Tsinghua Science and Technology

Volume 19 | Issue 5 Article 4

2014

Privacy Quantification Model Based on the Bayes Conditional Risk in Location-Based Services

Xuejun Zhang
Department of Computer Science and Technology, the Shaanxi Province Key Laboratory of Computer Network, Xi'an Jiaotong University, Xi'an 710049, China. School of Electronic and Information Engineering, Lanzhou Jiaotong University, Lanzhou 730070, China.

Xiaolin Gui
Department of Computer Science and Technology, the Shaanxi Province Key Laboratory of Computer Network, Xi'an Jiaotong University, Xi'an 710049, China.

Feng Tian
Department of Computer Science and Technology, the Shaanxi Province Key Laboratory of Computer Network, Xi'an Jiaotong University, Xi'an 710049, China.

Si Yu
Department of Computer Science and Technology, the Shaanxi Province Key Laboratory of Computer Network, Xi'an Jiaotong University, Xi'an 710049, China.

Jian An
Department of Computer Science and Technology, the Shaanxi Province Key Laboratory of Computer Network, Xi'an Jiaotong University, Xi'an 710049, China.

TSINGHUA SCIENCE AND TECHNOLOGY
ISSN 1007-0214 04/13 pp452-462
Volume 19, Number 5, October 2014

Privacy Quantification Model Based on the Bayes Conditional Risk in Location-Based Services

Xuejun Zhang*, Xiaolin Gui, Feng Tian, Si Yu, and Jian An

Abstract: The widespread use of Location-Based Services (LBSs), which allows untrusted service providers to collect large quantities of information regarding users' locations, has raised serious privacy concerns. In response to these issues, a variety of LBS Privacy Protection Mechanisms (LPPMs) have been recently proposed. However, evaluating these LPPMs remains problematic because of the absence of a generic adversarial model for most existing privacy metrics. In particular, the relationships between these metrics have not been examined in depth under a common adversarial model, leading to a possible selection of the inappropriate metric, which runs the risk of wrongly evaluating LPPMs. In this paper, we address these issues by proposing a privacy quantification model, which is based on Bayes conditional privacy, to specify a general adversarial model. This model employs a general definition of conditional privacy regarding the adversary's estimation error to compare the different LBS privacy metrics. Moreover, we present a theoretical analysis for specifying how to connect our metric with other popular LBS privacy metrics. We show that our privacy quantification model permits interpretation and comparison of various popular LBS privacy metrics under a common perspective. Our results contribute to a better understanding of how privacy properties can be measured, as well as to the better selection of the most appropriate metric for any given LBS application.

Key words: location-based services; Bayes decision estimator; privacy metric; adversarial model

1 Introduction

With the development of wireless communication technologies and mobile devices equipped with GPS chips, Location-Based Services (LBSs) provide personalized services to mobile users based on their geographical locations such as finding nearby Points Of Interest (POIs), tracking their physical fitness, and obtaining contextual information about their surroundings. Despite the great convenience LBSs providing to users, they also raise major privacy concerns when location information has to leave local devices to untrusted LBS provider. As location information included in the LBS queries is known to reveal the user's sensitive private information such as their home and work location, sexual preferences, political views, religious inclinations, and health conditions. As a result, a large number of LBS Privacy Protection Mechanisms (LPPMs) have been proposed to allow users to make use of the LBSs while mitigating concerns[1-6]. With the aim of assessing the effectiveness of these LPPMs, numerous privacy metrics have been investigated in prior research, for example, location k-anonymity,

• Xuejun Zhang, Xiaolin Gui, Feng Tian, Si Yu, and Jian An are with the Department of Computer Science and Technology, the Shaanxi Province Key Laboratory of Computer Network, Xi'an Jiaotong University, Xi'an 710049, China. E-mail: xlgui@mail.xjtu.edu.cn; tfft@stu.xjtu.edu.cn; yusi.computer@stu.xjtu.edu.cn; anjian@mail.xjtu.edu.cn.
• Xuejun Zhang is also with the School of Electronic and Information Engineering, Lanzhou Jiaotong University, Lanzhou 730070, China. E-mail: zxjlyl_new@stu.xjtu.edu.cn.
* To whom correspondence should be addressed.
Manuscript received: 2014-04-26; revised: 2014-07-14; accepted: 2014-08-13

location ℓ-diversity, location entropy, and expected estimator error. However, most of them have been specific to concrete systems and adversarial models and are difficult to generalize to other contexts. Once the adversarial model changes, the privacy metric is not sufficient for evaluating the effectiveness of the corresponding LPPMs. Furthermore, relationships between these different metrics are not investigated in depth, which results in a fragmented understanding of how privacy properties can be measured. Thus, there is a possibility of inappropriate privacy metrics being selected, which brings with it the risk of wrongly assessing users' privacy.

With the aim of addressing these issues, we propose a formal privacy quantification model for LBS privacy protection systems, which specifies a general adversarial model and uses the Bayes conditional risk as a privacy metric. We formalize the problem of establishing a unified measurement of privacy as an instance of a Bayes decision. In this decision, a user and an adversary interact strategically with each one's gain being the other's loss. On the one hand, the user strategically selects the optimal decision rule to maximize his privacy. On the other hand, the adversary, depending on his background knowledge, selects the optional decision rule to minimize the user's privacy. In the model, we formalize a generic adversarial model by his knowledge and inference attacks, and utilize the adversarial distortion function (the Bayes conditional risk) to establish the relationships between various popular LBS privacy metrics. We show that most LBS privacy metrics, such as location k-anonymity, ℓ-diversity, t-closeness, and ε-differential privacy, as well as location entropy, min-entropy, and mutual information or Kullback-Leibler (KL) divergence may be construed as particular cases of the Bayes conditional risk. Our major contributions are as follows:

(1) We provide a privacy quantification model that specifies a generic adversarial model and lays the foundation for the establishment of a unified measurement of privacy.

(2) We rely on well-established statistical methods and information theories to interpret that several well-known LBS privacy metrics may be construed as particular cases of the Bayes conditional risk metric.

