Un
Uncertainty

Uncertainty in AI: Probabilistic Reasoning

CSE 415: Introduction to Artificial Intelligence
University of Washington
Spring 2017

Outline

- Motivation
- Generalizing Modus Ponens
- Bayes' Rule
- Definitions and Laws of Probability
- The Monty Hall Problem

Motivation

Logical reasoning has limitations:

It requires that assumptions be considered "certain".

It typically uses general rules. General rules that are reliable may be difficult to come by.

Logical reasoning can be awkward for certain structured domains such as time and space.

Generalizing Modus Ponens

Modus Ponens:

P → Q	**If it's raining then I do my homework.**
P	**It's raining.**
Q	**I do my homework.**

Bayes' Rule: (general idea)

If P then sometimes Q	**If it's raining then I might do my homework.**
P	**It's raining.**
Maybe Q	**I might do my homework.**

(Bayes' rule lets us calculate the probability of Q, taking P into account.)

Bayes' Rule

E: Some evidence exists, i.e., a particular condition is true
H: some hypothesis is true.

P(E|H) = probability of E given H.
P(E|~H) = probability of E given not H.
P(H) = probability of H, independent of E.

$$P(H|E) = \frac{P(E|H)\,P(H)}{P(E)}$$

$$P(E) = P(E|H)\,P(H) + P(E|\sim H)(1 - P(H))$$

Applying Bayes' Rule

E: The patient's white blood cell count exceeds 110% of average.

H: The patient is infected with tetanus.

P(E|H) = 0.8	**class-conditional probability**
P(E|~H) = 0.3	**"**
P(H) = 0.01	**prior probability**

posterior probability:

$$P(H|E) = \frac{P(E|H)\,P(H)}{P(E)} = \frac{(0.8)\ (0.01)}{(0.8)\,(0.01) + (0.3)(0.99)} = \frac{0.008}{0.305} = 0.0262$$

$$P(E) = P(E|H)\,P(H) + P(E|\sim H)(1 - P(H))$$

Un Uncertainty

Odds

Odds are 10 to 1 it will rain tomorrow.

$$P(rain) = \frac{10}{10+1} = \frac{10}{11}$$

Suppose P(A) = 1/4
Then O(A) = (1/4) / (3/4) = 1/3

in general: $O(A) = \frac{P(A)}{P(\sim A)} = \frac{P(A)}{1 - P(A)}$

Un Uncertainty

Bayes' Rule reformulated...

$$\frac{P(H|E) = \frac{P(E|H)\,P(H)}{P(E)}}{P(\sim H|E) = \frac{P(E|\sim H)\,P(\sim H)}{P(E)}}$$

$$O(H|E) = \frac{P(E|H)}{P(E|\sim H)}\,O(H)$$

Un Uncertainty

Odds-Likelihood Form of Bayes' Rule

E: The patient's white blood cell count exceeds 110% of average.

H: The patient is infected with tetanus.

O(H) = 0.01/0.99

O(H|E) = λ O(H) **lambda is called the *sufficiency* factor.**

O(H|~E) = λ' O(H) **lambda prime is called the *necessity* factor.**

$\lambda = P(E|H)/P(E|\sim H) = 0.8/0.3 \approx 2.67$
$\lambda' = P(\sim E|H)/P(\sim E|\sim H) = 0.2/0.7 \approx 0.286$

Un Uncertainty

Definitions

A *random variable* X is a symbol that represents a class of events that may occur any number of times, and which take on values in a given set D called a *domain*.

Example: Let C be a coin-toss random variable with domain D={H, T}.

A *probability distribution* P for a random variable is a function that assigns to each domain element d_i a value p_i in the range 0 to 1. If D is finite then P is often given as a table.

d_i	p_i
H	0.5
T	0.5

Un Uncertainty

Definitions

A *joint distribution* over a set of random variables $X_1, X_2, \ldots, X_n$ is a function that assigns to each n-tuple of domain elements $(d_{1i1}, d_{2i2}, \ldots, d_{nin})$ a value $p_{i1i2\ldots in}$ in the range 0 to 1. If all the Dj are finite then P is often given as a table.

Example with n=2:

D_1	D_2	$P(d_{2i1}, d_{2i2})$
rain	no crash	1/4
rain	crash	1/8
clear	no crash	5/16
clear	crash	1/16
snow	no crash	1/8
snow	crash	1/8

← $P(X_1=d_{1i1}, X_2=d_{2i2})$

Un Uncertainty

Definitions

The *conditional probability* of $P(X=d_i \mid Y=d_j)$ is the probability of $X=d_i$, given $Y=d_j$.

Example:

Let D={0, 1, 2, 3, 4, 5, 6, 7, 8, 9}.

Let Odd be the set of outcomes X=1, X=3, ..., X=9.
Let Even " " X=0, X=2, ..., X=8.
Let Prime " " X=2, X=3, X=5, X=7.
P(Odd) = 0.5
P(Prime) = 0.4
P(Odd | Prime) = 0.75
P(Prime | Odd) = 0.6.

Un
Uncertainty

Laws of Probability

The sum rule: $\Sigma_{1<i<m} P(X = d_i) = 1$ Adding the probabilities of all the possible outcomes for a random variable must give a total of 1.0.

The product rule. P(x) P(y | x) = P(x, y).

Un
Uncertainty

The Monty Hall Problem

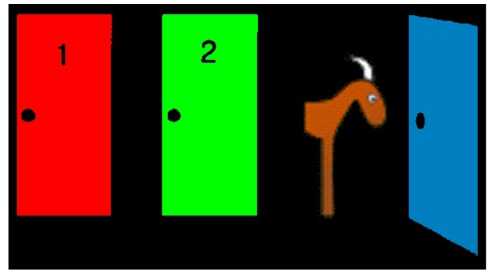

Un
Uncertainty

The Monty Hall Problem

There are three doors: a red door, green door, and blue door.
Behind one is a car, and behind the other two are goats.
You get to keep whatever is behind the door you choose.

You choose a door (say, red).

The host opens one of the other doors (say, green), which reveals a goat.

The host says, "Would you like to select the OTHER door?"

Should you switch?