

CMSC 671 (AI), Fall 2024 – Group Project Description

Please read this carefully and start thinking about the steps involved in designing and implementing your project. The project requirements are subject to changes and improvements (but hopefully not too many).

In this project, you will have the opportunity to explore an aspect of artificial intelligence that interests you. The specific problem you tackle and the approach you take do not have to be something we cover in class (though it may be easier if it is).

Because this project should be something that interests you, it is okay for it to overlap with your research interests (if applicable). Therefore, while in general you must form teams of between two and four people, solo projects that are heavily drawn on your existing research will be allowed. However, to do a solo project you must explicitly obtain written permission.

You have a lot of freedom in this project. This is not (necessarily) a research project, but you *must* do some investigation into how other people have tackled the same (or similar) problems, and compare and contrast your approach with past ones: that is, your effort does not exist in a vacuum, and you must acknowledge this and explain how your work fits in with previous/on-going work.

The following lists potential topics. You are not restricted to these.

1. Choosing a game, and creating an agent to play that game using artificial intelligence, for example, chess, bridge, Minecraft, Wordle;
2. Developing an agent designed to interact intelligently with people in some context, for example, a chatbot or virtual assistant;
3. Applying machine learning techniques to some existing corpus to draw conclusions, for example, a plagiarism detector, a COVID-19 outbreak predictor;
4. Developing an agent that (hypothetically) interacts with some real-world phenomenon, for example, the stock market;
5. Developing some toolkit for solving a standard type of AI problem, or extending such a toolkit with new capabilities (a software development project);
6. Developing a recommender system for some existing corpus, for example, to recommend Netflix suggestions;
7. Using NLP to analyze documents and draw intelligent conclusions, for example, a resume analyzer, a spoiler detector;
8. Formulating, implementing and comparing a novel solution to an existing problem;
9. Formulating, acquiring data for, and applying a sufficient baseline for a novel task.

You can implement your project in any language as long as it runs on a standard Linux system, and may use appropriate libraries and toolkits, assuming they do not implement the actual core of your project. While you do not have to discuss your project idea with us before the initial design, you are very much encouraged to do so, and we will be happy to talk with you about options.

Group Work

You do not have to, but may, work in teams of up to 4 people. While larger groups should tackle more ambitious projects, the deliverables (e.g., the length of the final paper) are the same regardless of group size; it is therefore in your best interests to work as a group where possible.

Deliverables

This project has four main deliverables, detailed below.

Project Design (20%)

You must submit a project design via Blackboard. This must be a PDF, and will probably take ~2 pages of text. It must be written as a document—in prose, not as a list or in bullet form—and should be divided into appropriate sections. The project design should use the AAAI formatting requirements, which means you will work in LaTeX (Overleaf), Word, or a Google Doc using Word-style formatting. You can find the [AAAI Author Kit](https://aaai.org/Conferences/AAAI-22/submission-guidelines/) at this page: <https://aaai.org/Conferences/AAAI-22/submission-guidelines/>. Using a conference template to format your work can be a little time consuming, so please make sure you download the author kit and start using it well in advance.

Only one project design should be submitted per group, and all members should have input into the project design. (Remember, submitting something with someone's name on it when they did not in fact work on it is a form of academic dishonesty.) All of the following points should be covered in as much text as necessary to give us a full idea about your project.

Your project design should contain:

1. The names of all participating team members, as authors.
2. A description and motivation of the project you have decided to do.
3. A description of the AI technique(s) you are going to use to implement your project.
4. A description of what you will implement for Phase I.
5. A description of what you will implement for Phase II.
6. How your implemented system draws on ideas from the AI literature:
 - Cite initial references for the AI concepts you mention (when citing the textbook, include section numbers). Do not consider just recent papers; try to find papers from the past 25 years. You may be surprised at the recurrence of ideas.
 - This can be focused on material and concepts that we covered in class, or—if your project uses techniques not covered in class—should provide sufficient references to external sources.
 - If you like, you may also discuss methods that you would/could use but don't expect to implement within the scope of the semester.
7. Your evaluation strategy. How will you test your system to see if it's working? How will you quantify how well it works?
 - Examples of possible evaluation strategies: Have your game player play against itself; have your game player play against software that solves the same problem, e.g., gnu chess; interact with a set of 5 humans and ask them how smooth the experience was; run your agent against a real-world problem and track how well it would have done; test against a held-out test set of data; compare results against human evaluation of a problem.

