Privacy Intrusion Detection Using Dynamic Bayesian Networks

Xiangdong An†‡& Dawn Jutla†
†Finance and Management Science Department
Saint Mary's University
Halifax, NS B3H 3C3, Canada
{xan,dawn.jutla}@smu.ca

Nick Cercone‡
‡Faculty of Computer Science
Dalhousie University
Halifax, NS B3H 1W5, Canada
{xan,nick}@cs.dal.ca

ABSTRACT

Concerns for personal information privacy could be produced during information collection, transmission and handling. In information handling, privacy could be compromised from both inside and outside of organizations. Within an organization, private data are generally protected by organizations' privacy policies and the corresponding platforms for privacy practices. However, private data could still be misused intentionally or unintentionally by individuals who have legitimate accesses to them. In general, activities of a database operator form a stochastic process, and at different time, privacy intrusion behavior may show different features. In particular, one's past activities can help determine the natures of his/her current practices. In this paper, we propose to use dynamic Bayesian networks to model such temporal environments and detect any privacy intrusions happened within them.

Categories and Subject Descriptors

H.2.7 [**Database Management**]: Database Administration—*Privacy protection*; I.2.3 [**Artificial Intelligence**]: Deduction and Theorem Proving—*Uncertainty, "fuzzy", and probabilistic reasoning*; I.5.4 [**Pattern Recognition**]: Applications—*Plan recognition*

General Terms

Management, Design, Security

Keywords

Privacy intrusion, intrusion detection, probabilistic reasoning, dynamic Bayesian networks

1. INTRODUCTIONS

Privacy is "the right to be left alone" [46]. Before the 1960s, people were mainly concerned about their (1) bodily privacy, (2) territorial privacy, and (3) communication privacy. With the advent of information technology (IT), concerns to information privacy appeared and increased quickly in the 1960s and 1970s [23]. Information privacy involves personal information collection and handling. Today, with rapid advancements in IT and its broad applications, privacy concerns have been produced in many aspects of our lives such as e-commerce, healthcare, financial services, wireless communication, video surveillance, biometrics, and data profiling [10]. It is expected that personal information must be: (1) obtained fairly and lawfully, (2) used only for the original specified purpose, (3) accessible to the subject, (4) kept secure and accurate, and (5) destroyed after its purpose is completed [16, 24, 17]. Private information processing following the five requirements is called fair information practices [35].

Information collection is essential for an enterprise to conduct its business. If privacy concerns are not properly solved, information losses may become financial losses. Therefore, information privacy has not only been a human rights issue, but has also been a social issue, an economic issue, as well as a science and technology issue [41].

The platform for Privacy Preferences (P3P) [11] has been widely used for encoding enterprises' privacy policies. P3P enables users to check websites' privacy policies before allowing websites to collect their personal information. However, the promises published using P3P may not be enforced throughout the business by enterprises. Especially, P3P is too coarse to express an internal enforceable privacy policy [40]. The Platform for Enterprise Privacy Practices (E-P3P) [27, 6] has been proposed to enable enterprises to enforce their policies properly. Enterprise Privacy Authorization Language (EPAL) [5] is the successor of E-P3P.

Though EPAL or E-P3P helps organizations enforce their policies, it is still the responsibility of the organizations, in particular their employees, to respect the applicable policies [7]. After one's information has been collected by an organization, privacy protection generally depends on how the organization and its employees are going to follow and enforce the organization's privacy policies. In particular, employees who have legitimate access to the database could use private information for illegal purposes. That is, users still need to trust the organization in protection of their personal data, especially when data could be across orga-

ICEC '06, August 14-16, 2006, Fredericton, Canada.

nization boundaries. On the other hand, the organization is accountable for misuse of private information by its employees (e.g. this is one of the ten basic principles outlined in Canada's PIPEDA — Personal Information Protection and Electronic Documents Act — fully enforced on January 1, 2004). Therefore, the organization is obliged to identify possible misuses of private data by its employees.

Methods have been proposed to enhance the protection of private data after they have been disclosed to organizations [30, 1, 21]. These methods are based on public key encryption and/or require users to trust the "trusted agents". However, the "trusted agents" could compromise private data, and the decrypted private data could still be misused. Another work [3] proposed a roadmap for a framework involving policy creation, enforcement, analysis and auditing. Regarding the auditing of the privacy policies, it suggested to develop a new policy language which should be able to specify auditing requirements for data accesses so that audit trails can be generated. In [8], such a policy language was indeed presented which supports the specification of conditions and obligations. However, how to audit was not discussed.

A privacy intrusion detection system [45] has been proposed to find and prevent any misuses of private data collected by an organization. Nevertheless, the authors of the paper simply suggested to detect abnormal activities by comparing an operator's activities with the normal user or system behavior profiles. They did not specify how to effectively handle uncertainties involved in operators' activities (e.g. an operator's occasionally assessing a database for an extended time could be demanded by the job, but could also be for illegal purposes) and how to effectively reduce such uncertainties using dependencies[1] among operators' activities (e.g. an operator needs to access a database for an extended time while he/she is processing tax returns at tax return peak period and he/she does process a lot of returns daily).

Bayesian networks (BNs) [36] can properly model such uncertainties and dependencies to effectively reason about the states of domains [48, 25, 28]. A Bayesian network model [2] has been proposed to detect internal privacy intrusion by organizations' employees. However, Bayesian networks assume the problem domains are static. They model domains with coarser granularity without recognizing temporal dependencies among domain events. That is, the effects of old domain states on the current domain state cannot be properly modeled using BNs.

Privacy intrusions happen in dynamic environments, and even they themselves could be time-series data. In this paper, we propose to use dynamic Bayesian networks (DBNs) [14] to detect internal privacy intrusion originated within organizations. Dynamic Bayesian networks can capture richer and more realistic domain dependencies [13]. They have been applied in dynamic domains for domain state monitoring [33, 32, 29], activity or plan recognition [4], forecasting [12], speech recognition [31, 13], medical diagnosis [22], and fault or defect detection [44, 39]. Nevertheless, their application to intrusion detection is rare reported [9, 26, 20]. Hidden Markov models (HMMs) [38] — considered simplest DBNs — have been applied in network intrusion detection by modeling time series on networks [47, 34]. However, HMMs use a single variable to represent domain states at a time instant, which cannot properly take advantage of conditional independencies among domain events. Hence, inference with HMMs would be significantly slow [42]. An extended HMM, called a factorial HMM [19], was once used for network traffic classification [49]. However, in the extended two Markov chain HMM, at each time instant, only 4 variables were used to model domain state, which still could not explore domain independencies fully. In general, DBNs are generalized HMMs being able to fully represent domain independencies and perform inference efficiently [51, 18, 4].

To our best knowledge, this is the first work of applying general dynamic Bayesian networks to intrusion detection, in particular, privacy intrusion detection. Besides online real-time privacy intrusion detection, the method can also be applied to offline privacy intrusion auditing.

The paper is organized as follows. In Section 2, we give a review on dynamic Bayesian networks. In Section 3, we describe a scenario where privacy intrusion could happen and propose a DBN model to detect any possible intrusion instances happened in it. In Section 4, we discuss the approach further. Conclusions are made in Section 5.

2. DYNAMIC BAYESIAN NETWORKS

Dynamic Bayesian networks (DBNs) are graphical models for probabilistic inference in dynamic domains, which are extended from Bayesian networks (BNs) for static domains. DBNs provide us an easy and compact way to specify the conditional independencies in dynamic domains.

A DBN consists of a finite number of BNs, each of which (called a *slice* of the DBN) corresponds to a particular time instant (or interval). BNs corresponding to successive instants are connected through arcs that represent how the state of the domain evolves over time. Like BNs, the structures of DBNs are directed acyclic graphs (DAGs), where each node represents a domain variable of interest at some time instant, and each directed arc represents the dependency between the two nodes it connects. The strength of each dependency is quantified by a conditional probability distribution (CPD) specifying the probabilities of the child being in specific values given the values of its parents. For simplicity, a DBN is generally assumed to satisfy the *Markovian property*: the state of the domain at time $t+1$ is independent of the states of the domain prior to time t, given the state of the domain at time t. In particular, like in a BN, each node in a DBN is conditionally independent of its nondescendants given its parents. These properties allow us to solve complex problems by cheaper local computations.

In a DBN, we also generally assume the nodes, the dependencies among nodes, and the strength of the dependencies at slice i are identical to those at slice j. In particular, we assume the dependency and its strength between a pair of nodes across two consecutive instants won't change over time. Hence, a DBN can be described by a two-slice DBN (2TBN) and the entire DBN can be obtained by unrolling the 2TBN.

Consider an example where an old man named Bob lives in a city far from the town his friend Peter lives. They talk with each other by phone every evening. Bob usually walks around if the weather of his city is *good*, and stays home otherwise. Peter does not know the weather condition of Bob's city, but can judge it from Bob's activities on that

[1] There may also exist uncertainties on dependencies, though.

day. The example can be represented by a DBN as shown in Figure 1. Note that the simplest DBN is actually an HMM for illustration purpose, which is not suitable for modeling complex domain as discussed above. We are going to show how general DBNs are used in privacy intrusion detection in Section 3.

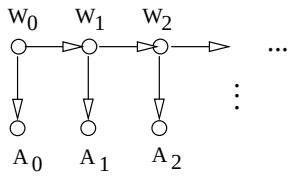

Figure 1: A DBN on the weather and an old man's activities.

In the DBN, each node represents a random variable, and each arrowed arc represents the causal dependencies between the two nodes connected. The subscript of a variable indicates the corresponding time instant. Hence, in the DBN, each instant is represented by two random boolean variables, W_i and A_i ($i = 0, 1, 2, ...$). The variable W_i represents the weather condition of Bob's city on day i, taking the value *good* when the weather is good and *bad* otherwise. The variable A_i represents Bob's activities on day i, taking the value *out* when Bob walks around, and *in* otherwise. Note that it is possible that Bob walks around while the weather is bad and vice versa. Nevertheless, Bob's activities are highly dependent on the weather conditions. The arrowed arc between two successive slices represents the evolution of the weather condition. The weather will very probably remain good next day if it is good today and bad if bad.

The DBN can be fully described by its first two slices: the structure and parameters of each new slice repeat those of the preceding one. The parameters are specified in Tables 1 and 2. The parameters specified in Table 2 will be repeated slice by slice. Therefore, to describe a DBN, the first one and a half slices are actually enough (i.e. 1.5TBN).

In Bayesian probability theory, probabilities are subjective corresponding to the degree of belief of the reasoners in the truth of the statements. The degree of belief is different from the degree of truth. People with different prior knowledge could *correctly* obtain different results from Bayesian probability theory.

