NYU Privacy Group

JHD

1 February 2017

Contents

Chapter 1

1 February 2017

1.1 Preliminaries

Introductions Kathy Sternburg directs ?? Institute. The rest of the audience are mostly from the law school or IIA: masters students, PhD, postdoc etc.

1.1.1 Executive Orders

There's another Executive Order[1], amending the Privacy Act (1974)[2], depriving non-US citizens of this. There are mixed records (US/non-US) which were created under previous administrations. The EO apparently instructs that these should not be protected if the person is not a US citizen. This has problems with the Umbrella Agreement, conditional on an Act, bilaterally approved. This was approved for "covered persons", and apparently the EO doesn't affect this. EU says "all OK", but the act requires the AG to designate countries. This was done by Lynch for EU countries. There is apparently a "court-stripping" clause, but it's not clear that's effective. "The administration could go hog-wild on this". Noted that there are legal challenges in the EU to the agreement even prior to this.

Q What about the AG?

A That's the legislation, not the EO.

***** The only rationale provided is bizarre.

KathyS This of course is a Federal action about Federal databases, and states have similar laws, which aren't directly affected. JHD privately notes that

[1] `https://www.whitehouse.gov/the-press-office/2017/01/25/presidential-executive-order-enhancing-public-safety-interior-united`: Enhancing Public Safety in the Interior of the United States. JHD later found `https://iapp.org/news/a/calm-down-trump-hasnt-tanked-privacy-shield-just-yet/`.

[2] This has rights of access, and correction.

states, of course, control Driving Licence registration, which is the *de facto* form of identification in the USA.

KathyS then switches the discussion to "sanctuary cities", which JHD didn't understand at the time[3].

1.1.2 Other Developments

Apparently NY can get locations from taxi drivers, and this is being extended to Uber.

The House has apparently passed a bunch of bills.

Action (class) against Facebook for undisclosed disclosure of birthdays.

The IRS is filing cases asking for users of BitCoin, as it's property.

1.2 Karanasiou (Bournemouth): The Intricacies of Machine Learning Algorithms

To appear in a journal, joint with Dimitris Panotis(?).

Shawn Buckle offered his personal data (consumer preferences etc.) for sale: 350 euros.

Federico Zannier including snapshots of him browsing, $2733.

CTRLIO (UK?) track your internet use, and provide you with details. Possibly you can sell these.

Cloud what happens when you add this. Also "data brokers".

Quotes Lighthill report saying human involvement is a *sine qua non* for Machine Learning. Fast forward to Microsoft's Tay.AI. Released on Twitter to see if it could learn. It became racist/xenophobic because of the tweets to it. But this is, says speaker, it "performing as it should".

Conceptual What is "foreseeable action", "causality" in tortious liability.

Definitional embodiment: what's the difference between product and service — software isn't "embodied".

Rule of Law/Due process See Pasquale and others. Algorithmic accountability. Not my concern.

Axiomatic issues There are several degrees of interaction between the human and these agents. For example, she (quoting something, it would seem) rates automatic driving on a scale of 0–5: with 2– holding the system responsible for at least some actions. We consider how the system is designed, but not necessarily how it performs.

[3]But does now. That EO is a direct attack on sanctuary cities, by name.

Opaque behaviour: we have provided data, but can't work out how it's being used. We are generally dealing with supervised learning.

Floor There are vast amounts of data available on drivers. There are also a lot of rules hand-programmed: after all we want them to be better than human drivers.

Combinatorial Explosion There are a vast number of scenarios.

Q What do we know about the background of the supervisors?

A Good question.

Deep Gives a six-layer model of the deep learning, e.g. one layer is "obstacle recognition". JHD is worried about this, as this seems to be retrospective rationalisation of the learning: how do we *know* that this is what that layer is doing?

Automated Decision Making ultimate aim is to act without human intervention.

Audit Claims that we can measure performance *ex post*, but also weights *ante*.

Q Good to hear you talk about algorithms. Pattern recognition and images are only part of the picture: cars have lots of sensors.

A I think we are discussing the algorithmic part too much, and we need to talk more about the simulated intelligence. The decision is based on data really drawn from human beings. A machine has no "intent", so what is the rôle of the machine.

Q' I agree that there's too much accountability.

Q There are a lot of angles. What aspect(s) of personhood are relevant? I like the separation between design and performance. The law has had to respond to mac/machine interactions for a long time. Europe has different concepts of personhood.

Q What did you mean by pedigree (A: see paper).

Q Personhood: this can apply to corporations. What does "personhood" buy you?

A This is my current work. ?? divides 'agents' into basic and 'hybrid'. Possibly we should treat machines as animals. Note that "deep pockets" practice will send lawyers after the manufacturers.

Q In the paper you conceptualise two of the dimensions, but not their relationships.

A I didn't want to look at "reasoning".

Q Copyright is a question of ownership, and can animals own these — "monkey selfie" case, with a Copyright Office opinion that it's not copyrightable.

A My point of departure is continental [Her phrase. JHD would have said Napoleonic] law, not common law, and I'm trying to understand how the notions I am familiar with transpose.

Chapter 2

8 February 2017

2.1 Preliminaries

Again introductions. Quite a lot of ILI fellows. One from Microsoft Research.

Q Apparently a firm has settled $2.2M a case in NJ over unauthorised access.

Q E-mail Privacy Act passed the House 2 days ago. "Fixes a loophole (in Stored Communication Act) allowing access without a warrant to e-mails over 180 days old". Last time it was derailed in the Senate: fingers crossed.

Q There's a murder case in which Amazon Echo recorded. Is this a witness, or evidence? Amazon has denied the request from information.

Q Professor reported the Facebook/Ireland case[1] over "model transfer clauses".

Q CBP collecting social media profiles, which was apparently "optional". Apparently green card holders, and in one case a US citizen, were asked for passwords.

2.2 Decision making, Machine Learning and the Value of Explanation: Katherine Strandburg (NYU)

I want to think about this from the point of view of due process. Protections vary, depending on the importance of the case being affected. Government interest, including cost, can be taken into account. Judges generally provide (not always written) explanation. Many agencies have rules about

[1] See `http://www.computing.co.uk/ctg/news/3004169/facebooks-data-transfer-mechanism-faces-legal-challenge-in-irish-court`. Note (see `http://www.reuters.com/article/us-eu-privacy-facebook-idUSKBN15M1K8`) that this is far from being just Facebook. It's in Ireland as that's Facebook's European HQ.

comments/responses. Note that jury decision, and legislation, don't require explanations. Also are cost-based decisions. Also investigations, e.g. "You are being audited by the IRS". Also when the activity would be undermined, e.g. search warrants and wire tapping (but here the warrant-granter has to have an explanation).

Descriptive explanations explain "on what basis did X arrive at Y". These may be critiqued if based on incorrect/irrelevant facts, logical mistakes (wrong rule), not credible or sincere.

Normative justification "why is outcome Y correct". Critiqued because of an inappropriate analytical approach, or not persuasive.

- Legal interpretation: almost never entirely straightforward (e.g. interpolating "except in an emergency" to a speed limit). Usually has normative aspects. Requires descriptive explanation and justification.
- Applying law to facts. This is a description task.

2.2.1 Reasons for explanations

I'm using the word "accuracy" deliberately: sometimes in ML it's used very narrowly.

1. Better decisions (correct decisions require a correct interpretation of the law, and correct application to relevant facts). Explaining helps decision makers themselves. It incentivises correct decision-making. Provide a basis fro dispute/review. Cumulatively, they promote robust legal development. Stored Communication Act is a problem, for example.
2. promote fair, unbiased decisions (which may be an intended consequence of a rule/law).
3. Promote legitimacy and trust (procedural justice literature).
4. Promote compliance (both the current person, and, if publicised, the wider community). But is "gaming the system', the evil twin here? Claims that "letter of the law" is the only requirement, and we have a right to know the law.
5. Promoting dignity and autonomy.

