PICO: Privacy through Invertible Cryptographic Obscuration

T. E. Boult
University of Colorado at Colorado Springs and Securics Inc.
tboult at vast . uccs . edu

Abstract

Significant research progress has been made in intelligent imaging systems, surveillance and biometrics-- improving robustness, increasing performance and decreasing cost. As a result, deployment of surveillance and intelligent video systems is booming and increasing the impact of these on privacy. For many, networked intelligent video systems, especially video surveillance and biometrics, epitomize the invasion of privacy by an Orwellian "big brother". While tens of millions in government funding have been spent on research improving video surveillance, virtually none has been invested in technologies to enhance privacy or effectively balance privacy and security.

This paper presents an example that demonstrates how using and adapting cryptographic ideas and combining them with intelligent video processing, technological approaches can provide for solutions addressing this critical trade-offs, potentially improving both security and privacy. After reviewing previous research in privacy improving technology in video systems, the paper then presents cryptographically invertible obscuration. This is an application of encryption techniques to improve the privacy aspects while allowing general surveillance to continue and allowing full access (i.e. violation of privacy) only with access to a decryption key.

1. Introduction and Background

This introduction discusses video surveillance, its effectiveness and some perceptions on privacy. Then it discusses video/privacy issues in multi-media systems. Though related, the different objectives of these two application areas create significant "privacy" issues. After discussing the two areas, the paper very briefly reviews the legal issues of privacy in multi-media systems. The introduction concludes with an overview of related research in privacy preserving/enhancing technologies for video systems.

There are many significant privacy concerns with vision and biometric systems. [Thieme-03] divides them into personal privacy and information privacy concerns. While some of the informational privacy concerns can be addressed with system level information security, the personal privacy concerns are inherent. At the boundary between them is the issue of long term storage of imagery and biometrics, which contain immutable personal data – if the DB is compromised or the person's data is compromised through other means (e.g. Insiders), the loss is permanent. The very long-term nature of privacy concerns requires we treat them with special care in designing systems.

Before the background we introduce our viewpoint, that privacy and security are not always adversarial goals. Traditional locks and keys for group resources provide a good example. Locks improve both the privacy and security of group members. The increase in security, compared to no locks, is obvious though many people greatly overestimate the true security. (The majority of locks used are compromised with easily obtained information.) From a privacy point of view, the members have increased confidence that only group

members have access, hence the actions inside the locked have increased privacy expectations. Furthermore, the physical key is almost always allocated such that it provides pseudo-anonymity, i.e. you know they are in a group but not which member. Even if the list of key holders was compromised, who used the key is not maintained and hence an individual's usage is private. Thus a master-lock or group-key facility increases both privacy and security. They traded convenience--it requires one to carry a key.

Compare this to a vision, e.g. face-based, "access control system". In this system the "data" is clearly traced to an individual (no anonymity/privacy), and for vaguely stated "security reasons" most of the systems maintain detailed logs of who entered when. At a security level, the added tracking may provide some value (though many consider it of little value), but the potential privacy implications are large, e.g. using it to determine if people are not at work on time or harass particular employees or subgroups. (It is interesting to note that many biometric systems are sold based on time and attendance rather than for improved security.) If the underlying vision-based database is compromised, data about the individual is (permanently) lost – unlike a physical lock/key, and the face cannot be changed. Here an increase in convenience (and potentially) security came at a loss of privacy.

Other abuses are the functional creep of imagery especially when combined with facial recognition. E.g. [Krause-01] talks of a government plan to take the facial image database (DB) of all Colorado drivers and sell them to be used for facial recognition. This clearly brings up issues of invasion of privacy and increase in power to the police (hence raising chances of its abuse). The same paper [Krause-01] also gives an instance where the police in Tampa, Florida used face recognition software on football fans at the Super Bowl without the knowledge of the people involved. [Agre-03] argues why face recognition systems should not be used in public places. Different organizations could easily be networked to use the faces captured by the face recognition system to find out all the places that a person has gone, or to scan over very large sets of surveillance images for particular individuals. There are times when the large-scale search is warranted, but without constraints the system is open to widespread abuse.

