Why Biometrics is not a Panacea

Peter Gutmann

University of Auckland

Two Usage Modes for Biometrics

Access control: Only these exact people are allowed in

- Exact-match biometric check weeds out 99.9% of users
  - Match exactly two people, and no-one else
- Backup PIN entry is used after the biometric check

Identification

- Inexact match used to find… uhh… things
  - Find one of 3 million people (DHS terrorist list) from a population of 6 billion
- The answer to all your terrorism problems

Biometrics for Identification

Biometrics will solve our political/ liquidity^H^H^H^Hterrorist problems

- Biometrics firms stock prices have tripled since 9/11
- US Government plans to spend $8B on biometrics

Biometrics 101

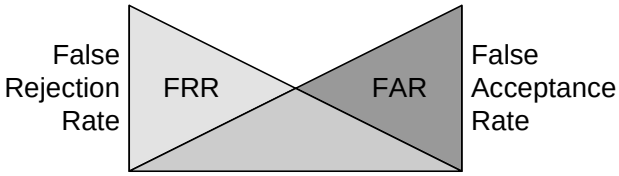

- Use is a tradeoff between FAR and FRR (think IDS)

Biometrics for Identification (ctd)

Medical research has standard terminology for these measures

- Sensitivity = $\frac{\text{true positives}}{\text{true positives + false negatives}}$
- Selectivity = $\frac{\text{true negatives}}{\text{true negatives + false positives}}$

Biometrics has barely considered the analysis techniques available from decades of medical research

- Much of it is still in the 19th-century patented cure-all medicine stage
- (There *is* serious research being done, but it rarely comes through when the biometric technology is being sold)

FAR/FRR Example

London City Airport

- Fingerprint scans for access control
- 1,600 employees
- 90,000 prints/day
- FAR = 1.5%, FRR = 0.001%
  - Vendor figures, these are always *exceedingly* optimistic (see later slides)
- Result: 1,500 false alarms per day
  - Statistical term for this is the base rate fallacy

Theoretical Background

Analysing false positives using the base-rate fallacy

- Apply a test for infection with the dreaded lurgy
- Test is 99% accurate
  - 99 of 100 ill patients will be detected
  - 99 of 100 healthy patients will be cleared
- Only 1 in 10,000 people have the disease
- Your doctor tells you that you've tested positive

What's the chance that you actually have this disease?

Theoretical Background (ctd)

From Bayesian statistics we have that

$$p(\mathrm{S}|\mathrm{P}) = \frac{p(\mathrm{S}) \times p(\mathrm{P}|\mathrm{S})}{p(\mathrm{S}) \times p(\mathrm{P}|\mathrm{S}) + p(\sim\mathrm{S}) \times p(\mathrm{P}|\sim\mathrm{S})}$$

where

- S = probability of being ill
- ~S = probability of not being ill
- P = probability of positive test result
- ~P = probability of negative test result

Theoretical Background (ctd)

Plugging in the numbers

$$p(\mathrm{S}|\mathrm{P}) = \frac{1/10000 \times 0.99}{1/10000 \times 0.99 + (1 - 1/10000) \times 0.01}$$

$$= 0.00980$$

$$= \sim 1\%$$

Even though the test is 99% certain, the fact that the population of healthy people is much larger than the population with the disease means that your chance of really having lurgy is only 1%

Biometrics and 9/11

Why didn't the US government use biometrics before 9/11?

> The government didn't have this stuff in place, precisely because it had been working on it and knew its limitations and didn't find any value for the costs involved. The government has been on top of this; the government's position hasn't changed
>
> — Jim Wayman, former director,
> US Government Biometrics Center

Problems with Biometrics

Biometric data is often sent around unprotected

- c.f. Plaintext password
- Example biometric system: Sensor outputs a binary yes/no signal on an external signal wire
- Capture biometric data with USB Snoop (software), USB Agent (hardware)

Once one biometric becomes the standard, everyone will know it

- Most biometric systems use open view traits

  I object to leaving my password on every glass I touch

  — Anon
- ICAO biometric passport design leaks data everywhere

Problems with Biometrics (ctd)

Biometric systems have never had to withstand serious attack

- Smart cards took 15 years of criminals walking all over them before vendors started taking security seriously

Fingerprint scanners work poorly with the elderly, manual workers

- 3-4% of the population (goats) have unstable biometric traits that can't be identified by sensors
- Attack: Train the system to accept less and less reliable images
  - Has happened (inadvertently) in real-world deployments as sensors were subject to wear and tear
  - System would accept anything (elbow, nose) as a valid print

Problems with Biometrics (ctd)

Can only compare traits against a database of trait characteristics

- Terrorists would have to register with the DHS in advance
- Typical terrorist photo is a grainy 10-year-old B&W shot at 150m distance
- Terrorism works because no-one knows who the grunts are
  - Biometrics can never catch disposable terrorists

No biometric system has ever caught a terrorist or serious criminal

The Problem with Biometrics as a Dragnet

When a highway patrolman is sent to his duty, he has to be given the authority to cite traffic violators. This cannot be done explicitly for each violator because at the time that the patrolman is sent to his duty, the traffic violator does not exist, and the identity of the future violators is not known, so that it is impossible to construct individual access rights for the violators at that time. The point is that the patrolman's authority has to do with the behavior of motorists, not their identity

— Naftaly Minsky, International Journal of Computer and Information Sciences, June 1978.

The only Effective ID for catching Terrorists

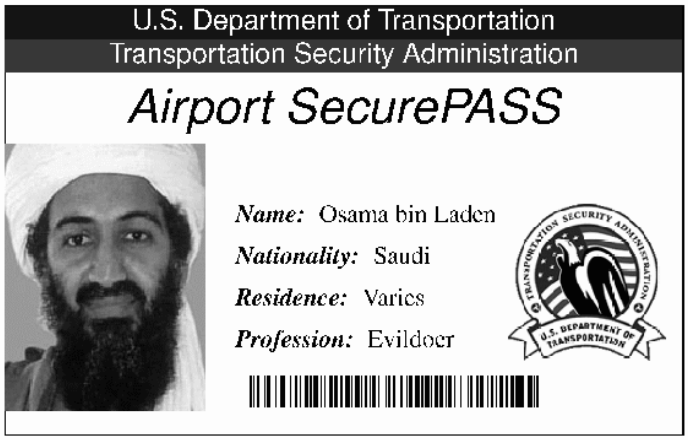

Image courtesy Steve Bellovin

How to apply this ID: Stop everyone who has given their occupation as "Evildoer"

Fingerprint Sensors put to the Test

Faking prints (pre-biometrics)

- Photograph a latent print using high-contrast film
- Develop the film using a process that leaves ridges of loops and whorls raised in relief
- Use sebum (oil from skin, in this case hair grease) to create "print"

Time to copy a print with the owner's co-operation

- 10 years ago: 2-3 hours
- Today: 15 minutes, $10

Time to copy a latent print

- 10 years ago: Several days
- Today: 30 minutes, $20

Fingerprint Sensors put to the Test (ctd)

Network Computing review, 1998

- Only two of six fingerprint readers could reject a plastic finger

Japanese researchers produced a 100% failure rate in fingerprint readers using $10 worth of raw materials

- "Gummy fingers" can be created from latent prints
- You leave a copy of your PIN on everything you touch
- Gummy finger use is almost impossible to detect

DansData did the same with Silly Putty

Fingerprint Sensors put to the Test (ctd)

Inspired by the Japanese, the Germans had a go too…

- Six capacitive fingerprint scanners (the most common type)
- Two optical fingerprint scanners
- One thermal fingerprint scanners
- Two face-recognition systems
- One iris scanner

All of them failed

- Some capacitive scanners could be fooled by breathing on the latent print (!!)

