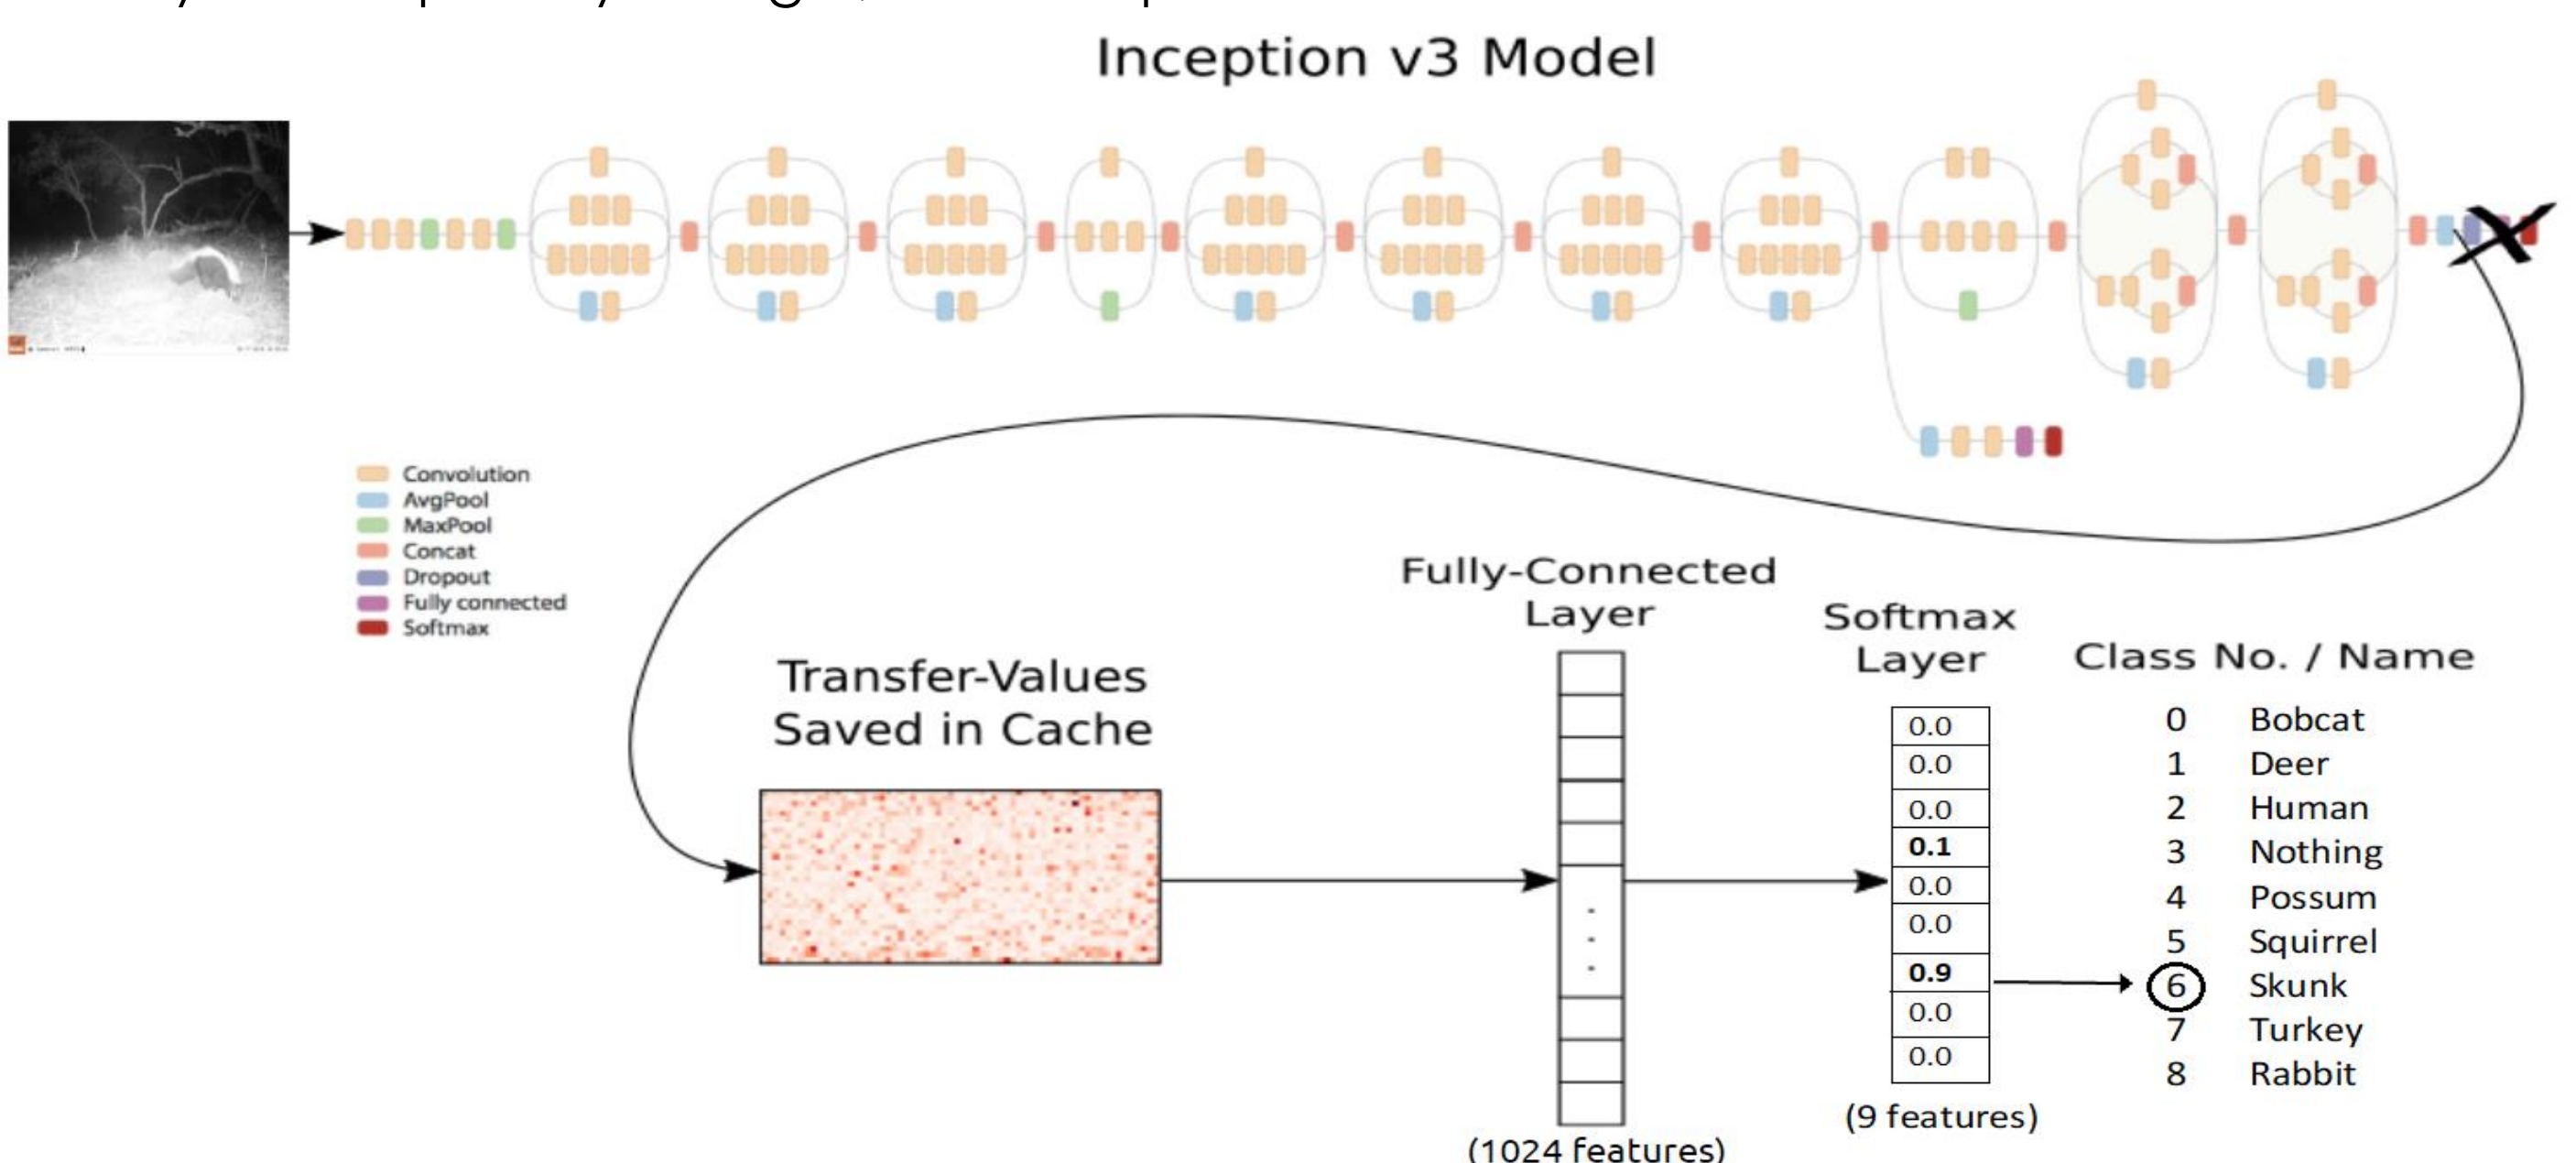
Dataset:

