

Assignment 7: Recognizing Objects in Images with Deep Learning Features

Apala Guha

Spring 2017

Readings You are highly recommended to read the Caffe Tutorials at <http://caffe.berkeleyvision.org/tutorial>.

Provided Data In this assignment, we will use 11,540 images from the Pascal VOC 2012 dataset¹. We will work with a partition of the Pascal VOC dataset which includes 11540 images.

1. `images`: A directory containing all 11540 images. Each image is named with a unique ID.
2. `train.txt`, `val.txt`: Text files containing the respective lists of image names.

1 Setting Up

Please follow the steps below to set up *Caffe* on your local machine:

1. Download *Caffe* from <https://github.com/BVLC/caffe>.
2. Follow this: http://caffe.berkeleyvision.org/install_apr.html. I tested it for Ubuntu 14.04.
3. On <http://caffe.berkeleyvision.org/installation.html>, follow the instructions for *CPU-only Caffe*, *Python*, and then *Caffe*.
4. In the *Caffe* directory, run:

```
make PyCaffe
```
5. `export LD_LIBRARY_PATH=path to caffe/build/lib:$LD_LIBRARY_PATH`

2 Image Classification

For this problem, we will be following the tutorial at <https://github.com/BVLC/caffe/blob/master/examples/00-classification.ipynb>. We will be using *CaffeNet* which is pre-trained. We will be applying *CaffeNet* to the validation images. For every validation image, display the label for that image, as determined by *CaffeNet*. Since there are thousands of validation images, you should do this only for a few images (about 10).

¹<http://host.robots.ox.ac.uk/pascal/VOC/voc2012/index.html>

Submission instructions. Submit the code file. Also submit a report which lists the layers used by *CaffeNet* in tabular format. The table should display the type and parameters for each layer in order. There should be a description column associated with each layer which describes its functionality in 1–2 sentences.

3 Model Visualization

Submit code to display any one kernel layer and any one activation layer.