


Countering Terrorism through Information and Privacy Protection Technologies

Security and privacy aren't dichotomous or conflicting concerns—the solution lies in developing and integrating advanced information technologies for counterterrorism along with privacy-protection technologies to safeguard civil liberties. Coordinated policies can help bind the two to their intended use.



ROBERT POPP
*National
Security
Innovations*

JOHN
POINDEXTER
*JMP
Consulting*

The terrorist attacks of September 11, 2001 transformed America like no other event since Pearl Harbor. The resulting battle against terrorism has become a national focus, and “connecting the dots” has become the watchword for using information and intelligence to protect the US from future attacks.

Advanced and emerging information technologies offer key assets in confronting a secretive, asymmetric, and networked enemy. Yet, in a free and open society, policies must ensure that these powerful technologies are used responsibly and that privacy and civil liberties remain protected. In short, Americans want the government to protect them from terrorist attacks, but fear the privacy implications of the government's use of powerful technology inadequately controlled by regulation and oversight. Some people believe the dual objectives of greater security and greater privacy present competing needs and require a trade-off; others disagree.¹⁻³

This article describes a vision for countering terrorism through information and privacy-protection technologies. This vision was initially imagined as part of a research and development (R&D) agenda sponsored by DARPA in 2002 in the form of the Information Awareness Office (IAO) and the Total Information Awareness (TIA) program. It includes a critical focus and commitment to delicately balancing national security objectives with privacy and civil liberties. We strongly believe that the two don't conflict and that the ultimate solution lies in utilizing information technologies for counterterrorism along with privacy-protection technologies to safeguard civil liberties, and twining them together with coordinated policies

that bind them to their intended use.

Background and motivation

Terrorists are typically indistinguishable from the local civilian population. They aren't part of an organized, conventional military force—rather, they form highly adaptive organizational webs based on tribal or religious affinities. They conduct quasi-military operations using instruments of legitimate activity found in any open or modern society, making extensive use of the Internet, cell phones, the press, schools, houses of worship, prisons, hospitals, commercial vehicles, and financial systems. Terrorists deliberately attack civilian populations with the objective to kill as many people as possible and create chaos and destruction. They see weapons of mass destruction not as an option of last resort but as an equalizer—a weapon of choice.

Of the numerous challenges to countering terrorism, none are more significant than being able to detect, identify, and preempt terrorists and terrorist cells whose identities and whereabouts are unknown a priori. (Alan Dershowitz's *Preemption: A Knife that Cuts Both Ways* [W.W. Norton & Company, 2006] offers an extensive discussion of preemption and the need for a legal structure.) In our judgment, if preemption is the goal, the key to detecting terrorists is to look for patterns of activity indicative of terrorist plots based on observations of current plots and past terrorist attacks, including estimates about how terrorists will adapt to avoid detection. Our fundamental hypothesis is if terrorists plan to launch an attack, the plot must involve people (the terrorists, their financiers, and so forth). The transactions all these people

conduct will manifest in databases owned by public, commercial, and government sectors and will leave a signature—detectable clues—in the information space. Because terrorists operate worldwide, data associated with their activities will be mixed with data about people who aren't terrorists. If the government wants access to this activity data, then it must also have some way to protect the privacy of those who aren't involved in terrorism.

This hypothesis has several inherent critical challenges. First, can counterterrorism analysts imagine and understand the numerous signatures that terrorist plans, plots, and activities will create? Second, if they do understand these signatures, can analysts detect them when they're embedded in a world of information noise before the attacks happen (in this context, *noise* refers to transactions corresponding to nonterrorists)? Finally, can analysts detect these signatures without adversely violating the privacy or civil liberties of nonterrorists? Ultimately, the goal should be to understand the level of improvement possible in our counterterrorism capabilities if the government could use advanced information technologies and access a greater portion of the information space; but also consider the impact—if any—on policies such as privacy, and then mitigate this impact with privacy-protection technology and corresponding policy.^{2,3}

Countering terrorism

Information technology plays a crucial role—and is a major tenet—of our counterterrorism strategy because it ultimately has to make sense out of and connect the relatively few and sparse dots embedded within the massive amounts of information potentially available to, and already flowing into, the government's intelligence and counterterrorism agencies.

Numerous information technologies can help intelligence analysts detect and understand the clues terrorists leave behind when plotting their next move. In the simplest terms, these technologies fall into one of two broad categories: collections and analytics. Figure 1 provides a simple illustration of this counterterrorism framework. For collections, we won't discuss the vast array of sensor technologies that fall within this category here; instead, see Table 1 (p. 27), which provides a sample of the authorization provided to the US intelligence community for its foreign and domestic intelligence and counterintelligence data collections.

