Ethics, Public Policy, and Technological Change

Rob Reich
Mehran Sahami
Head TA: Roberta Fischli

Housekeeping

- Assignment #2 (Philosophy Paper) was released last week
  - Due February 12
  - "Writing Guidelines for Philosophy Paper" handout on the Assignments webpage of class website
  - WIM (CS182W) students only: sign up for a slot with a TCP writing tutor as soon as possible
    - Link to sign-ups is available on class website
    - WIM revised draft due March 1st at 12noon

Today's Agenda

1. Perspectives on data privacy
2. Approaches to data privacy
   - Anonymization
   - Encryption
   - Differential Privacy
3. What can we infer from your digital trails?

Perspectives on Privacy

- Data privacy often involves a balance of competing interests
  - Privacy vs. national security/personal safety/innovation/convenience
- Making data available for meaningful analysis
  - For public goods
    - Auditing algorithmic decision-making for fairness
    - Medical research and health care improvement
    - Protecting national security
  - For private goods
    - Personalized advertising
- Protecting individual privacy
  - Personal value of privacy and respect for individual (thanks Rob!)
  - Freedom of speech and activity
  - Avoiding discrimination
  - Regulation: FERPA, HIPAA, GDPR, etc.
  - Preventing access from “adversaries”

Today's Agenda

1. Perspectives on data privacy
2. **Approaches to data privacy**
   - **Anonymization**
   - Encryption
   - Differential Privacy
3. What can we infer from your digital trails?

Anonymization

- Basic idea: drop personally identifying features from the data

Employer	Job Title	Nationality
Stanford University	Professor	Iran
Google Inc.	Senior Engineer	USA

D.O.B	Zipcode	Has Condition?
May 10,	94306	No
June 6,	94040	Yes

- What if we also drop gender and birth year?
  - Just to be safe, since gender and age are protected characteristics
- How well does that work?

That Worked Well…

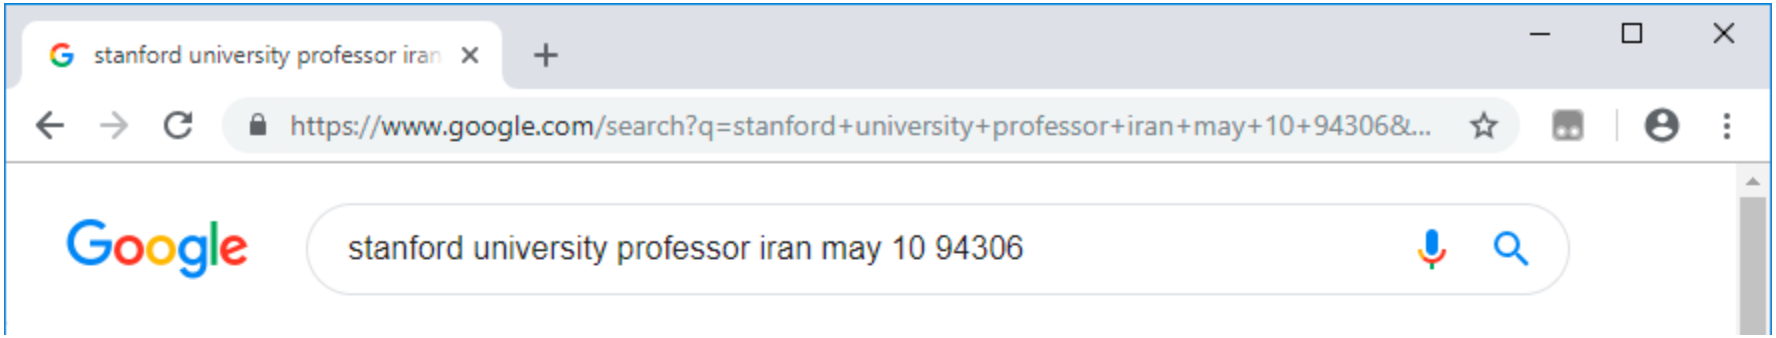

Massachusettes Hospital Visits

- Massachusetts Group Insurance Commission (GIC) released "anonymized" data of state employee hospital visits
  - 135,000 data records

