30 Knowledge Discovery as a Threat to Database Security

Daniel E. O'Leary
University of Southern California

Abstract

This chapter investigates the affect of knowledge discovery from databases on the security of databases. First, it examines the current concern with database systems for security from knowledge discovery. Second, this chapter discusses some of the potential security risks associated with knowledge discovery. Third, some potential structure for the development of controls for such systems is examined. It is suggested that the technology itself and the voluntary or involuntary nature of the unauthorized disclosure form a basis for analysis. Fourth, the quality of database security in general is stressed as an important set of controls to secure against knowledge discovery. Finally, it is noted that any set of controls should be compared with the benefit derived from these controls. Security can be established; however, it is important not to forget that ease for the decision maker is one of the primary reasons for creating and maintaining the database.

30.1 Introduction

Knowledge acquisition is regularly referred to as a bottleneck in the development of expert systems and other knowledge-based systems. When accomplished through the use of a knowledge engineer and expert interaction, substantial time and resources are often required to develop the appropriate knowledge base. Even then, there is always this concern: Does the expert say what s/he does or do what s/he says? Based on articles that have appeared in, for example, the *Wall Street Journal*, it might be that experts do not always make decisions using the information that they say they use. At any rate, the implication from a security perspective is that knowledge possessed by experts is secure to a certain extent. In addition, the expert has at least some control over the flow of knowledge to others.

As a result of increasing efforts to induce knowledge from existing data and, thus, lessen dependence on expressed expert knowledge, there have been improvements in the technology of knowledge acquisition from data. Although such advances might yield a whole new class of data-derived expert systems, they also raise potential security difficulties. The intent and the practice of knowledge discovery as discussed in detail throughout this book suggest that those who have access to a database might be able to extract from it (by discovery) knowledge that the database owners would consider to be proprietary or secret. From the point of view of the database owner, such knowledge could be regarded

as improperly acquired knowledge or unauthorized knowledge. (In the future, courts might deem it to be illegally acquired knowledge.) This gives rise to security concerns that are different from the concerns associated with traditional database security. These security risks and the corresponding controls are the focus of this chapter.

This chapter begins by demonstrating that knowledge acquisition from databases can be a major security problem by examining the consequences of unauthorized knowledge discovery. The Current Focus of Database Security section suggests that issues of knowledge acquisition from databases are not currently addressed. The fourth section, Security of Databases, examines current approaches to the security of database systems, including the use of knowledge discovery by these approaches. Methods and Limitations of Knowledge Discovery assesses some of the limitations of knowledge acquisition from data and then discusses the uses of these limitations as a basis for the development of controls. The next section suggests that the security concerns related to knowledge discovery from data can be investigated in terms of voluntary and involuntary disclosures and that the security controls associated with each might be different. Two subsequent sections treat some approaches to security with voluntary disclosures and examine the problems of involuntary disclosure. Finally, the last section offers a brief summary of the chapter and provides some extensions.

30.2 Consequences of Unauthorized Knowledge Discovery

Unauthorized knowledge discovery from a database can have several negative outcomes for the owner, depending on the contents of the particular database. These consequences derive primarily from an outside agent's ability to determine how the database is used in decision making.

First, given a database of characteristics and resulting decisions, an outsider might find that the owner's decisions are being made improperly or possibly illegally. For example, many applications gather information that is indirectly related to the characteristics of individual job applicants, and jobs can be granted or not granted on the basis of these characteristics. Typically, the name of an applicant contains a box for "Mr.," "Mrs.," or "Ms." Alternatively, names often indicate whether a person is male or female. If such information is captured in a database, then knowledge discovery techniques can be used to determine the relationship of the values to the ultimate disposition of applications. Although less offensive examples exist, this example does demonstrate the point: Knowledge discovery can be the basis for litigation or blackmail.

