

Detecting Privacy and Ethical Sensitivity in Data Mining Results

Peter Fule and John F. Roddick

School of Informatics and Engineering
Flinders University, GPO Box 2100
Adelaide, South Australia
Email:[roddick, peterf]@infoeng.flinders.edu.au

‘Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good; and for this reason the good has rightly been declared to be that at which all things aim. But a certain difference is found among ends...’

Aristotle, Ethica Nicomachea

Abstract

Knowledge discovery allows considerable insight into data. This brings with it the inherent risk that what is inferred may be private or ethically sensitive. The process of generating rules through a mining operation becomes an ethical issue when the results are used in decision making processes that effect people, or when mining customer data unwittingly compromises the privacy of those customers.

Significantly, the sensitivity of a rule may not be apparent to the miner, particularly since the volume and diversity of rules can often be large. However, given the subjective nature of such sensitivity, rather than prohibit the production of ethically and privacy sensitive rules, we present here an alerting process that detects and highlights the sensitivity of the discovered rules. The process caters for differing sensitivities at the attribute value level and allows a variety of sensitivity combination functions to be employed. These functions have been tested empirically and the results of these tests are reported.

Keywords: Data mining, Knowledge Discovery, Privacy, Ethics, Sensitivity, Association Rules.

1 Introduction

Knowledge discovery, in common with many powerful technologies, lends itself both to abuse and to great benefit. Moreover, like many technologies, the ability to harm or to cause offense can often be inadvertent.

The publication of a rule which subsequently has a negative impact on the community bears significant risks, through litigation, adverse publicity, loss of reputation and so on. However, the number and complexity of rules generated from many data mining systems means that the human post-processing of a data mining run can be long and potentially complex, leading to suspect rules being overlooked.

Copyright ©2004, Australian Computer Society, Inc. This paper appeared at Twenty-Seventh Australasian Computer Science Conference (ACSC2004), Dunedin, New Zealand. Conferences in Research and Practice in Information Technology, Vol. 26. Vladimir Estivill-Castro, Ed. Reproduction for academic, not-for profit purposes permitted provided this text is included.

For example, consider the following (fictional) rule:

$$\begin{aligned} &PostCode(5409), Age(18 - 25), Gender(Male) \quad (1) \\ &\rightarrow HepBStatus(Yes) \quad \gamma(20\%) \end{aligned}$$

This would be a worrying rule to discover for any segment of the population. However, consider the increased risk of offense if *Postcode(5409)* referred to, say, an indigenous community or the national parliament. Moreover, if this rule was a just one amongst several thousand, it could be difficult for a human observer to spot the potential problems. Nuances in specific attribute values, such as in the case here, can be easily missed without support for detecting privacy violations and ethically sensitive inferences. Since the mining process is inherently inductive, many rules may be overtly generalised (ie. sociological stereotypical) or may, because of inadequate statistical analysis and rule pruning, appear to be useful but instead be misleading.

This paper describes a process for evaluating a rule in terms of its perceived privacy and ethical sensitivity. As such, these measures provide an additional way in which to quantify a rule’s *interestingness* (Freitas 1999, Hilderman & Hamilton 1999, Hilderman & Hamilton 2001, Sahar 1999, Silberschatz & Tuzhilin 1995, Silberschatz & Tuzhilin 1996). Unlike most other methods which adopt objective statistical measures to determine interestingness, in this work we propose a subjective system for rating a rule’s interestingness.

To properly discuss issues of privacy and ethics in data mining the terms privacy and ethics need to be clearly defined.

Privacy will be referred to as an individual’s desire and ability to keep certain information about themselves hidden from others. Defining privacy in a legal context has historically been a difficult process which still hampers new privacy laws. Moreover, as discussed in (Wahlstrom, Roddick, Sarre, Estivill-Castro & de Vries 2002), complete privacy is not an inherent part of any society as participation in a society necessitates communication and negotiation, which renders absolute privacy unattainable (Gavison 1984).

Ethics will be referred to as a set of moral principles or a system of values which guides the behaviour of individuals and organisations. It is the *correct* way of doing things which as judged by society and often enforced through law (such as anti-discrimination legislation). To act ethically involves acting for the benefit of the community. It is entirely possible to act unethically yet legally.

The problems associated with rules such as that in example (1) affect both parties. For the objects of

such a rule there can be a negative impact through stereo-typing and an invasion of privacy. For a publisher, there is the risk of a loss of reputation and of litigation.

