CS 383: Artificial Intelligence

Probability

Prof. Scott Niekum — UMass Amherst

Today

- **Probability**
  - Random Variables
  - Joint and Marginal Distributions
  - Conditional Distributions
  - Product Rule, Chain Rule, Bayes' Rule
  - Inference
  - Independence

- **You'll need all this stuff A LOT for the next few weeks, so make sure you go over it now!**

Inference in Ghostbusters

- A ghost is in the grid somewhere

- Noisy sensor readings tell how close a square is to the ghost. Most likely observations:

  - On the ghost: red
  - 1 or 2 away: orange
  - 3 or 4 away: yellow
  - 5+ away: green

- 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

Ghostbusters, no probabilities

Uncertainty

- **General situation:**

  - **Observed variables (evidence)**: Agent knows certain things about the state of the world (e.g., sensor readings or symptoms)

  - **Unobserved 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 using beliefs and knowledge to perform *inference*

Random Variables

- A random variable is some aspect of the world about which we (may) have uncertainty

  - R = Is it raining?
  - T = Is it hot or cold?
  - D = How long will it take to drive to work?
  - L = Where is the ghost?

- We denote random variables with capital letters

- Like variables in a CSP, random variables have domains

  - R in {true, false}   (often write as {+r, -r})
  - T in {hot, cold}
  - D in [0, ∞)
  - L in possible locations, maybe {(0,0), (0,1), …}

Probability Distributions

- Associate a probability with each value

- Temperature:

$$P(T)$$

T	P
hot	0.5
cold	0.5

- Weather:

$$P(W)$$

W	P
sun	0.6
rain	0.1
fog	0.3
meteor	0.0

Probability Distributions

- Unobserved random variables have distributions

$P(T)$

T	P
hot	0.5
cold	0.5

$P(W)$

W	P
sun	0.6
rain	0.1
fog	0.3
meteor	0.0

Shorthand notation:

$$P(hot) = P(T = hot),$$
$$P(cold) = P(T = cold),$$
$$P(rain) = P(W = rain),$$
$$\cdots$$

- A discrete distribution is a table of probabilities of values

OK *if* all domain entries are unique

- A probability (lower case value) is a single number

$$P(W = rain) = 0.1$$

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

Joint Distributions

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

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

$P(T, W)$

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

  - 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?

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

Probabilistic Models

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

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

Distribution over T,W

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

Constraint over T,W

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

Events

- An *event* is a set E of outcomes

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

- 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=hot)

$P(T, W)$

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

Quiz: Events

- P(+x, +y) ?

$$P(X, Y)$$

X	Y	P
+x	+y	0.2
+x	-y	0.3
-x	+y	0.4
-x	-y	0.1

Quiz: Events

- P(+x, +y) ?

- P(+x) ?

iClicker:

A: 0.2

B: 0.3

C: 0.5

D: 1.0

$$P(X, Y)$$

X	Y	P
+x	+y	0.2
+x	-y	0.3
-x	+y	0.4
-x	-y	0.1

Quiz: Events

- P(+x, +y) ?

- P(+x) ?

- P(-y OR +x) ?

iClicker:

A: 0.2

B: 0.4

C: 0.6

D: 0.9

$P(X, Y)$

X	Y	P
+x	+y	0.2
+x	-y	0.3
-x	+y	0.4
-x	-y	0.1

Marginal Distributions

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

$P(T, W)$

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

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

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

$P(T)$

T	P
hot	0.5
cold	0.5

$P(W)$

W	P
sun	0.6
rain	0.4

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

Quiz: Marginal Distributions

$P(X,Y)$

X	Y	P
+x	+y	0.1
+x	-y	0.5
-x	+y	0.2
-x	-y	0.2

$$P(x) = \sum_y P(x,y)$$

$$P(y) = \sum_x P(x,y)$$

$P(X)$

X	P
+x	**?**
-x	

$P(Y)$

Y	P
+y	
-y	

iClicker:

A: 0.05

B: 0.1

C: 0.5

D: 0.6

Quiz: Marginal Distributions

$$P(X,Y)$$

X	Y	P
+x	+y	0.1
+x	-y	0.5
-x	+y	0.2
-x	-y	0.2

$$P(x) = \sum_{y} P(x,y)$$

$$P(y) = \sum_{x} P(x,y)$$

$$P(X)$$

| X	P
+x	0.6
-x	0.4

$$P(Y)$$

Y	P
+y	0.3
-y	0.7

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(a,b)

P(a   P(b)

$$P(T,W)$$

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

$$P(W = s|T = c) = \frac{P(W = s, T = c)}{P(T = c)} = \frac{0.2}{0.5} = 0.4$$

$$= P(W = s, T = c) + P(W = r, T = c)$$
$$= 0.2 + 0.3 = 0.5$$

Quiz: Conditional Probabilities

$$P(X,Y)$$

X	Y	P
+x	+y	0.2
+x	-y	0.3
-x	+y	0.4
-x	-y	0.1

- P(+x | +y) ?

iClicker:

A: 2/10

B: 1/4

C: 1/3

D: 1/2

Normalization Trick

$P(T, W)$

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

$$P(W = s | T = c) = \frac{P(W = s, T = c)}{P(T = c)}$$
$$= \frac{P(W = s, T = c)}{P(W = s, T = c) + P(W = r, T = c)}$$
$$= \frac{0.2}{0.2 + 0.3} = 0.4$$

$$P(W = r | T = c) = \frac{P(W = r, T = c)}{P(T = c)}$$
$$= \frac{P(W = r, T = c)}{P(W = s, T = c) + P(W = r, T = c)}$$
$$= \frac{0.3}{0.2 + 0.3} = 0.6$$

