Probability and Time: Hidden Markov Models (HMMs)

CPSC 322 Lecture 31

Lecture Overview

- **Recap**
- Markov Models
  - Markov Chain
  - Hidden Markov Models

Stationary Markov Chain (SMC)

A **stationary Markov Chain** : for all t >0

- $P(S_{t+1} | S_0,\ldots,S_t)$ = ________ and
- $P(S_{t+1} | S_t)$ is ________ for all t

We only need to specify ______ and __________

- Simple Model, easy to specify
- Often the natural model
- The network can extend indefinitely
- **Variations of SMC are at the core of most Natural Language Processing (NLP) applications!**

Learning Goals for today's class

You can:

- Specify the components of an Hidden Markov Model (HMM)
- Justify and apply HMMs to Robot Localization

Lecture Overview

- **Recap**
- Markov Models
  - **Markov Chain**
  - **Hidden Markov Models**

How do we minimally extend Markov Chains?

- **Maintaining** the **Markov** and **stationary** assumptions?

A useful situation to model is the one in which:

- the reasoning system **does not have access** to the **states**
- but can **make observations** that give some information about the current state

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:

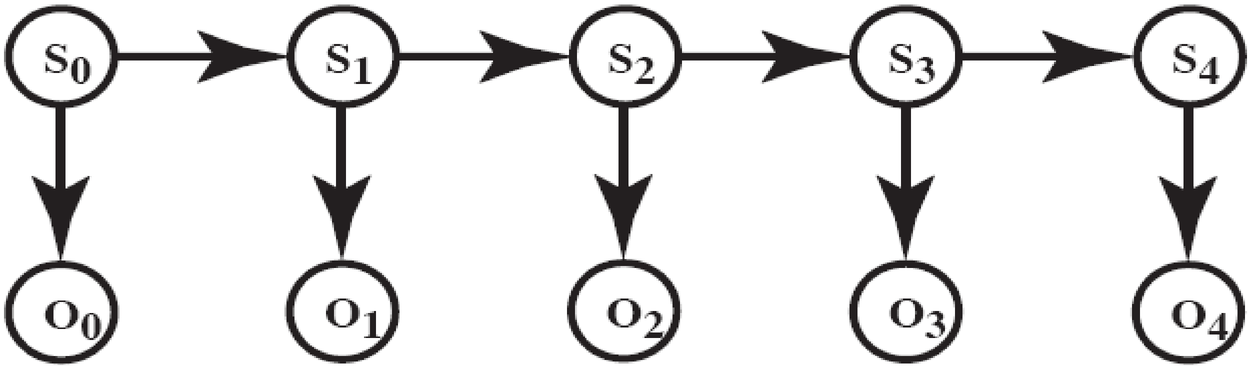

- $|\text{domain}(S)| = k$
- $|\text{domain}(O)| = h$

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

iclicker.

A. k

B. $h \times h$

C . $k \times h$

D. $k \times k$

Example: Localization for "Pushed around" Robot

- **Localization** (where am I?) is a fundamental problem in **robotics**
- Suppose a robot is in a circular corridor with 16 locations

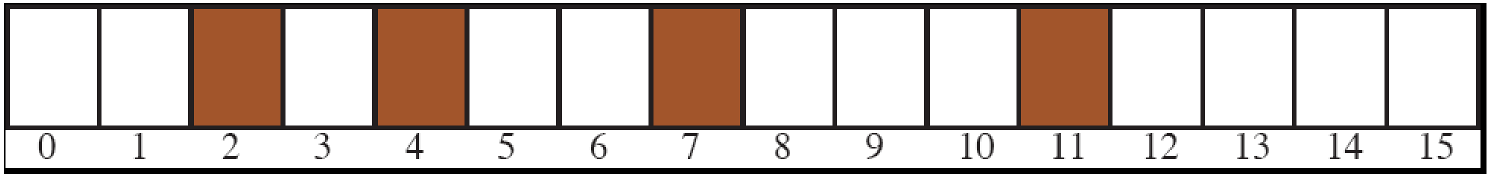

  - There are **four doors** at positions: 2, 4, 7, 11
  - The Robot initially doesn't know where it is
  - The Robot is **pushed around**. After a push it can stay in the same location, move left or right.
  - The Robot has a **Noisy sensor** telling whether it is in front of a door

This scenario can be represented as…

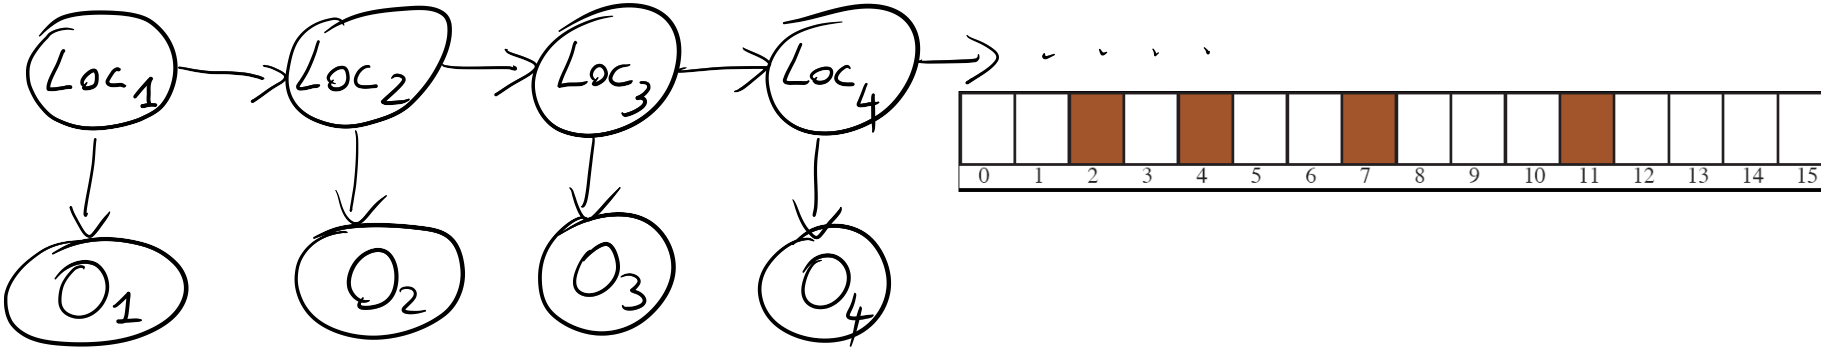

- **Example Stochastic Dynamics**: when pushed, it stays in the same location p=0.2, moves one step left or right with equal probability

$Loc_t = 10$

$P(Loc_{t+1} \mid Loc_t)$

iclicker.

A.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	.2	0	.4	0	.2	0	0	0

B.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	0	.4	.2	.4	0	0	0	0

C.

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0	0	0	0	0	0	0	0	.2	.2	.2	.2	.2	0	0	0

This scenario can be represented as…

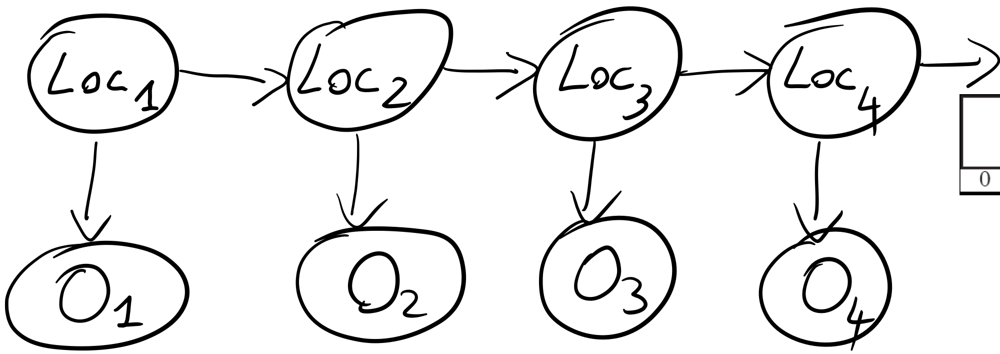

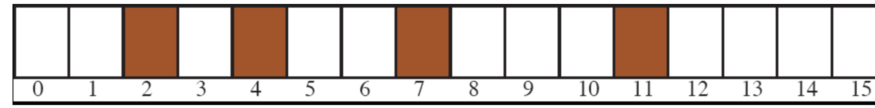

- **Example Stochastic Dynamics**: when pushed, it stays in the same location p=0.2, moves left or right with equal probability

$P(Loc_{t+1} \mid Loc_t)$

	0	1	2	…	15
0	.2	.4	0	…	.4
1	.4	.2	.4	…	0
2	0	.4	.2	…	0
…	…	…	…	…	…
15	.4	0	0	.4	.2

