Overall Framework:

Motivation

- Two separate entities namely Alice and Bob are working together to build a recommendation system for a specific group of people. Alice has the dataset of the people's specific activities for Bob to build the recommendation system which also includes private and sensitive data. Responsibilities of Alice's and Bob's are maintaining privacy and build the system accordingly. So while transmitting the data to Bob, Alice intentionally perturbed the data in such a way so that Bob can not infer some sensitive information. On the other hand, a minimum amount of correlation needed for Bob in order to infer required information to come up with a viable recommendation system. This paper has developed a statistical inference framework to minimize this tension between Alice and Bob.
- Though at the first glance, it looked like a *secrecy* problem in dataset but [3] outlined a significant difference between *secrecy* and *privacy* problem. In the privacy problem, revealing data provides required information for utility and possible loss of privacy at the same time. For example, in a credit card transaction environment receiver learns some public information (e.g. gender, weight, occupation), which is allowed and necessary for the correct transactions. But concurrently receiver is also capable of learning other private information such as income, health conditions which can enable receiver to build an unauthorized recommendation system. Thus, an optimal privacy preserving design should not decrease the *uncertainty* (entropy) of predicting cancer after revealing an individual's physical attributes.

Applications of the Proposed Framework [2]

Privacy Preserving Queries of a Database We have a database from n users, where a specific entry of each user, i.e., a vector $S \in \mathbb{R}^n$ consists of 1 discrete entry from all n users, considered *private*. The goal of a privacy preserving mapping is to present a query output Y such that S is hidden (the estimation capacity of Bob is minimized significantly), while still maintaining the required utility for the Bob in terms of a distortion constraint. Mathematically, given a function $f : \mathcal{S} \to \mathcal{X}$, with $x = f(S)$, a good privacy preserving design will provide a posterior probability $p_{X|Y}(x|y)$, such that, $p_{X|Y}(x|y) \propto |f^{-1}(x)|$.

Hiding Dataset Features Another application of the proposed framework can be *blurring* one or more particular entries from all users, resulting in difficulties in estimation from the output query, whereas maintaining the desired distortion based on the utility of the receiver.

Key Questions Being Answered

- A statistical inference framework has been proposed in order to understand the security threat an user faced by the utility provider, who is also a curious adversary. So the goal is to minimize the threat by limiting the amount of private (crucial) information received by the adversary. In order to achieve that, user must design a distortion (termed as **privacy preserving mapping**) between private and transmitted data.
- On the other hand, utility provider should not lack too much important information that it can not provide the necessary service. Thus, the optimization problem involves in designing the privacy preserving mapping such that user's privacy is maintained and also analyst can extract only the necessary information to deliver the service. Modeling this balances has been termed as **privacy-utility-trade-off** modeling. Thus the privacy preserving mapping should be such that distortion should not exceed a certain limit.
- To this end, authors have introduced two metrics, i.e., *average information leakage*, *maximum information leakage*. Intuitively, information leakage indicates how much "crucial" information curious adversary can gain by seeing the available output or by it's own capability without seeing output and thus the goal is to minimize this information leakage. Moreover, they showed that their minimization of information leakage can be framed as rate distortion theory by choosing the *log-loss* cost function and also can be cast as a convex problem.
- Finally they argued that so far celebrated *differential privacy* doesn't necessarily provide any privacy guarantee in terms of maximum and average information leakage. Thus a new metric *information privacy* has been introduced which captures both the concepts of privacy and differential privacy.

1 Comparison between differential privacy and mutual information privacy

ϵ-differential privacy: A privacy preserving mapping $p_{U|S}(.|.)$ provides ϵ-differential privacy if for all inputs s_1 and s_2 differing in at most one entry and all $B \subseteq \mathcal{U}$,

$$\Pr(U \in B|S = s_1) \leq \exp(\epsilon) \times \Pr(U \in B|S = s_2) \tag{1}$$

ϵ-information privacy: A privacy preserving mapping $p_{U|S}(.|.)$ provides ϵ-information privacy if $\forall s \subseteq \mathcal{S}^n$,

$$\exp(-\epsilon) \leq \frac{p_{S|U}(s|u)}{P_S(s)} \leq \exp(\epsilon) \ \forall u \in \mathcal{U} : p_U(u) > 0 \tag{2}$$

