CSE 573: Artificial Intelligence
Autumn 2012

Reasoning about Uncertainty
&
Hidden Markov Models

Daniel Weld

Many slides adapted from Dan Klein, Stuart Russell, Andrew Moore & Luke Zettlemoyer

Outline

- Overview
- Probability review
  - Random Variables and Events
  - Joint / Marginal / Conditional Distributions
  - Product Rule, Chain Rule, Bayes' Rule
- Probabilistic inference
  - Enumeration of Joint Distribution
  - Bayesian Networks – Preview
- Probabilistic sequence models (and inference)
  - Markov Chains
  - Hidden Markov Models
  - Particle Filters

Agent

Static vs. Dynamic

Environment

Fully vs. Partially Observable

Perfect vs. Noisy

Deterministic vs. Stochastic

Instantaneous vs. Durative

Percepts

Actions

Partial Observability

VS

4

Markov Decision Process (MDP)

S:	set of states
A:	set of actions
Pr(s'|s,a):	transition model
R(s,a,s'):	reward model
γ:	discount factor
s_0:	start state

Objective of a Fully Observable MDP

- Find a policy π: **S** → **A**
- which maximizes expected discounted reward
  - given an infinite horizon
  - assuming full observability

Partially-Observable MDP

S:	set of states
A:	set of actions
Pr(s'|s,a):	transition model
R(s,a,s'):	reward model
γ:	discount factor
s_0:	start state
E	set of possible pieces of evidence
Pr(e|s)	observation model

Ghostbusting Observations

- A ghost is in the grid somewhere
- Model 1: Sensor readings tell distance to ghost:
  - On top of pacman: red
  - 1 or 2 away: orange
  - 3 or 4 away: yellow
  - 5+ away: green

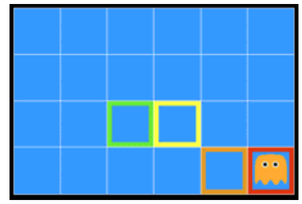

- Model 2: Sensors are noisy, but we know P(Color | Distance)

P(red | 3)	P(orange | 3)	P(yellow | 3)	P(green | 3)
0.05	0.15	0.5	0.3

Objective of a POMDP

- Find a policy
  - π: BeliefStates(**S**) → **A**
  - A belief state is a ***probability distribution*** over states
- which maximizes expected discounted reward
  - given an infinite horizon
  - assuming full observability

Particle Filtering

Planning in HW 4

- Map Estimate

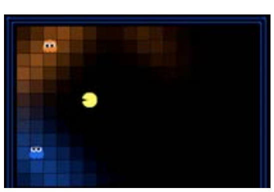

11

Projects

- You choose…
- Default 1
  - Extend Pacman reinforcement learning, eg UCT
- Default 2
  - Extend Pacman to real POMDP

12

Random Variables

- A *random variable* is some aspect of the world about which we (may) have uncertainty
  - R = Is it raining?
  - D = How long will it take to drive to work?
  - L = Where am I?
- We denote random variables with capital letters
- Random variables have domains
  - R in {true, false}
  - D in [0, 1)
  - L in possible locations, maybe {(0,0), (0,1), …}

Joint Distributions

- A *joint distribution* over a set of random variables: $X_1, X_2, \ldots X_n$ specifies a real number for each *outcome* (ie each assignment):

$$P(X_1 = x_1, X_2 = x_2, \ldots X_n = x_n)$$
$$P(x_1, x_2, \ldots x_n)$$

  - Must obey: $P(x_1, x_2, \ldots x_n) \geq 0$

$$\sum_{(x_1, x_2, \ldots x_n)} P(x_1, x_2, \ldots x_n) = 1$$

  - Size of distribution if n variables with domain sizes d?
- A *probabilistic model* is a joint distribution over variables of interest
- For all but the smallest distributions, impractical to write out

