Privacy for Tigers

Ross Anderson (Cambridge)
and Tanya Berger-Wolf (UIC)

Outline of talk

- Is tech helping the poacher more, or the ranger?
- How can we ensure that wildlife aggregation systems help rather than hinder?
- Lessons learned about AI security, usable compartmentation, and security economics
- Other recent work relevant to sustainability
- The broader questions about security and sustainability

Wildlife crime

- Half the world's wild animals are gone since 1970, mostly from habitat loss
- But where animals (and plants) are crime targets, it can be closer to 90%
- Poaching has been going on for generations
- But all of a sudden, there's tech: mobile phones, and large amounts of data
- The effects are starting to be noticed…

Early signs of concern

Namibia, 2014

Ajitora.asia

The market demand

Fighting wildlife crime

- Diplomacy: try to persuade China etc to ban imports of ivory, rhino horn, tiger parts
- Major agencies: try to bust international crime rings

- Nation states: fight smugglers, corruption; pay attention to local politics around wildlife
- National parks: ranger forces are the front line
- Will tech help the rangers more, or the poachers?

The threat from conservancy data

- 2016: Canada's Banff National Park bans VHF radio receivers after tourists track and harass animals with radio collars
- Some Indian tiger reserves have stopped using VHF radio collars
- Newer, smaller tags have GPS and crypto – including one that fits in rhinos' horns

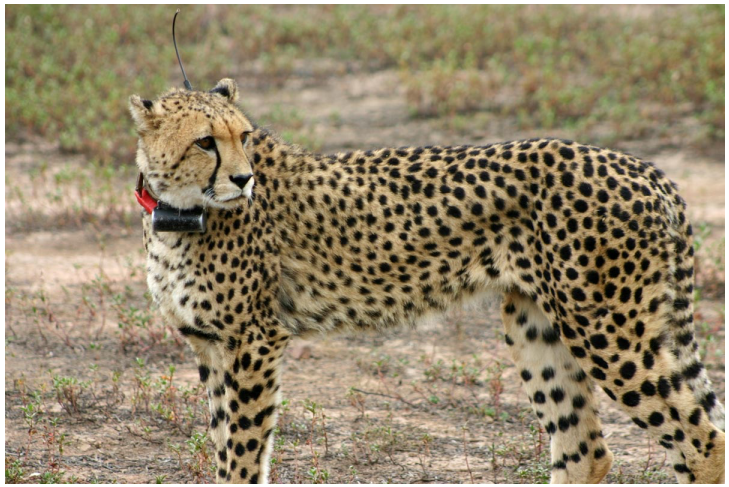

Conservancy data (2)

- Insider threat: the worst poachers are former rangers or former conservancy workers
- In rare cases, poaching gangs have colluded with a head of government
- Governments often blame their neighbors, making collaboration harder (e.g. snow leopards India / Pakistan / China …)
- Data from collars, camera traps, drones …
- Some NGOs are also mutually mistrustful

OR models of ranger / poacher

- Tambe et al 'PAWS' model from game theory
- The rangers want to know where the animals and the poachers are
- The poachers want to know where the rangers and the animals are
- The rangers want them to know nothing
- Goal: online wildlife data systems should not leak actionable intelligence to poachers

The new threat from open data

- SA's Knersvlakte nature reserve has 155 red-listed desert plant species
- 2 thieves arrested with 2000 specimens (+ data on earlier thefts, sales)
- Found via botanists' lists, JSTOR, geotagged tourist photos, iSpot
- Thieves fined $160k and deported (2015)

Wildlife aggregation sites

- As well as iSpot (exploited by the Knersvlakte poachers), there are several other websites
- The key innovation is image recognition, not just of species but of individuals
- Then link up sightings and combine behavioral, geographic, climate and other data
- I've been doing work since 2016 with Tanya Berger-Wolf who runs Wildbook
- What should the security policy be?