Q Is "this algorithm on these data" an explanation? Only understandable to a Data Scientist. Is there a parallel with DNA evidence,and the time this took to be accepted?

Q–HN Quote Trump's justification for the travel ban. Is this a good explanation?

A It's important to think about the context, and the audience for the explanation. For example, I might not understand a medical expert system, but I hope my doctor does.

2.3 Automated decision making

Do these [kinds of explanations?] have the same purposes here? Accuracy may be a motive for automated decision making, but this is a very limited meaning of accuracy. But training sets themselves may be biased or incomplete.

Do we need more and more data, but how does this conflict with data minimisation etc.?

2.4 Unrelated

(except that much of the discussion, see e.g. [Ang16], is over automated decision in the bail system.) There's an article about a recent change in the New Jersey bail system: `https://www.nytimes.com/2017/02/06/nyregion/new-jersey-bail-system.html`.

> Judge Caposela said the computer-generated scores did not tell the whole story, and were a guide, not a directive.

There's also a story about challenges to the manual bail system in Harris County Texas (essentially Houston): `https://www.nytimes.com/2017/03/09/us/houston-bail-reform-sheriff-gonzalez.html`.

Chapter 3

15 February 2017

3.1 Preliminaries

JHD Quoted [RA17] and "All training data contains biases".

? OMB has been removed from the White House site. Everything is now under the "Obama Archive". This is also not very user-friendly. There's no information about who's currently there.

? Danish Government has allotted money for an 'ambassador' to companies such as Amazon and Google.

? I met someone in Bay Area, very worried about privacy in the Cloud contracts she's working on for the Government.

HN Met a start-up who was looking for a way to provide "data stored that even we can't access".

? Google has mailed developers (someone said he got it a while ago) saying that all developers collecting sensitive data need to have a privacy policy. It appears that these policies aren't actually looked at.

? White House was looking at doing for software engineering (Cyber Independent Testing Lab) what is done, say, for fire equipment.

? Hack at company that helped FBI with San Bernardino iPhone. Apparently they used "illegal" zero-day attacks, and sold these.

3.2 Academic Institutions as Innovators and Data Collectors: Argyri Panezi (NYU)

Primary focus has been Libraries and Museums, and there is a strong connection with academia over digitisation. Learned law in Europe, where there's a strong

split between "public" and "private". Note the "Google books" story. The Google founders were connected with Stanford's digitisation projects.

Claims we are now moving into 3D connectivity, IoT, ML and AI. This leads to Smart Neighbourhoods and Smart Cities. Is this the 4th technological revolution (1=train/steam; 2=electricity; 3=cars). "People are voting with their feet in favour of city life". One of the main arguments in the Google case was over use by the blind.

The data collection consortia tend to be mixed public-private. See [Benkler-Frischmann]: layers of Infrastructure. But whom should one tell what to?

3.3

HN What are you looking at: collaboration public-private. Consider Facebook

Jessica Project ShotPlotter looking at real-time gunshot data, but also collecting a lot of conversations etc.

Prof/ NYC put together guidelines for IoT in smart cities, but there's a law enforcement exemption.

? You mentioned disability groups etc. There are issues of transparency.

A I am interested in the relationship between the data subjects and the collectors. I am not doing data research as such.

Hugo Had a paper (in health data, but general) on the relationship between subjects and collectors.

Jason So many of these projects look at the "data for good", but we need to look at the actual outputs. What data do we have on aspirations versus real outcomes. What were the unintended consequences?

? Jordan had linked ATMs to eye-scanners, which links to biometric data collected by the UN.

?? I don't know what "smart city" means. Why merge IoT, ML etc. Mentioned a SF novel by Bernard Finch; Rainbow Ends.

? It is interesting to have cost-benefit analysis directed as specific groups.

Afterwards JHD referred the speaker to [SEK14].

Chapter 4

15 February 2017: The Age of Algorithms

Speakers: Cathy O'Neil[1] and Julia Angwin[2]. This was a seminar at the NYU School of Journalism. The format was very much a conversation, with a host (whose name I didn't catch, marked 'Q') and the two speakers, and then questions from the floor.

Q What is an algorithm.

A–JA It's a recipe, whereas the data are the ingredients. We are using these to make a lot of decision that humans used to make.

A–CO'N A model is something more abstract, formalised as an algorithm. We all model, e.g. cooking dinner for my kids. Success="whether kids each vegetables", which is me projecting my agenda onto the model. No-one would even formalise this model, and my children would have a different model. By necessity, we simplify when the do this.

Q You've both made me hungry! But why are we so worried.

A–JA But these are used for real life: self-driving cars have points for what to hit — lamppost 1, mother with child 5 etc. Do we have a right to debate this algorithm? This is an important moral decision.

Q Later I want to talk about how we substitute statistics for moral decisions.

A–CO'N I have the example as innocuous as possible. The example is Facebook: it decides what to show on your wall based on its definition of success: "do people click on this". In fact, you may hate it. But in fact you're still engaged, and engagement is a (poor) proxy for Facebook's profits.

[1] Mathematician turned Hedge Fund Analyst turned freelance Algorithm Auditor. Author of "Weapons of Math Destruction".

[2] Pioneered study of the inequalities engendered by Big Data. Mathematician by training

Q Surely this is not the only algorithm.

A–JA Outrage is one of the best ways of keeping people online. We saw a WSJ that we should mix stories: serious, funny etc. Facebook chose a "neutral" algorithm, but the key question is "what are you optimising for". Almost every state is using software to predict likelihood of re-offending. Used, for example, in NY State for sentencing. When we analysed it, it was optimised to be black/white neutral. Question was re-arrested within two years. But the errors weren't the same: blacks were twice as likely to be scored high wrongly, and whites were twice as likely to be scored low wrongly. There's a real debate about what "fair' means. Quotes a Wisconsin example, where score overrode common sense. These have been is use for more than a decade, and it's crazy that I'm the first to look at it.

Q What's the equivalent of clinical trials for a drug.

A–CO'N None. Each of these were secret. NY Department of Education was using a proprietary algorithm to fire teachers. No validation. NY Justice was in fact better than most, in that it at least had a fairness goal.

A–JA The company Northpoint will show me the algorithm privately, but it's just a piece of algebra, with coefficients.

A–CO'N I would like to look at it differently. The data itself are biased. What is called "crime data" is actually "arrest data". Since the system is uneven, the system trained on data is unfair unless that mapping is unbiased. We don't have ground truth.

A–JA This has been spoken about before, but the debate didn't really kick off until we did our research. Now we can discuss morality, but only through the light of numbers. If you think about climate change, there's so much opportunity to sow doubt.

Q So many algorithms are using ML on millions of data points. If you have enough data, surely this all evens out.

A–JA Consider Obama vs Romney. The search of "X on guns" produced very different results on Google. It turned out they had trained on data before the election started. Example again of face training.

A–CO'N Google answered "Who won the popular vote" wrongly. Also "Did the Holocaust happen" gave No-4, Yes-2. Note that it takes a special kind of person to ask this question, so the ML algorithm is looking at the actions of people who ask that question.

A–JA Essentially this question was being trained on Holocaust deniers.

Q You are very articulate: how do you take this to the general public? How do you create a vocabulary that we can share.

A–JA What captured people is a human story. After my research, I spent a month searching for human stories. Bizarre previous arrests etc.