For analytics, key intelligence tools include collaboration; text analysis and decision aides; natural language processing (in particular, speech-to-text transcription and foreign-to-English translation); pattern analysis; and predictive (anticipatory) modeling. These technologies help analysts create models (and discover instances of new models) of terrorist activity patterns; search and exploit vast amounts of multimedia, multiformat, and multilingual speech and text; extract entities and entity relationships

from massive amounts of data; collaborate, reason, and share information and analyses so that analysts can hypothesize, test, and propose theories and mitigating strategies about plausible futures; and advise decision- and policy-makers on the impact of current or future policies and prospective courses of action. We don't discuss these technologies in detail here, but more information appears elsewhere.^{1,4,5}

In our view, modeling tools play a crucial role in countering terrorism. The analytical community first creates scenarios of terrorist plots and attacks using previous attacks, intelligence reports, red teams, war games, table-top exercises, and the like. These terrorism scenarios would consist of a range of transactions and steps that terrorists must perform in support of their plot to attack a specific target type using a specific mode of attack. Analysts then codify these scenarios in a set of quantitative and computational models based on a wide range of nonlinear mathematical and nondeterministic stochastic computational approaches for capturing social phenomena and pathological behavior. These models are essentially hypotheses about terrorist plots and would be translated into a series of questions about the types of transactions terrorists would need to execute, the types of evidence analysts would need to accrue, the keywords and patterns analysts would need to associate, and the like.

Terrorist activity isn't easily reduced or amenable to classical analytical methods; moreover, the associated data can be incredibly poor due to ambiguous, erroneous, and conflicting reports. No single theory or modeling approach is sufficient, so we must integrate an ensemble of models that have more information than any single model has to estimate a range of plausible futures and provide competing explanations as to what the information means. Robust adaptive strategies that hedge across these plausible futures will provide practical actionable options for the decision-maker to consider.^{1,4,5}

Early results show promise

The importance (and promise) of these information technologies has already emerged through experiments conducted with several entities in the intelligence com-

Numerous information technologies can help intelligence analysts detect and understand the clues terrorists leave behind.

munity. Experiments let us assess these technologies for utility and merit in the context of real-world problems before large amounts of funds are expended to fully implement them. Moreover, to push the envelope of what's

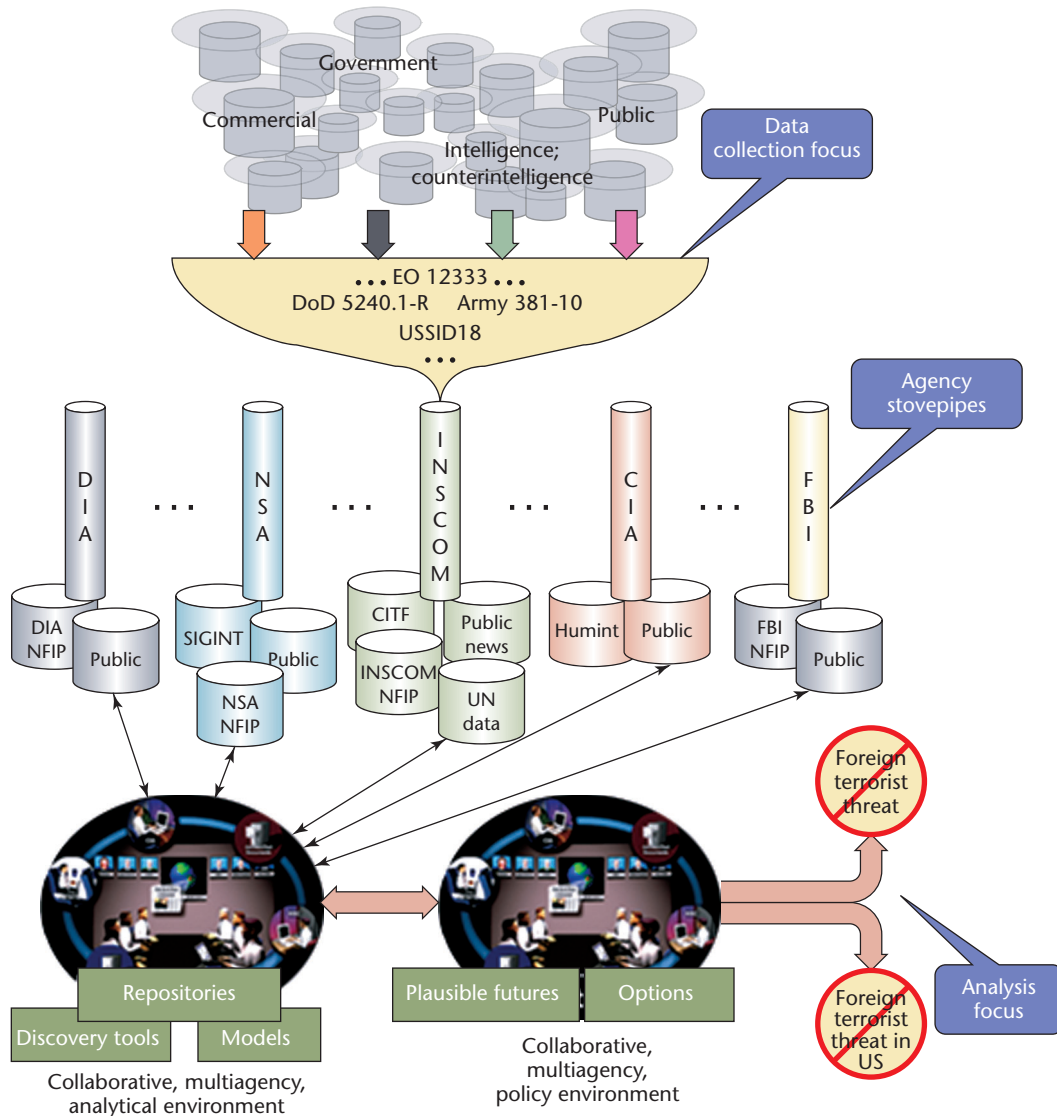


Figure 1. Counterterrorism framework. Information technologies fall into one of two broad categories: collections and analytics.

possible, failure in the experimental environment is an acceptable outcome for a particular technology.

