

Homework 3 (on Belief-space planning)
Due 10th November, 2010

Qn I. [Conformant planning]

Consider a domain with three literals, P, Q and R.

Suppose in the initial state, we only know that either P or Q is true (not both). Our goal is to make Q true.

Consider the following three actions with conditional effects:

A1: $P \Rightarrow [Q \vee R]$

A11: $P \Rightarrow [Q \vee \sim R]$

A2: $R \Rightarrow Q$

[Assume that the disjunction is exclusive--that is the effect $Q \vee R$ means either Q or R will be true--not both]

1. What is the size of the belief space for this problem? (i.e., number of belief states).
2. Show the initial belief state corresponding to this problem.
3. Show the set of actions that can be applied to this belief state and the resulting belief states.
4. Show one path (actions, and belief states) that lead to the goal.
5. Sketch the Labeled uncertainty graph for this problem, illustrating how labels propagate (the graph doesn't have to be complete--as long as you establish that you know how the technique works..)

Qn II. [Conditional Planning]

To the problem in III, suppose we added three sensing actions:

P?, Q?, R?

Each of which will tell us whether the corresponding literal is true or false.

We also have another effect for A1, which involves deleting Q, if it were true. So A1 is now:

A1: $P \Rightarrow [Q \vee R] ; Q \Rightarrow \sim Q$

The initial state and the goal state, as well as the rest of the actions, remain the same.

1. What is the size of the belief space now for this problem?
**What are the number of observation classes?
2. Show the complete set of actions that can be applied to the initial belief state, and for each one, the belief states that results

3. Show one path (actions, belief states) that lead to the goal.

Qn III. POMDPs

Consider the 2-state POMDP problem in Russell&Norvig. Suppose the initial belief state is $[.5, .5]$. Draw a 1-ply online look-ahead tree for this problem. The tree should show the two actions that can be done [stay] and [go], and for each action the observation [0] and [1]. You should show the belief state after just the action, and the belief state after the observation.

For the tree above, suppose you use the FOMDP approximation to compute the upper-bounds for the states. Show what the values of the leaf belief states are, and show how they are backed up to compute the value of the remaining states.