Table 1: The prior belief about the weather condition.

W_0	P(W_0)
good	0.75
bad	0.25

In the DBN as shown in Figure 1, variable A_i, representing the observation on day i, is often called *information*, *evidential* or *observable* variable. Variable W_i, representing the actual weather condition on day i, is called a *hypothesis* or *unobservable* variable. DBNs can, based on evidence collected from information variables, help efficiently evaluate the probabilistic state of hypothesis variables.

In our case, with the model, Peter is able to determine the probabilities of the weather being good on some days given Bob's activities for a few days. For instance, by talking over phone, Peter knows Bob's activities in last 3 days are as follows: walked around on day 0, stayed home on day 1, and walked around on day 2. By inference using the model, Peter knows it is 88.29% that the weather on day 0 was good, 65.35% on day 1 bad, and 78.52% day 2 good. The model can also predict that the weather on next day will be good by 66.41%, and Bob will walk around by 66.48%.

Table 2: The prior belief about the relationship between the weather condition and Bob's activities, and the evolution of the weather condition.

W_i	A_i	P(A_i|W_i)	W_i	W_{i+1}	P(W_{i+1}|W_i)
good	in	0.10	good	good	0.75
good	out	0.90	good	bad	0.25
bad	in	0.80	bad	good	0.65
bad	out	0.20	bad	bad	0.35

In the domain of privacy intrusion detection, information variables are associated with measurable properties of database and operators' activities (for online monitoring) or logged events (for auditing). Unobservable variables are associate with immeasurable properties of database or operators' activities. In this paper, we assume there are two hypothesis variables: the one used to denote whether a privacy intrusion is occurring (has occurred) or not (has not), and the other one used to denote what tasks an employee is (was) working on.

3. PRIVACY INTRUSION DETECTION

Privacy intrusion could happen in many industries where employees have chances to manipulate databases with private information such as health care, revenue agencies, financial services, etc. In different industries, the DBN privacy intrusion detection models could be different in variables and their dependencies due to different private data contents and data manipulation processes in different domains. However, the working principles of DBN privacy intrusion detection will be the same. In this section, as an example, we present a DBN model for privacy intrusion detection in the government's revenue agencies.

3.1 The Problem and the Method

In a government's revenue service, a representative is granted access to some information in the revenue database. For example, a representative usually has rights to retrieve taxpayers' personal information such as income, family composition, contact details, etc. Suppose the revenue service has a privacy policy that does not allow its employees to disclose such personal information to any third parties. Nevertheless, a representative could violate the policy by secretly disclosing such personal information to somebody else who is interested in that (e.g. a marketing agent). Current techniques such as access control or EPAL can do nothing in preventing such privacy invasion.

In [45], a set of features used to recognize anomaly activities were summarized as shown in Table 3, where "DB" represents database and "RDs" denotes records. Once the threshold regarding an activity is violated, the respective action will be taken to slow down or stop the possible anomaly activity.

However, simply observing the time an operator stays in

a database or the number of times an operator logins a database is not sufficient to determine privacy intrusion with high certainty, which may cause false alarms easily instead. In particular, in [45], all these features are evaluated individually. We believe the detection can be made with more certainty when all these features are considered in a comprehensive way. A Bayesian network (BN) model [2] has been proposed to properly combine contributions of all these features in detecting privacy intrusion. Especially, a measure about the degree of suspiciousness of intrusion is introduced through the BN model.

Table 3: Features used to recognize anomaly activities.

Features	Thresholds	Objects
Working Hours	8:00-17:00	Database
Duration in DB	10 minutes	Database
Duration on RDs	3 minutes	Records
Amount of RDs	100	Records
Modification	0	Records
Frequency on DB	10	Database
Frequency on RDs	3	Records

Nevertheless, BNs are proposed for modeling static domains, which lack facilities to model historical information. In general, the activities of a representative in a revenue agency form a stochastic process. For example, a representative could be assigned some job which needs significant access time to database in the following several days. In particular, in different days, the representative may need to handle different tasks in some order. Depending on the specific task, he/she may or may not need to access some records, or to perform some operations. Operations for a task may be causally related with each other. All these dynamic information may tell us much more about a representative's activities, hence help us reason about the privacy intrusion with more certainty.

Dynamic Bayesian networks (DBNs) provide effective facilities to model time series data to reason about the states of dynamic domains. In the next Subsection, with an example we show how DBNs can be used to detect possible internal privacy intrusion within a revenue agency.

3.2 An Example

Suppose John works in a revenue agency as a representative. His office hours are 8:00am-17:00pm every weekday. His duties include audit of tax returns and/or applications, and collection of outstanding accounts receivables and/or delivery of benefits. He usually first audit a set (less than 100) of returns/applications, then print and send out collection/delivery letters corresponding to the reviewed accounts.

In general, auditing takes more time than printing and sending collection/delivery letters. Hence, John spends more time on auditing than on collection/delivery. While auditing, he accesses databases and records more frequently, and spends more time on both databases and records. He may modify records as the consequences of auditing. He usually turns to printing/sending corresponding collection/delivery letters before 100 returns/applications have been audited. While on auditing, he is probably trying to intrude clients' privacy if he spends less than usual time on databases but accesses more than usual records.

While printing/sending collection/delivery letters, he accesses databases and records less frequently, and spends less time on databases and records. He usually does not need to modify any records at this time. In collection/delivery period, he is probably trying to intrude clients' privacy if he spends more than usual time on databases or accesses more than usual records.

