Heuristic Privacy Parameter Setting Strategies for Differential Privacy Models: Postprint

Authors: Ouyang Jia, Xiao Zhenghong, Liu Shaopeng, seal, Lin Piyuan

Date: 2018-05-20T00:00:00+00:00

Abstract

The differential privacy model is a strong privacy model that measures the level of privacy protection and the amount of noise via the privacy parameter , and has become a research hotspot in the field of privacy protection in recent years. However, the configuration of the privacy parameter relies solely on experiments or expert experience, which limits the adoption and dissemination of the differential privacy model. To address this problem, based on the (1, 2)-privacy model, we propose a heuristic privacy parameter setting strategy (limit privacy breaches in differential privacy, LPBDP), which analyzes the intrinsic relationship between the privacy parameter and (1, 2), thereby enabling the amount of noise added to be determined by (1, 2). LPBDP sets the privacy parameter through the following heuristic principle: if the attacker' s prior probability regarding the target victim is less than threshold 1, then after obtaining the noisy result returned by the differential privacy query mechanism, the attacker's posterior probability regarding the target victim must be less than threshold 2. Experimental results demonstrate that LPBDP can more intuitively configure the privacy parameter to satisfy differential privacy constraints.

Full Text

Preamble

http://www.arocmag.com/article/02-2019-01-037.html ChinaXiv 合作期刊计算机应用研究差分隐私模型的启发式隐私参数设置策略 * 欧阳佳 1，肖政宏 1，刘少鹏 1，印鉴 2，林丕源 3 (1. 广东技术师范学院计算机科学学院, 广州 510665; 2. 中山大学数据科学与计算机学院, 广州 510275; 3. 华南农业大学数学与信息学院, 广州 510642) 摘要：差分隐私模型是一种强隐私模型，用隐私参数 度量隐私保护程度及噪声量，近年来成为隐私保护领域的研究热点。但是隐私参数 的设置只能依赖于实验或专业人士经验，限制了差分隐私模型的使用与推广。针对这个问题，基于 (1，2)-隐私模型提出一种启发式的隐私参数 设置策略 (limit privacy breaches in differential privacy，LPBDP)，

分析隐私参数 与（1，2）的内在联系，实现噪声量的添加由（1，2）决定。LPBDP 通过如下启发式原则设置隐私参数

：如果攻击者关于目标受害者的先验概率小于阈值 1，攻击者得到差分隐私查询策略返回的加噪结果后，关于目标受害者的后验概率必须小于阈值 2。实验表明 LPBDP 能够更直观地设置隐私参数 以满足差分隐私约束。

关键词：隐私保护；差分隐私；隐私参数；隐私泄露中图分类号：TP309.2 doi: 10.3969/j.issn.1001-3695.2017.06.0649 Heuristic privacy parameter setting strategy for differential privacy model Ouyang Jia1, Xiao Zhenghong1, Liu Shaoeng1, Yin Jian2, Lin Piyuan3 (1. College of Computer Science Guangdong Polytechnic Normal University, Guangzhou 510665, China; 2. School of Data & Computer Science, Sun Yat-sen University, Guangzhou 510275, China; 3. College of Mathematics & Informatics, South China Agricultural University, Guangzhou 510642, China) Abstract: The differential privacy model is a kind of strong privacy model, which uses the privacy parameter to measure the degree of privacy protection and the amount of noise. In recent years, the privacy model has become a hotspot in the field of privacy protection. However, the setting of the privacy parameter can only depend on the experience of the lab or the professional experience, limiting the adoption and popularize of the differential privacy model. Aiming at this problem, a kind of heuristic privacy parameter setting strategy (limit privacy breaches in differential privacy, LPBDP) is proposed based on the (1, 2) -privacy model. The intrinsic relationship between the privacy parameter and (1, 2) is analyzed, and the addition of the noise quantity is determined by the parameters (1, 2) . LPBDP sets the privacy parameter by the following heuristic principle: If the attacker' s prior probability of the target victim is less than the threshold 1, then, the attacker' s posterior probability of the victim of the target must be less than threshold 2. Experiments show that LPBDP can more visually set the privacy parameter to meet the differential privacy constraints.

