Model for a Common Notion of Privacy Leakage on Public Database

Shinsaku Kiyomoto
KDDI R&D Laboratories, Inc.
2-1-15 Ohara, Fujimino, Saitama
356-8502, Japan
kiyomoto@kddilabs.jp

Keith M. Martin
Information Security Group,
Royal Holloway University of London
Egham, Surrey TW20 0EX, UK
keith.martin@rhul.ac.uk

Abstract

Privacy is an increasingly important aspect of data publishing services. If personal private information is leaked from the data, the service will be regarded as unacceptable by the original owners of the data. Two different approaches to defining a notion of database privacy, the generalization method and the perturbation method, have been independently studied. These two approaches have significantly differences, making it hard to compare related research. In this paper, we propose a unified model that is based on the perturbation method, but which is applicable to generalized data sets. In particular, this model applies the notion of differential privacy to data sets that satisfy *k*-anonymity. We demonstrate this approach through a simple case study. This is a first step towards a common notion for protecting database privacy.

1 Introduction

Privacy is an increasingly important aspect of data publishing. Sensitive data, such as medical records in public databases, are recognized as a valuable source of information for the allocation of public funds, medical research and statistical trend analysis [1]. However, if personal private information is leaked from the database, the service will be regarded as unacceptable by the original owners of the data.

There are two approaches to avoiding leaks of private information from public databases: generalization methods and perturbation methods. Generalization methods modify the original data to avoid identification of the records. These methods generate a common value for some records and replace identifying information in the records with the common value. However, detailed information is lost during this process. On the other hand, perturbation methods add noise to data. While perturbed data usually retains detailed information, it also normally includes fake data.

An important issue is how to evaluate these methods with regard to privacy leakage. In particular, when performing such an evaluation, it is difficult to model the background knowledge of an adversary trying to obtain private information from a database. Even if some fields of records in a database have been anonymized in some manner, an adversary may still be able to identify a record through background knowledge. For example, even if ZIP codes are generalized to include just the highest level of regional information in a medical database, this may still be enough to identify a record if there is only one case of a particular disease in that region and an adversary knows that a particular target has had that disease and lives in that region.

Since the generalization and perturbation methods take such different approaches, it is very difficult to compare them. In this paper, we propose a model that is based on the notion of differential privacy but is applicable to *k*-anonymized data sets. We explain the two methodologies in Section 3. Our unified model is presented in Section 4. Finally, Section 5 provides a case study that applies our model.

Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications, volume: 2, number: 1, pp. 50-62

2 Related Work

There are two major approaches to avoiding leaks of private information from public databases: perturbative methods and non-perturbative methods. Generalization methods are non-perturbative methods and they modify the original data to avoid identification of the records. These methods generate a common value for some records and replace identifying information in the records with the common value. However, detailed information is lost during this process. On the other hand, perturbative methods add noise to data. While perturbed data usually retains detailed information, it also normally includes fake data.

Differential Privacy [2, 3] is a notion of privacy for perturbative methods that is based on the statistical distance between two database tables differing by at most one element. The basic idea is that, regardless of background knowledge, an adversary with access to the data set draws the same conclusions, whether or not a person's data is included in the data set. That is, a person's data has an insignificant effect on the processing of a query. Differential privacy is mainly studied in relation to perturbation methods[4, 5, 6, 7, 8] in an interactive setting. Attempts to apply differential privacy to search queries were discussed in [9]. Li *et al.* proposed a matrix mechanism [10] that is applicable to predicate counting queries under a differential privacy setting. Computational relaxations of differential privacy has been discussed in [11, 12, 13]. There are some research papers that have discussed a relationship between an information-theoretic notion of privacy leakage and differential privacy. Alvim *et al.* showed how to model a query system in terms of an information-theoretic channel and compared the notion of differential privacy[14]. Barthe and Kopf performed an information-theoretic analysis of differential privacy and discussed upper bounds for the leakage of differentially private mechanisms[15].

Samarati and Sweeney [16, 17, 18] proposed a primary definition of privacy that is applicable to generalization methods. A data set is said to have *k-anonymity* if each record is indistinguishable from at least $k-1$ other records with respect to certain identifying attributes called *quasi-identifiers* [19]. Minimizing this information loss thus presents a challenging problem in the design of generalization algorithms. The optimization problem is referred to as the *k*-anonymity problem. Meyerson reported that optimal generalization in this regard is an NP-hard problem[20]. Aggarwal *et al.* proved that finding an optimal table including more than three attributes is NP-hard [21]. Nonetheless, *k*-anonymity has been widely studied because of its conceptual simplicity [22, 23, 24, 25, 26, 27]. Machanavajjhala *et al.* proposed another important definition of privacy in a public database [23]. The definition, called *l*-diversity assumes a strong adversary having certain background knowledge that allows the adversary to identify object persons in the public database.

Samarati proposed a simple binary search algorithm for finding a *k*-anonymous table[17]. A drawback of Samarati's algorithm is that for arbitrary definitions of minimality, it is not always guaranteed that this binary search algorithm can find the minimal k-anonymity table. Sun *et. al.* presented a hash-based algorithm that improves the search algorithm[27]. Aggarwal *et al.* proposed an $O(k)$-approximation algorithm [28] for the *k*-anonymity problem. A greedy approximation algorithm [29] proposed by LeFevre *et al.* searches optimal multi-dimensional anonymization. A genetic algorithm framework [30] was proposed because of its flexible formulation and its ability to allow more efficient anonymization. Utility-based anonymization [31, 32] makes *k*-anonymous tables using a heuristic local recoding anonymization. Moreover, the *k*-anonymization problem is viewed as a clustering problem. Clustering-based approaches [33, 34, 35, 36] search a cluster that has *k*-records.

In full-domain generalization, there are two heuristic approaches for generalization algorithms: the top-down approach and the bottom-up approach. Bayardo and Aggrawal proposed a generalization algorithm using the top-down approach [37]. The algorithm finds a generalization that is optimal according to a given fixed cost metric for a systematic search strategy, given generalization hierarchies for a single attribute. Incognito [38] is a bottom-up-based algorithm that produces all possible *k*-anonymous tables

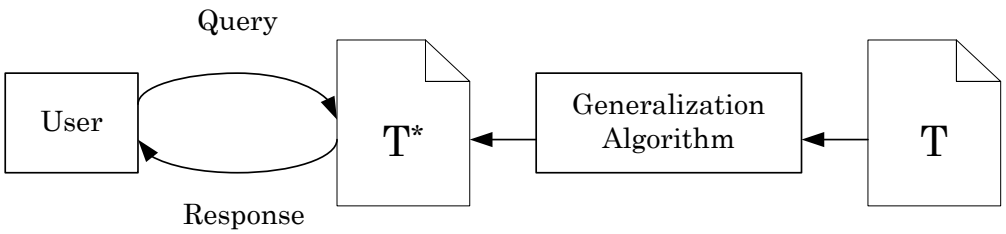

Figure 1: Generalization Method

for an original table.

Another approach for quantifying privacy leakage is an information theoretic definition proposed by Clarkson and Schneider [39]. They modeled an anonymizer as a program that receives two inputs: a user's query and database response for the query. The program acts as a noisy communication channel and produces an anonymized response as output. Hsu *et al.* provides a generalized notion [40] in decision theory for making a model of the value of privacy information. An alternative model for quantification of privacy information is proposed in [41]. In the model, the value of privacy information is estimated by the expected cost that the user has to pay for obtaining perfect knowledge from given privacy information. Furthermore, the sensitivity of different attribute values are taken into account in the average benefit and cost models proposed by Chiang *et al.*[42]. Krause and Horvitz presented utility-privacy tradeoffs in online services [43, 44].

3 Notion of Privacy

In this section we provide a brief introduction to the two methods and their related notions of privacy.

3.1 The Generalization method

As indicated in Figure 1, a generalization algorithm G transforms an original database table T into a modified database $\mathcal{G}(T) = T^*$. The generalization involves replacing a value with a less specific but semantically consistent value. For example, the original ZIP codes {02138, 02139} can be generalized to 0213*, thereby stripping the rightmost digit and semantically indicating a larger geographical area. A generalization algorithm has a hierarchy of attribute values, in which the upper node is a suppressed value of its lower nodes.

Generalization is a non-interactive method and is executed independently from queries to the public database. Responses to all queries are computed from the static table T^* that is generated in advance by the generalization algorithm. All users can access T^* without requiring access to the generalization algorithm itself.

At least k records must exist in the data set for each combination of attributes. Clearly any generalization algorithm that converts a database into one with k-anonymity involves a loss of information in that database. More precisely, suppose that a database table T has m records and n attributes $\{A_1, \ldots, A_n\}$. Each record $\mathbf{a^i} = (a_1^i, \ldots, a_n^i)$ can thus be considered as an n-tuple of attribute values, where d_j is the value of attribute A_j in record $\mathbf{a^i}$. The database table T itself can thus be regarded as the set of records $T = \{\mathbf{a^i} : 1 \leq i \leq m\}$.

Definition 1. (k-Anonymity) A database table T is said to have *k-anonymity* if and only if each n-tuple of attribute values $\mathbf{a} \in T$ appears at least k times in T.

Table 1: Example Table of *k*-anonymity $(k = 2, m = 11, n = 4, l = 1)$

Quasi-identifiers					Sensitive attributes
Birth	Gender	ZIP	Nationality		Problem
1985	male	0124*	Europe		shortness of breath
1985	male	0124*	Europe		chest pain
1985	female	0123*	Asia		hypertension
1985	female	0123*	Asia		hypertension
1984	female	0123*	Asia		obesity
1984	female	0123*	Asia		chest pain
1984	male	0123*	USA		chest pain
1984	male	0123*	USA		obesity
1984	male	0123*	USA		shortness of breath
1987	male	0123*	USA		chest pain
1987	male	0123*	USA		chest pain

The definition of *k*-anonymity does not on its own capture the concept of an adversary who has background knowledge that can help the adversary to distinguish records [23]. As a result, several extension to the basic idea have been proposed, including *l-diversity* and recursive (c,l)*-diversity*, as well as other suggestions in [45, 46, 25].

To further understand how *k*-anonymity can be applied, we consider the following example of a database whose attributes are divided into two classes. The data holder is expected to identify one class of attributes that may be used for linking with external information. Such attributes not only include explicit identifiers such as name, address and phone number, but also include attributes that in combination can uniquely identify individuals, such as date of birth and gender. Such attributes are often referred to as *quasi−identifiers* [19]. The remaining attributes are termed *sensitive attributes*, which represent the essential attributes in the table for responding to queries. Thus the database table T can be represented as $T = (T^q|T^s)$, where subtable T^q consists the quasi- identifiers and subtable T^s consists of the sensitive attributes. An example database table is shown in Table 1.

Since the sensitive attributes represent the essential information with regard to database queries, a generalization method is used to modify (anonymize) T^q in order to prevent the identification of the owners of the sensitive attributes, while retaining the full information in T^s. The generalization method modifies subtable T^q to T^{q*}, where T^{q*} has *k*-anonymity. The resulting modified database table is $T^* = (T^{q*}|T^s)$.