2 Related Works

In the LBS privacy protection community, numerous privacy metrics have been proposed to measure users' privacy. Among them, location k-anonymity is the most popular metric, and was first introduced by Gruteser and Grunwald[1] by extending the concept of k-anonymity in database privacy[7]. This metric refers to the situation in which the location precision contained in an LBS query is decreased to a much larger area where the query sender is indistinguishable from at least $k-1$ other users also present in that area. Because of its simplicity, location k-anonymity has been widely adopted in many different LPPMs, including spatial cloaking[2] and incremental clique-base cloaking[3]. To better quantify the privacy, Xue et al.[8] refined location k-anonymity and introduced location diversity to ensure that the cloaking region contains at least ℓ semantic locations. However, deeper understanding of location k-anonymity reveals its drawbacks. Shokri et al.[9] analyzed the effectiveness of location k-anonymity in different scenarios with regard to the adversary's background information and showed its flaws which the adversary can utilize to infer the user's current location. Riboni et al.[10] identified the threat in terms of location k-anonymity in recurrent LBS requests application and made use of the t-closeness metric to guarantee that the distance between the distributions over the queries from an issuer's cloaking region and that of the entire region is below a certain threshold. Freudiger et al.[11] used Kullback-Leibler divergence between prior and posterior distributions to measure the ability of the adversary to guess the probability of each user visiting specific POIs. This metric may be regarded as an averaged version of t-closeness and shown to be equal to the mutual information metric[12]. Inspired by this measure, Zhang et al.[4] proposed an LBS privacy-quantifying framework and used the mutual information metric to measure adversary's information gain in his inference attacks. Location entropy, stemming from Shannon entropy in information theory, is another widely used metric for measuring the uncertainty associated with locations contained in LBS queries. This metric quantifies the information that an adversary can obtain from one (or a series) of location update(s)[5, 13]. Hoh and Gruteser[6] indicated that the location entropy calculated using probabilities for different user identity assignments to observed locations in the queries is not sufficient because the metric does not contemplate differences among these locations as reported to the adversary. They used the expected distance error to quantify the degree of accuracy by which an adversary

can estimate the real location of a user by observing the obfuscated location and utilizing background knowledge. Shokri et al.[14] proposed a similar but more general use of the adversary's expected estimation error to quantify LBS privacy, taking into account the adversary's knowledge of user mobile patterns, LBS access pattern, and internal LPPM algorithm. Because of insensitivity to background knowledge, differential privacy[15] has been growing as a popular method for quantifying LBS privacy in recent years, compared to k-anonymity-based approaches. Dewri[16] proposed a mix of differential privacy and k-anonymity by fixing an anonymity set of k locations and requiring that the probability to report the same obfuscation location from any of these k locations should be similar. This method quantifies the privacy in terms of the relationship between an adversary's knowledge and their geographic uncertainty in correctly locating a user. However, no privacy is guaranteed outside the anonymity set. Andrés and Bordenabe[17] proposed a formal notion of geo-indistinguishability for location-based systems that protect users' exact locations while avoiding the need to fix an anonymity set by using the generalized variant of differential privacy from Ref. [18].

In this paper, we borrow the idea of Ref. [19] and use the Bayes decision theory to investigate relationship between these metrics in depth.

3 Privacy Qualification Model

In this section, we first present our formal framework and then define the adversarial mode, privacy evaluation metrics, and decision rules for both the user and adversary. Last, we provide a simple example of our formulation.

3.1 Formal framework

We consider N mobile users, who move within a geographical area partitioned into M discrete regions. Each user potentially employs their location-aware devices to submit local search queries to an untrusted LBS provider, at some frequency. We assume that time is discrete and is split into different time periods (for example, morning, afternoon). We define an LBS actual query as $q_{\mathrm{a}} = (u, r, t, C)$, where u is an issuer's id, r is the position where query q_{a} is issued at given time t, and C is the content of a query.

To protect their privacy, mobile users decide to choose an optimal LPPM that transforms each q_{a} into a distortion, q_{d}, which is then sent to service provider. We define an LBS distortion query as $q_{\mathrm{d}} = (u', r', t, C)$, where u' is an issuer's pseudonym, which can be null when an issuer's id is removed from his queries without being replaced by one of his pseudonyms; r' is a distortion position which is distorted by a selected LPPM, for instance, a pseudo-location or large cloaking region. This transformation is made according to a probability distribution $f(q_{\mathrm{d}} \mid q_{\mathrm{a}}) = \Pr(q_{\mathrm{d}} \mid q_{\mathrm{a}})$.

The adversary decides to run an optimal inference attack on distortion q_{d} to output estimations $\hat{q}_{\mathrm{a}}$ of q_{a}. We model the adversary's strategy as a probability distribution: $h(\hat{q}_{\mathrm{a}} \mid q_{\mathrm{d}}) = \Pr(\hat{q}_{\mathrm{a}} \mid q_{\mathrm{d}})$. This distribution describes the probability that, given observation q_{d}, estimated private information $\hat{q}_{\mathrm{a}}$ corresponds to the user's actual query q_{a}. We measure the privacy metric as the adversary's expected error in this estimation, $\hat{q}_{\mathrm{a}}$, given actual query q_{a}.

We abstract this process and propose a general framework that lays the Bayes decision principles for establishing a unified measurement of privacy, as shown in Fig. 1. The framework includes three key actors: (1) a mobile user, who wishes to use LBSs while preserving private information, (2) a (trusted) system, which utilizes LPPMs to guarantee the privacy of the mobile user, and (3) an untrusted LBS provider (who may also be the adversary), who strives to disclose the

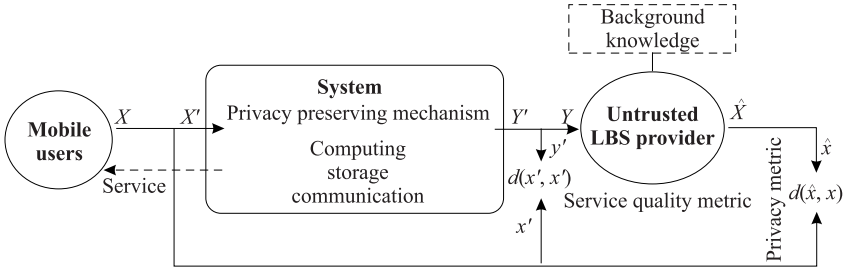

Fig. 1 LBS privacy quantification model based on the Bayes conditional risk.

mobile user's private information.

