Limiting Privacy Breaches in Differential Privacy

OUYANG Jia
Department of Computer Science
Sun Yat-Set University
Guangzhou, China
ouyangjia1@163.com

YIN Jian*
Department of Computer Science
Sun Yat-Set University
Guangzhou, China
issjyin@mail.sysu.edu.cn

LIU Shao-Peng
Department of Computer Science
Sun Yat-Set University
Guangzhou, China
l-shaopeng@live.cn

***Abstract*—In recently years, privacy-preserving data mining has become more import and attractedmore attention from data mining community. Among the existing privacy preserving models, ε -differential privacy provides the strongest privacy guarantees and has no assumption about the adversary's background information and compute ability. However, howto set ε to satisfy privacy is still an open problem. In this paper, we propose a tactic, named LPB (Limiting Privacy Breaches), to set the privacy parameter intuitively. LPB ensures that, if the prior belief about individual is bounded by some threshold, the posterior belief, after given the published randomized result, is no more than another threshold.**

Keywords-component; differential privacy;privacy breaches; privacy-preserving data mining

I. INTRODUCTION

Thegreat progress in computer science, networking,and storage technologies is resulting in an unprecedented amount of digitization of information.In concert with this explosive growth of data, data mining is gaining more attention from diverse areas such as healthcare, social network analysis, online search, and so on. However, concerns about privacy of person information have emerged globally.The current status in data mining research isdevising techniques that incorporate security and privacy issues.

The concept of privacy-preserving data mining has been recently been proposed in response to the above concerns derived from data mining algorithms.We can build many data mining models without disclosing the input data[1, 2].Specifically, the implementation of privacy-preserving data mining model is mainly considering the following two aspects: (1) How to preventprivacy leak in the data mining process; (2) How to make the data or result more utility.Currently, the field of privacy preserving data mining research work has focused on how to design privacy principles and algorithms to achieve a better balance between these two aspects.

In the last few years,ε-differential privacy[3-6]has emerged asa new criterion that provides a more robust privacy guarantee, regardless of the adversary's background knowledge.The basic idea ofε-differential privacy is to add enough noise to theanalysis results performed on a sensitive dataset before the result published. Specifically, given any two databases that differ on exactly onerecord*r*, a data mining algorithm that satisfiesε-differentialprivacy will output randomized results with almost identical probability distributions. The randomized results ensure that it ishard for the adversary to identify any individualrecord in the dataset, even if the adversary knows the informationof all remaining records.

However, there are two problems inε-differential privacy. The firstis that εjust limits how much one individual can affect the resulting model, not how much information is revealed about an individual[7].This issue will lead individuals to be easily identified by the adversary after he observes the randomized result. The second is thatthere is no intuitively policy to set the privacy parameter ε . Paper [7] proposed a definitionρ-differential identifiability that provides the same guarantees as differential privacy, but the parameter ρbounds the probability estimate that an individual contributed to the resulting model.Therefore, the policy makers can set the ε based the ρ -differential identifiability. But ρ-differential identifiabilityassumes

*Corresponding author: E-mail: issjyin@mail.sysu.edu.cn

that the prior probability of an individual being in the input dataset is must the same for all individual.It means that ρ -differential identifiability depend on prior distribution. However, it is usual that the prior probability for all individual is not equivalent, or even ishard to know.

In this paper, we propose a tactic which is independent on prior probability, named **LPB** (**L**imiting **P**rivacy **B**reaches), to set the ε privacy parameter intuitively. LPBformalize the privacy of an individual through(ρ_1, ρ_2)-privacy proposed in paper[8]. Informally, ρ_1, ρ_2-privacy, where$0 < \rho_1 < \rho_2 < 1$, means that if the adversary's prior belief (beforeseeing the randomized result) that an individual is in the input dataset is no more than ρ_1, his posterior belief(after seeing the randomized result) that an individual istruly in the input datasetis no more than ρ_2. In other words, publishing the randomized result changes the belief of the adversary by at most $\rho_2 - \rho_1$. We will give a good intuitively policy to set the ε for satisfying a given (ρ_1, ρ_2) -privacy requirement. In a word, LPB ensures that, if the prior belief about individual is bounded by some thresholdρ_1, the posterior belief, given the published randomized result, is no more than another threshold ρ_2.

The remainingsections of this paper are organized as follows. Related works are presented in Section 2. We introduce some basic notation in Section 3. In Section 4, we show how to limiting privacy breaches in differential privacy. Experiments and conclusion appear in Section 5 and Section 6.

II. RELATED WORK

Many privacy models consider individual identity disclosure.The key idea of these modelsis to protect an individual to be uniquelyre-identified.Samarati and Sweeney[9] proposed *k*-anonymity model.*K*-anonymity ensures thateach record in the table cannot be distinguishedfrom other *k*-1 records in the same group. We call these records which are not distinguished an equivalence class. Though*k*-anonymity prevents record linkage, no constraints are put on sensitive attribute. Therefore, the adversary can use homogeneity attack and background knowledge attack to infer the sensitive attribute value of the victim.Paper[10] proposes the diversity principle, called*l*-diversity, to prevent attribute linkage.The*l*-diversity requires an equivalence classto contain at least*l*distinct sensitive values.

K-anonymity and *l*-diversityare good privacy models to prevent the adversary from uniquely identifying an individual's record. However, these models are targeted to theparticular attack model,and they assume the adversary's knowledge is limited. Due to both their vulnerability to adversaries'backgroundinformation and their deterministic nature, many types of privacy attack have been proposed on these approaches derived using these models, leading to privacy compromise.

In contrast, ε-differential privacy, a new privacy model from the field of statistical disclosure control,is first proposed by D.Work[3]Differential privacy provides strong privacy guarantees that do not depend on an adversary's background knowledge. However, therehaveno good policies to set the privacy parameterε. Paper [7] proposed adefinition of ρ -differential identifiabilityas an alternate formulation, parameterized by the probability of individual identification. It is the first workthat provides a parameterization based on the risk of identifying an individual for differential privacy, while letting policy makers set parameters based on the concept of differential identifiability.

Our work isan extension version ofρ-differential identifiability.We propose a tactic which is independent on prior probability, and formalize the privacy of an individual through (ρ_1, ρ_2)-privacy.

III. PRELIMINARIES

In this paper, we assume that a database D contains n tuples, i.e., $D = \{x_1, x_2, \cdots, x_n\}$. Each x_i in D takes a fixed value from the universe U. Two databases D and D' are neighbor databases if they are differing by at most one exactly record. i.e., $D = \{x_1, x_2, \cdots, x_i, \cdots x_n\}$, and $D' = \{x_1, x_2, \cdots, x_i', \cdots x_n\}$.

A. Differential Privacy

Differential privacy requires that the removal or addition of a single database record does not significantly affect the outcome of any analysis. Formally, differential privacy is defined as follow.

Definition 1. (ε -differential privacy)[3]:we say a randomized algorithm M providesε-differential privacy, if for any output O of M and for any two neighbor databases D and D', we have

$$\Pr\left[M(D)=O\right]\le e^{\varepsilon}\cdot\Pr\left[M(D')=O\right]$$

The above inequality means that the randomized algorithm M always gives similar results on neighbor database. This can prevent the adversary to infer any records from the output O of M .The parameter ε is called *privacy budget*and is set to control the level of privacy.The lower values of ε is, the stronger privacyit gives, because ε limitsfurther the influence of a record on the outcome of analgorithm.

Laplace mechanism [3]is astandard method to achieve differential privacy. Laplace mechanism is based on the global sensitivity of function mapping underlying datasets to reals.

Definition 2. (Global Sensitivity): For any function $f:D\to\square^{d}$, the sensitivity of f is

$$\Delta f=\max_{D,D'}\left\|f(D)-f(D')\right\|_1$$

Where D and D' are any two neighbor databases.

Laplace Mechanism. For the analysis whose outputs are real, a way to achieve differential privacy is to add Laplace noise to the true output of a function.Given the sensitivity of a function f the addition of noise drawnfrom a calibrated Laplace distribution with the probability density function $p\left(x|\lambda\right)=\frac{1}{2\lambda}e^{-|x|/\lambda}$ provides ε -differential privacy.

Theorem 1[3]**.**For any function $f:D\to\square^{d}$ over an arbitrary domain D, the mechanism M

$$M(D)=f(D)+\left(Laplace\left(\frac{\Delta f}{\varepsilon}\right)\right)^{d}$$

Provides ε -differential privacy.

B. Privacy Model

In this paper, we assume the adversary is an informed adversary. The adversary has complete information about the universe U, every value in U is known. The adversary knows all information about every tuple in D expect one, i.e., he knowsthe D'.