3.2 The Perturbation method

Perturbation methods are interactive methods which dynamically compute answers to queries, as shown in Figure 2. For example, for a query involving a request for the average value of an attribute, a perturbation algorithm might randomly generate noise and then computes the average value of the data including this added noise. The values of responses are probabilistic because the computation involves random added noise.

The original table is kept secret and users can only access it via the perturbation algorithm. Several perturbation methods designed to prevent privacy leakage from responses to queries have been presented [4, 5].

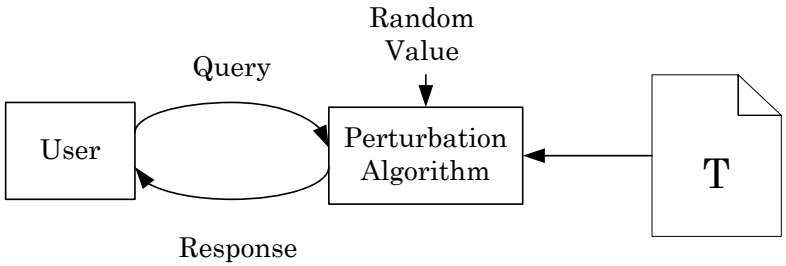

Figure 2: Perturbation Method

Differential privacy is mainly discussed in an interactive setting. Queries are modeled as *query functions*. Information in a database T is modified by the *perturbation algorithm* $\mathscr{K}$ in order to compute a response $\mathscr{K}(f,r,T)$, where r is a random input.

Definition 2. (Differential Privacy) A perturbation algorithm $\mathscr{K}$ provides ε-differential privacy for a query function f if, for all data sets T_1 and T_2 differing by at most one element, for any random inputs r_1 and r_2, and all possible responses t:

$$Pr[\mathscr{K}(f,r_1,T_1)=t] \leq e^{\varepsilon} Pr[\mathscr{K}(f,r_2,T_2)=t] \quad (e^{\varepsilon} \approx 1 \pm \varepsilon).$$

If ε is negligibly small, the response $\mathscr{K}(f,r_1,T_1)$ is equal to the response $\mathscr{K}(f,r_2,T_2)$ with high probability. This means that even if one data record was removed or changed from the data set, no response would become significantly more or less likely.

This definition is very strict and it is difficult to apply to existing perturbation mechanisms. Thus, Definition 2 can be relaxed [6]:

Definition 3. ((ε,δ)-differential privacy) A perturbation algorithm $\mathscr{K}$ provides (ε,δ)-differential privacy for a query function f if, for all data sets T_1 and T_2 differing by at most one element, for any random inputs r_1 and r_2, and all possible responses t:

$$Pr[\mathscr{K}(f,r_1,T_1)=t] \leq e^{\varepsilon} Pr[\mathscr{K}(f,r_2,T_2)=t] + \delta.$$

Other relaxation schemes [47, 48, 49] essentially also satisfy Definition 3 [50].

This concept of differential privacy can also be formalized as an indistinguishability game. The objective of the adversary is to distinguish two different games, **game 0** and **game i**, with non-negligible probability. In **game 0**, we prepare an original data set T and the adversary interacts with a perturbation algorithm $\mathscr{K}$ taking input (f,r_1,T_{-i}). In **game i**, data set T_{-i} denotes data set T where the i-th record has been replaced with with $\emptyset$, and the adversary interacts with $\mathscr{K}$ taking input (f,r_2,T_{-i}).

The adversary distinguishes the game based on the resulting responses. The indistinguishability depends on the distance between the distribution of the responses from the two different data sets. The perturbation algorithm $\mathscr{K}$ is said to be *ε-indistinguishable* with respect to data set T [7, 51], if for all $1 \leq i \leq m$:

$$Pr[\mathscr{K}(f,r_1,T)=t] \leq e^{\varepsilon} Pr[\mathscr{K}(f,r_2,T_{-i})=t].$$

(ε,δ)-indistinguishability can be defined in a similar way. The parameter δ (or δ/ε) is the advantage of the adversary distinguishing two games, where $Pr[\mathscr{K}(f,r_1,T)=t]$ is equal to 0.

