EE 519: Deep Learning Theory & Fundamentals **Spring 2023**

Homework 2
Due: May 12, 2023, 11:59PM PT

Student Name: *Instructor Name: John Lipor*

Problem 1 `conv.py` (8 pts)

For the first problem of this assignment, you will complete a convolutional layer for your MyTorch package. Your task is to complete `conv.py` as detailed by the documentation in the corresponding file and the lecture notes. In particular, you will implement the forward and backward passes of a two-dimensional convolutional layer. For simplicity, we will ignore the bias term in this layer.

Turn in the output of the notebook `Conv2DTester.ipynb`.

Problem 2 CNNs on CIFAR-10 (3 pts each)

The textbook focuses heavily on the Fashion-MNIST dataset since it is easy to quickly train a somewhat-accurate network on this dataset. However, this dataset may be too easy to develop an understanding of what works well and what doesn't in a CNN architecture. In this problem, we'll consider the more difficult (though still manageable) CIFAR-10 dataset, which is of similar size to Fashion-MNIST but has three input channels and is considered a much more challenging dataset. This dataset can be accessed through `torchvision.datasets` in a manner similar to Fashion-MNIST.

Your task will be to train two CNNs on this dataset—one that utilizes a linear output layer and a second that uses global average pooling. You may make use of any available literature or blogs to gain insight into effective network architectures for this dataset. Your goal is to achieve the best accuracy you can while allowing time for the rest of the problems, which will utilize your networks. You **must** try both dropout and batch normalization for each architecture and report their impact on performance.

(a) Implement a CNN in the vein of LeNet, AlexNet, or VGG that uses one or more convolutional layers, followed by one or more dense layers. Turn in:

  (a) a description of your final network architecture
  (b) how you arrived at this architecture
  (c) the impact of batch normalization
  (d) the impact of dropout
  (e) a plot of training and validation loss
  (f) your final validation accuracy.

(b) Implement a CNN in the vein of NiN, ResNet, etc. that integrates nonlinearities within the convolutional portions and utilizes global average pooling. Turn in the same items as for part (a).

Problem 3 Feature Maps and Filters (3 pts each)

As discussed in class, interpreting CNNs is largely performed by visualizing their different elements. In this problem, your task is to visualize both the intermediate feature maps (outputs of convolutional layers) as well as some of the learned filters for each of the two networks trained in Problem 2. For each network, turn in:

(a) images of subplots of feature maps for multiple convolutional layers. Each image should contain a number of subplots corresponding to the number of output channels for a given convolutional layer, and you should create one image per convolutional layer for no more than three layers.

(b) images of subplots of the learned filters for multiple convolutional layers, following the above specifications. Note whether your network ever learns filters similar to the well-known Gabor filters.

Problem 4 Class Activation Maps (5 pts)

In this problem, you will integrate class activation maps (CAMs) as defined in the paper Learning Deep Features for Discriminative Localization. You may use any code you find online, including the authors' own GitHub repository. The article here may also be helpful.

Turn in some sample CAM images from each architecture developed in Problem 2. Be sure to display images that are correctly classified as well as incorrectly classified, as well as a brief interpretation of what the CAMs show.

Problem 5 Guided Backpropagation (5 pts)

In this problem, you will produce saliency maps as defined in the paper Striving for Simplicity: The All Convolutional Net. You may use any code you find online. The article here may also be helpful.

Turn in some sample saliency map images from each architecture developed in Problem 2. Be sure to display images that are correctly classified as well as incorrectly classified, as well as a brief interpretation of what the saliency maps show.

Problem 6 Reflection on Datasets (5 pts)

Read Ch. 8 of the textbook *Patterns, Predictions, and Actions* available here, which discusses datasets and their impact on machine learning. In the `interesting-reading` channel on Slack, state (1) one thing you learned from a technical perspective regarding datasets (i.e., from the material up to the section *Harms associated with data*), (2) one thing you learned about the harms associated with data and how we might overcome these harms, and (3) one question you have after reading the chapter.