Customized Privacy Preserving for Classification Based Applications

Zaobo He
Department of Computer
Science
Georgia State University
Atlanta, Georgia 30302
zhe4@student.gsu.edu

Zhipeng Cai
Department of Computer
Science
Georgia State University
Atlanta, Georgia 30302
zcaic@gsu.edu

Yingshu Li
Department of Computer
Science
Georgia State University
Atlanta, Georgia 30302
yli@gsu.edu

ABSTRACT

The rise of sensor-equipped smart phones has enabled a variety of classification based applications that provide personalized services based on user data extracted from sensor readings. However, malicious applications aggressively collect sensitive information from inherent user data without permissions. Furthermore, they can mine sensitive information from user data just in the classification process. These privacy threats raise serious privacy concerns. In this paper, we introduce two new privacy concerns which are inherent-data privacy and latent-data privacy. We propose a framework that enables a data-obfuscation mechanism to be developed easily. It preserves latent-data privacy while guaranteeing satisfactory service quality. The proposed framework preserves privacy against powerful adversaries who have knowledge of users' access pattern and the data-obfuscation mechanism. We validate our framework towards a real classification-orientated dataset. The experiment results confirm that our framework is superior to the basic obfuscation mechanism.

Keywords

latent-data privacy; obfuscation, tradeoff, optimization

1. INTRODUCTION

Pervasiveness of mobile devices equipped with a variety of sensors necessitates various wireless applications providing convenient services to people, among which classification based applications are popular ones. Given a training data set and associated training labels, such applications determine the labels for an unlabeled user based on the multimodal sensory data collected by multiple sensors. For example, feature variables of customers may be used to mine buying interests of customers given some training examples [27] [6]; features extracted from health monitoring systems may be used to predict whether a user may carry a

disease in the future [5] [3].

However, these classification based applications raise serious privacy concerns, mainly because user data released to service providers contain sensitive information. Sensitive information, such as political preference [21], and sexual orientation [15], can be successfully inferred by either eavesdropper or malicious service providers. Such privacy issues are of frequent occurrence as many adversaries unrighteously collect personal data [9]. As shown in [9], among the 30 popular Android applications, 68 instances of 20 applications accessed users' sensory data and sent them to remote servers. Moreover, privacy leaks can be exacerbated if user data are analyzed together with other data sources that can be obtained from public social networks [20]. Thus, how to preserve user privacy without degrading the service quality of classification based applications is a crucial problem.

Many existing approaches preserve user privacy by filtering or perturbing user data when connecting to service providers [22] [4] [11]. Unfortunately, service quality may be degraded due to the distorted data. Obviously, a tradeoff between service quality and privacy preserving level is formed: a higher privacy preserving level results in more service quality loss. Moreover, the existing approaches overlook the fact that adversaries can not only directly access inherent user data but also mine latent sensitive information. We define these two kinds of privacy as *inherent-data privacy* and *latent-data privacy*, respectively. For example, *age* and *weight* are inherent-data explicitly reported by users to service providers; *gender*, *political preference*, and *sexual orientation* may be latent-data mined from user data.

In this paper, we mainly focus on latent-data privacy and investigate how to customize privacy preserving for classification based applications, so that the tradeoff between service quality and privacy preserving level can be optimized. Specifically, we employ a statistical analysis method as the classification technique, based on which we formalize the metrics to measure quality loss and latent-data privacy. For latent-data privacy, we expect adversaries cannot learn any sensitive information of users, that is, the cloaked data can guarantee that the classification results are also indistinguishable. We introduce a probability model for latent-data privacy and quantify it with level δ , *i.e.*, δ -latent-data privacy. Our framework can then comprehensively optimize the tradeoff between δ -privacy and service quality, *i.e.*, given user-specified thresholds regarding latent-data privacy, the best possible service quality can be obtained.

To the best of our knowledge, our work is the first one p-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ACM PAMCO'16 July 05 2016, Paderborn, Germany

© 2016 ACM. ISBN 978-1-4503-4346-6...\$15.00

DOI: /10.1145/2940343.2940345

reserving latent-data privacy towards multimodal data. The main contributions are summarized as follows:

- We introduce a novel framework to optimize the trade-off between privacy preserving level and service quality. In particular, we design a data obfuscation mechanism specifying how a user should perturb data to achieve the best possible service quality with latent-data privacy preserved.
- We also develop an analytical model which can assist system designers to construct an optimal privacy preserving strategy against adversaries, considering that adversaries may be aware of user access pattern and the underlying obfuscation mechanisms.
- We evaluate our proposed framework towards a real Marketing dataset. The results show that our framework can preserve latent-data privacy without sacrificing much service quality. Moreover, compared with the k -anonymity method, our framework provides better service quality with the same privacy preserving level and higher privacy preserving level with the same service quality.

