Privacy Attacks Practicum

Privacy & Fairness in Data Science

CS848 Fall 2019

UNIVERSITY OF WATERLOO | DSg Data Systems Group

Module 1: Intro to Privacy

1. Privacy Attacks Practicum
2. Differential Privacy
3. Basic Algorithms
4. Designing Complex Algorithms & Composition

Outline

- Recap Privacy Attacks
- Privacy Attack Exercises
- Desiderata of Privacy

The Massachusetts Governor Privacy Breach [Sweeney IJUFKS 2002]

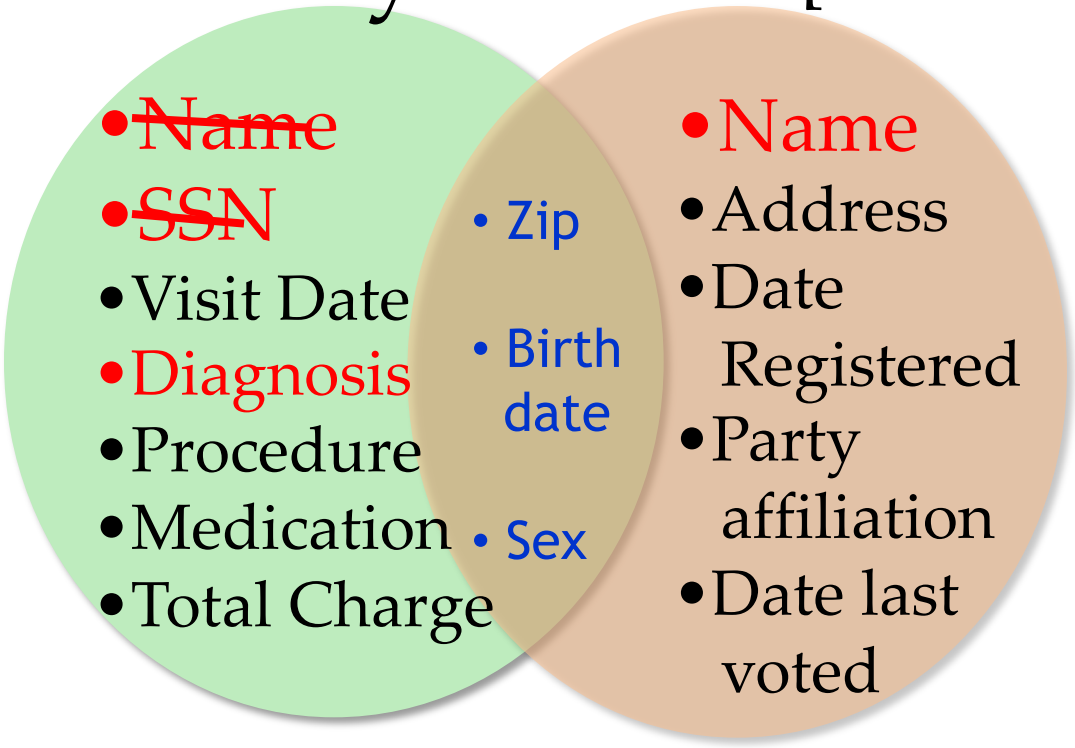

- Governor of MA **uniquely identified** using ZipCode, Birth Date, and Sex.

Name linked to Diagnosis

The Massachusetts Governor Privacy Breach [Sweeney IJUFKS 2002]

