Medical Test Accuracy and Statistics

IB Mathematics SL

Research Report - Specimen Paper A

Given the stimulus word, **Health**, I started thinking of the ways one can use probability (my favorite branch of Mathematics) in looking at medical topics. I decided to find out what it means for a medical test to be "97% accurate" or "99% accurate". We all know that no medical test is perfect.   How much error is acceptable, and what do these errors mean? What happens if a patient is diagnosed incorrectly?

To begin, we need to understand what the word "accurate" really means in terms of medical tests. The accuracy of a medical test refers to two things - sensitivity and specificity. Sensitivity is the ability of a test to correctly diagnose a condition in a patient who has that condition. Specificity is the ability of a test to correctly determine that a person does not have a condition. For example, if a very high percentage of people who have a disease test positive for that disease, the test has high sensitivity. If people who do not have a disease test negative, the test has high specificity. Therefore, a medical test which provides a correct result 99% of the time can be said to be "99% accurate".

Let's look at what medical test accuracy means to a general (fictional) population. Suppose that a new and terrible brain disease, *Mathiti*s, is known to currently affect approximately 5% of all students aged 16-18. A doctor has developed a medical test to diagnose this affliction. The test is "99% accurate". A large school district decides to test all of the 20000 students within this age group. The table below shows how the results would come out. The results in **green** represent a correct diagnosis, and the results in **red** are incorrect.

RESULTS FROM *MATHITIS* DETECTION TEST	test reads POSITIVE for disease	test reads NEGATIVE for disease	Totals
student really has *MATHITIS*	990	10	1000
student does not have *MATHITIS*	190	18810	19000
Totals	1180	18820	20000

In this particular school district, 19000 students do not have the disease, though only 18820 will test NEGATIVE. One way to look at this data is that 18810 out of 18820, or 99.95%, of the NEGATIVE test results are correct. These are called **true negatives**. 99.95% correct sounds pretty good, right? The table also shows, however, that although 1000 students (out of 20000) actually have the disease, 1180 will test positive. In other words, about 16.1% of the POSITIVE results are <u>incorrect</u>. These results are called **false positives**, and 16.1% incorrect doesn't sound very accurate!

We can also look at this (admittedly fictional) situation using conditional probability. The formula for conditional probability is $P(A|B) = \frac{P(A \cap B)}{P(B)}$. In words, we would say, "the probability of event A given that we know event B is equal to the probability of the intersection of events A and B, divided by the probability of B." In our example, let T be the event that a student tests positive for the disease, and let M be the event that a student has the dreaded disease *Mathitis*. What if we wanted to find the probability that a student who tests positive for the disease actually has the disease? We would use our formula for conditional probability:

$$P(M|T) = \frac{P(M \cap T)}{P(T)} = \frac{\left(\frac{990}{20000}\right)}{\left(\frac{1180}{20000}\right)} = \frac{990}{1180} \approx 0.839 = 83.9\%$$

In words, this means that about 83.9% of the students who test positive for *Mathitis* actually have the disease. This is easily seen using the values in the table.

What if you are among the students who test positive for this disease? Do you kiss your family and friends, not to mention your University acceptance, goodbye? Of course you don't. In most cases (real, not fictional), if a person tests positive for a serious disease, then a second, even more accurate, test is given. This raises the question, if a more accurate test exists, why didn't the doctor just administer that test in the first place? In many cases, this second test can be much more expensive and invasive than the initial test. It is just not economically feasible, nor physically comfortable, to give a more invasive test to the general population.

