Logic, Probability, and Privacy: A Framework for Specifying Privacy Requirements

Tsan-sheng Hsu, Churn-Jung Liau and Da-Wei Wang

Institute of Information Science, Academia Sinica, Taipei 115, Taiwan
tshsu@iis.sinica.edu.tw, liaucj@iis.sinica.edu.tw, wdw@iis.sinica.edu.tw

Abstract

In this paper, we propose a probabilistic hybrid logic for the specification of data privacy requirements. The proposed logic is a combination of quantitative uncertainty logic and basic hybrid logic with a satisfaction operator. We show that it is expressive enough for the specification of many well-known data privacy requirements, such as k-anonymity, l-diversity and its precursor logical safety, t-closeness, and δ-disclosure privacy. The main contribution of the work is twofold. On one hand, the logic provides a common ground to express and compare existing privacy criteria. On the other hand, the uniform framework can meet the specification needs of combining new criteria as well as existing ones.

Key words: Data privacy, information systems, probabilistic logic, hybrid logic, k-anonymity, logical safety, l-diversity, t-closeness, δ-disclosure privacy.

1 Introduction

To address the privacy concerns about the release of microdata, data is often sanitized before it is released to the public. For example, generalization and suppression of the values of quasi-identifiers are widely used sanitization methods. To assess the effect of sanitization methods, several data privacy criteria have been proposed. One of the earliest criteria was the notion of *k-anonymity*[13, 12, 14, 15]. Although k-anonymity is an effective way to prevent *identity disclosure*, it was soon realized that it was insufficient to ensure protection of sensitive attributes. To address the *attribute disclosure* problem, a *logical safety* criterion was proposed in [6]. The criterion was later expanded to the epistemic model in [17] and the well-known *l-diversity* criterion in [8, 9]. More recently, a variety of privacy criteria have been proposed[7, 2]. Due to the diversity of the privacy criteria, it is useful to have a flexible language for the specification of different privacy policies. The purpose of the paper is to provide such a formal specification language based on probabilistic hybrid logic.

Probabilistic hybrid logic is a fusion of a hybrid logic with a satisfaction operator[1] and a logic for reasoning about quantitative uncertainty[5]. The syntax of the proposed logic is comprised of well-formed formulas of both logics, and its semantics is based on epistemic probability structures with the additional interpretation of nominals. We show that the proposed probabilistic hybrid logic is expressive enough for the specification of data privacy requirements, such as k-anonymity, l-diversity and its precursor logical safety, t-closeness, and δ-disclosure privacy. Furthermore, the language is quite flexible so that we can specify personalized privacy requirements.

The remainder of this paper is organized as follows. In Section 2, we introduce the syntax and semantics of probabilistic hybrid logic. In Section 3, we define data representation and formulate the information systems as models of probabilistic hybrid logic. In Section 4, we explain how privacy requirements can be precisely specified with the proposed logic language. Finally, Section 5 contains some concluding remarks.

2 Probabilistic Hybrid Logic

2.1 Syntax

Hybrid logics are extensions of standard modal logics with *nominals* that name individual states in possible world models[1]. The simplest hybrid language is the extension of the basic modal language with nominals only. More expressive variants can include the existential modality $\mathtt{E}$, the satisfaction operator @, and the binder ↓. The simplest hybrid language is denoted by $\mathcal{H}$ and its extensions are named by listing the additional operators. For example, $\mathcal{H}(@)$ is the simplest hybrid language extended with the satisfaction operator @. On the other hand, the probabilistic logic $\mathcal{L}_n^{QU}$ proposed in [5] consists of *(linear) likelihood formulas* of the form

$$r_1 l_{a_1}(\varphi_1) + \cdots + r_k l_{a_k}(\varphi_k) > s,$$

where $r_1, \ldots, r_k, s$ are real numbers, $a_1, \ldots, a_k$ are (not necessarily distinct) agents, and $\varphi_1, \ldots, \varphi_k$ are well-formed formulas of the probabilistic language. The proposed probabilistic hybrid logic is a straightforward fusion of $\mathcal{H}(@)$ and $\mathcal{L}_n^{QU}$. The following definition gives the syntax of the resultant language.

Definition 1. *Let* $\mathtt{PROP} = \{p_1, p_2, \ldots\}$ *(the propositional symbols),* $\mathtt{AGT} = \{a_1, a_2, \ldots\}$ *(the agent symbols), and* $\mathtt{NOM} = \{i_1, i_2, \ldots\}$ *(the nominals) be pairwise disjoint, countably infinite sets of symbols. The well-formed formulas of the probabilistic hybrid logic* $\mathcal{PH}(@)$ *in the signature* $\langle \mathtt{PROP}, \mathtt{AGT}, \mathtt{NOM} \rangle$ *are given by the following recursive definition:*

$$\mathtt{WFF} ::= \top \mid p \mid i \mid \neg\varphi \mid \varphi \wedge \psi \mid \langle a \rangle \varphi \mid @_i \varphi \mid r_1 l_{a_1}(\varphi_1) + \cdots + r_k l_{a_k}(\varphi_k) > s,$$

where $p \in \mathtt{PROP}$*;* $i \in \mathtt{NOM}$*;* $a, a_1, \ldots, a_k \in \mathtt{AGT}$*;* $\varphi, \varphi_1, \ldots, \varphi_k \in \mathtt{WFF}$*; and* $r_1, \ldots, r_k, s$ *are real numbers.*

As usual, we abbreviate $\neg(\neg\varphi \wedge \neg\psi)$, $\neg(\varphi \wedge \neg\psi)$, and $\neg\langle a \rangle \varphi$ as $\varphi \vee \psi$, $\varphi \supset \psi$, and $[a]\varphi$ respectively. In addition, $(\varphi \supset \psi) \wedge (\psi \supset \varphi)$ is abbreviated as $(\varphi \equiv \psi)$; and several obvious abbreviations can be applied to likelihood formulas, e.g., $r_1 l_{a_1}(\varphi_1) + \cdots + r_k l_{a_k}(\varphi_k) < s$ denotes $(-r_1) l_{a_1}(\varphi_1) + \cdots + (-r_k) l_{a_k}(\varphi_k) > -s)$.

2.2 Semantics

The semantics of $\mathcal{PH}(@)$ is based on the *epistemic probability frame* introduced in [5].

Definition 2. *An epistemic probability frame is a tuple* $\mathfrak{F} = (W, (R_a)_{a \in \mathtt{AGT}}, (\mathcal{PR}_a)_{a \in \mathtt{AGT}})$*, where* W *is a set of possible worlds (states) and for each* $a \in \mathtt{AGT}$

- $R_a \subseteq W \times W$ *is a binary relation (the accessibility relation) on* W*, and*
- $\mathcal{PR}_a$ *is probability assignment, i.e., a function that associates a probability space* $(W_{w,a}, \mu_{w,a})$ *with each world* w*.*

Definition 3. *Let* $\mathfrak{F} = (W, (R_a)_{a \in \mathtt{AGT}}, (\mathcal{PR}_a)_{a \in \mathtt{AGT}})$ *be an epistemic probability frame. Then, an epistemic probability structure (or* $\mathcal{PH}(@)$ *model) based on* $\mathfrak{F}$ *is a pair* $\mathfrak{M} = (\mathfrak{F}, \pi)$*, where* $\pi : \mathtt{PROP} \cup \mathtt{NOM} \to 2^W$ *is an interpretation such that for all nominals* $i \in \mathtt{NOM}$*,* $\pi(i)$ *is a singleton. In this case, we also say that* $\mathfrak{F}$ *is the underlying frame of* $\mathfrak{M}$*.*

By slightly abusing the notation, we can identify a singleton and its element. Thus, when $\pi(i) = \{w\}$, we use $\pi(i)$ to denote both $\{w\}$ and w.

Definition 4. *Let $\mathfrak{M} = (W, (R_a)_{a \in \mathtt{AGT}}, (\mathcal{PR}_a)_{a \in \mathtt{AGT}}, \pi)$ be a $\mathcal{PH}(@)$ model and $w \in W$ be a possible world. Then, the satisfaction relation is defined as follows:*

1. $\mathfrak{M}, w \models \top$
2. $\mathfrak{M}, w \models p$ *iff* $w \in \pi(p)$ *for* $p \in \mathtt{PROP} \cup \mathtt{NOM}$
3. $\mathfrak{M}, w \models \neg\varphi$ *iff* $\mathfrak{M}, w \not\models \varphi$
4. $\mathfrak{M}, w \models \varphi \wedge \psi$ *iff* $\mathfrak{M}, w \models \varphi$ *and* $\mathfrak{M}, w \models \psi$
5. $\mathfrak{M}, w \models \langle a \rangle \varphi$ *iff there is a* w' *such that* $(w, w') \in R_a$ *and* $\mathfrak{M}, w' \models \varphi$
6. $\mathfrak{M}, w \models @_i \varphi$ *iff* $\mathfrak{M}, \pi(i) \models \varphi$
7. $\mathfrak{M}, w \models r_1 l_{a_1}(\varphi_1) + \cdots + r_k l_{a_k}(\varphi_k) > s$ *iff* $r_1 \mu_{w,a_1}(|\varphi_1| \cap W_{w,a_1}) + \cdots + r_k \mu_{w,a_k}(|\varphi_k| \cap W_{w,a_k}) > s$, *where* $|\varphi| = \{u \mid \mathfrak{M}, u \models \varphi\}$ *is the truth set of* φ *in the model* $\mathfrak{M}$.

