

CS 188: Artificial Intelligence

Hidden Markov Models



Instructors: Pieter Abbeel and Dan Klein --- University of California, Berkeley

[These slides were created by Dan Klein and Pieter Abbeel for CS188 Intro to AI at UC Berkeley. All CS188 materials are available at <http://ai.berkeley.edu>.]

Reasoning over Time or Space


- Often, we want to reason about a sequence of observations
 - Speech recognition
 - Robot localization
 - User attention
 - Medical monitoring
- Need to introduce time (or space) into our models

Probability Recap

- Conditional probability $P(x|y) = \frac{P(x,y)}{P(y)}$
- Product rule $P(x,y) = P(x|y)P(y)$
- Chain rule $P(X_1, X_2, \dots, X_n) = P(X_1)P(X_2|X_1)P(X_3|X_1, X_2) \dots = \prod_{i=1}^n P(X_i|X_1, \dots, X_{i-1})$
- X, Y independent if and only if: $\forall x, y : P(x, y) = P(x)P(y)$
- X and Y are conditionally independent given Z if and only if: $X \perp\!\!\!\perp Y | Z$
 $\forall x, y, z : P(x, y|z) = P(x|z)P(y|z)$


Markov Models

- Value of X at a given time is called the **state**



- Parameters: called **transition probabilities** or dynamics, specify how the state evolves over time (also, initial state probabilities)
- Stationarity assumption: transition probabilities the same at all times
- Same as MDP transition model, but no choice of action

Conditional Independence



- Basic conditional independence:
 - Past and future independent given the present
 - Each time step only depends on the previous
 - This is called the (first order) Markov property
- Note that the chain is just a (growable) BN
 - We can always use generic BN reasoning on it if we truncate the chain at a fixed length


Example Markov Chain: Weather

- States: $X = \{\text{rain, sun}\}$


- Initial distribution: 1.0 sun

- CPT $P(X_t | X_{t-1})$:

X_{t-1}	X_t	$P(X_t X_{t-1})$
sun	sun	0.9
sun	rain	0.1
rain	sun	0.3
rain	rain	0.7




Two new ways of representing the same CPT



Example Markov Chain: Weather

- Initial distribution: 1.0 sun



- What is the probability distribution after one step?

$$\begin{aligned}
 P(X_2 = \text{sun}) &= P(X_2 = \text{sun} | X_1 = \text{sun})P(X_1 = \text{sun}) + \\
 &\quad P(X_2 = \text{sun} | X_1 = \text{rain})P(X_1 = \text{rain}) \\
 &= 0.9 \cdot 1.0 + 0.3 \cdot 0.0 = 0.9
 \end{aligned}$$

Mini-Forward Algorithm


- Question: What's $P(X)$ on some day t ?



$P(x_1)$ = known

$$\begin{aligned}
 P(x_t) &= \sum_{x_{t-1}} P(x_{t-1}, x_t) \\
 &= \sum_{x_{t-1}} P(x_t | x_{t-1})P(x_{t-1})
 \end{aligned}$$

← Forward simulation



Example Run of Mini-Forward Algorithm

- From initial observation of sun

$$\begin{aligned}
 \begin{pmatrix} 1.0 \\ 0.0 \end{pmatrix}_{P(X_1)} &\begin{pmatrix} 0.9 \\ 0.1 \end{pmatrix}_{P(X_2)} &\begin{pmatrix} 0.84 \\ 0.16 \end{pmatrix}_{P(X_3)} &\begin{pmatrix} 0.804 \\ 0.196 \end{pmatrix}_{P(X_4)} &\longrightarrow &\begin{pmatrix} 0.75 \\ 0.25 \end{pmatrix}_{P(X_\infty)}
 \end{aligned}$$

- From initial observation of rain

$$\begin{aligned}
 \begin{pmatrix} 0.0 \\ 1.0 \end{pmatrix}_{P(X_1)} &\begin{pmatrix} 0.3 \\ 0.7 \end{pmatrix}_{P(X_2)} &\begin{pmatrix} 0.48 \\ 0.52 \end{pmatrix}_{P(X_3)} &\begin{pmatrix} 0.588 \\ 0.412 \end{pmatrix}_{P(X_4)} &\longrightarrow &\begin{pmatrix} 0.75 \\ 0.25 \end{pmatrix}_{P(X_\infty)}
 \end{aligned}$$