Now let's take a look at a real-life situation, the HIV/AIDS virus. It is estimated that about 33,000,000 people in the world are living with this virus. The world's population is currently about 6,770,000,000, which means that about 0.487% of the people on Earth have the HIV/AIDS virus. Let's say that an initial medical test for the virus is 99.7% accurate, which sounds very good. In fact, "A large study of HIV testing in 752 U.S. laboratories reported a sensitivity of 99.7% and specificity of 98.5% for enzyme immunoassay *[a type of detection test]*, and studies in U.S. blood donors reported specificities of 99.8% and greater than 99.99%"[1], so it seems reasonable to assume that an initial test might be 99.7% accurate. Here is what the results would look like if the test were given to every person on Earth:

INITIAL HIV/AIDS TEST RESULTS	test reads POSITIVE for HIV/AIDS	test reads NEGATIVE for HIV/AIDS	Totals
person really has the virus	32,901,000	99,000	33,000,000
person does not have the virus	20,211,000	6,716,789,000	6,737,000,000
Totals	53,112,000	6,716,888,000	6,770,000,000

Once again, we see that the NEGATIVE test results look pretty good. Only 0.00147% of these are incorrect, while 99.9985% of the NEGATIVE results are correct. The POSITIVE test results tell a different story, however. About 38.1% of the POSITIVE results are incorrect, or *false positives*. Only 61.9% of those who receive a POSITIVE diagnosis in the original medical test actually have the HIV/AIDS virus.

Does all this mean that we shouldn't take a POSITIVE result seriously? Is this so-called 99.7% "accurate" test flawed? The answer to both of these questions is NO. However, it is interesting that the results from this very serious real-life disease are somewhat similar to our fictional "*Mathitis*" situation. In both cases, the NEGATIVE test results seem pretty reliable, but the POSITIVE results are not so reliable.

Why do we seem to have such a high percentage of *false positives*? For one thing, not every person on Earth would generally be given an AIDS test, or a test for any other disease, for that matter. The people who are tested for most diseases are those who are showing symptoms or who are more at risk for that disease than the general population. In this case, we are looking at a relatively rare disease, yet the test is given to a huge number of people.   While anyone can get the HIV/AIDS virus, not every person on Earth is at equal risk, due to factors such as general health, age, lifestyle, geographic location, etc. Consider a disease like the recent "Swine Flu", or H1N1 virus. Not everyone needs to be tested for this flu virus. Those given the test are people who are already showing symptoms of the illness, such as a prolonged fever, or those who have been exposed to a person known to be infected.

In addition, the percentages would be different if we were to look at specific populations. Sadly, about two-thirds of the HIV/AIDS cases in the world are located in Sub-Sarahan Africa, where in some countries over 20% of the population is infected with the virus. We would expect many more *true positive* results to an HIV/AIDS test given in this part of the world than we would in areas like eastern Asia and parts of central Europe, where fewer than 0.1% of the populations are infected. If we were to conduct our conditional probability test on these populations, we would see quite different percentages than we would when looking at the world-wide population.

Of course, we have only been looking at our data in one way, using conditional probability. We all know that statistics can be used to tell the truth, or to distort the truth. One of my favorite quotes is, "Statistics are like bikinis. What they reveal is suggestive, but what they conceal is vital" *~Aaron Levenstein.*[2] Let's try looking at our data in two different ways, and see what it reveals.

If we look at the number of *false positives* from our HIV/AIDS test, we have **20,211,000**. This seems relatively large when compared to the total number of positive test results. Let T be the event that a person tests positive for the virus, and let V be the event that a person actually has the HIV/AIDS virus. So once again using conditional probability,

$$P(V'|T) = \frac{P(V' \cap T)}{P(T)} = \frac{\left(\frac{20211000}{6770000000}\right)}{\left(\frac{53112000}{6770000000}\right)} = \frac{20211000}{53112000} \approx 0.381 = 38.1\%$$

In words, this means that **if** a person tests POSITIVE for the virus in our initial test, there is a 38.1% probability that he does **not** actually have the virus.

Instead, now let's compare the number of false positives, **20,211,000**, to the total number of people given the test, **6,770,000,000** (i.e. the entire world population).

So we could say that the probability of any random person receiving a *false positive* result is equal to $\frac{20211000}{6770000000}$, which is approximately 0.298%. This supports the supposed 99.7% accuracy of the test. So although the conditional probability seems very high, there are actually a very small percentage of these *false positives* in the entire population.

Now let's look at what would happen when the people who receive a POSITIVE result are given a second, more expensive, but also more accurate, test. Let's suppose that this second test has an accuracy level of 99.9%.

