

Worst-Case Background Knowledge for Privacy-Preserving Data Publishing

David J. Martin Daniel Kifer Ashwin Machanavajjhala Johannes Gehrke
Joseph Y. Halpern
Cornell University
{djm, dkifer, mvnak, johannes, halpern}@cs.cornell.edu

Abstract

Recent work has shown the necessity of considering an attacker's background knowledge when reasoning about privacy in data publishing. However, in practice, the data publisher does not know what background knowledge the attacker possesses. Thus, it is important to consider the worst-case. In this paper, we initiate a formal study of worst-case background knowledge. We propose a language that can express any background knowledge about the data. We provide a polynomial time algorithm to measure the amount of disclosure of sensitive information in the worst case, given that the attacker has at most k pieces of information in this language. We also provide a method to efficiently sanitize the data so that the amount of disclosure in the worst case is less than a specified threshold.

1. Introduction

We consider the following situation. A data publisher (such as a hospital) has collected useful information about a group of individuals (such as patient records that would help medical researchers) and would like to publish this data while preserving the privacy of the individuals involved. The information is stored as a table (as in Figure 1) where each record corresponds to a unique individual and contains a sensitive attribute (e.g., disease) and some non-sensitive attributes (e.g., address, gender, age) that might be learned using externally available data (e.g., phone books, birth records). The data publisher would like to limit the disclosure of the sensitive values of the individuals in order to defend against an attacker who possibly already knows some facts about the table. Our goal in this paper is to quantify the precise effect of background knowledge possessed by an attacker on the amount of disclosure and to provide algorithms to check and ensure that the amount of disclosure is less than a specified threshold.

The problem we solve is of real and practical importance; an egregious example of a privacy breach was

the discovery of the medical records of the Governor of Massachusetts from an easily accessible and supposedly anonymized dataset. All that was needed was to link it to voter registration records [32]. To defend against such attacks, Samarati and Sweeney [29] introduced a privacy criterion called k -anonymity which requires that each individual be indistinguishable (with respect to the non-sensitive attributes) from at least $k - 1$ others. This is done by grouping individuals into *buckets* of size at least k , and then permuting the sensitive values in each bucket and sufficiently masking their externally observable non-sensitive attributes. Figure 2 depicts a table that is a 5-anonymous version of the table in Figure 1. Figure 3 depicts the permutation of sensitive values that was used to construct this table.

However, k -anonymity does not adequately protect the privacy of an individual;¹ for example, when all individuals in a bucket have the same disease, the disease of the individuals in that bucket is disclosed regardless of the bucket size. Even when there are multiple diseases in the same bucket, the frequencies of the diseases in the bucket still matter when an attacker has some background knowledge about the particular individuals in the table. Suppose the data publisher has published the 5-anonymous table as depicted in Figure 2. Consider an attacker Alice who would like to learn the diseases of all her friends and neighbors. One of her neighbors is Ed, a 27 year-old male living in Ithaca (zip code 14850). Alice knows that Ed is in the hospital that published the anonymized dataset in Figure 3, and she wants to find out Ed's disease. Using her knowledge of Ed's age, gender, and zip-code, Alice can identify the bucket in the anonymized table that Ed belongs to (namely, the first bucket). Alice does not know which disease listed within that bucket is Ed's since the sensitive values were permuted. Therefore, without additional knowledge, Alice's estimate of the probability that Ed has lung cancer is $2/5$. But suppose Alice knows that Ed had mumps as a child and is therefore extremely unlikely to get it again. After ruling out this possibility, the probability that Ed has lung can-

¹Indeed, the definition of k -anonymity does not even mention the sensitive attribute!

non-sensitive				sensitive
Name	Zip	Age	Sex	Disease
Bob	14850	23	M	Flu
Charlie	14850	24	M	Flu
Dave	14850	25	M	Lung Cancer
Ed	14850	27	M	Lung Cancer
Frank	14853	29	M	Mumps
Gloria	14850	21	F	Flu
Hannah	14850	22	F	Flu
Irma	14853	24	F	Breast Cancer
Jessica	14853	26	F	Ovarian Cancer
Karen	14853	28	F	Heart Disease

Figure 1. Original table

non-sensitive				sensitive
Name	Zip	Age	Sex	Disease
*	1485*	2*	M	Flu Lung Cancer Mumps Flu Lung Cancer
*	1485*	2*	F	Flu Breast Cancer Flu Heart Disease Ovarian Cancer

Figure 2. 5-anonymous table

non-sensitive				sensitive
Name	Zip	Age	Sex	Disease
Bob	14850	23	M	Flu
Charlie	14850	24	M	Lung Cancer
Dave	14850	25	M	Mumps
Ed	14850	27	M	Flu
Frank	14853	29	M	Lung Cancer
Gloria	14850	21	F	Flu
Hannah	14850	22	F	Breast Cancer
Irma	14853	24	F	Flu
Jessica	14853	26	F	Heart Disease
Karen	14853	28	F	Ovarian Cancer

Figure 3. Bucketized table

cer increases to $1/2$. Now, if Alice also somehow discovers that Ed does not have flu, then the fact that he has lung cancer becomes certain. Here, two pieces of knowledge of the form “Ed does not have X” were enough to fully disclose Ed’s disease. To guard against this, Machanavajjhala et al. [24] proposed a privacy criterion called ℓ -diversity that ensures that it takes at least $\ell - 1$ such pieces of information to sufficiently disclose the sensitive value of any individual. The main idea is to require that, for each bucket, the ℓ most frequent sensitive values are roughly equi-probable.

ℓ -diversity focuses on one type of background knowledge: knowledge of the form “individual X does not have sensitive value Y”. But an attacker might well have other types of background knowledge. For example, suppose Alice lives across the street from a married couple, Charlie and Hannah, who were both taken to the hospital. Once again, using her knowledge of their genders, ages and zip-codes, Alice can identify the buckets Charlie and Hannah belong to. Without additional background knowledge, Alice thinks that Charlie has the flu with probability $2/5$. But suppose that Alice knows that Hannah has had a flu shot recently but Charlie has not. Believing Hannah’s immunity to the flu to be much stronger than Charlie’s and knowing that they live together, Alice deduces that if Hannah has succumbed to the flu then it is extremely likely that Charlie has as well. This knowledge allows her to update her probability that Charlie has the flu to $10/19$. We show how these probabilities are computed in Section 3. ℓ -diversity does not guard against the type of background knowledge in this example.

