6.3 Bayes' Theorem

Introduction

There are many times when we want to assess the probability that a particular event occurs on the basis of partial evidence. For example, suppose we know the percentage of people who have a particular disease for which there is a very accurate diagnostic test. People who test positive for this disease would like to know the likelihood that they actually have the disease. In this section we introduce a result that can be used to determine this probability, namely, the probability that a person has the disease given that they test positive for it. To use this result, we will need to know the percentage of people who do not have the disease but test positive for it and the percentage of people who have the disease but test negative for it.

Similarly, suppose we know the percentage of incoming e-mail messages that are spam. We will see that we can determine the likelihood that an incoming e-mail message is spam using the occurrence of words in the message. To determine this likelihood, we need to know the percentage of incoming messages that are spam, the percentage of spam messages in which each of these words occurs, and the percentage of messages that are not spam in which each of these words occurs.

The result that we can use to answer questions such as these is called Bayes' Theorem and dates back to the eighteenth century. In the past two decades, Bayes' Theorem has been extensively applied to estimate probabilities based on partial evidence in areas as diverse as medicine, law, machine learning, engineering, and software development.

Bayes' Theorem

We illustrate the idea behind Bayes' Theorem with an example that shows that when extra information is available, we can derive a more realistic estimate that a particular event occurs. That is, suppose we know $p(F)$, the probability that an event F occurs, but we have knowledge that an event E occurs. Then the conditional probability that F occurs given that E occurs, $p(F \mid E)$, is a more realistic estimate than $p(F)$ that F occurs. In Example 1 we will see that we can find $p(F \mid E)$ when we know $p(F)$, $p(E \mid F)$, and $p(E \mid \overline{F})$.

EXAMPLE 1 We have two boxes. The first contains two green balls and seven red balls; the second contains four green balls and three red balls. Bob selects a ball by first choosing one of the two boxes at random. He then selects one of the balls in this box at random. If Bob has selected a red ball, what is the probability that he selected a ball from the first box?

Extra Examples

Solution: Let E be the event that Bob has chosen a red ball; $\overline{E}$ is the event that Bob has chosen a green ball. Let F be the event that Bob has chosen a ball from the first box; $\overline{F}$ is the event that Bob has chosen a ball from the second box. We want to find $p(F \mid E)$, the probability that the ball Bob selected came from the first box, given that it is red. By the definition of conditional probability, we have $p(F \mid E) = p(F \cap E)/p(E)$. Can we use the information provided to determine both $p(F \cap E)$ and $p(E)$ so that we can find $p(F \mid E)$?

First, note that because the first box contains seven red balls out of a total of nine balls, we know that $p(E \mid F) = 7/9$. Similarly, because the second box contains three red balls out of a total of seven balls, we know that $p(E \mid \overline{F}) = 3/7$. We assumed that Bob selects a box at random, so $p(F) = p(\overline{F}) = 1/2$. Because $p(E \mid F) = p(E \cap F)/p(F)$, it follows that $p(E \cap F) = p(E \mid F)p(F) = \frac{7}{9} \cdot \frac{1}{2} = \frac{7}{18}$ [as we remarked earlier, this is one of the quantities we need to find to determine $p(F \mid E)$]. Similarly, because $p(E \mid \overline{F}) = p(E \cap \overline{F})/p(\overline{F})$, it follows that $p(E \cap \overline{F}) = p(E \mid \overline{F})p(\overline{F}) = \frac{3}{7} \cdot \frac{1}{2} = \frac{3}{14}$.

We can now find $p(E)$. Note that $E = (E \cap F) \cup (E \cap \overline{F})$, where $E \cap F$ and $E \cap \overline{F}$ are disjoint sets. (If x belongs to both $E \cap F$ and $E \cap \overline{F}$, then x belongs to both F and $\overline{F}$, which is impossible.) It follows that

$$p(E) = p(E \cap F) + p(E \cap \overline{F}) = \frac{7}{18} + \frac{3}{14} = \frac{49}{126} + \frac{27}{126} = \frac{76}{264} = \frac{38}{63}.$$

We have now found both $p(F \cap E) = 7/18$ and $p(E) = 38/63$. We conclude that

$$p(F \mid E) = \frac{p(F \cap E)}{p(E)} = \frac{7/18}{38/63} = \frac{49}{76} \approx 0.645.$$

Before we had any extra information, we assumed that the probability that Bob selected the first box was 1/2. However, with the extra information that the ball selected at random is red, this probability has increased to approximately 0.645. That is, the probability that Bob selected a ball from the first box increased from 1/2, when no extra information was available, to 0.645 once we knew that the ball selected was red. ◂

Using the same type of reasoning as in Example 1, we can find the conditional probability that an event F occurs, given that an event E has occurred, when we know $p(E \mid F)$, $p(E \mid \overline{F})$, and $p(F)$. The result we can obtain is called **Bayes' Theorem;** it is named after Thomas Bayes, an eighteenth-century British mathematician and minister who introduced this result.