Throughout the paper, the measurable space where random variables take on values will be called an alphabet. We use italic uppercase letters to denote random variables, and lowercase letters to denote the particular values they take on. Probability Density Functions (PDFs) are represented by Pr and are subindexed by the corresponding random variable in case of ambiguity. For instance, $\Pr_X(x)$ denotes the value of function $\Pr_X$ at x, $\Pr_{X|Y}(x|y)$ denotes conditional probability values for x, given a status of nature in y.

Let X be the mobile user's private information that the adversary wishes to ascertain. The system's input is denoted by X', which refers to the user data required by the system to make a decision, in other words, a user's actual queries q_a. The system's decision is represented as Y', which refers to disclosed information — perhaps part of X' or a perturbation version of X', in other words, user's distortion query q_d. The adversary's input is denoted by Y, which models the information that serves as input for the adversary to ascertain X. In some cases, Y may directly be the information revealed by the system, in other words, $Y = Y'$. That is, the only information available to the adversary is exactly that disclosed by the system. In other circumstances, the adversary may observe a perturbed version of Y' — perhaps with any available background knowledge; in such cases, we have $Y \neq Y'$. The adversary's decision is denoted by $\hat{X}$, which represents the adversary's estimation of X from Y.

To simplify the discussion, we have $X = X'$ and $Y = Y'$ in the rest of paper unless otherwise specified.

3.2 Adversarial model

To compare various privacy metrics under the common conditions, a general adversarial model is necessary and an important element of the privacy quantification framework. We now proceed to present the adversarial model, which is characterized by his knowledge and attack(s).

3.2.1 Background knowledge

We assume an adversary can observe all the distortion Y' outputted by the system; furthermore, the adversary might have side information about X, as well as the knowledge of the decision rules used by the system. The side information can be modeled by a prior distribution, $\Pr_X(x)$, and the decision rules are made according to a distribution $\Pr_{Y|X}(y|x)$. Thus, the adversary can use the Bayes rule to induce a posterior distribution on the user's private information, x, conditional on observation y:

$$\Pr_{X|Y}(x|y) = \frac{\Pr_{Y|X}(y|x)\Pr_X(x)}{\sum_x \Pr_{Y|X}(y|x)\Pr_X(x)} \quad (1)$$

3.2.2 Inference attacks

Given observation on y, prior distribution $\Pr_X(x)$, and decision function $\Pr_{Y|X}(y|x)$, the adversary executes an inference attack to obtain estimations, $\hat{x}$, of the user's privacy information, x.

In this spirit, we consider an adversary who uses a Bayes decision estimator to infer users' privacy information. Namely, the adversary employs minimum-risk Bayes estimator or minimum-error-rate Bayes estimator to make a decision, $\hat{x}$, on X for every possible system decision, resulting in an observation y. Obviously, the adversarial decision will be accompanied by a cost, which is captured by the loss function $d_p: (\hat{x}, x) \to d_p(\hat{x}, x)$, which quantifies the adversary's loss caused by making decision $\hat{x}$ when the true state is x. The privacy loss depends on the location's semantics and also on the user's privacy requirements, and $d_p(\cdot)$ must be defined accordingly. For instance, if the user wants to hide just her exact current location, the appropriate distortion function could be the Hamming distance $d_p(x, \hat{x}) = 1$ if $x = \hat{x}$, otherwise $d_p(x, \hat{x}) = 0$. Alternatively, the user's privacy might depend on the physical distance, in other words, the squared-error distortion $d_p(x, \hat{x}) = (x - \hat{x})^2$.

We define the adversarial average loss associated with this decision as the Bayes conditional risk for estimator $\hat{x}$ in the discrete case:

$$R_c = E[d_p(\hat{x}, X)|y] = \sum_x d_p(\hat{x}, x)\Pr_{X|Y}(x|y) \quad (2)$$

where the expectation is taken under posterior distribution $\Pr_{X|Y}(x|y)$. Thus, the Bayes unconditional risk associated with that estimator is defined as the expectation of Bayes conditional risk over all possible observation y:

$$R_{uc} = E(R_c) = \sum_y \Pr_y(y) \sum_x d_p(\hat{x}, x)\Pr_{X|Y}(x|y) \quad (3)$$

where the expectation is taken under the jointly probability distribution of x and y.

The adversary's object is to choose $\hat{x}$ to minimize the Bayes unconditional risk among all possible estimators. For minimum-risk Bayes estimator, the

adversary's optimal estimator is to choose $\hat{x}$ to minimize the R_{uc}. It turns out that this optimal estimator is precisely.

$$\hat{x}_{\mathrm{Bayes}} = \arg\min_{\hat{x}}(R_{\mathrm{c}}) \tag{4}$$

for all y, in other words, the Bayes estimator is the one that minimizes the Bayes unconditional risk for every observation. For minimum-error-rate Bayes estimation, the adversary's optimal estimator is to select $\hat{x}$ to maximize the posterior distribution $\Pr_{X|Y}(x|y)$, which makes the average error rate minimum. In essence, minimum-error-rate Bayes estimation is maximum *a posterior* estimator (MAP). Mathematically,

$$\hat{x}_{\mathrm{MAP}} = \arg\max_{\hat{x}} \Pr_{X|Y}(x \mid y) \tag{5}$$

From Eqs. (2) and (5), we conclude that the Bayes estimator and MAP estimator coincide when the loss function is Hamming distance.

3.3 Evaluation metrics

Privacy metrics, accompanied with utility/service-quality metrics, provide a quantitative means of comparing the suitability of two or more LPPMs regarding the trade-off between privacy and utility. Therefore, we describe two main evaluation metrics to evaluate the utility loss caused by the distortion of the original query and the cost of the decision made by the adversary: the LBS privacy of users under inference attacks and the service-quality loss that they incur by using an LPPM.

3.3.1 Service quality metric

The LBS response quality depends on the system's decision Y instead of X. The distortion incurred by Y determines the service quality that the user experiences. We use distortion function $d_{\mathrm{q}} : (x, y) \rightarrow d_{\mathrm{q}}(x, y)$ to reflect the system's loss of quality due to an LPPM. The average quality distortion is computed as

$$Q = E\big[d_{\mathrm{p}}(X, Y)\big] = \sum_{x,y} \Pr_X(x)\Pr_{Y|X}(y \mid x) d_{\mathrm{p}}(x, y) \tag{6}$$

Function $d_{\mathrm{p}}(\cdot)$ determines the dissimilarity between actual location x and distorted location y. The semantics of this dissimilarity depend on the LBS under consideration, and also on the user's specific service-quality expectations. In many applications, the service-quality can be considered inversely proportions to the physical distance between x and y. For instance, applications of finding nearby POIs could have very different responses to x and to y even if they are only a few kilometers apart. On the contrary, in some LBS applications, the service-quality depends on other criteria, such as on whether y is within a region of interest. For a weather forecast application, for example, any distorted location y in the same city as the actual location x would result in a high quality LBS response.