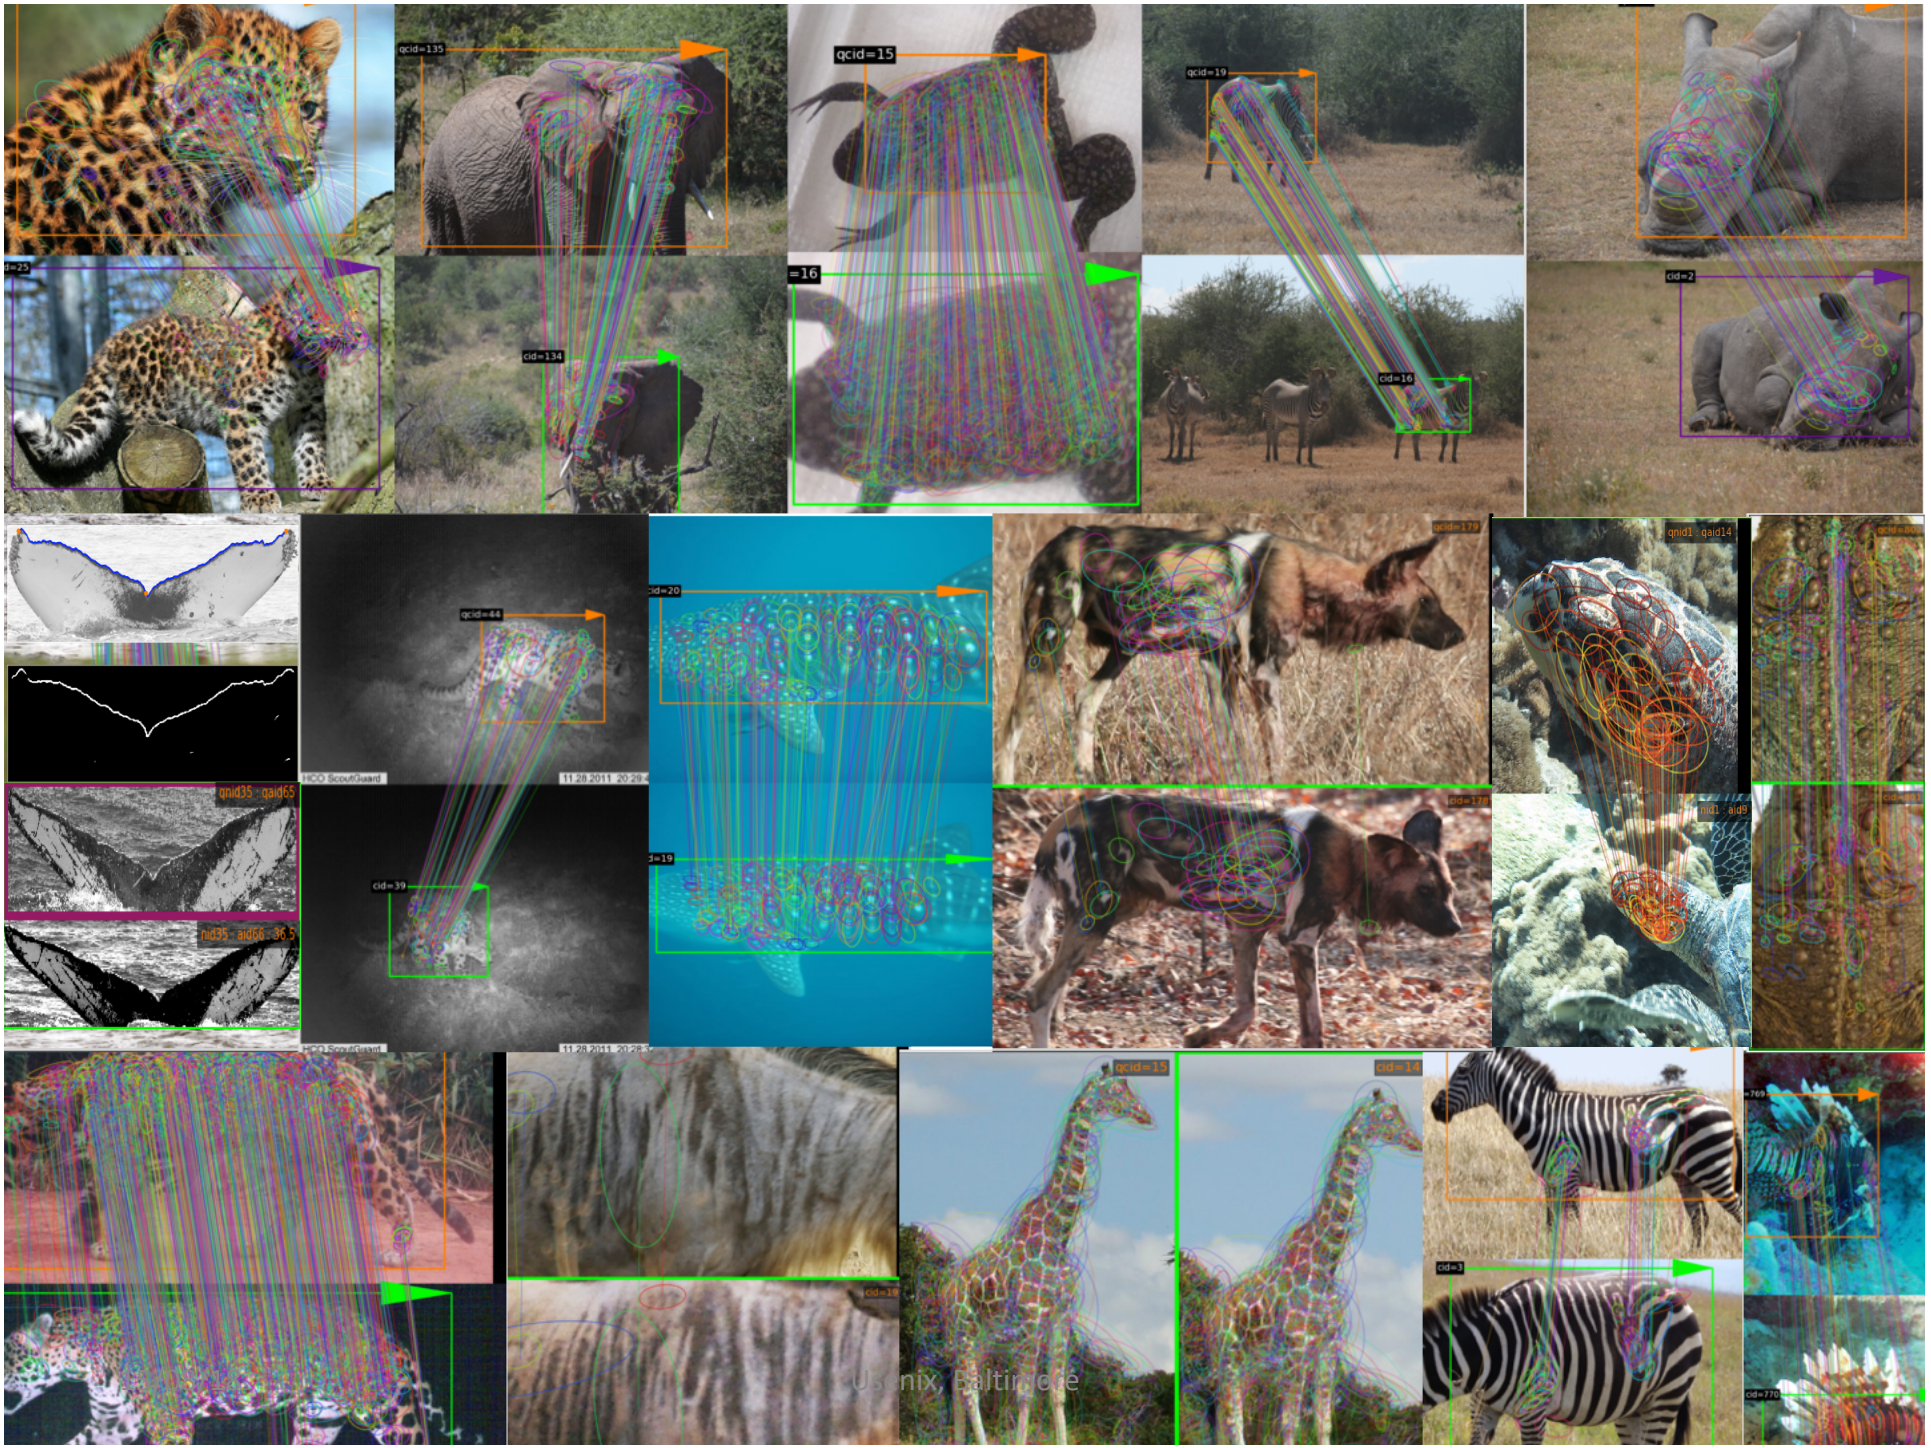
qcid=135
qcid=15
qcid=19
cid=134
=16
cid=16
cid=2
qcid=44
cid=20
qcid=179
qnid35 : qaid65
nid35 : aid66 : 36.5
cid=39
cid=19
qnid1 : qaid14
nid1 : aid9
qcid=1.5
cid=3
cid=770

How fast does the elephant population of Africa decline?

How far do whales travel?

How many bobcats are left in the world?

How many turtle hatchlings survive?

Sometimes open data can help..

- Thieves steal Coahuilan box turtles (*terrapene coahuila*) from Mexico and use forged papers to claim they were bred in captivity
- Solution: photograph all the wild ones and make the database available to the Fish and Wildlife Service
- Here, integrity's the issue

The threat model

- Former conservancy workers are usually the most competent poachers
- Even before Big Data, some would go over to the dark side with years of bush lore
- Incentives can be economic, social, political… so it's not just a technology problem
- US State Department funds multiple programs
- But now disaffected staff could leak gigabytes of tracking data, or the "lore" in AI models

The threat model (2)

- The eBird website started in the USA whose illegal hunting problem is tiny
- It now has 250,000 users and is spreading to countries where the problem is real
- Open data standards are being adopted by national parks, academia and others
- Second thoughts now: see Lindermeyer and Scheele “Do Not Publish”, *Science*, 2017
- And then there’s Flickr, Facebook, …

Flickr photos geotagged, by species

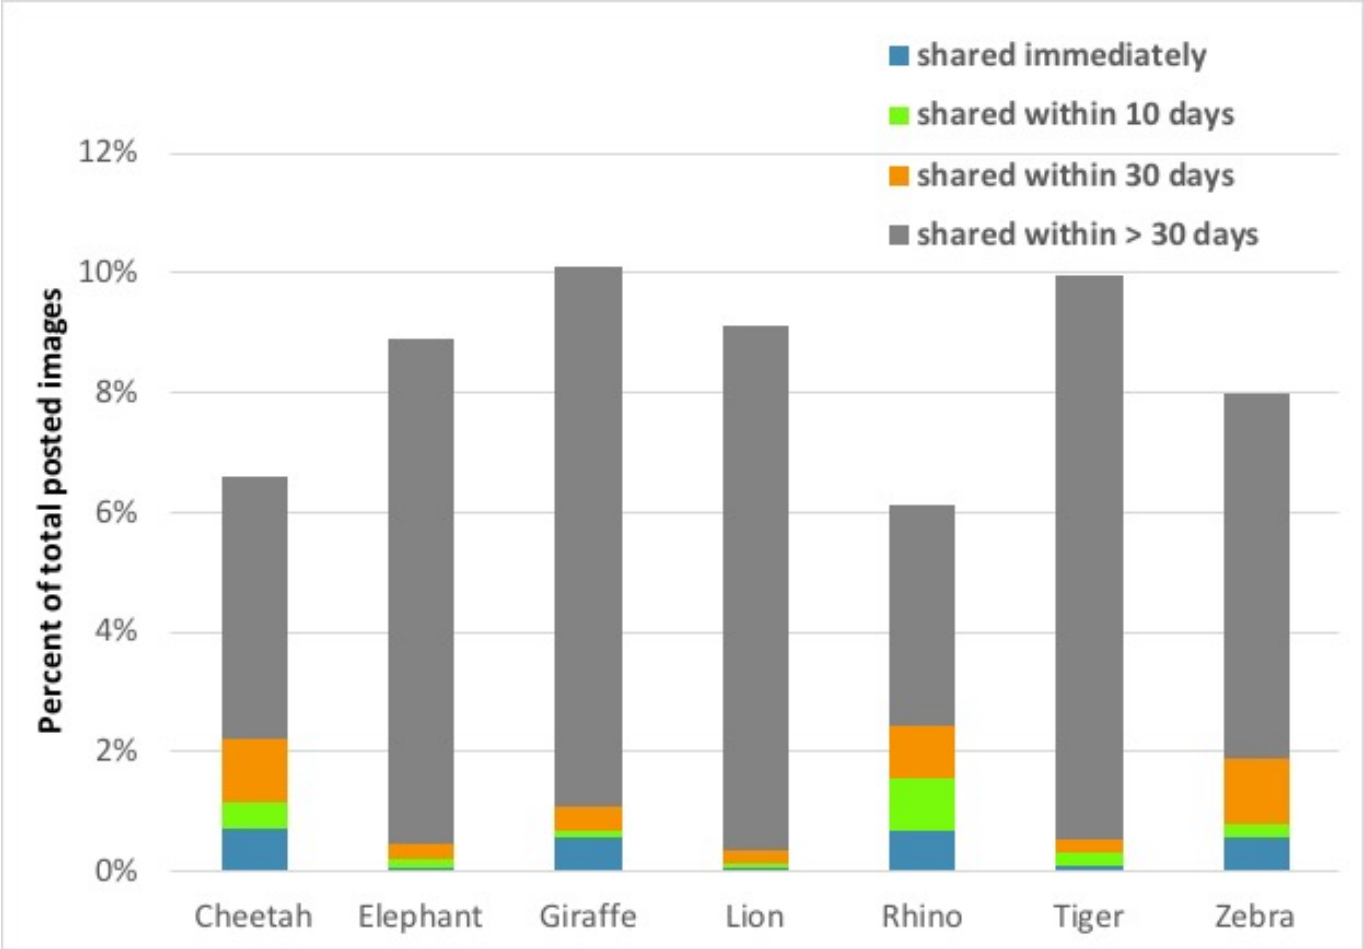

Aggregation and sensitivity

- First novel feature: while aggregation and analysis usually make medical records less sensitive, wildlife data become more so
- A single tagged photo says much less than that animal's location history
- Histories of all target animals in a reserve say "This is the time and place to hunt"
- AI enables the creation of skills and lore…

Emergent sensitivity

- Suppose you want to stop snow leopard poaching. AI can create tools such as
  - Hotspot models – where do the animals migrate?
  - Leopard spotters – graphics processing tools to spot an animal despite its camouflage
- We maybe don't want to put such lore and skills in the poachers' hands if we can avoid it
- Isn't there a philosophical gap here, between private data and public science?

The current system

- Second feature: context. Several dozen instances of Wildbook are deployed for different species / areas, so context is built in from deployment
- Zebras, leopards, cheetahs, rhinos, elephant, whale shark…
- Simple RBAC but rules vary by context (rhino location sensitive in Namibia but not in Kenya)
- Some data under conservancy control (camera traps, drones) and some external (tourist snaps)
- Inter-conservancy sharing is ad-hoc; chained sharing can lead to large compartments

The current system (2)

- A typical one-species instance of Wildbook has
  - Level 0: the Wildbook team of five people has access to everything
  - Level 1: perhaps 20 admins working for the conservancies that use the instance
  - Level 2: several hundred conservancy staff who use the system, including occasional poachers
  - Level 3: thousands of interested citizens, including probably all the tech-savvy poachers

The new system

- Wildbook has got a donation from Microsoft to “remove the duct tape” and redevelop the system with most instances on Azure
- So we can start to provide common services such as authentication and logging
- But how do we delegate administration to mutually mistrustful organisations in such a way as to keep incentives aligned?
- How do we make it scale, with only 5 staff?

Basic Concept of Operations

- As we consolidate, we have to add back context to the access control system so that both raw and derived data are classified by
  - Species
  - Location
  - Time
- The next question is how we make the admin scale despite mistrustful user groups

Context-based access control

- The direct inspiration is work on medical privacy in 1995–6 for the BMA (see my Oakland 96 paper or chapter 9 of SEv2)
- The first hospital system had rules like
  - 'A nurse can see the records of any patient who's been in her ward in the last 90 days'
  - 'A junior doctor can see the records of any patient who's been in her department's care'
- At the time we called these 'capabilities' and eventually they ended up on smartcards …

What can go wrong

- When UK hospital systems were thoughtlessly centralised a decade ago, the interplay between role and context was lost
- Access in UK health systems now depends on role + existence of a legitimate relationship between the user and the patient
- Result: all of a sudden, hospital receptionists could read psychiatric casenotes!
- So: the details really, really matter

Context-based access control (2)

- ‘Context-based access control’: see Covington, Fogla, Zhan and Ahamad, ACSAC 2002
- Their focus was non-intrusive authentication for IoT in the home; environment activates roles
- Wildlife crime gives a good new case study
- Two types of context: species, and location
- NGOs often focus on species, and rangers on location (with some mutual mistrust…)

Intrusion detection

- Third, it's not enough to find a dead rhino at a waterhole and then look at logs to see who was curious. We want to prevent crime!
- We need situational awareness – is any potential opponent taking an interest in us?
- So: who's motivated to stare at the logs?
  - The conservancy leaders (who want to know who's interested in their science)
  - The rangers (who may get shot at)

Why compartmentation is hard

- 1980s/90s: MLS systems ended up having millions of compartments, where nobody could get their work done, or one
- After 9/11: let analysts get on with it!
- In “Security Engineering” I predicted the move to bigger, fewer compartments would mean another Ames
- Instead, we got Manning and Snowden!
- But how might we do it better?

Can we do access control better?

