ISyE 6740 - Spring 2021
Project Report

Team Members: Laura Capalleja (902774307), Avani Reddy (902775377), Matt Sheahan (902812479)

Project Title: Image 'DNA Test' for Dogs

Problem Statement

Over the past year of the pandemic, many people have looked towards pets for company. If you happened to be without pet when March 2020 hit, you may have been out of luck. Many dog shelters were quickly emptied by eager new owners, putting both the dog and the owner in a better situation. But for many of the owners, there will always be a question of what breed their new dog is. For owners, this is of great interest because knowing the breed of your dog can help with understanding behavior in order to train the dog properly from youth onward. Owners can also use breed to identify potential health problems and risks and predict lifespan, helping to diagnose illnesses in the dog more easily. You can of course go and get your dog DNA tested, but perhaps there is a simpler way.

Our team is interested in being able to classify the breed of a dog using only an image. This could even help when dogs run away from their owners, and people need to easily identify what the breed of the dog is in order to inform others what type of dog they found. To do this is tricky, as all breeds of *Canis lupus familiaris* share similar features. Also within the same breed, you can have a wide range of sizes, shapes, and colors. All of this is only made more complicated by the natural inconsistencies in which photographs are taken. All of these difficulties taken together means that a successful classification for dog breeds could be easily modified to work for a large variety of other subjects.

An example of the difficulties can be seen in the figures below where the facial features are similar between a Basset Hound and a Blood Hound, but the size and color is more similar between a Basset Hound and a Beagle.

Figure 1: Example images of 3 similar breeds

(a) Basset Hound

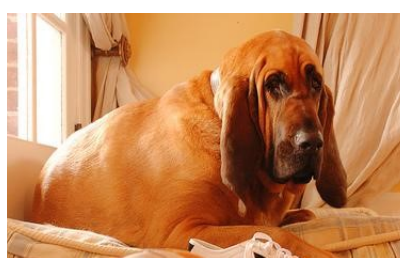

(b) Blood Hound

(c) Beagle

Data Source

The data set that we are using is the Stanford dogs data set which consists of 20,580 images of 120 different dog breeds from around the world[3]. These images come with a label and a bounding box to better identify the object to classify in the image. The data set is not distributed evenly between each breed. This can be seen in the below figure.

Figure 2: Number of Images by Breed

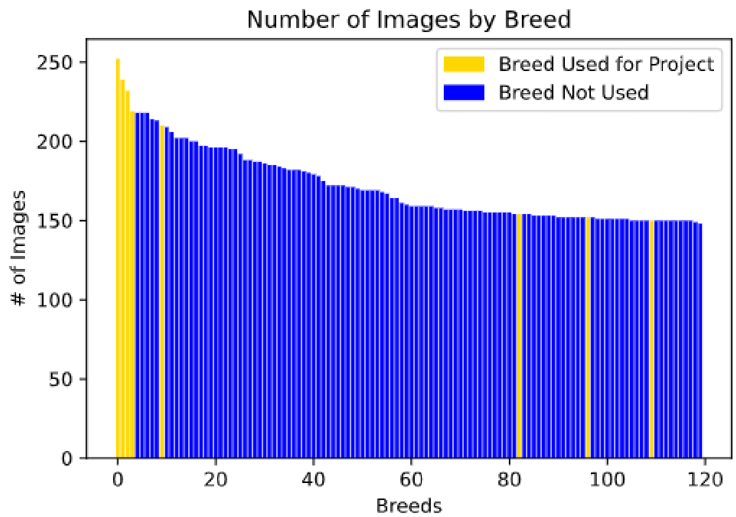

We scoped the data down to 8 breeds to more easily process the data: Afghan Hound, Border Collie, Germand Shephard, Leonberg, Maltese, Pomeranian, Leonberg, Scottish Deerhound. This still left 1608 images to train and test our model against which pushed up on our computational limits.

The next aspect of the data set that needed to be tackled was the various sizes of these images. Just looking at the 1608 images from our reduced data set there were many different sized images. In order to process the data properly all of these images needed to be the same size.

Figure 3: Image Sizes by Breed

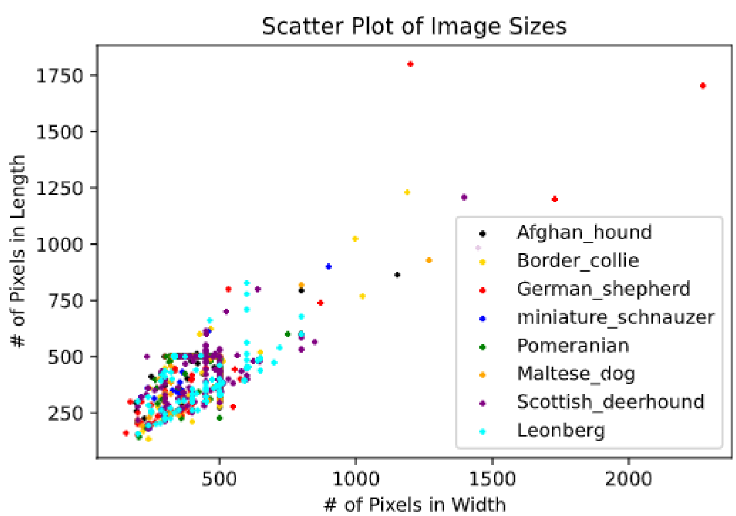

Methodology

Pre-processing: Several alterations to the images were required to prepare the images for classification. The first was cropping the images to the provided bounding boxes. This ensured that the image was zoomed into the dog and excluded background object, humans, etc. as much as possible. Background imagery can impact the accuracy of the classifier, and we did not want our model to identify features outside of the dog.

We then needed all of the images to be equal in size. To accomplish this, we resized the data to a predefined width or height equal to 255 (depending on if the photo is vertical or horizontal) and filled the rest of photo with black pixels in order to maintain pixel aspect ratio.

Additionally, we created mirror copies of each dog image, effectively duplicating our data, in order to increase our image count and improve the classification. We ran our model both with and without the mirrored data and saw significant accuracy improvements with the mirrored images. Moving forward, we refer to the model without our mirrored images as the base model, and the model with the mirrored images as the mirrored model.

To improve processing time, we downsampled the image by a factor of 4. This allowed for a more manageable dataset while still maintain enough clarity in the images. Below, in Figure 4, is an example of one of the images after pre-processing.

Figure 4: Processed image of Maltese dog

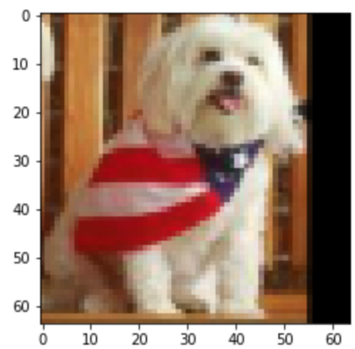

Dimensionality Reduction: As we have learned in class, many algorithms perform better when the number of features in the data is reduced. Reducing dimensionality can improve the training of large models and help focus on the important dimensions of the dataset. To do this, we experimented with ISOMAP, which is the preferred method for non-linear data. We plotted the similarity graphs and analyzed the resulting components. Below we plotted each breed of dog as a different color, and the ISOMAP components did not yield significant groupings for the dog breeds. For this reason, we chose to investigate other methodologies to best classify the dog images.

Figure 5: ISOMAP Component 1 vs. Component 2

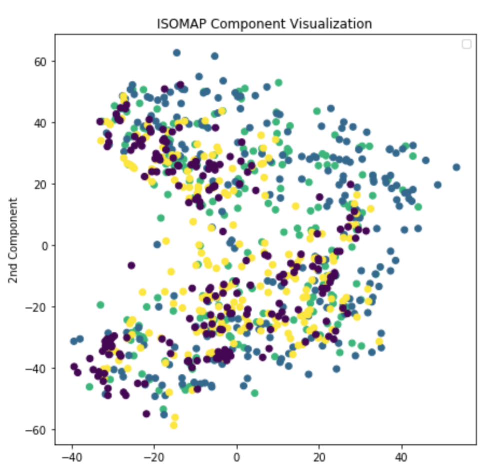

Classification: We investigated a wide range of classification methodologies, and solidified on a neural network. Neural network would likely be the best performing model (best accuracy score), but take the longest to perform. Neural networks have the ability to fit highly complex or non-linear data well, so we believed it would be the best approach for our data. Specifically, we used a CNN (Convolution Neural Network) model from Tensorflow[4]. CNN's are able to classify images well by detecting features independent of their location in the image by using a sliding window approach and scanning across an image, looking at a fixed sized block[2]. By executing multiple layers of this model, we look at very low level pixel features while subsequent layers learn higher level features. Lastly, in order to gauge the probability of the dog photo being a specific breed, we added a layer to our CNN using the softmax distribution - an activation for the last layer of a classification network so that the result can be interpreted as a probability distribution [6].

We trained our model against 67% of the data and tested the model against 33% of the data. We ran epoch=35 iterations of the model and tuned the filter parameter and feature detector parameter.

Evaluation and Final Results

In this report, we will compare the results of two models: with and without the mirrored data. In Figure 6, we observe the accuracy improvement over each iteration of the models. The orange line indicates the measured accuracy on the test data. In the base model, the test accuracy plateaus at around 49%. However, we see a large improvement in the mirrored model which reaches an accuracy of 66%.

Figure 6: Accuracy Over Epoch

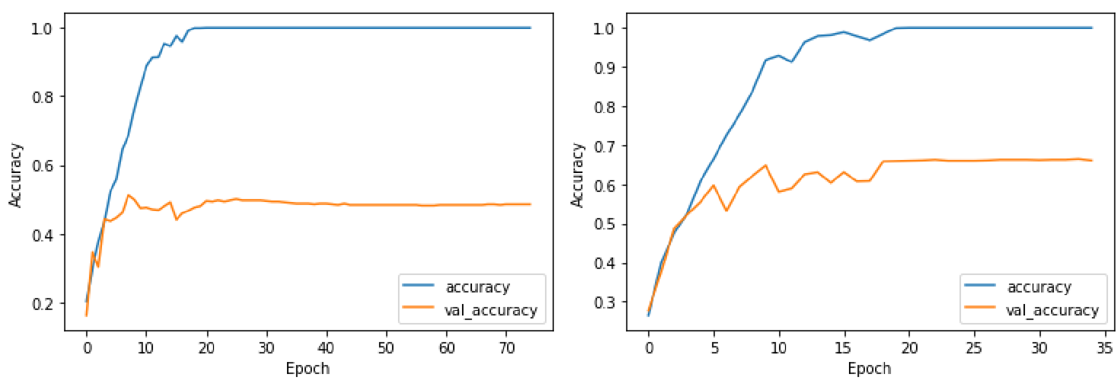

(a) Base Model (b) Mirrored Model

In Figure 7, we compare the resulting Confusion Matrix for the two models. Visually, we observe that a greater percentage of the dog images are classified correctly in the Mirrored Model. In the Base Model, German Shepherds and Scottish Deerhounds are often confused for other breeds, shown by the 15 misclassified German Shepherds in the confusion matrix. Maltese dogs are classified correctly at an even higher rate once we move from the base model to the mirrored model. This could be due to the fact that Maltese dogs are small, white fluffy dogs, differing greatly from the characteristics of other dogs in the mix.

Figure 7: Confusion Matrix

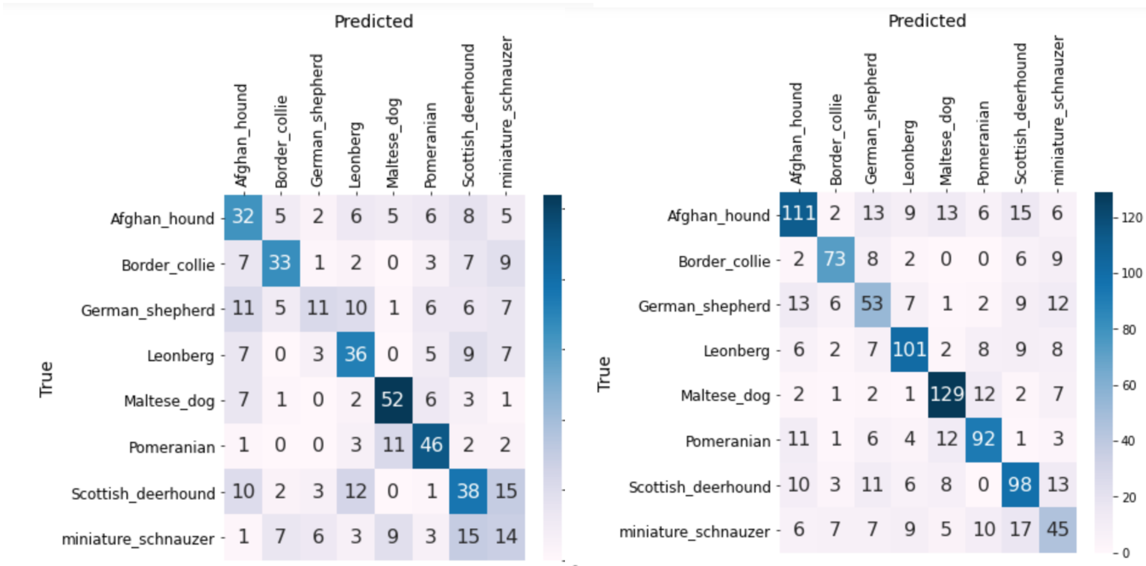

(a) Base Model

(b) Mirrored Model

In addition, we tested the various models on personal dog photos not included in the original dataset[1]. The purpose of using personal dog photos was to remove any bias potentially introduced in the dataset. In the dataset, all the dogs are assigned to a single breed, but in reality many dogs are a mix of various breeds. Since our purpose is to create a product that allows pet owners to identify unknown breeds, we tested on dogs with identified breeds. Max, Pepper, and Mizu are all owned by members of this team's families so we can identify their breeds accordingly. In the future we would like to test on dogs with unidentified breeds in order to provide more information to their owners about their breeds.

In Table 1 of the appendix, we see that all three dog photos were incorrectly classified in the Base Model, but correctly classified in the Mirrored Model. For example, in the base model, Max, a German Shepherd, was misclassified as a Leonberger. Leonbergers have similar attributes to German Shepherds, with similar color fur, dark noses, and similar size at full growth. Once we mirrored the data and retrained the model, we were able to correctly classify Max as a German Shepherd, along with Pepper as a Miniature Schnauzer, and Mizu as a Border Collie (the closest to Australian Shepherd in our mix of data).

Figure 8: Personal Dog Photos Used for Testing

(a) Miniature Schnauzer

(b) German Shepherd

(c) Australian Shepherd

Conclusion

In conclusion, CNN models can do a fairly good job of classifying images into a fairly small subset of classes even when there is not a ton of training data. We do believe, beyond computational limitations, if we had included all 120 breeds that the model would not perform well with the number of training images available. By flipping the images we were able to increase the size of our training images artificially, but this can lead to over-fitting in some situations[5]. An additional data augmentation step would be to try to augment the training data ”in-place” by random rotation. This would allow our model to better generalize while not duplicating images. In the end, our final model was able to correctly identify three dogs from an outside source leaving us satisfied that the model was not over-fit and with 3 happy doggos.

1 Appendix

1.1 References

[1] Jason Brownlee. *How to Make Predictions with Keras.* URL: https://machinelearningmastery.com/how-to-make-classification-and-regression-predictions-for-deep-learning-models-in-keras/. (accessed: 05.01.2021).

[2] Adit Deshpande. *A Beginner's Guide to Understanding Convolutional Neural Networks.* URL: https://adeshpande3.github.io/A-Beginner%5C%27s-Guide-To-Understanding-Convolutional-Neural-Networks/. (accessed: 04.30.2021).

[3] Aditya Khosla et al. "Novel Dataset for Fine-Grained Image Categorization". In: *First Workshop on Fine-Grained Visual Categorization, IEEE Conference on Computer Vision and Pattern Recognition.* Colorado Springs, CO, June 2011.

[4] Kishan Maladkar. *Overview Of Convolutional Neural Network In Image Classification.* URL: https://analyticsindiamag.com/convolutional-neural-network-image-classification-overview/#:~:text=CNNs%5C%20are%5C%20used%5C%20for%5C%20image,because%5C%20of%5C%20its%5C%20high%5C%20accuracy.&text=The%5C%20CNN%5C%20follows%5C%20a%5C%20hierarchical,and%5C%20the%5C%20output%5C%20is%5C%20processed.. (accessed: 04.15.2021).

[5] Adrian Rosebrock. *Keras ImageDataGenerator and Data Augmentation.* URL: https://www.pyimagesearch.com/2019/07/08/keras-imagedatagenerator-and-data-augmentation/. (accessed: 05.03.2021).

[6] Tensorflow. *f.Keras.Activations.Softmax. TensorFlow Core v2.4.1.* URL: www.tensorflow.org/api%5C_docs/python/tf/keras/activations/softmax. (accessed: 05.02.2021).

1.2 Team Name & Work

- Avani Reddy - worked on data manipulation, ISOMAP analysis, CNN algorithm, and paper
- Laura Capalleja - worked on data pre-processing, CNN algorithm, and paper
- Matt Sheahan - worked on data manipulation, CNN algorithm, and paper

Table 1: % of Breed Assignment for Personal Dog Photos

	Goal	**Afghan Hound**	**Border Collie**	**German Shepherd**	**Leonberg**	**Maltese**	**Pomer-anian**	**Scottish Deerhound**	**Miniature Schnauzer**
Base Model									
Max	German Shepherd	0.0000	0.0000	**0.0001**	0.9999	0.0000	0.0000	0.0000	0.0000
Pepper	Mini Schnauzer	0.0000	0.0000	0.0080	0.9902	0.0011	0.0008	0.0000	**0.0000**
Mizu	Border Collie	0.2655	**0.0000**	0.0031	0.7230	0.0000	0.0000	0.0084	0.0000
Mirrored Model									
Max	German Shepherd	0.0000	0.0000	**0.9930**	0.0070	0.0000	0.0000	0.0000	0.0000
Pepper	Mini Schnauzer	0.0000	0.0000	0.0012	0.0000	0.0000	0.0000	0.0000	**0.9988**
Mizu	Border Collie	0.0000	**1.0000**	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000