- “William Weld, then Governor of Massachusetts, assured the public that GIC had protected patient privacy by deleting identifiers.” (Paul Ohm, 2015)

Image: Wikimedia Commons, CC BY-SA 2.0

Massachusettes Hospital Visits

- Latanya Sweeney, now Professor at Harvard, then graduate student at MIT, investigates…
  - Sweeney knew Weld lived in Cambridge, MA
  - For $20, she bought Cambridge voter list, containing: name, address (including zip code), birth date, and gender of 54,805 voters in city
  - She joined voter data with GIC data, de-anonymizing (re-identifying) Weld
  - Only six people in Cambridge shared Weld's birth date, only three of those six were men, and only Weld lived in his zip code
  - She sent Weld's health records to his office

Image: Wikimedia Commons, CC BY 4.0, cropped

De-Anonymization Can Be Easy

- Sweeney (2000): 87 percent of all Americans could be uniquely identified using: zip code, birthdate, and sex.
  - ~42,000 zip codes
  - 365 birthdates per year for ~80 years
  - Most individuals identify as one of two sexes
  - 42,000 * 365 * 80 * 2 ≈ 2.5B. Population of US ≈ 330M
- Try it yourself: Data Privacy Lab: https://aboutmyinfo.org/identity

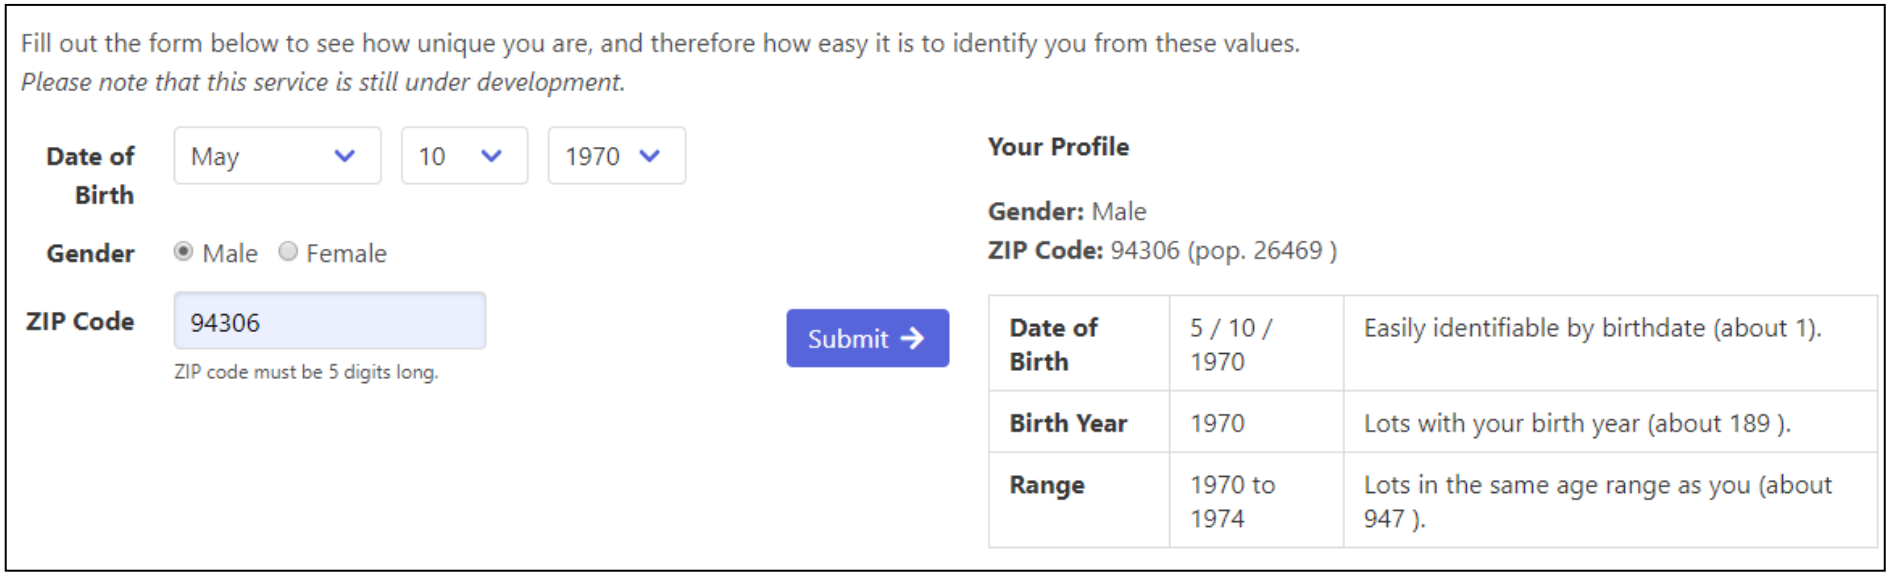

And People Can Get Fired For It

- America Online (AOL) releases an anonymized log of search queries for research purposes in 2006
  - 20M searchers
  - 650,000 users (tagged by “anonymous” ID)

- New York Times de-anonymizes several individuals in data
  - User #4417749: Thelma Arnold, a 62-year-old widow from Lilburn, GA
  - Queries included:
    “landscapers in Lilburn, Ga” (several people with last name *Arnold*)
    “homes sold in shadow lake subdivision gwinnett county georgia”
    “60 single men”
    “dog that urinates on everything”

- AOL Chief Technology Officer resigns, other researchers fired
  - AOL sued over privacy breach and agreed to pay $5M to resolve suit

Today's Agenda

1. Perspectives on data privacy
2. Approaches to data privacy
   - Anonymization
   - **Encryption**
   - Differential Privacy
3. What can we infer from your digital trails?

Encryption

- Encryption: process of encoding information so that only those with authorization can access it
- Why encrypt?
  - Keep data private/secure (e.g., iPhone)
  - Prevent eavesdropping (e.g., WhatsApp/Signal end-to-end encryption)
  - Guarantee data not altered in transport (e.g., database integrity)
- Data recipient must able to decrypt data for analytics
  - E.g., To share healthcare data for analytics, recipient must decrypt data
- Research direction: fully homomorphic cryptosystems
  - Allow arbitrary mathematical operations to be performed on *encrypted* data
  - Craig Gentry (Stanford CS PhD) showed first such system in 2009
  - Still too slow to be of practical use at scale

Symmetric Key Cryptography

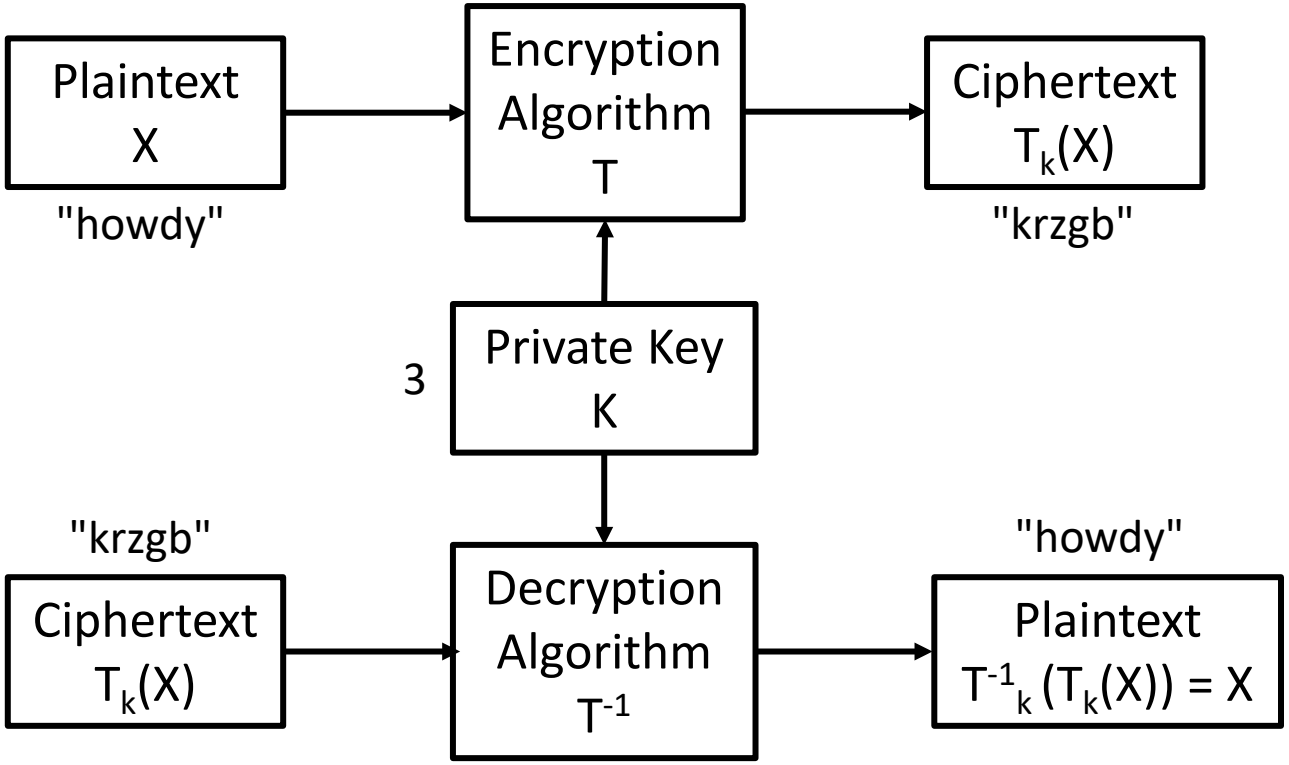

- Example: Caesar Cipher
  - Encrypt: shift letters *forward* K places in alphabet, wrapping from 'Z' to 'A'
  - Decrypt: shift letters *backward* K places, wrapping from 'A' back to 'Z'

Public Key Cryptography

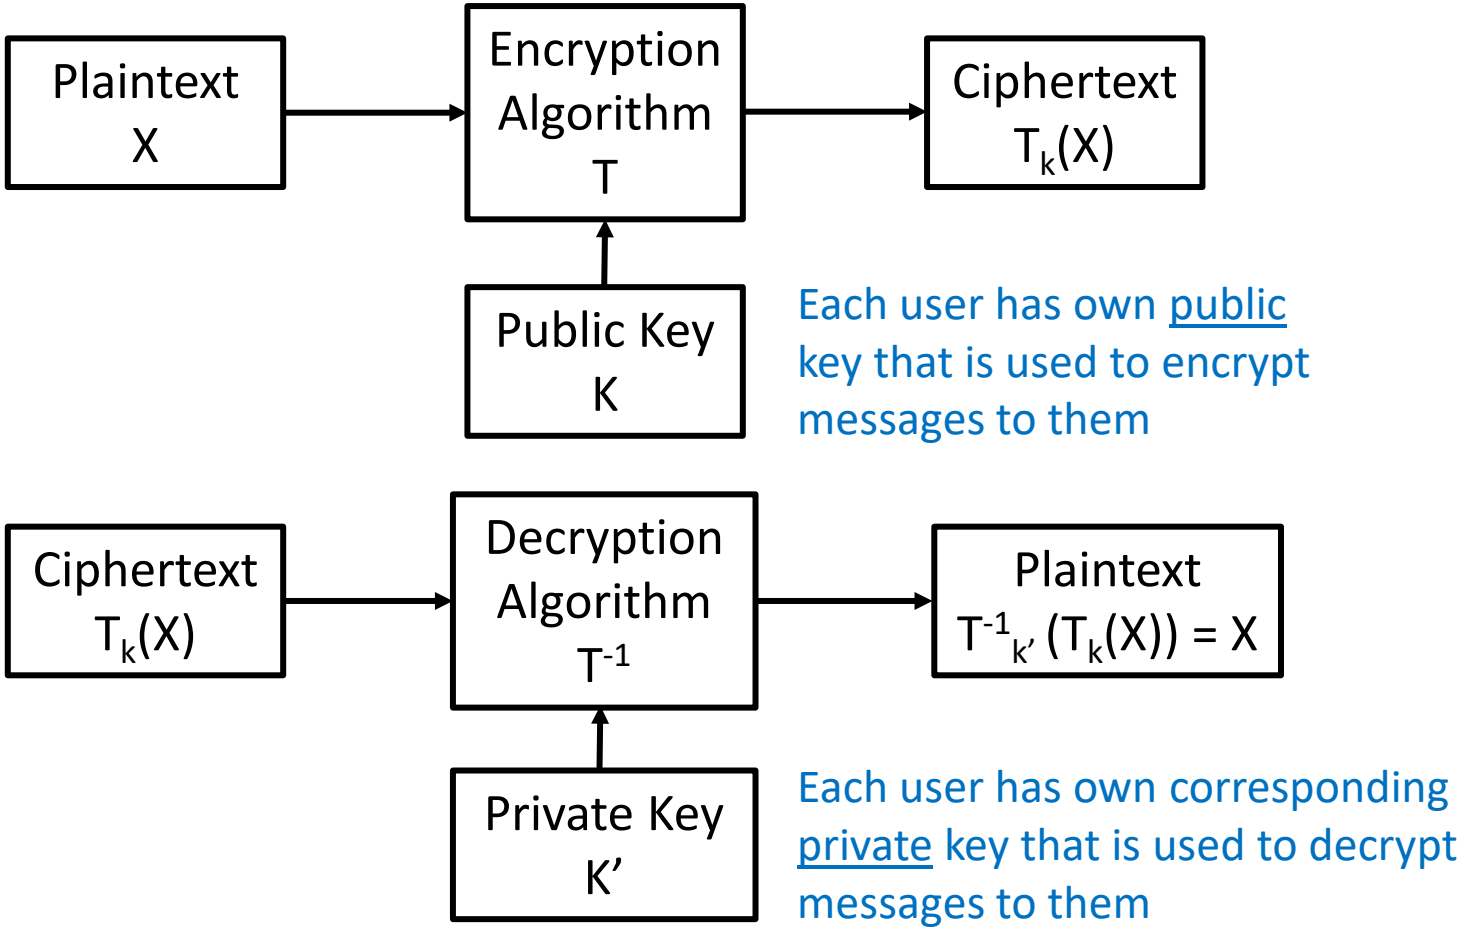

- Public key and private key somehow related to each other to make this possible

Public Key Cryptography

- 1976: Public key cryptography proposed by Diffie and Hellman
  - Initial proposal proved to be insecure
- 1978: Rivest, Shamir, and Adleman propose RSA algorithm
  - Secure (as far as we know) version of public key cryptography
- RSA cryptosystem
  - Public key (n) is based on product of two prime numbers p and q, $n = pq$
  - Private key (p and q) is based on knowing factorization of n into p and q
    - Factoring large numbers is hard, and our security relies on that!
  - Messages in this system represented as numbers
- Just to keep you up at night: practical quantum computing could break many common currently used cryptosystems (like RSA)

Public Key Cryptography

- Let’s just encrypt everything!
- Does that solve our privacy concerns?
- What happens if we want to apply analytics to data?

Today's Agenda

1. Perspectives on data privacy
2. Approaches to data privacy
   - Anonymization
   - Encryption
   - **Differential Privacy**
3. What can we infer from your digital trails?

Differential Privacy

- 2006: Dwork *et al* propose *differential privacy*
- Set-up: data curator wants to make access to database available for statistical analysis, but also wants to preserve privacy of subjects whose data is in database
  - E.g., census data, healthcare records, device usage, etc.
- Intuition: Consider two databases, one without your data and one with your data (i.e., one extra row of data). Result of a statistical query should be *almost indistinguishable* between two databases.
- Alternative intuition: a subject with data in the database is subject to *no more harm* from data analysis than if they were not in the database

Differential Privacy

- Formally: A randomized function *K* (operation on database) gives ε-differential privacy if for all data sets D and D' differing in at most one row (D' has extra row), and all S $\subseteq$ Range(*K*):

$$\frac{\Pr(K(\mathrm{D}) \in S)}{\Pr(K(\mathrm{D}') \in S)} \leq e^{\varepsilon}$$

  where the probability space is determined by the randomness used in *K* (i.e., *K* provides some randomness in result)