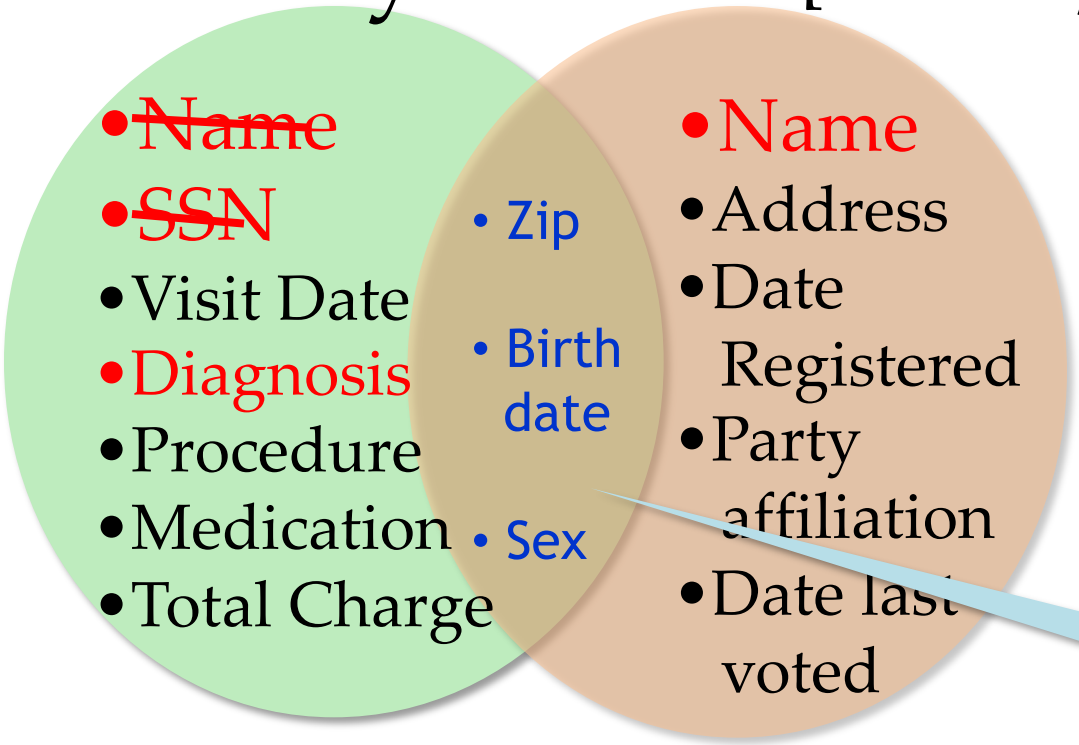

- 87 % of US population **uniquely identified** using ZipCode, Birth Date, and Sex.

Quasi Identifier

AOL data publishing fiasco

User IDs replaced with random numbers

865712345	Uefa cup
865712345	Uefa champions league
865712345	Champions league final
865712345	Champions league final 2013
236712909	exchangeability
236712909	Proof of deFinitti's theorem
112765410	Zombie games
112765410	Warcraft
112765410	Beatles anthology
112765410	Ubuntu breeze
865712345	Python in thought
865712345	Enthought Canopy

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

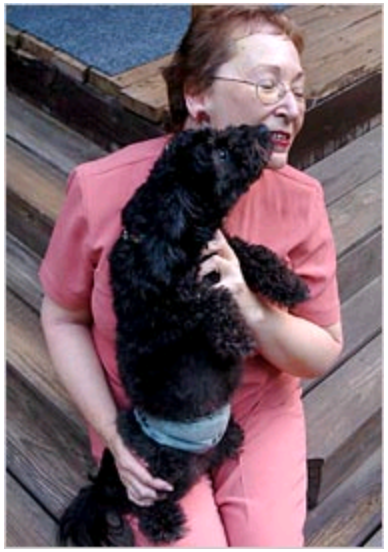

Your Turn!

- Divide into groups of 3

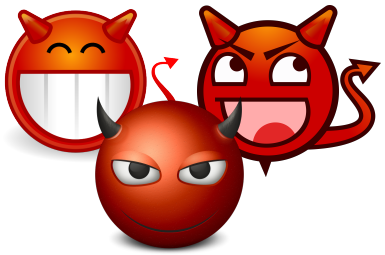

- Attack 4 problems as a group (15 mins)

Problem 1

- Social networks: graphs where each node represents a social entity, and each edge represents certain relationship between two entities

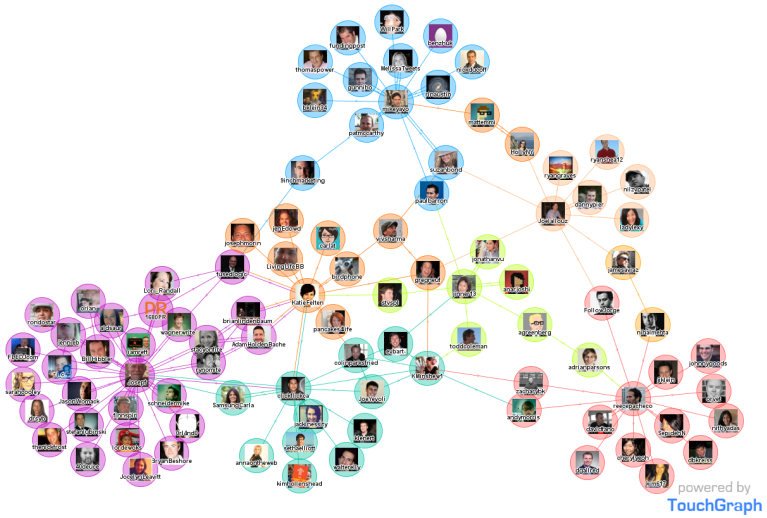

- Example: email communication graphs, social interactions like in Facebook, Yahoo! Messenger, etc.

Problem 1

- Anonymized email communication graph

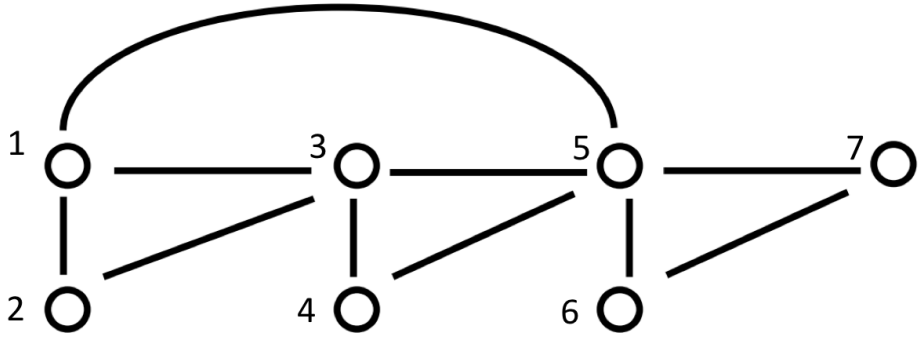

- Unfortunately for the email service providers, investigative journalists Alice and Cathy are part of this graph. What can they deduce?

Problem 2

- The email service provider also released perturbed records as per a linear function, but with *secret* parameters.

Node ID	Age (perturbed)
1	40
2	34
3	52
4	28
5	48
6	22
7	92

- What can Alice and Cathy deduce now?

Problem 3

- Releasing tables that achieve k-anonymity
  - At least k records share the same quasi-identifier
  - E.g. 4-anonymous table by generalization

	Non-Sensitive			Sensitive
	Zip code	**Age**	**Nationality**	**Condition**
1	130**	<30	*	AIDS
2	130**	<30	*	Heart Disease
3	130**	<30	*	Viral Infection
4	130**	<30	*	Viral Infection
5	130**	≥40	*	Cancer
6	130**	≥40	*	Heart Disease
7	130**	≥40	*	Viral Infection
8	130**	≥40	*	Viral Infection
9	130**	3*	*	Cancer
10	130**	3*	*	Cancer
11	130**	3*	*	Cancer
12	130**	3*	*	Cancer

(a)

Problem 3

- 2 tables of k-anonymous patient records

	Non-Sensitive			Sensitive
	Zip code	Age	Nationality	Condition
1	130**	<30	*	AIDS
2	130**	<30	*	Heart Disease
3	130**	<30	*	Viral Infection
4	130**	<30	*	Viral Infection
5	130**	≥40	*	Cancer
6	130**	≥40	*	Heart Disease
7	130**	≥40	*	Viral Infection
8	130**	≥40	*	Viral Infection
9	130**	3*	*	Cancer
10	130**	3*	*	Cancer
11	130**	3*	*	Cancer
12	130**	3*	*	Cancer

Hospital A (4-anonymous)

	Non-Sensitive			Sensitive
	Zip code	Age	Nationality	Condition
1	130**	<35	*	AIDS
2	130**	<35	*	Tuberculosis
3	130**	<35	*	Flu
4	130**	<35	*	Tuberculosis
5	130**	<35	*	Cancer
6	130**	<35	*	Cancer
7	130**	≥35	*	Cancer
8	130**	≥35	*	Cancer
9	130**	≥35	*	Cancer
10	130**	≥35	*	Tuberculosis
11	130**	≥35	*	Viral Infection
12	130**	≥35	*	Viral Infection

Hospital B (6-anonymous)

- If Alice visited both hospitals, can you deduce Alice's medical condition?

Problem 4


U.S. Department of Health & Human Services
About Us
Careers
Contact Us
Español
FAQ
Email Updates
AHRQ
Agency for Healthcare Research and Quality
Advancing Excellence in Health Care
HCUPnet
Healthcare Cost and Utilization Project
Home
Glossary
Methodology
Our Partners
Tutorial
Free Health Care Statistics
HCUPnet is a free, on-line query system based on data from the Healthcare Cost and Utilization Project (HCUP)
The system provides health care statistics and information for hospital inpatient, emergency department, and ambulatory settings, as well as population-based health care data on counties
Create a New Analysis
Get Quick Statistics Tables
Find out more about HCUP
What's new with HCUPnet
The HCUPnet Web site has been redesigned. The new site has a modernized look and feel, a simplified process for querying data, fewer clicks to reach the same information, and more flexibility in changing the content and display of data you are viewing.

Problem 4

- Publishes tables of counts, for counts that are less than 10, they are suppressed as *

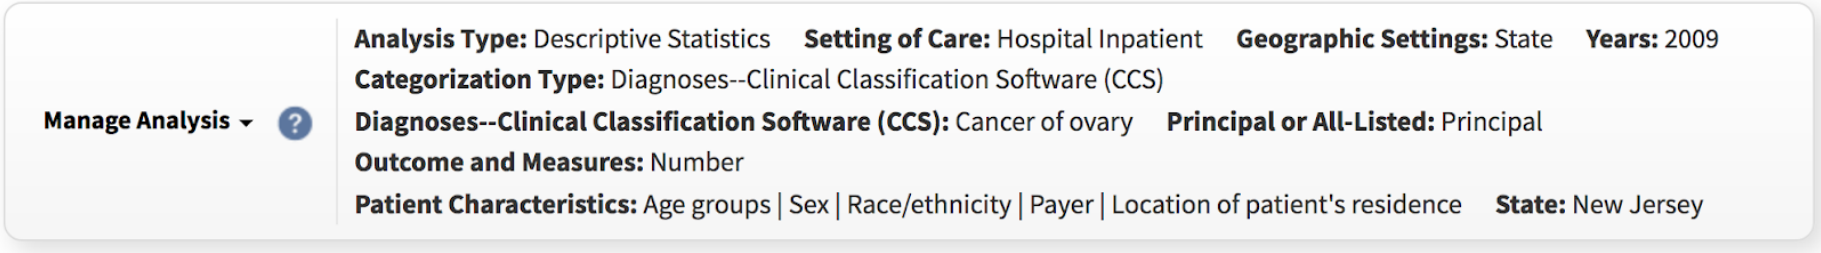

- Can you tell their values?

Let's begin! (15 mins)

Problem 1: Naïve Anonymization

- Auxiliary knowledge:
  - Alice has sent emails to Bob, Cathy, and Ed
  - Cathy has sent emails to everyone, except Ed

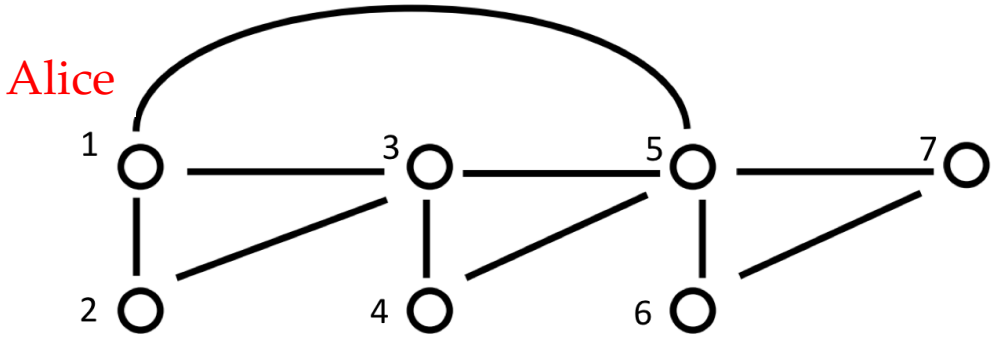

- Only one node has a degree 3 → node 1: Alice

Problem 1: Naïve Anonymization

- Auxiliary knowledge:
  - Alice has sent emails to Bob, Cathy, and Ed
  - Cathy has sent emails to everyone, except Ed

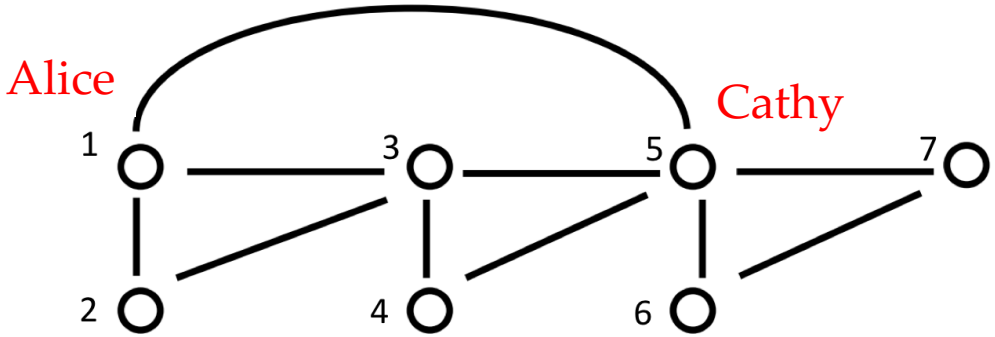

- Only one node has a degree 5 → node 5: Cathy

Problem 1: Naïve Anonymization

- Auxiliary knowledge:
  - Alice has sent emails to Bob, Cathy, and Ed
  - Cathy has sent emails to everyone, except Ed

