

Topics in Artificial Intelligence (CPSC 532L):

Assignment #2: *CNNs and Image Classification*

Due on Monday, January 22, 2017 at 9am PST

In this assignment you will experiment with VGG-16 one of the most common 16-layer Convolutional Neural Networks (CNNs) with approximately 138 million parameters. As part of the assignment you will learn how to instantiate a CNN network, load pre-trained weights, train your own network for classification, use the network to classify images, and how to utilize pre-trained network to extract generic image features.

This assignment is an adopted version of the Visual Recognition lab from Vicente Ordonez ([link](#)). However the original lab was written in Keras and TensorFlow, where as we are using PyTorch. Never the less the original lab maybe a good resource. You will also undoubtedly endup searching for PyTorch code samples and function use on the web. This is expected and is an important part of the learning experience with DNNs. However, if you do use any of the code from elsewhere you need to acknowledge the original source.

Programming environment

You will be using **PyTorch** for this assignment. If you are using Microsoft Azure credits provided, you will need to provision a Data Science Virtual Machine (VM) to use for the assignment. To do so, please follow the instructions on the Piazza and here ([clickable link](#)):

[Provisioning of the Data Science Virtual Machine for Linux on Microsoft Azure](#)

PyTorch should be pre-installed and available in the Azure Data Science VM. **Make sure you are keeping track of your Azure credits.** Remember that you are being charged as long as VM is running, even if it is idle and not executing any jobs. It maybe best to do the coding and debugging locally (in CPU mode) and run on the VM only when you verified that your code works and you need GPU support. If you opt to do this assignment on your personal computer, it is your responsibility to install PyTorch and ensure it is running properly. A number of tutorials for doing this are available on-line.


Data

For this assignment you will use approximately 100K images from the MS-COCO dataset for training and additional 100 validation images for testing the network. You will use the same data for future assignments. Each image in MS-COCO comes annotated with 5 descriptive sentences and bounding boxes (and segmentations) denoting each instance of one of the 80 object categories. For this assignment we will ignore sentences and object bounding boxes/segments. Instead we will convert annotated objects into an 80-dimensional classification vector which indicates whether a particular object class is present in the image or not. Note that in general, for most images, we will have few non-zero entries in this 80-dimensional vector. You will need to obtain the data as follows:
`wget http://ec2-52-41-153-66.us-west-2.compute.amazonaws.com:8000/data.zip` and upload it to the VM (see instructions in the notebook).

Assignment Instructions

You will need to follow instructions in the corresponding Jupyter notebook and submit the assignment as saved Jupyter notebook (with all the code included) when done. As part of the assignment we are providing the code to pre-process/re-size the images, to be of the required size, as well as the code that will generate 80-dimensional class label vectors from COCO annotations. The steps of the assignment are detailed below.

1. **Loading a Pre-trained Convolutional Neural Network (CNN).** We will work with the VGG-16 image classification CNN network first introduced in: [Very Deep Convolutional Neural Networks for Large-Scale Image Recognition](#) by K. Simonyan and A. Zisserman. The network architecture is illustrated in the Figure below. The provided Jupyter notebook provides the code to instantiate this



model and to pre-load the weights obtained by training the model to recognize 1,000 image categories as part of the [ImageNet](#) challenge.

2. **Making Predictions Using VGG-16 [10 points].** Given a pre-trained network, you will need to write the code to make predictions on the provided 100 validation images, this is equivalent of the forward pass through the network. Since the last layer of the VGG-16 (or most classification networks) is softmax layer, the output can be interpreted as probability of each of the 1,000 ImageNet categories being present in the image. As stated above, typically the final layer of VGG-16 is a softmax layer, however the pre-trained PyTorch model that we are using does not have softmax built into the final layer (instead opting to incorporate it into the loss function), therefore you will need to manually apply softmax to the output of the function. The code for this is provided for convenience (see notebook). For the first 10 validation images display the image itself and print (i) ground truth object classes present in the image and (ii) the top 5 predicted classes with corresponding probabilities (ordered by their probability in descending order).
3. **Computing Generic Visual Features using CNN [10 points].** It has been shown in a number of prior papers that output of fully connected layers in a CNN trained for image classification (ImageNet) can be used as effective features for other tasks. We will validate this here. First, write the code to remove the softmax prediction layer from the instantiated VGG-16 network with pre-trained weights. You will then be left with fc2 layer as the final layer of the network. Now, compute the features for all the 100K training and 100 validation images, saving them into the two corresponding files. Note that for every image i you should end up with $\mathbf{x}_i \in \mathbb{R}^{4,096}$ dimensional feature vector. This amounts to a forward pass using all of the data and may take some time.
4. **Visual Similarity [10 points].** One way to test if a given feature representation is “good” is to see if similar images end up being close in the feature space. Given the pre-computed features, for 10 images in the validation set compute the closest image in the training set (among 100K images) using Euclidian distance as the distance metric. In other words: $\arg \min_i \|\mathbf{x}_i - \hat{\mathbf{x}}\|^2$, where $\hat{\mathbf{x}}$ is a feature of the validation image and i is over the training images. Display the query and the nearest neighbor retrieved images.

5. **Training a Multi-label Classification Network [30 points]**. ImageNet classification task inherently assumes that only one image category is predominant in an image (single label multi-class), where as MS COCO images typically contain multiple objects. We will see if we can train a classification network that is better suited for predicting the object classes in this case. We will do so by building a simple 2-layer classification network which takes 4,096-dimensional feature vectors from above and outputs probability of the 80 categories present in MS COCO. The network will have two fully connected layers with sigmoid activations (same as in Assignment 1). The hidden dimension should be set to 512. For the loss, we will need to use multi-label loss (`MultiLabelSoftMarginLoss` in PyTorch). Once the classification network is trained, redo the prediction experiment in Part 2, *i.e.*, for the first 10 validation images display them and print (i) ground truth object classes present in the image and (ii) the top 5 predicted classes with corresponding probabilities as given by the trained classification network (in descending order). Briefly comment on the quality of the results.
6. **End-to-End Model Fine-tuning [30 points]**. Above we extracted features and then trained a classification model on top of the extracted features. This is obviously sub-optimal, so here you will need to integrate the feature extractor (VGG-16), and the classification model that predicts 80 categories. Instead of training the entire architecture from scratch we will train it starting from the ImageNet weights for the VGG-16 portion of the network. Again, once the model is trained on 100K images, redo the prediction experiment in Part 2 of the assignment, *i.e.*, for each one of the first 10 validation images display it and print (i) ground truth object classes present in the image and (ii) the top 5 predicted classes with corresponding probabilities as given by the trained classification network (in descending order). Briefly comment on the quality of the results.
7. **Hyper-parameter Tuning [10 points]**. The hyper-parameters (`learning rate`, `batch size`) are likely not the best and finding good hyper-parameters is an important step in training of neural networks. Do a 3x3 grid search over the hyper-parameters, training the model with 3 different learning rates (*e.g.*, $\frac{1}{10}$ and $10\times$ the original) and 3 different batch sizes (*e.g.*, $\frac{1}{2}$ and $2\times$ the original). This will result in running end-to-end learning 9 times. Each time plot the value of the loss on both training and validation sets as a function of the epochs and record the value of the training loss at the end of the training. Since you are learning on the same data the training loss values should be comparable across hyper-parameter choices and can be used to choose the best model (actually, in real experiments, you would choose the model based on validation loss, not the training loss). Briefly comment on the process and which hyper-parameters you found to be optimal.