No Free Lunch in Data Privacy

CompSci 590.03
Instructor: Ashwin Machanavajjhala

Outline

- Background: Domain-independent privacy definitions
- No Free Lunch in Data Privacy [Kifer-**M** SIGMOD '11]
- Correlations: A case for domain specific privacy definitions [Kifer-**M** SIGMOD '11]
- Pufferfish Privacy Framework [Kifer-**M** PODS'12]
- Defining Privacy for Correlated Data [Kifer-**M** PODS'12 & Ding-**M** '13]
  - Next class

Data Privacy Problem

Utility: $f_{private}$ *approximates* f

Privacy: No breach about any individual

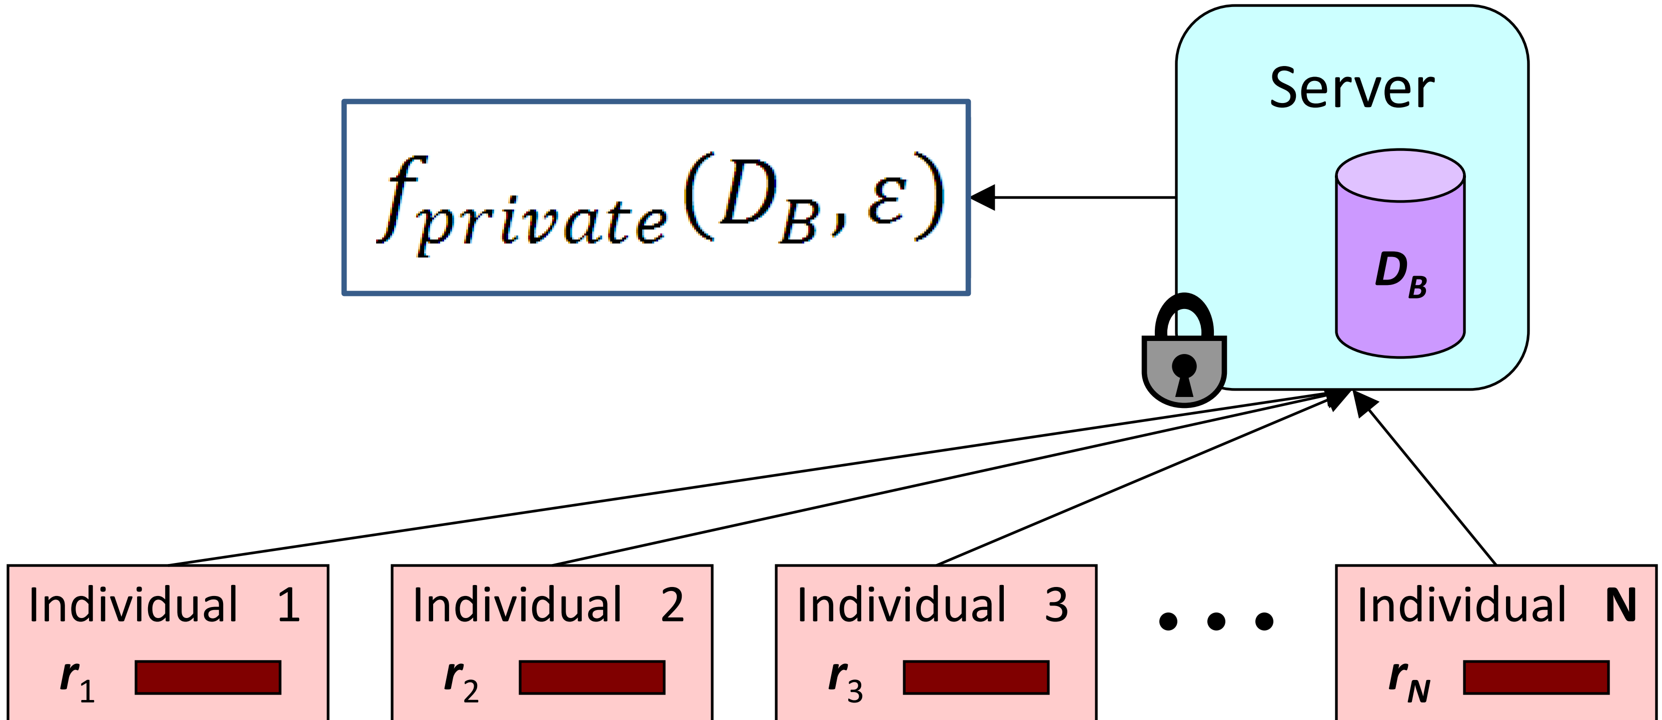

Data Privacy in the real world

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 Recommen-dations	Facebook	Another user	Friend links / profile	Recommend other users or ads to users based on social network

Semantic Privacy

... nothing about an individual should be learnable from the database that cannot be learned without access to the database.

T. Dalenius, 1977

Can we achieve semantic privacy?

- … or is there one *("precious…")* privacy definition to rule them all?

Defining Privacy

- In order to allow utility, a non-negligible amount of information about an individual must be disclosed to the adversary.

- Measuring information disclosed to an adversary involves carefully modeling the **background knowledge** already available to the adversary.

- … but we do not know what information is available to the adversary.

Many definitions & several attacks

K-Anonymity

Sweeney et al IJUFKS '02

L-diversity

Machanavajjhala et. al TKDD '07

T-closeness

Li et. al ICDE '07

E-Privacy

Machanavajjhala et. al VLDB '09

Differential Privacy

Dwork et. al ICALP '06

- Linkage attack
- Background knowledge attack
- Minimality /Reconstruction attack
- de Finetti attack
- **Composition attack**

Composability [Dwork et al, TCC 06]

Theorem **(Composability)**:

If algorithms A_1, A_2, …, A_k use independent randomness and each A_i satisfies $\boldsymbol{\varepsilon_i}$-differential privacy, resp.

Then, outputting all the answers together satisfies differential privacy with

$$\boldsymbol{\varepsilon = \varepsilon_1 + \varepsilon_2 + \ldots + \varepsilon_k}$$

Differential Privacy

- **Domain independent** privacy definition that is **independent of the attacker.**

- Tolerates many attacks that other definitions are susceptible to.
  - Avoids composition attacks
  - Claimed to be tolerant against adversaries with **arbitrary background knowledge.**

- Allows simple, efficient and useful privacy mechanisms
  - **Used in a live US Census Product** [**M** et al ICDE '08]

Duke UNIVERSITY

Outline

- Background: Domain independent privacy definitions.
- No Free Lunch in Data Privacy [Kifer-**M** SIGMOD '11]
- Correlations: A case for domain specific privacy definitions [Kifer-**M** SIGMOD '11]
- Pufferfish Privacy Framework [Kifer-**M** PODS'12]
- Defining Privacy for Correlated Data [Kifer-**M** PODS'12 & Ding-**M** '13]
  - Current research

No Free Lunch Theorem

It is not possible to guarantee *any* utility in addition to privacy, *without making assumptions about*

- *the data generating distribution* [Kifer-Machanavajjhala SIGMOD '11]
- *the background knowledge available to an adversary* [Dwork-Naor JPC '10]

Discriminant: Sliver of Utility

- Does an algorithm A provide any utility?

  $w(k, A) > c$ if there are k inputs $\{D_1, \ldots, D_k\}$ such that $A(D_i)$ give different outputs with probability $> c$.

- Example:
  If A can distinguish between tables of size <100 and size >1000000000, then $w(2,A) = 1$.

Duke UNIVERSITY

Discriminant: Sliver of Utility

Theorem: The discriminant of Laplace mechanism is 1.

Proof:

- Let Di = a database with n records and *n·i/k* cancer patients
- Let Si = the range [*n·i*/k – n/3k, *n·i*/k + n/3k]. All Si are disjoint

- Let M be the laplace mechanism on the query "how many cancer patients are there".
- Pr(M(Di) ε Si) = Pr(Noise < n/3k) > 1 – $e^{-n/3k\varepsilon}$ = 1 – δ

- Hence, discriminant w(k,M) > 1- δ
- As n tends to infinity, discriminant tends to 1.

Duke UNIVERSITY

Discriminant: Sliver of Utility

- Does an algorithm *A* provide any utility?

  $w(k, A) > c$ if there are k inputs $\{D_1, \ldots, D_k\}$ such that $A(D_i)$ give different outputs with probability $> c$.

- If $w(k, A)$ is close to 1
  - we *may* get some utility after using *A*.

- If $w(k, A)$ is close to 0
  - we *cannot* distinguish any *k* inputs – no utility.

Non-privacy

- D is randomly drawn from P_{data}.
- q is a sensitive query with k answers, s.t.,

 knows P_{data} but cannot guess value of q

- A is not private if:

 can guess q correctly based on P_{data} and A

No Free Lunch Theorem

- Let *A* be a privacy mechanism with $w(k,A) > 1- \varepsilon$

- Let *q* be a sensitive query with *k* possible outcomes.

- There exists a data generating distribution P_{data}, s.t.
  - $q(D)$ is uniformly distributed, but
  -  wins with probability greater than $1-\varepsilon$

Outline

- Background: Domain independent privacy definitions
- No Free Lunch in Data Privacy [Kifer-**M** SIGMOD '11]
- Correlations: A case for domain specific privacy definitions [Kifer-**M** SIGMOD '11]
- Pufferfish Privacy Framework [Kifer-**M** PODS'12]
- Defining Privacy for Correlated Data [Kifer-**M** PODS'12 & Ding-**M** '13]
  - Current research

Correlations & Differential Privacy

- When an adversary knows that individuals in a table are correlated, then (s)he can learn sensitive information about individuals even from the output of a differentially private mechanism.