Two approaches can be taken to mitigate the effects of ethical compromise. Firstly, privacy-preservation mechanisms can be put in place that limit access to data, restrict the scope of queries or perturb, hide or delete data so that undesired responses do not occur. Unfortunately, this can also affect the capacity of a mining system to generate beneficial results. The second approach is thus to allow unrestricted mining but to employ an alerting process to inform users to the potential sensitivities of rules, ie. to manage rather than eliminate the risk. A major problem that then needs to be overcome with this approach is that sensitivity is context dependent and thus global measures of sensitivity cannot be adopted. This is the problem tackled by this work.

This paper will discuss, in Section 2, the requirements of privacy protection and the enforcement of ethical principles as they pertain to knowledge discovery activities. The section also canvasses related research. In Section 3 we outline our process for providing context-sensitive alerting while in Section 4 we report on our empirical study relating to their use and the refinement of the associated sensitivity composition function (SCF). In Section 5 we outline some areas for future research and conclude the paper.

2 Discussion and Related Work

2.1 KDD and Ethics

Both inside and outside of the KDD community there is growing concern regarding the (ab)use of sensitive information (Boyens, Gunther & Teltzrow 2002, Cavoukian 1998, Clarke 1997, Clarke 1999, Gehrke 2002, Rachels 1975).

Estivill-Castro *et al.*, for example, cite recent surveys about public opinion surrounding personal privacy which show a raised level of concern about the use of private information (Estivill-Castro, Brankovic & Dowe 1999). There is some justification for this concern – a recent survey in *InfoWeek* (Wilder & Soat 2001) found that over 20% of companies store data on their customers with information about medical profile, a similar amount store customer demographics with salary and credit information, and over 15% store information about their customers' legal history. With this increasing level of storage of personal information there is a greater risk that misleading, erroneous or even defamatory rules might be generated.

To demonstrate the potentially misleading nature of data mining, Leinweber mined United Nations data combined with stock market data (Leinweber 1997). It was found that the best indicator for the S&P 500 Index was the estimated level of butter production in Bangladesh. It would be obvious that this is a statistical coincidence, but as other correlations are more difficult to refute, it is important to consider this difficulty in other situations. The use of more statistically appropriate interestingness measures can help address this problem. Moreover, the ability to judge that a generated rule is sensitive is highly dependent on the knowledge and experience of the domain expert, rather than the data miner. Since knowledge discovery techniques are increasing being applied in areas in which the data miner is unlikely to possess the required domain knowledge this is becoming an important aspect.

The first workshop focussing on privacy and data mining (Clifton & Estivill-Castro 2002) was recently

held in Japan. In common with much research in the area, the papers on the topic of privacy preservation in data mining generally focused on issues surrounding the sharing of data between organisations or on mechanisms to prohibit access to data during sharing. The difference in this work is instead to automate the alerting of users when data mining systems produce potentially sensitive results (as opposed to either screening potentially sensitive data or manually checking for sensitive rules), and to highlight these sensitive rules so that they can be reviewed before use/publication.

In data mining research, particularly in areas such as medical and health research, there are a considerable number of databases that could be considered ethically sensitive¹. Access to these datasets is usually tightly controlled with approval for the use of the data only granted where there is a clear and definable benefit to the research and a strong adherence to agreed research ethics. The problem for data mining researchers is that investigations using knowledge discovery tools are commonly open ended – it is not possible to know what will be found until it is discovered. Moreover, many useful investigations require the use of non-anonymised data (for example, to link episodes of treatment). It is hoped that the use of systems such as that described in this paper will help with relieving concerns about using data mining on ethically sensitive datasets and open them up for further research.

2.2 Related Work

Until recently, privacy protection and ethical alerting has received relatively little interest in mainstream KDD research. However, over the past few years there has been some important work, some which is discussed below. The recent concern over homeland defence, for example, has heightened the awareness for the need to find a balance between protecting the privacy of individuals and detecting terrorist threats. In addition, privacy protection for statistical databases is a related discipline and some of the techniques used here can be applied generally.

2.2.1 Privacy Preservation

In the literature, there are several situations in which *privacy preservation* is required:

- Secure sharing of data between organisations – Being able to share data for mutual benefit without compromising competitiveness (Clifton, Kantarcioglu, Vaidya, Lin & Zhu 2002).
- Confidentialisation of publicly available data – Ensuring that individuals are not identifiable from aggregated data and that inferences regarding individuals are disallowed (eg. from government census data) (Miller 1991).
- Anonymisation of private data – Individuals and organisations mutating or randomising information to preserve privacy.
- Access control – Privacy preservation has long been used in general database work to refer to the unauthorised extraction of data. This meaning has also been applied to data mining.

A number of techniques have been proposed including:

¹Our research programme has a focus on the mining of medical and health data as the field presents most dramatically the ethical dilemmas between the general good and individual privacy (Roddick, Fule & Graco 2003).

- *Authority control and cryptographic techniques* (Pinkas 2002). Such techniques effectively hide data from unauthorised access but do not prohibit inappropriate use by authorised (or naive) users.
- The *Anonymisation* of the data, in which any identifying attributes are removed from the source dataset. A variation on this can be a filter applied to the ruleset to suppress rules containing identifying attributes.
- *Query restriction*, which attempts to detect when statistical compromise might be possible through the combination of queries (Miller 1991, Miller & Seberry 1989).
- *Dynamic Sampling* and reducing the size of the available data set. This can be done by selecting a different set of source tuples for each query.
- Noise addition and data perturbation of individual entries in such a manner as to retain the accuracy of statistical queries. This can be done in two ways:
 - *Noise Addition* in which sets of values are changed such that common statistical and mining operations yield the same result. For example, Agrawal and Srikant explore the feasibility of privacy-preserving data mining by using techniques to perturb sensitive values in data (Agrawal & Srikant 2000). Two techniques are presented:
 - * *Value-class membership* in which values for an attribute are partitioned into a set of disjoint mutually exclusive classes, and
 - * *Value distortion* which returns a perturbed value. The perturbation is commonly a value calculated from either a uniform or Gaussian random distribution.
 - *Data swapping* where attribute values are interchanged in a way that maintains the results of statistical queries (Evfimievski, Srikant, Agrawal & Gehrke 2002).

These techniques, combined with sampling, create a trade off between accuracy and privacy

- *Multiparty Computation*. Clifton *et al.* discuss four methods in which multiple sites can generate rules without compromising each site's data (Clifton *et al.* 2002).

2.2.2 Alerting

For all these techniques, the emphasis is on prohibiting the production or viewing of non privacy-preserving rules.

These measures are not always suited or adequate for data mining techniques. For example, data anonymisation, while it is often regarded as the first and most minimal step toward protecting the privacy of data, often means that the common key is removed making the linking of cognate databases difficult. Moreover, there are often ways in which such mechanisms can be circumvented (for example, through inspecting the results of multiple sessions) and they do little to protect against naive or unethical uses of the data. It should be noted that in some circumstances, such as association rule generation, some aspects of




Figure 1: Hierarchy of Sensitivity for *Diagnosis*

anonymisation may not be necessary since the identifying fields are unique for each object and will not become a part of the ruleset as they cannot meet the support constraints.

The technique proposed here is the pragmatic ability to spot a sensitive rule when it is presented. Once a rule has been marked as sensitive there will be application-specific methods for dealing with it in a way that suits the situation. Instead of a manual inspection of the ruleset, we propose an automated process that rates each rule for its sensitivity in the areas of privacy and ethical compromise. The approach presented in this paper concentrates on the privacy and ethical issues of an organisation mining its own data. As such we aim to provide tools to alert users to the possibility of an ethical issue rather than employ filters and other techniques which prohibit the production of such rules. This is particularly important given our medical domain of application.

3 Process Description

A fundamental problem with determining rule sensitivity is that the sensitivity of particular attributes are subjective. What may be extremely sensitive to one person, culture or situation may be less sensitive in another. In the proposed system we address the problem of the subjective nature of ethics and privacy (as well as some of the concerns about the open ended nature of data mining) by automatically rating all generated results using user-defined sensitivity values.

The system works by storing ethical and privacy sensitivities associated with individual items separately to the data mining results. Sensitivities associated with fields can be created either by someone with expert knowledge of what is socially acceptable or through the gathering of societal perceptions using other means such as surveys. We separate the issues of privacy and ethical sensitivity because a rule can be ethically sensitive without being a privacy concern and vice versa.

The system operates by checking each rule in the result set against any sensitivities that may be associated with the rule's composite items using a sensitivity combination function or SCF. There are thus two important aspects to the system – the holding of item-level sensitivities and the manner in which these are combined to form a sensitivity rating for a rule.

3.1 Sensitivity Values and Sensitivity Hierarchies

We store the set of privacy and ethical sensitivity values for each attribute or attribute value in which we have a special interest. Assigning values to an attribute value level has the advantage of providing a more refined way in which to assign ratings. In our system we arbitrarily used a range 0...10 with 0 indicating no particular sensitivity.

These ratings are arranged in a *hierarchy of interest* (see the example in Figure 1) – a tree based structure which holds the items with their sensitivity rating. When the interest value for a particular item is required, the hierarchy is searched in a bottom-up manner with the most specialised value being used. Since most common items can be classified into some form of hierarchy this approach is widely applicable.

Item-level sensitivities can thus be specified in two ways.

