Some Privacy Issues in Knowledge Discovery: The OECD Personal Privacy Guidelines

Daniel E. O'Leary, University of Southern California

SEVERAL COUNTRIES HAVE GENERATED principles to protect individuals from the potential invasion of privacy that data collection and retrieval poses. The Organization for Economic Cooperation and Development has provided probably the best known set of guidelines. A number of countries have adopted these guidelines as statutory law, in whole or in part. The OECD has specific guidelines pertaining to data privacy that directly affect those performing knowledge discovery generally, and those who use so-called "personal data" in particular. In this article I will address such questions as

- What are the implications of the existing privacy guidelines, especially those of the OECD, for knowledge discovery?
- What are the limitations of these guidelines?
- How do the restrictions on knowledge discovery about individuals affect knowledge discovery on groups?
- How do legal systems influence knowledge discovery?

I hope that the answers I pose to these and other related issues will be helpful in generating a larger dialogue on this important subject.

There has been relatively little investigation into the privacy and security issues relevant to knowledge discovery, in particular, and intelligent systems in general. Developers have proposed and used intrusion-detection systems as the basis of security systems designed to protect privacy.[1-3] Typically, intrusion-detection systems determine if a user is an intruder or a legitimate user, generally by way of various internal system profiles. Earlier studies of security issues in intelligent systems included issues of privacy and the security of system knowledge.[4] There has been some concern about knowledge discovery as a different kind of threat to database security as well.[5]

THE OECD HAS SPECIFIC GUIDELINES PERTAINING TO DATA PRIVACY THAT DIRECTLY AFFECT THOSE PERFORMING KNOWLEDGE DISCOVERY GENERALLY, AND THOSE WHO USE SO-CALLED "PERSONAL DATA" IN PARTICULAR.

Risks to privacy and the principles of data protection

The classic definition of the invasion of privacy refers to the "abuse or disclosure of intimate personal data." Recently, there has been an effort to expand this definition to include other issues, such as the protection of general privacy and protection from the unauthorized use of one's "personal" data taken from computer databases.

Increasingly, companies and organizations are using computer-based systems to capture personal data. Although this method typically increases both efficiency and productivity, there are a number of risks to individual privacy. In particular, those risks include the following:[6]

- the data can be used for some purpose other than that for which it was collected;
- the data can be inaccurate, incomplete or irrelevant;
- there are risks of unauthorized access to personal information;
- individual databases can be linked, increasing the range of information about individuals;
- the increased ability to construct individual profiles from multiple sources may affect "decisions concerning the individual's qualifications, credit eligibility,

health, insurance consumption patterns, social security, employment and so on."[6]

As both the amount of information and number of users on the internet grows, these risks become increasingly likely to manifest themselves. This is particularly true for joining previously disparate databases.

Hence, many feel that additional guidelines and statutory-based controls are necessary to prevent the invasion of personal privacy. These concerns have led organizations to generate guidelines to mitigate these privacy risks, including the OECD and the Council of Europe. This article focuses on the OECD guidelines since many nations have adopted them as statutory law.

OECD Principles of Data Collection. The following are the OECD principles of data protection:[6]

(1) *Collection limitation*: Data should be obtained lawfully and fairly, while some very sensitive data should not be held at all.
(2) *Data quality*: Data should be relevant to the stated purposes, accurate, complete and up-to-date; proper precautions should be taken to ensure this accuracy.
(3) *Purpose specification*: The purposes for which data will be used should be identified, and the data should be destroyed if it no longer serves the given purpose.
(4) *Use limitation*: Use of data for purposes other than specified is forbidden, except with the consent of the data subject or by authority of the law.
(5) *Security safeguards*: Agencies should establish procedures to guard against loss, corruption, destruction, or misuse of data.
(6) *Openness*: It must be possible to acquire information about the collection, storage, and use of personal data.
(7) *Individual participation*: The data subject has a right to access and challenge the data related to him or her.
(8) *Accountability*: A data controller should be accountable for complying with measures giving effect to all these principles.

The OECD principles arose to help nations cope with the shipment of data outside the country of origin. They attempt to ensure that when data is transported across country borders the data subjects enjoy the same level of privacy as in the original country.