- Example 1: Contingency tables with pre-released exact counts

- Example 2: Social Networks

Contingency tables

Each tuple takes k=4 different values

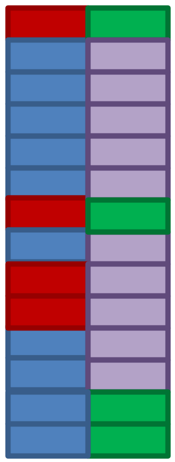

D

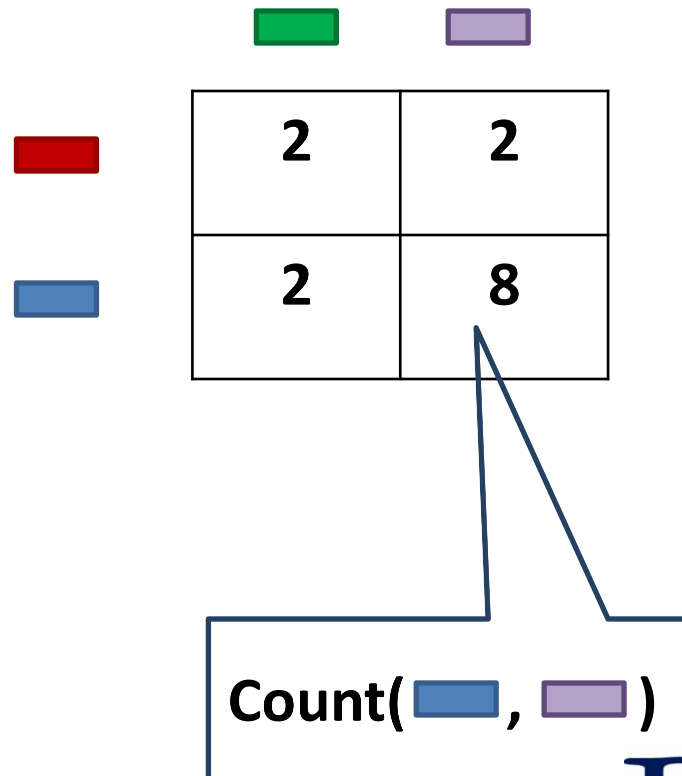

Duke UNIVERSITY

Contingency tables

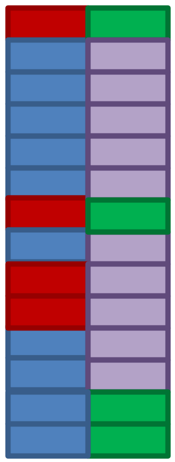

D

Want to release counts privately

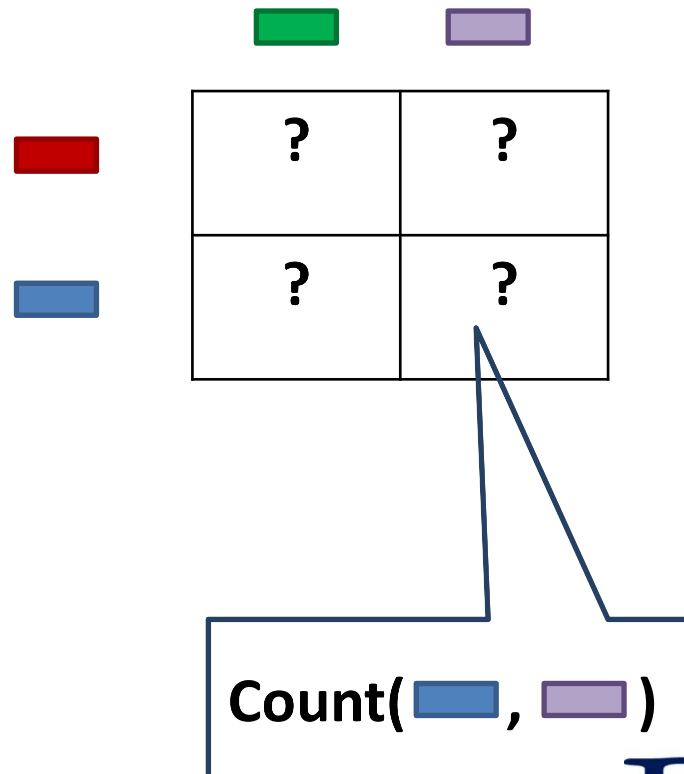

Laplace Mechanism

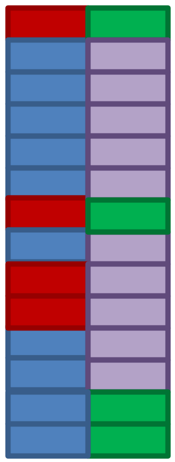

D

	$2 + Lap(1/\varepsilon)$	$2 + Lap(1/\varepsilon)$
	$2 + Lap(1/\varepsilon)$	$\underbrace{8 + Lap(1/\varepsilon)}$

Mean : 8
Variance : $2/\varepsilon^2$

Guarantees differential privacy.

