

UNIVERSITY OF MAINE
SCHOOL OF COMPUTING AND INFORMATION SCIENCE
COS 470/570: INTRODUCTION TO ARTIFICIAL INTELLIGENCE
Semester Project
Spring 2019

The class includes a semester project to tie together what you learn during the semester. This can be an individual or group (2–3 people) project. If you choose to work in a group, however, members must be all in COS 470 or all in COS 570.

Selecting a project

The first thing to do is to pick a topic for the project and clear it with me. The project must be large enough to be worth 25% of your grade—i.e., a substantial project—and must focus on AI content—e.g., though a pretty GUI is nice and takes time, the important thing is the AI, not other programming you may do.

I have included some project ideas at the end of this document in case you have trouble thinking of a topic. Keep in mind, though, that almost *anything* can be the topic of an AI project, so you can take something you're interested in and figure out how to incorporate AI into it.

When selecting a project topic, keep in mind:

- You will be stuck with this project all semester, so pick something you're interested in.
- The best way to start is to decide on a particular area of AI or a particular domain in which you are interested. Example AI areas are things like constraint satisfaction, planning, deep learning, etc. A *domain* is what you are applying AI techniques to—the area from which the problem you are solving is drawn. Example domains include controlling autonomous vehicles (autonomous underwater vehicles [AUVs], autonomous land vehicles [ALVs], unmanned aerial vehicles [UAVs], space vehicles, etc.), image labeling, medical diagnosis, controlling an automated house, factory scheduling, etc. The two are related. If you are interested in a particular AI area, then unless you are interested solely in the theoretical aspects, the next thing you would do is to decide on the domain in which you'd implement the AI technique(s). If you are interested in a domain, then the next thing you would do would be to decide which technique(s) you would use for your problem.
 - Although it's been my experience that most students do a programming (implementation) project, you may propose an implementation project, a design project, a theoretical project, or a combination. It's even possible to propose a literature review project, if there is some particular area we don't cover that you would like to delve into in more

detail (though such a project would entail a *lot* of work: reading, detailed analysis, and writing).

- Think big! You will not be penalized for not being able to accomplish what you set out to do if it was overly ambitious and you did a good job on a substantial portion of it. For example, if you are interested in cooperative systems or interface agents, you might come up with a project to create agents that can be used to schedule meetings for co-workers. You probably won't have a complete agent-based system at the end of the semester, but you may have quite a few of the modules completed as well as the overall design. For a truly large project, you may have *nothing* implemented, but rather have a detailed design.

You will initially propose what you think you can accomplish. As the semester progresses, it will become clear what can actually be done and what should be left out; you will need to keep me informed and get my agreement to changes from what you proposed. The amount and quality of your effort are what count, not the percentage to of the initially-proposed project that got done—as long as you keep me in the loop and get my approval.

- “Double-dipping” is, in general, not allowed. That is, you can't get paid or get other class credit for work you are doing in this class. If you are considering choosing something related to your Capstone project, you need to talk to me, the COS 497 (or 397) teacher, *and* your Capstone advisor to get approval.

If you are working on a masters or PhD, then your project in this course *can* be part of that—but you still need to talk to me (and ideally your advisor) to get permission. And there's nothing at all wrong with working on a project in this class that ultimately becomes a Capstone, MS, or PhD topic.

You *can* look for ways your project relates to other assigned work in class. For example, if your project needs a theorem prover, you should talk to me about using the theorem-proving programming assignment as part of your overall project.

- As you work on your project, make sure that *you are actually doing AI*. (See how I think this is important enough that I'm repeating myself here?) Do not get bogged down in interfaces, etc., that have nothing to do with the AI in your project. This is a common problem, and if you are at all unsure, talk to me as early as possible. I have seen some truly impressive projects that I had to give poor grades because, impressive as they were, there was little or no AI in them.

If you choose to work in a group, keep in mind:

- You'll be stuck with any project partners all semester, so choose wisely. Form the group either because of common interest in a topic or because you really get along with the other person/people. Keep in mind that you will be spending a large amount of time with them!
- I expect all members to do their equal share of the work—and by that, I mean the **AI-related** work: you can't split the project work up into, for example, the AI stuff and the GUI, etc.
- I will expect more of a group project than an individual project; twice the effort for a group of two, three times as much for a group of three.

Project components

Pre-proposal

This is a short (1–2 page) description of what you are thinking about doing for your project. This should include:

- the subject of the project;
- any AI technique(s) you think it may involve;
- what type of project it is (implementation, etc.; see below); and
- whether it is an individual or group project, and who is going to work on the project if it is a group project.

If you have not narrowed the project down to a single subject, you may list two or three projects you are considering and I can suggest which one to do (and/or which ones not to do).

I will get these back to you quickly with feedback and suggestions so you can finalize the project you are going to propose.

