CS 7150: Deep Learning — Fall 2024— Paul Hand

Project Planning
Due: TBD via Gradescope

Names: [Put Your Names Here]

For your final project, you will select a paper in which a neural network is trained to perform a particular task. Your task is to reproduce a result from that paper without consulting the code of that paper. You do not need to reproduce the whole paper. You will need to code the network 'from scratch,' as if you were the authors of this paper. You may use standard Deep Learning frameworks (e.g. PyTorch, TensorFlow, etc.). You may use code that is available on the internet from other authors as building blocks. You may run your algorithm in a slightly different context (e.g. CIFAR-10 instead of ImageNet).

You will need to write up a short (less than 4 page) report detailing: the paper you are replicated, the scientific context of the paper, the result from the paper that you are attempting to replicate, the details of your implementation, the results of your implementation. You are not required to get equal performance as the paper, but you should demonstrate that the network you trained was in some sense successful. Please use the NeurIPS Style files for your report.

You may work in groups of up to 3 people. You may work alone.

Every group will give a short (less than 3 minute) slide presentation to the entire class detailing the same items as the writeup. In your presentation, you should also describe the aspect of your project that was the most difficult to do (e.g. manipulating data, tuning hyperparameters, etc). After each presentation there will be a 1 minute period for questions.

If you want some ideas of projects (in no particular order), here are some ideas. You do not need to select one of these papers.

- Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks
- Deep Image Prior
- Deep Decoder: Concise Image Representations from Untrained Non-convolutional Networks
- Learning From Noisy Singly-labeled Data
- Understanding Deep Learning Requires Rethinking Generalization
- Deep Double Descent: Where Bigger Models and More Data Hurt
- Virtual Adversarial Training: A Regularization Method for Supervised and Semi-Supervised Learning

- NICE: Non-linear Independent Components Estimation
- Visualizing and Understanding Convolutional Networks
- How Does Batch Normalization Help Optimization?
- Learning without Forgetting
- Learning to Invert: Signal Recovery via Deep Convolutional Networks
- Explaining and Harnessing Adversarial Examples

Question 1. *Project Planning*

1. What paper are you planning to replicate part of?

   Response:

   The link to an arxiv copy of the paper is: [Put Paper Title Here]
   Authors: [Put Authors Here]

2. What figure or result of the paper are you planning on replicating?

   Response:

   Include a figure or table or other result here.

3. What aspect of the project do you anticipate being the most difficult and/or time consuming?

   Response: