

Project 5

Directions

- Expected time commitment: 2-4 hrs
- Visit this website to begin the tutorial:
 - https://pytorch.org/tutorials/beginner/deep_learning_60min_blitz.html
- Submit: An ipynb file that summarizes your findings
- Due: Friday, Nov. 30, 11:59PM

PyTorch Tutorial (3 points)

One reason for the recent renewed interest in neural networks is the creation of libraries, like PyTorch and TensorFlow, that can automatically compute complicated gradients. Since PyTorch is a hugely popular library, and easier to learn than TensorFlow, it makes sense to just dip our toes in the ocean of deep learning by learning PyTorch.

To use PyTorch, the **torch** and **torchvision** packages are required. Usually these are straightforward to install using **conda** or **pip**, but if you're on **stdlinux**, these python package management programs are not installed by default. So we will use [Google Colab](#) to install and use these packages.

Once you open a new notebook, start by running this command to install the packages:

```
!pip3 install torch torchvision
```

Then you should be good to go and complete the tutorial!

Experimentation (4 points)

Next, create a new network with some of the parameters changed. Do some experimentation to try and get better accuracy (on the image classification portion)! You should try at least 3 changes to the network architecture or other hyperparameters, keeping track of the results as you go.

Analysis (3 points)

Make an educated guess about why various changes had the effect they did on your results. Write these up in a text node (as part of the same ipynb file), so they will be recorded in your submission.