Figure 2 shows an approach to understanding how to measure the operational payoff of information technologies for counterterrorism. As the graphic shows, when doing traditional analysis, an analyst spends much of his or her time on the major processes broadly defined as research, analysis, and production. This *bathtub curve* shows that analysts spend much time doing research and production but too little time doing analysis. An objective of conducting experiments with this curve is to determine whether we can improve analyses via information technology by reversing this trend and inverting the curve.^{4,5}

Specifically, Figure 2 shows the results of an experiment in which the intelligence question posed to ana-

lysts was, “What is the threat posed by Al Qaeda’s weapons of mass destruction capabilities to several cities in the US?” The data were drawn from various classified intelligence sources, foreign news reports, and the Associated Press (AP) and other wire services.^{4,5} The information technologies used in the experiments included a peer-to-peer collaboration tool, a structured argumentation decision aide, a multilingual processing tool for audio phonetic searching/indexing as well as text filtering/categorization, and several graph-based link analysis tools. The results of the experiment show an inverted bathtub curve, allowing for more and better analysis in a shorter period of time, as a result of analysts using information technologies. The obvious significance is that analysts spend a greater percentage of their

Table 1. Sample of the US intelligence community’s legal authority for data collection.

| AUTHORITY | DESCRIPTION |
|--|---|
| Executive Order (EO) 12333 | Authorizes US intelligence activities |
| Foreign Intelligence Surveillance Act (FISA) of 1978 | Prescribes procedures for physical and electronic surveillance and collection of intelligence information between or among foreign powers |
| USA Patriot Act | Dramatically expands the authority of American law enforcement for fighting terrorism in the US and abroad |
| US Department of Defense (DoD) Directive 5240.1-R | Provides the DoD with implementation guidance for EO 12333 |
| Army regulation 381-10 | Provides the Army with implementation guidance for DoD Directive 5240.1-R |
| US Signals Intelligence Directive (USSID) 18 | Governs signal intelligence (SIGINT) for the National Security Agency (NSA) |

time doing what is most important in our view—namely, the critical-thinking tasks instead of the more mundane research and production tasks. The results also included an impressive savings in analyst labor (that is, half as many analysts participated in the IT-enhanced analysis) and an increase in the number of reports produced (that is, analysts created five reports in the time it took to create one manually).

Our explanation for the bathtub curve’s inversion for the intelligence question at hand includes

- The time spent in the research phase shrank dramatically by using the collaboration tool (Groove) across multiple agencies to harvest and share “all” pertinent data.
- The structured argumentation modeling tool (SEAS, for Structured Evidentiary Argumentation System) let analysts explicitly represent their hypotheses for comparison and assessment, and identify evidentiary data gaps for which data must be searched and harvested.
- The multilingual processing tool (FastTalk) let analysts phonetically index and search vast quantities of foreign audio streams and thereby reduce the time required to find pertinent data.
- The link analysis tools (Analyst Notebook) let analysts automatically capture portions of their analysis in an easy-to-understand visual format.^{4,5}

Figure 3 shows the utility of various information technologies in detainee operations support. In this scenario, actual government interrogators questioned actual detainees at the US military facility at Guantanamo Bay, Cuba, and wanted analytical support to make sense of the stacks of real reports from hundreds of interrogation sessions. The analysts used a link analysis tool to find nonobvious relationships between different entities (people, places, and things), a group detection tool to find nonobvious groupings among entities, an entity resolution tool to resolve entities and aliases in the inter-

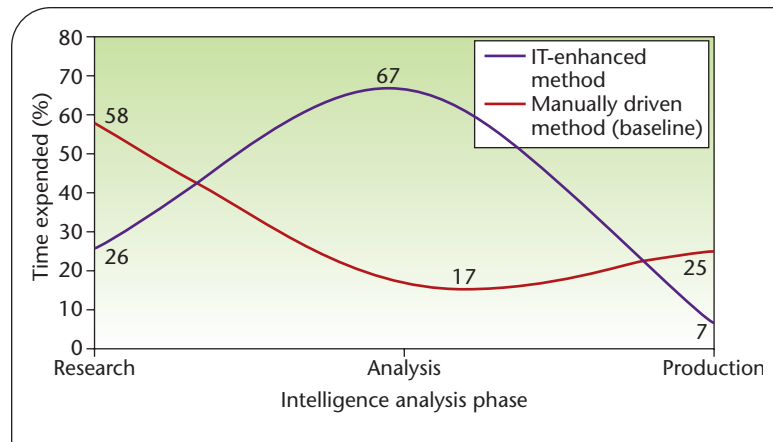


Figure 2. The analyst “bathtub” curve. The red curve represents the baseline distribution of time an analyst manually spends on research, analysis, and production; the blue curve represents the improvement due to information technology enhancements.