The dataset we used for the experiments is composed of 24973 images taken by trail cameras on SSU Preserves. The trail cameras take a series of three images every time they detect motion. The dataset contains twenty three classes of images across six camera groups. Classes with a minimum of 200 images were used in these experiments reducing the number of classes to nine. This is done to prevent a negative influence on the systems classification accuracy by excluding classes that have a low amount of images that can be trained on. All the images below are examples of the database and underneath the class name is the confidence the trained CNN had in its prediction



Deep Learning:

In this work, we explore using convolutional neural network (CNN) which are successive "layers" of neurons that learn to respond to image features such as color, texture, edges. These are combined in later layers to detect larger concepts such as faces, animals, etc. In particular we implement "transfer learning" a method where features pre-trained by the CNN are pulled for the new classification system to recognize images in your own database. In order to achieve this we used CNN model architecture called InceptionV3(Google Inc) which has been pre-trained on "imageNet" a database containing 1000 different categories. In order to complete the retraining process, we used an example script provided by TensorFlow, an open source machine learning software library developed by Google, and adapted it for our own case.



Experiment Setup:


Base Experiment: In this experiment we are setting up results in which to compare our other two tests. This is run using all six camera groups with class that contain more than 200 images.

Experiment 1: In this experiment we split the images among their individual camera groups but continue to use the same nine classes.

Experiment 2: In this experiment we apply a 200 image minimum to each animal class for the individual cameras.

Evaluating on Individual Cameras:

In order to try to improve model accuracy, we try to obtain classifications on individual cameras. The reasoning behind this is that even though the size of individual datasets would decrease, the influence by the various backgrounds would decrease as well. What we consider as background is anything else in the image that is not the animal we are trying to classify. In some cases animals of a single class may more likely be seen in certain areas thus keeping the larger portion of images on a specific camera hopefully allowing for better accuracy.



Results:

Base Experiment

The following results are the model's accuracy on the test subset across all cameras.