The remainder of the paper is organized as follows. Section 2 formalizes the investigated problem. Section 3 introduces some preliminary knowledge. The metrics for service quality and privacy are presented in Section 4. We illustrate the proposed framework in Section 5. The evaluation results are shown in Section 6. The related works are addressed in Section 7. Section 8 concludes the paper.

2. PROBLEM STATEMENT

2.1 User model

DEFINITION 1. Profiles. An arbitrary user u is described by a set of n -dimensional attribute vectors X_m^u ($1 \leq m \leq M_u$). For each X_m^u , $X_m^u = (c_1, \dots, c_t; s_{t+1}, \dots, s_n)$, where c_i ($1 \leq i \leq t$) and s_j ($t+1 \leq j \leq n$) is the non-sensitive and sensitive attribute, respectively.

DEFINITION 2. Access pattern. We call $\psi(X_m^u)$ as user u 's access pattern if each X_m^u is assigned a probability, obviously, $\sum_{1 \leq m \leq M} \psi(X_m^u) = 1$.

For example, in a health care application, a user may present different profiles in morning, afternoon, night in a day with different probability. Based on a user's X_m^u ($1 \leq m \leq M$), a service provider can classify the user into different groups (each group corresponds to one type of service), then the corresponding services are provided.

2.2 Adversary model

In classification based applications, a service provider is honest-but-curious. That is, it honestly classifies users into different groups and recommends corresponding services based on the classification results. However, it intends to discover sensitive information of users as much as possible.

Adversaries can mine sensitive information from the multimodal profiles of users just in the classification process. We assume there are two types of groups in classification based applications: non-sensitive groups $k_i \in G$, ($1 \leq i \leq K$) and sensitive groups $p_i \in G$, ($1 \leq i \leq P$), where $G =$

$(k_1, \dots, k_K; p_1, \dots, p_P)$ is the set of all the groups. Adversaries may classify users into sensitive groups privately. This malicious behavior seriously threatens user privacy. For example, a user of a healthcare application may not expect the service provider to know that she carries the AIDS disease. Thus, the classification results for sensitive groups should be inaccurate as much as possible. We denote this type of privacy as latent-data privacy.

2.3 Obfuscation mechanism

In our work, to preserve latent-data privacy in classification based applications, we assume users adopt an obfuscation mechanism that transforms sensitive attributes to pseudo-attributes. Thus, a real profile $X_m^u \in \Omega$ is transformed to a cloaked profile $X_m'^u \in \Omega'$, and users report $X_m'^u$ to the service provider, where Ω' is the value space of cloaked profiles. We use X/X' to represent $X_m^u/X_m'^u$ in the rest of paper without confusion. The obfuscation mechanism is a probability function that takes real profile $X = (c_1, \dots, c_t; s_{t+1}, \dots, s_n)$ as input and outputs the cloaked profile $X' = (c_1, \dots, c_t; s'_{t+1}, \dots, s'_n)$ with the following probability:

$$f(X'|X) = Pr(X' = (c_1, \dots, c_t; s'_{t+1}, \dots, s'_n) | X = (c_1, \dots, c_t; s_{t+1}, \dots, s_n)) \quad (1)$$

2.4 Problem definition

Users expect to be classified into non-sensitive groups as accurate as possible with their privacy preserved. Thus, the ultimate goal is to design a privacy preserving mechanism that can provide privacy guarantees, while users can control privacy preservation levels by trading off service quality. Specifically, given threshold on the preserving level of latent-data privacy, we want to develop a mechanism that can provide the best possible service quality. We want to contend against the adversaries who know the obfuscation mechanism ($f(\cdot)$) and the probability distributions of user's profile ($\psi(\cdot)$). Such information is valuable for adversaries to reverse-engineer the obfuscation mechanism to infer the real profiles of users. The investigated problem in this paper, δ -privacy preserving for classification based applications, is defined as follows:

Input:

- (1) User profile X .
- (2) User's Access pattern *i.e.*, $\psi(\cdot)$.
- (3) Level of latent-data privacy δ , specified by users.