3.3.2 Privacy metric

We quantify the user's privacy as the adversary's expected error in his Bayes error, denoted by Bayes conditional risk $P_{\mathrm{cond}} = R_{\mathrm{c}}$, which is the estimation error incurred by the adversary, conditioned on observation y.

In addition, we define worst-case privacy as $P_{\mathrm{worst}} = \min_y(R_{\mathrm{c}})$, and average-case privacy as $P_{\mathrm{avg}} = E(R_{\mathrm{c}}) = R_{\mathrm{uc}}$, which is the expectation of the adversary's estimation error over all possible observation y.

3.4 Decision strategies

In the decision process, the system selects the decision rule $\Pr_{Y|X}(y \mid x)$ that maximizes either the average-case privacy or the worst-case privacy without allowing the average-case privacy distortion to exceed a certain threshold. On the other hand, the adversary selects a Bayes estimator, which leads to the minimization of both measure of privacy.

Note that in a special cases, where the unknown variable X models the user's identity, our measure of privacy may be regarded as a measure of anonymity.

3.5 Example

Next, we present a simple example to illuminate the above formulation.

As a running example, we consider a user located at the location x who wishes to query an untrusted LBS provider for nearby restaurants in a private way. For the sake of simplicity, we consider $X = X'$ and assume that X is a binary random variable with $\Pr\{X = 0\} = \Pr\{X = 1\} = 1/2$. Namely, the system's input is the user's private information that needs to be protected. To avoid disclosing his exact location x, the user is solely responsible for protecting his location data by using the location perturbation $\Pr_{Y|X}$ (for example, adding noise). That is, the user becomes the system. Note that location perturbation presents the inherent trade-off between data privacy and utility.

According to randomized location perturbation rule $\Pr_{Y|X}$, the user could disclose a perturbed location, y', for each possible location data x. That is, the

user could decide to disclose the perturbed location of x with probability $\Pr_{Y|X}$, whereas no perturbation could be applied (for example, $y' = x$) with the probability $1 - \Pr_{Y|X}$. Note that, in this example, the system's decision rule is completely determined by $\Pr_{Y|X}$. Last, we assume that the adversary can only access to the disclosed information Y', namely, $Y = Y'$. Moreover, we consider the attacker's distortion function to be Hamming distance. This implies that the Bayes estimator matches the MAP estimator. Based on these reveal information and observation, the adversary uses a Bayes estimator to ascertain the user's actual location, X. It is easy to demonstrate that the attacker's best decision is $\hat{X} = Y$. Therefore, average privacy P_{avg} becomes

$$P_{\text{avg}} = \Pr\{X \neq \hat{X}\} = \Pr\{X \neq Y\} = \Pr\{X \neq Y'\} = \Pr_{Y|X}.$$

Similarly, we assume that the system's distortion function is also the Hamming distance. Therefore, average quality distortion Q becomes

$$Q = \Pr\{X' \neq Y'\} = \Pr\{X \neq Y'\} = \Pr_{Y|X}.$$

To describe the adversary's strategy, we define the average utility U as $1-Q$. According to this, the system would strive to maximize the average privacy with respect to $\Pr_{Y|X}$, subject to the constraint $U \geqslant u_0$. As can be seen from Fig. 2, the optimal value of average privacy is $P_{\text{avg}} = 1 - u_0$, for $1/2 < u_0 \leqslant 1$.

4 Theoretical Analysis

In the framework, we define our privacy metric as Bayes conditional risk P_{cond}. Of particular importance due to its generality and applicability is the adversary's distortion function $d_{\text{p}}(\cdot)$, which quantifies the adversary's performance[14]. We use the generalization of $d_{\text{p}}(\cdot)$ to establish a connection between most popular LBS privacy metrics and our metric. First, we contemplate the case where the adversarial distortion function $d_{\text{p}}(\cdot)$ is the Hamming distance. Second, we consider the more general case where $d_{\text{p}}(\cdot)$ can be any other distance function. In the special case of its being Hamming distance, we consider applications for both sporadic and recurrent LBSs. In the sporadic LBS scenario, we relate our metric to Hartley's entropy, location k-anonymity, ℓ-diversity, and min-entropy using MAP estimation, Bayes decision theory, and confidence set. In the recurrent LBS scenario, we use asymptotic equipartition property[20] to relate Shannon's entropy to our metric.

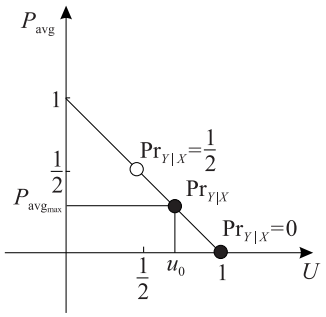

Fig. 2 The trade-off curve between the privacy and utility.

In the more general case where the adversary's distortion function $d_{\text{p}}(\cdot)$ is not Hamming distance, we examine two possible scenarios. First, we consider that the adversary's distortion function is known to the system. Assuming a Bayes attacker's strategy, we use the Bayes decision theory to justify the system's best decision rule. Second, we consider that the adversary's distortion function is unknown to the system. In this scenario, we use the concept of total variation to connect our metric to several privacy metrics, provided that the attacker uses the MAP estimation.

4.1 Hamming distortion distance

This section analyzes the connections between some popular privacy metrics and our metric when the adversary's distortion function is Hamming distance.

4.1.1 Sporadic LBS

The concept of the confidence set enables us establish a connection between our metric and some of the most popular privacy metrics[20]. A confidence set with confidence level p is a subset $C(X)$ of the alphabet X that depends only on variable X, such that $\Pr\{X \in C(X)\} = p$. In the case of continuous-valued random scalars, confidence sets commonly take the form of intervals. With this conception, we consider an adversarial model where the adversary only contemplates the shape of the PMF of the user's private information, X, to identify a confidence set for some desired confidence, p; beyond that, assume all included members equally relevant. MAP estimation within this set, provided it is uniformly distributed, leads to an estimation error of $1 - \dfrac{1}{|C(X)|}$.

