Computational Foundations of Cognitive Science

Lecture 19: Conditional Probability, Bayes' Theorem; Bayesian Models of Reasoning

Frank Keller

School of Informatics
University of Edinburgh
keller@inf.ed.ac.uk

February 23, 2010

1. Conditional Probability and Independence
   - Conditional Probability
   - Independence
   - Mid-lecture Problem
2. Bayes' Theorem
   - Total Probability
   - Bayes' Theorem
3. Bayesian Models of Reasoning
   - Background
   - Application to Diagnosis
   - Base Rate Neglect

Reading: Freund, Chs. 2.5–2.8.

Conditional Probability

Definition: Conditional Probability

If A and B are two events in a sample space S, and $P(A) \neq 0$ then the conditional probability of B given A is:

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

Intuitively, the conditional probability $P(B|A)$ is the probability that the event B will occur given that the event A has occurred.

Examples

The probability of having a traffic accident given that it snows: $P(\text{accident}|\text{snow})$.

The probability of reading the word *amok* given that the previous word was *run*: $P(\text{amok}|\text{run})$.

Conditional Probability

Example

A manufacturer knows that the probability of an order being ready on time is 0.80, and the probability of an order being ready on time and being delivered on time is 0.72. What is the probability of an order being delivered on time, given that it is ready on time?

R: order is ready on time; D: order is delivered on time.
$P(R) = 0.80$, $P(R \cap D) = 0.72$. Therefore:

$$P(D|R) = \frac{P(R \cap D)}{P(R)} = \frac{0.72}{0.80} = 0.90$$

Conditional Probability

From the definition of conditional probability, we obtain:

Theorem: Multiplication Rule

If A and B are two events in a sample space S, and $P(A) \neq 0$ then:

$$P(A \cap B) = P(A)P(B|A)$$

As $A \cap B = B \cap A$, it follows also that:

$$P(A \cap B) = P(B)P(A|B)$$

Example

Back to lateralization of language (see last lecture). Let $P(A) = 0.15$ be the probability of being left-handed, $P(B) = 0.05$ be the probability of language being right-lateralized, and $P(A \cap B) = 0.04$.

The probability of language being right-lateralized given that a person is left-handed:

$$P(B|A) = \frac{P(A \cap B)}{P(A)} = \frac{0.04}{0.15} = 0.267$$

The probability being left-handed given that language is right-lateralized:

$$P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{0.04}{0.05} = 0.80$$

Independence

Definition: Independent Events

Two events A and B are independent if and only if:

$$P(B \cap A) = P(A)P(B)$$

Intuitively, two events are independent if the occurrence of non-occurrence of either one does not affect the probability of the occurrence of the other.

Theorem: Complement of Independent Events

If A and B are independent, then A and $\bar{B}$ are also independent.

This follows straightforwardly from set theory.

Independence

Example

A coin is flipped three times. Each of the eight outcomes is equally likely. A: head occurs on each of the first two flips, B: tail occurs on the third flip, C: exactly two tails occur in the three flips. Show that A and B are independent, B and C dependent.

$$\begin{array}{ll} A = \{HHH, HHT\} & P(A) = \frac{1}{4} \\ B = \{HHT, HTT, THT, TTT\} & P(B) = \frac{1}{2} \\ C = \{HTT, THT, TTH\} & P(C) = \frac{3}{8} \\ A \cap B = \{HHT\} & P(A \cap B) = \frac{1}{8} \\ B \cap C = \{HTT, THT\} & P(B \cap C) = \frac{1}{4} \end{array}$$

$P(A)P(B) = \frac{1}{4} \cdot \frac{1}{2} = \frac{1}{8} = P(A \cap B)$, hence A and B are independent.
$P(B)P(C) = \frac{1}{2} \cdot \frac{3}{8} = \frac{3}{16} \neq P(B \cap C)$, hence B and C are dependent.

Mid-lecture Problem

The following figure shows a Venn diagram with probabilities assigned to its various regions. Show that A and B are independent, A and C are independent, B and C are independent, but A, B, and C are not independent.