Key Words: privacy-preserving; differential privacy; privacy parameter; privacy breaches

0 Introduction

With the development of computer science, networking, and storage technologies, human society has collected and stored data on an unprecedented scale. This explosive growth of data has, in turn, fueled tremendous advances in data mining, which has been successfully applied across various industries such as healthcare, social networks, and online search. However, since access to raw data is a prerequisite for data mining, and raw data often contains personal privacy information, individuals have become increasingly concerned about privacy protection. Privacy-preserving data mining has emerged to address these privacy concerns associated with data mining, with the goal of successfully building effective data mining models without leaking the original input data [**?**, **?**, **?**, **?**]. Specifically, privacy-preserving data mining needs to address two key issues: (a)

how to protect personal privacy during the data mining process, and (b) how to ensure the utility of data or results. Currently, research in privacy-preserving data mining primarily focuses on designing privacy criteria and algorithms that simultaneously satisfy these two key requirements.

The -differential privacy model [**?**, **?**, **?**, **?**] makes no assumptions about an attacker' s background knowledge and represents a very strong privacy protection model. Its basic idea is to add a small amount of noise to analysis results before publishing them to satisfy differential privacy requirements. The amount of noise is determined by the sensitivity of the analysis function or process and the privacy parameter , independent of the specific database type or size. The privacy parameter is a crucial parameter in the differential privacy model that determines the amount of noise added and measures the degree of privacy protection. As can be seen from the Laplace mechanism and exponential mechanism, larger values result in less noise added, while smaller values result in more noise. However, the differential privacy model faces two problems when determining how much noise to add: First, the privacy parameter only limits the influence of individual records on the results, rather than limiting how much information is leaked about individuals [**?**], which allows attackers to easily identify individuals' sensitive information after obtaining randomized results. Second, the setting of the privacy parameter can only rely on experiments or expert experience, lacking a more intuitive heuristic parameter setting method.

To address these two problems, the main contributions of this paper are as follows: a) To limit personal information leakage in differential privacy models, we propose a new attack model based on the (1, 2)-privacy model concept; b) We identify the relationship between the privacy parameter and (1, 2), and propose a heuristic privacy parameter setting strategy.

1 Related Work

2 Privacy Parameter Setting Based on (1, 2)-Privacy Model

2.1 Privacy Model

Literature [**?**] first introduced the concept of the (1, 2)-privacy model, defined as follows: When the prior probability of random variable X taking value x satisfies $\Pr[X = x] \leq \rho_1$, and after obtaining the perturbed result through privacy mechanism M, the posterior probability of x updates to $\Pr[x|M(D) = R] \leq \rho_2$, then x is said to satisfy the (1, 2)-privacy model constraint. From this definition, the (1, 2)-privacy model does not depend on prior probabilities—it means that if the prior probability does not exceed 1, the posterior probability must be less than 2, where 1 and 2 can be customized without depending on any background knowledge.

2.2 Attack Model

Assume the attacker' s background knowledge includes all possible values U and all tuples in the database except the nth tuple, denoted as D_{-n}. Additionally, the attacker knows all details of the privacy mechanism M and the probability density function of the added noise. To infer the value of the nth tuple, the attacker guesses all values in U with equal probability before obtaining the result from privacy mechanism M. When a user submits query f to privacy mechanism M and receives perturbed result $R = M_f(D)$, the attacker' s probability of guessing that the nth tuple has value i is given by:

$$\Pr[X = i|R = M_f(D)] = \frac{\Pr[R = M_f(D_i)]}{\sum_{j \in U} \Pr[R = M_f(D_j)]}$$

where $D_i = D_{-n} \cup \{i\}$. If $\Pr[X = i|R = M_f(D)] > \rho_2$, then privacy is breached.

2.3 Attack Model Example

The following example illustrates the above attack model process. Even though query mechanism M satisfies differential privacy constraints, attackers can still guess individual values with high posterior probability. Let f be an averaging query function, given dataset $D = \{1, 2, 3, 10\}$ where all values come from $U = \{1, 2, 3, 5, 10\}$. Assume the attacker knows $D_{-4} = \{1, 2, 3\}$ and wants to infer the 4th value. Since the 4th value could be any element from U, the sensitivity of f is $\Delta f = 9/4$. Let the differential privacy parameter be $\varepsilon = 5.0$, and the attacker obtains result $R = 5.401$ after submitting an averaging request. The missing value is one of the elements in U, and the attacker calculates posterior probabilities as shown in Table 1 .

Table 1 Attacker' s Guessing Values

Possible Value	$\Pr[M_f(D_i)] = 5.401$
1	1,2,3,1
2	1,2,3,2
3	1,2,3,3
5	1,2,3,5
10	1,2,3,10

We provide the posterior probability calculation process using value 10 as an example. If the privacy mechanism M returns result $R = 5.401$, then any dataset could potentially return this value. Therefore, based on the Laplace mechanism, we have:

$$\Pr[M_f(D_i) = 5.401] = \Pr[\text{mean}(D_i) + \text{Lap}(\lambda) = 5.401]$$

Applying the triangle inequality yields:

$$\Pr[M_f(D_i) = 5.401] = \frac{1}{2\lambda} e^{-\frac{|5.401 - \text{mean}(D_i)|}{\lambda}}$$

where $\lambda = \frac{\Delta f}{\varepsilon}$. First, based on the returned value 5.401, we calculate the probability $\Pr[M_f(D_{10}) = 5.401]$. The differential privacy mechanism M adds noise $\text{Lap}(\lambda)$ to the true value, allowing us to determine $\lambda = 1.1258$.

Theorem 1 [10]: If an -differential privacy mechanism M satisfies -amplification for all response values, where $\gamma \geq 1$, then M must satisfy (1, 2)-privacy constraints.

Based on Theorem 1, we can find the relationship between differential privacy parameter and (1, 2). From the definition of -amplification, we obtain:

$$\gamma = \max_{i,j} \frac{\Pr[R|M_f(D_i)]}{\Pr[R|M_f(D_j)]} \leq e^{\varepsilon \cdot \Delta f / \lambda}$$

Assuming equal prior probabilities for values in U, taking the natural logarithm of both sides yields:

$$\ln \gamma \leq \varepsilon \cdot \frac{\Delta f}{\lambda}$$

This gives us an important result: for any attacker, if we set the differential privacy parameter as:

$$\varepsilon = \frac{\lambda}{\Delta f} \ln \left(\frac{1 - \rho_1}{\rho_1} \cdot \frac{\rho_2}{1 - \rho_2} \right)$$

then the differential privacy mechanism M satisfies (1, 2)-privacy constraints. Since the Laplace distribution requires $\lambda > 0$, we have:

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

which implies $\rho_2 > \rho_1$. This means the posterior probability protecting individual privacy must exceed its prior probability; otherwise, protection would be meaningless. This conclusion obviously aligns with reality—if the probability of protecting an individual cannot exceed its prior probability, protection loses its purpose.

Note that if we set $\rho_1 = 1/m$ and $\rho_2 = \rho$, we obtain:

$$\varepsilon = \frac{\lambda}{\Delta f} \ln \left(\frac{m - 1}{\rho} - 1 \right)$$

This shows that -differential identifiability is a special case of our proposed LPBDP method.

3 Experimental Results and Analysis

This section first compares the differences and connections between LPBDP and -differential identifiability, as shown in Table 2 , and then experimentally analyzes the heuristic and semantic properties of LPBDP.

Table 2 Differences and Connections Between LPBDP and -Differential Identifiability

Aspect	LPBDP	-Differential Identifiability
Prior Probability	Greater than 1	Equal random guess probability 1/m
Posterior Probability	Less than 2	Less than
1 Setting	Customized based on actual needs	1/m (random guess probability)
Prior Knowledge Required	Basically no assumptions	Requires: (1) All possible values U and

On one hand, Table 2 shows that our proposed LPBDP method also satisfies differential privacy requirements and has advantages over -differential identifiability in terms of prior knowledge. LPBDP essentially requires no prior knowledge assumptions and offers better adaptability.

On the other hand, we analyze the practical application of LPBDP through experiments. To compare with -differential identifiability, we use the same aggregate query function (averaging: mean) and experimental data from the UCI Adult database, which contains 48,842 records with 14 attributes (9 categorical and 5 numeric). In this paper, we only use 3 numeric attributes. Table 3 describes the characteristics of the Adult database.

