

The Fallacy of the False Negative

Doug Winter

December 2003

Over the last few years techniques for computer pattern matching have advanced significantly. The same sorts of computer programs, generally termed *Biometrics*, enable software to perform some of the tasks that have traditionally been the preserve of humans.

A lot of these systems are now in successful use all over the world - facial recognition can be used to control access to sensitive areas. New Scotland Yard uses a computer to compare fingerprints. Voice recognition can be used to identify people, even over the telephone.

In the heightened tensions since 9/11 a number of opportunistic companies have started touting their computer software as a solution to terrorist threats. Politicians, aware of the short-term memory of their electorate, happily tout these technological solutions often, I suspect, in complete knowledge that they are never going to work.

These systems have escaped serious scrutiny on the whole because the apparent complexity of the software is off-putting. After all, who is going to argue with the boffins over whether their software works or not?

Unfortunately, many of the more audacious claims for these systems are never going to be achieved. One of the reasons for

this is something I'll call the "Fallacy of the False Negative".

One of the systems for which there was a huge amount of media hype about 18 months ago was facial recognition of terrorists at airports. It was claimed that a computer, given a number of photographs of (generally bearded, generally Arab) suspects, could then identify them automatically using airport cameras, perhaps even the cameras already present for surveillance.

This kind of arrangement is back-to-front from the normal way facial recognition would be used, and has been used successfully. Normally it would be used to control access to sensitive areas, for example bank vaults. Of the six billion people on earth, only a few hundred are allowed in, and everyone else should be turned away.

For the bank vault, a remarkably large error rate is acceptable. If the system is only 90% effective, it's still quite useable. If twenty people want to go into the vault each day, you'll only get two refusals a day, where you should allow people in. These are called "False Negatives" - negative recognitions when they should be positive. They can retry, go and fetch a colleague to let them in, or use their key. Whatever they do, this is still quite a useable (although I imagine quickly irritating)

system.

Of course, for every twenty intruders who try to get in, two will be allowed - this is a real problem. The odds aren't good for the intruder though, and they probably aren't good enough for someone to risk breaking in on the offchance they'll be allowed in. No system like this would be used without any other form of security in any case.

However, in the airport, everyone is to be allowed access except for a few hundred people who should be recognised. Used this way, the accuracy of the systems has a remarkably different effect.

Around 185,000 people pass through Heathrow airport each day, for example. None of these are terrorists. For these 185,000 non-terrorists, a system that is 90% effective will trigger for 18,500 of them. That's 18,500 false positives **each day** - on average, one every 5 seconds.

Even the most optimistic politician won't think that the airport police can deal with one false alarm every 5 seconds. Clearly any system that's only 90% effective is nowhere near good enough to use in this way. The systems that are currently in production are not even that good.

A system that only generates ten false alarms a day for Heathrow would probably be acceptable - that would need to be 99.995% effective. That is years away, if it is even achievable.

The same argument can be applied to a

whole range of these computer "solutions" to security problems. For example, with David Blunkett's Identity Card scheme one of the main aims is to stop people using multiple identities. This is supposed to work because the system will spot when the same iris print is used on more than one card.

This is another instance of the Fallacy of the False Negative. The most optimistic projections for iris recognition systems makes them 99%¹ effective. There are 60 million people in Britain. This means each print will 'match' 600,000 other prints even without any forgery. There's absolutely no way it's ever going to identify two cards with the same print.

It's hard to believe the politicians who push these systems so hard are unaware of these facts - the analysis is simple, and has surely been performed many times at the Home Office, to name but one.

There's a wider moral to this of course - technological solutions only work for technological problems. Terrorism and fraud are human problems, and they need human solutions to solve them.

This work is licensed under the Creative Commons Attribution License. To view a copy of this license, visit <http://creativecommons.org/licenses/by/1.0/>

Source versions of this work are always available at <http://www.britishsteal.com/articles/>

¹New Scientist, 22 November 2003