

# Ethical issues of ‘morality mining’

## Moral identity as a focus of data mining

**Markus Christen**

*Institute of Biomedical Ethics, University of Zürich, Switzerland & Psychology Department, University of Notre Dame, USA*

**Mark Alfano**

*Woodrow Wilson School and Center for Human Values, Princeton University, USA & Department of Philosophy, University of Oregon, USA*

**Endre Bangerter**

*Security Engineering Lab, Bern University of Applied Sciences, Switzerland*

**Daniel Lapsley**

*Psychology Department, University of Notre Dame, USA*

### **ABSTRACT**

When data mining aims to disclose information about the moral competences and values of individuals or groups – an undertaking we call ‘morality mining’ –, novel ethical problems emerge. These are only partially covered by the current debate on ethical data mining focusing on privacy with respect to discrimination, threats to autonomy, misuse of data, and the consequences of erroneous information. An ethics of morality mining is of particular relevance for research in social science and psychology that increasingly relies on data emerging from social networks, media portals, etc., where people act from or at least in accordance with their own values. In this conceptual contribution, we outline the basic idea of morality mining, explain why we believe that morality mining is associated with novel ethical problems, and suggest ways to address these problems.

### **INTRODUCTION**

Today, countless processes in the social and economic lives of individuals rely on the use of technological systems (Internet, cell phones, GPS, credit card pay systems, etc.) that generate large amounts of data. These data are increasingly the target of data mining (Fayyad et al. 1996, Nisbett et al 2009, Witten et al. 2011) that intends to uncover “hidden patterns” in large data sets. The results of data mining can be used in various practical ways, e.g., for market segmentation, customer profiling, recommender systems, credit rating, and fraud detection. Furthermore, they are increasingly used for understanding socio-economic

developments. A well-elaborated ethical debate on privacy violations through data mining deals with issues such as discrimination, threats to autonomy, misuse of data, and the consequences of erroneous information (Custers et al. 2012/forthcoming, Vaidya et al. 2006, and Zarksky 2003, among others). These issues are, however, not the focus of our contribution.

Our main interest is in an undertaking we call ‘morality mining’, which occurs when data mining aims to disclose information on the moral competences and values of individuals or groups. We claim that in this case novel ethical issues emerge, resulting from the fact that moral beliefs and convictions are central components of people’s identities (Narvaez & Lapsley 2009). A person with a moral identity constructs his or her self around moral categories, beliefs, and convictions that are chronically accessible for interpreting the interpersonal landscape. Presumably, the sensibilities and preferences of moral people would be on display in their interactions in virtual or Web-based environments, so social scientists and those with political or economic interests have strong incentives to collect data on the moral foundation of the behavior of individuals. The availability of “Big Data” (vast archives of digital text, speech, and video, along with new analysis technology and inexpensive computation; Williford & Henry 2012) for psychologists, sociologists, and ethicists working on moral issues offers novel opportunities to investigate the connection between the moral identity of a person and their behavior. For example, one may investigate mismatches between digital reputation in networks designed for professionals (like LinkedIn) and failures to comply with (moral) norms in work environments; one may identify basic ethical orientations (Narvaez 2008) that have a predictive value for people’s behavior; or one may track people’s potentials for moral hypocrisy (e.g., cheating behavior) based on their behavior in social networks. As political opinions are related to specific “moral worldviews” (Haidt 2012), this knowledge may become a tool for understanding people’s political behavior (e.g., in political science), but might also open up new avenues for manipulation according to hidden agendas (e.g., by providing biased information on personalized news portals that are tailored to the moral psychology of the reader).

Big science projects such as the EU flagship proposal FuturICT ([www.futurICT.eu](http://www.futurICT.eu)) aim to perform large-scale social data mining to forecast socio-economic crises. As “the lack of data, the lack of computational power, and the lack of computationally tested institutional designs” (Helbing & Balmelli 2011, p. 4) have been identified as major obstacles to scientifically addressing socio-economic crises, it seems plausible that governments, corporations, and social scientists will end up collecting information on both the moral values people hold (in order to identify desired states of individuals, groups, or societies) and the moral competences they have (in order to evaluate whether desired system states are achievable). Thus, seen from this very general point of view, it is probable that moral values and competences will become a focus of data mining activities; in particular, if they refer to the understanding of socio-economic development.

The increasing interest in morality mining raises novel ethical issues that go beyond the current discussion of the ethics of data mining. We mention just a few examples: A first issue concerns the importance of knowledge of the moral competences and values people hold for predicting their behavior. A model of “moral intelligence” (Tanner & Christen, in press) backed up by data gained through morality mining may reveal information that is more liable to misuse than, for instance, consumer data. A second issue

relies on the importance of the information on the moral identity of individuals to the targets of investigation. On the one hand, people have an interest in promoting their moral identities (and moral reputations) in public, so they may be more willing to disclose personal information of that kind, which diminishes the privacy problem. On the other hand, the desirability of disclosing what is morally important is context-sensitive: most people presumably do not want to expose their moral values to all parties in all contexts. This point is of particular importance when models and tools relying on the results of morality mining may in the future be used to test people's moral competences to aid in, for instance, job assignments. A third issue refers to the possibility of obtaining information concerning the immoral behavior of individuals or institutions based on morality mining. This point does not refer to the disclosing of explicit information on immoral behavior of persons and institutions, which has become easier using Web-based instruments (e.g., WikiLeaks), but to the possibility of disclosing the mechanisms of "hidden" immoral behavior, of which the agent is not aware of. The ethical critique of behavior is then settled on a new level, not referring to acts an agent did consciously, but refers to behavioral patterns that are revealed by morality mining. One final issue has to do with what Hacking (1999, 2006) calls the "looping effect" of human kinds. People are unlike most other objects of investigation insofar as they consciously react to both the process and the product of investigation. They behave differently when and because they are observed, and they intentionally modify their behavioral patterns and dispositions in response to the theories, concepts, and knowledge that are created about them. By investigating their morality, researchers may inadvertently change their morality – and not necessarily for the better. These ethical issues related to morality mining, which will be outlined in more detail below, may require adaptations of current standards with respect to privacy and data misuse.

Our contribution is structured as follows. In the Background section, we briefly introduce major data mining methods and the security issues involved from a conceptual point of view. We also outline the current state of knowledge of the connection between moral identity and behavior in order to identify links between moral psychology and morality mining. In the next section, we introduce the concept of morality mining and explore it by outlining potential sources for morality mining and research questions that this new field of inquiry may help us to answer. We also identify the major ethical issues in morality mining and relate them to the current debate on the ethics of data mining. Finally, we provide an ethical framework for morality mining and show how it might contribute to an improvement of the ethical awareness of individuals. The chapter ends with an outline of future research directions and conclusions.

## **BACKGROUND**

### **The New Faces of (Social) Data Mining**

We start with some conceptual and technical explanations with respect to data mining in general and its applications – in particular with respect to medical issues, as technical and ethical problems that emerged in this field bare some similarity to the ones we can expect in morality mining. We first note that the term 'data mining' is inconsistently used in the literature. The technical literature draws a clear distinction between data mining and knowledge discovery in databases (KDD), considering the former to be solely a technical step of KDD, i.e. the application of specific algorithms for extracting patterns from data (Fayyad

et al. 1996). This distinction emphasizes that knowledge is the end product of a data-driven discovery and that KDD involves various different questions, including how the data are stored and accessed, how algorithms can be scaled to massive data sets and still run efficiently, how results can be interpreted and visualized, and how the overall man-machine interaction can usefully be modeled and supported. In this process, many decisions are made by the user, whereas data mining is understood basically as a step for which automation should be aspired as much as possible.

This distinction, however, is often blurred outside the community of statisticians and data analysts, where the terms ‘data mining’ and ‘knowledge discovery’ are used synonymously. Various publications that were consulted by us did not differ between these two terms. Helbing and Ballesteri (2011) for example, use the term ‘social data mining’ when referring to the objective of increasing the knowledge about social and economic systems. This may also result from the fact that in particular in the German literature, ‘data mining’ and ‘knowledge discovery’ are used synonymously, as the entry of ‘data mining’ in the German Wikipedia reveals (<http://de.wikipedia.org/wiki/Data-Mining>). We will also not draw a sharp distinction between data mining and KDD, i.e. we refer to the whole process of database-driven knowledge discovery when using the term ‘data mining’. This means in particular, that data mining involves human decisions of various kinds that guide, shape and inform the process – but the enormous amount of data that serves as “raw material” is an incentive to automate decision steps.

Next, we briefly outline the main processes of data mining (or knowledge discovery). Fayyad et al. (1996) distinguish several conceptual steps that lead from data to knowledge. They can be summarized as follows: 1) Selection of the target data based on an understanding of the application domain and the goals of the analysis; 2) preprocessing of the data, which involves data cleaning, noise reduction, deciding on strategies for handling missing data fields, etc.; 3) data transformation involving, e.g., dimensionality reduction or finding invariant representations of the data; 4) pattern discovery (or data mining in the technical sense) including selection of the appropriate analysis tools, parameters, and models, as well as exploratory analysis; and 5) evaluation, interpretation, and visualization of the result (knowledge). For the technical step 4), Marakas (1999) distinguishes four major categories of processes: clustering (the identification of sets of objects grouped together in virtue of their similarity or proximity to each other), classification (the discovery of rules about whether an item or an event belongs to a particular subset or class of data), association (linkage analysis of events that have a high probability of repetition), and sequential analysis (identification of trends that relate events over time).

This brief and purely conceptual overview, which is independent of the kind of data analyzed, already shows that there are many potential obstacles to data mining, including both “naïve data mining”, which applies pattern recognition tools and statistical methods without a deeper understanding of the underlying assumptions and implications of the results, and “information pollution”, which refers to (potentially unknown) systematic distortions of data sets, in particular if they emerge from the use of Internet services (Helbing & Ballesteri 2011). Such problems will certainly also emerge in morality mining – and they may even be more difficult to resolve, as there is no common understanding of the domain of application with respect to accurate models and mechanisms that explain moral behavior even among experts (this is discussed further in the section “Moral Identity and Moral Behavior”).

Both interest in and uses of data mining have increased substantially since the 1990s. One reason for this is the pervasive digitization of various processes in business and everyday life. Beside (public) registries such as phone books, tax data, etc., that were traditionally object of data mining, various new sources for data are available today (Helbing & Baliatti 2011). They include data generated by electronic services (e.g., money transactions, consumer data), data generated by Internet activities (e.g., service provider logs, search queries), data from portable devices (e.g., mobile phone data, GPS), user-generated content (e.g., in social networks, microblogging services), security data (e.g., surveillance systems, security forms), unauthorized content captured by multimedia devices (e.g., Google street view, public webcams), and stolen data resulting from (criminal) malware use. This new dimension of data availability in combination with today's enormous computing power not only substantially decreased the cost of data mining, it also offered novel opportunities. In particular, the notion of 'social data mining' (Helbing & Baliatti 2011) emerged, in which enormous amounts of data are collected in order to garner insights into social and economic processes.

Another potential issue is that of criminal abuse based on novel insights and information gained through data mining. Cyber criminals have repeatedly shown to be skilled in capitalizing on public information available on the Internet. In fact, already in the early days of the Internet, hackers have relied on rather technical information on network infrastructures, Web servers, and names of systems administrators etc. to mount their social engineering attacks by first carrying out a reconnaissance phase. With the advent of social networks the amount of information, and in particular of personal information, has dramatically increased. Such information has been successfully abused in the past, for instance to commit identity theft by answering so-called security questions (i.e., questions related to a users' biography) for password recovery. It is relatively easy to come up with hypothetical abuse scenarios of morality mining (e.g., ransom, advanced deception techniques), but hard to tell which will materialize and to what extent. Nevertheless, we believe that the risks are real, since cybercriminals regularly have shown that they are able to monetize information.

Most current applications of data mining are still on the "phenomenological level", i.e. they refer purely to behavioral data and their correlations with socio-demographic factors (e.g., whether specific income-classes have a specific shopping behavior), and usually do not involve information on biological or psychological behavioral mechanisms. Explaining behavior still relies on the psychology of agents. But this is changing, as the reliance on psychological constructs for making behavioral prediction is now supplemented with the data of behavioral genetics. Below, we will argue that data referring directly to people's moral values and competences may have similar, if not greater, predictive power and reveal more of the mechanisms of the behavior of individuals than, for instance, data on shopping preferences. Thus, we briefly discuss data mining in genetics and medicine in order to have a reference point for our discussion.

## **Data Mining in Genetics and Medicine**

Data mining is a core technology for modern genetics. Without sophisticated data mining tools, the Human Genome Project would have proceeded at a snail's pace. Today, biomedical research increasingly

uses methods from data mining, machine learning, and text mining to investigate, for example, disease comorbidities, patient stratification, drug interactions, and clinical outcomes (Jensen et al. 2012). In particular, clinical data describing the phenotypes and treatment of patients represents a data source that has much greater research potential than is currently realized. Mining of electronic health records may enable researchers to identify new patient stratification principles and reveal unknown disease correlations. Integrating such data with genetic data is also expected to furnish a more fine-grained understanding of genotype-phenotype relationships.

Also the “new” data sources the Internet offers are increasingly applied to data mining in medicine. Already in use is social data mining relying on search queries, blogs, micro-blogging and social networking sites to form coherent representations of real-time health events (Boulos et al. 2010). An example is the monitoring of seasonal influenza through Web and social media. At first sight, those applications seem unproblematic with respect to ethics, but if they also include mining personal information from social networking sites to characterize health behaviors (e.g., Facebook, LinkedIn) or from sites with user-provided health data (e.g., PatientsLikeMe.com), the situation changes (Vayena et al. 2012). In addition, when knowledge gained through mining genetic and health record databases is combined with the possibilities the new “Big Data” sources open up, an even more precise picture of individuals may emerge. So called “reality mining” (the collection and analysis of machine-sensed environmental data pertaining to human social behavior) may provide new opportunities with respect to diagnosis, patient and treatment monitoring, health services planning, surveillance of disease and risk factors, and public health investigation and disease control (Pentland et al. 2009).

This has well-known consequences for privacy, which is the right of persons to be in control of their own information. A technical solution for protecting privacy is to de-identify research data, which allows researchers to circumvent costly consent regimes, but the lack of identifiers makes certain types of population-wide research impossible, as other information cannot be linked to data subjects (Jensen et al. 2012). Moreover, despite such precautions, re-identification has been shown on some occasions to be a genuine risk, especially when data on human DNA are involved, as even a relatively small set of markers can enable unique identification. A second ethical issue refers to group profiling or stereotyping. Data-mining technology facilitates the construction and use of group profiles, which are based on non-obvious patterns and statistical correlations that link together both sensitive and “trivial” information about persons, e.g. with respect to disease probabilities. Tavani (2004) argues that research subjects may unwittingly contribute to the construction of controversial or unfairly stigmatized new categories or groups of individuals. We will discuss this point in more detail below.

Given this short overview, one should, however, not underestimate the effort required to combine data gained through various sources in order to get a complete picture. An illustrative example is neuroscience, as understanding the brain requires a broad range of approaches and methods from the domains of biology, psychology, chemistry, physics, and mathematics. Doing this demands the acquisition and integration of vast amounts of data of many types, at multiple scales in time and in space, and gained by many disciplines (Akil et al. 2011). Researchers face huge difficulties when combining these data. One reason among many is that investigators often use terminology or spatial coordinate systems customized to their

own particular analysis approaches. Thus, understanding the context and content of the data, and determining the conditions under which they can be compared to other data sets of interest, is a huge endeavor even within a single field. One has to keep in mind these technical obstacles when discussing ethical issues of novel data mining applications, in order to assess the risk associated with specific unwanted developments.

## **Moral Identity and Moral Behavior**

Before exploring the concept of “morality mining”, we briefly review the current research on moral identity based on Narvaez & Lapsley (2005), which has emerged as a central construct for understanding the link between moral decision-making and moral behavior. For almost two generations the dominant paradigm of moral development research focused on the ability of individuals to adjudicate hard-case moral dilemmas that focused on issues of fairness (Kohlberg 1984). Making a decision about what is fair and what one ought to do when there are conflicting claims to justice was held to be the primary challenge of moral deliberation. Kohlberg’s research team argued that moral reasoning undergoes a series of developmental transformations that moves the quality of reasoning towards increasing cognitive and moral sophistication.

But as everyone recognizes, knowing the right thing to do and doing it can come apart. What would motivate moral action? For Kohlberg’s cognitive developmental approach, it was knowledge that a situation falls under the covering law of a moral rule, and moral rules are auto-motivating because of their prescriptive character. An alternative view emerged, however, that implicated the role of selfhood and identity in translating moral prescription into moral action. A moral self, or a person with moral identity, constructs the self around moral categories, beliefs, and convictions that are chronically accessible for interpreting the interpersonal landscape. A moral self cares about morality and identifies with it. Morality is essential, central, and important to self-understanding, and one is motivated to reduce any discrepancy between self-identity and actual behavior.

Hence the study of moral identity is a dominant focus of contemporary moral psychology research. It not only holds promise for better understanding the dynamics of moral behavior, but it also connects the study of moral cognition with constructs and theories of personality science and cognitive psychology, and introduces new possibilities for understanding moral behavior. It encourages reflection, for example, on emergent literatures that map out how cognitive automaticity influences moral judgment. It encourages integrative study of how trait psychology and social-cognitive approaches to personality underwrite moral selfhood and commitment to morality.

Such studies, as well as psychological studies in general, make increasingly greater use of the new opportunities provided by the pervasive use of information technology and the Internet – although mostly still in a “classical sense” such as large online surveys (e.g., [www.yourmorals.org](http://www.yourmorals.org)) and Web-based experiments. For example, the “Web Experiment List” and the “Web Survey List” (<http://wexlist.net>), two free Web services for researchers in psychology, help in the recruitment of participants and in the archiving of

studies. Also services like “Amazon Mechanical Turk” (<https://www.mturk.com>) are increasingly used to address a large and diverse subject pool.

These examples are not data mining applications – but this is about to change, as text data available on services like Facebook and Twitter allows for data mining with respect to psychological research questions. An example is iScience Maps (Reips & Garaizar 2011), a free Web service for researchers (<http://maps.iscience.deusto.es> and <http://tweetminer.eu>). This service allows researchers to assess via Twitter the effect of specific events in different places as they are happening and to make comparisons between cities, regions, and countries regarding psychological states and their evolution in the course of an event. A recent review of psychological studies based on Facebook data (Wilson et al. 2012) indicates not only a substantial increase of this kind of research, but also a shift in focus away from mere descriptive research questions (e.g., who is using Facebook?) to research focusing on explaining social behavior based on this data. Probably even more important is the growing use of smartphone apps for psychological research (Miller 2012). As they allow researchers to collect data of various modalities (visual, auditory, movement, location), they promise to generate data that may help to explain human behavior in real world-environments. Thus, we believe it’s only a matter of time until human moral behavior will become a new focus of data mining activities.

The recent case of David Petraeus, the disgraced former head of the United States Central Intelligence Agency (CIA), illustrates both the prospects and the pitfalls of such data mining. Petraeus apparently conducted an illicit affair, which was stumbled upon by FBI investigators and eventually divulged to the public. What began as an open-ended investigation of a few anonymous emails led, through the vast quantity of interconnected data available, to the embarrassing revelation that induced Petraeus to resign his post. It’s not hard to imagine that data mining by researchers rather than federal investigators could end up unearthing similarly titillating but also personally devastating connections.

## **MORALITY MINING**

### **A Definition of Morality Mining**

Morality is an important but complex determinant of human behavior. It involves norms, values, and virtues that have both a natural and cultural history and that serve to guide people with respect to right and wrong. Much of this guidance works in a quite automated way during our everyday live (Gigerenzer & Gaissmaier 2011), and ethical deliberation is probably the exception in daily routine decisions (Haidt 2001). Furthermore, humans may act contrary to their own values (due to systematic limitations in decision making, among many other things; Hastie & Dawes 2010), or may be hypocritical (i.e., pretend to be moral but avoid the costs of being moral when being confronted with temptations; Batson et al. 2002).

Nevertheless, human moral behavior always involves at least the demand to provide justifications for what has been done or what should be achieved – and people express their moral convictions both with respect to behaviors (e.g., abstinence from meat) and expressions (verbal, written). In a world character-



ized by the pervasive use of information technology, our moral behavior leaves traces of various kinds: such as:

- Opinions on moral issues expressed in e-mails, blogs, comments on news portals, tweets, etc.
- Information on affiliations to groups that stand for specified moral values (e.g., with respect to animal protection, abortion, global justice, etc.) that is sometimes explicit (membership to organizations) and sometimes implicit (revealed through analyzing the interaction partners of a person).
- Information on psychological abilities required to uphold moral behavior (e.g., motivation to resist temptations in shopping sites, moral sensibility with respect to reading behavior on news sites – which articles are overlooked –, etc.).
- Traces of behaviors that disclose the “immorality” of a person (e.g., with respect to cheating).

Accessing this kind of data allows researchers to construct knowledge of both the moral values and convictions a person (or a group of people) holds and the psychological competences of this person (or groups) – i.e. the “moral intelligence” (Tanner & Christen, in press) of a person. This knowledge has the following characteristics:

- **Relevance:** It concerns an aspect of a person that the person considers to be a decisive element of his or her self-understanding. It is knowledge that characterizes a person in a deeper sense than, for instance, his or her consumer preferences.
- **Context-relevant disclosure:** Because of its importance, the person has a motivation to present this knowledge to third persons, but this motivation is context-dependent (for example, someone may not want to disclose his or her position with respect to abortion in all situations).
- **Partially hidden:** This knowledge is not necessarily consciously known by the person him or herself, e.g., when the knowledge reveals that the current moral self-understanding is actually not expressed in the moral behavior of the person.
- **Character-enhancing:** The knowledge has the potential to change a person’s character; that is, it may provide insight for the person that improves self-understanding and may help to change the convictions a person holds.
- **Potentially stigmatizing:** This knowledge has an exceptionally strong stigmatization potential in the cases where it reveals an aspect of the person that is (considered) morally wrong. And compared to other potentially stigmatizing information (e.g. with respect to gender, race), the stigmatization effect may be justified and even approved by society.

It is the *combination* of these elements that makes knowledge gained through morality mining unique compared to other kinds of knowledge that results from data mining. Health-related knowledge is probably the closest neighbor. For example, knowledge that person X will develop a disease with a certain probability due to his or her genetic profile and behavioral habits is indeed relevant for X, was hidden to X, may have the potential to stigmatize X, and X would probably prefer to disclose this information only in specified contexts (e.g., when talking to his or her doctor). It may even have a character-enhancing potential, if changes in behavioral habits may reduce the probability of developing a disease or temper its symptoms.

To summarize, morality mining is defined as data mining of information that reflects the moral identity of a person and that generates knowledge that is relevant for a person, that may be hidden, character-enhancing and/or stigmatizing for the person, and that a person only wants to disclose in specified contexts.

## **Ethical Issues in Morality Mining**

The ethical discussion of data mining in general usually refers to the value of “privacy” – a cluster concept that unifies a number of moral considerations in support of data protection. According to most scholars, this involves three types of privacy (Tavani 2004). *Accessibility privacy* is freedom from intrusion. Historically, this is the first notion of privacy to be codified into law. The Fourth Amendment to the U.S. Constitution protects citizens from unreasonable searches and seizures. *Decisional privacy* is freedom from interference in your personal choices. *Informational privacy* is a person’s ability to restrict access to and to control the flow of his or her private information. But there are also other problematic issues with ethical implications, e.g. the costs of obtaining and protecting information, the possibility that incorrect conclusions are drawn from the data, or that the data can be used for other than the original purposes for which they were collected (Payne & Trumbach 2009).

However, despite the prominence of privacy in the ethical debate on information and communication technology (ICT) applications (Van der Hoven et al. 2012), it is surprising that users often underestimate risks of their information privacy on, for instance, online social network sites, and that younger users in particular are much less concerned about potential privacy threats than adults (Hugl 2011). There seems to be a gap between the emphasis on privacy as a guiding principle on the one hand, and the actual expression of privacy concerns by users in their everyday behavior on the other hand. This may also explain why most textbooks on data mining deal insufficiently with ethics, and one has to suspect that unless individual instructors make an effort to discuss ethical issues in their courses, students will not be exposed to them through their textbooks (Lawler & Molluzzo 2006). However, we can expect that as people become aware of the threats inherent in social networking, such as the potential for bullying, an increasing awareness of privacy and data protection also from the side of the users and producers of ICT services – and we expect that an increasing interest in morality mining will support this trend.

Following Van der Hoven et al. (in press), ethical concerns of data mining gain their ethical relevance through their relation to the following moral topics: (a) the prevention of harm to data-subjects, (b) informational justice and non-discrimination, (c) respect for moral autonomy, (d) information inequality and fairness in markets for personal data. We will briefly introduce those topics and explore the extent to which morality mining specifically addresses them. Finally, we will discuss a novel field of ethical concern that is uniquely related to morality mining.

The *avoidance of harm* is a basic orientation within ethics that has various occurrences in practical fields (e.g., in medical ethics, it relates to the principle of non-maleficence; Beauchamp & Childress 2009). In an information society, people can be harmed in many ways on the basis of the information that is availa-

ble about them. Obvious examples include information that allows illegitimate access to bank accounts, identity theft, or blackmailing using compromising information. However, the abuse of such information rarely involves the use of elaborate data mining tools and procedures since it is mostly predicated on “raw data” rather than patterns or clusters gained through data mining (Zarsky 2002). An important exception concerns mistakes in the data mining process that may, for example, result in fallacious personal credit reports. This could result in the inability to secure a loan, open a bank account, or receive a charge card. Another exception may be nuisances such as tailored advertisements and the like, resulting from data mining applications.

The knowledge produced by morality mining certainly has the potential to cause direct harm. Obvious examples emerge when the information refers to “immoral” behavior and when one is able to relate this knowledge to an individual person. Knowing that person X has, e.g., a tendency to cheat is information that can compromise X. Such examples are, however, not completely comparable to “classical” bullying or blackmailing, as one may argue that disclosing immoral behavior of person X is actually desirable. For example, several States in the U.S. publish the domicile of convicted child molesters, thus valuing the right to be informed of the presence of “immorality” higher than the right of privacy of the individual. We do not discuss whether this specific case is ethically justified, but we suppose that mining data of moral behavior of people may increase such tendencies of valuing “transparency” higher than privacy also with respect to moral violations that are more controversial. For example, some communities may want to disclose the “moral status” of a person with respect to topics like abortion, gay-rights, or capital punishment and may use this information to harm this person. Beside such obvious cases, the issue of harm may also arise *without* relating mined information to specified individuals, although the effects may be ambiguous. We will discuss this in more detail below when outlining the looping effect.

Closely related to harm avoidance is the issue of *information justice and non-discrimination*. This topic has a practical and a more fundamental aspect in the context of data mining. The practical aspect refers to the issue of group profiling and the so-called “inference problem”, where sensitive information is derived from non-sensitive data and meta-data (Farkas & Jajodia 2002). Information gathered in data mining is usually implicit in patterns in the data. These patterns suggest new associations about people, which place them into new categories. Such group profiles may be valid for the group and for individuals as members of that group, though not for individuals as such. When individuals are judged by group characteristics they do not possess as individuals, this may strongly influence the advantages and disadvantages of using group profiles (Custers 2003). If an individual is being judged based on information that was wrongly ascribed to him, most legal systems provide opportunities to have the information changed or deleted, possibly combined with compensation of damages. But group data is often anonymous data which does not fall under data protection laws. Besides that, most people are unaware of the group profiles they are being judged by. This could result in what Appiah (2011) calls ‘probabilistic harm’ (the harming of people by decreasing their chances of getting some good).

The specific ethical problem with respect to group profiling is not the creation of group profiles *per se*, as not all practices of discrimination are *prima facie* forbidden and automatized data mining can also create “fairer” results (Zarsky 2002). For example, insurance companies use actuarial tables to determine which

socio-demographic groups have a higher incidence of risk, when developing criteria for eligibility for insurance coverage. It is well known that in the U.S. (and elsewhere) teenage boys are placed in a high-risk category for automobile insurance because of the statistical rate of automobile accidents involving male teenage drivers. However, this is a fairly well known correlation that is considered to be legitimate knowledge for defining insurance costs. Furthermore, this policy is transparent and thus open to public debate. But many recent correlations discovered by data mining, based on hitherto unnoticed (and often-times non-obvious) relationships between characteristics and features of persons are not transparent – in particular, if they suggest “new facts” about individuals (Tavani 2004b). In addition, Vedder (2001) also notes that data in group profiles are often used as if they were personal data (even though they are not), which increases the practical problem related to group profiling and the inference problem. Finally, one has to take into account that the process of knowledge discovery involves human decisions that may (sub-consciously) bias profiling. For example, in video surveillance the “unequal gaze” problem has arisen (Zarsky 2002). There, it has been claimed that when surveillance tools are controlled manually, examining their recordings leads at times to the finding that the surveying device is not gazing evenly, but tends to focus on minorities, even when these individuals are not exhibiting suspicious behavior. The inevitable result of such unequal gazing is a biased increase in the incidence of events involving minorities, as they are the people who are constantly being monitored.

All these problems are expected to play a role in morality mining, too. For example, a group profile describing a “coherent moral profile” may emerge, although no single person of this group actually fulfills all these properties. One could imagine that research reveals a typical “Kantian” pattern of moral identity that goes along with values that are understood in a specific way (e.g., autonomy) and moral emotions (e.g., shame, guilt) that are expressed in a characteristic way. A person may then be classified as “Kantian” and implicitly attributed with specific moral emotions, although the individual may actually not have them. Current research using neuroscientific methods have a tendency to create such group profiles (e.g., Greene 2008). In combination with (disputed) claims that some of these moral identities may be related to morally fallacious argumentative patterns (Singer 2005), such group profiles may have a discriminatory effect. An interesting point to consider would be, whether such group profiles would have a tendency to mask actual controversies in ethics. One could also speculate that ethical positions held by researchers may have a similar effect on decisions (that are inevitable in the process of data mining) in the process of creating knowledge compared to the unequal gaze problem described above. This would be a point to consider when assessing the credibility of, e.g., moral group profiles.

The more fundamental aspect with respect to information justice refers to the work of Walzer (1983), who outlined that justice considerations are “sphere dependent”, i.e. the definition of goods and the appropriateness of a distribution mechanism are restricted to specific spheres. This analysis also applies to information (Van der Hoven et al., in press): the meaning and value of specified information is local, and allocation schemes and local practices that distribute access to information should accommodate those local meanings and should therefore be associated with specific spheres. For example, people usually do not object to the use of their personal medical data for medical purposes – for their own personal health affairs, but often also in relation to their families or communities (e.g. with respect to epidemiological research). However (set aside some well-defined exceptions), they do object to their medical data being

used to classify them or disadvantage them socio-economically. Thus, Walzer's framework translated to the information domain may actually explain why we consider specific cases of information misuse as "statistical harm". When translating this picture into the realm of morality mining, however, the problem may emerge that the knowledge mined seems not necessarily to be bound to a specific sphere. Psychologists have identified "protected values" (Tanner et al. 2008) that people uphold independent of the context. The motivation to disclose or not to disclose such values is probably not dependent of the sphere the person is in, but dependent of the detailed context even within the same sphere, i.e. an analysis along Walzer's work may not be helpful in all cases.

*Autonomy* refers to the human capacity to shape our own moral biographies, to present ourselves in a way that fits our self-understanding, to reflect on our moral careers, and to evaluate and identify with our moral choices without the critical gaze and interference of others and without pressure to conform to the 'normal' or socially desired identities. In his analysis of shame and privacy, Velleman (2006) draws attention to self-presentation as a constitutive feature of moral persons. People want to outline their moral personality, but experience the normative pressures that public opinion and moral judgments of others impose. This is why moral knowledge is context sensitive, because when this knowledge about a person becomes available, it facilitates the formation of beliefs and judgments about the person. Context is decisive with respect to the effect of this information. If, for example, a person favoring the "pro-choice" position in the abortion controversy is invited to a church service where he recognizes that abortion is seen as the ultimate sin by the other participants, he probably would not actively promote his position in that context, knowing that his position could lead to immediate stereotyping, i.e., the other participants would attribute unfavorable characteristics to him. Thus, disclosing such information in these contexts is not only an issue of potentially harming someone, it also undermines the self-presentation of a person with respect to a central aspect of his or her personality, namely the moral identity of the person.

Respect for autonomy concerns not only control over one's own moral profile, but also the capacities needed for autonomy (Zarsky 2002). If others know information about the habits and behavior of a person, they can more easily manipulate him or her. Take the example of a cheating husband, who uses a commercial dating site for that purpose, but finally decides to quit this behavior, because he comes to the conclusion that this is not fair vis-à-vis his wife. Changes in the surfing behavior of the person captured by a tracking cookie of the dating site, however, provides a "warning signal" to the site that a customer may be lost – and the site operators send him targeted offers involving verbal and visual stimuli that have been previously determined to be most effective in getting him to visit their website. Given the known difficulties to uphold motivation to act upon one's own moral values, morality mining may indeed help to undermine autonomy.

A more subtle effect with respect to personal autonomy relates to the increasingly common practice of tailored content generation, which generates a "vicious circle" that is according to Zarsky (2002) described as follows: (a) Individuals inform the information providers which types of knowledge and information they are interested in and provide (both implicitly and explicitly) personal information as about their dispositions and interests; (b) The content providers supply individuals with specific information "tailored" to the needs of every person, according to each provider's specific strategy, and chosen on the

basis of the personal information previously collected; (c) The individuals require additional information. This time, however, the request is affected by the information previously provided; (d) Again, the information providers supply information, in accordance with their policies and discretion. Pariser (2011) describes this phenomenon in which a website algorithm selectively guesses what information a user would like to see based on information about the user as “filter bubble”. Such bubbles separate users from information that disagrees with their viewpoints, effectively isolating them in their own cultural or ideological bubbles. With respect to morality mining, one may speculate that such processes will play a role as well and enhance differences between ethical positions. One then could expect that the process of finding a compromise in actual moral conflicts will be complicated, as the number of “contact points” between such position decreases.

The issues of *information inequality and fairness* relate to the increasing importance of data-mined knowledge for business, for example in mass media communication and advertising. As ICT is pervasive in business relations, consumers start to realize that every time they buy something, they give away something more than just money: the information about their purchase or transaction. The particularly high valuation of some Internet and ICT companies on the financial markets relies mainly on the assumption that they have access to a unique source of knowledge accessible through data mining that allows understanding customers in a much more precise way than previously. Given the fact that collecting this information is bound to the use of a costly infrastructure provided by these companies (servers, storage) as well as complex tools to mine the data, a considerable “information inequality” builds up. Ironically, the Internet engenders both this information inequality and potential for injustice in parallel with the well-known “information equalization effect,” which refers to the fact that Internet users have vast informational resources at their fingertips that were unavailable even two decades ago. This information inequality results at least partially from the fact that (Internet) users provide information without knowing the implications of what they are consenting to when they, for instance, sign a contract or agree to terms of use. An additional aspect comes into play when taking into account that this knowledge may actually be useful for addressing social problems, which leads to the notion that this kind of knowledge may even be considered a public good to which many more people and institutions than just a few companies should have access (Van der Hoven et al., in press). Knowledge gained through morality mining could indeed be of this type, as it may increase our understanding of the society. But it is actually questionable whether the effect of making such knowledge a “public good” that is available to everyone would indeed have positive effects – a point which we will discuss next.

Thus, beside these moral topics where morality mining already may have unique implications, we believe that morality mining involves an additional ethical issue that is related to the unique nature of moral knowledge: It may yield insights about a person – sometimes to that very person – that have the potential to change this person. The ideal of “becoming a morally better person” is probably among the oldest topics in practical philosophy and pedagogy, and an immense literature deals with it. But morality mining may reveal a novel source for understanding oneself, as the possibility to trace one’s own behavior in time both in virtual environments as well as related to the use of electronic devices allows an unprecedented access to one’s own history. The timeline function of Facebook could be understood as a first step towards “personalized” data mining intended to uncover hidden aspects of one’s own behavior with moral

consequences. Doing this may actually be a desideratum of many people – and raises the question of the kind, quality, and use of this knowledge.

An analogy may help clarify this point. For almost a decade, the Web-based vote advice application Smartvote (<http://www.smartvote.ch>) has been established in Switzerland and has become a widely used service in elections. This tool is based on the idea of preference matching; i.e., any Smartvote user (voter) may answer the same set of questions as the candidates and then gets a list of candidates that indicates the political distance between the user and the candidates. Its output is a “political profile” on basic political issues like “environment”, “migration”, and “social security” that conveys an easy and catchy visual message on what kind of “political person” both the candidates is as well as the voter himself. This visual message has even become a standard in communicating political issues in the media and frames the political discourse. This surely has positive effects as it allows for a more transparent choice in elections, unbiased by campaign rhetoric and propaganda – but it also involves the question of how these profiles are generated, and whether they really refer to what is on stake in an election. One could now imagine creating an analogous instrument for “moral profiling”, i.e. an instrument that visually represents in a simple way the “moral identity” of a person. This may indeed be a communication tool one would like to use, e.g. in social networks, to find “appropriate” friends. It might also be a tool that creates moral insight. But similar questions both with respect to the societal effects of such instruments as well as the methodology behind creating such messages will emerge.

Furthermore, one has to question the epistemological status of this kind of knowledge. Morality involves a tension between “is” and “ought” – some of the knowledge that characterizes a moral person refers to what a person actually does and is, and some to what the person would like to be and do. This tension is not necessarily of a hypocritical kind (i.e., the person just does not do what she considers to be right, although she upholds the image of following this value), but may reflect the struggle every moral person has when trying to become a “better person”. Thus, the question emerges in what way knowledge that is gained through morality mining actually represents the morality of the person. It may make sense to distinguish between knowledge that refers to the moral reputation of a person and the knowledge that this reputation is not in all situations actually uphold.

Besides this ambivalence on the individual level, we may also expect seemingly contradicting effects for the case, where knowledge gained through morality mining is not related to an individual person. Many, if not most people are conditionally cooperative (Bicchieri 2005). That is, they will do the right thing, but only if they think enough other people are doing it too. Aggregate data from morality mining could influence the future behavior of conditionally cooperative people in two directions. If they currently think that enough others are doing the right thing (e.g. paying their taxes), but they find out that this assumption is false, they may cease to do the right thing themselves. But if they currently think that not enough others are doing the right thing, then find out that this assumption is false, they might actually improve their behavior.

This is just one example of how morality mining might trigger what Hacking (1999, 2006) calls the “looping effect” for human kinds. Unlike, say, molecules, when humans are the object of investigation

they consciously react to both the process and the product of investigation. If people learn at time  $t$  that the product of investigation is that 80% of people shirk their tax duties, then the 20% of people who currently do pay taxes are likely to stop paying taxes after  $t$ . But if they learn that 90% of people do pay their taxes, perhaps some of the 10% who previously had not would start to. Thus, the product of investigation (the statistical knowledge, theories, and concepts generated by investigation) influences the object of investigation. What we think we know about people changes as they learn what we think we know about them.

A nice example of this phenomenon is the Goldstein, Cialdini, and Griskevicius's (2008) investigation of social proof. They found that guests at a hotel were more likely to reuse their towels if they read a message that said that 75% of the guests at that hotel reused their towels than if they read a message that merely exhorted them to do the right thing. Moreover, they were even more likely to reuse their towels if they read that 75% of the guests *in their very room* had reused their towels than if they were merely told that 75% of the guests in the hotel did so.

Another example of the looping effect, this time in response to the process of investigation and monitoring, is discussed in more detail in Alfano (2013): people are more disposed to share resources when they feel that they are being watched or monitored by a poster that depicts a face than when they make the same decision in the presence of a poster that depicts flowers. They are also less inclined to steal and less inclined to litter when they feel watched in this way. If such subtle manipulations of the feeling of being watched can modulate how people behave, it stands to reason that morality mining – which, though subtle, is sure to be noticed – may produce similar effects. We need to bear in mind such effects when we think about morality mining, as they impinge on both its epistemology and its morality. Results in morality mining could end up being quite frail if the looping effect continually changes the behavior of the objects of investigation, and there is simply no guarantee that those changes will always be for the best.

## **Solutions and Recommendations**

So far, we have outlined that knowledge gained through morality mining raises ethical issues that are partially captured by the current ethical discussion with respect to data mining (although sometimes in a unique way) – but also involve novel problems. The first insight allows for a straightforward recommendation – namely that the various techniques in order to enhance privacy by design that are currently developed (Van der Hoven et al., in press) should also be applied in morality mining. As several of the problematic points that were outlined so far result from the connection of mined moral knowledge to the individual person, a particular emphasis would be to prevent such a connection. Another point to mention is that some data mining applications can actually have a clearly positive effect, e.g. as a powerful aid to the anti-discrimination analyst, capable of automatically discovering the patterns of discrimination that emerge from the available data with stronger evidence (Ruggieri et al. 2010).

There are, however, also recommendations that go beyond the current discussion on enhancing privacy through technological means. They result from the special status of moral knowledge as a potential instrument to enhance people and societies. For example, knowing that person X does not necessarily fol-



low its own moral convictions can mean several things: First, it can refer to the possibility that X is a hypocrite, i.e. deliberately does not follow own convictions. The knowledge could then be used to disclose a hypocrite – although it is not clear whether this strategy is actually optimal to diminish immoral or hypocrite behavior in a society (Christen, submitted). Second, it can also express a tension the person is well aware of and he or she is actually trying to overcome (i.e. the person is in the process of moral change), and is hindered in doing this by a lack of psychological competences required to follow his or her own moral convictions. Then, the use of this knowledge would be different, as it may motivate the person to enhance the necessary psychological competences (e.g., willpower). Third, this knowledge can also express the fact that the person is actually not aware of the fact that he or she is acting against his own convictions. The reason may be that the person is uninformed about the explication of an abstract value to a specific situation, or that he or she is unaware that a specific act actually violates this conviction. The use of this knowledge would then be to generate “ethical awareness” or “moral mindfulness” (DesAutels 2004), i.e. a deliberate insight into a gap between one’s “is” and “ought”. Indeed, it might turn out that the lack of moral mindfulness leads many otherwise good people to do what they themselves would not reflectively endorse. Their bad actions would then be revealed not as motivated by ill-will but as enabled by ignorance or lack of attention. This might not expunge the badness of the actions completely, but it would go some way towards helping people to understand themselves and each other.

In that sense, one could also construe an argument in favor of morality mining, namely as an obligation to use the novel data sources in order to increase the ethical insight of individuals (e.g., by providing tools as an analogue of “smartvote” that gives insights about political values a person holds) or societies. We sometimes lack obligations because we lack relevant knowledge: you have no obligation to help your choking neighbor if you do not know that he is choking and have no way of finding out. But in some cases this point leads to a second, which is the obligation to inquire or learn, to acquire moral knowledge. You might not have an obligation to improve your character if you have no way of knowing that it is defective, but you might have an obligation to find out whether and how it is defective, and other people – the morality miners – might have at least a duty of imperfect obligation to help you.

Finally, we indicate that the generation of tools to increase the awareness on a person’s own morality (not necessarily involving morality mining) may also be accompanied with practical ethical issues. For example, some companies may promote the use of such tools in job assessments in order to ensure that the “right” person will get a job. This leads to questions like the credibility of such results, the stigmatization potential of “bad test results” or a violation of moral autonomy of persons. Although we cannot discuss these points further, we remind that an increasing use of morality mining may help to promote such practical applications and requires a sensibility for such practical ethical problems.

## **FUTURE RESEARCH DIRECTIONS**

The project of “morality mining” is embedded in two major scientific and technological developments that characterize the conditions of human moral behavior in the modern world: On the one hand, findings in moral psychology, neuroscience, anthropology and other fields increasingly outline the limitations of human moral agency. Personal development, neuronal conditions, and social context frame moral actions,

leading to questions about the limitations of human responsibility. On the other hand, modern life with its pervasive use of information and communication technology creates an unprecedented amount of information on a person's (moral) actions through "electronic traces" of various kinds, allowing – in principle – in-depth inquiry into individuals' personal morality. Tools are emerging out of science and technology that enable us to understand our own moral behavior, and perhaps even to enhance it. This enhancement is not pharmacological but technological, generated as it is by gaining insights facilitated by technological instruments, and the use of them to improve moral behavior as well as to create optimal conditions for moral flourishing of, e.g., one's children.

Thus, science and technology create knowledge both with respect to limitations as well as potential improvements of human moral behavior. Future research must therefore involve both of these directions and should try to find connections between them. This involves in particular answers to the following questions:

- 1) To what extent does the use of tools provided by information technology shape and change the moral psychology of individuals?
- 2) To what extent is "morality mining" able to make solid and reliable statements about the moral identity of individuals. What are the methodological pitfalls in this endeavor?
- 3) What are the potential unwanted side-effects of morality mining, i.e. how exactly could a negative "looping effect" be prevented?
- 4) What would instruments that increase the ethical awareness of individuals and institutions look like? What information should they provide in order to be useful?
- 5) To what extent are moral dispositions properties of individual people versus relations between an individual and a social network?
- 6) Is self-knowledge positively and linearly correlated with moral behavior, or are there instead limits to the benefits to be had from moral self-knowledge?

## **CONCLUSIONS**

In summary, we suggest that morality mining will become an important aspect in future social data mining applications that raises particular issues. Those issues will not only require being included in current work on privacy by design. They will also require research in order to create novel instruments allowing for ethical awareness generation.

Finally, we remind that this undertaking we call morality mining is still in its infancy. Although large research ventures like FuturICT are on their way that probably will include mining of data on people's moral identities, there are – at least to our knowledge – not yet any real world applications. Also large companies like Google or Facebook, who probably have access to data allowing for morality mining, probably are not focusing on this aspect, unless it has a commercial application (e.g., helps to tailor advertising). Finally, reconstructing moral identities from mined data will not be an easy undertaking and faces similar challenges of current approaches to combine various databases. But the emerging science of social data mining will, to our understanding, take the moral behavior of people increasingly into its focus, as

there is a deep motivation to understand the moral foundation of our society in order to understand and shape the socio-economic development of our society.

## REFERENCES

Akil, H., Martone, M. E., Van Essen, D. C. (2011). Challenges and Opportunities in Mining Neuroscience Data. *Science*, 331, 708-712.

Alfano, M. (2013). *Character as Moral Fiction*. Cambridge: Cambridge University Press.

Appiah, A. (2011). "Group Rights" and Racial Affirmative Action. *The Journal of Ethics*, 15(3), 265-280.

Batson, C.D., Thompson, E.R., & Chen, H. (2002). Moral hypocrisy: addressing some alternatives. *Journal of Personality and Social Psychology*, 83(2), 330-339.

Beauchamp, T. L., & Childress, J. F. (2009). *The principles of biomedical ethics* (6. ed.). Oxford: Oxford University Press.

Bicchieri, C. (2005). *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge: Cambridge University Press.

Boulos, M.N.K., Sanfilippo, A.P., Corley, C.D., Wheeler S. (2010). Social Web mining and exploitation for serious applications: Technosocial Predictive Analytics and related technologies for public health, environmental and national security surveillance. *Computer Methods and Programs in Biomedicine*, 100, 16-23.

Christen, M. (submitted): Overcoming Moral Hypocrisy in a Virtual Society. In: Complexity and Human Experiences. Pan Stanford Publishing.

Custers, B., Zarsky, T., Schermer, B., & Calders, T. (Eds.) (2012/forthcoming). *Discrimination and Privacy in the Information Society: Effects of Data Mining and Profiling Large Databases*. Berlin/Heidelberg: Springer.

Custers, B. (2003). Effects of unreliable group profiling by means of data mining. In G. Grieser, Y. Tanaka, A. Yamamoto (eds.): *Discovery Science. Lecture Notes in Artificial Intelligence Vol. 2843* (pp. 291-296). Berlin: Springer.

DesAutels, P. (2004). Moral mindfulness. In DesAutels & Walker (eds.), *Moral Psychology*, pp. 69-82. New York: Rowman & Littlefield.

Farkas, C., & Jajodia, S. (2002): The Inference Problem: A Survey. *SIGKDD Explorations*, 4(2), 6-11.

Fayyad, U.M., Piatetsky-Shapiro, G., & Uthurusamy, R. (1996). *Advances in Knowledge Discovery and Data Mining*. Cambridge MA: MIT Press.

Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic Decision Making. *Annual Reviews in Psychology*, 62, 451-482.

Goldstein, N., Cialdini, R., & Griskevicius, V. (2008). A room with a viewpoint: Using social norms to motivate environmental conservation in hotels. *Journal of Consumer Research*, 35.

Greene, J.D. (2008). The Secret Joke of Kant's Soul. In: W. Sinnott-Armstrong (Ed.): *Moral Psychology*, Vol. 3 (pp. 35-79). Cambridge: MIT Press.

Hacking, I. (1999). *The Social Construction of What?* Harvard University Press.

Hacking, I. (2006). Making up people. *London Review of Books*, 28:16.

Haidt, J. (2012). *The Righteous Mind: Why Good People are Divided by Politics and Religion*. London: Penguin Books.

Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108, 814-834

Hastie, R.K., & Dawes, R.M. (2009). *Rational Choice in an Uncertain World: The Psychology of Judgment and Decision Making*, 2nd ed. Thousand Oaks, CA: Sage.

Helbing, D., & Baliotti, S. (2011). From social data mining to forecasting socio-economic crises. *European Physics Journal – Special Topics*, 195, 3-68.

Hugl, U. (2011). Reviewing person's value of privacy of online social networking. *Internet Research* 21(4), 384-407.

Jensen, P.B., Jensen, L.J., & Brunak, S. (2012). Mining electronic health records: towards better research applications and clinical care. *Nature Reviews – Genetics*, 13, 395-405.

Kohlberg, L. (1984): *The Psychology of Moral Development: The Nature and Validity of Moral Stages (Essays on Moral Development, Volume 2)*. San Francisco: Harper & Row.

Lawler, J., Molluzzo, J.C. (2005). A Study of Data Mining and Information Ethics in Information Systems Curricula. *Information Systems Education Journal*, 4(34), 3-15.

- Marakas, G.M. (1999). *Decision Support Systems in the 21st Century*. Upper Saddle River, NJ: Prentice Hall.
- Miller, G. (2012). The Smartphone Psychology Manifesto. *Perspectives on Psychological Science*, 7, 221-237.
- Narvaez, D. (2008). Triune Ethics: The Neurobiological Roots of Our Multiple Moralities. *New Ideas in Psychology*, 26, 95-119.
- Narvaez, D., & Lapsley, D.K. (Eds.) (2009). *Personality, Identity, and Character*. Cambridge: Cambridge University Press.
- Nisbet, R., Elder, J., & Miner, G. (2009). *Handbook of Statistical Analysis and Data Mining Applications*. Burlington MA: Academic Press.
- Pariser, E. (2011). *The Filter Bubble. What the Internet Is Hiding from You*. New York: Penguin Press.
- Payne, D., & Trumbach, C.C. (2009). Data mining: proprietary rights, people and proposals. *Business Ethics: A European Review*, 18(3), 241-252.
- Pentland, A., Lazer, D., Brewer, D., & Heibeck, T. (2009). Using reality mining to improve public health and medicine. *Studies in health technology and informatics*, 149, 93-102.
- Reips, U.-D., & Garaizar P. (2011). Mining twitter: A source for psychological wisdom of the crowds. *Behavioral Research*, 43, 635-642.
- Ruggieri, S., Pedreschi, D., & Turini F. (2010). Data Mining for Discrimination Discovery. *ACM Transactions on Knowledge Discovery from Data*, 4(2), Article 9.
- Singer, P. (2005): Ethics and Intuitions. *The Journal of Ethics*, 9, 331-352.
- Tanner, C., & Christen, M. (in press): Moral Intelligence – A Framework for Understanding Moral Competences. In: M. Christen, J. Fischer, M. Huppenbauer, C. Tanner, & C. van Schaik (Eds.): *Empirically Informed Ethics*. Library of Ethics and Applied Philosophy. New York: Springer.
- Tanner, C., Medin, D.L., Iliev, R. (2008): Influence of deontological vs. consequentialist orientations on act choices and framing effects: When principles are more important than consequences. *European Journal of Social Psychology*, 38: 757-769
- Tavani, H.T. (2004): *Ethics and Technology: Ethical Issues in an Age of Information and Communication Technology*. Hoboken, New Jersey: John Wiley and Sons.

Tavani, H.T. (2004b): Genomic research and data-mining technology: Implications for personal privacy and informed consent. *Ethics and Information Technology*, 6, 15–28.

Vaidya, J., Clifton, C., & Zhu, M. (2006). *Privacy Preserving Data Mining*. New York: Springer.

Vayena, E., Mastroianni, A., Kahn J. (2012). Ethical Issues in Health Research With Novel Online Sources. *American Journal of Public Health*, 102(12), in press.

Van den Hoven, J., Helbing, D., Pedreschi, D., Domingo-Ferrer, J., Gianotti, F., Christen, M. (in press). FuturICT – The Road towards Ethical ICT. *European Physical Journal - Special Topics*.

Vedder, A.H. (2001). KDD, Privacy, Individuality, and Fairness. In R.A. Spinello & H.T. Tavani (Eds.): *Readings in CyberEthics* (pp. 404–412). Sudbury, MA: Jones and Bartlett.

Velleman, J.D. (2006). *Self to Self. Selected essays*. Cambridge: Cambridge University Press.

Walzer, M. (1983). *Spheres of Justice*. Oxford: Blackwell

Williford, C., & Henry, C. (2012). *One Culture. Computationally Intensive Research in the Humanities and Social Sciences. A Report on the Experiences of First Respondents to the Digging Into Data Challenge*. Washington, D.C.: Council on Library and Information Resources. Retrieved September 11, 2012, from <http://www.clir.org/pubs/reports/pub151>

Wilson, R.E., Gosling, S.D., Graham, L.T. (2012). A Review of Facebook Research in the Social Sciences. *Perspectives on Psychological Science*, 7, 203-220.

Witten, I.H., Frank, E., & Hall, M.A. (2011). *Data Mining: Practical Machine Learning Tools and Techniques*. Burlington MA: Morgan Kaufmann

Zarsky, T.Z. (2003). “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, Article 1.

## KEY TERMS AND DEFINITIONS

**Moral identity:** Moral identity is a concept which maintains that people use moral issues such as fairness, kindness, compassion, etc. to define themselves. A person with a moral identity constructs his or herself around moral categories, beliefs, and convictions that are chronically accessible for interpreting the interpersonal landscape.

**Morality mining:** Morality mining is defined as data mining of information that reflects the moral identity of a person and that generates knowledge that is relevant for a person, that may be hidden, character-enhancing and/or stigmatizing for the person, and that a person only wants to disclose in specified contexts.

## **KEYWORDS**

Morality mining; moral identity; moral behavior; privacy; social data mining; ethical awareness; psychology; ethics