Marginal counts

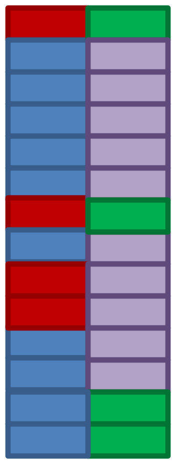

D

2 + Lap(1/ε)	2 + Lap(1/ε)	4
2 + Lap(1/ε)	8 + Lap(1/ε)	10
4	10	

Auxiliary marginals published for following reasons:

1. **Legal**: 2002 Supreme Court case Utah v. Evans
2. **Contractual**: Advertisers must know exact demographics at coarse granularities

Does Laplace mechanism still guarantee privacy?

Duke UNIVERSITY

Marginal counts

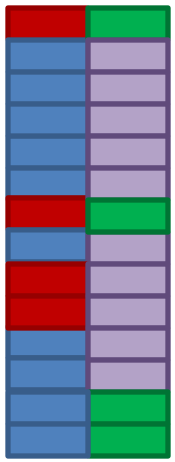

D

$2 + Lap(1/\varepsilon)$	$2 + Lap(1/\varepsilon)$	4
$2 + Lap(1/\varepsilon)$	$8 + Lap(1/\varepsilon)$	10
4	10	

Count (,) = $8 + Lap(1/\varepsilon)$

Count (,) = $8 - Lap(1/\varepsilon)$

Count (,) = $8 - Lap(1/\varepsilon)$

Count (,) = $8 + Lap(1/\varepsilon)$

Duke UNIVERSITY

Marginal counts

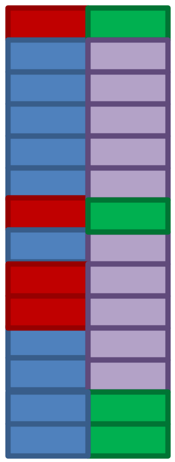

D

$2 + Lap(1/\varepsilon)$	$2 + Lap(1/\varepsilon)$	4
$2 + Lap(1/\varepsilon)$	$8 + Lap(1/\varepsilon)$	10
4	10	

Mean : 8

Variance : $2/ke^2$

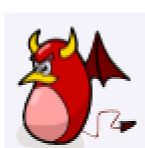

can reconstruct the table with high precision for large k

Reason for Privacy Breach

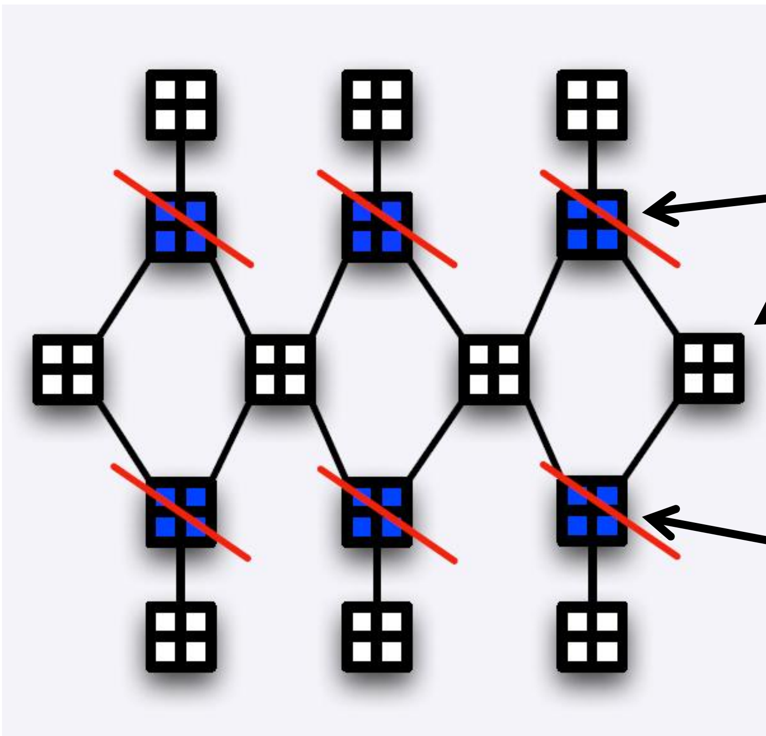

Space of all possible tables

- Pairs of tables that differ in one tuple
- 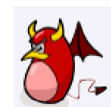 cannot distinguish them