Output:

User's profile obfuscation mechanism $f(\cdot)$ as defined in (1) satisfying that:

- a. The user can be classified into non-sensitive groups as accurate as possible;
- c. User's latent-data privacy can be preserved with level δ ;

3. PRELIMINARIES

3.1 Classification method

There are many classification approaches such as Bayes, decision trees, neural networks, and *etc.* [12]. In this paper, we focus on the statistical analysis classification method that is based on the concept of Mahalanobis distance. We choose this classification method, as it is widely utilized in the literature. We briefly introduce it here and more details can

be found in [23].

One can say that a user belongs to group η (k_i or p_i) if the distance between the user's profile and the centroid of group η is less than a threshold. Let G_η be the training dataset of group η , namely, G_η is the set of user profiles that belong to group η . $\overline{G}_\eta$ is the mean vector of G_η , namely, $\overline{G}_\eta$ is the centroid of group η . Let $G_\eta(i)$ be the i th row of G_η , and $\widehat{G}_\eta$ is a matrix where $\widehat{G}_\eta(i) = G_\eta(i) - \overline{G}_\eta$. Thus, the distance from the user with profile X to the centroid of group k can be computed as follows:

$$D_\eta^2(X) = (X - \overline{G}_\eta)^T C_\eta^{-1} (X - \overline{G}_\eta) + \ln |C_\eta|$$

$$C_\eta = \frac{1}{n_\eta - 1} \widehat{G}_\eta^T \widehat{G}_\eta$$

where n_η is the number of training vectors of G_η . For user u with profile X , if $D_\eta^2(X)$ is less than the threshold of group η , i.e., $D_\eta^2(X) \leq \Delta_\eta$, this user belongs to group η denoted as $u(X) \in \eta$.

4. METRICS

In this section, we introduce how to measure service quality and privacy.

4.1 Service quality metric

We want to evaluate the service quality loss Q_{loss} resulted by the obfuscation mechanism $f(\cdot)$. We use $d_q(\cdot)$ to denote the service quality metric, such that $d_q(X, X')$ measures how much quality loss there is if the service provider classifies a user with obfuscated profile X' rather than the real profile X . Thus, given $f(\cdot)$, $\psi(X)$, and $d_q(\cdot)$, we can define the quality loss as the sum of $d_q(X, X')$ over all X and X' for a user:

$$Q_{loss}(f(\cdot), \psi(X), d_q) = \sum_{X, X'} \psi(X) f(X'|X) d_q(X, X') \quad (2)$$

From (2), we find that the distortion induced by $f(X'|X)$ determines the service quality. We say that there is no quality loss if the cloaked profile X' does not change the classification result for non-sensitive groups. Otherwise, there exists quality loss. The next problem is how to formalize $d_q(X, X')$ in classification based applications.

Note that $d_q(\cdot)$ depends on data semantics and user-specified privacy requirements. Euclidean metric is a typical choice in location based services. However, in classification based applications, two different multimodal profiles may produce the same classification results so that taking the distance (such as Euclidean distance) between real profile X and cloaked profile X' as a metric is not advisable. Since whether a user belongs to a group is determined by the distance between the profile and the centroid of this group, we can evaluate service quality loss through the difference between the distance from the real profile to the centroid of a group and the distance from the cloaked profile to the centroid of a group: the less the difference, the better the service quality. Thus, quality loss can be measured by the expected $d_q(X, X')$ over all X , X' and all non-sensitive groups $k_i \in G$, which can be specified as follows:

$$\sum_{X, X' \in \mathcal{X}} \psi(X) f(X'|X) \sum_{k_i \in G} d_q^{k_i}(X, X') \quad (3)$$

where

$$d_q^{k_i}(X, X') = |D_{k_i}^2(X) - D_{k_i}^2(X')|,$$

and the distances to the centroid from X and X' , i.e., $D_{k_i}^2(X)$ and $D_{k_i}^2(X')$, are computed as follows:

$$D_{k_i}^2(X) = (X - \overline{G}_{k_i})^T C_{k_i}^{-1} (X - \overline{G}_{k_i}) + \ln |C_{k_i}|$$

$$D_{k_i}^2(X') = (X' - \overline{G}_{k_i})^T C_{k_i}^{-1} (X' - \overline{G}_{k_i}) + \ln |C_{k_i}|$$

4.2 Latent-data privacy metric

Adversaries may observe the cloaked profile X' outputted by the obfuscation mechanism together with $\psi(\cdot)$ and $f(\cdot)$. Then the posterior probability can be computed as follows:

$$Pr(X|X') = \frac{Pr(X, X')}{Pr(X')} = \frac{f(X'|X)\psi(X)}{\sum_X f(X'|X)\psi(X)}$$

towards the true profile X , conditional on the cloaked profile X' .

The adversaries can then choose $\hat{X}$ such that the difference between the distance from X to the centroid of a sensitive group and the distance from $\hat{X}$ to the centroid of a sensitive group is as small as possible. With minimized $\hat{X}$, the user's conditional expected privacy is

$$\min_{\hat{X}} \sum_X Pr(X|X') d_p(X, \hat{X}) \quad (4)$$

which is similar to $d_q(X, X')$, where $d_p(X, \hat{X})$ is the difference between the distance from the real profile to the centroid of a sensitive group and the distance from the guessed profile to the centroid of a sensitive group.

For a given X' , the conditional latent-data privacy is given by Equation (4). However, the probability of X' outputted by the obfuscation mechanism is $Pr(X') = \sum_X f(X'|X)\psi(X)$. Thus, the user's unconditional expected privacy is

$$\sum_{X'} Pr(X') \min_{\hat{X}} \sum_X Pr(X|X') d_p(X, \hat{X})$$

$$= \sum_{X'} \min_{\hat{X}} \sum_X \psi(X) f(X'|X) d_p(X, \hat{X})$$

Then the δ -latent-data privacy can be defined as follows:

DEFINITION 3. δ -latent-data privacy. An obfuscation mechanism $f(\cdot)$ satisfies δ -latent-data privacy if for any user u with profile $X = (c_1, \dots, c_t; s_{t+1}, \dots, s_n)$, for all $X' = (c_1, \dots, c_t; s'_{t+1}, \dots, s'_n)$ and all sensitive groups $p_i \in G$,

$$\sum_{X'} \min_{\hat{X}} \sum_X \psi(X) f(X'|X) \sum_{p_i \in G} d_p^{p_i}(X, \hat{X}) \geq \delta \quad (5)$$

From the above metrics, we found that both the service quality metric and the latent-data privacy metric computation require that the users have the values of $\overline{G}_\eta$ and C_η , where η is k_i or p_i . Hence, we introduce another dimension to the privacy-service quality tradeoff: the information the service provider disclosed to users. Suppose the service provider give the user the classification method, as well as the full training dataset, i.e., G_η . Then, the users can execute classification method locally, i.e., which groups are belonged to. Obviously, this would provide best privacy (the user does not reveal his profile to service provider) as well as best service quality (no obfuscation for the user's profile).

However, this is inadvisable from the view of service provider for commercial concerns such as the service is not free generally. Hence, the service provider just need to disclose $\bar{G}_\eta$ and C_η to users, based on which the users can measure service quality and the latent-data privacy, moreover, the users cannot deduce G_η from the $\bar{G}_\eta$ and C_η as indicated in the distance equation in III.A.

5. (ϵ, δ) -PRIVACY PRESERVING WITH OPTIMAL SERVICE QUALITY

We can develop an optimal mechanism as follows:

Minimize:

$$\sum_{X, X' \in \chi} \psi(X) f(X'|X) \sum_{k_i \in G} d_q^{k_i}(X, X')$$

Subject to:

$$\begin{aligned} \sum_{X'} \min_{\hat{X}} \sum_X \psi(X) f(X'|X) \sum_{p_i \in G} d_p^{p_i}(X, \hat{X}) &\geq \delta \\ \sum_{X' \in \chi} f(X'|X) &= 1, \quad X \in \chi \\ f(X'|X) &\geq 0 \quad X, X' \in \chi \end{aligned} \quad (6)$$

Following equation (6) in Definition 3, it is nontrivial to determine whether an obfuscation probability $f(\cdot)$ preserves δ -latent-data privacy. We transform the non-linear condition in Definition 3 to a linear problem. We define

$$Z_{X'} = \min_{\hat{X}} \sum_X \psi(X) f(X'|X) \sum_{p_i \in G} d_p^{p_i}(X, \hat{X}) \quad (7)$$

The above equation is equal to

$$Z_{X'} \leq \sum_X \psi(X) f(X'|X) \sum_{p_i \in G} d_p^{p_i}(X, \hat{X}), \quad \forall \hat{X} \quad (8)$$

