

## Homework 4

---

Using the assigned reading listed on the course page, answer the questions below with a short response. Note that we are looking for concise statements that show understanding, not quantity. The total discussion should be less than a page.

### Getting familiar with Deep Learning

1. Work through [Pytorch 60-minute Blitz](#). You can use Google Colab to run the code so that you don't have to install anything on your machine (Colab links are at the top of each page of the Pytorch tutorial). The last part of the tutorial involves training a net on the CIFAR10 dataset to classify images. Running the code they provide should give around 50% accuracy. Your task for this question is to make a change to the code that results in higher accuracy. Explain your change, and explain what it is doing and why it helps. Provide a plot comparing the accuracy before and after the change (a line plot with epoch on the x-axis and accuracy on the y-axis). Try to see if you can make any relevant connection to the things we have learned in class already (e.g., double descent).

(**Note:** To enable GPU/TPU acceleration on Colab  
Runtime → Change runtime type → Hardware accelerator.)

### Interpreting neural networks

1. The work of Bau et al. tests the causal role of neurons in two settings: within classifiers, and within generative models. How is the causal role of a neuron tested? What different kinds of conclusions can be drawn when doing these tests in the two different types of networks?

**Submission:** Upload a PDF of your response through Canvas by **10/12 at 1pm**. For Question 1, please submit a PDF of the Google Colab or Jupyter Notebook along with the answers.