rogation reports, a Bayesian classification tool to classify detainees of unknown status as either statistically more likely to resemble known terrorists or nonterrorists, and a link chart visualization tool to pull everything together. These tools showed the interrogators web-like diagrams of connections (or relationships) among different entities that weren’t readily apparent, inconsistencies in detainee stories, salient relationships across detainees, useless data to disregard, and data that could be most informative for follow-up interrogations. The tools’ output also included a rank-ordered list of detainees with the likelihood that each had attributes resembling known terrorists or nonterrorists.⁴⁻⁶ (It should be noted that officials at Guantanamo Bay established the “ground truth” in terms of which detainees were terrorists and which ones weren’t.) Based on conversations with the intelligence analysts who performed this work, anecdotal evidence suggests that the detainees classified as “likely a terrorist” were in fact terrorists, and no cases

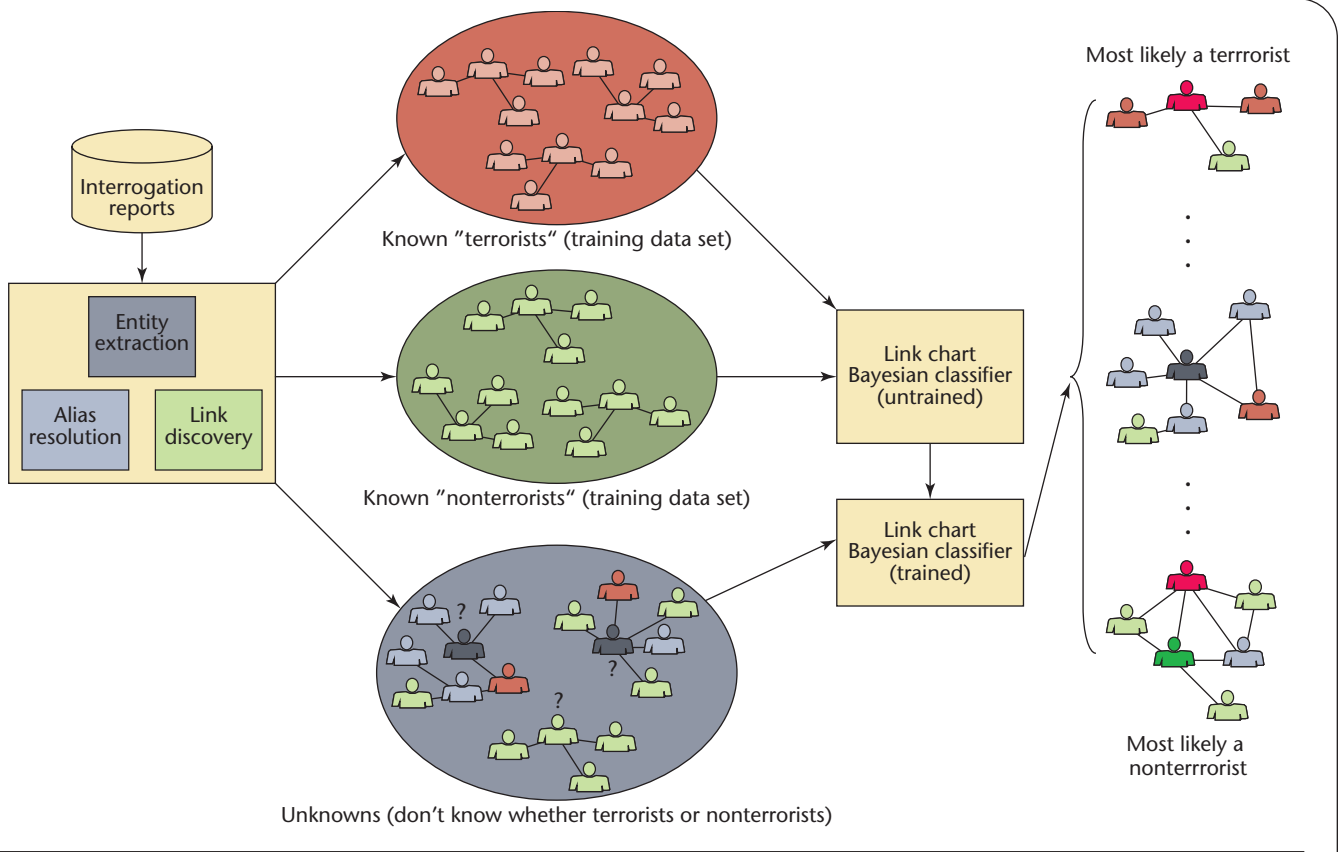


Figure 3. Detainee operations support at Guantanamo Bay. The analysts used information technologies to build web-like diagrams of relationships between entities that weren't immediately apparent.

were found in which detainees who weren't terrorists were classified as "likely a terrorist."

Figure 4 shows an experiment in which a novel multilingual IT front-end system automatically ingests, transforms, extracts, and autopopulates in near real time the back-end analytical models from massive amounts of text data. In this experiment, the problem concerned understanding and forecasting the preconditions and root causes that give rise to instability in nation states. Failed states are important because they offer a safe haven and potential breeding ground for terrorists. The challenge posed to analysts here was to assess and forecast the level of instability in two specific countries in Southeast Asia. The data came from a variety of open sources and included more than 1 million English documents and 2,300 non-English documents. The information technologies used included a back-end rebel activity model (RAM) based on a Bayesian network and hidden Markov models (HMMs) that measured the amount of rebel activity (on the part of separatists, insurgents, terrorists, Islamic extremists, and so forth); a front-end language-independent text-based transformation and categorization tool based on a Hilbert engine (a technology that numerically encodes ASCII text into vec-

tors in Hilbert space); and a linguistic pattern analyzer (LPA) that automatically populates the HMMs in the RAM model.