So as we look at vision systems, which often contain data/imagery about individuals, we believe it is the fundamental responsibility of the vision system designer to consider privacy issues and see if they can improve both security and privacy or at least make clear the tradeoffs implemented in the system.

1.1. Video Surveillance Background

Video surveillance is becoming more and more common all over the world. People are forever under the watchful 'eye' of the camera even as they go through their day-to-day activities. CCTV is widely used for surveillance in banks, parking lots, shopping malls, airports, and other public places. In the past decade, the use of CCTV has grown to unprecedented levels. A decade ago, Britain was spending between $225M and $450M US dollars *per year* installing estimated 300,000 cameras. Growth in the market is estimated at fifteen to twenty per cent annually, [Davies-96].

Despite massive adoption of such technology, there has been no solid evidence that shows surveillance cameras have had an overall deterrent effect. In fact, surveillance may only serve to displace crime. Richard Thomas, Acting Deputy Chief Constable for Gwent, U.K., recently told the BBC that he believed video surveillance simply pushed some crime beyond the range of the cameras [Flaherty-98]. The formal studies which have been published have been characterized as "...post hoc shoestring efforts by the untrained and self interested practitioner" [Pawson-Tilley-94].

The idea of displacement is not new to video surveillance; it's a long-standing issue. But it does mean that when designing vision-systems for security, broader questions need to be asked. While many schools have installed video surveillance, it has had little impact. A 1993 USA Weekend survey reported that 2,000 students were physically attacked each hour of the school day [Ansley-93]. The areas that are being watched (physically or by video) have reduced the incidents, but the areas that are not (because of cost) or cannot (because of privacy) have become havens. In the USA survey, nearly half of those surveyed said they avoided school restrooms out of fear [Ansley-93]. In addition to addressing improving vision-based security systems performance, reducing cost and increasing functionality, it is also important that we develop solutions that will solve the privacy problems or we will just push the problems into the unprotected areas thus reducing privacy overall with little overall security gain.

On the positive side for security, however, the cameras are also creating a vastly increased rate of conviction after crimes are detected. Once people know they have been videotaped, many admit the offense immediately, [Priv-05]. For this to be effective, and to lead to a significant deterrent, the video must have sufficient resolution/quality for prosecution. This means the systems must be designed to provide "evidence" not just data.

There are many privacy issues in surveillance. While some "invasion" is unintentional, or even just potential, as stated in [Senior-et-al-03,] the personnel who are in-charge of scanning these video images are often either ignorant about their job or tend to misuse their powers, for example engaging in voyeurism. While cameras in airport or school bathrooms might improve security, the potential abuses prohibit their use.

Not all privacy invasions are about video in special settings. Many people also have concerns to simply being in the video, even in their workplace or the city street, because they have with no control over who can use the data or for what. The extent of concern was highlighted in a survey commissioned by the UK Home Office [Honess-Chaman-92], which found that more than fifty per cent of people felt neither government nor private security firms should be allowed to make decisions to allow the installation of CCTV in public places; 72 per cent agreed "these cameras could easily be abused and used by the wrong people"; 39 per cent felt that people who are in control of these systems could not be "completely trusted to use them only for the public good"; 37 per cent felt that "in the future, cameras will be used by the government to control people". While this response could be interpreted a number of ways, it goes to the heart of the privacy and civil rights dilemma. More than one respondent in ten believed that CCTV cameras should be banned.

1.2. Privacy issues in multi-media environments

Interactive multi-media (video and audio) environments raise their own privacy concerns. Of course the same legal issues apply (so audio should be used very carefully). While at one level many of the systems have an implied "consent" since the user is interacting, the real issue of privacy is how the data is used/stored outside the local operational context.

As reported in [Adams-Sasse-01], the CHI99 panel, *Trust me, I'm accountable: trust and accountability online* provided clear separation of two positions within the community:

1. "As the new technology environments develop, users will adapt their privacy expectations and behaviors"

2. "Privacy is a complex problem, but it will not go away. To design successful applications, we have to acknowledge the problem and start tackling it, proactively".

The first position, loosely described as either ignore it or let the users adapt, is done with some risk – HIPPA has shown that long-term privacy issues tend to be resolved in favor of privacy and if privacy rules change after systems are deployed it can have significant ramification and cost. It also ignores trends in user interfaces, where people adapt when the interface is convenient and resist when it is not. A few examples of privacy failure, even at the level of jokes about others' performance, can doom a system. The importance of user feedback on, and control of, potentially invasive/private information in interactive systems is well documented, e.g. [Bellotti-96, Lee-et-al-97; Smith-Hudson-95]. Yet, as noted by [Davies-97], most privacy research to date has focused on policies and mechanisms around the concept of personal information - data that can be used to identify an individual. [Adams-Sasse-01] argues "such a data-centric approach cannot work well in the domain of multimedia communications. The majority of data in this field allows identification of a person (e.g. video image, voice patterns). Labeling all audio and video data as personal information -- and thus declaring it to be off limits -- is hardly practical." What is needed is an approach that balances the need to use the data with the privacy expectations.

[Adams-Sasse-99] report an example where those installing a multimedia application judged the situation (staff common room) as public, and thus saw no problem with broadcasting images over the Internet. The users, however, regarded the situation as private or semi-private and felt their privacy was being invaded through the installation of a camera. The result was an emotive rejection of the technology and decreased trust in those who had introduced it.

In settings where management has access to this type of data, there is often concern that the multi-media data can be used, out of context, for surveillance and to secretly assess employees.

At a "government" level there is also room for serious concern. As The Patriot Act showed, the government can play its "security" card and often trump the privacy of individuals. The more aggressive activities where the executive branch ordered the NSA to conduct electronic surveillance of American citizens without court approval, feeds the concerns citizens and justify the concerns of privacy advocates.

At a minimum, clearly articulated and publicly posed privacy policies are needed. But realistically, technology that can help enforce that policy will increase trust in the system. To take the prevailing attitude that the data is freely available within the collecting organization but not available outside ignores the potential for internal privacy violations and intentional or duplicitous disregard for that policy.

1.3. Legal issues around privacy and multi-media

There is a plethora of legal issues related to multi-media surveillance, with different issues at the country/federal level and different rules for each state with the US. We briefly summarize them here and use California as a state example. I am not a lawyer and this is not legal advice and should not be used as conclusive decisions or interpretations.

In the US, Title I of the Electronic Communications Privacy Act ("Title I", which replaced the older Title III) limits wiretaps by the government, including law enforcement. Under Title I, government must obtain warrants prior to secretly intercepting some communications. Any video surveillance that also has an audio component must comply with the Title I. If video surveillance device can intercept sound, and the surveillance constitutes a search, the police must first obtain a warrant prior to the installation of the

device. Title I also limits government use of recordings collected by individuals or companies who may have installed the device for other reasons.

There is a general rule, however, that applies to conversations on video camera tapes, even in the workplace. Regardless of the state, it is almost always illegal to record or disclose a conversation to which you are not a party, do not have consent to tape, and could not naturally overhear. According to [RCFP-05], that is pretty much the definition of "eavesdropping", and site includes summaries of each state's standing.

The Fourth Amendment prohibits unreasonable searches and seizures. In Katz v. United States, the Supreme Court declared, "the Fourth Amendment protects people, not places." … "What a person knowingly exposes to the public, even in his own home or office, is not a subject of Fourth Amendment protection," but, "what he [that person] seeks to preserve as private, even in an area accessible to the public may be constitutionally protected." [US369]

It is important to note that the courts in interpreting this have found "Generally, one walking along a public sidewalk or standing in a public park cannot reasonably expect that his activity will be immune from the public eye or from observation by the police" [McCray-State-90]. Thus for the most part, for "silent video", it is interpreted as allowing video systems in public.

A search that reveals information within a private place that could not be discerned by the naked eye violates the Fourth Amendment protection against unreasonable searches. In United States v. Karo, the Supreme Court held unconstitutional the monitoring of the movement of a container of chemicals inside various houses by the police. The Court concluded "[in]discriminate monitoring of property that has been withdrawn from public view would present far too serious a threat to privacy interests in the home to escape entirely some sort of Fourth Amendment oversight."[US468] (Reasonable zoom lenses have been found acceptable, but zero-visible light systems have been questioned with mixed results in lower courts).

The Supreme Court has adopted a two-part test to determine whether or not policy activity constitutes a search of an individual: (1) Has the individual manifested a subjective expectation of privacy? and, (2) Is society prepared to recognize that expectation as reasonable or legitimate? This test balances the privacy interests of individuals against society's desire to maintain effective law enforcement. It is why, in general, a public bathroom is still off-limits while the hallway directly outside is not. But the balance also depends on the available options and societies expectations – as we become increasingly numb to video surveillance invasions of privacy, the second part of the test above is continuously weakened. This is also why most "interactive" video systems or access control systems are currently subject to "privacy" legal standing, because the user expectation is that it is not private. Technology that improves privacy while supporting security could further improve privacy by resulting in the courts seeing the potential of the new technology as changing societal expectations.

Federal rules only limit the use of these technologies by the government itself. For corporate/state use, the rules are state specific and include government restrictions, see [RCFP-05], and the broader tort liability issues. California courts have recognized an action under tort law for the invasion of an individual's privacy, e.g. see [Dieteman-Time-71]. In California Tort law, the right to privacy encompasses four distinct torts: (1) unreasonable intrusion into a person's solitude or into her private affairs; (2) publicity which places an individual in false light in the public eye; (3) public disclosure of true, embarrassing facts; (4) unauthorized commercial use of an individual's personage, e.g. see [Johnson-Hartcort-75]. Though there has been no major case-addressing tort (3), it is expected to provide protection of some aspects of privacy even in the "public" areas of a workplace.

There are also specific state laws related to camera, especially those with audio abilities, and 24 states do have specific hidden camera laws. Also, some states have laws REQUIRING a notice or posting that there is surveillance equipment. And regardless of whether a state has a criminal law regarding cameras, undercover recording in a private place can prompt civil lawsuits for invasion of privacy.

Outside the US, the legal issues have a wide range, but generally similar characteristics exist. Post 9/11, there has been a generally shifting, by most governments, to allow even greater uses of surveillance of their citizens.

In summary, the legal issues are a mixture balancing who can record and what can be done with the data. Most countries allow vision systems in public places with few restrictions on what can be done by the government. Tort penalties limit what is done with the data outside of the government. Vision researchers developing systems should be aware of the implications, but today, unfortunately, the issues are more of social acceptance than legal limitations

1.4. Privacy enhancing vision research.

One way to enhance the privacy in video related work is designing applications that simply do not retain enough information to invade privacy. Most readers are probably familiar with TV shows were individual faces are blurred/masked beyond recognition. If the data is not there, the privacy concerns are greatly reduced.

In our previous surveillance work, [Boult-03] (including the commercially deployed versions of the system), the majority of the "tracking" is done with so few pixels on person that people there, including the dockworkers, did not see it as an invasion of privacy. The added benefit was that a system that has only a few pixels on person, cover larger areas of a facility with fewer cameras and thus decrease costs.

A similar effect happens with systems that automatically detect events and store or at least present limited or no video/image data. Event detection provides the added security, but limits its value for use in prosecution, but especially limits its value for defense against liability claims, since if not all video is recorded there is a potential for the court to blame the software for missing it.

Though no commercial system does it, a mixture of event detection for real-time monitoring combined with encrypted storage of the full video provide a simple but privacy enhancing technology. If the event detection has even only a moderate "potential event", it might provide a degraded video to the monitors (e.g. very highly compressed) while storing high quality encrypted video. This idea is a significant simplification of the ideas discussed in [Senior-et-al-05]. The key management would need to be similar to that discussed herein.

The paper [Newton-et-al-05] discusses an algorithm called k-same to "de-identify" facial images and hence make the face(s) inappropriate for being used with face recognition software. The paper states, without significant discussion or justification, that blacking out the face is unacceptable. For some multi-media work it may be, but I would argue that for most surveillance it is perfectly acceptable. That paper presents a new privacy-enabling algorithm, named k-Same, that limits the ability of face recognition software to reliably recognize faces while maintaining facial details in the images.

