Privacy-Protective Facial Recognition:

Biometric Encryption Proof of Concept

November 2010

Information and Privacy Commissioner,
Ontario, Canada

Acknowledgements

The authors gratefully acknowledge the help of Alex Stoianov, Vance Lockton, Ken Anderson and Michelle Chibba of the IPC, Les DeSouza, Klaus Peltsch and Geoff Truscott of OLG, and Soren Frederiksen of iView Systems in the preparation of this paper.

Information and Privacy Commissioner, Ontario, Canada

2 Bloor Street East
Suite 1400
Toronto, Ontario
M4W 1A8
Canada

416-326-3333
1-800-387-0073
Fax: 416-325-9195
TTY (Teletypewriter): 416-325-7539
Website: www.ipc.on.ca
Privacy by Design: www.privacybydesign.ca

Privacy-Protective Facial Recognition:

Biometric Encryption Proof of Concept

A Research Report on the use of Biometric Encryption to Limit "Self-Excluded" Problem Gambler Access to Gaming Venues

Ann Cavoukian, Ph.D. and Tom Marinelli

Table of Contents

1. Introduction

The rapid, accurate authentication of individuals has become a challenge across many sectors and jurisdictions, as organizations express a need to know who they are dealing with. Current security models allow for three primary forms of authentication: something you know (e.g. a password or other shared secret), something you have (e.g. an identification card), or something you are (e.g. biometrics). Increasingly, the third type of authentication — biometrics — is being viewed as the ultimate means of verification or identification, and many agencies begin to deploy biometric systems (such as fingerprinting or facial recognition) across a broad range of applications.

In the summer of 2007, the Ontario Lottery and Gaming Corporation (OLG)[1] approached the Information and Privacy Commissioner of Ontario, Canada (IPC) to discuss the use of facial biometrics to enhance their ability to identify individuals entering gaming sites who had enrolled in OLG's voluntary 'self-exclusion' program. Although the program is entirely voluntary (opt-in), seeking to recognize only those individuals who have provided positive consent, the increased use of facial recognition technology raises a number of privacy and security concerns. Given their mutual interest in respecting the privacy of all casino patrons, the IPC and OLG agreed that the application of an emerging Privacy-Enhancing Technology — Biometric Encryption (BE) — to a facial recognition system at an OLG casino would be an ideal "win-win" project.

The IPC has long had an interest in Biometric Encryption [1, 2]. It was hypothesized that by incorporating BE as part of a multi-layered approach to privacy, OLG's facial recognition system could ensure that the use and storage of problem gambler records would receive a high degree of privacy assurance. This use of *Privacy by Design* — in which privacy protections are designed directly into technologies, from the outset — would make it possible to achieve a "positive-sum" outcome, in which both the functionality of the biometric system **and** the privacy of individuals are respected.

In this paper, we describe the innovative proof of concept research and development work of a collaborative team consisting of OLG, IPC, members of the University of Toronto (U of T)'s Electrical and Computer Engineering Department, and video surveillance/tracking and biometrics firm iView Systems. This project looked to integrate a "Made in Ontario" BE algorithm developed by the University of Toronto researchers [3] Kostas Plataniotis, Ph.D and Karl Martin, Ph.D. into a commercially-available facial recognition system. The end goal of this collaboration was to develop a technology that could function in a real-world environment, and would offer dramatically improved privacy protection over simple facial recognition, without compromising functionality, security or performance — the hallmarks of a positive-sum, *Privacy by Design* application.

1 The Ontario Lottery and Gaming Corporation is designated as an institution for the purposes of the *Freedom of Information and Protection of Privacy Act*, R.S.O. 1990, c. F.31

2. Facial Recognition and Voluntary Self-Exclusion Programs

2.1 Voluntary Self-Exclusion

In a commitment to fostering an environment of responsible gambling within their gaming sites, many partnerships and programs have been developed by Canadian (and global) gaming authorities to provide both patrons and employees with support and information about addiction issues. In many jurisdictions, these initiatives include a program called "voluntary self-exclusion," which allows individuals the opportunity to opt for a self-imposed ban from one or more gaming sites. All Canadian casinos offer some form of self-exclusion program, though these programs vary in scope (individual casino vs. all gaming sites overseen by a particular gaming authority), length (six months to indefinite), and penalty for breaches (removal from premises, trespassing notice, fine, escalation to a non-voluntary ban, etc.).[2]

In Ontario, the self-exclusion program offered by the Ontario Lottery and Gaming Corporation allows individuals to voluntarily have their names removed from OLG marketing and promotional databases. Enrollees in this program, if found by OLG staff at a gaming site, will also be escorted from the premises by Security staff and issued a trespass notice. An OLG self-exclusion extends for an indefinite period, with a minimum length of six months, after which an individual can submit a written application for reinstatement. Though it is first and foremost the responsibility of self-excluded individuals to remain away, OLG looks to provide assistance, as necessary, by detecting enrollees who attempt to enter a gaming site. Improving detection helps OLG create a key disincentive for self-excluded individuals returning to a site. Of course, this presents a challenge – how best to identify self-excluded individuals amongst a large number of regular patrons?

2.2 Detecting Self-Excluded Individuals

Until recently, OLG's process of detecting self-excluded individuals was largely manual. Enrollees were voluntarily photographed and personal information about them was collected, at their request, to be used in subsequent identification. These photos and associated information were then distributed to OLG gaming sites where they were printed and stored in secure binders accessible by security personnel who, among other responsibilities, would undertake the arduous task of trying to match faces in the casino with photos in the binders. Such a process of manual facial recognition suffers many obvious challenges, due to the limits of staff (and human) capability. Of particular note, humans are not generally good at recognizing the faces of people they don't know [4], and

2 For information on casino self-exclusion programs in Canada, see:
British Columbia: http://www.bclc.com/cm/gamesense/voluntary-self-exclusion.htm
Alberta: http://www.aglc.gov.ab.ca/responsiblegambling/selfexclusionprogram.asp
Saskatchewan (Indian Gaming Authority): http://www.siga.sk.ca/Self%20Exlusion%20Broch.pdf
Manitoba: http://www.mlc.mb.ca/MLC/content.php?pageid=420&langdir=E
Ontario: http://www.olg.ca/about/responsible_gaming/practices.jsp
Quebec: http://lotoquebec.com/corporatif/nav/en/responsible-gaming/self-exclusion-program
Nova Scotia: http://www.nsgc.ca/rgVoluntary.php
PEI: http://www.alc.ca/PlayResponsibly.aspx?id=2041
Others: Look for the 'Responsible Gaming' section of individual casinos' websites, or visit the Responsible Gaming Information Centre within the casino property.

may quickly be overloaded by the task of reviewing the many faces that appear in a busy casino environment (particularly as staff are not searching for a single individual, but instead must watch for *any* self-enrolled person).

As there are thousands of self-identified problem gamblers enrolled in the program, OLG wanted to examine whether technological tools could aid them in more efficiently and effectively meeting their objectives for the self-exclusion program. Such a tool would be required to:

- Reliably detect most self-excluded problem gamblers;
- Not interfere with the smooth flow of other patrons into the casino;
- Be cost-effective; and
- Respect all casino patrons' privacy.

An automated facial recognition system was thought to be an attractive tool to enhance and support the manual inspection process. Computers and algorithms are well-suited to objectively processing enormous amounts of data, in a reliable manner. Furthermore, such a solution seemed feasible within a casino environment, in which patrons are already aware of and accustomed to the presence of video surveillance. From a business perspective, the general patron experience would also be unaffected by such an approach since facial images may be captured at a distance, with no requirement for any direct physical interaction.

2.3 Facial Recognition: One-to-One vs. One-to-Many

Facial recognition systems analyze human faces — a biometric — for the purpose of identifying or verifying them. In general, there are two key processes in a facial recognition system — enrollment and authentication. The enrollment process involves taking a photo of an individual, and collecting any necessary identity-related information. The photo is then automatically analyzed and biometric features are algorithmically extracted, creating what is called a "biometric template." The template is a unique numeric representation of the individual's facial features, and may be stored in a database or on a device such as a smart card or similar token.

The steps involved in the authentication process depend on how the system will be used. Biometric systems such as facial recognition can be deployed in 1:1 ("one-to-one") or 1:many ("one-to-many") modes, depending on the application. 1:1 comparisons are generally used to provide access control (e.g. controlling access to a locked room or a computer system), while 1:many systems are generally applied as a watch list system (e.g. a system designed to find one or more particular individuals in a crowd) or a system preventing multiple enrollments. Access control systems are typically concerned with letting the correct person in; watch list systems are typically concerned with keeping specific people out.

Facial recognition, like most biometric systems, is easier to deploy for a 1:1 application. The authentication process for a 1:1 matching system involves two stages. First, the person requiring access makes an identity claim (for example, by presenting an employee ID), which is used as an index to retrieve a single template from a database of biometrics collected by the organization during the enrollment phase. The system then captures the individual's live biometric (e.g. an image of his

or her face), and compares it against the retrieved template to verify whether the person is who he or she claims to be.

In 1:many watch list mode the system must compare each 'live' biometric captured against a full list of stored templates (a "watch list") — in effect, rapidly performing a matching task against all individuals in the database. The identification process for a 1:many facial recognition system is as follows: once a watch list has been created, and the system is installed, an image of each individual within range of a camera is temporarily captured. Biometric features are extracted from these new images, which are then compared against each of the templates collected during the enrollment phase. Each comparison yields a "matching score," which represents the degree of similarity between the image of the patron and a stored biometric template. The process concludes by determining whether any of the scores are high enough to be included on the list of top matches (the number of top matches is called the "rank"), which may then be followed by manual inspection for purpose of confirmation.

In biometric systems, scoring rules (e.g. the minimum score required to declare a match) are generally administrator configurable. This allows for the management of the rates of false acceptance (i.e. wrongly matching captured images with those of others) and false rejection (i.e. failing to positively identify images of individuals who **are** on the watch list), to which biometric systems are subject. Typically, there is an inverse relationship between the false acceptance rate (FAR) and the false rejection rate (FRR), in which the reduction of one causes an increase in the other. Most biometric systems (including most generic face recognition systems reported in the literature) are required to maintain a very low (e.g., 1 in 10,000) false acceptance rate (FAR) and an acceptable false rejection rate (FRR). In an access control system, for example, it will generally be less problematic to correct a false rejection (via a secondary access mechanism, or a re-scan of the biometric) than a false acceptance (which permits access to an unauthorized individual).

3. Biometric Encryption — The *Privacy by Design* Approach to Biometric Systems

OLG was aware that a decision regarding the adoption of facial recognition technology to aid its self-exclusion program could not be made based on technical considerations alone – sound technology does not necessarily mean good public policy. For instance, privacy advocates have long held that surveillance and biometric systems represent significant privacy concerns. Potential issues that have been identified include [1, 5]:

1. **Function creep** — When personal data is collected, organizations often face suggestions as to why they ought to do something more with it, or temptations to expand the scope of a system – in this case, to use the biometric data for purposes other than those initially intended and described upon collection of the information.

2. **Data linkage** — The uniqueness of biometric templates across individuals allows for the possibility that biometric databases, even if they store only templates (i.e. no images) and are anonymous, can be algorithmically linked for data mining, profiling, investigation, and other purposes.

3. **Data misuse** — Unlike tokens and passwords, biometrics are not the sort of things that can be replaced or reset. Care must be taken to ensure that they are not vulnerable to threat or abuse.

4. **Security vulnerabilities** — Biometric systems are potentially vulnerable to a range of attacks, including: spoofing, interception, replay, substitution, masquerade and Trojan horse attacks; tampering; overriding Yes/No response, etc.

The desire to develop a privacy-protective facial recognition system presented an excellent opportunity for OLG to practice "*Privacy by Design*." *Privacy by Design* (*PbD*) is predicated on the notion that technology can be enlisted to *protect* privacy, rather than encroaching upon it. Practicing *PbD* requires embedding internationally-accepted fair information practices and the 7 Foundational Principles of *PbD* directly into the design of technologies, at the architecture level [6]. *PbD* emphasizes the "positive-sum paradigm," in which it is recognized that embedding privacy measures need not weaken security, functionality or performance — quite the opposite. As opposed to a zero-sum paradigm, which brings unnecessary trade-offs and false dichotomies, *Privacy by Design* serves to enhance the overall design by creating technologies that achieve strong privacy without compromising performance — a doubly-enabling "win-win" outcome.

Biometric Encryption (BE) — explained in detail in the IPC white paper [1] — uses *Privacy by Design* to directly address the privacy and security concerns associated with biometric systems. BE is a process that securely binds a key to, or extracts a key from, a biometric, such that neither the key nor the biometric can be retrieved from the "helper data" (also called a "private template") created by this process and stored by the application, except upon presentation of the correct live biometric sample for verification. In essence, the key is "encrypted" with the biometric — a 'fuzzy' process due to the natural variability of biometric samples. The key can represent any value required by the particular application— for instance, it may be a cryptographic key or a pointer into a related private information database.

