

## CS 7150: Deep Learning — Spring 2021 — Paul Hand

Day 10 — Preparation Questions For Class

Due: Wednesday 2/24/2021 at 2:30pm via [Gradescope](#)

Names: [Put The Names Of Your Group Here]

You may consult any and all resources in answering the questions. Your goal is to have answers that are ready to be shared with the class (or on a hypothetical job interview) as written. Your answers should be as concise as possible. When asked to explain a figure, your response should have the following structure: provide context (state what experiment was being run / state what problem is being solved), state what has been plotted, remark on what we observe from the plots, and interpret the results.

Submit one document for your group and tag all group members. We recommend you use Overleaf for joint editing of this TeX document.

**Directions:** Read '[ImageNet Classification with Deep Convolutional Neural Networks](#)' (AlexNet).

- Read the entire paper

**Question 1.** *Explain all of the terms in the optimization algorithm presented in Section 5.*

**Response:**

**Question 2.** *The AlexNet paper used a learning rate schedule where the learning rate was lowered when validation error stopped improving. Why is it reasonable to have a schedule where learning rate decreases? Why wait until validation error stops improving (as opposed to imposing a specific schedule based on epoch number)?*

**Response:**

**Question 3.** *Explain Figure 1.*

**Response:**

Context:

What is plotted:

What we observe:

Interpretation:

**Question 4.** *What is dropout? What evidence is there that it works? Why does it work?*

**Response:**

**Question 5.** *Explain the different data augmentation strategies used in the AlexNet paper. What do each of these strategies accomplish?*

**Response:**

**Directions:** Read 'Deep Residual Learning for Image Recognition' (ResNets).

- Read Section 1, 3, 4.1

**Question 6.** *The ResNet paper reports 3.57% error on the ILSVRC. Some people would claim this performance is superhuman. Look up the rate of error achieved by humans. Why is the human error rate not 0%? (After all, wasn't it labelled by humans?) Do you think it is fair to say that this net can achieve superhuman performance at image classification?*

**Response:**

**Question 7.** *Explain the right column of Figure 3. Include the meaning of the text in each of the boxes, what the solid arrows mean, what the dashed arrows mean, what "pool, /2" and "avg pool" mean.*

**Response:**

**Question 8.** *Estimate the number of weight parameters in the three nets depicted in Figure 3 of the ResNet paper. Identify where most of the parameters are in each of the three nets.*

**Response:**

**Question 9.** *Explain Figure 4 of the ResNet paper. Make sure to explain why there are two sudden steep drops in error % in both plots.*

**Response:**

Context:

What is plotted:

What we observe:

Interpretation: