A Little Opening Fun

Three prisoners, A, B, and C, are locked in separate cells. Everyone knows that one of them will be executed the next day and that the other two will be pardoned. Only the governor knows which one will be executed, and he likes it that way.

Prisoner A asks the prison guard for a favor. “Please ask the governor who will be executed, and then take a message to one of my friends, either B or C, to let him know that he will be pardoned in the morning.” The guard agrees to do that.

The Question

Later, the guard stops by A's cell. He tells A that he gave the good news to B.

What are A's chances of being executed,
given this information?

(Try to answer this logically or mathematically,
... not by energetic waving of hands.)

Imperfect Systems...

inadequate sensors and effectors

the nature of the environment

the pragmatics of knowledge representation

... Uncertain Results

Data are uncertain.

Conclusions are uncertain.

Actions are possibly inappropriate.

The **results of actions** are uncertain.

The **results of learning** are uncertain.

How to Think?

Logic and rules are not enough.

Statistics and probability become useful tools.

An intelligent agents may want to base its decisions on *preferences* and *utilities* rather than merely on inferences.

Assertions Become Contingent

Each statement is given a label that indicates the agent's *degree of belief* in the assertion.

Instead of saying:

The patient has a cavity.

We say:

*It is **highly likely** that the patient has a cavity.*

```
import 2000.session13.certaintyfactors;
```

Probability as a Tool

An agent's degree of belief can change over time:

- before seeing any evidence *prior probability*

$$P(x)$$

- after seeing evidence **e** *posterior probability*

$$P(x \mid e)$$

Example: Your Dentist

before seeing any evidence

$$P(x) = 0.2$$

after evidence: phone call on tootache

$$P(x|e_1) = 0.7$$

after evidence: visual exam

$$P(x|e_1 \wedge e_2) = 0.95$$

after evidence: X-ray

$$P(x|e_1 \wedge e_2 \wedge e_3) = 0.999$$

An Exercise

I have to fly to Dallas on March 1. My internal planner has created a plan that is *99% certain* to get me to the airport on time for my flight.

Should I use this plan or try to find another one?

Why or why not?

Context Matters

What is the cost of executing the plan?

How averse to risk am I?

How easily can I generate alternative plans?

*This is an **economic decision**.*

Probability and Rationality

In the real world, agents have preferences for some states of the world over other states of the world.

Consider some examples...

A person donates 10% of her gross income to charity.

You are playing in the last round of a chess tournament. You have 4.5 points, and your next closest competitor has 4.0 points. The winner of the tournament earns \$10,000, and second place wins \$5,000.

You are playing Garry Kasparov a game of chess. He offers a draw in a position that you think you might be able to win.

You are an injured Cliff Lee. You feel good enough to pitch in today, though the team doctor tells you that there is a 50% chance that you will worsen your injury by pitching today.

Today is June 4.

You are an injured Cliff Lee. You feel good enough to pitch in today, though the team doctor tells you that there is a 50% chance that you will worsen your injury by pitching today.

Today is Game 7 of the World Series.

Rationality

Economists and AI researchers use a notion of ***utility theory*** to measure the expected outcomes of an action and choose among them.

"acting rationally"

==

choosing the action with the highest expected value

Probability is merely a way to compute expectations.

Probability 101

x	an assertion
$P(x)$	degree of belief in x
$P(x \mid E)$	degree of belief in x given that we know E

Probability 101

$0 \leq P(x) \leq 1$ for all x

P(true) == 1

(false) == 0

P(a or b) == P(a) + P(b) - P(a and b)

if a and b are independent

P(x | E) == P(x and E) / P(E)

Probability in Practice

We know that:

$$P(x \mid E) == P(x \text{ and } E) \;/\; P(E)$$

So:

$$P(x \text{ and } E) == P(x \mid E) * P(E)$$

Probability in Practice

But evidence is relative, so:

$$P(E \text{ and } x) == P(E \mid x) * P(x)$$

And so:

$$P(E \mid x) * P(x) == P(x \mid E) * P(E)$$

Bayes' Law

And so:

$$P(x \mid E) == \frac{P(E \mid x) * P(x)}{P(E)}$$

... the work horse of probabilistic AI.

Bayes' Law at Work

You recently read that doctors have reported several cases of meningitis in your town. You wake up one morning with a sore throat, so you decide to call your doctor. He encourages you to come in for an appoint-ment but calms your fears of having meningitis. It turns out that meningitis strikes only 1 in 50,000 people, and the probability that a person will have a sore throat if they do have meningitis is only 50%. And, besides, the chance that you will wake up with a sore throat on any given day is 5%.

What are the chances that you have the disease?

$$P(\text{HaveSoreThroat} \mid \text{HaveDisease}) = 0.50$$
$$P(\text{HaveDisease}) = 0.00002$$
$$P(\text{HaveSoreThroat}) = 0.05$$

$$P(\text{HaveDisease} \mid \text{HaveSoreThroat}) =$$

$$\frac{P(\text{HaveSoreThroat}|\text{HaveDisease}) * P(\text{HaveDisease})}{P(\text{HaveSoreThroat})}$$

$$= (0.50 * 0.00002) / 0.05$$
$$= 0.0002$$

Bayes' Law and Arithmetic

The arithmetic required is tedious and unforgiving.

Humans are not good at reasoning in this way.

Computers are quite good at reasoning in this way.

Bayes' Law and Data

It can be quite expensive to collect and maintain all the data required to reason with Bayes’ Law.

We could ask human experts for the data.
We could collect empirical data.

Computers are quite good at reasoning in this way.

Bayes' Law and Humans

Humans survive just fine without reasoning well in this way.

Instead, humans reason *qualitatively*. They are usually concerned more with “good enough” than with “optimal”.

Computers are not good at reasoning in this way. Yet.

Bonus Solution to "A Little Opening Fun"

Three prisoners, A, B, and C, are locked in their cells. ... Later, the guard tells A that he gave the pardon message to prisoner B.

What are A's chances of being executed, given this information?

Incorrect Solution

F_X = “<x> will be freed” E_X = “<x> will be executed”

$$P(E_A \mid F_B) = \frac{P(F_B \mid E_A) * P(E_A)}{P(F_B)}$$

$$= \frac{1 * (1/3)}{(2/3)} = \frac{1}{2}$$

Correct Solution

F^*_x = "**The guard said that** <x> will be freed"

$$P(E_A \mid F^*_B) = \frac{P(F^*_B \mid E_A) * P(E_A)}{P(F^*_B)}$$

$$P(E_A \mid F_B) = \frac{(1/2) * (1/3)}{(2/3)} = \frac{1}{3}$$