

Privacy in a Mobile-Social World

CompSci 590.03

Instructor: Ashwin Machanavajjhala

Administrivia

<http://www.cs.duke.edu/courses/fall13/compsci590.3/>

- Wed/Fri 1:25 – 2:40 PM
- “Reading Course + Project”
 - No exams!
 - Every class based on 1 (or 2) assigned papers that students *must* read.
- Projects: (60% of grade)
 - Individual or groups of size 2
- Class Participation (other 40%)
 - May be one simple assignment ... 1 short (20 min) presentation
- Office hours: by appointment

Administrivia

- Projects: (60% of grade)
 - Theory/algorithms for privacy
 - Implement/adapt existing work to new domains
 - “**Break**” an existing privacy algorithm
- Goals:
 - Literature review
 - Some original research/implementation
- Timeline (details will be posted on the website soon)
 - **Sep 27**: Choose Project (ideas will be posted ... new ideas welcome)
 - **Oct 11**: Project proposal (1-4 pages describing the project)
 - **Nov 8**: Mid-project review (2-3 page report on progress)
 - **Dec 4**: Final presentations and submission (6-10 page conference style paper + 10-15 minute talk)

Why you should take this course?

1. Privacy is (one of) the most important grand challenges in managing today's data!
 1. *“What Next? A Half-Dozen Data Management Research Goals for Big Data and Cloud”, Surajit Chaudhuri, Microsoft Research*
 2. *“Big data: The next frontier for innovation, competition, and productivity”, McKinsey Global Institute Report, 2011*

Why you should take this course?

1. Privacy is (one of) the most important grand challenges in managing today's data!
2. Very active field and tons of interesting research.

We will read papers in:

- *Data Management (SIGMOD, VLDB, ICDE)*
- *Theory (STOC, FOCS)*
- *Cryptography/Security (TCC, SSP, NDSS)*
- *Machine Learning (KDD, NIPS)*
- *Statistics (JASA)*

Why you should take this course?

1. Privacy is (one of) the most important grand challenges in managing today's data!
2. Very active field and tons of interesting research.
3. Intro to research by working on a cool project
 - *Read scientific papers about an exciting data application*
 - *Formulate a problem*
 - *Perform a scientific evaluation*

Today

- Bird's-eye view introduction to big-data and privacy
 - Privacy attacks in the real-world
 - (In)formal problem statement
 - Course overview
-
- (If there is time) A privacy preserving algorithm

INTRODUCTION

Lecture 1 : 590.03 Fall 13

Data Explosion: Internet

Data Points

Share and Share Alike

The amount of data shared online in one week tops what the Hubble Space Telescope collected in its first 20 years. Here's a snapshot of our habits.

Estimated User Data Generated per day [Ramakrishnan 2007]

- 8-10 GB public content
- ~4 TB private content


Data Explosion: Social Networks

- 91% of online users ...
- 25% of all time spent online ...
- 200 million tweets a day ...
- millions of posts a day ...
- 6 billion photos a month ...



Data Explosion: Mobile

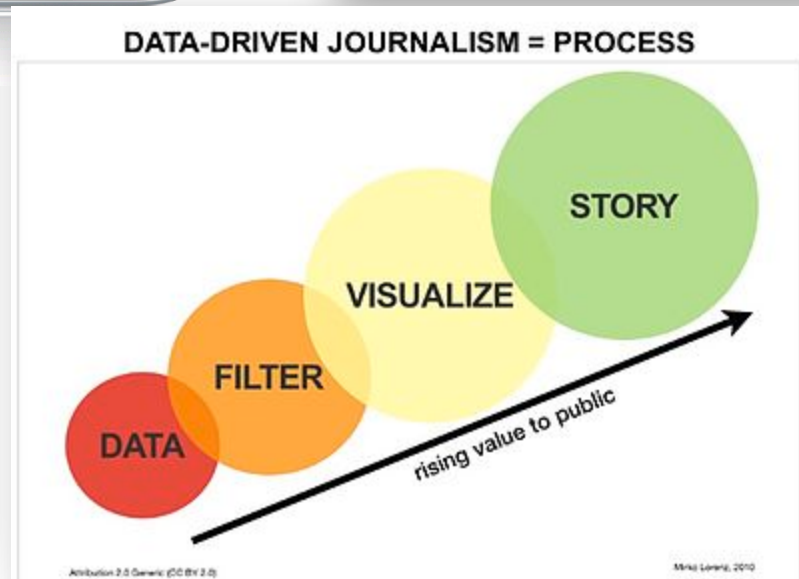

- ~5 billion mobile phones in use!



Source: Cisco VNI Mobile, 2011

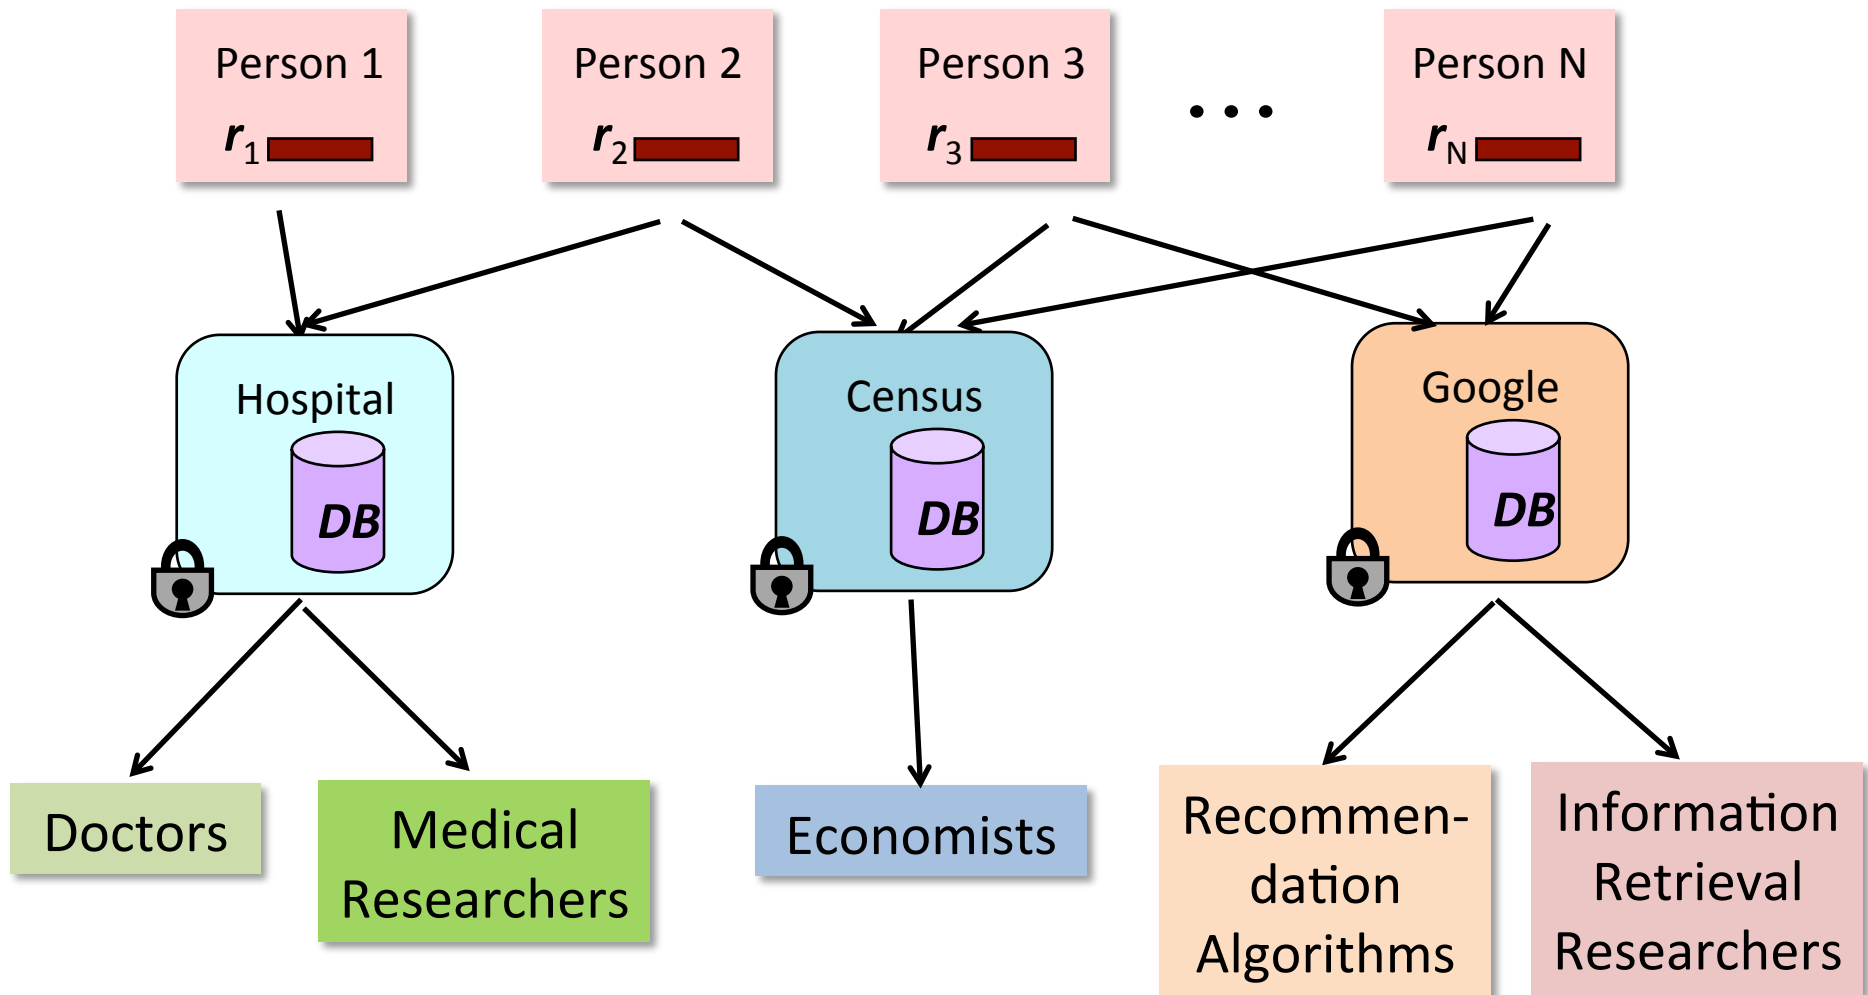
Big-Data impacts all aspects of our life

The collage shows a Yahoo! search page with a search bar and various search results. Below it is a Facebook page for 'Celebs on Facebook' with a public figure profile and a post about Lady Gaga's GRAMMY awards. To the right is a snippet of an IMDb page for 'Jurassic Park'.



Lecture 1 : 590.03 Fall 13

Personal Big-Data



Sometimes users can control and know who sees their information ...

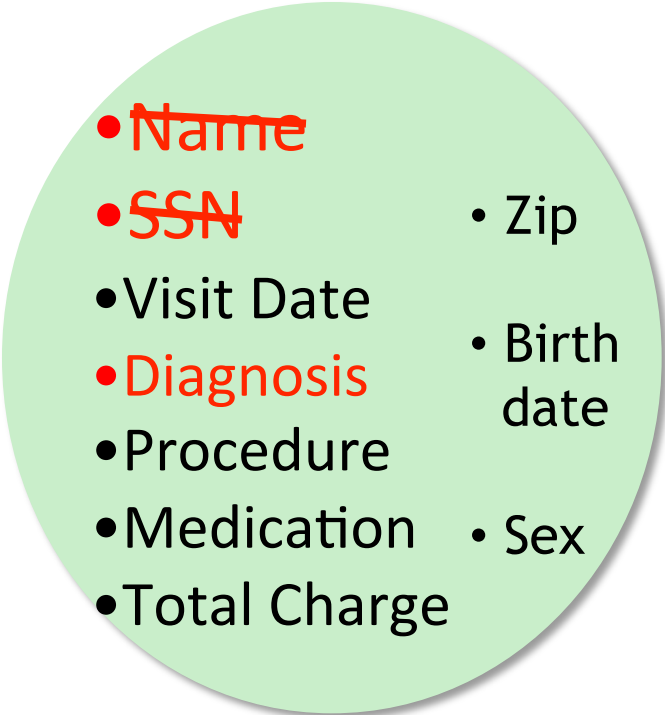
Who Can View My Full Profile	
<input type="radio"/>	My Friends Only
<input checked="" type="radio"/>	Public
<input type="radio"/>	Only Users <i>Over 18</i>

Privacy Settings	
<input type="checkbox"/>	Friend Requests - Require email or last name
<input type="checkbox"/>	Comments - approve before posting
<input type="checkbox"/>	Hide Online Now
<input type="checkbox"/>	Show My Birthday to my Friends 🎂
<input type="checkbox"/>	Photos - No Forwarding
<input type="checkbox"/>	Blog Comments - Friends Only
<input type="checkbox"/>	Friend Requests - No Bands
<input type="checkbox"/>	Block Users Under 18 From Contacting Me

... but not always !!

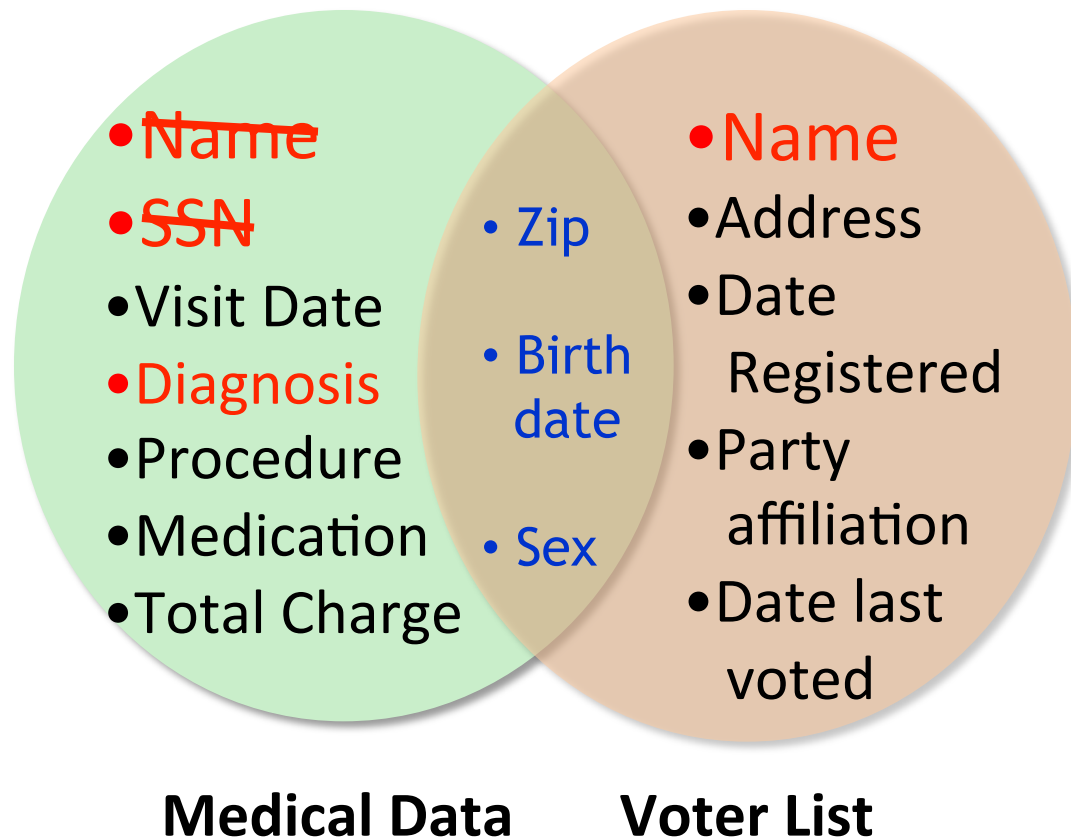
A graphic featuring a cartoon burglar with a beard, wearing a tan cap and a black mask, carrying a green sack with a dollar sign. To the right of the burglar, the text 'PLEASE ROB ME' is written in large, bold, red, sans-serif capital letters. Below this text is a blue banner with the text 'Raising awareness about over-sharing' in white. Underneath the banner, it says 'Check out our [guest blog post](#) on the CDT website.' To the right of the banner is a map graphic with two red location pins, each containing a white 'X'. Below the banner are social media interaction buttons: 'Like', 'Send', and '29,523 people like this.' At the bottom left of the graphic, it says 'Check your own Twitter timeline for checkins'. On the right side, there is a 'More Info' section with a 'Home' link and a partially visible 'Why' link.

