Ethics and Information Technology (2005) 7:211–220
DOI 10.1007/s10676-006-0010-6

Data mining to combat terrorism and the roots of privacy concerns

Frans A. J. Birrer
Leiden Institute for Advanced Computer Science (LIACS), P.O. Box 9512, 2300 RA, Leiden, Netherlands
E-mail: birrer@liacs.nl

Abstract. Recently, there has been a heavy debate in the US about the government's use of data mining in its fight against terrorism. Privacy concerns in fact led the Congress to terminate the funding of TIA, a program for advanced information technology to be used in the combat of terrorism. The arguments put forward in this debate, more specifically those found in the main report and minority report by the TAPAC established by the Secretary of Defense to examine the TIA issue, will be analysed to trace the deeper roots of this controversy. This analysis will in turn be used as a test case to examine the adequacy of the usual theoretical frameworks for these kinds of issues, in particular the notion of privacy. Whereas the dominant theoretical framing of the notion of privacy turns around *access* to information, most of the core arguments in the debate do not fit in this kind of framework. The basic disagreements in the controversy are not about mere access, they involve both access *and use*. Furthermore, whereas the issue of access by itself refers to a more or less static situation, the real disagreements much more concern the organisational dynamics of the use of information, the mechanisms in the organisation that control these dynamics, and the awareness present within the organisation of the 'social risks' these dynamics represent. The bottom line question is whether the assessment of these gives sufficient reason for trust.

Key words: data mining, ethics, privacy, risk, security, systems of subliminal enticement, terrorism

Introduction

Data mining is an emerging technology, that is perceived as highly promising in a number of areas, and that is increasingly used and developed. Recently data mining techniques for combating terrorism raised extensive discussion in the US. The discussion provides a good starting point for exploring some analytical categories to identify what are the main ethical issues connected with data mining in this context. Furthermore, such an analysis can produce clarifying observations where the discussion is confused.

Data mining techniques were and are involved in several governmental programs in the US such as Terrorist Information Awareness (TIA, involving an integrated database system to identify potential foreign terrorists) and Multistate Anti-terrorism Information Exchange (MATRIX, connecting databases of the states participating). A column by William Safire in the New York Times[1] triggered a public discussion on data mining under TIA and other programs. Its strong impact is illustrated by the fact that in January, 2003, Congress imposed a moratorium on data mining under TIA as well as under similar programs until more information would be provided on these programs, and in September of the same year decided to terminate the funding of TIA.

In this article various arguments in this controversy are analysed, in particular the arguments that can be found in the main and minority report from the Technology and Privacy Advisory Committee (TAPAC) that was specifically installed to examine this matter. These arguments will be confronted with the common theoretical framing of such issues, particularly with respect to privacy concerns. It will turn out that the main origins of disagreement do not fit into the usual frameworks.

What is data mining?

Broadly conceived, data mining is a field of computer science that can be described as concerned with 'the extraction of implicit, previously unknown, and potentially useful information from data'[2] or

[1] William Safire. You are a suspect. *New York Times*, November 14, 2002, p. A 35.

[2] Ian A. Witten, Eibe Frank. *Data mining*. Morgan Kaufmann, San Francisco, 2000, p. 1.

'extracting useful information from large data sets or databases'.[3] Whereas the term Knowledge Discovery in Databases (KDD) commonly is used to cover the whole trajectory from data preparation up to implementation, the term 'data mining' tends to be restricted to the actual extraction process itself. Since there is some confusion concerning the precise meaning of the term 'data mining', and since it is important to have a proper understanding of what kind of basic techniques are involved, a brief explanation is due regarding to the meaning of a phrase like 'implicit information'.

Databases are constructed to answer specific questions ('queries'). A data base is set up in such a way that the queries that are specific for the data base can be processed easily, by means of direct linkages between the items that the queries may connect. For instance, a company collects some personal data about its employees, such as name, address, marital status, salary, bank account, etc.; on entering the name (or personal identification number) of an employee, the administration database should then be able to answer queries like 'what is this person's address?', or 'what is this person's salary?' With some more effort, however, it may also be possible to get answers to slightly less obvious questions like 'which employee is living at this address?', or even 'which male employees are unmarried and frequently absent?' Of course, if these are not standard queries, some additional programming or combination of answers to standard queries will be necessary. Even more effort may be needed when answers are sought that require information from several separate databases. Information like this exists in the database only in an *implicit* form, in the sense that the database was not set up to answer such questions, and is not structured in such a way as to find the answer in the most straightforward possible way (i.e., through a standard query). This, then, is a first type of activity that could be described as 'data mining': searching a (large) database (or a set of coupled databases) for items with a specific combination of characteristics that does not correspond to a standard query. Although in common language the term 'data mining' certainly seems appropriate here, this kind of activity is usually not included under the term 'data mining' in the computer science literature.

A second mode of operation is that we let the computer itself search for (frequently occurring or otherwise significant) combinations of characteristics in a database or collection of databases. For instance, a supermarket database could be searched to find products that are often bought together. Or the police may want to map networks of criminals or criminal activities. For such searches clustering and other algorithms exist or can be developed. This mode of operation is often called '*descriptive data mining*', in contrast to the third type to be discussed next.

A third mode of operations is that patterns are searched for and used with the aim of predicting certain characteristics. For instance, the police may be interested in characteristics that could be indicative for criminal activities. Such patterns may be discovered from a small subset of those activities that are known to be connected to criminality; the predictive value is then tested on a different subset of activities known to be connected to criminality; finally, the pattern may be used as indicative of the potential criminality of activities that were not already known to be so. These kinds of searches are called '*predictive data mining*'.

Although the distinction between descriptive and predictive data mining frequently occurs in the computer science literature, the semantics again is not always entirely clear-cut. A 'descriptive' pattern such as products that are often bought together in a supermarket can be used to change the display, grouping those articles together; this could in a sense already be said to be used predictively, namely based on the assumption that new purchases will follow the same pattern. More frequently, the distinction between descriptive and predictive data mining seems to refer to the difference between interpreting a pattern at an overall level, or as a predictor for individual items (persons). As we will see later, from an ethical point of view, the most vital distinction is whether or not individuals or groups of individuals are treated differently on the basis of statistical (i.e., uncertain) inferences.

It is the second and third type of data mining that tend to get most emphasis in the computer science literature, and in fact are often identified with the term, while the much broader term KDD tends to include the first type. Since there is no general agreement on the precise definitions of terms like 'data mining' or 'KDD', their meaning often remains a bit fuzzy. The book by Hand, Mannila and Smyth mentioned earlier, for instance, presents the following 'working definition' of data mining: ''Data mining is the analysis of (often large) observational data sets to find unsuspected relationships and to summarise the data in novel ways that are both understandable and useful to the data owner',[4] which seems considerably more narrow than the description they presented just

[3] David Hand, Heikki Mannila, Padhraic Smyth. *Principles of data mining*. MIT Press, Cambridge (Mass.), 2001, p. xxvii.

[4] Hand, Mannila, Smyth, p. 1.

a few pages earlier and that was already quoted above. No matter how we name them, it is essential to be aware that different types of operations can be involved.

Privacy and ethics

The great majority of the concerns that have been raised regarding the use of data mining to combat terrorism are commonly described as issues of privacy. Privacy is another notion that needs to be handled with care, since it easily triggers misleading representations of what is at stake. Originally, the term privacy referred to issues regarding intrusion by journalists, photographers etc. into the private life of (well-known) individuals. In this original context, it gave rise to the conception of an unassailable private sphere where strangers should not have uninvited access. This idea very much has left its imprint on the term in its more recent use, regarding personal data. Unfortunately, the analogue in the form of a clearly demarcated sphere of personal data where strangers are not granted access is a misleading one, for at least three reasons.[5]

First, the spectrum of parties potentially involved in informational privacy issues is so wide, and their relations so different from case to case, that a uniformly defined demarcation line for information not to be accessed by outsiders is unlikely to exist. Instead, for issues of informational privacy a tradeoff will have to be made for each case, or at least for each type of case, between the needs, interests and wishes of all parties involved. For an issue like the combat of terrorism, the tradeoff and the demarcation line will be different from that for an activity like direct marketing. This also implies that in general, information privacy arguments cannot be based solely on unalienable rights. It all depends on the identity and the objectives of the information gatherer/user, and, of course, on the individual the information is about. Furthermore, these tradeoffs need not necessarily be decidable by ethics, they will often be subject of political negotiation. Ethics can only attempt to specify extreme boundaries of definitely unacceptable outcomes, and at the meta-level it can try to specify when the negotiation process is fair.

Second, the majority of the issues that in common parlance are labelled as issues of privacy cannot be reduced to matters of access. Acceptability depends on the *use* that is made of that information as much as on the nature of the information. It will already be clear that what is called *predictive* data mining involves much more than just access to the data, it implies (and is inextricably connected with) decisions with respect to individuals or groups of individuals. It is not accidental that in much privacy regulation the admittance to create and use a certain data collection is strictly coupled to the declared *purpose* of those activities. Often the greatest threat to the individual whose records are being viewed is not just in the access to the data, but much more in the *conclusions* that will be drawn from those data, in the *actions* that will be undertaken as a result of these conclusions, and in the *consequences* of those actions for that particular individual. When the record data are transformed into patterns as indicator for something like (increased potential of) criminal activity, it is no longer the data as such, but their specific framing and use in the context of certain decisions and actions, that is at stake when acceptability is concerned. Even relatively innocent looking information can thus become used for highly disputable purposes.