Thus, the condition in Definition 3 is transformed to

$$\begin{aligned} \sum_{X'} Z_{X'} &\geq \delta \\ \text{S.t. } Z_{X'} &\leq \sum_X \psi(X) f(X'|X) \sum_{p_i \in G} d_p^{p_i}(X, \hat{X}), \quad \forall \hat{X} \end{aligned}$$

6. EVALUATION

6.1 Setup

Dataset. We employ the Marketing dataset [1] to evaluate our mechanism. It contains profile information of 6876 users, recorded by the questionnaires filled out by the shopping mall customers in the San Francisco Bay area. Each user's profile contains 13 attributes such as sex, marital status, age, education, occupation, *etc.* Each attribute takes discrete values from a finite space, for example, attribute *type of home* has 7 values taking from the space (house, condominium, apartment, mobile, *etc.*). The goal of this questionnaire is to predict the annual household income from the 13 attributes and classify users into 9 groups (less than \$10,000, \$10,000 to \$14,999, \$15,000 to \$19,999, and so on). We randomly select 100 non-redundant user profiles to simulate a user's profiles in a period of time, and the remaining profiles serve as training items.

Comparison. We compare our mechanism with the basic obfuscation mechanism which is widely utilized in the literatures. Consistent with differential privacy, the basic obfuscation mechanism is constructed with a specific protection level $k = 1, 2, 3, \dots$. For each profile X , we search its $k - 1$ nearest profiles (using Mahalanobis distance, Hamming distance or Euclidean distance as dissimilarity metrics). For an arbitrary profile X , the cloaked profile X' will be uniformly selected from the $k - 1$ nearest profiles together with X . Hence, X is replaced by any of the k profiles with probability $\frac{1}{k}$, and by any of the remaining profiles with probability 0.

Privacy configuration. Since different users can regard different attributes and classification results as sensitive information, we have two ways of choosing sensitive attributes and classification groups. Unless otherwise stated, we select *Sex*, *Marital status*, *Age*, *Occupation*, *Householder status* as sensitive attributes and the last three groups with higher household income as sensitive groups. Alternatively, we randomly select 5 attributes and 3 groups as sensitive information and set the remaining be as non-sensitive information.

6.2 Results

In order to perform a fair comparison between OptQL and BasicObf, we compare their latent-data privacy levels for a same quality loss. We select a δ , then compute the quality loss of our mechanism: $ql = \text{QL}(\text{OptQL}, \delta)$. Meanwhile, we select a proper k such that the quality loss of the basic obfuscation mechanism is $ql = \text{QL}(\text{BasicObf}, k)$. Hence, the two mechanisms have a same quality loss. The correspondence between δ and k is shown in Table 1, where δ is from 0.71 to 0.82.

With the selected k , we compute the latent-data privacy level respectively. With the increase of δ or k , we expect the level of latent-data privacy increases as well. We measure the latent-data privacy level in terms of adversary error when it infers user's latent sensitive information, denoted as *AdvErrLa*:

$$\text{AdvErrLa} = \min_h \sum_{\hat{X}, X', X} \psi(X) f(X'|X) h(\hat{X}|X') \sum_{p_i \in G} d_p^{p_i}(X, \hat{X})$$

We conduct two sets of experiments. In the first set, denoted as *Type 1*, we regard *Sex*, *Marital status*, *Age*, *Occupation*, *Householder status* as sensitive attributes and the last three groups as sensitive groups. In the second set, denoted as *random*, we randomly select 5 attributes and 3 groups as sensitive information and the remaining as non-sensitive information. As shown in Fig.2, for different privacy configurations, our mechanism clearly has a higher level of latent-data privacy than the basic obfuscation mechanism when both of them have a same service quality. Meanwhile, as we can see from Fig.1, with the increase of the latent-data privacy level by increasing δ , the quality loss increases as well.