THEOREM 1 **BAYES' THEOREM** Suppose that E and F are events from a sample space S such that $p(E) \neq 0$ and $p(F) \neq 0$. Then

$$p(F \mid E) = \frac{p(E \mid F)p(F)}{p(E \mid F)p(F) + p(E \mid \overline{F})p(\overline{F})}.$$

Proof: The definition of conditional probability tells us that $p(F \mid E) = p(E \cap F)/p(E)$ and $p(E \mid F) = p(E \cap F)/p(F)$. Therefore, $p(E \cap F) = p(F \mid E)p(E)$ and $p(E \cap F) = p(E \mid F)p(F)$. Equating these two expressions for $p(E \cap F)$ shows that

$$p(F \mid E)p(E) = p(E \mid F)p(F).$$

Dividing both sides by $p(E)$, we find that

$$p(F \mid E) = \frac{p(E \mid F)p(F)}{p(E)}.$$

To complete the proof, we show that $p(E) = p(E \mid F)p(F) + p(E \mid \overline{F})p(\overline{F})$. First, note that $E = E \cap S = E \cap (F \cup \overline{F}) = (E \cap F) \cup (E \cap \overline{F})$. Furthermore, $E \cap F$ and $E \cap \overline{F}$ are disjoint, because if $x \in E \cap F$ and $x \in E \cap \overline{F}$, then $x \in F \cap \overline{F} = \emptyset$. Consequently, $p(E) =$

$p(E \cap F) + p(E \cap \overline{F})$. We have already shown that $p(E \cap F) = p(E \mid F)p(F)$. Moreover, we have $p(E \mid \overline{F}) = p(E \cap \overline{F})/p(\overline{F})$, which shows that $p(E \cap \overline{F}) = p(E \mid \overline{F})p(\overline{F})$. It follows that

$$p(E) = p(E \cap F) + p(E \cap \overline{F}) = p(E \mid F)p(F) + p(E \mid \overline{F})p(\overline{F}).$$

To complete the proof we insert this expression for $p(E)$ into the equation $p(F \mid E) = p(E \mid F)p(F)/p(E)$. We have proved that

Links

$$p(F \mid E) = \frac{p(E \mid F)p(F)}{p(E \mid F)p(F) + p(E \mid \overline{F})p(\overline{F})}. \qquad \triangleleft$$

Bayes' Theorem can be used to assess the probability that someone testing positive for a disease actually has this disease. The results obtained from Bayes' Theorem are often somewhat surprising, as Example 2 shows.

EXAMPLE 2 Suppose that one person in 100,000 has a particular rare disease for which there is a fairly accurate diagnostic test. This test is correct 99% of the time when given to someone with the disease; it is correct 99.5% of the time when given to someone who does not have the disease. Given this information can we find

(a) the probability that someone who tests positive for the disease has the disease?
(b) the probability that someone who tests negative for the disease does not have the disease?

Should someone who tests positive be very concerned that he or she has the disease?

Solution: (a) Let F be the event that a person has the disease, and let E be the event that this person tests positive for the disease. We want to compute $p(F \mid E)$. To use Bayes' Theorem to compute $p(F \mid E)$ we need to find $p(E \mid F)$, $p(E \mid \overline{F})$, $p(F)$, and $p(\overline{F})$.

We know that one person in 100,000 has this disease, so that $p(F) = 1/100{,}000 = 0.00001$ and $p(\overline{F}) = 1 - 0.00001 = 0.99999$. Because someone who has the disease tests positive 99% of the time, we know that $p(E \mid F) = 0.99$; this is the probability of a true positive, that someone with the disease tests positive. We also know that $p(\overline{E} \mid F) = 0.01$; this is the probability of a false negative, that someone who has the disease tests negative. Furthermore, because someone who does not have the disease tests negative 99.5% of the time, we know that $p(\overline{E} \mid \overline{F}) = 0.995$. This is the probability of a true negative, that someone without the disease tests negative. Finally, we know that $p(E \mid \overline{F}) = 0.005$; this is the probability of a false positive, that someone without the disease tests positive.

The probability that someone who tests positive for the disease actually has the disease is $p(F \mid E)$. By Bayes' Theorem, we know that

$$\begin{aligned} p(F \mid E) &= \frac{p(E \mid F)p(F)}{p(E \mid F)p(F) + p(E \mid \overline{F})p(\overline{F})} \\ &= \frac{(0.99)(0.00001)}{(0.99)(0.00001) + (0.005)(0.99999)} \approx 0.002. \end{aligned}$$

This means that only 0.2% of people who test positive for the disease actually have the disease. Because the disease is extremely rare, the number of false positives on the diagnostic test is far greater than the number of true positives, making the percentage of people who test positive who actually have the disease extremely small. People who test positive for the diseases should not be overly concerned that they actually have the disease.

(b) The probability that someone who tests negative for the disease does not have the disease is $p(\overline{F} \mid \overline{E})$. By Bayes' Theorem, we know that

$$p(\overline{F} \mid \overline{E}) = \frac{p(\overline{E} \mid \overline{F})p(\overline{F})}{p(\overline{E} \mid \overline{F})p(\overline{F}) + p(\overline{E} \mid F)p(F)}$$

$$= \frac{(0.995)(0.99999)}{(0.995)(0.99999) + (0.01)(0.00001)} \approx 0.9999999.$$

Consequently, 99.99999% of the people who test negative really do not have the disease. ◀

Note that in the statement of Bayes' Theorem, the events F and $\overline{F}$ are mutually exclusive and cover the entire sample space S (that is, $F \cup \overline{F} = S$). We can extend Bayes' Theorem to any collection of mutually exclusive events that cover the entire sample space S, in the following way.

THEOREM 2 **GENERALIZED BAYES' THEOREM** Suppose that E is an event from a sample space S and that $F_1, F_2, \ldots, F_n$ are mutually exclusive events such that $\bigcup_{i=1}^{n} F_i = S$. Assume that $p(E) \neq 0$ and $p(F_i) \neq 0$ for $i = 1, 2, \ldots, n$. Then

$$p(F_j \mid E) = \frac{p(E \mid F_j)p(F_j)}{\sum_{i=1}^{n} p(E \mid F_i)p(F_i)}.$$

We leave the proof of this generalized version of Bayes' Theorem as Exercise 17.