In the IBM research paper [Senior-et-al-05], the researchers have discussed their method of rendering face images unusable by face identification software. They suggest methods to obscure some facial features or alter the statistics of some facial features such that face recognition software cannot recognize the faces. The paper seeks to use computer vision to understand the video so they can leave "just enough" of the information contained in a video

stream to allow video-based tasks (including both surveillance and other "person aware" applications) to be accomplished, while hiding superfluous details, particularly identity, that can contain privacy-intrusive information. The technology was implemented as a privacy console that manages operator access to different versions of the video-derived data according to access control lists. It is unclear if the original copy of the original data is maintained.

Sony has a patent [Berger-00] in which they have proposed a method of detecting skin in images and replacing it with other colors, hence making it impossible to determine the race of the individual. Matsushita's patent [Wada-et-al-01] talks of a method to obscure a "privacy region" of an image as seen on camera.

In the remainder of the paper we present our approach that changes the way faces/people/targets appear in surveillance video to protect privacy. In addition to this, our method also makes it possible for authorized personnel, e.g. after obtaining a court order, to convert this back to the original image(s). The same concept can easily be adapted for "wiretaps" or voice recording, where the resulting voice data could not be decrypted until the court order was provided, setting up a two-stage process, getting permission to begin collection of potentially important surveillance data while building the case to get the actual warrant to view/listen to the surveillance data.

2. Cryptographic obscuration

All of the aforementioned methods serve the basic purpose of privacy enhancement by obscuring the face images obtained from surveillance video. However they have done so at the sake of security; they all lack a method to revert the transformed image back to its original form if there is sufficient reason to warrant it. The proposed Privacy through Invertible Cryptographic Obscuration (PICO) can be used to obscure faces, but the original face data can recovered. The authorized personnel are given the necessary encryption keys and parameters.

As a first example of cryptographically invertible obscuration, we use face detection software to detect the faces in an image or video. An application with some mixture of "skin" detection, text detection, motion detection and/or "voice" detection would equally apply. The basic concept is cryptographic extension of the obscuration idea that has been explored by many and might be viewed as a special transform in the sense of [Senior-et-al-05]. While Senior-et-al did address encryption, their approach requires a special "privacy console" and the encrypted data is out-of-band reprocessed data requiring special equipment. In contrast, in our case the potential private data is detected and the associated component of the media file is modified with the sensitive data encrypted in place. Unlike just blanket encryption, which would leave the data useless, the goal is for the un-encrypted data in the media to still be useful for general surveillance, e.g. in the door monitoring example shown here, the unencrypted data is sufficient to detect suspicious behaviors or people leaving with packages without knowing who is in the scene. In the parking lot example, a guard could probably tell if the person was trying to break into the vehicle without knowing who they are. Since the encrypted data appears as basically random numbers, the associated part of the file simply appears (or sounds) like noise. This idea, leaving most of the data in its original form and only protecting the "private" regions, improves privacy over just blank encryption because it allows the system to bed in settings that require actual observation. The partial encoding to provide privacy is shared with past work such as [Senior-et-al-05],

Figure 1: Two examples of Privacy through Invertible Cryptographic Obscuration (PICO), where the "privacy" model is based on face detection. The imagery is generally sufficient for security analysis without the need to decode the face region.

with the major advancement here focusing on the developing an approach can operate with minimal changes to the existing infrastructure.

Unlike earlier work, where the information was permanently destroyed, this PICO approach maintains the majority of security objectives because, if there was sufficient reason, the "private" data could be decrypted. E.g. if the car was subsequently stolen, a "warrant" might be issued to decrypt the face and hence identify the subject. By improving the "security" value, it can actually improve privacy for the general public by increasing the operations that are willing to deploy the privacy-enhanced technology. If the point of the cameras were for security or liability protection, the two biggest factors used in Return-on-investment justification for video systems, previous "obscuration" or de-identification systems would simply not be acceptable because they totally undermine the objective. Hence, while previous work locally appears to improve privacy, non-invertible obscuration approaches don't provide any advantage if they are not deployed. While PICO formally supports recovery (thereby locally violating privacy), it can improve the overall privacy of the general public by increasing the acceptability and deployments of the approach. The argument is, of course, a bit of a slippery slope, as it depends on securing the "keys" and trusting that they will be provided only when the situation warrants the privacy violation. In ongoing work, we are proposing to build a distributed PKI-based infrastructure that would allow an external organization, e.g. privacy oriented group of the US justice system, to control the public keys used.

