POLS/CSSS 510 · Maximum Likelihood Methods

REVIEW OF BAYES RULE

Christopher Adolph

Department of Political Science
Center for Statistics and the Social Sciences
University of Washington, Seattle

Deriving Bayes Rule

Using the definitions of conditional probability and joint probability,
we can derive one of the most famous and useful results in probability theory

This result was first proved by Thomas Bayes,
an otherwise obscure 18th century minister living in Kent

$$\text{conditional probability} \quad = \quad \frac{\text{joint probability}}{\text{marginal probability}}$$

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

$$\mathrm{P}(a|b) = \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(b)}$$

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

$$\mathrm{P}(a|b) = \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(b)}$$

$$\mathrm{P}(b|a) = \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(a)}$$

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

$$P(a|b) = \frac{P(a \cap b)}{P(b)}$$

$$P(b|a) = \frac{P(a \cap b)}{P(a)}$$

$$P(b|a)P(a) = P(a \cap b)$$

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

$$\begin{aligned}
\mathrm{P}(a|b) &= \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(b)} \\
\mathrm{P}(b|a) &= \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(a)} \\
\mathrm{P}(b|a)\mathrm{P}(a) &= \mathrm{P}(a \cap b) \\
\mathrm{P}(a|b)\mathrm{P}(b) &= \mathrm{P}(a \cap b)
\end{aligned}$$

$$\text{conditional probability} = \frac{\text{joint probability}}{\text{marginal probability}}$$

$$\begin{aligned}
\mathrm{P}(a|b) &= \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(b)} \\
\mathrm{P}(b|a) &= \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(a)} \\
\mathrm{P}(b|a)\mathrm{P}(a) &= \mathrm{P}(a \cap b) \\
\mathrm{P}(a|b)\mathrm{P}(b) &= \mathrm{P}(a \cap b) \\
\mathrm{P}(a|b)\mathrm{P}(b) &= \mathrm{P}(b|a)\mathrm{P}(a)
\end{aligned}$$

$$
\begin{aligned}
\text{conditional probability} &= \frac{\text{joint probability}}{\text{marginal probability}} \\
\mathrm{P}(a|b) &= \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(b)} \\
\mathrm{P}(b|a) &= \frac{\mathrm{P}(a \cap b)}{\mathrm{P}(a)} \\
\mathrm{P}(b|a)\mathrm{P}(a) &= \mathrm{P}(a \cap b) \\
\mathrm{P}(a|b)\mathrm{P}(b) &= \mathrm{P}(a \cap b) \\
\mathrm{P}(a|b)\mathrm{P}(b) &= \mathrm{P}(b|a)\mathrm{P}(a) \\
\mathrm{P}(a|b) &= \frac{\mathrm{P}(b|a)\mathrm{P}(a)}{\mathrm{P}(b)}
\end{aligned}
$$

This is Bayes Rule

It says the probability of a given b depends on the probability of b given a and the marginal probabilities of a and b

During a routine checkup, your doctor tells you some bad news:
you tested positive for a rare disease.

The disease affects 1 in 10,000 people,
and the test has a 99% effectiveness rate.

What are the chances you have the disease?

Steps to solve a probability problem

1. Identify the possible events
2. Identify the quantity of interest in terms of probability of specific events
3. Collect all the probabilities you know
4. Use the rules of probability to calculate what you want to know from what you do know

Identify the possible events

You either have the disease or not have the disease

You either test positive or negative

This leads to four possible combinations which comprise the whole sample space:

- have the disease and test positive
- have the disease and test negative
- have no disease and test positive
- have no disease and test negative

