Chapter 1
Data Dilemmas in the Information Society: Introduction and Overview

Bart Custers

Abstract. This chapter provides and introduction to this book and an overview of all chapters. First, it is pointed out what this book is about: discrimination and privacy issues of data mining and profiling and solutions (both technological and non-technological) for these issues. A large part of this book is based on research results of a project on how and to what extent legal and ethical rules can be integrated in data mining algorithms to prevent discrimination. Since this is an introductory chapter, it is explained what data mining and profiling are and why we need these tools in an information society. Despite this unmistakable need, however, data mining and profiling may also have undesirable effects, particularly discriminatory effects and privacy infringements. This creates dilemmas on how to deal with data mining and profiling. Regulation may take place using laws, norms, market forces and code (i.e., constraints in the architecture of technologies). This chapter concludes with an overview of the structure of this book, containing chapters on the opportunities of data mining and profiling, possible discrimination and privacy issues, practical applications and solutions in code, law, norms and the market.

1.1 The Information Society

Vast amounts of data are nowadays collected, stored and processed. These data are used for making a variety of administrative and governmental decisions. This may considerably improve the speed, effectiveness and quality of decisions. However, at the same time, it is common knowledge that most databases contain errors. Data may not be collected properly, data may be corrupted or missing, and data may be biased or contain noise. In addition, the process of analyzing the data might include biases and flaws of its own. This may lead to discrimination. For instance,

Bart Custers
eLaw, Institute for Law in the Information Society, Leiden University, The Netherlands
e-mail: bartcusters@planet.nl

B. Custers et al. (Eds.): Discrimination & Privacy in the Information Society, SAPERE 3, pp. 3–26.
springerlink.com

when police surveillance takes place only in minority neighborhoods, their databases would be heavily tilted towards such minorities. Thus, when searching for criminals in the database, they will only find minority criminals.

As databases contain large amounts of data, they are increasingly analyzed in automated ways. Among others, data mining technology is applied to statistically determine patterns and trends in large sets of data. The patterns and trends, however, may easily be abused, as they often lead to unwanted or unjustified selection. This may result in the discrimination of particular groups.

Furthermore, processing huge amounts of data, often personal data, may cause situations in which data controllers know many of the characteristics, behavior and whereabouts of people. Sometimes to the extent of knowing (often based on statistics) more about individuals than these individuals know about themselves. Examples of such factors are life expectancies, credit default risks and probabilities of involvement in car accidents. Ascribing characteristics to individuals or groups of people based on statistics may create a digital world in which every person has several digital identities.[1] Whether these digital identities derived from data processing are a correct and sufficiently complete representation of natural persons or not, they definitely shed different light on our views of privacy. This book addresses the issues arising as a result of these practices.

In this chapter I will provide an introduction to this book and an overview of the chapters that will follow. In this first section I will briefly introduce the premise of this book and what triggered us to write it. Next, in Section 1.2, I will explain briefly what data mining and profiling are and why we need these tools in an information society. This is not a technical section: a more detailed overview of data mining techniques can be found in Chapter 2. In Section 1.3, I will explain why this book focuses on discrimination and privacy issues. In this section, I will also point out that this book is not only about identifying and describing possible problems that data mining and profiling tools may yield, but also about providing both technical and non-technical solutions. This will become clear in Section 1.4, where I sketch the structure of this book.

1.1.1 What This Book Is About

This book will deal with the ways in which new technologies, particularly data mining, profiling and other technologies that collect and process data, may prevent or result in discriminatory effects and privacy infringements. Focus of the book will also be on the question how and to what extent legal and ethical rules can be integrated into technologies, such as data mining algorithms, to prevent such abuse. Developing (legally and ethically) compliant technologies is increasingly important because principles such as “need to know” and “select before you collect” seem difficult to implement and enforce. Such principles focusing on access controls are increasingly inadequate in a world of automated and interlinked databases and information networks, in which individuals are rapidly losing grip on who is using

[1] Solove, D. (2004).

their information and for what purposes, particularly due to the ease of copying and disseminating information. A more integrated approach, not merely focusing on the collection of data, but also on the use of data (for instance using concepts like transparency and accountability) may be preferable.

Because of the speed with which many of the technological developments take place, particularly in the field of data mining and profiling, it may sometimes be difficult for people without a technological background to understand how these technologies work and what impact the may have. This book tries to explain the latest technological developments with regard to data mining and profiling in a manner which is accessible to a broad realm of researchers. Therefore, this book may be of interest to scientists in non-technical disciplines, such as law, ethics, sociology, politics and public administration. In addition, this book may be of interest to many other professionals who may be confronted with large amounts of information as part of their work.

1.1.2 Responsible Innovation

In 2009 the Netherlands Organization for Scientific Research (NWO) commenced a new research program on responsible innovation.[2] This program (that is still running) focuses on issues concerning technological developments that will have a dramatic impact (either positive or negative) on people and/or society. The program contributes to responsible innovation by increasing the scope and depth of research into societal and ethical aspects of science and technology.

A key element of the program is the interaction between research of technological sciences (such as computer science, mathematics, physics and chemistry) and non-technological sciences (such as law, ethics and sociology), to generate cooperation between these disciplines from the early stages of developing new technologies. When it comes to legal, ethical and social effects of new technologies, parties involved are sometimes tempted to shun specific responsibilities.[3] It is often the case that engineers and technicians assert that they only build a particular technology that others can use for better or for worse. The end users, however, often state from their perspective that they only use technologies for the purposes for which they were intended or designed. A value-sensitive design approach may contribute to incorporating legal, ethical and social aspects in the early stages of developing new technologies.[4]

Another key element of the program is the use of valorization panels. Valorization is the concept of disseminating and exploiting the results of scientific (particularly academic) research results to society (particularly industries and governments) to ensure the value of this knowledge is used in practice. For this purpose, research results of the projects are discussed with valorization panels, consisting of representatives of industries and governments.

As part of the NWO program, a project team which consisted of the editors of this book was granted funding for research with regard to responsible innovation of data

[2] http://www.nwo.nl/nwohome.nsf/pages/NWOA_73HBPY_Eng
[3] Vedder, A.H., and Custers, B.H.M. (2009).
[4] Friedman, B., Kahn, P.H., Jr., and Borning, A. (2006).

mining and profiling tools.[5] The aim of this project was to investigate how and to what extent legal and ethical rules can be integrated into data mining algorithms to prevent discrimination. For the practical testing of theories this project developed, data sets in the domain of public security made available by police and justice departments, were used for testing. The project's focus was on preventing an outcome according to which selection rules turn out to discriminate particular groups of people in unethical or illegal ways. Key questions were how existing legal and ethical rules and principles can be translated into formats understandable to computers and in which way these rules can be used to guide the data mining process. Furthermore, the technological possibilities were used as feedback to formulate concrete guidelines and recommendations for formalizing legislation. These concrete tasks also related to broader and abstract themes, such as clarifying how existing ethical and legal principles are to be applied to new technologies and what the limits of privacy are. Contrary to previous scholarly attempts to examine privacy in data mining, this project did not focus on (a priori) access limiting measures regarding input data. The project's focus rather was on (a posteriori) responsibility and transparency. Instead of limiting the access to data, which is increasingly hard to enforce, questions as to how data can and may be used were stressed.

The research project was scheduled to run from October 2009 to October 2010 and conclude at that time. In reality, it never did. The research results encouraged us to engage in further research, particularly when we discovered that simply deleting discrimination sensitive characteristics (such as gender, ethnic background, nationality) from databases still resulted in (possibly) discriminating patterns. In other words, things were far more complicated than everyone initially thought. New algorithms were developed to prevent discrimination and violations of privacy. Thus far, the research results were presented in several internationally acclaimed scientific journals, at international conferences in seven countries and in technical reports, book chapters and popular journals. A complete overview of the research results can be found at the wiki of the project.[6]

During one of the meetings with the valorization panel, the panel members suggested that the research results, particularly the more technical results, are very interesting for people with a non-technical background. Thus, the valorization panel asked us whether it would be possible to combine the research results in a book that explains the latest technological developments with regard to data mining and profiling in a manner which is comprehensible to a crowd which lacks a technological background. This book tries to achieve this. This book presents the research results of our project together with contributions of leading authors in this field, all written in a non-technical language. Complicated equations were avoided as much as possible or moved to the footnotes. Technological terminology is avoided in some places and carefully explained in other places. Similarly, the jargon of the legal and other non-technical chapters is avoided or carefully explained. All this should help non-technical readers to understand what is technologically already possible (or impossible) and how exactly it works. At the same time it should help technical readers to understand how end users really view, use and judge these technological tools and why they are sometimes

[5] http://www.nwo.nl/nwohome.nsf/pages/NWOP_8K6G4N_Eng
[6] http://wwwis.win.tue.nl/~tcalders/dadm/doku.php

criticized. A more thorough understanding of all these disciplines may help responsible innovation and technology use.

1.2 Data Mining and Profiling

This book addresses the effects of data mining and profiling, two technologies that are no longer new but still subject to constant technological developments. Data mining and profiling are often mentioned in the same breath, but they may be considered separate technologies, even though they are often used together. Profiling may be carried out without the use of data mining and vice versa. In some cases, profiling may not even involve (much) technology, for instance, when psychologically profiling a serial killer. There are many definitions of data mining and profiling. The focus of this book is not on definitions, but nevertheless, a description of what we mean by these terms may be useful.

Before starting, it is important to note that data mining refers to actions that go beyond a mere statistical analysis. Although data mining results in statistical patterns, it should be mentioned that data mining is different from traditional statistical methods, such as taking test samples.[7] Data mining deals with large databases that may contain millions of records. Statisticians, however, are used to a lack of data rather than to abundance. The large amounts of data and the way the data is stored make straightforward statistical methods inapplicable. Most statistical methods also require clean data, but, in large databases, it is unavoidable that some of the data is invalid. For some data types, some statistical operations are not allowed and some of the data may not even be numerical, such as image data, audio data, text data, and geographical data. Furthermore, traditional statistical analysis usually begins with an hypothesis that is tested against the available data. Data mining tools usually generate hypotheses themselves and test these hypotheses against the available data.

1.2.1 Data Mining: A Step in the KDD-Process

Data mining is an automated analysis of data, using mathematical algorithms, in order to find new patterns and relations in data. Data mining is often considered to be only one step, the crucial step though, in a process called Knowledge Discovery in Databases (KDD). Fayyad et al. define Knowledge Discovery in Databases as the nontrivial process of identifying valid, novel, potentially useful, and ultimately understandable patterns in data.[8] This process consists of five successive steps, as is shown in Figure 1.1. In this section, it is briefly explained how the KDD process takes place.[9] A more detailed account on data mining techniques is provided in Chapter 2.

---

[7] Hand, D.J. (1998).

[8] Fayyad, U.M., Piatetsky-Shapiro, G. and Smyth, P. (1996b), p. 6.

[9] Distinguishing different steps in the complex KDD process may also be helpful in developing ethical and legal solutions for the problems of group profiling using data mining.

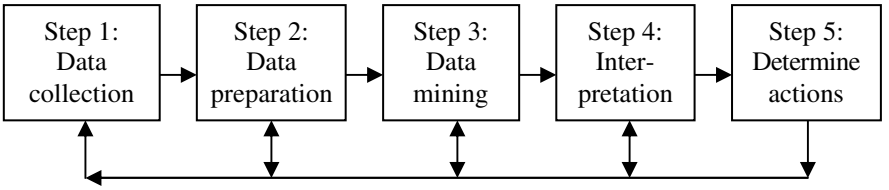

Fig. 1.1 Steps in the KDD process

Step 1: Data Collection

The first step in the KDD process is the collection of data. In the case of information about individuals, this may be done explicitly, for instance, by asking people for their personal data, or non-explicitly, for instance, by using databases that already exist, albeit sometimes for other purposes. The information requested usually consists of name, address and e-mail address. Depending on the purpose for which the information will be used, additional information may be required, such as credit card number, occupation, hobbies, date of birth, fields of interests, medical data, etc.

It is very common to use inquiries to obtain information, which are often mandatory in order to obtain a product, service, or price reduction. In this way, a take-it-or-leave-it situation is created, in which there is often no choice for a consumer but to fill in his personal data.[10] In most cases, the user is notified of the fact that privacy regulations are applied to the data. However, research shows that data collectors do not always keep this promise, especially in relation to information obtained on the Internet.[11] The same research also shows that customers are often not informed about the use that is made of the information, and in general much more information is asked for than is needed, mainly because it is thought that such data may be useful in the future.

Step 2: Data Preparation

In the second step of the KDD process, the data is prepared by rearranging and ordering it. Sometimes, it is desirable that the data be aggregated. For instance, zip codes may be aggregated into regions or provinces, ages may be aggregated into five-year categories, or different forms of cancer may be aggregated into one disease group. In this stage, a selection is often made of the data that may be useful to answer the questions set forth. But in some cases, it may be more efficient to make such a selection even earlier, in the data collection phase. The type of data and the structure and dimension of the database determine the range of data-mining tools that may be applied. This may be taken into account in selecting which of the available data will be used for data mining.

[10] These take-it-or-leave-it options are sometimes referred to as *conditional offers*.
[11] Artz, M.J.T. and Eijk, M.M.M. van (2000).

Step 3: Data Mining

The third step is the actual data-mining stage, in which the data are analyzed in order to find patterns or relations. This is done using mathematical algorithms. Data mining is different from traditional database techniques or statistical methods because what is being looked for does not necessarily have to be known. Thus, data mining may be used to discover new patterns or to confirm suspected relationships. The former is called a 'bottom-up' or 'data-driven' approach, because it starts with the data and then theories based on the discovered patterns are built. The latter is called a 'top-down' or 'theory-driven' approach, because it starts with a hypothesis and then the data is checked to determine whether it is consistent with the hypothesis.[12]

There are many different data-mining techniques. The most common types of discovery algorithms with regard to group profiling are clustering, classification, and, to some extent, regression. Clustering is used to describe data by forming groups with similar properties; classification is used to map data into several predefined classes; and regression is used to describe data with a mathematical function. Chapter 2 will elaborate on the data mining techniques.

In data mining, a *pattern* is a statement that describes relationships in a (sub)set of data such that the statement is simpler than the enumeration of all the facts in the (sub)set of data. When a pattern in data is interesting and certain enough for a user, according to the user's criteria, it is referred to as *knowledge*.[13] Patterns are interesting when they are novel (which depends on the user's knowledge), useful (which depends on the user's goal), and nontrivial to compute (which depends on the user's means of discovering patterns, such as the available data and the available people and/or technologies to process the data). For a pattern to be considered knowledge, a particular certainty is also required. A pattern is not likely to be true across *all* the data. This makes it necessary to express the certainty of the pattern. Certainty may involve several factors, such as the integrity of the data and the size of the sample.

Step 4: Interpretation

Step 4 in the KDD process is the interpretation of the results of the data-mining step. The results, mostly statistical, must be transformed into understandable information, such as graphs, tables, or causal relations. The resulting information may not be considered knowledge by the user: many relations and patterns that are found may not be useful in a specific context. A selection may be made of useful information. What information is selected, depends on the questions set forth by those performing the KDD process.

An important phenomenon that may be mentioned in this context is *masking*. When particular characteristics are found to be correlated, it may be possible to use trivial characteristics as indicators of sensitive characteristics. An example or this is indirect discrimination using redlining. Originally redlining is the practice

[12] SPSS Inc. (1999), p. 6.
[13] Adriaans, P. and Zantinge, D. (1996), p. 135.

of denying products and services in particular neighborhoods, marked with a red line on a map to delineate where not to invest. This resulted in discrimination against black inner city neighborhoods. For instance, when people living in a particular zip code area have a high health risk, insurance companies may use the zip code (trivial information) as an indication of a person's health (sensitive information), and may thus use the trivial information as a selection criterion. Note that refusing insurance on the basis of a zip code may be acceptable, as companies may choose (on the basis of market freedom) the geographic areas in which they operate. On the other hand, refusing insurance on the basis of sensitive data may be prohibited on the basis of anti-discrimination law. Masking may reduce transparency for a data subject, as he or she may not know the consequences of filling in trivial information, such as a zip code. In databases redlining may occur not necessarily by geographical profiling, but also by profiling other characteristics

Step 5: Acting upon Discovered Knowledge

Step 5 consists of determining corresponding actions. Such actions are, for instance, the selection of people with particular characteristics or the prediction of people's health risks. Several practical applications are discussed in Part III of this book. During the entire knowledge discovery process, it is possible –and sometimes necessary– to feedback information obtained in a particular step to earlier steps. Thus, the process can be discontinued and started over again when the information obtained does not answer the questions that need to be answered.

1.2.2 From Data to Knowledge

The KDD-process may be very helpful in finding pattern and relations in large databases that are not immediately visible to the human eye. Generally, deriving patterns and relations are considered creating added value out of databases, as the patterns and relations provide insight and overview and may be used for decision-making. The plain database may not (or at least not immediately) provide such insight. For that reason, usually a distinction is made between the terms data and knowledge. Data is a set of facts, the raw material in databases usable for data mining, whereas knowledge is a pattern that is interesting and certain enough for a user.[14] It may be obvious that knowledge is therefore a subjective term, as it depends on the user. For instance, a relation between vegetable consumption and health may be interesting to an insurance company, whereas it may not be interesting to an employment agency. Since a pattern in data must fulfill two conditions (*interestingness* and *certainty)* in order to become knowledge, we will discuss these conditions in more detail.

Interestingness

According to Frawley et al. (1991), interestingness requires three things: novelty, usefulness and non-triviality. Whether a pattern is *novel* depends on the *user's*

[14] Frawley, W.J., Piatetsky-Shapiro, G. and Matheus, C.J. (1993).

knowledge. A pattern that is not new may not be interesting. For instance, when a pattern is found according to which car accidents occur only in the group of people of over 18 years of age, this is not surprising, since the user may have already expected this.[15] Whether a pattern is already known to other people does not matter; what matters is that the pattern is new to the user.

A pattern is *useful* when it may help in achieving the *user's goals*. A pattern that does not contribute to achieving those goals may not be interesting. For instance, a pattern that indicates groups of people who buy many books is of no interest to a user who wants to sell CDs. Usefulness may be divided into an efficacy component and an efficiency component. *Efficacy* is an indication of the extent to which the knowledge contributes to achieving a goal or the extent to which the goal is achieved. *Efficiency* is an indication of the speed or easiness with which the goal is achieved.

Non-triviality depends on the *user's means*. The user's means have to be proportional to non-triviality: a pattern that is too trivial to compute, such as an average, may not be interesting. On the other hand, when the user's means are too limited to interpret the discovered pattern, it may also be difficult to speak of 'knowledge'. Looking at Figure 1.1 again, where the KDD process is illustrated, may clarify this, as a certain insight is required for Step 4, in which the results of data mining are interpreted.

Certainty

The second criterion for knowledge, certainty, depends on many factors. The most important among them are the integrity of the data, the size of the sample, and the significance of the calculated results. The *integrity* of the data concerns corrupted and missing data. When only corrupted data are dealt with, the terms *accuracy* or *correctness* are used.[16] When only missing data are dealt with, the term *completeness* is used. Integrity may refer to both accuracy and the completeness of data.[17]

Missing data may leave blank spaces in the database, but it may also be made up, especially in database systems that do not allow blank spaces. For instance, the birthdays of people in databases tend to be (more often than may be expected) on the 1st of January, because 1-1 is easiest to type.[18] Sometimes, a more serious effort is made to construct the values of missing data.[19]

The *sample size* is a second important factor influencing certainty. However, the number of samples that needs to be taken may be difficult to determine. In general, the larger the sample size, the more certain the results. Minimum sample sizes for acceptable reliabilities may be about 300 data items. These and larger samples, sometimes running up to many thousands of data items, used to be problematic for statistical research, but current databases are usually large enough to provide for enough samples.[20]

[15] In Europe, driving licenses may generally be obtained from the age of 18.
[16] Berti, L., and Graveleau, D. (1998).
[17] Stallings, W. (1999).
[18] Denning, D.E. (1983).
[19] Holsheimer, M., and Siebes, A. (1991).
[20] Hand, D.J. (1998).

A third important factor influencing certainty is *significance*. Significance indicates whether a discovered result is based on coincidence. For instance, when a coin is thrown a hundred times, it may be expected that heads and tails will each occur fifty times. If a 49-51 ratio were to be found, this may be considered a coincidence, but if a 30-70 ratio were found, it may be difficult to assume this is coincidental. The latter result is significantly different from what is expected. With the help of confidence intervals (see below), it is possible to determine the likelihood of whether a discovered result may be considered a coincidence or not.

Once the certainty of particular knowledge has been determined using a chosen mathematical method, it is up to the user to decide whether that certainty is sufficient for further use of that knowledge. The standard technique for calculating certainty in the case of regression techniques is the calculation of the *standard error*. The standard error indicates to what extent the data differs from the regression function determined. The larger the value of the standard error, the larger the spreading of the data. Using standard errors, it is possible to calculate *confidence intervals*. A confidence interval is a range of values with a given chance of containing the real value. In most cases, the user's confidence interval is chosen in such a way that confidence is fixed at 95 or 99 per cent.

Finally, it should be mentioned that for profiles, certainty is closely related to reliability. The reliability of a profile may be split into (a) the reliability of the profile itself, which comprises certainty, and (b) the reliability of the use of the profile. This distinction is made because a particular profile may be entirely correct from a technological perspective, but may still be applied incorrectly. For instance, when data mining reveals that 80 % of all motels are next to highways, this may be a result with a particular certainty. When all motels were counted, the certainty of this pattern is 100 %, but when a sample of 300 motels were taken in consideration, of which 240 turned out to lie next to highways, the certainty may be less because of the extrapolation. However, if a motel closes or a new motel opens, the reliability of the pattern decreases, because the pattern is based on data that are no longer up to date, yielding a pattern that represents reality with less reliability. The reliability of the use of a particular profile is yet another notion. Suppose a particular neighborhood has an unemployment rate of 80 %. When a local government addresses all people in this neighborhood with a letter regarding unemployment benefits, their use of the profile is not 100 % reliable, as they also address people who are employed.

1.2.3 Profiles of Individuals and of Groups

Profiling is the process of creating profiles. Although profiles can be made of many things, such as countries, companies or processes, in this book we focus on profiles of people or groups of people. Hence, we consider a profile a property or a collection of properties of an individual or a group of people. Several names exist for these profiles. Personal profiles are also referred to as *individual profiles* or *customer profiles,* while group profiles are also referred to as *aggregated profiles.* Others use the terms *abstract profiles* and *specific profiles* for group

profiles and personal profiles, respectively.[21] Another common term is *risk profiles*, indicating the some kind of risk of an individual or group of people (such as the risk of getting a heart-attack, of not paying your mortgage or of being a terrorist).

A personal profile is a property or a collection of properties of a particular individual. A *property*, or a *characteristic*, is the same as an *attribute*, a term more used often in computer sciences. An example of a personal profile is the personal profile of Mr John Doe (44), who is married, has two children, earns 25,000 Euro a year, and has two credit cards and no criminal record. He was hospitalized only twice in his life, once for appendicitis and last year because of lung cancer.

A group profile is a property or a collection of properties of a particular group of people.[22] Group profiles may contain information that is already known; for instance, people who smoke live, on average, a few years less than people who do not. But group profiles may also show new facts; for instance, people living in zip code area 8391 may have a (significantly) larger than average chance of having asthma. Group profiles do not have to describe a causal relation. For instance, people driving red cars may have (significantly) more chances of getting colon cancer than people driving blue cars. Note that group profiles differ from individuals with regard to the fact that the properties in the profile may be valid for the group and for individuals as members of that group, though not for those individuals as such. If this is the case, this is referred to as *non-distributivity* or non-distributive properties.[23] On the other hand, when properties are valid for each individual member of a group as an individual, this is referred to as *distributivity* or distributive properties.

Several data mining methods are particularly suitable for profiling. For instance, classification and clustering may be used to identify groups.[24] Regression is more useful for making predictions about a known individual or group. More on these and other techniques can be found in Chapter 2.

1.2.4 Why We Need These Tools

The use of data mining and profiling is still on the increase, mainly because they are usually very efficient and effective tools to deal with the (also) ever increasing amounts of data that we collect and process in our information society. According to Moore's Law, the number of transistors on an integrated circuit (a 'chip' or 'microchip') for minimum component costs doubles every 24 months.[25] This more or less implies that storage capacity doubles every two years (or that data storage costs are reduced by fifty percent every two years). This empirical observation by Gordon Moore was made in 1965; by now, this doubling speed is approximately 18 months. From this perspective there is hardly any need to limit the amounts of

[21] See Bygrave, L.A. (2002), p. 303, and Clarke, R. (1997).
See www.anu.edu.au/people/roger.clarke/dv/custproffin.html.
[22] Note that when the group size is 1, a group profile boils down to a personal profile.
[23] Vedder, A.H. (1999).
[24] SPSS Inc. (1999), p. 131.
[25] Schaller, R.R. (1997).

data we are collecting and processing. However, the amounts of data are enormous, so we do need tools to deal with these huge amounts of data. Data mining and profiling are exactly the type of technologies that may help us with analyzing and interpreting large amounts of data.

It is important to stress that due to Moore's Law we cannot get around the need for data mining and profiling tools. These tools, along with other tools for data structuring and analysis, are extremely important and it would be very difficult for an information society like ours if they would not be available. To stress this point we will provide here some major advantages of profiling. The advantages of profiling usually depend on the context in which they are used. Nevertheless, some advantages may hold for many or most contexts. At times group profiles may be advantageous compared to individual profiles. Sometimes profiling, whether it is individual profiling or group profiling, may be advantageous compared to no profiling at all. The main advantages of profiling, particularly of group profiling, concern *efficacy*, i.e., how much of the goal may be achieved, and *efficiency*, i.e., how easily the goal may be achieved. Data mining and profiling may process huge amounts of data in a short time; data that is often too complex or too great for human beings to process manually. When many examples are present in databases, (human) prejudices as a result of certain expectations may be avoided.

Profiling may be a useful method of finding or identifying target groups. In many cases, group profiling may be preferable to individual profiling because it is more cost efficient than considering each individual profile. This *cost efficiency* may concern lower costs in the gathering of information, since less information may be needed for group profiles than for individual profiles. Remember that if a group profile is based on less information, it is usually less reliable (see Section 1.2.2). But higher costs may also be expected in the time-consuming task of approaching individuals. While individuals may be approached by letter or by phone, groups may be approached by an advertisement or a news item. Take as an example baby food that is found to be poisoned with chemicals. Tracing every person who bought the baby food may be a costly process, it may take too much time, and some people may not be traced at all. A news item and some advertisements, for instance, in magazines for parents with babies, may be more successful.

Another advantage of group profiling over individual profiling is that group profiles may offer more possibilities for selecting targets. An individual may not appear to be a target on the basis of a personal profile, but may still be one. Group profiles may help in tracking down potential targets in such cases. For instance, a person who never travels may not seem an interesting target to sell a travel guide to. Still, this person may live in a neighborhood where people travel frequently. She may be interested in travel guides, not so much for using them for her own trips, but rather to be able to participate in discussions with her neighbors. A group profile for this neighborhood predicts this individual's potential interest in travel guides, whereas an individual profile may not do so. Such selection may also turn out to be an advantage for the targets themselves. For instance, members of a

high-risk group for lung cancer may be identified earlier and treated, or people not interested in cars will no longer receive direct mail about them.

Profiling, regardless of whether individuals or groups are profiled, may be more useful than no profiling at all. Without any profiling or selection, the efficiency or 'hit ratio' is usually poor. For instance, advertising using inadequately defined target groups, such as on television, is less efficient than advertising only to interested and potentially interested customers.

1.3 Discrimination, Privacy and Other Issues

Despite all the opportunities described in the previous section, there are also concerns about the use of data mining and profiling. This book deals with the effects of data mining and profiling. By effects, we refer to a neutral term of what the use of these tools may result in. These effects can be positive (or at least positive to some people), as illustrated in the previous section and will be illustrated in Part III of this book. However, these effects can also be negative (or at least negative to some people). This book will deal with two major potentially negative effects of data mining and profiling, namely discrimination and privacy invasions. That is not to say that these are the only possible negative effects. Other negative effects, such as de-individualization,[26] possible loss of autonomy, one-sided supply of information, stigmatization and confrontation with unwanted information may be other examples of possible negative effects.[27] However, this book will focus on discrimination and privacy issues regarding data mining and profiling, since most progress has been made in the development of discrimination-aware and privacy preserving data mining techniques. Furthermore, even though discrimination and privacy may sometimes be difficult notions in law and ethics, they are still easier to grasp than notions like de-individualization and stigmatization, for which there hardly any legal concepts. For instance, most countries have laws regarding equal treatment (non-discrimination) and privacy, but laws against de-individualization or stigmatization are unknown to us.

1.3.1 Any News?

A New Book

Over the last years, many books and papers have been written on the possible effects of data mining and profiling.[28] What does this book to add to all this knowledge already available? First of all, most of these books focus on privacy issues, whereas this book explicitly takes discrimination issues into account. Second, we tried to include more technological background in this book, in a way that should be understandable to readers with a non-technical background. Third,

[26] Vedder, A.H. (1999).
[27] Custers, B.H.M. (2004), p. 77.
[28] Hildebrandt, M. and Gutwirth, S. (2008); Harcourt, B.E. (2007); Schauer, F. (2003); Zarsky, T. (2003); Custers, B.H.M. (2004).

this book provides technological solutions, particularly discrimination aware and privacy preserving data mining techniques. Fourth, this book explains state of the art technologies, an advantage over books published before, even though we realize that technological developments are very fast, outdating this book also within a few years.

A New Technology

Profiles were used and applied in the past without data mining, for instance, by (human) observation or by empirical statistical research. Attempts were often made to distinguish particular individuals or groups and investigate their characteristics. Thus, it may be asked what is new about profiling by means of data mining? Is it not true that we have always drawn distinctions between people?

Profiling by means of data mining may raise problems that are different from the problems that may be raised by other forms of statistical profiling such as taking test samples, mainly because data mining generates hypotheses itself. Empirical statistical research with self-chosen hypotheses may be referred to as *primary data analysis*, whereas the automated generating and testing of hypotheses, as with data mining, may be referred to as *secondary data analysis*. In the automated generating of hypotheses, the known problems of profiling may be more severe and new types of problems may arise that are related to profiling using data mining.[29] There are four reasons why profiling using data mining may be different from traditional profiling.

The first reason why profiling using data mining may cause more serious problems is a scale argument. Testing twice as much hypotheses with empirical research implies doubling the amount of researchers. Data mining is an automated analysis and does not require doubling the amount of researchers. In fact, data mining enables testing large numbers (hundred or thousands) of hypotheses (even though only a very small percentage of the results may be useful). There may be an overload of profiles.[30] Although this scale argument indicates that the known problems of group profiling are more severe, it does not necessarily imply new problems.

A second difference is that, in data mining, depending on the techniques that is used, every possible relation can be investigated, while, in empirical statistical research, usually only causal relationships are considered. The relations found using data mining are not necessarily causal. Or they may be causal without being understood. In this way, the scope of profiles that are discovered may be much broader (only a small minority of all statistical relations is directly causal) with unexpected profiles in unexpected areas. Data mining is not dependent on coincidence. Data-mining tools automatically generate hypotheses, independent of whether a relationship is (expected to be) causal or not.

---

[29] A distinction may be made between technology-specific and technology-enhanced concerns, because technology-specific concerns usually require new solutions, while conventional solutions may suffice for the technology-enhanced concerns. See also Tavani, H. (1999).

[30] See also Mitchell, T.M. (1999) and Bygrave, L.A. (2002) , p. 301.

Profiles based on statistical (but not necessarily causal) relationships may result in problems that are different from the problems of profiles based on causal relations, such as the aforementioned masking. Statistical results of data mining are often used as a starting point to find underlying causality, but it is important to note that merely statistical relations may already be sufficient to act upon, for instance, in the case of screening for diseases. The automated generation of hypotheses contributes to the scale argument as well: the number of profiles increases largely because non-causal relations can be found as well.

A third difference between data mining and empirical statistical research is that with the help of data mining trivial information may be linked (sometimes unintentionally) to sensitive information. Suppose data mining shows a relation between driving a red car and developing colon cancer. Thus, a trivial piece of information, the color of a person's car, becomes indicative of his or her health, which is sensitive information. In such cases the lack of transparency regarding data mining may start playing an important role: people who provide only trivial information may be unaware of the fact that they may also be providing sensitive information about themselves when they belong to a group of people about whom sensitive information is known. People may not even know to what groups they belong.

A fourth difference lies in a characteristic of information and communication technology that is usually referred to as the 'lack of forgetfulness of information technology'.[31,32] Once a piece of information has been disclosed, it is practically impossible to withdraw it. Computer systems do not forget things, unless information is explicitly deleted, but even then information can often be retrieved.[33] Since it is often difficult to keep information contained, it may spread through computer systems by copying and distribution. Thus, it may be difficult to trace every copy and delete it. This technological characteristic requires a different approach to finding solutions for the problems of profiling and data mining.

1.3.2 Problems and Solutions

This is a book about discrimination and privacy. That makes it a book on problems. However, instead of only discussing problems, we also provide solutions or directions for solutions to these problems. If data mining and profiling have undesirable effects, it may be regulated in several ways. Lessig distinguishes four different elements that regulate.[34] For most people, the first thing that comes to mind is to use legal constraints. Laws may regulate where and when and by whom data mining and profiling are allowed and under which conditions and circumstances. They operate as a kind of constraint on everyone who wants to use data mining and profiling.

---

[31] Blanchette, J.F., and Johnson, D.G. (1998).

[32] For this argument it should be noted that data mining is regarded as an information technology, contrary to empirical statistical research.

[33] It may be argued that paper files do not 'forget' either, but paper files are, in general, less accessible and thus there is generally less spreading of the information they contain.

[34] Lessig, L. (2006).

But laws are not the only, and often not the most significant constraint, to regulate something. Sometimes, things may be legal, but nevertheless considered unethical or impolite. Lessig mentions the example of smoking, something that is not illegal in many places, but may be considered impolite, at least without asking permission of others present in the same room. Examples of ethical issues that are strictly speaking not illegal that we will come across in this book are stigmatization of people, polarization of groups in society and de-individualization. Such norms have a certain constraint on behavior.

Apart from laws and norms, a third force is the market. Price and quality of products are important factors here. When the market supplies a wide variety of data mining and profiling tools (some of these tools may be less discriminating or more privacy friendly than others), there is more to choose from, reducing constraints. However, when there are only one or two options available, the market constrains the options. High prices (for instance, for data mining tools that do not discriminate or are privacy friendly) that may limit what you can buy.

The fourth and last constraint is created by technology. How a technology is built (its architecture) determines how it can be used. Walls may constrain where you are can go. A knife can be used for good purposes, like cutting bread, or for bad purposes, like hurting a person. Sometimes these constraints are not intended, but sometimes they are explicitly included in the design of a particular technology. Examples are copy machines that refuse to copy banknotes and cars that refuse to start without keys and, in some cases, without alcohol tests. In our case of data mining and profiling technologies, there are many constraints that can be built into the technologies. That is the reason why we separated these 'solutions in code' (Part IV of this book) from the other solutions (Part V of this book). Although this book has a strong focus on technological solutions, this does not mean, however, that this is the only (type of) solution. In some cases, what is needed are different attitudes, and in some cases new or stricter laws and regulations.

1.4 Structure of This Book

1.4.1 Part I: Opportunities of Data Mining and Profiling

Part I of this book explains the basics of data mining and profiling and discusses why these tools are extremely useful in the information society.

In Chapter 2, Calders and Custers explain what data mining is and how it works. The field op data mining is explored and compared with related research areas, such as statistics, machine learning, data warehousing and online analytical processing. Common terminology regarding data mining that will be used throughout this book is discussed. Calders and Custers explain the most common data mining techniques, i.e., classification, clustering and pattern mining, as well as some supporting techniques, such as pre-processing techniques and database coupling.

In Chapter 3, Calders and Žliobaitė explain why and how the use of data mining tools can lead to discriminative decision procedures, even if all

discrimination sensitive data in the databases is removed or suppressed before the data mining is commenced. It is shown how data mining may exhibit discriminatory behavior towards particular groups based, for instance, upon gender or ethnicity. It is often suggested that removing all discrimination sensitive attributes such as gender and ethnicity from databases may prevent the discovery of such discriminatory relationships.[35] Without sensitive data it is impossible to find sensitive patterns or relations, it is argued. Calders and Žliobaitė show that this is not necessarily true. They carefully outline three realistic scenarios to illustrate this and explain the reasons for this phenomenon.

1.4.2 Part II: Possible Discrimination and Privacy Issues

Part II of this book explains the basics of discrimination and privacy and discusses how data mining and profiling may cause discrimination and privacy issues.

In Chapter 4, Gellert, De Vries, De Hert and Gutwirth compare and distinguish between European anti-discrimination law and data protection law. They show that both rights have the same structure and increasingly turn to the same mode of operation in the information society, even though their content is far from identical. Gellert, De Vries, De Hert and Gutwirth show that this is because both rights are grounded in the notion of negative freedom as evidenced by I. Berlin[36], and thus aim at safeguarding the autonomy of the citizen in the information society. Finally, they analyze two cases where both rights apply, and draw conclusions on how to best articulate the two tools.

In Chapter 5, Pedreschi, Ruggieri and Turini address the problem of discovering discrimination in large databases. Proving discrimination may be difficult. For instance, was a job applicant turned down because she was pregnant or because she was not suited for the job? In a single case, this may be difficult to prove, but it may be easier if there are many cases. For instance, if a company with over one thousand employees has no employees from ethnic minorities, this may be due to discrimination. Similarly, when all top management boards in a country consist of 90% of males, this may indicate possible discrimination. In Chapter 5, the focus is on finding discriminatory situations and practices hidden in large amounts of historical decision records. Such patterns and relations may be useful for anti-discrimination authorities. Pedreschi, Ruggieri and Turini discuss the challenges in discovering discrimination and present an approach for finding discrimination on the basis of legally-grounded interesting measures.

In Chapter 6, Romei and Ruggieri present an annotated bibliography on discrimination analysis. Literature on discrimination discovery and prevention is mapped in the areas of law, sociology, economics and computer sciences. Relevant legal and sociological concepts such as prejudices, racism, affirmative action (positive discrimination) and direct versus indirect discrimination are

[35] For instance, article 8 of the European Data Protection Directive (95/46/EC) explicitly limits the processing of special categories of data that is considered especially sensitive to data subjects, such as personal data revealing racial or ethnic origin, political opinions, religious or philosophical beliefs, trade-union membership, health and sex life.

[36] Berlin, I. (1969).

introduced guided by ample references. Furthermore, literature on economic models of labor discrimination, approaches for collecting and analyzing data, discrimination in profiling and scoring and recent work on discrimination discovery and prevention is discussed. This inventory is intended to provide a common basis to anyone working in this field.

In Chapter 7, Schermer maps out risks related to profiling and data mining that go beyond discrimination issues. Risks such as de-individualization and stereotyping are described. To mitigate these and other risks, traditionally the right to (informational) privacy is invoked. However, due to the rapid technological developments, privacy and data protection law have several limitations and drawbacks. Schermer discusses why it is questionable whether privacy and data protection legislation provide adequate levels of protection and whether these legal instruments are effective in balancing different interests when it comes to profiling and data mining.

1.4.3 Part III: Practical Applications

Part III of this book sets forth several examples of practical applications of data mining and profiling. These chapters intend to illustrate the added value of applying data mining and profiling tools. They also show several practical issues that practitioners may be confronted with.

In Chapter 8, Kamiran and Žliobaitė illustrate how self-fulfilling prophecies in data mining and profiling may occur. Using several examples they show how models learnt over discriminatory data may result in discriminatory decisions. They explain how discrimination can be measured and show how redlining may occur. Redlining originally is the practice of denying products and services in particular neighborhoods, marked with a red line on a map to delineate where not to provide credit. This resulted in discrimination against black inner city neighborhoods. In databases this effect may also occur, not necessarily by geographical profiling, but also by profiling other characteristics. Kamiran and Žliobaitė present several techniques to preprocess the data in order to remove discrimination, not by removing all discriminatory data or all differences between sensitive groups, but by addressing differences unacceptable for decision-making. With experiments they demonstrate the effectiveness of these techniques.

In Chapter 9, Schakel, Rienks and Ruissen focus on knowledge discovery and profiling in the specific context of policing. They observe that the positivist epistemology underlying the doctrine of information-led policing is incongruent with the interpretive-constructivist basis of everyday policing, and conclude that this is the cause of its failure to deliver value at the edge of action. After shifting focus from positivist information-led policing to interpretive-constructivist knowledge-based policing, they illustrate how profiling technologies can be used to design augmented realities to intercept criminals red-handedly. Subsequently, Schakel, Rienks and Ruissen discuss how the processing of data streams (rather than databases) can meet legal requirements regarding subsidiarity, proportionality, discrimination and privacy.

In Chapter 10, Van den Braak, Choenni and Verwer discuss the challenges concerning combining and analyzing judicial databases. Several organizations in the criminal justice system collect and process data on crime and law enforcement. Combining and analyzing data from different organizations may be very useful, for instance, for security policies. Two approaches are discussed, a data warehouse (particularly useful on an individual level) and a dataspace approach (particularly useful on an aggregated level). Though in principle all applications exploiting judicial data may violate data protection legislation, Van den Braak, Choenni and Verwer show that a dataspace approach is preferable with regard to taking precautions against such data protection legislation violations.

1.4.4 Part IV: Solutions in Code

Part IV of this book provides technological solutions to the discrimination and privacy issues discussed in Part II.

In Chapter 11, Matwin provides a survey of privacy preserving data mining techniques and discusses the forthcoming challenges and the questions awaiting solutions. Starting with protection of the data, methods for identity disclosure and attribute disclosure are discussed. However, adequate protection of the data in databases may not be sufficient: privacy infringements may also occur based on the inferred data mining results. Therefore, also model based identity disclosure methods are discussed. Furthermore, methods for sharing data for data mining purposes while protecting the privacy of people who contributed the data are discussed. Specifically, the chapter presents scenarios in which data is shared between a number of parties, either in a horizontal a or vertical partition. Then the privacy of individuals who contributed the data is protected by special-purpose cryptographic techniques that allow parties performing meaningful computation on the encrypted data. Finally, Matwin discusses new challenges like data from mobile devices, data from social networks and cloud computing.

In Chapter 12, Kamiran, Calders and Pechenizkiy survey different techniques for discrimination-free predictive models. Three types of techniques are discussed. First, removing discrimination from the dataset before applying data mining tools. Second, changing the learning procedures by restricting the search space to models that are not discriminating. Third, adjusting the models learned by the data mining tools after the data mining process. These techniques may significantly reduce discrimination at the cost of accuracy. The authors' experiments show that still very accurate models can be learned. Hence, the techniques presented by Kamiran, Calders and Pechenizkiy provide additional opportunities for policymakers to balance discrimination against accuracy.

In Chapter 13, Hajian and Domingo-Ferrer address the prevention of discrimination that may result from data mining and profiling. Discrimination prevention consists of inducing patterns that do not lead to discriminatory decision, even if the original data in the database is inherently biased. A taxonomy is presented for classifying and examining discrimination prevention methods. Next, preprocessing discrimination prevention methods are introduced and it is discussed how these methods deal with direct and indirect discrimination

respectively. Furthermore, Hajian and Domingo-Ferrer present metrics that can be used to evaluate the performance of these approaches and show that discrimination removal can be done at a minimal loss of information.

In Chapter 14, Verwer and Calders show how positive discrimination (also known as affirmative action) can be introduced in predictive models. Three solutions based upon so-called Bayesian classifiers are introduced. The first technique is based on setting different thresholds for different groups. For instance, if there are income differences between men and women in a database, men can be given a high income label above $90,000, whereas women can be given a high income label above $75,000. Instead of income figures, the labels high and low income could be applied. This instantly reduces the discriminating pattern. The second techniques focuses on learning two separate models, one for each group. Predictions from these models are independent of the sensitive attribute. The third and most sophisticated model is focused on discovering the labels a dataset should have contained if it would have been discrimination-free. These latent (or hidden) variables can be seen as attributes of which no value is recorded in the dataset. Verwer and Calders show how decisions can be reverse engineered by explicitly modeling discrimination.

1.4.5 Part V: Solutions in Law, Norms and the Market

Part V of this book provides non-technological solutions to the discrimination and privacy issues discussed in Part II. These solutions may be found in legislation, norms and the market. Many of such solutions are discussed in other books and papers, such as (to name only a few) the regulation of profiling,[37] criteria for balancing privacy concerns and the common good,[38] self-regulation of privacy,[39] organizational change and a more academic approach,[40] and valuating privacy in a consumer market.[41] We do not discuss these suggested solutions in this book, but we do add a few other suggested solutions to this body of work.

In Chapter 15, Van der Sloot proposes to use minimum datasets to avoid discrimination and privacy violations in data mining and profiling. Discrimination and privacy are often addressed by implementing data minimization principles, restricting collecting and processing of data. Although data minimization may help to minimize the impact of security breaches, it has also several disadvantages. First, the dataset may lose value when reduced to a bare minimum and, second, the context and meaning of the data may get lost. This loss of context may cause or aggravate privacy and discrimination issues. Therefore, Van der Sloot suggests an opposite approach, in which minimum datasets are mandatory. This better ensures adequate data quality and may prevent loss of context.

In Chapter 16, Finocchiaro and Ricci focus on the opposite of being profiled, which is building one's own digital reputation. Although people have some

[37] See, for instance, Bygrave, L.A. (2002).
[38] Etzioni, A. (1999), p. 12/13.
[39] Regan, P.M. (2002).
[40] See, for instance, Posner, R.A. (2006), p. 210.
[41] See, for instance, Böhme (2009) and Böhme and Koble (2007).

choices in what information they provide about themselves to others (so-called informational self-determination),[42] this choice is limited to the data in databases and usually does not pertain to any results of data mining and profiling. Furthermore, due to the so-called lack of forgetfulness of information technology,[43] people have even less influence on their digital reputation. In order to reinforce the informational self-determination of people, Finocchiaro and Ricci propose the inverse of the right not to know,[44] which is the right to oblivion,[45] providing for the deletion of information which is no longer corresponds to an individual's identity.

In Chapter 17, Zarsky addresses the commonly heard complaint that there is a lack of transparency regarding the data that is collected by organizations and the ways in which these data are being used. Particularly in the context of data mining and profiling, transparency and transparency enhancing tools have been mentioned as important policy tools to enhance autonomy.[46] Transparency may also forward democracy, enhance efficiency and facilitate crowdsourcing, but it may also undermine policies and authority and generate stereotypes. While acknowledging that transparency alone cannot solve all privacy and discrimination issues regarding data mining and profiling, Zarsky provides a policy blueprint for analyzing the proper role and balance for transparency in data mining and profiling.

In Chapter 18, Zarsky considers whether the use of data mining can be conceptualized as a search (possibly an illegal search) and how this perspective can be used for policy responses. Illegal search is a common concept in criminal law, but applying this concept in the setting of data mining is novel. Three normative theories are introduced on illegal searches: these may be viewed as unacceptable psychological intrusions, as limits to the force of government or as limits to ´fishing expeditions´, i.e., looking through data of people who raise no suspicion. Zarsky shows how these theories can be used to understand data mining as illegal searches and how regulators and policymakers can establish which data mining practices are to be allowed and which must be prohibited.

1.4.6 Part VI: Concise Conclusions

Part VI of this book provides some concise conclusions. In Chapter 19, some general conclusions are drawn and the way forward is discussed. Throughout the book it becomes clear that a powerful paradigm shift is transpiring. The growing use of data mining practices by both government and commercial entities leads to both great promises and challenges. They hold the promise of facilitating an

---

[42] Westin, A. (1967).
[43] Blanchette, J.F., and Johnson, D.G. (1998).
[44] Chadwick, R., Levitt, M., and Shickle, D. (1997).
[45] The right to oblivion is sometimes referred to as the right to be forgotten. This right was also included in the EU proposal for revision of the EU data protection legislation that leaked end of 2011. See: https://www.privacyinternational.org/article/quick-review-draft-eu-data-protection-regulation
[46] Hildebandt, M. (2009).

information environment which is fair, accurate and efficient. At the same time, it might lead to practices which are both invasive and discriminatory, yet in ways the law has yet to grasp.

Chapter 19 starts with demonstrating this point by showing how the common measures for mitigating privacy concerns, such as a priori limiting measures (particularly access controls, anonymity and purpose specification) are mechanisms that are increasingly failing solutions against privacy and discrimination issues in this novel context.

Instead, we argue that a focus on (a posteriori) accountability and transparency may be more useful. This requires improved detection of discrimination and privacy violations as well as designing and implementing techniques that are discrimination-free and privacy-preserving. This requires further (technological) research.

But even with further technological research, there may be new situations and new mechanisms through which privacy violations or discrimination may take place. This is why Chapter 19 concludes with a discussion on the future of discrimination and a discussion on the future of privacy. With regard to discrimination, it is worth mentioning that a shift to automated predictive modeling as means of decision making and resource allocation might prove to be an important step towards a discrimination-free society. Discriminatory practices carried out by officials and employees could be detected and limited effectively. Nevertheless, two very different forms of discrimination-based problems might arise in the future. First, novel predictive models can prove to be no more than sophisticated tools to mask the "classic" forms of discrimination of the past, by hiding discrimination behind new proxies for the current discriminating factors. Second, discrimination might be transferred to new forms of population segments, dispersed throughout society and only connected by one or more attributes they have in common. Such groups will lack political force to defend their interests. They might not even know what is happening.

With regard to privacy, the adequacy of the current legal framework is discussed with regard to the technological developments of data mining and profiling discussed in this book. The European Union is currently revising the data protection legislation. The question whether these new proposals will adequately address the issues raised in this book is dealt with.

References

Adriaans, P., Zantinge, D.: Data mining. Addison Wesley Longman, England (1996)

Artz, M.J.T., van Eijk, M.M.M.: Klant in het web. In: Privacywaarborgen voor Internettoegang. Achtergrondstudies en verkenningen, vol. 17. Registratiekamer, Den Haag (2000)

Berlin, I.: Four Essays on Liberty. Oxford University Press, Oxford (1969)

Berti, L., Graveleau, D.: Designing and Filtering On-line Information Quality: New Perspectives for Information Service Providers. In: Ethicomp 1998, 4th International Conference on Ethical Issues of Information Technologies, Rotterdam (1998)

Blanchette, J.F., Johnson, D.G.: Data retention and the panopticon society: the social benefits of forgetfulness. In: Introna, L. (ed.) Computer Ethics: Philosophical Enquiry

(CEPE 1998). Proceedings of the Conference held at London School of Economics, December 13-14, pp. 113–120. London ACM SIG/London School of Economics (1998)

Böhme: Valuating Privacy with Option Pricing Theory. In: Berthold, S. (ed.) Workshop on the Economics of Information Security (WEIS 2009), June 24-25. University College London, London (2009)

Böhme, Koble: On the Viability of Privacy-Enhancing Technologies in a Self-regulated Business-to-consumer Market: Will Privacy Remain a Luxury Good? In: Proceedings of Workshop on the Economics of Information Security (WEIS), June 7-8. Carnegie Mellon University, Pittsburgh (2007)

Bygrave, L.A.: Data protection law; approaching its rationale, logic and limits. Information law series, vol. 10. Kluwer Law International, The Hague (2002)

Chadwick, R., Levitt, M., Shickle, D.: The right to know and the right not to know. Avebury Ashgate Publishing Ltd., Aldershot (1997)

Clarke, R.: Customer profiling and privacy implications for the finance industry (1997)

Custers, B.H.M.: The Power of Knowledge; Ethical, Legal, and Technological Aspects of Data Mining and Group Profiling in Epidemiology, p. 300. Wolf Legal Publishers, Tilburg (2004)

Denning, D.E.: Cryptography and Data Security. Addison-Wesley, Amsterdam (1983)

Etzioni, A.: The Limits of Privacy. Basic Books, New York (1999)

Fayyad, U.M., Piatetsky-Shapiro, G., Smyth, P.: From Data Mining to Knowledge Discovery: An Overview. In: Fayyad, U.M., Piatetsky-Shapiro, G., Smyth, P., Uthurusamy, R. (eds.) Advances in Knowledge Discovery and Data Mining. AAAI Press/The MIT Press, Menlo Park, California (1996b)

Frawley, W.J., Piatetsky-Shapiro, G., Matheus, C.J.: Knowledge Discovery in Databases; an overview. In: Piatetsky-Shapiro, G., Frawley, W.J. (eds.) Knowledge Discovery in Databases. AAAI Press/The MIT Press, Menlo Park, California (1993)

Friedman, B., Kahn Jr., P.H., Borning, A.: Value Sensitive Design and information systems. In: Zhang, P., Galletta, D. (eds.) Human-Computer Interaction in Management Information Systems: Foundations, pp. 348–372. M.E. Sharpe, Armonk (2006)

Hand, D.J.: Data mining: statistics and more? The American Statistician 52(2), 112–118 (1998)

Harcourt, B.E.: Against Prediction: Profiling, Policing and Punishing in an Actuarial Age. University of Chicago Press, Chicago (2007)

Hildebandt, M.: Behavioral Biometric Profiling and Transparency Enhancing Tools. Scientific Report of EU-program Future of Identity in the Information Society (FIDIS), WP7-D7.12 (2009), http://www.fidis.net

Hildebrandt, M., Gutwirth, S.: Profiling the European Citizen. Springer, Heidelberg (2008)

Holsheimer, M., Siebes, A.: Data Mining: the Search for Knowledge in Databases. Report CS-R9406 Centrum voor Wiskunde en Informatica, Computer Science/Department of Algorithmics and Architecture (1991)

Lessig, L.: Code Version 2.0. Basic Books, New York (2006)

Mitchell, T.M.: Machine Learning and Data Mining. Communications of the ACM 42(11) (1999)

Posner, R.A.: Uncertain Shield. Rowman & Littlefield Publishers, Inc., New York (2006)

Regan, P.M.: Privacy and commercial use of personal data: policy developments in the United States. Paper Presented at the Rathenau Institute Conference on Privacy, Amsterdam (January 17, 2002)

Schaller, R.R.: Moore's Law: Past, Present and Future. IEEE Spectrum 34, 52–59 (1997)

Schauer, F.: Profiles, Probabilities and Stereotypes. Harvard University Press, Cambridge (2003)

Solove, D.: The Digital Person; Technology and Privacy in the Information Age. University Press, New York (2004)
SPSS Inc. Data Mining with Confidence. SPSS Inc., Chicago (1999)
Stallings, W.: Cryptography and Network Security; principles and practice. Prentice Hall, Upper Saddle River (1999)
Tavani, H.: Internet Privacy: some distinctions between Internet-specific and Internet-enhanced privacy concerns. In: Proceedings of the 4th Ethicomp International Conference on the Social and Ethical Impacts of Information and Communication Technologies, Ethicomp 1999 (1999)
Vedder, A.H.: KDD: The challenge to individualism. Ethics and Information Technology 1, 275–281 (1999)
Vedder, A.H., Custers, B.H.M.: Whose responsibility is it anyway? Dealing with the consequences of new technologies. In: Sollie, P., Düwell, M. (eds.) Evaluating New Technologies: Methodological Problems for the Ethical Assessment of Technology Developments, pp. 21–34. Springer, New York (2009) (The international library of ethics, law and technology, 3)
Westin, A.: Privacy and Freedom. Bodley Head, London (1967)
Zarsky, T.: Mine 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 and Technology 5, 57 (2002-2003)