

BIOS 600: Principles of Statistical Inference

Case Study: Breast Cancer Screening

Fall 2012

Reading

- ▶ Pagano and Gauvreau, Chapter 6, Section 6.3-6.4

What Does a Positive Mammogram Mean?

Suppose that 1% of women at age forty who participate in routine screening have breast cancer. 85% of women with breast cancer will have positive mammographies. 10% of women without breast cancer will also get positive mammographies. A woman in this age group had a positive mammography in a routine screening.

What is the probability that she actually has breast cancer?

Write your best guess here:

Breast Cancer Case Study

Most doctors get the same wrong answer on this problem - usually, only around 15% of doctors get it right (Casscells, Schoenberger, and Grayboys 1978; Eddy 1982; Gigerenzer and Hoffrage 1995; and many other studies). Most doctors estimate this probability to be in the range of 70-85%.

Breast Cancer Case Study

The correct answer is 7.9%, obtained as follows.

- ▶ Out of 10,000 women, 1% or 100 have breast cancer; we expect 85 of those 100 (85%) to have positive mammograms.

Breast Cancer Case Study

The correct answer is 7.9%, obtained as follows.

- ▶ Out of 10,000 women, 1% or 100 have breast cancer; we expect 85 of those 100 (85%) to have positive mammograms.
- ▶ From the same 10,000 women, 99% or 9900 will not have breast cancer, and of those 9900 women, 990 (10%) will also get positive mammograms.

Breast Cancer Case Study

The correct answer is 7.9%, obtained as follows.

- ▶ Out of 10,000 women, 1% or 100 have breast cancer; we expect 85 of those 100 (85%) to have positive mammograms.
- ▶ From the same 10,000 women, 99% or 9900 will not have breast cancer, and of those 9900 women, 990 (10%) will also get positive mammograms.
- ▶ Total number of women with positive mammograms:
 $990 + 85 = 1075$

Breast Cancer Case Study

The correct answer is 7.9%, obtained as follows.

- ▶ Out of 10,000 women, 1% or 100 have breast cancer; we expect 85 of those 100 (85%) to have positive mammograms.
- ▶ From the same 10,000 women, 99% or 9900 will not have breast cancer, and of those 9900 women, 990 (10%) will also get positive mammograms.
- ▶ Total number of women with positive mammograms:
 $990 + 85 = 1075$
- ▶ Of the 1075 women with positive mammograms, only 85 have cancer, so expressed as a proportion, this is $\frac{85}{1075} = 0.079$ or 7.9%. This is the *conditional probability* of breast cancer given a positive mammogram.

Breast Cancer Case Study

The correct answer is 7.9%, obtained as follows.

- ▶ Out of 10,000 women, 1% or 100 have breast cancer; we expect 85 of those 100 (85%) to have positive mammograms.
- ▶ From the same 10,000 women, 99% or 9900 will not have breast cancer, and of those 9900 women, 990 (10%) will also get positive mammograms.
- ▶ Total number of women with positive mammograms:
 $990 + 85 = 1075$
- ▶ Of the 1075 women with positive mammograms, only 85 have cancer, so expressed as a proportion, this is $\frac{85}{1075} = 0.079$ or 7.9%. This is the *conditional probability* of breast cancer given a positive mammogram.
- ▶ Thus the 40 year-old woman with a positive mammogram is still much more likely to learn that she is cancer-free than that she has cancer during follow-up.

Statistical Terminology

Let A be the event that someone has disease and let B be the event that a screening test is positive

- *Prevalence*: proportion of subjects who have a disease at a given point in time (1% for breast cancer at 40), $\Pr(A)$

Statistical Terminology

Let A be the event that someone has disease and let B be the event that a screening test is positive

- ▶ *Prevalence*: proportion of subjects who have a disease at a given point in time (1% for breast cancer at 40), $\Pr(A)$
- ▶ *Sensitivity*: probability of a positive test result given that the individual really has the disease, $\Pr(B | A)$

Statistical Terminology

Let A be the event that someone has disease and let B be the event that a screening test is positive

- ▶ *Prevalence*: proportion of subjects who have a disease at a given point in time (1% for breast cancer at 40), $\Pr(A)$
- ▶ *Sensitivity*: probability of a positive test result given that the individual really has the disease, $\Pr(B | A)$
- ▶ *False positive rate*: probability of a positive test result given that the individual is disease-free, $\Pr(B | \bar{A})$

Statistical Terminology

Let A be the event that someone has disease and let B be the event that a screening test is positive

- ▶ *Prevalence*: proportion of subjects who have a disease at a given point in time (1% for breast cancer at 40), $\Pr(A)$
- ▶ *Sensitivity*: probability of a positive test result given that the individual really has the disease, $\Pr(B \mid A)$
- ▶ *False positive rate*: probability of a positive test result given that the individual is disease-free, $\Pr(B \mid \bar{A})$
- ▶ *Specificity*: probability of a negative test result given that the individual is disease-free (1-false positive rate), $\Pr(\bar{B} \mid \bar{A})$

Statistical Terminology

Let A be the event that someone has disease and let B be the event that a screening test is positive

- ▶ *Prevalence*: proportion of subjects who have a disease at a given point in time (1% for breast cancer at 40), $\Pr(A)$
- ▶ *Sensitivity*: probability of a positive test result given that the individual really has the disease, $\Pr(B | A)$
- ▶ *False positive rate*: probability of a positive test result given that the individual is disease-free, $\Pr(B | \bar{A})$
- ▶ *Specificity*: probability of a negative test result given that the individual is disease-free (1-false positive rate), $\Pr(\bar{B} | \bar{A})$
- ▶ At the time of screening, we do not know A but only know B , so we are interested in the *positive predictive value*: probability of disease given a positive test result, $\Pr(A | B)$

Risk applet online

Breast Cancer Screening Animation

Breast Cancer Case Study

How do we formalize our thought process mathematically?

- ▶ To start, we know that a 40 year old woman has a 1% chance of having breast cancer (*prior probability*).

Breast Cancer Case Study

How do we formalize our thought process mathematically?

- ▶ To start, we know that a 40 year old woman has a 1% chance of having breast cancer (*prior probability*).
- ▶ Once we saw her positive mammogram, we updated that chance to 7.9% (*posterior probability*).

Breast Cancer Case Study

How do we formalize our thought process mathematically?

- ▶ To start, we know that a 40 year old woman has a 1% chance of having breast cancer (*prior probability*).
- ▶ Once we saw her positive mammogram, we updated that chance to 7.9% (*posterior probability*).
- ▶ This reasoning is formalized in *Bayes's Theorem*, which describes how we update our *prior* probabilities (the 1%) once we obtain additional data (the mammogram result), and is given by

$$\Pr(A \mid B) = \frac{\Pr(B \mid A)\Pr(A)}{\Pr(B)} = \frac{\Pr(A \cap B)}{\Pr(B)}.$$

Bayes's Theorem

$$\Pr(A \mid B) = \frac{\Pr(B \mid A)\Pr(A)}{\Pr(B)}$$

- ▶ A=breast cancer, B=positive mammogram

Bayes's Theorem

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

- ▶ A=breast cancer, B=positive mammogram
- ▶ Out of 10,000 women, 1% (prevalence) or 100 have breast cancer $\rightarrow \Pr(A) = 0.01$ and $\Pr(\bar{A}) = 0.99$

Bayes's Theorem

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

- ▶ A=breast cancer, B=positive mammogram
- ▶ Out of 10,000 women, 1% (prevalence) or 100 have breast cancer $\rightarrow \Pr(A) = 0.01$ and $\Pr(\bar{A}) = 0.99$
- ▶ We expect 85% of women with cancer to have positive mammograms $\rightarrow \Pr(B | A) = 0.85$; this is called the *sensitivity* of the mammogram

Bayes's Theorem

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

- ▶ A=breast cancer, B=positive mammogram
- ▶ Out of 10,000 women, 1% (prevalence) or 100 have breast cancer $\rightarrow \Pr(A) = 0.01$ and $\Pr(\bar{A}) = 0.99$
- ▶ We expect 85% of women with cancer to have positive mammograms $\rightarrow \Pr(B | A) = 0.85$; this is called the *sensitivity* of the mammogram
- ▶ We expect 10% of the women without cancer to have positive mammograms $\rightarrow \Pr(B | \bar{A}) = 0.10$; the fraction of women without cancer who have negative mammograms (0.90) is called the *specificity* of the mammogram