$P(Loc_1) =$

0	1	2	…	15
1/16	1/16	1/16	…	1/16

This scenario can be represented as…

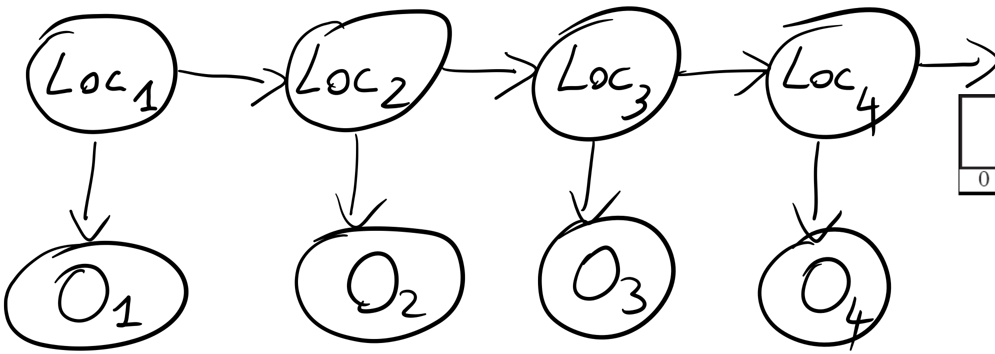

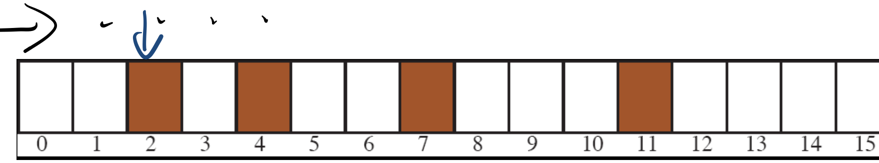

$dom(Loc_i) = \{0, 1, \ldots, 15\}$

Example of **noisy sensor** telling whether it is in front of a door.

- If it is in front of a door $P(O_t = T) = .8$
- If not in front of a door $P(O_t = T) = .1$

$P(O_t \mid Loc_t)$

Loc_t	$P(O_t=T)$	$P(O_t=F)$
0	.1	.9
1	.1	.9
2	.8	.2
3	.1	.9
4	.8	.2
…	…	…

Useful inference in HMMs

- **Localization**: Robot starts at an unknown location and it is pushed around *t* times. It wants to determine where it is

$$P(Loc_t \mid o_1 \ldots\ldots o_t)$$

- **In general:** compute the posterior distribution over **the current state** given **all evidence to date**

$$P(S_t \mid O_0 \ldots O_t)$$

Example : Robot Localization

- Suppose a robot wants to determine its location based on its actions and its sensor readings
- Three actions: *goRight, goLeft, Stay*
- This can be represented by an **augmented HMM**