$P(W | T = c)$

W	P
sun	0.4
rain	0.6

Normalization Trick

$$P(W = s | T = c) = \frac{P(W = s, T = c)}{P(T = c)}$$
$$= \frac{P(W = s, T = c)}{P(W = s, T = c) + P(W = r, T = c)}$$
$$= \frac{0.2}{0.2 + 0.3} = 0.4$$

$P(T, W)$

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

SELECT the joint probabilities matching the evidence

⟶

$P(c, W)$

T	W	P
cold	sun	0.2
cold	rain	0.3

NORMALIZE the selection (make it sum to one)

⟶

$P(W | T = c)$

W	P
sun	0.4
rain	0.6

$$P(W = r | T = c) = \frac{P(W = r, T = c)}{P(T = c)}$$
$$= \frac{P(W = r, T = c)}{P(W = s, T = c) + P(W = r, T = c)}$$
$$= \frac{0.3}{0.2 + 0.3} = 0.6$$

Normalization Trick

$P(T, W)$

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

SELECT the joint probabilities matching the evidence

→

$P(c, W)$

T	W	P
cold	sun	0.2
cold	rain	0.3

NORMALIZE the selection (make it sum to one)

→

$P(W|T = c)$

W	P
sun	0.4
rain	0.6

- Why does this work?

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

To Normalize

- (Dictionary) To bring or restore to a normal condition

  All entries sum to ONE

- Procedure:

  - Step 1: Compute Z = sum over all entries
  - Step 2: Divide every entry by Z

- Example 1

W	P
sun	0.2
rain	0.3

Normalize →

Z = 0.5

W	P
sun	0.4
rain	0.6

- Example 2

T	W	P
hot	sun	20
hot	rain	5
cold	sun	10
cold	rain	15

Normalize →

Z = 50

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

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

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

  $\left.\begin{array}{l}\end{array}\right\}$ $X_1, X_2, \ldots X_n$

  All variables

- We want:

  * Works fine with multiple query variables, too

  $$P(Q|e_1 \ldots e_k)$$

- Step 1: Select the entries consistent with the evidence

x	$P(x)$
-3	0.05
-1	0.25
0	0.07
1	0.2
5	0.01
2	0.15

- Step 2: Sum out H to get joint of Query and evidence

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

- Step 3: Normalize

$$\times \frac{1}{Z}$$

$$Z = \sum_q P(Q, e_1 \cdots e_k)$$

$$P(Q|e_1 \cdots e_k) = \frac{1}{Z} P(Q, e_1 \cdots e_k)$$

Inference by Enumeration

- ## P(W)?

  p(W=sun) = 0.3 + 0.1 + 0.1 + 0.15 = 0.65
  p(W=rain) = 0.05 + 0.05 + 0.05 + 0.2 = 0.35

- ## P(W | winter)?

  p(W=sun , winter) = 0.1 + 0.15 = 0.25
  p(W=rain , winter) = 0.05 + 0.2 = 0.25
  p(W=sun | winter) = 0.25 / 0.25 + 0.25 = 0.5
  p(W=rain | winter) = 0.25 / 0.25 + 0.25 = 0.5

- ## P(W | winter, hot)?

  p(W=sun , winter, hot) = 0.1
  p(W=rain , winter, hot) = 0.05
  p(W=sun | winter, hot) = 0.1 / 0.1 + 0.05 = 2/3
  p(W=rain | winter, hot) = 0.05 / 0.1 + 0.05 = 1/3

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

- Obvious problems:

  - Worst-case time complexity $O(d^n)$

  - Space complexity $O(d^n)$ to store the joint distribution

  - What about continuous distributions?

The Product Rule

- Sometimes have conditional distributions but want the joint

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

The Product Rule

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

■ 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	
dry	sun	
wet	rain	
dry	rain	

The Chain Rule

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

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

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

- Why is this always true?

$$P(x_{1+2}, x_3) = P(x_3|x_{1+2})P(x_{1+2})$$

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

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

Bayes' Rule

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

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

That's my rule!

- Dividing, we get:

likelihood

prior

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

normalization

- 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

- 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: meningitis, S: stiff neck

$$\left.\begin{array}{l} P(+m) = 0.0001 \\ P(+s|+m) = 0.8 \\ P(+s|-m) = 0.01 \end{array}\right\} \begin{array}{l}\text{Example}\\\text{givens}\end{array}$$

$$P(+m|+s) = \frac{P(+s|+m)P(+m)}{P(+s)} = \frac{P(+s|+m)P(+m)}{P(+s|+m)P(+m) + P(+s|-m)P(-m)} = \frac{0.8 \times 0.0001}{0.8 \times 0.0001 + 0.01 \times 0.999}$$

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

  =0.0008

Quiz: Bayes' Rule

- Given:

$$P(W)$$

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

iClicker:

A: .06 / .72

B: .06 / .78

C: .72 / .78

D: .72

- What is P(sun | dry) ?

Quiz: Bayes' Rule

- Given:

$P(W)$

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

- What is P(W | dry) ?

p(sun | dry) = p(dry | sun) p(sun) / p(dry) = 0.9 * 0.8 / Z = .72 / Z

p(rain | dry) = p(dry | rain) p(rain) / p(dry) = 0.3 * 0.2 / Z   = 0.06 / Z

Z = .72 + .06   = .78

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:

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

Ghostbusters with Probability