- Firstly, the system can set sensitivities for groups of items as non-leaf nodes in the hierarchy. For example, a rating for

- `Diagnosis.Any` or
- `Diagnosis.Cardiac.Any`

can be specified to indicate that *Diagnosis* or *Cardiac* has a sensitivity regardless of its specific value.

- Secondly, at leaf nodes, we can set sensitivities for specific values within an attribute. For example, a rating for

- `Diagnosis.Respiratory.URTI` or
- `Diagnosis.Cardiac.HIA`

can specify elevated or depressed levels of interest.

In our process, different ratings can be given for the perceived rating of an attribute for privacy and for ethical sensitivity (or, indeed, for any other subjective measure of interest). Note that the process degrades gracefully when not provided with sensitivity values, ie. no special interest is assigned to any rule.

3.2 Sensitivity Combination Function

A Sensitivity Combination Function (SCF) is used to calculate a rule's rating based on each item's privacy and ethical values, their position in either the antecedent or consequent, the number of items in the itemset, and so on. It can readily be seen that the manner in which the SCF functions is central to the item-based ratings being accurately translated into ratings for the resulting rules.

To test some candidate functions we undertook a survey² to determine, first, how people rate individual items in terms of privacy and ethical sensitivity and, second, how people then rate a rule that contains these items. Note that we were not interested in the ratings given to individual items but in how these item-level ratings translated in the minds of our participants to those rules they found sensitive. In particular,

- Is there a simple way in which item-level sensitivities can be mathematically translated to a rating for a rule? If so, is this formula robust?
- Does the position (ie. in the antecedent or consequent) of an ethically sensitive item in the rule affect a rule's rating?
- Does the number of items in the rule affect a rule's rating?
- Are there other structural aspects that should be considered, such as non-leaf values within a hierarchy?

²FUSA SBS Ethics Committee Approval #2654.

The manner in which item level ratings are combined is a complex issue. In our work thirteen candidate functions were considered. We assumed a maximum rating of 10 and a minimum of 0. The SCFs that were tested are :

1. Average : $\sum_i^n \frac{d_i}{n}$
2. Antecedents Average : $\frac{\sum_i^p a_i}{p}$
3. Consequents Average : $\frac{\sum_i^q c_i}{q}$
4. Non-zero Average : $\frac{\sum_i^n d_i}{n - \sum_i^n (d_i=0)}$
5. Weighted Average : $\frac{\sum_i^p a_i + 2 \sum_i^q c_i}{p + (2q)}$
6. Weighted Non-zero Average : $\frac{\sum_i^p a_i + 2 \sum_i^q c_i}{p - \sum_i^p (a_i=0) + (2(q - \sum_i^q (c_i=0)))}$
7. Heavily Weighted Average : $\frac{2}{3} \sum_i^p \frac{a_i}{p} + \frac{1}{3} \sum_i^q \frac{c_i}{q}$
8. Highest Value : $Max(d_i | i \in n)$
9. Weighted Highest Value : $\frac{2}{3} Max(a_i | i \in p) + \frac{1}{3} Max(s_i | i \in q)$
10. Average of Exponentially Shifted Values : $\frac{\sum_i^n d_i^{1.1}}{1.259}$
11. Average Increased by Count of Antecedents < 2 : $Min(10, \frac{10+p}{12} \sum_i^n \frac{d_i}{n})$
12. Average Decreased by Count of Antecedents > 3 : $Min(10, \frac{13}{10+p} \sum_i^n \frac{d_i}{n})$
13. Random : $random\ value \in 0 \dots 10$

Where p is the number of antecedents, q the number of consequents, $n = p + q$. a_i is the item rating for the i^{th} antecedent, c_i is the item rating for the i^{th} consequent and d_i is the item rating for the i^{th} item in the rule.

The weighted SCFs (numbers 2, 3, 5, 6, 7 and 9) were used to test if either the antecedents or consequents are more important for judging sensitivity. The hypothesis was that one might be emphasized more heavily than the other due in a person's judgment of the sensitivity of the rule. SCFs 11 and 12 explored the effect of the rule item count. The hypothesis here was that longer rules would be more sensitive because they are more specific or specialised. The highest value SCFs (7, 8 and 9) explored the hypothesis that emphasizing items with higher rating may be a good model. The random value algorithm was included as a base reference.

For example, a rule $A, B \rightarrow C$ with A, B and C, having sensitivity ratings of 3, 5 and 7 respectively, would yield a rule rating of 7 if the *Highest Value* SCF was used, 4 for *Antecedents Average* and so on. The results of the experiments done using the data collected in the survey are given in Section 4.

