

Premise

The final project is our best way to evaluate what you've learned throughout the course. The goal is to evaluate you exactly how deep learning practitioners are evaluated in industry. You will collaborate in groups of 2-4 students and come up with a real data-driven problem that you'd like to model. You'll then:

1. Pull and explore the data.
2. Decide on a data hypothesis (e.g. I believe there exists a good embedding of natural language questions; I believe noses are a salient feature in human faces but not in Pokemon; I believe that earnings call audio is a signal for company performance).
3. Determine your architecture priors.
4. Run at least one non-deep learning benchmark (e.g. random forest, logistic regression, ARIMA, etc.).
5. Run at least one base deep model (feedforward net, CNN, RNN).
6. Run at least one advanced deep model (augmenting with autoencoders, ResNets, GANs, etc).
7. Visualize your findings.
8. Write up your results in a technical report and in a blog post aimed at a non-technical audience.

Part 1: Proposal (5 pts), due Friday, March 22nd

Turn in a document with the following:

1. (0 pts) All team members' names (3-4 students per group)
2. (3 pts) A description of the data set (where it's from or how it was made) as well as a sample of the data (a 5-row header of the data or a visualization of 5 examples should suffice).
3. (2 pts) A description of the data hypothesis and your inductive biases / priors.

Note that it is your responsibility (and part of your evaluation) to choose a dataset that you think will be well-addressed by a deep learning solution. Your team should schedule a brief conference with a member of the teaching team to get their proposal approved before submission – the hope is that all projects will have a high a priori likelihood of an interesting deep learning analysis.

Part 2: Project Implementation (35 pts), due Monday, April 29th

Turn in a zip file with the following:

1. All code associated with the project
 - (1 pts) Code readability
 - (5 pts) Correct implementation of the non-deep benchmark.
 - (5 pts) Correct implementation of the base deep learning model.
 - (10 pts) Thoughtful selection and correct implementation of the advanced model(s).
 - (extra credit, 10 pts) Using a novel architecture with better results than the baselines.
2. The technical report (10 pts)
 - (5 pts) Well-designed visualizations of the results.
 - (5 pts) Analysis and evaluation; conclusions.
 - (extra credit, 5 pts) Interpretation of the model (e.g. using DISTILL).
3. The blog post (4 pts)
 - (4 pts) General readability
 - (extra credit, 5 pts) Pizzaz

The team will present their project in class on either Monday, 4/29 or Wednesday, 5/1.