  The products in the versions made available to us were more of the nature of toys than serious security measures

  — c't Magazine

Fingerprint Sensors put to the Test (ctd)

Swedish student tested the latest products at CeBIT 2004 trade show

- All 9 failed

Clarkson University did it with "digits from cadavers and fake fingers molded from plastic, or even something as simple as Play-Doh or gelatin"

The machines could not distinguish between a live sample and a fake one

— Stephanie Schuckers, Associate Professor of Electrical and Computer Engineering

Face recognition put to the test

Palm Beach Airport (Florida) facial recognition system was "less accurate than a coin toss"

- 10,000 face captures/day
- Worked only 47% of the time, under ideal testing conditions
- Eyeglasses or turning your head slightly would fool it

US DoD biometrics test found Visionics system could identify an individual 51% of the time

- Iridian iris scans only failed 6% of the time (Iridian owns the patents on iris scanning)
- That's 40,000 false alarms for 1M daily air travellers

NIST found a 43% FAR for face recognition *under perfect conditions*

Face recognition put to the test (ctd)

Tampa, Florida two-year pilot yielded zero results

- Large numbers of false positives
  - Confused male and female subjects, different ages, weights
  - Woman identified as male sex offender
- Initially, officers increased threshold to get rid of false alarms
  - Result: No-one was matched any more
- Eventually, police just stopped using it
  - Logs show it was still active, but no-one paid any attention to it
  - Number of extra police that could have been employed for the projected cost: About 20

It was of no benefit to us. It served no real purpose

— Tampa Police Captain Bob Guidana

Face recognition put to the test (ctd)

2001 Super Bowl ("Snooper Bowl") scanned 70,000 participants for terrorists/criminals

- Zero matches
- Did match a ticket scalper, but he vanished before police got there

INS stopped using face recognition at Mexico/US border because it didn't work

Sydney Airport SmartGate system couldn't detect when two Japanese tourists swapped passports

Face recognition put to the test (ctd)

DARPA face recognition test, under perfect conditions, yielded

- FRR = 33%
- FAR = 10%

To detect 90% of terrorists we'd need to raise the alarm for one in every three people… it's absolutely inconceivable that any security system could be built around this kind of performance

— Image Metrics COO Gareth Edwards

Face recognition put to the test (ctd)

Boston Logan Airport trial

- 10 of the 19 September 11 hijackers boarded at Logan
- Only detected 37% of "volunteer" terrorists
  - This was under artificially good conditions: Only 40 participants going through 2 checkpoints
- FaceIt spokesperson said their figures showed an 85.7% success rate
  - 53 out of 153 "terrorist" entries = 37%

Logan Airport results will not affect plans to use face recognition to enhance passport security

— Kelly Shannon, US State Department

Face recognition put to the test (ctd)

Fresno Airport trial

- US Army Research Lab test showed that the Visionics FaceIt system correctly identifies individuals only 51% of the time

True matches (picking criminals out of a crowd) appear to be a crapshoot

— Thomas Claburn, "Smart Business"

Use of Multiple Biometrics

Some countries have proposed the use of multiple biometrics to reduce problems with individual biometrics

- This makes things *less* secure, not more

Huh?

- Combining two measures makes them stronger
- However, dual biometrics have four measures (2 × FAR and FRR), not two

When combining error rates

- One becomes stronger than the better of the two
- One becomes weaker than the worse of the two

Unless the FAR/FRR are perfectly matched, two biometrics are significantly worse than one

Use of Multiple Biometrics (ctd)

The chances of two biometrics being perfectly matched is slim

- Little chance of a match when the underlying technologies are fundamentally different

Even if (somehow) there are two well-matched technologies, their working environment will cause a mismatch

- Poor lighting will affect facial recognition but not fingerprints
- Sweat/dirt/airline food remnants will affect fingerprints but not facial recognition

Closing thoughts

The industry has set up [a] negative perception, because it has claimed biometrics are foolproof in terms of identification for security purposes. That's the wrong approach. Number one, it's not foolproof. And number two, it really isn't for security — it's for convenience

— Jim Wayman, former director,
US Government Biometrics Center

- Users frequently forget passwords, but rarely forget fingers

People like biometrics because they believe in them

— Tom Rowley, Veridicom CEO