Proposal

In this longer document, you will *really* propose your project, much like we AI researchers do when approaching a funding agency (e.g., the National Science Foundation) for a grant. The proposal requires a bit of work, which is why there is a pre-proposal so you don't waste your time. This is a written document that will follow the general outline for, say, an NSF proposal. It must include:

1. Project summary: a very short (< 1 page) overview of what you are proposing.
2. Project description: this (the “narrative”, as it’s called) is the bulk of the proposal. Although the subsections don’t need to be named like this, the narrative should include:
 - (a) Motivation: Why are you doing the project? What is the need/problem to be solved? Why is the problem important (to you, to society, etc.)?
 - (b) Background and related work: discussion of the problem area and what others have done. For example, if you are working on automating a home, you might talk about what’s needed to do that and how it’s currently done (Philips Hue, Apple Homekit, etc.). For the related work:
 - i. COS 470 students: This will be a *very* light “literature review”: read the related areas of the textbook and discuss them briefly as best you can. If you can’t yet understand the concepts because we haven’t gotten to the background needed yet, then you can say how you *think* the areas are related to what you want to do.
 - ii. COS 570 students: You should do what the COS 470 students do, but also do a brief literature search (using, e.g., CiteSEER or Google Scholar), read two or three of the most on-point papers, and discuss them briefly as to how they relate to your project.

- (c) Proposed approach: what you plan to do, briefly, but in as much detail as possible. Here is where you will, for example (for an implementation project), describe the AI techniques you plan to use, how you plan to use them, the software architecture, etc.
 - (d) Project plan/schedule: Here, you will tell me what tasks are needed to complete the project and when you expect those to be done. You should at least specify milestones (including dates), but you can also get fancier (start/end times, a Gantt chart, etc.).
 - (e) Deliverables: what will be the result of having done this project—a program, a paper, etc.?
3. Equipment and facilities needed: No, I'm not going to provide equipment for you,¹ but I want you think a little bit about if you have the resources to actually do the project. For design and other non-implementation projects, this will just consist of a statement saying you don't need equipment.
- For most projects, the “facilities” part may be just be a statement saying that none are needed. For some implementation projects, though, you may need to think a bit about this. If you are planning to work on home automation, for example, then you need a home to try it out in; this will get you thinking about who's house/apartment you're going to use.
- For any project, if you need meeting space or space to work, list it here. I may be able to let some of you work in our lab, for example, if the need arises, and there is student meeting space on the first floor of Boardman.
4. Budget: No, I'm *really* not going to give you money to work on this!² I just want you to think about what you actually need to get the project done. For most projects, this will be a statement saying that no financial resources are required for the project. For some, however, an actual budget spreadsheet will be appropriate. For example, if you are going to build a mobile robot, then you would list the supplies you will have to buy and their approximate cost. This will help you decide if you can afford to do the project.
5. Biographical sketches: In a real proposal, each investigator lists a brief form of their *curriculum vitae* (the academic version of a résumé. For your proposal, this is where you tell me who is working on the project and anything (briefly) you think I should know about them related to the project. For example, if you are doing something related to home automation and one of your group members has some expertise in that area, this is where you would tell me that.

Status reports

Each week, you will provide a *brief* status report for your project. This is due on Fridays and will be part of your journal. You should say what work was done during the week, if you are still on schedule, and any milestones you have met (or missed!).

Written report

This will be due at the end of the semester. An ideal format for this for most projects is to write it as a scientific paper, including an abstract, introduction, background/related work, methods, results,

¹Unless you need a mobile robot; we have some of those that you can use.

²Pretty much what NSF does for 90% of the projects proposed to it!

discussion, and bibliography. (This will obviously not be an appropriate format for all projects, for example, for a literature survey project.) In the “methods” section (which probably won’t be named that in most papers), you will discuss what you did and how you did it. For a programming (implementation) project, this will contain a description of the design as well as how your program was written, what AI techniques were used, the knowledge representation used (if applicable), etc. For a design project, this will just be a detailed description of the design. For a theoretical project, this might be where you give the details of a proof, for example, that you did. The results section will describe just that, results: Did your program run? What data do you have from running it? What errors occurred or unexpected things happened? etc. The discussion section is where you step back and do analysis of the results and give conclusions about your work.

Presentation/demonstration

I *may* require a public presentation of your project and/or a demonstration, depending on timing, room availability, etc. I will give you more information about this nearer the end of the semester.

Types of projects

There are five kinds of projects you can do: implementation, design, mixed, theoretical, and literature survey. An *implementation project* requires you to implement something, i.e., an intelligent artifact of some kind (software or software + hardware) and is pretty self-explanatory. It should be small enough to be implemented during the semester.

Design projects are free to be of larger scope than you could actually implement in a semester. For example, you might choose to design an intelligent house controller that learns your preferences and controls the household systems accordingly. It is unlikely that you could build a complete system in a semester, but you *could* design one. The design would need to be carefully worked out, including a specification of the architecture of the system, a detailed description of the AI techniques needed and how they would work, a detailed description of any knowledge required by the program, etc.

A *mixed project* is a hybrid of design + implementation. For example, you may not be able to build an entire intelligent house controller, but you might be able to build some portions of it as well as design the entire thing.

With respect to these three types of projects, I encourage you to make your projects as ambitious as possible. What you will be graded on is not so much whether you were able to solve the problem you selected, but rather your design, your implementation (if any), and the quality of your work overall. If you find that you have selected an implementation project that is too much to get done, don’t worry: talk to me about it, and you can probably switch it to a mixed or design project, even late in the semester.

A *theoretical project* would examine some aspect of AI theory or an AI technique itself in detail. For example, you might try to improve an existing search or constraint satisfaction technique, create a new kind of neural network, or develop a new kind of knowledge representation. Some such projects might require some programming, for example, to test the search technique or to implement and test the neural network; however, this would not be the main focus of the project. Other projects might require proving something; for example, you might prove the completeness and time complexity of a new type of search you’ve come up with.

A *literature survey project* requires you to search the scientific literature for papers on the chosen topic, create a synopsis of those papers, and write a thoughtful analysis of the papers in comparison to one another and as a group (e.g., what is the state of the art as represented by the papers?). *All of the papers must be AI papers*, and there should be enough of them to cover the most important approaches to the topic.

In general, I would not recommend this type of project, unless you are a grad student or Capstone student whose project needs such a literature review. (In which case, talk to me first.)

Sample projects

Intelligent robot control. The robot in question could be land-, air-, sea-, or space-based, or it could be a factory robot.

Example projects might be:

- Create an intelligent high-level planner to control the robot to carry out goals provided by the user.
- Implement a behavior-based controller (e.g., the subsumption architecture) for the robot.
- Path planning and following for the robot in a real environment.
- Homeostatic control for the robot: handling unanticipated events, keeping the robot “healthy”, etc.

This project could be done in simulation or using a real robot. My lab has some small mobile robots available for use, and there are many available for purchase online. (To give you some idea of the cost, the GoPiGo robots we are using in another class, which are Raspberry Pi-based and come with a camera and sonar, are < \$200, Pi included.)

Numerous simulators are available for mobile robots as well, for example, Player/Stage/Gazebo.

My lab works on intelligent robot control, so it is possible to work on one of our projects. See the MaineSAIL web page for information about our projects and talk to me if interested.

Intelligent house controller. A staple of science fiction is a house that is intelligent and can take care of its occupants. For this project, you would design such an agent, possibly implemented some portions and testing them in simulation. You would have to decide what you want the house to be able to do—e.g., observe the occupants to decide when to turn lights on/off, turn the heat up, fix breakfast, etc.—and then decide how you would get it to do it.

Deep learning. Build and test a deep-learning system for some task. For example, you might investigate a deep convolutional network or a long short-term memory network for image recognition/labeling, implement it using TensorFlow or Theano, then train and test it using one of the image repositories available on line. (Since this has already been done for image recognition, maybe you would apply deep learning to some other area of interest for you.)

Diagnosis. Given a set of symptoms, what is wrong? This could be diagnosing a disease or a fault in a piece of equipment. You could use rule-based reasoning, a Bayesian network, some sort of abductive reasoning, model-based reasoning, deep learning, etc.

Intelligent spam filter. You might use a Bayesian network, automated reasoning techniques, or deep learning, for example.

Cybersecurity. This could be knowledge-based or from deep learning, e.g.

Data mining. Finding patterns of interest in large datasets. A special type of this is scientific discovery: finding meaningful patterns in scientific data. For example, programs have been written that “discovered” Kepler’s laws of planetary motion from data similar to what Kepler had.

Physical security. This might entail sensors and/or a mobile robot watching site for intruders, fire, etc.

Multiagent system. This is a huge area of AI that involves getting multiple intelligent agents to work together, either to produce a more intelligent whole or to accomplish tasks a single agent cannot. Our lab works on projects in this area; see the [MaineSAIL web page](#) for information about our projects and talk to me if you are interested.

Intelligent chatbot. Write a chatbot that is more intelligent than the ones usually found on the Web.

Intelligent video game non-player agent. Build a better non-player agent for a video game.

Scheduling meetings. Scheduling is NP-hard; it has been approached in many ways, including constraint satisfaction. You might design and build a scheduling system for meetings, for example, that uses constraint-based techniques along with knowledge-based techniques, or you might apply deep learning to the problem.

Game playing. Game playing has a long history in AI, dating back to its very beginnings. You might build a chess-playing program, for example, or one to play some other game in which you are interested. (Games of chance are particularly difficult, although some good progress has recently been made in poker.) You could use traditional minimax-based approaches, maybe augmented with knowledge-based reasoning techniques, or you might try a deep-learning approach.

Creativity. Use a generative adversarial network (GAN) or other technique to create a system that does something creative. This could include creating artworks, texts, music, or really, anything you are interested in. See Prof. Sofian Audry’s website, [For example](#).

Important dates

Project assigned: 1/29/2019

Pre-proposal due: 2/12/2019

Full proposal due: 3/5/2010

Status reports due: Weekly on Fridays, starting 2/19/2019

Final project due: Last day of classes.

Presentation/demonstration (possibly): TBA