5-fold Validation

Average Accuracy = correctly classified in a class/ total images of class
Average Accuracy overall: 0.8579 | Standard Deviation: 0.0063

Average F1 Score overall = (2 * precision * recall)/(precision + recall)
Average F1 Score overall: 0.8026 | Standard Deviation: 0.0887

Confusion Matrix for the 5th-fold Validation and Number of Images Per class. The confusion matrix shows counts for actual vs predicted labels (bobcat, deer, human, nothing, possum, rabbit, skunk, squirrel, turkey). The number of images per class is: bobcat (575), deer (4272), human (2200), nothing (10416), possum (515), rabbit (1806), skunk (590), squirrel (2397), turkey (643).

Experiment 1:

Results of Individual Cameras:

LowerTrailCamera: Average Accuracy: 0.8884 | Standard Deviation: 0.0096
UpperMostROWCamera: Average Accuracy: 0.7827 | Standard Deviation: 0.0237

Confusion matrices for LowerTrailCamera and UpperMostROWCamera. LowerTrailCamera shows high accuracy across all classes. UpperMostROWCamera shows lower accuracy, particularly for bobcat and squirrel.

NorthernTowerMeadowCamera: Average Accuracy: 0.9064 | Standard Deviation: 0.0050
UpperROWWoodChipFieldCamera: Average Accuracy: 0.8726 | Standard Deviation: 0.0075

Confusion matrices for NorthernTowerMeadowCamera and UpperROWWoodChipFieldCamera. NorthernTowerMeadowCamera shows high accuracy. UpperROWWoodChipFieldCamera shows lower accuracy, particularly for bobcat and squirrel.

UpperTrailCamera: Average Accuracy: 0.6053 | Standard Deviation: 0.0401
SODPlotCamera: Average Accuracy: 0.8562 | Standard Deviation: 0.0088

Confusion matrices for UpperTrailCamera and SODPlotCamera. UpperTrailCamera shows low accuracy across all classes. SODPlotCamera shows high accuracy across all classes.

Experiment 2:

Due to some cameras capturing lower numbers of certain kinds of animals implementing a minimum image count prevents negative influences by these classes on the overall accuracy.

LowerTrailCamera: Average Accuracy: 0.8901 | Standard Deviation: 0.0090
UpperMostROWCamera: Average Accuracy: 0.7953 | Standard Deviation: 0.0167

Confusion matrices for LowerTrailCamera and UpperMostROWCamera. LowerTrailCamera shows high accuracy. UpperMostROWCamera shows lower accuracy, particularly for bobcat and squirrel.

NorthernTowerMeadowCamera: Average Accuracy: 0.9064 | Standard Deviation: 0.0050
UpperROWWoodChipFieldCamera: Average Accuracy: 0.9560 | Standard Deviation: 0.0126

Confusion matrices for NorthernTowerMeadowCamera and UpperROWWoodChipFieldCamera. NorthernTowerMeadowCamera shows high accuracy. UpperROWWoodChipFieldCamera shows high accuracy across all classes.