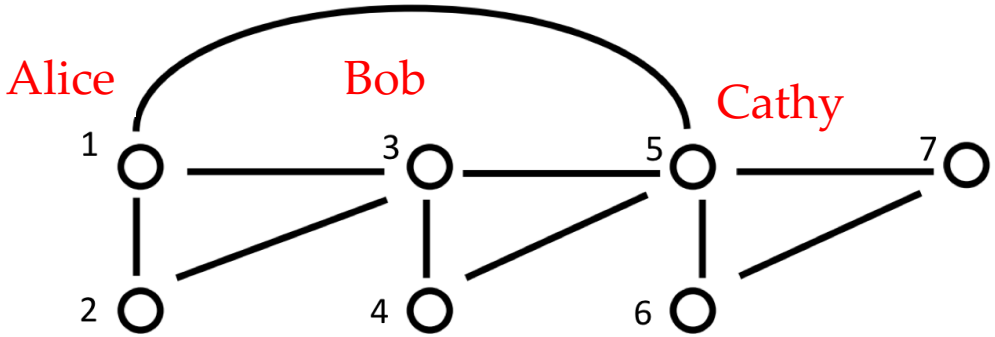

- Alice and Cathy know that only Bob has sent emails to both of them → node 3: Bob

Problem 1: Naïve Anonymization

- Auxiliary knowledge:
  - Alice has sent emails to Bob, Cathy, and Ed
  - Cathy has sent emails to everyone, except Ed

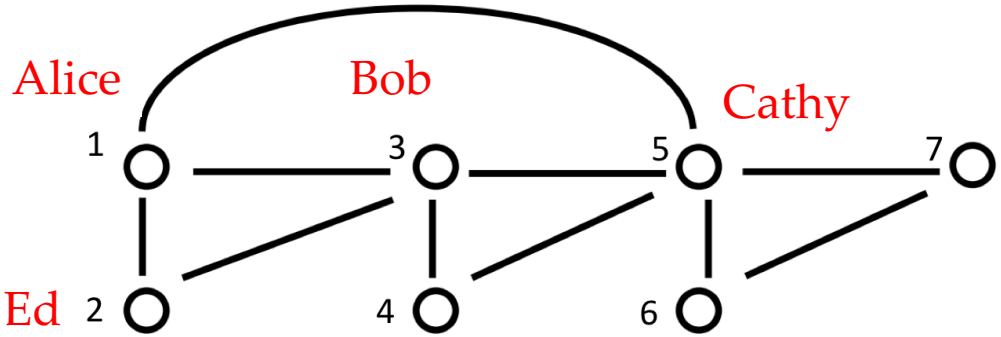

- Alice has sent emails to Bob, Cathy, and Ed only
  → node 2: Ed

Attacks using Background Knowledge

- Degrees of nodes [Liu and Terzi, SIGMOD 2008]
- The network structure, e.g., a subgraph of the network. [Zhou and Pei, ICDE 2008, Hay et al., VLDB 2008]
- Anonymized graph with labeled nodes [Pang et al., SIGCOMM CCR 2006]

Desiderata for a Privacy Definition

1. Resilience to background knowledge
   - A privacy mechanism must be able to protect individuals' privacy from attackers who may possess background knowledge

