

Final Project

Important dates:

[due date here] – Form your teams

In class [shortly after teams] – Project conferences

[0.5-1 weeks after conferences] – Proposal

[one week after proposal] – Progress Report 1 (implementation)

[one week after progress 1] – Progress Report 2 (results)

[1-1.5 weeks after progress 2] – Complete Draft

[one week after draft] – Paper Review

[during final exam slot] Presentations (**no late submissions**)

[after presentations] – Final Paper Submission (**no late submissions**)

0. Introduction

In the final project for this class, you will have the opportunity to dig deeper into an artificial intelligence problem that particularly interests you. You and your partner will select a problem, develop and evaluate a solution approach, and report your findings through a paper and presentation at the [course name here] Conference on AI Research. All told the project will span over a month of work and will account for 20% of your final grade.

Forming Groups

You may choose your partner for this project if you wish, or you may choose to be assigned a partner. In the case of an odd class size I will permit one 3-person group (first come, first serve). I will post an “assignment” where you can tell me the members of your team and some other basic information.

I would like you to register your group by **[deadline]**. After that I will assign partners to anyone not already in a registered group (in order to ensure that everyone has a partner in time to begin!).

Choosing Your Topic

Your project should be closely related to the AI problems/methods we have discussed (or will shortly discuss) this semester: heuristic search, local search, game playing, MDPs, reinforcement learning, constraint satisfaction problems, and logical or probabilistic inference. If you are highly motivated to pursue a topic that we have not discussed in class, please discuss it with me as soon as possible.

This is a simplification, but in the AI literature, there tend to be two broad categories of research paper:

- *Application focused* papers present an approach (or approaches) for obtaining good performance on a challenging and/or important problem. The main goal of such a paper is to clearly (and completely) describe the approach, evaluate its effectiveness in comparison to alternative approaches, and to investigate why it is as effective as it is in hopes of extracting more general lessons. Typically multiple approaches are compared in order to find the most effective.
- *Method focused* papers present a new algorithm/method (or a novel combination of existing methods) that addresses some shortcomings of existing work, or that can solve problems that could not be solved previously. In these papers the goal is to clearly describe the new approach and evaluate its strengths and weaknesses in comparison to existing approaches. Typically the new method is evaluated in multiple example problems that help to illustrate when it is most effective, and when it is not the best choice.

Your project may be of either flavor (or some combination). Note that you should not expect to break any records in a domain or to develop a revolutionary new algorithm in a few short weeks! It's okay if your idea has already been covered in the literature, or if your approach doesn't turn out to work very well in the end. The important thing is that the conceptualization, execution, evaluation, and communication of your project are of high quality, reflecting your understanding and engagement with the ideas at hand.

Your project should not be trivial (e.g. make an agent that plays Rock-Paper-Scissors) but should also not be too ambitious (e.g. make an agent that plays Jeopardy). A good rule of thumb is to aim for a project that is about at the scope of the smaller projects from earlier in the semester, except working with a problem or a method that we have not studied. I will try to help you find a reasonable scope for the given timeframe. If you are having trouble thinking of a project, feel free to come meet with me. I'm sure we can brainstorm some project ideas together.

The main products of your project will be your final report and final presentation, but there are several other components to the project as well. These will be described below.

Project Conferences

I want to meet with each team to discuss your project ideas, so I can help you develop a project with an appropriate scope for the available time. I may be able to anticipate significant challenges you haven't foreseen or I may be able to identify that a problem is simpler than you first estimate. Please come to this conference with at least one concrete idea for a project including problems you want to study *and* approaches you may apply.

I will hold project conferences **during class on [dates]**. I will assign a 15 minute timeslot to each team during that time. In some cases we might want to follow up on that conversation with a second meeting.

1. Project Proposal (2 pts)

Due **[deadline]**

The [course name here] Conference on AI Research has announced the availability of small grants for young researchers (CAIRY grants), in hopes of stimulating submissions to the conference. Submit a proposal to the CAIRY grant committee outlining your project.

Your proposal should be at most 2 pages long (single spaced). It must include:

- A title for your project (might not end up being the title of your paper)
- A clear description of the problem you propose to solve
- The motivation for solving this problem (what makes it challenging and/or important?)
- Acknowledgement of any existing work in the literature regarding this problem that you intend to reference (include full citations)
- Acknowledgement of any existing code, simulators, data, etc. that you plan to make use of
- The approach you plan to take
- How you plan to evaluate the effectiveness of your approach

Most grant proposals would also include a budget, but since the CAIRY grant is for small projects and for a fixed amount (*ahem* \$0), the committee did not deem this necessary, so don't worry about that.

