

Defining Privacy for Data Mining

Chris Clifton Murat Kantarcioglu Jaideep Vaidya
Department of Computer Sciences
Purdue University
1398 Computer Science Building
West Lafayette, IN 47907-1398
{clifton, kanmurat, jsvaidya}@cs.purdue.edu

Abstract

Privacy preserving data mining – getting valid data mining results without learning the underlying data values – has been receiving attention in the research community and beyond. It is unclear what privacy preserving means. This paper provides a framework and metrics for discussing the meaning of privacy preserving data mining, as a foundation for further research in this field.

1. Introduction

There has recently been a surge in interest in privacy preserving data mining[2, 1, 17, 14, 22, 18]. Even the popular press has picked up on this trend[10, 16]. However, the concept of what is meant by privacy isn't clear. In this paper we outline some of the concepts that are addressed in this line of research, and provide a roadmap for defining and understanding privacy constraints.

Generally when people talk of privacy, they say “keep information about me from being available to others”. However, their real concern is that their information not be *misused*. The fear is that once information is released, it will be impossible to prevent misuse. Utilizing this distinction – ensuring that a data mining project won't enable *misuse* of personal information – opens opportunities that “complete privacy” would prevent. To do this, we need technical and social solutions that ensure data will not be released.

Another view is *corporate privacy* – the release of information about a collection of data rather than an individual data item. I may not be concerned about someone knowing my birthdate, mother's maiden name, or social security number; but knowing all of them enables identity theft. This collected information problem scales to large, multi-individual collections as well. A technique that guarantees no individual data is revealed may still release information

describing the collection as a whole. Such “corporate information” is generally the goal of data mining, but some results may still lead to concerns (often termed a secrecy, rather than privacy, issue.)

1.1. Secure multiparty computation

As an example of a definition of privacy, and its limitations, let us look at Secure Multiparty Computation (SMC)[23, 11]. The idea of SMC is that the parties involved learn nothing but the results. Informally, imagine we have a trusted third party to which all parties give their input. The trusted party computes the output and returns it to the parties.

SMC enables this *without* the trusted third party. There may be considerable communication between the parties to get the final result, but the parties don't learn anything from this communication. The computation is secure if given just one party's input and output from those runs, we can *simulate* what would be seen by the party. In this case, to simulate means that the distribution of what is actually seen and the distribution of the simulated view over many runs are computationally indistinguishable. We may not be able to exactly simulate every run, but over time we cannot tell the simulation from the real runs.

Since we could simulate the runs from knowing only our input and output, it makes sense to say that we don't learn anything from the run other than the output. This seems like a strong guarantee of privacy, and has been used in privacy preserving data mining work[15, 7, 8].

We must be careful when using Secure Multiparty Computation to define privacy. For example, suppose we use a SMC technique to build a decision tree from databases at two sites[15], classifying people into high and low risk for a sensitive disease. Assume that the non-sensitive data is public, but the sensitive data (needed as training data to build the classifier) cannot be revealed. The SMC computation won't reveal the sensitive data, but the resulting classifier

will enable all parties to estimate the value of the sensitive data. It isn't that the SMC was "broken", but that the result itself violates privacy.

1.2. Limitations on results

How can we constrain the results of data mining? There has been work in this area, addressing specific problems such as hiding specific association rules [3, 20] or limiting the confidence in *any* data mining[5]. While these provide some specific techniques, the means available to constrain results can be quite limited. What is needed is a general way to specify what is and is not allowed.

One possible approach is constraint-based data mining [4]. This line of research is concerned with improving the efficiency of algorithms and understandability of results through providing up-front constraints on what results would be of interest. Would the languages used to describe these constraints also serve to define what results are *acceptable* from a privacy standpoint? While the current methods do not *enforce* that nothing outside the constraints can be learned, they could provide a starting point for further research.

The rest of this paper provides some specific suggestions for methods to specify privacy constraints in ways that still allow data mining. We start with a discussion of individual privacy, and methods to protect it. We then discuss methods for corporate privacy, or constraining what is learned from a collection. We conclude with several orthogonal metrics for defining and measuring privacy.

2. Individual privacy

Typically people think of privacy as protecting individual data. Most legal efforts have been directed to this end. For example, the European Community regulates *personal data*[9]:

'personal data' shall mean any information relating to an identified or identifiable natural person ('data subject'); an identifiable person is one who can be identified, directly or indirectly, in particular by reference to an identification number or to one or more factors specific to his physical, physiological, mental, economic, cultural or social identity;

and specifies that data can be

kept in a form which permits identification of data subjects for no longer than is necessary for the purposes for which the data were collected or for which they are further processed. Member States shall lay down appropriate safeguards

for personal data stored for longer periods for historical, statistical or scientific use.

The key element here is "identifiable": As long as the data cannot be traced to an individual, the regulations do not apply.

The U.S. HIPAA rules[13] are similar – they apply to *protected health information*, defined as individually identifiable health information:

Individually identifiable health information is information that is a subset of health information, including demographic information collected from an individual, and:

1. Is created or received by a health care provider, health plan, employer, or health care clearinghouse; and
2. Relates to the past, present, or future physical or mental health or condition of an individual; the provision of health care to an individual; or the past, present, or future payment for the provision of health care to an individual; and
 - (a) That identifies the individual; or
 - (b) With respect to which there is a reasonable basis to believe the information can be used to identify the individual.

Any information that cannot be traced to a specific individual falls outside the scope of most privacy laws. This provides one solution to privacy and data mining projects: As long as the data used is not individually identifiable, there should be no problems.

Another factor in individual data is how the data is collected and held. The U.S. HIPAA rules assume that data is created and held by a healthcare provider. This gives a *corporate* model – individually identifiable information first "appears" within a collection held by someone other than the individual. An alternative is the "world wide web" model, where individuals provide the data in electronic form themselves. These models lead to different solutions. We first look at some generally applicable solutions, then discuss some that are only relevant in the corporate model.

2.1. Obscuring data

One approach to privacy is to obscure data: making private data available, but with enough noise added that exact values (or approximations sufficient to allow misuse) cannot be determined. One approach, typically used in census data, is to aggregate items. Knowing the average income for a neighborhood is not enough to determine the actual income of a resident of that neighborhood. An alternative

is to add random noise to data values, then mine the distorted data. While this lowers the accuracy of data mining results, research has shown that the loss of accuracy can be small relative to the loss of ability to estimate an individual item. We can reconstruct the original distribution of a collection of obscured numeric values, enabling better construction of decision trees[2, 1]. This would enable data collected from a web survey to be obscured at the source – the correct values would never leave the respondent’s machine – ensuring that exact (misusable) data doesn’t exist. A technique has also been developed for association rules, enabling valid rules to be learned from data where items have been randomly added to, or removed from, individual transactions[18].

Data obscuration is effective both in the web and corporate model. Obscuration can be done by the individual (if the receiver isn’t trusted), or by the holder of data (to reduce concerns about breached security.) However, obscuring data falls into a legal gray area. Rules such as EC 95/46 and HIPAA would probably view individually identifiable data with values obscured as protected, even if the exact values are unknowable. However, obscuration could be as or more effective than aggregation methods used on publicly available census data at protecting actual data values. Demonstrating the effectiveness of data obscuration in comparison with census data could improve public acceptance, and lead to changes in legal standards.

Data obscuration techniques could also be used to ensure that otherwise identifiable data isn’t individually identifiable. Re-identification experiments have shown that data that might be viewed as non-identifiable, such as birth date and postal code, can in combination allow identification of an individual[21]. Obscuring the data could make re-identification impossible, thus meeting both the spirit and letter of privacy laws.

2.2. Perfect privacy

One problem with the above is the tradeoff between privacy and accuracy of the data mining results. Can we do better? Using the concept of Secure Multiparty Computation, the answer is clearly yes – in the “web survey” example, the respondents can engage in a secure multiparty computation to obtain the results, and reveal no information that is not contained in the results. However getting thousands of respondents to participate synchronously in a complex protocol is impractical. While useful in the corporate model, it is not appropriate for the web model. Here we present a solution based on a moderately trusted third party – the party is not trusted with exact data, but trusted only not to collude with the “data receiver”.

There are various means of achieving privacy, both technical and nontechnical. Part of the problem is the need to

create a solution which is feasible in terms of efficiency, security, and without limitations in usability. Technical solutions can be formulated without restrictions in usability, by making suitable assumptions. By a judicious use of non-technical mechanisms, we can realize these assumptions in real life.