7. RELATED WORKS

Threats. Privacy threats towards user data have been extensively documented. On the one hand, a large body of location-privacy preserving mechanisms (LPPMs), diverse in techniques or goals, have been studied. They generally deal with inherent-data privacy. On the other hand, inference attacks from users' public data also incur stringent concerns, and they are referred as latent-data privacy. In

Table 1: correspondence between δ and k , for level of latent-data privacy.

ql	3.2516	3.2840	3.2925	3.2925	3.3020	3.3067	3.3091	3.4481	3.4491	3.5240	3.5748	3.6177
δ	0.71	0.72	0.73	0.74	0.75	0.76	0.77	0.78	0.79	0.80	0.81	0.82
k	85	87	88	88	88	88	89	97	97	102	104	106

Figure 1: Quality loss of OptQL for different values of δ .

such a case, user’s sensitive information such as *gender* [21], *political preference* [16], *sexual orientation* [16] or even *drug use* could be inferred. In [?], the authors show that how the seemingly harmless published interests can leak sensitive information. The work in [2] shows that two New York Time journalists could identify sensitive information of individuals from the published search logs made available by AOL, including *name*, *age*, *sex*, *location* and so forth. Apart from inferring demographic information, deanonymizing is another type of inference attack. It is shown in [24] that medical visits of individuals have been successfully identified from the anonymized medical visit data made available by the Group Insurance Commission (GIC). The work in [20] shows that individuals’ can be identified based on movie rating data released by Netflix, even though the data are anonymous. In [13], the authors show that the spreading rumor in mobile social networks could also violate user privacy. Hence, the spate of privacy threats has spurred a large body of research in protecting inherent-data privacy and latent-data privacy. Next, we review the literatures from two aspects.

Privacy preserving data publishing. Many definitions for privacy have been proposed in the area of privacy preserving data publishing. The most prominent techniques include *t*-closeness [17], *l*-diversity [19] and *k*-anonymity [25]. As indicated in Section 6, *k*-anonymity guarantees that adversaries cannot distinguish real data from at least other $k - 1$ obfuscated data. However, as argued in [8], the data privacy preserving problems should be considered in a more rigorous way. Differential privacy, which is a complete framework, then becomes a popular technique for privacy preserving data publishing. For instance, the work in [4] proposes a mechanism that provides optimal service quality under certain data constraints based on differential privacy. However, differential privacy is employed to preserve aggregate single-modal data generally such as location privacy [14] [7]. It is not suitable for multimodal data. Moreover, [26] proposes an privacy preserving incentive mechanism in mobile crowdsourcing systems to improve system

(a)

(b)

Figure 2: Tradeoff between latent-data privacy and service quality, when privacy is configured by (a): Specific attributes and groups ; (b): random attributes and groups (x-axis just shows δ).

utility.

Privacy preserving data mining. Many privacy preserving data mining techniques distort or hide data before release data [10] [18]. Hence, the tradeoff between data obfuscation and service quality has been studied in several statistical based techniques. However, most of them only protect data privacy, rather than the privacy of the mining results, *i.e.*, latent information.

In contrast, our work improves the aforementioned works in considering latent-data privacy and service quality all together, rather than just the tradeoff between one type of privacy and service quality. To the best of our knowledge, our work is the first one to investigate the tradeoff among these three aspects.

8. CONCLUSIONS

Classification based applications incur privacy risks where sensitive attributes may be disclosed or latent sensitive information may be inferred based on the released user profiles. We define these two types of privacy threats as *Inherent*-

data privacy and *Latent-data privacy*, respectively. These privacy concerns necessitate the need of robust privacy preserving mechanisms, based on which the tradeoff between privacy and service quality can be optimized. In this paper, we propose an optimal profile-obfuscation mechanism preserving latent-data privacy while guaranteeing satisfactory service quality. Our mechanism guarantees privacy against powerful adversaries who have knowledge of user's access pattern and the obfuscation mechanism. The experiment results towards a real classification dataset demonstrate that our work can find the optimal mechanism for given privacy constraints, and that our work is superior to the basic obfuscation mechanism. Moreover, our work does not have to sacrifice much privacy (respectively, service quality) in order to guarantee service quality (respectively, privacy).

9. ACKNOWLEDGMENTS

This work is partly supported by the NSF under grant No.CNS-1252292, and the NSFC under grants NOs 61373083, 61370084 and 61502116.