A–CO'N There I was, working in finance and disgusted by the AAA ratings. Said "the public needs to know", so wrote the book. What we need is a literacy about what algorithms are and aren't. We have to stop journalists writing down what the PR people in Silicon Valley are spouting. "Facebook friends" is a proxy for "your neighbourhood when growing up". Child abuse risk scored are another topic. Child services have a really difficult job. They want to make it "more efficient". Both FP and FN are really bad. To make an algorithm, we need to quantify the harm of both events.

Q Isn't the search for an algorithm a wish to avoid this?

A–CO'N Right. People want to avoid moral decisions. "Did Justice ever tell you what attributes to use?". "No — they trust me, because I have a PhD".

A–JA The place where we need it. Federal Law says that you can dispute your credit score[3]. The process is far from perfect, but at least it's an example.

A–CO'N Yes, you can do this with traditional scores. but you can't do "what if", on an app, which would be my definition of transparency.

Host Should we be inherently hostile? Example was environmental remediation. There was a great deal of prejudice before.

A–JA I am not against algorithms, only unaccountable algorithms. I view my job as using the scientific method: hypothesis, enough data?

Floor Are there areas where it should never be used?

A–JA I have a lot of concern over sentencing. There's a philosophical question about sentencing over the basis of a hypothetical future event.

A–CO'N It's important to decouple recidivism from sentencing. There is no scientific meta-study. There ought to be an opportunity for a real debate, which we aren't having.

Floor I'm in hedge fund myself. When we look at our algorithms, we usually have a high success rate, e.g. pawn shops.

A–CO'N Every time you're talking about a human outcome, you have multiple points of view.

Q Presumably in Dan's environmental remediation, there was a public debate. In most, the public has no say, or can find out. Private companies with machine learning secret algorithms are taking over regulatory duties, e.g. sentencing.

[3] JHD: see [Ang16].

A–JA Public institutions have a great problem hiring competent people in this area.

Floor Asset management systems, e.g. Blackrock. This shifts sums in 10^{12} order. Doesn't this have a massive influence, and shouldn't it have a moral viewpoint.

A–CO'N The financial system is so opaque. It was an open secret in finance that the AAA was broken.

Twitter When a social worker/judge has to make such a decision, how much information are they given.

A–JA In court, there's a number, and an auto-generated narrative.

A–CO'N For mine, one question was "did anyone in your family ever go to prison". But wouldn't a judge say "Your Honour, please send this person away for longer because his father went to prison" was unconstitutional.

Floor I'm a bad cook: how do I know what questions to ask.

A–JA A good start is to look at the questions, as above. The questionnaire was not public. Mine had 25 agree/disagree questions. I looked at 18,00 scores, but you don't have to do that: take a small sample, and see if it throws up questions.

A–CO'N Agree, but also ask simple questions. How often do these mistakes work against black people? Algorithms need to be made accountable, but the people who commission them have to do the accountability exercises.

A–JA NJ has just replaced their sentencing system. I am not a fan of the bail system.

A–CO'N Personality tests are another example. They have replaced HR people. We should ask "what's the evidence that your algorithm is legal?". There's a "Fair Hiring Act" apparently. This question should be asked, not of the test producers, but of the users, WalMart etc.

Floor Isn't there an incentive problem. If we paid Google to be fair, wouldn't they be fair? How much is a market capture problem? Will Blockchain help?

A–CO'N Blockchain will never help here. We don't have a model for "fair". Google was based on the premise that people on the Internet tell the truth, so now they're screwed. Also, fairness costs money — cited an economics paper.

floor In medicine, we tend to use algorithms. But when there's a bad outcome, people often say "I had a feeling". I'm a midwife. We use them, say, for depression medication.

A–CO'N The ideal data scientist would ask "what are your instincts" and build these into the model. When I worked in travel, the staff had a lot of instincts.

A–JA These automated systems are imperfect. Look at Scully: decision versus algorithm. There's then a liability issue. People are afraid to override. We should recalibrate the algorithm every few months.

Floor How much responsibility should the implementation decision-maker have? how many of these get into trouble?

A–JA This system is designed to remove accountability. Look at "never again tech" where Silicon Valley people swore not to work on removal of Muslims etc.

A–CO'N Google has an anti-ISIS program, influencing people via advertisements. This is really scary.

Q Isn't this engineering, where there are ethical standards?

A–CO'N There's no professional society out there. We could learn a lot from the people who think about Privacy Ethics. Anyone can slap "Data Scientist" on their CV.

Q What is the obligation on the data scientist to talk to the public.

A–CO'N The journalist's job is to talk to the public, the data scientist's job is to build an ethical model. I talked with some-one who built a personality test. He had a bunch of caveats, but the company overrode them, and sold it as "one size fits all".

Floor I consider myself an ethical data scientist. The algorithms are open source. That's true, but unhelpful. As a taxpayer etc., what can we do to urge our companies to be fair.

A–CO'N Law suits. There are laws not being enforced.

A–JA Facebook let you buy a house ad, with "no blacks will see this ad", but I couldn't get evidence of this. It's very hard to find these cases.

A–CO'N When JA couldn't find these ads, they built an ad for housing.

A–JA We advertised our event in the housing category, and were able to get this through with a "no minorities".

Same Floor Why not release the algorithms, and get the community to assess.

A–JA OS is not a magic bullet: consider Heartbleed. OS relies on volunteers, and there's no incentive to do boring auditing.

Same But hackathons solve some of these problems.

A–CO'N We don't necessarily know "what is fair".

A-JA We did publish our data, since criminal records are public.

Floor I'm a free software developer. The tools are very dependent on the data. Recidivism is possibly public, but housing, credit have privacy issues, and others have corporate issues. Zip codes are a proxy for race, so what do we do about this?

A–CO'N That's why we need audit.

Q How are *you* going to bring the public into this?

A–CO'N I get people e-mailing me "This happens to me: should I get a lawyer", to which I tend to say 'yes'. I don't even know what the standards ought to be for fairness.

A–JA Maybe we're asking too much of these algorithms. Are we ready to put these decisions onto a computer?

Floor I couldn't take down verbatim, but was essentially "humans are innocent until proved guilty, whereas drugs, bridges etc. are unsafe until proved safe. Should we move algorithms into the second category?".

A–both Yes.

Q Great place to close.

JHD's own view With the caveat that he really meant "algorithmic decision making", I think I agree with him, at least for decisions where wrong decisions have a certain impact. As noted in this paper [Ang16] by JA, we have that right for refusal of credit decisions, so maybe that should be the bar.

4.1 JHD Postscript

See Angwin's article in [ALMK16], the data analysis in [LMKA16], and the rejoinder in [FBL16].

Chapter 5

22 February 2017

5.1 Preliminaries

DHS Collecting passwords from non-citizens entering US. KS and I have signed the petition.

Other Looked at one guy's 'phone, decided he was a gay escort because of the dating app., and when he came back later, and his 'phone didn't have the app, refused him entry on the ground he wasn't working.

Other One citizen refused, and had all his electronic data seized.

KS Law, but mine are pre-Riley (?), and Riley is very important.

? NL Parliament rushed through surveillance laws, heavily criticised, but "reasons". Same as UK rushing.

French (being French, I sympathise with rushed anti-terror laws — JHD did he mean sympathise with the need, or the complaint? Unclear) Facebook announcement. He wants to make Facebook the critical piece of infrastructure for political engagement. You can promote your group, and pay for it to be advertised. Also gives prominence (power) to "Group Leaders".

? IoT in Germany with child's doll. This was "abolished" under laws about spying devices. Apparently no refund, as people purchased a spying device. JHD: see [BBC17] and `https://www.nytimes.com/2017/02/17/technology/cayla-talking-doll-hackers.html`, which also says that the "Heelo Barbie" doll is not for sale in Germany, where it is known as "Stasi Barbie". (17 February 2017):

> A spokesman for the federal agency [Bundesnetzagentur] told Sueddeutsche Zeitung daily that Cayla amounted to a "concealed transmitting device", illegal under an article in German telecoms law (in German).

? Trump administration will roll back Title IX protection for transgender. Apparently infighting between Education and Justice here. (Education=DeVos apparently wanted to keep the protection, but Justice=Sessions won)

JHD Mentioned Chapter 4

5.2 Privacy and Innovation

Paper by Grace Ha, Yafit Lev-Aretz and KS.

Motivation Consider the argument "Privacy Regulation Stifles Regulation". But what does this mean? Also a more inclusive cost-benefit analysis. When we looked at, say, drugs, these arguments were much more nuanced: "this particular regulation feature harms this sort of innovation" (environment), or "balanced with patient welfare" (health). There was an argument for no regulation for data in driverless cars, for example. [Goldfarkber+?2012] "but also between data-based innovation and protecting customer privacy", but argument still very general.

See Consumer Privacy Bill of Rights: 2012/15/16: no longer on White House website.

Trade Very general arguments against previous.

Innovation? Maybe some products shouldn't be built. [Tarsky: Privacy-Innovation Conundrum]: five categories. A more privacy-preserving ecosystem generates more commitment. Barriers to data sharing reduce barriers to entry. But much of this was about "privacy" rather than "privacy regulation". There are technologies that employ personal information to provide services (personalisation etc.), also Uber-like use of person information. So a 2x2: Pro/Anti and Products versus Business model. Note that Facebook now makes you (appear to) read their new policies.

Must Discuss (but briefly) privacy-related market failures. Safe Harbor etc. How do we handle the equivalent of Cap and trace versus C&C, as in environment. Process versus substantive regulation. Auditing?

Models "Notice and Consent" (no stifling effect, compliance costs are wasted, incentive to perverse innovation) and "Outlawing Data Collection" will certainly have a stifling effect, and will have to trigger different innovation.

Q Past CB analyses?

A Good question.

Ira Five categories from Tarsky: regulators say enhance innovation.

A Good point.

Q What about trust-enhancing regulatory design?

A Tarsky's argument?

HN The argument has gained great traction, so the paper will be helpful. We should see why.

Q–Carl How can we formalise this? No such devices exist yet. Hence we should favour innovation.

KS Or adopt a precautionary principle.

Q US DoD money is sums any EU university can only dream of.

Chapter 6

NYU Tandon Seminar 25 February

Introduction by a co-founder of the Center for CyberSecurity.

- There's a lot of journalism in NYU, and New York generally, so the next two points are relevant.
- How can journalists protect their sources (major research question).
- Training for journalists on cyber hygiene (short term).
- Links with NYU Abu Dhabi.
- We have 100 alumni from "Cyber ROTC Program" in the U.S. Government.
- Strong outreach program: CyberSecurity Outreach Week.

Q Silicon backdoors?

A Yes: introduction in design phase, testing (pro/con), etc.

This workshop was born out of one in Wisconsin, where we realised that half the participants were from NY, so why travel so far?

6.1 Approaches to Cryptographic Obfuscation

Kurt Rohloff (Associate Professor, New Jersey Institute of Technology). Noted that NJIT≠Stevens.

I got into CyberSecurity by a backdoor: really a control theorist by training.

- Only accessed by me: AES etc.
- Only accessed by those I intend: PKI etc.

- Outsource processing, e.g. spam filtering without reading
- FHE
- But, if your IP is an expensively trained neural net (or indeed a chip), how do you protect it? FHE won't help. Note that tactical hardware, such as a UAV, is cheap, and may well be captured. So can be convert Java bytecode to obfuscated. Theory [Sahai, Waters] but insane runtimes. Now plausible. Consider an acceptor for `1?01`. Have an offline keygen, which produces "keys" for an obfuscator, which takes this and produces an obfuscated program (again offline). This is then run online by an evaluator.

6.1.1 Our work

We are *not* using multilinear maps, though that's a great paper generator. Think of this pattern as being an FSM (Moore machine in my terminology). Input bits drive state transitions, which we represent as multiplications of ring elements. The ring elements are encoded as lattice trapdoor functions. We are (or should be) known for our lattice encryption implementations. For every time step i, we have two R matrices $R_{\{0,1\},i}$ (corresponding to 0 or 1) and two S matrices. Maybe 1000×1000 with 200-bit entries. Let the bit stream be a_i. Then compute $\prod_i R_{a_i,i}$ and $\prod_i S_{a_i,i}$ [relatively easy to parallelise], and use these and the public keys, and see if the result is "small'.

Claims that "lattice schemes are the 'new' phase of PKE. NIST is apparently pushing these." "FHE schemes are lattice-based". IBM, Microsoft have some good work in this area. We have a library that I like.

- Plaintext are integer vectors modulo small p
- Ciphertext is vectors modulo large q

"If you tell them [Commerce] that security is not defined by key length, it's much easier to get an export licence". Modern lattice encryption is based on the hardness of SVP. Note that it is a *conjecture* that this is quantum-resistant. PALISADE is our general C++library. Multiple Math backends, and hardware acceleration capability. "GPL is the kiss of death, we use 2-clause BSD". Developing an FPGA-based accelerator. Targeting a public release at Financial Crypto (April, Malta).

Q Query about Obfuscation? What about patching?

A It's like compiling: if you change the program, you have to recompile it! But in practice there's a long way still to go, at least for keygen.

Q Proof of correctness?

A We assume input program is correct.

6.2 Whys, Whens, and Hows of Achieving Privacy through Data Obfuscation

Helen F Nissenbaum (Professor, Cornell Tech and NYU Steinhardt). Director of Information Law Institute.

Working definition of data obfuscation, "the production, inclusion dots" (from her paper). Shows picture of RAF plane dropping radar chaff over Germany as an example of obfuscation. Also pictures of *C. mulmeinensis* (orb weaving spider) to produce decoy egg sacs (from her garden: she couldn't tell which was genuine).

Look at TrackMeNot 0.9 (`https://cs.nyu.edu/trackmenot/`) and [HN09].

Claims that you need "just ends" and "proportionate response". "law is lagging behind in protection of privacy", hence the load we are imposing on the search engine(s) is moral.

Also Adnauseum, that clicks on every link on a page. But what malicious ads? Claims there's a solution here. Adnauseum also presents a picture of the adverts being presented, so you can see how the ad industry is tracking you. We were approved for the Chrome store, and then kicked off and actively deactivated. "Your extension is still not compliant with our single purpose policy. We consider your extension's functionality that blocks malware to be a distinct purpose from [hiding ...]".

Obfuscation workshop 7–8 April at NYU.

Q Any lawsuits over violating terms of service [of the search engines]?

A We were worried about this, and click-fraud accusation, but no effort.

Q What about contextual e-mail: I write about my beagle, and gets bombed with ads about dogs.

A Note Google's new federated policy. Need to put pressure on companies, by showing companies that we'll resist.

Q How does TrackMeNot generate queries?

A Random selection of terms from an RSS feed that *you* provide. Therefore the database is not shared between users.

Q You are asking us to trust these applications instead: why should we trust you?

A We do have a privacy policy, and the code is open on GitHub.

Q Can you generate "undesirable clicks": pornography, Al Quaida etc.?

A We initially had complaints both about the danger and the blandness of our (TrackMeNot) generated queries. You choose the RSS feed, and there is a black-listing option.

6.3 Practical Privacy-Preserving Computations over Encrypted Data: Database Access and Publish-Subscribe Protocols

Giovanni Di Crescenzo (Senior Scientist at Applied Communication Sciences, a Vencore company; Adjunct Professor at NYU Tandon School of Engineering). JHD notes [DCMO00].