A wff φ is said to be *true* in a model, denoted by $\mathfrak{M} \models \varphi$, if $\mathfrak{M}, w \models \varphi$ for all $w \in W$.

3 $\mathcal{PH}(@)$ Models of Information Systems

3.1 Information systems

In database applications, microdata, such as medical records, financial transaction records, and employee data, are typically stored in information systems. Am *information system* or *data table* is formally defined as follows[10]:

Definition 5. *An information system or a data table*[1] *is a tuple $T = (U, A, \{V_f \mid f \in A\})$, where U is a nonempty finite set, called the universe, and A is a nonempty finite set of attributes such that each $f \in A$ is a total function $f : U \to V_f$, where V_f is the domain of values for f.*

Let $B \subseteq A$ be a subset of attributes. Then, the *indiscernibility relation* with respect to B is defined on U as follows:

$$ind_T(B) = \{(x, y) \mid \forall f \in B f(x) = f(y)\}. \tag{1}$$

Usually, we omit the symbol T in the indiscernibility relation when the underlying information system is clear from the context. We also abbreviate an equivalence class of the indiscernibility relation $[x]_{ind(B)}$ as $[x]_B$.

The attributes of an information system can be partitioned into three subsets [4, 11]. First, we have a subset of *quasi-identifiers*, the values of which are known to the public. For example, in [14, 15], it is noted that certain attributes like birth-date, gender, and ethnicity are included in some public databases, such as census data or voter registration lists. These attributes, if not appropriately generalized, may be used to re-identify an individual's record in a medical data table, thereby causing a violation of privacy. The second kind is the subset of *confidential attributes*, the values of which we have to protect. It is often the case that an asymmetry exists between the values of a confidential attribute. For example, if the attribute is a HIV test result, then the revelation of a '+' value may cause a serious invasion of privacy, whereas it does not matter to know that an individual has a '−' status. Note that confidential attributes can also serve as quasi-identifiers in some cases. However, since the values of confidential attributes

[1] Also called knowledge representation systems or attribute-value systems in [10].

are not easily accessible by the public, in this paper, we simply assume that the set of quasi-identifiers is disjoint with the set of confidential attributes. The remaining attributes are *neutral attributes* that are neither quasi-identifying, nor confidential. Hereafter, we assume that the set of attributes $A = Q \cup C \cup N$, where Q, C, N are pairwise disjoint, Q is the set of quasi-identifiers, C is the set of confidential attributes, and N is the set of neutral attributes. Sometimes, the set of attributes is defined such that it contains identifiers that can be used to identify a person's data record. However, for simplicity, we equate each individual with his/her identifier, so the universe U can be considered as the set of identifiers. Furthermore, since identifiers are always removed in a released data table, U simply denotes a set of serial numbers for a de-identified information system.

Example 1. Table 1 is a simple example of an information system. The quasi-identifiers of the information systems are "Date of Birth" and "ZIP". The confidential attributes are "Income" and "Health Status". The values of "Health Status" indicate"normal"(0), "slightly ill"(1), and "seriously ill"(2). "Height" is a neutral attribute. ■

U	Date of Birth	ZIP	Height	Income	Health Status
i_1	24/09/56	24126	160	100K	0
i_2	06/09/56	24129	160	70K	1
i_3	23/03/56	10427	160	100K	0
i_4	18/03/56	10431	165	50K	2
i_5	20/04/55	26015	170	30K	2
i_6	18/04/55	26032	170	70K	0
i_7	12/10/52	26617	175	30K	1
i_8	25/10/52	26628	175	50K	0

Table 1: An information system in a data center

A common technique for protecting privacy is to release the information system in a sanitized form. Formally, we define *sanitization* as an operation on information systems.