The concept of Biometric Encryption (BE) was first introduced in the mid-'90s by Tomko et al. [7]. In subsequent works, many BE solutions were proposed (more information on BE and related technologies can be found in [2, 8, 9]). It should be noted that while some of the BE solutions (see, for example, [10]) hinted at the possibility of a secure application, no explicit treatment of this type of construction has been considered by the existing solutions. In most cases, the cryptographic key was simply assumed to be the output of a BE verification algorithm. To the best of our knowledge, a BE application in a watch list scenario has never been discussed.

In general, Biometric Encryption schemes offer a number of advantages over traditional biometric systems, including [1]:

1. **Images, biometric templates and keys are not retained** — With BE, the user is always in control of his or her biometric — it is not stored (in either raw or template form) and therefore, can't be compromised. Further, the original biometric cannot be recreated (ideally) from the information that has been stored — it is untraceable.

2. **Multiple / cancellable / revocable identifiers** — BE allows a single biometric to be associated with any number of accounts and keys. Importantly, though, the nature of the helper data means that there should be no way to derive a common biometric which would allow someone to link and associate the accounts.

3. **Improved authentication security** — By securely binding account identifiers to a user's biometric, BE allows for these identifiers to be stronger (of greater length and complexity) and randomly generated, as there is no need for the user to remember them.

4. **Greater public confidence, acceptance, and use; greater compliance with privacy laws** — By including BE from the outset, user biometric data will remain under the exclusive control of the individual, minimizing the potential for identity theft and unwarranted surveillance, thereby increasing public confidence in the system.

5. **Suitable for large-scale applications** — Traditional large-scale biometric systems frequently entail storage of templates on centralized databases, which present a tempting target to prospective identity thieves. There is considerably less risk associated with storing private templates, however, since even if compromised the biometric cannot be derived (i.e. the private template becomes largely useless to a hacker).

In [3], University of Toronto researchers studied a range of issues with regard to the application of BE to a facial recognition system, including image pre-processing, feature extraction, cryptography, error correcting, and key binding, among others. Results of their simulation testing showed that BE could, in theory, be effectively integrated into a watch list facial recognition system. What remained was the practical development and deployment of such a system — an opportunity presented through a partnership with OLG and its self-exclusion program.

4. OLG FR + BE Application

As previously mentioned, for the OLG's self-exclusion program, an automated facial recognition system was determined to be the best technology to enhance the effectiveness of a manual inspection/detection process. First, such a system captures facial images at a distance, with no need for user interaction. This is an important consideration as, in Ontario, casino visitors do not generally need to provide identification upon entry — thus, a remote system is needed in order to preserve the current entrance experience for non-enrolled individuals. Secondly, a facial recognition system will be able to operate in conjunction with the legacy, photograph-based system, without the need for re-enrollment of individuals (which would require a visit by each enrollee to a gaming site). Other remote biometric modalities, such as iris-on-the-move or gait recognition, are not yet sufficiently advanced for OLG's application, and would not satisfy the legacy requirement.

It should be mentioned, however, that facial recognition at a distance is quite challenging: there are illumination, camera position, pose, etc. problems that seriously impact system accuracy. As a result, the accuracy numbers for live facial recognition significantly vary in the literature: from FRR ~ 1% — 3% at FAR = 0.1% in the controlled FRVTE NIST test [11] to FRR ~ 40% — 70% in the German Federal Criminal Police Office study at a railway station [12]. The conditions of the latter

test are much closer to the OLG environment. The accuracy of facial recognition has, however, significantly improved over the past decade [11], increasing the chance for a successful deployment (compared, for example, to the failed facial recognition test in Tampa in 2001 [13, 14]).

In relation to the above challenges, the U.S. Federal Bureau of Investigation (FBI) has blasted facial recognition technology in general for its alleged failure "to deliver the highly reliable verification required" [15]. While this opinion is not shared by most biometrics experts, it is true that live facial recognition does not have accuracy levels comparable to fingerprints, DNA, or iris scan for FBI applications, which require searches through a database of tens of millions of records. However, in the context of the OLG application, where the database size is much smaller (about 20,000 records), facial recognition is expected (based on the German test [12]) to identify at least two out of three self-exluded persons with a manageable rate of false alarms. While this level is less than ideal, it must be recalled that the current system of manual identification is significantly less accurate. As well, it will be shown in this section that a novel system design alongside a gradual approach of several field tests can help to bring the overall performance of facial recognition to a level acceptable for the OLG application. Finally, OLG was careful not to fully remove the human element of the identification system – in line with the U.S. National Academy of Sciences statement that "no biometric technology is infallible" [16] – by ensuring that the final decision of whether a match was declared was made by a human operator, and that a number of manual or alternative recognition methods (checking for license plates of enrollees, monitoring for use of enrollees' 'frequent player' cards, etc.) were kept in place. This recognizes, and mitigates, the potential fallibility of any biometrics-based system.

Along with the challenges associated with properly deploying a biometric system, the integration of Biometric Encryption into a facial recognition system, as required by the application proposed by OLG, is far from a trivial task, as it requires a re-engineering of the underlying architecture of a commercial facial recognition product. In this section, we also describe the issues faced by the collaborative team of researchers from OLG, the University of Toronto, iView Systems, and the IPC, the means by which they were addressed, the privacy protections made available through use of *Privacy by Design*, and briefly discuss the results of proof of concept testing of the system.

4.1 System Overview

The self-exclusion context at OLG is a 1:many (watch list) scenario in which the system must identify self-excluded individuals amongst a crowd of other patrons. Biometric Encryption alone is not recommended in a pure watch list scenario such as this, as the computing power required to perform the 1:many comparisons would be daunting. However, a standard facial recognition system can be used to reduce the normal 1:many comparison to a near-1:1 comparison, by filtering out all but the top few matches[3]. Thus, a system was developed which was composed of two distinct components (see Figure 1):

- A **watch list module**, that uses traditional facial recognition technology in a 1:many mode to produce a top-matches list for every patron walking into a casino. The list can contain zero or more potential matches, but typically has fewer than five; and,

3 A. Stoianov, private communication, August 2007.

- A **BE module,** that attempts to release keys for each of the subjects on the top matches list. If a key is successfully released, a match alert is generated for review by an administrator.

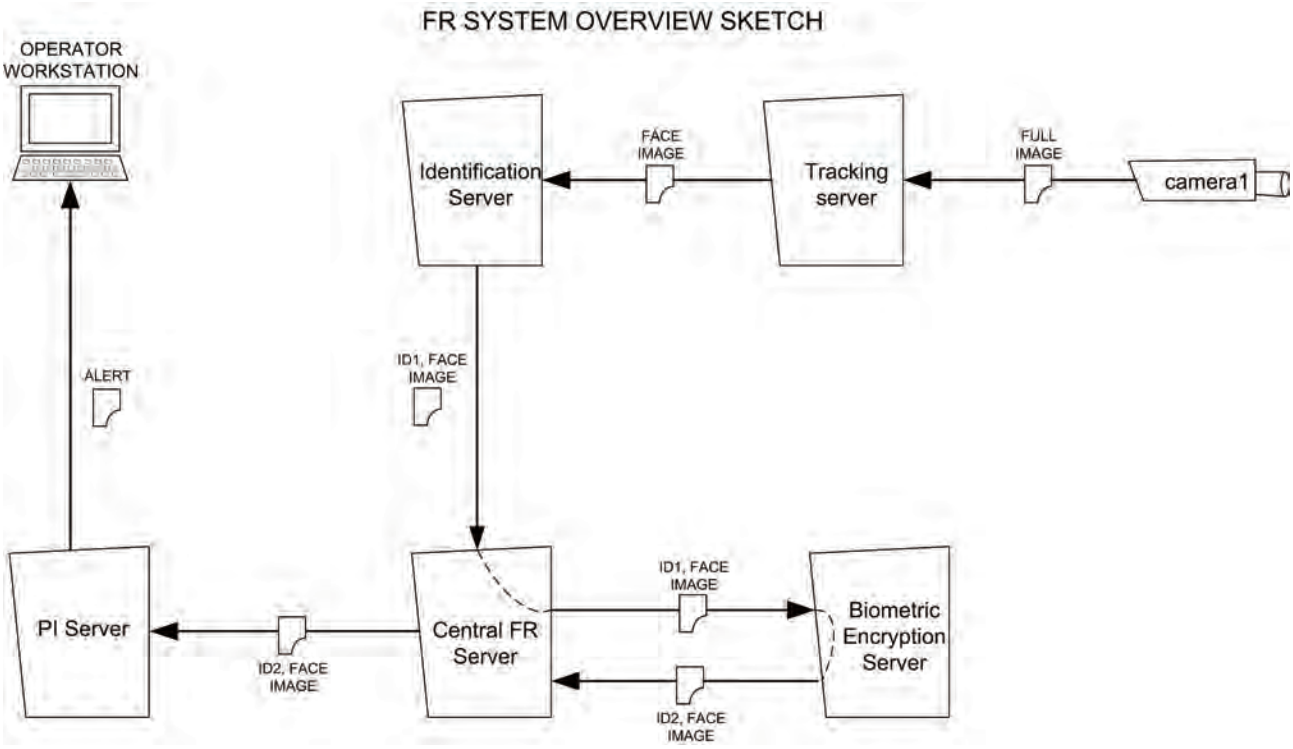

Figure 1: System Overview

Such a configuration has accuracy benefits on top of the privacy enhancements generally gained through the use of BE. For instance, in contrast to many biometric systems, the OLG operating scenario requires a minimized FRR (since this represents the rate at which enrolled self-exclusion subjects would go undetected and allowed into the gaming premises), while maintaining an acceptable FAR (as large numbers of false matches may increase staff frustration and inspire distrust in the system). In proof of concept testing (as described later in this document) of a watch list facial recognition system with the described configuration, the FAR results improved when compared to the watch list system without BE — with a minimal (or zero) increase in FRR. This can be understood by the fact that the BE module receives and evaluates a candidate list of identities from the watch list module — and thus cannot introduce any additional false acceptances (i.e., add to the list of potential matches), but may reject some. This is inherent in any system design that has a watch list identification module in series with the BE module, which acts as a second classifier. In all simulation cases, the BE module in fact rejected many imposter candidates, thus reducing the FAR.

4.2 Enrollment and Identification

In the proposed system design, the "self-excluded" subject identification is performed using a vendor-supplied facial recognition system - iGWatch from iView Systems, which uses an FR algorithm SDK from Cognitec Systems, Germany. A biometrics-based cryptosystem (another term for BE) is implemented in tandem to offer privacy protection of the subject's personal information by way of a bound pointer key.

As shown in Figure 2, in the enrollment phase, the subject's facial image is captured, and he or she is assigned a non-meaningful, unique enrollee ID (**id**). A commercial facial recognition system then extracts biometric features and generates a template (**t1**) that is stored in the face recognition database (**FR database**), indexed by the enrollee ID. Another set of biometric features (**t2**) is sent to the BE key binding algorithm, which creates BE helper data (or a 'private template') from the biometric data and a pointer key. This pointer key represents the location of the subject's facial image and other personal information (PI) within a database of self-exclused individuals (**SE database**), and is normally generated at random. Finally, the BE helper data, bk(k, t2), is stored in another database (**helper database**), again indexed by the same enrollee ID. For the OLG implementation t1 and t2 use different facial algorithms to extract features and are not interoperable, which is important from a security standpoint.

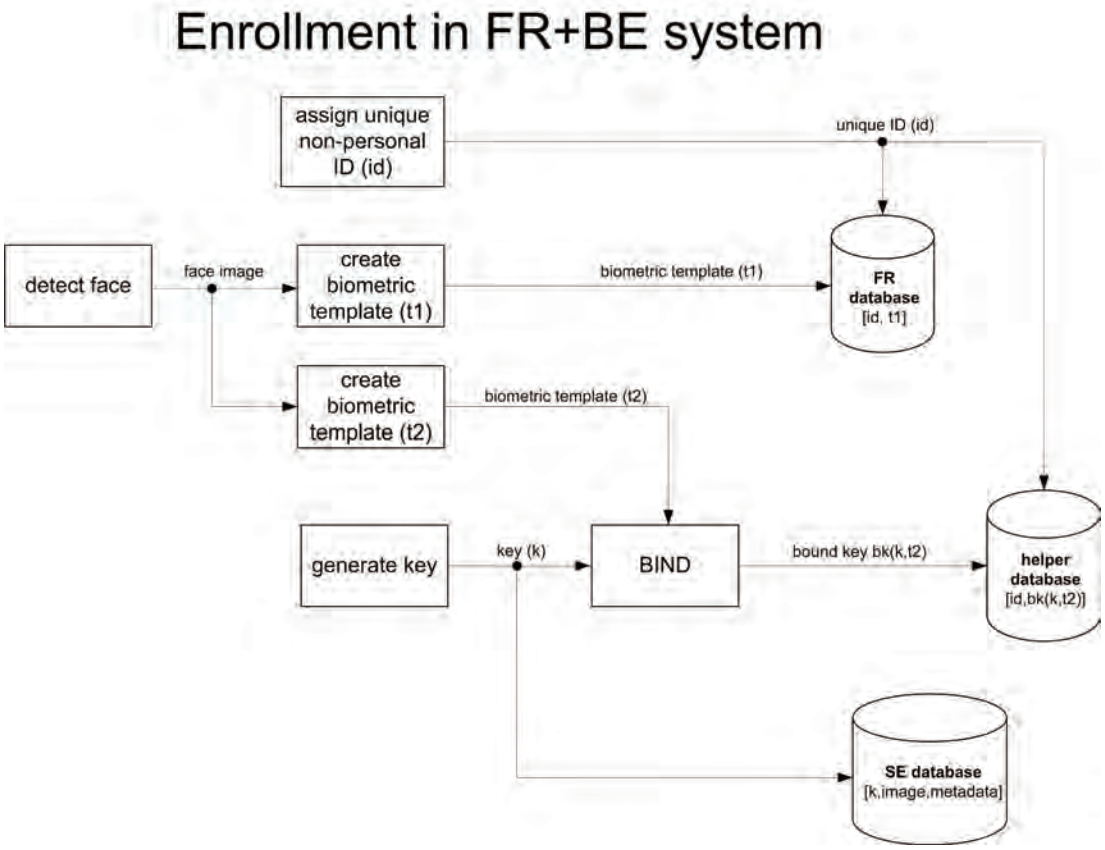

Figure 2: Enrollment in FR + BE system

During identification (Figure 3), the vendor-supplied system will attempt to do a one-to-many match of subjects entering the monitored facility to those in the FR database. This "Preliminary Identification" stage is typical of a "watch list" scenario. The enrollee IDs of the top matches are then output to the BE key retrieval algorithm. If a key can be retrieved from the BE helper data associated with one of the potential matches, the final verification stage is entered. Here, the pointer to the stored personal information (including photo) associated with the potential match is regenerated (from the BE helper data), and the record at that location is retrieved. An operator then manually compares the retrieved facial image with the image of the casino patron in question. As such, the

final decision of whether to approach a person and ask to confirm his or her identity is left to a human operator. It is important to note, with regard to the privacy of those patrons not registered in the self-exclusion program, that though an image of each individual entering the facility will be captured and analysed by the facial recognition system, no captured or derived information (e.g. images or biometric templates) is stored by this system should there be no match identified. In the event that a match is declared by a human operator, the captured image is planned to be securely stored in the system for one year to be provided in cases of legal challenges.

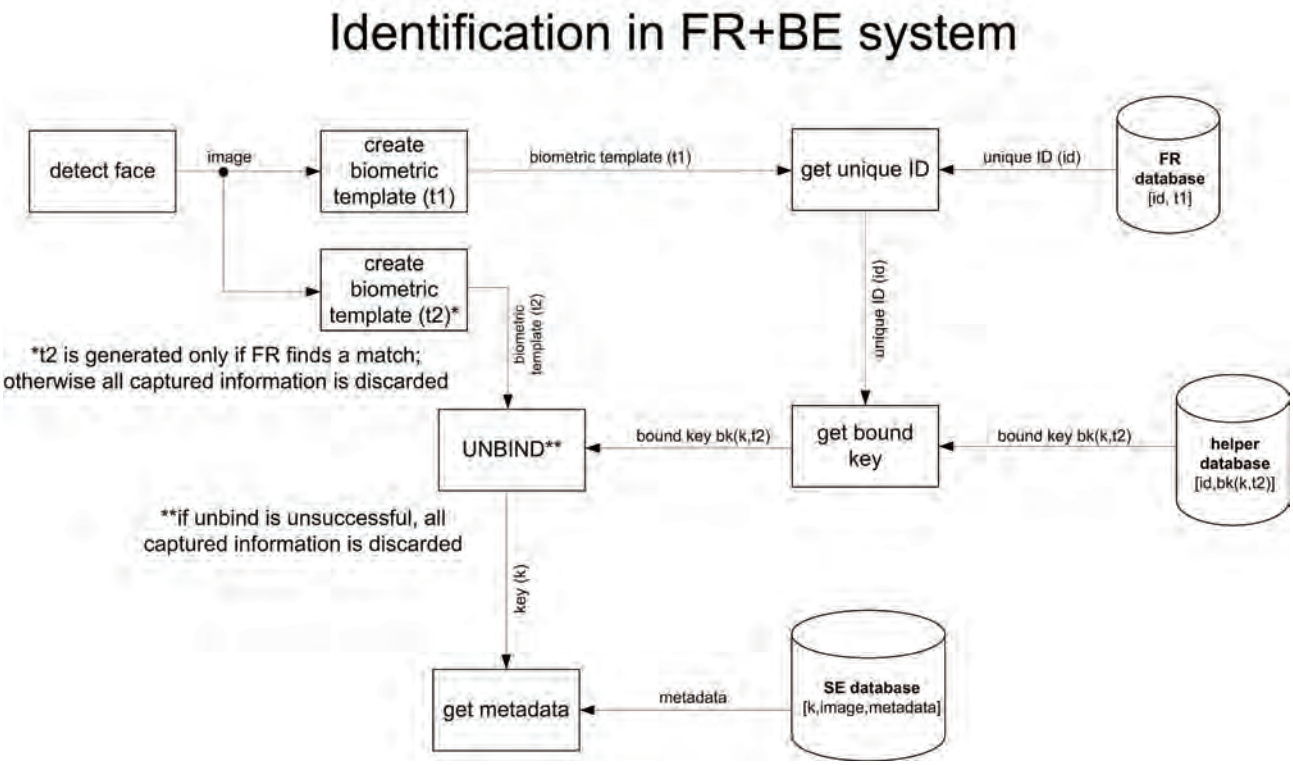

Figure 3: Identification in FR + BE system

4.3 Privacy Protections

As stated, this is the first BE application proposed for a biometric watch list scenario. This scenario, though, differs from most other watch lists. A common application of a watch list involves the detection of subjects who have been identified as posing a risk to public safety or security. In such a system, the primary privacy issues are associated with the general public. It is important that any information — such as captured images or biometric templates — related to non-watch listed individuals is not stored. Such a system must also have a sufficiently low false acceptance rate (FAR) that non-watch listed individuals would rarely be approached by security personnel to confirm their identity. The privacy situation is quite different, however, for the OLG Self-Exclusion Program, which consists of individuals who have voluntarily put themselves on the watch list. In addition to maintaining the privacy of non-enrolled individuals (which OLG accomplishes by not storing any captured or derived information when the system does not find a match), the personal information of enrollees should enjoy the highest possible level of privacy protection (similar to that of health records) while it is in OLG's custody. BE can be a significant aid in achieving this important standard.

The above-described system architecture uses several techniques to increase the privacy and security of the enrollee records throughout the system (as shown in Figure 4). Conventional cryptography is used to encrypt all images in the PI database; these images must be stored in order for a small set of authorized users to use non-biometric means (e.g. visual comparisons) to spot or verify a self-excluded patron. The databases themselves are encrypted to protect data while it is 'at rest,' while communication between clients and the databases also occur through encrypted pathways, to protect data 'in motion.' Both the databases and the application clients are access controlled, and the entire system is deployed on a secure internal network, which is protected from the external world.

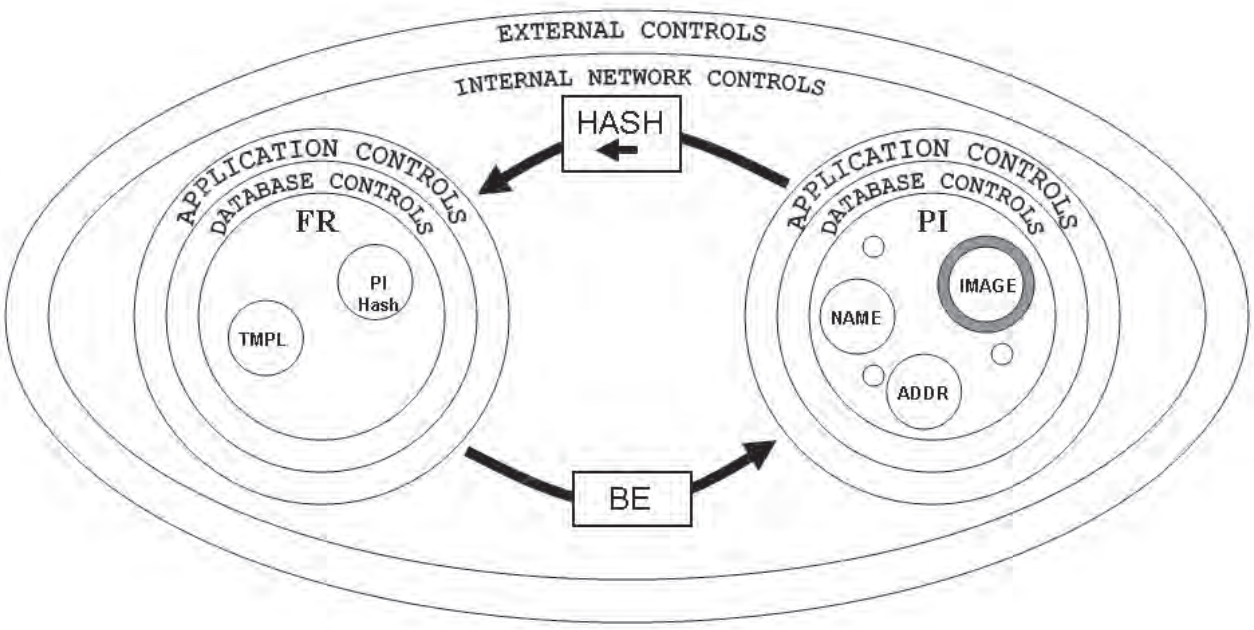

Figure 4: *Privacy by Design*: A multi-layered approach to privacy

One of the principal privacy protections of this system, though, is that the link to the photo and other personal information of a self-excluded person can only be determined by accessing a key that is bound to the person's biometric. In order to reveal the stored information, the BE key retrieval algorithm must be able to regenerate a biometrically encrypted pointer key. To achieve this, the person's live facial image is required — control, thus, rests with the individual. This control also makes it much more difficult for the information to be linked with other, third party databases without the user's consent. The OLG system further uses different template generation algorithms which are independent of each other, to ensure that the two biometric templates (used by the facial recognition and BE modules, respectively) are not the same or interoperable. This prevents the possibility that the vendor's template could be transformed and used to retrieve a key from the corresponding BE helper data.

In the OLG application, the connection between a PI record and the FR database must also be accessible, to allow for updates to records in the FR database (for example, to de-enroll someone from face recognition). To ensure that BE cannot be circumvented to reveal the link from an FR record to the PI record, a one-way hashing algorithm is used to reveal the link between a PI record and a corresponding FR record (see Figure 5).

