Introducing Privacy-Protective Surveillance:

Achieving Privacy ***and*** Effective Counter-Terrorism

Ann Cavoukian[1], Khaled El Emam[2], et al[3]

Executive Summary

A new concept for surveillance – Privacy-Protective Surveillance (PPS), is being advanced in this paper – a positive-sum (opposite of zero-sum) alternative to current counter-terrorism surveillance systems, with a methodology developed for its implementation.

As long as the threat of terrorism exists and the global conditions that instantiate those threats continue, effective measures will be needed to counteract terrorism. At the same time, in order for a free and open society to function properly, civil liberties must be protected. Above all, privacy, as the ability of law-abiding individuals to control the collection, use, and disclosure of personal information about themselves – referred to at times as "informational self-determination," must be protected.

Most approaches to protecting privacy, while ensuring measures to counteract terrorism, seek to strike a "balance" between these two interests. This often leads to engaging in a zero-sum paradigm of giving up what is perceived to be the "less important value," namely privacy, in favor of the "more significant value," namely public safety (imagine a see-saw – the more that one side goes up, the other side must go down). This zero-sum trade-off is invariably destructive in free and open societies. It is not only inappropriate, it is unnecessary. Privacy and counter-terrorism measures can indeed co-exist, with both values being respected, instead of being positioned as opposing forces requiring unnecessary trade-offs, or false dichotomies.

Building on *Privacy by Design* (the international framework recognized as "an essential component of fundamental privacy protection" by Data Protection and Privacy Commissioners in 2010) is Privacy-Protective Surveillance (PPS) – a positive-sum, "win-win" alternative to most counter-terrorism surveillance systems. An extension of Artificial Intelligence, by embedding privacy directly into its design and architecture, through the use of such technologies as intelligent virtual agents, homomorphic encryption,

1 Information and Privacy Commissioner, Ontario, Canada (commissioner.ipc@ipc.on.ca).

2 Associate Professor, Faculty of Medicine and the School of Electrical Engineering and Computer Science, University of Ottawa (kelemam@ehealthinformation.ca).

3 Our deepest thanks go to Dr. George Tomko, whose work inspired this paper and formed the basis of the methodology for PPS; Michelle Chibba; Alex Stoianov; and David Weinkauf.

and machine-learning data analysis networks, PPS allows for privacy and counter-terrorism to co-exist in tandem, without diminishing the intelligence-gathering abilities of the systems involved. Specifically, PPS offers the development of a new system design of privacy-protective "feature detection." This has the ability to scan the Web and related databases to detect digital evidence relating to terrorist activity, while ensuring that any personally identifying information (PII) on unrelated, law-abiding individuals is not collected. In those cases associated with targeted activity, PII will automatically be encrypted upon collection, analyzed securely and effectively within the "space of cipher text," and only divulged to the appropriate authorities with judicial authorization (a warrant).

One of the most attractive elements of PPS is the fact that its intelligent agents will only collect data that is considered to be "significant." Significant data is defined by transactions or events that are believed to be associated with terrorist-related activities. For example, purchasing fertilizer capable of bomb-making or accessing a bomb-making website.

An important consequence of PPS's collection of significant data is that its virtual intelligent agents would effectively be "blind" to "seeing" any other information they may run across during their searches. Since each intelligent agent (of which there would be thousands) would only be configured to search for a single "feature of interest," it would be "blind" to everything else – the agent would be oblivious to any other "non-features" such as additional personal information. This would avoid exposing the personal information of millions of people who were not considered to be persons of interest – leaving their privacy intact, and dramatically expected to reduce the harmful incidence of false positives.

In addition, the use of homomorphic encryption will allow PPS to make computations or engage in data analytics on encrypted values – data that cannot be "read" because it does not appear in plain text. This provides additional assurance to individuals that no "prying eyes" would be able to record or monitor their actions within the system.

Finally, the intelligence gathered by PPS will be context-specific. In order to become information of value, data must be placed in the appropriate context. The ability of intelligent agents to acquire substantive knowledge of the topic surrounding the particular feature they are designed to collect, in order to serve as the appropriate frame of reference, will both improve the assessment of terrorist-related activities as well as mitigate any additional privacy burdens. For example, if the feature detected was purchasing fertilizer capable of bomb-making, then the agent would seek to determine the occupation of the individual – student, farmer, banker – to aid in producing a conditional probability table for the likelihood of that activity being related to terrorism, on the basis of a probabilistic graphical model (PGM).

In order for PPS to produce knowledge about terrorist threats, a PGM will be structured beforehand by intelligence experts comprising: (1) all of the features of interest in determining potential terrorist activity (treated as nodes in the model); and (2) the connections between those features. The graphical model will highlight features that need to be detected by artificial agents. Once developed, the agents' task of determining what actual features were triggered by an individual will serve to prune the PGM into semantics, so to speak, so that conditional probabilities for each feature/node can be assigned and then treated as evidence to infer the probability of terrorist activity given the detected features. It is this probability inference that will, in part, be used by the court to decide whether or not to issue a warrant to release the encryption key to decipher the identity of the individual in question.

By illustrating the organizing methodology behind PPS, we hope to demonstrate that, contrary to appearances, it *is* possible to have both privacy and effective counter-terrorism. Indeed, we possess the technology and can develop the system design to achieve this doubly-enabling end result. By doing so, we will be able to implement strong counter-terrorism measures, while ensuring the future of freedom and liberty – a win/win proposition!