Perfect privacy in the SMC sense implies that there is absolutely no release of any meaningful information to any third party. Current e-commerce transactions have a trusted (central) third party with access to all the information. The “trust” is governed by legal contracts enjoining the improper release of information. In some cases, the third party is dispensed with and contracts exist between the interested parties themselves. This is obviously insecure from the technical perspective. Though it has been proven that a SMC solution does exist for any functionality, the computation and/or communication required may be high. Other factors, such as the need for continual online availability of the parties, create further restrictions and problems in real-world settings such as a web-based survey.

However, if we jettison the idea of using *only* the interested parties, we can obtain a middle ground solution that does not require a fully trusted third party. We can instead use a fixed number of *untrusted, noncolluding* parties/sites to do the computation.

Assume the existence of k *untrusted, noncolluding* sites.

- *Untrusted* implies that none of these sites should be able to gain any useful information from any of the inputs of the local sites.
- *Noncolluding* implies that none of these sites should collude with any other sites to obtain information beyond the protocol.

Then, all of the local parties can split their local inputs into k random shares which are then split across the k untrusted sites. Each of these random shares are meaningless information by themselves. However, if any of the parties combined their data, they would gain some meaningful information from the combined data. For this reason, we require that the sites be noncolluding. We believe this assumption is not unrealistic. Each site combines the shares of the data it has received using a secure protocol to get the required data mining result.

The following is a brief description of this approach. Every party is assumed to have a single bit of information x_i , identified by some key i . Each party locally generates a random number r_i and then sends $(i, \bar{x}_i = x_i \oplus r_i)$ to one site and (i, r_i) to the second site. Note that neither site will be able to predict the x_i . Due to the xor \oplus (operation), the input they see is indistinguishable from any uniformly generated random sequence. Given any data mining task(f) defined on $X = [x_1, x_2, \dots, x_n]$, it suffices to evaluate $f(\bar{X} \oplus R) = f(X)$ since $R = [r_1, r_2, \dots, r_n]$ and $\bar{X} \oplus R =$

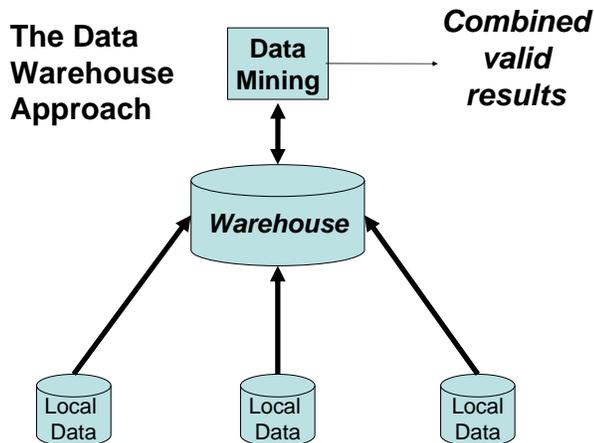


Figure 1. The Data Warehouse approach to mining distributed sources.

$[\bar{x}_1 \oplus r_1, \bar{x}_2 \oplus r_2, \dots, \bar{x}_n \oplus r_n]$. It is a known fact that with the assumption of existence of trapdoor permutations (RSA is assumed to be a trapdoor permutation), any functionality g , ($g : \{0, 1\}^* \times \{0, 1\}^* \mapsto \{0, 1\}^* \times \{0, 1\}^*$) can be evaluated privately in the semi-honest model.[11] Since the initial xor operation can be easily represented as a circuit, given functionality f , we can define a functionality $g(X, R) = f(X \oplus R)$. Thus, any data mining functionality can be evaluated privately without revealing any information other than the final result.

While the architecture is not especially efficient, indeed not even necessarily very practical for large quantities of data, it does demonstrate a method of maintaining perfect privacy while computing the required data mining function.

3. Corporate privacy

In the corporate model, the privacy restrictions need not be just on individually identifiable data. Often it is the body of data that must be protected. Protecting individual data items may not be enough – we may need to protect against learning from the collection. This provides an interesting conflict. Generally the goal of data mining is to learn trends and patterns that hold across a collection of data. How can we reconcile this with a desire to keep overall patterns and trends secret?