- From yet another initial distribution $P(X_1)$:

$$\begin{aligned}
 \begin{pmatrix} p \\ 1-p \end{pmatrix}_{P(X_1)} &\dots &\longrightarrow &\begin{pmatrix} 0.75 \\ 0.25 \end{pmatrix}_{P(X_\infty)}
 \end{aligned}$$

[Demo: L13D1,2,3]

Video of Demo Ghostbusters Basic Dynamics



Video of Demo Ghostbusters Circular Dynamics



Video of Demo Ghostbusters Whirlpool Dynamics



Stationary Distributions

- For most chains:

- Influence of the initial distribution gets less and less over time.
- The distribution we end up in is independent of the initial distribution

- Stationary distribution:


- The distribution we end up with is called the **stationary distribution** P_∞ of the chain
- It satisfies

$$P_\infty(X) = P_{\infty+1}(X) = \sum_x P(X|x)P_\infty(x)$$



Example: Stationary Distributions

- Question: What's $P(X)$ at time $t = \text{infinity}$?



$$P_\infty(\text{sun}) = P(\text{sun}|\text{sun})P_\infty(\text{sun}) + P(\text{sun}|\text{rain})P_\infty(\text{rain})$$

$$P_\infty(\text{rain}) = P(\text{rain}|\text{sun})P_\infty(\text{sun}) + P(\text{rain}|\text{rain})P_\infty(\text{rain})$$


$$P_\infty(\text{sun}) = 0.9P_\infty(\text{sun}) + 0.3P_\infty(\text{rain})$$

$$P_\infty(\text{rain}) = 0.1P_\infty(\text{sun}) + 0.7P_\infty(\text{rain})$$

$$P_\infty(\text{sun}) = 3P_\infty(\text{rain})$$

$$P_\infty(\text{rain}) = 1/3P_\infty(\text{sun})$$

Also: $P_\infty(\text{sun}) + P_\infty(\text{rain}) = 1$ \Rightarrow $P_\infty(\text{sun}) = 3/4$
 $P_\infty(\text{rain}) = 1/4$




X_{t-1}	X_t	$P(X_t X_{t-1})$
sun	sun	0.9
sun	rain	0.1
rain	sun	0.3
rain	rain	0.7

Application of Stationary Distribution: Web Link Analysis


- PageRank over a web graph

- Each web page is a state
- Initial distribution: uniform over pages
- Transitions:
 - With prob. c , uniform jump to a random page (dotted lines, not all shown)
 - With prob. $1-c$, follow a random outlink (solid lines)



- Stationary distribution

- Will spend more time on highly reachable pages
- E.g. many ways to get to the Acrobat Reader download page
- Somewhat robust to link spam
- Google 1.0 returned the set of pages containing all your keywords in decreasing rank, now all search engines use link analysis along with many other factors (rank actually getting less important over time)



Application of Stationary Distributions: Gibbs Sampling*

- Each joint instantiation over all hidden and query variables is a state: $\{X_1, \dots, X_n\} = H \cup Q$


- Transitions:

- With probability $1/n$ resample variable X_i according to

$$P(X_i | x_1, x_2, \dots, x_{i-1}, x_{i+1}, \dots, x_n, e_1, \dots, e_m)$$

- Stationary distribution:


- Conditional distribution $P(X_1, X_2, \dots, X_n | e_1, \dots, e_m)$
- Means that when running Gibbs sampling long enough we get a sample from the desired distribution
- Requires some proof to show this is true!



Hidden Markov Models



Pacman – Sonar (P4)




[Demo: Pacman – Sonar – No Beliefs(L14D1)]

Video of Demo Pacman – Sonar (no beliefs)




Hidden Markov Models

- Markov chains not so useful for most agents
 - Need observations to update your beliefs
- Hidden Markov models (HMMs)
 - Underlying Markov chain over states X
 - You observe outputs (effects) at each time step



Example: Weather HMM





An HMM is defined by:

- Initial distribution: $P(X_1)$
- Transitions: $P(X_t | X_{t-1})$
- Emissions: $P(E_t | X_t)$

R_{t-1}	R_t	$P(R_t R_{t-1})$	R_t	U_t	$P(U_t R_t)$
+r	+r	0.7	+r	+u	0.9
+r	-r	0.3	+r	-u	0.1
-r	+r	0.3	-r	+u	0.2
-r	-r	0.7	-r	-u	0.8

Example: Ghostbusters HMM

- $P(X_1)$ = uniform
- $P(X | X')$ = usually move clockwise, but sometimes move in a random direction or stay in place
- $P(R_t | X)$ = same sensor model as before: red means close, green means far away.