Thus far 24 countries have adopted the OECD guidelines to varying degrees, including Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Iceland, Ireland, Italy, Japan, Luxembourg, the Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, Turkey, the UK, and the US. Not all countries employ the OECD guidelines as statutory law, and not all countries have adopted all eight guidelines. Instead, the level of participation—the number of guidelines adopted—varies from country to country.

Twelve nations have adopted all eight of the principles in statutory law; Japan adopted seven of the principles (excluding #7) and the UK has adopted six of the principles (excluding #7 and #8), as statutory law. Alternatively, Australia, Canada, New Zealand and the US do not offer protection to personal data handled by private corporations. However, those four countries have similar statutory constraints on personal data held in the public sector.

Scope of application: personal data. The OECD developed the primary protective guidelines for personal data. Consequently, the KDD community must determine what kinds of data fall under the heading of personal. According to Neisingh and de Houwer, personal data is data gathered by corporations and government, including financial, educational, economic, social, political, medical, criminal, welfare, business, and insurance data.[6] As a result, it is easy to see that these principles affect many different kinds of data sets.

"WE HAVE TO BE FORTHRIGHT WITH THE PUBLIC. WE HAVE TO HAVE THEIR CONFIDENCE. WE HAVE TO CONVINCE THEM WE'RE WORKING FOR THE COMMON GOOD. THEN WE CAN INVADE THEIR PRIVACY."

The OECD guidelines and knowledge discovery

In the following section I discuss the impact, implications, and limitations of the guidelines for knowledge discovery.

Collection limitation. This principle states that "some very sensitive data should not be held at all," thereby limiting the scope of knowledge discovery from data. If the data is "very sensitive," knowledge discovery researchers should probably not have access to the data, as it could lead to repercussions. Such sensitive data is likely to include information about religious beliefs, race, national origin, and other issues.

However, it is not clear what it means for data to be sensitive. What may be deemed sensitive depends on the context and country in which the authorities develop the legislation. What is sensitive in one country may not be sensitive in another, suggesting that knowledge discovery could differ from country to country. Accordingly, such cultural differences could form the international differences in computer science practices.

Data quality. Knowledge discovery may influence the data quality principle. For example, knowledge discovery may lead to questions about additional categories of information, such as derived data. The data quality principle suggests that researchers differentiate derived data from original data and not include it in the database, since its accuracy could not be assured. Over time the data on which derived data is based may change, thereby changing the derived data as well. As a result, researchers should not store this data, as it could be outdated. If the derived data is kept, researchers should treat it with the same concerns as the original data.

Also, this principle's recommendation that proper precaution be taken suggests that there be quality standards in knowledge discovery. However, since the discipline is still evolving, it may be premature to talk about generating standards.

Purpose specification limitations. This principle indicates that databases be used only for the declared purposes. Goals for the use of data should be generated, and the data should be used to accomplish those goals exclusively. Any other uses would require the consent of the data subject. Consequently, it is critical that if a database is planned for knowledge discovery, then the use of knowledge discovery is specified.

In addition, if knowledge discovery is only done on databases for which knowledge discovery has been declared, then only those databases generated since the gathering of purpose information began are available for this activity. Accordingly, legacy and existing databases are probably outside the scope of knowledge discovery. Users may have to declare the specific knowledge discovery tasks when gathering the data, instead of declaring anticipated knowledge discovery for some general purpose.

The purpose principle is critical for knowledge discovery using multiple databases. If the data was gathered for use in a single database, analysis across multiple databases generally would violate the purpose principle. This could limit knowledge discovery using individual personal data to particular databases.

IT IS NOT CLEAR WHAT IT MEANS FOR DATA TO BE SENSITIVE. WHAT MAY BE DEEMED SENSITIVE DEPENDS ON THE CONTEXT AND COUNTRY IN WHICH THE AUTHORITIES DEVELOP THE LEGISLATION.

This principle threatens knowledge discovery's potential to expand on its on discoveries as well. Feedback can play a very important role in knowledge discovery tasks. As the system generates more knowledge, that knowledge can form the basis of the search for additional knowledge. Therefore, if the principle limits the knowledge discovery task to the first level findings specified in the original purpose, it limits the power of knowledge discovery significantly.

Another possible limitation is the required level of detail in the statement of purpose. It is possible, in an extreme case, that authorities would require researchers to elicit each specific knowledge discovery activity, not just the fact that knowledge discovery would be done.

Use limitation. This principle is closely related to the purpose specification principle, as it specifies that if data is to be used for some purpose other than the originally specified purpose, the data subject must provide consent. By extension, the data subject will need to provide consent when his or her personal data is to be used for knowledge discovery. The purpose specification principle requires users to identify the original use of the information, and the use limitation principle constrains data use to the original purpose. Both principles require data subject consent if changes in the use of the data occur.

The use limitation principle has a direct impact on performing knowledge discovery from related databases. Generally, expanding the analysis of knowledge discovery from one database to multiple databases would require data subject consent, since the interaction of multiple, previously unconnected databases would suggest alternative uses beyond the original scope.

Acquiring data subject consent may be very difficult, in part because most data subjects would have difficulty understanding the technology of knowledge discovery. Further, it is unclear what level of detail data subjects would need. For example, would the awareness that knowledge discovery was being done be sufficient, or would users need to explain the individual task level?

Security safeguards. This principle calls for establishing safeguards against the misuse of data. In some cases, knowledge discovery may qualify as a misuse of data, especially if unauthorized users perform knowledge discovery, or if knowledge discovery occurs without gathering consent. As a result, authorization procedures for knowledge discovery must be established.

The limitations associated with the statement of purpose also influence the security safeguard principle. A particular concern is how to secure a database from knowledge discovery without eliminating access to virtually all users.

Openness. Taken to one extreme, the openness principle suggests that data subjects should be able to acquire information about the uses of knowledge discovery and the specific knowledge discovered about them. Requiring analysts to inform individuals about particular derived data could limit the general use of knowledge discovery and thereby inhibit its use. If knowledge discovery does not lead to inferences about individual data

subjects, there would not necessarily be an openness issue.

However, since it is virtually impossible to deter users of a database from performing knowledge discovery, it will be equally difficult to know for certain whether knowledge discovery is being done using information about a particular data subject. Thus, the individual participation and accountability principles play a critical role in controlling inappropriate knowledge discovery.

Individual participation. This principle suggests that data subjects should be able to challenge knowledge discoveries related to them. These discoveries might pertain to the individual only or to the individual's relationship to specific groups. The knowledge discovered may directly influence how the users perceive and treat the data subject, possibly adversely affecting that person's available options.

In light of the right to challenge knowledge discoveries, it is critical to document the development of conclusions. Substantiating the quality of different knowledge discovery approaches and algorithms will become increasingly important. The development and use of standards will help mitigate the challenges to knowledge discovery findings.

Accountability. This principle calls for a data controller who is accountable for user compliance with the OECD measures. Thus, a knowledgeable data controller should authorize and be responsible for the adherence of knowledge discovery activity to the OECD measures. In addition, the data controller should inform data subjects of the use and findings from knowledge discovery.

However, due to the decentralization of databases and the difficulty of controlling knowledge discovery activity by those who have access to databases, data controllers will have great difficulty monitoring knowledge discovery effectively. Accordingly, it will be important for the data controller to inform database users and maintenance personnel about the policies regarding knowledge discovery activities, including the consequences of inappropriate use.

Knowledge discovery about groups

This article has thus far focused primarily on privacy issues associated with individual personal data. The OECD guidelines do not refer explicitly to knowledge discovery about groups. As a result, unless the knowledge discovered directly affects the individual personal data, there would be no general application of the guidelines. Instead, alternative legislation or guidelines could be used to guide knowledge discovery about groups. For example, in the US it is illegal to discriminate against certain groups based on sex, race, religion, or national origin. Knowledge discovery about groups, then, could comply with these laws in its use of knowledge pertaining to these categories.

Further, the OECD guidelines suggest that individuals have the right to control the use of data about themselves, even in apparently innocuous knowledge discovery about groups. As a result, individuals could request that they not be included in the generation of knowledge about groups of which they may be a member.

THE OECD GUIDELINES COULD NOT ANTICIPATE OR ADDRESS MANY IMPORTANT ISSUES REGARDING KNOWLEDGE DISCOVERY, AND THUS SEVERAL PRINCIPLES ARE TOO GENERAL OR UNENFORCEABLE.

One drawback of these individual privacy constraints is that they could interfere with important knowledge discoveries. For example, certain diseases seem to strike some groups and not others. As a result, information relating to groups could be the key to the discovery of cures, or other important kinds of knowledge.

Legal systems and other guidelines

The OECD guidelines form one basis of analysis. This article could also be extended to investigate alternative sets of guidelines and statutory laws. The Council of Europe issued a similar set of guidelines for the European community that included the eight OECD principles and some additional constraints relating to so-called transborder dataflows. As alternative legal structures develop, researchers could analyze them for their impact on knowledge discovery.

Legal systems offer bases for the interpretation of different terms and situations in knowledge discovery as well. Many states, for the purposes of protecting litigants from undue invasions of privacy by adverse parties, have statutes defining personal or consumer information. For example, the California Code of Civil Procedure, section 1985.3(1), provides detailed definitions about personal records: they are the "original or any copy of books, documents, or other writings pertaining to a consumer and which are maintained by any 'witness' which is a physician, chiropractor...."

In the specific case of litigation, there are laws regarding the disclosure of information. For example, the California Code of Civil Procedure, section 1985.3, deals with "Subpoena for production of personal records," while 1985.4 summarizes the law regarding "production of consumer records maintained by state or local agency."

Further, in many cases different levels of government regulate certain industries, to a certain extent. Such industries include insurance, law, accounting, and medicine. As a result, these industries are likely to have regulations on limitations of disclosure of information. In other cases, the industries are self-regulated.

THE LIMITATIONS TO THE USE OF knowledge discovery that I've discussed are not limited to the new methods of knowledge discovery developed by the artificial intelligence community. Rather, they apply to all methods used to generate knowledge, including more traditional statistical and database approaches. The OECD guidelines limit the knowledge that can be obtained using any process, including direct examination, classic database updates, queries, or statistical methods.

This article provides some insight into the problems concerning personal privacy and data faced by those who wish to employ knowledge discovery. When it comes to personal data, there often are statutory limitations that vary from country to country. As a result, it suggests that the practice of com-

puter science and artificial intelligence varies from country to country, based on different cultural and legal differences. However, it is clear that there are some general principles of data collection and maintenance that a number of countries adhere to. Those principles influence what data can be used in knowledge discovery and how users process and maintain discovered data.

Many of the above limitations are a result of the OECD personal privacy legislation predating knowledge discovery's widespread use in the artificial intelligence community. The OECD guidelines could not anticipate or address many important issues regarding knowledge discovery, and thus several principles are too general or unenforceable.

Acknowledgments

I acknowledge the comments of the anonymous referees and Lance Hoffman on earlier versions of this article. I also thank B. Chandrasekaran and Gregory Piatetsky-Shapiro for their efforts in developing and coordinating this forum.

References

1. D. Denning, "An Intrusion-Detection Model," *IEEE Trans. on Software Engineering*, Vol. SE 13, No. 2, Feb. 1987, pp. 222–232.
2. W. Tenor, "Expert Systems for Computer Security," *Expert Systems Review*, Vol. 1, No. 2, 1988, pp. 3–6.
3. D. O'Leary, "Intrusion Detection Systems," *J. Information Systems*, Vol. 6, No. 1, 1992, pp. 63–74.
4. D. O'Leary, "Expert System Security," *IEEE Expert*, Vol. 5, No. 3, 1990, pp. 59–70.
5. D. O'Leary, "Knowledge Discovery as a Threat to Database Security," in *Knowledge Discovery in Databases*, G. Piatetsky-Shapiro and W. Frawley, eds., MIT Press, Cambridge, Mass., 1991, pp. 507–516.
6. A. Neisingh, and J. de Houver, *Grensoverschrijdend Gegevenssverker*, Klynveld, Brussels, Belgium, 1987. Translated as *Transborder Data Flows*, KPMG, New York, 1988.

Daniel E. O'Leary is an associate professor at the School of Business of the University of Southern California. Dan received his BS from Bowling Green State University (Ohio), his masters from the University of Michigan, and his PhD from Case Western Reserve University. He has served as the Program and General Chair of the IEEE Conference on Artificial Intelligence Applications and as the chair of the AAAI Workshop on Verification and Validation of KBS. Dan is a member of AAAI, ACM and IEEE. He can be reached at the University of Southern California, 3660 Trousdale Parkway, Los Angeles, CA 90089-1421; oleary@RCF.usc.edu.