

Profiling: From Data to Knowledge

The challenges of a crucial technology

Mireille Hildebrandt

Profiling is not about data but about knowledge. It provides a crucial technology in a society that is flooded with noise and information. Profiling is another term for sophisticated pattern recognition, and the enabling technology for Ambient Intelligence. It confronts us with a new type of inductive knowledge, inferred by means of automated algorithms. To the extent that decisions that impact our lives are based on such knowledge, we need to develop the means to make this knowledge accessible for individual citizens and provide them with the legal and technological tools to anticipate and contest such knowledge or challenge its application.

Introduction: Saved by profiling technologies?

Profiling seems to be the only viable technology that can save us from *the* 2 problems of Information Society: (1) the overload of information and (2) the blurring of the borders between noise, information and knowledge. More data does not necessarily mean more knowledge or information and the increasing use of computer technologies may flood us with data out of which no single human mind can filter what is relevant. We may end up lost instead of knowledgeable. In the most optimistic scenario, profiling promises a dynamic, contextual selection of relevant information and the inference of pertinent knowledge. One could say that a profile is knowledge to the extent that it turns data into information, allowing one to discriminate between relevant and irrelevant data (in a specific context). However, this type of knowledge does not deliver metaphysical truths, proof of causality or conclusive reasoning. Instead it builds on mathematical correlations between aggregated machine-readable data, which correlations are indicative of expected future behaviour. As with all overly optimistic faith in technological progress, reality is sobering but nevertheless profiling does make a difference. Sobering because profiling is not just a technology (combining hardware with software), but also a practice, in need of professional knowledge to decide which are spurious or otherwise irrelevant correlations. Machines are helping us to sort out noise from information and to infer knowledge from data, but in the end we write the software and check whether it does the job. However, profiling does make a difference, because the type of knowledge produced by profiling practices is of a different nature compared to traditional scientific knowledge that starts out with hypotheses to be tested in search of causes (empiricism) or reasons (rationalism). Here we have a pragmatic type of

knowledge that predicts future behaviour on statistical grounds, often building on what seems trivial data. The results are manifold. Genetic profiling can deliver effective correlations between individual genes and diseases, without having much of a clue as to the underlying causal chain. Profiling of keystroke behaviour can serve to identify a person as the same person whenever she goes online, thus allowing web profiling even without any other identifier. The possibility to identify a person in many different ways and for many different purposes leads to the classification of profiling as one of three basic types of identity management [HaMe06].

Integrating the data of a variety of databases enables the construction of refined group profiles that should adequately represent categories of peoples, providing a detailed picture of e.g. their probable earning capacities (e.g. credit storing) and proneness to risk (e.g. health, criminal recidivism, victimisation). Hereunder we will explain how FIDIS understands profiling and indicate some of the challenges it evokes, especially as it is the enabling technology for Ambient Intelligence (AmI) and *The Internet of Things*. In the concluding remarks we will argue for a well balanced tool kit of technological and legal instruments to protect some of the basic tenets of constitutional democracy, as these may face serious provocation in an information society that renders its citizens virtually transparent.

1 The process of profiling

FIDIS is interested in the process of *automated* profiling, which involves [HiBa05]:

- recording of data (taking note of them in a computable manner):
- storing data (in a way that renders them accessible, aggregated in a certain way)
- tracking data (recording and storing over a period of time, linking data to the same data subject)



Dr Mireille Hildebrandt

Senior Researcher at LSTS, Vrije Universiteit Brussel, teaching at Faculty of Law Erasmus University Rotterdam.

Focus on issues of identity in constitutional democracy.

E-Mail: hildebrandt@frg.eur.nl

- identifying patterns in the data (by running algorithms through the data base)
- monitoring data (checking whether new data fit the pattern or produce outliers)

The process is often described in terms of knowledge discovery in data bases (KDD), involving the collection and storage of data, data mining, interpretation and decision making [Cu04]. For the data mining process a set of non-proprietary guidelines is freely available, partly funded by the European Commission, developed in conjunction with practitioners and vendors, called the Cross-Industry Standard Process for Data Mining (CRISP-DM).¹ This model emphasises the feedback between the different stages of data mining, consisting of business understanding, data understanding, data preparation, modelling, evaluation and deployment. What is important to keep in mind is the fact that automated profiling depends on adequate recording and storage of digitalised data. Between the events, transactions or movements and their storage a translation takes place that transforms a fluid moment into machine-readable data. The data are recorded as a type of brute facts, de-contextualised and – as far as the machine is concerned – without meaning. Most data may be trivial in themselves, acquiring new meaning after having been connected with other trivial data and used for decision making. The crucial instances in the process of data mining are the emergence of correlations and their interpretation. Profiling is basically a matter of pattern recognition: the knowledge inferred from the data consists of association, classification or clustering. For instance, the process of KDD may produce patterns that correlate a certain gait to specific learning disabilities. However, the interpretation – the meaning of the pattern that is found – depends on practical wisdom or professional knowledge, putting the newly found correlations in a specific professional context. In this case experts specialised in learning disabilities would be involved to assess the relevance of the correlations. If the automated profile is considered as relevant knowledge by the experts, it also defines certain data as relevant, thus transforming them into information. After all, data can be noise or information depending on such knowledge. Applications of profiling technologies can be found in marketing, criminal investigation, detection of fraud or money laundering. However, in the context of autonomic pro-

filings involving real-time adaptation of an environment to a user's inferred preferences, the interpretation may be done by machines. In a reiterating process of checking for outliers while applying generated profiles, a machine-learning process may evolve, even if the software will be checked and adjusted by means of human intervention. Autonomic computing will become especially important if the vision of Ambient Intelligence becomes a reality, to which we will devote some special attention below.

2 Some pertinent distinctions

When speaking of profiling one may refer to a host of different phenomena, which can lead to a Babylonian confusion. For this reason FIDIS discriminates between group profiles and personalised profiles and between the construction of profiles and the application of profiles. In both cases the distinction is analytically salient, while in practice the phenomena intermingle.

2.1 Group profiles and personalised profiles

A group profile identifies and represents a group (community or category), of which it describes a set of attributes. The group can consist of people that think of themselves as a community, like a class of students, adherents to a specific religion or members of an association. The group can also consist of a category of people that have no connections amongst them, other than the fact that profiling has established them as a category. For instance, data mining may produce a correlation between left-handed people and a certain disease or a certain propensity towards artistic endeavour. This correlation is probabilistic and does not depend on the fact that left handed people form any sort of community. The fact that one can be identified as a member of this category does not necessarily mean that one shares the attributes of this group. This will be discussed hereunder in reference to non-distributive group profiles.

A personalised profile identifies and represents a person, of which it describes a set of attributes. The profile can be based entirely on the recorded data of one individual person, for instance his keystroke behaviour or a combination of different types of correlated data like keystroke

behaviour and surfing habits. Because the profile is directed to one individual person of whom it may disclose intimate knowledge, a personalised profile seems to have a direct impact on privacy. However, this depends on how we understand privacy, since it may be the case that the profile identifies a person over a period of time as the same person, disclosing his surfing habits, without having access to his name or other personal data. At the same time we should take note that group profiles can be highly specific in a particular context, providing a very rich profile that comes very close to a personalised profile. This indicates the blurring of the border between the two types of profiles. Furthermore, personalised profiles can be aggregated to produce a group profile.

2.2 Distributive and non-distributive group profiles

Group profiling identifies and represents a group, that is, a community or a category. In the case of a distributive group profile, the attributes of the group are also the attributes of all the members of the group [Ve99]. For instance, the attribute of 'not being married' for the group of bachelors, but also for any member of that group. However, in the case of a non-distributive group profile, matters are complicated. Imagine if a person is included in the group of people with blue eyes and red hair and imagine that it is the case that a group profile is constructed for this category that indicates 88 % probability of a specific type of skin disease. This does not mean that this particular person has an 88% chance to have this disease, because this may depend on other factors (like age, sunlight, eating habits, use of skin lotions). It does mean, however, that belonging to the category allows what Schauer [Sch03] calls a non-universal generalisation.

In real life, most generalisations are non-universal, meaning that we learn to cope with the abundance of detail by imposing some order in the form of generalisations or categories that provide adequate standards to assess a new situation. These generalisations are seldom universal; they are short hand for more complex standards that incorporate the fact that not all members of a category share the same features. If I say that people who smoke will end up with lung cancer I will probably be aware of the fact that this may be the case for a number of people but not for all. This 'goes without

¹ See www.crisp-dm.org

saying'. Group profiling seems to produce such generalisations on the basis of sophisticated algorithms, based on mathematical inferences instead of intuitive rules of thumb. This, of course, means that all the dangers of non-universal generalisations apply equally to non-distributive profiles: one cannot presume for any member of a category that the group profile applies without access to additional information. As soon as decisions – with for instance legal consequence or other serious impact – are taken on the basis of such a wrongful presumption, we find ourselves in the realm of illegitimate discrimination.

2.3 Construction and application of profiles

When discussing profiling we should distinguish between the construction of a profile, by means of data mining techniques, and the application of a profile, for instance to inform the decision which category of people should be offered (or refused) a specific service. As indicated above, the processes of construction and application are intermingled. Applying a profile may consist of checking outliers that suggest unusual – or undesirable – behaviour. This checking of outliers is at the same time a test for the profile, allowing adjustment to curb the number of outliers or to change some of the parameters that have turned out to be spurious.

3 The Internet of Things and Autonomic Profiling

Profiling is *the* enabling technology for Ambient Intelligence (ISTAG 2001), [SchHi05] and what has been called *The Internet of Things* [ITU05]. Without adequate profiling we will not be able to handle the innumerable data recorded when the real world goes online, we will miss the means to make sense of the data, mistaking noise for information or information for knowledge. In fact, we would be engulfed by noise, lacking the tools to discriminate between what is relevant at which moment in which context.

Imagine all *things* to be RFID-tagged and part of an RFID-system that allows reading and online storing of their status, location and other data, while at the same time all spaces are provided with sensor

devices and CCTV-cameras that detect movement, temperature, and other data. When this vision materialises we will find ourselves in the *everywhere* of a networked environment that seamlessly integrates real time monitoring with real time proactive adjustments of the environment. Ambient Intelligence implies that the environment is able to anticipate a user's wishes, even before he becomes aware of them. This is expected to move well beyond anticipating how you like your coffee or room temperature, as it may cater to your specific health needs, travel plans or your preferred professional infrastructure. To allow real time adjustment of the environment we need autonomic profiling. Autonomic profiling goes one step further than automated profiling. The term is derived from what Paul Horn, IBM's senior vice-president, has named autonomous computing [KeCh03]. This is a type of computing that not only performs algorithmic functions on incoming data, but also takes a number of decisions that amount to a kind of self-management. Horn compares autonomic computing to the autonomic function of our nervous system, claiming it should provide for a continuously readjusted environment without disturbing us with complex decision-making processes. Just like your nervous system does not ask for your consent to adjust your body temperature or heart rate, autonomic computing should unobtrusively work out the right fit with your surroundings. Autonomic profiling thus implies that adaptive environments function smoothly without too much intervention of the end-user, meaning that machines take all the necessary decisions, based on their profiling activities [Hi07]. This is meant to unburden the human person, but it may obviously also disempower citizens regarding the choices that are made for them [BoCo04].

4 Beyond privacy and security?

The discourse on the dangers or threats faced by further development of information society seems locked in a debate about the balance between privacy and security. Citizens are persuaded that in times of international terrorism and transnational organised crime, they are better off when trading a bit of their privacy for security and it seems that apart from privacy advocates not many citizens have sleepless nights over this trade-off. The exchange of cross-

disciplinary perspectives within the FIDIS research community has led to broader perspectives on these issues, with special regard to the implications of profiling for democracy and the rule of law [HiGu05].

First of all the debate seems entirely focused on data, while many of these are trivial and most of the time they are assessed by machines rather than by humans. Apart from abuse there seems little reason to fear the collection and storage of these data.

However, in the case of profiling we are not dealing with data, but with inferred knowledge. For two reasons this is more worrying: (1) non-distributive group profiles are based on probabilities, this means that the group profile does not automatically apply to each member of the group, (2) profiles may reveal sophisticated knowledge about a person that is more intimate than sensitive personal data.

Solove [So04] warns that we may develop a general fear that anything we do will be recorded and can be used against us at any point in time, on the basis of knowledge produced by indifferent anonymous machines. He suggests that the metaphor of Big Brother does not cover the distributed spying generated by a host of private and public organisations. Instead he refers us to the metaphor of Kafka's *The Trial*, because it saliently articulates the vagueness of the accusations and the indifference of the prosecuting bureaucracy.

Second, as a consequence of the focus on data instead of knowledge, the debate seems to be directed to anonymisation, or the use of pseudonyms, in order to protect personal data.

However, citizens may rather need protection against the application of profiles, or at least access to such profiles and transparency concerning their use.

Data Protection is a tool of transparency that aims to guarantee access to the processing of personal data, but precisely when personal data are anonymised data protection legislation is no longer applicable. This means that citizens have no legal right to even access the knowledge that is inferred from these anonymised data and may be used in ways that impact their lives. Once a profile is linked to an identifiable person – for instance in the case of credit scoring – it may turn into a personal data, thus reviving the applicability of data protection legislation. This protection, however, comes after

the fact [SchHi07], not providing access to the dynamic group profiles available to the service provider, who may even protect the profiles by means of intellectual property rights. Art. 15 of the data protection directive (D 46/95 EC) does provide some protection in the case a 'decision which produces legal effects concerning him or significantly affects him and which is based solely on automated processing of data intended to evaluate certain personal aspects relating to him, such as his performance at work, creditworthiness, reliability, conduct'. As Bygrave [By01] explains in his analysis of art. 15, the article provides data subjects with a right not be subjected to such decisions, but we may expect that if the right is not exercised, these types of decisions will continue to proliferate. Also in this case, one does not have access to the dynamic group profiles that may or may not be applied. This produces an asymmetry of knowledge between profiler and profiled subject.

Third, in an AmI environment people may be identified on the basis of behavioural biometric profiling, which renders identification in the sense of art 2 (a) of D 46/95 EC unnecessary, again ruling out applicability of data protection legislation as it stands now.

Fourth, in the context of ICT privacy is often reduced to control over the disclosure of personal data, while in fact privacy is a good or value that concerns both more and less than the exchange of data. Privacy concerns the capacity to continuously reconstruct one's identity and to control the borders between self and others [Ag01].

For this reason it is pertinent to distinguish between data protection and privacy, because the one is a tool that aims for transparency, while the latter refers the opacity of the personal sphere that should enable the positive and negative freedom of individual citizens, empowering them to partake in private and public life without undue interference.

Reducing privacy to control over personal data mistakes data protection, which actually aims for a free flow of information, for the protection of essential rights and liberties, which may be at stake at the moment of application of group profiles rather than at the moment of data collection.

Fifth, the prevalent focus on privacy and security issues seems to distract attention from far-reaching consequences of advanced profiling technologies for equality,

fairness and due process in the wider societal context.

Profiling shifts the balance of power between those that can afford profiling (mostly large organisations) and those that are being profiled (mostly individual citizens), because the profilers have a certain type of knowledge to which those profiled have no effective access. This particular lack of transparency is not only a matter of the non-applicability of data protection regimes in the case of anonymised data. At this moment we also lack the technological tools to anticipate the type of profiles that may be constructed and applied to us.

Sixth, the value of privacy is often understood as if privacy is a private good, one that can be traded at will against other private goods, or even disowned in order to protect public goods like intra- or international security. Without denying that privacy is also a private good, we should not forget the value of privacy as a public good that is preconditional for a viable constitutional democracy.

Like other public goods, such as security, equality, fairness and due process, privacy needs protection beyond the arbitrary decisions of individual citizens. Profiling may not impact our *sense* of privacy, or even our expectation of privacy, because we are not aware of it. But it may still invade our privacy to a much greater extent than unauthorised use of personal data. The threat of autonomic profiling is the unobtrusive ubiquitous disclosure of patterns that define our most intimate habits, beyond our awareness. It may provide us with a golden cage, in as far as AmI caters to our inferred wishes. But, like Sunstein claims, even if we would prefer this in our capacity of private citizens, it will undermine our capacities to function as public citizens. This is the case because we will lack the confrontation with what the machines expect us to dislike, thus reducing our confrontation with 'unplanned, unexpected encounters [that] are central to democracy itself'. [Su01:8].

5 Ambient Law: Integration of Data Protection, TETs and PETs

Anonymisation may protect against abuse of personal data, but it will not protect against application of group profiles, inferred from anonymised data or from the application of a group profile to a person that is identified only by means of behavioural biometric profiling. This is not to claim that the application of group profiles is a bad thing in itself [Kr86], or to claim that anonymisation or the use of pseudonyms makes no sense. PETs can provide much needed means for identity management, as discussed in [BaMeHa05]. The point is that we need to find ways to render the processing of data transparent, after the data have been anonymised and before they are applied. Citizens must be able to anticipate the profiles that may be applied to them and be given the legal and technological tools to resist the validity and relevance of the profile in their particular case.

For this reason FIDIS aims to develop a cross-disciplinary perspective between computer scientists, technologists and lawyers to prepare a technological infrastructure that would:

- integrate the mandatory aspects of Data Protection legislation and
- facilitate machine-to-machine communication between citizens' personal digital assistants and networked environments, allowing adequate anticipation of automatic profiling.

In this case the focus of Privacy Enhancing Technologies (PETs) would be on what has been called the 'principle of minimum asymmetry' [Ji02], combining the data minimisation principle that restricts the flow of information of data subjects to data processors, with a maximisation of the feedback from data processors. Not just to find out what happened to your personal data, but first of all to find out which profiles may be inferred that will impact you as a member of a certain category. To allow such 'counter profiling' we need to develop Transparency Enhancing Technologies (TETs).

Ambient Law would imply that the use of PETs (and TETs) is not left to individual preference but is part and parcel of a legal-technological framework that is preconditional for the exercise of individual preference. It involves clear thinking about the normative impacts of technological artifacts and technological infrastructure and demands political choice about the kind of information society we want to inhabit.

ditional for the exercise of individual preference. It involves clear thinking about the normative impacts of technological artifacts and technological infrastructure and demands political choice about the kind of information society we want to inhabit.

Summary

Data Protection is focused on data. It takes a proactive perspective by demanding that data are collected in a restricted manner, pointedly expressed in the data minimisation principle.

Profiling is not about data but about knowledge. However it feeds on data and in the context of an Internet of Things or an Ambient Intelligent environment it demands as many data as possible. Even though the protection of personal data can limit profiling by limiting the input of data, anonymisation will not limit but rather facilitate large scale group profiling. The protection needed at this point is not just of our own data but the protection of our capacity to anticipate which group profiles may affect our personal lives. For this we need to create a legal-technological infrastructure that provides us with the legal-technological means to minimise the leaking of data, to anticipate which profiles may affect us, to contest the inherent knowledge claims they entail and to challenge their application if necessary.

Literature

- Ag01 Agre, P. E. 'Introduction', *Technology and Privacy: The New Landscape*. P. E. Agre and M. Rotenberg. Cambridge, Massachusetts, MIT, 2001
- BaMeHa05 Bauer, M., Meints, M., Hansen, M. (Hrsg.), *FIDIS Deliverable D3.1 – Structured Overview on Prototypes and Concepts of Identity Management Systems*, Frankfurt a.M. 2005. Download: <http://www.fidis.net/486.0.html>.
- BoCo04 Bohn, J., V. Coroama, et al. *Social, Economic, and Ethical Implications of Ambient Intelligence and Ubiquitous Computing*, Institute for Pervasive Computing, ETH Zurich, Zurich, 2004. Download: www.vs.inf.ethz.ch/publ/papers/socialambient.pdf.
- By01 Bygrave, L. 'Minding the Machine. Art 15 and the EC Data Protection Directive and automated profiling.' *Computer Law & Security Report*, 17: pp. 17-24, 2001
- Cu04 Custers, B. *The Power of Knowledge. Ethical, Legal, and Technological Aspects of Data Mining and Group Profiling in Epidemiology.*, Wolf Legal Publishers, Nijmegen 2004.
- HaMe06 Hansen, M., Meints, M., 'Digitale Identitäten – Überblick und aktuelle Trends', in this issue.
- Hi07 Hildebrandt, M. 'Defining Profiling: A New Type of Knowledge', in *Profiling the European Citizen. A Cross-disciplinary Perspective*. M. Hildebrandt and S. Gutwirth (Eds.), Springer 2007.
- HiBa05 Hildebrandt, M., Backhouse, J. *FIDIS Deliverable D7.2 – Descriptive analysis and inventory of profiling practices*. Brussels 2005. Download via: www.fidis.net
- HiGu05 Hildebrandt, M. and S. Gutwirth, (Eds.) *FIDIS Deliverable D7.4 – Implications of profiling practices on democracy and rule of law*. Brussels 2005. Download via <http://www.fidis.net>
- ISTAG01 ISTAG., *Scenarios for Ambient Intelligence in 2010, Information Society Technology Advisory Group 2001*. Download: <http://www.cordis.lu/istag-reports.htm>
- ITU05 International Telecommunications Union (ITU), *The Internet of Things*. Geneva 2005.
- Ji02 Jiang, X. *Safeguard Privacy in Ubiquitous Computing with Decentralized Information Spaces: Bridging the Technical and the Social*. Privacy Workshop September 29, 2002, University of California, Berkeley. Berkeley 2002. Download: <http://guir.berkeley.edu/pubs/ubicomp2002/privacyworkshop/papers/jiang-privacyworkshop.pdf>
- KeCh03 Kephart, J. O. and D. M. Chess 'The Vision of Autonomic Computing.' Computer January 2003.
- Kr86 Kranzberg, M., 'Technology and History: 'Kranzberg's Laws'' *Technology and Culture* 27: pp. 544-560, 1986
- Sch03 Schauer, F. *Profiles Probabilities and Stereotypes*. Cambridge, Massachusetts, London, England, Belknap Press of Harvard University Press 2003.
- SchHi05 Schreurs, W., M. Hildebrandt, et al. (Eds.), *FIDIS Deliverable D7.3 – Report on Actual and Possible Profiling Techniques in the Field of Ambient Intelligence*, p. 68, Brussels 2005.
- SchHi07 Schreurs, W., M. Hildebrandt, et al. 'Cogitas ergo sum. The role of data protection law and non discrimination law in group profiling in the private sector, in: *Profiling the European Citizen. A Cross-disciplinary Perspective*, M. Hildebrandt and S. Gutwirth (Eds.), Springer 2007
- So04 Solove, D. J., *The Digital Person. Technology and Privacy in the Information Age*. New York, New York University Press 2004.
- Su01 Sunstein, C., *Republic.com*. Princeton and Oxford, Princeton University Press 2001.
- Ve99 Vedder, A. „KDD: The challenge to individualism.“ *Ethics and Information Technology* 1, pp. 275-281, 1999.