Furthermore, we use the Rényi entropy as a measure of uncertainty. Rényi entropy, as the more general definition of information measures, generalizes Hartley entropy, Shannon entropy, and min-entropy. Specifically, for $\alpha > 0$ and $\alpha \neq 1$, Rényi

entropy of order α of discrete random variable X with alphabet $\chi = \{x_1, x_2, \cdots, x_n\}$ is defined as

$$H_\alpha(X) = \frac{1}{1-\alpha} \log \Big(\sum_{i=1}^{n} P_X(x_i)^\alpha \Big).$$

First, in the limit for $\alpha \to 0$, H_α is just the logarithm of the size of X' support set. Provided the probabilities are nonzero, the Rényi entropy is Hartley entropy:

$$H_0(X) = \log |X| = \log n \tag{7}$$

Second, as $\alpha \to 1$, the limiting value of H_α is Shannon's entropy:

$$H_1(X) = -\sum_i \Pr_x(x_i) \log \Pr_x(x_i) \tag{8}$$

Last, as $\alpha \to \infty$, the limiting value of H_α converges to min-entropy:

$$H_\infty(X) = \min_i[-\log \Pr_x(x_i)] = -\log[\max_i \Pr_x(x_i)] \tag{9}$$

As $H_0(X) \geqslant H_1(X) \geqslant H_\infty(X)$ with equality—if and only if X is uniformly distributed—we define them as best-case, average-case, and worst-case measurements of privacy, respectively. Specifically, Hartley entropy $H_0(X)$ measures the cardinality of all possible values of X regardless of their probabilities, whereas Shannon entropy $H_1(X)$ is a weighted average of such logarithms, and the $H_\infty(X)$ is the minimum of the self-information $-\log \Pr_X(x_i)$.

In the following, we use Rényi entropy to explore the relations between our metric and other metrics. Location k-anonymity is a popular privacy metric among the LBS privacy protection researchers. With this technique, users' private information X are submitted to the LBS provider via a privacy proxy. The privacy proxy performs a cloaking procedure to augment a mobile user's query location to a cloaking region (which geographically covers not only the user who issues the query but also at least $k-1$ other users) and then transmits the distortion query Y to the LBS provider. Since all users in the cloaking set report the same cloaking regions in their distortion queries Y', the adversary cannot distinguish the location or query attribute of any user from the received queries. This suggests that the adversary has no knowledge from which to infer user's private information. Mathematically, this is reflected by assuming that the posterior probability $\Pr_{X|Y}(\cdot|\, y)$, with respect to sending any queries, has uniform distribution across all users. In reality, the adversary may have some background knowledge from which to infer the real issuer. Therefore, we consider the more general case in which Y consists of Y' plus any background knowledge. We assume that the adversary uses a MAP estimator to infer user's LBS privacy. Under this model, our Bayes conditional privacy P_{cond} becomes

$$P_{\text{cond}} = \Pr_{X|Y}\{(X \neq \hat{x})|y\} = 1 - \max_x \Pr_{X|Y}(x|y) \tag{10}$$

which is the MAP estimation error e_{MAP}, conditioned on observation y. We substitute Eq. (9) into Eq. (10) and obtain

$$P_{\text{cond}} = e_{\text{MAP}} = 1 - 2^{\log \max_x \Pr_{X|Y}(x|y)} = 1 - 2^{-H_\infty(X|x)},$$

which shows that the concept of min-entropy is related to MAP estimator. If we apply aforementioned uniformity condition of $\Pr_{X|Y}(\cdot\,|\,y)$ and assume that PMF is the uniform distribution on a group of exactly k individuals in k-anonymity, the adversary cannot link the distortion query to the actual issuer with any probability larger than $1/k$, and then the user's privacy is

$$P_{\text{cond}} = 1 - \frac{1}{k} = 1 - 2^{\log |k|} = 2^{-H_0(X|y)},$$

which expresses the conditional privacy with respect to Hartley's entropy. This means that the k-anonymity metric may be a special case of our privacy measure, determined by this Rényi entropy.

As mentioned in Section 2, the location k-anonymity metric has some limitations and is insufficient at ensuring privacy in some cases. To address these limitations, several privacy metrics have been proposed. In the remainder of this section, we mainly focus on the ℓ-diversity metric[8], which builds on the k-anonymity principle and aims at conquering the problem of group privacy violations. As can be shown in Fig. 3a, although k-anonymity is satisfied, the group privacy is violated when all users inside Acme's headquarters are (with high probability) Acme's employees. Because they belong semantically to the same group. Motivated by this, the conception of location ℓ-diversity has been proposed to improve

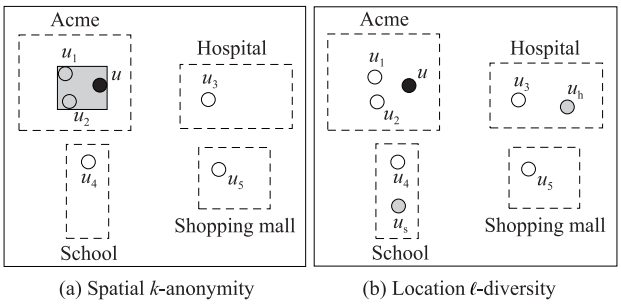

(a) Spatial k-anonymity (b) Location ℓ-diversity

Fig. 3 k-anonymity vs. location ℓ-diversity.

upon location k-anonymity by guaranteeing that each LBS query can be associated with at least ℓ semantically different locations. This is shown in Fig. 3b, where the proxy sends a packet of $\ell = 3$ queries originating at u, u_{h}, and u_{s}. The semantic locations of these users are Acme, Hospital, and School, respectively. Therefore, an attacker cannot identify with probability larger than $1/\ell$ that the query comes from Acme. The objective of this metric is to prevent the group users' query location from being disclosed. Therefore, we consider that the adversary's unknown X represents the query location. Other variables remain the same as in the previous interpretation. Given distortion query y, we assume that the ℓ-diversity requirement is met by ensuring that the probability distribution $P_{X|Y}(\cdot \mid y)$ of the query location with the group of semantic locations is the uniform distribution on a set of at least ℓ locations.