The Massachusetts Governor Privacy Breach [Sweeney IJUFKS 2002]

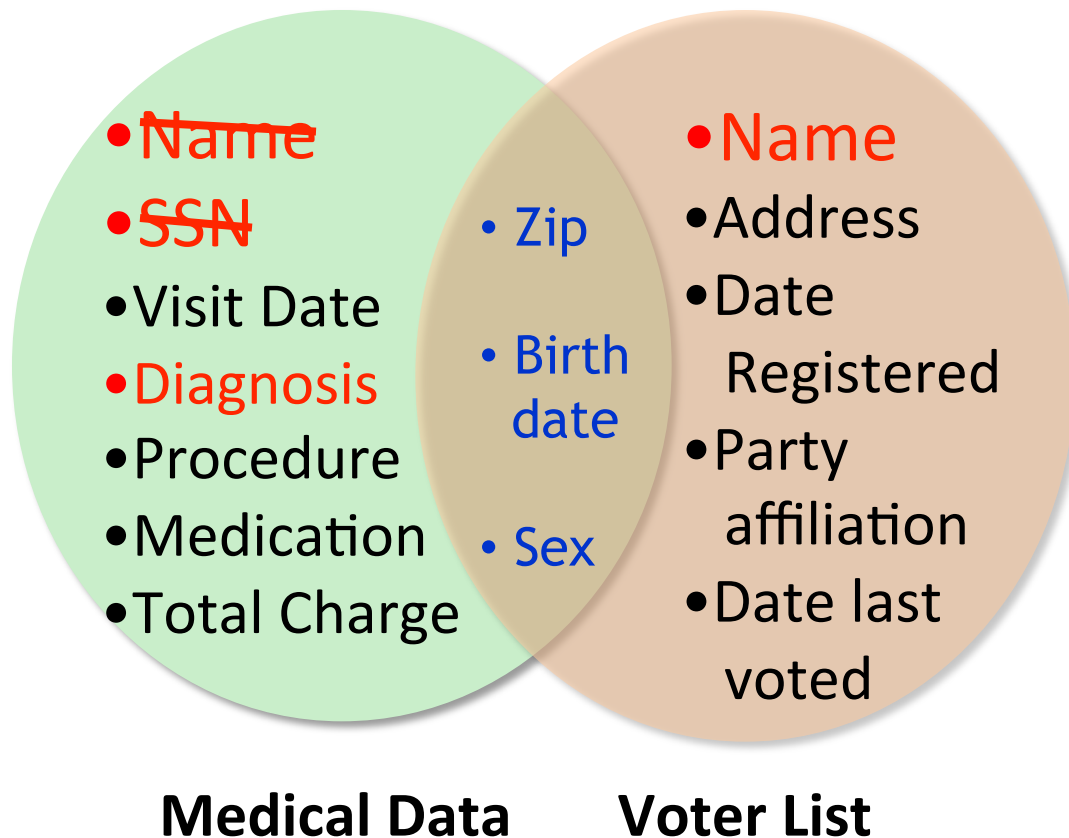
- 
- ~~Name~~
 - ~~SSN~~
 - Visit Date
 - ~~Diagnosis~~
 - Procedure
 - Medication
 - Total Charge
 - Zip
 - Birth date
 - Sex

Medical Data

The Massachusetts Governor Privacy Breach [Sweeney IJUFKS 2002]



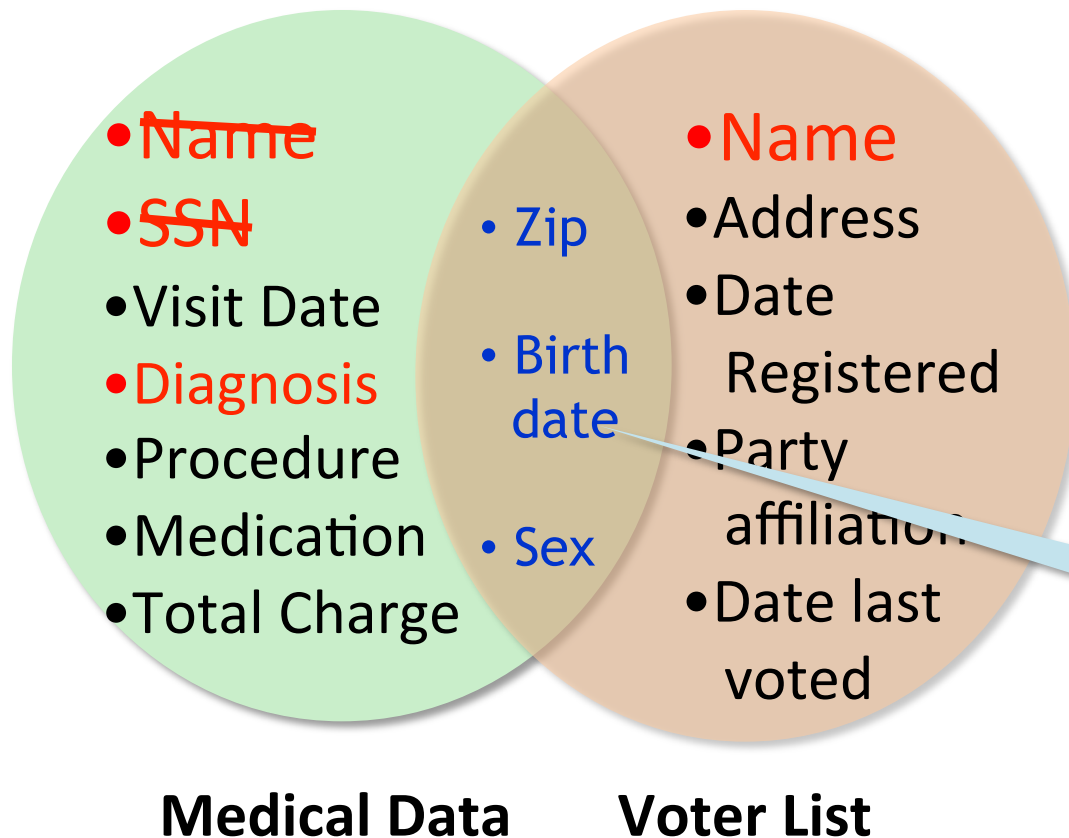
The Massachusetts Governor Privacy Breach [Sweeney IJUFKS 2002]



- Governor of MA uniquely identified using ZipCode, Birth Date, and Sex.

Name linked to Diagnosis

The Massachusetts Governor Privacy Breach [Sweeney IJUFKS 2002]



- 87 % of US population **uniquely identified** using ZipCode, Birth Date, and Sex.

Quasi Identifier

AOL data publishing fiasco ...

*“... Last week AOL did another stupid thing ...
... but, at least it was in the name of science...”*

Altnet, August 2006

AOL data publishing fiasco ...

AOL “anonymously” released a list of 21 million web search queries.

Ashwin222	Uefa cup
Ashwin222	Uefa champions league
Ashwin222	Champions league final
Ashwin222	Champions league final 2007
Pankaj156	exchangeability
Pankaj156	Proof of deFinitti’s theorem
Cox12345	Zombie games
Cox12345	Warcraft
Cox12345	Beatles anthology
Cox12345	Ubuntu breeze
Ashwin222	Grammy 2008 nominees
Ashwin222	Amy Winehouse rehab

AOL data publishing fiasco ...

AOL “anonymously” released a list of 21 million web search queries.

UserIDs were replaced by random numbers ...


865712345	Uefa cup
865712345	Uefa champions league
865712345	Champions league final
865712345	Champions league final 2007
236712909	exchangeability
236712909	Proof of deFinitti’s theorem
112765410	Zombie games
112765410	Warcraft
112765410	Beatles anthology
112765410	Ubuntu breeze
865712345	Grammy 2008 nominees
865712345	Amy Winehouse rehab


Privacy Breach

[NYTimes 2006]

A Face Is Exposed for AOL Searcher No. 4417749

By MICHAEL BARBARO and TOM ZELLER Jr.
Published: August 9, 2006


 SIGN IN TO E-
THIS



Privacy breaches on the rise...



Why 'Anonymous' Data Sometimes Isn't


By Bruce Schneier  12.13.07

Last year, Netflix published 10 million movie rankings by 500,000 customers, as part of a challenge for people to come up with better recommendation systems than the one the company was using.

The New York Times Business Day
Technolo


WORLD | U.S. | N.Y. / REGION | BUSINESS | TECHNOLOGY | SCIENCE | HE

Marketers Can Glean Private Data on Facebook



Facebook Ads

Reach the exact audience you want with relevant targeted ads.