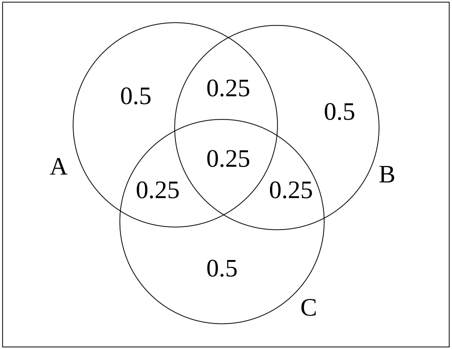

Total Probability

Theorem: Rule of Total Probability

If events $B_1, B_2, \ldots, B_k$ constitute a partition of the sample space S and $P(B_i) \neq 0$ for $i = 1, 2, \ldots, k$, then for any event A in S:

$$P(A) = \sum_{i=1}^{k} P(B_i)P(A|B_i)$$

$B_1, B_2, \ldots, B_k$ form a *partition* of S if they are pairwise mutually exclusive and if $B_1 \cup B_2 \cup \ldots \cup B_k = S$.

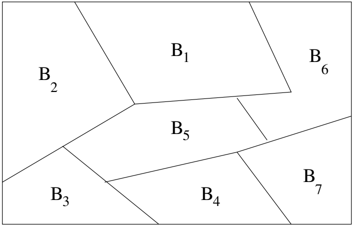

Total Probability

Example

In an experiment on human memory, participants have to memorize a set of words (B_1), numbers (B_2), and pictures (B_3). These occur in the experiment with the probabilities $P(B_1) = 0.5$, $P(B_2) = 0.4$, $P(B_3) = 0.1$.

Then participants have to recall the items (where A is the recall event). The results show that $P(A|B_1) = 0.4$, $P(A|B_2) = 0.2$, $P(A|B_3) = 0.1$. Compute $P(A)$, the probability of recalling an item.

By the theorem of total probability:

$$\begin{aligned} P(A) &= \textstyle\sum_{i=1}^{k} P(B_i)P(A|B_i) \\ &= P(B_1)P(A|B_1) + P(B_2)P(A|B_2) + P(B_3)P(A|B_3) \\ &= 0.5 \cdot 0.4 + 0.4 \cdot 0.2 + 0.1 \cdot 0.1 = 0.29 \end{aligned}$$

Bayes' Theorem

Bayes' Theorem

If $B_1, B_2, \ldots, B_k$ are a partition of S and $P(B_i) \neq 0$ for $i = 1, 2, \ldots, k$, then for any A in S such that $P(A) \neq 0$:

$$P(B_r|A) = \frac{P(B_r)P(A|B_r)}{\sum_{i=1}^{k} P(B_i)P(A|B_i)}$$

This can be simplified by renaming $B_r = B$ and by substituting $P(A) = \sum_{i=1}^{k} P(B_i)P(A|B_i)$ (theorem of total probability):

Bayes' Theorem (simplified)

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

Bayes' Theorem

Example

Reconsider the memory example. What is the probability that an item that is correctly recalled (A) is a picture (B_3)?

By Bayes' theorem:

$$\begin{aligned} P(B_3|A) &= \frac{P(B_3)P(A|B_3)}{\sum_{i=1}^{k} P(B_i)P(A|B_i)} \\ &= \frac{0.1 \cdot 0.1}{0.29} = 0.0345 \end{aligned}$$

The process of computing $P(B|A)$ from $P(A|B)$ is sometimes called *Bayesian inversion.*

Background

Let's look at an application of Bayes' theorem to the analysis of cognitive processes. First we need to introduce some data.

Research on human decision making investigates, e.g., how physicians make a *medical diagnosis* (Casscells et al. 1978):

Example

If a test to detect a disease whose prevalence is 1/1000 has a false-positive rate of 5%, what is the chance that a person found to have a positive result actually has the disease, assuming you know nothing about the person's symptoms or signs?

