

CPSC 522 — AI-2 — Reasoning and Acting Under Uncertainty

Course Projects

David Poole

February 2013

A major component of CPSC 522 will be a course project. Projects can take one of the following forms (or perhaps some suitable combination):

- 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 (or Ph.D.) thesis. It's 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; instead you need to solve a small part, and list the rest as future work (perhaps as a thesis topic).
- 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. A good selection of articles is crucial.
- An implementation: You should choose an existing approach or existing system described in a research paper that tackles some problem you find interesting, and construct a prototype implementation. The aim of the implementation is to test a hypothesis (e.g., the one in the paper); your implementation should be as minimal as possible. You will be required to critique the results of the research paper and your implementation, identify strengths and weaknesses, possible extensions, and difficulties that are impossible to overcome in your chosen framework. Are the results in the research paper able to be reproduced?
- An experimental project: This will also require an implementation, and the aim to experiment with different approaches to the same problem. The goal will be to evaluate and compare different approaches; the implementation will be a tool to achieve this goal.

Note that the goal of any implementation is to test some hypothesis. Either your own hypothesis for a research project, an existing hypothesis or a class of hypotheses for the experimental project. You need only build enough of the implementation to be able to test the hypothesis or hypotheses. You only need a non-trivial user interface if that is required to test you hypothesis (e.g., if the project is to identify a method of preference elicitation). Your implementation won't be evaluated; only the resulting paper.

Projects will be done individually. You are encouraged to collaborate with other students to find complementary synergistic projects (e.g., sharing the same code to test different hypotheses), but your written

project must be the your own. If you want to use the same work for more than one course, you need to get explicit permission from each instructor, and you will typically have to emphasize a different aspect of the work for each course. We treat plagiarism very seriously. You must reference any work you use or assistance you receive. If you want to use someone else's words, you should put them in quotations and attribute them properly.

You will be required to hand in a short proposal in the middle of the term (see schedule below), and then four copies of your project in the second last week of classes. You will be expected to peer review three other projects and give a conference-style presentation of your project. You will have a chance to revise your project based on the reviews.

The following are the deadlines:

- March 1: project proposals due (email David).
- April 5 11:00 am: projects due (electronically) and distributions of projects for review.
- Week of April 8-12: project presentations (exact times to be arranged).
- April 12 11:00 am: reviews finished and returned to authors.
- April 22: Final revised projects due.

All projects must be approved by David. A written (1 page) project-proposal is due on March 1; but you are *strongly* encouraged to have a project in mind before (and hopefully a proposal prepared) before then. The key is to get started early and come talk me about it before hand.

You will submit an electronic version of your project, and will be expected to review 3 projects. (I will review the other copy). Reviewers will be confidential (the reviewers will know the authors, but the authors will not know who reviewed your paper). I will know who reviewed what papers, and reviewing will be a component in the final grade. I will distribute the review form before the projects are due.

One way to get ideas is to start looking through the course readings on certain topics - don't wait until we get to the middle or end of the course, for instance, to look at articles on decision-theoretic planning, reinforcement learning and preference elicitation. *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, so make sure you talk with me *early and often*). For readings for the rest of the course look at the course schedule and <http://s.ubc.ca/~poole/cs522/2013/readings.html>, which will be continually updated.

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 to see me and I'll point you in the right direction.

See http://cs.ubc.ca/~poole/cs522/2013/project/project_ideas.pdf for some samples of possible projects. This is just to give you a sense of the expected size of the project, and possibly some ideas, not to limit the choice of topics. If you are interested, or just want to find out more about any of these (or related topics), ask me before or after class, or drop by my office and we can chat about it.