Venn Diagram of disease and test events

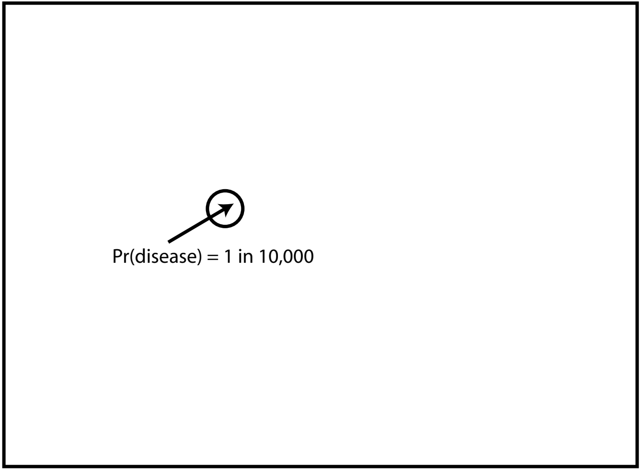

Let's start with the event of having the disease

Venn Diagram of disease and test events

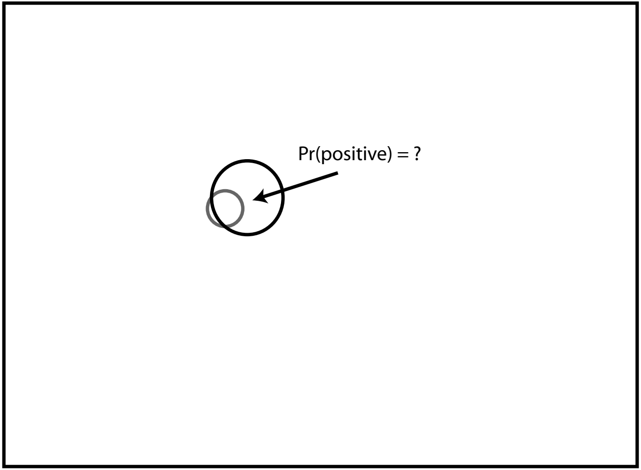

Testing positive is a separate, possibly joint event with having the disease

Venn Diagram of disease and test events

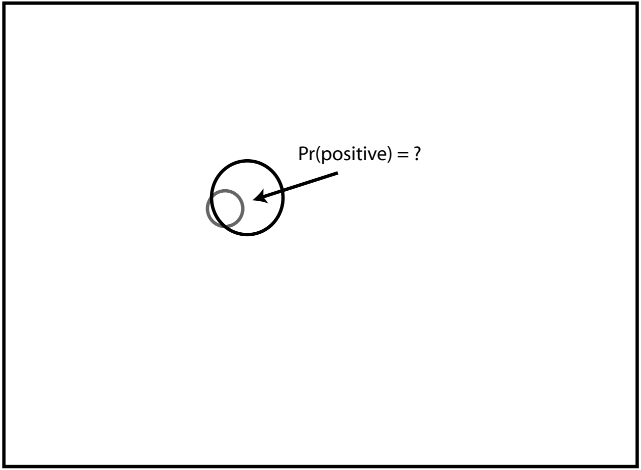

Testing positive is a separate, possibly joint event with having the disease

We don't yet know the marginal probability of testing positive

Venn Diagram of disease and test events

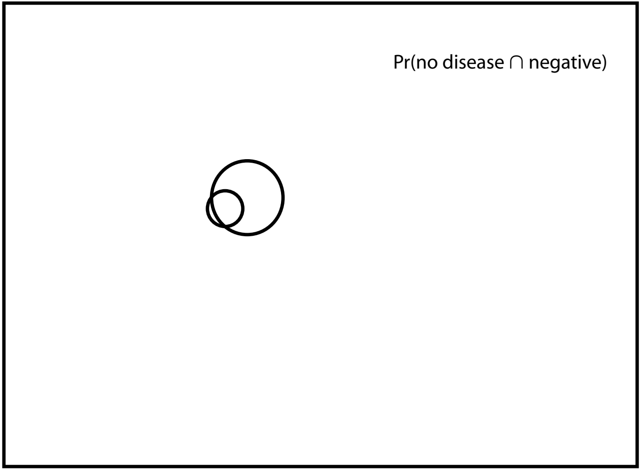

The area outside both circles is the probability of neither having the disease nor testing positive

Venn Diagram of disease and test events

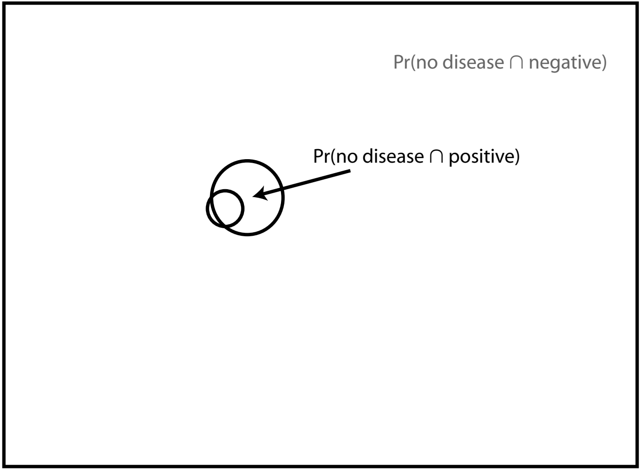

This large crescent shape is the probability of testing positive even though you don't have the disease

