Paxson
Spring 2011

CS 161
Computer Security

Discussion 11

April 13, 2011

Question 1 ***Detecting attacks*** **(7 min)**

Suppose that S is a network-based intrusion detector that works by passively analyzing individual UDP and TCP packets. Suppose that A is a host-based intrusion detector that is a component of the browser that processes and analyzes individual URLs before they are loaded by the browser.

Your company decides to build a hybrid scheme for detecting malicious URLs. The hybrid scheme works by combining scheme S and scheme A, running both in parallel on the same traffic. The combination could be done in one of two ways. Scheme H_E would generate an alert if for a given network connection either scheme S or scheme A generates an alert. Scheme H_B would generate an alert only if both scheme S and scheme A generate an alert for the same connection. (Assume that there is only one URL in each network connection.)

(a) Assuming that decisions made by S and A are well-modeled as independent processes, and ignoring any concerns regarding evasion, what can you say about the false positives and false negatives of H_B and H_E?

> **Solution:** The key insight here is that alarms by H_B will be a subset of the alarms generated by H_E. Since H_B will generate fewer alarms for non-malicious activities, it will have less false positives. On the other hand, because it generates fewer alarms, it might miss more malicious activity, implying more false negatives.

(b) If deploying the hybrid scheme in a new environment, is one of H_E and H_B clearly better?

> **Solution:** In the absence of more data, particularly the cost of false positive and false negatives, as well as the rate of malicious and non-malicious activity, it is impossible to make any decision.

Question 2 ***Base Rate Fallacy*** **(7 min)**

The Department of Homesand Security is really concerned about the threat of "sanders," basically people who come to American Shores and leave with sand, a precious national resource. Deeming this unacceptable, the senate authorizes a project to automatically identify sanders crossing into the country, tentatively called "Operation Sanders Crossing" or OSX for short. As part of OSX, you are tasked to choose a Sander Detection

System from the many offered by defense contractors.

(a) Let S denote the event that a sander is crossing the border, and let C denote a border crossing. Let A denote the event that the sander detection system generates an alarm. In what follows, identify which is the term that is the correct notation for false positive, false negative, true positive and true negative.

$$\mathbb{P}[A \mid S]$$
$$\mathbb{P}[\neg A \mid S]$$
$$\mathbb{P}[A \mid \neg S]$$
$$\mathbb{P}[\neg A \mid \neg S]$$

Solution: In order, they are: true positive, false negative, false positive, true negative.

(b) One of the system advertises a false negative rate of 1 in 10,000 and a false positive rate of 1 in 10,000. Your boss thinks this rate is pretty darn good and thinks they should be bought. What do you think?

Solution: In the absence of base rate, it is impossible to make an informed decision.

(c) In terms of the notation above, write down the notation for the *Bayesian detection rate*, i.e., the probability that an alarm is actually for a sander.

Solution: $\mathbb{P}[S \mid A]$

In the IDS scenarios discussed in lecture, this is the probability that there is actually an attack, when your IDS generates an alarm. Note that this is different from the true positive rate, which is the probability of an alarm if an attack takes place. Since attacks are often uncommon, the true positive rate can give the wrong idea to the reader.

(d) If a doctor tells you that you tested positive for something and that the test is 99% accurate, does that mean you have a 99% chance of having the disease? What is the chance of having a disease if the disease is really rare, say only 1 in 1 million

people have the disease ? You can assume that the false negative and false positive rate for this test are both 1%.

Solution: No. The chances are more like 1 in 10,000.

Let's say 1 million people take the test. We know that only 1 person has the disease out of the 1 million. How many will test positive? Well, out of the 999,999 people without the disease, 1% will test positive, which is approximately 10,000. Since only 1 person has the disease, the probability of you having the disease if you test positive is approximately 1 in 10,000.

Your chance of having that disease would be 99% only if it were the case that the base rate of people having the disease and not having the disease is equal. This is not the case here—the probability of having a disease is a million times less. Keep on the lookout for base-rate fallacies anywhere this difference in base rate is huge. This is particularly relevant to security because the base rate of attacks is usually much much lower than that of normal non-malicious activity.

We can also solve this using Bayes Theorem.

$$\begin{aligned}
&\mathbb{P}\left[\text{have disease} \,|\, \text{tested positive}\right] \\
&= \frac{\mathbb{P}\left[\text{tested positive} \,|\, \text{have disease}\right] \mathbb{P}\left[\text{have disease}\right]}{\mathbb{P}\left[\text{tested positive}\right]} \\
&= \frac{0.99 * (\frac{1}{1{,}000{,}000})}{\frac{0.99*1+0.01*999{,}999}{1{,}000{,}000}} \\
&= \frac{0.99}{10,000.98} \\
&= \frac{1}{10,102}
\end{aligned}$$

The base-rate fallacy is assuming that $\mathbb{P}\left[\text{have disease} \,|\, \text{tested positive}\right] = \mathbb{P}\left[\text{tested positive} \,|\, \text{have disease}\right]$.

Question 3 ***Side and Covert Channels*** **(5 min)**

(a) What is the difference between side channels and covert channels?

Solution: A side channel is a channel that leaks information due to the physical implementation. It's a *side* channel in the sense that it is not a theoretical weakness in a system, but rather an effect of its physical implementation. Side channels do not involve two cooperating parties; they instead are used by a single party to extract information they are not meant to have.

A covert channel is a channel that allows information transfer between two cooperating parties that aren't supposed to be able to communicate.

(b) Consider implementing the RSA cryptography algorithm. The typical way is to go through the 'key' bit by bit. The pseudo-code looks something like this:

```
foreach (bit in key) {
  if (bit) {
    // do multiplication and all hard work if bit is 1
  }
  // do other simpler stuff that you need to do regardless
}
```

Recall the cable box with a tamper resistant private key inside it that Prof. Paxson talked about in the lecture. Can you imagine a side-channel attack on the above implementation to find the private key? HINT: Can you do something with a multimeter?

Solution: The length of time the power is at its peak can give you a clear indication of the bit pattern of the key. Multiplication usually requires more power, and this is noticeable in embedded systems. For example, see the graph on Wikipedia [1].

References

[1] Power attack. https://secure.wikimedia.org/wikipedia/en/wiki/File:Power_attack.png.