TECH | 2/16/2012 @ 11:02AM | 837,678 views

How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did

Privacy Breach: Informal Definition

A data sharing mechanism **M**

that allows

an unauthorized party



to learn sensitive information about any individual,

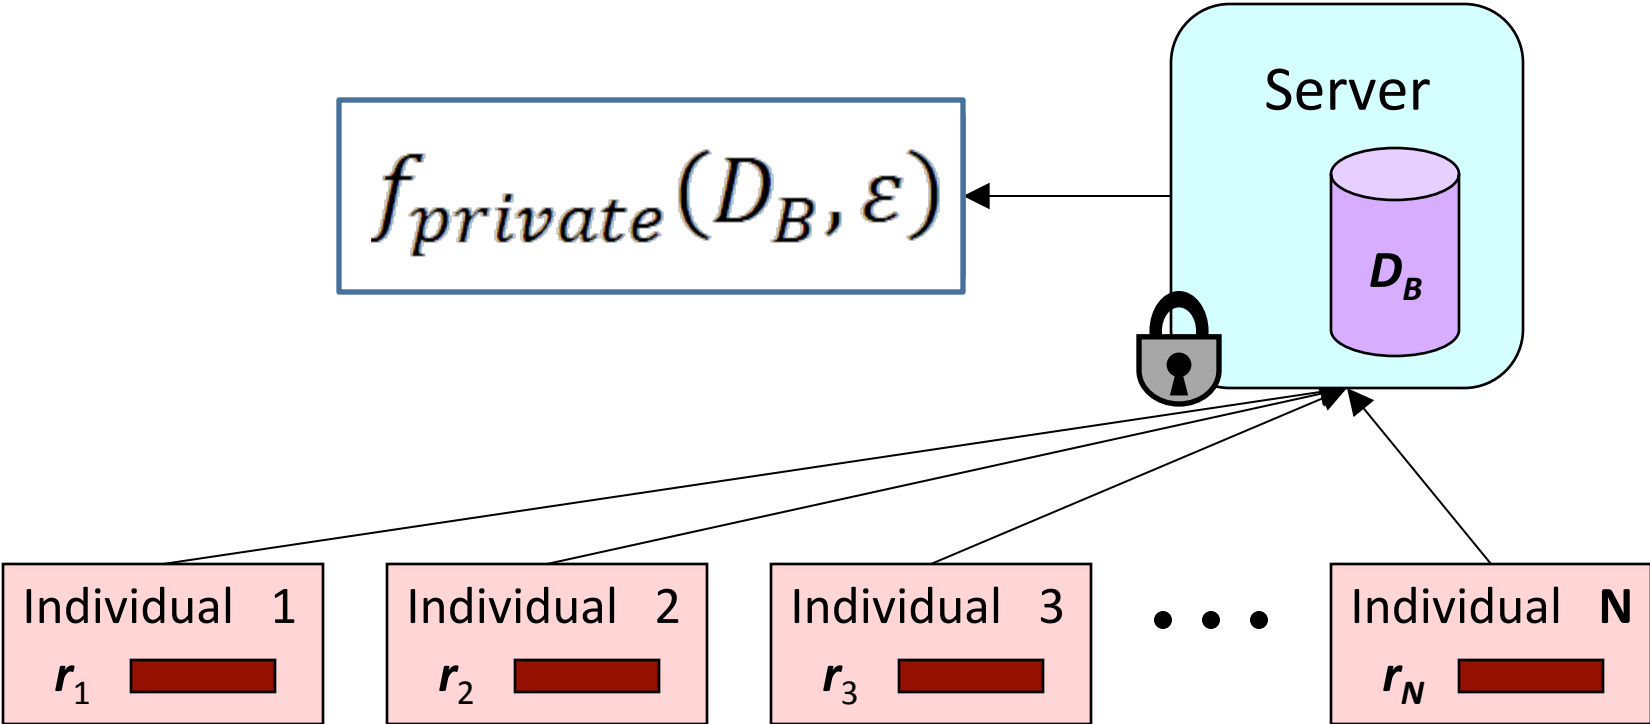
which



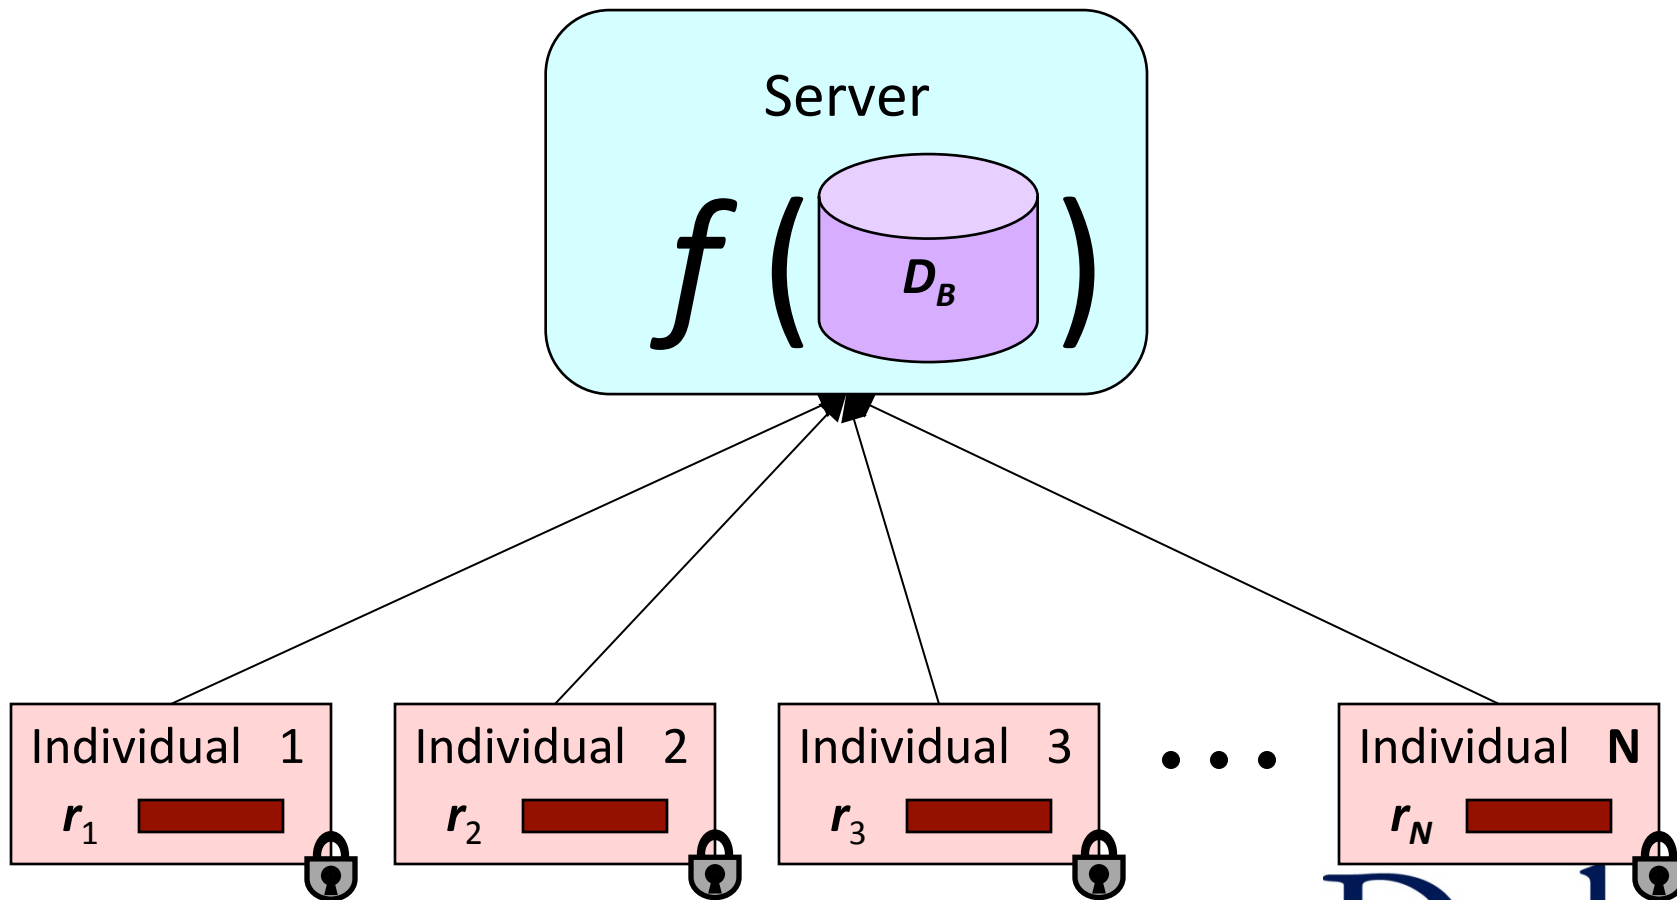
could not have learnt without access to **M**.