Venn Diagram of disease and test events

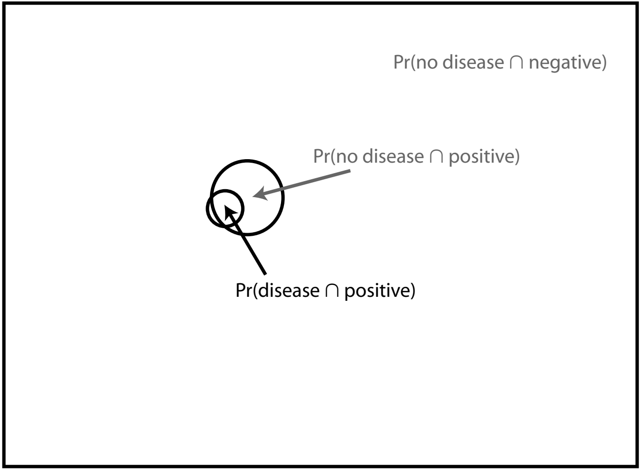

This intersection is the probability of testing positive and having the disease

Venn Diagram of disease and test events

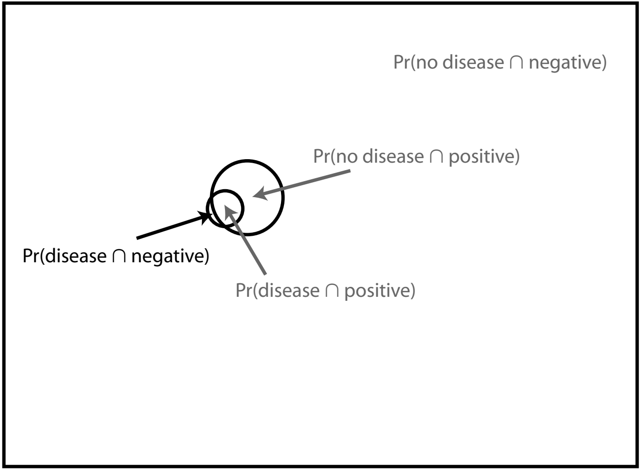

And this crescent is the probability of having the disease but failing to detect it

Identify the quantity of interest

Clearly, we want to know the possibility you have the disease *given* a positive test result

The simple probability of having a disease isn't enough –
you are worried because you have new information:
you had a positive test

In formal terms, we want to find:

$$\Pr(\text{disease}|\text{positive})$$

This is not the same as either $\Pr(\text{disease})$ or $\Pr(\text{positive})$

How do we find it? Let's start with what we know

Collect all the probabilities you know

We know how likely a random person is to have the disease:

$$\begin{aligned} \text{Pr}(\text{disease}) &= \text{1 in 10,000} \\ &= 0.01\% \\ &= 0.0001 \end{aligned}$$

Collect all the probabilities you know

We know how likely a random person is to have the disease:

$$\begin{aligned}\Pr(\text{disease}) &= \text{1 in 10,000} \\ &= 0.01\% \\ &= 0.0001\end{aligned}$$

From this, we can calculate the complement:

$$\begin{aligned}\Pr(\text{no disease}) &= 1 - \Pr(\text{disease}) \\ &= \text{9,999 in 10,000} \\ &= 99.99\% \\ &= 0.9999\end{aligned}$$

Collect all the probabilities you know

We also know some conditional probabilities. The test is "99% effective," meaning it has a 99% probability of detecting the correct disease status, so:

$$
\begin{aligned}
\Pr(\text{positive}|\text{disease}) &= 0.99 \\
&= 99 \text{ in } 100
\end{aligned}
$$