- Note, for small values of ε: $e^{\varepsilon} \approx 1 + \varepsilon$
  - Smaller values of ε lead to "more privacy" (i.e., difference from using either database is smaller), but less accuracy in data analysis

Randomized Response

- Consider getting data on a sensitive question: have you ever violated the Honor Code at Stanford?
  - Unlikely that students will tell the truth
- Instead, I ask you to secretly flip a fair coin
  - If coin is heads, then you answer truthfully
  - If coin is tails, then you flip coin again: heads → "yes", tails → "no"
  - Each individual has deniability for answering "yes"
- Can still estimate true percentage of students that have violated the Honor Code (call that p)
  - Let y = Pr(answer = "yes") = $(1/2)p + (1/2)(1/2) = (1/2)p + 1/4$
  - Solve for $p = 2y - (1/2)$
- Satisfies (local) differential privacy
  - Using a biased coin (different probability of "heads") impacts ε

Noise Injection

- When database is queried, add random noise to the answer
- Formally, if *f*(x) is function we wanted to compute from database, we return *f*(x) + C, where C is random value
  - To satisfy differential privacy, we can determine value C using a Laplace or Normal distribution

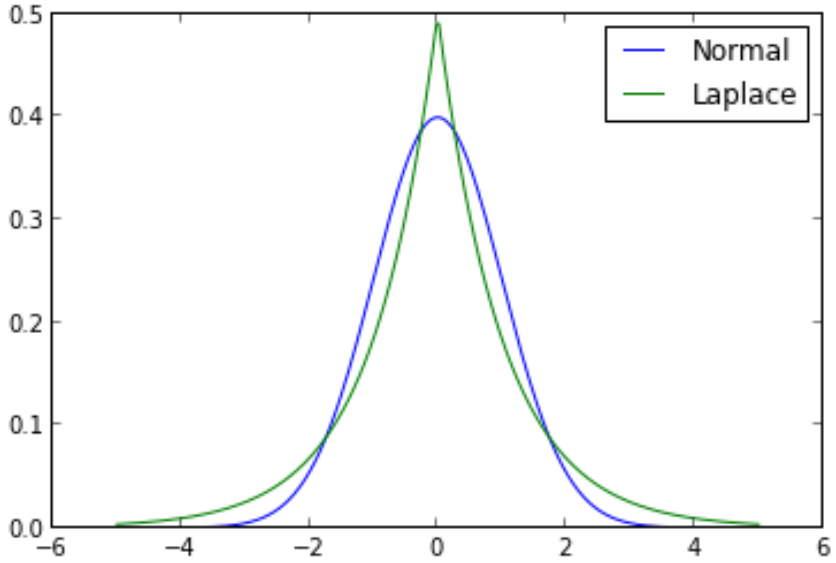

Differential Privacy in Industry

- Apple

Apple has adopted and further developed a technique known in the academic world as local differential privacy to do something really exciting: gain insight into what many Apple users are doing, while helping to preserve the privacy of individual users.

From https://www.apple.com/privacy/docs/Differential_Privacy_Overview.pdf

- Google

Randomized Aggregatable Privacy-Preserving Ordinal Response, or RAPPOR, is a technology for crowdsourcing statistics from end-user client software, anonymously, with strong privacy guarantees. This paper describes and motivates RAPPOR, details its differential-privacy and utility guarantees…

From https://research.google/pubs/rappor-randomized-aggregatable-privacy-preserving-ordinal-response/

Differential Privacy Concerns

- Let's apply differential privacy to all data gathering/access!
- Does that solve our privacy concerns?
- Where this implemented in the system matters:
  - At the client (when data is gathered)
  - In the database (when data is stored)
  - By the analyst (when data is queried for use)
- Decision has significant implications for privacy and future usage

Differential Privacy Concerns

- Querying repeatedly can undermine privacy guarantees
  - E.g., Statistical sampling techniques for mean values
- Randomness can cause small groups to be "washed out"
  - This is a concern for 2020 US census