Statistical Database Privacy (Trusted Collector)

Utility: $f_{private}$ approximates f
Privacy: No breach about any individual



Statistical Database Privacy (Untrusted Collector)



Statistical Databases in real-world applications

- Trusted Data Collectors

Application	Data Collector	Third Party (adversary)	Private Information	Function (utility)
Medical	Hospital	Epidemiologist	Disease	Correlation between disease and geography
Genome analysis	Hospital	Statistician/ Researcher	Genome	Correlation between genome and disease
Advertising	Google/FB/Y!	Advertiser	Clicks/ Browsing	Number of clicks on an ad by age/region/gender ...
Social Recommendations	Facebook	Another user	Friend links / profile	Recommend other users or ads to users based on social network

Statistical Databases in real-world applications

- Untrusted Data Collectors

Application	Data Collector	Private Information	Function (utility)
Location Services	Verizon/AT&T	Location	Local Search
Recommendations	Amazon/Google	Purchase history	Product Recommendations
Traffic Shaping	Internet Service Provider	Browsing history	Traffic pattern of groups of users

Statistical Database Privacy is not ...

- Encryption:

Statistical Database Privacy is not ...

- Encryption:
Alice sends a message to Bob such that Trudy (attacker) does not learn the message. Bob should get the correct message ...
- Statistical Databases:
A set of individuals want Bob (attacker) to learn aggregate properties about the data, but not properties about individuals.

Statistical Database Privacy is not ...

- Computation on Encrypted Data:

Statistical Database Privacy is not ...

- Computation on Encrypted Data:
 - Alice stores encrypted data on a (malicious) server.
 - Server performs computation on the encrypted data and returns encrypted answer to Alice.

- Statistical Databases:
 - Alice wants the server to learn aggregate properties from the database.

Statistical Database Privacy is not ...

- The Millionaires Problem:

Statistical Database Privacy is not ...

- Secure Multiparty Computation:
 - A set of agents each having a private input x_i ...
 - ... Want to compute a function $f(x_1, x_2, \dots, x_k)$
 - Each agent must learn no other information than what can be inferred from their private input and the answer.

- Statistical Database:
 - Should not learn any private input
 - (... function output can disclose some inputs ...)

Statistical Database Privacy is not ...

- Access Control:

Statistical Database Privacy is not ...

- Access Control:
 - A set of agents want to access a set of resources (could be files or records in a database)
 - Access control rules specify who is allowed to access (*or not access*) certain resources.
 - 'Not access' usually means no information must be disclosed
- Statistical Database:
 - A single database and a single agent
 - Want to release aggregate statistics about a set of records without allowing access to individual records

Privacy Problems

- In today's cloud context a number of privacy problems arise:
 - Encryption when communicating data across a unsecure channel
 - Secure Multiparty Computation when different parties want to compute on a function on their private data without using a centralized third party
 - Computing on encrypted data when one wants to use an unsecure cloud for computation
 - Access control when different users own different parts of the data
- Statistical Database Privacy:
Quantifying (and bounding) the amount of information disclosed about individual records by the output of a valid computation.

Statistical Database Privacy: Key Problems

What is a right definition of privacy?

How to develop mechanisms that
trade-off privacy for utility?

What is Privacy?

- *“... the ability to determine for ourselves when, how, and to what extent information about us is communicated to others ...”*

Westin, 1967

- *Privacy intrusion occurs when new information about an individual is released.*

Parent, 1983

Anonymity

- The property that an individual's record is indistinguishable from *many* other individual's records.
- K-Anonymity : popular definition where *many* = $k-1$
- Used for
 - *Social network anonymization*
 - *Location privacy*
 - Anonymous routing

Privacy is not Anonymity

- Bob's record is indistinguishable from records of other Cancer patients
 - We can infer Bob has Cancer !
- “New Information” principle
 - Privacy is breached if releasing D (or $f(D)$) allows an adversary to learn sufficient new information.
 - *New Information = distance(adversary's prior belief, adversary's posterior belief after seeing D)*
 - *New Information* can't be 0 if the output D or $f(D)$ should be useful.

Privacy Definitions

- Many privacy definitions
 - L-diversity, T-closeness, M-invariance, ϵ - **Differential privacy**, **Pufferfish**, ...
- Definitions differs in
 - What information is considered sensitive
 - Specific attribute (disease) vs all possible properties of an individual
 - What is the adversary's prior
 - All values are equally likely vs Adversary knows everything about all but one individuals
 - How is new information measured
 - Information theoretic measures
 - Pointwise absolute distance
 - Pointwise relative distance

No Free Lunch

- Why can't we have a single definition for privacy?
 - For every adversarial prior and every property about an individual, new information is bounded by some constant.
- No Free Lunch Theorem: For every algorithm that outputs a D with even a sliver of utility, there is some adversary with a prior such that privacy is not guaranteed.