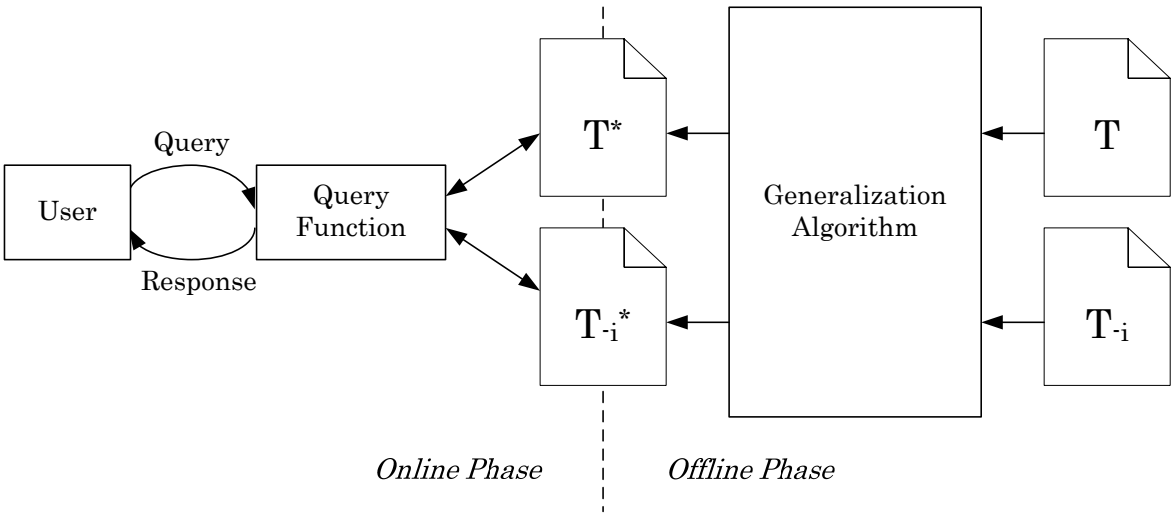

Figure 3: New Model for Generalization Methods

4 Model for Generalization Methods

Our goal is to define a unified notion for the two different approaches: generalization methods and perturbation methods. We can compare privacy leakage in these two different models if we have a unified notion of what privacy leakage means. To realize a unified notion of privacy, we need to describe generalization methods in the same manner as differential privacy.

Generalization methods based on *k*-anonymity are non-interactive methods that generate a static table from the original table. Responses to queries are deterministic because the response is computed from the static table. In this section, we present a new probabilistic model for generalization methods, as shown in Figure 3.

In the original notion of the perturbation method, the perturbation algorithm is defined for each query function. On the other hand, non-interactive generalization methods are independent of query functions and the adversary can use several types of query for the anonymized (generalized) tables. Therefore, we will define a query function f to be independent from the generalization function.

We will thus sequentially execute a generalization function $\mathscr{G}$ and a query function f for the table T:

$$\mathscr{G} : \mathscr{G}(T) \rightarrow T^*$$

$$f : f(T^*, query) \rightarrow t,$$

where the anonymized table T^* is the output of $\mathscr{G}(T)$, and *query* is query information for the query function f. The generalization function $\mathscr{G}$ computes an anonymized table T^* from the original table T. The query function produces a response t from *query* and table T^*. We later consider a model for the query function.

In order to be able to apply the indistinguishability notions from differential privacy, we must first construct an interactive game between an adversary and a public database. To do so, anonymized tables in which one element is different have to be prepared. We now propose one model for making such tables.

We first make tables where the i-th element is different from the original table T, then we generate anonymized tables. The symbols T_{-i}, T^*, and $(T_{-i})^*$ denote a table in which the i-th element is different from T, an anonymized table of T, and an anonymized table of T_{-i}, respectively. The query function

executes responses from the anonymized tables. That is, the functions $\mathscr{K}(f,r,T)$ and $\mathscr{K}(f,r,T_{-i})$ in the original notion of differential privacy are replaced by $f(T^*,query)$, and $f((T_{-i})^*,query)$, respectively.

Using the above model, we are able to calculate the probability that the query function provides different outputs between **game 0** and **game i**. This is the same concept as indistinguishability for differential privacy described in Section 2.3. This probability implies the probability that the adversary will successfully distinguish two games; a game using the original table and a game using a table in which one record is modified.

A query function is selected according to the use case of the table and then the adversary accesses the anonymized table via the query function. In these games, the adversary can obtain the output of the query function; however, the adversary cannot obtain the anonymized table itself.

5 Case Study

In this section, we apply our new model to a *k*-anonymized table. As well as being illustrative, the purpose of this case study is to provide an instantiation of the query function and to show that the probability that the adversary can distinguish the two games can be calculated.

5.1 Anonymized table

For simplicity, we consider the case where T^s is a matrix of m records with just one attribute. We assume that all records have an index value d that provides a unique identifier for that record.

We define T_{-i} to be the table where the sensitive attribute of the i-th record in T is replaced with a fixed value $\emptyset$. This generates m different tables and thus defines m games between the original table T and tables in which one record has been modified. We then analyze the probability that the adversary can tell the difference between the anonymized table T^* and the table $(T_{-i})^*$.

Let $|S|$, $|Q|$, and $|M|$ be, respectively, the number of possible values for each set of sensitive attributes in table T^*, variations of quasi- identifiers q_x in T^*, and the total number of the records in T^*. It is assumed that the adversary knows $|S|$, $|Q|$ and $|M|$.

5.2 Query function

Deciding how to model the query function is an important issue that we have not yet discussed. In this subsection, we define a query function f for a *k*-anonymized table. The query function receives a user's query and responds with results based on the anonymized table T^* or $(T_{-i})^*$. An anonymized table that satisfies *k*-anonymity is divided into small subgroups; each group has at least *k*-records and the same quasi-identifier q_x.

We assume the use of an anonymized table and that a user tries to obtain sensitive information for a quasi-identifier q_x. For example, a user requests the values of sensitive information for $\{1984, male, *, USA\}$ in Table 1. We model the query function using two functionalities.

Let d be an index of the d-th record, q_x be a set of m attribute values in T^{q*}, and s be a value for the sensitive attribute. The two functionalities are defined as follows:

- **read.** For input of an index value d, the function outputs the d-th record. That is, $f(T^*, query = \{\textbf{read}, d\}) \rightarrow \{d, q_x^d, s^d\}$, where q_x^d and s^d are values of the quasi-identifier and the sensitive attribute in the d-th record. If the d-th record does not exist, then the function outputs $failed$.

- **search.** For input q_x and/or s, the function outputs the number u of records and index values which have a quasi-identifier q_x and/or sensitive attribute s. That is, $f(T^*, query = \{\textbf{search}, q_x, s\}) \rightarrow$

u, D, where u and D are the number of records and a sequence of index values that have the same quasi-identifier and/or sensitive attribute. If s or q_x do not exist, the function outputs $failed$.

