

This is a draft version. Final version published in B. Custers, T. Calders, B. Schermer & T. Zarsky (eds.), *Discrimination and Privacy in the Information Society. Data Mining and Profiling in Large Databases*, Springer: Heidelberg 2012, p. 273-287.

DRAFT

From Data Minimization to Data Minimummization:

Preserving contextuality in data mining & profiling

Abstract Data mining and profiling offer great opportunities, but also involve risks related to privacy and discrimination. Both problems are often addressed by implementing data minimization principles, which entail restrictions on gathering, processing and using data. Although data minimization can sometimes help to minimize the scale of damage that may take place in relation to privacy and discrimination, for example when a data leak occurs or when data are being misused, it has several disadvantages as well. Firstly, the dataset loses a rather large part of its value when personal and sensitive data are filtered from it. Secondly, by deleting these data, the context in which the data were gathered and had a certain meaning is lost. This chapter will argue that this loss of contextuality, which is inherent to data mining as such but is aggravated by the use of data minimization principles, gives rise to or aggravates already existing privacy and discrimination problems. Thus, an opposite approach is suggested, namely that of data minimummization, which requires a minimum set of data being gathered, stored and clustered when used in practice. This chapter argues that if the data minimummization principle is not realized, this may lead to quite some inconveniences; on the other hand, if the principle is realized, new techniques can be developed that rely on the context of the data, which may provide for innovative solutions. However, this is far from a solved problem and it requires further research.

1. Introduction

Gathering, processing and distributing data, distilling patterns, aggregated profiles and statistical or causal relationships from datasets and applying the gathered rules and profiles in practical decisions all have huge opportunities to offer in relation to both the discovery, the application and the dissemination of knowledge. Data mining and (group) profiling are techniques that have been used since long, but with the emergence of new technical possibilities and processing capacities, these have become the dominant modes of data analyses. Through these techniques, profiles of terrorists are created so as to forestall criminal activities, relationships between specific characteristics and diseases may be discovered so as to prevent them or treat them in an early stage and business profiles are fine tuned to meet consumer interests. However, there are some dangers attached to the use of data mining and profiling. The two major issues regard privacy and discrimination problems.

Privacy might be in danger when personal data of an individual are gathered, used to profile him or used in practical decisions and practices. The discrimination of a particular

person or group may occur when personal characteristics, relating to such information as gender, sexual preferences, political and religious beliefs or ethnicity, are gathered, analyzed and used to bestow upon a person or group a different, disadvantageous treatment. A much used solution in relation to the privacy aspects, but which may also be of use in relation to discriminatory practices, is the implementation of so called privacy enhancing technologies. The technical framework for data processing may be built in such a way that it prevents privacy and discrimination problems, such as by data minimization, which entails a minimum set of sensitive¹ data gathered, stored and used.

Although data minimization sometimes helps to minimize the scale of danger or damage, it has several disadvantages as well. First and most prominently, when valuable data are excluded from the database, it decreases in value and usefulness. Secondly, by deleting these data, the context in which the information was gathered and had a certain meaning is lost. This chapter will argue that from this loss of context, a tendency which is inherent to data mining as such but is aggravated by the use of data minimization principles, problems related to privacy and discrimination arise. Thus, another, opposite approach is suggested, namely that of data *minimummization*. This principle requires a minimum set of data being gathered, stored and clustered. Instead of requiring that certain data is not collected, the principle rule of data minimization, the data *minimummization* principles requires that the context of the data in the form of metadata is collected along with the data. By requiring and clustering a minimum set of (contextual) information, the value of the dataset is retained or even increased, and the privacy and discrimination problems following from the loss of context might be better addressed than by the data minimization principle.

This chapter will proceed as follows. The first section will shortly distinguish four phases of knowledge discovery in databases. The second and third section will point out some general rules relating to privacy and discrimination, with which these may come into conflict. The fourth section will put forward one of the most prominent solutions for these problems, namely that of privacy enhancing technologies and especially the concept of data minimization. The fifth section will analyze some of the problems relating to this technique. The sixth section will offer an alternative solution: data *minimummization*.

2. Data Mining and Profiling Techniques

Data mining is commonly used as an umbrella concept for knowledge discovery in databases, though more correct, it is only one of several phases.² The first step of knowledge discovery in databases is the gathering of data. Gathering information may be done for example through fieldwork, queries, harvesting the internet and personal observations, but also through interconnecting databases and merging them together. Secondly, storing the data and organizing the material. The latter may be necessary not only in relation to making it

¹ In this Chapter, the term 'sensitive data' will refer to both privacy and discriminatory sensitive data, unless where indicated.

² Custers (2004); Skillicorn (2009); Westphal (2009); Larose (2006).

computer readable, but also to enable correct analyses of the data and making them comparable. The third phase is that of actual data mining. Data mining refers to the discovery, most commonly with the use of (mathematical) algorithms, of hidden patterns and subtle relationships in data and the inference of rules that allow for the prediction of future results.³ The patterns and relationships need not to be causal, but may also be statistical. Also, these patterns may be indirect, so that the direct relationship between for example race and solvency is replaced by the relationship between a racially determined zip code and solvency. This is called redlining or masking.⁴ The final stage in the process is applying the knowledge and patterns in real life decisions. This is often done with the assistance of either individual or group profiles.⁵ A pattern obtained through data mining will commonly show the probability that characteristic A is combined with characteristic B. For example, it might be discovered that 67% of the people with curly hair use hair products to style their hairdo or that 86% of the people having a certain zip code possess an expensive car. Thus, targeting such groups most commonly entails a certain margin of error.

3. Data Protection Legislation

Knowledge discovery in databases may among others come into conflict with two legal values: privacy and equality. To provide for some basic fundamentals for assessing the (il)legality of such practices, this section will address the topic of privacy and data protection legislation, the next one will do so with regard to anti-discrimination laws. The main focus will be on European legislation.

Privacy refers to the right to respect for one's private and family life, home and communications, while data protection refers to the right to the protection of personal data concerning a person. The right to privacy is most prominently protected by the European Convention on Human Right and is a moral concept, seen as instrumental in relation to the realisation of autonomy, negative freedom and dignity. If these values are violated or endangered, for example through the use of data mining, then this practice is prohibited unless it is prescribed by law, it is necessary in a democratic society and the infringement is proportional in relation to the goal it serves.

Even more relevant in relation to knowledge discovery in databases is the right to data protection. The European Data Protection Directive, the most important text in this respect, is applicable when personal data⁶ are being processed (entailing both the gathering, processing, use and dissemination of data)⁷ and spells out several obligations for the so called 'data

³ <<http://www.gao.gov/new.items/d07293.pdf>>.

⁴ Squires (2003); Kuhn (1987); LaCour-Little (1999).

⁵ Hildebrandt & Gutwirth (2008).

⁶ Article 2(a) Data Protection Directive 95/46/EC (hereafter: DPD).

⁷ Article 2(b) DPD.

controller', who determines the purpose and means of processing,⁸ in relation to the 'data subject', the one to which the data refer. The directive distinguishes between non-sensitive personal data, with which a person may be identified either directly or indirectly, and sensitive data, relating to information concerning race, ethnicity, political, religious and philosophical beliefs, trade-union membership and data concerning health and sex life with which a person may be either directly or indirectly identified.⁹ The requirements for processing sensitive personal data are stricter than for non-sensitive data.

One of the core doctrines in the directive is that of 'informed consent'. The data controller has certain transparency obligations,¹⁰ correlating with the information rights of the data subject,¹¹ which relate to information regarding the identity of the data controller, the data processed by him and the purposes for which this is done. Furthermore, the directive requires a legitimate purpose for the data processing, the most prominent possibility being the consent of the data subject;¹² subsequently, the data subject has the ability to object to the processing of his data¹³ and to request the erasure or blocking of his personal data.¹⁴ The concept of 'informed consent', relating to the consent or objection to data processing on the basis of adequate and complete information,¹⁵ is instrumental in safeguarding the autonomy of the individual. Besides the doctrine of 'informed consent', two other important principles figure in the directive. Firstly, the so called privacy enhancing principles, regarding the security of processing techniques and data minimization rules, which will be discussed in the fifth section, and secondly, the quality principles, relating to the quality of decision making, the quality of the data themselves and the quality of data processing, which will be elaborated on in the seventh section. Both privacy and data protection problems shall be referred to in this chapter under the umbrella concept 'privacy problems'. First, the general fundamentals of anti-discrimination laws will be outlined in the next section.

4. Anti-Discrimination Legislation

The European legislation regarding discrimination is a bit more scattered. Most importantly, both the Charter of Fundamental Rights and the European Convention on Human Rights contain a general prohibition on the discrimination upon grounds such as gender, race, colour, language, religion, political opinion, nationality, ethnic or social origin,

⁸ Article 2(d) DPD.

⁹ Article 8.1 DPD.

¹⁰ Article 10 DPD.

¹¹ Article 12 DPD.

¹² Article 7 & 8 DPD.

¹³ Article 14 DPD.

¹⁴ Article 12 DPD.

¹⁵ Article 2(h) DPD.

association with a national minority, property, birth genetic features, language, disability, age or sexual orientation. Then, there are also some specific European directives, such as the Employment Equality Directive, prohibiting discrimination on the basis of sexual orientation, religious belief, age and disability in the area of employment, the Racial Equality Directive, among others prohibiting discrimination on the basis of race or ethnicity in the context of employment, the Gender Goods and Services Directive, expanding the scope of sex discrimination regulation to the area of goods and services, and the Gender Social Security Directive, guarantying equal treatment in relation to social security.¹⁶

Generally, these texts make a distinction between direct discrimination and indirect discrimination. The former is usually described as the situation where one person or group is treated less favorably on one of the above mentioned grounds, while the latter is commonly described as the situation where an apparently neutral provision, criterion or practice would put persons of one group at a particular disadvantage compared with persons of the other group.¹⁷ Two exceptions figure repeatedly in the different legal texts. The first is the case of positive discrimination¹⁸ and the second is the case in which the discrimination on the basis of one of the mentioned grounds is objectively justifiable.¹⁹ Positive discrimination involves specific measures taken with a view to ensuring full equality in practice, that aim to prevent or compensate for disadvantages linked to racial or ethnic origin, sex or any other of the above described characteristics. Proportionate differences in individuals' treatment on the basis of sensitive characteristics may be objectively justifiable if such a characteristic constitutes a genuine and determining requirement or factor, provided that the objective is legitimate and the requirement is proportionate.

5. Data Minimization Principles

Knowledge discovery in databases may come into conflict with both privacy and discrimination legislation on several points. These will not be covered extensively, but an example of a privacy violation may be found in the case where personal data are being gathered without a legitimate purpose, where these data are being processed in an 'unsafe' manner, leading to for example data leaks, or where these data are used to undermine the autonomy of the individual. Violations of anti-discrimination laws may for example occur when data regarding gender, religious beliefs, ethnicity and the likes are directly used to bestow on a person or a group a discriminatory treatment or when this is done indirectly, using for example the technique of redlining or masking. Dissemination of such data or knowledge and patterns distilled from them may also lead to a violation of the right to

¹⁶ <http://www.echr.coe.int/NR/rdonlyres/DACA17B3-921E-4C7C-A2EE-3CDB68B0133E/0/182601_FRA_CASE_LAW_HANDBOOK_EN.pdf>.

¹⁷ Article 2(a) & (b) directives **2000/43/EC & 2004/113/EC**.

¹⁸ Article 5 directive **2000/43/EC**. Article 6 directive **2004/113/EC**.

¹⁹ Article 5.2 directive **2004/113/EC**. Article 4 directive **2000/43/EC**.

privacy or to stigmatization of individuals and groups. Especially among privacy scholars, one of the most commonly suggested solutions for such problems is the use of so called privacy enhancing technologies.

(1) Firstly, the Data Protection Directive holds that the controller must implement appropriate technical and organizational measures to protect personal data against accidental or unlawful destruction or accidental loss, alteration, unauthorized disclosure or access, in particular where the processing involves the transmission of data over a network, and against all other unlawful forms of processing.²⁰ Thus, privacy enhancing technologies may be used to minimize the risk of data security breaches by controlling the access to the data, for example through the use of passwords, by encrypting the data and by protecting databases against cyber-attacks. This way, the risk of privacy violations is minimized.

(2) Secondly, both the danger and the scale of the possible damage are minimized through the use of so called data minimization techniques. Concepts such as privacy by design and privacy preserving data mining are closely aligned to this approach. (2a) The Data Protection Directive holds that personal data may only be processed where they are adequate, relevant and not excessive in relation to the specific purpose for which they are collected.²¹ Thus the data controller must specify a specific goal for data processing and the data used should be necessary and proportional in relation to satisfying this objective.

(2b) Another data minimization principle contained in the directive refers to the length of time in which the gathered data may be kept. The directive holds that personal data may be kept in a form which permits identification of data subjects for no longer than is necessary for the specific purpose for which the data were collected.²² For example, there has been some controversy surrounding Google Street View. Google gathers photographs with cars and people on it. It blurs the faces and the license plates before publishing them on the website. This process takes Google up to a year, but the members of the leading advisory organ of the European Union with regard to data protection (the Article 29 Working Party) have asked Google to limit the period it keeps the non-blurred photographs to six months, since they feel that the period Google maintains is excessive.²³

(2c) A final data minimization principle embedded in the directive refers to the way in which the data are kept. The principles of the directive do not apply on data rendered anonymous in such a way that the data subject is no longer identifiable. To determine whether a person is identifiable or not, account should be taken of all the means likely reasonably to be used either by the data controller or by any other person to identify the data subject.²⁴ Thus, anonymous data often refers to data originally able to identify a person, but being stripped of all identifiers, no longer do so. Whether data are able to identify a person must be assessed on a case by case basis. The Article 29 Working Party holds that such assessment '[] should be carried out with particular reference to the extent that the means are likely reasonably to be used for identification []. This is particularly relevant in the case of

²⁰ Article 17 DPD.

²¹ Article 6.1(c) DPD.

²² Article 6.1(d) DPD

²³ <<http://www.edri.org/edriagram/number8.5/article-29-wp-google-street-view>>.

²⁴ Recital 26 DPD.

statistical information, where despite the fact that the information may be presented as aggregated data, the original sample is not sufficiently large and other pieces of information may enable the identification of individuals.’²⁵ This refers among others to techniques used in the data mining process.

The data minimization principles are often referred to in technical literature as well. The abovementioned principles are often caught in the phrase ‘input privacy data mining’. First, a limitation may be posed on the inclusion in databases of information related to privacy or discrimination sensitive data. Second, limitations may be posed on the use of such data for data mining practices, among others through the use of cell suppression and restricting access to statistical queries that may reveal confidential information.²⁶ The main goal of ‘input privacy data mining’ is to minimize the amount of sensitive data, but still allow for an equally valuable data mining process: the so called ‘no-outcome-change’ property.²⁷

Somewhat less well-known and less practiced is the concept of ‘output privacy data mining’.²⁸ This does not refer to the inclusion of data in the database or the use of particular data in data mining processes, but refers to the use of data in the outcome of this process, for example in the rule, pattern or profile distilled from the data.²⁹ The reason for this additional instrument is that ‘input privacy data mining’ is not always sufficient to exclude privacy violations or discriminatory results.³⁰ This may either be caused by masking, indirect discrimination or re-identification, but may also be due to the fact that even although no sensitive data was used in the data mining process, the eventual outcome may still be discriminatory or violate someone’s privacy.³¹ To address outcome based problems, technical solutions may be implemented to prevent particular data from being used in actual practices and decisions.

6. Loss of Contextuality

The principles of data minimization described above help to minimize both the risk and the scale of damage if for example data is misused or a data leak occurs. Also, it may limit the use of particular compromising data in actual practices and decisions. There are however several downsides to using this technique. Firstly, the dataset may lose part of its value through this process. ‘From a data mining perspective the primary issue with informational privacy is that by limiting the use of (particular) personal data, we run the risk

²⁵ Working Party (2007), p. 21.

²⁶ Ruggieri, Pedreschi & Turini (2010); Pedreschi, Ruggieri & Turini (2008); Custers (2004).

²⁷ Bu et al. (2007).

²⁸ Wang & Liu (2008).

²⁹ Verykios et al. (2004).

³⁰ Kantarcioglu, Jin & Clifton (2004).

³¹ Porter (2008).

of reducing the accuracy of the data mining exercise. So while privacy may be protected, the utility of the data mining exercise is reduced'.³² Secondly, knowledge discovery in databases in general and data minimization in particular undermines the context in which data play a role and have a certain meaning, which may create or aggravate (the risk of) privacy violations and discriminatory practices.

Firstly, to retain the value and the meaning of the data, the data itself should be correct and accurate. This may also entail the inclusion of contextual information. However, this principle is often undermined in knowledge discovery in databases, among others since a margin of error is commonly accepted.³³ It also involves a simplification and a decontextualization of reality, since an analysis of few but determining categories is often easier, yields to more direct and concrete correlations and is thus more valuable, than a model which tries to approximate reality's complexity.³⁴ Last but not least, there are costs involved with accurate and complete data gathering, costs which not all parties involved in data mining are willing to bear because a particular threshold in reliability is often sufficient.

Secondly, the data should be updated so that changed facts or changed contexts are incorporated in the database. Typically however, data mining and profiling are used to predict the behavior of people on the bases of old information. Furthermore, when storing the data, one or more of four weaknesses commonly occurs. 'The data may be incomplete, missing fields or records. It may be incorrect, involving non-standard codes, incorrect calculations, duplication, linkage to the wrong individual or other mistaken inputting; the initial information provided may have been incorrect. It may be incomprehensible, involving (for example) bad formatting or the inclusion of multiple fields in one field. It may be inconsistent, involving overlapping codes or code meanings that change over time. Furthermore, even if data is recorded accurately and properly, different databases may use different formatting standards, making data sharing or the "interoperability" of different databases difficult.'³⁵

Thirdly, to retain the value and the meaning of the data, the context of data should be preserved in the process of data analyses and mining. However, harvesting different databases or merging databases together, which is often the case with regard to data mining, may give rise to a problem. '[W]hen data is used in a new context, it may not be interpreted in the same way as previously used, because the new party using the data may not understand how the data was originally classified.'³⁶ By using data for reasons and purposes not envisaged when gathered, data may be taken and judged out of context. For example, the '[] data which circulate on the web were "issued" by people concerned with a precise objective, or in a particular context. The exchanges of data of all kinds and the possibilities to use search engines with any key words engender the risk that we be judged "out of context". [This also refers to] the question of contextual integrity; the person provides his/her data in a

³² Schermer (2011), p. 49.

³³ Ramasastry (2006).

³⁴ Larose (2006), p. 1-2.

³⁵ Renke (2006), p. 791-792.

³⁶ Ramasastry (2006), p. 778.

given context and expects reasonably that it will be processed in this same context, at the risk of it being judged “out of context”.³⁷

Finally, contextuality is important in assessing the value of the outcomes of the data mining process, either in patterns, profiles or concrete decisions. This may be especially important since, as has been said, automatically processed profiles and decisions usually do not evaluate the outcome and result of the data mining process in specific contexts, effecting specific individuals. Again, there is a tendency in knowledge discovery in databases to disregard the context of data.

The tendency in data mining processes to disregard the context of data are aggravated by the use of data minimization techniques³⁸ and cannot be addressed if stuck to this principle, since what is needed is gathering a minimum rather than a minimized amount of data, the data must be updated every now and then, which requires a continued search for data, and the context in which the patterns, profiles and rules acquired by data mining are applied must be evaluated after the process is done.³⁹ Although the principle of data minimization aims at excluding or at least minimizing the risk of privacy and discrimination problems, it may sometimes only aggravate these problems.

For example, if police surveillance mostly takes place in particular neighbourhoods with a lot of immigrants or ethnical minorities, then the gathered data about criminal activities would be heavily tilted towards these groups in society. Incorporation of the methodology of the research in the metadata is thus essential to avoid discrimination and stigmatization towards these minorities.⁴⁰ Furthermore, not keeping data accurate and up to date may lead to privacy and discrimination problems. If a person has decided to quit smoking, but a cigarette company keeps on profiling a consumer as a smoker, this might violate his autonomy and privacy.

Subsequently, the data mining and harvesting process must respect the context of the data. First, disregard of the purpose for which the data were gathered, the purpose limitation principle, may not only lead to a loss of the contextuality of data, but may also undermine the autonomy of the individual as his informed consent with regard to data processing for a specific purpose is transgressed.⁴¹ Secondly, data minimization is not always able to exclude privacy violating or discriminatory results⁴² given the redlining effect.⁴³ Data minimization not only offers no adequate solution in this respect, it might also make it difficult to assess whether a rule is indirectly discriminating or privacy violating.⁴⁴

³⁷ Pouillet & Rouvroy (2008), p. 10 & 14.

³⁸ Guzik (2009); Müller (2009).

³⁹ The only principle that safeguards the contextuality in data mining that is not in tension with data minimization techniques is the purpose limitation principle, which both limits the use of data and ensures that the context of the data is retained.

⁴⁰ Custers (2004).

⁴¹ Taviani (2004).

⁴² Calders & Verwer (2010); Ruggieri, Pedreschi & Turini (2010).

⁴³ Calders & Verwer (2010).

⁴⁴ Pedreschi, Ruggieri & Turini (2008).

Finally, during the stage in which the acquired patterns and profiles are used in practice it is vital to assess the context in which they are applied. Even although rules and profiles are not obtained from analysing sensitive data, they may still have a violating effect in terms of privacy and discrimination. Thus, it would be useful to incorporate background knowledge about the context in which rules and profiles are applied to assess whether such problems or dangers exist.⁴⁵ Again, to avoid privacy or discrimination problems, a larger set of data regarding the context in which rules, patterns and profiles are applied is needed rather than a small or a minimal set.

7. Data Minimization

The loss of contextuality in data mining and profiling leads to privacy and discrimination problems. Implementing the data minimization principle often leads to a further loss of context. A contrary principle might offer a more satisfactory approach. Not minimizing the amount of data gathered, stored and used, but requiring a certain minimum set of (meta)data to be gathered, stored and used when applying the results. In short, the shift from data minimization to data *minimummization*.

There are already several legal provisions that safeguard the correct interpretation of data and their context, among others to be found in the Data Protection Directive. These may provide useful building blocks for the data *minimummization* principle. The existing safeguards can be summarized as the principles of quality, both of the data themselves, the processing of the data and in the use of the data. These may come in tension with the data minimization principles from the same directive, since the principles of quality may often require additional information, not strictly necessary for the satisfaction of the specific purpose for data processing.

Firstly, the Data Protection Directive spells out that the data must be kept accurately and, where necessary, kept up to date; every reasonable step must be taken to ensure that data which are inaccurate or incomplete, having regard to the purposes for which they were collected or for which they are further processed, are erased or rectified.⁴⁶ As data regarding the context of information may be vital for correct interpretations, the first data quality principle may require the collection of such data in the database.

Secondly, the data and the context in which they play a role must be regularly updated, so that a change in facts, their significance and their context will be incorporated in the database. This relates to the second phase in the process of knowledge discovery in databases, as distinguished in section two of this chapter.

Thirdly, the Data Protection Directive spells out that data should be collected for specified, explicit and legitimate purposes and not further processed in a way incompatible with those purposes.⁴⁷ This rule entails two separate duties. The purpose for processing data

⁴⁵ Ruggieri, Pedreschi & Turini (2010).

⁴⁶ Article 6.1(d) DPD. Also see article 12 (b) DPD.

⁴⁷ Article 6.1(b) DPD.

must be explicit and specified. For example, the purpose ‘commercial interest’ will be insufficiently specific. Secondly, further processing, which means the use of data already gathered by the data controller or by a third party for another purpose than the original one, is prohibited when the purpose for processing is incompatible with the original purpose. This provision prohibits the so called function creep of data processing, which signifies the tendency to use already collected data, either by governments or by market parties, for all kinds of purposes and functions not originally intended. The third principle of quality restricts the processing of data to one specified sphere, namely the context of and purpose for which the data were originally gathered.

Finally, the Data Protection Directive contains a restriction on the use of personal data and on making of decisions on the basis of such data. The limitation regards decisions which produce legal effects concerning a person or that significantly affect him, which are based solely on the automated processing of data and which are intended to evaluate certain personal aspects relating to him, such as his performance at work, creditworthiness, reliability, conduct, etc. Such automated decision making, which is quite common in data mining processes, entails the danger of reducing a person to a number and so undermines his individuality and his autonomy. This is partially overcome by granting the data subject the right to knowledge of the logic involved in any automatic processing of data concerning him.⁴⁸ However, this leaves the problem that automatic, computer based analyses and decisions tend to be viewed by humans as absolute and that the data mining process and the outcome thereof only seldom take into account particular contexts and specific individual characteristics.⁴⁹ This risk of contextually detached decision-making is addressed in the directive by granting the individual the right to object to automatic processed decisions, thus granting him the right to be individually judged by another human.⁵⁰

From these existing provisions, a more coherent approach to data minimumization can be developed. Four data minimumization principles can be distinguished, relating to the four stages of knowledge discovery in databases distinguished in the section two.

1. Gathering data: firstly, metadata should be registered and conserved about which data was gathered where and when. This makes it easier to assess for example whether databases are tilted towards criminal activities by minorities due to an over analysis of certain neighbourhoods. Furthermore, the methodology of the process of obtaining the data, among others what data was gathered, by whom and how, should be incorporated in the metadata as well. Finally, the purpose for the gathering of data must be clear.
2. Storing data: the data gathered in the databases should be both accurate and complete. This means for example that relevant contextual data, which are vital for the correct assessment of gathered data, should be incorporated and clustered in the database. This preserves the context of the data in the further course of the data mining process. Attached to this cluster of information should be the metadata described in the previous point. Furthermore, the gathered data must be kept up to date on a regular basis. Finally, decisions on categorisation and organisation of gathered material should be clear and metadata about the

⁴⁸ Article 12 DPD.

⁴⁹ Com(90) final – syn 287 and 288, Brussels, 13 September 1990. Com(92) 422 final – Syn 287, Brussels, 15 October 1992.

⁵⁰ Article 15 DPD.

database itself should be included, for example about who owns it, where it is located, why and when it was build, when the data were included and when they were updated.

3. Analysing data: when analysing data, the previous cluster of data and the metadata about the gathering of information, the database and the organisation and categorization of the material should be preserved. Added should be metadata about the process of analyses, the algorithms used, the databases harvested and the methodology of mining. This may ensure that it can be assessed from hindsight whether patterns, profiles and rules distilled from the data are (indirectly) discriminating or privacy violating. Finally, the context for which the data were gathered, i.e. the purpose limitation, must be respected.
4. Using (aggregated) data: when using the patterns, profiles and rules obtained through data mining, the metadata regarding the gathering of the data, the database, the organisation and categorization of the material and the used analysing techniques as well as the clustered set of data should be accessible. Finally, data must be gathered about in what context the patterns, profiles and rules will be applied and used, so as to assess whether this may lead to privacy violations or discriminatory practices. This may also help to assess whether a discriminatory rule may lead to positive discrimination or is objectively justifiable.

As previously argued, the loss of context may lead to or aggravate privacy and discrimination problems. Inherent to current data mining and profiling practices seems a loss of contextuality, a loss which is not restored, but only aggravated by the data minimization principle. The four data minimummization principles, on the other hand, may be used and implemented to preserve the contextuality of data in data mining and profiling practices. How this should be done is beyond the scope of this chapter.

8. Conclusion

A common definition of autism is context blindness.⁵¹ People suffering from autism treat data, rules and knowledge as isolated facts, as absolute, and thereby disregard the context in which they play a role. Thus, an autistic person may stop at the middle of a zebra-crossing if the traffic light turns red. To him, 'red' signifies 'stop' and nothing else, independent of the given context, while for non-autistic persons, a red traffic light when at the middle of a zebra-crossing signifies 'walk faster', rather than 'stop'. Thus, a set of rules and facts beget a different meaning in different contexts.

Data always signify a certain meaning in a specific context. If this context changes, the information may lose its or beget another meaning. With regard to indexical words such as 'I', 'You', 'Here', 'There', 'This', 'That', 'Now', 'Today', 'Yesterday' and 'Tomorrow', one needs to know where, when and by whom a phrase was uttered to determine the meaning of the phrase. More generally, all data is contextuality determined in time and location, the so called spatio-temporal context. The phrase 'It is cold here' might signify different things in different contexts. If it is uttered after a long trip through the dessert, it might signify a positive feeling, while if it is uttered in a room with an open window, it might signify 'Could

⁵¹ Vermeulen (2009).

you please close the window'. Likewise, the time at which a phrase is uttered is significant.⁵² Furthermore, the context may change over time. The phrase 'A bald man living on Abbey Road 4 in London', may originally signify only person A, but over some time could relate to both person A and B, to person B only or to no one at all. Reference can also be made to so called contextual and conversational implicatures. Suppose just after a job interview, the employee would contact one of the persons on the list of references and were to ask that person whether the applicant would be fit for an university job as researcher and the answer would be 'Well, I can tell you for sure that he makes good coffee'. Since the presumption is that a speaker will provide the maximum relevant information and this information is not relevant at all in this specific context, this would presumably mean 'no'.⁵³ (Again, this changes if uttered when applying for a job in the canteen). Contextuality is essential to understanding and interpreting data and information.

In a way, data mining, profiling and knowledge discovery in data bases give rise to a form of collective autism. Knowledge discovery in databases has the tendency to disregard the contextuality of information. Data are sometimes incorrect, incomplete and out of date, the data set may be tilted towards a certain group of people due to the research methodology, the data may be analyzed and used in a different context and for a different purpose than was originally intended and it's not uncommon that the context in which rules and profiles are put to work in practice are disregarded.

Knowledge discovery in databases may conflict with legal provisions regarding discrimination and privacy. A currently widely propagated solution is that of data minimization, which entails a restriction on the amount of sensitive data gathered, analyzed in the data mining process and used in practical decisions based on the data mining results. The tendency in knowledge discovery in data bases to disregard the context of the data is only aggravated by the data minimization principle.

The loss of contextuality leads to loss of value of the database and the outcome of the data mining process. Moreover, this chapter has argued, the loss of contextuality may give rise to or aggravate already existing privacy and discrimination problems. Thus, sometimes, the data minimization principle may have a counterproductive effect.

Therefore, rather than minimizing the amount of data, this chapter has argued for a minimum amount of data. This replaces the data minimization principle with the data *minimum*mization principle. The latter principle requires a minimum set of data being gathered, stored and clustered when used in practice. First, with regard to the gathering of data, the methodology with, the context in and the reasons for which the data were gathered should be included. With regard to storing data, the data must be correct, accurate and kept up to date; the decisions on categorization and organization of the data should be incorporated. With regard to analyzing data, metadata should be incorporated about the process of analyses, the algorithms used, the databases harvested and the methodology of mining. Finally, with regard to using (aggregated) data, data must be gathered about in what context the patterns, profiles and rules will be applied and used.

By requiring and clustering a minimum set of (contextual) information, the value of the dataset is retained or even increased, and the privacy and discrimination problems

⁵² Grice (1975).

⁵³ This example refers to the maxim of relevance.

following from the loss of context might be better addressed than by the data minimization principle. Nevertheless, it has to be stressed that not all privacy and discrimination problems are caused by a loss of contextuality, nor can all privacy and discrimination problems be solved by the data minimummization principles. Moreover, the data minimummization principles are neither totally new to the technical, nor to the juridical doctrine. Finally, no efforts have been made in this chapter to outline how data minimummization principles may be put into practice or be implemented in data mining rules.

9. Literature

- S. Bu, et al. (2007). Preservation of Patterns and Input-Output Privacy. *Proceedings of ICDE'2007*, 696-705.
- T. Calders, & S. Verwer (2010). 'Three Naive Bayes Approaches for Discrimination-Free Classification'. *Data Mining and Knowledge Discovery* 21(2), 277-292.
- B.H.M. Custers (2004). *The Power of Knowledge; Ethical, Legal, and Technological Aspects of Data Mining and Group Profiling in Epidemiology*. Tilburg: Wolf Legal Publishers
- A. Evfimievski, R. Srikant, R. Agrawal & J. Gehrke (2002). Privacy preserving mining of association rules. *Proceedings of the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'02)*, 217–228.
- J.S. Fulda (2000). Data Mining and Privacy. *Alb. L.J. Sci. & Tech.* (11), 105-113.
- H.P. Grice (1975). Logic and conversation. In: P. Cole & J. Morgan (eds.) *Syntax and semantics* (3). New York: Academic Press, 41-58.
- K. Guzik (2009). Discrimination by Design: Data Mining in the United States's 'War on Terrorism'. *Surveillance & Society* (7), 1-17.
- M. Hildebrandt & S. Gutwirth (eds.) (2008). *Profiling the European Citizen Cross-Disciplinary Perspectives*. New York: Springer.
- M. Kantarcioglu, J. Jin & C. Clifton (2004). When do data mining results violate privacy? In: *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data mining (KDD'04)*. ACM, New York, 599–604.
- P. Kuhn (1987). Sex discrimination in labor markets: The role of statistical evidence. *The American Economic Review* (77), 567-583.

- M. LaCour-Little (1999). Discrimination in mortgage lending: A critical review of the literature. *Journal of Real Estate Literature* (7), 15-50.
- D.T. Larose (2006). *Data mining methods and models*. New Jersey: John Wiley & Sons, Inc. All.
- V.C. Müller (2009). Would you mind being watched by machines? Privacy concerns in data mining. *AI & Soc* (23), 529–544.
- D. Pedreschi, S. Ruggieri & F. Turini (2008). Discrimination-aware Data Mining. *KDD*, 560-568.
- C.C. Porter (2008). De-Identified Data and Third Party Data Mining: The Risk of Re-Identification of Personal Information. *Shidler i.L. Com. & Tech.* (30), article no. 3.
- Y. Pouillet & A. Rouvroy (2008). General introductory report. <http://portal.unesco.org/ci/en/files/27268/12145631033Intro_gen_rapporteur_Y-Pouillet_en.pdf/Intro_gen_rapporteur_Y-Pouillet_en.pdf>.
- A. Ramasastry (2006). Lost in translation? Data mining, national security and the “adverse inference” problem. *Santa Clara Computer & High Tech. L.J.* (22), 757-796.
- W.N. Renke (2006). Who controls the past now controls the future: counter-terrorism, data mining and privacy. *Alta. L. Rev.* (43), 779-823.
- S. Ruggieri, D. Pedreschi & F., Turini (2010). Data Mining for Discrimination Discovery. *Transactions on Knowledge Discovery from Data* 4(2), 9:1-9:40.
- B.W. Schermer (2011). The limits of privacy in automated profiling and data mining. *Computer law & security review* 2 (7), 45-52.
- D. Skillicorn (2009). *Knowledge Discovery for Counterterrorism and Law Enforcement*. Boca Raton: Taylor & Francis Group, LLC.
- G.D. Squires (2003). Racial profiling, insurance style: Insurance redlining and the uneven development of metropolitan areas. *Journal of Urban Affairs* 25(4), 391-410.
- H.T. Tavani (2004). Genomic research and data-mining technology: Implications for personal privacy and informed consent. *Ethics and Information Technology* (6), 15–28.
- P. Vermeulen (2009). *Autisme als Context Blindheid*. Berchem: EPO.
- V.S. Verykios, et al. (2004). State-of-the-art in Privacy Preserving Data Mining. *Sigmod Record* 33 (1), 50-57.

T. Wang & L. Liu (2011). Output Privacy in Data Mining. *Transactions on Database Systems* 36 (1), 1-37.

C. Westphal (2009). *Data mining for Intelligence, Fraud & Criminal Detection*. Boca Raton: Taylor & Francis Group, LLC.

Working Party (2007). *Opinion 4/2007 on the concept of personal data*. WP 136: 01248/07/EN.

T.Z. Zarsky (2003). Mini your own business!: making the case for the implications of the data mining of personal information in the forum of public opinion. *Yale Journal of Law & Technology* (5), 1-56.

DRAFT