abling Policies for Counter-Terrorism, R. Popp and J. Yen, eds., Wiley–IEEE Press, 2005.

6. "The Limits of Hindsight" (editorial), *Wall St. J.*, 28 July 2003, p. A10.

Privacy-Preserving Surveillance Using Databases from Daily Life

Latanya Sweeney, *Carnegie Mellon University*

As the price of disk storage continues to plummet, the cost of capturing and sharing data approaches zero, making it economical to collect more and more information on individuals' daily lives, often without any particular purpose.[1]

One proposed use for all this information is homeland security (law enforcement and intelligence). When fragments of captured information are combined, they provide person-specific, population-based data for profiling individuals. Database systems might use the data to find behavioral patterns of individuals engaged in illegal activity or planning terrorist acts.

Privacy concerns

American programs that sought to use databases for surveillance include CAPS II (computer-assisted passenger screening) and TIA (Total Information Awareness).[2] Both programs faced serious turmoil over privacy concerns. Such concerns include the following:

- The bulk of people whose information is in the database have done nothing to warrant suspicion.
- Surveillance on databases tends to exasperate privacy expectations and personal protections. While American courts have historically ruled that a person in a public space should have no expectation of privacy,[3] information stored in databases can be so invasive as to remove private enclaves within public spaces. For example, on a crowded bus, you can orient a document to limit what others can see. But limiting what a hidden camera with a zoom lens can see is difficult because its existence and viewing angle are unknown.
- Information in a database can be gathered from private spaces. For example, a private inquiry made on a home phone can become part of a database, making it indistinguishable from inquiries made at a public shop.
- Organizations using databases for surveillance purposes don't tend to implement Fair Information Practices (www3.ftc.gov/reports/privacy3/fairinfo.htm) because they don't want criminals and terrorists to alter their information or behavior. Therefore, no individual whose information is contained in the data has control over his information. The organizations don't seek consent from subjects or give notice to those included. (Arguably, doing so would be impractical.) So typically, subjects don't know their information is being held, and they have no right or means of correcting errors in it.

As the price of disk storage continues to plummet, it becomes economical to collect more and more information on individuals' daily lives, often without any particular purpose.

- No judicial review or impartial oversight exists to weigh societal benefits against individual risks. No independent third party limits fishing expeditions unwarranted inquiries, snooping on friends, or other kinds of "fishing expeditions."

My goal is to guarantee (or at least maximize) privacy protection while making data useful for surveillance. This work introduces a framework that addresses database privacy conditions in surveillance databases such that

- no person whose information is contained in the database can be reidentified without permission,
- investigators can access necessary information contained in the database freely and easily, and
- results from qualified inquiries are equivalent to results found in the absence of privacy protection.

Methods

One way to satisfy these privacy conditions is to model the probable cause predicate in American jurisprudence. A law officer wanting to intrude on a person's private life or affairs needs a search warrant, which a judge can issue. The officer appears before the judge and reports either facts for which he or she has first-hand knowledge or facts that he or she learned through an informant. Typically, the judge uses a two-prong test to make a decision: what is the basis of the knowledge, and is the source believable (see figure 4a)? We can model this process in technology by replacing the officer with anomaly or data-mining algorithms and the informant with data from various sources. We can replace the human judge with a combination of contracts and certifications from the original data collectors and a technology-enforceable policy statement with preset levels that match the identifiability of provided information with the minimal information the algorithm needs (see figure 4b). The technology capable of enforcing the policy is called *selective revelation*.

The first step in constructing a selective-revelation system requires identifying the algorithms to be used and the kinds of data involved. The person setting up the system performs analyses to provably anonymize the data and to verify that the algorithms remain useful with the anonymized data.

Once the initial step is complete, the person maps related regulations, policies, best practices, laws, and data certifications onto the scale of identifiability—from anonymous to identifiable—to specify the authority by which data can be accessed at each status (see figure 5). Finally, boundaries of algorithmic utility are established to identify the algorithmic circumstances under which more identifiable data is necessary.

Figure 5 shows how identifiability maps to investigation status. During normal operation, the surveillance agency uses anonymized data. If the agency encounters unusual activity, as evidenced by algorithmic results, then the system lowers the identifiability of related cases to "de-identified." De-identified data has no explicit identifiers but isn't provably anonymous. As the investigation status shifts downward, the provided information be-

