What is privacy?

Mainack Mondal

CS 60081
Autumn 2024

Roadmap

- What is privacy?
  - Privacy theories
- How to protect privacy?

Recall CIA model

- Is confidentiality == privacy?

Recall CIA model

- Is confidentiality == privacy? NO

Recall CIA model

- Is confidentiality == privacy? NO
- Secrecy, confidentiality, privacy, anonymity are different
  - Secrecy: Keep data hidden (e.g., incriminating evidence)
  - Confidentiality: Keep data hidden from unauthorized entities
  - Privacy: Both disclose and hide a person's data depending on correct context
  - Anonymity: Keep identify of a person secret

Recall CIA model

- Is confidentiality == privacy? NO
- Secrecy, confidentiality, privacy, anonymity are different
  - Secrecy: Keep data hidden (e.g., incriminating evidence)
  - Confidentiality: Keep data hidden from unauthorized entities
  - Privacy: Both disclose and hide a person's data depending on correct context
  - Anonymity: Keep identify of a person secret

Privacy: definitions

- Have a very extensive history
  - 1890: Warren and Brandeis (Law)
  - 1967: Alan Westin (Law)
  - 1975: Irwin Altman (Anthropology)
    - 1992: Sandra Petronio (CPM theory)
    - 2003: Palen and Dourish's interpretation
  - 2008: Daniel Solove (Solove's taxonomy
  - 2011: Helen Nissenbaum (Contextual integrity theory)

  - Privacy laws around the world

Privacy: definitions

- Have a very extensive history
  - 1890: Warren and Brandeis (Law)
  - 1967: Alan Westin (Law)
  - ~~1975: Irwin Altman (Anthropology)~~
    - ~~1992: Sandra Petronio (CPM theory)~~
    - ~~2003: Palen and Dourish's interpretation~~
  - 2008: Daniel Solove (Solove's taxonomy
  - 2011: Helen Nissenbaum (Contextual integrity theory)

  - Privacy laws around the world

Warren and Brandeis's theory (1890)

- "the protection afforded to thoughts, sentiments, and emotions, … enforcement of the more general *right of the individual to be let alone*"
  - Libel and slander are insufficient in considering only damage to reputation
  - The right to prevent, rather than profit from, publication of personal information
  - What about information of public figures / incidents?

Warren and Brandeis's theory (1890)

- "the protection afforded to thoughts, sentiments, and emotions, … enforcement of the more general *right of the individual to be let alone*"
  - Libel and slander are insufficient in considering only damage to reputation
  - The right to prevent, rather than profit from, publication of personal information
  - What about information of public figures / incidents? (Does not consider them)

Westin: Privacy as control (1967)

- "Privacy is the claim of individuals, groups or institutions to determine for themselves when, how, and to what extent information about them is communicated to others."

--- Alan Westin

- Four states of privacy
  - Solitude: not observed by others
  - Intimacy: communicate with a small group
  - Anonymity: free from identification/surveilance
  - Reserve: limit information disclosure to others and others respecting the desire

Westin: Privacy as control (1967)

- Question:
  - X and Y are sitting in a restaurant and X was talking about his personal life.
  - Z, an eavesdropper sitting in the next table, are listening to them, although X did not realize it.
  - Can you explain, using Alan Westin's privacy definition and privacy states, if X's privacy is being violated in this scenario?

Solove's pluralistic notion of privacy

- Uses Wittgenstein's concept of 'family resemblances'
  - capture the notion of privacy people have in their mind
  - Privacy has many meanings
  - Like family resemblances, they are all related

- Focuses on data lifecycle
  - Different disruptions in each phase (data collection, processing, dissemination and invasion)
  - https://wiki.openrightsgroup.org/wiki/A_Taxonomy_of_Privacy

Solove's pluralistic notion of privacy

Information collection	surveillance (watching, listening to, or recording an individual's activities)
	interrogation (pressuring of individuals to divulge information)
Information processing	aggregation (gathering together information about a person)
	identification (connecting information to an individual)
	insecurity (problems caused by the way information is handled and protected)
	secondary use (use of data for purposes unrelated to the purposes for which the data was initially collected without the data subject's consent)
	exclusion (failure to provide individuals with notice and input about their records)

Solove's pluralistic notion of privacy

Information Dissemination	breach of confidentiality (breaking a promise to keep a person's information confidential)
	disclosure (revealing true information about a person to others)
	exposure (revealing another's nudity, grief, or bodily functions)
	increased accessibility (amplifying the accessibility of information)
	blackmail (threat to disclose personal information)
	appropriation (use of the data subject's identity to serve the aims and interests of another)
	distortion (the dissemination of false information about a person)
Invasion	intrusion (invasive acts that disturb one's tranquility or solitude)
	decisional interference (the government's incursion into the data subject's decisions regarding her private affairs)

An interesting application

a person should not worry about government or surveillance if they have "nothing to hide."

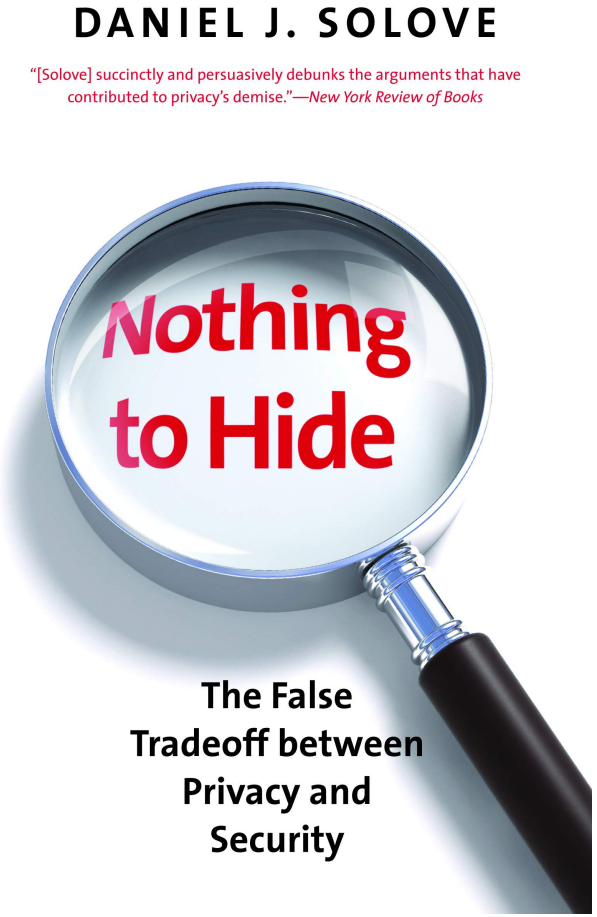

Nothing to hide

"When engaged directly, the nothing-to-hide argument can ensnare, for it forces the debate to focus on its narrow understanding of privacy. But when confronted with the plurality of privacy problems implicated by government data collection and use beyond surveillance and disclosure, the nothing-to-hide argument, in the end, has nothing to say."

Nothing to hide

"When engaged directly, the nothing-to-hide argument can ensnare, for it forces the debate to focus on its narrow understanding of privacy. But when confronted with the plurality of privacy problems implicated by government data collection and use beyond surveillance and disclosure, the nothing-to-hide argument, in the end, has nothing to say." (aggregation, secondary use)

Privacy as contextual Integrity (CI)

- A “normative” model of privacy
  - How privacy should be

- Considers “appropriate flow” of information
  - Appropriate flows conform to contextual informational norms
  - Each norm is : <Data subject, sender, recipient, information type, and transmission principle>
  - There are socially appropriate *norms*
  - Useful to systematically understand social norms: Evaluating the Contextual Integrity of Privacy Regulation: Parents' IoT Toy Privacy Norms Versus COPPA, N. Apthorpe, S. Varghese, N. Feamster, USENIX Security Symposium, 2019

Basic approach of CI

A framework to argue about privacy violation

Privacy is preserved by **appropriate flows of information**

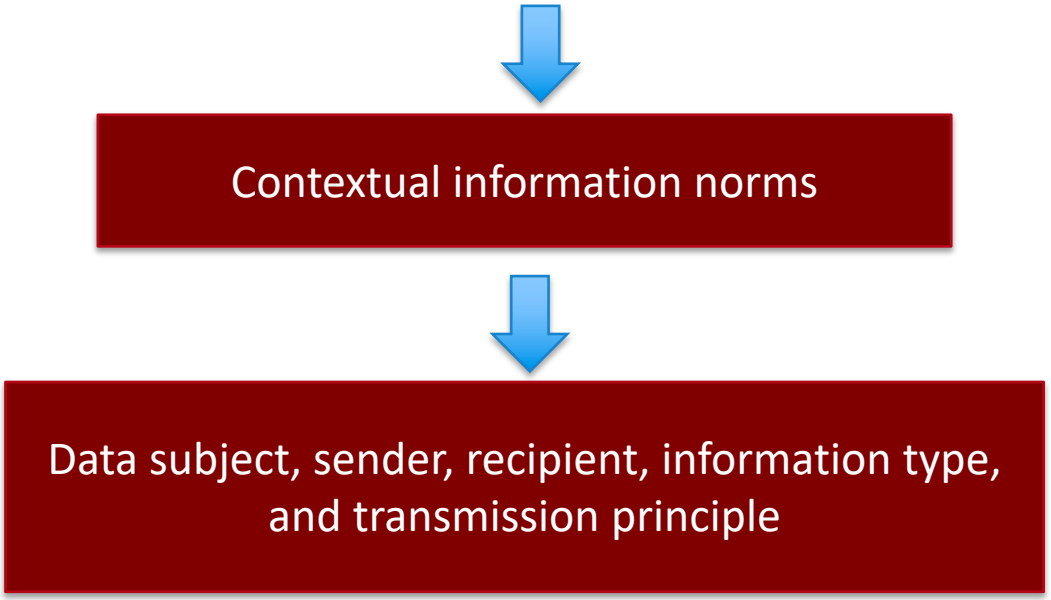

Conceptions of privacy are based on **dynamic ethical concerns**

Roadmap

- What is privacy?
  - Privacy theories

- How to protect privacy?