Bayes's Theorem

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

- ▶ A=breast cancer, B=positive mammogram
- ▶ Out of 10,000 women, 1% (prevalence) or 100 have breast cancer $\rightarrow \Pr(A) = 0.01$ and $\Pr(\bar{A}) = 0.99$
- ▶ We expect 85% of women with cancer to have positive mammograms $\rightarrow \Pr(B | A) = 0.85$; this is called the *sensitivity* of the mammogram
- ▶ We expect 10% of the women without cancer to have positive mammograms $\rightarrow \Pr(B | \bar{A}) = 0.10$; the fraction of women without cancer who have negative mammograms (0.90) is called the *specificity* of the mammogram
- ▶ How do we get $\Pr(B)$?

Bayes's Theorem

$$\Pr(A \mid B) = \frac{\Pr(B \mid A)\Pr(A)}{\Pr(B)}$$

- How do we get $\Pr(B)$?

Bayes's Theorem

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

- ▶ How do we get $\Pr(B)$?
 - ▶ B is the event of a positive mammogram

Bayes's Theorem

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

- ▶ How do we get $\Pr(B)$?
 - ▶ B is the event of a positive mammogram
 - ▶ The event B occurs both in the group with breast cancer and in the cancer-free group

Bayes's Theorem

$$\Pr(A \mid B) = \frac{\Pr(B \mid A)\Pr(A)}{\Pr(B)}$$

- ▶ How do we get $\Pr(B)$?
 - ▶ B is the event of a positive mammogram
 - ▶ The event B occurs both in the group with breast cancer and in the cancer-free group
 - ▶ So we can write $B = (B \cap A) \cup (B \cap \bar{A})$, which means everyone with a positive mammogram either has breast cancer and a positive mammogram (true positive) or no breast cancer and a positive mammogram (false positive)

Bayes's Theorem

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

- ▶ How do we get $\Pr(B)$?
 - ▶ B is the event of a positive mammogram
 - ▶ The event B occurs both in the group with breast cancer and in the cancer-free group
 - ▶ So we can write $B = (B \cap A) \cup (B \cap \bar{A})$, which means everyone with a positive mammogram either has breast cancer and a positive mammogram (true positive) or no breast cancer and a positive mammogram (false positive)
 - ▶ Because A (cancer) and \bar{A} (no cancer) are mutually exclusive, then we can use the additive rule to get $\Pr(B) = \Pr(B \cap A) + \Pr(B \cap \bar{A})$

Bayes's Theorem

$$\Pr(A \mid B) = \frac{\Pr(B \mid A)\Pr(A)}{\Pr(B)}$$

- How do we get $\Pr(B)$?

Bayes's Theorem

$$\Pr(A \mid B) = \frac{\Pr(B \mid A)\Pr(A)}{\Pr(B)}$$

- ▶ How do we get $\Pr(B)$?
 - ▶ $\Pr(B) = \Pr(B \cap A) + \Pr(B \cap \bar{A})$

Bayes's Theorem

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

- ▶ How do we get $\Pr(B)$?
 - ▶ $\Pr(B) = \Pr(B \cap A) + \Pr(B \cap \bar{A})$
 - ▶ Recall the multiplicative rule of probability states that
$$\Pr(A \cap B) = \Pr(B \cap A) = \Pr(A)\Pr(B | A) = \Pr(B)\Pr(A | B)$$

Bayes's Theorem

$$\Pr(A \mid B) = \frac{\Pr(B \mid A)\Pr(A)}{\Pr(B)}$$

- ▶ How do we get $\Pr(B)$?
 - ▶ $\Pr(B) = \Pr(B \cap A) + \Pr(B \cap \bar{A})$
 - ▶ Recall the multiplicative rule of probability states that $\Pr(A \cap B) = \Pr(B \cap A) = \Pr(A)\Pr(B \mid A) = \Pr(B)\Pr(A \mid B)$
 - ▶ So $\Pr(B \cap A) = \Pr(A)\Pr(B \mid A) = 0.01(0.85) = 0.0085$

