Topics in Data Science, Data Protection
CS 4364

Spring 2020
L. Longpré

Class notes 3/10/2020

Differential Privacy

Differential privacy: database disclosure reveals essentially the same information about an individual whether or not the individual is part of the database.
Definition to classify algorithms that release statistics about data. The goal is for an observer seeing the algorithm's output not being able to tell if a particular individual's information was used by the algorithm. In other words, the result of the algorithm should not be differentiable if an individual is included or not.
Definition:
A randomized function K gives ε-differential privacy if for all data sets D_1 and D_2 differing on at most one element,

$$\Pr[K(D_1) \in S] \leq \exp(\varepsilon) \cdot \Pr[K(D_2) \in S]$$

Model of computation

Trusted and trustworthy database curator. offline model: sanitized database. online model: adaptive queries.

Examples

"Smoking causes cancer" example comparison with cryptosystem definitions what is auxiliary information?

Formalizing differential privacy

Statistical information about embarassing behavior:

1. Flip a coin
2. If tails, respond truthfully
3. If heads, flip a second coin and respond yes if heads and no if tails.

Exercise: From survey result, compute the relation between the number of yes answers and the number of actual embarassing behavior.

Randomization is essential.

Randomized algorithm: From domain A to discrete set B. Each element of A induces a probability distribution on B.

Database can be seen as a histogram. Example: finite set X of possible records, database is 1,0,2,0,0,2,1,0,0,...

$$||x||_1 = \sum_{i=1}^{|X|} |x_i|$$

The ℓ_1 distance between two databases x and y is $||x - y||_1$.

(Other way to formalize databases: multisets of rows, or ordered lists of rows.)
Let D a domain of possible databases

Differential Privacy definition

A randomized algorithm M with domain D is (ϵ, δ)-differentially private if for all $S \subseteq \mathrm{Range}(M)$ and for all $x, y \in D$ such that $||x - y||_1 \leq 1$:

$$\Pr[M(x) \in S] \leq \exp(\varepsilon)\Pr[M(y) \in S] + \delta$$

If $\delta = 0$, we say that M is ε-differentially private.

Properties

1. Protection against arbitrary risks
2. Neutralization of linkage attacks
3. Quantification of privacy loss
4. Composition
5. Group privacy
6. Post-Processing