Privacy laws around the world

- US has sector-specific laws, minimal protections
  - FTC investigates fraud & deceptive practices
  - FCC regulates telecommunications
- EU GDPR (general data protection regulation)
  - Later in the course

Fair Information practice principles (FIPP)

- Notice/Awareness: Consumers should be given notice of an entity's information practices before any personal information is collected from them

Fair Information practice principles (FIPP)

- Notice/Awareness: Consumers should be given notice of an entity's information practices before any personal information is collected from them
- Choice/Consent: Choice and consent in an on-line information-gathering sense means giving consumers options to control how their data is used

Fair Information practice principles (FIPP)

- Notice/Awareness: Consumers should be given notice of an entity's information practices before any personal information is collected from them
- Choice/Consent: Choice and consent in an on-line information-gathering sense means giving consumers options to control how their data is used
- Access/Participation: Not only a consumer's ability to view the data collected, but also to verify and contest its accuracy in inexpensive and timely manner

Fair Information practice principles (FIPP)

- Notice/Awareness: Consumers should be given notice of an entity's information practices before any personal information is collected from them
- Choice/Consent: Choice and consent in an on-line information-gathering sense means giving consumers options to control how their data is used
- Access/Participation: Not only a consumer's ability to view the data collected, but also to verify and contest its accuracy in inexpensive and timely manner
- Integrity/Security: Information collectors should ensure that the data they collect is accurate and secure

Fair Information practice principles (FIPP)

- Notice/Awareness: Consumers should be given notice of an entity's information practices before any personal information is collected from them
- Choice/Consent: Choice and consent in an on-line information-gathering sense means giving consumers options to control how their data is used
- Access/Participation: Not only a consumer's ability to view the data collected, but also to verify and contest its accuracy in inexpensive and timely manner
- Integrity/Security: Information collectors should ensure that the data they collect is accurate and secure
- Enforcement/Redress: In order to ensure that companies follow the Fair Information Practice Principles, there must be enforcement measures (self-regulation, sue by users, Government regulation)

Privacy enhancing tools

- Encryption (bitlocker)
- Anonymity (Tor, VPN)
- Tracker Blockers (ublock, adblocker)
- Opt-out tools (consent form, cookie banners)
- Social network privacy controls / Access control

How to measure privacy?

- Internet Users' Information Privacy Concerns (IUIPC)
  - 10 multiple choice questions divided into 3 sections: control, awareness, collection
  - Options are 7-point scale for each questions
    - From strongly disagree to strongly agree

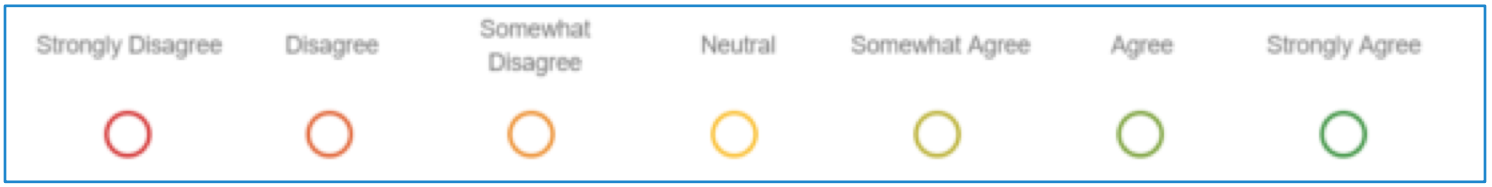

- Malhotra, N.K., Kim, S.S., Agarwal,J., 2004. Internet users' information privacy concerns (IUIPC): the construct, the scale, and a causal model. Information Systems Research15(4), 336–355.

IUIPC control scale questions

1. Consumer online privacy is really a matter of consumers' right to exercise control and autonomy over decisions about how their information is collected, used, and shared.

2. Consumer control of personal information lies at the heart of consumer privacy.

3. I believe that online privacy is invaded when control is lost or unwillingly reduced as a result of a marketing transaction.

IUIPC awareness scale questions

1. Companies seeking information online should disclose the way the data are collected, processed, and used.

2. A good consumer online privacy policy should have a clear and conspicuous disclosure.

3. It is very important to me that I am aware and knowledgeable about how my personal information will be used.

IUIPC collection scale questions

1. It usually bothers me when online companies ask me for personal information.
2. When online companies ask me for personal information, I sometimes think twice before providing it.
3. It bothers me to give personal information to so many online companies.
4. I'm concerned that online companies are collecting too much personal information about me.

Protecting privacy of data

- We want to protect privacy of sensitive data attributes

- What if the attacker capability is not known?

Privacy of databases

- Mechanisms for hiding privacy sensitive attributes in databases
  - K-anonymity
  - Differential privacy

- Slides heavily borrowed from
  - Vitaly Smatikov from Cornell
  - Li Xiong from Emory
  - Aaron Roth from Upenn
  - Sebastian Benthall from NYU

Public Data Conundrum

- Health-care datasets
  - Clinical studies, hospital discharge databases …
- Genetic datasets
  - $1000 genome, HapMap, deCode …
- Demographic datasets
  - U.S. Census Bureau, sociology studies …
- Search logs, recommender systems, social networks, blogs …
  - AOL search data, social networks of blogging sites, Netflix movie ratings, Amazon …

What About Privacy?

- First thought: anonymize the data
- How?
- Remove "personally identifying information" (PII)
  - Name, Social Security number, phone number, email, address… what else?
  - Anything that identifies the person directly
- Is this enough?

Re-identification by Linking

Microdata

ID	QID			SA
Name	Zipcode	Age	Sex	Disease
Alice	47677	29	F	Ovarian Cancer
Betty	47602	22	F	Ovarian Cancer
Charles	47678	27	M	Prostate Cancer
David	47905	43	M	Flu
Emily	47909	52	F	Heart Disease
Fred	47906	47	M	Heart Disease

Voter registration data

Name	Zipcode	Age	Sex
Alice	47677	29	F
Bob	47983	65	M
Carol	47677	22	F
Dan	47532	23	M
Ellen	46789	43	F

Latanya Sweeney's Attack (1997)

Massachusetts hospital discharge dataset

Medical Data Released as Anonymous

SSN	Name	...nicity	Date Of Birth	Sex	ZIP	Marital Status	Problem
			09/27/64	female	02139	divorced	hypertension
			09/30/64	female	02139	divorced	obesity
		asian	04/18/64	male	02139	married	chest pain
		asian	04/15/64	male	02139	married	obesity
		black	03/13/63	male	02138	married	hypertension
		black	03/18/63	male	02138	married	shortness of breath
		black	09/13/64	female	02141	married	shortness of breath
		black	09/07/64	female	02141	married	obesity
		white	05/14/61	male	02138	single	chest pain
		white	05/08/61	male	02138	single	obesity
		white	09/15/61	female	02142	widow	shortness of breath

Voter List

Name	Address	City	ZIP	DOB	Sex	Party	.....
.....	.....	.....	.....	.....	.....	.....	
.....	.....	.....	.....	.....	.....	.....	
Sue J. Carlson	1459 Main St.	Cambridge	*02142*	*9/15/61*	*female*	democrat	.....
.....	.....	.....	.....	.....	.....	.....	

Figure 1: ...e-identifying anonymous data by linking to external data

Public voter dataset

Quasi-Identifiers

- Key attributes
  - Name, address, phone number - uniquely identifying!
  - Always removed before release
- Quasi-identifiers
  - (5-digit ZIP code, birth date, gender) uniquely identify 87% of the population in the U.S.
  - Can be used for linking anonymized dataset with other datasets

Classification of Attributes

- Sensitive attributes
  - Medical records, salaries, etc.
  - These attributes is what the analysts need, so they are always released directly

Key Attribute	Quasi-identifier			Sensitive attribute
Name	DOB	Gender	Zipcode	Disease
Andre	1/21/76	Male	53715	Heart Disease
Beth	4/13/86	Female	53715	Hepatitis
Carol	2/28/76	Male	53703	Brochitis
Dan	1/21/76	Male	53703	Broken Arm
Ellen	4/13/86	Female	53706	Flu
Eric	2/28/76	Female	53706	Hang Nail

K-Anonymity: Intuition

- The information for each person contained in the released table cannot be distinguished from at least k-1 individuals whose information also appears in the release
  - Example: you try to identify a man in the released table, but the only information you have is his birth date and gender. There are k men in the table with the same birth date and gender.
- Any quasi-identifier present in the released table must appear in at least k records

Generalization

- Goal of k-Anonymity
  - Each record is indistinguishable from at least k-1 other records
  - These k records form an equivalence class
- Generalization: replace quasi-identifiers with less specific, but semantically consistent values

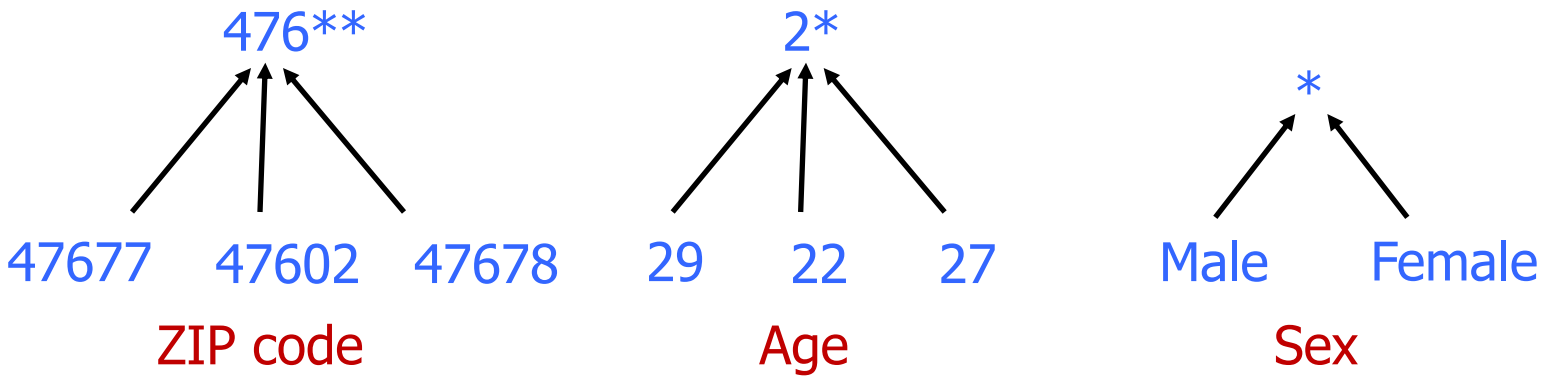

Achieving k-Anonymity

- Generalization
  - Replace specific quasi-identifiers with less specific values until get k identical values
  - Partition ordered-value domains into intervals

Example of a k-Anonymous Table

	Race	Birth	Gender	ZIP	Problem
t1	Black	1965	m	0214*	short breath
t2	Black	1965	m	0214*	chest pain
t3	Black	1965	f	0213*	hypertension
t4	Black	1965	f	0213*	hypertension
t5	Black	1964	f	0213*	obesity
t6	Black	1964	f	0213*	chest pain
t7	White	1964	m	0213*	chest pain
t8	White	1964	m	0213*	obesity
t9	White	1964	m	0213*	short breath
t10	White	1967	m	0213*	chest pain
t11	White	1967	m	0213*	chest pain

Figure 2 Example of *k*-anonymity, where *k*=2 and QI={*Race*, *Birth*, *Gender*, *ZIP*}

At least two people
With same attributes

QI = quasi identifier tuple

Curse of Dimensionality

[Aggarwal VLDB '05]

- Generalization fundamentally relies on spatial locality
  - Each record must have k close neighbors
- Real-world datasets are very sparse
  - Many attributes (dimensions)
    - Amazon customer records: several million dimensions
  - Not possible to create k close neighbors
- Projection to low dimensions loses all info ⇒ k-anonymized datasets are useless

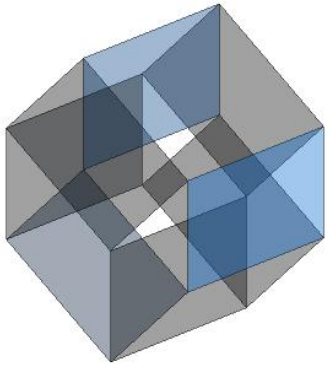

Two (and a Half) Interpretations

- Membership disclosure: Attacker cannot tell that a given person in the dataset
- Sensitive attribute disclosure: Attacker cannot tell that a given person has a certain sensitive attribute
- Identity disclosure: Attacker cannot tell which record corresponds to a given person

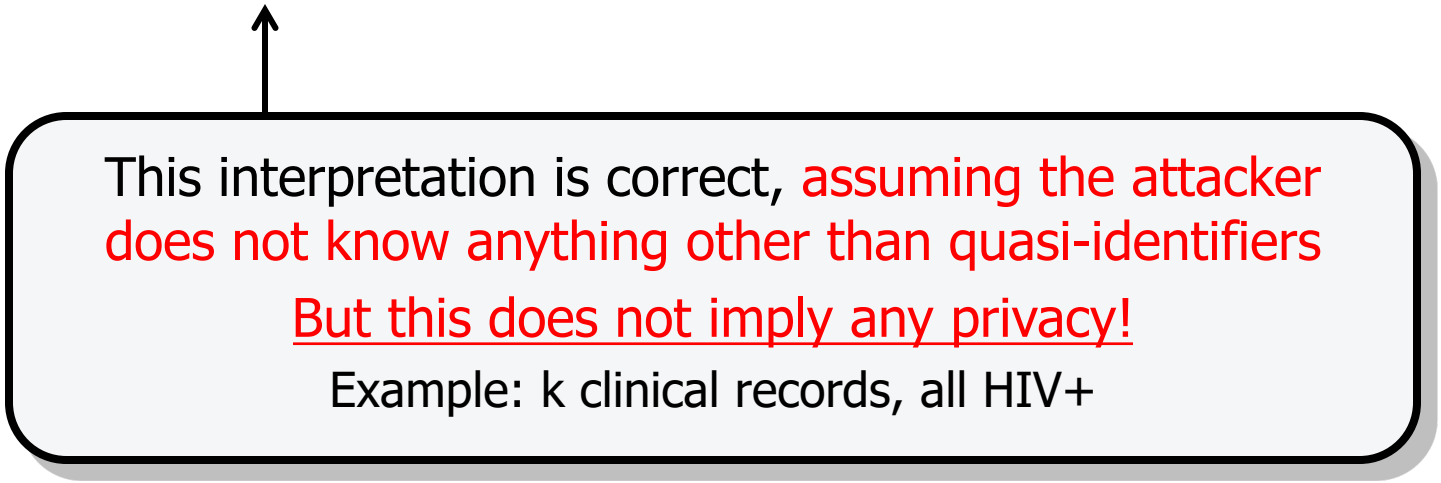

Attacks on k-Anonymity

- k-Anonymity does not provide privacy if
  - Sensitive values in an equivalence class lack diversity
  - The attacker has background knowledge

Homogeneity attack

Bob	
Zipcode	***Age***
47678	27

Background knowledge attack

Yoshiko		
Zipcode	***Age***	***Race***
47673	36	Japanese

Low chance of heart disease

A 3-anonymous patient table

Zipcode	Age	Disease
476**	2*	Heart Disease
476**	2*	Heart Disease
476**	2*	Heart Disease
4790*	≥40	Flu
4790*	≥40	Heart Disease
4790*	≥40	Cancer
476**	3*	Heart Disease
476**	3*	Cancer
476**	3*	Cancer

k-Anonymity Considered Harmful

- Syntactic
  - Focuses on data transformation, not on what can be learned from the anonymized dataset
  - "k-anonymous" dataset can leak sensitive information
- "Quasi-identifier" fallacy
  - Assumes a priori that attacker will not know certain information about his target
- Relies on locality
  - Destroys utility of many real-world datasets

What are we going to talk about?

- Mechanisms for hiding privacy sensitive attributes in databases
  - K-anonymity
  - Differential privacy

- Slides heavily borrowed from
  - Vitaly Smatikov from Cornell
  - Li Xiong from Emory
  - Aaron Roth from Upenn
  - Sebastian Benthall from NYU
  - Roger Grosse from University of Toronto

First an intuition

Randomized response

A way to ensure *some* privacy

- Have you ever dodged your taxes?

A way to ensure *some* privacy

- Have you ever dodged your taxes?
- Flip a coin.
- If the coin lands Heads, then answer truthfully.
- If it lands Tails, then flip it again.
  - If it lands Heads, then answer Yes.
  - If it lands Tails, then answer No.

Probability of responses?

A way to ensure *some* privacy

- Have you ever dodged your taxes?
- Flip a coin.
- If the coin lands Heads, then answer truthfully.
- If it lands Tails, then flip it again.
  - If it lands Heads, then answer Yes.
  - If it lands Tails, then answer No.

Probability of responses?

	Yes	No
Dodge	3/4	1/4
No Dodge	1/4	3/4

A way to ensure *some* privacy

Tarun the Tax Investigator assigns a prior probability of 0.02 to Shiv the businessman having dodged his taxes. Then he notices he answered Yes to the survey. What is the posterior probability?

A way to ensure *some* privacy

Tarun the Tax Investigator assigns a prior probability of 0.02 to Shiv the businessman having dodged his taxes. Then he notices he answered Yes to the survey. What is the posterior probability?

$$\begin{aligned}\Pr(\text{Dodge} \mid \text{Yes}) &= \frac{\Pr(\text{Dodge})\Pr(\text{Yes} \mid \text{Dodge})}{\Pr(\text{Dodge})\Pr(\text{Yes} \mid \text{Dodge}) + \Pr(\text{NoDodge})\Pr(\text{Yes} \mid \text{NoDodge})} \\ &= \frac{0.02 \cdot \frac{3}{4}}{0.02 \cdot \frac{3}{4} + 0.98 \cdot \frac{1}{4}} \\ &\approx 0.058\end{aligned}$$

So Tarun's beliefs have not shifted too much

Lets improve this idea: blending in the crowd

Blending into a Crowd

Frequency in DB or frequency in underlying population?

- Intuition: "I am safe in a group of k or more"
  - k varies (3… 6… 100… 10,000?)
- Many variations on theme
  - Adversary wants predicate g such that $0 < \#\{i \mid g(x_i)=true\} < k$
- Why?
  - Rare property helps re-identify someone
  - Implicit: information about a large group is public (prior distribution)
    - E.g., liver problems more prevalent among diabetics

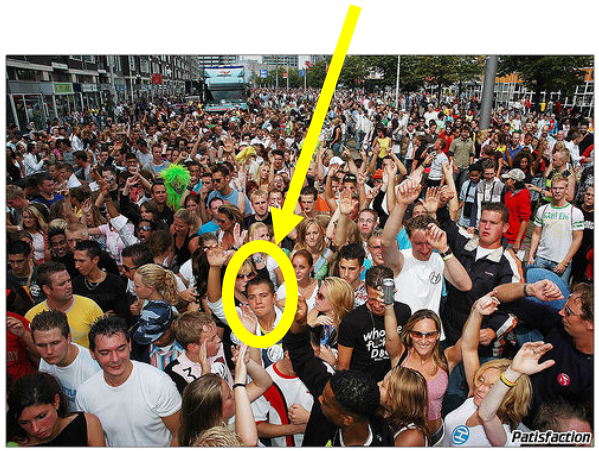

Clustering-Based Definitions

- k-anonymity
  - Partition D into bins
  - Safe if each bin is either empty, or contains at least k elements
- Cell bound methods
  - Release marginal sums

	brown	blue	Σ
blond	2	10	12
brown	12	6	18
Σ	14	16	

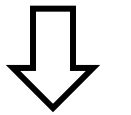

	brown	blue	Σ
blond	[0,12]	[0,12]	12
brown	[0,14]	[0,16]	18
Σ	14	16	

Issues with Clustering

- Purely syntactic definition of privacy
- What adversary does this apply to?
  - Does not consider adversaries with side information
  - Does not consider probability
  - Does not consider adversarial algorithm for making decisions (inference)

"Bayesian" Adversaries

- Adversary outputs point $z \in D$
- Score = $1/f_z$ if $f_z > 0$, 0 otherwise
  - f_z is the number of matching points in D
- Sanitization is safe if $E(\text{score}) \leq \varepsilon$
- Procedure:
  - Assume you know adversary's prior distribution over databases
  - Given a candidate output, update prior conditioned on output (via Bayes' rule)
  - If $\max_z E(\text{score} \mid \text{output}) < \varepsilon$, then safe to release

Issues with “Bayesian” Privacy

- Restricts the type of predicates adversary can choose
- Must know prior distribution
  - Can one scheme work for many distributions?
  - Sanitizer works harder than adversary
- Conditional probabilities don’t consider previous iterations

Basic setup of differential privacy

- There is a database D which potentially contains sensitive information about individuals.
- The database curator has access to the full database. We assume the curator is trusted.
- The data analyst wants to analyze the data.
  - She asks a series of queries to the curator, and the curator provides a response to each query.
- The way in which the curator responds to queries is called the mechanism.
  - We'd like a mechnism that gives helpful responses but avoids leaking sensitive information about individuals.

Differential Privacy set up

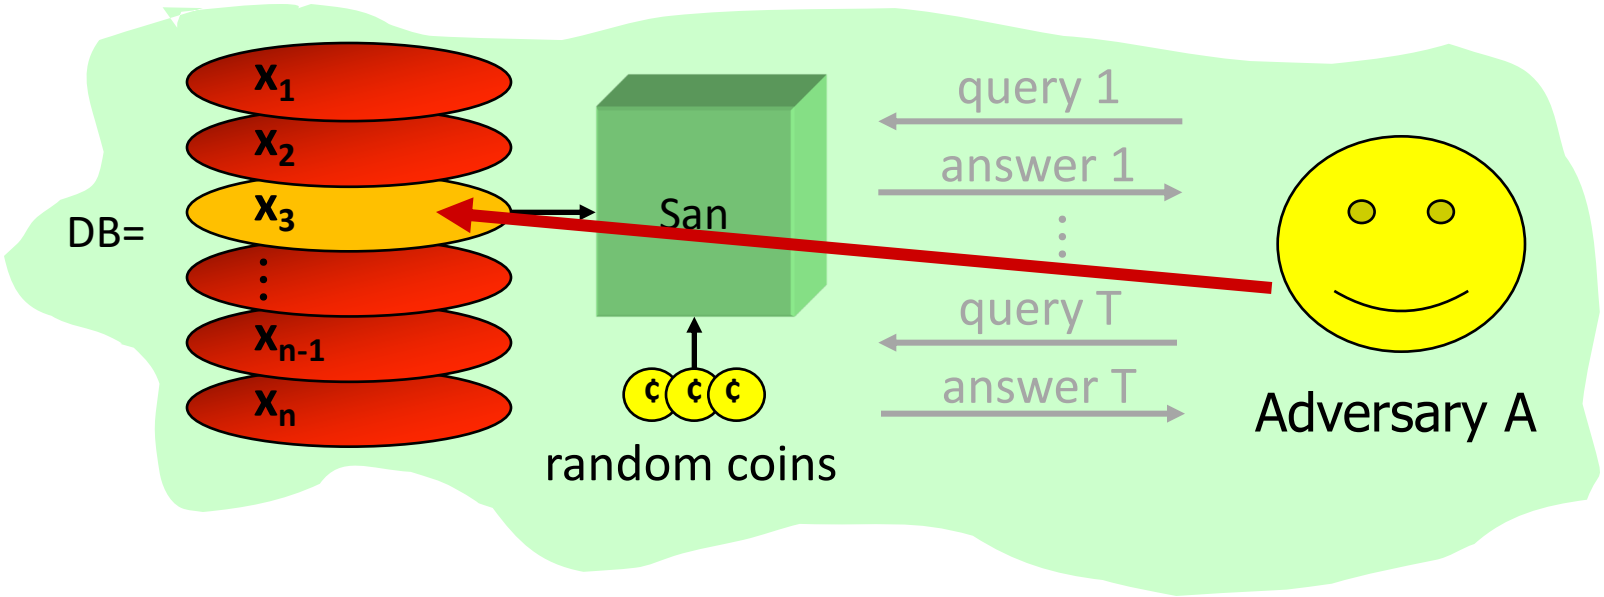

Idea: Whatever is already known, situation won't get worse

Differential Privacy

A mechanism $\mathcal{M}$ is ε-differentially private if for any two neighbouring databases $\mathcal{D}_1$ and $\mathcal{D}_2$, and any set $\mathcal{R}$ of possible responses

$$\Pr(\mathcal{M}(\mathcal{D}_1) \in \mathcal{R}) \leq \exp(\varepsilon) \Pr(\mathcal{M}(\mathcal{D}_2) \in \mathcal{R}).$$

An approximation for small ε…

Differential Privacy

A mechanism $\mathcal{M}$ is ε-differentially private if for any two neighbouring databases $\mathcal{D}_1$ and $\mathcal{D}_2$, and any set $\mathcal{R}$ of possible responses

$$\Pr(\mathcal{M}(\mathcal{D}_1) \in \mathcal{R}) \leq \exp(\varepsilon) \Pr(\mathcal{M}(\mathcal{D}_2) \in \mathcal{R}).$$

An approximation for small ε, exp(ε) ≈ 1 + ε

Also for any response y

$$\exp(-\varepsilon) \leq \frac{\Pr(\mathcal{M}(\mathcal{D}_1) = y)}{\Pr(\mathcal{M}(\mathcal{D}_2) = y)} \leq \exp(\varepsilon)$$

An example

- Anna is an attacker who wants to figure out if Patrick (x) is in the cancer database D. Her prior probability for him being in the database is 0.4. D is ε-differentially private. She makes a query and gets back y = M(D).
- She's narrowed it down to two possible databases D_1 and D_2, which are identical except that $x \in D_1$ and $x \notin D_2$.
- After observing y, what are bounds on posterior probability $\Pr(x \in D \mid y)$ using bayes rule?

An example

$$\begin{aligned}
\Pr(x \in \mathcal{D} \mid y) &= \frac{\Pr(x \in \mathcal{D}) \Pr(y \mid x \in \mathcal{D})}{\Pr(x \in \mathcal{D}) \Pr(y \mid x \in \mathcal{D}) + \Pr(x \notin \mathcal{D}) \Pr(y \mid x \notin \mathcal{D})} \\
&\geq \frac{\Pr(x \in \mathcal{D}) \Pr(y \mid x \in \mathcal{D})}{\Pr(x \in \mathcal{D}) \Pr(y \mid x \in \mathcal{D}) + \exp(\varepsilon) \Pr(x \notin \mathcal{D}) \Pr(y \mid x \in \mathcal{D})} \\
&= \frac{\Pr(x \in \mathcal{D})}{\Pr(x \in \mathcal{D}) + \exp(\varepsilon) \Pr(x \notin \mathcal{D})} \\
&\geq 0.4 \exp(-\varepsilon)
\end{aligned}$$

- Similarly $\Pr(x \in \mathcal{D} \mid y) \leq 0.4 \exp(\varepsilon)$
- So, Anna has not learned much about Patrick

Differential Privacy

[Dwork ICALP 2006]

For every pair of inputs that differ in one row

For every output …

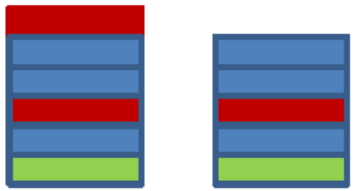

D_1 D_2

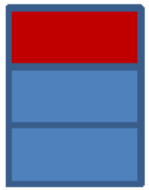

O

Adversary should not be able to distinguish between any D_1 and D_2 based on any O

$$\log\left(\frac{\Pr[A(D_1) = O]}{\Pr[A(D_2) = O]}\right) < \varepsilon \quad (\varepsilon > 0)$$

Can deterministic algorithms satisfy differential privacy?

Non trivial deterministic algorithms do not satisfy differential privacy

Space of all inputs

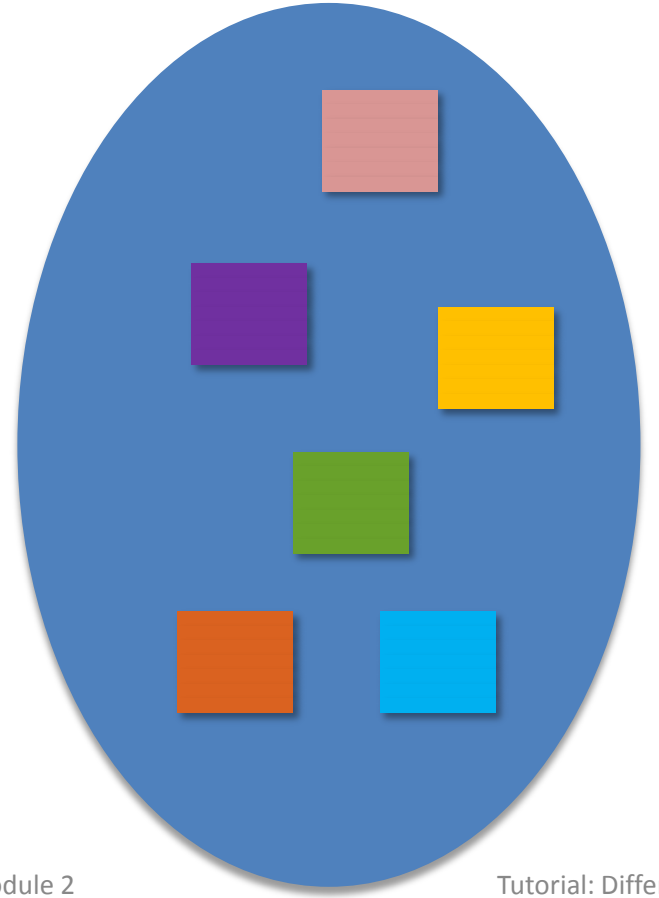

Space of all outputs
(at least 2 distinct outputs)

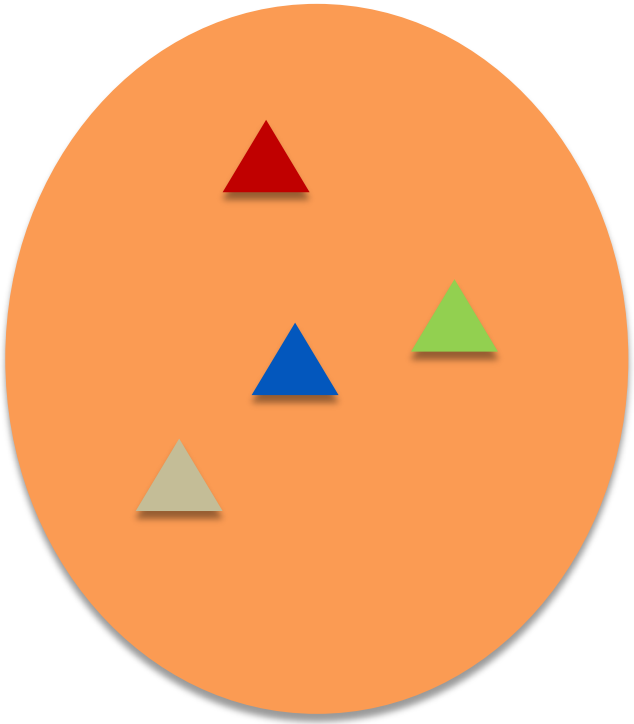

Non-trivial deterministic algorithms do not satisfy differential privacy

Each input mapped to a distinct output.