The experiment's results were impressive—given a corpus of 1,236,300 documents, a human would need 117 man years to read it all (assuming it took 12 minutes to read each document), or 280 humans to read the documents in six months. The automated front-end system based on LPA, the Hilbert engine, and RAM would take a mere 0.05 man years with a one-time cost of 0.76 man years to configure LPA with the numerous multilingual scripts. Assuming it cost US\$100K per man year, the automated front-end would provide a savings of US\$11,695,141 over the human method.

Signatures in silos

One of the major criticisms leveled against an approach such as ours is that what we're describing is mass data-veillance—warehousing massive amounts of data in a megadatabase and using data mining techniques that will lead to multiple false positives and a massive invasion of Americans' privacy. We disagree. Although we appreciate the significant information policy challenges concerning data analysis in actual transaction spaces, we

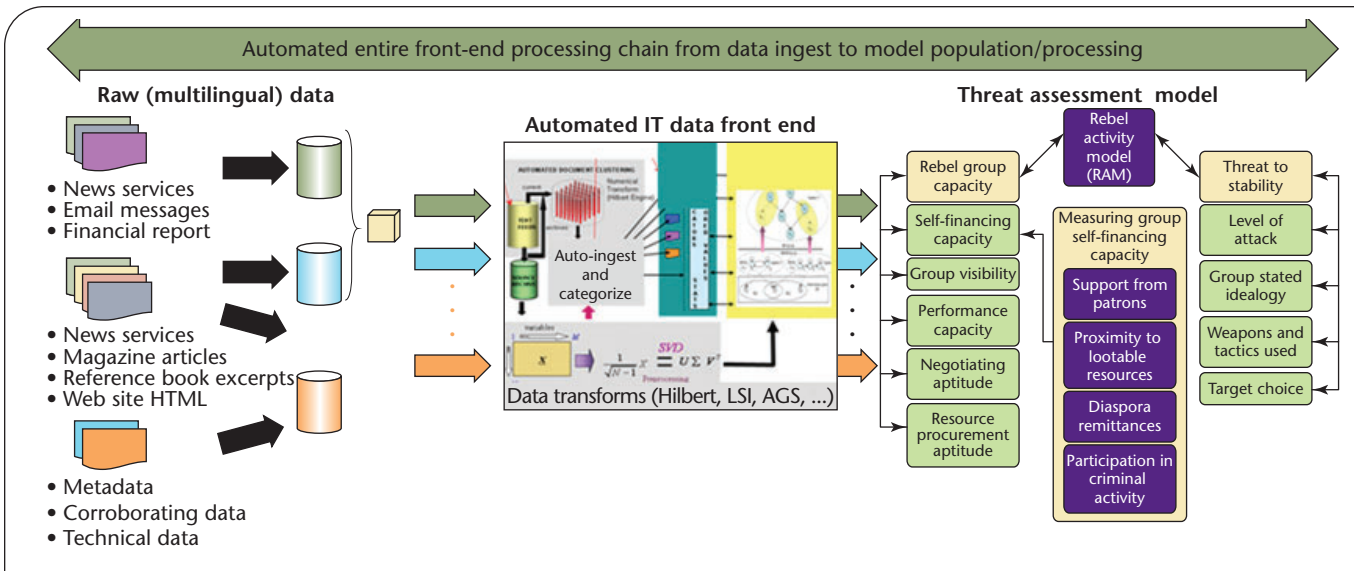


Figure 4. A multilingual IT front-end system. This tool automatically ingests, transforms, extracts, and autopopulates in near real time the back-end analytical models from massive amounts of text data.

Table 2. Data mining vs. terrorism detection.

| COMMERCIAL DATA MINING | TERRORISM DETECTION |
|--|--|
| Discover comprehensive models of databases to develop statistically valid patterns | Detect connected instances of rare patterns |
| No starting points | Known starting points or matches with patterns estimated by analysts |
| Apply models over entire data | Reduce search space; results are starting points for human analysis |
| Independent instances (records) | Linked transactions (networks) |
| No correlation between instances | Significant autocorrelation |
| Minimal consolidation needed | Consolidation is key |
| Dense attributes | Sparse attributes |
| Sampling okay | Sampling destroys connections |
| Homogenous data | Heterogeneous data |
| Uniform privacy policy | Nonuniform privacy policy |

believe technology and enabling policies can help preserve civil liberties and protect the privacy of those people who aren't terrorists while keeping us all safer from attack.