John almost does not work out of his office hours. If he does, the probability he intrudes clients' private information becomes higher than that he does within his office hours. His probability of privacy intrusion will be adjusted by his working behaviours: the probability will become lower and lower if he has been working without suspected intrusion operations, and become higher if he performs suspected operations.

The example can be modeled with a DBN. As we discussed above, a DBN can be described by its first 2 (or 1.5) slices. The first two slices of the DBN is shown in Figure 2, where each dotted box denotes one slice, and arcs across boxes indicate the evolution of domain states. In each slice, "F_d" denotes "usage frequency of databases", "F_r" "usage frequency of records", "T_r" "time spent on records", "M_r "modification of records", "Tk" "task" (audit or collection/delivery), "$Intr$" "intrusion of privacy" (true or false), "Hrs" "working hours" (8am-17pm or not), "A_r" "amount of records", and "T_d" "time on databases". The digits (0 or 1) following these labels indicate the respective slice number (e.g. $Tk1$ represents "the task at slice 1"). The arc between $Tk0$ and $Tk1$ represents the evolution of John's tasks: John will highly probably (80%) audit in the next period if he audits currently, and will more probably (75%) collect/deliver if he collects/delivers currently. The arc between $Intr0$ and $Intr1$ represents the evolution of John's privacy intrusion: John will more probably intrude clients' privacy in the future if he does currently (or did in the past), and will less probably intrude clients' privacy if he does not currently (or did not in the past). The more recent the intrusion happened, the more the intrusion probability will be raised. The longer the intrusion has not happened, the less the intrusion probability will become. The parameters are specified in Tables 4, 5, 6, 7 and 8.

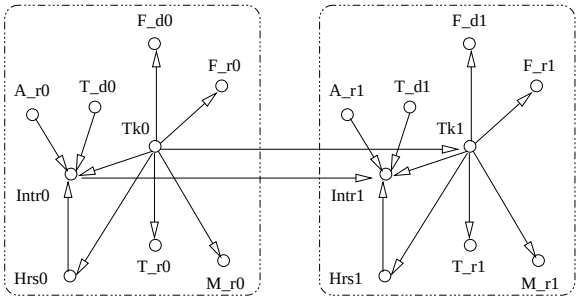

Figure 2: A Dynamic Bayesian network for privacy intrusion detection.

In these tables, "aud" denotes "audit", and "col" "collection/delivery". "10-" denotes less than 10, and "10+" more than 10. Similarly, "3-" and "100-" denotes less than 3 and 100 respectively, and "0+", "3+" and "100+" more than 0, 3 and 100 respectively. "8-17" indicates the office hours between 8:00am and 5:00pm, and "¬8-17" indicates the time other than office hours. The dots (.'s) in $P(Intr_0|.)$

Table 6: Conditional probability distribution table for F_d_i, F_r_i, M_r_i, T_r_i, and Hrs_i ($i \geq 1$).

Tk_i	F_d_i	$P(F\_d_i\|Tk_i)$	F_r_i	$P(F\_r_i\|Tk_i)$	M_r_i	$P(M\_r_i\|Tk_i)$	T_r_i	$P(T\_r_i\|Tk_i)$	Hrs_i	P($Hrs_i\|Tk_i$)
aud	10-	0.45	3-	0.35	0	0.10	3-	0.25	8-17	0.95
aud	10+	0.55	3+	0.65	0+	0.90	3+	0.75	¬8-17	0.05
col	10-	0.99	3-	0.95	0	0.99	3-	0.90	8-17	0.99
col	10+	0.01	3+	0.05	0+	0.01	3+	0.10	¬8-17	0.01

Table 7: Prior probabilistic knowledge about $Intr_0$.

Tk_0	Hrs_0	T_d_0	A_r_0	$Intr_0$	$P(Intr_0\|.)$	Tk_0	Hrs_0	T_d_0	A_r_0	$Intr_0$	$P(Intr_0\|.)$
aud	8-17 (¬8-17)	10-	100-	true	0.60 (0.80)	col	8-17 (¬8-17)	10-	100-	true	0.15 (0.65)
aud	8-17 (¬8-17)	10-	100-	false	0.40 (0.20)	col	8-17 (¬8-17)	10-	100-	false	0.85 (0.35)
aud	8-17 (¬8-17)	10-	100+	true	0.70 (0.90)	col	8-17 (¬8-17)	10-	100+	true	0.25 (0.70)
aud	8-17 (¬8-17)	10-	100+	false	0.30 (0.10)	col	8-17 (¬8-17)	10-	100+	false	0.75 (0.30)
aud	8-17 (¬8-17)	10+	100-	true	0.25 (0.75)	col	8-17 (¬8-17)	10+	100-	true	0.60 (0.80)
aud	8-17 (¬8-17)	10+	100-	false	0.75 (0.25)	col	8-17 (¬8-17)	10+	100-	false	0.40 (0.20)
aud	8-17 (¬8-17)	10+	100+	true	0.40 (0.85)	col	8-17 (¬8-17)	10+	100+	true	0.65 (0.85)
aud	8-17 (¬8-17)	10+	100+	false	0.60 (0.15)	col	8-17 (¬8-17)	10+	100+	false	0.35 (0.15)

Table 4: Prior probabilistic knowledge about Tk_0, A_r_i and T_d_i ($i \geq 1$).

Tk_0	P(Tk_0)	A_r_i	P(A_r_i)	T_d_i	P(T_d_i)
aud	0.75	100-	0.75	10-	0.50
col	0.25	100+	0.25	10+	0.50

Table 5: Probabilistic evolution knowledge about Tk_i ($i \geq 1$).

Tk_{i-1}	Tk_i	P($Tk_i\|Tk_{i-1}$)
aud	aud	0.80
aud	col	0.20
col	aud	0.25
col	col	0.75

and $P(Intr_i|.)$ denote the abbreviation of the corresponding conditions. The conditional probability values in brackets correspond to the conditions where values in brackets are used if applicable. For example, in the first row of left part of Table 7, the probability 0.80 in the bracket corresponds to $P(Intr_0 = true|Hrs_0 = \neg 8-17, T\_d_0 = 10-, A\_r_0 = 100-)$.

Before any observation (evidence) is entered, the DBN shows that John, currently, is mostly probably (75%) auditing returns/applications, and the probability he audits reduces with time going on (e.g. from 75% on the first period to 58.80% on the fourth period). That is, the longer he audits, the more probable he could turn to collect/deliver. He usually works within his office hours (with a probability of about 96% on each day). His privacy intrusion probability is 45.44% initially (assuming we are not too confident about if he could intrude clients' private information), and reduces slightly with time going on (e.g. from 45.44% on the first period to 42.84% on the fourth period). That the privacy intrusion probability reduces with time going on is because we assume previous privacy intrusion will increase one's future intrusion probability, and one's non-intrusion credits will reduce his future intrusion probability.

Assume we start to observe and reason about John's activities from a certain day. John's activities on one day corresponds to one slice of the DBN. As mentioned above, we assume John's working task Tk_i ($i \geq 0$), and the nature of his activities (whether he intrudes clients' privacy $Intr_i$ ($i \geq 0$)) are not observable. In the first two days, all evidence indicates he is auditing returns and/or applications ($F\_d_{0,1}$=10+,$F\_r_{0,1}$=3+,$M\_r_{0,1}$=0+, $T\_r_{0,1}$=3+). With such evidence entered, the model indicates that John is highly probably (99.99% on each day) auditing returns/applications. Nevertheless, without further intrusion-specific evidence entered, the model cannot tell us too much about privacy intrusion (the model shows the intrusion probabilities for the two days are 47.34% and 46.81%, respectively). In this model, the intrusion-specific variables are A_r_i, T_d_i, and Hrs_i ($i \geq 0$). With evidence on these variables for the first two days obtained and entered ($A\_r_{0,1}$=100-,$T\_d_{0,1}$=10+, and $Hrs_{0,1}$=8-17), the model indicates that John has very low probabilities for privacy intrusion (25% and 20% respectively for each day). The intrusion probabilities become lower because the evidence indicates none of the suspected activities.

On the third day, suppose we observe F_d_2=10-,F_r_2=3-,M_r_2=0, and T_r_2=3-. With such evidence entered, the model indicates that John is highly probably (98.16%) doing collection/delivery job. Before intrusion-specific evidence is entered, the model shows that John could intrude private information with a probability of 33.86% on the day. With intrusion-specific evidence entered (A_r_2=100+, T_d_2=10+, and Hrs_2=¬8-17), the model indicates that John could intrude clients' private information with a probability of 79%. This is because the evidence indicates highly suspected activities. This high intrusion probability will raise John's future privacy intrusion probability through intrusion variable dependency between two successive DBN slices. For example, his intrusion probability on the fourth day becomes 47.78%, a rise from the initial 42.84% (before any evidence is entered), and 38.48% before intrusion-specific evidence on the third day is entered. This rise could significantly increase the model's sensitivity to John's privacy intrusion-

Table 8: Prior probabilistic knowledge about $Intr_i$ ($i \geq 1$).

Tk_i	Hrs_i	T_d_i	A_r_i	$Intr_{i-1}$	$Intr_i$	$P(Intr_i\|.)$	Tk_i	Hrs_i	T_d_i	A_r_i	$Intr_{i-1}$	$Intr_i$	$P(Intr_i\|.)$
aud	8-17 (¬8-17)	10-	100-	true	true	0.70 (0.90)	col	8-17 (¬8-17)	10-	100-	true	true	0.25 (0.75)
aud	8-17 (¬8-17)	10-	100-	true	false	0.30 (0.10)	col	8-17 (¬8-17)	10-	100-	true	false	0.75 (0.25)
aud	8-17 (¬8-17)	10-	100-	false	true	0.50 (0.70)	col	8-17 (¬8-17)	10-	100-	false	true	0.05 (0.55)
aud	8-17 (¬8-17)	10-	100-	false	false	0.50 (0.30)	col	8-17 (¬8-17)	10-	100-	false	false	0.95 (0.45)
aud	8-17 (¬8-17)	10-	100+	true	true	0.80 (0.95)	col	8-17 (¬8-17)	10-	100+	true	true	0.35 (0.80)
aud	8-17 (¬8-17)	10-	100+	true	false	0.20 (0.05)	col	8-17 (¬8-17)	10-	100+	true	false	0.65 (0.20)
aud	8-17 (¬8-17)	10-	100+	false	true	0.60 (0.85)	col	8-17 (¬8-17)	10-	100+	false	true	0.15 (0.60)
aud	8-17 (¬8-17)	10-	100+	false	false	0.40 (0.15)	col	8-17 (¬8-17)	10-	100+	false	false	0.85 (0.40)
aud	8-17 (¬8-17)	10+	100-	true	true	0.35 (0.85)	col	8-17 (¬8-17)	10+	100-	true	true	0.70 (0.90)
aud	8-17 (¬8-17)	10+	100-	true	false	0.65 (0.15)	col	8-17 (¬8-17)	10+	100-	true	false	0.30 (0.10)
aud	8-17 (¬8-17)	10+	100-	false	true	0.15 (0.65)	col	8-17 (¬8-17)	10+	100-	false	true	0.50 (0.70)
aud	8-17 (¬8-17)	10+	100-	false	false	0.85 (0.35)	col	8-17 (¬8-17)	10+	100-	false	false	0.50 (0.30)
aud	8-17 (¬8-17)	10+	100+	true	true	0.50 (0.50)	col	8-17 (¬8-17)	10+	100+	true	true	0.75 (0.95)
aud	8-17 (¬8-17)	10+	100+	true	false	0.50 (0.50)	col	8-17 (¬8-17)	10+	100+	true	false	0.25 (0.05)
aud	8-17 (¬8-17)	10+	100+	false	true	0.30 (0.30)	col	8-17 (¬8-17)	10+	100+	false	true	0.55 (0.75)
aud	8-17 (¬8-17)	10+	100+	false	false	0.70 (0.70)	col	8-17 (¬8-17)	10+	100+	false	false	0.45 (0.25)

related activities. For example, with the same evidence as that on the third day observed on the fourth day, the model shows that John could intrude private information with a probability of 90.8%.

4. DISCUSSION

The approach presented in the paper aims to deal with the general internal attacks on databases for stealing large amount of private data. It is not generally intended to detect privacy breach of an or a few individuals, though the model can be made to detect such invasions if the data are irrelevant to the database operator's (current) job. It is difficult to automate the detection of the disclosure of such small amount of data generally.

The effectivity of the framework counts on the applicability of the features used. In different problem domains, the features used to detect internal privacy intrusion will be different. The features proposed in the paper for privacy intrusion detection in a revenue agency is not complete and may not be the most effective depending on the specific operator and his task. Features can be obtained manually or statistically. To obtain features manually, we need to have a good knowledge about the database operator's activities. The more we know about an operator, the more certain our knowledge about the operator is, the more effective the framework will be. To increase and make certain our knowledge about an operator's activities, we could make some specific regulations regarding procedures and operations he needs to follow in his routine business. For example, he should not access or cannot copy some records; print-outs will be monitored, etc. Obtaining features statistically is a problem of DBN learning from live data, which will be investigated in the future. Anyway, feature creation and selection is an interesting aspect of the future work.

The proposed technique will be particularly effective if the large amount of data being accessed is irrelevant to the operator's job since the irrelevancy is an outstanding feature. More features that can determine the nature of one's job will help determine the irrelevance.

Detection of irrelevance between data being accessed and one's job can also help prevent users from inferring sensitive information from (statistical) databases, which is known as inference control [15, 37, 50]. A user may be able to deduce sensitive data by making a sequence of non-sensitive queries that is irrelevant to his job. DBNs can help detect such irrelevances, and determine if these queries together form an inference channel [43].

5. CONCLUSION

Private information is usually protected by secure techniques (such as access control and authentication). Data privacy will be compromised if there is no data security. However, even when data are "secure", private information could still be infringed on. In privacy protection, we are more concerned about how to protect data when data have to be disclosed or even when data are "secure". Many techniques such as P3P and EPAL have been proposed to protect private information from being unfairly used or processed. Nevertheless, these techniques may not prevent persons misusing their legitimate access rights from invading data privacy.

It turns out that database operators' activities could tell us if they are trying to invade data privacy. Since database operators' activities are generally stochastic processes (the task of a database operator, and hence his specific operations, could change from time to time; the previous activities of an operator could be related with his/her current activities; one's privacy intrusion history may indicate he/she is highly interested in private information, etc), we propose to use dynamic Bayesian networks (DBNs) to detect privacy intrusion in such temporal environments. An example is given in the paper to demonstrate how DBNs can be applied to privacy intrusion detection.

With DBNs, the past (historical) activities of an operator can be used to help determine the nature of his/her current activities or even help forecast the nature of his/her future activities. The effects of (contributions from) all historical and current events to the intrusion probability are coherently combined using the Bayesian probability theory. The DBN method can be applied to both online privacy intrusion detection and offline privacy intrusion auditing. To our best knowledge, this is the first work to apply dynamic Bayesian

networks to the topic.

In the future, we are going to study how to derive a DBN privacy intrusion detection model from real data.

6. ACKNOWLEDGEMENT

We thank anonymous reviewers for their helpful comments.

7. REFERENCES

[1] G. Aggarwal, M. Bawa, P. Ganesan, H. Garcia-Molina, K. Kenthapadi, N. Mishra, R. Motwani, U. Srivastava, D. Thomas, J. Widom, and Y. Xu. Enabling privacy for the paranoids. In M. A. Nascimento, M. T. Özsu, D. Kossmann, R. J. Miller, J. A. Blakeley, and K. B. Schiefer, editors, *Proceedings of the 13th International Conference on Very Large Data Bases*, pages 708–719, Toronto, Canada, August 31-September 3 2004. Morgan Kaufmann.

[2] X. An, D. Jutla, and N. Cercone. Bayesian network privacy intrusion detection. Technical report, Faculty of Computer Science, Dalhousie University, NS, Canada, 2006.

