James Lenfestey

6.834J

Problem Set 1

Part A

i. Reward-based planning under uncertainty. What are effective ways to plan given that state favorability is expressed with a known reward function but the world is not completely observable? I.e., how do we marginalize the PO in POMDP?

ii. Bayesian inference distributed over multiple agents. Given a collection of agents with partial knowledge of the world (derived from limited observations), how do we aggregate the information into a form suitable for statistical inference?

iii. Rule-based actions and logical description of environment. What are suitable ways of logically describing an environment and the robots actions? What are robust

Part D

Friedman et al., "Learning Probabilistic Relational Models",

http://citeseer.ist.psu.edu/cache/papers/cs/7400/http:zSzzSzrobotics.stanford.eduzSzuserszSzgetoorzSzpaperszSzlprm.pdf/friedman99learning.pdf

This paper provides an introduction to the concept of a Probabilistic Relational Model (PRM). As I mentioned in the previous part, PRMs are like Bayesian networks except that they admit a relational structure over the variables that are being described. Thus, we have objects with certain types and attributes, where the values a certain object's attribute assumes are correlated with other objects to which it is related. Again, the conditional probability relationships can be expressed in a graphical structure like a Bayesian network. This paper describes how to find the parameters of the model given a scheme and corpus of relational data. To do so, it builds on existing methods for training Bayesian networks.

This paper is important because it is widely cited and seems to have started the field. It is important for our purposes because it is necessary to understand PRMs before one can move on to DPRMs, the structure that we will be using to model state as it evolves through time.

Sanghai et al., "Dynamic Probabilistic Relational Models",

http://www.cs.washington.edu/homes/weld/papers/sanghai-ijcai03.pdf

In this paper DPRMs are defined. Put simply, DPRM is to PRM as a DBN is to a conventional Bayesian net. PRMs have schemas, descriptions of their relational structure (like in a database). An instantiation of a PRM is to fill in the schema with particular objects and values for its attributes. The idea of a DPRM is to describe, probabilistically, how an instantiation at a given instant in time evolves to the next time-step. Since an instantiation is basically a set of objects (and associated values) and instantiations change over time, this setup permits objects to enter or leave the world (or, our model of it).

The authors go on to describe a Monte-Carlo procedure for performing inference in their model based on the application of simplifying assumptions. The technique is based on the observation that DPRM can be unrolled into DBN, albeit a very large one. Rao-Blackwellising the relational attributes reduces the size of the state space, without which the procedure wouldn't work. The authors show, by experiment, that this method is able to perform inference accurately in domains consisting of over 1000 objects.

The significance of this paper is that it gives us a model that supports statistical inference of structured data over unbounded domains. For our purposes, this means that

that number of friendly robots about which we have to reason is immaterial, so long as they all fit the same relational description.

Doucet, et al. "Rao-Blackwellised Particle Filtering for Dynamic Bayesian Networks", http://www.cs.ubc.ca/~murphyk/Papers/rbpf_uai00.pdf

This article provides an account of the development of Particle Filtering (PF) techniques up to the use of Rao-Blackwellisation, an analytical technique used by Sanghai et al. to make a tractable sampling technique for DPRMs. In its own right, PF is a Monte-Carlo technique for doing inference in DBNs. They work by propagating a collection of sample points from one time-step to the next via sampling in accordance with the transition probabilities. The new points are then reweighted in accordance with the observations at the new times step. Finally, the points are resampled according to their weight, so that a greater number of samples lie in the region that is most probably the correct value of the hidden variable. This process is repeated until convergence.

The Rao-Blackwellised PF (RBPF) is an advance in that it uses analytic methods like the junction tree algorithm to marginalize away some of the hidden variables, thus reducing the size of the sample space and hence the number of particles necessary to perform accurate inference. The RBPF algorithm is described in this paper, which I will use in the implementation of DPRM inference.

Part E

I propose to simulate a collection of mapping robots that explore some virtual terrain. They will implement the reasoning capabilities described in Part C in the sense that they describe themselves, other robots, and the world around using a DPRM. I will simulate noisy and unreliable sensors, thus forcing the robot to consult its failure models.

To preserve salient features of the multi-agent model, the ability for two robots to communicate, thus exchanging map data, will be left intact. Further, each robot will maintain estimates of its peers' regions of expertise.

Although a fair amount of intelligence can be imbued into the executive, at present it is not the aim of this project. Instead, it will likely be hard coded with what seems to be desirable behavior, nonetheless relying heavily on the robot's estimates of world state.