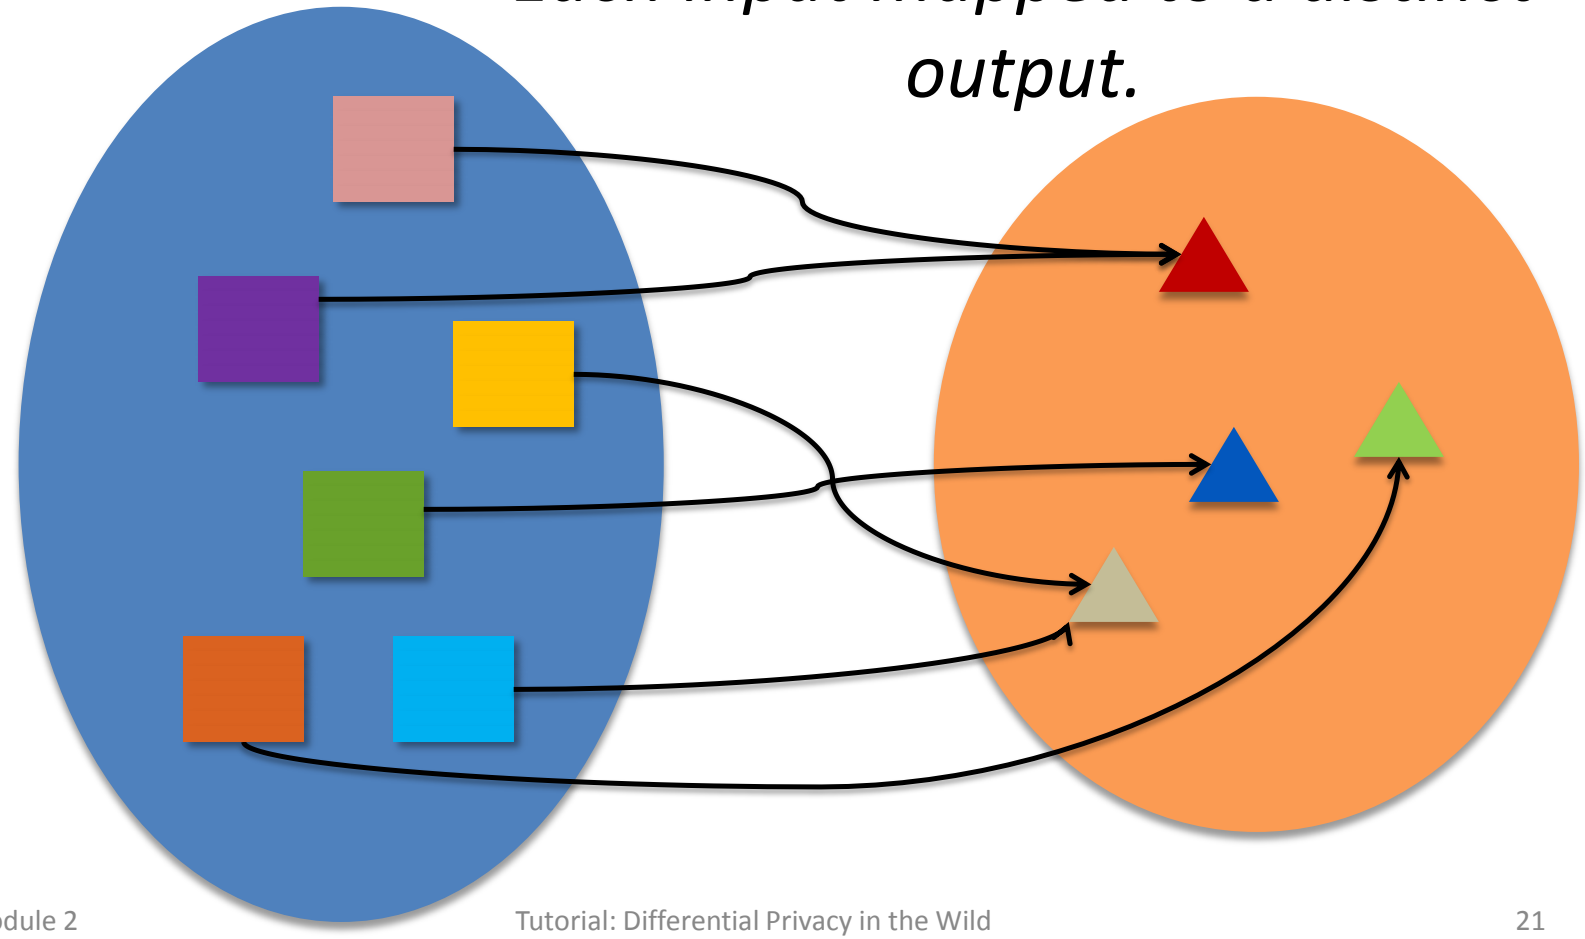

There exist two inputs that differ in one entry mapped to different outputs.

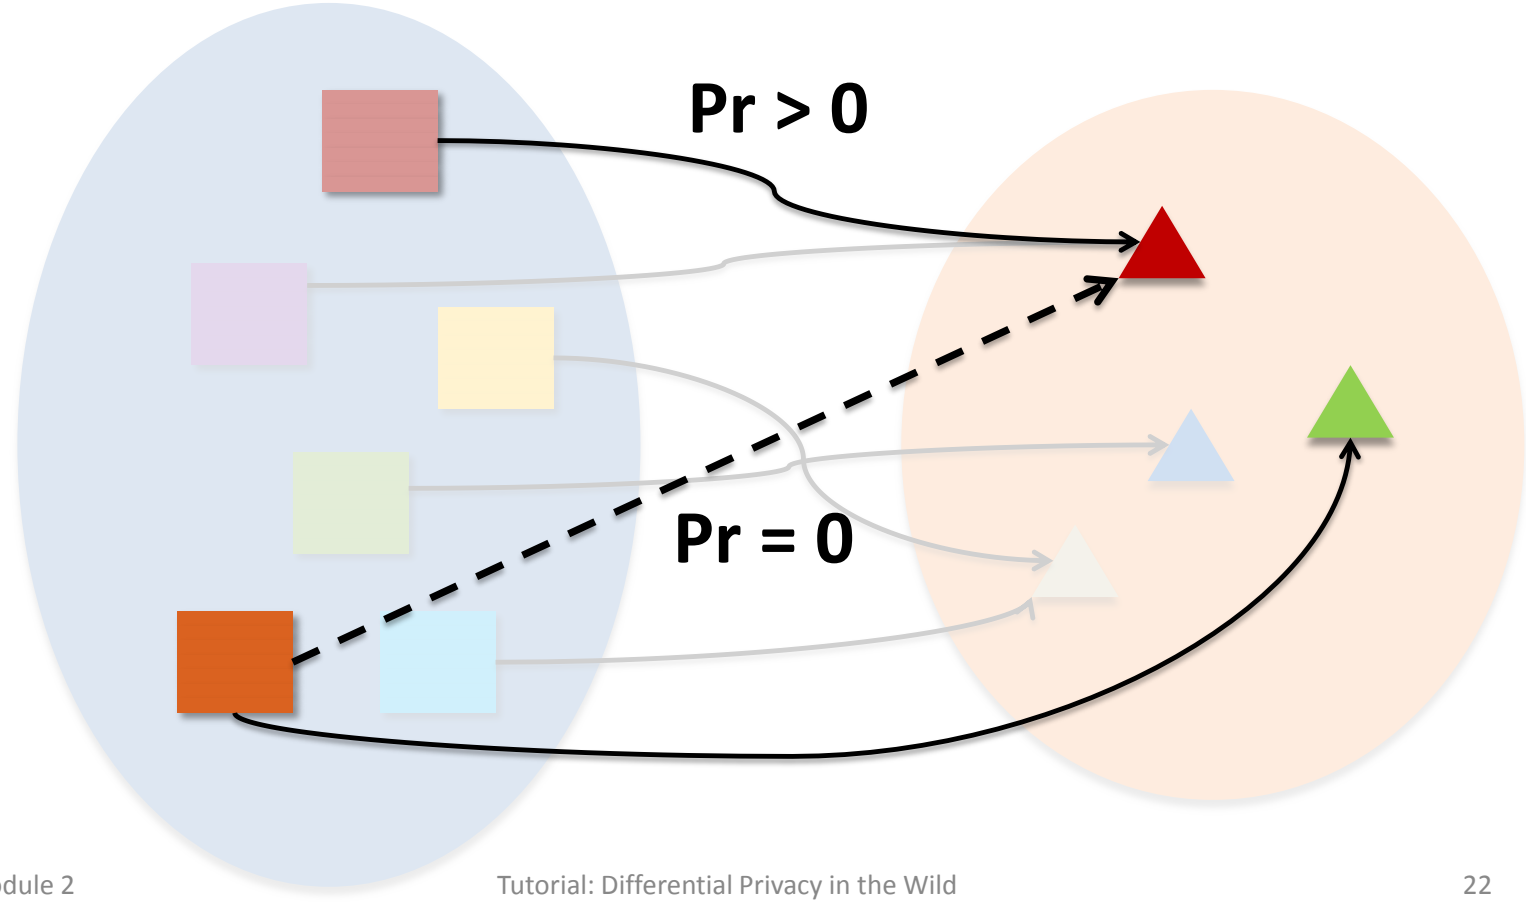

Output Randomization

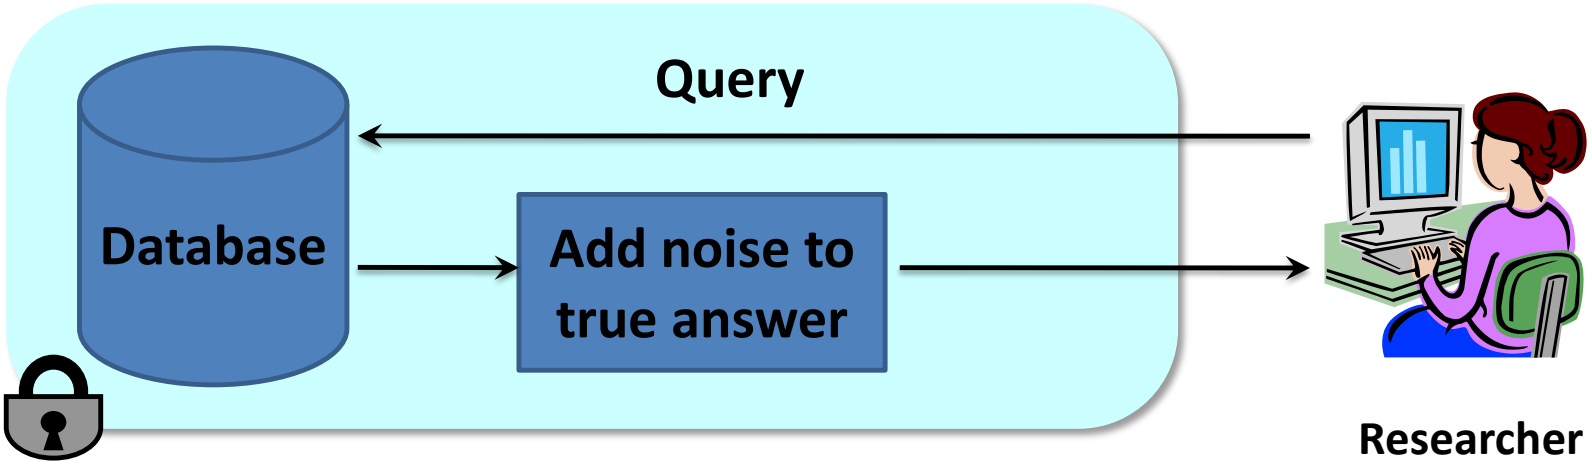

- Add noise to answers such that:
  - Each answer does not leak too much information about the database.
  - Noisy answers are close to the original answers.