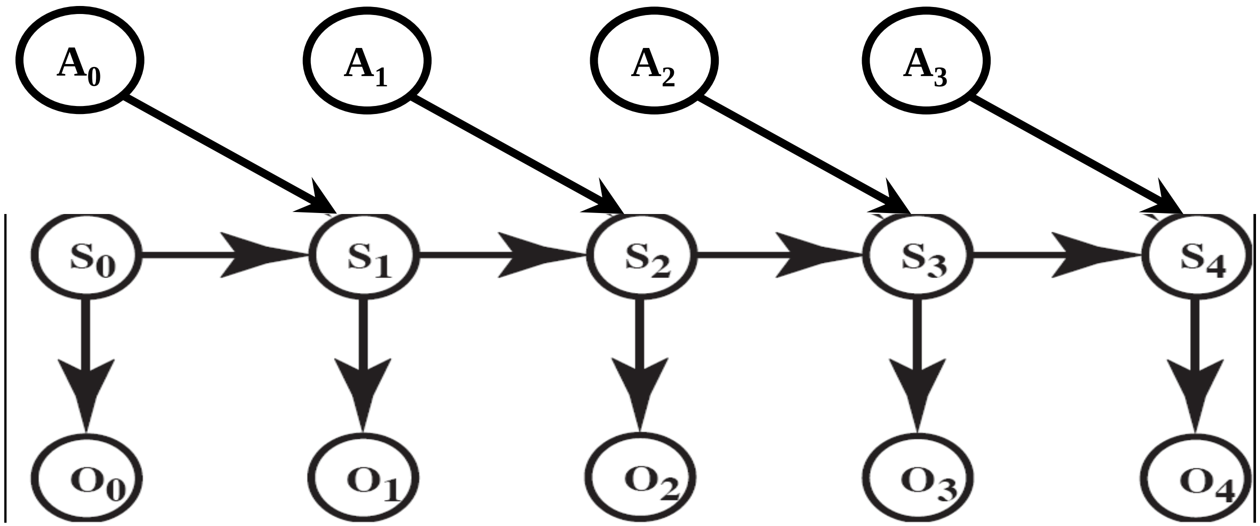

Robot Localization Sensor and Dynamics Model

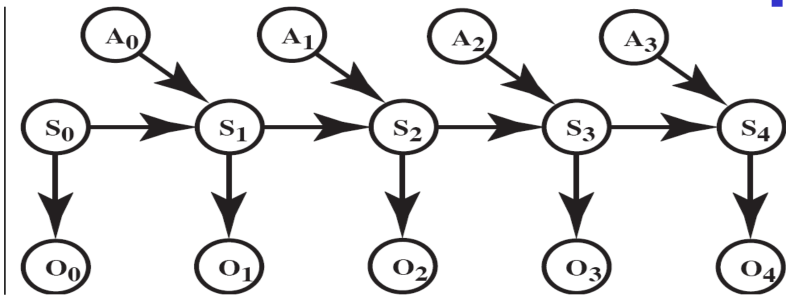

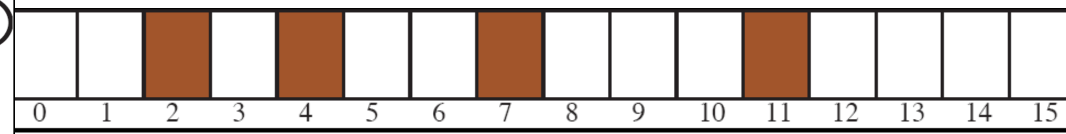

- **Sample Sensor Model** (assume same as for pushed around)
- **Sample Stochastic Dynamics**: $P(Loc_{t+1} \mid Action_t , Loc_t)$

$$P(Loc_{t+1} = L \mid Action_t = goRight , Loc_t = L) = 0.1$$

$$P(Loc_{t+1} = L+1 \mid Action_t = goRight , Loc_t = L) = 0.8$$

$$P(Loc_{t+1} = L + 2 \mid Action_t = goRight , Loc_t = L) = 0.074$$

$$P(Loc_{t+1} = L' \mid Action_t = goRight , Loc_t = L) = 0.002 \text{ for all other locations } L'$$

- All location arithmetic is modulo 16
- The action *goLeft* works the same but to the left

Dynamics Model - More Details

- **Sample Stochastic Dynamics**: $P(Loc_{t+1} \mid Action, Loc_t)$

$P(Loc_{t+1} = L \mid Action_t = goRight, Loc_t = L) = 0.1$

$P(Loc_{t+1} = L+1 \mid Action_t = goRight, Loc_t = L) = 0.8$

$P(Loc_{t+1} = L+2 \mid Action_t = goRight, Loc_t = L) = 0.074$

$P(Loc_{t+1} = L' \mid Action_t = goRight, Loc_t = L) = 0.002$ *for all other locations L'*

goRight

	0	1	2	3	…	15
0	.1	.8	.074	.002	…	.002
1	.002	.1	.8	.074	…	.002
2	.002	.002	.1	.8	…	.002
3	.002	.002	.002	.1	…	.002
…	…	…	…	…	…	…
15	.8	.074	.002	.002	…	.1

goLeft

	0	1	2	3	…	15
0	.1	.002	.002	.002	…	.8
1	.8	.1	.002	.002	…	.074
2	.074	.8	.1		…	.002
3	.002	.074	.8	.1	…	.002
…	…	…	…	…	…	…
15	.002	.002	.002	.002	…	.1

stay (deterministic)