A powerful new disclosure avoidance system (DAS) designed to withstand modern re-identification threats will protect 2020 Census data products….
The 2020 DAS is based on a framework for assessing privacy risk known as differential privacy. ***It is the only solution that can respond to this threat while maximizing the availability and utility of published census data.***

Source: https://www.census.gov/programs-surveys/decennial-census/decade/2020/planning-management/process/disclosure-avoidance.html

We study the impact of the U.S. Census Bureau's latest disclosure avoidance system (DAS) on a major application of census statistics, the redrawing of electoral districts. We find that the DAS ***systematically undercounts the population in mixed-race and mixed-partisan precincts, yielding unpredictable racial and partisan biases.***

Source: Christopher T. Kenny *et al*, The use of differential privacy for census data and its impact on redistricting: The case of the 2020 U.S. Census, Science Advances 7(41), 2021

Today's Agenda

1. Perspectives on data privacy
2. Approaches to data privacy
   - Anonymization
   - Encryption
   - Differential Privacy
3. **What can we infer from your digital trails?**

Aggregation and Inference

- Aggregation of public information
  - Each time you are outside, that's essentially public information
  - Aggregation: someone following you outdoors all the time

  - Purchasing something in store, people next to you can see what you buy
  - Aggregation: compilation of everything you have bought at all stores

  - So, at what point do we feel this becomes a privacy violation?

- In the on-line world, aggregation and inference is the standard
  - How well you can do it is the difference between Yahoo and Google
  - Yahoo sold to Verizon for $4.8 billion. Google market cap > $4 trillion.

- Let's play a game: Privacy Chicken!

https://www.nytimes.com/interactive/2020/01/21/opinion/privacy-chicken-game.html

As a Bloomberg journalist explained in 2006, ***"Google maximizes the revenue it gets from that precious real estate by giving its best position to the advertiser who is likely to pay Google the most in total****, based on the price per click multiplied by Google's estimate of the likelihood that someone will actually click on the ad."*

That pivotal multiplier was the result of Google's advanced computational capabilities trained on its most significant and secret discovery: behavioral surplus.

From this point forward, ***the combination of ever-increasing machine intelligence and ever-more-vast supplies of behavioral surplus would become the foundation of an unprecedented logic of accumulation****.*

