

# CPSC 522 — AI-II — Reasoning and Acting Under Uncertainty Course Projects

David Poole

February 2006

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. 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.
- 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, 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.

Note that the goal of any implementation is to test some hypothesis. Either your own hypothesis for a research project or an existing hypothesis. You need only build enough

of the implementation to be able to test the hypothesis. You only need, for example, a non-trivial user interface if that is required to test your hypothesis (e.g., if the project is to identify a method of preference elicitation). Your implementation won't be evaluated; only the resulting paper.

Although 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 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 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.

Keep in mind the following are the deadlines:

- February 23: project proposals due.
- April 3: projects due (4 copies of your project).
- April 4: distributions of projects for review.
- April 10-13: project presentations (exact times to be arranged).
- April 18: reviews finished and returned to authors.
- April 24: Final revised projects due.

All projects must be approved by David. A written (1 page) project-proposal is due on February 23; 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 to me about it before hand.

You will have to submit 4 copies of your project, and will be expected to review 3 projects. (I will review the other copy). Reviewers will be confidential (i.e., 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 the articles on planning or decision theory. *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 will I, so make sure you talk with me *early and often*). For a general overview of much of the second part of the course, look at Boutilier, Dean and Hanks "Decision Theoretic Planning: Structural Assumptions and Computational Leverage", JAIR, Vol 11, 1-94, 1999 (see <http://www.jair.org>), or Kaelbling, L.P., Littman, M.L., and Moore, A.W. (1996) "Reinforcement Learning: A Survey", JAIR, Volume 4, pages 237-285.

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. If your research area is something other than reasoning (e.g., vision, natural language, complexity theory, graphics, robotics, systems, etc.), there might be suitable projects that tie your primary interests into the topics for the course. Other related areas include verification and diagnosis.

## Possible Projects

Here are some representative 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.

### Belief Networks

- In-depth survey of recent methods for finding good elimination orderings (triangulations), perhaps with some empirical evaluation
- Implement a version of the variable elimination that uses structured conditional probability tables. How much time and space does it use compared to tabular method? (Is the overhead worthwhile?) Can you combine different methods for compact tables? For example, allowing procedural definitions of the conditional probabilities (e.g., as needed for the arithmetic problem of assignment 1)
- (Researchy) What is a good elimination ordering heuristic for contextual variable elimination, or one of the other methods whose inference time and space is not just measured by tree width.
- How can conditioning methods handle context-specific independence or causal independence efficiently?
- Investigate combining recursive conditioning (using and/or trees) with search methods that just enumerate the most likely worlds.
- Implement a stochastic simulation algorithm; experiment with some aspect of it (see also MDPs). Investigate how to combine exact inference with stochastic simulation.
- Investigate how different inference methods can be combined (e.g., search, stochastic simulation, variable elimination). Can you identify which aspects of a problem most effectively use each method?
- Develop a detailed belief net model of a problem in your area of interest. Why is it useful/necessary? Implement it. Does it actually work?

- Investigate the problem of existence uncertainty: what is the probability that an object with a certain description actually exists.
- Investigate the interaction of probabilities and ontologies (e.g., as in the Web Ontology Language, OWL). What are the possibilities for how these can be combined?
- Survey learning methods for belief nets. Compare some methods.
- (Researchy) Develop a novel method for representing CPTs that combine aspects of causal independence (e.g., noisy or) and context-specific independence.
- (Researchy) Suggest a way to handle structured variables, e.g., organized in an abstraction hierarchy. Or ways to handle semantic network type relationships with uncertainty (e.g., with structured models and descriptions).
- Suggest an extension and/or survey extensions of belief networks to handle some representational task, such as repeated structures, object-oriented belief networks or first-order representations. (First survey what exists first, then either concentrate on the survey or on the new method).
- Suggest a representation that may be a concise specification to represent uncertainty in some particular domain (e.g., intelligent tutoring systems, robotics); what is particular to this domain? Suggest how the conciseness can be exploited computationally. Test it, either by implementing some non-trivial aspects of the domain, or showing how well the algorithms work in practice.
- Investigate ways to reason about the probability of existence and identity (the probability that two descriptions refer to the same individual). Investigate how such reasoning can be represented in IBAL or some other framework.
- Carry out an evaluation of IBAL and/or ICL by using them for some non-trivial application. Compare them for the same domain; what are each one's strengths and weaknesses.