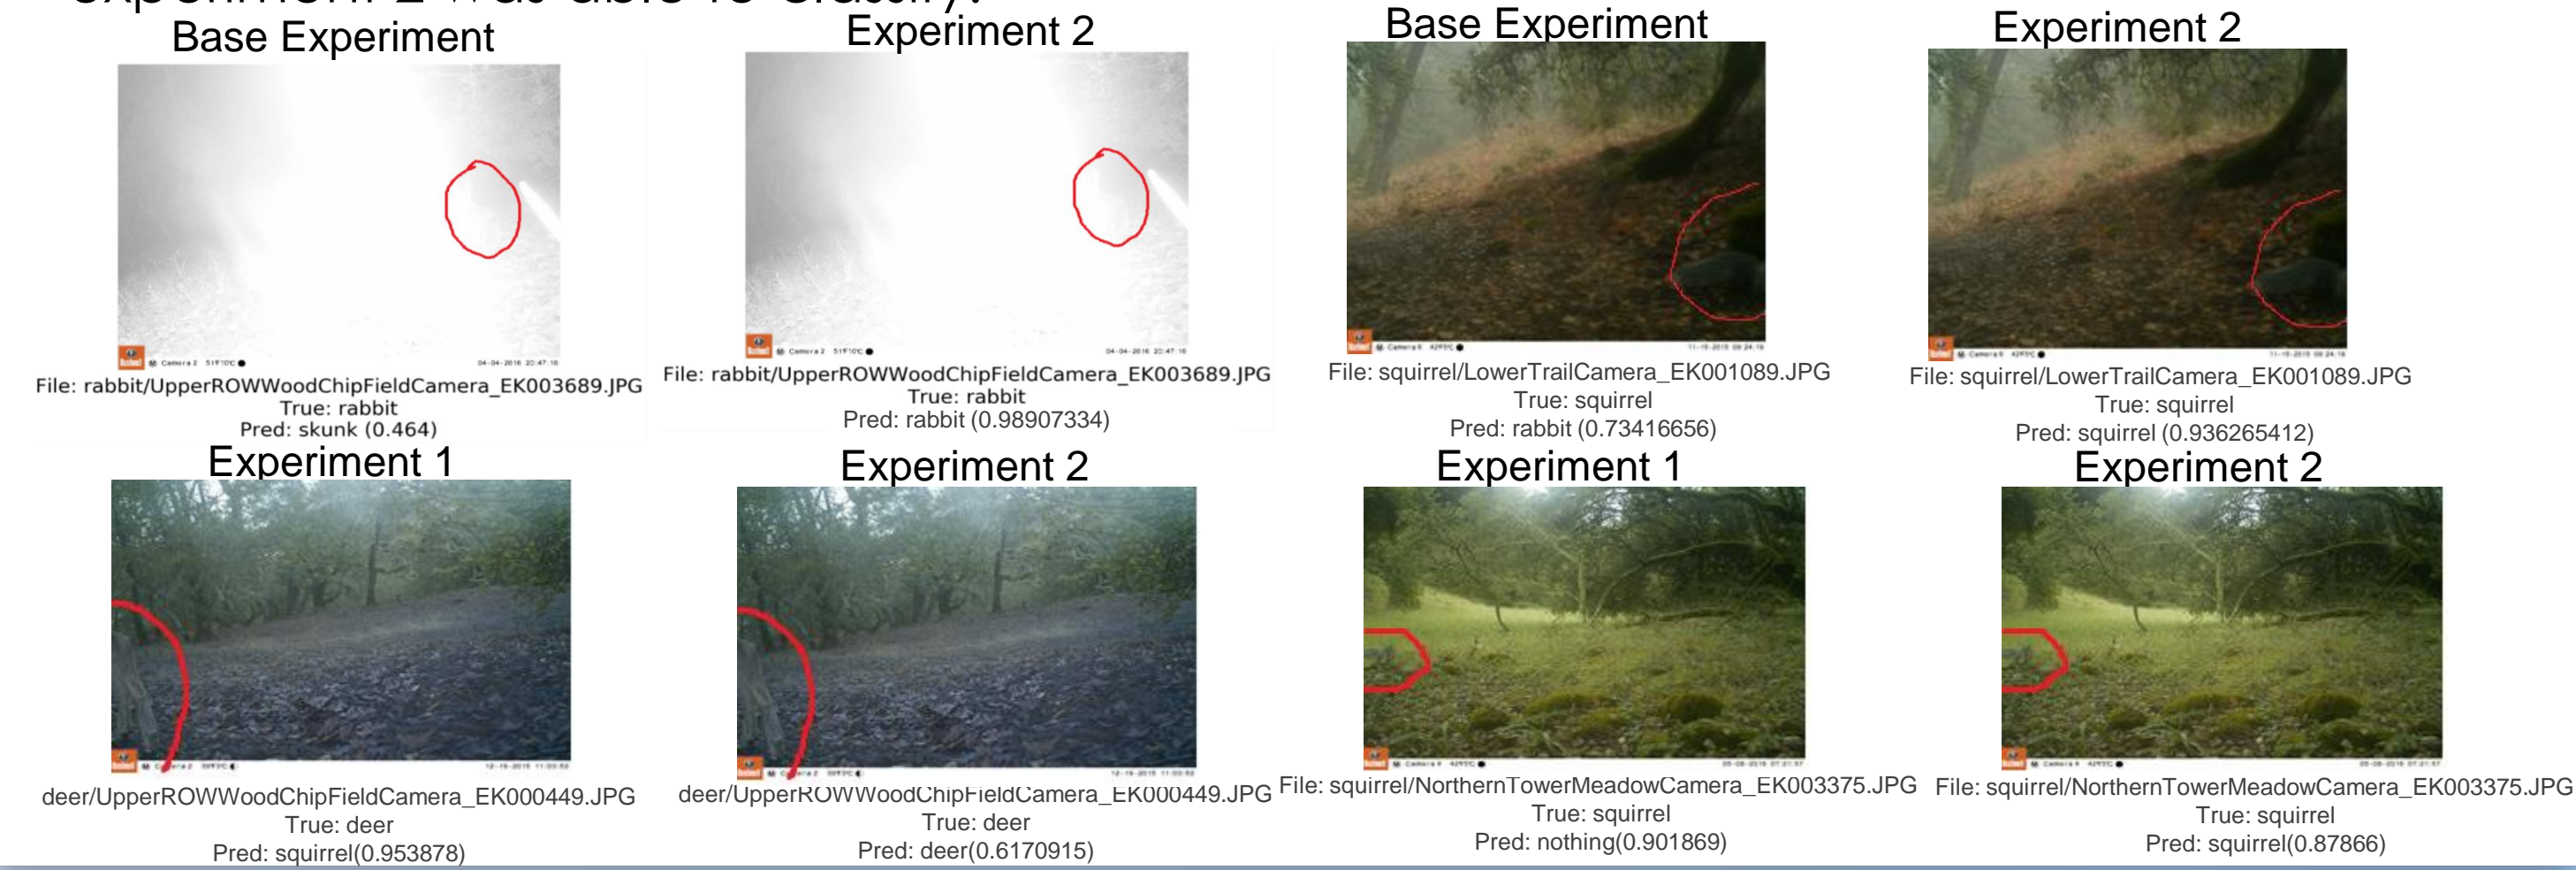
SODPlotCamera: Average Accuracy: 0.8904 | Standard Deviation: 0.0160
UpperTrailCamera No Classes Contain at Least 200 Images