It is thus clear that we need a more general-purpose framework that can capture knowledge of *any* property of the underlying table that an attacker might know. Moreover, unlike in the two examples above where we knew Alice’s background knowledge, we will not assume that we know exactly what the attacker knows. We therefore take the following approach. In Section 2, we propose a language that is expressive enough to capture any property of the sensitive values in a table. This language enables us to decompose background knowledge into basic units of information. Then, given an anonymized version of the table,

we can quantify the worst-case disclosure risk posed by an attacker with k such units of information; k can be thought of as a bound on the power of an attacker. In Section 3, we show how to efficiently preserve privacy by ensuring that the *worst-case* (i.e., maximum) disclosure for *any* k pieces of information is less than a specified threshold. Furthermore, we show to integrate our techniques into existing frameworks to find a “minimally sanitized” table for which the maximum disclosure is less than a specified threshold. We present experiments in Section 4, related work in Section 5, and we conclude in Section 6.

To the best of our knowledge, this is the first such formal analysis of the effect of unknown background knowledge on the disclosure of sensitive information.

2. Framework

We begin by modeling the data publishing situation formally. Let P be a (finite) set of people. For each $p \in P$, we associate a tuple t_p which has one sensitive attribute S (e.g., disease) with finite domain and one or more non-sensitive attributes. We overload notation and use S to represent both the sensitive attribute and its domain. The data publisher has a table T , which is a set of tuples corresponding to a subset of P . The publisher would like publish T in a form that protects the sensitive information of any individual from an attacker with background knowledge that can be expressed in a language \mathcal{L} . (We propose such a language to express background knowledge in Section 2.2.)

2.1. Bucketization

We first need to carefully describe how the published data is constructed from the underlying table if we are to correctly interpret this published data. That is, we need to specify a sanitization method. We briefly describe two popular sanitization methods.

- The first, which we term *bucketization* [34], is to partition the tuples in T into *buckets*, and then to separate the sensitive attribute from the non-sensitive ones

by randomly permuting the sensitive attribute values within each bucket. The sanitized data then consists of the buckets with permuted sensitive values.

- The second sanitization technique is *full-domain generalization* [32], where we coarsen the non-sensitive attribute domains. The sanitized data consists of the coarsened table along with generalization used. Note that, unlike bucketization, the exact values of the non-sensitive attributes are not released; only the coarsened values are released.

Note that if the attacker knows the set of people in the table and their non-sensitive values, then full-domain generalization and bucketization are equivalent. In this paper, we use bucketization as the method of constructing the published data from the original table T , although all our results hold for full-domain generalization as well. We plan to extend our algorithms to work for other sanitization techniques, such as data swapping [10] (which, like bucketization, also permutes the sensitive values, but in more complex ways) and suppression [29], in the future.

We now specify our notion of bucketization more formally. Given a table T , we partition the tuples into buckets (i.e., horizontally partition the table T according to some scheme), and within each bucket, we apply an independent random permutation to the column containing S -values. The resulting set of buckets, denoted by \mathcal{B} , is then published. For example, if the underlying table T is as depicted in Figure 1, then the publisher might publish bucketization \mathcal{B} as depicted in Figure 3. Of course, for added privacy, the publisher can completely mask the identifying attribute (Name) and may partially mask some of the other non-sensitive attributes (Age, Sex, Zip).

For a bucket $b \in \mathcal{B}$, we use the following notation.

P_b	set of people $p \in P$ with tuples $t_p \in b$
n_b	number of tuples in b
$n_b(s)$	frequency of sensitive value $s \in S$ in b
s_b^0, s_b^1, \dots	sensitive values in decreasing order of frequency in b

2.2. Background Knowledge

We pessimistically assume that the attacker has managed to obtain complete information about which individuals have records in the table, what their non-sensitive data is, and which buckets in the bucketization these records fall into. That is, we assume that the attacker knows P_b , the set of people in bucket b , for each $b \in \mathcal{B}$, and knows $t_p[X]$ for every person p in the table and every non-sensitive attribute X . We call this *full identification information*. One way of obtaining identification information in practice is to

link quasi-identifying non-sensitive attributes published in the bucketization (e.g., address, gender, age) with publicly available data (e.g., phone directories, birth records) [32].

We make the standard random worlds assumption [6]: in the absence of any further knowledge, we consider all tables consistent with this bucketization to be equally likely. That is, the probability of $t_p \in b$ having s for its sensitive attribute is $n_b(s)/n_b$ since each assignment of sensitive attributes to tuples within a bucket is equally likely.

We now need to consider knowledge beyond the identification information that an attacker might possess. We assume that this further knowledge is the knowledge that the underlying table satisfies a given *predicate* on tables. That is, the attacker knows that the underlying table is among the set of tables satisfying the given predicate. This is a rather general assumption. For example, “the average age of heart disease patients in the table is 48 years” could be one such predicate. In order to quantify the power of such knowledge, we use the notion of a *basic unit* of knowledge, and we propose a language which consists of finite conjunctions of such basic units. Given full identification information, we desire that any predicate on tables be expressible using a conjunction of the basic units that we propose. We employ a very simple propositional syntax.

Definition 1 (Atoms) An atom is a formula of the form $t_p[S] = s$, for some value $s \in S$ and person $p \in P$ with tuple $t_p \in T$. We say that atom $t_p[S] = s$ involves person p and value s .

The interpretation of atoms is obvious: $t_{\text{Jack}}[\text{Disease}] = \text{flu}$ says that the Jack’s tuple has the value flu for the sensitive attribute Disease.

The basic units of knowledge in our language are *basic implications*, defined below.

Definition 2 (Basic implications) A basic implication is a formula of the form

$$(\wedge_{i \in [m]} A_i) \rightarrow (\vee_{j \in [n]} B_j)$$

for some $m \geq 1, n \geq 1$ and atoms $A_i, B_j, i \in [m], j \in [n]$ (note that we use the standard notation $[n]$ to denote the set $\{0, \dots, n - 1\}$).

The fact that basic implications are a sufficiently expressive “basic unit” of knowledge is made precise by the following theorem.²

Theorem 3 (Completeness) Given full identification information and any predicate on tables, one can express the knowledge that the underlying table satisfies the identification information and the given predicate using a finite conjunction of basic implications.

²See [25] for proofs.

Hence we can model arbitrarily powerful attackers.³ Consider an attacker who knows the disease of every person in the table except for Bob. Then publishing any bucketization will reveal Bob’s disease. To avoid pathological and unrealistic cases like this, we need to assume a bound on the power of an attacker. We model attackers with bounded power by limiting the number of basic implications that the attacker knows. That is, the attacker knows a single formula from language $\mathcal{L}_{\text{basic}}^k$ defined below.

Definition 4 $\mathcal{L}_{\text{basic}}^k$ is the language consisting of conjunctions of k basic implications. That is, $\mathcal{L}_{\text{basic}}^k$ consists of formulas of the form $\bigwedge_{i \in [k]} \varphi_i$ where each φ_i is a basic implication.

k can thus be viewed as a bound on the attacker’s power and can be increased to provide more conservative privacy guarantees.

Note that our choice of basic implications for the “basic unit” of our language has important consequences on our assumptions about the attacker’s power. In particular, some properties of the underlying table might require a large number of basic implications to express. Since basic implications are essentially CNF clauses with at least one negative atom, our language suffers from an exponential blowup in the number of basic units required to express arbitrary DNF formulas. It may be that other choices of basic units may lead to equally expressive languages while at the same time requiring fewer basic units to express certain natural properties, and we consider this an important direction for future research. Nevertheless, many natural types of background knowledge have succinct representations using basic implications. For example, Alice’s knowledge that “if Hannah has the flu, then Charlie also has the flu” is simply the basic implication

$$t_{\text{Hannah}}[\text{Disease}] = \text{flu} \rightarrow t_{\text{Charlie}}[\text{Disease}] = \text{flu}$$

And the knowledge that “Ed does not have flu” is

$$t_{\text{Ed}}[\text{Disease}] = \text{flu} \rightarrow t_{\text{Ed}}[\text{Disease}] = \text{ovarian cancer}$$

In general, we can represent $\neg t[S] = s$ by $(t[S] = s) \rightarrow (t[S] = s')$ for any choice of $s' \neq s$ since each tuple has exactly one sensitive attribute value.

Note that maintaining privacy when there is dependence between sensitive values, especially *across buckets*, is a problem that has not been previously addressed in the privacy literature. The assignments of individuals to sensitive values in different buckets are not necessarily independent. As we saw in the example with Hannah and Charlie, fixing a particular assignment in one bucket could affect what

³A major shortcoming of the ℓ -diversity definition was that its choice of “basic unit” of knowledge was essentially negated atoms (i.e., $\neg t_p[S] = s$) which cannot capture all properties of the underlying table. For example, negations cannot express basic implications in general.

assignments are possible in another. One of the contributions of this paper is that we provide a polynomial time algorithm for computing the maximum disclosure even when the attacker has knowledge of such dependencies.

2.3. Disclosure

Having specified how the bucketization \mathcal{B} is constructed from the underlying table T and how an attacker’s knowledge about sensitive information can be expressed in language $\mathcal{L}_{\text{basic}}^k$, we are now in a position to define our notion of disclosure precisely.

Definition 5 (Disclosure risk) The disclosure risk of bucketization \mathcal{B} with respect to background knowledge represented by some formula φ in language $\mathcal{L}_{\text{basic}}^k$ is

$$\max_{t_p \in T, s \in S} \Pr(t_p[S] = s \mid \mathcal{B} \wedge \varphi)$$

That is, disclosure risk is the likelihood of the most highly predicted sensitive attribute assignment.

Definition 6 (Maximum disclosure) The maximum disclosure of bucketization \mathcal{B} with respect to language $\mathcal{L}_{\text{basic}}^k$ that expresses background knowledge is

$$\max_{t_p \in T, s \in S, \varphi \in \mathcal{L}_{\text{basic}}^k} \Pr(t_p[S] = s \mid \mathcal{B} \wedge \varphi)$$

By our assumptions in 2.2, we compute $\Pr(t_p[S] = s \mid \mathcal{B} \wedge \varphi)$ by considering the set of all tables consistent with bucketization \mathcal{B} and with background knowledge φ and then taking the fraction of those tables that satisfy $t_p[S] = s$. Using this, the maximum disclosure of the bucketization in Figure 3 with respect to $\mathcal{L}_{\text{basic}}^1$ turns out to be $\frac{5}{8}$ and occurs when φ is $t_{p'} = s' \rightarrow t_p = s$ where p and p' are distinct people in the first bucket and s and s' are flu and lung cancer or vice versa. Our goal is to develop general techniques to:

1. efficiently calculate the maximum disclosure for any given bucketization, and
2. efficiently find a “minimally sanitized” bucketization⁴ (or the set of all minimally sanitized bucketizations) for which the maximum disclosure is below a specified threshold (if any exist).

3. Checking And Enforcing Privacy

In Section 2.2, we defined basic implications as the “unit of knowledge” and showed that this was a fully expressive

⁴We will make precise the notion of “minimally sanitized” in Section 3.4; we want “minimal sanitization” in order to preserve the utility of the data.

(in the presence of full identification information) and reasonable choice. We now show how to efficiently calculate and limit maximum disclosure against an attacker who has full identification information and has up to k additional pieces of background knowledge (i.e., up to k basic implications). In order to do this, we will show in Theorem 9 that there is a set of k basic implications that maximizes disclosure with respect to $\mathcal{L}_{\text{basic}}^k$. Furthermore, each such implication has *only one atom in the antecedent and one atom in the consequent*. This motivates the following definition.

Definition 7 (Simple implications) A simple implication is a formula of the form $A \rightarrow B$ for some atoms A, B .

3.1. Hardness of computing disclosure risk

Unfortunately, naive methods for computing the maximum disclosure will not work – in fact, we can show that computing the disclosure risk of a given bucketization with respect to a given set of k simple implications is #P-hard. Note that k simple implications can be written in 2-CNF, for which satisfiability is easily checkable. Complexity is introduced in trying to *simultaneously* satisfy the k implications *and* the given bucketization. In fact, deciding whether a given bucketization is consistent with a set of k simple implications is NP-complete.

Theorem 8 Given as input bucketization \mathcal{B} and a conjunction of simple implications φ , the problem of deciding if \mathcal{B} and φ are both satisfiable by some table T is NP-complete. Moreover, given an atom C as further input, the problem of computing $\Pr(C \mid \mathcal{B} \wedge \wedge_{i \in [k]} \varphi_i)$ is #P-complete.

3.2. A special form for maximum disclosure

It turns out that, despite the hardness results above, computing the *maximum* disclosure with respect to language $\mathcal{L}_{\text{basic}}^k$ can be done in polynomial time. The key insight is summarized in Theorem 9.

Theorem 9 For any bucketization, there is a set of k simple implications, all sharing the same consequent, such that the conjunction of these k simple implications maximizes disclosure with respect to $\mathcal{L}_{\text{basic}}^k$.

This insight is tremendously useful in devising a polynomial-time dynamic programming algorithm for computing the maximum disclosure with respect to $\mathcal{L}_{\text{basic}}^k$ as it allows us to restrict our attention to sets of k simple implications of the form $(t_{p_i}[S] = s_i) \rightarrow (t_p[S] = s)$ for people $p, p_i \in P$, and values $s, s_i \in S$, $i \in [k]$. The proof of Theorem 9 follows from the following two lemmas.

Lemma 10 For any formulas $\psi, \varphi, \theta_i, \varphi_i$,

$$\begin{aligned} \Pr(\varphi \mid \psi \wedge (\wedge_{i \in [k]} (\theta_i \rightarrow \varphi_i))) \\ \leq \Pr(\varphi \mid \psi \wedge (\wedge_{i \in [k]} (\theta_i \rightarrow \varphi))) \end{aligned}$$

Starting with any set of k basic implications that maximize disclosure,⁵ Lemma 10 enables us to replace the consequent in all the basic implications by a single common atom (namely the atom corresponding to the highest predicted assignment of sensitive value to an individual), while still maintaining maximum disclosure.

Lemma 11 For any formulas ψ, B, θ_i , where B is an atom and θ_i is a conjunction of atoms, there exist atoms A_i such that

$$\begin{aligned} \Pr(B \mid \psi \wedge (\wedge_{i \in [k]} (\theta_i \rightarrow B))) \\ \leq \Pr(B \mid \psi \wedge (\wedge_{i \in [k]} (A_i \rightarrow B))). \end{aligned}$$

Next, Lemma 11 allows us to replace the antecedent of each of the resulting implications by an atom (possibly with a different atom for each implication), while still maintaining maximum disclosure.

In both Lemmas 10 and 11, we use ψ to represent the attacker’s knowledge about the bucketization \mathcal{B} . However, it is worthwhile pointing out that neither lemma places any restriction on ψ or on the underlying probability distribution. This makes the results presented here extremely general and powerful because *they characterize the form of background knowledge that maximizes disclosure risk for any form of anonymization and for any additional background knowledge*.

The main idea behind the proof of Lemma 10 (and also Lemma 11) can be illustrated as follows. Consider a bucketization \mathcal{B} . Let $(t_{p_i}[S] = s_i) \rightarrow (t_{p'_i}[S] = s'_i)$, for $i \in \{0, 1\}$, be two simple implications which maximize the disclosure of \mathcal{B} with respect to $\mathcal{L}_{\text{basic}}^2$. For convenience, we let A_i denote the atom $t_{p_i}[S] = s_i$ and B_i the atom $t_{p'_i}[S] = s'_i$. Let C be the atom $t_p[S] = s$ such that $\Pr(C \mid \mathcal{B} \wedge (\wedge_{i \in [2]} (A_i \rightarrow B_i)))$ is the maximum disclosure.

Now let us restrict our attention to the set of tables consistent with \mathcal{B} . Let \mathcal{T}_1 be the set of tables satisfying the simple implications $A_0 \rightarrow B_0$ and $A_1 \rightarrow B_1$, and let \mathcal{T}_2 be the set of tables satisfying $A_0 \rightarrow C$ and $A_1 \rightarrow C$. Figure 4 is a diagrammatic representation of \mathcal{T}_1 and \mathcal{T}_2 . Each row in the the truth table on the left (resp., right) in Figure 4 represents a subset of \mathcal{T}_1 (resp., \mathcal{T}_2). The variables a, b, c, d, e, f, g, h in the left-most (resp., a, b, d', f', h' in the right-most) column represents the size of the corresponding set. For example, the set of tables represented by the second row is the set of tables that satisfy the atom C but do not satisfy A_0 and A_1 , and the number of such of tables is b .

It is now clear from Figure 4 that the implications $A_0 \rightarrow C$ and $A_1 \rightarrow C$ also produce the maximum disclosure as follows. $\Pr(C \mid \wedge_{i \in [2]} A_i \rightarrow B_i) = \frac{b+d+f+h}{a+b+c+d+e+f+g+h}$ and $\Pr(C \mid \wedge_{i \in [2]} A_i \rightarrow C) = \frac{b+d'+f'+h'}{a+b+d'+f'+h'}$. Also

⁵There always exists some set of k basic implications that maximize disclosure since there are only finitely many atoms and therefore $\mathcal{L}_{\text{basic}}^k$ is finite.

	$\wedge_{i \in [2]}(A_i \rightarrow B_i)$						$\wedge_{i \in [2]}(A_i \rightarrow C)$					
	A_0	A_1	B_0	B_1	C		A_0	A_1	B_0	B_1	C	
a	0	0	*	*	0	=	0	0	*	*	0	a
b	0	0	*	*	1	=	0	0	*	*	1	b
c	0	1	*	1	0							
d	0	1	*	1	1	\subseteq	0	1	*	*	1	d'
e	1	0	1	*	0							
f	1	0	1	*	1	\subseteq	1	0	*	*	1	f'
g	1	1	1	1	0							
h	1	1	1	1	1	\subseteq	1	1	*	*	1	h'

Figure 4. Truth tables

$\frac{b+d+f+h}{a+b+c+d+e+f+h} \leq \frac{b+d+f+h}{a+b+d+f+h} \leq \frac{b+d'+f'+h'}{a+b+d'+f'+h'}$ since $d \leq d'$, $f \leq f'$, and $h \leq h'$. Thus $\Pr(C \mid \wedge_{i \in [2]} A_i \rightarrow B_i) \leq \Pr(C \mid \wedge_{i \in [2]} A_i \rightarrow C)$.

3.3. Computing maximum disclosure efficiently

Having reduced our search space from sets of basic implications that could lead to maximum disclosure to sets of simple implications with the same consequent, we are now in a position to create an efficient algorithm to compute the maximum disclosure. We want to *maximize* $\Pr(A \mid \mathcal{B} \wedge \wedge_{i \in [k]}(A_i \rightarrow A))$ over all atoms $A, A_i, i \in [k]$. Notice that for any atoms $A, A_i, i \in [k]$ such that A and $\wedge_{i \in [k]} A_i \rightarrow A$ are consistent with bucketization \mathcal{B} we have:

$$\begin{aligned}
& \Pr(A \mid \mathcal{B} \wedge (\wedge_{i \in [k]} A_i \rightarrow A)) \\
&= \frac{\Pr(A \wedge (\wedge_{i \in [k]}(A_i \rightarrow A)) \mid \mathcal{B})}{\Pr((\wedge_{i \in [k]}(A_i \rightarrow A)) \mid \mathcal{B})} \\
&= \frac{\Pr(A \mid \mathcal{B})}{\Pr((\neg A \wedge (\wedge_{i \in [k]} \neg A_i)) \vee A \mid \mathcal{B})} \\
&= \frac{\Pr(A \mid \mathcal{B})}{\Pr(\neg A \wedge (\wedge_{i \in [k]} \neg A_i) \mid \mathcal{B}) + \Pr(A \mid \mathcal{B})}
\end{aligned}$$

So it suffices to construct an efficient algorithm to *minimize*, over all atoms $A, A_i, i \in [k]$,

$$\frac{\Pr(\neg A \wedge (\wedge_{i \in [k]} \neg A_i) \mid \mathcal{B})}{\Pr(A \mid \mathcal{B})}. \quad (1)$$

In Section 3.3.1, we show how to minimize $\Pr(\wedge_{i \in [k]} \neg A_i \mid \mathcal{B})$ over atoms A_i involving individuals in the same bucket. We use this in Section 3.3.2 to provide a dynamic programming algorithm MINIMIZE1 that minimizes Formula (1) over atoms $A, A_i, i \in [k]$ involving individuals in the same bucket. Finally, in Section 3.3.3, we use MINIMIZE1 to construct another dynamic programming algorithm MINIMIZE2 to minimize Formula (1) jointly over the entire bucketization.

3.3.1 Minimizing $\Pr(\wedge_{i \in [k]} \neg A_i \mid \mathcal{B})$ for one bucket

Consider all sets of k atoms involving people whose tuples are in a single $b \in \mathcal{B}$. Each set of k atoms is asso-

Algorithm 1: MINIMIZE1(b, i, \hat{k}_i, \hat{k})

Input: b is the bucket under consideration

Input: i is the index of the next person p_i for which k_i (i.e., the number of atoms involving person p_i) is to be determined (initially 0)

Input: \hat{k}_i is the upper bound for k_i (initially k)

Input: \hat{k} is the number of atoms for which the people involved have yet to be determined (initially k)

```

1:  $p_{\min} \leftarrow 1$ 
2: for  $k_i = 1, 2, \dots, \min(\hat{k}_i, \hat{k})$  do
3:    $p \leftarrow \text{MINIMIZE1}(b, i+1, k_i, \hat{k} - k_i)$ 
4:    $p \leftarrow \frac{n_b - i - \sum_{j \in [k_i]} n_b(s_b^j)}{n_b - i} \times p$ 
5:    $p_{\min} \leftarrow \min(p_{\min}, p)$ 
6: end for
7: return  $p_{\min}$ 

```

ciated with a tuple (l, k_0, \dots, k_{l-1}) , where l is the number of people involved in the k atoms, and k_i is the number of atoms involving the i -th person. We label the k atoms $A_{i,j}$ for $i \in [l]$ and $j \in [k_i]$ such that atom $A_{i,j}$ is the j -th atom (out of k_i atoms) involving the i -th person. Lemma 12 provides a closed form for the minimum value of $\Pr(\wedge_{i \in [k]} \neg A_i \mid \mathcal{B})$ over all sets of k atoms associated with a particular (l, k_0, \dots, k_{l-1}) .

Lemma 12 *Let $b \in \mathcal{B}$ be any bucket. Let k, l , and k_0, k_1, \dots, k_{l-1} be such that $k = \sum_{i \in [l]} k_i$ and $k_i \geq k_{i+1}$ for all $i \in [l-1]$. Let $s_b^0, s_b^1, s_b^2, \dots$ be the sensitive values arranged in descending order of frequency in b . Then $\Pr(\wedge_{i \in [l], j \in [k_i]} \neg A_{i,j} \mid \mathcal{B})$ is minimized over all atoms $A_{i,j}$ when, $A_{i,j}$ is $t_{p_i}[S] = s_b^j$, for all $i \in [l]$ and all $j \in [k_i]$, where $p_0, p_1, \dots, p_{l-1} \in P_b$ are distinct. Consequently, the minimum probability is given by:*

$$\prod_{i \in [l]} \frac{n_b - i - \sum_{j \in [k_i]} n_b(s_b^j)}{n_b - i} \quad (2)$$

Note that $l \leq k$ and $k = \sum_{i \in [l]} k_i$ since each atom involves exactly one person. So the question of minimizing $\Pr(\wedge_{i \in [k]} \neg A_i \mid \mathcal{B})$ over all atoms A_i that mention only tuples in b becomes one of minimizing $\prod_{i \in [l]} \frac{n_b - i - \sum_{j \in [k_i]} n_b(s_b^j)}{n_b - i}$ over all $l \leq k$ and all k_0, \dots, k_{l-1} such that $\sum_{i \in [l]} k_i = k$.

This can easily be done using Algorithm 1. Thus, calling MINIMIZE1($b, 0, k, k$) minimizes $\Pr(\wedge_{i \in [k]} \neg A_i \mid \varphi_{\mathcal{B}})$ over all atoms A_i that involve people with tuples in bucket b . It is easy to modify the algorithm to remember the minimizing values of k_0, \dots, k_{l-1} , and thus we can even reconstruct the set of minimizing atoms according to Lemma 12.

Algorithm complexity. Note that the parameters of MINIMIZE1 are bounded. That is, for every recursive call MINIMIZE1(b, i, k_i, \hat{k}) that occurs inside the initial call to MINIMIZE1($b, 0, k, k$), parameter b does not change, and parameters i, \hat{k}_i, \hat{k} are all bounded by k (i.e., the number of implications we allow the attacker to know). So we can easily turn this into an $O(k^3)$ time and space algorithm using dynamic programming.

3.3.2 Minimizing Formula (1) within one bucket

Let us now minimize $\frac{\Pr(\neg A \wedge (\bigwedge_{i \in [k]} \neg A_i) | \mathcal{B})}{\Pr(A | \mathcal{B})}$ over all $k + 1$ atoms A and A_i , for $i \in [k]$, that only mention tuples in bucket b . Clearly any A, A_i that simultaneously minimize the numerator and maximize the denominator will work. We know that $\text{MINIMIZE1}(b, 0, k + 1, k + 1)$ will minimize the numerator. According to Lemma 12, at least one of these minimal $k + 1$ atoms mention the most frequent sensitive value. So, taking this atom to be A , we maximize the denominator as well. Thus, the minimum value is

$$\text{MINIMIZE1}(b, 0, k + 1, k + 1) \times \frac{n_b}{n_b(s_b^0)}.$$

3.3.3 Minimizing Formula (1) over all buckets

We look again at minimizing $\frac{\Pr(\neg A \wedge (\bigwedge_{i \in [k]} \neg A_i) | \mathcal{B})}{\Pr(A | \mathcal{B})}$, except this time, we allow A and A_i for $i \in [k]$ to mention tuples in possibly different buckets. To do this, we make use of the independence between buckets. Suppose that the $k + 1$ minimizing atoms (including A) are such that k_i of them mention tuples in bucket b_i , for each $i \in [l]$ for some $l \leq k + 1$. Let b_j be the bucket containing the tuple mentioned by A . Then, since the permutation of sensitive values for each bucket was picked independently, we can compute the minimum as

$$\frac{n_{b_j}}{n_{b_j}(s_{b_j}^0)} \times \prod_{i \in [l]} \text{MINIMIZE1}(b_i, 0, k_i, k_i).$$

So we need to minimize the above for all choices of $l \leq k + 1$, j , and k_0, k_1, \dots, k_{l-1} (which we can assume without loss of generality to be in descending order). Assuming buckets in \mathcal{B} are labeled as b_0, b_1, b_2, \dots , this is done by the MINIMIZE2 .

So $\text{MINIMIZE2}(0, k, \text{true})$ minimizes $\frac{\Pr(\neg A \wedge (\bigwedge_{i \in [k]} \neg A_i) | \mathcal{B})}{\Pr(A | \mathcal{B})}$ over all atoms $A, A_i, i \in [k]$. It is easy to modify the algorithm to remember the i 's and h_i 's, and hence reconstruct the minimizing atoms.

Algorithm complexity. Note that the parameters of MINIMIZE2 are bounded. That is, for every recursive call to $\text{MINIMIZE2}(i, h_i, a)$ that occurs inside the initial call to $\text{MINIMIZE2}(0, k, \text{true})$, parameter i is bounded by the number of buckets, parameter k_i is bounded by the total number of implications k , and a is either *true* or *false*. Thus, assuming that we first memoize (i.e., precompute all possible calls to) MINIMIZE1 (which we can do in time $O(|\mathcal{B}| \times k^3)$), we can modify the MINIMIZE2 algorithm using dynamic programming to take an additional $O(|\mathcal{B}| \times k)$ time and space. So the whole algorithm can be made to run in $O(|\mathcal{B}| \times k^3)$ time and space.

Incidentally, if one had two bucketizations \mathcal{B} and \mathcal{B}^* that differed only in that \mathcal{B}^* was the result of removing some

Algorithm 2 : $\text{MINIMIZE2}(i, h_i, a)$

Input: i is the current bucket b_i (initially 0)

Input: h_i is number of atoms $A_j, j \in [k]$ that we have yet to determine (initially k)

Input: a is a flag representing whether atom A involves a person in an earlier bucket $b_j, j < i$ (initially *false*)

```

1:  $r_{\min} \leftarrow \infty$ 
2: if  $i = |\mathcal{B}|$  then
3:   // Finished all buckets
4:   return  $r_{\min}$ 
5: end if
6: for  $h_{i+1} = 0, 1, 2, \dots, h_i$  do
7:    $u \leftarrow \text{MINIMIZE1}(b_i, 0, h_{i+1}, h_{i+1})$ 
8:    $x \leftarrow \text{MINIMIZE2}(i + 1, h_i - h_{i+1}, \text{true})$ 
9:   if  $a = \text{false}$  then
10:    // Atom  $A$  does not involve an earlier bucket  $b_j, j < i$ 
11:    // So either  $A$  involves  $b_i \dots$ 
12:     $v \leftarrow \text{MINIMIZE1}(b_i, 0, h_{i+1} + 1, h_{i+1} + 1)$ 
13:     $r_{\min} \leftarrow \min(r_{\min}, v \times x \times \frac{n_{b_i}}{n_{b_i}(s_{b_i}^0)})$ 
14:    // ... or else  $A$  involves a later bucket  $b_j, j > i$ 
15:     $r_{\min} \leftarrow \min(r_{\min}, u \times \text{MINIMIZE2}(i + 1, h_i - h_{i+1}, \text{false}))$ 
16:  else
17:    // Atom  $A$  involves an earlier bucket  $b_j, j < i$ 
18:     $r_{\min} \leftarrow \min(r_{\min}, u \times x)$ 
19:  end if
20: end for
21: return  $r_{\min}$ 

```

buckets from \mathcal{B} and adding x new buckets to \mathcal{B} , then, after we run the algorithm for \mathcal{B} , we memoize MINIMIZE1 for the x new buckets; so the incremental cost of running the algorithm for \mathcal{B}^* is $O(|\mathcal{B}^*| \times k + x \times k^3)$ -time. Moreover, if one knew in advance which buckets were going to be removed, one could order the buckets b_0, b_1, \dots appropriately to reuse much of the memoization of MINIMIZE2 as well.

3.4. Finding a safe bucketization

Armed with a method to compute the maximum disclosure, we now show how to efficiently find a “minimally sanitized” bucketization for which maximum disclosure is below a given threshold. Intuitively, we would like a minimal sanitization in order to preserve the utility of the published data. Let us be more concrete about the notion of minimal sanitization. Given a table, consider the set of bucketizations of this table. We impose a partial ordering \preceq on this set of bucketizations where $\mathcal{B} \preceq \mathcal{B}'$ if and only if every bucket in \mathcal{B}' is the union of one or more buckets in \mathcal{B} . Thus the bucketization \mathcal{B}_\top that has all the tuples in one bucket is the unique top element of this partial order, and the bucketization \mathcal{B}_\perp that has one tuple per bucket is the unique bottom element of this partial order. Our notion of a “minimally sanitized” bucketization is one that is as low as possible in the partial order (i.e., as close to \mathcal{B}_\perp) while still having maximum disclosure lower than a specified threshold.

Definition 13 ((c, k)-safety) Given a threshold $c \in [0, 1]$,

we say that \mathcal{B} is a (c, k) -safe bucketization if the maximum disclosure of \mathcal{B} with respect to $\mathcal{L}_{\text{basic}}^k$ is less than c .

If the maximum disclosure is *monotonic* with respect to the partial ordering \preceq , then finding a \preceq -minimal (c, k) -safe bucketization can be done in time logarithmic in the height of the bucketization lattice (which is at most the number of tuples in the table) by doing a binary search. The following theorem says that we do indeed have monotonicity.

Theorem 14 (Monotonicity) *Let \mathcal{B} and \mathcal{B}' be bucketizations such that $\mathcal{B} \preceq \mathcal{B}'$. Then the maximum disclosure of \mathcal{B} is at least as high as the maximum disclosure of \mathcal{B}' with respect to $\mathcal{L}_{\text{basic}}^k$.*

Another approach is to find *all* \preceq -minimal (c, k) -safe bucketizations, and return the one that maximizes a specified utility function. The monotonicity property allows us to make use of existing algorithms for efficient itemset mining [4], k -anonymity [7, 22] and ℓ -diversity [24].⁶ For example, we can modify the Incognito [22] algorithm, which finds all the \preceq -minimal k -anonymous bucketizations, by simply replacing the check for k -anonymity with the check for (c, k) -safety from Section 3.3. We can thus find the bucketization that maximizes a given utility function subject to the constraint that the bucketization be (c, k) -safe.

4. Experiments

In this section, we present a case-study of our framework for worst-case disclosure using the Adult Database from the UCI Machine Learning Repository [27]. We only consider the projection of the Adult Database onto five attributes – Age, Marital Status, Race, Gender and Occupation. The dataset has 45,222 tuples after removing tuples with missing values. We treat Occupation as the sensitive attribute; its domain consists of fourteen values. We use pre-defined generalization hierarchies for the attributes similar to the ones used in [22]. Age can be generalized to six levels (un-suppressed, generalized to intervals of size 5, 10, 20, 40, or completely suppressed), Marital Status can be generalized to three levels, and Race and Gender can each either be left as is or be completely suppressed. We consider all the possible anonymized tables using those generalizations.

We computed the maximum disclosure for k pieces of background knowledge, for k ranging from 0 (i.e., no background knowledge) to 12 (since we know that maximum disclosure certainly reaches 1 at $k = 13$ because there are only fourteen possible sensitive values). Figure 5 plots, for one anonymized table, the number of pieces of knowledge available to an adversary against the maximum disclosure

⁶While these algorithms typically have worst-case exponential running time in the height of the bucketization lattice, they have been shown to run fast in practice.

for both negated atoms (ℓ -diversity) and basic implications. In the anonymized table used, all the attributes other than Age were suppressed and the Age attribute was generalized to intervals of size 20. The solid line corresponds to implication statements and the dotted line corresponds to negated atoms. This graph agrees with our earlier observation that implication-type background knowledge subsumes negation; the maximum disclosure for k negated atoms is always smaller than the maximum disclosure for k implications. However, note that, for a given k , the difference between the maximum disclosure for negated atoms and for basic implications is not too large. This means that an anonymized table which tolerates maximum disclosure due to k negated atoms need not be anonymized much further to defend against k implications.

Intuitively, if all the buckets in a table have a nearly uniform distribution, then the maximum disclosure should be lower, but the exact relationship is not obvious. To get a better picture, we performed the following experiment. We fixed a value k for the number of pieces of information. For every entropy value h , we looked at all tables $\mathcal{T}(h)$ for which the minimum entropy of the sensitive attribute over all buckets was equal to h . Amongst $\mathcal{T}(h)$ we found the table $T(h)$ with the least maximum disclosure for k implications. Let the worst case disclosure for $T(h)$ given k pieces of knowledge be denoted by $w(T(h), k)$. We plotted h versus $w(T(h), k)$ for $k = 1, 3, 5, 7, 9, 11$ in Figure 6. We see a behavior which matches our intuition. For a given k , the disclosure risk monotonically decreases with increase in h . This is because increasing h means that we are looking at tables with more and more entropy in their buckets (and, consequently, less skew). We plotted an analogous graph (which we do not show here) for negation statements and observed very similar behavior.

5. Related Work

Many metrics have been proposed to quantify privacy guarantees in publishing anonymized data-sets. ‘Perfect privacy’ [12, 26] guarantees that published data does not disclose any information about the sensitive data. However, checking whether a conjunctive query discloses any information about the answer to another conjunctive query is shown to be very hard (Π_2^p -complete [26]). Subsequent work showed that checking for perfect privacy can be done efficiently for many subclasses of conjunctive queries [23]. Perfect privacy places very strong restrictions on the types of queries that can be answered [26] (in particular, aggregate statistics cannot be published). Less restrictive privacy definitions based on asymptotic conditional probabilities [11] and certain answers [30] have been proposed. Statistical databases allow answering aggregates over sensitive values without disclosing the exact value [1]. De-




Figure 5. Disclosure vs # pieces of background knowledge

identification, like k -anonymity [28, 32] and “blending in a crowd” [8], ensures that an individual cannot be associated with a unique tuple in an anonymized table. However, under both of those definitions, sensitive information can be disclosed if groups are homogeneous.

Background knowledge can lead to disclosure of sensitive information. Su et al. [31] and Yang et al. [35] limit disclosure when functional dependencies in the data are known to the data publisher upfront. The notion of ℓ -diversity [24] guards against limited amounts of background knowledge unknown to the data publisher. Farkas et al. [16] provide a survey of indirect data disclosure via inference channels.

There are several approaches to anonymizing a dataset to ensure privacy. These include generalizations [7, 22, 29], cell and tuple suppression [9, 29], adding noise [1, 5, 8, 15], publishing marginals that satisfy a safety range [14], and data swapping [10], where attributes are swapped between tuples so that certain marginal totals are preserved. Queries can be posed online and the answers audited [20] or perturbed [13]. Not all approaches guarantee privacy. For example, spectral techniques can separate much of the noise from the data if the noise is uncorrelated with the data [17, 19]. Anatomy [34] is a recently proposed anonymization technique that corresponds exactly to the notion of bucketization that we use in this paper. When the attacker knows full identification information, then generalization provides no more privacy than bucketization. However, we recommend generalizing the attributes before publishing the data since this will prevent attackers that do not already have full identification information from reidentifying individuals via linking attacks [32]. In many cases, the fact that a particular individual is in the table is considered sensitive information [8].

The utility of data that has been altered to preserve privacy has often been studied for specific future uses of the data. Work has been done on preserving association rules while adding noise [15]; reconstructing distributions of continuous variables after adding noise with a known distri-




Figure 6. Entropy vs Maximum Disclosure Risk

bution [5, 3]; reconstructing data clusters after perturbing numeric attributes [8]; and maximizing decision tree accuracy while anonymizing data [18, 33]. There have also been some negative results for utility. Publishing a single k -anonymous table can suffer from the curse of dimensionality [2] - large portions of the data need to be suppressed to ensure privacy. Subsequent work [21] shows how to publish several tables instead of a single one to combat this.

6. Conclusions

In this paper, we initiate a formal study of the worst-case disclosure with background knowledge. Our analysis does not assume that we are aware of the exact background knowledge possessed by the attacker. We only assume bounds on the attacker’s background knowledge in terms of the number of basic units of knowledge that the attacker possesses. We propose basic implications as an expressive choice for these units of knowledge. Although computing the probability of a specific disclosure from a given set of k basic implications is intractable, we show how to efficiently determine the worst-case over all sets of k basic implications. In addition, we show how to search for a bucketization that is robust (to a desired threshold ϵ) against any k basic implications by combining our check for (ϵ, k) -safety with existing lattice-search algorithms. Finally, we demonstrate that, in practice, ℓ -diversity has similar maximum disclosure to our notion of (ϵ, k) -safety, which guards against a richer class of background knowledge.

Since we chose basic implications as our units of knowledge, our algorithms will clearly yield very conservative bucketizations if we try to protect against an attacker who knows information that can only be expressed using a large number of basic implications. One way to reduce the number of basic units required is to add more powerful atoms to our existing language. Finding the right language for basic units of knowledge is an important direction of future work.

Other directions for future work include extending our

framework for probabilistic background knowledge, studying cost-based disclosure (since it was observed in [24] that not all disclosures are equally bad), and extending our results to other forms of anonymization, such as data-swapping and collections of anonymized marginals [21].

Acknowledgments. This work was supported by the National Science Foundation under Grants IIS-0541507, IIS-0636259, CNS-0627680, IIS-0534064, and ITR-0325453; by the DoD Multidisciplinary University Research Initiative (MURI) program administered by the ONR under grants N00014-01-1-0795 and N00014-04-1-0725; by AFOSR under grant FA9550-05-1-0055; by a Sloan Foundation Fellowship; and by gifts from Yahoo! and Microsoft. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

References

- [1] N. R. Adam and J. C. Wortmann. Security-control methods for statistical databases: a comparative study. *ACM Comput. Surv.*, 21(4):515–556, 1989.
- [2] Charu C. Aggarwal. On k -anonymity and the curse of dimensionality. In *VLDB*, pages 901–909, 2005.
- [3] D. Agrawal and C. C. Aggarwal. On the design and quantification of privacy preserving data mining algorithms. In *PODS*, 2001.
- [4] R. Agrawal and R. Srikant. Fast algorithms for mining association rules in large databases. In *VLDB*, 1994.
- [5] R. Agrawal and R. Srikant. Privacy preserving data mining. In *SIGMOD*, 2000.
- [6] F. Bacchus, A. J. Grove, J. Y. Halpern, and D. Koller. From statistical knowledge bases to degrees of belief. *A.I.*, 87(1-2), 1996.
- [7] R. J. Bayardo and R. Agrawal. Data privacy through optimal k -anonymization. In *ICDE*, 2005.
- [8] S. Chawla, C. Dwork, F. McSherry, A. Smith, and H. Wee. Toward privacy in public databases. In *TCC*, 2005.
- [9] L. H. Cox. Suppression, methodology and statistical disclosure control. *Journal of the American Statistical Association*, 75, 1980.
- [10] T. Dalenius and S. Reiss. Data swapping: a technique for disclosure control. *Journal of Statistical Planning and Inference*, 6, 1982.
- [11] N. Dalvi, G. Miklau, and D. Suciu. Asymptotic conditional probabilities for conjunctive queries. In *ICDT*, 2005.
- [12] A. Deutsch and Y. Papakonstantinou. Privacy in database publishing. In *ICDT*, 2005.
- [13] I. Dinur and K. Nissim. Revealing information while preserving privacy. In *PODS*, pages 202–210, 2003.
- [14] A. Dobra. *Statistical tools for disclosure limitation in multiway contingency tables*. PhD thesis, Carnegie Mellon University, 2002.
- [15] A. Evfimievski, J. Gehrke, and R. Srikant. Limiting privacy breaches in privacy preserving data mining. In *PODS*, 2003.
- [16] C. Farkas and S. Jajodia. The inference problem: a survey. *SIGKDD Explor. Newsl.*, 4(2), 2002.
- [17] Z. Huang, W. Du, and B. Chen. Deriving private information from randomized data. In *SIGMOD*, 2004.
- [18] Vijay S. Iyengar. Transforming data to satisfy privacy constraints. In *KDD*, pages 279–288, 2002.
- [19] H. Kargupta, S. Datta, Q. Wang, and K. Sivakumar. On the privacy preserving properties of random data perturbation techniques. In *ICDM*, pages 99–106, 2003.
- [20] K. Kenthapadi, N. Mishra, and K. Nissim. Simulatable auditing. In *PODS*, 2005.
- [21] Daniel Kifer and Johannes Gehrke. Injecting utility into anonymized datasets. In *SIGMOD*, 2006.
- [22] K. LeFevre, D. DeWitt, and R. Ramakrishnan. Incognito: Efficient full-domain k -anonymity. In *SIGMOD*, 2005.
- [23] A. Machanavajjhala and J. Gehrke. On the efficiency of checking perfect privacy. In *PODS*, 2006.
- [24] A. Machanavajjhala, J. Gehrke, D. Kifer, and M. Venkatasubramanian. ℓ -diversity: Privacy beyond k -anonymity. In *ICDE*, 2006.
- [25] D. Martin, D. Kifer, A. Machanavajjhala, J. Gehrke, and J. Halpern. Worst-case background knowledge in privacy. Technical report, Cornell University, 2006.
- [26] G. Miklau and D. Suciu. A formal analysis of information disclosure in data exchange. In *SIGMOD*, 2004.
- [27] U.C. Irvine Machine Learning Repository. <http://www.ics.uci.edu/mllearn/mlrepository.html>.
- [28] P. Samarati. Protecting respondents' identities in microdata release. In *IEEE Transactions on Knowledge and Data Engineering*, 2001.
- [29] P. Samarati and L. Sweeney. Protecting privacy when disclosing information: k -anonymity and its enforcement through generalization and suppression. Technical report, CMU, SRI, 1998.
- [30] K. Stoffel and M. Studer. Provable data privacy. In *DEXA*, 2005.
- [31] T. Su and G. Ozsoyoglu. Controlling fd and mvd inferences in multilevel relational database systems. *IEEE TKDE*, 3(4), 1991.
- [32] L. Sweeney. k -anonymity: a model for protecting privacy. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, 10(5):557–570, 2002.
- [33] K. Wang, B. C. M. Fung, and P. S. Yu. Template-based privacy preservation in classification problems. In *ICDM*, November 2005.
- [34] X. Xiao and Y. Tao. Anatomy: Simple and effective privacy preservation. In *VLDB*, 2006.
- [35] X. Yang and C. Li. Secure xml publishing without information leakage in the presence of data inference. In *VLDB*, pages 96–107, 2004.