

CS224R Spring 2023 Deep Reinforcement Learning Final Project

1 Final Project Requirements

The CS224R final project requires implementing, evaluating, and documenting a new research idea that pertains to the main topics of the course. Students are expected to prepare a proposal, milestone report, and final report, as well as present their work in the final poster session, with specific details as outlined in this document. This document should serve as your go-to resource for any and all information regarding the final project. If and when updates are made to this document, the course staff will notify the entire class immediately and detail the exact changes.

1.1 Important Information

1.1.1 Grading

The project will count for 50% of the course grade, broken down as follows:

- Project proposal – 10% (see Section 3)
- Project milestone – 5% (see Section 4)
- Poster presentation – 15% (see Section 5)
- Final project report – 20% (see Section 6)

1.1.2 Deadlines

Each component of the final project will be due to [Gradescope](#) by 11:59 p.m. (with the exceptions of the project survey through Google Forms and the in-person poster session) on the following dates:

- Project survey – 4/17/23 (see Section 2)
- Project proposal – 4/26/23 (see Section 3)
- Project milestone – 5/24/23 (see Section 4)
- Poster presentation – 6/7/23, 4-7p.m., Gates AT&T Lawn (see Section 5)
- Final project report – 6/12/23 (see Section 6)

1.1.3 Project Mentors

All final projects will be assigned a single mentor CA to serve as a point of contact. Your project mentor should be your primary point of contact throughout the quarter as you work on your project. This individual will also be the person who ultimately grades your project milestone and final reports. Once your project mentor has been assigned, you should not visit other CAs at office hours to discuss your final project since much of your time will likely be spent simply bringing them up to speed; instead, you should visit your project mentor's office hours as issues, challenges, or questions arise.

In your response to the project survey (Section 2), you may choose to indicate specific members of the teaching staff who you think compliment your interests and/or are well-suited to guide your particular proposed project. To help you with this, listed below are the broad specialities of each CA:

- **Rafael Rafailov:** Offline RL, Model-Based RL, Imitation Learning.
- **Dilip Arumugam:** Bayesian RL, Information-theoretic RL, Efficient Exploration, Hierarchical RL.
- **Annie Xie:** Continual/Lifelong RL, Autonomous RL, Multi-Agent Systems.
- **Regina Wang:** Imitation Learning, Inverse RL, Robotics, Human-in-the-Loop RL.
- **Ansh Khurana:** Offline RL, Computer Vision, Robustness, Meta Learning.
- **Saurabh Kumar:** Continual/Lifelong RL, Efficient Exploration, Hierarchical RL.
- **Jonathan Yang:** Offline RL, Exploration, Imitation Learning, Autonomous RL.
- **Max Sobol Mark:** Offline RL, Model-Based RL, Robotics, Computer Vision.

Clearly, this list of reinforcement-learning topics is not exhaustive and, furthermore, each CA likely has expertise that extends beyond what is listed. That said, even if you find yourself matched with a project mentor whose specialities do not overlap with your chosen project area, your mentor is still broadly well-versed in reinforcement learning, well aware of what constitutes a solid final project for the course, and is ready to help you to the best of their abilities. We anticipate that some final projects will fall outside our areas of expertise as a teaching staff, and that is okay so long as there is clear connection to and application of core course topics.

You should think of your project mentor as an individual who can provide you with feedback as you hit checkpoints in your final project; it would not be in your best interest to seek out advising from your project mentor on a week-by-week basis. Instead, consult your mentor when you have uncertainty on the path ahead or there are major decisions to be made which can dramatically impact the contents of your final report.

1.2 Practicalities

1.2.1 Project Groups

Final projects for the course may be completed by individual students or in teams of no more than 3 people. We encourage everyone to work in groups with the expectation that the overall contributions of the final project be commensurate with the group size.

While the project survey will ask that you list any potential collaborators, project groups are mutable up to and until the deadline of the project proposal. In other words, your group is locked in once you submit your project proposal.

1.2.2 Contributions

To help us (and yourselves) keep track of individual workloads, we ask that project proposals include an anticipated breakdown of individual roles and responsibilities. We will also ask for a similar breakdown of individual contributions in the final project report as well. It is perfectly fine for roles to adapt and change over the course of the project including deleted roles and newly added roles between the proposal and final report. If there are any changes in responsibilities, we ask that they be documented in the final report so we can ensure an equitable division of labor among group members.

Project groups are welcome to work with people not enrolled in the class, however we expect the contribution breakdowns in the project proposal and final report to reflect the work of all members involved (both internal and external). Final projects may be shared between CS224R and another class however, in this case, you should clearly indicate this in your project proposal and we will expect a more ambitious project.

1.2.3 Using Late Days

With the exception of the project poster (which **cannot be submitted late** due to the live poster session), entire project groups may apply a maximum of two late days to any other graded component of the final project, including the final report. Groups cannot aggregate collective late days and late days will be applied to all group members; this means that, to apply two late days, each group member must have two unused late days to spare.

1.2.4 Computing Resources

We will be providing \$20 in cloud computing resources to each student for use on homework assignments and the final project, which should amount to more than 40 hours of GPU time (and much more CPU time). Details about how to access these resources are forthcoming.

2 Project Survey (Due 4/17/23)

All projects should evaluate novel ideas that pertain to deep reinforcement learning, potentially spanning foundational algorithm design to novel applications. For students conducting research in a lab on campus, you are encouraged to pursue a project related to your research area, but are not required to. It's also a good idea to think early about the data (simulated or real) that you'll need to collect, and the computational resources you'll need.

Brainstorming – You may discuss the topic of your final project with course staff by private message on Ed or in office hours. If you are not sure about the topic, we encourage you to speak with us. If you are looking for ideas for your project, we will soon post a document on Ed that contains some ideas we have collected from the AI community, though we also encourage you to come up with your own.