SECONDARY HIV/AIDS TEST RESULTS	test reads POSITIVE for HIV/AIDS	test reads NEGATIVE for HIV/AIDS	Totals
person really has the virus	**32,868,099**	**32,901**	32,901,000
person does not have the virus	**20,211**	**20,190,789**	20,211,000
Totals	32,888,310	20,223,690	**53,112,000**

We can look at these results in a few different ways, and consider the implications of these results. First, let's consider the number **32,868,099**, which represents the number of *true positives*. Remember, these people have now tested positive for the virus twice. So we can say that **32,868,099** out of the estimated 33,000,000 people who actually have the virus, or about 99.6%, have received a correct diagnosis twice. This is a huge percentage of people who can now begin receiving treatment and can hopefully live longer, fuller lives. I think most of us would agree that this is the whole purpose of medical testing, and that being able to properly treat 99.6% of the patients with this or any other disease is a good thing.

What about the patients who have been diagnosed incorrectly in either test? We have **99,000** people who received a *false negative* result from the first test and were not given the second test. We have **32,901** people who received a false negative result on the second test. We also have **20,211** people who received *false positive* results twice. This means that we have a total of 152,112 people who have been mis-diagnosed. 131,901 people have the virus and may not receive the treatment they need. Even worse, they might unknowingly pass the virus on to others. And the **20,211** people who do no have the virus, yet have tested positive twice? They may be receiving unnecessary treatment and living in fear.

These are just the whole numbers, let's take a look at the statistics. We have 152,112 people who have been mis-diagnosed n at least one of the tests. This represents about 0.00225% of the world population. Is this acceptable? Is it OK for 20211 people (or

0.00000298% of the world population) to live with the mistaken notion that they have a deadly virus? In this case, it is not so easy for us to say yes, this is acceptable. We are looking at an extremely small percentage of the population, yet it doesn't seem fair.

But what can be done? Let's consider a few things. First, we need to remember that an HIV/AIDS test, or any other test, would not be given to the entire population. So it is unlikely that there would even be as many as 0.00000298% "*double false positives*" in any test for a serious disease. Second, if a person who received a *false negative* continued to display symptoms of a disease, it is assumed that that person would receive some type of medical attention, and would be tested again. And finally, we can hope that anyone (both *true* and *false positives*) who was receiving any type of treatment would be monitored closely by a medical professional.

To me, the statistical nature of this topic is quite interesting and fun, from a mathematical perspective. The personal and ethical implications, however, are somewhat staggering. In looking at the numerical data, I was constantly asking myself, "What if I were one of the false positives?" I guess the best thing for me to remember is that at least I am lucky enough to have access to the medical testing and treatment I need in any type of health situation. Unfortunately, there are too many people in the world who do not have this luxury.

BIBLIOGRAPHY

1. Annals of Internal Medicine; "Screening for HIV: A Review of the Evidence for the U.S. Preventive Services Task Force"; Roger Chou, MD; Laurie Hoyt Huffman, MS; Rongwei Fu, PhD; Ariel K. Smits, MD, MPH; and P. Todd Korthuis, MD, MPH; http://www.annals.org/cgi/content/full/143/1/55

2. "Quotations about Statistics"; http://www.quotegarden.com/statistics.html

3. UNAIDS: Joint United Nations Programme on HIV/AIDS; "2008 Report on the Global AIDS Epidemic"; http://www.unaids.org/en/KnowledgeCentre/HIVData/GlobalReport/2008/2008_Global_report.asp

4. CancerQuest - Emory University; "Sensitivity and Specificity of Medical Tests"; http://www.cancerquest.org/index.cfm?page=3305

5. CancerQuest - Emory University; "Frequently Asked Questions: False Positives and Negatives"; http://www.cancerquest.org/index.cfm?page=4923

6. Tau Topics; "Accuracy of Diagnostic Laboratory Medical Tests: Specificity and Sensitivity"; http://www.michaelandjudystouffer.com/judy/articles/specsen.htm