Cognitive Modeling

Lecture 11: Models of Decision Making

Sharon Goldwater

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

February 15, 2010

Reading: Cooper (2002: Ch. 6, Secs. 6.1,6.2).

Decision Making

How do people make decisions? For example,

- Medicine: Which disease to diagnose?
- Business: Where to invest? Whom to trust?
- Law: Whether to convict?
- Admissions/hiring: Whom to accept?

In all these cases, two kinds of information is used:

- Background knowledge (prevalence of disease, previous experience with business partner, historical rates of return in market, etc).
- Specific information about this case (test results, facial expressions and tone of voice, company business reports, etc)

Decision Making

Example question from a study of decision-making for 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?

Decision Making

Most frequent answer: 95%

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

Correct answer: about 2%

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

Only 12% of subjects give the correct answer.

Mathematics underlying the correct answer: Bayes' Theorem.

Bayes' Theorem

To analyze the answers that subjects give, we need:

Bayes' Theorem

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*

We also need the *rule of total probability.*

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 $B_1 \cup B_2 \cup \ldots \cup B_k = S$.

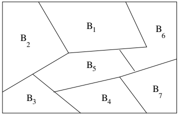

Application of Bayes' Theorem

In Casscells et al.'s (1978) example, we have:

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

$P(h) = 1/1000 = 0.001 \quad P(\bar{h}) = 1 - P(h) = 0.999$
$P(D|\bar{h}) = 5\% = 0.05 \quad P(D|h) = 1$ (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., $P(h)$).

Base rate neglect: people tend to ignore/discount base rate information (as in Casscells et al.'s (1978) experiments).

- has been demonstrated in a number of experimental situations;
- often presented as a fundamental bias in decision making.

Does this mean people are irrational/sub-optimal?

Base Rates and Experience

Casscells et al.'s (1978) study is abstract and artificial. Other studies show that

- data presentation affects performance (1 in 20 vs. 5%).
- direct experience of statistics (through exposure to many outcomes) affects performance.
- task description affects performance ("psychological test" versus "statistics test" when assessing personal profiles).

Suggests subjects may be interpreting questions and determining priors in ways other than experimenters assume.

- Ex: is it reasonable to assume that a medical test is given if there is no evidence of disease?

Base Rates and Experience

First, evidence that subjects can use base rates: diagnosis task of Medin and Edelson (1988).

- *Training phase:*
  - 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;
  - different diseases had different base rates;
  - ended when subjects had achieved perfect diagnosis accuracy.
- *Transfer phase:*
  - subjects were tested on single symptoms and combinations they had not seen in the training phase.

Experimental Data

Structure of Medin and Edelson's (1988) experiment:

Symptoms	Disease	No. of trials
a & b	1	3 trials
a & c	2	1 trial
d & e	3	3 trials
d & f	4	1 trial
g & h	5	3 trials
g & i	6	1 trial

Symptoms a, d, g are *imperfect predictors;* symptoms b, c, e, f, h, i are *perfect predictors.*

Diseases 1, 3, 5 are *high frequency,* diseases 2, 4, 6 are *low frequency.*

Experimental Results

Results in transfer phase:

- when presented with a high frequency perfect predictor (e.g., b), 81.2% responses for correct disease (e.g., 1);
- when presented with a low frequency perfect predictor (e.g., c), 92.7% responses for correct disease (e.g., 3).

Indicates: symptom/disease associations acquired correctly.

- when presented with a high freq. imperf. predictor (e.g., a), 78.1% responses for correct high freq. disease (e.g., 1), 14.6% responses for correct low freq. disease (e.g., 2).

Indicates: base rate information is used.

Modeling Decision Making

Medin and Edelson's (1988) results suggest that decision-making could be based on Bayesian reasoning.

Cooper (2002: Ch. 6) presents a Cogent model:

- knowledge base contains frequency information about symptoms and diseases, acquired by counting.
- computes predictions using Bayes' Rule.

Problems: no plausible model of learning, prediction fails in transfer phase when symptoms conflict.

Cooper's (2002) Model

In transfer phase, subjects are presented with symptoms s and have to predict a disease d. Model does so using Bayes' Rule:

$$P(d|s) = \frac{P(s|d)P(d)}{P(s)}$$

$P(s|d)$, $P(d)$, and $P(s)$ are determined from frequencies observed in the training phase.