Evaluating Ideas – Here are some examples of *weak* project proposals that do not satisfy the project requirements, and how they can be improved:

1. Weak: run an existing algorithm out of the box on a new dataset.
Strong: develop a modified algorithm that is particularly suited for the challenges of a new application
2. Weak: re-implement an existing reinforcement-learning algorithm
Strong: re-implement a recent paper and investigate an extension of the method that may have been mentioned as future work in the paper.
3. Weak: sweep hyper-parameters, do architecture search of some algorithm
Strong: investigate the weaknesses of a particular algorithm when tested in new ways and pursue a solution

Submission – Everyone, either individually or as part of a larger group, should submit the [Project Survey](#) which will allow us to collect preliminary information on final project topics, determine the total number of groups, and assign project mentors. While project groups should, ideally, be finalized by this stage, we will allow group changes up to the project proposal deadline.

Mentor Assignments – The course staff will release a spreadsheet detailing project mentor assignments for each group within a few days of the survey submission deadline.

3 Project Proposal (Worth 10% – Due 4/26/23)

The project proposal should be an extended abstract motivating and outlining the project you plan to complete. Your proposal should have the following structure:

1. **Objective** 1/4 page. Explain the objectives of the project and why the objectives are relevant and important.
2. **Related Work** 1/4 – 1/2 page. Briefly review the most relevant prior work, and highlight where those works fall short of meeting the objectives described above.
3. **Technical Outline** 1/2 page. Explain your approach at a high-level, making clear the novel technical contribution, and describe the evaluation plan.
4. **Team Contributions** 1/4 page. Include a brief, high-level breakdown of the planned contributions of each individual group member.

Submission & Grading – Submit one proposal PDF per group to [Gradescope](#) under the “Project Proposal” assignment that includes the (finalized) list of project group members. While the deadline for the project proposal is in late-April, we strongly encourage you to begin developing your idea for the project sooner. The grading for the proposal will be light and primarily intended to give you feedback on how to adjust/pivot your proposal to meet the project requirements.

Modifying Contributions – It is perfectly fine for items in your breakdown of team member contributions to remain, adapt, or even disappear between submitting the proposal and the final project. Our goal is to get as even a distribution of workload as possible amongst group members. Any changes to this list of contributions that you do make after submitting the proposal should be documented in your final project report.

4 Milestone Report (Worth 5% – Due 5/24/23)

Your milestone report should be one-page document that answers the following questions:

1. What experiments have you conducted so far? Tell us about one of these experiments.
2. Are there any changes to the research hypothesis or objective from the proposal based on your initial findings?
3. What are the concrete steps that need to happen in order to bring the project from where it is now to completion?

Required Experiment – The milestone report must provide details on at least one experiment that you have done since the proposal. The experiment need not be successful, but you should have attempted something. If it did not work as expected, you should briefly discuss why. You are encouraged to include a plot or figure.

Submission & Grading – Submit one one-page milestone report PDF per group, with the names of all project members, to [Gradescope](#) under the assignment “Milestone Report.” The milestone report will be graded by your project mentor. The grading of the milestone report will be light as the primary goal is to provide you feedback so that there is a clear path going forward to the poster presentation and final report.

5 Poster Presentation (Worth 15% – Due 6/7/23, 4-7p.m.)

Students are expected to prepare a research poster that includes the high-level motivation, the problem that the project is aiming to solve, any prior work needed to understand the approach, an outline of the approach, a summary of the experimental results, and a description of takeaways and future work. The poster should also include any work that is planned to be done before completion of the final report.

Poster Dimensions – We recommend printing posters in landscape orientation with size 24" × 36". While other sizes are acceptable, this size will work best with the poster boards and easels we will provide for you at the poster session, so please make sure your poster isn't too much bigger or smaller than this.

Poster Printing – Please note that you are not required to do commercial poster printing; you can print the poster on smaller sheets (letter size or A4 size) using regular printers and tape them together. Adobe Acrobat Reader automates this — look for the Poster options in the Print dialog. That said, both FedEx and Walgreens have deals where, if it's your first time printing or you make a new account with another email, you can print a poster for around \$15-20. You are expected to print out the posters yourself, so please plan ahead! Some other options that have been suggested in other CS courses include: Lathrop Library's Tech Desk, CVS Photo, Biotech Productions, and Walmart.

Poster Exemptions – At least one group member must be present at the poster session, unless you have reached out via email to the course staff and received special permission to submit a video presentation. The only exception is if the project group is made up entirely of (remote) SCPD students, in which case, the group will be expected (by default) to submit a five-minute video presentation of their poster instead. A Google Drive or YouTube link pointing to the video should be uploaded to [Gradescope](#) under the "Project Poster Video" assignment.

Location & Time – The poster session will be held on 6/7/23, 4-7p.m. at the Gates AT&T Lawn.

6 Final Report (Worth 20% – Due 6/12/23)

The final report should be in the style of a research paper, preceded by a one-page extended abstract. The one-page extended abstract should summarize the main findings and accomplishments of your final project, while the main paper should describe and motivate the method in detail, discuss relevant prior work and how the project makes new contributions relative to these prior works, and discuss the results, including any relevant figures or plots. The extended abstract should be attached as the first page of the full report.

Formatting & Contributions – Successful reports will have a main body that is about eight pages in length, but there are no hard length requirements or formatting requirements, except the one-page extended abstract. For project groups, please include a fully-updated breakdown of the contributions of individual team members and highlight why these adjustments were necessary from your original breakdown in the project proposal; this is as much a reflection on the research process for you as it is a tool for us to track equitable division of labor.

Submission & Grading – Submit one final report per group with the names of all project members to [Gradescope](#) under the assignment “Final Project Report.”