

CS 188: Artificial Intelligence

Probability

Pieter Abbeel – UC Berkeley
Many slides adapted from Dan Klein.

Our Status in CS188

- We're done with Part I Search and Planning!
- Part II: Probabilistic Reasoning
 - Diagnosis
 - Tracking objects
 - Speech recognition
 - Robot mapping
 - Genetics
 - Error correcting codes
 - ... lots more!
- Part III: Machine Learning

2

Part II: Probabilistic Reasoning

- Probability
- Distributions over LARGE Numbers of Random Variables
 - Representation
 - Independence
 - Inference
 - Variable Elimination
 - Sampling
 - Hidden Markov Models

3


Probability

- Probability
 - Random Variables
 - Joint and Marginal Distributions
 - Conditional Distribution
 - Inference by Enumeration
 - Product Rule, Chain Rule, Bayes' Rule
 - Independence
- You'll need all this stuff A LOT for the next few weeks, so make sure you go over it now and know it inside out! The next few weeks we will learn how to make these work computationally efficiently for LARGE numbers of random variables.

4

Inference in Ghostbusters

- A ghost is in the grid somewhere
- Sensor readings tell how close a square is to the ghost
 - On the ghost: red
 - 1 or 2 away: orange
 - 3 or 4 away: yellow
 - 5+ away: green



- Sensors are noisy, but we know $P(\text{Color} | \text{Distance})$

$P(\text{red} 3)$	$P(\text{orange} 3)$	$P(\text{yellow} 3)$	$P(\text{green} 3)$
0.05	0.15	0.5	0.3

Uncertainty

- General situation:
 - **Evidence:** Agent knows certain things about the state of the world (e.g., sensor readings or symptoms)
 - **Hidden variables:** Agent needs to reason about other aspects (e.g., where an object is or what disease is present)
 - **Model:** Agent knows something about how the known variables relate to the unknown variables
- Probabilistic reasoning gives us a framework for managing our beliefs and knowledge

0.11	0.11	0.11
0.11	0.11	0.11
0.11	0.11	0.11

0.17	0.16	0.10
0.09	0.17	0.10
0.04	0.09	0.17

-0.04	-0.01	0.03
-0.01	0.05	0.05
-0.01	0.05	0.01

6

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
- Like variables in a CSP, random variables have domains
 - R in {true, false} (sometimes write as {+r, -r})
 - D in $[0, \infty)$
 - L in possible locations, maybe $\{(0,0), (0,1), \dots\}$

7

Probability Distributions

- Unobserved random variables have distributions

T	P
warm	0.5
cold	0.5

W	P
sun	0.6
rain	0.1
fog	0.3
meteor	0.0

- A distribution is a TABLE of probabilities of values
- A probability (lower case value) is a single number

$$P(W = \text{rain}) = 0.1 \quad P(\text{rain}) = 0.1$$

- Must have: $\forall x P(x) \geq 0 \quad \sum_x P(x) = 1$

8

Joint Distributions

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

$$P(X_1 = x_1, X_2 = x_2, \dots, X_n = x_n) \quad P(T, W)$$

- Size of distribution if n variables with domain sizes d?

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

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

- For all but the smallest distributions, impractical to write out

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

9

Probabilistic Models

- A probabilistic model is a joint distribution over a set of random variables

Distribution over T,W

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

- Probabilistic models:
 - (Random) variables with domains. Assignments are called *outcomes*
 - Joint distributions: say whether assignments (outcomes) are likely
 - Normalized*: sum to 1.0
 - Ideally: only certain variables directly interact

Constraint over T,W

T	W	P
hot	sun	T
hot	rain	F
cold	sun	F
cold	rain	T

- Constraint satisfaction probs:
 - Variables with domains
 - Constraints: state whether assignments are possible
 - Ideally: only certain variables directly interact

10

Events

- An *event* is a set E of outcomes

$$P(E) = \sum_{(x_1, \dots, x_n) \in E} P(x_1 \dots x_n)$$

- From a joint distribution, we can calculate the probability of any event

- Probability that it's hot AND sunny?
- Probability that it's hot?
- Probability that it's hot OR sunny?