Projects from IARPA project on Security and Privacy Assurance Research.

6.3.1 Old problem: privacy Information Retrieval

Correctness client gets true answer (with high probability)

semi-privacy messages don't reveal query content to server

Efficiency sublinear in database size, for both communication/time complexity.

Let (E, FD) be an xor-homomorphic scheme with $E(a)^b = E(a \wedge b)$. Client sends $\sqrt{n}$ ciphertexts (can be lowered using recursion). "Relatively efficient by the standards of theoretical cryptography, but ..." (large load on server).

6.3.2 TA1 Problem: database retrieval

Complex queries, and client can't distinguish between a non-compliant query and a query with no matches.

data dynamic, and no re-initialisation needed.

Privacy on both sides, i.e. only relevant data sent to client, as well as above.

Performance ≤ 10 cost factor on communication over standard query with TLS. Datasets in 10TiB range.

No idea why this is a US Government project: conjecture was no-fly lists.

ACS solution. Distinguish data owner from the server, with index structure allowing operations over encrypted data (equality-preserving, symmetric encryption). There is an off-line protocol between Client and Owner exchanging metadata information. Offline protocols between Server and Owner to database structure management.

Client asks "attribute j is $f(k, v)$". Then server looks for $f(k, v)$ (we used a B-tree), and returns encrypted record (which was encrypted by owner).

For multi-occurrence, dynamic databases, need the owner to pad keywords with multiplicity counter.

But have to support range queries etc. The key idea is rank characterisation. Client gets $\text{rank}(v_0)$, $\text{rank}(v_1)$, and then queries as above for each intermediate rank. Rank characterisation doesn't work if there are multiple occurrences:

fixed as $\text{rank}_{lower}(v_0)$ and $\text{rank}_{upper}(v_1)$. Also client learns the rank, so we add a nonce shift for the query.

Also more problems with unordered domains, which JHD didn't understand.

6.3.3 TA3 Problem: publish-subscribe protocols

Data items come with topics, and subscriptions come with interests. Should allow keywords, conjunctions etc. in subscriptions. Privacy requirement is that no-one knows more than the number of subscriptions. For each item, subscribers get the item, rest get at most the number of topics and the length of the data-item. Publisher learns nothing. Same ≤ 10 requirement.

ACS solution, again a three-party one. Linear in number of matching subscribers. Again there's a "matching server", using equality-homomorphic encryption. The major technical tool is Conditional Oblivious Transfer, a variant of Rabin.

Push-type protocol works, we prove pull-type is equivalent to database retrieval.

Q1 can we relax the trust assumptions on the server? Note that we are resistant against simple intrusion on the server.

Q2 What are the next problems with practical 2/3 party solutions?

Q–HN Who has to co-operate?

A In the PIR problem, clearly the two have to cooperate.

Q–HN Did you say that there is no proof?

A We don't have a proof of privacy against the server.

6.4 Cyber Threats in Energy Sector: Challenges and Possible Solutions

Mooi Choo Chuah (Professor, Lehigh University). Motivation: large attack in Ukraine, 225K customers form 3 distribution companies on December 15 216. Spearphish (Word document), credential theft, VPN access, Workstation Remote, Control and Operate, Tools and Tech. I fell victim to a Spearphish in the name of our new President, until it wanted login, when I stopped. Malware was a variant of BlackEnergy3. The attackers knew how to write custom malicious firmware, and render field devices such as serial-to-Ethernet inoperable (and irrecoverable). Generate tons of 'phone calls to prevent customers from reporting outage. Note that they also took control of operator workstation. Took out 27 substations. Same techniques used on mining company and rail operators.

Note that Internet-based IDS are not sufficient for SCADA network, as traffic flows very different.

Use OMNET++ to simulate a SCADA network.

Also see HPFortify 2014 (October) attack on Dyn scored 1Tbpos

Q Firewalls? ISP requires them.

A But it happened: even simple IP address analysis would have shown this.

Note smart thermostats, that can be hacked to believe the temperature is wrong, can cause power outages. Displayed "motherhood" list of "to do"s. Note a website called Shodan that had a database.

6.5 Editing the Immutable Blockchain

Giuseppe Ateniese (David[1] and GG Farber Endowed Chair in Computer Science and Department Director at Stevens Institute of Technology). Developed with Accenture.

The Blockchain remains decentralised and immutable, but there's a Plan B if things go wrong. This presented the benefits, but makes it viable for enterprise use. Much press. There are proposals to use Blockchain for medical records, identity management, IoT devices, smart grid, Real Estate Title Insurance, supply chain management, Post-trade services (this one I actually believe in). Is this hype? Most people don't understand it: I had to read multiple papers, and then teach a class on it, to understand it myself.

Note, this is a decentralised database (nothing new here). Every node stores a copy. How do we know which is right? BitCoin uses proof-of-work. In a private permissionless one, just use voting. For practical reasons, the database contains blocks. The BitCoin genesis (first block) was generated 8 years ago. After nodes have collected a block of data, a consensus mechanism selects a node to write the new block onto the chain. Block B_i contains $H(B_{i-1})$. Everyone then stores this block as well, even if it contains pornography etc. Note that BitCoin have have hard forks: two nodes containing $H(B_{i-1})$. BitCoin says you always select the longest chain. You can only add a block every 10 minutes, which means that you can't change distant history.

Note that GDPR requires *deletion* of information, which is not soluble by a retraction record. smart contracts are going to be amazing. But there are technical issues currently. The crowdsourced capital fund DAO, recorded on the immutable Blockchain Ethereum, had $60M cryptocurrency stolen.[2]

The BitCoin database is currently 100GiB. I would rather not do this, but then I don't have the full security. So what about "Blockchain for IoT" claims. Note that our invention is intended for permissioned systems. "Only trusted administrators acting of agreed rules can edit".

In our model, there's a padlock on each link, and the key is in an "emergency box". Easier said than done, but we have a padlock design that works for all consensus systems. The key is also an issue (note that we currently trust our PKI's keys' security). Divide the key into shares.

[1] That is the Internet Farber!

[2] JHD: see [Ano16].

6.5.1 Details

Standard Blockchain.

1. $s'' = H(ctr', G(s', x', r'))$
2. HashPrev (s')
3. Transactions (x')
4. Nonce (ctr')
5. Randomness (r')

G has to be a collision-resistant hash function. We use a G which is a "chameleon hash"[3], i.e. with the key, one can find a collision. Use HyperLedger's Open Source implementation. However, we do (or at least might) want to leave a "scar". And not even the trusted party can remove this.

The "right to be forgotten" is challenged: an ECJ case about a Dutchman who has found that his identity was uploaded into the BitCoin Blockchain is stalled.[4]

Currently working on details of key storage using trusted hardware, and combining this with partitioning. Note that if the key is destroyed, we re back with classic Blockchain.

Q How do you guarantee that everyone will update?

A Good question — there's nothing you can do, since it's decentralised. But the point is that I *can* remove the offending block in the new system, so my "memory" is now compliant.

6.5.2 Subsequently

Q You've patented the chameleon hash?

A We patented the whole thing.

Q–JHD You keep saying "collision", but don't you mean "second pre-image"?

A Details.

6.6 Securing the Software Supply Chain

Justin Cappos (Assistant Professor, NYU Tandon School of Engineering; Member of the Center for CyberSecurity) Basically Git + GCC, maybe with Travis CI, various tests, then ship a Debian package. Apparently the back door in

[3] JHD: see, e.g. [ADM04].