Second, the relationship between characteristics and decisions can be used by outsiders to develop unauthorized insights into decision making by database users. For example,

auditors and judges are likely to base decisions on certain characteristics when determining whether a financial document should receive further attention or whether a guilty verdict should be issued. If an audit client has access to such a database, insights could be obtained to camouflage activity or change a decision. Auditor clients could then work around the rules established by auditors: If an auditor would audit a document only if it concerned an expenditure of more than $1000, then any illegal transaction would likely be for less than $1000. Further, cases could be constructed by attorneys so that judges would only see the appropriate responses. Thus, unauthorized or illegal knowledge acquisition can destroy established control relationships in various organizational settings.

Third, access to a database of decisions could provide insights that would allow the user to exploit future decision making. For example, transactions by a stockbroker could be anticipated if it were known which variables affected which decisions. Similarly, business strategies could be preempted by competitors who knew which decision variables were used. Thus, unauthorized knowledge acquisition can help to destroy a "level playing field."

These and, possibly, other consequences indicate the importance of considering knowledge discovery from databases as a major security issue. Unfortunately, maybe because of the relatively recent advent of knowledge-acquisition technology, current thinking about databases does not explicitly consider this issue in respect to security.

30.3 The Current Focus of Database Security

Ullman (1982) indicates that "we need to protect against both undesired modification or destruction of data and against unauthorized reading of data" (p. 355). With increasing knowledge discovery capabilities, a third type of protection is also deserving of attention: protection against undesired or unauthorized extraction of knowledge from data.

Protection against unauthorized extraction differs from protection against undesired modification because those interested in unauthorized knowledge acquisition would not likely change or destroy the data because it could lead to discovering the unauthorized knowledge extraction. Further, there is no reason to assume that protection against database changes or destruction will secure the system against all types of unauthorized knowledge acquisition.

Protection against unauthorized extraction also differs from protection against the reading of data. Legitimate users must be able to read the data; for example, they might need to read the data to accomplish their job responsibilities. However, at some point, they might be interested in unauthorized knowledge acquisition. Thus, eliminat-

ing the ability to read data might not be feasible in protecting against unauthorized knowledge acquisition because the threat might be posed by authorized personnel. This possibility suggests that additional security arrangements (controls) should be made to accommodate the need for protection against unauthorized knowledge extraction. This is not to understate the importance of general database security. In fact, as discussed later, general database security is an important control against unauthorized knowledge acquisition from data.

However, because knowledge acquisition is a different issue, it should be examined as its own source of control problems. As such, there are at least two sets of issues to examine. First, the technology associated with knowledge acquisition from data can be the source of limitations that can form the basis of some controls (Methods and Limitations of Knowledge Discovery: The Case of Rule Induction). Second, knowledge acquisition from data can be different from other database problems for a variety of reasons: For example, knowledge can be acquired from either voluntarily or involuntarily released data (Voluntary and Involuntary Disclosures).

30.4 Security of Databases: Alternative Approaches

The nature of the security risks discussed earlier suggests that once the damage is done, it might be too late to correct the problem of unauthorized knowledge acquisition. Although corrective steps such as litigation might be initiated, major efforts should be aimed at preventing and further detecting unauthorized knowledge acquisition.

To some extent, controls aimed at preventing unauthorized access can function as preventative controls against unwanted knowledge acquisition. Thus, some control is afforded by general database security that uses traditional methods, such as passwords and other approaches that are relatively new. In some cases, these new approaches use knowledge discovery to assist in the adaptive securing of the database system.

30.4.1 Views for Multilevel Database Security

Denning et al. (1987) develop a multilevel secure relational database model. This model has several critical aspects. First, they "explicitly allow the specification of derived data in a database schema so that the relationships between stored and derived data can be formally expressed" (p. 129). Second, they distinguish between views that retrieve data (access views) and views that classify data (classification constraints). Third, aggregation constraints serve to define and control the access to aggregates of information.

The design presented in Denning et al. (1987) is a sophisticated database security system to prevent unauthorized database access and change. Unfortunately, in published

versions of this model, it does not appear that the system is designed to learn. Thus, the design does not include a knowledge-acquisition component. Further, the model does not appear to provide explicit concern for knowledge discovery from the database. Instead, the focus is on the protection of information and information about the information.

30.4.2 Intrusion-Detection Models

Intrusion-detection models are generally aimed at determining either when an unauthorized user is attempting to access the system or has entered the system or when an authorized user has exceeded the authorization level or is attempting to exceed the authorization level. Thus, the models are primarily preventative or detective. For example, if a user is trying to break into a system, then s/he would generate an abnormally high rate of password failures; the system would thus be alerted. The objective of of intrusion-detection systems (Denning 1987; Tenor 1988) is to detect a wide range of security violations, from unauthorized abuses by insiders to attempted break-ins by outsiders. These systems typically assume that the system or database is being exploited if use is abnormal.

This approach requires that a system can determine what is normal and abnormal use. Typically, these systems use a set of characteristics to define normal behavior profiles. The system then compares a user's actual behavior with these profiles to determine normality or abnormality. Based on its findings, it then takes some type of corrective action. Typically, these systems use both rule-based and statistical methods in their examination of user and system behavior. In addition, some (Tenor 1988) use a learning component designed to keep the system adaptive.

30.5 Methods and Limitations of Knowledge Discovery: The Case of Rule Induction

As previously noted, one of the potential variables in database security is the actual technology used in knowledge discovery. Limitations in the technology can form the basis for controls aimed to exploit these limitations. One of the primary approaches to knowledge discovery is *rule induction*, or learning from examples. Rule induction is the process of discovering decision-making rules from data. Rule induction from databases can be divided into at least two different categories: statistically oriented approaches (Quinlan 1983) and symbolic manipulation approaches (Mitchell 1977; Michalski and Chilausky 1980).

30.5.1 Limitations of Techniques

One of the primary limitations of any rule-induction system is the degree of nonstationarity in the database: Do the relationships in the data change over time? If the relationships change, then rules extracted from the data should also change.

There are also limitations associated with specific methods of rule induction (Goodman and Smyth 1988). First, statistical techniques cannot easily deal with incremental learning. Typically, given new information, a new decision tree must be designed each time the same data are examined. Second, although symbolic techniques typically include learning from new information as a basic mechanism, they do not handle noise in the data well.

30.5.2 Implications for Database Security

These and other limitations can form the basis for controls that assist in the development of databases that are more secure against knowledge discovery. First, consider the limitation of being nonstationary. Some types of decision-making information are secure if the process is nonstationary, depending on the type of security consequence (Consequences of Unauthorized Knowledge Discovery). Because of the ongoing change in such a system, competitors would rarely be able to obtain enough knowledge to build a system that could anticipate current decision-making rules for competitive advantage.

Although it might assist decision making, removing nonstationarity from data might have a negative affect on the security of this database. As a result, it is important to build smoothing approaches into specific application programs, rather than the databases, thus putting another level of security into the database.

In addition, in some situations, it might be beneficial to take steps to make the database appear to be nonstationary when it isn't. Such an approach might make it seem that the database is inappropriate for rule induction. Implementation of this feature could be placed in the application software used to access the database, which would protect the system when the offenders access the database directly. In other cases, however, the lack of stationarity might not provide a level of security because historical insights might be derived. These, too, can be evidence for, or point to, illegalities or anomalies in management.

Second, that incremental learning is difficult for statistical methods implies that nonstationarity increases the cost of unauthorized knowledge discovery and, thus, discourages such activity. If the database is frequently updated, and it is nonstationary, then it can be costly to update the discovered knowledge using statistical techniques. Alternatively, if the knowledge base is not frequently updated, and the database is nonstationary, then the knowledge discovery will often be outdated. However, if the database is stationary,

then the frequency of updating will not likely have any impact on the data relationships that an intruder is trying to elicit, except that the intruder would not necessarily know whether the relationships are nonstationary. Thus, infrequent updates might discourage knowledge acquisition from data.

Third, as previously noted, symbolic techniques do not handle noise well. Thus, in some situations it might be appropriate to build noise into the database to establish system security, possibly by using false but identifiable records, or performing some form of operation on database elements to disguise the contents. The controls suggested by these limitations can be combined with others so that a collection of approaches is used to secure individual databases.

30.6 Voluntary and Involuntary Disclosures

Another major factor related to securing databases against unauthorized knowledge discovery is the nature of the ultimate disclosure of the data by the user: voluntary or involuntary. Whether the data are disclosed voluntarily or involuntarily can affect the security measures taken to mitigate the risks.

If the data are disclosed on a voluntary basis, perhaps to fulfill regulatory requirements (as is the case with accounting data disclosed to the Securities and Exchange Commission), it can be a signal to some group or subgroup of disclosees of some behavior of a firm or it might be explicit evidence of a firm's actions. In the first instance, it is assumed that the data have some meaning to the user. Typically, the extent and format of the data presentation are dictated. In the last two instances, the firm is dependent on the user to determine meaning from the data. Effectively, the user of the data must obtain some knowledge from it. The building of expert systems or knowledge-based systems is a relatively new type of voluntary data disclosure. Here, data are provided to a group of knowledge engineers or statisticians, and they are expected to find knowledge in the data.

If the data disclosure is involuntary, the disclosure can be accidental, or legitimate users might purposely obtain the data in an illegal manner. In any case, involuntary data disclosure can be used by the acquiring party to obtain insight into decision-making processes. Whether disclosure of data is involuntary or voluntary, if the data are used to build an expert system, the effect of the disclosure is limited to the extent to which technology exists for knowledge discovery from data.

30.6.1 Voluntary Disclosure

The voluntary disclosure of information is done for a number of reasons, which are largely beyond the scope of this chapter. However, a number of issues are of concern as a result of the voluntary nature of disclosure, particularly in those cases where regulatory concerns do not spell out disclosure requirements in detail, and substantial discretion is given to the discloser. First, the level of detail disclosed might provide the user of the data with too many insights and too much knowledge. Second, the voluntarily disclosured information can be assembled to determine undisclosed information. At least one system has been built to assist with these concerns, although it primarily deals with the second issue.

Rule-Based Systems: Edaas Edaas (expert disclosure analysis and avoidance system) is an expert system used by the Environmental Protection Agency (EPA). Edaas (Feinstein and Siems 1985) is designed to advise on the disclosure of confidential business information (CBI).

Chemical manufacturers, importers, and processors submit detailed information to EPA on thousands of chemical substances used in commerce. EPA has instituted security procedures to prevent the direct release of this information. However, if a request for information is not sensitive, then EPA tries to honor the request unless the information can be combined with other nonsensitive information to too closely estimate sensitive data protected under federal nondisclosure law. For example, this indirect disclosure could be used to estimate particular expenditures and, thus, might be used to estimate corporate strategies or research and development plans.

Edaas contains two separate knowledge bases, each represented in a different way. Rules are used to represent the specific law concerning information release. Another knowledge base contains known relationships between pieces of company-related CBI and non-CBI data.

Edaas includes about 60,000 chemicals in the database and has 30 categories of chemical data. For each class of chemical and category of data, there are approximately 11 rules; the system has about 200,000 rules. Clearly, the maintenance associated with such a large system of rules takes on monstrous proportions because an induction approach was not used.

Limitations of Rule-Based Approaches Systems such as Edaas suggest that rules for monitoring the security of databases could be developed. Unfortunately, such an approach might require a different set of rules for each database. Further, a rule-based approach can rapidly become outdated as new knowledge discovery techniques are elicited. Finally, such an approach can require substantial resources to develop and maintain such

a system.

30.6.2 Involuntary Disclosure

As previously noted, one of the primary security concerns with illegal knowledge discovery lies in the involuntary disclosure of information. Control over such disclosure can take either a traditional or a nontraditional approach (for example, intrusion-detection methods).

Traditional Approaches As previously discussed, potentially important in establishing databases that are secure against knowledge discovery are traditional approaches, such as limiting physical or logical access using passwords or other vehicles (Denning et al. 1987) to establish hierarchical access to databases.

Use of Intrusion-Detection Models To date, no intrusion-detection models have been developed to protect databases against knowledge discovery intrusions. However, an intrusion-detection approach could be used to assist in this process.

Certain behavior might suggest that a user is going to try to extract knowledge from a database. For example, if a user is planning to induce knowledge from data, then one of the following approaches might be likely:

First, normal use of any database is unlikely to require examination of an entire record and its hierarchically related contents for a large number of records, yet such a database perusal would be one way to manually generate sufficient data for an induction program.

Second, a user interested in unauthorized induction can dump the database contents to hard copy for later use. In this case, the profiles developed from previous models would be useful. For example, Denning (1987) suggests that the leakage of data by a legitimate user can be detected by noting which users log onto the system at unusual times or route data to remote printers not normally used. In addition, the dumping of a database might be reason enough to arouse suspicion.

Third, a user could dump the data to a file for further processing. Because the database has the information in it, such large-scale dumping would be considered abnormal; thus, the security system could track and monitor such events.

Further, as is currently the case with some systems, not only should users be required to have authorization but also programs. An authorized user could not then use the database in conjunction with an induction program to tease knowledge from the database. Thus, unauthorized program use of a database would not be permitted, or it would be tracked by the system.

30.7 Summary and Extensions

This chapter examined the kinds of risks that might differentiate knowledge acquisition from databases from other security risks. Existing and potential opportunities for the security of database systems from knowledge acquisition were also studied.

Clearly, the sources of controls include general database security controls, primarily those aimed at the prevention and detection of unauthorized use. However, some additional controls can come from the technology of knowledge discovery. Limitations in the technology can be used against the technology.

Additionally, a useful structure for the overall set of controls is based on the voluntary or involuntary nature of the database disclosures. Systems can be developed to focus on the particular needs that derive from the voluntary or involuntary nature of the disclosures. AI-based systems and learning systems can be developed as controls.

In any case, a set of several controls is likely to result in the best security for the system. Unfortunately, some steps that ensure the security of the database might have a negative impact on decision makers. As a result, a cost-benefit analysis is clearly necessary.

30.8 References

Denning, D. 1987. An Intrusion-Detection Model. *IEEE Transactions on Software Engineering* SE-13(2).

Denning, D.; Akl, S.; Heckman, M.; Lunt, T.; Morgenstern, M.; Neumann, P.; and Schell, R. 1987. Views for Multilevel Database Security. *IEEE Transactions on Software Engineering* SE-13(2).

Feinstein, J., and Siems, J. 1985. EDDAS: An Experimental System at the U.S. Environmental Protection Agency for Avoiding Disclosure of Confidential Business Information. *Expert Systems* 2(2): 72–85.

Goodman, R. and Smyth, P. 1988a. Automated Rule Acquisition. Presented at the Third AAAI Knowledge Acquisition Workshop, St. Paul, Minn.

Michalski, R. and Chilausky, R. 1980. Learning from Being Told and Learning from Examples. *International Journal of Policy Analysis and Information Systems* 4: 125–161.

Mitchell, T. 1977. Version Spaces: A Candidate Elimination Approach to Rule Learning. In Proceedings of the Fifth International Joint Conference on Artificial Intelligence, 305–310. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Quinlan, R. 1983. Learning Efficient Classification Procedures and Their Application to Chess Endgames. In *Machine Learning: An Artificial Intelligence Approach*, eds. R. Michalski, J. Carbonell, and T. Mitchell. San Mateo, Calif.: Morgan Kaufman.

Tenor, W. 1988. Expert Systems for Computer Security. *Expert Systems Review* 1(2).

Ullman, J. 1982. *Principles of Database Systems*. Rockville, Md.: Computer Science.