- Following 9/11, the Mitre report on 'Horizontal Integration' (JASON JSR 04-132 2004) analysed failures of the US system of classifications & clearances
- Classifications are driven by risk but behavior by flows; classification of metadata very murky …
- Misalignment of info owner/protector roles leads to Stalinist bureaucracy to workarounds
- Proposal: risk-based dynamic compartments

Wildbook comes quite close …

- Data sources both owned (camera traps etc) and derived (from owned + public data)
- All protected by those who care: conservancy for species and rangers for territory
- Rate limits for silent alarm and for blocking
- Outstanding problems / opportunities:
  - Open-source access control tools don't quite fit
  - Usable management, by biologists, is the focus

Second-order issues

- Many security-engineering topics turn up, e.g.
  - Side channels: you can track snow leopards by tracking the goats they feed on
  - Social: tons of photos on flickr, facebook etc
  - Malware: most Androids in Africa seem rooted
  - …
- In fact protecting wildlife data has the law enforcement / intel issues as a subproblem but must be solved with vastly less resource
- Usable compartmentation is the key

Lessons learned

- A fascinating halfway house between defence and health-privacy system approaches
- Derived data makes everything more complex, as do heterogeneous ownership and mistrust
- Stable compartments can arise from context
- Usability is paramount
- You need capable motivated defenders, and they must have useful things to do

Broader lessons

- We have more and more projects now that aim to contribute to sustainability, or at least may have some impact on it.
- Two more examples:
  - Tracing stolen bitcoin: we've been making better tools which show up regulatory failures
  - Security patching of durable goods: if security patching means that cars have to be traded in after five years like laptops, the carbon cost will be huge

Second example: bitcoin

- Cryptocurrency mining burns +- 7GW
- It supports a lot of crime, from ransomware to money laundering
- It leads the young astray – I get research proposals like “fix world peace by putting elections on the blockchain”

Second example: bitcoin

- Cryptocurrency mining burns +- 7GW
- It supports a lot of crime, from ransomware to money laundering
- It leads the young astray – I get research proposals like “fix world peace by putting elections on the blockchain”
- So is there any way we could help kill it?

Second example: bitcoin

- Cryptocurrency mining burns +- 7GW
- It supports a lot of crime, from ransomware to money laundering
- It leads the young astray – I get research proposals like “fix world peace by putting elections on the blockchain”
- So is there any way we could help kill it?
- Can we trace stolen bitcoin better and thus undermine fungibility?

Hacking bitcoin fungibility

- People assumed “haircut taint” – a wallet with 3 stolen bitcoin and 7 freshly mined is deemed to have 10 marked “30% stolen”
- But an 1816 court case says you should use FIFO to trace mixed assets through an account
- Linode: Btc 46,653 stolen in 2012, haircut now taints 16,855,619 addresses (93% of total)
- FIFO only taints 254,120 (1.3%)
- Our FIFO tool is now online at taintchain.org

It ain’t what it seems

- Once we have a flashlight, we find all sorts of bad and scary things in the cave
- Most transactions are now off-blockchain, with the exchanges having morphed into a shadow banking system
- Try reading their balance sheets, and look at the regulators’ excuses for inaction
- See “Bitcoin redux”, my web page

Third example: IoT and safety

- The EU regulates safety of all sorts of devices
- They asked Éireann Leverett, Richard Clayton and me to examine what IoT means for this
- Once there's software everywhere, safety and security get entangled
- How will we have to update safety regulation (and safety regulators) to cope?
- We studied cars, medical devices and grid equipment but the lessons are much broader

The punch line

- Phones, laptops: patch them monthly, but make them obsolete quickly so you don't have to support 100 different models

The punch line

- Phones, laptops: patch them monthly, but make them obsolete quickly so you don't have to support 100 different models
- Cars, medical devices: we test them to death before release, but don't connect them to the Internet, and almost never patch

The punch line

- Phones, laptops: patch them monthly, but make them obsolete quickly so you don't have to support 100 different models
- Cars, medical devices: we test them to death before release, but don't connect them to the Internet, and almost never patch
- So what happens to support costs now we're starting to patch cars?

Security and citizenship

- 1990s: crypto war 1 on key escrow, export control and government access to data
- 2000s: security economics provides a framework for thinking about policy
- 2012 Menlo report mandating ethical process for security research involving humans
- James Mickens' splendid keynote this morning
- Now: consider environmental impacts too?

Sustainability more broadly

- As sustainability moves up political and personal agendas, it may well become relevant to ever more security research
- There's a pull from some applications
- And moral pressure from some students
- But there's no community or SIG yet, or even a 'sustainability' line in the ethics process
- Who else here is interested?