Tables that do not satisfy background knowledge

Reason for Privacy Breach

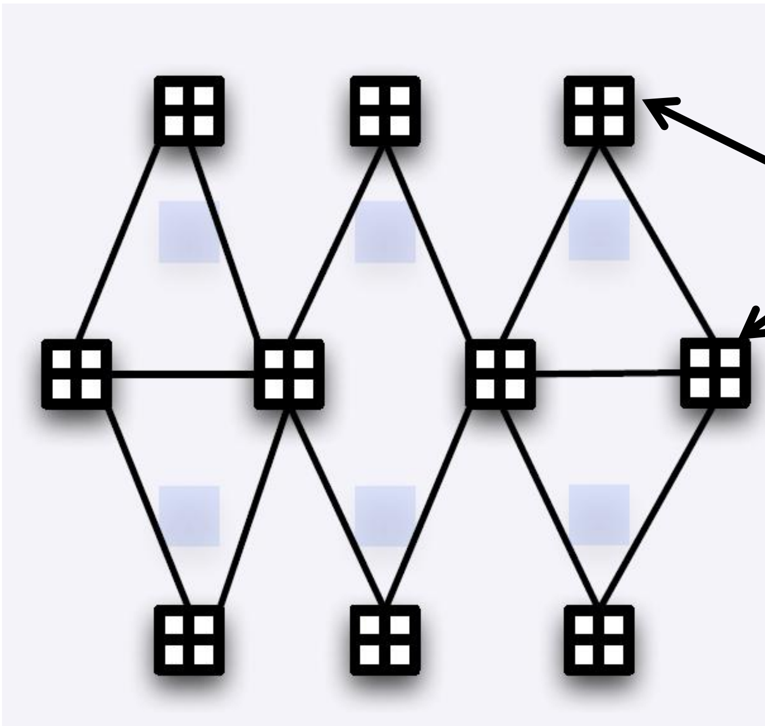

 can distinguish between every pair of these tables based on the output

Space of all possible tables

Duke UNIVERSITY

Correlations & Differential Privacy

- When an adversary knows that individuals in a table are correlated, then (s)he can learn sensitive information about individuals even from the output of a differentially private mechanism.

- Example 1: Contingency tables with pre-released exact counts

- Example 2: Social Networks

A count query in a social network

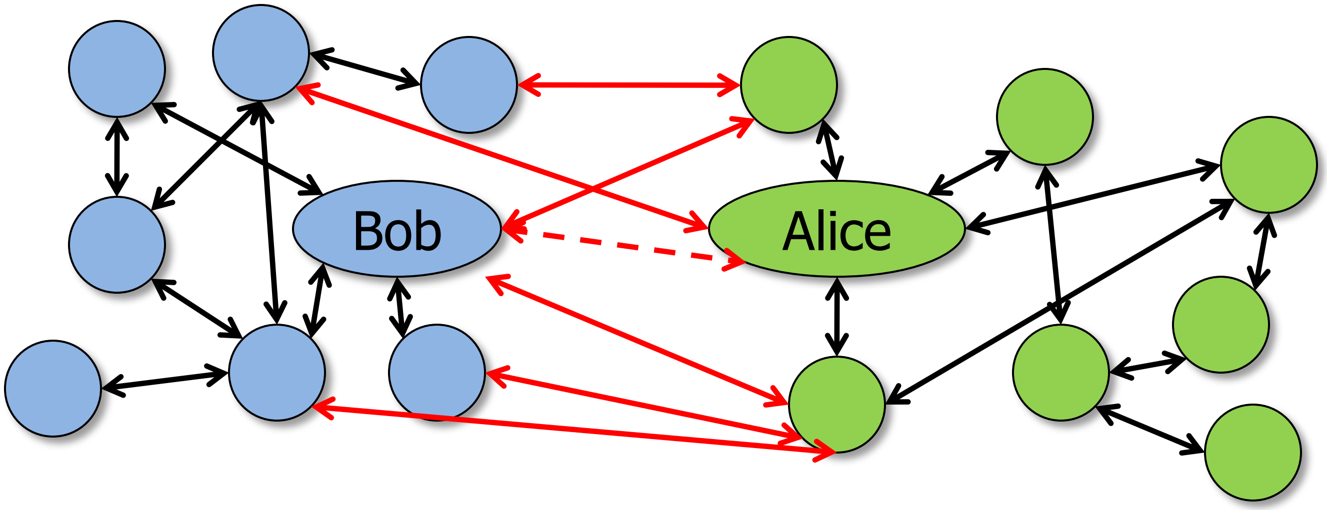

- Want to release the number of edges between **blue** and **green** communities.
- Should not disclose the presence/absence of Bob-Alice edge.