- Typically, the events we care about are *partial assignments*, like $P(T=\text{hot})$

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

11

Marginal Distributions

- Marginal distributions are sub-tables which eliminate variables
- Marginalization (summing out): Combine collapsed rows by adding

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

$$P(t) = \sum_s P(t, s)$$

T	P
hot	0.5
cold	0.5

$$P(s) = \sum_t P(t, s)$$

W	P
sun	0.6
rain	0.4


$$P(X_1 = x_1) = \sum_{x_2} P(X_1 = x_1, X_2 = x_2)$$

12

Conditional Probabilities

- A simple relation between joint and conditional probabilities
 - In fact, this is taken as the *definition* of a conditional probability

$$P(a|b) = \frac{P(a, b)}{P(b)}$$



$$P(T, W)$$

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

$$P(W = r | T = c) = ???$$

13

Conditional Distributions

- Conditional distributions are probability distributions over some variables given fixed values of others

Conditional Distributions		Joint Distribution																						
$P(W T)$	$P(W T = hot)$	$P(T, W)$																						
	<table border="1"> <thead> <tr> <th>W</th> <th>P</th> </tr> </thead> <tbody> <tr> <td>sun</td> <td>0.8</td> </tr> <tr> <td>rain</td> <td>0.2</td> </tr> </tbody> </table>	W	P	sun	0.8	rain	0.2	<table border="1"> <thead> <tr> <th>T</th> <th>W</th> <th>P</th> </tr> </thead> <tbody> <tr> <td>hot</td> <td>sun</td> <td>0.4</td> </tr> <tr> <td>hot</td> <td>rain</td> <td>0.1</td> </tr> <tr> <td>cold</td> <td>sun</td> <td>0.2</td> </tr> <tr> <td>cold</td> <td>rain</td> <td>0.3</td> </tr> </tbody> </table>			T	W	P	hot	sun	0.4	hot	rain	0.1	cold	sun	0.2	cold	rain
W	P																							
sun	0.8																							
rain	0.2																							
T	W	P																						
hot	sun	0.4																						
hot	rain	0.1																						
cold	sun	0.2																						
cold	rain	0.3																						
	$P(W T = cold)$																							
	<table border="1"> <thead> <tr> <th>W</th> <th>P</th> </tr> </thead> <tbody> <tr> <td>sun</td> <td>0.4</td> </tr> <tr> <td>rain</td> <td>0.6</td> </tr> </tbody> </table>	W	P	sun	0.4	rain	0.6																	
W	P																							
sun	0.4																							
rain	0.6																							

14

Normalization Trick

- A trick to get a whole conditional distribution at once:
 - Select the joint probabilities matching the evidence
 - Normalize the selection (make it sum to one)

$$P(T, W)$$

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

Select

$$P(T, r)$$

T	R	P
hot	rain	0.1
cold	rain	0.3

Normalize

$$P(T|r)$$

T	P
hot	0.25
cold	0.75

- Why does this work? Sum of selection is P(evidence)! (P(r), here)

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

15

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*

16

Inference by Enumeration

- P(sun)?
- P(sun | winter)?
- P(sun | winter, warm)?

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

17

Inference by Enumeration

- General case:
 - Evidence variables: $E_1 \dots E_k = e_1 \dots e_k$
 - Query* variable: Q
 - Hidden variables: $H_1 \dots H_r$
- $\left. \begin{array}{l} X_1, X_2, \dots, X_n \\ \text{All variables} \end{array} \right\}$
- 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:

$$P(Q, e_1 \dots e_k) = \sum_{h_1 \dots h_r} P(Q, h_1 \dots h_r, e_1 \dots e_k)$$

X_1, X_2, \dots, X_n
 - 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

* Works fine with multiple query variables, too

Inference by Enumeration Example 2: Model for Ghostbusters

- Reminder: ghost is hidden, sensors are noisy
- T: Top sensor is red
B: Bottom sensor is red
G: Ghost is in the top

- Queries:
P(+g) = ??
P(+g | +t) = ??
P(+g | +t, -b) = ??

- Problem: joint distribution too large / complex

0.50
0.50

Joint Distribution

T	B	G	P(T,B,G)
+t	+b	+g	0.16
+t	+b	-g	0.16
+t	-b	+g	0.24
+t	-b	-g	0.04
-t	+b	+g	0.04
-t	+b	-g	0.24
-t	-b	+g	0.06
-t	-b	-g	0.06

The Product Rule

- Sometimes have conditional distributions but want the joint

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

- Example:

P(W)	
R	P
sun	0.8
rain	0.2

P(D W)		
D	W	P
wet	sun	0.1
dry	sun	0.9
wet	rain	0.7
dry	rain	0.3

P(D,W)		
D	W	P
wet	sun	0.08
dry	sun	0.72
wet	rain	0.14
dry	rain	0.06

The Chain Rule

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

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

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

- Why is this always true?

- Can now build a joint distributions only specifying conditionals!

- Bayesian networks essentially apply the chain rule plus make conditional independence assumptions.

21

Bayes' Rule

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

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

That's my rule!

- Dividing, we get:

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



- Why is this at all helpful?

- Lets us build one conditional from its reverse
- Often one conditional is tricky but the other one is simple
- Foundation of many systems we'll see later (e.g. ASR, MT)

- In the running for most important AI equation!