$P(T, W)$

T	W	P
hot	sun	0.4
hot	rain	0.1
cold	sun	0.2
cold	rain	0.3

Terminology

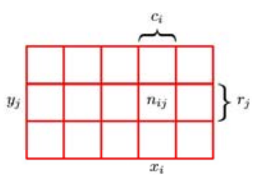

Marginal Probability

$$p(X = x_i) = \frac{c_i}{N}.$$

Joint Probability

$$p(X = x_i, Y = y_j) = \frac{n_{ij}}{N}$$

Conditional Probability

$$p(Y = y_j | X = x_i) = \frac{n_{ij}}{c_i}$$

X value is given

Independence

P(A∧B) = P(A)P(B)

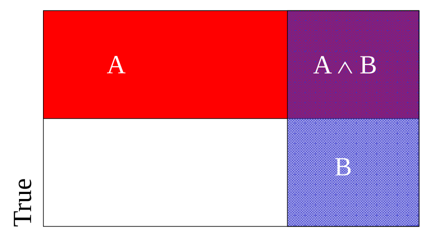

Conditional Independence

Are A & B independent? P(A|B) ≤ P(A)

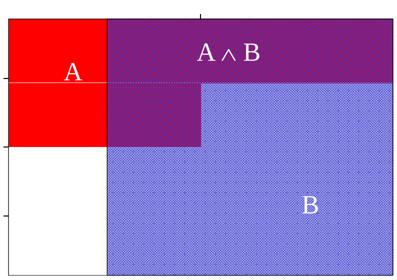

P(A)=(.25+.5)/2
= .375

P(B)= .75

P(A|B)=(.25+.25+.5)/3
=.3333

A, B Conditionally Independent Given C

P(A|B,C) = P(A|C) C = spot free

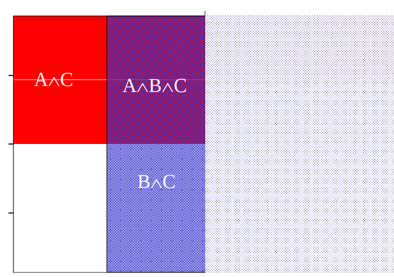

P(A|¬C) =.5
P(A|B,¬C)=.5

Probabilistic Inference

- *Probabilistic inference*: compute a desired probability from other known probabilities (e.g. conditional from joint)
- We generally compute conditional probabilities
  - P(on time | no reported accidents) = 0.90
  - These represent the agent's *beliefs* given the evidence
- Probabilities change with new evidence:
  - P(on time | no accidents, 5 a.m.) = 0.95
  - P(on time | no accidents, 5 a.m., raining) = 0.80
  - Observing new evidence causes *beliefs to be updated*

Inference by Enumeration

P(hot | winter)?

S	T	W	P
summer	hot	sun	0.30
summer	hot	rain	0.05
summer	cold	sun	0.10
summer	cold	rain	0.05
winter	hot	sun	0.10
winter	hot	rain	0.05
winter	cold	sun	0.15
winter	cold	rain	0.20

Inference by Enumeration

- General case:
  - Evidence variables: $E_1 \dots E_k = e_1 \dots e_k$
  - Query* variable: Q
  - Hidden variables: $H_1 \dots H_r$

  } $X_1, X_2, \dots X_n$ *All variables*
- We want: $P(Q|e_1 \dots e_k)$
- First, select the entries consistent with the evidence
- Second, sum out H to get joint of Query and evidence:

$$P(Q, e_1 \dots e_k) = \underbrace{\sum_{h_1 \dots h_r} P(Q, h_1 \dots h_r, e_1 \dots e_k)}_{X_1, X_2, \dots X_n} \; / \; P(e_1 \dots e_k)$$

- Finally, normalize the remaining entries to conditionalize
- Obvious problems:
  - Worst-case time complexity $O(d^n)$
  - Space complexity $O(d^n)$ to store the joint distribution