[4] JHD: this seems to be `https://medium.com/@hankmoonie/mans-right-to-be-forgotten-case-stalls-after-he-is-found-on-the-bitcoin-blockchain-1a32c4fc0963#.byfrieoxv`.

Juniper routers was added this way, ? by NSA. Also Linux kernel. See [Tho84]. There have been problems with Git's signing (now improved). Stress on "reproducible builds", which isn't as easy as it seems. See work at the ToR project, which is currently at 80% reproducible! Existing distribution systems also have problems — see our previous work [SMCD10].

6.6.1 In Toto

(see `https://in-toto.github.io/`) is our project to look at the entire chain Consider a Project owner, who produces a "layout" describing the processed. Various functionaries produce signed metadata. Then users can verify this.

> That's obvious, I could design this.

Until you try it, yes. I have a talk "87 things that don't work". Things we need are rules like "`.o` doesn't go to package".

Alice project owner

Bob is the developer, to clone the repository and make the changes.

Carl the packager.

Usual steps, with in-toto recording added. In practice we would add these to our scripts. Ran a demo, and screwed up the packaging, so the verifier complained. This is integrated into various communities, including Haskell.

Q What happens if the build server is compromised?

A In our sample layout there was only one party doing the build. But you should be doing reproducible builds, which is the answer.

Q Most significant code bases have external includes.

A You can track the includes as far as they go using our system, after that you merely know what it is you depend on.

Q Define "integrity".

A Only what is intended is allowed to happen. Analogy: your aspirin is not tampered with.

Chapter 7

1 March 2017

7.1 Preliminaries

Ira There's a tool that uses geolocation to find social media postings that are "relevant" to a case.

KS A lot of privacy case law comes from Workers Compensation cases.

? Amazon has files a brief to quash their subpoena, saying it infringes 'Alexa's' First Amendment rights'.

KS This will surely end up

? Cloudpets. JHD added the "qwe password" training video [The17].[1]

KS Foreign Surveillance Act: to be reviewed, e.g. section 702. Came up in vetting of new NS Advisor.

? GDPR versus erasure.

? A story of foreign writers. Also French businessman using Uber sues.[2]

7.2 The effects of filtering, Luise Papcke (Columbia)

I come from policy theory.

[1] There's more subsequently, and there seems to be a server-side problem as well. [O'N17, Hun17]. In the light of the fact that it's a mis-configured publically-indexed MongoDB database that's the source of the leak, and the actual break is a server-side break not a client-side break (even though the client side is weak), JHD finds it very hard not to see the case for regulation and auditing.

[2] `https://www.washingtonpost.com/news/morning-mix/wp/2017/02/13/a-frenchman-sued-uber-for-48-million-a-telltale-app-glitch-says-he-made-his-wife-suspect-he-was-cheating/?utm_term=.34394f26c1dc`.

Discrimination Notes that 'Big Data' allows us to re-introduce discrimination that was previously forbidden "questions you can't ask". Problems of spurious correlation and categorisation.

Targeting This is claimed to produce better services — "smart cities". but this leads to "DIY surveillance" (JHD thinks she means "DIY counter-surveillance") and image management. But what happens to "Targeting of voters".

Personalisation "Filter Bubbles" [Pariser2011]. We don't know how these filtering algorithms work.

Q How does Helen's theory [of obfuscation[fit in.

A Doesn't that imply things have already gone wrong.

Q See Durham UK Louise Amoore. [Amo11]

A Noted

Q ??

A I come from "All categorisation is bad".

Q From ML point of view, there are real problems with going into the depths of these. Collaborative filtering is one issue. How does one know one isn't discrimination against a protected category by the back door?

Q If we're all forced to do this "DIY surveillance", doesn't give us empathy with the marginalised classes.

HN I was struck by the fact you mentioned that the ML machine can identify non-historical classifications, and these aren't understood/protected.

Q Gary Marks paper on online/offline surveillance. Also, were you really talking about methods of filtering?

Q Frank Pasquale "The Black Box Society" [Pas15], "The Policy of Algorithms".

A Thanks. I am looking at ....

Q I'm a computational social scientist. That article (?) is dated: Lazer is more modern. [LKKV14].

KS In law, some classifications are deemed "invidious", some not. Can we maintain this? If we don't know what the categorisations are, how can we engage in (counter-)surveillance.

A It's bad enough that we know there's unequal treatment for race, but we can at least try to counteract that. What about other kinds?

Q "To say that we're citizens is to say that we're agreeing to give equal chances etc., not just favour the rich etc." So insurance is a problem in this context.

Chapter 8

22 March 2017: current topics

General introduction: we didn't do "current topics" in view of the subject.

8.1 Smart Toys: Privacy and Data Security

Children's Privacy and Data Security are the main issues. Cayla was the case considered.

1. Parents' notice and consent
2. data retention
3. purpose limitation (using audio files to improve ML algorithm, for the toy and other products/services[1]). This is in the fine print, but not clear.
4. sharing with third-party service providers
5. advertising embedded in the smart toy

1. password protection
2. encryption
3. data breach notification (or lack thereof), slow reaction to notification.
4. risks for children: inappropriate content, grooming, identity theft.
5. risks for the home: entry-point into home network.

? There's some work at U. Washington on "toys that listen".

?(bis) These toys can be used to spy on a custodial parent.

[1] The company also works for law enforcement and intelligent services.

Q:KS A couple of years ago one of our students wrote a paper about smart toys: "which parent": kid goes to another's to play, and is recorded.

Q:HN Isn't it the smartness that's the real concern.

A There was a toy with ML that learned to speak like a child. "Amazon echo".

Q I can't see parents getting used to this new form of advertising.

Q Is there a notice that these record?

A In principle, but certainly not in plain language or on the box.

Q Discriminatory: affluence, and language (non-native). We have a rubber duck that interfaces with other IoT devices.

Q:KS COPPA: Children's Online Privacy Protection Act?

? "hello Barbie" was marketed to get round COPPA on age grounds. This has a lot of backlash, though.

KS Getting Congress to amend COPPA, if necessary, is one of the few doable things in the privacy area.

Q But what can the app on the 'phone share? Is this also a backdoor into the parents' phones?

Post Meeting Note: JHD spoke briefly to the speaker (Dutch?). She had mostly gone on what she knew from the EU side. But she had heard of CloudPets [FB17b, FB17a, The17, Hun17, Zor17].

8.2 Wikileaks on CIA Hacking Revelations

- Wikileaks will not distribute the "weaponised" cyberweapons.
- Claimed CIA has developed and absorbed malware targeting smart devices (smart 'phones, TVs, cars etc.) — "Weeping Angel" is a Samsung smart TV exploit to watch viewers
- CIA hoarded zero day exploits
- CIA knows "fingerprints" and repackages their code to look like other people's
- US Consulate in Frankfurt is the spybase.
- Snowden tweeted that CIA hasn't cracked encryption, merely subverted the apps.
- "Fine Dining" is a close-range menu of hacks, concealed under a Prezzi front-end.

- Ricky bobby uses DLL exploits
- root kits, and McAfee has, since this revelation, released a root kit detector.
- Wikileaks apparently conducting a Twitter poll. Claim is all the leakers "are millennials".

Q:Ira There's nothing really new here: Malware kits are readily available on Darknet. The Russians and Chinese are already in this market. A lot of the coverage is pretty naïve.

A It does seem that CIA had not been alerting companies.

Q There are agencies like DHS in this space as well. Is the CIA actually creating a market for zero-day? Wikileaks seem to be stepping into a regulatory rôle: if you promise to fix it in 90 days, we'll tell you.

Q Note that this "fingerprinting" means that "Russian fingerprints" can be attacked as not genuine.

