Markov chain

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

What probabilities need to be specified? What Independence assumptions are made?

Markov chain

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

- $P(S_0)$ specifies initial conditions
- $P(S_{t+1}|S_t)$ specifies the dynamics
- $P(S_{t+1}|S_0, \ldots, S_t) = P(S_{t+1}|S_t)$.
- 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 future is independent of the past 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) is a belief network:

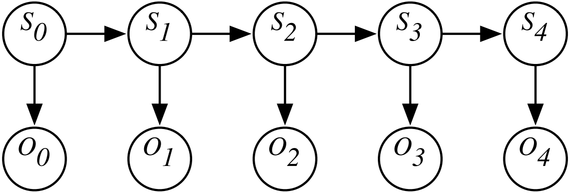

The probabilities that need to be specified:

Hidden Markov Model

- A Hidden Markov Model (HMM) is a belief network:

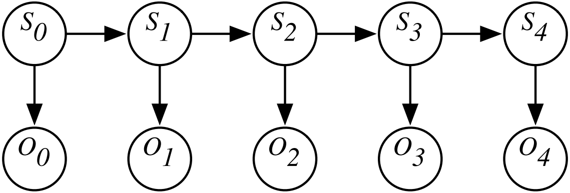

The probabilities that need to be specified:

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

Filtering

Filtering:

$$P(S_i | o_1, \ldots, o_i)$$

What is the current belief state based on the observation history?

Filtering

Filtering:

$$P(S_i | o_1, \ldots, o_i)$$

What is the current belief state based on the observation history?

$$\begin{aligned} P(S_i | o_1, \ldots, o_i) &\propto P(o_i | S_i o_1, \ldots, o_{i-1}) P(S_i | o_1, \ldots, o_{i-1}) \\ &= ??? \sum_{S_{i-1}} P(S_i S_{i-1} | o_1, \ldots, o_{i-1}) \\ &= ??? \end{aligned}$$

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:

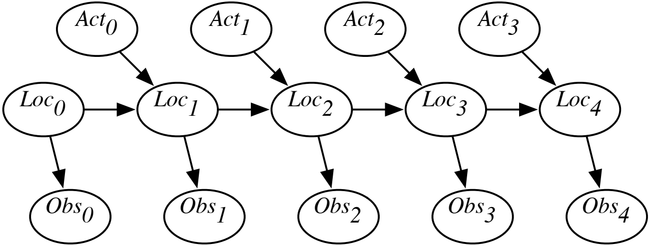

Example localization domain

- Circular corridor, with 16 locations:

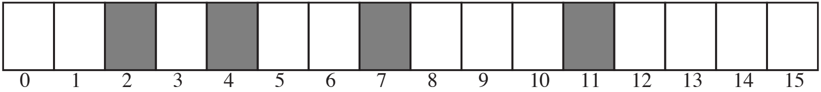

- Doors at positions: 2, 4, 7, 11.
- Noisy Sensors
- Stochastic Dynamics
- Robot starts at an unknown location and must determine where it is.

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

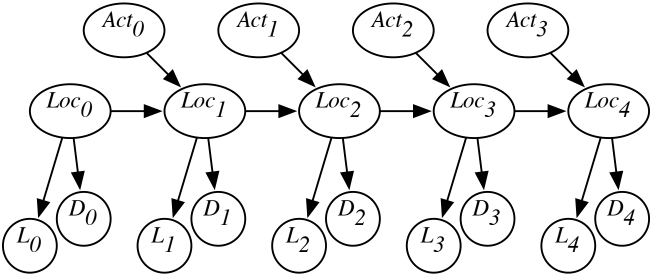

S_t robot location at time t
D_t door sensor value at time t
L_t light sensor value at time t