Bayes' Nets: Big Picture

- Two problems with using full joint distribution tables as our probabilistic models:
  - Unless there are only a few variables, the joint is WAY too big to represent explicitly
  - Hard to learn (estimate) anything empirically about more than a few variables at a time
- Bayes' nets: a technique for describing complex joint distributions (models) using simple, local distributions (conditional probabilities)
  - More properly called graphical models
  - We describe how variables locally interact
  - Local interactions chain together to give global, indirect interactions

Bayes' Net Semantics

- Let's formalize the semantics of a Bayes' net
- A set of nodes, one per variable X
- A directed, acyclic graph
- A conditional distribution for each node
  - A collection of distributions over X, one for each combination of parents' values

$$P(X|a_1 \dots a_n)$$

  - CPT: conditional probability table

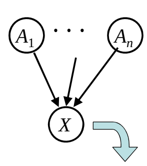

$$P(X|A_1 \dots A_n)$$

A Bayes net = Topology (graph) + Local Conditional Probabilities

Example Bayes' Net: Car

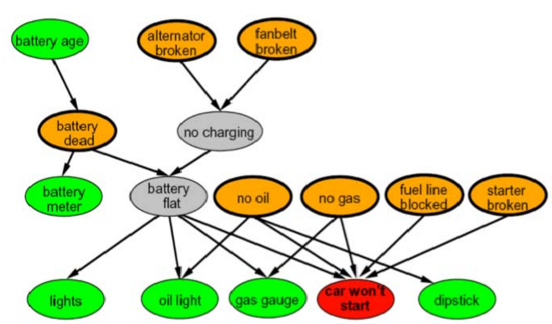

The Chain Rule

- More generally, can always write any joint distribution as an incremental product of conditional distributions

$$P(x_1, x_2, x_3) = P(x_1)P(x_2|x_1)P(x_3|x_1, x_2)$$

$$P(x_1, x_2, \ldots x_n) = \prod_i P(x_i|x_1 \ldots x_{i-1})$$

Bayes' Rule

- Two ways to factor a joint distribution over two variables:

$$P(x, y) = P(x|y)P(y) = P(y|x)P(x)$$

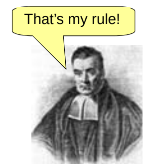

- Dividing, we get:

$$P(x|y) = \frac{P(y|x)}{P(y)}P(x)$$

- Why is this at all helpful?
  - Lets us build a conditional from its reverse
  - Often one conditional is tricky but the other one is simple
  - Foundation of many systems we'll see later
- In the running for most important AI equation!

Inference with Bayes' Rule

- Example: Diagnostic probability from causal probability:

$$P(\text{Cause}|\text{Effect}) = \frac{P(\text{Effect}|\text{Cause})P(\text{Cause})}{P(\text{Effect})}$$

- Example:
  - m is meningitis, s is stiff neck

$$P(s|m) = 0.8$$
$$P(m) = 0.0001$$
$$P(s) = 0.1$$

Example givens

$$P(m|s) = \frac{P(s|m)P(m)}{P(s)} = \frac{0.8 \times 0.0001}{0.1} = 0.0008$$

Ghostbusters, Revisited

- Let's say we have two distributions:
  - Prior distribution over ghost location: P(G)
    - Let's say this is uniform
  - Sensor reading model: P(R | G)
    - Given: we know what our sensors do
    - R = reading color measured at (1,1)
    - E.g. P(R = yellow | G=(1,1)) = 0.1

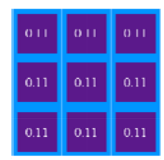

- We can calculate the posterior distribution P(G|r) over ghost locations given a reading using Bayes' rule:

$$P(g|r) \propto P(r|g)P(g)$$

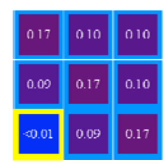

Markov Models (Markov Chains)

- A Markov model includes:
  - Random variables X_t for all time steps t (the state)
  - Parameters: called transition probabilities or dynamics, specify how the state evolves over time (also, initial probs)

$P(X_1)$ and $P(X_t|X_{t-1})$

- Later we'll see that a Markov model is just a chain-structured Bayesian Network (BN)