[DMNS 06]

Laplace Mechanism

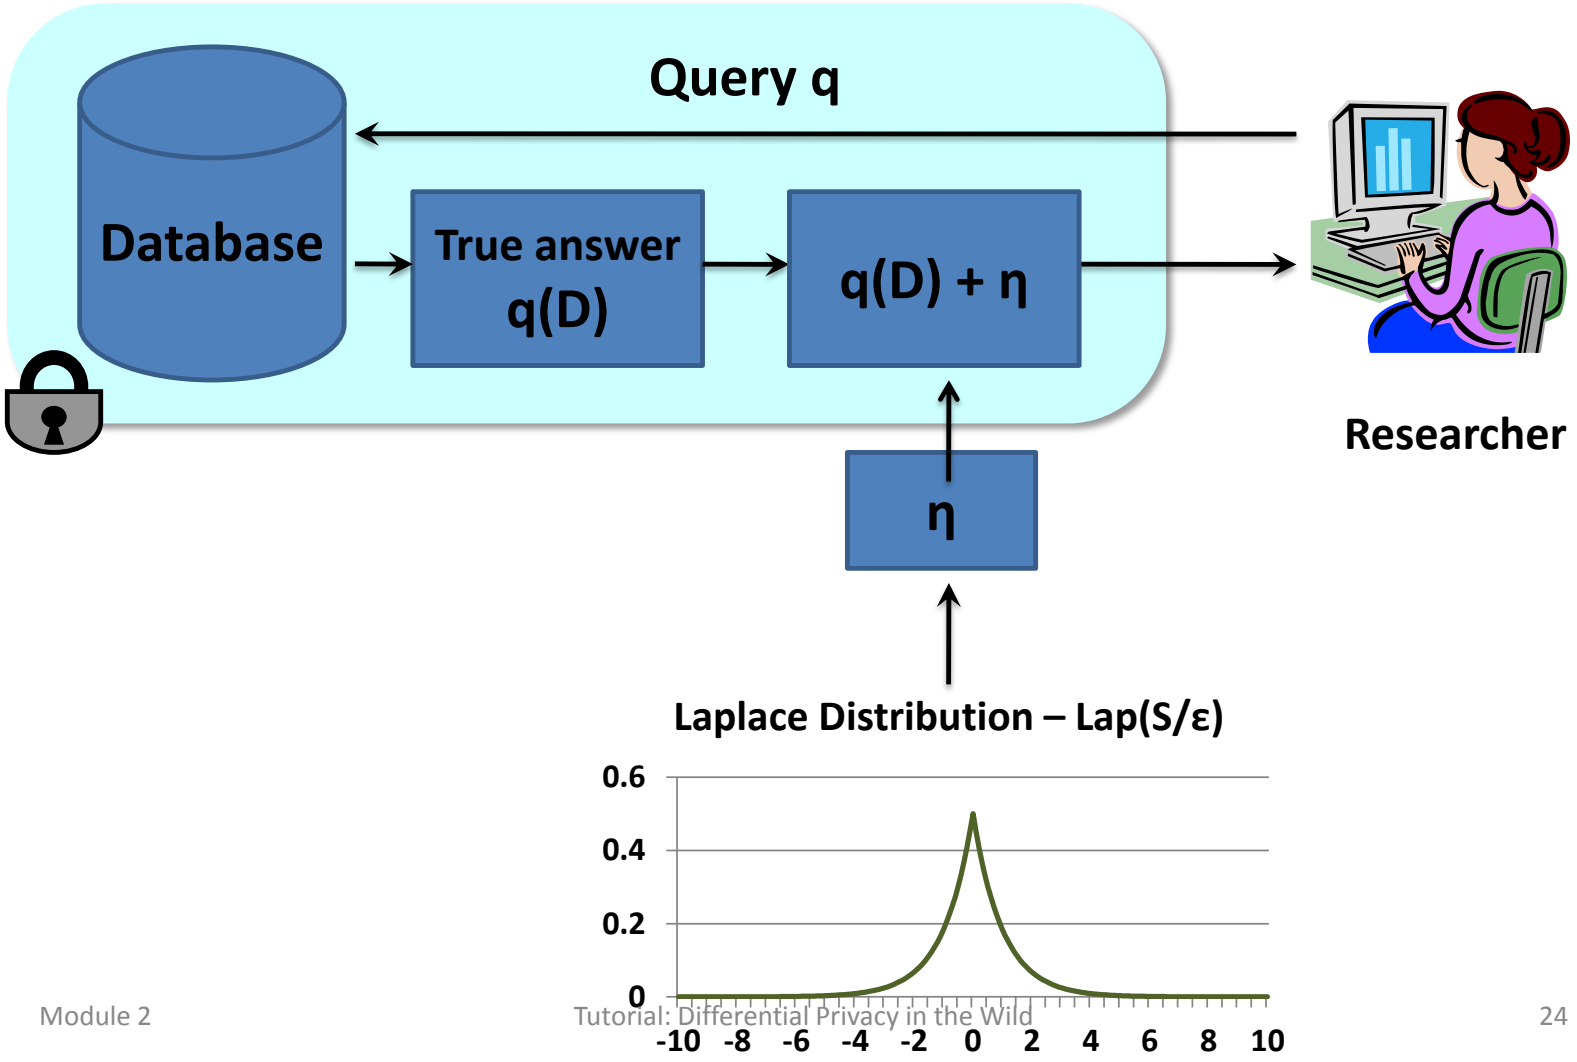

Laplace Distribution

- PDF: $f(x \mid \mu, b) = \frac{1}{2b} \exp\left(-\frac{|x-\mu|}{b}\right)$
- Denoted as Lap(b) when u=0
- Mean u
- Variance $2b^2$

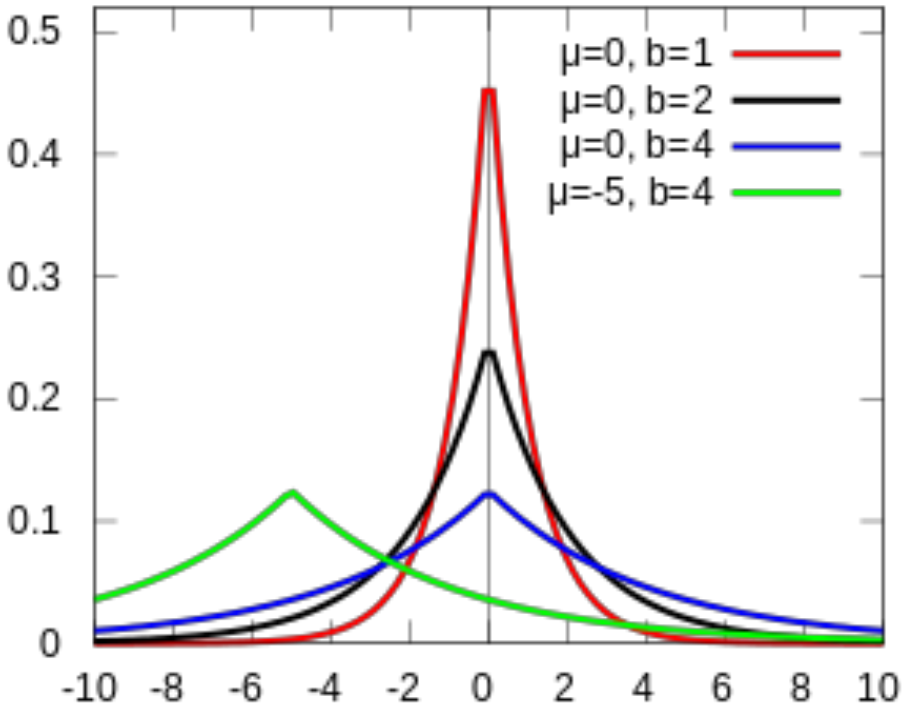

How much noise for privacy?

