

CS 7150: Deep Learning — Fall 2024 — Paul Hand

Day 7 — Preparation Questions For Class

Due: Tuesday 10/1/2021 at 12:30 pm 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 ‘[Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift](#)’ by Ioffe and Szegedy.

- Read entire paper.

Question 1. *In the context of CNNs, what is Batch Normalization? At test time, there may be only a single image passed into the network, and the variance of a quantity over only a single datapoint is undefined. How is this issue dealt with?*

Response:

Question 2. *Explain Figure 1.*

Response:

Context:

What is plotted:

What we observe:

Interpretation:

Question 3. *Explain Figure 2.*

Response:

Context:

What is plotted:

What we observe:

Interpretation:

Directions: Read ‘[How Does Batch Normalization Help Optimization?](#)’ by Santurkar et al.

- Read entire paper, except Sections 4-5.

Question 4. *Explain Figure 1.*

Response:

Context:

What is plotted:

What we observe:

Interpretation:

Question 5. *Explain Figure 2.*

Response:

Context:

What is plotted:

What we observe:

Interpretation:

Question 6. *Explain Figure 3.*

Response:

Context:

What is plotted:

What we observe:

Interpretation:

Question 7. *Explain Figure 4.*

Response:

Context:

What is plotted:

What we observe:

Interpretation:

Question 8. *Having read both papers how would you respond to the question “Why does BatchNorm work”?*

Response:

Question 9. *What is the point of having an explanation for why BatchNorm works?*

Response: