

CSC 2534 — Decision Making Under Uncertainty

Course Projects

Craig Boutilier

Fall 2014

A major component of CSC 2534 will be a course project. Projects can take one of the following four forms (or perhaps some suitable combination):

1. A research project: You should identify some relevant research problem and propose and investigate a solution. It's important to remember that this is only a course project, not a Master's thesis. Its scope should be circumscribed to reflect this fact. Also, the nature of "research" is such that success cannot be guaranteed (at least within some fixed amount of time). It is more important that you propose something plausible and show some reasonable progress. If the project is very ambitious, you shouldn't expect success within the time allotted for the course. But if you find it interesting, consider its potential for extension into a thesis topic.
2. A critical literature survey: You should pick a suitably circumscribed research topic and do an in-depth survey of literature in the area. You must demonstrate a good understanding of the area, and the ability to evaluate critically competing approaches to the problems in the area. Demonstrating this level of understanding should include identification of interesting directions for future research or open problems in the area surveyed, so a project of this type should include at least some broad discussion along these lines. A good selection of articles is crucial.
3. An implementation: You should choose an existing approach or system that tackles some problem you find interesting, and construct a prototype implementation. You will be required to critique your implementation, identify strengths and weaknesses, possible extensions, and difficulties that are impossible to overcome in your chosen framework. Keep it relatively simple though.
4. An experimental project: This will also require an implementation, but the aim is not to see if you can implement a particular theory; but rather to experiment with different approaches to the same problem. The goal will be to evaluate and compare the different approaches, not the implementation itself.

All projects must be approved in advance. You should submit a written project-proposal (length: one-half to one page) by November 4, 2014; but you are *strongly* encouraged to have a project in mind before then. If you submit a proposal by October 21, 2014, I will get you feedback by October 28, 2014. Projects will be due on some date to be determined (the latest date possible given whatever marks deadline exists). The key is to get started early. If you want to discuss a project topic with me before submitting your proposal, send me email and we'll arrange a time. Because of the number of students in the class, I'll likely set specific "project office hours" and ask you to sign up using Doodle for slots. Times TBA.

Because of the number of students registered in the class, project presentations are not feasible (due to the sheer amount of time required for everyone to do a 15 or 20 minute presentation). So we'll skip presentations unless enrolment changes significantly before the end of September.

One way to get ideas is to start looking through the course readings on certain topics. *Skim* through all the course material keeping an eye open for something that looks interesting to you. Even if it's just a general topic or area, it will provide you with pointers to the literature on related issues (so can I). For certain general course topics, the course syllabus lists a number of useful readings (texts and papers) you should look at early to see if you find this general area interesting: Also take a scan through other papers among the course readings.

One way to identify useful research projects is to start a very small literature survey in your area: this should point out lots of small problems of suitable scope. If you have a general area in mind, come see me and I'll point you in the right direction. If your research area is something other than reasoning (e.g., learning, vision, natural language, theory, graphics, verification, diagnosis, robotics, systems, etc.), there will no doubt be suitable projects that tie your primary interests into the topics covered in the course. Take advantage of that fact to do something that interests you.