Bayesian Statistics

VALENCIA 8

by Alyssa Curran

ISBA 2006

Coming Up…

- Background and history of Bayes' theorem and Bayesian statistical inference
- Definition and explanation of Bayes' theorem
- Conflict between Bayesians and Frequentists
- Derivation from conditional probabilities
- Alternative forms of Bayes' theorem:
  - ~ odds and liklihood ratio
  - ~ probability densities
  - ~ extension: more than two variables
- Examples
  - ~ conditional probabilities
  - ~ the Monty Hall problem

Background and History

- Bayes' theorem was devised by Reverend Thomas Bayes (1702-1761)
- He studied how to compute a distribution for the parameter of a binomial distribution
- His work was published in *An Essay towards solving a Problem in the Doctrine of Chances*, made public by his friend Richard Price after his death
- These results were replicated by Pierre Simon Laplace in an essay in 1774, though he was unaware of Bayes' work
- Bayes' theorem does not mention the order in which the events occur, it measures their correlation rather than cause and effect
- The preliminary results of Bayes' essay imply the theorem, but Bayes did not actually focus on that result

Conflict between Bayesians and Frequentists

* Frequentists and Bayesians disagree about the types of quantities to which probabilities should be assigned in applications

* Frequentists assign probabilities to random events according to their frequencies of occurrence or to subsets of populations as proportions of the whole

* Bayesians assign probabilities to propositions that are uncertain

* Research has been done in which to develop new procedures to allow an agreement between the Bayesians' and Frequentists' approaches to testing hypotheses

Definition and an Explanation of Bayes' Theorem

- Bayes' theorem is a result of probability theory which relates conditional probabilities
- If A and B denote two events, $P(A|B)$ denotes the conditional probability of A occuring, given that B occurs
- The two conditional probabilities $P(A|B)$ and $P(B|A)$ are related with Bayes' theorem
- An application of Bayes' theorem is statistical inference, in which evidence or observations are used to update or to newly infer the probability that a hypothesis may be true. Bayes' theorem provides a rule for strengthening the evidence-based beliefs
- The theorem relates the conditional and marginal probabilities of events A and B:

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

~ $P(A)$ is the prior (marginal) probability of A

~ $P(A|B)$ is the conditional probability of A, given B (posterior probability)

~ $P(B|A)$ is the conditional probability of B, given A

~ $P(B)$ is the prior (marginal) probability of B (normalizing constant)

Derivation from Conditional Probabilities

To derive Bayes' theorem, start with the definition of conditional probability:

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

Rearranging, we are given:

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

~ this is called the product rule for probabilities

Divide both sides of the equation by P(B):

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

And we have Bayes' theorem!

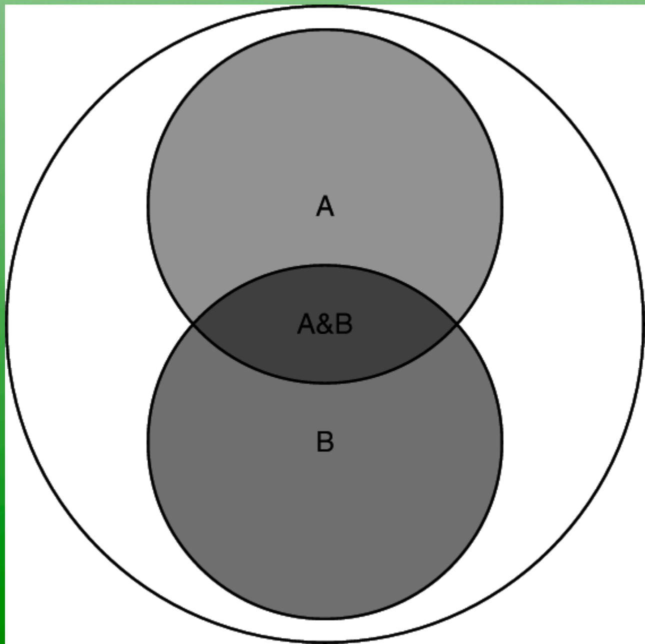

Alternative forms of Bayes' Theorem

* Bayes' law in terms of an odds and liklihood ratio:

Odds: $O(A|B) = O(A) \cdot \Lambda(A|B)$

Liklihood: $\Lambda(A|B) = \dfrac{L(A|B)}{L(A^c|B)} = \dfrac{P(B|A)}{P(B|A^c)}$

* In terms of probability densities:

$$f(x|y) = \frac{f(y,x)}{f(y)} = \frac{f(y|x)\, f(x)}{f(y)}$$

* Extension to problems with more than 2 variables:

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

Examples

Conditional Probabilities:

~ We have two bowls of m&m's – #1 has 10 purple and 30 green and #2 has 20 of each color

~ Russ picks a bowl randomly, and then picks an m&m randomly. The color turns out to be green. How probable is it that Russ picked out of bowl #1? (the probability Russ picked bowl #1 given he has a green m&m?)

~ Event A is that Russ picked out of #1, and event B is that he picked a green m&m

$P(A) = 0.5$

$P(B) = 50/80 = 0.625$

$P(B|A) = 30/40 = 0.75$

~ We can compute the probability of Russ selecting bowl #1, given he got a green m&m by substituting in the values:

$$P(A|B) = \frac{P(B|A)\,P(A)}{P(B)} = \frac{0.75 \times 0.5}{0.625} = 0.6$$

~ As we would intuitively expect, the probability is more than ½ ☺

Examples…continued

The Monty Hall Problem:

~ We have 3 doors (red, green, blue) behind one of which is a prize

~ We pick the red door, which is not opened until later

~ The presenter opens the green door (who is not permitted to open the door with the prize behind it or the door we have picked) to reveal no prize

~ Should we change our mind about our initial choice of red to blue?

~ We need to find the probabilities of the prize being behind the red, green, and blue doors (A_r, A_g, and A_b):

- $P(A_r) = P(A_g) = P(A_b) = 1/3$
- B: presenter opens green door, $P(B) = ½$
- If prize is behind red door, $P(B|A_r) = ½$
- If prize is behind green door, $P(B|A_g) = 0$
- If prize is behind blue door, $P(B|A_b) = 1$

So, $P(A_r|B) = \dfrac{P(B|A_r)\,P(A_r)}{P(B)} = \dfrac{½ \times 1/3}{½} = 1/3$

$P(A_g|B) = \dfrac{P(B|A_g)\,P(A_g)}{P(B)} = \dfrac{0 \times 1/3}{½} = 0$

$P(A_b|B) = \dfrac{P(B|A_b)\,P(A_b)}{P(B)} = \dfrac{1 \times 1/3}{½} = 2/3$

The correct choice based on probability is to switch to the blue door ☺

In Summary...

We learned about:

- **The background and history of Bayes' theorem and Bayesian statistical inference**
- **The definition and explanation of Bayes' theorem**
- **The conflict between Bayesians and Frequentists**
- **The derivation from conditional probabilities**
- **3 Alternative forms of Bayes' theorem**
- **2 Examples**

The End