Table 3 Adult Database Characteristics

Attribute	Max	Min	Sensitivity	Random Guess Probability
age (AG)	90	17	73	0.0137
education-num (EN)	16	1	15	0.0625
hours-per-week (HW)	99	1	98	0.0101

To determine the Laplace distribution function for added noise, we must calculate the sensitivity of the averaging function: $\Delta f = \frac{U_{\max} - U_{\min}}{|D|}$. For example, if an attacker knows all records in the database except one, the attacker' s guessing range for age is 1-99, so the function sensitivity is $\Delta f = 90/48842 = 0.0015$, and the random guess probability is $\rho_1 = 1/73 = 0.0137$, as shown in Table 3 under RG (random guess).

LPBDP shows that noise addition is influenced by both prior and posterior probabilities. We first verify through experiments how added noise is affected by prior probability, setting posterior probability to $\rho_2 = 0.2$. The differential privacy mechanism must satisfy (1, 2)-privacy requirements, where ε is calculated using our derived formula. As shown in Figure 1 [Figure 1: see original paper], the noise amount λ increases with prior probability, meaning more noise is needed as ρ_1 increases—consistent with -differential identifiability.

To validate LPBDP' s practicality, we submitted 1,000 averaging queries for each of the 4 attributes. Figures 2 [Figure 2: see original paper]-4 show the impact on noise rate, calculated as NR = $\frac{\text{range}(R)}{U_{\max} - U_{\min}}$, where range(R) is the interval of perturbed query results. Q1 is the first quartile, Q3 is the third quartile, and Q3-Q1 is the interquartile range. Figures 3 [Figure 3: see original paper]-4 show that all response values cluster near the true values. When ρ_2 is fixed at 0.2 and ρ_1 increases, more noise is required to satisfy privacy constraints. When ρ_1 is fixed and ρ_2 increases, λ also increases, requiring more noise.

Figure 5 [Figure 5: see original paper] examines the impact of privacy parameter ε on differential privacy. The results show that LPBDP achieves the same effect as -differential privacy. However, for -differential privacy, setting ε is a major challenge, typically done through experience or experiments. LPBDP offers better semantic meaning for privacy parameter setting, requiring only that $\rho_2 > \rho_1$.

4 Conclusion

Previous methods for setting privacy parameters in differential models primarily relied on experiments or expert experience. This paper proposes a heuristic differential privacy parameter setting strategy. -Differential identifiability is another parameter setting strategy for differential privacy, but it depends on two assumptions: (1) knowing each value' s prior probability and assuming knowledge of all possible values U and |U|; (2) equal prior probabilities for all possible values. However, some application scenarios cannot satisfy these assumptions. Our proposed method addresses this limitation by introducing a new privacy parameter setting strategy LPBDP based on the (1, 2)-privacy model. The advantage of this strategy is that the (1, 2)-privacy model does not depend on prior probabilities, does not require knowledge of U, and LPBDP still satisfies differential privacy constraints.

References

[1] Ouyang Jia, Yin Jian, Liu Shaopeng. An effective differential privacy transaction data publishing strategy [J]. Journal of Computer Research and Development, 2014, 51(10): 2195-2205.

[2] Ouyang Jia, Yin Jian, Liu Shaopeng. A distributed transaction data differential privacy publishing strategy [J]. Journal of Software, 2015, 26(6): 1457-1472.

[3] Goethals B, Laur S, Lipmaa H, et al. On private scalar product computation for privacy-preserving data mining [C]//Proc of the 7th International Conference on Information Security and Cryptology. 2004: 104-120.

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

[5] Dwork C, McSherry F, Nissim K, et al. Calibrating noise to sensitivity in private data analysis [C]//Proc of the 3rd Conference on Theory of Cryptography Theory of Cryptography. 2006: 265-284.

[6] Dwork C. Differential privacy [C]//Proc of International Colloquium on Automata, Languages and Programming. 2006: 1-12.

[7] Dwork C. Differential privacy in new settings [C]//Proc of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms. Society for Industrial and Applied Mathematics. 2010: 174-183.

[8] Dwork C. Differential privacy: a survey of results [C]//Proc of the 5th Conference on Theory and Applications of Models of Computation. 2008: 1-19.

[9] Lee J, Clifton C. Differential identifiability [C]//Proc of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data mining. [S. l.]: ACM Press, 2012. 1041-1049.

[10] Evfimievski A, Gehrke J, Srikant R. Limiting privacy breaches in privacy preserving data mining [C]//Proc of the 22nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems. [S. l.]: ACM Press, 2003: 211-222.

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

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

Note: Figure translations are in progress. See original paper for figures.