The NRC Takes on Data Mining, Behavioral Surveillance, and Privacy

Susan Landau
Sun Microsystems

In mid 2000, *The Wall Street Journal* reported that the US Federal Bureau of Investigation (FBI) was developing a tool for wiretapping at an Internet service provider (ISP).[1] Carnivore, later renamed DCS 1000, was built to capture communications content—email, Web pages, and so forth—or the transactional information in the communications of targeted suspects. The project raised much concern. What exactly was Carnivore acquiring—communications only of targeted suspects or messages between nontargeted people as well? IP headers typically reveal some communications content, so what was Carnivore collecting when executing pen registers or trap-and-trace orders, which in the telephone world were simply signaling information?

The US Department of Justice authorized a study of Carnivore's "technical" aspects, insisting on the right to edit the report before release. The best-known university security research groups declined to participate; instead, a team at the less well-known Illinois Institute of Technology examined the FBI's wiretapping tool. Their report[2] raised concerns about potential overcollection and did little to alleviate public concern. In mid 2001, it seemed likely that Congress would seek to limit Carnivore's use. In response, Attorney General John Ashcroft announced that a senior Department of Justice official, Daniel Collins, would study the privacy risks raised by DCS 1000.

Collins's report never appeared. Instead, in the seven years after September 11, 2001, privacy protection took a back seat to the US effort of defending the country against further terrorist attack. Indeed, the aspects of DCS 1000 that most concerned civil-liberties groups—in particular, the potential overcollection of transactional data—was made legal under §216 of the USA PATRIOT Act. In *Protecting Individual Privacy in the Struggle Against Terrorists: A Framework for Program Assessment*,[3] the National Research Council (NRC) examines two technologies used to combat terrorism—data mining and behavior and psychological surveillance (the latter involves monitoring involuntary actions to reveal a person's "true" thoughts)—and their impact on privacy. The study highlights the issue of effectiveness, and carefully but powerfully makes the point that data mining and behavioral surveillance programs that fail the effectiveness test protect neither the nation's privacy nor its security. This report comes none too soon.

The NRC is the research arm of the US National Academies of Science and Engineering and the Institute of Medicine. It produces reports on current concerns at the intersection of science, technology, and policy. These reports don't present new research per se—rather, they synthesize other work to provide direction to the government on how to address complex scientific and technological issues. The studies are typically insightful but have had mixed impact; some gather dust on shelves while others, such as the 1996 report on cryptography policy,[4] have had a substantial effect on national policy. This new report tackles an important issue on the boundary between security and civil liberties; it does so succinctly and with insight, and it provides valuable and important recommendations. It should be read, paid attention to, and followed.

The Major Recommendations

Changes in societal habits—the ubiquitous electronic records we leave behind in every action we perform—combined with the decreasing costs of computer storage and increasing capabilities of search tools have made data mining, the analysis of massive data sets for "interesting" patterns, an irresistible tool of government antiterrorism efforts. The use of these tools by the private sector

for fraud detection and sales tracking and by the government for counterterrorism purposes has exploded in recent years. The NRC report examines the role of data mining and behavioral surveillance technologies in counterterrorism programs and presents a framework for the government to use in evaluating these programs.

Above all, this report provides the common sense so far lacking in appraising these government counterterrorism programs. As mentioned earlier, it homes in on effectiveness, an issue that should have been central in discussions about such programs as the Transportation Security Administration's (TSA's) no-fly list, DARPA's Total Information Awareness, and the Department of Homeland Security's (DHS's) Analysis, Dissemination, Visualization, Insight, and Semantic Enhancement (ADVISE) programs—but which, instead, appears to have been completely ignored [NRC report, p. 88].

Protecting Individual Privacy in the Struggle Against Terrorists relies heavily on an earlier NRC report on privacy[5] and makes two major policy recommendations:

- US government agencies should have a systematic process in place to examine the effectiveness, lawfulness, and consistency with US values of every information-based program for detecting and countering terrorism before such a system can be deployed and periodically thereafter, especially when the program enters a new phase in deployment.
- The US government should periodically assess whether the nation's laws, policies, and procedures protecting individual privacy are sufficient given new technologies [NRC report, pp. 5–6].

When programs are expanded, the effect on citizens' civil liberties is often ignored, making the issue of "mission creep"—extending a program to new purposes after its original deployment—particularly important. That the NRC report has called this point out in its major recommendations is invaluable.

The fact that the report's co-chairs are Charles Vest, president of the National Academy of Engineering and president emeritus of the Massachusetts Institute of Technology, and William Perry, former Secretary of Defense—and that other participants included members of the law enforcement and national security communities, lawyers, terrorism experts, and research scientists from Google and Microsoft—should increase the report's impact. Attention should be paid, but the question is, will it?

The Main Issues

Protecting Individual Privacy in the Struggle Against Terrorists begins by clarifying data mining, splitting it into two types: subject-based data mining, which follows the digital tracks of someone whose data has appeared somewhere "interesting" (such as a license plate near a terrorist bombing), and pattern-based data mining, which uses pattern-matching techniques to uncover odd behaviors that might lead to "persons of interest" [NRC report, pp. 22–23]. The former, which is nothing more than a modern technological enhancement of older law enforcement methods, doesn't raise the privacy concerns that the latter technique does. Consequently, the NRC report focuses on privacy concerns raised by pattern-based data mining and behavioral profiling.

Trenchant Observations

As the report examines the basis for making decisions regarding data mining and behavioral profiling programs, it presents a number of important observations. Many of these have rarely been part of the public discussion on counterterrorism efforts:

- "[T]here is no way to make personal information in databases fully anonymous" [NRC report, p. 4]. That is to say, technical solutions alone cannot solve the privacy problems created by data mining.
- "[T]he technical reality is that the number of false negatives can never be zero" [NRC report, p. 40]. Once it's acknowledged that no predictive system can catch every potential terrorist—an obvious conclusion—the issue of effectiveness and the need for such evaluation of all programs becomes clear.
- "The threshold consideration of any privacy-sensitive technology is whether it is effective towards a clearly defined national security or law enforcement purpose. The question of effectiveness must be assessed through rigorous testing guided by scientific principles" [NRC report, p. 42]. Again, this is an obvious point (but one that has been ignored for much of the past decade).
- "Too frequently the argument is heard that national security is too important and the terrorist threat too great to pause to ask hard questions of the systems to be deployed to protect the nation. In the committee's view, that is the wrong approach. It is precisely because national security is important and the threats to it are great that it is so important to ensure that the systems to be deployed to protect the nation are effective and are consistent with US values" [NRC report, p. 47]. Programs enacted without proper consideration of their value and effectiveness are solutions that are ultimately unworkable; they erode privacy without providing security.
- "When such [an information-based] system uses personally identifiable information or otherwise affects privacy, the

documentation should be examined by an entity, such as an independent scientific review committee, that is capable of evaluating the scientific evidence of effectiveness outside the agency promoting the new system" [NRC report, p. 51]. Such dispassionate evaluation is clearly necessary. Without it, the government will fund programs of dubious merit. Some of those, which impinge upon civil rights, may even add to security risks by alienating certain groups of citizens. As the report observes, scientific evaluation is critical in determining a program's value, yet no evidence has come to light that any such evaluations have been taking place.

- "[B]ecause technology and events usually outpace law, it is necessary to constantly consider what types of information-based programs should be lawful. In short, are they consistent with the values of US society?" [NRC report, p. 52]. The 1974 Privacy Act, which deals with data that is part of "a system of records," is woefully out of date. But even recent law on surveillance can be difficult to interpret in the face of such technologies as voice over IP, massively multiplayer online role-playing games, and the like—thus, the report's second major recommendation emphasizing the need for reviewing and updating the nation's privacy laws. A longer 2007 NRC study on privacy[5] details many examples of the need to modernize US privacy law. The current NRC report, which, at just over 100 pages, is more narrowly focused, includes an appendix with an excellent description of how law has lagged behind technology [NRC report, Appendix F].
- "[Total Information Awareness]-style data mining was, and still is, possible because there are few restrictions on government access to third-party business records" [NRC report, p. 248]. The ability of the US government to legally buy citizenry data from third-party brokers that it can't amass itself is probably the largest gaping hole, of the many, in the outdated 1974 Privacy Act. Because of the vast amount of information available from these data brokers, the US government is in a position to obtain essentially all information about individuals without going through the courts [NRC report, p. 57].
- "Data of poor quality limit the value of data mining in a number of ways" [NRC report, p. 74]. The reason is that, "error-prone data are, of course, both a threat to privacy (as innocent individuals are mistakenly associated with terrorist activity) and a threat to effectiveness (as terrorists are overlooked because they have been hidden by errors in the data that would have suggested a terrorist connection)" [NRC report, p. 77]. Effectiveness will improve law enforcement and national security efforts while simultaneously protecting privacy.

The report, which includes detailed appendices, has many other acute observations, far too numerous to list here.

Above all, this report provides the common sense so far lacking in appraising these government counterterrorism programs.

Exploring That Which Is Only Thought

Nearly a century ago, US Supreme Court Justice Louis Brandeis wrote in a dissenting opinion in a wiretapping case,[6]

> Ways may some day be developed by which the government, without removing papers from secret drawers, can reproduce them in court, and by which it will be enabled to expose to a jury the most intimate occurrences of the home. Advances in the psychic and related sciences may bring means of exploring unexpressed beliefs, thoughts and emotions.

That time has come. Of course, for years, law enforcement and national security have relied on lie detectors despite their dubious ability to determine whether a suspect is telling the truth.[7] But more recently, behavioral surveillance has received increased attention; such surveillance includes research in monitoring facial expressions for involuntary muscle movement that betrays a subject's "real" thoughts, research in vocalization (including pitch, timbre, tempo, and so on) to determine a subject's emotional state, observing a subject's autonomic nervous system to detect deception, and so forth. Some of this monitoring could, in fact, reveal a subject's real thoughts. The same concerns surrounding data mining are much stronger regarding behavioral surveillance: is the monitoring valid, accurate, effective? Whereas data mining simply uses previously collected data and analyzes it, behavioral surveillance collects new data in a highly intrusive manner [NRC report, pp. 250–262]. Thus, the questions of validity, accuracy, and effectiveness are even more significant in this domain.

And in Practice

NRC studies have a remarkable ability to pull together

government documents that clarify what's happening behind the scenes. By briefly discussing several surveillance programs, including Computer-Assisted Passenger Prescreening System II (CAPPS II), the Multistate Anti-Terrorism Information Exchange (MATRIX), Able Danger, ADVISE, the Automated Targeting System, the Electronic Surveillance Program (also known as the Terrorist Surveillance Program), the Novel Intelligence from Massive Data Program, the Enterprise Data Warehouse, the ICE Pattern Analysis and Information Collection System, Intelligence and Information Fusion, the Fraud Detection and National Security Data System, and the Financial Crimes Enforcement Network (FinCEN), this report accomplishes such clarification in a particularly murky program area. A pattern of privacy violations emerges from the NRC compendium—that is, although the TSA stated that it wouldn't use data mining techniques in CAPPS II, the program included checking several databases (a distinction that the public would be hard-pressed to understand) [NRC report, p. 220], nothing in the MATRIX system prevented the monitoring of political activities [NRC report, p. 224].

The report's focus was on policy, but it also touched on various technical concerns. One of particular interest to this magazine's readers is the distance between academic-style research on privacy-enhancing technologies, which have focused on ideal cases, and the problems actually faced in the real world. Thus, for example, *k*-anonymity is of little use in protecting privacy if the user is allowed to aggregate data from multiple databases [NRC report, p. 245]. This ties back to the committee's central concern: effectiveness. How effective are privacy protection tools? Effectiveness is key to ensuring the development of good technology, privacy protection, and good security.

The Framework for Assessing Effectiveness

The report's authors didn't feel that a bureaucratic checklist for privacy would be useful. The US government's Privacy Impact Assessments—annual assessments by federal agencies on whether data-handling procedures conform to laws, regulations, and policies regarding privacy—have little actual impact on privacy, suggesting the limited value of such an approach. (Indeed, the report discusses the DHS's ADVISE program, which the DHS Privacy Office didn't include in a report on data mining: the Privacy Office "considered [ADVISE] a tool or technology and not a specific implementation of data mining."[8]) Instead, the committee provided a "framework" that people at all levels of government—judges, policy makers and implementers, and legislators—as well as the public and the press could use in evaluating counterterrorism programs. With technology outpacing law in this domain, it's crucial that new programs respect tomorrow's laws as much as today's. Thus, two questions lie at the heart of this framework: is the information-based program effective? And does it comply with the law and societal values, especially in protecting the data subject's civil liberties? [NRC report, p. 46]

I've included from the report the following abbreviated list of measures by which the information-based programs should be judged:

- a clearly stated purpose for the information-based program;
- a sound rational basis for the program and its components;
- a sound experimental basis for the program and its components;
- the ability to scale;
- existence of a clear operational and business process for the program;
- a ready ability to interoperate with systems and tools inside and outside organizational boundaries;
- robust (that is, reliable in the field as well as in the face of user error and/or countermeasures);
- appropriate and accurate data used in the program;
- data used properly handled and accounted for (that is, appropriate "data stewardship");
- objective systems testing and evaluation;
- ongoing assessments; and
- documented evidence of the program's effectiveness and compliance with key requirements [NRC report, pp. 48–51].

The framework follows with a set of criteria for evaluating the agency, the program, and the data for adherence to US laws and values, as well as a similar set of criteria for developing new laws and policies [NRC report, pp. 59–61].

The framework's merit lies in its simplicity—and its emphasis on program effectiveness and adherence to US values—and the fact that it has no classified arcana hanging around the edges. By designing around values, the framework applies to today's laws as well as tomorrow's.

Can the Framework Work?

The proposed framework for evaluation of such programs is good, but who will actually see to its implementation? What branch of government will make such evaluations happen? Unfortunately, the report drops the ball and gives no practical recommendations for implementation, a major gap in an otherwise excellent piece of work. There's simply a lack of thinking through the practicalities of how the US government might carry out the proposed work.

One possibility is that Congress passes legislation requiring that any data mining or behavioral

profiling program be funded only if the funding agency has implemented an evaluation framework. An administrative route to achieve the same ends might be an Office of Management and Budget (OMB) requirement that agencies perform such evaluations prior to funding programs. The latter approach might be simpler to achieve: it would take only the director of the OMB to agree to it to make it happen (rather than a majority of the US House and Senate in the case of a legislative solution). In any case, it behooves the NRC committee to take action to see that its recommendations are actually carried out; its proposals are too important to simply collect dust.

I didn't agree with everything I read in the report—for example, the comment that the single most powerful motivation for terrorism is revenge [NRC report, p. 115] seemed speculative to me, and the claim that the anonymity of the Web enabled recruitment of women to terrorist Islamic activities needed more explanation in view of the participation of Iraqi women as suicide bombers [NRC report, p. 118]. But these are minor quibbles. The report hits the big issues, does so with clarity and insight, and has made some unequivocal and valuable recommendations about how such data mining and behavioral profiling programs should be conducted. The framework's simplicity and emphasis on values are extremely useful and should help with adoption. I see only one major gap in the report—namely, the lack of a practical road map for enacting the report's recommendations.

The report is balanced and authoritative, and I also found it quite readable. It isn't written in the chatty style that Bruce Schneier employs in his analyses of security, but neither does it use turgid legal prose (even when discussing legal issues).

The past seven years have seen a remarkable erosion of civil liberties in the name of national security. This NRC analysis, written by a balanced panel well versed in security concerns, asks not what we can do to balance privacy and security while protecting ourselves against terrorists, but asks instead what we can do to achieve security *and* privacy while protecting ourselves. This report should be read by researchers, law enforcement and national security, legislators and their staff and—most importantly—by the new administration. And then it should be implemented. □

References

1. N. King Jr. and T. Bridis, "FBI's Wiretaps to Scan E-Mail Spark Concern," *The Wall Street Journal*, 11 July 2000, p. A3.
2. S.P. Smith et al., *Independent Technical Review of the Carnivore System*, IIT Research Inst., 8 Dec. 2001.
3. US Nat'l Research Council, *Protecting Individual Privacy in the Struggle Against Terrorists: A Framework for Program Assessment*, Nat'l Academies Press, 2008.
4. K. Dam and H. Lin, *Cryptography's Role in Securing the Information Society*, Nat'l Academies Press, 1996.
5. J. Waldo, H. Lin, and L. Millet, *Engaging Privacy and Information Technology in a Digital Age*, Nat'l Academies Press, 2007.
6. *Olmstead v. United States*, 277 U.S. 438 (1928), Brandeis, Louis, dissenting, p. 473.
7. US Nat'l Research Council, *The Polygraph and Lie Detection*, Nat'l Academies Press, 2003.
8. *Data Mining: Early Attention to Privacy in Developing a Key DHS Program Could Reduce Risks*, GAO-07-293, US Government Accountability Office, Feb. 2007, p. 3.

Susan Landau *is a distinguished engineer at Sun Microsystems, where she works on security, cryptography, and public policy, including surveillance issues, digital rights management, and identity management. She is coauthor (with Whitfield Diffie) of* Privacy on the Line: the Politics of Wiretapping and Encryption, *updated and expanded edition (MIT Press, 2007). Landau has a PhD in theoretical computer science from MIT. She is the 2008 winner of the Anita Borg Institute for Women and Technology Women of Vision—Social Impact, a fellow of the American Association for the Advancement of Science, and an ACM distinguished engineer. Contact her at susan.landau@sun.com.*