Uncertainty

Russell & Norvig Chapter 13

Uncertainty

Let A_t be the action of leaving for the airport t minutes before your flight
Will A_t get you there on time?

A purely logical approach either
1. risks falsehood: “A_{120} will get me there on time”, or
2. leads to conclusions that are too weak for decision making:

   “A_{120} will get me there on time if there's no accident and it doesn't rain and my tires remain intact etc.”

   (A_{1440} might reasonably be said to get me there on time but I'd have to stay overnight in the airport …)

Uncertainty

Let A_t be the action of leaving for the airport t minutes before your flight
Will A_t get you there on time?

Uncertainty results from:
1. partial observability (road state, other drivers' plans, etc.)
2. noisy sensors (traffic reports)
3. uncertainty in action outcomes (flat tire, etc.)
4. complexity of modeling traffic

Questions

- How to represent uncertainty in knowledge?
- How to perform inference with uncertain knowledge?
- Which action to choose under uncertainty?

Dealing with uncertainty

- Implicit
  - Ignore what you are uncertain of when you can
  - Build procedures that are robust to uncertainty
- Explicit
  - Build a model of the world that describes uncertainty about its state, dynamics, and observations
  - Reason about the effect of actions given the model

Methods for handling uncertainty

- Default Reasoning:
  - Assume the car does not have a flat tire
  - Assume A_{120} works unless contradicted by evidence
- Issues: What assumptions are reasonable? How to handle contradictions?
- Worst case reasoning (the world behaves according to Murphy’s law).
- Probability
  - Model agent's degree of belief
  - Given the available evidence, A_{120} will get me there on time with probability 0.95

Probability

- Probabilities relate propositions to agent's own state of knowledge
  e.g., $P(A_{120}$ | no reported accidents) = 0.96
- Probabilities of propositions change with new evidence:
  e.g., $P(A_{120}$ | no reported accidents, 5 a.m.) = 0.99

Making decisions under uncertainty

Suppose I believe the following:

$P(A_{60}$ gets me there on time | …)	= 0.001
$P(A_{90}$ gets me there on time | …)	= 0.70
$P(A_{120}$ gets me there on time | …)	= 0.95
$P(A_{150}$ gets me there on time | …)	= 0.99
$P(A_{1440}$ gets me there on time | …)	= 0.9999

- Which action to choose?
  Depends on my preferences for missing flight vs. time spent waiting, etc.
  - Utility theory is used to represent and infer preferences
  - Decision theory = probability theory + utility theory

Axioms of probability

- For any events *A*, *B* in a space of events Ω
  - $0 \le P(A) \le 1$
  - $P(\Omega) = 1$ and $P(\phi) = 0$
  - $P(A \vee B) = P(A) + P(B) - P(A \wedge B)$

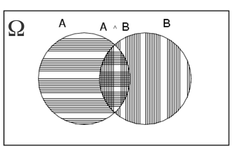

Axioms of probability

- $0 \le P(\omega) \le 1$ $\quad \sum_{\omega \in \Omega} P(\omega) = 1$
- $P(A \vee B) = P(A) + P(B) - P(A \wedge B)$
  [inclusion-exclusion principle]

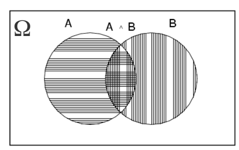

Example

- Consider a deck of cards (52 cards) and the following events:
  - A king
  - A face card
  - A spade
  - A face card or a red suit
  - A card
- What is the probability of each of the above events?

Where do probabilities come from

Two camps:

- Frequentist interpretation
- Bayesian interpretation

Frequentist interpretation

- Draw a ball from an urn containing n balls of the same size; r are red, the rest black.
- The probability of the event "the ball is red" corresponds to the relative frequency with which we expect to draw a red ball

  P(red) = ?

Subjective probabilities

There are many situations in which there is no objective frequency interpretation:

- You have worked hard on your AI class and you believe that the probability that you will get an A is 0.9
- There are theoretical justifications for subjective probabilities!

The Bayesian viewpoint

- Probability is "degree-of-belief".
- To the Bayesian, probability lies subjectively in the mind, and can be different for people with different information
- In contrast, to the frequentist, probability lies objectively in the external world.

Random Variables

- A random variable can be thought of as an unknown value that may change every time it is inspected.
- Suppose that a coin is tossed three times and the sequence of heads and tails is noted. The sample space for this experiment is:

  S={HHH, HHT, HTH, HTT, THH, THT, TTH, TTT}.

  X - the number of heads in three coin tosses. X assigns each outcome in S a number from the set {0, 1, 2, 3}.

Outcome	**HHH**	**HHT**	**HTH**	**THH**	**HTT**	**THT**	**TTH**	**TTT**
X	3	2	2	2	1	1	1	0

- We can now ask the question – what is the probability for observing a particular value for X (the distribution of X).