Collect all the probabilities you know

We also know some conditional probabilities. The test is "99% effective," meaning it has a 99% probability of detecting the correct disease status, so:

$$\begin{aligned} \Pr(\text{positive}|\text{disease}) &= 0.99 \\ &= 99 \text{ in } 100 \\ \Pr(\text{negative}|\text{disease}) &= 0.01 \\ &= 1 \text{ in } 100 \end{aligned}$$

Collect all the probabilities you know

We also know some conditional probabilities. The test is "99% effective," meaning it has a 99% probability of detecting the correct disease status, so:

$$
\begin{aligned}
\Pr(\text{positive} \mid \text{disease}) &= 0.99 \\
&= 99 \text{ in } 100 \\
\Pr(\text{negative} \mid \text{disease}) &= 0.01 \\
&= 1 \text{ in } 100 \\
\Pr(\text{positive} \mid \text{no disease}) &= 0.01 \\
&= 1 \text{ in } 100
\end{aligned}
$$

Collect all the probabilities you know

We also know some conditional probabilities. The test is "99% effective," meaning it has a 99% probability of detecting the correct disease status, so:

$$
\begin{aligned}
\text{Pr}(\text{positive}|\text{disease}) &= 0.99 \\
&= 99 \text{ in } 100 \\
\text{Pr}(\text{negative}|\text{disease}) &= 0.01 \\
&= 1 \text{ in } 100 \\
\text{Pr}(\text{positive}|\text{no disease}) &= 0.01 \\
&= 1 \text{ in } 100 \\
\text{Pr}(\text{negative}|\text{no disease}) &= 0.99 \\
&= 99 \text{ in } 100
\end{aligned}
$$

What we still don't know

$$\begin{aligned} \Pr(\text{disease}|\text{positive}) &= \textbf{?} \\ \Pr(\text{no disease}|\text{positive}) &= \textbf{?} \end{aligned}$$

How are we going to calculate these?

What we still don't know

$$\begin{aligned} \Pr(\text{disease}|\text{positive}) &= \textbf{?} \\ \Pr(\text{no disease}|\text{positive}) &= \textbf{?} \end{aligned}$$

How are we going to calculate these?

Let's use Bayes Rule:

$$\Pr(\text{disease}|\text{positive}) = \frac{\Pr(\text{disease})\Pr(\text{positive}|\text{disease})}{\Pr(\text{positive})}$$

Recall that $\Pr(\text{disease}) = 0.0001$, $\Pr(\text{positive}|\text{disease}) = 0.99$, $\Pr(\text{no disease}) = 0.9999$, and $\Pr(\text{positive}|\text{no disease}) = 0.01$

$$\Pr(\text{disease}|\text{positive}) = \frac{\Pr(\text{disease})\Pr(\text{positive}|\text{disease})}{\Pr(\text{positive})}$$

Recall that $\Pr(\text{disease}) = 0.0001$, $\Pr(\text{positive}|\text{disease}) = 0.99$, $\Pr(\text{no disease}) = 0.9999$, and $\Pr(\text{positive}|\text{no disease}) = 0.01$

$$\begin{aligned}
\Pr(\text{disease}|\text{positive}) &= \frac{\Pr(\text{disease})\Pr(\text{positive}|\text{disease})}{\Pr(\text{positive})} \\
\Pr(\text{disease}|\text{positive}) &= \frac{0.0001 \times 0.99}{\Pr(\text{positive})}
\end{aligned}$$

Recall that $\Pr(\text{disease}) = 0.0001$, $\Pr(\text{positive}|\text{disease}) = 0.99$, $\Pr(\text{no disease}) = 0.9999$, and $\Pr(\text{positive}|\text{no disease}) = 0.01$

$$
\begin{aligned}
\Pr(\text{disease}|\text{positive}) &= \frac{\Pr(\text{disease})\Pr(\text{positive}|\text{disease})}{\Pr(\text{positive})} \\
\Pr(\text{disease}|\text{positive}) &= \frac{0.0001 \times 0.99}{\Pr(\text{positive})}
\end{aligned}
$$

We don't know $\Pr(\text{positive})$. How do we find it?