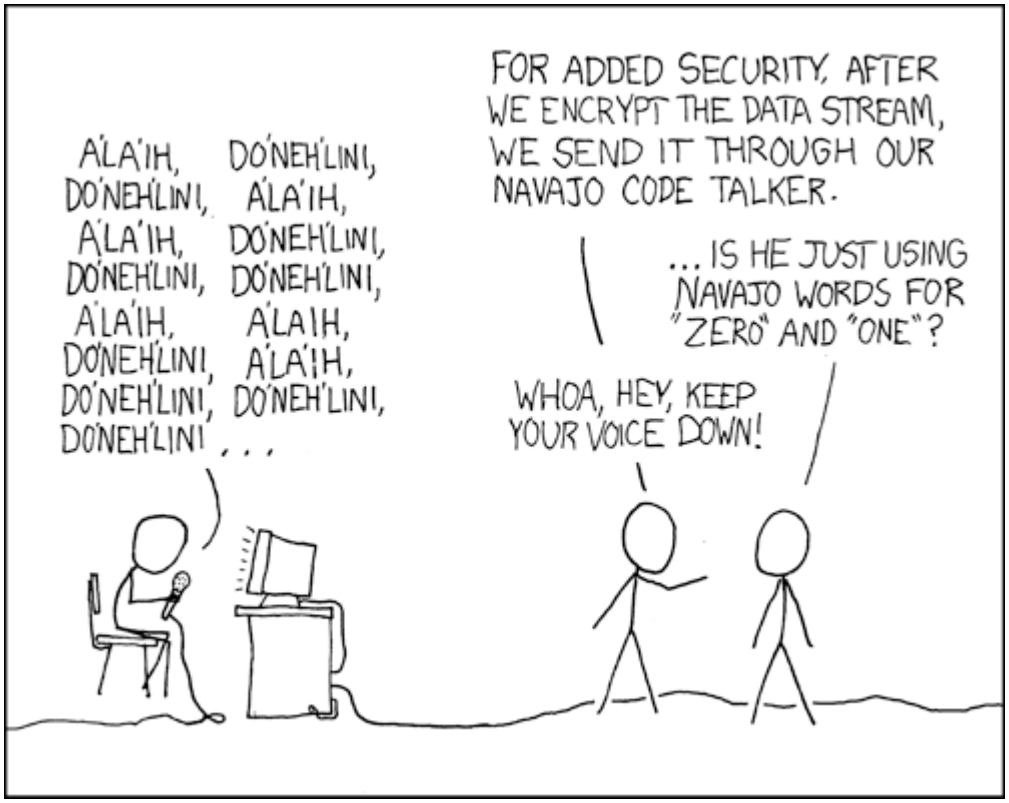
A'LA'IH, DO'NEH'LINI,
DO'NEH'LINI, A'LA'IH,
A'LA'IH, DO'NEH'LINI,
DO'NEH'LINI, DO'NEH'LINI,
A'LA'IH, A'LA!H,
DO'NEH'LINI, A'LA'!H,
DO'NEH'LINI, DO'NEH'LINI,
DO'NEH'LINI . . .
FOR ADDED SECURITY, AFTER WE ENCRYPT THE DATA STREAM, WE SEND IT THROUGH OUR NAVAJO CODE TALKER.
... IS HE JUST USING NAVAJO WORDS FOR "ZERO" AND "ONE"?
WHOA, HEY, KEEP YOUR VOICE DOWN!

Problem 2: Privacy by Obscurity

- Many organization think their data are private because they perturb the data and make the parameters of perturbation secret.

Problem 2: Privacy by Obscurity

Node ID	Name	Age ($\alpha x + \beta$)	True Age
1	Alice	40	25
2	Ed	34	
3	Bob	52	
4		28	
5	Cathy	48	29
6		22	
7		92	

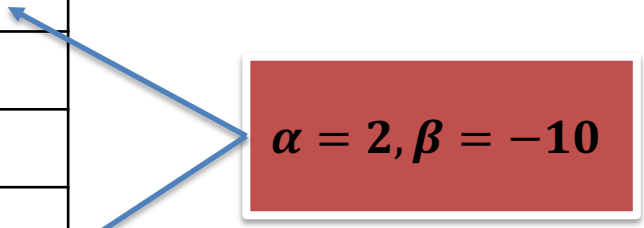

Problem 2: Privacy by Obscurity

Node ID	Name	Age ($\alpha x + \beta$)	True Age
1	Alice	40	25
2	Ed	34	22
3	Bob	52	31
4		28	19
5	Cathy	48	29
6		22	16
7		92	51

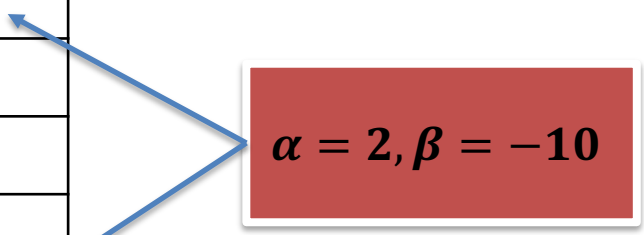

Desiderata for a Privacy Definition

1. Resilience to background knowledge
   - A privacy mechanism must be able to protect individuals' privacy from attackers who may possess background knowledge

2. Privacy without obscurity
   - Attacker must be assumed to know the algorithm used as well as all parameters [MK15]

Problem 4: Post-processing

Counts less than k are suppressed achieving k-anonymity

Age	#discharges	White	Black	Hispanic	Asian/ Pcf Hlnder	Native American	Other	Missing
#discharges	735	535	82	58	18	*	19	22
1-17	*	*	*	*	*	*	*	*
18-44	70	40	13	*	*	*	*	*
45-64	330	236	31	32	*	*	11	*
65-84	298	229	35	13	*	*	*	*
85+	34	29	*	*	*	*	*	*