Bayes's Theorem

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

- ▶ How do we get $\Pr(B)$?
 - ▶ $\Pr(B) = \Pr(B \cap A) + \Pr(B \cap \bar{A})$
 - ▶ Recall the multiplicative rule of probability states that $\Pr(A \cap B) = \Pr(B \cap A) = \Pr(A)\Pr(B | A) = \Pr(B)\Pr(A | B)$
 - ▶ So $\Pr(B \cap A) = \Pr(A)\Pr(B | A) = 0.01(0.85) = 0.0085$
 - ▶ and $\Pr(B \cap \bar{A}) = \Pr(\bar{A})\Pr(B | \bar{A}) = 0.99(0.10) = 0.099$

Bayes's Theorem

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

- ▶ How do we get $\Pr(B)$?
 - ▶ $\Pr(B) = \Pr(B \cap A) + \Pr(B \cap \bar{A})$
 - ▶ Recall the multiplicative rule of probability states that $\Pr(A \cap B) = \Pr(B \cap A) = \Pr(A)\Pr(B | A) = \Pr(B)\Pr(A | B)$
 - ▶ So $\Pr(B \cap A) = \Pr(A)\Pr(B | A) = 0.01(0.85) = 0.0085$
 - ▶ and $\Pr(B \cap \bar{A}) = \Pr(\bar{A})\Pr(B | \bar{A}) = 0.99(0.10) = 0.099$
 - ▶ So $\Pr(B) = \Pr(B \cap A) + \Pr(B \cap \bar{A}) = 0.0085 + 0.099 = 0.1075$
(positive mammograms are fairly common occurrences)

Bayes's Theorem

So figuring out the probability of cancer when the mammogram is positive, we have

$$\begin{aligned}\Pr(A | B) &= \frac{\Pr(B | A)\Pr(A)}{\Pr(B)} \\ &= \frac{0.85(0.01)}{0.1075} = 0.079\end{aligned}$$

or 7.9%

Example 1: Positive Predictive Value

- ▶ Of 10,000 people in the US, we expect 1 to have HIV and 9999 to be HIV-free
- ▶ Of 10,000 people with HIV, we expect 9980 to test positive when given a diagnostic test for HIV and 20 to test negative
- ▶ Of 10,000 people without HIV, we expect 1 to test positive and 9999 to test negative

What is the positive predictive value of this HIV test?

Example 2: Positive Predictive Value

- ▶ We expect 1.5% of US IV drug users to have HIV
- ▶ Of 10,000 people with HIV, we expect 9980 to test positive when given a diagnostic test for HIV and 20 to test negative
- ▶ Of 10,000 people without HIV, we expect 1 to test positive and 9999 to test negative

What is the positive predictive value of this HIV test in the population of IV drug users?

Example 3: Positive Predictive Value

- ▶ We expect 24.8% of adults aged 15-49 in Botswana to have HIV
- ▶ Of 10,000 people with HIV, we expect 9980 to test positive when given a diagnostic test for HIV and 20 to test negative
- ▶ Of 10,000 people without HIV, we expect 1 to test positive and 9999 to test negative

What is the positive predictive value of this HIV test in this population in Botswana?

Example 4: Positive Predictive Value

Investigators in China conducted a study of an HPV self-test for high-grade lesions in a high risk population.

- ▶ The prevalence of high-grade lesions in this population was 4.3%
- ▶ The sensitivity of the test is 83%
- ▶ Specificity is 86%

What is the probability that someone who has a positive result on the HPV self-test in this population actually has high-grade lesions?

Example 5: Kill the Wabbit

A certain brand of home pregnancy test claims 90.6% sensitivity and 99.9% specificity on the first day of missed menses.

Assuming roughly 10% of test users are actually pregnant, what is the positive predictive value of this pregnancy test?

[Link to the Operatic Allusion](#)

What's Up with Rabbits?

The rabbit test, or Aschheim-Zondek test, was an early pregnancy test developed in 1927. The original test actually used mice and was based upon the observation that when urine from a female in the early months of pregnancy is injected into immature female mice, the ovaries of the mice enlarge and show follicular maturation. The rabbit test consisted of injecting the tested woman's urine into a female rabbit, then examining the rabbit's ovaries a few days later, which would change in response to a hormone only secreted by pregnant women. The hormone, hCG, is produced during pregnancy and indicates the presence of a fertilized egg; it can be found in a pregnant woman's urine and blood. The rabbit test became a widely used bioassay (animal-based test) to test for pregnancy. The term "rabbit test" was first recorded in 1949 but became a common phrase in the English language.