Note that for an exhaustive and mutually exclusive (disjoint) set of events $\mathcal{B}$, the marginal probability of some other event a can be found using

$$\Pr(a) = \sum_{\forall b_i \in \mathcal{B}} \Pr(b_i)\Pr(a|b_i)$$

Recall that $\Pr(\text{disease}) = 0.0001$, $\Pr(\text{positive}|\text{disease}) = 0.99$,
$\Pr(\text{no disease}) = 0.9999$, and $\Pr(\text{positive}|\text{no disease}) = 0.01$

$$\begin{aligned}
\Pr(\text{disease}|\text{positive}) &= \frac{\Pr(\text{disease})\Pr(\text{positive}|\text{disease})}{\Pr(\text{positive})} \\
\Pr(\text{disease}|\text{positive}) &= \frac{0.0001 \times 0.99}{\Pr(\text{positive})} \\
\Pr(\text{disease}|\text{positive}) &= \frac{0.0001 \times 0.99}{\begin{array}{l}\Pr(\text{disease})\Pr(\text{positive}|\text{disease}) \\ \quad + \Pr(\text{no disease})\Pr(\text{positive}|\text{no disease})\end{array}}
\end{aligned}$$

We don't know $\Pr(\text{positive})$. How do we find it?

Note that for an exhaustive and mutually exclusive (disjoint) set of events $\mathcal{B}$, the marginal probability of some other event a can be found using

$$\Pr(a) = \sum_{\forall b_i \in \mathcal{B}} \Pr(b_i)\Pr(a|b_i)$$

Recall that $\Pr(\text{disease}) = 0.0001$, $\Pr(\text{positive}|\text{disease}) = 0.99$, $\Pr(\text{no disease}) = 0.9999$, and $\Pr(\text{positive}|\text{no disease}) = 0.01$

$$
\begin{aligned}
\Pr(\text{disease}|\text{positive}) &= \frac{\Pr(\text{disease})\Pr(\text{positive}|\text{disease})}{\Pr(\text{positive})} \\
\Pr(\text{disease}|\text{positive}) &= \frac{0.0001 \times 0.99}{\Pr(\text{positive})} \\
\Pr(\text{disease}|\text{positive}) &= \frac{0.0001 \times 0.99}{\begin{array}{l}\Pr(\text{disease})\Pr(\text{positive}|\text{disease}) \\ \quad + \Pr(\text{no disease})\Pr(\text{positive}|\text{no disease})\end{array}} \\
\Pr(\text{disease}|\text{positive}) &= \frac{0.0001 \times 0.99}{0.0001 \times 0.99 + 0.9999 \times 0.01}
\end{aligned}
$$

Recall that $\Pr(\text{disease}) = 0.0001$, $\Pr(\text{positive}|\text{disease}) = 0.99$, $\Pr(\text{no disease}) = 0.9999$, and $\Pr(\text{positive}|\text{no disease}) = 0.01$

$$\begin{aligned}
\Pr(\text{disease}|\text{positive}) &= \frac{\Pr(\text{disease})\Pr(\text{positive}|\text{disease})}{\Pr(\text{positive})} \\
\Pr(\text{disease}|\text{positive}) &= \frac{0.0001 \times 0.99}{\Pr(\text{positive})} \\
\Pr(\text{disease}|\text{positive}) &= \frac{0.0001 \times 0.99}{\begin{array}{l}\Pr(\text{disease})\Pr(\text{positive}|\text{disease}) \\ \quad + \Pr(\text{no disease})\Pr(\text{positive}|\text{no disease})\end{array}} \\
\Pr(\text{disease}|\text{positive}) &= \frac{0.0001 \times 0.99}{0.0001 \times 0.99 + 0.9999 \times 0.01} \\
\Pr(\text{disease}|\text{positive}) &= 0.0098 \approx 1\%
\end{aligned}$$

Even though the test is "99% effective",
only 1% of those diagnosed with the disease have it!

Should you be worried your doctor will get this wrong?

Kahnemann & Tversky (1972, 1973) argued that most people make errors when facing Bayesian updating problems like this one

Eddy (1982) posed this very problem to a sample of physicians and found most made massive errors, averaging an order of magnitude

Anecdotally, I have encountered physicians who make this error

RECOMMENDATIONS

1. Choose a doctor with statistics texts on his shelf!

2. Report probability information as ratios ("1 in 5"), not proportions ("0.20")

Gigerenzer and Hoffrage (1995) found this drastically improved inference, including among doctors faced with our example