

CS 7150: Deep Learning — Fall 2024 — Paul Hand

Day 5 — Preparation Questions For Class

Due: Tuesday 9/24/2024 at 12:30pm Eastern Time 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 ‘[Visualizing Higher-Layer Features of a Deep Network](#)’.

- Read only Section 3.

Question 1. *Explain the visualization method the authors propose. What are its drawbacks?*

Response:

Directions: Read ‘[Visualizing and Understanding Convolutional Networks](#)’ (ZFNet).

- Read the whole paper. You can skip Sections 4.3 and 5.3.

Question 2. *Summarize the contributions of ZFNet paper. Be as succinct yet comprehensive as possible.*

Response:

Question 3. *Explain Figure 2. In your explanation, you do not need to summarize the visualization algorithm. Make sure that your response explains why the structure of the plot for Layer 1 is different than the structure of the plot for Layers 2-5.*

Response:

Context: Your answer should begin with something like “A ... net was trained to perform ... ”

What is plotted:

What we observe:

Interpretation:

Question 4. *In Figure 2, why even bother showing the visualizations instead of the raw image 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 5. *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 6. *Explain Figure 4.*

Response:

Context: Your answer should begin with something like “A ... net was trained to perform ... ”

What is plotted:

What we observe:

Interpretation:

Question 7. *Explain Figure 5c. In the interpretation, make sure to explain what is going on with the category ‘entertainment center.’ What could you do to improve the rotational invariance of the classifier?*

Response:

Context: Your answer should begin with something like “A ... net was trained to perform ... ”

What is plotted:

What we observe:

Interpretation:

How could you improve rotational invariance of classifier?

Question 8. *What is the process by which the authors improved a neural network architecture using their visualization technique? What are the challenges of this as a general process? Can you envision alternative processes 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. Justify the value of all the derived parameters.*

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: