Using keras to distinguish cats and dogs Mathematical Programming with Python

MATH 2604: Advanced Scientific Computing 4
Spring 2025
Monday/Wednesday/Friday, 1:00-1:50pm

https://people.sc.fsu.edu/~jburkardt/classes/python_2025/cats/cats.pdf

1 What do we want?

A common machine learning task is to "classify" data, that is, to make a judgment about what class it belongs to. An especially difficult classification problem involves images. A special kind of technique for this problem is known as a convolutional neural network, which is able to identify specific features that show up, perhaps in different locations, in related images.

This project involves running a keras program that tries to determine whether a picture is a cat or a dog. This is an unusual project because I think you can simply grab the original data and code and run it. Assuming you can get the code to run, I would want you to describe the problem to the class, discuss what you observed in the program, and say something about neural networks, and the image classification problem.

2 Stepping through the online article

A description. the datasets, and the code are available at

`https://keras.io/examples/vision/image_classification_from_scratch/`

In the following text, I have outlined the steps described in the online discussion. The online discussion seems to contain the code, broken up into pieces. I think the entire code is also available there, separately, as a single file.

The database includes images of cats and dogs. Each image is stored as a `jpeg` file, containing 180×180 pixels and 3 color values (R,G,B). Your first task is to locate, copy, and unpack the database of images. If you

follow the directions in the website, you should be able to create a folder called `PetImages` with subfolders labeled `Cat` and `Dog`.

You are then suggested to filter out some corrupted files, that is, `jpeg` files that do not begin with the string `JFIF` in their header. A code is given for doing this.

The next task is to split each set of images into two groups, one for training (creating the model), and one for validation (testing the model). Again, a short code is supplied.

Now it's possible to take a look at a few of the pictures. This is done with a code in the secton of the article called *Visualize the data*.

In the *Using image data augmentation* section, the author suggests that the data should be slightly modified, to improve the accuracy of the model. One code sets up this process, and the next code allows you to see some examples of the augmented data.

The author suggests that the data should be rescaled, and gives two ways of doing so. The first way works best on a GPU, so I think the second way is best for you, which is described as *Option 2: apply it to the dataset*.

In the section on *Configure the dataset for performance*, the author shows how to set up the computation so that it runs more smoothly, with fewer delays in fetching data.

Now in `Build a model`, the actual neural network processing is set up. This is a big code, and it is followed by a very long flow diagram that tries to suggest what is happening.

In the `Train the model` section, your neural network is given the data, and tries to determine a set of internal weights that will allow it to approximately distinguish the dog and cat pictures.

In the final section on `Run inference on new data`, you demonstrate how your neural network classifies an image it has not seen before.

3 A reference for keras

You can find an introduction to keras in the book:

```
Francois Chollet,
Deep Learning with Python,
Second Edition,
Manning, 2021,
ISBN: 9781617296864
https://www.manning.com/books/deep-learning-with-python-second-edition
```

The cats and dogs exercise is covered in chapter 8, *Introduction to deep learning for computer vision*.