Background

Most frequent answer: 95%

Reasoning: if false-positive rate is 5%, then test will be correct 95% of the time.

Correct answer: 2%

Reasoning: assume you test 1000 people; the test will be positive in 50 cases (5%), but only one person actually has the disease. Hence the chance that a person with a positive result has the disease is $1/50 = 2\%$.

Only 12% of subjects give the correct answer.

Mathematics underlying the correct answer: Bayes' Theorem.

Bayes' Theorem

We need to think about Bayes' theorem slightly differently to apply it to this problem (and the terms have special names now):

Bayes' Theorem (for hypothesis testing)

Given a hypothesis h and data D which bears on the hypothesis:

$$P(h|D) = \frac{P(D|h)P(h)}{P(D)}$$

$P(h)$: independent probability of h: *prior probability*
$P(D)$: independent probability of D
$P(D|h)$: conditional probability of D given h: *likelihood*
$P(h|D)$: conditional probability of h given D: *posterior probability*

Application to Diagnosis

In Casscells et al.'s (1978) examples, we have the following:

- h: person tested has the disease;
- $\bar{h}$: person tested doesn't have the disease;
- D: person tests positive for the disease.

The following probabilities are known:

$$P(h) = 1/1000 = 0.001 \qquad P(\bar{h}) = 1 - P(h) = 0.999$$
$$P(D|\bar{h}) = 5\% = 0.05 \qquad P(D|h) = 1 \text{ (assume perfect test)}$$

Compute the probability of the data (rule of total probability):

$$P(D) = P(D|h)P(h) + P(D|\bar{h})P(\bar{h}) = 1 \cdot 0.001 + 0.05 \cdot 0.999 = 0.05095$$

Compute the probability of correctly detecting the illness:

$$P(h|D) = \frac{P(h)P(D|h)}{P(D)} = \frac{0.001 \cdot 1}{0.05095} = 0.01963$$

Base Rate Neglect

Base rate: the probability of the hypothesis being true in the absence of any data (i.e., prior probability).

Base rate neglect: people have a tendency to ignore base rate information (see Casscells et al.'s (1978) experimental results).

- base rate neglect has been demonstrated in a number of experimental situations;
- often presented as a fundamental bias in decision making;
- however, experiments show that subjects use base rates in certain situations;
- it has been argued that base rate neglect is only occurs in artificial or abstract mathematical situations.

Base Rates and Experience

Potential problems with in Casscells et al.'s (1978) study:

- subjects were simply told the statistical facts;
- they had no first-hand experience with the facts (through exposure to many applications of the test);
- providing subjects with experience has been shown to reduce or eliminate base rate neglect.

Medin and Edelson (1988) tested the role of experience on decision making in medical diagnosis.

Base Rates and Experience

Medin and Edelson (1988) trained subjects on a diagnosis task in which diseases varied in frequency:

- subjects were presented with pairs of symptoms and had to select one of six diseases;
- feedback was provided so that they learned symptom/disease associations;
- base rates of the diseases were manipulated;
- once subjects had achieved perfect diagnosis accuracy, they entered the transfer phase;
- subjects now made diagnoses for combinations of symptoms they had not seen before; *made use of base rate information.*

Summary

- Conditional probability: $P(B|A) = \frac{P(A \cap B)}{P(A)}$;
- independence: $P(B \cap A) = P(A)P(B)$.
- rule of total probability: $P(A) = \sum_{i=1}^{k} P(B_i)P(A|B_i)$;
- Bayes' theorem: $P(B|A) = \frac{P(B)P(A|B)}{P(A)}$;
- there are many applications of Bayes' theorem in cognitive science (here: medical diagnosis);
- base rate neglect: experimental subjects ignore information about prior probability.

References

Casscells, W., A. Schoenberger, and T. Grayboys. 1978. Interpretation by physicians of clinical laboratory results. *New England Journal of Medicine* 299(18):999–1001.

Medin, D. L. and S. M. Edelson. 1988. Problem structure and the use of base-rate information from experience. *Journal of Experimental Psychology: General* 117(1):68–85.