A grant proposal is a form of persuasive writing. The CAIRY grant is competitive, so your proposal will have to stand out to a reviewer who may be reading dozens of submissions. Your proposal should be designed to quickly engage that reviewer, and to communicate your ideas clearly and efficiently. Your paragraph structure and flow, your wording, your formatting, and even your title can all be used to make your proposal as effective as possible. Essentially your goal is to convince your reader that your project is a good investment.

It is understood of course that you will not have all the details worked out and the reviewer doesn't want to hear all the details anyway. After reading your proposal the reader should be convinced that your project is worth pursuing and that you have thought through the key issues and have a solid plan for making progress and for evaluating your success.

2. Progress Report 1 – Implementation (1 pt)

Due [deadline]

By this point you should have made significant progress on implementing, testing, and debugging your agent, your experimental domains, and any comparison methods. You will submit a brief report (no more than 1 page) on your implementation progress, as well as the code you have so far (I will not be grading your code – this is just for a sanity check). Your report should discuss:

- What you have completed so far
- What you still have left to do
- Any particular challenges you have faced and how you have (or plan to) overcome them
- Any changes you have had to make to your original proposed project
- How you have been testing your implementation for correctness

Your report can be informal (you do not need to spend a lot of time drafting/revising it), but it should nevertheless be clear and comprehensive. Your grade for this part will be primarily based on your progress and the thoughtfulness reflected in your report. Note that if you have run into unexpected difficulties that have prevented you from finishing your implementation as quickly as you'd hoped, that does not necessarily mean you will be penalized as long as you are dealing with the challenges in a reasonable way. Uncovering unforeseen barriers *is* progress, though it doesn't always feel that way!

3. Progress Report 2 – Results (1 pts)

Due [deadline]

By this point you should have made significant progress on generating experimental results. You will submit a brief report (at most 1 page of text plus additional figures as appropriate) on the results you have gathered. Your report should discuss:

- What results you have gathered thus far (include graphs/tables as appropriate)
- What results you still plan to gather
- A brief description of what the results say so far
- Any changes you have had to make to your original proposed project (e.g. experiments that turned out to be infeasible, or additional experiments that have been suggested by the results)

Your report can be informal (you do not need to spend a lot of time drafting/revising it), but it should nevertheless be clear and comprehensive. You might want to think about this as a dry run of the results section in your final report! As with the first progress report, your grade for this part will be primarily based on your progress and the thoughtfulness reflected in your report. Again, if you have run into unexpected difficulties that have prevented you from gathering results as

quickly as you'd hoped, that does not necessarily mean you will be penalized if you are responding reasonably to the challenge.

4. Complete Draft (1 pt)

Due [deadline]

You will submit a draft of your paper about two weeks before the final due date. **I will not be grading your draft** except to simply check for completeness (if you have something specific that you'd like my advice about, I'd be happy to discuss it with you one-on-one). However, as in a real conference, you will receive paper reviews from your peers, which should hopefully help you improve your paper for your final submission (see below). Note that this *should not be your first draft!* You should polish your paper as much as you can so your reviewers are not bogged down by surface-level issues and can suggest revisions that you might not have otherwise seen.

Your paper should follow the Association for the Advancement of Artificial Intelligence style guidelines (a LaTeX style file and a Word template are available) and should be no more than 6 pages, the length of a conference paper at the real AAI conference (you may have references on a seventh page if necessary). You should imagine that you are writing this paper for an undergraduate research conference and as such, you may assume a basic shared knowledge of AI methods, but must provide any additional background necessary to understand your project. You should adopt the formal tone of a research paper: avoid conversational language and aim for precision, conciseness, and clarity. Your paper should include:

- An “abstract”: a 1-paragraph summary of the project and your findings
- A clear description of the problem and its importance
- A summary of any relevant related work you referred to, and how it relates to your approach
- A description of your algorithm/method/approach – ideally it should be clear and precise enough that an informed reader could re-implement your project
- A description of the experimental methodology you used to evaluate your approach – ideally a reader should be able to reproduce your results
- A clear presentation and analysis of your results
- A discussion of the implications of your results for others working on related problems
- A brief summary of the conclusions of your project

Ultimately, your grade will be based on the substance of your project (whether the project was conceptualized in way that reflects understanding of the core ideas, executed correctly, and thoroughly evaluated in a principled manner) and the effectiveness of your paper's presentation (whether it is clear, engaging, and comprehensive, and whether the style is appropriate for the stated audience/venue). Notably **your grade does not depend on the success of your**

approach! In fact, given the available time, it is likely that what you try will not work very well. I am looking at the quality of your scientific methodology and communication.