[3] A. I. Antón, E. Bertino, N. Li, and T. Yu. A roadmap for comprehensive online privacy policy. Technical report, CERIAS, Purdue University, West Lafayette, CERIAS-2004-47, 2004.

[4] L. Ardissono, P. Brna, and A. Mitrovic, editors. *A comparison of HMMs and dynamic Bayesian networks for recognizing office activities*, volume 3538 of *Lecture Notes in Computer Science (LNCS)*, Edinburgh, Scotland, UK, July 24-29 2005. Springer.

[5] P. Ashley, S. Hada, G. Karjoth, C. Powers, and M. Schunter. Enterprise privacy authorization language (EPAL 1.2). Technical report, W3C Member Submission, http://www.w3.org/Submission/2003/SUBM-EPAL-20031110, November 2003.

[6] P. Ashley, S. Hada, G. Karjoth, and M. Schunter. E-P3P privacy policies and privacy authorization. In *Proceedings of Workshop on Privacy in the Electronic Society (WPES'02)*, pages 103–109, Washington, DC, USA, November 21 2002.

[7] A. Barth, J. C. Mitchell, and J. Rosenstein. Conflict and combination in privacy policy languages. In *Proceedings of Workshop on Privacy in the Electronic Society (WPES'04)*, Washington, DC, USA, October 28 2004. ACM Press.

[8] J. G. Cederquist, R. Corin, M. A. C. Dekker, S. Etalle, and J. I. den Hartog. An audit logic for accountability. In A. Sahai and W. H. Winsborough, editors, *Proceedings of the 6th IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY'05)*, pages 34–43, Stockholm, Sweden, June 6-8 2005. IEEE Computer Society.

[9] S. Chebrolu, A. Abraham, and J. P. Thomas. Feature deduction and ensemble design of intrusion detection systems. *Computers & Security*, 24(4):295–307, 2005.

[10] P. R. ClearingHouse. Privacy today: A review of current issues. Technical report, Privacy Rights Clearinghouse, San Diego, CA, USA. http://www.privacyrights.org, 2002.

[11] L. Cranor, M. Langheinrich, M. Marchiori, M. Presler-Marshall, and J. Reagle. The platform for privacy preferences 1.0 (P3P 1.0) specification. Technical report, W3C Recommendation, http://www.w3.org/TR/P3P, April 2002.

[12] P. Dagum, A. Galper, and E. Horvitz. Dynamic network models for forecasting. In D. Dubois, M. P. Wellman, B. D'Ambrosio, and P. Smets, editors, *Proceedings of the 8th Conference on Uncertainty in Artificial Intelligence (UAI-1992)*, pages 41–48, Stanford, CA, USA, 1992. Morgan Kaufmann Publishers.

[13] P. Dagum, A. Galper, E. Horvitz, and A. Seiver. Uncertain reasoning and forecasting. *International Journal of Forecasting*, 11(1):73–87, 1995.

[14] T. Dean and K. Kanazawa. Probabilistic temporal reasoning. In *Proceedings of the 7th National Conference on Artificial Intelligence (AAAI-1988)*, pages 524–528, St. Paul, Minnesota, 1988. AAAI Press.

[15] D. E. Denning and J. Schlörer. Inference control for statistical databases. *IEEE Computer*, 16(7):69–82, 1983.

[16] EPIC and PI. Privacy & human rights: An international survey of privacy laws and developments. Technical report, Electronic Privacy Information Center (EPIC), Washington, DC, USA, http://www.epic.org, and Privacy International (PI), London, UK, http://www.privacyinternational.org, 2004.

[17] EXOCOM. Privacy technology review. Technical report, Business Strategy and IT Consulting Division, The EXOCOM Group Inc. (for the Office of Health and the Information Highway, Health Canada), 2001.

[18] Z. Ghahramani. An introduction to hidden Markov models and Bayesian networks. *International Journal of Pattern Recognition and Artificial Intelligence*, 15(1):9–42, 2001.

[19] Z. Ghahramani and M. I. Jordan. Factorial hidden Markov models. *Machine Learning*, 29:245–273, 1997.

[20] V. Gowadia, C. Farkas, and M. Valtorta. Paid: A probabilistic agent-based intrusion detection system. *Computers & Security*, 24(7):529–545, October 2005.

[21] J. A. Halderman, B. Waters, and E. W. Felten. Privacy management for portable recording devices. In *Proceedings of Workshop on Privacy in the Electronic Society (WPES'04)*, pages 16–24, Washington, DC, USA, October 28 2004.

[22] S. Hanks, D. Madigan, and J. Gavrin. Probabilistic temporal reasoning with endogenous change. In P. Besnard and S. Hanks, editors, *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence (UAI-1995)*, Montréal, Québec, Canada, August 18-20 1995. Morgan Kaufmann Publishers.

[23] L. J. Hoffman. Computers and privacy: A survey. *Computing Surveys*, 1(2):85–103, June 1969.

[24] S. Jajodia. Database security and privacy. *ACM Computing Surveys*, 28(1):129–131, March 1996.

[25] K. Johansen and S. Lee. Network security: Bayesian network intrusion detection. Technical report, Department of Computer Science, Johns Hopkins University, Baltimore, MD, USA, 2003.

[26] P. Kabiri and A. A. Ghorbani. Research on intrusion detection and response: A survey. *International Journal of Network Security*, 1(2):84–102, 2005.