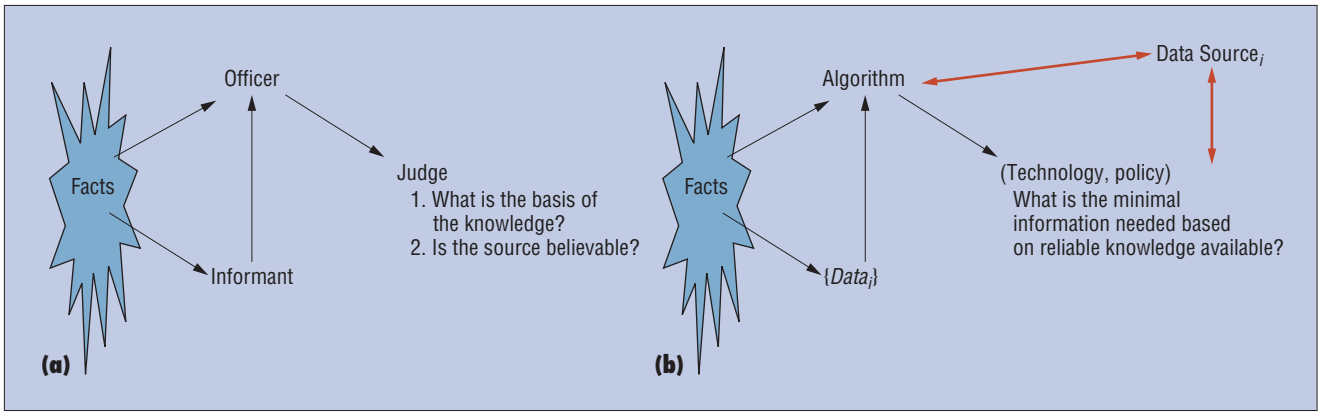

Figure 4. Probable cause predicate as conducted by (a) a human judge and (b) technology.

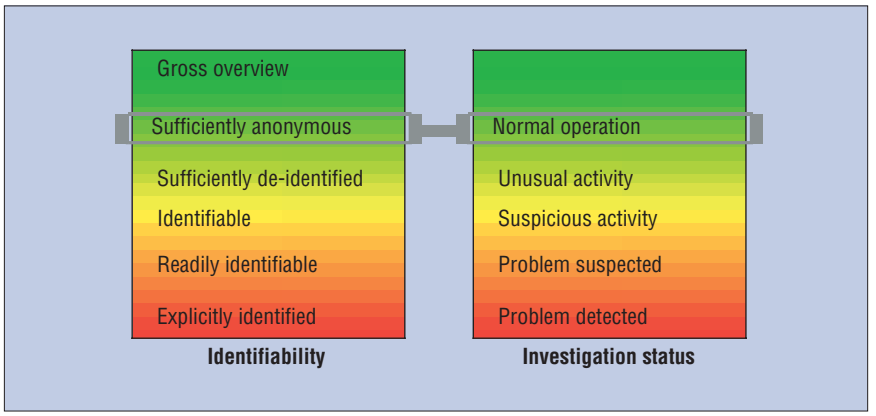

Figure 5. Selective revelation scales matching the identifiability of the data (left) to the operational mode (right).

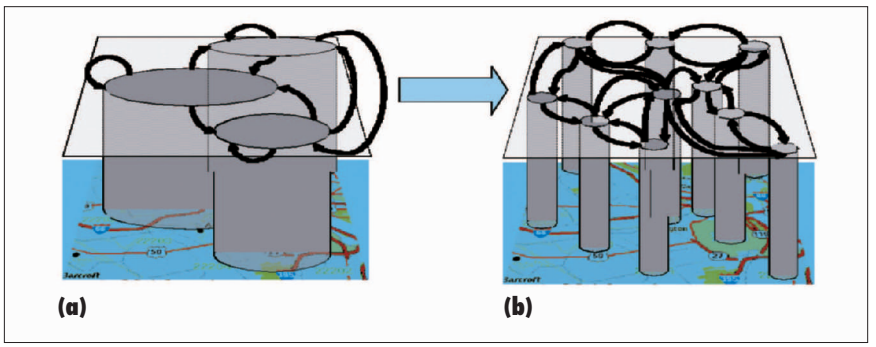

Figure 6. Dynamically augmenting data access as surveillance warrants. (a) Crude relationships are derived from sufficiently anonymous data. (b) More details are revealed using identifiable data.

comes increasingly more identifiable, until the agency meets the criteria for providing explicitly identified data. Figure 6 demonstrates the effect of lowering identifiability.

Example

Earlier, I constructed a selective-revelation system for bioterrorism surveillance was constructed in which hospitals, physicians, and labs provided medical data to a public health agency to determine whether an unusual number of respiratory cases were presented.[4] I anonymized the data under the scientific standard of the medical regulation known as the Health Insurance Portability and Accountability Act (HIPAA) (Privacert Compliance was used; www.privacert.com). The early aberration re-porting system algorithm from the Centers for Disease Control was used with the anonymized data. If it found evidence of unusual activity, the system automatically lowered anonymity. If further evidence emerged that an outbreak was underway, fully identified data under the Public Health Law was provided by the system. This selective revelation system provided impartial, automated oversight to database inquiries. It demonstrates how the American public can enjoy both safety and privacy.

References

1. L. Sweeney, "Information Explosion," *Confidentiality, Disclosure, and Data Access*, P. Doyle et al., eds., Elsevier, 2001, pp. 43–74.
2. "Your Papers, Please," World Net Daily, 16 Jan. 2003; www.worldnetdaily.com/news/article.asp?ARTICLE_ID=30523.
3. Rest. 2d, §652B.
4. L. Sweeney, "Privacy-Preserving Bio-terrorism Surveillance," *Proc. AAAI Spring Symp. AI Technologies for Homeland Security*, 2005.

The Changing Face of Privacy Policy and the New Policy-Technology Interface

Paul Rosenzweig, *Heritage Foundation*

America's rules- and regulation-driven model of privacy protection is undergoing a major transition. Driven partly by needs spawned in the wake of 9/11, the traditional