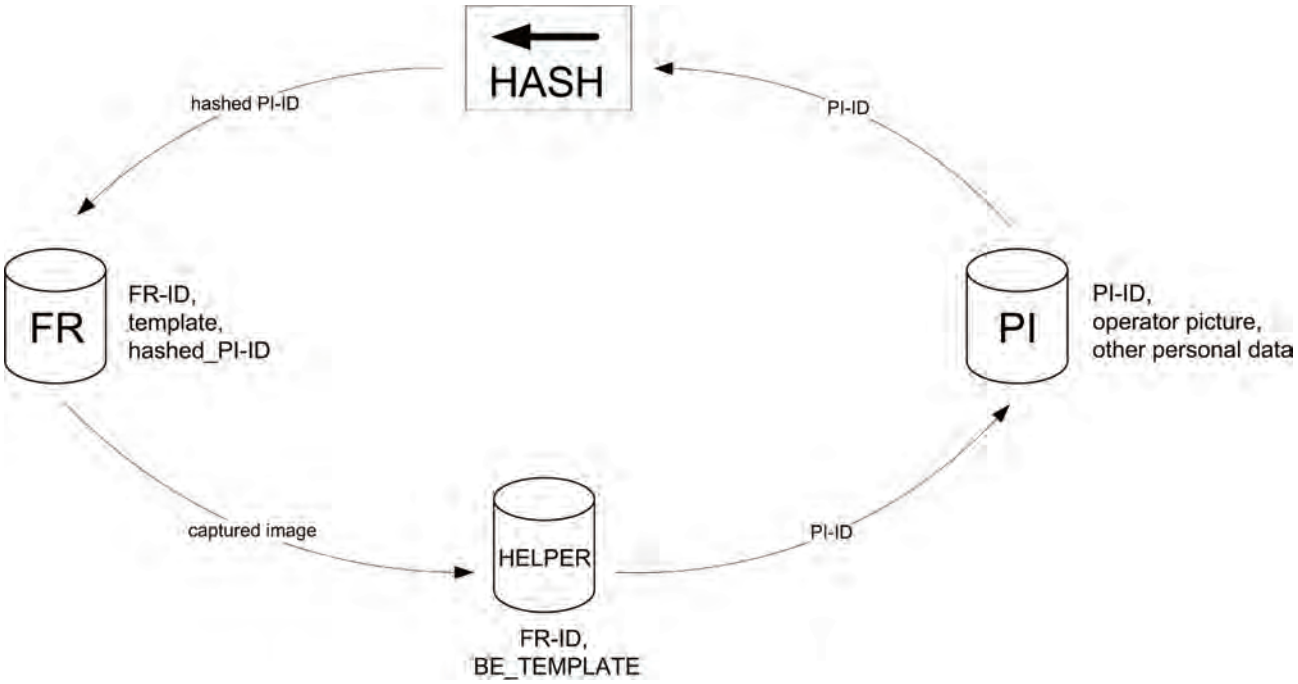

Figure 5: BE in the proposed OLG application

This architecture has the benefit of allowing for the sharing the personal data with some user-authorized third parties, such as a self-help workgroup of gambling addicts. In this situation, the self-excluded person retains full control over the data, since accessing it would require his physical presence at the workgroup. Further, by providing an extra security layer, the system with BE offers better safeguards over information either stored or transmitted to a third party. To access and decrypt personal data, a would-be attacker would face the added task of cracking the BE helper data. Even if successful, such an attack would be limited in scope, as though a 'crack' would cause a breach of one person's information, the rest of the system would remain uncompromised (as there is no single key or template to decrypt all records).

Finally, BE may also enhance privacy protection for the casino patrons who are not self-excluded — the general public. BE works as a second classifier that follows the vendor's biometric engine. As shown in [3], this can substantially reduce the system FAR without a significant impact on FRR. In other words, fewer legitimate users will be inconvenienced by identity checks, as compared to the use of a facial recognition system without BE.

It is acknowledged that a facial recognition BE solution for a watch list system, such as that proposed for OLG's application, does not necessarily capture all the privacy benefits that could be provided by BE in a 1:1 application [1, 2]. As a biometric modality, facial recognition is weaker in terms of entropy than, for example, iris and fingerprints [2, 17]. This means that only a relatively short key can be bound to the biometric. Overall, face-based BE is also more susceptible to offline attacks. As a caution, developers and system administrators should be aware that facial images are more generally available than other biometrics (through social networking sites, etc.). As such, the potential should be addressed that such an image could be reformatted and submitted to the facial recognition system in order to retrieve the stored information.

With regard to the OLG application in particular, privacy concerns may arise from the fact that personal data, including facial images, from the OLG self-exclusion program are already, and will likely continue to be, held by OLG in a central database(s), if only for data protection and backup purposes. Access to these databases will be limited and strictly regulated, however. Further, access to personal data which bypasses the BE system must be available, for a number of purposes. Should a self-excluded individual need to update his or her personal information (such as address, to remain free from marketing materials), it would be a poor practice to require the individual to visit a gaming site to present his or her biometrics. As such, a mechanism is required to allow non-BE access to a self-excluded person's own data. OLG may also use some non-biometric means to spot a self-excluded person, such as automobile license plate numbers or monitoring use of the person's loyalty card number (some self-excluded people still try to use the card after sneaking into the facilities). In this case, as well, the system operator will need to be able to effectively bypass the facial recognition system (with or without BE) to access some of the self-excluded person's data.

Overall, though, it is felt that BE in the OLG watch list scenario, even given any real or potential challenges, provides significant *privacy protection* both to self-excluded persons and to the general public. This privacy protection is achieved mainly on the operational level. BE also offers a more *secure system architecture* — an attacker must still penetrate all of the standard security safeguards in order to access BE helper data. BE may even improve the *overall accuracy* of the watch list system. In other words, BE could bring a "triple-win" advantage to a conventional watch list system by transforming its surveillance nature.

4.4 Proof of Concept

The system described above has moved beyond the conceptual phase. In fact, proof of concept testing has allowed OLG, in collaboration with iView Systems, U of T and the IPC, to demonstrate that a face recognition application with BE is viable for development and deployment in a casino environment. As this was a first-of-its-kind effort, expectations were unknown. The proof of concept testing had the following main steps and results:

- Facial recognition technology was proven effective in a casino environment through several tests in real field conditions (up to 20,000 visitors on Saturday) involving OLG control group participants (simulating the self-excluded persons) and the general public. The Correct Identification Rate (CIR)[4] was increased to a maximum of approximately 91% from a starting point of approximately 30%. The CIR of 91% is a best case result achieved by controlling participant pose; a more realistic CIR range for the field is between 60% and 80%. The advances in CIR were achieved using a measured approach of several field tests and were mainly due to corrections in lighting, camera position and subject pose through the use of "attention-getting" devices like marketing screens. This compares positively to, for instance, a 2007 German Federal Criminal Police Office study which achieved a 30-60% recognition rate for a facial recognition watch list field tested at a railway station [12].

4 CIR = 1 - FRR

- Biometric Encryption did not decrease the efficiency of face recognition in a pipeline test. Positive matches from face recognition were fed into the BE system; the BE system marginally affected the original CIR (by less than 1%) while reducing the FAR by 40% to 50%. This result was an unexpected benefit, which, as described prior, occurs due to the status of the BE module as a secondary classifier.
- Face recognition was field tested using the actual OLG self-excluded images to determine the degree to which detection rates would improve. Preliminary results show that FR is a valuable tool in the overall goal of assisting self-excluded patrons from staying out of gaming facilities.
- The system architecture was successfully created to integrate BE into a commercial face recognition product (iGWatch from iView Systems) while maintaining OLG's core requirements. This architecture treated BE as one component in a multi-layered approach to privacy and security of the overall system.

5. Conclusions

The Ontario Lottery and Gaming Corporation's self-exclusion program was identified as an ideal opportunity to deploy, for the first time, Biometric Encryption (BE) as a Privacy-Enhancing Technology in a biometric watch list scenario. The system, as designed, sequentially combined a commercial, one-to-many face recognition system with a BE module. This use of BE as a secondary classifier was shown to enhance patron privacy (both for those on the watch list, and regular patrons), system security, and even overall accuracy of the watch list system within the context of the OLG self-exclusion program. Based on the results of a field test of the system, it was also shown that facial recognition technology can contribute to OLG's program objectives. This technology would serve as one part of OLG's Responsible Gaming program, to assist self-excluded patrons to keep their self-expressed commitment not to enter gaming sites.

The development of a facially-oriented, Biometric Encryption application that may be integrated with commercially-available facial recognition systems holds great promise. We firmly believe that this exciting and innovative project will generate considerable interest for other applications at OLG, and from other casinos across the country, and around the world. This leading-edge research should foster the development of a commercially-available product that will facilitate the conduct of responsible management with respect to gaming and privacy — a positive-sum approach.