Sometimes this involves limiting results, as described in Section 1.2. Another situation where such corporate privacy issues arise is with distributed data sources. Here the problem is one site learning results that hold at another site (except, of course, for results that can be derived from the global results.) Figure 1 illustrates this scenario. The typical approach to data mining of distributed data is to build a

data warehouse containing all the data, then mine the warehouse. This requires that the warehouse be trusted to maintain the privacy of all parties - since it knows the source of data, it learns site-specific information as well as global results. What if there is no such trusted authority? In a sense this is a scaled-up version of the individual privacy problem, however it is an area where the Secure Multiparty Computation approach is more likely to be applicable.

Addressing these issues requires understanding the reasons behind them. We now discuss two issues that lead to privacy concerns in collections of data, and ways to understand those that enable data mining to proceed.

3.1. Secrecy

Individual privacy concerns can lead to corporate privacy concerns. The holder of a collection of individual data may be trusted by those individuals, but if that data is revealed, this trust (often protected contractually) is broken. The collection holder may be willing to participate in a distributed data mining project, but only if it can ensure that its own data items are not revealed. Secure Multiparty Computation would seem to provide a solution to this, however the problem of results revealing private information still remains.

Another issue is protecting the data holder. For example, a medical study may want to use data mining to establish overall trends from hospital data. Even if the techniques used protect patient privacy, they may reveal hospital-specific information. Rules establishing conditions that lead to a high complication rate for a particular operation would be useful study results. However, if these conditions are tied to a particular hospital, there may be liability (or at least public relation) implications. Such implications may limit willingness to participate in such a study. We can develop efficient techniques for data mining that protect such study[14].

In other cases, participants in the data mining project may have specific secrets they wish to withhold, such as trade secrets. An industry consortium may want to mine data to find ways that all members can use to improve their processes. However, learning a particular member's trade secret, and sharing that with the rest of the consortium, is inappropriate. The ability to protect secrets while otherwise participating in a data mining exercise will expand the applicability of data mining.

3.2. Limitations on collaboration

At other times, it may not be the participants that are concerned about sharing data, but external parties. As an example, U.S. antitrust regulations limit the ability of companies to collaborate. The basic premise is protecting the consumer: collaboration that harms the consumer (e.g., price-

fixing cartels) are illegal. The problem is establishing that collaboration *is* to the consumers benefit. If the CEOs of Ford and General Motors meet privately, the public doesn't know if there is illegal collaboration, likely triggering an investigation. Privacy preserving data mining techniques provide a solution to this. If we can prove that no information is shared other than the results, it is easier to justify that the collaboration is for legal purposes. As we discuss in the next section, this leads to a “need to know” metric for determining acceptable information sharing.

4. Metrics

We have discussed some issues imposing privacy constraints on data mining. How do we translate these into solutions that address the issues? The key is an ability to measure privacy. Since privacy has many meanings, we require a set of metrics. Several suggestions are given in this section.

4.1. Bounded knowledge

The data obscuration approach leads to a *bounded knowledge* metric. Here we state that some information about a protected attribute may be revealed, but that the actual value can only be estimated. This may be based on hard bounds (e.g., by adding noise from a random variable uniformly distributed within the bounds), or probabilistic estimates (e.g., by adding noise from a gaussian distribution).

A good measure for quantifying privacy based on such bounded estimation is given by Agrawal and Aggarwal[1]. They propose a measure based on the *differential entropy* of a random variable. Specifically, if we add noise from a random variable A , the privacy is:

$$\Pi(A) = 2^{-\int_{\Omega_A} f_A(a) \log_2 f_A(a) da}$$

where Ω_A is the domain of A .

This metric has several nice features. It is intuitively satisfying in simple cases, e.g., if U is a uniformly distributed random variable between 0 and a . Furthermore, if a sequence of random variables A_i converges to B , then $\Pi(A_i)$ converges to $\Pi(B)$. These properties provides a meaningful way to relate non-uniform distributions to the easily understandable uniform distribution case.

The authors extend this definition to *conditional privacy*, which captures the possibility that the inherent privacy from obscuring data may be reduced by what we can learn from a collection. They show how this can be applied to measure the actual privacy after reconstructing distributions of the original data to improve the accuracy of decision trees build on the obscured data[2, 1]. Similar analyses on other data obscuration techniques would provide an effective way to compare those techniques.