[27] G. Karjoth and M. Schunter. A privacy policy model for enterprises. In *Proceedings of the 15th IEEE Computer Security Foundations Workshop*, June 24-26 2002.

[28] C. Kruegel, D. Mutz, W. Robertson, and F. Valeur. Bayesian event classification for intrusion detection. In *Proceedings of the 19th Annual Computer Security Applications Conference*, LasVegas, Nevada, USA, December 08-12 2003.

[29] X. Li and Q. Ji. Active affective state detection and user assistance with dynamic Bayesian networks. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 35(1):93–105, 2005.

[30] M. C. Mont, S. Pearson, and P. Bramhall. Towards accountable management of identity and privacy: Sticky policies and enforceable tracing services. In *Proceedings of the 14th International Workshop on Database and Expert Systems Applications (DEXA'03)*, pages 377–382, September 1-5 2003.

[31] A. V. Nefian, L. Liang, X. Pi, and K. Murphy. Dynamic Bayesian networks for audio-visual speech recognition. *EURASIP Journal on Applied Signal Processing*, 11:1–15, 2002.

[32] A. E. Nicholson. Fall diagnosis using dynamic belief networks. In *Proceedings of the 4th Pacific Rim International Conference on Artificial Intelligence (PRICAI-96)*, pages 206–217, 1996.

[33] A. E. Nicholson and J. M. Brady. Dynamic belief networks for discrete monitoring. *IEEE Transactions on Systems, Man, and Cybernetics, special issue on Knowledge-Based Construction of Probabilistic and Decision Models*, 24(11):1593–1610, 1994.

[34] H.-J. Park and S.-B. Cho. Privilege flows modeling for effective intrusion detection based on HMM. In *Proceedings of the 2nd International Workshop on Chance Discovery (CDWS2) in the 7th Pacific Rim International Conference on Artificial Intelligence (PRICAI-02)*, Tokyo, Japan, August 19 2002.

[35] R. G. Parker. Privacy issues: Business impacts and responsibilities. Technical report, CAAA/SAP AG Technology and Accounting Education Seminar Series, Canadian Academic Accounting Association (CAAA), 2005.

[36] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, San Franciso, CA, USA, 1988.

[37] X. Qian, M. Stickel, P. Karp, T. Lunt, and T. Garvey. Detection and elimination of inference channels in multilevel relational database systems. In *Proceedings of the 1993 IEEE Symposium on Security and Privacy*, pages 110–116, Oakland, CA, May 24-26 1993.

[38] L. R. Rabiner. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of IEEE*, 77(2):257–286, February 1989.

[39] A. B. Salem, L. Bouillaut, P. Aknin, and P. Weber. Dynamic Bayesian networks for classification of rail defects. In *Proceedings of the Fourth International Conference on Intelligent Systems Design and Applications (ISDA'04)*, Budapest, Hungary, August 26-28 2004.

[40] M. Schunter and P. Ashley. The platform for enterprise privacy practices. In *Proceedings of the 2002 Information Security Solutions Europe Conference (ISSE'02)*, Paris, France, October 2-4 2002.

[41] M. Shroff. Annual report of privacy commissioner 2003-2004. Technical report, Office of the Privacy Commissioner, New Zealand, 2004.

[42] P. Smyth, D. Heckerman, and M. Jordan. Probabilistic independence networks for hidden Markov probability models. *Neural Computation*, 9(2):227–269, 1997.

[43] J. Staddon. Dynamic inference control. In M. J. Zaki and C. C. Aggarwal, editors, *Proceedings of the 8th ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery (DMKD'03)*, pages 94–100, San Diego, CA, June 13 2003. ACM Press.

[44] R. Sterritt, A. Marshall, C. Shapcott, and S. McClean. Exploring dynamic bayesian belief networks for intelligent fault management systems. In *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, pages 3646–3652, September 2000.

[45] H. S. Venter, M. S. Olivier, and J. H. P. Eloff. PIDS: A privacy intrusion detection system. In S. M. Furnell and P. S. Dowland, editors, *Proceedings of the 4th International Network Conference*, pages 255–262, Plymouth, UK, July 2004.

[46] S. Warren and L. Brandeis. The right to privacy. *Harvard Law Review*, 4:193–220, 1890.

[47] C. Warrender, S. Forrest, and B. Pearlmutter. Detecting intrusions using system calls: Alternative data models. In *Proceedings of 1999 IEEE Symposium on Security & Privacy*, pages 133–145, Berkeley, CA, May 9-12 1999.

[48] W.-K. Wong, G. Cooper, and M. Wagner. Bayesian network anomaly pattern detection for disease outbreaks. In *Proceedings of the 20th International Conference on Machine Learning (ICML-2003)*, Washington DC, USA, 2003.

[49] C. Wright, F. Monrose, and G. M. Masson. HMM profiles for network traffic classification. In *Proceedings of the Workshop on Visualization and Data Mining for Computer Security (VizSEC/DMSEC'04)*, pages 9–15, Washington, DC, USA, October 29 2004.

[50] R. W. Yip and K. N. Levitt. Data level inference detection in database systems. In *Proceedings of the 11th IEEE Computer Security Foundations*, pages 179–189, Rockport, MA, June 9-11 1998.

[51] G. Zweig and S. Russell. Speech recognition with dynamic Bayesian networks. In *Proceedings of the 15th National Conference on Artificial Intelligence (AAAI-1998)*, pages 173–180, Madison, WI, USA, 1998. AAAI Press.