Modern pregnancy tests still operate on the basis of testing for the presence of the hormone hCG. Due to medical advances, use of a live animal is no longer required.

It is a common misconception that the injected rabbit would die only if the woman was pregnant. This led to the phrase "the rabbit died" being used as a euphemism for a positive pregnancy test. In fact, all rabbits used for the test died, because they had to be surgically opened in order to examine the ovaries. While it was possible to do this without killing the rabbit, it was generally deemed not worth the trouble and expense.

–Courtesy of Wikipedia

How might you design a better test?

A perfect test has sensitivity=specificity=1. Numerous factors need to be considered when deciding how to design a better test for cancer, including

- ▶ Human costs

How might you design a better test?

A perfect test has $\text{sensitivity} = \text{specificity} = 1$. Numerous factors need to be considered when deciding how to design a better test for cancer, including

- ▶ Human costs
 - ▶ HIV test for purposes of screening blood supply must have high sensitivity

How might you design a better test?

A perfect test has sensitivity=specificity=1. Numerous factors need to be considered when deciding how to design a better test for cancer, including

- ▶ Human costs
 - ▶ HIV test for purposes of screening blood supply must have high sensitivity
 - ▶ HIV test for diagnosing infection in humans should have high sensitivity and high specificity

How might you design a better test?

A perfect test has sensitivity=specificity=1. Numerous factors need to be considered when deciding how to design a better test for cancer, including

- ▶ Human costs
 - ▶ HIV test for purposes of screening blood supply must have high sensitivity
 - ▶ HIV test for diagnosing infection in humans should have high sensitivity and high specificity
- ▶ Costs of follow-up if test is positive

How might you design a better test?

A perfect test has sensitivity=specificity=1. Numerous factors need to be considered when deciding how to design a better test for cancer, including

- ▶ Human costs
 - ▶ HIV test for purposes of screening blood supply must have high sensitivity
 - ▶ HIV test for diagnosing infection in humans should have high sensitivity and high specificity
- ▶ Costs of follow-up if test is positive
- ▶ Whether follow-up is detrimental for false positives

How might you design a better test?

A perfect test has sensitivity=specificity=1. Numerous factors need to be considered when deciding how to design a better test for cancer, including

- ▶ Human costs
 - ▶ HIV test for purposes of screening blood supply must have high sensitivity
 - ▶ HIV test for diagnosing infection in humans should have high sensitivity and high specificity
- ▶ Costs of follow-up if test is positive
- ▶ Whether follow-up is detrimental for false positives
- ▶ Bear in mind that there is a trade-off in general between sensitivity and specificity (can improve one at the cost of the other)

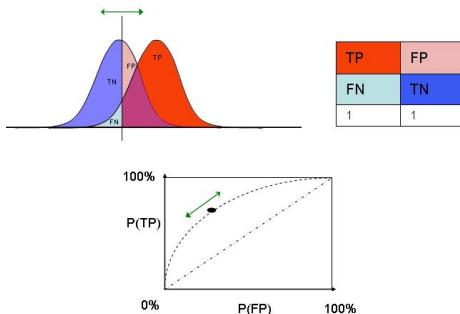
Sensitivity-specificity tradeoff

- ▶ True positive: has both positive test result and disease
- ▶ False positive: has positive test result and no disease
- ▶ True negative: has negative test result and no disease
- ▶ False negative: has negative test result but has disease

Click to explore the **sensitivity-specificity tradeoff**

Receiver operating characteristic (ROC) curves

The receiver operating characteristic (ROC) curve is a plot of the sensitivity (true positive rate) against the false positive rate (1 minus the specificity). The area under the ROC curve is often used to describe the quality of a test.



ROC curves

Click to explore **ROC curves for varying distributional separation**

Reading for Next Time

- ▶ Pagano and Gauvreau, Chapter 7, Sections 7.1-7.2