References

[1] A. Cavoukian and A. Stoianov, "Biometric Encryption: A Positive-Sum Technology that Achieves Strong Authentication, Security AND Privacy", March 2007. http://www.ipc.on.ca/images/Resources/bio-encryp.pdf

[2] A. Cavoukian and A. Stoianov, "Biometric Encryption: The New Breed of Untraceable Biometrics". Chapter 26 in N. V. Boulgouris, K. N. Plataniotis, E. Micheli-Tzanakou, eds.: Biometrics: fundamentals, theory, and systems. Wiley - IEEE Press, pp. 655 - 718, 2009.

[3] Martin, K., Lu, H., Bui, F., Plataniotis, K. N. and Hatzinakos, D., "A biometric encryption system for the self-exclusion scenario of face recognition," *IEEE Systems Journal: Special Issue on Biometrics Systems*, vol. 3, no. 4, pp. 440-450, 2009.

[4] P. Hancock, V. Bruce, and A.M. Burton. "Recognition of Unfamiliar Faces". Trends in Cognitive Science, v. 4, No. 9, pp. 330-337, 2000.

[5] N. K. Ratha, J. H. Connell, and R. M. Bolle. "Enhancing security and privacy in biometrics-based authentication systems". IBM Systems Journal, v. 40, No. 3, pp. 614–634, 2001.

[6] A. Cavoukian. "The 7 Foundational Principles: Implementation and Mapping of Fair Information Practices", May 2010. http://www.ipc.on.ca/images/Resources/pbd-implement-7found-principles.pdf

[7] G.J. Tomko, C. Soutar, and G.J. Schmidt. "Fingerprint controlled public key cryptographic system". U.S. Patent 5541994, July 30, 1996 (Priority date: Sept. 7, 1994).

[8] A. K. Jain, K. Nandakumar, and A. Nagar. "Biometric Template Security". EURASIP Journal on Advances in Signal Processing, v. 2008, Article ID 579416, pp. 1-17, 2008.

[9] P. Tuyls, B. Škorić, and T. Kevenaar, eds. "Security with Noisy Data: Private Biometrics, Secure Key Storage and Anti-Counterfeiting." Springer-Verlag, London, 2007.

[10] N. Delvaux, J. Bringer, J. Grave, K. Kratsev, P. Lindeberg, J. Midgren, J. Breebaart, T. Akkermans, M. van der Veen, R. Veldhuis, E. Kindt, K. Simoens, C. Busch, P. Bours, D. Gafurov, B. Yang, J. Stern, C. Rust, B. Cucinelli, and D. Skepastianos, "Pseudo identities based on fingerprint characteristics". In IEEE 4th international conference on intelligent information hiding and multimedia signal processing (IIH-MSP 2008), August 15 - 17, Harbin, China, 2008.

[11] P. J. Phillips, W. T. Scruggs, A. J. O'Toole, P. J. Flynn, K. W. Bowyer, C. L. Schott, and M. Sharpe. "FRVT 2006 and ICE 2006 Large-Scale Results". NIST Report NISTIR 7408, March 29, 2007.

[12] Bundeskriminalamt. "Face recognition as a search tool." Final Report. http://www.eucpn.org/download/?file=GER%20FAce%20Recognition.pdf&type=14

[13] K. Bonson and R. Johnson. "How Facial Recognition Systems Work." http://electronics.howstuffworks.com/facial-recognition.htm

[14] S. Olsen. "ACLU decries face-recognition tools." http://news.cnet.com/ACLU-decries-face-recognition-tools/2100-1023_3-800864.html

[15] David Moss, "FBI techs shy away from facial recognition: spends 40 years losing face", Nov. 3, 2009. http://www.theregister.co.uk/2009/11/03/fbio_face_recognition/

[16] "Biometric Recognition: Challenges and Opportunities". J. N. Pato and L. I. Millett, Editors. Whither Biometrics Committee, Computer Science and Telecommunications Board, Division on Engineering and Physical Sciences, National Research Council. Sept 24, 2010. http://www.nap.edu/catalog/12720.html

[17] R. Youmaran, A. Adler, and S. Loyka, "Towards a measure of biometric information ". Can. Conf. Computer Elec. Eng. (CCECE), Ottawa, Canada, May 7-10, 2006.

About the Authors

Ann Cavoukian, Ph.D., Information and Privacy Commissioner of Ontario, Canada

Dr. Ann Cavoukian is recognized as one of the leading privacy experts in the world. Noted for her seminal work on Privacy Enhancing Technologies (PETs) in 1995, her concept of *Privacy by Design* seeks to embed privacy into the design specifications of technology, thereby achieving the strongest protection. An avowed believer in the role that technology can play in protecting privacy, Dr. Cavoukian's leadership has seen her office develop a number of tools and procedures to ensure that privacy is strongly protected, not only in Canada, but around the world.

Dr. Cavoukian is also one of the foremost proponents of Biometric Encryption. She has co-authored, with Dr. Alex Stoianov, several works on the topic, including a book chapter (in 'Biometrics: Theory, Methods and Applications') and an article for the Encyclopaedia of Biometrics. Dr. Cavoukian sits on the European Biometrics Forum's International Biometric Advisory Council, and serves as the Chair of the Identity, Privacy and Security Institute at the University of Toronto, Canada. Recently reappointed as Commissioner for an unprecedented third term, Dr. Cavoukian has made promotion of the development and commercialization of Biometric Encryption — a positive-sum, *Privacy by Design* technology — a priority.

Tom Marinelli, Ontario Lottery and Gaming Corporation

Tom Marinelli is the Acting Chief Executive Officer and most recently Chief Information Officer for the Ontario Lottery and Gaming Corporation (OLG). Mr. Marinelli, P.Eng., has served OLG in a number of progressively senior positions. His prior position was Vice President, Gaming Support. Mr. Marinelli has previously served as Vice President, Corporate Business Optimization and, prior to that, Chief Technology Officer. As Chief Information Officer, Mr. Marinelli was accountable for the effective and secure use of OLG's information technology and the enablement of business opportunities through information and information technology.

Information and Privacy Commissioner of Ontario
2 Bloor Street East, Suite 1400
Toronto, Ontario
Canada M4W 1A8
Telephone: (416) 326-3333
Fax: (416) 325-9195
E-mail: info@ipc.on.ca
Website: www.ipc.on.ca

Ontario Lottery and Gaming Corporation
4120 Yonge Street, Suite 500
Toronto, Ontario
M2P 2B8
Telephone: (416) 224-1772
Fax: (416) 224-7000
E-mail: olgcontactus@olg.ca

November 2010

http://www.olg.ca | http://www.privacybydesign.ca

Information and Privacy Commissioner, Ontario, Canada