Q The US has had a "responsible disclosure" policy for some time, and there's an inter-agency process. If a zero-day is being actively used by one agency, we don't want another agency disclosing it. Note that the "proximity" requirement limits the scale. Fundamentally, how did these get disclosed: looks like very poor CIA security.

JHD Many memory stick studies [Sam16, TDF+16]. But shouldn't the CIA be doing this?

Q Aren't there a lot of restrictions, e.g. DMCA, which get in the way of legitimate researchers. Case of a disclosure resulting in hacking prosecution.

Sebastian I am also a defense contractor working on vulnerabilities: there is a national vulnerability database. There's a history of vulnerabilities being known but not fixed.

Q Who's the target? Look at Snowden.

A There's no explicit target mentioned.

8.3 First Amendment rights for Amazon's Alexa

This listens passively for "Alexa", not storing until (1 second before) that is said. Does it record during the playing of the music? Amazon claims not. Government wanted 48 hours of data from the device. Amazon had a 90-page brief in response, described by KS as "short on precedents".

- User speech is protected

- The responses contain user speech, and also Amazon-generated speech
- Too general
- Riley v California (2014) reveal more than any isolated record
- Monica Lewinsky/Kramerbooks case
- The case *Search King v Google Tech*: ranking of search results is a constitutionally protected (First Amendment) opinion; also Zhang v Baido Inc., being argued over currently (but also has jurisdictional issues).

Defendant (like Kramerbooks) gave in: so there's no actual case. But wouldn't every criminal defendant will probably do the same, so this becomes permission by default.

Q The Government gets 1A warrants all the time, so what's new?

KS This is a lot like wiretap laws, which were sometimes too vague. There's a 4A/1A conflict here. This is also a subpoena-ish case.

A

Chapter 9

29 March 2017

JHD noted yesterday's House of Representatives vote to roll back ISP rules: great for the ISPs, but also for VPN services [Mer17]. The NY Times[1] opinion piece (written by an ex-FCC chairman) points out that this gives bizrre asymmetry between cell phone calls and internet access.

9.1 Preliminaries

Toys §8.1 I'll go back and look at these other suggestions.

Microsoft Academic liaison was present: long-standing friend.

HN What's the privacy (se above): I hear disaster, and "no big deal". JHD cited NY Times, "disaster".

Also EU proposes rules on encrypted communication. It was also pointed out that the recent attacks didn't use encrypted communication.

Google maps has changed its settings, allowing the option to share your location.

Supreme Court has instituted a FA rule re an NYU law saying that you have to advertise the higher price if there are cash/credit card differences. This seems to extend the definition of "speech".

KS Registration page up for conference on Algorithms and Explanation. 27/28

[1] `https://www.nytimes.com/2017/03/29/opinion/how-the-republicans-sold-your-privacy-to-internet-providers.html`.

9.2 The Privacy/security Tradeoff in the Era of Surveillance Capitalism abd Targeting-and-convincing Infrastructures

Speaker is a French Engineer, who wanted to understand the engineer/policy interface. At Kennedy school. Their policy fellows thought at first that I was a dangerous libertarian. I'm now at Columbia.

The state wants personal data for security reasons, and citizens want privacy. In France there was a Government database being built (merging various), which caused the Loi informatique et libertés (1973–78).

Ban on encryption export, blown by PGP.

Clipper chip, designer by

Four horsemen of the infocalypse: child pornography, terrorists, organized prime, software pirates.

Snowden moved the debate from law enforcement to intelligence agencies.

Now debate about backdoors etc., quote from Keith Alexander (2013): "everyone understands that if we give up a capability that is critical to the defense of this nation, then people will die".

Vance *et al.* op.-ed. in NY Times, "On behalf of crime victims the world over, we are asking whether this encryption is truly worth the cost".

Apple/FBI wanted "this phone", not data in general.

Npw we interact with companies, Google etc. (JHD: post FCC-neutering) and ISPs. So the Government needn't colect, as the companies do it for them. Therefore, do we want privacy form the state or from companies?

Zubuff2015 "Big other: surveillance capitalism ...".

T&C Infra This is the 'fake news" debate. This is not new, [Albright2016] investiagted the fake news sites. A network triggered to spread this news in a very targeted way. Bolongs tothe Google ads/facebook/witter newsfeed ecosystem. So maybe companies now realise that holding personaldata is a risk as well as an asset.

CIA Can be viewed as privacy principles, even though thought of as information security principles.

IoT will only increase data collection, and will give security issues.

9.3 Q&A

Jason There's a difference between National Security arguments, and privacy arguments. No privacy argument can trump a warrant. The Apple case was Security versus Security. Also, the whole "company bubble" has a much longer history over cables etc. But maybe the companies are shifting. Perhaps Snowden forced companies to take a side.

Ira It's more nuanced. "Victory over export controls" was shown by Snowden to be camouflage for other controls. I don't think the "transparency reports" are meaningful. I liked the "convincing infrastructure" point. But why si privacy a security argument?

A I don't buy the Russia story in full. The NSA are taking this as a security question. DNI was condemning a targeting operation. It's both a security and privacy threat. I don't think this will work in France because of our data protection regulations. Also fewer people use Twitter/Facebook.

Microsoft Stanford has been studying the polarisation of political speech in the US.

A The Yahoo breach was 33cents. But this cost is likely to increase in the future. GDPR.

Q Kate Barberg (?) Seattle Times today has been tracking crisis communication which has morphed into fake news. There's also a panel on fake news creators, who earned $200K/month from clicks. It's a business as well. "Surveillance Intermediaries" by Rozen. . . .

Q Framing this in securitisation terms might create discrimination.

Q Filing showed that Trump used a variety of information agents. Accuracy on targeting is not as high as you might think. There's also a spectrum, not fake/nonfake.

A I'm not trying to draw a spectrum of "fake", just looking at technology.

Q Look at Judy Cohen's work.

JHD noted that there was a dialogue of the deaf: the Americans couldn't hear "security" other than as "national security".

Chapter 10

5 April 2017

10.1 Preliminaries

HN NYTimes article: Banks are fighting non-regulated providers over consumer data. Example reminds her of MOOCs, FitBit etc., of people who sidestep regulators by describing themselves as service providers. Someone compared this with Amazon cash (newly launched).

?? Mobile health data: adverts targeting women near specialist health facilities.

JHD notes [Ger17] and [Bur17]. Also JHD writing response to `http://www.parliament.uk/business/committees/committees-a-z/commons-select/science-and-technology-committee/news-parliament-2015/algorithms-in-decision-making-inquiry-launch-16-17/`. Also French inquiry.

? ND passed a law targeting the pipeline protestors. ?CLAM?

Ira FCC head said that ISPs had never intended to sell data, and that now that Congress has acted, FCC and FTC can work together.

KS+others This one seems to have affected more mainstream people.

?? Uber using beharoiral studies on their drivers.

KS "extreme vetting" seems to include all social media information, also a bill requiring a warrant to look at US citizens cellphones at the border.

10.2 Privacy as Commons: Melanie Santippo(?)

"Commons" really means community management. Structures within communities. So privacy is "appropriateness of flow". Looking at how processes work in practice, rather than formal rules.

10.2.1 Example: Chatham House Rule

> "Neither the identity not the affiliation of the speakers, nor of any other participant, may be revealed" [1927; 1992; 2002].

Who decides? Does it establish enough of a sense of security? Examples: UNESCO, WEF (expert networks)[1], Brookings.

Craig Health Nwtwoks, especically multi-state ones.

A See book on Medical Commons with KS.

Tobias Measures taken to create the authority. What enforces them? Divorcing the statement from the speaker can change it.

A Inhibition is a worry. We have essentially chosen successful networks. "Indigenous commons" have been less successful: UNESCO-sponsored indigenous language studies sometimes failed.