4.2. Need to know

The *Need to know* concept is well established in controlling access to data. In the U.S., access to classified data requires both a security clearance and a justification of why the data should be accessed. The same concept appears in the EC95/46 directive:

Member States shall provide that personal data may be processed only if:

- (a) the data subject has unambiguously given his consent; or
- (b) processing is necessary for the performance of a contract to which the data subject is party or in order to take steps at the request of the data subject prior to entering into a contract; or ...

Note clause (b): It is acceptable to use individual data if it is needed to achieve a result requested by the individual.

This same standard also applies to corporate privacy problems. A need to know standard can be used to decide if collaboration between companies falls afoul of antitrust regulations. For example, airlines in the U.S. make their fares available on shared reservation systems used by travel agents (Sabre, Apollo). This allows easy comparison of prices by both consumers and the airlines – and a quick check will show that the airlines generally offer the same prices on the same routes. This gives the appearance of illegal price fixing, and the airlines were found to engage in such price fixing. Airlines would put out a notice of proposed prices, and if other airlines didn't match the price the notice would be removed.

Operations today are similar – the shared reservations systems exist and prices are the same on the same routes. The only change is that the proposed prices no longer exist. Instead, an airline must actually change its prices, *then* see if other airlines go along. Why is this allowed?

The key is that the current system provides two key benefits to consumers:

1. Price competition among the airlines – they are *allowed* to offer as low a price as they want; and
2. Easy comparison shopping – consumers can check prices, and take the lowest.

These benefits cannot be accomplished without allowing airlines to change their prices. Giving consumers the ability to see and compare prices also enables the airlines to see and compare prices. Sharing currently available prices thus meets the *need to know* standard. The previous “notice of proposed pricing” approach did not, as this information was not shared with or useful to consumers.

Secure Multiparty Computation [23, 11] provides a basis for data mining that meets the need to know standard, but it

is not sufficient. We still need to ensure that the data mining results (or potential results) will meet the standard.

Need to know is interesting in that it is a binary measure. A result is either acceptable (it is required to accomplish the end goal), or unacceptable. It is also difficult in that it can apply to a collection of results rather than a single result; two different result sets may each be sufficient to achieve a goal (either alone meets the need to know standard), but knowing both result sets exceeds need to know. Formal definitions of this metric are probably domain specific, however the legal and social acceptability of this “metric” suggests that further research is warranted.

4.3. Protected from disclosure

Another problem is when we have specific items we want to protect. This may be individual data items, specific rules we don’t want disclosed, or even general classes of knowledge that must be protected. The database security community has developed effective techniques for *inference prevention*[6, 12, 24]. This work is for “provable facts”, or inferences that are always true. Data mining finds inferences that are interesting, but do not always hold. Methods have been proposed to alter data to bring the support or confidence of specific rules below a threshold[3, 20], but choice of an appropriate threshold is still difficult. Alternatively, we may want to protect against association rules that involve a particular outcome (e.g., any rules that pertain to equipment failures), but the problem of defining when a rule is considered strong enough to violate privacy concerns is still a problem.

One option is to use classification as a measure. Many data mining problems can be expressed in terms of classification, e.g., association rules can be used as decision rules. If the ability to classify is limited, many other types of data mining will be limited as well. This idea has been used to evaluate the risk posed by data mining when the knowledge to be protected is not known[5]. However, this assumes the goal is prevention of *any* data mining. Use of classification as a metric to prohibit learning specific facts is ripe for exploration. For example, we could state “it should be impossible to learn a classifier from the data that can predict a person having AIDS with $P(\text{falsehit}) < .9$ ” – any classifier suggesting someone had AIDS would be wrong at least 90% of the time. This would alleviate concerns that even if the individually identifiable data didn’t contain the protected attribute (has AIDS), the data might enable someone to *learn* that attribute.

4.4. Anonymity

The goal during the privacy-preserving data mining process is to achieve results without revealing the identity of the

individual users or any information that may result in identifying different people. As discussed in Section 1.1, current definitions of security such as secure multi-party computation may not be enough. Ensuring that the *results* do not violate privacy requires more.

One method for protecting privacy is *k*-anonymity[19]. Roughly, the goal of *k*-anonymity is to only release data where for all possible queries, at least *k* results will be returned. To achieve this result, generalization and suppression techniques are used. In generalization techniques, some attributes are replaced with more general values so that *k* people will be found with any attribute value. For example, exact ages are replaced by some age ranges. In suppression techniques, data points that may cause too much generalization may be eliminated or a column that has identifying information can be deleted. Although this approach works well for individual data, it is not directly applicable to restricting privacy-preserving data mining *results*.

We propose the following definition of individual privacy that maintains the spirit of *k*-anonymity, but protects against data mining *results* that can be used to predict information about an individual.

Definition 4.1 *Two records ($X_1 \in X, X_2 \in X$) that belong to different individuals are p -indistinguishable if for every function $f : X \mapsto \{0, 1\}$ (X is the domain of the user provided data) that can be evaluated in polynomial-time $|Pr\{f(X_1) = 1\} - Pr\{f(X_2) = 1\}| \leq p$ where $0 < p < 1$*

Definition 4.2 *A data mining process said to be p -individual privacy preserving if at every step of the process, any two individual records are p -indistinguishable.*

Example: Assume that we are using the model described in Section 2.2 to find the count of particular attribute value, e.g., we are looking for the number of people that have a particular cancer type, and that this information is represented as binary attribute. According to architecture described in Section 2.2, each user *i* (*i* is assigned by the Original site) sends $(i, X_i \oplus r_i)$ to the original site and sends (i, r_i) to the Non-colluding Storage Site, where r_i is a random bit and \oplus is the xor operation. Clearly the data gathering process is individual privacy preserving, because what both sites see about individuals is indistinguishable from random data. So the probability of any particular individual having the cancer is the same. Now let us find the count of the individuals who have the cancer. To do this, assume the Original site and the Non-colluding Storage Site add $\frac{d(1-p)}{p}$ (*d* is number of actual data entries) noisy data items to actual data and count the amount of of fake support. The processing site joins the $(i_j, X_i \oplus r_i)$ with (i_j, r_i) . Due to potential re-use of the data actual id’s *i* is replaced by random permutation and fake data is randomly distributed among

the original data. When the original site receives the count, it subtracts the amount of fake support and learns the actual result. Clearly, the original site does not learn anything from this process. After the process, the probability that any given individual has the cancer is the same, even though there may be huge difference between prior and posterior probabilities.

We now show that p -Indistinguishability holds for any two data entries on the processing site. Assume that the processing site compares two randomly ordered data items. Since with probability p any item compared is fake, the probability that any statement is true is less than p . In other words, $Pr\{f(X_{org}) = 1|X_{seen}\} = p * Pr\{f(X_{org}) = 1|X_{seen} \text{ is true}\} + (1 - p) * 0 \leq p$. We can see that $|Pr\{f(X_1) = 1|X_1 \text{ is seen}\} - Pr\{f(X_2) = 1|X_2 \text{ is seen}\}| \leq p$. (In the worst case, one probability will be zero and one will be p .) Since p -Indistinguishability is satisfied for any given data item pairs on the processing site, we can conclude that the data mining process is individual privacy preserving under the semi-honest model (no site colludes to reveal the original values).

5. Conclusions

Privacy preserving data mining has the potential to increase the reach and benefits of data mining technology. However, we must be able to justify that privacy is preserved. For this, we need to be able to communicate what we mean by “privacy preserving”. The current mixture of definitions, with each paper having its own definition of what “privacy” is maintained, will lead to confusion among potential adopters of the technology.

This paper presents some suggestions for defining and measuring privacy preservation. We show how these relate to both privacy policy and practice in the wider community, and to techniques in privacy preserving data mining.

This is by no means the definitive word on the subject. While some measures, such as the differential entropy metric of [1], have clear mathematical foundations and applications, others (such as using classification accuracy as a means of protecting rules from disclosure) have strong potential for further development. Adopting a common framework for discussion of privacy preservation will enable next generation data mining technology to make substantial advances in alleviating privacy concerns.

References

[1] D. Agrawal and C. C. Aggarwal. On the design and quantification of privacy preserving data mining algorithms. In *Proceedings of the Twentieth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 247–

255, Santa Barbara, California, USA, May 21-23 2001. ACM.

[2] R. Agrawal and R. Srikant. Privacy-preserving data mining. In *Proceedings of the 2000 ACM SIGMOD Conference on Management of Data*, pages 439–450, Dallas, TX, May 14-19 2000. ACM.

[3] M. Atallah, E. Bertino, A. Elmagarmid, M. Ibrahim, and V. Verykios. Disclosure limitation of sensitive rules. In *Knowledge and Data Engineering Exchange Workshop (KDEX'99)*, pages 25–32, Chicago, Illinois, Nov. 8 1999.

[4] Special issue on constraints in data mining. *SIGKDD Explorations*, 4(1), June 2002.

[5] C. Clifton. Using sample size to limit exposure to data mining. *Journal of Computer Security*, 8(4):281–307, Nov. 2000.

[6] H. S. Delugach and T. H. Hinke. Wizard: A database inference analysis and detection system. *IEEE Transactions on Knowledge and Data Engineering*, 8(1), Feb. 1996.

[7] W. Du and M. J. Atallah. Privacy-preserving cooperative scientific computations. In *14th IEEE Computer Security Foundations Workshop*, pages 273–282, Nova Scotia, Canada, June 11-13 2001.

[8] W. Du and M. J. Atallah. Privacy-preserving statistical analysis. In *Proceeding of the 17th Annual Computer Security Applications Conference*, New Orleans, Louisiana, USA, December 10-14 2001.

[9] Directive 95/46/ec of the european parliament and of the council of 24 october 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. *Official Journal of the European Communities*, No L(281):31–50, Oct. 24 1995.

[10] A. Eisenberg. With false numbers, data crunchers try to mine the truth. *New York Times*, July 18 2002.

[11] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game - a completeness theorem for protocols with honest majority. In *19th ACM Symposium on the Theory of Computing*, pages 218–229, 1987.

[12] T. H. Hinke, H. S. Delugach, and R. P. Wolf. Protecting databases from inference attacks. *Computers and Security*, 16(8):687–708, 1997.

[13] Standard for privacy of individually identifiable health information. *Federal Register*, 66(40), Feb. 28 2001.

[14] M. Kantarcioglu and C. Clifton. Privacy-preserving distributed mining of association rules on horizontally partitioned data. In *The ACM SIGMOD Workshop on Research Issues on Data Mining and Knowledge Discovery (DMKD'02)*, pages 24–31, June 2 2002.

[15] Y. Lindell and B. Pinkas. Privacy preserving data mining. In *Advances in Cryptology – CRYPTO 2000*, pages 36–54. Springer-Verlag, Aug. 20-24 2000.

[16] S. Lohr. Online industry seizes the initiative on privacy. *The New York Times on the web*, Oct. 11 1999.

[17] K. Muralidhar, R. Sarathy, and R. A. Parsa. An improved security requirement for data perturbation with implications for e-commerce. *Decision Science*, 32(4):683–698, Fall 2001.

[18] S. J. Rizvi and J. R. Haritsa. Maintaining data privacy in association rule mining. In *Proceedings of 28th International Conference on Very Large Data Bases. VLDB*, Aug. 20-23 2002.

- [19] P. Samarati and L. Sweeney. Protecting privacy when disclosing information: k -anonymity and its enforcement through generalization and suppression. In *Proceedings of the IEEE Symposium on Research in Security and Privacy*, May 1998.
- [20] Y. Saygin, V. S. Verykios, and C. Clifton. Using unknowns to prevent discovery of association rules. *SIGMOD Record*, 30(4):45–54, Dec. 2001.
- [21] L. Sweeney. *Computational Disclosure Control: A Primer on Data Privacy Protection*. PhD thesis, Massachusetts Institute of Technology, 2001.
- [22] J. S. Vaidya and C. Clifton. Privacy preserving association rule mining in vertically partitioned data. In *The Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, July 23-26 2002.
- [23] A. C. Yao. How to generate and exchange secrets. In *Proceedings of the 27th IEEE Symposium on Foundations of Computer Science*, pages 162–167. IEEE, 1986.
- [24] R. Yip and K. Levitt. The design and implementation of a data level database inference detection system. In *Proceedings of the Twelfth Annual IFIP WG 11.3 Working Conference on Database Security*, Chalkidiki, Greece, July 15–17 1998.