10. REFERENCES

- [1] <http://sci2s.ugr.es/keel/dataset.php?cod=163>.
- [2] M. Barbaro and T. Zeller. A face is exposed for AOL searcher no. 4417749. *New York Times*, 9:2008, 2006.
- [3] S. Bashir, U. Qamar, and M. Younus Javed. An ensemble based decision support framework for intelligent heart disease diagnosis. In *Information Society (i-Society), 2014 International Conference on*, pages 259–264, Nov 2014.
- [4] N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi. Optimal geo-indistinguishable mechanisms for location privacy. *CCS '14*, pages 251–262, New York, NY, USA, 2014. ACM.
- [5] M.-C. Chen, H.-C. Liao, and C.-L. Huang. Predicting breast tumor via mining dna viruses with decision tree. volume 5, pages 3585–3589, Oct 2006.
- [6] C. Cumby, A. Fano, R. Ghani, and M. Krema. Predicting customer shopping lists from point-of-sale purchase data. *KDD '04*, pages 402–409, New York, NY, USA, 2004. ACM.
- [7] R. Dewri. Local differential perturbations: Location privacy under approximate knowledge attackers. *IEEE Transactions on Mobile Computing*, 12(12):2360–2372, Dec. 2013.
- [8] C. Dwork. An ad omnia approach to defining and achieving private data analysis. *PinKDD'07*, pages 1–13, Berlin, Heidelberg, 2008.
- [9] W. Enck, P. Gilbert, B.-G. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. Taintdroid: An information-flow tracking system for realtime privacy monitoring on smartphones. *OSDI'10*, pages 1–6. USENIX Association, 2010.
- [10] B. C. M. Fung, K. Wang, R. Chen, and P. S. Yu. Privacy-preserving data publishing: A survey of recent developments. *ACM Comput. Surv.*, 42(4):14:1–14:53, June 2010.
- [11] M. Götz, S. Nath, and J. Gehrke. Maskit: Privately releasing user context streams for personalized mobile applications. *SIGMOD '12*, pages 289–300, New York, NY, USA, 2012. ACM.
- [12] J. Han, M. Kamber, and J. Pei. *Data Mining: Concepts and Techniques*. San Francisco, CA, USA, 3rd edition, 2011.
- [13] Z. He, Z. Cai, and X. Wang. Modeling propagation dynamics and developing optimized countermeasures for rumor spreading in online social networks. In *Distributed Computing Systems (ICDCS), 2015 IEEE 35th International Conference on*, pages 205–214, June 2015.
- [14] S.-S. Ho and S. Ruan. Differential privacy for location pattern mining. *SPRINGL '11*, pages 17–24, New York, NY, USA, 2011. ACM.
- [15] M. Kosinski, D. Stillwell, and T. Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 110(15):5802–5805, 2013.
- [16] M. Kosinski, D. Stillwell, and T. Graepel. Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, February 2013.
- [17] N. Li, T. Li, and S. Venkatasubramanian. t-closeness: Privacy beyond k-anonymity and l-diversity. pages 106–115, 2007.
- [18] T. Li, N. Li, J. Zhang, and I. Molloy. Slicing: A new approach for privacy preserving data publishing. *IEEE Trans. on Knowl. and Data Eng.*, 24(3):561–574, Mar. 2012.
- [19] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkatasubramanian. L-diversity: Privacy beyond k-anonymity. *ACM Trans. Knowl. Discov. Data*, 1(1), Mar. 2007.
- [20] A. Narayanan and V. Shmatikov. Robust de-anonymization of large sparse datasets. *SP '08*, pages 111–125, Washington, DC, USA, 2008. IEEE Computer Society.
- [21] S. Salamatian, A. Zhang, F. Calmon, S. Bhamidipati, N. Fawaz, B. Kveton, P. Oliveira, and N. Taft. How to hide the elephant or the donkey in the room: Practical privacy against statistical inference for large data. pages 269–272, Dec 2013.
- [22] R. Shokri, G. Theodorakopoulos, C. Troncoso, J.-P. Hubaux, and J.-Y. Le Boudec. Protecting location privacy: Optimal strategy against localization attacks. *CCS '12*, pages 617–627, New York, NY, USA, 2012.
- [23] J. Stevens. *Applied Multivariate Statistics for the Social Sciences*. L. Erlbaum Associates Inc., Hillsdale, NJ, USA, 1986.
- [24] L. Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 10(5):557–570, 2002.
- [25] L. Sweeney. K-anonymity: A model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 10(5):557–570, Oct. 2002.
- [26] Y. Wang, Z. Cai, G. Yin, Y. Gao, X. Tong, and G. Wu. An incentive mechanism with privacy protection in mobile crowdsourcing systems. *Computer Networks*, 102:157 – 171, 2016.
- [27] Y. Zhang and M. Pennacchiotti. Predicting purchase behaviors from social media. *WWW '13*, pages 1521–1532, 2013.