KS Itmatters who the members are. Doctors only, or doctors and patients.

A Note a rare diseases group of patients who feel that they need more sharing.

Q This was my first conversation with Brett. We were saying that contextual integrity is not enough.

A Layers of contexts.

Q I think we discuss trust some more.What makes people pull back from optimal sharing? Expert groups in the EU (DG Agriculture), where the experts have "double shoes".

A EU expert groups are interesting. We aren't explicitly dealing with trust, but governance is included. We have seen examples with newspaper editors.

Ira Not sure what you're looking at.

A Ideas are in some sense personal. So this is what we are examining. Look at the constraints on commercial applications of genome commons.

Q Fascinating. Local or global networks. Do cultural differences matter. You ere distinguishing procuding knowledge from producing rules. Does this affect behaviour. Wikipedia versus WEF

Q Nice to see an application of Knowledge Commons. Is "legitimacy" from the participants. Film from David Burnett following the production of GDPR. The silences are interesting as well.

[1] Looking only at finished networks. One of HN's former students is in charge of configuring the infrastructure.

A US Congress as a Knowledge Commons: who can input? This has nothing in common with the other cases. Legitimacy to outsiders is interesting, e.g. medical commons.

Q Literature on participatory democracy. Does participation have inherent value. Look at "occupy".

A My previous work dealt with whether PD affected legitimacy.

HN Two things: contextual integrity talks about the legitimacy of norms, there's also procedural legitimacy. what happens if doctors share among themselves breaking HIPAA.

floor There's an instagram of doctors, which has been infiltrated by journalists: "we have a case like this that we've never seen before".

Chapter 11

12 April 2017

11.1 Preliminaries

Q

A

11.2

Chapter 12

17 April 2017

12.1 Preliminaries

Q

A

12.2

Bibliography

[ADM04] G. Ateniese and B. De Medeiros. Identity-based chameleon hash and applications. *International Conference on Financial Cryptography*, pages 164–180, 2004.

[ALMK16] J. Angwin, J. Larson, S. Mattu, and L. Kirchner. Machine Bias. There is software that is used across the county to predict future criminals. And it is biased against blacks. `https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing`, 2016.

[Amo11] L. Amoore. Data Derivatives. *Theory*, 28:24–43, 2011.

[Ang16] J. Angwin. Make Algorithms Accountable. `http://mobile.nytimes.com/2016/08/01/opinion/make-algorithms-accountable.html`, 2016.

[Ano16] Anonymous. Ethereum/TheDAO hack simplified. `http://blog.erratasec.com/2016/06/etheriumdao-hack-similfied.Ethereum/TheDAO%20hack%20simplified`, 2016.

[BBC17] BBC. German parents told to destroy Cayla dolls over hacking fears. `http://www.bbc.com/news/world-europe-39002142`, 2017.

[Bur17] G. Burton. Indiegogo-funded IoT start-up bricks customer's gadget over negative review and 'rudeness'. `http://www.computing.co.uk/ctg/news/3007879/indiegogo-funded-iot-start-up-bricks-customer-s-gadget-over-negative-review-and-rudeness`, 2017.

[DCMO00] G. Di Crescenzo, T. Malkin, and R. Ostrovsky. Single Database Private Information Retrievel Implies Oblivious Transfer. In B. Preneel, editor, *Proceedings EUROCRYPT 2000*, pages 122–138, 2000.

[FB17a] L. Franceschi-Bicchierai. How This Internet of Things Stuffed Animal Can Be Remotely Turned Into a Spy Device. `https://motherboard.vice.com/en_us/article/how-this-internet-of-things-teddy-bear-can-be-remotely-turned-into-a-spy-device`, 2017.

[FB17b] L. Franceschi-Bicchierai. Internet of Things Teddy Bear Leaked 2 Million Parent and Kids Message Recordings. https://motherboard.vice.com/en_us/article/internet-of-things-teddy-bear-leaked-2-million-parent-and-kids-message-recordings, 2017.

[FBL16] A.W. Flores, K. Bechtel, and C.T. Lowenkamp. False Positives, False Negatives, and False Analyses: A Rejoinder to Machine Bias: There's Software Used across the Country to Predict Future Criminals. And It's Biased against Blacks. *Federal Probation*, 80:38–46, 2016.

[Ger17] J. Gershman. Court: Warrant Needed for Vehicle 'Black Box' Data. *Wall Street Journal 31 March 2017 Issue 74*, 269:3–3, 2017.

[HN09] D.C. Howe and H. Nissenbaum. TrackMeNot: resisting surveillance in web search. *Lessons from the Identity Trail: Anonymity*, pages 417–436, 2009.

[Hun17] T. Hunt. Data from connected CloudPets teddy bears leaked and ransomed, exposing kids' voice messages. https://www.troyhunt.com/data-from-connected-cloudpets-teddy-bears-leaked-and-ransomed-exposing-kids-voice-messages/, 2017.

[LKKV14] D. Lazer, R. Kennedy, G. King, and A. Vespignani. The parable of Google Flu: traps in big data analysis. *Science 6176*, 343:1203–1205, 2014.

[LMKA16] J. Larson, S. Mattu, L. Kirchner, and J. Angwin. How we analyzed the COMPAS recidivism algorithm. https://www.propublica.org/article/how-weanalyzed-the-compas-recidivism-algorithm, 2016.

[Mer17] C. Merriman. US web users' browsing habits to go up for sale as Congress votes to roll-back FCC privacy rules. http://www.computing.co.uk/ctg/news/3007437/us-web-users-browsing-habits-to-go-up-for-sale-as-congress-votes-to-roll-back-fcc-privacy-rules, 2017.

[O'N17] P. O'Neill. Internet-connected teddy bear company hacked; 2 million parent-child voice messages held ransom. https://www.cyberscoop.com/internet-connected-teddy-bear-company-hacked-2-million-parent-child-voice-messages-held-ransom/, 2017.

[Pas15] F. Pasquale. The black box society: The secret algorithms that control money and information. *Harvard University Press*, 2015.

[RA17] L. Rainie and J. Anderson. Code-Dependent: Pros and Cons of the Algorithm Age. http://www.pewinternet.org/2017/02/08/code-dependent-pros-and-cons-of-the-algorithm-age/, 2017.

[Sam16] M. Samarati. 17% of US employees would use a USB stick found in the street. http://www.itgovernanceusa.com/blog/17-of-us-employees-would-use-a-usb-stick-found-in-the-street/, 2016.

[SEK14] J.Y. Song, J.K. Eom, and S.I. Kim. Evaluation of Elderly Mobility Based on Transit Card Data in Seoul. *Promet - Traffic&Transportation*, 26:281–290, 2014.

[SMCD10] J. Samuel, N. Mathewson, J. Cappos, and R. Dingledine. Survivable key compromise in software update systems. *In Proceedings of the 17th ACM conference on Computer and communications security*, pages 61–72, 2010.

[TDF+16] M. Tischer, Z. Durumeric, S. Foster, S. Duan, A. Mori, E. Bursztein, and M. Bailey. Users Really Do Plug in USB Drives They Find. https://zakird.com/papers/usb.pdf, 2016.

[The17] The Guardian. CloudPets stuffed toys leak details of half a million users. https://www.theguardian.com/technology/2017/feb/28/cloudpets-data-breach-leaks-details-of-500000-children-and-adults, 2017.

[Tho84] K. Thompson. Reflections on Trusting Trust. *Comm. ACM*, 27:761–763, 1984.

[Zor17] Z. Zorz. CloudPets connected toys can be turned into remote surveillance devices. https://www.helpnetsecurity.com/2017/03/01/cloudpets-security/, 2017.