3.3 The Process

The process can be implemented to function in one of the following stages of the data mining process:

- As a part of the itemset generation algorithm. This allows, in some cases, the use of the sensitivity rating to prune the itemsets,

- As part of rule generation. Rules can be tested against a maximum sensitivity and pruned accordingly,
- In post processing. Filtering or visualising the sensitivity of rules after they have been created. This allows different users to have either a restricted or unrestricted view of the rules.

Each of these integrates, to differing extents, the process of assessing sensitivity more tightly into the process of generating the rules. It should be noted that at present the only options are either to forbid the use of an attribute before data mining occurs (the option often adopted by ethics committees) or to impose post-rule generation methods such as those described in Section 2.2.

The use of the SCF in the itemset generation is similar to the manner in which the support value for an itemset is calculated. During itemset generation, if the itemset's interestingness rating is above some user-specified threshold the itemset could be culled (as with itemsets with low support). In practice, to be able to function in this role the generating function would need to monotonically increase as the itemset grows and thus is dependent on the SCF function used. In our experiments, the best performing SCF was non-monotonic.

Integrating the function into the rule generation stage ensures that rules that are above the threshold value are never created. A major drawback to this implementation is the limitation of the allowable functions. Integration of the SCF at the rule generation phase operates by accepting a set of itemsets and assigning a rating for each of the rules created from them. Rules can then be filtered or flagged as appropriate. Using the system at this stage gives some flexibility to the generating function used, allowing non-monotonic functions to be used but does not allow sensitivity ratings to influence itemset support thresholds.

Using the system as a post processor involves taking rules already created and rating them using the SCF. Note that interestingness measures might combine statistical thresholds and the supplied subjective ratings. At this stage, however, it is not possible to do this for rules which lay outside the supplied thresholds and which have thus already been culled.

One drawback to using the system at the post processor stage is that the rules have already been generated. As discussed elsewhere (Wahlstrom & Roddick 2001), once information has been revealed it is difficult to hide. However, the system can have a legitimate role as a guide to help data miners assess the sensitivity of their findings.

In our experimentation we found that using the system to generate rules from itemsets provided the best balance between the three options. Rules that are deemed too sensitive can be set to not be revealed which reduces the risk of revealing sensitive information while the confidence thresholds can be influenced by the sensitivity ratings. In addition, the list of itemsets is complete, which reduces the risk of losing important information through premature pruning of itemsets. Finally, the itemsets used for input do not reveal as much information as rules, which gives this method some security over the post processor implementation.

3.4 Visualisation Tools

Once the potentially sensitive rules have been found it is important to appropriately pass these alerts to the user. In common with the findings of many researchers (Ceglar, Roddick & Calder 2003), we have

found that simply adding the sensitivity rating to a text based list of the discovered rules is inadequate. The final stage of the process is thus to provide an easy to use visualisation tool which, *inter alia*, highlights these ratings.

The ISetNav tool (shown in Figure 2) is one such implementation. A full discussion of the tool is tangential to this paper but in common with many rule and itemset visualisers, it uses colour to highlight interesting aspects of an itemset/rule which, in the case of ISetNav, includes sensitivity ratings.

4 Experiments and Results

A survey instrument was developed to determine the appropriate SCF³. The performance of each SCF was evaluated by the extent to which the predicted results from the generation function, using as input the item-level sensitivities provided by each respondent, correlated with the selected rule-level rating given by that respondent in the survey.

The survey was constructed to test a range of issues and topics and thus aimed to provoke the best range of reactions. Rules were carefully constructed to test our hypotheses on rule structure: groups of rules created from the same itemset, rules with only a single item difference and rules with (expected) highly sensitive antecedent and (expected) low sensitivity consequents.

The performance of each algorithm is shown in Figure 3. An analysis of the results showed that the differences between the performance of the different SCFs was statistically significant. Figure 3 displays the three measures used to test the suitability of each SCF. The false negative response shows the extent to which an SCF under-reported the sensitivity of a rule as compared to the observed response. Conversely, false positives measure the over-reporting of the sensitivity. These two measures are important for judging how often an SCF will miss sensitive rules or give false alerts.

In addition to the measures shown in Figure 3, the number of times that each SCF predicted a result that was within ± 1 of the observed result was also recorded. This measured whether the SCF accurately predicts the result, without emphasizing the severity of incorrect results. The results were interesting in that several of the SCF's that performed poorly in the average and error measures performed considerably better in this measure.

As always, it is difficult to simulate human behavior. We tend not to be exact in our responses, especially in areas such as privacy which can impinge on many aspects of our higher reasoning. In fact many respondents expressed difficulties in giving numerical responses to the survey. However, notwithstanding these comments, the results shown indicate that an SCF can be useful.

Despite the complexity of some of the algorithms, analysis of the results indicated that using the *Non-zero Average* algorithm provided the most accurate model of human behavior from the algorithms tested. Figure 3 shows that it provides the best average deviation from the observed results and provides an acceptable false negative rate. The results also show that both *Antecedent Average* and *Consequent Average* performed more poorly than straight *Average* indicating that all items play a role in a user's perception of sensitivity.

³The survey was given to academics of all disciplines, members of industry and lay-people. The demographics of the survey respondents was spread fairly evenly for gender and age. Most respondents were Australian (of varying ethnic backgrounds).

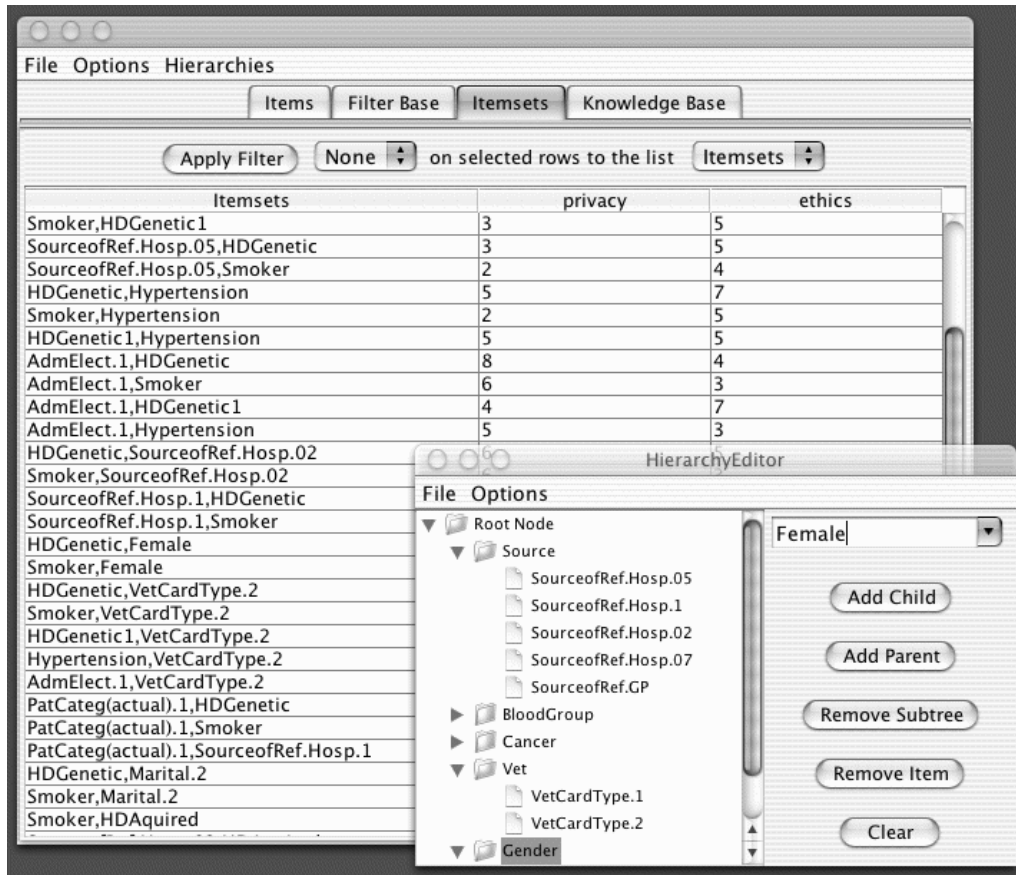


Figure 2: ISetNav Rule Visualisation Software

5 Conclusion

There are currently no systems, that the authors are aware of, that is available to data miners who are concerned about the potential sensitivity of the information that they are extracting from a database. The system described here provides a partial solution by addressing the issues of stereo-typing, privacy protection and the use of ethically sensitive data in an informed knowledge discovery environment. It does this by empowering the users with alerts which can be accepted or dismissed by the user as appropriate. Ie. it allows the user to better manage the risk presented by sensitive rules.

While the paper outlines our empirical testing of the Sensitivity Combination Function, further work in this area could include some refinement of the SCF. It has also been suggested that such alert functions could be used in other contexts in which subjective measures of interestingness are required, such as in a homeland security context by deliberately encoding as sensitive, the characteristics which are of most interest. This, too, is the subject of further research. Further work may also include application of the system in other domains.

If decisions and actions in our society were carried out ethically in all cases there would be little need for individual privacy and most processes would happen in an appropriate manner. However, even if the aim is to abide by accepted ethical principles, any sensitivities embedded in the data might be overlooked resulting in an inadvertent compromise of these principles.

The value of data mining to organisations is considerable but such benefits can be negated if the results of the process are abused. We hope that systems such as the one described here can assist in avoiding at least inadvertent misuse.

6 Acknowledgment

We would like to record our thanks to all those that took part in the survey.

References

- Agrawal, R. & Srikant, R. (2000), Privacy-preserving data mining, *in* W. Chen, J. Naughton & P. A. Bernstein, eds, 'ACM SIGMOD Conference on the Management of Data', ACM Press, Dallas, TX, pp. 439–450.
- Boyens, C., Gunther, O. & Teltzrow, M. (2002), Privacy conflicts in CRM services for online shops: A case study, *in* C. Clifton & V. Estivill-Castro, eds, 'Privacy, Security and Data Mining', Vol. 14 of *Conferences in Research and Practice in Information Technology*, ACS, Maebashi City, Japan, p. 27.
- Cavoukian, A. (1998), 'Data mining: Staking a claim on your privacy'.
- Ceglar, A., Roddick, J. & Calder, P. (2003), Guiding knowledge discovery through interactive data mining, *in* P. C. Pendharkar, ed., 'Managing Data Mining Technologies in Organisations: Techniques and Applications', Idea Group Pub., Hershey, PA, pp. 45–87. Ch. 4.
- Clarke, R. (1997), Privacy and dataveillance, and organisational strategy, *in* 'Region 8 EDPAC'96 Information Systems Audit and Control Assoc. Conf', Perth. Australia.
- Clarke, R. (1999), Person-location and person-tracking: Technologies, risks and policy implications, *in* '21st International Conference on Pri-

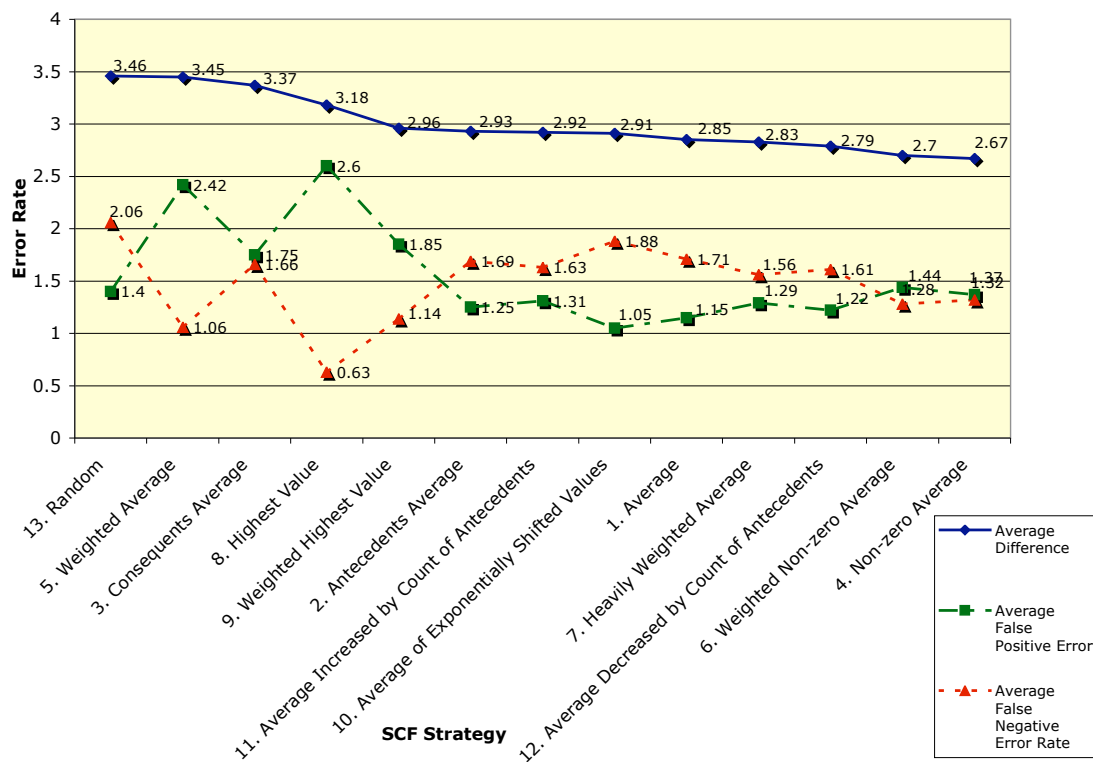


Figure 3: Results of SCF Evaluation

vacy and Personal Data Protection’, pp. 131–150.

Clifton, C. & Estivill-Castro, V., eds (2002), *Privacy, Security and Data Mining - Proc. IEEE International Conference on Data Mining Workshop on Privacy, Security, and Data Mining*, Vol. 14 of *Conferences in Research and Practice in Information Technology*, ACS, Maebashi City, Japan.

Clifton, C., Kantarcioglu, M., Vaidya, J., Lin, X. & Zhu, M. (2002), ‘Tools for privacy preserving data mining’, *SigKDD Explorations* 4(2), 28–34.

Estivill-Castro, V., Brankovic, L. & Dowe, D. (1999), ‘Privacy in data mining’, *Privacy - Law and Policy Reporter* 6(3), 33–35.

Evfimievski, A. (2002), ‘Randomization in privacy-preserving data mining’, *SigKDD Explorations* 4(2), 28–34.

Evfimievski, A., Srikant, R., Agrawal, R. & Gehrke, J. (2002), Privacy preserving mining of association rules, in ‘Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining’, ACM.

Freitas, A. (1999), ‘On rule interestingness measures’, *Knowledge Based Systems* 12(5-6), 309–315.

Gavison, R. (1984), Privacy and the limits of the law, in ‘(?)’, pp. 332–351.

Gehrke, J., ed. (2002), *Special Issue on Privacy and Security*, Vol. 4 of *SigKDD Explorations*, ACM.

Hilderman, R. J. & Hamilton, H. J. (1999), Heuristic measures of interestingness, in J. Zytkow & J. Rauch, eds, ‘3rd European Conference on

Principles and Practice of Knowledge Discovery in Databases (PKDD’99)’, Vol. 1704 of *Lecture Notes in Artificial Intelligence*, Springer, Prague, pp. 232–241.

Hilderman, R. J. & Hamilton, H. J. (2001), Evaluation of interestingness measures for ranking discovered knowledge, in D. W.-L. Cheung, G. J. Williams & Q. Li, eds, ‘5th Pacific-Asia Conference on Knowledge Discovery and Data Mining - PAKDD 2001’, Vol. 2035 of *Lecture Notes in Computer Science*, Springer, Hong Kong, China, pp. 247–259.

Johnson, D. G. & Nissenbaum, H. (1995), *Computers, ethics and social values*, Prentice-Hall, New Jersey.

Leinweber, D. (1997), ‘Stupid data mining tricks: Over-fitting the S&P 500’, *First Quadrant Monograph*.

Miller, M. (1991), A model of statistical database compromise incorporating supplementary knowledge, in B. Srinivasan & J. Zeleznikow, eds, ‘Second Australian Database-Information Systems Conference’, World Scientific, Sydney.

Miller, M. & Seberry, J. (1989), ‘Relative compromise of statistical databases’, *Australian Computer Journal* 21(2), 56–61.

Pinkas, B. (2002), ‘Cryptographic techniques for privacy-preserving data mining’, *SigKDD Explorations* 4(2), 12–19.

Rachels, J. (1975), ‘Why privacy is important’, *Philosophy and Public Affairs* 4(4).

Roddick, J. F., Fule, P. & Graco, W. J. (2003), ‘Exploratory medical knowledge discovery : Experiences and issues’, *SigKDD Explorations* 5(1).

- Sahar, S. (1999), Interestingness via what is not interesting, in S. Chaudhuri & D. Madigan, eds, 'Fifth International Conference on Knowledge Discovery and Data Mining', ACM Press, San Diego, CA, USA, pp. 332–336.
- Silberschatz, A. & Tuzhilin, A. (1995), On subjective measures of interestingness in knowledge discovery, in U. M. Fayyad & R. Uthurusamy, eds, 'First International Conference on Knowledge Discovery and Data Mining (KDD-95)', AAAI Press, Menlo Park, CA, USA, Montreal, Quebec, Canada, pp. 275–281.
- Silberschatz, A. & Tuzhilin, A. (1996), 'What makes patterns interesting in knowledge discovery systems?', *IEEE Transactions on Knowledge and Data Engineering* **8**(6), 970–974.
- Wahlstrom, K. & Roddick, J. F. (2001), On the impact of knowledge discovery and data mining, in J. Weckert, ed., 'Selected Papers from the Second Australian Institute of Computer Ethics Conference', Vol. 1 of *Conferences in Research and Practice in Information Technology*, ACS, Canberra, pp. 22–27.
- Wahlstrom, K., Roddick, J. F., Sarre, R., Estivill-Castro, V. & de Vries, D. (2002), On the ethics of data mining, Technical Report KDM-02-006, Flinders University.
- Wilder, C. & Soat, J. (2001), 'Information week research'.