	0	1	2	3	…	15
0	1				…	
1		1			…	
2			1		…	
3				1	…	
…	…	…	…	…	…	…
15					…	1

Robot Localization additional sensor

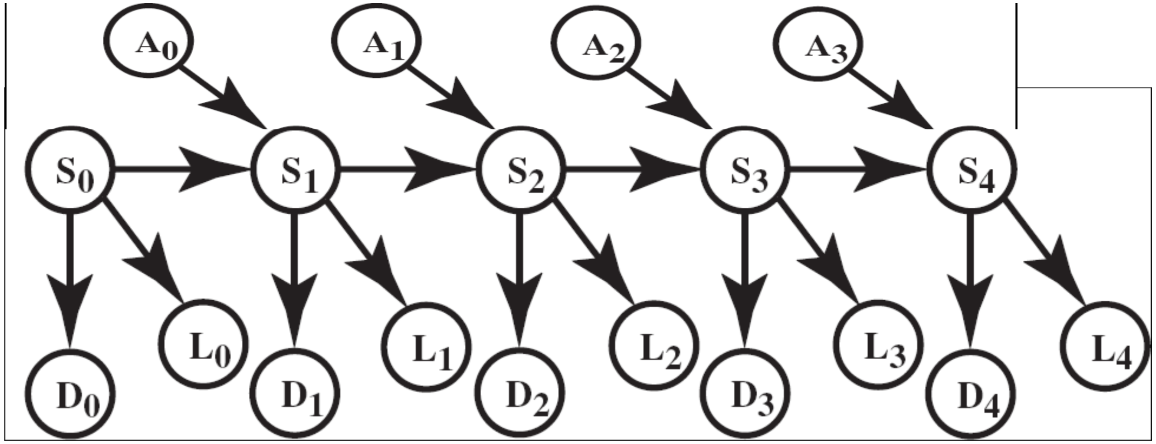

$L_t = T$

the Robot senses light

- **Additional Light Sensor:** there is light coming through an opening at location 10

$P(L_t \mid Loc_t)$

$P(L_t=F)$ — — — — — . .2 .05 .01 .05 .2 .4 . . . .

$P(L_t=T)$. — — — . . .8 .95 .99 .95 .8 .6 — —

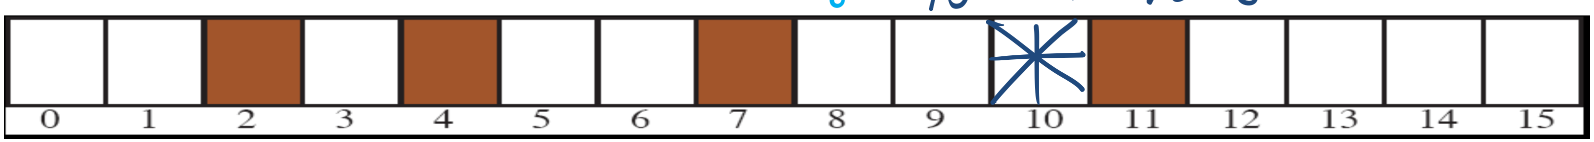

- **Info from the two sensors is combined :"Sensor Fusion"**

The Robot starts at an unknown location and must determine where it is

The model appears to be too ambiguous

- Sensors are too noisy
- Dynamics are too stochastic to infer anything

But inference actually works pretty well.

You can check it at :

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

You can use standard BNet inference. However you typically take advantage of the fact that time moves forward (not in 322)

Sample scenario to explore in demo

- Keep making observations without moving. What happens?
- Then keep moving without making observations. What happens?
- Assume you are at a certain position alternate moves and observations
- ….

HMMs have many other applications….

Natural Language Processing: e.g., Speech Recognition

- *States:* phoneme \ word
- *Observations*: acoustic signal \ phoneme

HMM 1 HMM 2

My name is

Bioinformatics: Gene Finding

- *States:* coding / non-coding region **--YYY--**
- *Observations:* DNA Sequences **ATCGGAA**

For these problems the critical inference is:

find the most likely sequence of states given a sequence of observations **(Viterbi algorithm, CPSC 422)**