Note that the functionality must satisfy the limitation that the adversary cannot obtain the table itself. The adversary can obtain sensitive attributes using only the above two functionalities.

5.3 Calculating the probability

We calculate the probability that two games have a different response t. From the definition of the query function, we find three possible scenarios that use two types of queries **read** and **search**, as follows;

- **Scenario 0: read a record**. The adversary requests results of $f(T^*, query = \{\mathbf{read}, d\})$ for all $d \in T^*$ and receives $\{d, q_x^d, s^d\}$. In this scenario, the adversary finds the difference when the **read** query hits q_x of the modified i-th record, because the i-th record has a different sensitive attribute value. Thus, the probability is simply calculated as $\leq \frac{1}{|M|}\pi$, where π is the maximum number of queries that the adversary is allowed.

- **Scenario 1: search by quasi-identifier**. The adversary requests the result of $f(T^*, query = \{\mathbf{search}, q_x, *\})$ and receives u and D. Then, the adversary requests results of $f(T^*, query = \{\mathbf{read}, d\})$ for all $d \in D$ and receives $\{d, q_x^d, s^d\}$. In this scenario, the adversary finds the difference when the **search** query hits q_x of the modified i-th record and the **read** query hits the modified i-th record, because the i-th record has a different sensitive attribute value. That is, the case of the response $t = \{i, q_x^i, \emptyset\}$ has the maximum probability. The number of records that have the same quasi-identifier q_x is at least k. Thus, the probability is $\leq \frac{1}{k|Q|}\pi$, where π is the maximum number of queries that the adversary is allowed. For example, if we consider the anonymized table as Table 2, the probability is $\frac{1}{8}\pi$.

- **Scenario 2: search by sensitive attribute**. The adversary requests the result of $f(T^*, query = \{\mathbf{search}, *, s\})$ and receives u and D. In this scenario, the adversary finds the difference when the **search** query hits s of the modified i-th record because of the difference of u. The case of the response $t = \{u, D\}$ has the maximum probability. Thus, the probability is $\leq \frac{1}{|S|}\pi$, where π is the maximum number of queries that the adversary is allowed.

- **Scenario 3: search by both quasi-identifier and sensitive attribute**. The adversary requests the result of $f(T^*, query = \{\mathbf{search}, q_x, s\})$ and receives u and D. In this scenario, the adversary finds the difference when the **search** query hits q_x and s of the modified i-th record because of the difference of u. The case of the response $t = \{u, D\}$ has the maximum probability. Thus, the probability is $\leq \frac{1}{|Q||S|}\pi$, where π is the maximum number of queries that the adversary is allowed. Compared to the previous two scenarios, this scenario is rather inefficient.

Generally, the table satisfies the condition $|Q| \leq |S| \leq |M|$. From the analysis of the above scenarios, the most efficient scenario is Scenario 2, and the maximum probability that the difference between $\mathscr{G}(T)$ and $\mathscr{G}(T_{-i})$ is calculated as $\leq \frac{1}{|S|}\pi$, where $|S| \leq k|Q|$. This means that there is no efficient strategy with a probability less than $\frac{1}{|S|}\pi$, where π is considered as a computational bound on the adversary. The success probability $\frac{1}{|S|}\pi$ is the same as the success probability when the adversary simply guesses the value of the sensitive information without access to the table. Thus, in our model, the anonymized table should satisfy the condition that no efficient strategy with a probability less than $\frac{1}{|S|}\pi$ exists.

Next, we consider an adversary who knows q_x of a target person as background knowledge. In this case, the most efficient scenario is Scenario 1 and the adversary uses a fixed q_x for all games according to

Table 2: Example Modified Table

Quasi-identifiers					Sensitive attribute
Birth	Gender	ZIP	Nationality		Problem
1985	male	0124*	Europe		shortness of breath
1985	male	0124*	Europe		$\emptyset$
1985	female	0123*	Asia		hypertension
1985	female	0123*	Asia		hypertension
1984	female	0123*	Asia		obesity
1984	female	0123*	Asia		chest pain
1984	male	0123*	USA		chest pain
1984	male	0123*	USA		obesity
1984	male	0123*	USA		shortness of breath
1987	male	0123*	USA		chest pain
1987	male	0123*	USA		chest pain

the background knowledge. The maximum probability is $\leq \frac{1}{k}\pi$ for a k-anonymized table. On the other hand, if we use the original table T instead of T^*, and the adversary knows the correct quasi-identifier of the person, the probability goes to 1. Thus, a generalization algorithm producing the k-anonymized table seems to be an effective way of protecting against leakage of privacy information. However, if we consider the worst case where records of the same q_x have the same value of sensitive attribute s, such as {1987, male, 0123*, USA} in Table 1, we obtain a new maximum probability equal to 1 using the following scenario:

- **Scenario 3': search by quasi-identifier, then search by sensitive attribute**. The adversary requests the result of $f(T^*, query = \{\textbf{search}, q_x, *\})$ and receives u and D, the adversary requests results of $f(T^*, query = \{\textbf{read}, d\})$ for a $d \in D$ and receives $\{d, q_x^d, s^d\}$, then the adversary requests the result of $f(T^*, query = \{\textbf{search}, q_x^d, s^d\})$ and receives u and D. In this scenario, if the adversary knows background knowledge q_x (such as q_x ={1987, male, 0123*, USA} in Table 2), the adversary can search the record i by the q_x. Thus, the probability that the attacker can find the difference is 1.