Data mining commonly refers to using techniques rooted in statistics, rule-based logic, or artificial intelligence to comb through large amounts of data to discover previously unknown but statistically significant patterns. However, the general counterterrorism problem is much harder because unlike commercial data mining applications, we must find extremely rare instances of patterns across an extremely wide variety of activities and hidden relationships among individuals. Table 2 gives a series of reasons for why commercial data mining isn't the same as terrorism detection in this context. We call our technique for counterterrorism activity *data analysis*, not data mining.⁷

Instead of warehousing data in one megadatabase, we believe data must be left distributed over the large number of heterogeneous databases residing with their data owners. In an intelligence context, agency silos and stovepipes aren't necessarily bad—they allow analysts from different agencies to create alternative competing hypotheses, and they also protect agency-specific sources and methods. In our judgment, the goal shouldn't be to tear down these silos, but to punch holes in them and enable collaboration across agencies when appropriate and advantageous.

Advanced search and discovery tools should be used to search and query relevant databases—under rigorous access control and privacy protections—with the results of the search/query added to the analytical models. Because finding evidence of a suspicious terrorist plot isn't

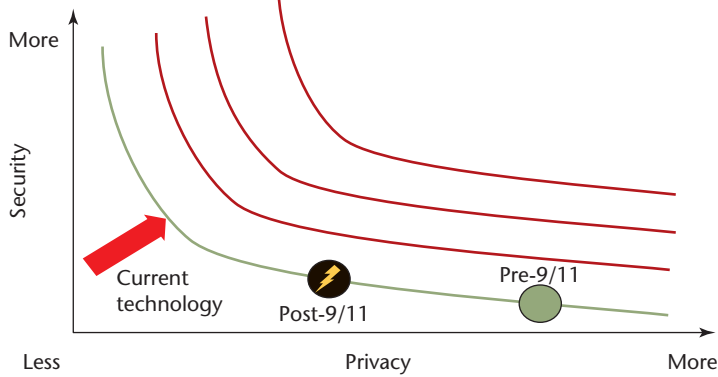


Figure 5. Security vs. privacy curves. Laws and policies dictate where we are on a curve; new privacy technology can create new curves.

easy, we believe two basic types of queries are necessary: subject-based queries (sometimes referred to as *particularized suspicion*) and pattern-based queries (sometimes referred to as *nonparticularized suspicion*).

Subject-based queries let analysts start with known suspects, look for links to other suspects, people, places, things, or suspicious activities, and do so within well-defined and practiced sets of legal and regulatory protocols. Law enforcement personnel have used this technique successfully for years as part of their background investigations and as a forensic tool. In the previous section, we gave examples of how subject-based queries can be beneficial for counterterrorism purposes (such as in the Guantanamo Bay detainee example), but this might not be enough. To get ahead of the terrorism problem, we need to consider pattern-based queries that don't require a subject's prior identification.

Pattern-based queries let analysts take a predictive model and create specific patterns that correspond to anticipated terrorist plots, and use (largely existing) discovery tools and advanced search methods to find instances of these patterns in the information space. This latter approach becomes essential because it can provide clues about terrorist sleeper cells made up of people who have never engaged in activity that would link them to known terrorists. Nonparticularized suspicion raises even higher the question of civil liberties, though—currently, no well-defined or practiced legal or regulatory protocols govern its operation, so a new privacy policy framework for management and oversight is needed (we'll briefly discuss this later).

With respect to false positives, some of our critics have stated that pattern-based queries create more false positives than they help resolve. Dealing with false positives—which are a legitimate concern given that the government might get it wrong and stigmatize or inconvenience nonterrorists—requires pattern-based

queries to be issued iteratively in a privacy-sensitive manner (specifically, via anonymization and selective revelation techniques). Handling them also requires multiple stages of human-driven analysis in which analysts can't act on the results of such queries until a third-party legal authority has established sufficient probable cause. Analysts would refine queries in stages, seeking to gain more confirmation while invoking numerous privacy-protection techniques in the process. This isn't unlike the tried and proven signal-processing analysis techniques found in antisubmarine warfare, in which human-driven analysis addresses false positives at various stages in a similar manner.⁸

Safeguarding civil liberties

Americans expect their government to protect them from enemy attack as well as safeguard (or at least not violate) their civil liberties and privacy. We believe these two ideals aren't mutually exclusive: Figure 5 shows how our goal (and challenge) is to maximize security at an acceptable level of privacy. In other words, we can pick acceptable levels of privacy and through the development and use of technology, create new level of privacy versus security curves, thus increasing security. A full discussion of what privacy means from a legal and regulatory context is beyond this article's scope, but for a working definition, we would argue that personal privacy is only violated if the violated party suffers some tangible loss, such as unwarranted arrest or detention, for example. The right balance between the two must be understood, as well as the corresponding social costs, benefits, and roles played by the public, government, and private sectors.

As discussed earlier, analysts must systematically use information technologies to detect and discover instances of known or emerging terrorist signatures, but they must also be able to exploit the permitted information sources they need to access and do so while protecting the privacy of nonterrorists. Privacy-protection technology is a key part of the solution not only to protect privacy but also to encourage the intelligence, law enforcement, and counterterrorism communities to share data without fear of compromising sources and methods. However, the American public has legitimate concerns about whether protections for privacy are adequate to address the potential negative consequences of increased government use of permitted information sources. These concerns are heightened because there is little understanding or knowledge about how the government might use this data.

The R&D community has explored several promising privacy-protection technologies, especially those that are most relevant to the pattern-based query approach. We briefly describe some of them here, but more detailed information appears elsewhere.⁹⁻¹¹




Figure 6. Privacy appliance concept. A tamper-resistant, cryptographically protected device serves as a trusted privacy-enforcing guard between the user and the database.