It is fine (and likely) to make changes to the design after you submit this document. However, because the project report both guides your project and is a significant percentage of the project grade, you should start working on it well in advance of the due date, and turn in something that covers all the points above. Your document should be written in clear, understandable English. You will want to reuse parts of this text for the final report, so plan accordingly.

Phase I (25%)

For Phase I, you should submit an updated version of your project description (~3–4 pages), along with the code you have developed to date. The primary goal is to have something that is working at a basic level, so that you can improve its performance in the next phase—for example, an agent that can play a game of bridge, but doesn't yet choose its cards as intelligently as it will for the final project.

Your updated project description should be an edited version of the original design you turned in, but your Phase I section should now be completely filled out with progress to date, and your evaluation section should have evaluations of the current functionality of your project. You can use the project design to discuss major design changes that occurred, expand on your bibliography and references, and otherwise provide information about the progress of the project.

Your code should be compatible with standard Linux machines, and please make it easy to run. You may use any external libraries you wish to use. Your turnin must include everything necessary to run your project, including a dataset if appropriate. (If your dataset is large, please make arrangements with us to provide it separately from your Blackboard turnin.) Please provide detailed instructions, in a README file, on how we should run everything.

Phase II (25%)

For Phase II, you should submit your final project code, showing your agent or intelligent system completing its task. At this stage you should have completed programming, and have a complete system that you can demonstrate working on a task. What exactly this will entail will depend on your project and how you have chosen to evaluate your project.

Your code should be compatible with standard Linux machines, and please make it easy to run. You may use any external libraries you wish to use. Your turnin must include everything necessary to run your project, including a dataset if appropriate. Please provide detailed instructions, in a README file, on how we should run everything.

Final Report (30%)

A final report, in AAAI style. The report should be roughly 6-8 pages, *not* including references (which will usually be another 1-2 pages). In writing the final report, you may reuse what you wrote for the progress report, but be sure to specifically and thoroughly address the final report components. Please use sections appropriately to organize your report; I suggest possible section headers below, but you may organize the report in any reasonable, coherent, readable way.

Your final report should contain:

1. A title.
2. The names of all participating team members, as authors.
3. A brief (≤ 250 word) abstract.
4. An **Introduction** section containing a description and motivation of the project.
5. A **Related Work** section (longer than in the progress update) of how your solution fits in with what has been previously done on the problem (or related problems).

- Cite references for the AI concepts you mention (when citing the textbook, include section numbers). Do not consider just recent papers; try to find papers from the past 25 years. You may be surprised at the recurrence of ideas.
- 6. An **Approach** section describing the AI technique(s) you used to implement your project.
 - This is the core of the document, describing what you did, and is likely to be broken up into subsections. This should revolve around a description of what you implemented.
- 7. A **Results** section containing your evaluation strategy.
 - If your project involves experimentation, proper experimental methodology.
 - A clear description and analysis of your results.
 - An analysis of the limitations of your work.
- 8. A **Conclusion** section containing any further discussion of the work done and identifying potential follow-up work, both short- and long-term.
 - You may want to talk about better methods that you weren't able to implement. If you ended up with something relatively simple, but had some ambitious/interesting ideas that you weren't able to get working in the scope of the semester, you should talk about it here.

Do not write just to fill up space. While six pages may seem like a lot at first, you can fill it up very quickly. Be sure to cite appropriately and follow all academic honesty standards. You may include figures (your own); be sure to only use figures you have made. However, make the figures count: do not include them simply to pad the paper.

A Note on Cooperation

Teams can and should be working together closely, and you are welcome to use any external libraries that do not implement the core of what you propose to do. However, it is *critical* that you thoroughly document any external code that you use, whether it be an included library, something found on Stack Exchange, or something from the book. Using code from a source that you don't properly acknowledge is plagiarism.

If you receive assistance from someone external to your project team (other than the instructional staff), even if it is someone else in the class, you must acknowledge that assistance in your code (if it is coding help) *and* in the project writeups. Provide specifics about how that individual helped.

Hints and Tips

Google Scholar is an easy way to find linked and cited papers. The AAAI digital library also offers an extensive listing of AI-based conferences and proceedings. Of particular relevance are the flagship AAAI, ICML (International Conference on Machine Learning), and KDD (Knowledge Discovery and Data Mining) proceedings.

We are very much available to help you with any aspect of your project, during office hours, over email, or by appointment. Please feel free to take advantage of this to discuss topics, finding relevant papers, and the direction of your project.

This document will be updated over the course of the semester as people have questions, etc. New versions will be added to the schedule with version numbers appended.