We want to add enough noisy to the data mining outcome to achieve differential privacy to prevent the adversary infer the unknown value. This goal is not achievable if the value in *U* is overly popular among the universe. Therefore, privacy protection should be relative to the "prior" of the value in *U*. To capture such a relative notion of privacy, we adapt the (ρ_1,ρ_2)-privacy originally proposed by paper[8].Consider two neighbor databases D and D'. The adversary don't known the value of x_n.Let's define X as a random variable for the description of the adversary's prior knowledge about $x\in U$. For the adversary, prior to randomization, each possible value $x\in U$ has probability $\Pr(X=x)$ .Now, suppose that therandomized algorithm M returns the randomized result $R=M_f(D)$. R is a random variable for the randomized result. Let's define $\Pr\left(X=x|R=r\right)$ as the posterior probability that, given the randomized result r, the true value of X is x. So, we can call $\Pr\left(X=x|R=r\right)$ is the adversary's ability to infer the value victim. We limit this ability by (ρ_1,ρ_2)-privacy.

Definition3. ((ρ_1,ρ_2)-privacy) [8]:We say that there is a upward (ρ_1,ρ_2)-privacy breach with respect to a value x for X, iffor all $r\in Range\left(M_f(D)\right)$,

$$\Pr\left[X=x\right]\le\rho_1\text{ and }\Pr\left[X=x|M_f(D)=r\right]\ge\rho_2$$

Here $0<\rho_1<\rho_2<1$ and $\Pr\left[M_f(D)=r\right]>0$, (ρ_1,ρ_2) -privacy holds if upward privacy breach if eliminated.

Essentially, (ρ_1,ρ_2)-privacy means that whenever the prior does not exceed ρ_1, the posterior must not exceed ρ_2. The value of ρ_1 and ρ_2 are set by policy maker.

To illustrate, we give an example where the query function mean M satisfyingε-differential privacy enables the adversary to guess the missing value with high

TABLE I. POSSIBLE VALUE

Possible Value	Possible Data Set	True Mean	NoisyAdded	Pr[$M_f(D_i)$]=5.401	Posterior Probability
1	1,2,3,1	7/4	3.291	0.0238	0.0751
2	1,2,3,2	8/4	3.401	0.0216	0.0682
3	1,2,3,3	9/4	2.791	0.0372	0.1174
5	1,2,3,5	11/4	2.291	0.0580	0.1831
10	**1,2,3,10**	**16/4**	**1.041**	**0.1762**	**0.5562**

posterior probability.Given $D=\{1,2,3,10\}$, the value in$D$ is drawn from $U=\{1,2,3,5,10\}$, the sensitivity of the query function mean on D is $16/4-7/4=9/4$. Assume the adversary already knows $\{1,2,3\}\subset D$, he wants to infer the missing value.Assume that $\varepsilon=2$ and the response result $r=5.041$. The missing value may be a value from U. The adversary computes the posterior probability $\Pr\left[X=x\middle|M_f(D)=r\right]$ as theTable I showed.

We take the possible value 10 for example to show the process of computing the posteriorprobability,

$$\begin{aligned}&\Pr\left[X=10\middle|M_f(D_i)=5.401\right]\\&=\frac{\Pr[X=10]\cdot\Pr\left[M_f(D_i)=5.401\middle|X=10\right]}{\sum_{i\in U}\Pr[X=i]\cdot\Pr\left[M_f(D_i)=5.401\right]}\\&=\frac{\Pr[X=10]\cdot\Pr\left[M_f(D_{10})=5.401\right]}{\sum_{i\in U}\Pr[X=i]\cdot\Pr\left[M_f(D_i)=5.401\right]}\end{aligned}$$

Where $D_i=\{D'\cup i\}$ .Now, we compute $\Pr\left[M_f(D_{10})=5.401\right]$ firstly.As $mean(D_{10})=4$, then $R-mean(D)=1.041$ is the noisy added to the real mean value. $\lambda=\frac{\Delta}{\varepsilon}=\frac{9}{8}=1.125$ yields

$$\Pr\left[M_f(D_{10})=5.401\right]=\frac{1}{2\cdot1.125}\cdot e^{-\frac{|1.041|}{1.125}}=0.1762$$

Assume that the value in U has the same prior probability $\rho_1=1/5$,

$$\Pr\left[X=10\middle|M_f(D_i)=5.401\right]=\frac{\Pr\left[M_f(D_{10})=5.401\right]}{\sum_{i\in U}\Pr\left[M_f(D_i)=5.401\right]}=0.5562$$

If $\rho_1=0.2,\rho_2=0.5$, then this mean M is not achieve (ρ_1,ρ_2)-privacy.

IV. LIMITING PRIVACY BREACHES IN DIFFERENTIAL PRIVACY

This section we describe how to **L**imit **P**rivacy **B**reaches (**LPB**) in differential privacy.The core idea is to set the privacy parameter ε to provide (ρ_1,ρ_2) -privacy. Laplace mechanism add noisy Y to every function response $R=f(D)+Y$, where Y is an i.i.d. random variable drawn from a Laplace distribution.To achieve (ρ_1,ρ_2)-privacy, paper[8] proposed an approach of*amplification.*

Definition 5[8]. A randomization operator M is at most γ -amplifying for all result $r\in M_f(D)$ if

$$\forall D_i,D_j:\frac{\Pr[D_i\to r]}{\Pr[D_j\to r]}\le\gamma \tag{1}$$

Where $\gamma\ge1$, $D_i=\{D'\cup i|i\in U\}$, $D_j=\{D'\cup j|j\in U\}$. If M return the result r, then any dataset $D=\{D'\cup i|i\in U\}$ may return this result. Therefore, we get the following equation according to Laplace mechanism.

$$\begin{aligned}\frac{p[D_i\to r]}{p[D_j\to r]}&=\frac{\frac{1}{2\lambda}\cdot e^{-\frac{|r-f(D_i)|}{\lambda}}}{\frac{1}{2\lambda}\cdot e^{-\frac{|r-f(D_j)|}{\lambda}}}\\&=e^{\frac{|r-f(D_j)|-|r-f(D_i)|}{\lambda}}\end{aligned} \tag{2}$$

Since $|f(D_i)-f(D_j)|\le\Delta$,simple application of thetriangle inequality yields

$$\frac{p[D_i\to r]}{p[D_j\to r]}\le e^{\frac{|f(D_i)-f(D_j)|}{\lambda}}\le e^{\frac{\Delta}{\lambda}}$$

The next theorem 2, also due to [8], relates the γ -amplifying condition to (ρ_1,ρ_2)-privacy.

Theorem 2[8]. (ρ_1,ρ_2) -privacy is guaranteed if the ε -differential privacysatisfies the γ -amplifying condition for all response values r of R, where $\gamma\le\frac{\rho_2}{\rho_1}\cdot\frac{1-\rho_1}{1-\rho_2}$.

The proof can be found in paper [8]. Let us derive the connection between λ and (ρ_1,ρ_2)-privacy from theorem2.Follow the equation(2), we get

$$\frac{p[D_i\to r]}{p[D_j\to r]}\le e^{\frac{\Delta}{\lambda}}\le\frac{\rho_2}{\rho_1}\cdot\frac{1-\rho_1}{1-\rho_2} \tag{3}$$

Since $0<\rho_1<\rho_2<1$, taking the natural log of both sides yields

$$\frac{\Delta}{\lambda}\le\ln\left(\frac{\rho_2}{\rho_1}\cdot\frac{1-\rho_1}{1-\rho_2}\right),$$

$$\lambda\ge\frac{\Delta}{\ln\left(\frac{\rho_2}{\rho_1}\cdot\frac{1-\rho_1}{1-\rho_2}\right)}.$$

Therefore, we have an import conclusion, for any informed adversary if $\lambda=\Delta\Big/\ln\left(\frac{\rho_2}{\rho_1}\cdot\frac{1-\rho_1}{1-\rho_2}\right)$, M is achieving (ρ_1,ρ_2) -privacy.Laplace distribution requires $\lambda>0$,yields

$$\ln\left(\frac{\rho_2}{\rho_1}\cdot\frac{1-\rho_1}{1-\rho_2}\right)>0,$$

$$\frac{\rho_2}{\rho_1}\cdot\frac{1-\rho_1}{1-\rho_2}>1,$$

$$\rho_2 > \rho_1 .$$

This implies that it is impossible to protect the privacy of individuals in the database with the probability less than the prior probability ρ_1.

Observing that, if we set $\rho_1 = 1/m$, where $m = |U|$, we can get $\lambda \geq \Delta \Big/ \ln \frac{(m-1)\rho_2}{1-\rho_2}$. It means thatρ-differential identifiability is the instance of LPB proposed in this paper.

V. EXPERIMENTS

This section experimentally evaluates the practical applicability of LPB. For comparing withρ-differential identifiability, weuse the same aggregate queries:**mean**andthe same database Adult Database from the UCI, comprised of 48,842 individuals.There are 14 attributes in the database which is contains 9 categorical and 5 numerical attributes. Only 3numerical attributes are used in this experiment.The Table**II** shows the description of the database.