Problem 4: Post-processing

Age	#discharges	White	Black	Hispanic	Asian/ Pcf Hlnder	Native American	Other	Missing
#discharges	735	535	82	58	18	**1**	19	22
1-17	**3**	**1**	*	*	*	*	*	*
18-44	70	40	13	*				*
45-64	330	236	31	32			1	*
65-84	298	229	35	13	*	*	*	*
85+	34	29	*	*	*	*	*	*

= 535 – (40+236+229+29)

Problem 4: Post-processing

Age	#discharges	White	Black	Hispanic	Asian/ Pcf Hlnder	Native American	Other	Missing
#discharges	735	535	82	58	18	**1**	19	22
1-17	**3**	**1**	[0-2]	[0-2]	[0-2]	[0-2]	[0-2]	[0-2]
18-44	70	40	13	*	*	*	*	*
45-64	330	236	31	32	*	*	11	*
65-84	298	229	35	13	*	*	*	*
85+	34	29	*	*	*	*	*	*

Problem 4: Post-processing

Age	#discharges	White	Black	Hispanic	Asian/ Pcf Hlnder	Native American	Other	Missing
#discharges	735	535	82	58	18	**1**	19	22
1-17	**3**	**1**	[0-2]	[0-2]	[0-2]	[0-2]	[0-2]	[0-2]
18-44	70	40	13	*	*	*	*	*
45-64	330	236	31	32	*	*	11	*
65-84	298	229	35	13	*	*	*	*
85+	34	29	[1-3]	*	*	*	*	*

Can Construct Tight Bounds on Rest of Data

[VSJO 13]

Age	#discharges	White	Black	Hispanic	Asian/ Pcf Hlnder	Native American	Other	Missing
#discharges	735	535	82	58	18	1	19	22
1-17	3	1	[0-2]	[0-2]	[0-1]	[0]	[0-1]	[0-1]
18-44	70	40	13	[9-10]	[0-6]	[0]	[0-6]	[1-8]
45-64	330	236	31	32	[10]	[0]	11	[10]
65-84	298	229	35	13	[2-8]	[1]	[2-8]	[4-10]
85+	34	29	[1-3]	[1-4]	[0-1]	[0]	[0-1]	[0-1]

Can Construct Tight Bounds on Rest of Data

[VSJO 13]

In fact, when linked with queries giving other statistics, we can figure out that exactly 1 Native American woman diagnosed with ovarian cancer went to a privately owned, not for profit, teaching hospital in new Jersey with more than 435 beds in 2009. Furthermore, the woman did not pay by private insurance, had a routine discharge, with a stay in the hospital of 33.5 days, with her home residence being in a county with 1 million plus residents (large fringe metro, suburbs), and her age was exactly 75 years.

Desiderata for a Privacy Definition

1. Resilience to background knowledge
   - A privacy mechanism must be able to protect individuals' privacy from attackers who may possess background knowledge

2. Privacy without obscurity
   - Attacker must be assumed to know the algorithm used as well as all parameters [MK15]

3. Post-processing
   - Post-processing the output of a privacy mechanism must not change the privacy guarantee [KL10, MK15]

Problem 3: Multiple Releases

- 2 tables of k-anonymous patient records [GKS08]

	Non-Sensitive			Sensitive
	Zip code	Age	Nationality	Condition
1	130**	<30	*	AIDS
2	130**	<30	*	Heart Disease
3	130**	<30	*	Viral Infection
4	130**	<30	*	Viral Infection
5	130**	≥40	*	Cancer
6	130**	≥40	*	Heart Disease
7	130**	≥40	*	Viral Infection
8	130**	≥40	*	Viral Infection
9	130**	3*	*	Cancer
10	130**	3*	*	Cancer
11	130**	3*	*	Cancer
12	130**	3*	*	Cancer

Hospital A (4-anonymous)

	Non-Sensitive			Sensitive
	Zip code	Age	Nationality	Condition
1	130**	<35	*	AIDS
2	130**	<35	*	Tuberculosis
3	130**	<35	*	Flu
4	130**	<35	*	Tuberculosis
5	130**	<35	*	Cancer
6	130**	<35	*	Cancer
7	130**	≥35	*	Cancer
8	130**	≥35	*	Cancer
9	130**	≥35	*	Cancer
10	130**	≥35	*	Tuberculosis
11	130**	≥35	*	Viral Infection
12	130**	≥35	*	Viral Infection

