

CS 333 Final Project

1 Project Overview

For your final project, you will build a system to solve an NLP task of your choice. The project should involve a substantial programming component, equivalent to one CS 333 assignment.

You will have access to cloud computing resources from the NCSA Delta cluster.

If you are working on a thesis or research project, I am also willing to consider a topic that is related. But you **cannot** propose a topic that overlaps with a class project in another course.

Component	Points	Due Date
Project proposal	(part of HW 7)	11/17
Lit review	(part of HW 8)	11/24
Presentation	15 points	12/9
Code	30 points	12/18 4pm
Paper	55 points	12/18 4pm

All project components must be submitted by 4pm on December 18th. This is a college-wide deadline. I cannot accept any work after this deadline.

2 Project checkpoints

2.1 Project Proposal

As part of HW 7, you will write a brief description and timeline for your project.

2.2 Literature Review

As part of HW 8, you will read at least 3 papers related to your topic. You should write a paragraph summarizing each paper. The papers that you read should be cited in your final report. You are also welcome to read more.

3 Deliverables

3.1 Codebase

You will submit your code at the end of the semester. Your code should be organized and commented. **You must also submit a README text file that explains how to run your code.**

3.2 Presentation

We will have short presentations on the final day of class. You will have **3 minutes** to briefly present your project. You are not required to have results to share, but if you do have preliminary results, you can discuss them.

You should design 1 slide to use in your presentation.

3.3 Report

You will write a short paper about your project. The report should be **single-spaced and at least 6 pages**. There is no page limit.

Your report should be structured as follows:

- **Introduction:** introduce and motivate your task. You should discuss and cite related work.
- **Data:** illustrate your task and describe the data that you are using.
- **Model:** explain your model.
- **Metric:** present your evaluation metric(s) and justify why it is appropriate.
- **Results:** present your results. You should analyze any trends or patterns you notice in how the model performs. You should include at least two figures visualizing model performance, and you should make it clear which results you are treating as reliable.
- **Conclusion:** summarize your findings and discuss any threats to their validity.
- **References:** provide citations. This does not count towards the required page length.

4 Rubrics

Codebase Rubric (30pt)

- Does the project involve a substantial engineering effort?
- Does the code successfully run the models?
- Is the evaluation metric appropriate?
- Does the code evaluate model performance?
- Is the code commented and organized?
- Is there a README that describes how to run the code?

Presentation Rubric (15pt)

- **Talk (10pt)**
 - Is the task explained well?
 - Does the talk make good use of the slide, without merely reading off of it?
 - Is it clear how model performance will be measured?
- **Slide (5pt)**
 - Does the slide contain an example to illustrate the task?
 - Is the information presented clearly?
 - Are figures captioned and sources cited?

Report Rubric (55pt)

- **Introduction (10pt)**

- Is the research question clearly explained?
- Is the research situated with respect to previous work?
- Is previous work cited properly?
- Is the phenomenon of interest explained clearly?
- Are there examples of the task?
- Is the task clearly articulated?

- **Data (10pt)**

- Is the data clearly explained?
- Was care taken to select high-quality data?
- Is it clear how the data is going to be used?

- **Model (10pt)**

- Is the model design clearly explained?
- Is the model appropriate to the task?
- Are the design decisions related to the model explained clearly and thoroughly?

- **Metric (5pt)**

- Is the evaluation paradigm clear?
- Is it clear how model success or failure will be measured, for each model?
- Is the evaluation metric(s) used to assess model performance clearly explained?
- Is the proposed evaluation metric appropriate?

- **Results (10pt)**

- Is the discussion of model performance clear and thorough?
- Is the model performance contextualized appropriately by discussing baselines and/or previous work?
- Are trends in the model performance highlighted and discussed?
- Are there at least two visualizations of model performance?

- **Conclusion (5pt)**

- Are the findings summarized in a concise and clear way?
- Are the claims about model performance made clear?
- Are threats to the validity of the findings discussed?

- **General (5pt)**

- Is the report well-organized?
- Is it easy for a reader to follow?
- Has it been proofread?