To determine the noise distribution, we must calculate the Δf. For example, assume the adversary knows age of every individual except one. The possible value the adversary would guess is 1~99. Therefore, the sensitive $\Delta f = \left| f(D_{90}) - f(D_{17}) \right| = \frac{90-17}{48842} = 0.0015$. An adversary's probability of a correct random guess is 1/|*U*| which is presented in the Table II as RG.

TABLE II. DESCRIPTION OF ADULT DATABASE

Attribute	Max	Min	Sensitive	RG
age(AG)	90	17	0.0015	0.0137
education-num(EN)	16	1	0.0031	0.0101
hourse-per-week(HW)	99	1	0.0020	0.0625

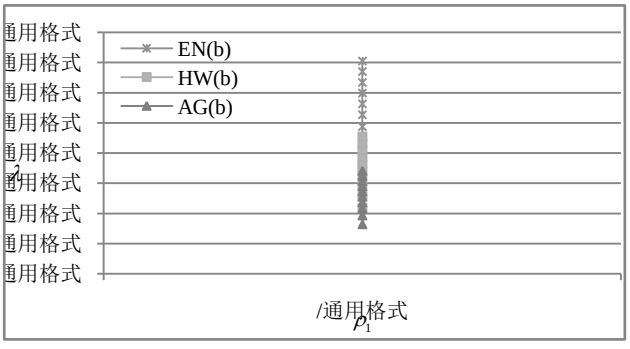

Figure 1. Themagnitude of thenoisy added versus breach level pho1with pho2 = 0.5

In our first set of experiments, we examine how much noisy should be added when the prior probability ρ_1 varying with $\rho_2 = 50\%$ .LPB shows that the magnitude of the noisy added not only depend on the posterior probability, but also depend on the prior probability. Figure 1 shows how noisy depend on the privacy requirement. We require that there are no breaches with the prior below ρ_1 and posterior at $\rho_2 = 50\%$, where $\rho_1 = 1\% \cdots 10\%$. As we see in the figure 1,the magnitude of the noisyadded (λ)increases when the prior probability ρ_1 increases. That means the smaller the value $\rho_2 - \rho_1$ is, the more noisy is needed. Therefore, for some prior probability more than the probability of random guess $\rho_1 > 1/|U|$, differential identifiability gives less noisy to protect (ρ_1, ρ_2)-privacy.

To show thereliabilityof results fromLPB, the results of 1000 queries for mean on attributes of age, hours-per-week, and education number as ρ_1 and ρ_2 varied are demonstrated in Figure 2~4.The vertical axis gives the noise ratio:

$$\text{Noise ratio}(\text{NR}) = \frac{R - f(D)}{U_{range}}$$

Where R is a response and $U_{range} = \max - \min$ is the range of the domain. The legend is presented in the Figure 2(b), other figures2 ~ 4are the same. Q1 means the 25th percentile, and Q3 is the 75th percentile.

As the figures show, all the responses are close to the true answer.When ρ_1 is fixed, increasing ρ_2, more noisy is needed to protect (ρ_1, ρ_2)-privacy, this is the same as differential identifiability. However, differential identifiability is dependent on the same prior probabilitywhich is set to the random guess 1/|*U*| while our work is not.Setting ρ_2 to 0.2, we can see that (from left to right) when ρ_1 increases, the noisy needed also increases because λ getting bigger (showed in Figure 1).

Figure 5 studies the information for differential privacy as the values of ε varies.Figure 5 shows that LPB can achieve the same effect as the ε-differential privacy.However, for ε-differential privacy, how to set the privacy parameter ε is a difficult problem.Most worksetparameters throughexperiences or experiments, while the way of setting privacy parameters by LPB has a better semantic interpretation which is that we must set ε to achieve (ρ_1, ρ_2)-privacy.

VI. CONCLUSIONS

In this paper, we propose a tactic to set the privacy parameter for ε -differential privacyintuitively. Compared to differential identifiability which assumes that the prior probability of an individual being in the input dataset is the same for all individual, our work provides a more flexible privacy parameter setting methods which do not depend on the prior probability.In the future work, we intend to investigate other version differential privacy such as probability differential privacy [11]support (ρ_1, ρ_2) -privacy as well.