Hospital B (6-anonymous)

- Alice is 28 and she visits both hospitals

Problem 3: Multiple Releases

- 2 tables of k-anonymous patient records [GKS08]

	Non-Sensitive			Sensitive
	Zip code	Age	Nationality	Condition
1	130**	<30	*	AIDS
2	130**	<30	*	Heart Disease
3	130**	<30	*	Viral Infection
4	130**	<30	*	Viral Infection
5	130**	≥40	*	Cancer
6	130**	≥40	*	Heart Disease
7	130**	≥40	*	Viral Infection
8	130**	≥40	*	Viral Infection
9	130**	3*	*	Cancer
10	130**	3*	*	Cancer
11	130**	3*	*	Cancer
12	130**	3*	*	Cancer

Hospital A (4-anonymous)

	Non-Sensitive			Sensitive
	Zip code	Age	Nationality	Condition
1	130**	<35	*	AIDS
2	130**	<35	*	Tuberculosis
3	130**	<35	*	Flu
4	130**	<35	*	Tuberculosis
5	130**	<35	*	Cancer
6	130**	<35	*	Cancer
7	130**	≥35	*	Cancer
8	130**	≥35	*	Cancer
9	130**	≥35	*	Cancer
10	130**	≥35	*	Tuberculosis
11	130**	≥35	*	Viral Infection
12	130**	≥35	*	Viral Infection

Hospital B (6-anonymous)

- 4-anonymity + 6-anonymity ⇏ k-anonymity , for any k

Desiderata for a Privacy Definition

1. Resilience to background knowledge
   - A privacy mechanism must be able to protect individuals' privacy from attackers who may possess background knowledge

2. Privacy without obscurity
   - Attacker must be assumed to know the algorithm used as well as all parameters [MK15]

3. Post-processing
   - Post-processing the output of a privacy mechanism must not change the privacy guarantee [KL10, MK15]

4. Composition over multiple releases
   - Allow a graceful degradation of privacy with multiple invocations on the same data [DN03, GKS08]

Why Composition?

- Reasoning about privacy of a complex algorithm is hard.

- Helps software design
  - If building blocks are proven to be private, it would be easy to reason about privacy of a complex algorithm built entirely using these building blocks.

Dinur Nissim Result [DN03]

- A vast majority of records in a database of size n can be reconstructed when $n \log(n)^2$ queries are answered by a statistical database …

  … even if each answer has been arbitrarily altered to have up to $o(\sqrt{n})$ error

A Bound on the Number of Queries

- In order to ensure utility, a statistical database must leak some information about each individual

- We can only hope to bound the amount of disclosure

- Hence, there is a limit on number of queries that can be answered

Desiderata for a Privacy Definition

1. Resilience to background knowledge
   - A privacy mechanism must be able to protect individuals' privacy from attackers who may possess background knowledge
2. Privacy without obscurity
   - Attacker must be assumed to know the algorithm used as well as all parameters [MK15]
3. Post-processing
   - Post-processing the output of a privacy mechanism must not change the privacy guarantee [KL10, MK15]
4. Composition over multiple releases
   - Allow a graceful degradation of privacy with multiple invocations on the same data [DN03, GKS08]

Summary

- Privacy attacks on naïve approaches
- Desiderata include resilience to background knowledge, privacy without obscurity, closure under post-processing, and composition.
- Next, how to define privacy and design privacy-preserving mechanism that achieve these desiderata?
  - Differential Privacy
  - Basic Algorithms and Composition

References

- [S02] Sweeney, “K-anonymity”, IJFUKS 2010
- [LT08] Liu and Terzi, “Towards Identity Anonymization on Graphs”, SIGMOD 2008
- [ZP08] Zhou and Pei, “Preserving Privacy in Social Networks Against Neighborhood Attacks”, ICDE 2008
- [HMJTW08] Hay et al, “Resisting Structural Reidentification Anonymized Social Networks”, VLDB 2008
- [PAPL06] Pang et al , “The devil and packet trace anonymization”, SIGCOMM CCR 2006
- [VSJO13] Vaidya et al., “Identifying inference attacks against healthcare data repositories”, AMIA 2013
- [GKS08] Ganta et al. “Composition Attacks and Auxiliary Information in Data Privacy”, KDD 2008
- [DN03] Dinur, Nissim, “Revealing information while preserving privacy”, PODS 2003
- [KL10] Kifer, Lin, “Towards an Axiomatization of Statistical Privacy and Utility.”, PODS 2010
- [MK15] Machanavajjhala, Kifer, “Designing statistical privacy for your data”, CACM 2015