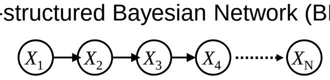

Conditional Independence

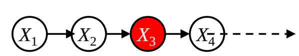

- Basic conditional independence:
  - Each time step only depends on the previous
  - Future conditionally independent of past given the present
  - This is called the (first order) Markov property
- This chain is just a (growing) BN
  - We could use generic BN reasoning on it if we truncate the chain at a fixed length

Markov Models (Markov Chains)

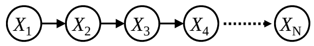

- A Markov model defines
  - a joint probability distribution:

$$P(X_1, \ldots, X_n) = P(X_1) \prod_{t=2}^{N} P(X_t | X_{t-1})$$

- One common inference problem:
  - Compute marginals $P(X_t)$ for some time step, t

States *vs.* Random Variables

Weather Probabilistic FSM:

- States = {rain, sun}
- Transitions:

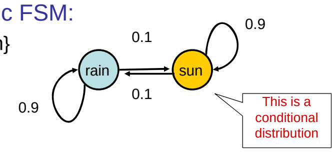

Markov Chain

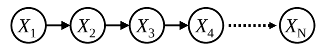

Example: Markov Chain

- Weather:
  - States = {rain, sun}
  - Transitions:

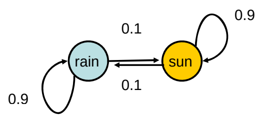

- Initial distribution: 1.0 sun
- What's the probability distribution after one step?

$$P(X_2 = \text{sun}) = P(X_2 = \text{sun}|X_1 = \text{sun})P(X_1 = \text{sun}) + P(X_2 = \text{sun}|X_1 = \text{rain})P(X_1 = \text{rain})$$

$$0.9 \cdot 1.0 + 0.1 \cdot 0.0 = 0.9$$

Markov Chain Inference

- Question: probability of being in state x at time t?
- Slow answer:
  - Enumerate all sequences of length t which end in s
  - Add up their probabilities

$$P(X_t = sun) = \sum_{x_1 \ldots x_{t-1}} P(x_1, \ldots x_{t-1}, sun)$$

$$P(X_1 = sun)P(X_2 = sun|X_1 = sun)P(X_3 = sun|X_2 = sun)P(X_4 = sun|X_3 = sun)$$

$$P(X_1 = sun)P(X_2 = rain|X_1 = sun)P(X_3 = sun|X_2 = rain)P(X_4 = sun|X_3 = sun)$$

$$\vdots$$

Mini-Forward Algorithm

- Question: What's P(X) on some day t?
  - We don't need to enumerate all 2^t sequences!

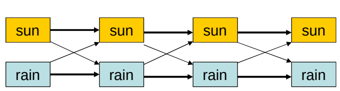

$$P(x_t) = \sum_{x_{t-1}} P(x_t|x_{t-1})P(x_{t-1})$$

$$P(x_1) = \text{known}$$

Forward simulation

Example

- From initial observation of sun

$$\begin{pmatrix} 1.0 \\ 0.0 \end{pmatrix} \quad \begin{pmatrix} 0.9 \\ 0.1 \end{pmatrix} \quad \begin{pmatrix} 0.82 \\ 0.18 \end{pmatrix} \quad \Rightarrow \quad \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}$$

$P(X_1)$ $P(X_2)$ $P(X_3)$ $P(X)$

- From initial observation of rain

$$\begin{pmatrix} 0.0 \\ 1.0 \end{pmatrix} \quad \begin{pmatrix} 0.1 \\ 0.9 \end{pmatrix} \quad \begin{pmatrix} 0.18 \\ 0.82 \end{pmatrix} \quad \Rightarrow \quad \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}$$

$P(X_1)$ $P(X_2)$ $P(X_3)$ $P(X)$

Stationary Distributions

- If we simulate the chain long enough:
  - What happens?
  - Uncertainty accumulates
  - Eventually, we have no idea what the state is!
- Stationary distributions:
  - For most chains, the distribution we end up in is independent of the initial distribution
  - Called the stationary distribution of the chain
  - Usually, can only predict a short time out

Pac-man Markov Chain

Pac-man knows the ghost's initial position, but gets no observations!