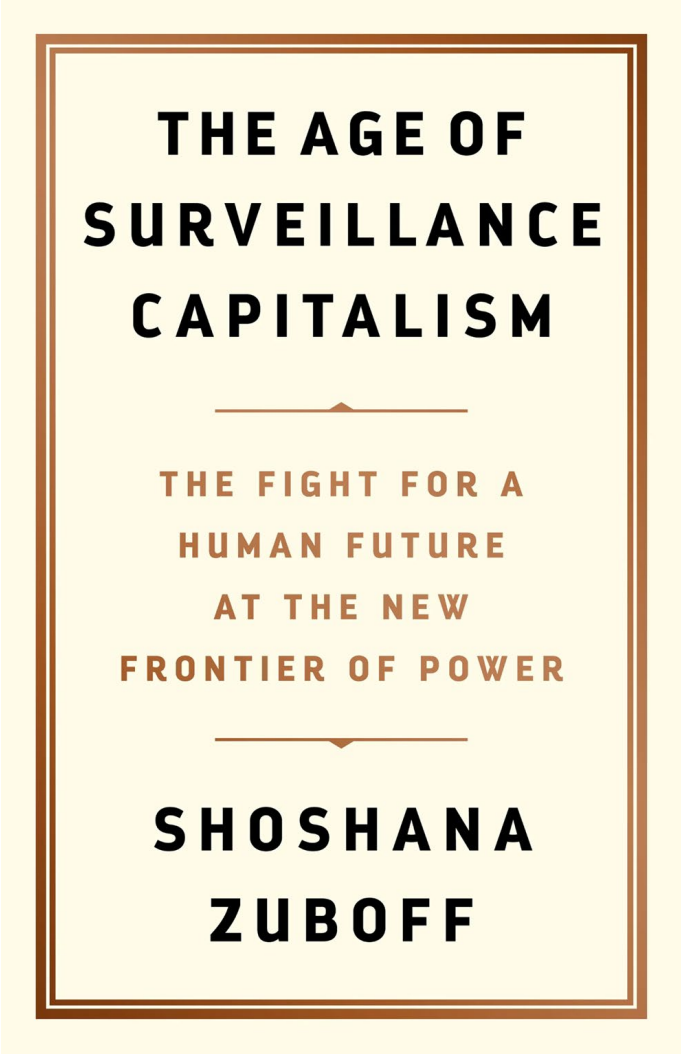

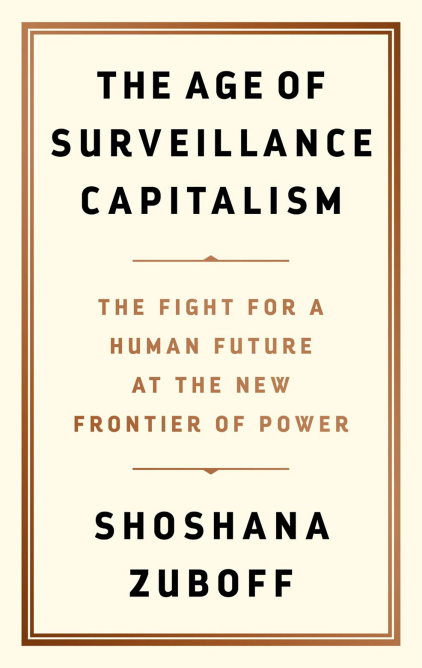

Google's many patents filed during those early years illustrate the explosion of discovery, inventiveness, and complexity detonated by the state of exception that led to these crucial innovations and the firm's determination to advance the capture of behavioral surplus.[43] Among these efforts, I focus here on one patent submitted in 2003 by three of the firm's top computer scientists and titled "Generating User Information for Use in Targeted Advertising."[44] The patent is emblematic of the new mutation and the emerging logic of accumulation that would define Google's success. Of even greater interest, it also provides an unusual glimpse into the "economic orientation" baked deep into the technology cake by reflecting the mindset of Google's distinguished scientists as they harnessed their knowledge to the firm's new aims.[45] In this way, the patent stands as a treatise on a new political economics of clicks and its moral universe, before the company learned to disguise this project in a fog of euphemism.

43. For example, consider this exemplary sample of Google patents filed during this general time frame: Krishna Bharat, Stephen Lawrence, and Mehran Sahami, Generating user information for use in targeted advertising, US9235849 B2, filed December 31, 2003, and issued January 12, 2016,

44. Three distinguished computer scientists, Krishna Bharat, Stephen Lawrence, and Meham Sahami, invented the technologies and techniques described in this patent (Generating user information for use in targeted advertising).

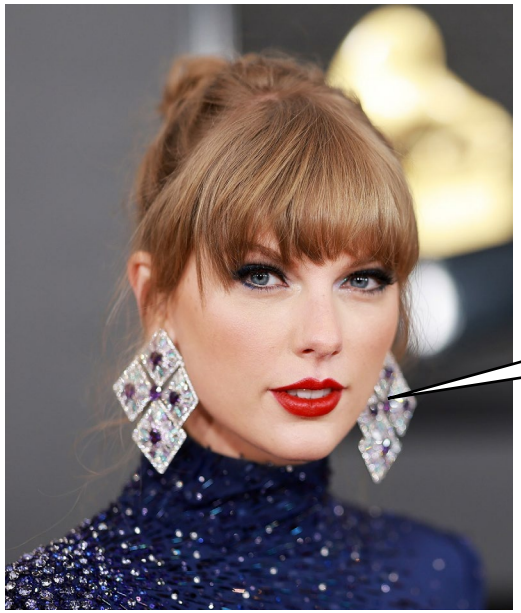

It's you. Hi. You're the problem, it's you!

I blame Meham!

This is the part where I tell you stories that won't show up in the slides and are another reason why you really should be coming to class.