To circumvent this threat, l-diversity was proposed by Machanavajjhala *et. al.* [23]. If the table satisfies k-anonymity and l-diversity, it is ensured that the number of sensitive information values in records with the same q_x is more than l. The probability of Scenario 3' becomes $\leq \frac{1}{l}\pi$.

Our evaluation results fit an intuitive understanding of the l-diversity definition as follows. l-diversity is required in order to prevent background knowledge attacks; an adversary can easily obtain the value for sensitive information in the worst case, where the adversary has background knowledge about a quasi-identifier of a target person. It is not sufficient for privacy protection that the table satisfies k-anonymity.

Thus, we think that our case study provides evidence that privacy leakage can be modeled in this way based on a k-anonymized table.

6 Conclusion

In this paper, we have proposed a new model of differential privacy for evaluating tables with k-anonymity. Furthermore, we presented a case study of a k-anonymized table based on our model.

In the evaluation, we calculated the probability that the adversary could distinguish two games: **game 0** and **game i** in our case study. The adversary can distinguish the two games when a response in **game 0** is different from a response in **the game i**. There are two cases, corresponding to when the adversary executes either a **read** query for record i, or a**search** query for sensitive information in record i. In each case, the probability that **game 0** has the same response t as the response of **game i** is 0 ($Pr[f((T)^*, query) = t] = 0$), and $Pr[f((T_{-i})^*, query) = t]$ gives δ (or δ/e^{ε}) in our model, because of the relation $Pr[f((T)^*, query) = t] \leq e^{\varepsilon} Pr[f((T_{-i})^*, query) = t] + \delta$. Thus, we found that the probability is comparable with the parameter δ (or δ/e^{ε}) in (ε, δ)-differential privacy.

One open question is whether the definition of the query function is sufficient to define services using an anonymized table in the real world. We will continue to consider whether the functionalities are suitable for different database systems. Furthermore, existing differential privacy models are assumed to have a limited query function that is a part of a perturbation function. Existing models raise a similar question about the modeling of the query function. We will also consider this problem in future research.

Even though our ultimate goal is to define a new privacy notion applicable to all methods of privacy protection, our research is intended to be a first step in considering the relationship between the two different security notions. We believe our work will trigger more research in this area.

References

[1] N. R. Adam and J. C. Wortmann, "Security-control methods for statistical databases: A comparative study," *ACM Computing Surveys*, vol. 21, no. 4, pp. 515–556, 1989.

[2] C. Dwork, "Differential privacy," in *Proc. of the 33rd International Colloquium on Automata, Languages and Programming (ICALP'06), Venice, Italy, LNCS*, vol. 4052. Springer-Verlag, July 2006, pp. 1–12.

[3] C. Dwork, "Differential privacy: A survey of results," in *Proc. of the 5th Annual Conference on Theory and Applications of Models of Computation (TAMC'08), Xi'an, China, LNCS*, vol. 4978. Springer-Verlag, December 2008, pp. 1–19.

[4] S. E. Fienberg and J. McIntyre, "Data swapping: Variations on a theme by dalenius and reiss," in *Proc. of Privacy in Statistical Databases 2004 (PSD'04), Barcelona, Spain, LNCS*, vol. 3050. Springer-Verlag, June 2004, pp. 14–29.

[5] W. E. Winkler, "Masking and re-identification methods for public-use microdata: Overview and research problems," in *Proc. of Privacy in Statistical Databases 2004 (PSD'04), Barcelona, Spain, LNCS*, vol. 3050. Springer-Verlag, June 2004, pp. 231–246.

[6] C. Dwork, K. Kenthapadi, F. McSherry, I. Mironov, and M. Naor, "Our data, ourselves: Privacy via distributed noise generation," in *Proc. of Eurocrypt 2006, Saint Petersburg, Russia, LNCS*, vol. 4004. Springer-Verlag, May-June 2006, pp. 486–503.

[7] C. Dwork, F. McSherry, K. Nissim, and A. Smith, "Calibrating noise to sensitivity in private data analysis," in *Proc. of the 3rd Theory of Cryptography Conference (TCC'06), New York, USA, LNCS*, vol. 3876. Springer-Verlag, March 2006, pp. 265–284.

[8] C. Dwork, G. N. Rothblum, and S. Vadhan, "Boosting and differential privacy," pp. 51–60, October 2010.

[9] P. Kodeswaran and E. Viegas, "Applying differential privacy to search queries in a policy based interactive framework," in *Proc. of ACM First International Workshop on Privacy and Anonymity for Very Large Datasets (PAVLAD'09), Hong Kong*. ACM Press, November 2009, pp. 25–32.

[10] C. Li, M. Hay, V. Rastogi, G. Miklau, and A. McGregor, "Optimizing linear counting queries under differential privacy," in *Proc. of the 29th ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems (PODS'10), Indiana, USA*. ACM Press, June 2010, pp. 123–134.

[11] I. Mironov, O. Pandey, O. Reingold, and S. Vadhan, "Computational differential privacy," in *Proc. of CRYPTO 2009, Santa Barbara, California, USA, LNCS*, vol. 5677. Springer-Verlag, August 2009, pp. 126–142.

[12] A. McGregor, I. Mironov, T. Pitassi, O. Reingold, K. Talwar, and S. Vadhan, "The limits of two-party differential privacy," in *Proc. of the 51st IEEE Annual Symposium on Foundations of Computer Science (FOCS'10), Las Vagas, USA*. IEEE, October 2010, pp. 81–90.

[13] A. Groce, J. Katz, and A. Yerukhimovich, "Limits of computational differential privacy in the client/server setting," in *Proc. of the 8th Theory of Cryptography Conference (TCC'11), to appear, Brown University, USA, LNCS*. Springer-Verlag, March 2011.

[14] M. S. Alvim, K. Chatzikokolakis, P. Degano, and C. Palamidessi, "Differential privacy versus quantitative information flow," *CoRR, abs/1012.4250*, 2010.

[15] G. Barthe and B. Kopf, "Information-theoretic bounds for differentially private mechanisms," in *IACR Cryptology ePrint Archive: Report 2011/071, http://www.springerlink.com/content/c288812n57277k7r/*, 2011.

[16] P. Samarati and L. Sweeney, "Generalizing data to provide anonymity when disclosing information," in *Proc. of the 17th ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems (PODS'98), Seattle, Washington*. ACM Press, June 1998, p. 188.

[17] P. Samarati, "Protecting respondents' identities in microdata release," *IEEE Trans. on Knowledge and Data Engineering*, vol. 13, no. 6, pp. 1010–1027, 2001.

[18] L. Sweeney, "Achieving *k*-anonymity privacy protection using generalization and suppression," *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, vol. 10, no. 5, pp. 571–588, 2002.

[19] T. Dalenius, "Finding a needle in a haystack —or identifying anonymous census record," *Journal of Official Statistics*, vol. 2, no. 3, pp. 329–336, 1986.

[20] A. Meyerson and R. Williams, "On the complexity of optimal *k*-anonymity," in *Proc. of the 23rd ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'04), Paris, France*. ACM Press, June 2004, pp. 223–228.

[21] G. Aggarwal, T. Feder, K. Kenthapadi, R. Motwani, R. Panigrahy, D. Thomas, and A. Zhu, "Anonymizing tables," in *Proc. of the 10th International Conference on Database Theory (ICDT'05), Edinburgh, Scotland, LNCS*, vol. 3363. Springer-Verlag, January 2005, pp. 246–258.

[22] S. S. Al-Fedaghi, "Balanced *k*-anonymity," *World Academy of Science, Engineering and Technology (WASET)*, vol. 1, no. 6, pp. 179–182, June 2005.

[23] A. Machanavajjhala, J. Gehrke, and D. Kifer, "*l*-diversity: Privacy beyond *k*-anonymity," in *Proc. of the 22nd International Conference on Data Engineering (ICDE'06), Atlanta, USA*. IEEE, April 2006, pp. 24–35.

[24] A. Machanavajjhala and J. Gehrke and D. Kifer, "*t*-closeness: Privacy beyond *k*-anonymity and *l*-diversity," in *Proc. of the 23rd IEEE International Conference on Data Enginering (ICDE'07), Istanbul, Turkey*. IEEE, April 2007, pp. 106–115.

[25] R. C.-W. Wong, J. Li, A. W.-C. Fu, and K. Wang, "(α,k)-anonymity: an enhanced *k*-anonymity model for privacy preserving data publishing," in *Proc. of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (SIGKDD'06), Philadelphia, USA*. ACM Press, August 2006, pp. 754–759.

[26] T. M. Truta and B. Vinay, "Privacy protection: *p*-sensitive *k*-anonymity property," in *Proc. of the 22nd IEEE International Conference on Data Engineering (ICDE'06), Atlanta, USA*. IEEE, April 2006, pp. 94–103.

[27] X. Sun, H. Wang, J. Li, T. M. Truta, and P. Li, "(p^+,α)-sensitive *k*-anonymity: a new enhanced privacy protection model," in *Proc. of the 8th IEEE International Conference on Computer and Information Technology (CIT'08), Sydney, Australia*. IEEE, July 2008, pp. 59–64.

[28] G. Aggarwal, T. Feder, K. Kenthapadi, R. Motwani, R. Panigrahy, D. Thomas, and A. Zhu, "Approximation algorithms for k-anonymity," *Journal of Privacy Technology*, November 2005.

[29] K. LeFevre, D. J. DeWitt, and R. Ramakrishnan, "Mondrian multidimensional k-anonymity," in *Proc. of the 22nd International Conference on Data Engineering (ICDE'06), Atlanta, USA*. IEEE, April 2006, pp. 25–35.

[30] V. S. Iyengar, "Transforming data to satisfy privacy constraints," in *Proc. of the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (SIGKDD'02), Edmonton, Canada*. ACM Press, July 2002, pp. 279–288.

[31] J. Xu, W. Wang, J. Pei, X. Wang, B. Shi, and A. W.-C. Fu, "Utility-based anonymization using local recoding," in *Proc. of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*

(SIGKDD'06), Philadelphia, USA. ACM Press, August 2006, pp. 785–790.

[32] J. Xu and W. Wang and J. Pei and X. Wang and B. Shi and A. W.-C. Fu, "Utility-based anonymization for privacy preservation with less information loss," *SIGKDD Explor. Newsl.*, vol. 8, no. 2, pp. 21–30, 2006.

[33] J.-W. Byun, A. Kamra, E. Bertino, and N. Li, "Efficient k-anonymity using clustering technique," in *Proc. of the 12th International Conference on Database Systems for Advanced Applications (DASFAA'07), Bangkok, Thailand*, April 2007, pp. 188–200.

[34] T. M. Truta and A. Campan, "K-anonymization incremental maintenance and optimization techniques," in *Proc. of the 2007 ACM symposium on Applied computing (SAC'07), Seoul, Korea*. ACM Press, March 2007, pp. 380–387.

[35] J.-L. Lin and M.-C. Wei, "An efficient clustering method for k-anonymization," in *Proc. of the 2008 International Workshop on Privacy and Anonymity in Information Society (PAIS'08), Nantes, France*. ACM Press, March 2008, pp. 46–50.

[36] H. Zhu and X. Ye, "Achieving k-anonymity via a density-based clustering method," in *Proc. of Advances in Data and Web Management (APweb/WAIM'07), Huang Shan, China, LNCS*. Springer-Verlag, June 2007, pp. 745–752.

[37] B. Bayardo and R. Agrawal, "Data privacy through optimal *k*-anonymity," in *Proc. of 21st International Conference on Data Engineering (ICDE'05), Tokyo, Japan*, April 2005, pp. 217–228.

[38] K. LeFevre, D. J. DeWitt, and R. Ramakrishnan, "Incognito: Efficient full-domain *k*-anonymity," in *Proc. of the 24th ACM SIGMOD International Conference on Management of Data (SIGMOD'05), Baltimore, Maryland, USA*. ACM Press, June 2005, pp. 49–60.

[39] M. R. Clarkson and F. B. Schneider, "Quantification of integrity," in *Proc. of the 23rd IEEE Computer Security Foundations Symposium (CSF'10), Edinburgh, UK*. IEEE, July 2010, pp. pp.28–43.

[40] T.-S. Hsu, C.-J. Liau, D.-W. Wang, and J. K.-P. Chen, "Quantifying privacy leakage through answering database queries," in *Proc. of the 5th International Conference on Information Security (ISC'02), Sao Paulo, Brazil, LNCS*, vol. 2433. Springer-Verlag, Septmeber-October 2002, pp. 162–176.

[41] Y. C. Chiang, T.-S. Hsu, S. Kuo, and D.-W. Wang, "Preserving confidentially when sharing medical data," in *Proc. of Asia Pacific Medical Information Conference*, 2000.

[42] Y. T. Chiang, Y. C. Chiang, T.-S. Hsu, C.-J. Liau, and D.-W. Wang, "How much privacy? - a system to safe guard personal privacy while releasing database," in *Proc. of the 3rd International Conference on Rough Sets and Current Trends in Computing (RSCTC'02), Malvern, USA, LNCS*, vol. 2475. Spriger-Verlag, October 2002, pp. 226–233.

[43] A. Krause and E. Horvitz, "A utility-theoretic approach to privacy and personalization," in *Proc. of AAAI'08, Chicago, USA*, vol. 2. AAAI Press, July 2008, pp. 1181–1188.

[44] A. Krause and E. Horvitz, "A utility-theoretic approach to privacy in online services," *Journal of Artificial Intelligence Research*, vol. 39, pp. 633–662, 2010.

[45] A. Machanavajjhala, J. Gehrke, and D. Kifer, "*t*-closeness: Privacy beyond *k*-anonymity and *l*-diversity," in *Proc. of the 23rd IEEE International Conference on Data Engineering (ICDE'07), Istanbul, Turkey*. IEEE, April 2007, pp. 106–115.

[46] X. Sun, H. Wang, J. Li, T. M. Truta, and P. Li, "(p^+, α)-sensitive *k*-anonymity: a new enhanced privacy protection model," in *Proc. of the 8th IEEE International Conference on Computer and Information Technology (CIT'08), Sydney, Australia*. IEEE, July 2008, pp. 59–64.

[47] K. Chaudhuri and N. Mishra, "When random sampling preserves privacy," in *Proc. of CRYPTO 2006, Santa Barbara, California, USA, LNCS*, vol. 4117. Springer-Verlag, August 2006, pp. 198–213.

[48] I. Dinur and K. Nissim, "Revealing information while preserving priovacy," in *Proc. of the 22nd ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'03), San Diego, USA*. ACM Press, June 2003, pp. 202–210.

[49] K. Nissim, S. Raskhodnikova, and A. Smith, "Smooth sensitivity and sampling in private data analysis," in *Proc. of the 39th ACM Symposium on Theory of Computing (STOC'07), San Diego, USA*. ACM Press, June 2007, pp. 75–84.

[50] S. P. Kasiviswanathan and A. Smith, *A Note on Differential Privacy: Defining Resistance to Arbitrary Side Information*. CoRR, abs/0803.3946, 2008.

[51] V. Shmatikov, "Differential privacy," *Lecture Slides, The University of Texas at Austin Department of Computer*.

Shinsaku Kiyomoto received his B.E. in Engineering Sciences and his M.E. in Materials Science from Tsukuba University, Japan, in 1998 and 2000, respectively. He joined KDD (now KDDI) and has been engaged in research on stream ciphers, cryptographic protocols, and mobile security. He is currently a senior researcher at the Information Security Lab. of KDDI R&D Laboratories Inc. He was a visiting researcher of the Information Security Group, Royal Holloway University of London from 2008 to 2009. He received his Doctorate in Engineering from Kyushu University in 2006. He received the IEICE Young Engineer Award in 2004. He is a member of JPS.

Prof. Keith M. Martin is Director of the Information Security Group. He joined the ISG as a lecturer in January 2000. He received his BSc (Hons) in Mathematics from the University of Glasgow in 1988 and a PhD from Royal Holloway in 1991. Between 1992 and 1996 he held a Research Fellowship at the University of Adelaide, investigating mathematical modeling of cryptographic key distribution problems. In 1996 he joined the COSIC research group of the Katholieke Universiteit Leuven in Belgium, working on security for third generation mobile communications. Keith's current research interests include key management, combinatorial cryptography, applications of cryptography and wireless sensor network security. Keith became Director of the ISG in 2010. He is an Associate Editor of IEEE Transactions on Information Theory in the area of Complexity and Cryptography.