## **Actions**

- (Researchy) Develop a representation for factored actions (actions with multiple aspects or components).
- (Researchy) Develop a representation for actions with continuous parameters or effects (possibly describe how they could be used in certain planning or MDP algorithms).
- (Researchy) Consider how richer action representations could be used in planning, such as continuous time, concurrent actions, actions that take time to complete (such as filling a bath), hierarchical representations.

- (Researchy) Suggest how hierarchical control can be incorporated into decision-theoretic planning. For example, how can the continuous and discrete controls be combined?

## Preferences

- (Researchy) Develop some model of preference elicitation (interacting with a user to determine his/her preferences). For example, letting people specify their preferences in the real-estate domain, and inferring utilities to recommend new houses to look at.
- (Researchy) Develop some model for learning user preferences based on observation of their actions in specific situations.
- (Researchy) How can preferences of one user be inferred from comparing that person with other people whose partial preferences are known (one simple case of this called collaborative filtering for recommender systems)
- How can you learn preferences and explain to the user what model of preferences you are using?
- Given a set of preferences (e.g., as in a CP-net), and a set of constraints, give an algorithm to find a most-preferred outcome.
- How can we acquire the values of a user and build a user interface so that the computer does what the user wants. See e.g., International Conference on Intelligent User Interfaces: <http://www.iuiconf.org>.

## Decision making & reinforcement learning

- How can additive utility or multiplicative utility or other models of combining utilities be exploited in influence diagram evaluation?
- (Researchy) Extend structured model (e.g., Boutilier, Dearden, Goldszmidt 1995) for MDPs to deal with additive or multiplicative reward functions.
- (Researchy) Generalize influence diagrams to the multi-agent setting: develop an algorithm or investigate existing algorithms for equilibrium computation that exploits the influence diagram structure
- How can various extensions to belief network inference (e.g, context-specific independence, or causal independence) be exploited in decision-theoretic planning?
- (Researchy) Learning methods to compose "decomposed" value functions for an MDP with independent reward components
- Develop a detailed structured MDP model of a problem in your area of interest. (Why is it useful/necessary? Implement it perhaps.)

- Investigate how Bayes net inference can be used for reinforcement learning.
- Investigate ways to do exploration for reinforcement learning.
- Investigate how stochastic simulation (e.g., particle filtering) can be used to find good policies.
- (Researchy) Develop a simulation algorithm for policy evaluation. Evaluate it.
- Consider some aspect of multiple-agent MDPs.
- Develop a program for some aspect of the trading agents competition.

### **Classical, Conditional, Stochastic Planning**

- Research into possible representations of actions and observations suitable for a decision-theoretic planner (with an eye toward issues like the frame problem, etc.)
- Research into strategies for generating the "conditions" used in conditional planning. How can the uncertainty be incorporated? [Conditional planning is an extension of classical planning that allows incomplete knowledge, and usually information gathering actions such as "Determine whether door is locked". The future plan can be contingent on the "outcome" of such an "observation" (e.g., if locked, do nothing; if unlocked, go get keys, return, then lock the door). There has not been a lot of research (in the classical, deterministic setting as to when to insert such "branches" into plans in a sensible way.)
- Implement a conditional planner; or a simple stochastic planner.
- Research into (simple) planning strategies that are appropriate when utilities are used to judge plans rather than goal achievement (e.g., how could you take an existing planner and modify it's basic strategy to deal with utilities/preferences?)
- Investigate how you can incorporate individuals and relations into probabilistic reasoning, sequential decision making, or multi-agent systems.

....or anything else you may be interested in that is relevant to the course!