CSE 473: Introduction to Artificial Intelligence

Hanna Hajishirzi HMMs Inference, Particle Filters

slides adapted from Dan Klein, Pieter Abbeel ai.berkeley.edu And Dan Weld, Luke Zettelmoyer



Recap: Reasoning Over Time

- Markov models $x_1 \rightarrow x_2 \rightarrow x_3 \rightarrow x_4 \rightarrow \cdots \rightarrow P(X_1) \qquad P(X|X_{-1})$
- Hidden Markov models





P(E|X)

Х	E	Р
rain	umbrella	0.9
rain	no umbrella	
sun	umbrella	0.2
sun	no umbrella	0.8



Sensor model: can read in which directions there is a wall, never more than 1 mistake

Motion model: may not execute action with small prob.



required 1 mistake





















Inference: Find State Given Evidence

We are given evidence at each time and want to know

$$B_t(X) = P(X_t | e_{1:t})$$

- Idea: start with P(X₁) and derive B_t in terms of B_{t-1}
 - equivalently, derive B_{t+1} in terms of B_t

Inference: Base Cases





 $P(X_2)$

Inference: Base Cases



 $P(X_2)$

$$P(x_2) = \sum_{x_1} P(x_1, x_2)$$
$$= \sum_{x_1} P(x_1) P(x_2 | x_1)$$

Passage of Time

Assume we have current belief P(X | evidence to date)

 $B(X_t) = P(X_t | e_{1:t})$

Then, after one time step passes:

$$P(X_{t+1}|e_{1:t}) = \sum_{x_t} P(X_{t+1}, x_t|e_{1:t})$$

= $\sum_{x_t} P(X_{t+1}|x_t, e_{1:t}) P(x_t|e_{1:t})$
= $\sum_{x_t} P(X_{t+1}|x_t) P(x_t|e_{1:t})$



• Or compactly:

$$B'(X_{t+1}) = \sum_{x_t} P(X'|x_t) B(x_t)$$

- Basic idea: beliefs get "pushed" through the transitions
 - With the "B" notation, we have to be careful about what time step t the belief is about, and what evidence it includes

Example: Passage of Time



As time passes, uncertainty "accumulates"

T = 2









Inference: Base Cases



$P(X_{1}|e_{1})$ $P(x_{1}|e_{1}) = P(x_{1},e_{1})/P(e_{1})$ $\propto_{X_{1}} P(x_{1},e_{1})$ $= P(x_{1})P(e_{1}|x_{1})$

Observation

Assume we have current belief P(X | previous evidence):

 $B'(X_{t+1}) = P(X_{t+1}|e_{1:t})$

• Then, after evidence comes in:



$$P(X_{t+1}|e_{1:t+1}) = P(X_{t+1}, e_{t+1}|e_{1:t}) / P(e_{t+1}|e_{1:t})$$

$$\propto_{X_{t+1}} P(X_{t+1}, e_{t+1}|e_{1:t})$$

 $= P(e_{t+1}|e_{1:t}, X_{t+1})P(X_{t+1}|e_{1:t})$

$$= P(e_{t+1}|X_{t+1})P(X_{t+1}|e_{1:t})$$

• Or, compactly:

 $B(X_{t+1}) \propto_{X_{t+1}} P(e_{t+1}|X_{t+1})B'(X_{t+1})$

- Basic idea: beliefs "reweighted" by likelihood of evidence
- Unlike passage of time, we have to renormalize

Example: Observation

As we get observations, beliefs get reweighted, uncertainty "decreases"

0.05	0.01	0.05	<0.01	<0.01	<0.01
0.02	0.14	0.11	0.35	<0.01	<0.01
0.07	0.03	0.05	<0.01	0.03	<0.01
0.03	0.03	<0.01	<0.01	<0.01	<0.01

Before observation

<0.01	<0.01	<0.01	<0.01	0.02	<0.01
<0.01	<0.01	<0.01	0.83	0.02	<0.01
<0.01	<0.01	0.11	<0.01	<0.01	<0.01
<0.01	<0.01	<0.01	<0.01	<0.01	<0.01

After observation





 $B(X) \propto P(e|X)B'(X)$



Filtering: P(X_t | evidence_{1:t})

Elapse time: compute P(X_t | e_{1:t-1})

$$P(x_t | e_{1:t-1}) = \sum_{x_{t-1}} P(x_{t-1} | e_{1:t-1}) \cdot P(x_t | x_{t-1})$$
Observe: compute P(X_t | e_{1:t})

$$P(x_t | e_{1:t}) \propto P(x_t | e_{1:t-1}) \cdot P(e_t | x_t)$$



 X_2

 E_2



Example: Weather HMM



R _t	R _{t+1}	$P(R_{t+1} R_t)$
+r	+r	0.7
+r	-r	0.3
-r	+r	0.3
-r	-r	0.7

R _t	Ut	$P(U_t R_t)$
+r	+u	0.9
+r	-u	0.1
-r	+u	0.2
-r	-u	0.8

Pacman – Sonar (P4)



Approximate Inference

- Sometimes |X| is too big for exact inference
 - |X| may be too big to even store B(X)
 - E.g. when X is continuous
 - |X|² may be too big to do updates
- Solution: approximate inference by sampling
- How robot localization works in practice

Approximate Inference: Sampling



Sampling

- Sampling is a lot like repeated simulation
 - Predicting the weather, basketball games, ...
- Basic idea
 - Draw N samples from a sampling distribution S
 - Compute an approximate probability

- Why sample?
 - Learning: get samples from a distribution you don't know
 - Inference: getting a sample is faster than computing the right answer



Sampling

- Sampling from given distribution
 - Step 1: Get sample *u* from uniform distribution over [0, 1)
 - E.g. random() in python
 - Step 2: Convert this sample *u* into an outcome for the given distribution by having each target outcome associated with a sub-interval of [0,1) with sub-interval size equal to probability of the outcome

Example



 $\begin{array}{l} 0 \leq u < 0.6, \rightarrow C = red \\ 0.6 \leq u < 0.7, \rightarrow C = green \\ 0.7 \leq u < 1, \rightarrow C = blue \end{array}$

- If random() returns u = 0.83, then our sample is C = blue
- E.g, after sampling 8 times:



Particle Filtering



Particle Filtering

- Filtering: approximate solution
- Sometimes |X| is too big to use exact inference
 - |X| may be too big to even store B(X)
 - E.g. X is continuous
- Solution: approximate inference
 - Track samples of X, not all values
 - Samples are called particles
 - Time per step is linear in the number of samples
 - But: number needed may be large
 - In memory: list of particles, not states
- This is how robot localization works in practice
- Particle is just new name for sample

0.0	0.1	0.0
0.0	0.0	0.2
0.0	0.2	0.5





Representation: Particles

- Our representation of P(X) is now a list of N particles (samples)
 - Generally, N << |X|</p>
 - Storing map from X to counts would defeat the point
- P(x) approximated by number of particles with value x
 - So, many x may have P(x) = 0!
 - More particles, more accuracy
- For now, all particles have a weight of 1



Particles: (3,3)

(2,3) (3,3) (3,2) (3,3)

(3,2) (1,2) (3,3) (3,3) (2,3)

Particle Filtering: Elapse Time

Each particle is moved by sampling its next position from the transition model

 $x' = \operatorname{sample}(P(X'|x))$

- Samples' frequencies reflect the transition probabilities
- Here, most samples move clockwise, but some move in another direction or stay in place
- This captures the passage of time
 - If enough samples, close to exact values before and after (consistent)



(3,3)(2,3)(3,3)(3,2)

(3,3)(3,2)(1,2)(3,3)

(3,3) (2,3)

(3,2)(2,3)(3,2)

(3,1)

(3,3)(3,2)

(1,3)

(2,3)(3,2) (2,2)

Particle Filtering: Observe

Slightly trickier:

- Don't sample observation, fix it
- Similar to likelihood weighting, downweight samples based on the evidence

w(x) = P(e|x) $B(X) \propto P(e|X)B'(X)$

 As before, the probabilities don't sum to one, since all have been downweighted (in fact they now sum to (N times) an approximation of P(e))



Particle Filtering: Resample

Particles:

(3,2) w=.9

(2,3) w=.2 (3,2) w=.9 (3,1) w=.4

(3,3) w=.4

(3,2) w=.9 (1,3) w=.1

(2,3) w=.2 (3,2) w=.9 (2,2) w=.4

(New) Particles

(3,2) (2,2)

(3,2) (2,3)

(3,3) (3,2) (1,3) (2,3) (3,2) (3,2)

- Rather than tracking weighted samples, we resample
- N times, we choose from our weighted sample distribution (i.e. draw with replacement)
- This is equivalent to renormalizing the distribution
- Now the update is complete for this time step, continue with the next one

		-
:		



Recap: Particle Filtering

Particles: track samples of states rather than an explicit distribution

		Elapse		Weight		Resample		
•							•	
Particle	<i>م</i> د.		Particles:		Particles:		(New) Parti	cles:
(3 3)			(3.2)		(3.2) w= 9		(3.2)	01001
(2,3)			(2,3)		(2,2) w= 2		(2,2)	
(2,3)			(3,2)		(2,3) w = 9		(2,2)	
(3,3)			(3,2)		(3,1) w= 4		(2, 3)	
(3,2)			(3,2)		(3,2) w=4		(2,3)	
(3,3)			(3,2)		(3,2) w= 9		(3,2)	
(1.2)			(1.3)		(1.3) w=.1		(1.3)	
(3.3)			(2.3)		(2,3) w=.2		(2.3)	
(3.3)			(3.2)		(3.2) w=.9		(3.2)	
(2,3)			(2,2)		(2,2) w=.4		(3,2)	

 $x' = \operatorname{sample}(P(X'|x))$ w(x) = P(e|x)

Video of Demo – Moderate Number of Particles



Which Algorithm?

Particle filter, uniform initial beliefs, 25 particles



Which Algorithm?

Exact filter, uniform initial beliefs



Which Algorithm?

Particle filter, uniform initial beliefs, 300 particles



Robot Localization

In robot localization:

- We know the map, but not the robot's position
- Observations may be vectors of range finder readings
- State space and readings are typically continuous (works basically like a very fine grid) and so we cannot store B(X)
- Particle filtering is a main technique





Particle Filter Localization (Sonar)



[Video: global-sonar-uw-annotated.avi]

Particle Filter Localization (Laser)



Robot Mapping

- SLAM: Simultaneous Localization And Mapping
 - We do not know the map or our location
 - State consists of position AND map!
 - Main techniques: Kalman filtering (Gaussian HMMs) and particle methods





[Demo: PARTICLES-SLAM-mapping1-new.avi]