

Final Project

CS 6501 Reinforcement Learning (Spring 2024)*

1 Goals

- Give students the chance to apply what's learned in the class to problems they care about.
- Encourage students to delve into specific topics, discovering something not covered in the course.
- Promote collaborations and knowledge exchange between classmates from various backgrounds.

2 Types of Projects

The types of projects are widely open. Below are some, but not exhaustive, examples.

- **Application:** Pick a decision-making task. Try to use RL techniques (or combining other techniques) to complete the task.
- **Algorithm Design:** Identify weakness of existing algorithms in specific tasks. Try to improve existing algorithms or propose new algorithms.
- **Systematic Comparison:** Consider a specific decision-making task that can be dealt with by different methods. Provide a systematic comparison on the advantages/disadvantages/trade-offs among different methods.
- **Theoretical Understanding:** Identify phenomenons or algorithms that are never theoretically understood or analyzed in the literature. Try to use mathematical tools to justify or analyze them. Another possibility is to identify a unified framework to understand different algorithms.
- **Literature Survey:** Pick an area in decision making. Provide a detailed and clear overview and comparison of existing techniques. Identify important open problems in the area.

3 Grouping

1-3 students in a group.

*This final project specification is inspired by those of Shangtong Zhang's CS 6316 and Yen-Ling Kuo's CS 6501.

4 Breakdown

4.1 Proposal (5%) – due on February 16

- Team members
- Proposed problem and motivations
- Preliminary plans on the approach
- Related works to read
- Timeline plan

4.2 Milestone (5%) – due on March 29

Use NeurIPS format: ([download here](#)). No more than 4 pages excluding references.

- Title, Authors
- Introduction: Introduce the problem and its motivation.
- Related works: Describe the related works and how they relate to your project.
- Approach: Describe the methods you tried or intend to apply to your problem.
- Preliminary Results: Describe the current state of the project and your results so far.

4.3 Presentation (10%) – April 15, 17, 22, 24

Give a comprehensive presentation on your final project in 15mins, plus 3mins of Q&A. The presentation should be easy to understand, have a clear motivation, and include necessary background if the methodology is outside the course materials. Please clearly present your methods and experiments if there are. You do not need to have all the results before the presentation, but just present what you have done so far and the plans before the final report due day.

4.4 Final report (15%) – due on May 5

Use NeurIPS format: ([download here](#)). No more than 9 pages of main texts excluding references. You can have additional supplementary materials with unbounded number of pages. The content should be comparable to a standard conference or workshop paper, typical with the following content (of course you can add to or minus from them based on your need):

- Abstract: Briefly describe the problem you are solving, your methods, and the key results.
- Introduction: Introduce the problem, its importance, and a summary of your results.
- Related Work: Describe the existing works (published papers or existing systems) that are related to your project. How is your method similar to or different from those works?

- Preliminaries: Introduce necessary background knowledge of your approach. For example, if your method is adapted from the algorithm of some previous work, you can introduce their algorithm in more detail in this part.
- Methods: Describe the methods used to solve the problem you presented in the introduction. Why do you design the method/model this way? Have you considered other designs? You should include figures and/or tables to help illustrate your method or compare it with other methods.
- Experiments: Describe the experiments you set up to evaluate your proposed method. Discuss the results you got and analysis to show why the method works or doesn't work. The expected experiments will vary for different projects. Here are some items you may consider: explain your metrics, what methods you compare with (might be a previous work), any (hyper)parameters or architecture that may affect the results, ablation results, visualization of models/results, qualitative examples, etc.
- Conclusion: Summarize your key findings and what you have learned. Suggest ideas for future extensions/work if someone is continuing your idea.

(The above breakdown is borrowed from Yen-Ling Kuo's website)

Notice that May is the season of NeurIPS conference paper deadline and ICML workshop paper deadlines. You're encouraged to submit your project to them without listing me as an author. If you need any suggestions on the submission, feel free to discuss with me.

5 Late Policy

Every late day costs 20% deduction in the grade.