Confusion matrix for SODPlotCamera and a list of classes for UpperTrailCamera that do not contain at least 200 images: bobcat (83), deer (195), human (43), nothing (167), possum (69), rabbit (120), skunk (83).

Results Analysis:

Splitting the database up amongst the individual cameras as performed in experiment 1 showed that certain cameras were better at classifying specific animal groups then others. UpperMostROWCamera(UMRC) is better at classifying bobcats than NorthernTowerMeadowCamera(NTMC), but NTMC is better at classifying turkeys than UMRC. This can be contributed by the differences in image quantities for these cameras. NTMC is also better at classifying skunks than UMRC even though their image quantities for these classes are similar. Meaning that the camera backgrounds are influencing these cameras differently.

Experiment 2 demonstrates an increase in classification accuracy due to the loss of influence by classes containing small image amounts. Individual camera groups were able to classify harder to classify images that the base experiment was not able to do. The image of the rabbit to the left is barely noticeable due to the glare. The base experiment classified the image as skunk while experiment 2 correctly classified the image with high confidence. To the right is a squirrel nearly out of the image yet again experiment 2 was able to classify.



Conclusion and Future Research:

- Each of these cameras receive different inputs so processing them as a group hurts accuracy.
- Camera backgrounds are not similar so data pulled from these will alter confidence levels in other cameras.
- Some cameras record more images of a specific animal class than others so treating the cameras separately can improve classification confidence for these animals.
- Splitting amongst individual cameras does hurt its ability to classify animal groups that have a low number of images.
- Future work would involve testing data augmentation for classes that comprise of small datasets as well as feature extraction from earlier layers in the CNN.

Acknowledgements:

The database used in these test were provided by Dr. Chris Halle. This project has been funded by the Norwick award and a special thanks goes out to the family.