Markov Models

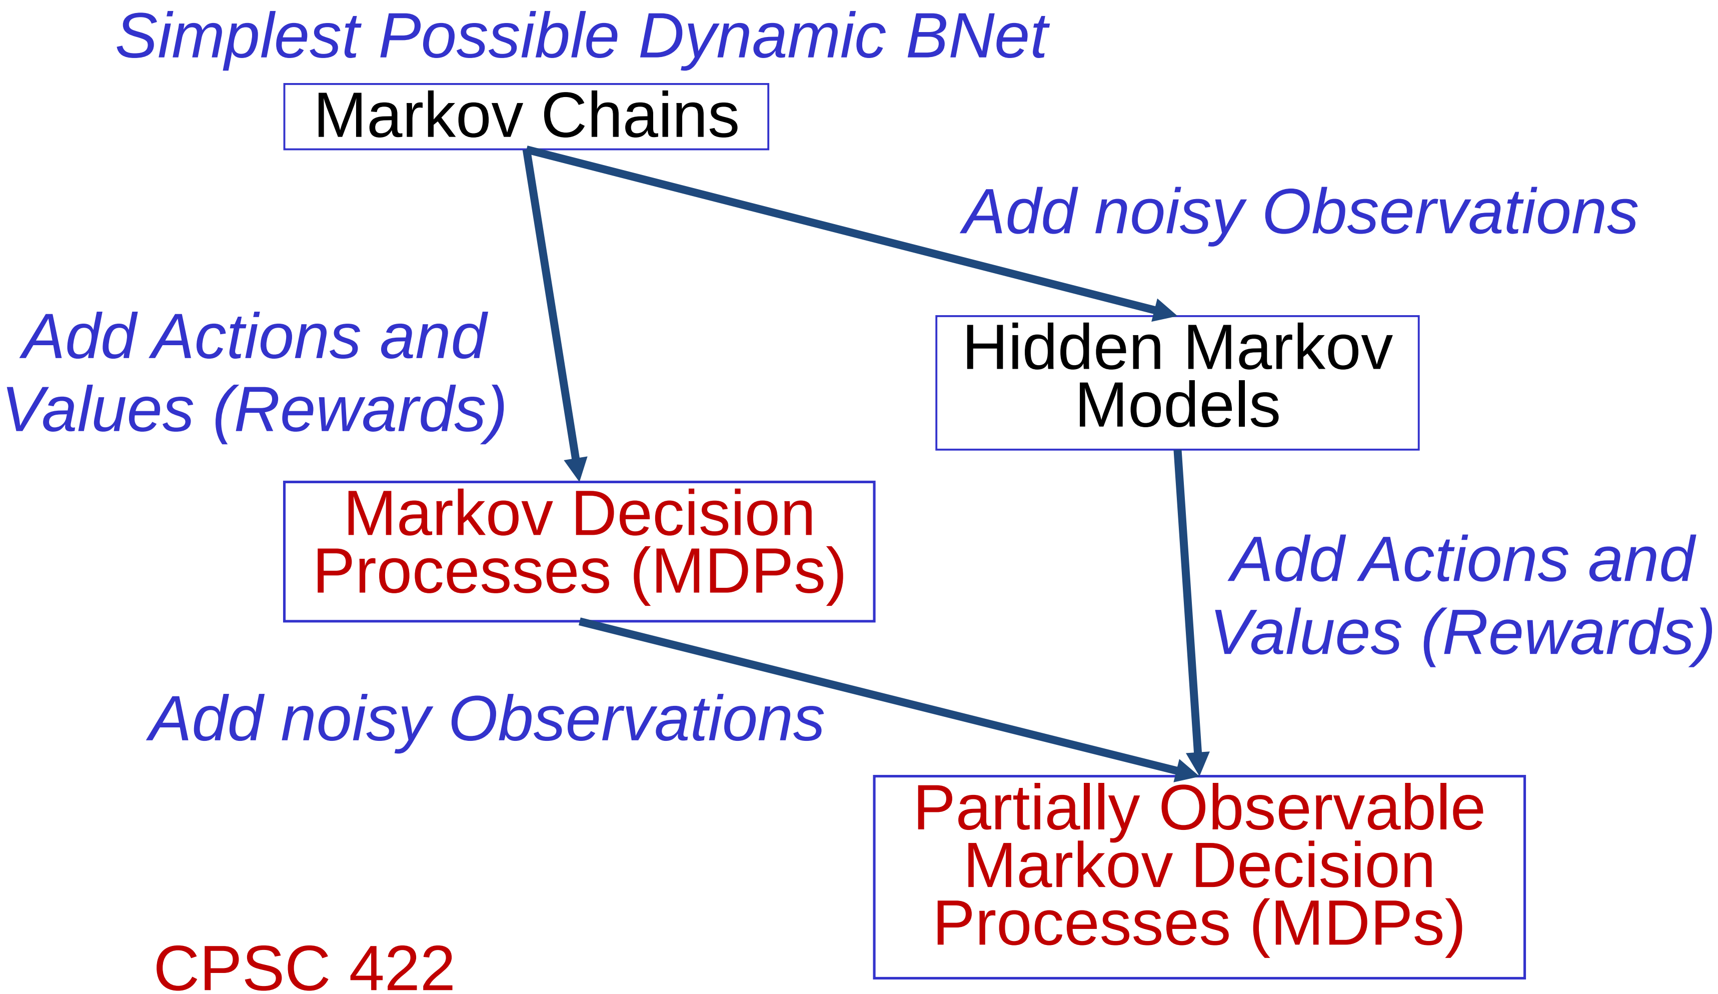

Next Class

- **One-off decisions***(Textbook 9.2)*
- **Single Stage Decision networks** *(9.2.1)*