Privacy appliance

Our privacy appliance concept involves the use of a separate tamper-resistant, cryptographically protected device placed on top of databases. The appliance would be a trusted, guarded interface between the user and the database analogous to a firewall, smart proxy, or a Web accelerator. It would implement several privacy functions and accounting policies to enforce access rules established between the database owner and the user. It would also explicitly publish the details of its technology, verify the user's access permissions and credentials (packaged with the query in terms of specific legal and policy authorities), and filter out queries not permitted or that illegally violate privacy. Finally, it would create an immutable audit log that captures the user's activity and transmits it to an appropriate trusted third-party oversight authority to ensure that abuses are detected, stopped, and reported. (Granted, our privacy appliance concept assumes the third party is trusted, which is often the hardest problem to solve.) The privacy appliance's operation must be automated to respond to the dynamic, time-sensitive nature and scale of the problem and to ensure the privacy policy's implementation. Figure 6 illustrates the privacy appliance concept in terms of some of its key privacy functions as well as how it would work operationally.

Data transformation

Used within the privacy appliance, data transformation employs well-known mathematical encoding tech-

niques to transform data from a plaintext representation to cipher, thus making computer processing more efficient and the data unintelligible to humans. Once transformed, analysts could apply a plethora of data analysis functions to understand the data's significance, keeping the identities of subjects hidden from analysts but still allowing the detection of terrorist activity patterns, such as data searching, alias and entity resolution, and pattern-query matching. Because the data is represented in unintelligible cipher, no personally identifiable data is disclosed to the analyst, thus privacy protection is maintained.

Anonymization

Similar to data transformation, anonymization is a technique used within the privacy appliance: it generalizes or obfuscates data, providing the system with a guarantee that any personally identifiable information in the released data can't be determined, yet the data still remains useful from an analytical viewpoint. As an example, instead of releasing to an analyst a database record consisting of [name(first, last); telephone #(area code, exchange, line number); address(street, town, state, zip code)], an anonymized version of this database record could be [name(first); telephone #(area code); address(state)]. For this approach to work, analysts will have to make connections between queries and thus will require some sort of anonymized unique identifier as well. Much more thorough treatment of various anonymization techniques and applications for privacy appears elsewhere.^{10,11}

Selective revelation

Another technique to employ in the privacy appliance is selective revelation, which gives incremental access to and analysis of increasingly personally identifiable data. In this approach, what an analyst gets back in response to a pattern-based query varies in depth and specificity depending on the analyst, the investigation's status, and other criteria. The analyst's knowledge of an individual's identity would occur only after a sufficient level of suspicion and appropriate legal threshold were met. The approach proceeds incrementally by requiring data owners to release subsets of data— anonymized, filtered, or statistically characterized—to an analyst's pattern-based query. Initially, no personally identifiable data is provided in response to the query. If the results turn out to be meaningful after iteration and refined patterns or queries—say, only an acceptably few individuals match the query, or the level of suspicion or probable cause has been heightened—then additional permissions and authorization through an appropriate (yet currently nonexistent) legal framework would need to be secured to release personally identifiable data of individuals under suspicion.

Immutable audit

Another technique to be used in the privacy appliance is an immutable audit, which automatically and permanently records all accesses to data, with no possibility of undetected alteration or tampering. To prevent potential abuses by malicious agents, audit logs would be designed so that any misdeeds or corruption are detected with the highest probability. Audit logs would be cryptographically protected and transmitted to a trusted third-party oversight authority. Privacy tools to query and analyze audit logs are also critical. The contents of the audit log could contain fields such as the analyst's identity and credentials, the authorizations and permissions allowed, the date and time of the data access, the data requested, and the data returned.

Self-reporting data

An important technology that isn't directly related to the privacy appliance but is important from a civil liberties perspective is self-reporting data. This is a method for truth maintenance as well as for reporting on the data's distribution. Data used in analysis should be active (that is, it should report back to a central authority about where it is and for what it's being used); this point is essential to correct any information that's later proved to be false.

Privacy laws

Government access to personally identifiable data raises legitimate concerns about the protection of civil liberties, privacy, and due process. Given the limited applicability of current privacy laws to the modern digital era, practical policies for new information technology use, redress, and

oversight are vital. New privacy policy can help ensure that controls and protections accompany the use of the information technologies we discussed earlier. Here, we list several basic principles as examples of the types of policy to consider:¹²

- *Neutrality.* New information technology should build in existing legal and policy limitations about access to personally identifiable or third-party data.
- *Minimize intrusiveness.* Personally identifiable data is voluntary but might be required as a condition of service (such as driver's licenses), thus it should be anonymized or rendered pseudonymous and disaggregated (when possible).
- *Intermediate not ultimate consequence.* Personal identification by a new information technology shouldn't directly lead to ultimate consequence (such as arrest); instead, analysts should view it as cause for additional investigation.
- *Audits and oversight.* New information technology should have strong built-in technological safeguards such as audit and oversight mechanisms to detect and deter abuse.
- *Accountability.* New information technology should be used in a manner that ensures accountability of the executive branch to the legislative branch for its use.
- *Necessity of redress mechanisms.* Robust legal mechanisms for the correction of false positives should be in place.
- *People and policy.* Internal policy controls, training, administrative oversight, enhanced congressional oversight, and civil and criminal penalties for abuse should all be in place.

We hope these considerations will be taken into account along with legal protocols for pattern-based searches; technology does and can play a key role in the careful balance of security with privacy.

Information and privacy-protection technologies are powerful tools for counterterrorism, but it's a mistake to view technology as the complete solution to the problem. Rather, the solution is a product of the whole system—the people, culture, policy, process, and technology. Technological tools can help analysts do their jobs better, automate some functions that analysts would otherwise have to perform manually, and even do some early sorting of masses of data. But in the complex world of counterterrorism, the technologies alone aren't likely to be the only source for a conclusion or decision.

Ultimately, the goal should be to understand the level of improvement possible in our counterterrorism operations using advanced tools such as those described here but also to consider their impact—if any—on privacy. If research shows that a significant improvement to detect

and preempt terrorism is possible while still protecting the privacy of nonterrorists, then it's up to the government and the public to decide whether to change existing laws and policies. However, research is critical to prove the value (and limits) of this work, so it's unrealistic to draw conclusions about its outcomes prior to R&D completion. As has been reported,⁶ research and development continues on information technologies to improve national security; encouragingly, the Office of the Director of National Intelligence (ODNI) is embarking on an R&D program to address many of the concerns raised about potential privacy infringements. □

Acknowledgments


The views expressed herein are the authors' alone and don't reflect the views of any private-sector or governmental entity.

References

1. R. Popp and J. Yen, eds., *Emergent Information Technologies and Enabling Policies for Counter-Terrorism*, Wiley & Sons/IEEE Press, 2006.
2. *Report to Congress Regarding the Terrorism Information Awareness Program*, DARPA, May 2003; [response to Consolidated Appropriations Resolution, Pub. L. no.108-7, div. M, sec. 111(b), 2003].
3. J. Poindexter, "Overview of the Information Awareness Office," *DARPA Tech 2002*, DARPA, 2002; www.fas.org/irp/agency/dod/poindexter.html.
4. R. Popp et al., "Countering Terrorism through Information Technology," *Comm. ACM*, vol. 47, no. 3, 2004, pp. 36-43.
5. E. Jonietz, "Total Information Overload," *MIT Tech. Rev.*, vol. 106, no. 6, 2003, p. 68.
6. S. Harris, "Signals and Noise," *Nat'l J.*, vol. 38, no. 24, 2006, pp. 50-58.
7. M. DeRosa, *Data Mining and Data Analysis for Counterterrorism*, CSIS Press, 2004.
8. T. Senator, "Multi-Stage Classification," *Proc. 5th IEEE Int'l Conf. Data Mining (ICDM 05)*, IEEE CS Press, 2005, pp. 386-393.
9. "Security with Privacy," DARPA's Information Systems Advanced Technology (ISAT) study, Dec. 2002; www.cs.berkeley.edu/~tygar/papers/ISAT-final-briefing.pdf.
10. L. Sweeney, "Weaving Technology and Policy Together to Maintain Confidentiality," *J. Law, Medicine and Ethics*, vol. 25, nos. 2-3, 1997, pp. 98-110.
11. L. Sweeney, "k-Anonymity: A Model for Protecting Privacy," *Int'l J. Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 5, 2002, pp. 557-570.
12. P. Rosenzweig, "Privacy and Consequences: Legal and Policy Structures for Implementing New Counter-Terrorism Technologies and Protecting Civil Liberty," *Emergent Information Technologies and Enabling Policies for Counter-Terrorism*, R. Popp and J. Yen, eds., Wiley & Sons/IEEE Press, 2006, pp. 421-438.

Robert Popp, now CEO of National Security Innovations, recently served as a senior government executive within the US Department of Defense as deputy director of the Information Awareness Office (IAO) at DARPA and assistant deputy under-secretary of defense for advanced systems and concepts in the Office of the Secretary of Defense. He serves on the Defense Science Board, is a senior associate for the Center for Strategic and International Studies, and is an associate editor of IEEE Transactions on Systems, Man, and Cybernetics. Popp holds two patents, authored numerous journal and conference papers, and edited *Emergent Information Technologies and Enabling Policies for Counter-Terrorism* (Wiley & Sons/IEEE Press, 2006). Popp has a PhD in electrical engineering from the University of Connecticut, and a BA/MA in computer science from Boston University. Contact him at rpopp@nationalsecurityinnovations.com.

John Poindexter, now a private consultant, most recently served as director of the Information Awareness Office (IAO) at DARPA. He also serves on the board of directors for Saffron Technology, a computer software company that produces associative memory applications. Prior to working at DARPA, Poindexter served as National Security Advisor and Deputy National Security Advisor under President Ronald Reagan from 1983 to 1986. He has a PhD and an MS in physics, both from the California Institute of Technology. Contact him at john@jmpconsultant.com.



Stay on Track

IEEE Internet Computing reports emerging tools, technologies, and applications implemented through the Internet to support a worldwide computing environment.

In 2007, we'll look at

- Autonomic Computing
- Roaming
- Distance Learning
- Dynamic Information Dissemination

... and more!

IEEE
Internet Computing

www.computer.org/internet