Differential Privacy

Differential Privacy: A Survey of Results

kyunghee KIM

Department of Statistics

Sungkyunkwan University

May 2, 2022

Overview

1. Privacy attack and protection
2. Differential Privacy

Privacy Attack

Linkage(differencing) attack

- Attack on obtaining information about a particular individual by aggregating multiple auxiliary information about a particular individual.
- If you add a lot of naive queries, privacy can be attacked.

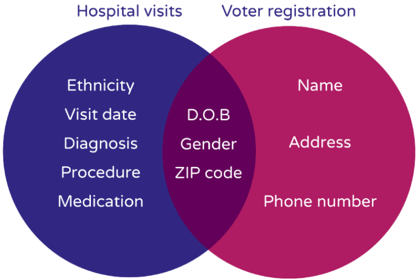

Solution for privacy attack

There are lots of solutions to prevent privacy attack,

- k-anonimity

  Def. Manipulate the data so that the number of elements corresponding to the query response is not one.

  PROB. Risk varies depending on the auxiliary information, and there is always an auxiliary info that can obtain information through k-anonymity.

- Not answering

  Def. If there is only one element in a query's response, do not respond.

  PROB. It is equivalent to admitting that the query is related to critical information

  PROB. computational problem.

Solution for privacy attack(Cont'd)

Better solution than existing cryptosystem.

1. Synthetic data
2. Homomorphic Encryption (HE)
3. Differential Privacy (DP)

DP Definition

Differencial Privacy (DP)

- A concept that prevents advertisers from finding information about a particular element through queries when there are multiple elements in the DB.
- DP is a definition not an algorithm!

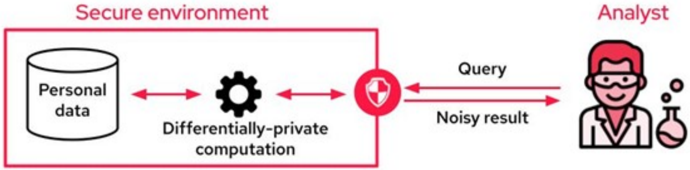

DP Motivation

DP motivation

- Regardless of background(auxiliary) information, there is a risk as much as ϵ, and the rest is protected. Now, it is possible to define a mathematical and clear standard for the level of privacy protection.

DP Advantages

- Ensure privacy against any threat
- Numericalize privacy loss
- Once DP is guaranteed, it is guaranteed through any post-processing.

How does DP work?

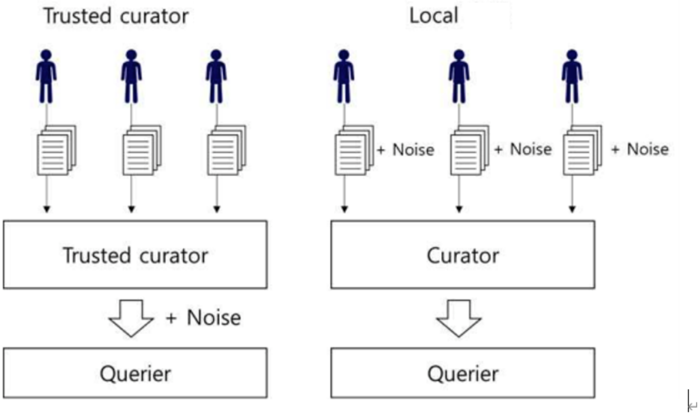

How does DP work?

- Add noise to the true value and assumed to be a real response.
- Ensure that the query's response value does not change significantly as a particular element is included in the DB. Avoid inference for specific elements.

Little change in DB $\rightarrow$ A big difference in results $\rightarrow$ Stability issues

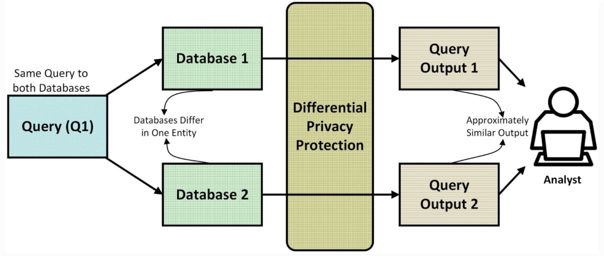

DP Framework

Definition 1. Randomized Algorithm.

M: randomized algorithm

On input $a \in A$, the algorithm M outputs $M(a) = b$ with probability $(M(a))_b$ for each $b \in B$.

Definition 2.ϵ-Differential Privacy. for all $S \in Range(M)$, $||d_1 - d_2||_1 \leq 1$:

$$\frac{P[M(D1) \in S]}{P[M(D2) \in S]} \leq e^{\epsilon}$$

M can be average, sum, ML, count, $\cdots$. ϵ can be 0.01, 0.1, ln2, $\cdots$, $e^{0.01} \approx 1.01$.

DP Framework

Sensitivity type in ϵ-Differential Privacy

- Global Sensitivity

  Maximum value of change due to insertion or deletion of a particular individual

$$\Delta f = max_{D1,D2}||M(D1) - M(D2)||_1$$

  Vulnerable to Outlier and likely to degrade overall performance

- Local Sensitivity

$$\Delta f = max_{D1}||M(D1) - M(D2)||_1$$

DP Framework(Cont'd)

Definition 3. $(\epsilon,\ \delta)$ Differential Privacy.

$$P[M(D1) \in S] \leq e^{\epsilon} P[M(D2) \in S] + \delta.$$

Definition 4. $(k\epsilon,\ 0)$ Differential Privacy. $||d_1 - d_2||_1 \leq k$:

$$P[M(D1) \in S] \leq e^{k\epsilon} P[M(D2) \in S]$$

If δ gets smaller, credible gets bigger.

