

CS294-082 Fall2021 Project Requirements

- Form teams of 3 people.
- Pick a past, current, or future machine learning project, a famous machine learning paper, your machine-learning-based thesis, a machine-learning project for another class.
- Send a 1 paragraph proposal to Eleanor by October 13th, 2021 11:59am (hard deadline) stating the participants, their roles and the project idea.
- Write an article (3-10 pages) discussing the experimental design of the machine learning approach chosen. Deadline is the last day of finals week.
- The article must include responses to all questions below as applicable to your project. If the article does not respond to the questions and is not based on measurements taught in the class, the project will fail.
- Articles must clearly indicate the contribution of each co-author.

9-ish Questions to Verify when Designing Machine Learning Experiments

- 1) What is the variable the machine learner is supposed to predict? How accurate is the labeling? What is the annotator agreement (measured)?
- 2) What is the required accuracy metric for success? How much data do we have to train the prediction of the variable? Are the classes balanced? How many modalities could be exploited in the data? Is there temporal information? How much noise are we expecting? Do we expect bias?
- 3) What is the Memory Equivalent Capacity for the data (as a dictionary). What is the expected Memory Equivalent Capacity for a neural network?
- 4) What is the expected generalization in bits/bit and as a consequence the average resilience in dB? Is that resilience enough for the task? How bad can adversarial examples be? Do we expect data drift?
- 5) Is there enough data? How does the capacity progression look like?
- 6) Train your machine learner for accuracy at memory equivalent capacity. Can you reach near 100% memorization? If not, why (diagnose)?

- 7) Train your machine learner for generalization: Plot the accuracy/capacity curve. What is the expected accuracy and generalization ratio at the point you decided to stop? Do you need to try a different machine learner? How well did your generalization prediction hold on the independent test data? Explain results. How confident are you in the results?
- 8) Comment on any other quality assurance measures possible to take/ the authors should have taken. Are there application-specific ones? If time is present: How did you deal with it?
- 9) How does your experimental design ensure repeatability and reproducibility?