

CS 175 (W25): Project in Artificial Intelligence

Project Progress Report

Due date: Friday, February 21, 2025 (Pacific Time)

Roy Fox

<https://royf.org/crs/CS175/W25>

The report of your project's progress will consist of a prototype, "proof-of-concept", working implementation of your proposed idea. "Working" here means that, while the environment and algorithm may be limited in complexity, you've implemented a complete and (seemingly) bug-free AI pipeline (e.g. training and deploying an RL agent), you've run it to completion, measured quantitative and qualitative experimental results, and analyzed them (i.e. looked at them with a critical eye). The results themselves may or may not be showing the AI succeeding at the task (or your other metrics), but if they don't, then you should have explicitly formed (but perhaps not yet tested) clear ideas of the reason(s) for that. The idea is for the progress report to be a full draft of the final report, with all parts in place and only missing content within these parts that depends on your further work and results.

Part 1 Clean Up the Website (30 points)

Given that you are now progressing on your project, it is time to update the website to look better.

- **Project/Team Name:** Some of the teams picked very unimaginative titles for their teams and repos. Now is a good time to pick a short name that is relevant, representative, interesting, or fun. Rename the group in Canvas and the repository you are using for your code, and update your website to match (in `_config.yml`). Submit your website's URL on Canvas, making sure to submit the new one if it has changed.
- **Improving the First Page:** Most of you haven't change the first page, `index.md`, as it wasn't required. Well, now it is. Change the main page of the project to be something more interesting, by including things like a link to your source code repository, images and screenshots that are appropriate for your project, links to online resources that you used or found relevant, and anything else you can think of. Near the top of the page, include a short (2–3 sentences) outline of your project.
- **Team page:** Make sure your team page is updated with the names and any additional info (e.g. photos, personal websites, etc. — none of this is mandatory, it's up to you) of all members of your team.

Part 2 Technical Description (35 points)

Add a progress report to your website by adding a filename called `status.md` to your repo's `/docs` folder. This file should start exactly with the following lines:

```
1 ---
2 layout: default
3 title: Status
4 ---
```

After that you'll add your Markdown content, with level-2 headers (starting with `##`) for the following sections:

- **Project Summary:** A short paragraph outlining the main idea of the project (see further instructions here: <https://royf.org/crs/CS175/W25/proposal.pdf>). You now have a better sense of what your project is about than you did in the proposal, so update and clarify beyond that version. Do not change your proposal page! (A few of you have already changed it to reflect a change of topic, but don't change it further, rather keep it as it was at the time you locked down your topic).
- **Approach:** Give a detailed description of your approach, in a few of paragraphs (at least a couple). You should summarize the main algorithm you are using, such as by writing out how it samples data and the equations of the loss(es) it optimizes (you can copy this information from scientific publications or online resources, in which case cite them clearly. The default GitHub Pages we shared includes an example of redering math within Markdown). You should also give details about the approach as it applies to your scenario, such as how you set up inputs and outputs (e.g. states / observations, actions, and rewards), how much data you use (e.g. for how many interaction steps you train), and the values of any hyperparameters (cite your source for default hyperparameter values, and for any changed values detail if and how you tune them and the numbers you end up using). A good guideline is to incorporate sufficient details so that most of your approach is reproducible by a reader. You're encouraged to use figures for this, as appropriate, e.g. as we used in the exercises.
- **Evaluation:** An important aspect of your project, as we mentioned in the beginning, is evaluating your project. Be clear and precise about describing the evaluation setup, for both quantitative and qualitative results. Present the results to convince the reader that you have a working implementation. Use plots, charts, tables, screenshots, figures, etc. as appropriate. For each type of evaluation that you perform, you'll probably need at least 1 or 2 paragraphs (excluding figures etc.) to describe it.
- **Remaining Goals and Challenges:** In a few paragraphs, describe your goals for the remainder of the quarter. At the very least, describe how you consider your prototype to be limited, and what you want to add to make it a complete contribution. Note that if you think your algorithm is working well enough, but have not performed sufficient evaluation to gain insight, doing this should be a goal. Similarly, you may propose comparing with other algorithms / approaches / manually tailored solutions (when feasible) that you did not get a chance to implement, but can enrich your discussion in the final report. Finally, given your experience

so far, describe some of the challenges you anticipate facing by the time your final report is due, to what degree you expected to become obstacles, and what you might try in order to overcome them.

- **Resources Used:** Mention all the resources that you found useful in writing your implementation. This should include everything like code documentation, AI/ML libraries, source code that you used, StackOverflow, etc. You do not have to be comprehensive, but it is important to report the sources that are crucial to your project. One aspect that does need to be comprehensive is a description of any use you made of AI tools.

Part 3 Video Summary (35 points)

Create a ~3 minute video summarizing your progress so far. The video should contain a brief problem description (using images, screenshots, or video captures), as well as “before” and “after” performance captures in an early stage and (respectively) end of the agent’s training. If you identify failure modes (cases where the AI fails), include them as well. You are free to include any other parts of the technical description above, such as a summary of your approach or remaining goals, but no specific part of this is required. You should embed the video on the status report page on the website. For example, if you upload the video on YouTube, add the embed video code provided by YouTube to the top of `status.md` (ask on the forum if any technical issues arise). The video should be of reasonably high quality (minimum resolution of 720p), and the speech, if any, should be comprehensible; but it doesn’t need to be polished at all, as it’s meant as a rough draft toward the final report.