Hidden Markov Models and Decision Theory Intro

CPSC 322 Lecture 30

March 28, 2007
Textbook §9.5

Lecture Overview

Recap

Hidden Markov Models

Decision Theory Intro

Variable elimination algorithm

To compute $P(Q|Y_1 = v_1 \wedge \ldots \wedge Y_j = v_j)$:

- Construct a factor for each conditional probability.
- Set the observed variables to their observed values.
- For each of the other variables $Z_i \in \{Z_1, \ldots, Z_k\}$, sum out Z_i
- Multiply the remaining factors.
- Normalize by dividing the resulting factor $f(Q)$ by $\sum_Q f(Q)$.

One Last Trick

One last trick to simplify calculations: we can repeatedly eliminate all leaf nodes that are neither observed nor queried, until we reach a fixed point.

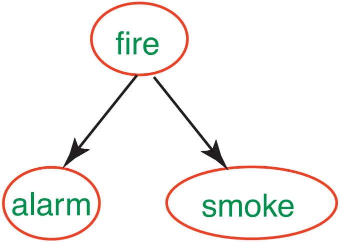

Can we justify that on a three-node graph—Fire, Alarm, and Smoke—when we ask for:

- $P(Fire)$?
- $P(Fire \mid Alarm)$?

Lecture Overview

Recap

Hidden Markov Models

Decision Theory Intro

Markov chain

- A Markov chain is a special sort of belief network:

- Thus $P(S_{t+1}|S_0,\ldots,S_t) = P(S_{t+1}|S_t)$.
- Often S_t represents the state at time t. Intuitively S_t conveys all of the information about the history that can affect the future states.
- "The past is independent of the future given the present."

Stationary Markov chain

- A stationary Markov chain is when for all $t > 0$, $t' > 0$, $P(S_{t+1}|S_t) = P(S_{t'+1}|S_{t'})$.
- We specify $P(S_0)$ and $P(S_{t+1}|S_t)$.
  - Simple model, easy to specify
  - Often the natural model
  - The network can extend indefinitely

Hidden Markov Model

- A Hidden Markov Model (HMM) starts with a Markov chain, and adds a noisy observation about the state at each time step:

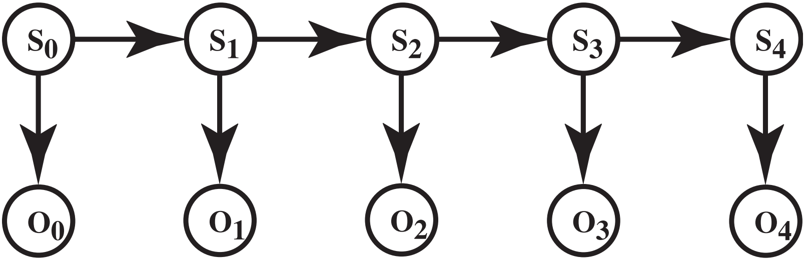

- $P(S_0)$ specifies initial conditions
- $P(S_{t+1}|S_t)$ specifies the dynamics
- $P(O_t|S_t)$ specifies the sensor model

Example: localization

- Suppose a robot wants to determine its location based on its actions and its sensor readings: Localization
- This can be represented by the augmented HMM:

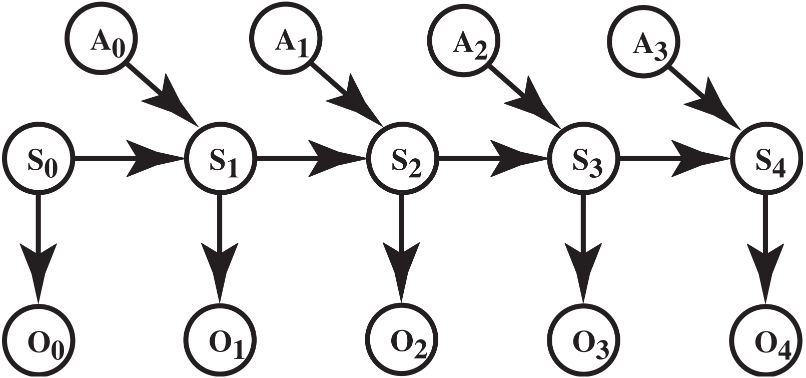

Example localization domain

- Circular corridor, with 16 locations:

- Doors at positions: 2, 4, 7, 11.
- Noisy Sensors
- Stochastic Dynamics
- Robot starts at an unknown location and must determine where it is.

Example Sensor Model

- $P(Observe\ Door \mid At\ Door) = 0.8$
- $P(Observe\ Door \mid Not\ At\ Door) = 0.1$

Example Dynamics Model

- $P(loc_{t+1} = L | action_t = goRight \wedge loc_t = L) = 0.1$
- $P(loc_{t+1} = L + 1 | action_t = goRight \wedge loc_t = L) = 0.8$
- $P(loc_{t+1} = L + 2 | action_t = goRight \wedge loc_t = L) = 0.074$
- $P(loc_{t+1} = L' | action_t = goRight \wedge loc_t = L) = 0.002$ for any other location L'.
  - All location arithmetic is modulo 16.
  - The action $goLeft$ works the same but to the left.

Combining sensor information

- Example: we can combine information from a light sensor and the door sensor: "Sensor Fusion"

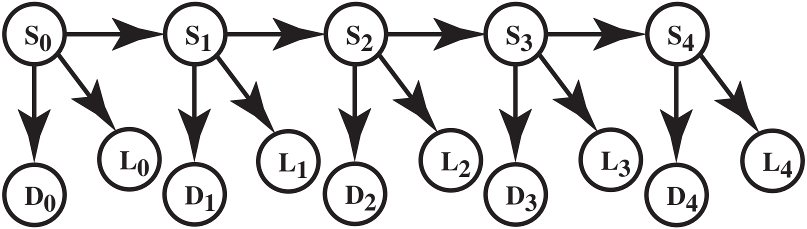

- S_t: robot location at time t
- D_t: door sensor value at time t
- L_t: light sensor value at time t

Localization demo

- http://www.cs.ubc.ca/spider/poole/demos/localization/localization.html

Lecture Overview

Recap

Hidden Markov Models

Decision Theory Intro

Decisions Under Uncertainty

- In the first part of the course we focused on decision making in domains where the environment was understood with certainty
  - Search/CSPs: single decisions
  - Planning: sequential decisions

- In uncertain domains, we've so far only considered how to represent and update beliefs
- What if an agent has to make decisions in a domain that involves uncertainty?
  - this is likely: one of the main reasons to represent the world probabilistically is to be able to use these beliefs as the basis for making decisions

Decisions Under Uncertainty

- An agent's decision will depend on:
  1. what actions are available
  2. what beliefs the agent has
     - note: this replaces "state" from the deterministic setting
  3. the agent's goals

- We've spoken quite a lot about (1) and (2).
  - today let's consider (3)
  - we'll move from all-or-nothing goals to a richer notion: rating how happy the agent is in different situations