

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

Course Project Ideas

David Poole

February 2013

A major component of CPSC 522 will be a course project. See <http://cs.ubc.ca/~poole/cs522/2013/project/project.pdf> for details of the project.

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.

Relational Probabilistic Models

- Inference: how can we approximate inference in large models? How well does it work in practice?
- Learning: how can we effectively learn probabilities, e.g., for hidden variables?
- Structure Learning: how can we learn new rules?
- Inference: for some cases, we can infer the probability of hidden variables directly, and for some hidden variables not all data is relevant, how can we automatically determine which data is relevant to which variables?
- Representation: how are the probabilities affected by the population size? How can parameters be set so that they have desired dependence on population size?
- Lifted inference: how much can we do inference without grounding the representation?
- Implement and test some lifted inference algorithm. What are its limitations?

- Develop a model for some real-world problem. What prior knowledge is available? What data is available? How could it be evaluated?
- Investigate how a model can be learned from heterogeneous data sets.
- Investigate various aggregation schemes. What (desirable/undesirable) properties do they have? Which work better in practice?
- Knowledge acquisition for relational models: how can we make the models easily designable?
- Modularity: for various representations; What if you want to conjoin models that are developed independently? If there is a big model, can someone take a subset of the model and use it?
- Representation: how to build hierarchical representations, so that the general knowledge can be transferred to new domains?
- How can causal knowledge be incorporated into RPMs?
- How can ontologies be incorporated into RPMs?

Belief Networks

- 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.
- Investigate combining recursive conditioning (using and/or trees) with search methods that just enumerate the most likely worlds. How can you estimate the errors without searching all of the space?
- 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?
- Develop a novel method for representing CPTs that combine aspects of causal independence (e.g., noisy or) and context-specific independence.
- 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 ICL, 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

- Develop a representation for factored actions (actions with multiple aspects or components).
- Develop a representation for actions with continuous parameters or effects (possibly describe how they could be used in certain planning or MDP algorithms).
- 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.
- Suggest how hierarchical control can be incorporated into decision-theoretic planning. For example, how can the continuous and discrete controls be combined?

Preferences

- 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.
- Integration of preferences and ontologies: using the vocabulary of an ontology for specifying preferences.
- Develop some model for learning user preferences based on observation of their actions in specific situations.

- 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?
- 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 or MDP evaluation?
- How can various extensions to belief network inference (e.g, context-specific independence, or causal independence) be exploited in decision-theoretic planning?
- 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.
- Are there better representations of belief states that can be used for POMDPs? (Better than a belief space with dimension the number of world states.) What are representations that can effectively be used by an agent and can be optimized/learned?
- Research into strategies for generating the "conditions" used in conditional planning. How can the uncertainty be incorporated?
- 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.

Integrating Individuals, Relations, Ontologies and Uncertainty

- Investigate how you can incorporate individuals and relations into probabilistic reasoning, sequential decision making, or multi-agent systems.
- How to learn probabilistic models from data that are described using formal ontologies.
- How to handle the case where models refer to roles and the observed individuals do not have roles assigned. What is the difference between model individuals and observed individuals?
- Semantic science: making predictions about data that is described using rich ontologies. There is lots of work in making scientific data adhere to formal ontologies. There is much less work on allowing for hypotheses that make predictions on such data that also take into account the ontologies.
- A semantic science search engine. Given a hypothesis, how can we find data that is relevant to learn from? Given data, how can we find hypotheses that make predictions on that data?
- 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?
- Existence and identity uncertainty. Investigate a particular instance and show how such knowledge can be represented, learned and reasoned with.
- Rule-based learning and inference. How can rules be learned and how can we do inference efficiently?

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