Cooper's (2002) Model

Compute probabilities from frequency counts:

$$\begin{array}{lll} P(d_1) = 3/12 & P(a|d_1) = 3/3 & P(a) = 4/12 \\ P(d_2) = 1/12 & P(b|d_1) = 3/3 & P(b) = 3/12 \\ \dots & P(a|d_2) = 1/1 & P(c) = 1/12 \\ & P(c|d_2) = 1/1 & \dots \\ & \dots & \end{array}$$

Compute predictions given a single symptom:

$$P(d_1|a) = \frac{P(a|d_1)P(d_1)}{P(a)} = \frac{(3/3)(3/12)}{4/12} = .75$$

$$P(d_1|b) = \frac{P(b|d_1)P(d_1)}{P(b)} = \frac{(3/3)(3/12)}{3/12} = 1$$

Similarly, $P(d_2|a) = .25$, $P(d_2|c) = 1$.

Cooper's (2002) Model

What about conflicting symptoms?

$$P(d_1|b,c) = \frac{P(b,c|d_1)P(d_1)}{P(b,c)} = \frac{(0)(3/12)}{0} = ??$$

- Cooper uses this problem with conflicting symptoms to argue against the Bayesian model.
- However, Cooper's implementation takes a naive view of probability.

Uncertainty

In probabilistic models, there are two sources of uncertainty.

1. Given a known distribution $P(X)$, the *outcome* is uncertain.

   e.g., $P(X = a) = .3, P(X = b) = .7$

2. In general, the *distribution* itself is uncertain, as it must be estimated from data.

   e.g., $P(X = a) \approx .3$ or $P(X = a) = .3 \pm .01$

Cooper's model fails to consider the second kind of uncertainty.

Probability $\neq$ Counting

Thought experiment: what is a good estimate of $P(H)$ in each case?

1. I pick up a coin off the street, and start flipping.
   a. Flip 10 times: 4T, 6H.
   b. Flip 100 times: 40T, 60H.
2. I have a coin in my pocket, and I tell you it's weighted. I pull it out and start flipping.
   a. Flip 10 times: 4T, 6H.
   b. Flip 100 times: 40T, 60H.

Frequentist Statistics

Standard *frequentist* statistics strives to be objective. Interprets probabilities as proportions of infinite number of trials.

- Probabilities are estimated from repeated observations.
- More observations $\rightarrow$ more accurate estimation.
- Focuses on ruling out hypotheses, not estimating their probabilities.
- Ex: Data = (3T, 7H). Estimate $P(H) = .7$, but margin for error is large, does not rule out $P(H) = .5$.

Used widely in controlled scientific experiments.

Bayesian Statistics

Bayesian interpretation of probabilities is that they reflect *degrees of belief*, not frequencies.

- Belief can be influenced by frequencies: observing many outcomes changes one's belief about future outcomes.
- Belief can be influenced by other factors: structural assumptions, knowledge of similar cases, complexity of hypotheses, etc.
- Hypotheses can be assigned probabilities.

Works much better for cognitive modeling.

Bayes' Theorem, Again

Bayesian interpretation of Bayes' theorem:

Bayes' Theorem

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

$P(h)$: *prior probability* reflects plausibility of h regardless of data.
$P(D|h)$: *likelihood* reflects how well h explains the data.
$P(h|D)$: *posterior probability* reflects plausibility of h after taking data into account.

Note that $P(h)$ may differ from the "base rate" (which implies simply counting).

Discussion

Reconsider modeling and experimental evidence:

- Cooper's model fails not because of Bayes' rule, but because probabilities are equated with relative frequencies.
- Similarly, evidence of base rate neglect fails to consider factors besides frequency that might affect prior probabilities.
- Next class: more detail on Bayesian methods and relationships to cognitive modeling.

Summary

- Bayes' theorem can be applied to human decision making;
- early experimental results seemed to indicate that subjects ignore prior probabilities: base rate neglect;
- however, more recent studies show that subject can learn base rate information from experience;
- rational analysis using Bayesian view suggests that equating probabilities with relative frequencies is the problem;
- subjects may use additional information to determine prior probabilities.

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.

Cooper, Richard P. 2002. *Modelling High-Level Cognitive Processes*. Lawrence Erlbaum Associates, Mahwah, NJ.

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.