1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

$P(X_1)$

1/6	1/6	1/2
0	1/6	0
0	0	0

$P(X | X' = \langle 1, 2 \rangle)$


[Demo: Ghostbusters – Circular Dynamics – HMM (L14D2)]

Video of Demo Ghostbusters – Circular Dynamics -- HMM



Conditional Independence

- HMMs have two important independence properties:
 - Markov hidden process: future depends on past via the present
 - Current observation independent of all else given current state



- Quiz: does this mean that evidence variables are guaranteed to be independent?
 - [No, they tend to be correlated by the hidden state]

Real HMM Examples


- Speech recognition HMMs:
 - Observations are acoustic signals (continuous valued)
 - States are specific positions in specific words (so, tens of thousands)
- Machine translation HMMs:
 - Observations are words (tens of thousands)
 - States are translation options
- Robot tracking:
 - Observations are range readings (continuous)
 - States are positions on a map (continuous)

Filtering / Monitoring

- Filtering, or monitoring, is the task of tracking the distribution $B_t(X) = P_t(X_t \mid e_1, \dots, e_t)$ (the belief state) over time
- We start with $B_1(X)$ in an initial setting, usually uniform
- As time passes, or we get observations, we update $B(X)$
- The Kalman filter was invented in the 60's and first implemented as a method of trajectory estimation for the Apollo program


Example: Robot Localization

Example from
Michael Pfeiffer




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.

Example: Robot Localization




Lighter grey: was possible to get the reading, but less likely b/c
required 1 mistake


Example: Robot Localization




Example: Robot Localization



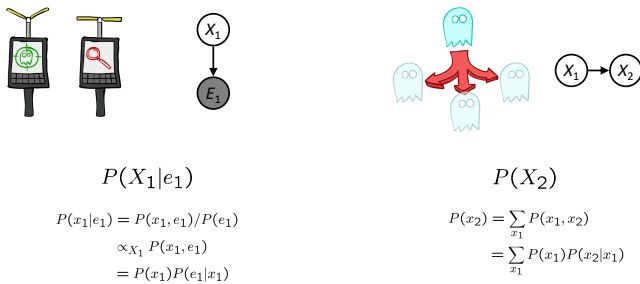
Example: Robot Localization



Example: Robot Localization



Inference: Base Cases



Passage of Time

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

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

- Then, after one time step passes:

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

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


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

- 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”

(Transition model: ghosts usually go clockwise)



Observation

- Assume we have current belief $P(X | \text{previous evidence})$:

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

- Then, after evidence comes in:

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

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


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

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

- 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

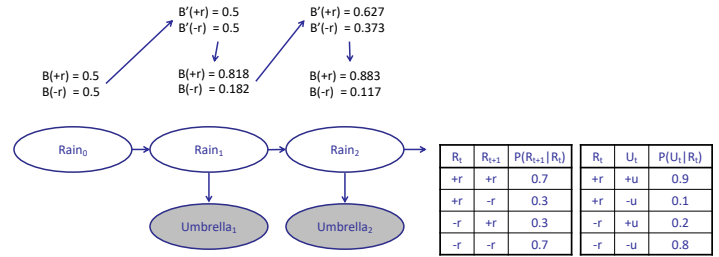


Example: Observation

- As we get observations, beliefs get reweighted, uncertainty “decreases”



Example: Weather 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

$$\begin{aligned}
 P(x_t|e_{1:t}) &\propto_X P(x_t, e_{1:t}) \\
 &= \sum_{x_{t-1}} P(x_{t-1}, x_t, e_{1:t}) \\
 &= \sum_{x_{t-1}} P(x_{t-1}, e_{1:t-1}) P(x_t|x_{t-1}) P(e_t|x_t) \\
 &= P(e_t|x_t) \sum_{x_{t-1}} P(x_t|x_{t-1}) P(x_{t-1}, e_{1:t-1})
 \end{aligned}$$

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 $P(X|evidence)$
- 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)



Pacman – Sonar (P4)



[Demo: Pacman – Sonar – No Beliefs(L14D1)]

Video of Demo Pacman – Sonar (with beliefs)



Next Time: Particle Filtering and Applications of HMMs