Third, ethical considerations, if they are to be practically relevant, cannot be detached from the social context and social dynamics in which they arise and to which they apply. The issue is not just whether certain acts are ethical in an abstract ethical evaluation space; a question that is at least as pressing is which types of ethical evaluation systems are socially viable, that is, could be stabilised in our social system, and what arrangements would be necessary for this stabilisation. Nor can we confine ourselves to ask what acts would be ethical; of equal importance is to ask what are the drives that direct people towards unethical behaviour, and how these drives can be tempered and diverted.[6] Too often, proposals for measures to enhance ethical behaviour are based on an analysis that simply assumes that all actors are willing to, and actually will, behave ethically. But real world ethics cannot completely ignore the chains of acts and consequences that occur in a world in which not everyone will automatically do what has been shown to be ethical. Even more than the first point, this implies that questions of ethics become inextri-

[5] Cf. Frans A.J. Birrer. Applying ethical and moral concepts and theories to IT contexts: some key problems and challenges. In Richard A. Spinello, Herman T. Tavani, editors, *Readings in cyberethics*, Jones & Bartlett Computer Science, Sudbury (MA), 2001, pp. 91–97.

[6] I have entered into such questions in more detail in Birrer, 2000, and elsewhere Frans A.J. Birrer. Computer technology, subliminal enticement, and the collectivisation of ethics, In Deborah G. Johnson, Jim H. Moor, Herman T. Tavani, editors, *Computer ethics: Philosophical enquiry* (CEPE 2000 Proceedings). Dartmouth College, Dartmouth, 2000.

cably intertwined with questions pertaining to the social science dimension. To the 'ethical purist' that may sound discomforting, and as an intrusion and defilement of veritable ethics. But as has been recognised in many areas of scholarship today, practical problems do not stop at disciplinary boundaries; compulsory attempts to squeeze a practical problem into the confinement of a single discipline usually end up in serious misrepresentation and irrelevance.

Some might want to maintain that we should stick to the original meaning of 'privacy', and that it is its conflation with other issues that should be blamed for creating confusion. Academically speaking, this argument has a point. Practically speaking, however, we are forced to observe that the common use of the term 'privacy' has extended beyond this original notion to include issues of use, and that it is not feasible for philosophers to reverse this trend. Moreover, it might seriously harm the societal discussion on privacy to insert statements based on a notion of privacy that does not correspond to the way it is presently used in the public debate. It seems a wiser and more realistic strategy to use a more restricted term to refer to privacy in its original sense.

Data mining, privacy, and ethics

As we already observed, privacy issues are not merely about whether one is allowed to know something about another person, but also about how one uses that knowledge. That is, the context of use is as important as the (content of the) information itself. When we then focus on the use of information, and on the ethics of that use, it seems to me that the following distinction is crucial to the very nature of the ethical problem:

> Sometimes, the data are used just as they are, that is, without any further interpretation that they did not have from the very start. Validation of the decisions and choices made on the basis of the data then depends only on the adequacy of the data. I will call conclusions that are drawn in this way '*direct inferences*'.
>
> On the other hand, as we have already seen exemplified in the form of pattern searches, information (data) can also be used as an indicator for something else, usually for something that itself is difficult to measure directly; e.g., certain patterns of behaviour might become employed as an indicator of increased possibility of being a terrorist, even though such behaviour could also be displayed by completely innocent persons, and by no means necessarily implies that one is a terrorist. I will call such inferences '*correlative inferences*'. Here, inferences go beyond the meaning of the data as such.

For correlative inferences, validation of the decisions or choices involves a statistical correlation between the indicator and what it is supposed to indicate. In ethical terms, it particularly implies the responsibility for 'false positives' and 'false negatives'.[7] False positives refers to items or cases that are categorised in the target group but do no actually belong to the target group. False negatives are items or cases that are not categorised as belonging to the target group but that actually do belong to the target group. In contradistinction to direct inference, correlative inferences do not merely depend on the accuracy of the data for their reliability; in fact, indirect interpretation implies that a certain amount of miscategorisations is deliberately accepted. That makes this kind of inference ethically very distinct from the first.

Data mining for the combat of terrorism in the US

Data mining techniques are explored and used in several governmental programs in the US, not just in the combat against terrorism, but also against other forms of criminality, as support for financial accounting, and for other purposes.[8] Given the nature of data mining, it is not surprising that this technique was taken up as a promising instrument against terrorism, particularly after the events of September 11, 2001 had swung this issue to the top of the agenda.[9]

Early 2002, DARPA announced the Total Information Awareness (TIA) program, to be conducted by the newly created Information Awareness Office (IAO). Its aim was to explore and develop a wide range of techniques in the area of information processing and communication, data mining being one of

[7] Cf. Frans A.J. Birrer. Statistical evidence: Responsibilities and the burden of proof. In Cor van Dijkum, Jörg Blasius, Henk Kleijer, Branko van Hilten, editors, *Recent developments and applications is social research methodology*, SISWO, Amsterdam, 2004.

[8] For an overview of current Federal efforts in data mining see: GAO (General Accounting Office). *Data mining. Federal efforts cover a wide range of uses*. GAO Report 04-548, GAO, Washington, DC, 2004.

[9] For a recent overview of the possibilities of data mining techniques in a homeland security context see: Jesus Mena. *Homeland security. Techniques and technologies*. Charles River Media, Hingham (Mass.), 2004.

those techniques. Some members of Congress, as well as some members of the general public, found the program's description insufficiently clear, and accordingly felt uneasy regarding the potential consequences of the program. The public discussion on data mining under TIA and other programs was very much triggered by a column by William Safire in the New York Times in November 2002.[10] An internal review by the Department of Defense (DoD) suggested that there was nothing wrong with the TIA, but this was by no means enough to silence the critics.[11]

Already on January 16, 2003, Congress imposed a moratorium on data mining under TIA as well as under similar programs until more information would be provided on these programs.[12]

In February 2003, Donald Rumsfeld, Secretary of Defense, established the TAPAC to examine the use of 'advanced information technologies to identify terrorists before they act'.[13]

On May 20, 2003, DARPA produced a more extensive report on TIA.[14] In the mean time, the name of the program had been changed into 'Terrorist Information Awareness'. Again, the report assured that TIA would not endanger the privacy of American citizens. But again, resistance and suspicions were not laid to rest.

On September 25, 2003, Congress decided, by a vote of 407-15 in the House and a vote of 95-0 in the Senate,[15] to terminate the funding of TIA, and the dissolution of IAO; a few specific subprograms of TIA would be continued elsewhere.[16]

In March 2004, the TAPAC committee published its report (TAPAC, 2004). The recommendations of the main report included the establishment of a regulatory framework, oversight mechanisms, yearly public reports by the DoD, requirement of a written finding by a federal magistrate or judge for access to federal databases, acceptable rates of false positives and a system for dealing with false positives. The report was a reflection of the controversy itself, since it contained a minority report, by committee member William T. Coleman, Jr., that significantly departed from the opinions expressed in the main text. The points where Coleman disagreed included that DoD should report to a number of committees instead of writing public reports, and that a written finding for access to federal databases would not be necessary. The primary concern of this paper will not be with the details of these or other proposed measures; instead, it will focus on the general arguments that are put forward in the discussion by each side to support its position.

An analysis of the US discussion: the TAPAC report examined

With a minority report commenting on the main report, and the main report commenting on the minority report, the TAPAC report presents a comprehensive microcosm where a wide range of the arguments in the broader public discussion can be found. I will therefore take the report as the primal source for arguments pro and contra, while occasionally referring to other sources where similar arguments are put forward.

The main report starts off its discussion with a number of juridical arguments, suggesting the importance of privacy considerations in American law, and even in the Constitution, particularly the Fourth Amendment.[17] The protection by the Fourth Amendment is not absolute, there are exceptions. One of these exceptions is that the Supreme Court held in 1976 that the Fourth Amendment does not apply to information held by a third party. The main report interprets this exception as being based on the assumption that the information is provided voluntarily by the citizen in question; since the provision of much information to the government actually is involuntarily, or at least on a basis of which the voluntary character is questionable (as is in fact true of much information provided by citizens to private organisations as well), it would still fall under the Fourth Amendment. Coleman, in his minority report, sticks to a more absolute interpretation of the Supreme Court's decision, and holds that the use of

[10] Saffire, 2002.

[11] Technology and Privacy Advisory Committee (TAPAC). *Safeguarding privacy in the fight against terrorism*. Department of Defense, Washington, DC, 2004, p. 16

[12] 108th Congress, 1st session, s. 188

[13] TAPAC, 2003, p. 1

[14] Defense Advanced Research Projects Agency (DARPA). *Report to Congress regarding the Terrorist Information Awareness Program (in response to Consolidated Appropriations Resolution, 2003, Pub. L. No. 108-7, Division M, § 111 (b))*, May 20, 2003.

[15] Nancy Kranich. *Commentary: MATRIX and the new surveillance states: the multistate anti-terrorism information exchange*. Free Expression Policy Project, October 16, 2003.

[16] Department of Defense Appropriations Act, 2004, Pub. L. No. 108-84, September 25, 2003.

[17] 'The right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures, shall not be violated, and no Warrants shall issue, but upon probable cause, supported by Oath or affirmation, and particularly describing the place to be searched, and the persons or things to be seized.' (US Constitution, Amendment IV).

third party information should not be restricted in the TIA program. More in general, Coleman believes that the main report is unjustified in lifting the issues to the level of the Constitution:

> The report (...) wrongly elevates the concept of privacy and protection thereof (...) to the same constitutional level as the fundamental values of liberty, free speech, religion, the political process and racial discrimination issues (...).[18] (Coleman, 2004, p. 67).

These arguments do not, of course, constitute an outright ethical discussion, since their primary frame of reference is whether something is or is not in fact protected by actual law, rather than what arguments would or would not make certain activities morally or even politically acceptable. Nevertheless, already in this juridical discussion some contours of ethical/normative background positions emerge. For Coleman, the goal of fighting terrorism is of such overriding urgency, that he is prepared to accept less restrictions than the rest of the committee. Although both sides now and then tend to picture the opponent in slightly black-and-whitish terms, neither is taking a black and white position itself: Coleman *does* acknowledge that *some* restrictions are appropriate, and the rest of the committee *does* see the battle against terrorism as a highly urgent one. Here the trade-off character of the problem already shows up: the point is not whether the information under discussion does or does not belong to a universally defined private sphere not to be accessed by others, nor can the matter be decided by universal deontological rules of behaviour; rather, the interests of various parties will have to be weighted against each other.

The core of disagreement comes further to the surface in the section where the main report describes what it sees as the most important privacy risks presented by government data mining. Six main categories of risks are distinguished (listed in the order of the report): 'chilling effects and other surveillance risks'; 'data aggregation risks'; 'data inaccuracy risks'; 'false positives'; 'mission creep'; and 'data processing risks'. I prefer to present them in a slightly different order here.

Data inaccuracy risks. This is the only one of the six risks mentioned that refers to merely technical mistakes. It includes things like data errors, mistakenly identifying two persons with the same name (or not recognising different spellings of a name to belong to the same person), etc. These risks are problems of 'direct inference' in the sense discussed earlier.

False positives. This problem was already mentioned above. In this type of case, the conclusion concerning an individual is not wrong because of some technical mistake (that could have been eliminated by more accuracy); rather it results from the fact that when using 'indirect inference', a percentage of 'mistakes' is *necessarily* to be accepted.

Data processing risks. These are uses of data by people with authorised access who nevertheless use these data in a way not intended by the organisation. Examples are misuse, and undue disclosure of information to outsiders.

Mission creep. This term refers to the phenomenon that the goals of an organisation or parts of an organisation may creep away from their original point of reference. Due to such a shift, the organisation may start using the data it possesses for new purposes that were not originally foreseen.

Chilling effects and other surveillance risks. This category contains effects on social relations in terms of general atmosphere. Through cultural chilling, people's behaviour and lives could of course be seriously affected.

Data aggregation risks. These are risks due to the combination of information from different databases, transnational data flows, etc. This category does not seem quite at the same typological level as the other five, since it does not refer to some kind of *effect*; rather, it refers to a number of factors that increase the complexity and the extent of the problem, and thus could aggravate the other risks. The report presents this category as specifically enhancing the chilling effect, but one would say that these are factors also intensify privacy problems more in general.

The types of risks identified in the TAPAC report figure in numerous other writings on the TIA case. Just a few will be mentioned here as illustration. The problem of false positives is almost always present in some form, like in the aforementioned column by Safire,[19] in Taipale's extensive juridical analysis[20], in the reaction by the American Civil Liberties Union,[21] in a letter to two members of the Senate by the US Public Policy Committee of the Association for Computer Machinery[22] (ACM is the largest, US based professional organisation for computer

[18] William T. Coleman, Jr. *Separate statement of William T. Coleman*. in TAPAC, 2004, p. 67.

[19] Safire, 2001.

[20] Kim A. Taipale. Data mining and domestic security: connecting the dots to make sense of data. *Columbia Science and Technology Law Review* 5 (2): 1–83, 2003.

[21] American Civil Liberties Union (ACLU). *Total information compliance: The TIA's burden under the Wyden Amendment. A preemptive analysis of the government's proposed super surveillance program*, May 19, 2003.

[22] US Public Policy Committee of the Association for Computing Machinery (USACM). *Letter to John Warner and Carl Levin*. (resp. chairman and member of the Senate Committee on Armed Services). January 23, 2003.

professionals), and in the analysis by Mary DeRosa.[23] It is interesting to note that another committee of the ACM, worried that data mining might get a bad name, in a public statement pointed out that data mining as such is not against civil liberties.[24] Taipale has also rightly observed that the problem of false positives is not specific to computer searches or data mining, but is characteristic of all criminal investigation.[25] The term 'mission creep' is mentioned explicitly by DeRosa,[26] but themes of similar nature abound elsewhere as well.

When we examine the risks in the TAPAC list, we see that most of them actually are *social* risks, that is, they do not originate in technical or natural phenomena as such, but reside in human behaviour; in this particular case, they relate to the ways in which those who are in control of certain information might deal with that information. The risks that humans will behave in an undesirable way can be rated from a wide range of perspectives, from wildly naive and optimistic to highly critical and suspicious. It is on this axis that a main source of disagreement can be found. Coleman is considerably more at the optimistic side than the rest of the commission. In his minority report several statements can be found that confirm and explain his position in this respect. On the one hand, he points to the high qualifications of the DARPA personnel, and the obviousness (or perhaps even the moral obligation) that one should trust the authorities:

> 'DARPA uses its own highly skilled personnel, but it also seeks out universities and other scientific, intelligence and engineering experts to help it develop new ideas, new technologies which work.'[27]

Rather ironically, he remarks:

> 'Perhaps I am still misled by the fact that in my youth my parents taught me that policemen on the beat and other law enforcement officers are friends, not enemies, and in my life, most often, it has turned out that way.'[28]

On the other hand, he expounds the threat of terrorism, a threat that in his view the other Committee members underestimate:

> 'The report does not sufficiently reflect appreciation of the extent of the security and national defence problems – many novel and new – facing the Nation today in the early stages of the war on terrorism, nor of the important, desirable, beneficial, needed results DARPA was attempting to achieve.'[29] 'The report does not set forth or emphasise sufficiently the nature of the enemy facing us, not one of a foreign nation state, not one easily identified by uniform, but one who wears civilian clothes, one who purposely mixes with innocent US persons, using their facilities (banks, airplanes, flying schools, etc.) and who otherwise immerses himself or herself in our free, open, friendly society (...)'[30] 'There is fair evidence that if governmental officials, agents and employees had real-time, meaningful access on September 10, 11, 2001, to the knowledge then in governmental files – all such information being originally obtained legally, and there was no statutory restriction of which governmental agents could look at it – the terrorist attacks of September 11 probably could have been prevented.'[31]

Coleman also expresses his disagreement with respect to the analysis of the origin of the controversy. The main report states that

> '[TIA] was flawed by its perceived insensitivity to critical privacy issues, the manner in which it was presented to the public, and the lack of clarity and consistency with which it was described. DARPA stumbled badly in its handling of TIA, for which the agency has paid a significant price in terms of its credibility in Congress and with the public.'[32]

The report also mentions the slightly sinister big-brotherish appearance of the project's original naming ('Total Information Awareness'), and of the IAO's logo (a pyramid with an al seeing eye on top) and the accompanying motto 'Scientia est potentia' (science is power).[33] Coleman finds some of these judgements about DARPA and TIA too negative:

> '(...) I think failure came about because the Congress did not really understand the TIA program and did not appreciate that it was a research venture rather than the use, operation and application

[23] Mary DeRosa. *Data mining and data analysis for counterterrorism*. Center for Strategic and International Studies (CBIS) Report, Washington, DC, 2004.

[24] Executive Committee Association For Computing Machinery, Special Interest Group on Knowledge, Discovery & Data Mining (ACM SIGKDD). *'Data mining' is NOT against civil liberties*. June 30, 2003 (revised July 28, 2003).

[25] Taipale, 2003, p. 69.

[26] DeRosa, 2004.

[27] Coleman, 2003, p. 73.

[28] Coleman, 2003, p. 74.

[29] Coleman, 2003, p. 67.

[30] Coleman, 2003, p. 67.

[31] Coleman, 2003, p. 71.

[32] TAPAC, 2003, pp. viii.

[33] TAPAC, 2003, p.18.

of what such research showed would be worthwhile to develop, produce and put into use by others. Congress ignored that privacy issues and protection thereof were also being considered by TIA and DARPA in the research. This in part occurred because DARPA did not explain the project to Congress in the informing way one would expect of DOD. To that extent, DARPA is at fault.'[34]

Once more the core issue seems to be whether or not one is prepared to trust that the TIA personnel and the TIA organisation will avoid the risks mentioned above, here with the emphasis on the extent to which DARPA's communicative strategy provides reasons either for trust or for doubt. Coleman apparently believes that brief information should be sufficient, whereas the main report suggests that only much more elaborate information and a thorough address of privacy concerns could have produced sufficient reassurance.

It is interesting to compare Coleman's remarks on this issue with DoD's own report, that was published somewhat later, in December 2003. The following paragraph is sufficiently significant to justify full quoting:

> Although the DARPA development of TIA-type technologies could prove valuable in combating terrorism, DARPA could have better addressed the sensitivity of the technology to minimise the possibility for Governmental abuse of power and to help ensure the successful transition of the technology into an operational environment. Several factors contributed to the condition.
>
> - DARPA did not implement the best business practice of performing a privacy impact assessment.
> - USD (ATL) initially provided limited oversight of the TIA development and did not ensure that DARPA included in the effort the appropriate DoD policy, privacy, and legal experts.
> - DARPA efforts historically focused on development of new technology rather than on the policies, procedures, and legal implications associated with the operational use of technology.
> - The DARPA position was that planning for privacy in the operational environment was not its responsibility because TIA research and experiments used synthetic artificial data or information obtained through normal intelligence channels.
>
> As a result, DoD risks spending funds to develop systems that may not be either deployable or used to their fullest potential without costly revision.'[35]

Apparently, DARPA largely agrees with the main report on this issue.

Social risks and their analysis in terms of 'systems of subliminal enticement'

It was already observed that the risks outlined by the TAPAC committee are to a major extent *social risks*. 'False positives', 'data processing risks' and 'mission creep' can be conceived as the *core social risks*, whereas the other risks are either social but derived (chilling), social but enhancing rather than primary ('data aggregation risks') or technical ('data inaccuracy risks'). Social risks are not due to technical or natural inevitabilities, but to the ways in which people who are in control of certain information might deal with that information. Such risks can be generally characterised by what I have elsewhere[36] called 'systems of subliminal enticement'. Briefly summarised, the idea is that unethical or undesirable behaviour often takes place in the presence of the following three elements:

- *enticement*: the unethical or undesirable options present an enticement to the actor, such as the enticement to optimise one's individual interest rather than the common interest, or to sacrifice long term considerations for short term satisfaction.
- *subliminal*: the actor does not want to face that the option chosen is less desirable or less ethical, and devises all kinds of excuses and rhetoric to conceal this, to him(her)self and to others.
- *systems*: there are systemic dependencies between different actors that lead to prisoners' dilemmas, tragedies of the commons and similar configurations that enhance choices sacrificing the common good for individual benefit.

Systems of subliminal enticement are at stake in all social risks mentioned above. If we confine ourselves to a few illustrations from the core social risks:

- *false positives*: Will the interests of those persons who will be falsely identified as positives not be

[34] Coleman, 2003, p. 81.

[35] Office of the Inspector General, Department of Defense. *Information technology management terrorism information Awareness program (D-2004-033)*. Department of Defense, December 2003, p. 4.

[36] Frans A.J. Birrer, 2000.

sacrificed to the wish to make searches as wide as possible? (enticement) Will those within the organisation who would plead for more restricted searches not easily be put aside as not being sufficiently 'tough' or 'committed' (rhetoric, prisoners' dilemma)?
- *data processing risks*: Will organisation members not misuse or leak information? (enticement) Will they not be inclined to justify such actions for themselves by excuses like 'Others also do it.' or 'If I don't do it someone else will'? (rhetoric, prisoners' dilemma)
- *mission creep*: Will there not be a shift in the organisation's goals and culture such that the balance between security and liberty more and more becomes dominated by the first?

In more general terms the social risks originate in an intertwinement of the sincere concern about terrorism on the one hand, and individuals seeking personal opportunities, organisations and parts of organisations craving to expand in terms of size, territory and power, on the other. Each of these motives drives towards a perspective in which the consequences for innocent individual citizens and for society in general easily become downplayed or outright ignored. In the resulting competition for being the most ambitious and the most undaunted fighter against terrorism, those members of the organisation who express doubts are likely to be put aside as 'not tough enough'. Strong group think may obliterate critical or balancing inputs. The high degree of individual and group autonomy that is the natural mode of organisation for enterprises like TIA, further increases the risk of unperceived gradual shifts towards deviation from the originally intended goals and morale. Numerous complexities, such as the lack of a uniform definition of the notion of terrorism,[37] and the vague and slippery nature of dividing lines such as that between subject based searches (starting from individuals who are already subject) and pattern based searches[38] all add to a level of untransparency that provides an eager substrate for processes such as mission creep.

One may ask, of course, whether the Congressional decisions were the right ones to block these undesirable processes. Taipale[39] believes that by terminating the funding of TIA, Congress in fact destroyed the opportunity that a program would be developed while held accountable to the Congress, a program that included the development of privacy enhancing techniques (in the Genesys subproject).

Warrants for trust

The notion of systems of subliminal enticement incorporates a number of concerns that are not whole-heartedly pleasant to face. Most of us would no doubt prefer that we could assume that everyone or at least almost everyone naturally contributes to the common good. Most of us probably believe of ourselves that we generally aim to do so. But good intentions do not necessarily work out right collectively. More importantly, dubious motives can be concealed with seemingly impeccable ones; in such situations, it can be hard to prove whether an individual is motivated by the common good or by individual benefits. It is therefore important to notice that I have not proposed the analysis of systems of subliminal enticement primarily as a means to put moral blame on individual actors (although for extreme, clear-cut cases this may be a consequence, as it would be of any ethical analysis); its main aim lies at a meta-level, namely to analyse the structure and content of interaction and communicative exchanges in order to identify and seeking to replace those structures unnecessarily prone to such undesirable effects. It is a framework for *analysing social risks*, for the purpose of *prevention* rather than *attributing blame after the event*.[40] If blame is to be attributed, it should be for conducting organisational practices that are needlessly vulnerable to such risks, because of a failure to consider them in advance.

Analysis of social risks is relevant and necessary, not only for TIA or for data mining, but in all those cases where decisions have potentially harmful consequences for society or for groups of individuals, especially when there are barriers for those potentially affected individuals to identify in advance the risks that they are exposed to, and to influence the decisions that generate those risks. Such barriers particularly exist when decisions are based on expert advice (be it scientific experts or otherwise persons who have knowledge that is not easily assessed by outsiders), and of course also when there are security reasons for not revealing every detail of what is going on, as in the case of TIA. We see this type of issues increasingly emerging in public discussion today. The controversies that arose in various countries on the

[37] Yonah Alexander. Terrorism: a definitional focus, in Yonah Alexander, Edgar H. Brenner, *Terrorism and the law*, Transnational Publishers, Ardsley (NY), 2001 David J. Whittaker (ed.). *The terrorism reader*. Routledge, London, 2001.

[38] TAPAC, 2003, p. 45.

[39] Taipale, 2003.

[40] Cf. Thompson , who discusses privacy concerns in terms of security. Paul B. Thompson. Privacy, secrecy and security. *Ethics and Information Technology*, 3: 13–19, 2001.

reliability and sincerity of governmental information concerning the Iraq war, and the manipulation of information in stock market and other financial scandals are just a few examples of this. Like in the case of TIA, such events may trigger severe public outrage, that is hard to put to rest.

What is desperately needed in such cases where not everything that is going on can be scrutinised by everyone is, of course, trust. But it would be ill advised to ask for *blind* trust here. What we need are *warrants for trust*, that is, clearly articulated and actively implemented policies that cover the most important social risks, on the basis of a systematic analysis of these risks. This leaves the question whether it will always be possible to define policies that suffice as warrants for trust. That in part depends on the amount of 'social risk' that we are prepared to run. And that, of course, is a matter for public deliberation. What is clear is that the current practice does *not* always suffice to put the public debate and the public resistance at rest. New equilibria will have to be found.

Summarising the general arguments

In this paper, I started by explaining that the notion of privacy, as it is used in public debate today, cannot be restricted to the question who has or has not certain information on other individuals. At least as important is how this information will be used.

Second, this use often involves correlative inferences, which bring in specific ethical issues that arise when dealing with uncertainty, such as false positives and false negatives.

Third, as a consequence, the main risks involved are social risks, that is, risks pertaining to how those people who control or have access to information will actually deal with that information.

Finally, privacy will only be deemed sufficiently guaranteed in as far as the institution that controls certain information and its use can show a system of measures that is likely to contain these risks to such a degree that it is found acceptable by the public.

References

ACLU (American Civil Liberties Union). *Total information compliance: The TIA's burden under the Wyden Amendment. A preemptive analysis of the government's proposed super surveillance program*, May 19, 2003.

Y. Alexander. Terrorism: a Definitional Focus. In Y. Alexander and E.H. Brenner, editors, Terrorism and the Law, . Transnational Publishers, Ardsley, NY, 2001.

F.A.J. Birrer. Applying Ethical and Moral Concepts and Theories to IT Contexts: Some Key Problems and Challenges. In R.A. Spinello and H.T. Tavani, editors, Readings in Cyberethics, pp. 91–97. Jones & Bartlett Computer Science, Sudbury, MA, 2001.

F.A.J. Birrer. Statistical Evidence: Responsibilities and the Burden of Proof. In C. van Dijkum , J. Blasius, H. Kleijer and B. van Hilten, editors, Recent Developments and Applications is Social Research Methodology, . SISWO, Amsterdam, 2004.

F.A.J. Birrer. Computer Technology, Subliminal Enticement, and the Collectivisation of Ethics. In D.G. Johnson, J.H. Moor and H.T. Tavani, editors, Computer Ethics: Philosophical Enquiry (CEPE 2000 Proceedings). Dartmouth College, Dartmouth, 2000.

W.T. Coleman, Jr. *Separate Statement of William T. Coleman*. TAPAC, 2004.

M. DeRosa. *Data Mining and Data Analysis for Counterterrorism*. Center for Strategic and International Studies (CBIS) Report, Washington, DC, 2004.

DARPA (Defense Advanced Research Projects Agency). *Report to Congress regarding the Terrorist Information Awareness Program (in response to Consolidated Appropriations Resolution, 2003, Pub. L. No. 108-7, Division M, § 111 (b))*, May 20, 2003.

Executive Committee ACM SIGKDD (Association For Computing Machinery, Special Interest Group on Knowledge, Discovery & Data Mining). *'Data mining' is NOT against civil liberties*. June 30, 2003 (revised July 28, 2003).

GAO (General Accounting Office). *Data mining. Federal efforts cover a wide range of uses*. GAO Report 04-548, GAO, Washington, DC, 2004.

D. Hand, H. Mannila and P. Smyth, *Principles of Data Mining*. MIT Press, Cambridge, Mass., 2001.

N. Kranich, *Commentary: MATRIX and the new surveillance states: the multistate anti-terrorism information exchange*. Free Expression Policy Project, October 16, 2003.

J. Mena. *Homeland Security. Techniques and Technologies*. Charles River Media, Hingham, Mass., 2004.

Office of the Inspector General, Department of Defense. *Information technology management: Terrorism Information Awareness program (D-2004-033)*. Department of Defense, December 2003.

W. Safire, You are a suspect. *New York Times*, November 14, 2002, p. A 35.

TAPAC (Technology and Privacy Advisory Committee). *Safeguarding Privacy in the Fight Against Terrorism*. Department of Defense, Washington, DC, 2004.

K.A. Taipale. Data Mining and Domestic Security: Connecting the Dots to Make Sense of Data. *Columbia Science and Technology Law Review*, 5(2): 1–83, 2003.

P.B. Thompson. Privacy, Secrecy and Security. *Ethics and Information Technology*, 3 13–19, 2001.

USACM (US Public Policy Committee of the Association for Computing Machinery). *Letter to John Warner and Carl Levin*. (resp. chairman and member of the Senate Committee on Armed Services). January 23, 2003.

D.J. Whittaker (Ed.) *The Terrorism Reader*, Routledge, London, 2001.

I.A. Witten and E. Frank, *Data Mining*. Morgan Kaufmann, San Francisco, 2000.