ϵ-information privacy implies directly 2ϵ- differential privacy

$$\frac{\Pr(U \in B|S = s_1)}{\Pr(U \in B|S = s_2)} = \frac{\Pr(S = s_1|U \in B)\Pr(S = s_2)}{\Pr(S = s_2|U \in B)\Pr(S = s_1)} \leq \exp(2\epsilon) \tag{3}$$

and maximum information leakage of at most $\epsilon/\ln 2$ bits

$$H(S) - H(S|U = u) = \sum_{s \in \mathcal{S}^n} p_{S|U} p_U(u) \log \frac{p_{S|U}(s|u)}{p_S(s)} \tag{4}$$

$$\leq \sum_{s \in \mathcal{S}^n} p_{S|U} p_U(u) \frac{\epsilon}{\ln 2} \tag{5}$$

$$= \frac{\epsilon}{\ln 2} \tag{6}$$

In **Theorem 4**, authors state that $\forall \epsilon > 0$ and $\forall \delta \geq 0$, there exists an $n \in \mathbb{Z}_+$, sets $\mathcal{S}^n$ and $\mathcal{U}$, a prior $p_S(.)$ over $\mathcal{S}^n$ and a corresponding privacy preserving mapping $p_{U|S}(.|.)$ which leaks δ bits on average despite being ϵ-differential private.

Consequently, authors have constructed an example of ϵ-differential private system, where there is a possibility of leaking an arbitrarily large amount of information on average. Following the footsteps of [1] authors perturbed query output with Laplacian noise N in order to achieve ϵ-differential privacy. Furthermore, after observing U, for a given ϵ and constraining the amount of information leakage, receiver (Bob) does a maximum a posteriori estimation of Y. Corresponding equations are as follows,

$$U = Y + N, \;\; N \sim \text{Lap}(1/\epsilon) \tag{7}$$

$$p_N(r; \epsilon) = \frac{\epsilon}{2} \exp(-|r|\epsilon) \tag{8}$$

$$p_Y(y) = \begin{cases} \frac{1}{1+n/K} & \text{if } y \mod k = 0 \\ 0 & \text{otherwise} \end{cases} \tag{9}$$

To this end, authors made an interesting observation that, if E is a binary random variable which indicates the event that receiver makes an incorrect estimation of Y given U. Then $I(Y;U)$ can be made arbitrarily large by choosing suitable parameters. Associated equations are eq.(10 - 13).

$$I(Y;U) \geq I(E,Y;U) - 1 \tag{10}$$

$$\geq I(Y;U|E) - 1 \tag{11}$$

$$\geq \Pr\{E = 0\} \geq I(Y;U|E = 0) - 1 \tag{12}$$

$$= (1 - e^{-\frac{k\epsilon}{2}}) \log(1 + \frac{n}{k}) - 1 \tag{13}$$

Thus we can comment on comparison between differential privacy and information privacy by the following observations

- Theorem 4 and associated counterexample shows that a large amount of information can be leaked in the ϵ-differential privacy design.
- Furthermore, even differential privacy can not perform better in terms of average information leakage when data samples is sufficiently large.
- On the other hand, one apparent benefit of differential privacy is it's independence of prior distribution.

Comments on the Paper [3] has taken care some important issues which has not been addressed in this paper

- The statistical assumptions on the data that allow information-theoretic analysis.
- Different features of a dataset has different level of *information* content. During perturbation of dataset of how different level of importance can be exploited is not clear.
- It is unclear how proposed added Laplace noise would be helpful in perturbation when data is categorical. For example, in various cases, service provider need to have the precise last few digits of one's credit card number. In those cases, how additive noise will keep output precise is not clear.
- [4] pointed out that external correlation (with variables not in database but publicly accessible) play a very crucial role in privacy context. How those correlation play their role in privacy-utility trade-off is not clear.

References

[1] C Dwork. "Differential privacy, in automata, languages and programming". In: *ser. Lecture Notes in Computer Scienc* 4052 (2006), p. 112.

[2] Salman Salamatian et al. "Privacy-Utility Tradeoff and Privacy Funnel". In: (2020).

[3] Lalitha Sankar, S Raj Rajagopalan, and H Vincent Poor. "Utility-privacy tradeoffs in databases: An information-theoretic approach". In: *IEEE Transactions on Information Forensics and Security* 8.6 (2013), pp. 838–852.

[4] Latanya Sweeney. "k-anonymity: A model for protecting privacy". In: *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10.05 (2002), pp. 557–570.