UIUC, CS498, Section EA - Autumn 04 - Homework #4

Due Date: Tuesday, December 16, 2004, 12pm

December 2, 2004

1 Situation Calculus (30%)

A robot is expected to deliver mail to different offices in the 3rd floor of the Siebel Center. It can use the actions *turn*(α),*go-fwd(d)*,*give(p)*, for angle α, distance d, and paper p. The building is modeled with a qualitative map (a planar graph) that connects the rooms with their entrances (in a corridor or another room), with angles and distances associated with edges and nodes in the graph. The robot can travel only along the edges of this graph.

1. Represent the problem above in Situation Calculus. Write down successor-state axioms assuming that actions are deterministic.
2. Use a theorem prover of your choice (e.g., the one you used for HW#2) to find a plan that will deliver two papers to two different rooms.
3. Download GOLOG (http://www.cs.toronto.edu/cogrobo/cogrobo.1/systems.html), and use it and your Situation Calculus formulation to find a plan for delivering 20 papers to 10 rooms.
4. Now assume that the actions are stochastic with a Gaussian distribution (for turning and going forward) or a binomial distribution of failure (for giving the paper). Represent this additional information using nature actions (see [Boutilier etal. '03][1]).

2 Factored Dynamic Bayesian Networks (25%)

1. Prove contraction for transition model T with respect to distributions $\sigma^{(t\bullet)}, \hat{\sigma}^{(t\bullet)}$

$$D[T[\sigma^{(t\bullet)}]||T[\hat{\sigma}^{(t\bullet)}]] \leq D[\sigma^{(t\bullet)}||\hat{\sigma}^{(t\bullet)}]$$

2. Prove the single process contraction theorem (Theorem 3 in [Boyen & Koller '98]): For process Q, anterior (prior to transition) distributions φ, ψ, ultirior distributions φ', ψ', and minimal mixing rate γ_Q,

$$D[\varphi'||\psi'] \leq (1 - \gamma_Q)D[\varphi||\psi]$$

3. Give a DBN algorithm that finds the most likely explanation for observations, and performs the same factoring as done for filtering (introducing an approximation error at every stage). What bound do we get for the error at every step? What bound do we get for the overall process?
4. Represent your robot example from question 1.4 in a DBN. Use the Matlab BN package of HW#3 to find the probability that all papers are delivered at the end of the sequence you found in question 1.3.

[1] This paper was given in class, and can be downloaded from the class syllabus web page.

3 Particle Filtering (20%)

1. Particle filtering is a method by which we are given n independent samples from distribution σ_t, compute their probability after transition with T, update with observations (observation model O), form a new distribution $\hat{\sigma}_{t+1}$ and re-sample n independent samples from this distribution. Prove that $\hat{\sigma}_{t+1} \rightarrow \sigma_{t+1}$ when $n \rightarrow \infty$, for $\sigma_{t+1} = O_h(T(\sigma_t))$.
2. Derive an MCMC-based (e.g., Gibbs) algorithm for filtering observations in a DBN of t time steps. Give a proof that your algorithm converges to the correct distribution.
3. Give an example in which your algorithm is better (converges faster) than particle filtering. Give an example in which particle filtering is better (converges faster) than your algorithm.

4 Logical Filtering (25%)

1. Generalize the NNF filtering algorithm to multi-valued fluents.
2. Prove that your algorithm is correct for actions that map states 1:1 (you can use the proofs in the paper on the syllabus as guidance).