[Dwork et al., TCC 2006]

Sensitivity: Consider a query q: $I \rightarrow R$. S(q) is the smallest number s.t. for any neighboring tables D, D',

$$| q(D) - q(D') | \leq S(q)$$

Theorem: If **sensitivity** of the query is **S**, then the algorithm $A(D) = q(D) + Lap(S(q)/\varepsilon)$ guarantees ε-differential privacy

Example: COUNT query

- Number of people having disease
- Sensitivity = 1

- Solution: $3 + \eta$,
  where η is drawn from $Lap(1/\varepsilon)$
  - Mean = 0
  - Variance = $2/\varepsilon^2$

D

Disease (Y/N)
Y
Y
N
Y
N
N

Example: SUM query

- Suppose all values x are in [a,b]
- Sensitivity = b

Privacy of Laplace Mechanism

- Consider neighboring databases D and D’
- Consider some output O

$$\frac{\Pr\left[A(D) = O\right]}{\Pr\left[A(D') = O\right]} = \frac{\Pr\left[q(D) + \eta = O\right]}{\Pr\left[q(D') + \eta = O\right]}$$

$$= \frac{e^{-|O - q(D)|/\lambda}}{e^{-|O - q(D')|/\lambda}}$$

$$\leq e^{|q(D) - q(D')|/\lambda} \leq e^{S(q)/\lambda} = e^{\varepsilon}$$

Utility of Laplace Mechanism

- Laplace mechanism works for **any function** that returns a real number
- Error: $E(\text{true answer} - \text{noisy answer})^2$

$$= \text{Var}(\text{Lap}(S(q)/\varepsilon))$$

$$= 2*S(q)^2 / \varepsilon^2$$

- Error bound: very unlikely the result has an error greater than a factor (Roth book Theorem 3.8)

Properties of differential privacy (1)

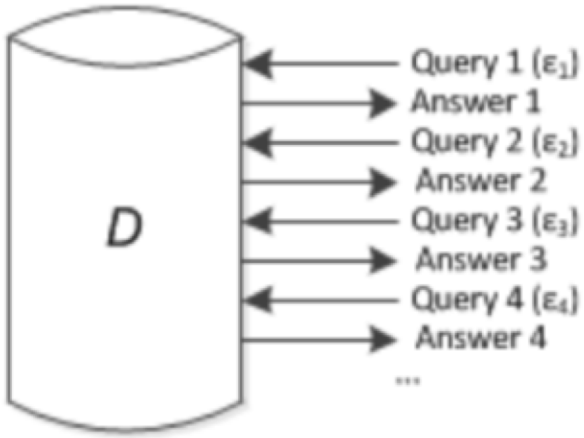

Sequential composition

Properties of differential privacy (1)

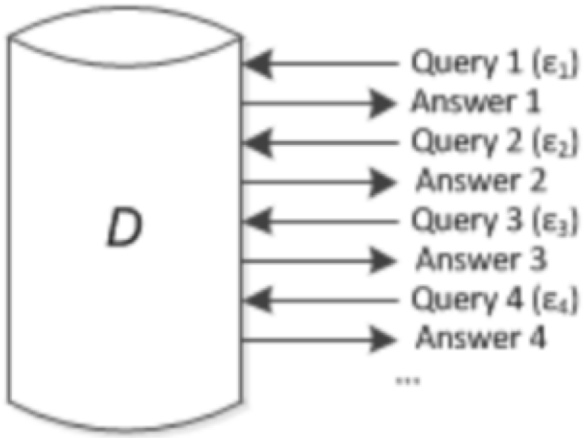

Sequential composition

If M_1, M_2, ..., M_k are algorithms that access a private database D such that each M_i satisfies ε_i - differential privacy,

then the combination of their outputs satisfies ε-differential privacy with $\varepsilon=\varepsilon_1+...+\varepsilon_k$

Properties of differential privacy (2)

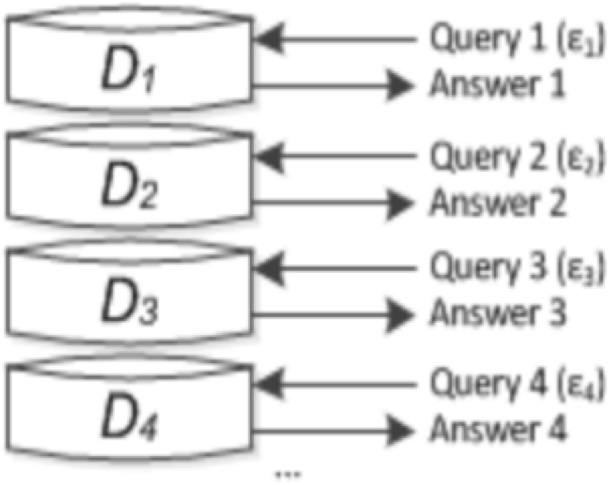

Parallel composition

Properties of differential privacy (2)

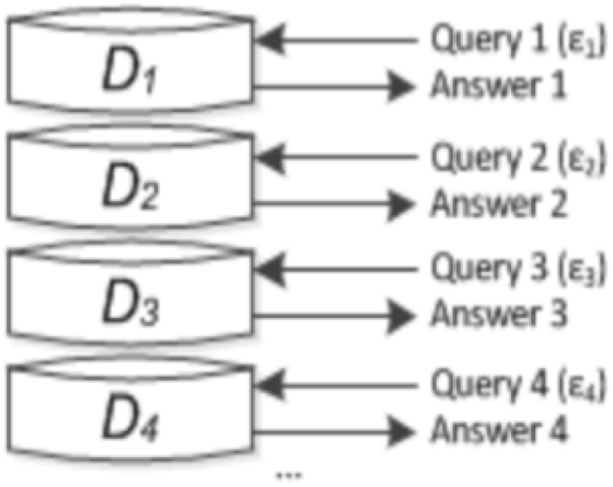

Parallel composition

If M_1, M_2, ..., M_k are algorithms that access disjoint databases D_1, D_2, ..., D_k such that each M_i satisfies ε_i -differential privacy,

then the combination of their outputs satisfies ε-differential privacy with $\varepsilon = \max\{\varepsilon_1,...,\varepsilon_k\}$

Properties of differential privacy (3)

- If M*1* is an ε differentially private algorithm that accesses a private database D,
- then outputting M*2*(M*1*(D)) also satisfies ε- differential privacy.

This is called postprocessing property

Postprocessing

- Suppose a data analyst takes the result y = M(D) and further processes it with some algorithm f (without peeking at the data itself). Is it still private?

Postprocessing

- Suppose a data analyst takes the result y = M(D) and further processes it with some algorithm f (without peeking at the data itself). Is it still private?
- Let R be a set of possible outputs, and R′ be the pre-image under f , i.e. $R' = \{y : f(y) \in R\}$.

$$\Pr(f(M(D1)) \in R) = \Pr(M(D1) \in R')$$

$$\leq \exp(\varepsilon) \Pr(M(D2) \in R')$$

$$= \exp(\varepsilon) \Pr(f(M(D2)) \in R)$$

Postprocessing

- Suppose a data analyst takes the result y = M(D) and further processes it with some algorithm f (without peeking at the data itself). Is it still private?
- Let R be a set of possible outputs, and R′ be the pre-image under f , i.e. $R' = \{y : f(y) \in R\}$.

$$\Pr(f(M(D1)) \in R) = \Pr(M(D1) \in R')$$

$$\leq \exp(\varepsilon) \Pr(M(D2) \in R')$$

$$= \exp(\varepsilon) \Pr(f(M(D2)) \in R)$$

Hence, the composition f (M) is also ε-differentially private. No matter how clever the analyst is, or the resources she throws at it, she can't learn more than ε about an individual entry!

Before we finish…

There is definition of approximate differential privacy

$$\Pr[M(X) \in \mathcal{S}] \leq e^{\epsilon} \Pr[M(X) \in \mathcal{S}] + \delta$$

Does sequential, parallel and postprocessing properties still hold?

Connection between DP and CI

- Integrating Differential Privacy and Contextual Integrity
  - https://www.usenix.org/conference/pepr22/presentation/benthall

Summary

- What of privacy
  - Right to be let alone
  - Westin's definition
  - Solovey's taxonomy
  - CI theory
- How of privacy
  - K-anonymity
  - Differential privacy

Practical resources

- A easy to digest book
  - https://programming-dp.com/
- A hard to digest book
  - https://www.cis.upenn.edu/~aaroth/Papers/privacybook.pdf
- A Privacy-Integrated Query Language (PINQ)
  - http://research.microsoft.com/en-us/projects/pinq/
- Fuzz: a typed functional language for differentially private computations
  - http://privacy.cis.upenn.edu/software.html