Definition 6. *Let $T = (U, A, \{V_f \mid f \in A\})$ be an information system. Then, a sanitization operation $\sigma = (\iota, (s_f)_{f \in A})$ is a tuple of mappings such that*

- *$\iota : U \to U'$ is a 1-1 de-identifying mapping, where $|U'| = |U|$, and*
- *for each $f \in A$, $s_f : V_f \to V'_f$ is a sanitizing mapping, where V'_f is the domain of sanitized values for f.*

The application of σ on T results in a sanitized information system $\sigma T = (U', A', \{V'_f \mid f \in A\})$ such that $A' = \{f' \mid f \in A\}$; and for each $f \in A$, $f' = s_f \circ f \circ \iota^{-1}$, where $\circ$ denotes the functional composition. Note that the de-identifying mapping ι is invertible because it is a bijection.

The universe U' in a sanitized information system is regarded as the set of *pseudonyms* of the individuals. A sanitization operation $\sigma = (\iota, (s_f)_{f \in A})$ is *truthful* if for each $f \notin Q$, $s_f = id$ is the identity function; and it is *proper* if $\iota(ind_T(Q)) = \{(\iota(x), \iota(y)) \mid (x, y) \in ind_T(Q)\}$ is a proper subset of $ind_{\sigma T}(Q)$. In this paper, we only consider truthful sanitization operations. Moreover, in most cases, proper sanitization is necessary for the protection of privacy. A special sanitization, called *trivial* sanitization, is commonly used as the baseline of privacy assessment[2]. Formally, a sanitization operation is trivial if, for all $f \in Q$, $|V'_f| = 1$. The suppression of all quasi-identifiers can achieve the effect of trivial sanitization.

Example 2. In privacy research, generalization is a widely-used sanitization operation. For example, the date of birth may only be given as the year and month, or only the first two digits of the ZIP code may be given. A concrete generalization of the information system in Table 1 is presented in Table 2. The first column of the table shows the pseudonyms of the individuals. Note that the sanitization is truthful and proper. ■

d_1	09/56	24***	160	100K	0
d_2	09/56	24***	160	70K	1
d_3	03/56	10***	160	100K	0
d_4	03/56	10***	165	50K	2
d_5	04/55	26***	170	30K	2
d_6	04/55	26***	170	70K	0
d_7	10/52	26***	175	30K	1
d_8	10/52	26***	175	50K	0

Table 2: A sanitized information system

When a sanitized information system is released, the sanitizing mappings are usually known to the public, but the de-identifying mapping must be kept secret. In fact, when a sanitization is truthful and the adversary knows the values of the quasi-identifiers, the adversary can easily infer the sanitizing mappings. For example, in the previous sanitized information system, it is easy to see how "ZIP" and "Date of Birth" are generalized.

3.2 Models of sanitized information systems

To specify an information system and its sanitization, we have to use a fixed language. Let us consider an information system $T = (U, A, \{V_f \mid f \in A\})$, where $A = Q \cup N \cup C$ and a truthful sanitization operation $\sigma = (\iota, (s_f)_{f \in A})$. In addition, let $\sigma T = (U', A', \{V'_f \mid f \in A\})$ be defined as above. We assume that $U = \{i_1, \cdots, i_n\}$ and $U' = \{d_1, \cdots, d_n\}$. Then, the signature of our language comprises

- $\mathtt{PROP} = \{(f, v) \mid f \in N \cup C, v \in V_f\}$,
- $\mathtt{AGT} = \{a_0, a_1\}$, and
- $\mathtt{NOM} = U \cup U'$.

In Pawlak's decision logic[10], a propositional symbol (f, v) is called a *descriptor*, which means that the value of attribute f of an individual is v. Here, we only specify neutral and confidential attributes with the language. We consider two agents a_0 and a_1; and we assume that agent a_0 only receives the trivially sanitized information system, and a_1 receives the system σT. The set of nominals is partitioned into two subsets such that each i_j denotes an individual's identifier and each d_j represents the individual's pseudonym. The $\mathcal{PH}(@)$ models compatible with the sanitization of an information system are then defined as follows.

Definition 7. *Let $T = (U, A, \{V_f \mid f \in A\})$ be an information system, $\sigma = (\iota, (s_f)_{f \in A})$ be a truthful sanitization, and $\sigma T = (U', A', \{V'_f \mid f \in A\})$ be the sanitized system, where $A = Q \cup N \cup C$, $U = \{i_1, \cdots, i_n\}$ and $U' = \{d_1, \cdots, d_n\}$. Then, a $\mathcal{PH}(@)$ model $\mathfrak{M} = (W, R_0, R_1, \mathcal{PR}_0, \mathcal{PR}_1, \pi)$ with the above-mentioned signature is a model of σT if it satisfies the following conditions:*

- *$W = \{w_1, \cdots, w_n\}$;*
- *for R_0 and R_1:*
  - *$R_0 = W \times W$,*
  - *$R_1 = \{(w_j, w_k) \mid (d_j, d_k) \in ind_{\sigma T}(Q), 1 \leq j, k \leq n\}$;*
- *for the probability assignments:*
  - *$\mathcal{PR}_0$ associates a probability space (W, μ_0) with each world w such that $\mu_0(\{w\}) = \frac{1}{n}$ for each $w \in W$,*
  - *$\mathcal{PR}_1$ associates a probability space $(\pi([d_j]_Q), \mu_{w_j,1})$ with each world w_j such that $\mu_{w_j,1}(\{w\}) = \frac{1}{|\pi([d_j]_Q)|}$ for each $w \in \pi([d_j]_Q)$, where $[d_j]_Q$ is the equivalence class of d_j with respect to $ind_{\sigma T}(Q)$;*
- *and for the interpretation π:*
  - *$\pi(d_j) = w_j$ for $d_j \in U'$,*
  - *$\pi(i_j) \in \pi([d_j]_Q)$ for $i_j \in U$ and $\pi(i_j) \neq \pi(i_k)$ if $j \neq k$ for $1 \leq j, k \leq n$,*
  - *$\pi((f, v)) = \{w_j \mid f'(d_j) = v\}$ for $f \in N \cup C$ and $v \in V_f$.*

The models of σT reflect the adversary's uncertainty about the identities of the individuals. The possible worlds stand for the individuals. Although, the pseudonym of each individual is fixed, as specified by the interpretation π, the adversary is uncertain about the identifiers of the individuals. The information that an adversary can obtain is determined by the values of the individuals' quasi-identifiers, so an identifier may refer to any individual in a class of individuals that are indiscernible with respect to the quasi-identifiers. This is specified by the second clause of the interpretation π. With trivial sanitization, all individuals are indiscernible, so the accessibility relation R_0 is the universal relation. On the other hand, the sanitization operation σ results in the indiscernibility relation $ind_{\sigma T}(Q)$, so the relation R_1 is its isomorphic copy over the domain of possible worlds. Furthermore, we assume that the *indifference principle* applies to individuals, so both probability assignments associate a unform distribution with each possible world. Since the two probability assignments are characterized completely by the accessibility relations and R_0 is simply the universal relation, we can omit these three components from a model of σT and write it as a simple hybrid model (W, R_1, π). By the definition of π, there may be more than one $\mathcal{PH}(@)$ model for a given σT. Hence, a wff φ is *valid* in σT, denoted by $\Vdash_{\sigma T} \varphi$, if it is true in all models of σT.

4 The Specification of Data Privacy Requirements

In this section, we explain how the language of $\mathcal{PH}(@)$ can be used to specify different data privacy policies such as k-anonymity, l-diversity, and t-closeness. As in the preceding section, let $T = (U, A, \{V_f \mid f \in A\})$ denote an information system, where $A = Q \cup N \cup C$. In addition, let $U = \{i_1, \cdots, i_n\}$, $\sigma = (\iota, (s_f)_{f \in A})$ be a truthful sanitization, and $\sigma T = (U', A', \{V'_f \mid f \in A\})$ be the sanitized information system, where $U' = \{d_1, \cdots, d_n\}$.

4.1 Specification of k-anonymity

According to [13, 12, 14, 15], σT satisfies the k-anonymity criterion if $|[d]_Q| \geq k$ for any $d \in U'$. This is easily expressed in $\mathcal{PH}(@)$ language by the following formula:

$$l_{a_1}(i) \leq \frac{1}{k}$$

for $i \in \mathtt{NOM}$. Formally, we have the following theorem.

Theorem 1. *A sanitized information system σT satisfies the k-anonymity criterion iff $\Vdash_{\sigma T}$ $(l_{a_1}(i) \leq \frac{1}{k})$ for $i \in U$.*

The formal specification means that an individual can be identified with probability at most $\frac{1}{k}$. In particular, it can be derived that $@_d(l_{a_1}(i) \leq \frac{1}{k})$ is valid in σT for any $d \in [\iota(i)]_Q$, which means that, given any record whose quasi-identifiers are indiscernible from i's quasi-identifiers, the adversary will be able to recognize i with probability at most $\frac{1}{k}$.

4.2 Specification of logical safety

The logical safety criterion was proposed in [6] to prevent homogeneity attacks. Subsequently, it was articulated into an epistemic model for privacy protection in the database linking context [17]. Here, we consider a simplified version of the logical safety criterion. Recall that, in modal logic, the modality-free formulas are called *objective* formulas. Let Γ denote the set of all nominal-free objective formulas, i.e., the set of descriptors closed under Boolean combinations. The logical safety criterion allows a flexible personalized privacy requirements, so each individual can specify the information that he/she wants to keep confidential. More precisely, $Sec : U \to 2^\Gamma$ is such a specification function. According to the semantics of decision logic[10], a pseudonym d satisfies a descriptor (f, v) with respect to σT, denoted by $d \models_{\sigma T} \varphi$, if $f'(d) = v$, and the satisfaction relation is extended to all formulas in Γ as usual. We normally omit the subscript σT. It is said that the adversary knows the individual i has property φ, denoted by $i \models K\varphi$ if, for $d \in [\iota(i)]_Q$, $d \models_{\sigma T} \varphi$. Then, σT satisfies the logical safety criterion if $Sec(i) \cap \{\varphi \mid i \models K\varphi\} = \emptyset$ for $i \in U$.

Theorem 2. *A sanitized information system σT satisfies the logical safety criterion iff $\Vdash_{\sigma T}$ $@_i\neg[a_1]\varphi$ (or equivalently $\Vdash_{\sigma T} @_i l_{a_1}(\varphi) < 1$) for $i \in U$ and $\varphi \in Sec(i)$.*

4.3 Specification of l-diversity

In the same spirit of logical safety, the principle of l-diversity is formulated in [8, 9].

Definition 8. *Let f be a fixed confidential attribute. Then, an equivalence class E of $ind_{\sigma T}(Q)$ is l-diverse if $f'(E) = \{f'(dj) \mid d \in E\}$ contains at least l "well-represented" values, and σT is l-diverse if each of its equivalence classes is l-diverse.*

We consider two instances of l-diversity that are proposed in [8, 9] to articulate the notion of "well-represented" values:

1. Distinct l-diversity. This is the simplest instance of l-diversity. It requires that there are at least l distinct values in $f'(E)$, i.e., $|f'(E)| \geq l$, for each equivalence class E.

2. Recursive (c, l)-diversity. Let $|f'(E)| = m$ and let $k_j (1 \leq j \leq m)$ be the number of times the j^{th} most frequent confidential value appears in the records of E. Then, E satisfies (c, l)-diversity if $k_1 < c(k_l + k_{l+1} + \cdots + k_m)$, and σT satisfies (c, l)-diversity if every equivalence class of $ind_{\sigma T}(Q)$ satisfies it.

For the specification of distinct l-diversity, let us define an (positive) *f-clause* of length m as a disjunctive formula $\bigvee_{j=1}^{m}(f, v_j)$ such that $v_j \neq v_k$ for any $j \neq k$. An f-clause of length 1 is also called an *f-atom*. Then, we have the following result.

Theorem 3. *A sanitized information system σT satisfies the distinct l-diversity iff $\Vdash_{\sigma T} \neg[a_1]\varphi$ for any f-clause φ of length less than l.*

A direct corollary of the theorem shows that distinct l-diversity can be seen as a special case of logical safety.

Corollary 1. *A sanitized information system satisfies the distinct l-diversity iff it satisfies the logical safety criterion with $Sec(i)$ being the set of all f-clauses of length less than l.*

Example 3. This example shows that logical safety is more general and flexible than distinct l-diversity. Let us consider the sanitized information system in Example 2. We assume that the average income of individuals in the community is between 50K and 70K, so any income above this range is considered confidential by an individual. On the other hand, for the health status attribute, an individual may consider serious illness as confidential. Now, the system obviously satisfies distinct 2-diversity for each confidential attribute. However, it may cause problems for an individual if it is known that his income is 100K or he is seriously ill. In such cases, the system would violate the logical safety criterion if $Sec(i)$ includes the wff $(f_{ic}, 100K) \vee (f_{hs}, 2)$ because it would be known that both i_3 and i_4 have this disjunctive property if the system is released to the public. ■

Our logic can also specify recursive (c, l)-diversity, although the specification is a little complicated.

Theorem 4. *A sanitized information system σT satisfies recursive (c, l)-diversity iff for any f-clause $\bigvee_{j=1}^{m} \varphi_j$,*

$$\Vdash_{\sigma T} (\psi_1 \wedge \psi_2 \wedge \psi_3) \supset cl_{a_1}(\varphi_l) + \cdots + cl_{a_1}(\varphi_m) > l_{a_1}(\varphi_1),$$

where $\psi_1 = [a_1] \bigvee_{j=1}^{m} \varphi_j$, $\psi_2 = \bigwedge_{j=1}^{m-1} l_{a_1}(\varphi_j) \geq l_{a_1}(\varphi_{j+1})$, and $\psi_3 = l_{a_1}(\varphi_m) > 0$.

4.4 Specification of t-closeness and δ-disclosure privacy

It is recognized that criteria like k-anonymity and l-diversity are purely *syntactic* in the sense that they only consider the distribution of attribute values in a sanitized system, without measuring how much information an adversary may learn from the publication of the system[2]. On the other hand, several *semantic* criteria, such as the average benefit model[3, 16], t-closeness[7], and δ-disclosure privacy[2] have been proposed to capture the incremental gain in the adversary's knowledge. The common feature of these criteria is that they compare the distribution of attribute values in the sanitized system with that in the trivially sanitized system. The semantic criteria are formulated as the t-closeness principle in [7].

Definition 9. *An equivalence class of $ind_{\sigma T}(Q)$ is said to exhibit t-closeness if the distance between the distribution of a sensitive attribute in that class and the distribution of the attribute in the whole table is no more than a threshold t, and σT satisfies t-closeness if each of its equivalence classes exhibits t-closeness.*

To implement the t-closeness criterion, the distance between two probability distributions must be specified precisely. Let $\alpha = (\alpha_1, \ldots, \alpha_m)$ and $\beta = (\beta_1, \ldots, \beta_m)$ denote two probability distributions over a sample space with m outcomes. The variational distance is defined as follows([7]):

$$D_{var}(\alpha, \beta) = \sum_{j=1}^{m} \frac{1}{2}|\alpha_j - \beta_j| = \sum_{\alpha_j > \beta_j} (\alpha_j - \beta_j) = -\sum_{\alpha_j < \beta_j} (\alpha_j - \beta_j),$$

where the second and third equations hold because $\sum_{j=1}^{m} \alpha_j = \sum_{j=1}^{m} \beta_j = 1$.

Theorem 5. *A sanitized information system σT satisfies the t-closeness criterion based on the variational distance iff for any f-clause $\bigvee_{j=1}^{m} \varphi_j$ and $0 \leq k \leq m$,*

$$\Vdash_{\sigma T} (\psi_1 \wedge \psi_2 \wedge \psi_3) \supset \sum_{j=1}^{k} (l_{a_0}(\varphi_j) - l_{a_1}(\varphi_j)) \leq t,$$

where $\psi_1 = [a_0] \bigvee_{j=1}^{m} \varphi_j$, $\psi_2 = \bigwedge_{j=1}^{k} l_{a_0}(\varphi_j) > l_{a_1}(\varphi_j)$, and $\psi_3 = \bigwedge_{j=k+1}^{m} l_{a_0}(\varphi_j) \leq l_{a_1}(\varphi_j)$.

The difference between syntactic and semantic privacy criteria is easily observed by comparing the above theorem with the preceding ones, since the baseline agent a_0 with the trivial sanitization information does not appear in the logical specification of k-anonymity and l-diversity; however, it plays a crucial role in the formulation of the t-closeness criterion.

The δ-disclosure criterion proposed in [2] is another semantic privacy criterion. It is similar to the average benefit criterion in [3, 16], although the latter is only defined for two-valued attributes. Given a set of records E and a confidential attribute value v, let $p(E, v)$ denote the fraction of records in E whose confidential attribute value is v. Then, an equivalence class E of $ind_{\sigma T}(Q)$ is *δ-disclosure-private* with regard to the confidential attribute f if, for all $v \in V_f$,

$$|\log \frac{p(E, v)}{p(U', v)}| < \delta,$$

and σT is δ-disclosure-private if each equivalence class of $ind_{\sigma T}(Q)$ is δ-disclosure-private.

Theorem 6. *A sanitized information system σT is δ-disclosure-private iff*

$$\Vdash_{\sigma T} (l_{a_1}(\varphi) < 2^{\delta} l_{a_0}(\varphi)) \wedge (l_{a_1}(\varphi) > 2^{-\delta} l_{a_0}(\varphi))$$

for all f-atom φ.

5 Concluding Remarks

In this paper, we propose a probabilistic hybrid logic for the specification of data privacy policies. The logic is expressive and flexible enough to represent many existing privacy criteria, such as k-anonymity, logical safety, l-diversity, t-closeness, and δ-disclosure.

The main contribution of the logic is twofold. On one hand, the *uniformity* of the framework explicates the common principle behind a variety of privacy requirements and highlights their differences. For example, as mentioned in Section 4.4, the difference between syntactic and semantic privacy criteria is easily observed by using the logical specifications. On the other hand, the *generality* of the framework extends the scope of privacy specifications. In particular, we can specify heterogeneous requirements between different individuals, so it is possible to achieve personalized privacy specification. For example, we can use $@_i\neg[a_1]\varphi \wedge @_j\neg[a_1]\psi$ to express different privacy requirements of individuals i and j.

Moreover, the logic allows arbitrary combinations of existing privacy requirements, so we can express compound privacy criteria. For example, we can use $@_i\neg[a_1]\varphi \wedge l_{a_1}(i) \leq \frac{1}{k}$ to express that both logical safety and k-anonymity are required for the individual i. Since unexpected attacks may occur occasionally, existing criteria may be inadequate; hence, it may be necessary to specify new criteria. For example, the logical safety criterion may be combined with δ-disclosure to require formulas in $Sec(i)$, instead of simply f-atoms, to satisfy the δ-disclosure privacy criterion. In addition, it is possible to consider the weight of a secret in order to measures the seriousness of revealing the secret. Thus, $Wsec : U \times \Gamma \rightarrow [0, 1]$ is defined as the weight function for each individual and secret. Then, we can combine the weight with existing privacy criteria to obtain new privacy protection models. This may facilitate a more effective tradeoff between privacy protection and data utility. Our logic language provides a uniform framework to meet the specification needs of such new criteria as well as existing ones.

References

[1] C. Areces and B. ten Cate. Hybrid logics. In P. Blackburn, J. van Benthem, and F. Wolter, editors, *Handbook of Modal Logic*, pages 821–868. Elsevier, 2007.

[2] J. Brickell and V. Shmatikov. The cost of privacy: destruction of data-mining utility in anonymized data publishing. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 70–78, 2008.

[3] 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 *Proceedings of the 3rd International Conference on Rough Sets and Current Trends in Computing*, LNCS 2475, pages 226–233. Springer-Verlag, 2002.

[4] T. Dalenius. Finding a needle in a haystack - or identifying anonymous census records. *Journal of Official Statistics*, 2(3):329–336, 1986.

[5] J. Halpern. *Reasoning about Uncertainty.* The MIT Press, 2003.

[6] T.-s. Hsu, C.J. Liau, and D.W. Wang. A logical model for privacy protection. In *Proceedings of the 4th International Conference on Information Security*, LNCS 2200, pages 110–124. Springer-Verlag, 2001.

[7] N. Li, T. Li, and S. Venkatasubramanian. t-closeness: Privacy beyond k-anonymity and l-diversity. In *Proc. of the 23rd International Conference on Data Engineering (ICDE)*, pages 106–115, 2007.

[8] A. Machanavajjhala, J. Gehrke, D. Kifer, and M. Venkitasubramaniam. l-diversity: Privacy beyond k-anonymity. In *Proc. of the 22nd IEEE International Conference on Data Engineering (ICDE)*, page 24, 2006.

[9] A. Machanavajjhala, J. Gehrke, D. Kifer, and M. Venkitasubramaniam. l-diversity: Privacy beyond k-anonymity. *ACM Transactions on Knowledge Discovery from Data*, 1(1), 2007.

[10] Z. Pawlak. *Rough Sets–Theoretical Aspects of Reasoning about Data.* Kluwer Academic Publishers, 1991.

[11] P. Samarati. Protecting respondents' identities in microdata release. *IEEE Transactions on Knowledge and Data Engineering*, 13(6):1010–1027, 2001.

[12] P. Samarati and L. Sweeney. Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression. Technical report SRI-CSL-98-04, Computer Science Laboratory, SRI International, 1998.

[13] L. Sweeney. Guaranteeing anonymity when sharing medical data, the datafly system. A.I. Working Paper AIWP-WP344, MIT AI Lab., 1997.

[14] L. Sweeney. Achieving k-anonymity privacy protection using generalization and suppression. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, 10(5):571–588, 2002.

[15] L. Sweeney. k-anonymity: a model for protecting privacy. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, 10(5):557–570, 2002.

[16] D.W. Wang, C.J. Liau, and T.-s. Hsu. Medical privacy protection based on granular computing. *Artificial Intelligence in Medicine*, 32(2):137–149, 2004.

[17] D.W. Wang, C.J. Liau, and T.-s. Hsu. An epistemic framework for privacy protection in database linking. *Data and Knowledge Engineering*, 61(1):176–205, 2007.