CS 7150: Deep Learning — Summer-Full 2020 — Paul Hand

Week 2 — Preparation Questions For Class
Due: Monday April 11, 2020 at 12:00 PM Eastern time via Gradescope

Name: [Put Your Name Here]
Collaborators: [Put Your Collaborators 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.

Directions: Read the articles 'Deep Learning' (Three Giants Paper) and 'Visualizing and Understanding Convolutional Networks'. (ZFNet)

Questions on Three Giants:

Question 1. *Briefly explain the selectivity-invariance dilemma.*

Response:

Question 2. *Provide a sentence (or two) from the Three Giants Paper that you find interesting and would like to discuss in class.*

Response:

Questions on ZFNet:

Question 3. *Summarize the contributions of the paper. Be as succinct yet comprehensive as possible.*

Response:

Question 4. *Precisely explain what is being plotted in Figure 2. Make sure to explain why the organization/presentation of layer 1 is different than layers 2 and higher. You do not need to summarize the visualization algorithm in your explanation.*

Response: Your answer should be begin with something similar to: A ... net was trained to solve [some task].

Question 5. *In Figure 2, why even bother showing the visualizations instead of the patches? Provide two examples where the visualization is needed (and the image patches are insufficient) to understand the qualitative interpretation of a particular feature map.*

Response:

Question 6. *What do the authors mean by "invariances" of a feature map? Why does the method presented in this paper better show the invariances of the feature maps in comparison to gradient- or Hessian-based analysis of the provided neuron?*

Response:

Question 7. *The authors saw that their classifier's performance was not invariant to image rotation except for the category 'entertainment center.' Explain what is going on with this category. What could you do to improve the rotational invariance of the classifier?*

Response:

Question 8. *How useful/general do you find the process the authors used for improving the architecture of a network by visualizing the activations of its inner layers? What are the challenges of this as a general process? What alternative procedures can you envision for improving neural network architectures?*

Response:

Question 9. *Identify which parameters in Figure 3 are user specified, and which are derived from those user-specified parameters. Verify that the derived parameters are correct.*

Response:

Question 10. *The authors make the claim that their result brings into question the utility of benchmarks with small (i.e. $< 10^4$) training sets. Why do they say this? Do you agree? Are small datasets of no value anymore?*

Response: