CS 7150: Deep Learning — Spring 2024— Paul Hand

HW 1
Due: Friday October 4, 2024 at 9:00 PM Eastern time via Gradescope

Names: [Put Your Names Here]

You will submit this homework in groups of up to 3. You may consult any and all resources. Note that some of these questions are somewhat vague by design. Part of your task is to make reasonable decisions in interpreting the questions. Your responses should convey understanding, be written with an appropriate amount of precision, and be succinct. Where possible, you should make precise statements. For questions that require coding, you may either type your results with figures into this tex file, or you may append a pdf of output of a Jupyter notebook that is organized similarly. You may use code available on the internet as a starting point.

Question 1. *Is a* 1×1 *convolution operation the same as scaling the input by a single scalar constant? Explain. If the answer is sometimes yes, then make sure to explain when it is and when it isn't.*

Response:

Question 2. *In this problem, you will train a classifier on the MNIST dataset. You can find this dataset in TorchVision. Train a fully-connected neural network with 2 hidden layers and ReLU activations.*

(a) Clearly convey the architecture (by providing a figure) and training procedure you used.

Response:

(b) Plot training loss, training accuracy, and test accuracy vs. iteration count.

Response:

Question 3. *Download an AlexNet and a ResNet101 model that have been trained on ImageNet. You can find these models in TorchVision. Using a camera, take a picture of an object that belongs to one of the ImageNet classes. Use both models to classify the image. Output the top 5 predicted classes and their corresponding probabilities, according to each model. Were the models correct?*

Response: