

CMSC 478 Machine Learning - Spring 2018

Homework Assignment 3

Due at the start of class on Thursday, April 12th

1. Exercises 3.10, 3.11, and 4.2 from the online version of Sutton and Barto.
The complete book can be found here:

<http://www.incompleteideas.net/book/ebook/the-book.html>

The first two exercises are here:

<http://www.incompleteideas.net/book/ebook/node34.html>

and the third one is here:

<http://www.incompleteideas.net/book/ebook/node41.html>

2. In this exercise you'll experiment with CNNs for image classification using the Torch framework in Python. **It is important that you start on this early because you'll be installing code and data, and if you run into problems you may not be able to get them resolved at the last minute.**

Read through the PyTorch 60 Minute Blitz here:

http://pytorch.org/tutorials/beginner/deep_learning_60min_blitz.html

You can skip the optional section on data parallelism.

At the end of the "Training a Classifier" section you can download `ci-far10_tutorial.py`. Do that. To run it you'll need to install the `torch` and `torchvision` packages for python. The easiest way to do that is with `pip`:

- `pip install torch`
- `pip install torchvision`

Instructions on how to install `pip` can be found here:

<https://pip.pypa.io/en/stable/installing/>

When you run the CIFAR10 tutorial code the first time, it will download the CIFAR10 dataset, train a simple CNN, and report classification accuracy.

Your task is to try to improve that accuracy and document the results (include everything you try regardless of whether it worked or not). Things to consider are:

- Change filter sizes and/or numbers in the convolutional layers
- Add a convolutional layer
- Change the size(s) of the fully connected layer(s)
- Try an activation function other than ReLU
- Use different learning rates or momentum
- Train for more epochs
- Look around on the web for advice on other things to try

Turn in the output of the CIFAR10 code from your initial run before making any changes. Then write a short document that lists the things you did and the resulting test accuracy, with output from the CIFAR10 demo to document that accuracy. Explain why you think each step helped or hurt accuracy. Also, turn in the code that produced the best result you were able to get along with the output.