Some teams might not have completed the work of gathering results by this deadline. In that case, you should still write as complete a draft as you can, leaving room for the pending results and describing what you expect from them. Depending on the extent of the missing content, you might not receive full credit for completeness but this is intentionally a small component of the overall project grade.

5. Paper Review (2 pts)

Due [deadline]

Your paper will be read and reviewed by at least two of your classmates; hopefully their feedback will be useful to you as you develop your final version. In turn, you will read and (anonymously) review a paper. Your review should be similar to the one you wrote in Project 4. It should be 1-2 pages long and address the following:

- First, give a 1-paragraph summary of the paper (in your own words) that describes the problem being solved, the approach taken (at a high level), and the key findings.
- What are the main claims of the paper and are they well supported?
 - If a claim is supported empirically, does the experiment provide compelling evidence?
 - If a claim is supported theoretically, is the argument/proof sound/correct?
 - Does the paper discuss *both* the strengths and weaknesses of the approach? How effectively are the strengths/weaknesses illustrated and evaluated?
- Is the paper well-written and clear?
 - Were you able to follow the logic of the paper?
 - Were important concepts clearly introduced?
 - Are both the high-level ideas and the low-level details communicated effectively?
 - Could an informed reader reproduce the results presented in the paper?
- How significant are the contributions of the paper likely to be?
 - Did the paper do a good job of motivating the problem? Are you convinced that it is important?
 - Did the paper do a good job of motivating its approach? Are you convinced that this is a promising direction?
 - If the paper discusses existing work, does it clearly describe the relationship to the results presented in the paper?
- Lastly give a brief summary of your review and an overall assessment of the quality of the paper.

As before, you should keep in mind your multi-faceted audience, though in this case you do not have to pretend! You should clearly explain your reasoning and support your arguments so your review can be understood by a “program chair” (me!) who has not carefully read the paper (true story!). For the sake of the authors (your classmates!) you should also make sure your review takes a respectful tone, demonstrates your understanding of the paper, and (most importantly) offers a *constructive* critique by discussing both positive aspects of the paper and specific, concrete ways that it could be improved.

Remember, your classmates are counting on you to help them improve their paper, just as you are counting on your reviewers, so take your reviewing responsibility seriously. That said, a reviewer is not a copy editor – if you spot a few type-os it could be helpful to point them out, but you should focus your energy on higher-level concerns. Do not try to re-write the paper or identify every tiny flaw for your classmates.

As before, for grading I am far more interested in the quality of your writing than your opinions about the paper. If you take a patently absurd position then you may lose points, but generally I will focus on whether your points are well supported and stated clearly and not on whether I agree with your assessment.

6. Oral Presentation (3 pts)

[date] (no late submissions)

The [course name here] Conference on AI Research will be held during our scheduled final exam slot. **Your project should be effectively complete by this point.** A presentation that contains placeholders for pending results is not very effective! You will give a **10-minute presentation** on your project, followed by 5 minutes for questions. Remember that your audience consists of your classmates: they are familiar with many general AI concepts but *they are not familiar with your project*. Your presentation should be polished, clear, and accessible. It should communicate:

- The problem you worked on
- Why the problem is challenging and worth solving
- What approach you took
- How you evaluated the effectiveness of your approach
- How well your approach performed

Note that 10 minutes is *very little time*. Presentations at real research conferences are typically at most 15 minutes long, and can be as short as 1 minute! Giving a good, short presentation is a difficult (and valuable!) skill and *takes practice*. You cannot expect to present every detail of your project, so you'll have to pick out the key points to get across.

You are very unlikely to write a good presentation of the appropriate length in your first try. I strongly urge you to develop your presentation and then to *practice it*. Practicing will reveal whether it is the right length and will also help you identify points that are difficult to clearly get across and might need more work.

Note on PowerPoint: Visual aids are an important part of an effective presentation. However, there are ways to use slides effectively and ways to use them ineffectively. Your slides are a visual support for what you are saying; they are not a script. If someone is reading your slides, they are not listening to you! If *you* are reading your slides then you are not connecting with your audience. You should minimize text on your slides. No text is often best! If you must have text to emphasize key points, you should avoid full sentences; stick to pithy statements of only a few words. Instead of text, your slides should contain what you want to point to while you are talking. If you find yourself drawing in the air as you describe your project, you should actually draw that on a slide!

Your grade will be based on the content and effectiveness of your presentation.

7. Final Draft (10 pts)

Due [**deadline**] (no late submissions)

The final version of your paper will be due at the end of the final exam period. In addition to the criteria listed above, I will also take into account the extent to which you incorporated the feedback from the reviews you received. You are not beholden to follow every specific piece of advice from your reviewers, but you should thoughtfully consider their suggestions and address at least the spirit of their concerns. They are, after all, your audience!