Adversary knows how social networks evolve

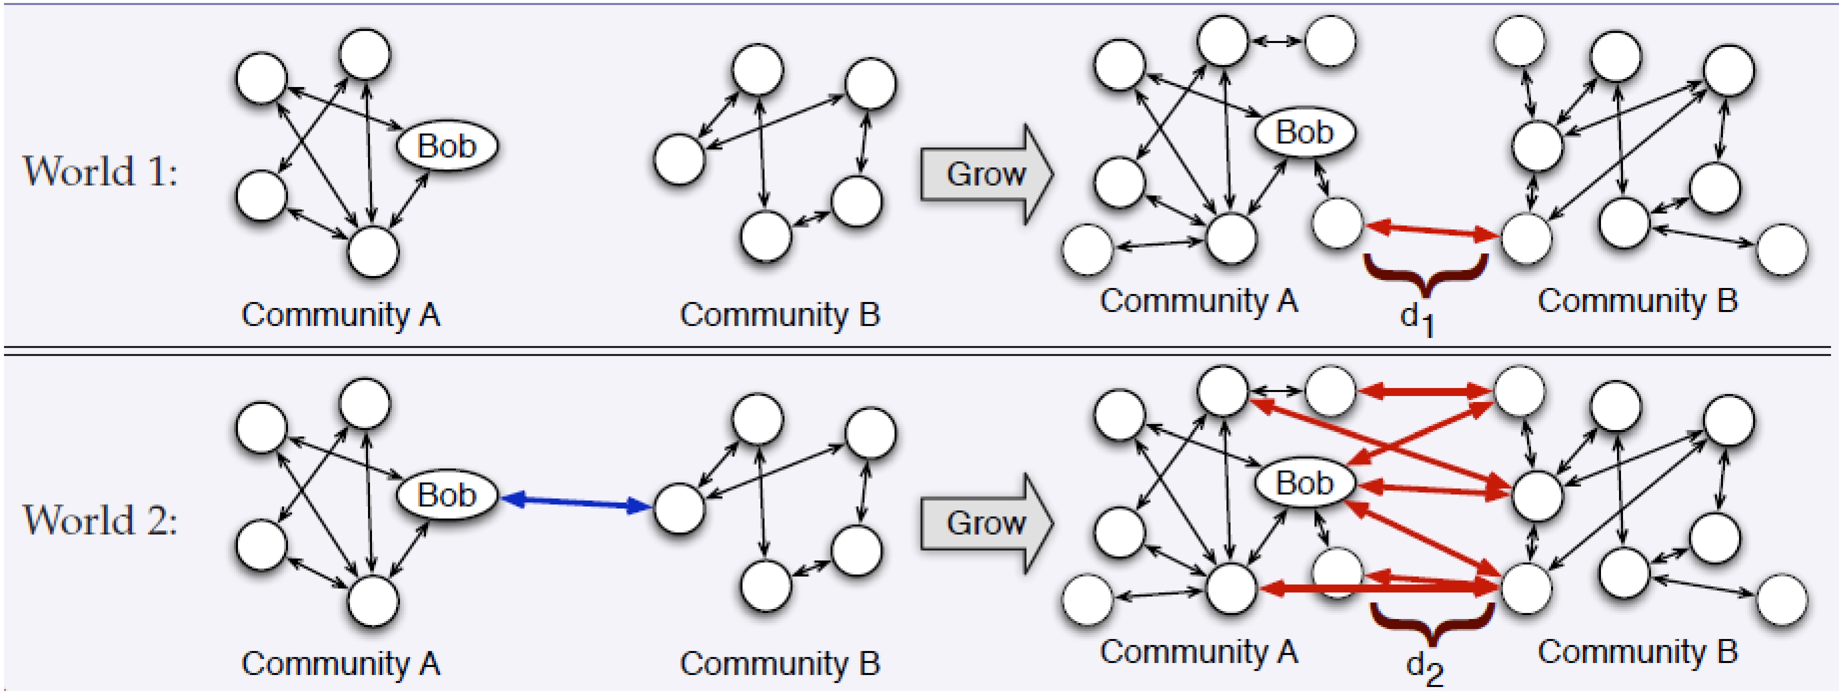

- Depending on the social network evolution model, *(d_2-d_1)* is *linear* or even *super-linear* in the size of the network.