Algorithms for Privacy

- Basic Building Blocks
 - Generalization or coarsening of attributes
 - Suppression of outliers
 - Perturbation
 - Adding noise
 - Sampling

Algorithms for Privacy

- Build complex algorithms by piecing together building blocks.
- But, each building block leads to some information disclosure. And, information disclosure may not add up linearly.
 - If A1 releases the fact that Bob's salary is $\leq 50,000$, while A2 releases the fact that Bob's salary is $\geq 50,000$; then we know Bob's salary is exactly 50,000.
 - **Composition of Privacy**
- Algorithms may be reverse-engineered.
 - If algorithm perturbs x by adding 1, then x can be reconstructed.
 - **Simulatability of Algorithms**

Algorithms for Privacy

- Anonymous/Private Data Publishing
 - Medical/Census Data, Search Logs, Social Networks, Location GPS traces
- Answering Statistical Counting Queries
 - Number of students enrolled in this class categorized by gender, nationality
 - Data Cubes (database), Marginals (statistics)
- Social Network Analysis
 - Measures of centrality (what is the degree distribution? How many triangles?)
- Streaming Algorithms
 - Continuously monitor number of cars crossing a toll booth.
 - Location Privacy, Health ...

Algorithms for Privacy

- Game Theory
 - Can I participate in an auction without the output of the auction revealing my private utility function?
 - Modern advertising is based on auction design.
 - Auctions and Mechanism Design
- Machine Learning
 - Regress disease and gender/location/age
 - Inside tip: Big open area. Much theory – doesn't work in practice
- Recommendations
 - Think netflix, amazon ...
- Advertising

Course Outline

<http://www.cs.duke.edu/courses/fall13/compsci590.3/>

Theory/Algorithms (Lectures 1-18)

Applications (Lectures 19-26)

Project Presentations (Lecture 27)

RANDOMIZED RESPONSE

Lecture 1 : 590.03 Fall 13

Case Study: Census Data Collection

- N respondents asked a sensitive “yes/no” question.
- Surveyor wants to compute fraction π who answer “yes”.
- Respondents don’t trust the surveyor.
- What should the respondents do?

Randomized Response

- Flip a coin
 - heads with probability p , and
 - tails with probability $1-p$ ($p > \frac{1}{2}$)
- Answer question according to the following table:

	True Answer = Yes	True Answer = No
Heads	Yes	No
Tails	No	Yes

Utility Analysis

- π : True fraction of respondents answering “yes”
- p : Probability coin falls heads
- $Y_i = 1$, if the i^{th} respondent says “yes”
= 0, if the i^{th} respondent says “no”

	Yes	No
Heads	Yes	No
Tails	No	Yes

$P(Y_i = 1) = (\text{True answer} = \text{yes AND coin} = \text{heads}) \text{ OR}$
 $(\text{True answer} = \text{no AND coin} = \text{tails})$

$$= \pi p + (1-\pi)(1-p) = p_{\text{yes}}$$

$$P(Y_i = 0) = \pi(1-p) + (1-\pi)p = p_{\text{no}}$$

Utility Analysis

- Suppose n_1 out of N people replied “yes”, and rest said “no”
- What is the best estimate for π ?
- Likelihood: $L = {}^n C_{n_1} p_{\text{yes}}^{n_1} p_{\text{no}}^{(n-n_1)}$
- Most likely value of π : (by setting $dL/d\pi = 0$)

$$\pi_{\text{hat}} = \{n_1/n - (1-p)\}/(2p-1)$$

Privacy

- Adversary's prior belief: $P(\text{Bob's true answer is "yes"}) = \theta$

- Suppose Bob answers "yes".

$P(\text{Bob's true answer is "yes" | Bob says "yes"})$

$= P(\text{Bob says "yes" AND Bob's true answer is "yes"}) / P(\text{Bob says yes})$

$= \frac{P(\text{Bob says "yes" | Bob's true answer is "yes"})P(\text{Bob's true answer is "yes"})}{P(\text{Bob says "yes" | Bob's true answer is "yes"})P(\text{Bob's true answer is "yes"}) + P(\text{Bob says "yes" | Bob's true answer is "no"})P(\text{Bob's true answer is "no"})}$

$= p\theta / p\theta + (1-p)(1-\theta) \leq \mathbf{p/(1-p) \theta}$

Privacy

- Adversary's prior belief:
 $P(\text{Bob's true answer is "yes"}) = \theta$
- Suppose Bob answers "yes".
Adversary's posterior belief:
 $P(\text{Bob's true answer is "yes"} \mid \text{Bob says "yes"}) \leq p/(1-p) \theta$

Adversary's posterior belief is always bounded by $p/1-p$ times the adversary's prior belief (irrespective of what the prior is)

Privacy vs Utility tradeoff

- When $p = 1$ (return truthful answer)
 - $p/1-p = \text{infinity}$: no privacy
 - $\pi_{\text{hat}} = n1/n = \text{true answer}$
- When $p = \frac{1}{2}$ (return random answer)
 - $p/1-p = 1$: perfect privacy
 - We cannot estimate π_{hat} since the answers are independent of the input.
 - $P_{\text{yes}} = \pi p + (1-\pi)(1-p) = \frac{1}{2}(\pi + 1 - \pi) = \frac{1}{2} = P_{\text{no}}$

Next Class

- Attacks on naively anonymized data
 - Netflix recommendations
 - Social networks