

CS224n: Project Report Instructions

Last updated on March 13, 2019

Each team (both default and custom) submits **one** project report, which is worth 30% of your grade. This document specifies what information you should include in your report – it applies to both default and custom projects.

1 Report contents

Your final report should be written in the same style as a NLP / Deep Learning research paper, and ideally written in a way that a fellow CS224n student could understand. It should be a PDF created using the NeurIPS 2018 LaTeX template.¹ Make sure to turn **on** the camera-ready setting i.e. remove the row numbers, and **remove** the NeurIPS 2018 footer (see Piazza post @2101 for instructions). This will make your report look nicer, but we won't penalize if you don't do it.

Your report should be **6-8 pages**² (not including references). It should contain the following sections (though you can use a different structure if you prefer). Sections with an asterisk (*) were *not* part of the milestone.

Title: The title of your project

Team members: Provide the names and @stanford.edu email addresses of all of your team members.

Abstract: An abstract should concisely (less than 300 words) motivate the problem, describe your aims, describe your contribution, and highlight your main finding(s).

***Introduction:** The introduction explains the problem, why it's difficult, interesting, or important, how and why current methods succeed/fail at the problem, and explains the key ideas of your approach and results. Though an introduction covers similar material as an abstract, the introduction gives more space

¹<https://www.overleaf.com/latex/templates/template-for-neural-information-processing-systems-nips-2018/zxpdbfpqgxs>

²We expect that default project reports may generally be shorter than custom project reports. If you're doing a relatively straightforward default project, make sure to thoroughly describe your approach and experimental details, and thoughtfully discuss your results.

for motivation, detail, references to existing work, and to capture the reader’s interest.

***Related work:** This section helps the reader understand the research context of your work, by providing an overview of existing work in the area.

- You might discuss: papers that inspired your approach, papers that you use as baselines, papers proposing alternative approaches to the problem, papers applying your methods to different tasks, etc.
- This section shouldn’t go into deep detail in any one paper (for example, there probably shouldn’t be any equations) – instead it should explain how the papers relate to each other, and how they relate to your work.

Approach: This section details your approach(es) to the problem. For example, this is where you describe the architecture of your neural network(s), and any other key methods or algorithms.

- You should be specific when describing your main approaches – you probably want to include equations and figures.
- You should also describe your baseline(s). Depending on space constraints, and how standard your baseline is, you might do this in detail, or simply refer the reader to some other paper for the details. Default project teams can do the latter when describing the provided baseline model.
- If any part of your approach is original, make it clear (so we can give you credit!). For models and techniques that aren’t yours, provide references.
- If you’re using any code that you didn’t write yourself, make it clear and provide a reference or link. When describing something you coded yourself, make it clear (so we can give you credit!).

Experiments: This section contains the following.

- **Data:** Describe the dataset(s) you are using (provide references). If it’s not already clear, make sure the associated task is clearly described.
- **Evaluation method:** Describe the evaluation metric(s) you use, plus any other details necessary to understand your evaluation.
- **Experimental details:** How you ran your experiments (e.g. model configurations, learning rate, training time, etc.)
- **Results:** Report the quantitative results that you have found so far. Use a table or plot to compare multiple results and compare against baselines.
 - If you’re a default project team, you should **report the F1 and EM scores you obtained on the test leaderboard** in this section. Make it clear whether you are on the non-PCE or PCE leaderboard. You can also report dev set results if you like.

- Comment on your quantitative results. Are they what you expected? Better than you expected? Worse than you expected? Why do you think that is? What does that tell you about your approach?

***Analysis:** Your report should include some *qualitative evaluation*. That is, try to understand your system (how it works, when it succeeds and when it fails) by measuring or inspecting key characteristics or outputs of your model.

- Types of qualitative evaluation include: commenting on selected examples, error analysis, measuring the performance metric for certain subsets of the data, ablation studies, comparing the behaviours of two systems beyond just the performance metric, and visualizing attention distributions or other activation heatmaps.
- The *Practical Tips for Final Projects Lecture Notes* has a detailed section on qualitative evaluation – you may find it useful to reread it.

***Conclusion:** Summarize the main findings of your project, and what you have learnt. Highlight your achievements, and note the primary limitations of your work. If you like, you can describe avenues for future work.

(Optional) Additional information: Include the following, as necessary:

- **(Custom projects only) Mentor:** Give the name of your mentor. If you have an external mentor, additionally list your CS224n staff mentor.
- **(Optional) External collaborators:** If you have any collaborators who are not CS224n students, list them.
- **(Optional) Sharing project:** If you are sharing your project between CS224n and another class, indicate it here.

References: Your references section should be produced using BibTeX.

***Appendix (optional):** If you wish, you can include an appendix, which should be part of the main PDF, and does not count towards the 6-8 page limit. Appendices can be useful to supply extra details, examples, figures, results, visualizations, etc., that you couldn't fit into the main paper. However, your grader *does not* have to read your appendix, and you should assume that you will be graded based on the content of the main part of your paper only.

2 Improving your technical writing

As a reminder, the *Milestone Instructions* included extensive resources to help you improve your technical writing. You can use these, and previous feedback you've received, to improve your technical writing.

3 Grading and feedback

Your project report will be graded holistically, taking into account many criteria: originality, performance of your methods, complexity of the techniques you used, thoroughness of your evaluation, amount of work put into the project, analysis quality, writeup quality, demonstrating strong understanding, etc.

Your report will be graded by two staff members, whose scores will be combined into your final score. You will also receive some brief feedback on your report. If you are doing a custom project, your CS224n staff mentor will be one of the graders.

4 Code

We ask you to submit your code as a zip file (up to 1MB) to Gradescope.

- **Do** include all project code written or adapted by you.
- **Don't** include the whole source code for off-the-shelf packages that you used without adapting (e.g. CoreNLP or Pytorch)
- **Don't** include model checkpoints or data.

Your code will not be graded – we collect it solely so that we can investigate honor code issues if necessary.

5 Team contributions

If you are a multi-person team, we ask you to submit (on Gradescope) a brief summary of what each team member did for the project (about 1 or 2 sentences per person). We will read these descriptions. For almost all teams, it will have no effect (team members all receive same grade), but for teams with very unequal contribution, we may investigate and/or give different grades to team members.

6 Submission instructions

To summarize, here are the instructions:

- Submit your report to Gradescope under **Project Report [Custom Only]** or **Project Report [Default Only]**. *Make sure to tag all of your team members* – only tagged team members will receive credit.
- Submit your code to Gradescope under **Project Code**.
- If you are a multi-person team, submit a brief description of team contributions to Gradescope under **Project: Team Contributions**.

As previously announced, the due date is 4:30pm on Sunday March 17, and teams can spend up to three late days on the project report – meaning some teams may submit as late as 4:30pm on Wednesday March 20. **However, you cannot submit later than 4:30pm on Wednesday March 20** (even if you're willing to take a penalty). All reports must be uploaded to Gradescope by 4:30pm on Wednesday March 20 in order to be graded – **if you miss this deadline, your report might go ungraded.**

7 Posting reports online

All final reports will be posted on the CS224n website. If you *do not* want your report to be published online, please let us know in a private Piazza post, and we won't upload it.