Random Variables

- Boolean random variables
  e.g., *Cavity* (do I have a cavity?)
  Distribution characterized by a number p.
- Discrete random variables
  e.g., *Weather* is one of <*sunny,rainy,cloudy,snow*>
- Domain values must be exhaustive and mutually exclusive

- The (probability) distribution of a random variable X with m values $x_1, x_2, \ldots, x_n$ is:

  $(p_1, p_2, \ldots, p_m)$

  with $P(X=x_i) = p_i$ and $\Sigma_i\, p_i = 1$

Joint Distribution

- Given n random variables $X_1, \ldots, X_n$
- The joint distribution of these variables is a table in which each entry gives the probability of one combination of values of $X_1, \ldots, X_n$
- Example:

	Toothache	¬Toothache
Cavity	0.04	0.06
¬Cavity	0.01	0.89

P(¬Cavity∧Toothache) P(Cavity∧¬Toothache)

It's all in the joint

	Toothache	¬Toothache
Cavity	0.04	0.06
¬Cavity	0.01	0.89

- P(Toothache) = P((Toothache ∧Cavity) ∨ (Toothache∧¬Cavity))
  = P(Toothache ∧Cavity) + P(Toothache∧¬Cavity)
  = 0.04 + 0.01 = 0.05

  We summed over all values of Cavity: **marginalization**
- P(Toothache ∨ Cavity) =
  P((Toothache ∧Cavity) ∨ (Toothache∧¬Cavity)
  ∨ (¬Toothache ∧Cavity)) =
  0.04 + 0.01 + 0.06 = 0.11

  These are examples of **inference by enumeration**

Conditional Probability

- Definition:
  P(A|B) =P(A ∧ B) / P(B) (if P(B) > 0)
- Read: probability of A given B
  - Example: P(snow) = 0.03 but P(snow | winter) = 0.06, P(snow | summer) = 1e-6
- can also write this as:
  P(A ∧ B) = P(A|B) P(B)
- called the **product rule**

Example

	Toothache	¬Toothache
Cavity	0.04	0.06
¬Cavity	0.01	0.89

P(Cavity|Toothache) = P(Cavity∧Toothache) / P(Toothache)
= 0.04/0.05 = 0.8

Product rule

- P(A ∧ B ∧ C) = P(A|B,C) P(B|C) P(C)

Bayes' Rule

P(A ∧ B) = P(A|B) P(B)
= P(B|A) P(A)

$$\boxed{P(B|A) = \frac{P(A|B)\,P(B)}{P(A)}}$$

Example

	Toothache	¬Toothache
Cavity	0.04	0.06
¬Cavity	0.01	0.89

- Given:
  P(Cavity) = 0.1
  P(Toothache) = 0.05
  P(Cavity|Toothache) = 0.8
- Using Bayes' rule:
  P(Toothache|Cavity) = (0.8x0.05)/0.1
  = 0.4

The Monty Hall Problem

Suppose you're on a game show, and you're given the choice of three doors: Behind one door is a car; behind the others, goats. You pick a door, say No. 1, and the host, who knows what's behind the doors, opens another door, say No. 3, which has a goat. He then says to you, "Do you want to pick door No. 2?" Is it to your advantage to switch your choice?

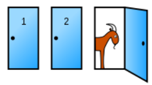

source: http://en.wikipedia.org/wiki/Monty_Hall_problem

Solution

P(C2|O3) = P(O3|C2)P(C2)/P(O3) =
= 1 * 1/3 / 1/2 = 2/3

P(C1|O3) = P(O3|C1)P(C1)/P(O3)
= 1/2*1/3 / 1/2
= 1/3

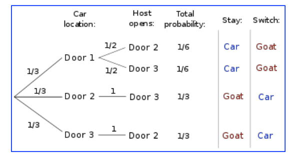

Solution

P(C2|O3) = P(O3|C2)P(C2)/P(O3) =
= 1 * 1/3 / 1/2 = 2/3

P(O3) = P(O3|C1)P(C1) +
P(O3|C2)P(C2) +
P(O3|C3)P(C3) =
= 1/2*1/3 + 1*1/3 + 0*1/3 1/2
= 1/2

Probabilities in the wumpus world

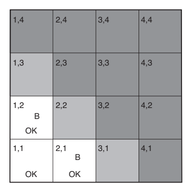

There is no safe choice at this point!
But are there squares that are less likely to contain a pit?

Probabilities in the wumpus world

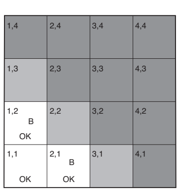

There is no safe choice at this point!
But are there squares that are less likely to contain a pit?

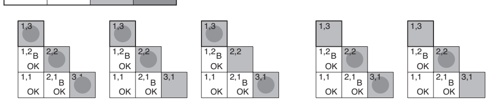

Generalization of Bayes' rule

$P(A \wedge B \wedge C) = P(A \wedge B|C)\,P(C)$
$= P(A|B,C)\,P(B|C)\,P(C)$

$P(A \wedge B \wedge C) = P(A \wedge B|C)\,P(C)$
$= P(B|A,C)\,P(A|C)\,P(C)$

$$P(B|A,C) = \frac{P(A|B,C)\,P(B|C)}{P(A|C)}$$

It's all in the joint but…

- The naïve representation runs into problems.
- Example:
  - Patients in a hospital are described by attributes such as:
    - Background: age, gender, history of diseases, …
    - Symptoms: fever, blood pressure, headache, …
    - Diseases: pneumonia, heart attack, …
- A probability distribution needs to assign a number to each combination of values of these attributes
  - Size of table is exponential in number of attributes

Bayesian networks

- Provide an efficient representation that relies on independence relations between variables.

Naïve Bayesian

- In many cases, effects of a cause are conditionally independent from each other.
- Simplify the rule in a naïve Bayes model

$$P(Cause, Effect_1, \ldots Effect_n) = P(Cause)\prod_i P(Effect_i \mid Cause)$$