DP Framework(Cont'd)

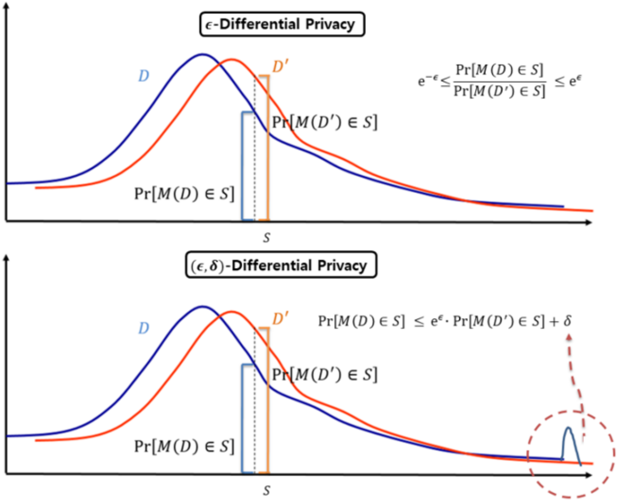

Mechanism to add noise

For ϵ-Differential Privacy,

- Laplace mechanism

$$L(D) = f(D) + Z, Z \sim Lap(0, b), where P(Z|\mu, b) = \frac{1}{2b} e^{-\frac{|z-\mu|}{b}}$$

  z is proportional to $e^{-\epsilon|z|/\Delta f}$ and $b = \Delta f/\epsilon$. It is not applicable to categorical data.

- Exponential mechanism

  It is applicable to categorical data, but takes much longer time.

- Geometric mechanism

  Discrete variant for laplace mechanism.

data type

Data type

- histogram query

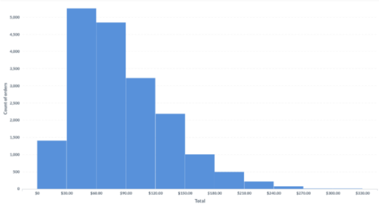

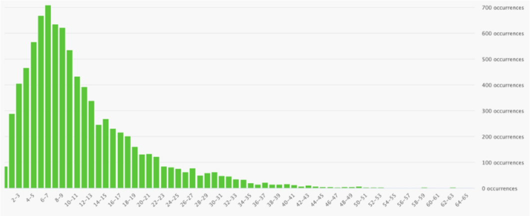

- Partitioning

  Find optimal subdivision. e.g., Tree base

Composibility

More queries result in additional information disclosure, so it is managed with a privacy budget. Budget is determined by the number of queries and ϵ.

- Sequential composition

  Satisfy $\sum_i \epsilon_i$-differential privacy for successive queries

- Parallel composition

  Satisfy Max ϵ_i-differential privacy if a continuous query occurs to disjoint domain Di.

Application

- Data mining and ML

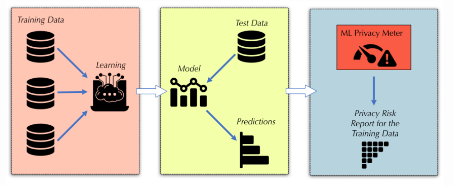

- Biometrics

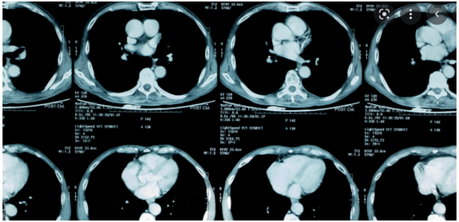

DP Future work(Cont'd)

- Side channel attack
- How much noise?
- Utility?
- Cost?

Side channel attack

Side channel attack type

- Timing attack
- State attack
- Privacy budget attack

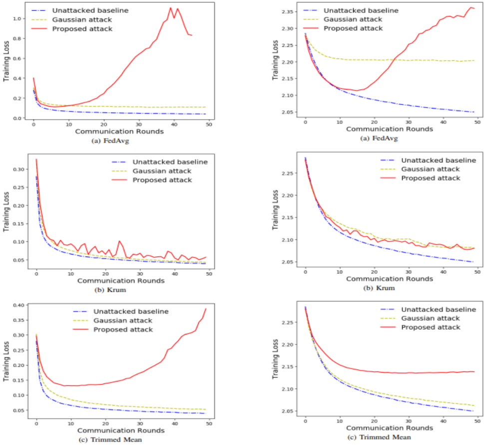

DP Future work(Cont'd)

How much noise?

1. We cannot tell.
2. We can consider sensitivity.

Sensitivity: if output changes a lot, it needs to set a lot of noise to make DB indistinguishable.

DP Future work(cont'd)

Utility?

1. Data has to be huge (e.g., census)
2. Give incentive(ϵ) almost 10.
3. Not directly using data, utility loss occured.

Cost?

1. For each data, needed to make DP framework.
2. Needs expert to make framework.
3. To fix this problem, use correlated sensitivity has proposed.

RFP 16

1. Demonstration application achieves two goals: privacy and data availability

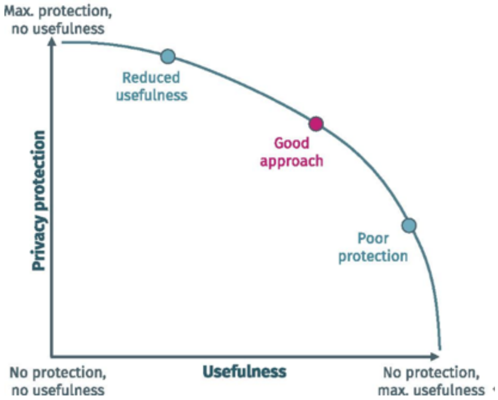

The End