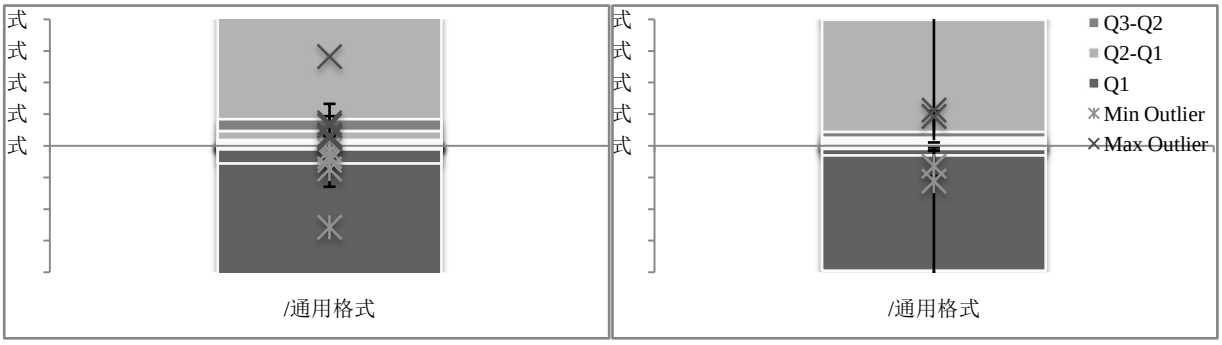

(2-a) $\rho_1 = 0.0137$ (2-b) $\rho_1 = 0.049$

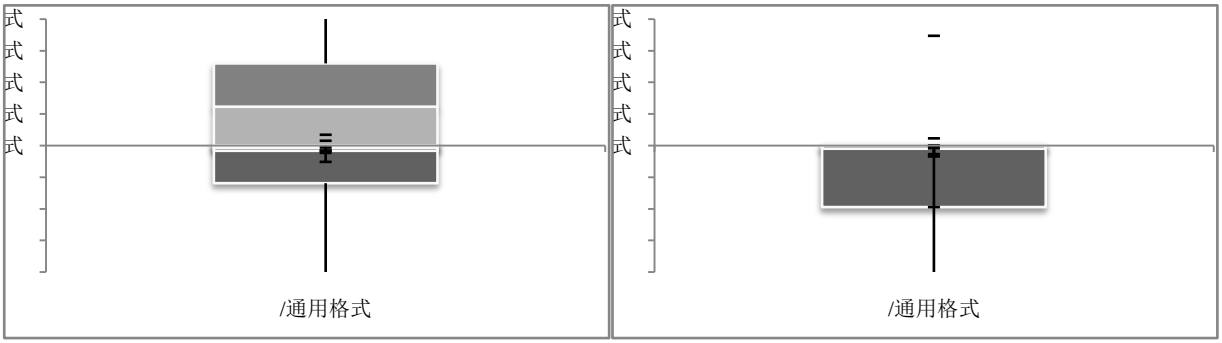

(2-c) $\rho_1 = 0.09$ (2-d) $\rho_1 = 0.19$

Figure 2. Noise Ratio for Age

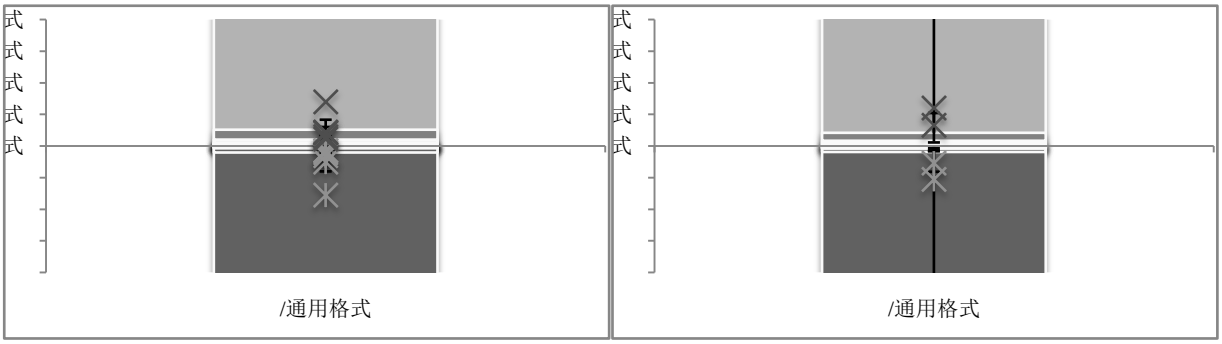

(3-a) $\rho_1 = 0.0101$ (3-b) $\rho_1 = 0.049$

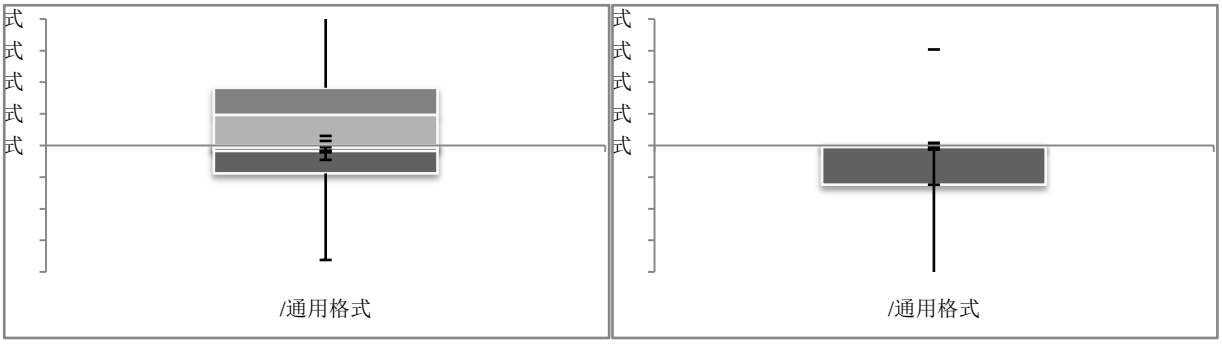

(3-c) $\rho_1 = 0.09$ (3-d) $\rho_1 = 0.19$

Figure 3. Noise Ratio for Hours-per-week

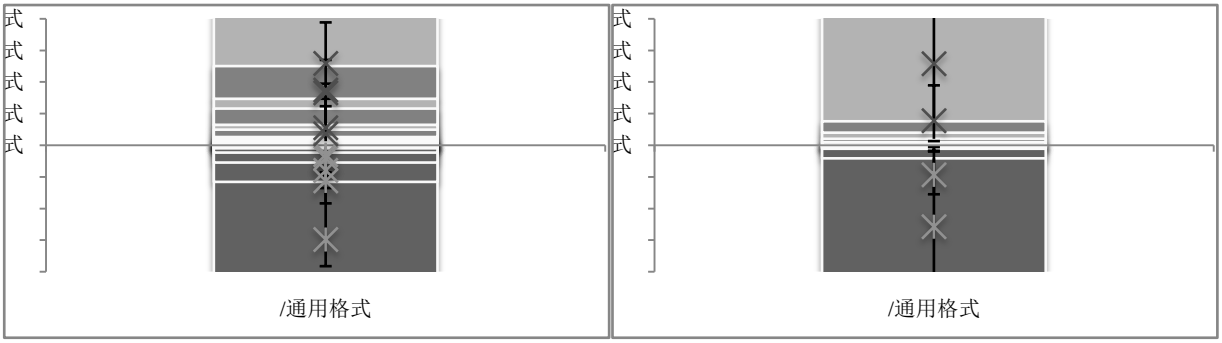

(4-a) $\rho_1 = 0.0625$ (4-b) $\rho_1 = 0.079$

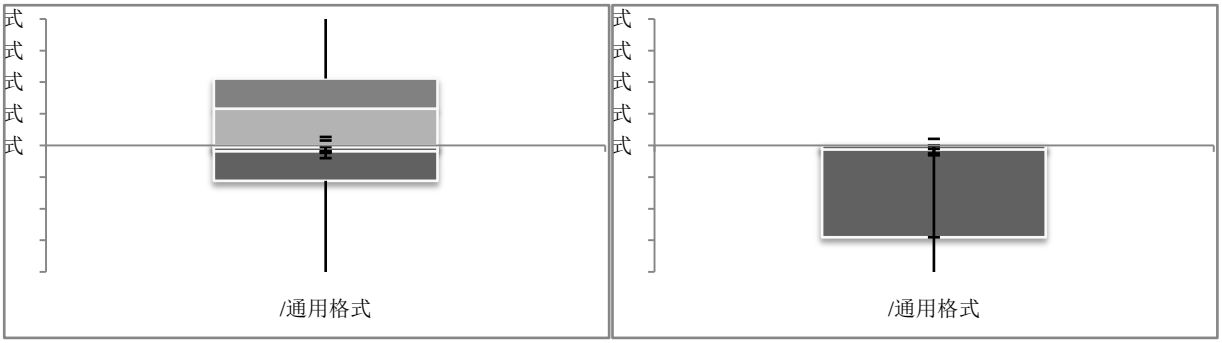

(4-c) $\rho_1 = 0.09$ (4-d) $\rho_1 = 0.19$

Figure 4. Noise Ratio for Education Number

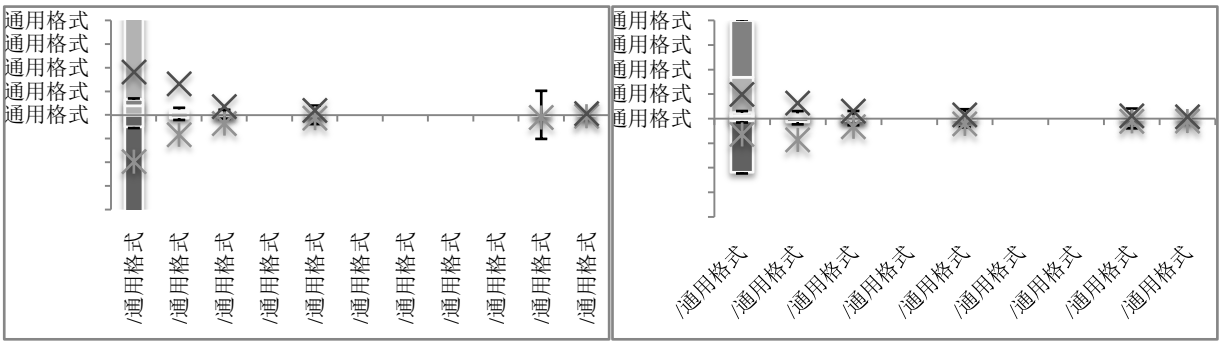

(5-a)Age(5-b) Hours-per-week

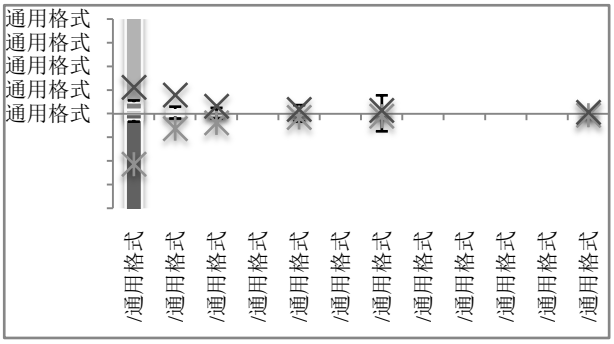

(5-c)Education Number

