



# Recap: Reasoning Over Time

- Markov models




$$P(X_1)$$

$$P(X|X_{-1})$$



$$P(E|X)$$

- Hidden Markov models




X	E	P
rain	umbrella	0.9
rain	no umbrella	0.1
sun	umbrella	0.2
sun	no umbrella	0.8

# Passage of Time

- Assume we have current belief  $P(X \mid \text{evidence to date})$

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



- Then, after one time step passes:

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

- Or, compactly:

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

- 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”

<0.01	<0.01	<0.01	<0.01	<0.01	<0.01
<0.01	<0.01	<0.01	<0.01	<0.01	<0.01
<0.01	<0.01	1.00	<0.01	<0.01	<0.01
<0.01	<0.01	<0.01	<0.01	<0.01	<0.01

T = 1

<0.01	<0.01	<0.01	<0.01	<0.01	<0.01
<0.01	<0.01	0.06	<0.01	<0.01	<0.01
<0.01	0.76	0.06	0.06	<0.01	<0.01
<0.01	<0.01	0.06	<0.01	<0.01	<0.01

T = 2

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

T = 5

$$B'(X') = \sum_x P(X'|x)B(x)$$

Transition model: ghosts usually go clockwise

# 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

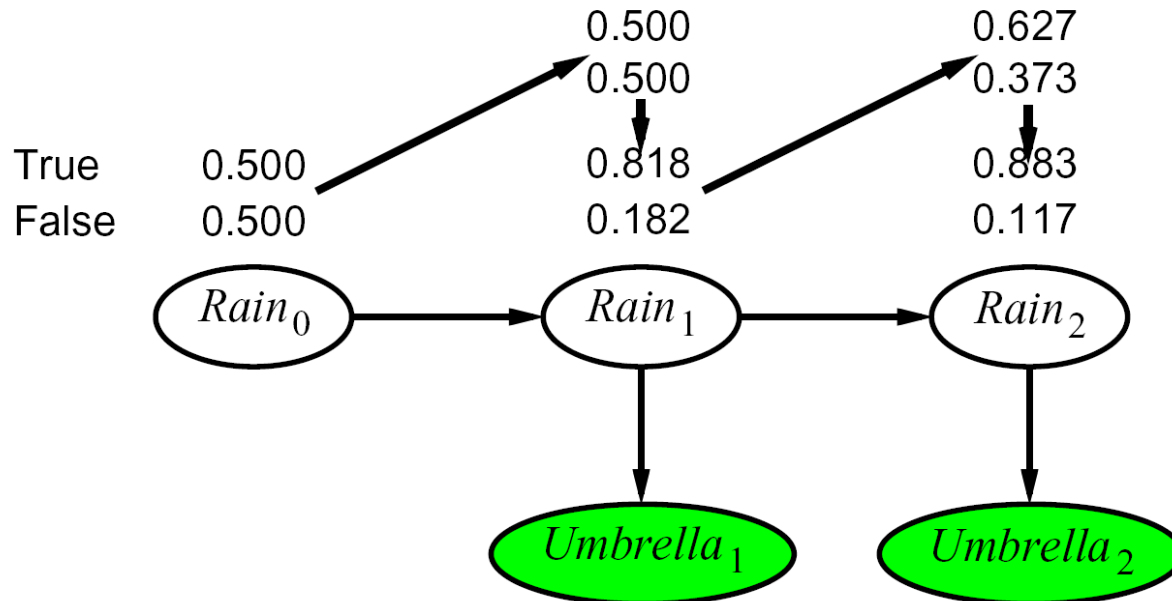
Before observation

<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

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

# Example HMM



# The Forward Algorithm

- We are given evidence at each time and want to know

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

- We can derive the following updates

$$P(x_t|e_{1:t}) \propto_X P(x_t, e_{1:t})$$

$$= \sum_{x_{t-1}} P(x_{t-1}, x_t, e_{1:t})$$

$$= \sum_{x_{t-1}} P(x_{t-1}, e_{1:t-1}) P(x_t|x_{t-1}) P(e_t|x_t)$$


$$= P(e_t|x_t) \sum_{x_{t-1}} P(x_t|x_{t-1}) P(x_{t-1}, e_{1:t-1})$$

We can normalize as we go if we want to have  $P(x|e)$  at each time step, or just once at the end...

# Online Belief Updates

- Every time step, we start with current  $P(X | \text{evidence})$
- We update for time:

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



- We update for evidence:

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



- The forward algorithm does both at once (and doesn't normalize)
- Problem: space is  $|X|$  and time is  $|X|^2$  per time step

- Voice Recognition:

<http://www.youtube.com/watch?v=d9gDcHBmr3I>




# Filtering

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

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



**Belief:**  $\langle P(\text{rain}), P(\text{sun}) \rangle$

$P(X_1)$        $\langle 0.5, 0.5 \rangle$       *Prior on  $X_1$*



$P(X_1 | E_1 = \text{umbrella})$        $\langle 0.82, 0.18 \rangle$       *Observe*

$P(X_2 | E_1 = \text{umbrella})$        $\langle 0.63, 0.37 \rangle$       *Elapse time*

$P(X_2 | E_1 = \text{umb}, E_2 = \text{umb})$        $\langle 0.88, 0.12 \rangle$       *Observe*


# Particle Filtering

- Sometimes  $|X|$  is too big to use exact inference
  - $|X|$  may be too big to even store  $B(X)$
  - E.g.  $X$  is continuous
  - $|X|^2$  may be too big to do updates
- 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



# Representation: Particles

- Our representation of  $P(X)$  is now a list of  $N$  particles (samples)
  - Generally,  $N \ll |X|$
  - Storing map from  $X$  to counts would defeat the point
- $P(x)$  approximated by number of particles with value  $x$ 
  - So, many  $x$  will 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)

(2,1)

(3,3)

(3,3)


(2,1)

# Particle Filtering: Elapse Time

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

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

- This is like prior sampling – samples' frequencies reflect the transition probs
  - Here, most samples move clockwise, but some move in another direction or stay in place
- This captures the passage of time
    - If we have enough samples, close to the exact values before and after (consistent)




# Particle Filtering: Observe

- Slightly trickier:
  - Don't do rejection sampling (why not?)
  - We don't sample the observation, we fix it
  - This is similar to likelihood weighting, so we downweight our samples based on the evidence

$$w(x) = P(e|x)$$

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

- Note that, as before, the probabilities don't sum to one, since most have been downweighted (in fact they sum to an approximation of  $P(e)$ )



# Particle Filtering: Resample



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

Old Particles:

(3,3)  $w=0.1$   
(2,1)  $w=0.9$   
(2,1)  $w=0.9$   
(3,1)  $w=0.4$   
(3,2)  $w=0.3$   
(2,2)  $w=0.4$   
(1,1)  $w=0.4$   
(3,1)  $w=0.4$   
(2,1)  $w=0.9$   
(3,2)  $w=0.3$


New Particles:

(2,1)  $w=1$   
(2,1)  $w=1$   
(2,1)  $w=1$   
(3,2)  $w=1$   
(2,2)  $w=1$   
(2,1)  $w=1$   
(1,1)  $w=1$   
(3,1)  $w=1$   
(2,1)  $w=1$   
(1,1)  $w=1$



# 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 often used



<http://www.youtube.com/watch?v=INLja6Ya3Ig&feature=related>

<http://www.youtube.com/watch?v=kqJpuMNHFG&feature=related>

# Ghostbusters

- **Plot:** Pacman's grandfather, Grandpac, learned to hunt ghosts for sport.
- He was blinded by his power, but could hear the ghosts' banging and clanging.
- **Transition Model:** All ghosts move randomly, but are sometimes biased
- **Emission Model:** Pacman knows a “noisy” distance to each ghost

Noisy distance prob  
True distance = 8