Last, we also suppose again that the adversary uses MAP estimation. Under the premise of a MAP attacker, our measure of Bayes conditional privacy boils down to a MAP error. If we also apply the assumption above about the uniformity of $P_{X|Y}(\cdot \mid y)$ and suppose that this distribution is uniform on a group of ℓ locations, then the Bayes conditional privacy yields $P = 1-1/\ell = 1-2^{-H_0(X|y)}$ which expresses our metric again in terms of Hartley's entropy. In short, the ℓ-diversity metric lends itself to be interpreted as a particular case of our more general privacy measure.

4.1.2 Recurrent LBS

In recurrent LBS applications, the k-anonymity or ℓ-diversity metric is not sufficient for describing users' privacy requirements[21]. This is due to the fact that the time-series associated with an LBS query contains user's location X and forms a trajectory $T = \{X_1, X_2, \cdots, X_n\}$, denoted by X^n, that may reveal the real identity of the issuer if, for example, it links to the user's home or office.

The location entropy metric, however, is effective for measuring the uncertainty associated with location information contained in LBS queries and for quantifying the information an adversary can obtain from a series of location updates. To interpret, under the perspective of our framework, the Shannon entropy as a measure of privacy, we consider an adversary utilizing the observations about the sequence Y^k of k perturbed locations that the system submits to the LBS provider to thereby determine sequence X^k of k unknown locations visited by user. Furthermore, the adversary's distortion function is Euclidean distance. Alternatively, if the adversary's interest in profiling a user's behavior lies in whether that user is at home, or the movies, or any given sensitive location, then Hamming distance is more appropriate. According to the concept of the asymptotic equipartition property[20], among all possible n^k sequences in which each sequence X^k contains k independent, identically distributed random variables that are drawn according to $\Pr_X(\cdot)$ with alphabet size n, there exists a typical subset $S(X)$ of sequences almost certain to occur. More precisely, for any $\varepsilon > 0$, there exists a k sufficiently large that $\Pr(S(X)) > 1-\varepsilon$ and $|S(X)| \leqslant 2^{k(H_1(X)+\varepsilon)}$. A similar result occurs for those sequences X^k that are jointly typical with a given typical sequence Y^k. That is, the set of all these sequences X^k is defined as the conditionally typical set $S(X^k|Y^k)$ and satisfies that $\Pr(S(X^k|Y^k)) > 1-\varepsilon$ for large k; its cardinality is bounded by Shannon's conditional entropy, $|S(X^k|Y^k)| \leqslant 2^{k(H_1(X)+\varepsilon)}$. Further, it turns out that these conditionally typical sequences are likely approximately equal in the exponents with probability $2^{-kH_1(X|Y)}$ (Note that two sequences a_k and b_k are approximately equal in the exponent if $\lim_{k\to\infty} \frac{1}{k} \log \frac{a_k}{b_k} = 0$). While the most likely sequence may in fact not belong to the typical set, the set of typical sequences involves a sufficiently large number of sequences that is equivalent to a probability arbitrarily close to certainty.

Leveraging the above consideration, we regard $S(X^k|Y^k)$ as set of arbitrarily high confidence with cardinality $2^{kH_1(X|Y)}$, approximately equal in the exponent. This shows that the Shannon entropy of unknown X, given observation Y, is an approximate measure of the size of a high confidence set—a measure suitable for adversarial models based on the estimation of sequence rather than individual samples. In reality, the most likely sequence may in fact be atypical; thus, Shannon entropy is not directly applicable to MAP estimation over the entire set of sequences. Nevertheless, because the most likely sequence is simply a repetition of the most likely symbol, MAP estimation on sequence is a trivial extension of the argument on min-entropy.

4.2 Other distortion distance

In this section, we consider the two possible alternatives regarding the system's knowledge of the distortion function: either $d_{\mathrm{p}}(\cdot)$ is known to the system or it is

not. Under the former assumption, the system utilizes the Bayes estimator to find optimal decision rules $\Pr_{Y|X}(y|x)$ that maximize either worst-case privacy or average-case privacy and satisfies constraints on average distortion. The latter assumption, however, describes a more general and realistic scenario. The reminder of this subsection precisely interprets several privacy metrics under this assumption. The only piece of information that is known to the system is the maximum values attained by function, namely, $d_{\max} = \max_{x,\hat{x}}(d_{\mathrm{p}}(x,\hat{x}))$.

4.2.1 KL divergence and mutual information

In Ref. [11], the notion of KL divergence is used to measure the ability of the adversary to guess the probability of each user's visiting specific POIs in an LBS scenario. Under the assumption that the adversary uses a MAP strategy, we use the concept of total variation[20] to relate Bayes conditional privacy to the KL divergence as:

$$\Delta P = E_{\mathrm{po}}\{d_{\mathrm{p}}(X,\hat{x}_{\mathrm{pr}})\} - E_{\mathrm{po}}\{d_{\mathrm{p}}(X,\hat{x}_{\mathrm{po}})\} \leqslant 2\sqrt{2}d_{\max}\sqrt{D_{\mathrm{KL}}\left(\Pr_{X|Y}(\cdot|y) \parallel \Pr_X(x)\right)} \tag{11}$$

where ΔP denotes the reduction in Bayes conditional privacy, which is determined by using prior distribution and posterior distribution; E_{po} denotes that the expectation has taken over the posterior distribution $\Pr_{X|Y}(\cdot\,|\,y)$; $\hat{x}_{\mathrm{pr}}$ and $\hat{x}_{\mathrm{po}}$ denote the adversary's estimation when using prior and posterior distribution; $D_{\mathrm{KL}}(\cdot)$ represents the KL divergence. From this inequality, we can see that the minimization of KL divergence leads to the minimization of an upper bound on ΔP. This allows us to establish a connection between our metric and t-closeness[10]. The t-closeness boils down to defining a maximum discrepancy between the prior and posterior distributions:

$$t = \max_y\{D_{\mathrm{KL}}\left(\Pr_{X|Y}(\cdot\,|\,y) \parallel \Pr_X(x)\right)\}.$$

Under this definition and on account of Formula (11), we can get

$$\Delta P \leqslant 2\sqrt{2}d_{\max}\sqrt{t}.$$

Therefore, t-closeness is essentially equivalent to bounding the decrease in Bayes conditional privacy.