Differential privacy fails to avoid breach

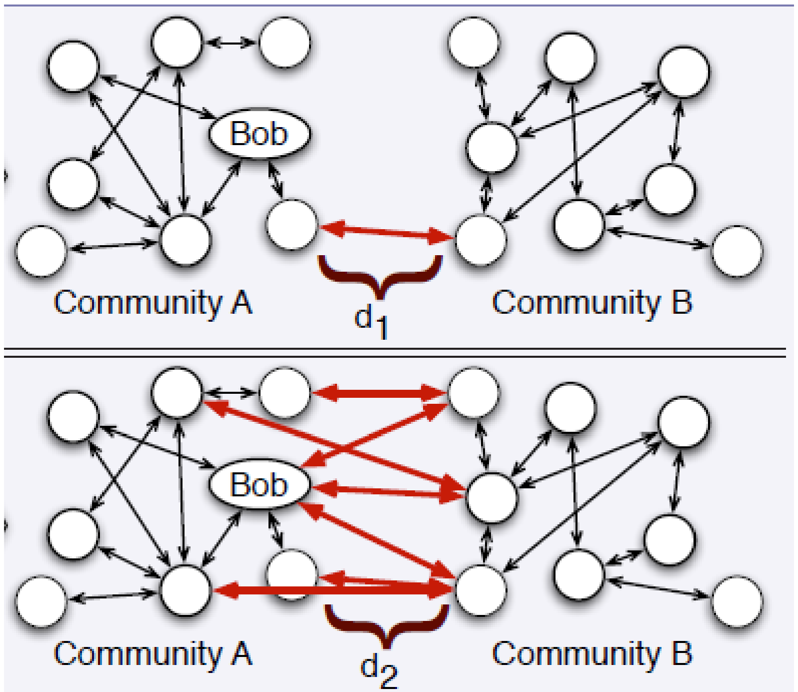

Output $(d_1 + \delta)$

$\delta \sim \text{Laplace}(1/\varepsilon)$

Output $(d_2 + \delta)$

Adversary can distinguish between the two worlds if $d_2 - d_1$ is large.

Outline

- Background: Domain independent privacy definitions
- No Free Lunch in Data Privacy [Kifer-**M** SIGMOD '11]
- Correlations: A case for domain-specific privacy definitions [Kifer-**M** SIGMOD '11]
- Pufferfish Privacy Framework [Kifer-**M** PODS'12]
- Defining Privacy for Correlated Data [Kifer-**M** PODS'12 & Ding-**M** '13]
  - Current research

Why we need domain specific privacy?

- For handling correlations
  - Prereleased marginals & Social networks [Kifer-**M** SIGMOD '11]
- Utility driven applications
  - For some applications existing privacy definitions do not provide sufficient utility [**M** et al PVLDB '11]
- Personalized privacy & aggregate secrets [Kifer-**M** PODS '12]

Qn: How to design principled privacy definitions customized to such scenarios?

Pufferfish Framework

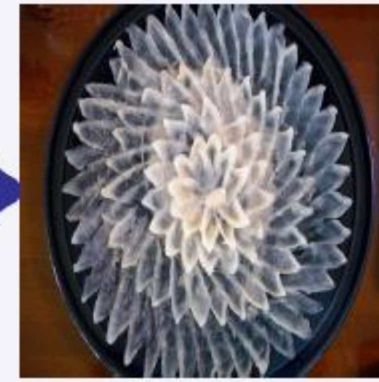

- Pufferfish (data):
  - contains tetrodotoxin (sensitive information).
- Toxin is everywhere:
  - Liver
  - Intestines
  - Skin / Muscles
- Removing all toxin = removing fish

- Chef (algorithm):
  - Processes the fish.
- Certification and license (privacy definition):
  - Rules chef must follow / restrictions on algorithm
  - Guarantees output is (relatively) safe.

- Fugu (sanitized data):
  - Tasty (high utility)
  - Minimal toxins
  - Minimal leakage of sensitive information

Pufferfish Semantics

- What is being kept secret?
- Who are the adversaries?
- How is information disclosure bounded?

Sensitive Information

- **Secrets**: S be a set of potentially sensitive statements
  - "individual j's record is in the data, and j has Cancer"
  - "individual j's record is not in the data"

- **Discriminative Pairs**: Spairs is a subset of SxS. Mutually exclusive pairs of secrets.
  - ("Bob is in the table", "Bob is not in the table")
  - ("Bob has cancer", "Bob has diabetes")

Adversaries

- An adversary can be completely characterized by his/her prior information about the data
  - We do not assume computational limits

- **Data Evolution Scenarios**: set of all probability distributions that could have generated the data.
  - *No assumptions*: All probability distributions over data instances are possible.