22

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
- | | | |
|-----------------|---|----------------|
| $P(s m) = 0.8$ | } | Example givens |
| $P(m) = 0.0001$ | | |
| $P(s) = 0.1$ | | |

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

- Note: posterior probability of meningitis still very small
- Note: you should still get stiff necks checked out! Why?

23

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

0.11	0.11	0.11
0.11	0.11	0.11
0.11	0.11	0.11

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

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

0.17	0.10	0.10
0.09	0.17	0.10
<0.01	0.09	0.17

24

Independence

- Two variables are *independent* if:

$$\forall x, y : P(x, y) = P(x)P(y)$$

- Says their joint distribution *factors* into a product two simpler ones.
- Usually variables are not independent!
- Equivalent definition of independence:

$$\forall x, y : P(x|y) = P(x)$$

- We write: $X \perp\!\!\!\perp Y$

- Independence is a simplifying *modeling assumption*
 - Empirical* joint distributions: at best "close" to independent
 - What could we assume for {Weather, Traffic, Cavity, Toothache}?
- Independence is like something from CSPs, what?

25

Example: Independence?

T	P
warm	0.5
cold	0.5

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

W	P
sun	0.6
rain	0.4

T	W	P
warm	sun	0.3
warm	rain	0.2
cold	sun	0.3
cold	rain	0.2

26

Example: Independence

- N fair, independent coin flips:

H	0.5
T	0.5

H	0.5
T	0.5

...

H	0.5
T	0.5

$$P(X_1, X_2, \dots, X_n)$$

27

Conditional Independence

- $P(\text{Toothache}, \text{Cavity}, \text{Catch})$
- If I have a cavity, the probability that the probe catches in it doesn't depend on whether I have a toothache:
 - $P(+\text{catch} | +\text{toothache}, +\text{cavity}) = P(+\text{catch} | +\text{cavity})$
- The same independence holds if I don't have a cavity:
 - $P(+\text{catch} | +\text{toothache}, -\text{cavity}) = P(+\text{catch} | -\text{cavity})$
- Catch is *conditionally independent* of Toothache given Cavity:
 - $P(\text{Catch} | \text{Toothache}, \text{Cavity}) = P(\text{Catch} | \text{Cavity})$
- Equivalent statements:
 - $P(\text{Toothache} | \text{Catch}, \text{Cavity}) = P(\text{Toothache} | \text{Cavity})$
 - $P(\text{Toothache}, \text{Catch} | \text{Cavity}) = P(\text{Toothache} | \text{Cavity}) P(\text{Catch} | \text{Cavity})$
 - One can be derived from the other easily

28

Conditional Independence

- Unconditional (absolute) independence very rare (why?)

- Conditional independence* is our most basic and robust form of knowledge about uncertain environments:

$$\forall x, y, z : P(x, y|z) = P(x|z)P(y|z) \quad X \perp\!\!\!\perp Y | Z$$

$$\forall x, y, z : P(x|z, y) = P(x|z)$$

- What about this domain:

- Traffic
- Umbrella
- Raining

- What about fire, smoke, alarm?

29

The Chain Rule Revisited

$$P(X_1, X_2, \dots, X_n) = P(X_1)P(X_2|X_1)P(X_3|X_1, X_2) \dots$$

- Trivial decomposition:

$$P(\text{Traffic}, \text{Rain}, \text{Umbrella}) = P(\text{Rain})P(\text{Traffic}|\text{Rain})P(\text{Umbrella}|\text{Rain}, \text{Traffic})$$

- With assumption of conditional independence:

$$P(\text{Traffic}, \text{Rain}, \text{Umbrella}) = P(\text{Rain})P(\text{Traffic}|\text{Rain})P(\text{Umbrella}|\text{Rain})$$

- Representation size: 1 + 2 + 4 versus 1 + 2 + 2

- Bayes' nets / graphical models are concerned with distributions with conditional independences

30

Ghostbusters Chain Rule

- Each sensor depends only on where the ghost is

$$P(T,B,G) = P(G) P(T|G) P(B|G)$$

- That means, the two sensors are conditionally independent, given the ghost position

- T: Top square is red
B: Bottom square is red
G: Ghost is in the top

- Givens:
 $P(+g) = 0.5$
 $P(+t | +g) = 0.8$
 $P(+t | -g) = 0.4$
 $P(+b | +g) = 0.4$
 $P(+b | -g) = 0.8$

T	B	G	P(T,B,G)
+t	+b	+g	0.16
+t	+b	-g	0.16
+t	-b	+g	0.24
+t	-b	-g	0.04
-t	+b	+g	0.04
-t	+b	-g	0.24
-t	-b	+g	0.06
-t	-b	-g	0.06