In Ref. [12], privacy risk metric R is defined as the conditional KL divergence between the prior and posterior distributions, which turns out to coincide with the mutual information between the adversary's unknown X and observation Y:

$$R = E_Y\left\{D_{\mathrm{KL}}\left[\Pr_{X|Y}(\cdot\,|\,y) \parallel P_X(x)\right]\right\} = E\left\{\log\frac{P_{X|Y}(X|Y)}{P_X(X)}\right\} = I(X;Y) = H(X) - H(X|Y) = H(Y) - H(Y|X) \tag{12}$$

where $I(X;Y)$ denotes the mutual information between X and Y and $H(\cdot)$ denotes the entropy.

With Eq. (12), we can rewrite Formula (11) as

$$\frac{E(\Delta P^2)}{8d_{\max}^2} \leqslant I(X;Y) \tag{13}$$

As can be seen from Formula (13), the minimization of mutual information contributes to the minimization of an upper bound on ΔP. Concretely, to protect users' privacy, the system can compute users' optimal decision rules $\Pr_{Y|X}(y|x)$ with the aim of minimizing mutual information between X and Y and simultaneously need to ensure that the loss of service quality does not exceed a certain threshold. That is to say, the system strives to solve optimization problem:

$$\underset{Ed_{\mathrm{q}}(X,Y)\leqslant Q,\Pr_{Y|X}(y|x)\geqslant 0}{\text{Minimize}} \quad I(X;Y) \tag{14}$$

which bears a strong resemblance to the rate distortion problem and has been well solved in the field of information theory. The importance of this lies in the fact that some of the information theoretic results and methods for the rate distortion problem can be extended to the problem (14).

Note that the very same metric, or conceptually equivalent variations, may be interpreted under different perspectives. For instance, "mutual information" is the discrepancy between conditional entropy and unconditional entropy — effectively, the posterior uncertainty modeled by Shannon entropy, normalized with respect to its prior correspondence. Under this respective, mutual information might also be connected to Shannon's entropy.

4.2.2 Difference privacy

Researchers have indicated multiple times that any notion of privacy is incomplete without explicit statements on the capabilities of an adversary. Because of insensitivity to background knowledge, differential privacy has been growing in popularity to quantify LBS privacy. Andrés and Bordenabe[17] proposed a formal privacy definition of geo-indistinguishability for LBSs, as well as a randomized perturbation rule $f(Y|X) = \Pr(Y|X)$ that allows a user to disclose enough location information to obtain the desired service. Formally, we say that a perturbation rule satisfies ϵ-geo-indistinguishability if and only if any two different points of interest (typically the user's

possible locations) x, $\hat{x}$:

$$\sup |\ln \frac{\Pr(x)}{\Pr(\hat{x})}| \leqslant \epsilon d_{\mathrm{p}}(x, \hat{x}).$$

Equivalently, the definition can be formulated as

$$\Pr(y|x) \leqslant \mathrm{e}^{\epsilon d_{\mathrm{p}}(x,\hat{x})} \Pr(y|\hat{x}) \tag{15}$$

for all x, $\hat{x}$, and observation on y. If $d_{\mathrm{p}}(\cdot)$ is Hamming distance between databases x and $\hat{x}$ (in other words, the number of individuals in which they differ), the above definition is an instance of a generalized variant of differential privacy, using an arbitrary metric between secrets. After a slight manipulation to Formula (15), we may say that a randomized perturbation rule provides ϵ-differential privacy when

$$\epsilon = \max_{y,x,\hat{x}} \ln \frac{\Pr(y|x)}{\Pr(y|\hat{x})}.$$

Although this formulation does not quite match the problem in terms of prior and posterior distributions, this manipulation enables us to still establish a loose relation with our conditional privacy: $R \leqslant t \leqslant \epsilon$. Thus we may understand differential privacy metric as a worst-case privacy criterion. Note that although there is a formal similarity between the metrics, there are also substantial differences regarding assumptions, objectives, models, and privacy guarantees.

5 Conclusions

In the LBS scenario, there exists a number of privacy metrics for evaluating the effectiveness of various LPPMs. However, most of these privacy metrics have been conceived for concrete systems and adversarial models and are difficult to generalize into other contexts. Even for a specific application, we often find that various privacy metrics are available.

In this paper, we have presented a privacy quantification model for specifying the relationship between these different metrics in LBS applications, considering a general adversarial model, as well as the cost of the adversarial decision. The core of our proposal is a more general form of an adversary's distortion function $d_{\mathrm{p}}(\cdot)$, which quantifies the adversarial cost caused by making a decision. Based on this, we define a general LBS privacy metric that is denoted by Bayes conditional risk, meaning that our assessment is conditioned on an adversarial observation. Moreover, we employ some conception of information theory and Bayes decision theory to interpret and compare a variety of well-known LBS privacy metrics under a common perspective. Our analytic results shed new light on the understanding of those metrics and their suitability when it comes to applying them to specific applications and also assist designers in choosing a more grounded and systematic approach to accurately evaluating the effectiveness of the LPPMs.

Acknowledgements

This work was supported in part by the National Science and Technology Major Project (No. 2012ZX03002001-004), the National Natural Science Foundation of China (Nos. 61172090, 61163009, and 61163010), the PhD Programs Foundation of Ministry of Education of China (No. 20120201110013), and the Scientific and Technological Project in Shaanxi Province (Nos. 2012K06-30 and 2014JQ8322).

References

[1] M. Gruteser and D. Grunwald, Anonymous usage of location-based services through spatial and temporal cloaking, in *Proc. the First Int'l Conf. on Mobile Systems, Application, and Services*, San Francisco, USA, 2003, pp. 31-42.

[2] M. F. Mokbel, C. Y. Chow, and W. G. Aref, The new Casper: Query processing for location services without compromising privacy, in *Proc. the 32nd Int'l Conf. on Very Large Data Bases*, Seoul, Korea, 2006, pp. 763-774.

[3] X. Pan, J. L. Xu, and X. F. Meng, Protecting location privacy against location-dependent attacks in mobile services, *IEEE Transaction on Knowledge and Data Engineering,* vol. 24, no. 8, pp. 1506-1519, 2012.

[4] X. J. Zhang, X. L. Gui, Z. C. Feng, F. Tian, S. Yu, and J. Q. Zhao, A quantifying framework of query privacy in location-based service, (in Chinese), *Journal of Xi'an Jiaotong University,* vol. 48, no. 2, pp. 8-13, 2014.

[5] T. Xu and Y. Cai, Feeling-based location privacy protection for location-based services, in *Proc. the 16th ACM Conf. on Computer and Communications Security*, Chicago, USA, 2009, pp. 348-357.

[6] B. Hoh and M. Gruteser, Protecting location privacy through path confusion, in *Proc. the First Int'l Conf. on Security and Privacy for Emerging Areas in Communications Networks*, Boppard, Germany, 2005, pp. 194-205.

[7] L. Sweeney, k-anonymity: A model for protecting privacy, *International Journal of Uncertainty, Fuzziness and Knowledge-Based System,* vol. 10, no. 5, pp. 557-570, 2002.

[8] M. Q. Xue, P. Kalnis, and H. K. Pung, Location diversity: Enhanced privacy protection in location based services, in *Proc. the 4th Symposium on Location and Context Awareness*, Tokyo, Japan, 2009, pp. 70-87.

[9] R. Shokri, C. Troncoso, C. Diaz, J. Freudiger, and J.-P. Hubaux, Unraveling an old cloak: k-anonymity for location privacy, in *Proc. the 2010 ACM Workshop on Privacy in the Electronic Society*, Chicago, USA, 2010, pp. 115-118.

[10] D. Riboni, L. Pareschi, C. Bettini, and S. Jajodia, Preserving anonymity of recurrent location-based queries, in *Proc. the 16th Int'l Symposium on Temporal Representation and Reasoning*, Bressanone, Italy, 2009, pp. 62-69.
[11] J. Freudiger, R. Shokri, and J.-P. Hubaux, Evaluating the privacy risk of location-based services, in *Proc. the 15th Int'l Con. on Financial Cryptography and Data Security*, Gros Islet, Saint Lucia, 2012, pp. 31-46.
[12] D. Rebollo-Monedero, J. Forné, and J. Domingo-Ferrer, From t-closeness-like privacy to postrandomization via information theory, *IEEE Trans. on Knowledge and Data Engineering,* vol. 22, no. 11, pp. 1623-1636, 2010.
[13] B. Hoh, M. Gruteser, H. Xiong, and A. Alrabady, Achieving guaranteed anonymity in gps traces via uncertainty-aware path cloaking, *IEEE Trans. on Mobile Computing,* vol. 9, no. 8, pp. 1089-1107, 2010.
[14] R. Shokri, G. Theodorakopoulos, J.-Y. L. Boudec, and J.-P. Hubaux, Quantifying location privacy, in *Proc. the 32nd IEEE Symposium on Security and Privacy (SP)*, Oakland, USA, 2011, pp. 247-262.
[15] C. Dwork, Difference privacy, in *Proc. the International Colloquium on Automata, Languages and Programming*, Venice, Italy, 2006, pp. 1-12.
[16] R. Dewri, Local differential perturbations: Location privacy under approximate knowledge attackers, *IEEE Transactions on Mobile Computing,* vol. 12, no. 12, pp. 2360-2372, 2013.
[17] M. E. Andrés and N. E. Bordenabe, Geo-indistinguishability: Differential privacy for location-based system, in *Proc. the 20th ACM Conf. on Computer and Communications Security*, Berlin, Germany, 2013, pp. 901-914.
[18] K. Chatzikokolakis, M. E. Andrés, N. E. Bordenabe, and C. Palamidessi, Broadening the scope of differential privacy using metric, in *Proc. the Privacy Enhancing Technologies*, Bloomingto, USA, 2013, pp. 82-102.
[19] R. Shokri, G. Theodorakopoulos, C. Troncoso, J.-P. Hubaux, and J.-Y. L. Boudec, Protecting location privacy: Optimal strategy against localization attacks, in *Proc. the 19th ACM Conf. on Computer and Communication Security*, Raleigh, USA, 2012, pp. 617-626.
[20] D. Rebollo-Monedero, J. Parra-Arnau, C. Diaz, and J. Forné, On the measure of privacy as an attacker's estimation error, *International Journal of Information Security,* vol. 12, pp. 129-149, 2013.
[21] A. Pingley, N. Zhang, X. W. Fu, H.-A. Choi, S. Subramaniam, and W. Zhao, Protection of query privacy for continuous location based services, in *Proc. the 30th IEEE Communications Society Conf. on Computer Communications*, Shanghai, China, 2011, pp. 1710-1718.

Xuejun Zhang is currently a PhD candidate in Department of Computer Science and Technology, Xi'an Jiaotong University. He is also a lecturer in School of Electronic and Information Engineering, Lanzhou Jiaotong University. He was a visiting scholar at Wuhan University during 2010 to 2011. He received his MS degree from Southeast University in 2008. His research interests include location-based privacy protection and quantification and service-computing security.

Xiaolin Gui is currently a professor of School of Electronic and Information Engineering, Xi'an Jiaotong University, China. He is also the Director of the Shaanxi Province Key Laboratory of Computer Network and the Director of Research Institute of Computer Network and Engineering. His research covers secure computation of open network system including Grid, P2P, and cloud computing, dynamic trust management theory, data and privacy protection, and Internet of Things. He got his PhD, MEng, and BEng degrees from Xi'an Jiaotong University in 2001, 1993, and 1988, respectively. He has published over 130 academic papers and books. He was the recipient of New Century Excellent Talents in University of China in 2005.

Feng Tian is a PhD candidate in Department of Computer Science and Technology at Xi'an Jiaotong University. He got his BEng degree from Xi'an Jiaotong University in 2009. His research interests include cloud computing, location privacy protection, and data outsourcing security.

Si Yu is currently a PhD candidate at Xi'an Jiaotong University. He got his BEng degree from the Shaanxi University of Science and Technology in 2008. His research interests include cloud computing and virtualization security.

Jian An is currently a lecturer in Department of Computer Science and Technology, Xi'an Jiaotong University. He received his PhD degree from Xi'an Jiaotong University in 2013. He got his BEng and MEng degrees from Xinjiang University in 2006 and 2009, respectively. His research interests include services computing, social networking, and privacy and security.