Figure 5. Noise Ratio for Differential Privacy

REFERENCE

[1] Agrawal R, Srikant R. Privacy-preserving data mining[J]. ACM Sigmod Record. 2000, 29(2): 439-450.

[2] Aggarwal C C, Philip S Y. A general survey of privacy-preserving data mining models and algorithms[M]. Springer, 2008.

[3] Dwork C, Mcsherry F, Nissim K, et al. Calibrating noise to sensitivity in private data analysis[J]. Theory of Cryptography. 2006: 265-284.

[4] Dwork C. Differential privacy[J]. Automata, languages and programming. 2006: 1-12.

[5] Dwork C. Differential privacy in new settings[C]. In: SODA '10 Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms.Society for Industrial and Applied Mathematics, 2010. 174-183.

[6] Dwork C. Differential privacy: A survey of results[J]. Theory and Applications of Models of Computation. 2008: 1-19.

[7] Lee J, Clifton C. Differential identifiability[C]. In: Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining.Beijing, China: ACM, 2012. 1041-1049.

[8] Evfimievski A, Gehrke J, Srikant R. Limiting privacy breaches in privacy preserving data mining[C]. In: ACM, 2003. 211-222.

[9] Sweeney L. k-anonymity: A model for protecting privacy[J]. International Journal of Uncertainty Fuzziness and Knowledge Based Systems. 2002, 10(5): 557-570.

[10] Machanavajjhala A, Kifer D, Gehrke J, et al. l-diversity: Privacy beyond k-anonymity[J]. ACM Transactions on Knowledge Discovery from Data (TKDD). 2007, 1(1): 3.

[11] Machanavajjhala A, Kifer D, Abowd J, et al. Privacy: Theory meets practice on the map[C]. In: Data Engineering, 2008. ICDE 2008. IEEE 24th International Conference on.IEEE, 2008. 277-286.