  - *I.I.D.*: Set of all f such that: $P(data = \{r_1, r_2, \ldots, r_k\}) = f(r_1) \times f(r_2) \times \ldots \times f(r_k)$

Duke UNIVERSITY

Information Disclosure

- Mechanism M satisfies ε-Pufferfish(S, Spairs, D), if for every
  - w ε Range(M),
  - (s_i, s_j) ε Spairs
  - Θ ε D, such that $P(s_i \mid \theta) \neq 0$, $P(s_j \mid \theta) \neq 0$

$$P(M(\text{data}) = w \mid s_i, \theta) \leq e^{\varepsilon} P(M(\text{data}) = w \mid s_j, \theta)$$

Pufferfish Semantic Guarantee

$$e^{-\epsilon} \leq \frac{P(s_i \mid \mathfrak{M}(\mathfrak{D}ata) = \omega, \theta)}{P(s_j \mid \mathfrak{M}(\mathfrak{D}ata) = \omega, \theta)} \Big/ \frac{P(s_i \mid \theta)}{P(s_j \mid \theta)} \leq e^{\epsilon}$$

Posterior odds of s_i vs s_j

Prior odds of s_i vs s_j

Duke UNIVERSITY

Assumptionless Privacy

- Suppose we want to make protect against any adversary
  - No assumptions about adversary's background knowledge

- Spairs:
  - "record j is in the table with value x" vs "record j is not in the table"
- Data Evolution: All probability distributions over data instances are possible.

A mechanism satisfies ε-Assumptionless Privacy
if and only if
for every pair of database D1, D2, and every output w
$P(M(D1) = w) \leq e^{\varepsilon} P(M(D2) = w)$

Assumptionless Privacy

A mechanism satisfies ε-Assumptionless Privacy
if and only if
for every pair of database D1, D2, and every output w
$P(M(D1) = w) \leq e^{\varepsilon} P(M(D2) = w)$

- Suppose we want to compute the number of individuals having cancer.
  - D1: all individuals have cancer
  - D2: no individual has cancer
  - For assumptionless privacy, the output w should not be too different if the input was D1 or D2
  - Therefore, need O(N) noise (where N = size of the input database).
  - Hence, not much utility.

Applying Pufferfish to Differential Privacy

- Spairs:
  - "record j is in the table" vs "record j is not in the table"
  - "record j is in the table with value x" vs "record j is not in the table"

- Data evolution:
  - Probability record j is in the table: π_j
  - Probability distribution over values of record j: f_j
  - For all $\theta = [f_1, f_2, f_3, \ldots, f_k, \pi_1, \pi_2, \ldots, \pi_k]$

  - $P[\text{Data} = D \mid \theta] = \prod_{r_j \text{ not in } D} (1-\pi_j) \times \prod_{r_j \text{ in } D} \pi_j \times f_j(r_j)$

Applying Pufferfish to Differential Privacy

- Spairs:
  - "record j is in the table" vs "record j is not in the table"
  - "record j is in the table with value x" vs "record j is not in the table"
- Data evolution:
  - For all $\theta = [f_1, f_2, f_3, \ldots, f_k, \pi_1, \pi_2, \ldots, \pi_k]$
  - $P[\text{Data} = D \mid \theta] = \prod_{r_j \text{ not in } D} (1-\pi_j) \times \prod_{r_j \text{ in } D} \pi_j \times f_j(r_j)$

A mechanism M satisfies differential privacy

if and only if

it satisfies Pufferfish instantiated using Spairs and {θ}

(as defined above)

Differential Privacy

- Sensitive information:
  All pairs of secrets “individual j is in the table with value x” vs “individual j is not in the table”

- Adversary:
  Adversaries who believe the data is generated using *any* probability distribution that is *independent* across individuals

- Disclosure:
  ratio of the prior and posterior odds of the adversary is bounded by e^{ε}

Characterizing "good" privacy definition

- We can derive conditions under which a privacy definition resists attacks.

- For instance, any privacy definition that can be phrased as follows **composes** with itself.

$$\forall w, P(M(D_1) = w) \leq e^{\varepsilon} P(M(D_2) = w)$$
$$\forall (D_1, D_2) \in \mathfrak{D} \subseteq 2^I$$

where *I* is the set of all tables.

Summary of Pufferfish

- A semantic approach to defining privacy
  - Enumerates the information that is secret and the set of adversaries.
  - Bounds the odds ratio of pairs of mutually exclusive secrets

- Helps understand assumptions under which privacy is guaranteed

- Provides a common framework to develop theory of privacy definitions
  - General sufficient conditions for composition of privacy (see paper)

Next Class

- Application of Pufferfish to Correlated Data

- Relaxations of differential privacy
  - E-Privacy
  - Crowd-blending privacy

References

[**M** et al PVLDB’11]

A. Machanavajjhala, A. Korolova, A. Das Sarma, “*Personalized Social Recommendations – Accurate or Private?*”, PVLDB 4(7) 2011

[Kifer-**M** SIGMOD’11]

D. Kifer, A. Machanavajjhala, “*No Free Lunch in Data Privacy*”, SIGMOD 2011

[Kifer-**M** PODS’12]

D. Kifer, A. Machanavajjhala, “*A Rigorous and Customizable Framework for Privacy*”, PODS 2012

[Ding-**M** ‘13]

B. Ding, A. Machanavajjhala, “Induced Neighbors Privacy(*Work in progress)*”, 2012