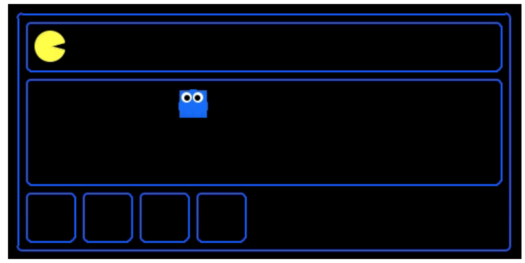

Web Link Analysis

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

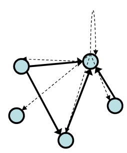

Hidden Markov Models

- Markov chains not so useful for most agents
  - Eventually you don't know anything anymore
  - Need observations to update your beliefs
- Hidden Markov models (HMMs)
  - Underlying Markov chain over states S
  - You observe outputs (effects) at each time step
  - As a Bayes' net:

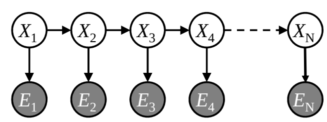

Example

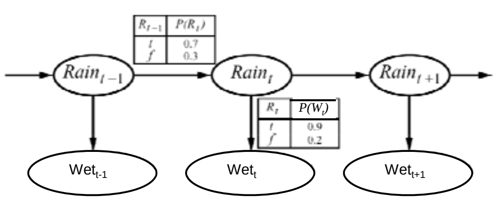

- An HMM is defined by:
  - Initial distribution: $P(X_1)$
  - Transitions: $\mathbf{P}(\mathbf{X}_t|\mathbf{X}_{t-1})$
  - Emissions: $P(E|X)$

Example

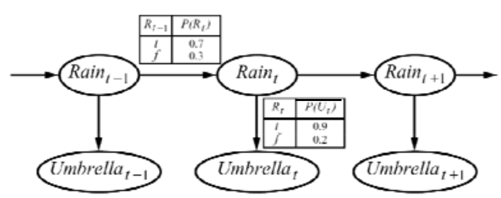

- An HMM is defined by:
  - Initial distribution: $P(X_1)$
  - Transitions: $\mathbf{P}(\mathbf{X}_t|\mathbf{X}_{t-1})$
  - Emissions: $P(E|X)$

Hidden Markov Models

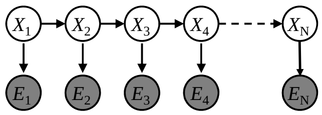

- Defines a joint probability distribution:

$$P(X_1,\dots,X_n,E_1,\dots,E_n) = P(X_{1:n},E_{1:n}) = P(X_1)P(E_1|X_1)\prod_{t=2}^{N} P(X_t|X_{t-1})P(E_t|X_t)$$

Ghostbusters HMM

- $P(X_1)$ = uniform
- $P(X'|X)$ = usually move clockwise, but sometimes move in a random direction or stay in place
- $P(E|X)$ = 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

$P(X_1)$

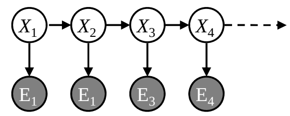

1/6	1/6	1/2
0	1/6	0
0	0	0

P(X'|X=<1,2>)

P(E|X)	P(red | 3)	P(orange | 3)	P(yellow | 3)	P(green | 3)
	0.05	0.15	0.5	0.3

HMM Computations

- Given
  - joint $P(X_{1:n},E_{1:n})$
  - evidence $E_{1:n} = e_{1:n}$
- Inference problems include:
  - Filtering, find $P(X_t|e_{1:t})$ for all t
  - Smoothing, find $P(X_t|e_{1:n})$ for all t
  - Most probable explanation, find

    $x^*_{1:n} = \text{argmax}_{x_{1:n}} P(x_{1:n}|e_{1:n})$