In these examples, regions around detected faces are encrypted and the encryption key and other details are saved as a structured comment. The overall process is described in Figure 2. We propose to use Public Key encryption to protect the DES or AES *session key* used to encrypt the multi-media data that is encoded in place. Then the PK key, the encrypted session key and the region definitions are stored in a comment (and/or a XML descriptor file). Because the encryption is in place, the viewer is then just your standard image/video viewing software. The decryption details are not publicly known, hence maintaining the privacy of the individuals in the video. But if the need arises, then all the details of the original face can be provided to authorized personnel. When a "warrant" is provided, the system could be given the private key (if it is used only once), or it could

provide the encrypted session key(s), which will then be decrypted and returned. (The DES/AES session keys can be regularly changed say every 1 min, as they can just be "random" numbers.) This aspect of re-obtaining the original images from the transformed images is what makes our method unique from prior work. A digital signature of the image or face could be added to the comments ensuring traceability.

In order to demonstrate our idea, we used the OpenCV computer vision library's code to do real-time detection of the faces in an image or video, and then we encrypt the relevant regions. For performance reasons, we recommend using a Public Key algorithm only to protect the AES/DES key, and using AES/DES to encrypt the actual data. Since we are encrypting data in place, padding is not always a viable option. There are important details about "rounding" the regions used so that the data to be encrypted is an appropriate size. Both AES and 3DES encryption are block encryption algorithms that encrypt 64-byte blocks and handle larger data block-by-block. This requires either padding or alignment of data boundaries. The examples here used 3DES to encode the actual data, and thus we had to make changes to our detected regions so they were defined as a set of 64-byte blocks. This can be done in the spatial domain if using ppm, gif, Tiff or other lossless image formats. If using lossless compression, encryption is done in the spatial domain before the compression.

If the image is stored in a JPEG, png, Tiff or other lossy compression formats, particularly common in today's IP-based cameras, then encryption is done during (or after) compression and must work on the block format used in the encoding, e.g. the DCT blocks (usually 8x8 or 16x16) of the JPEG file format. One approach is for this to be done after the quantization of the DCT but before the lossless Huffman encoding. Because of the 64-byte size requirement, combined with the inherent high-entropy of the result of compression, the per block encoding will no longer be as efficient as before. While one could drop terms and then consider padding, it would not significantly reduce the size of the image (the padded numbers are, after encoding, approximately random). The alternative is to apply the encoding after the quantization and Huffman encoding that would then require padding.

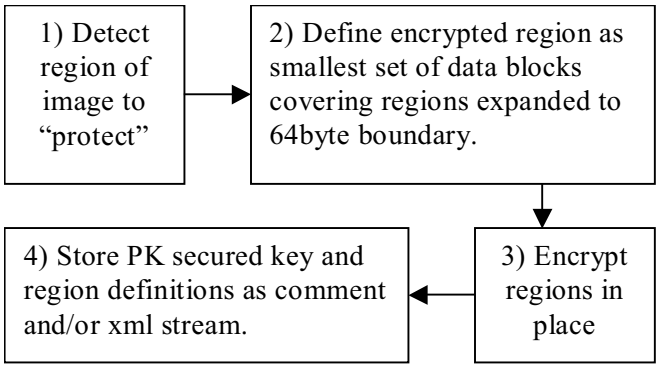

Figure 2: Steps in PICO processing

In both compressed and uncompressed cases, the encryption key and other required parameters (e.g. Bounding box of each regions) are saved as comments in the beginning of the image file, though an added XML-based descriptor file can also be used. For video, M-Jpeg is a trivial extension of the JPEG approach while for Mpeg video, similar approaches could be applied per I-frame. While it might seem sufficient to do this on the I frames, the work of [Li-et-al-96] shows that for real privacy all the frames needing encoding as a scene can produce a few frames of "visible" data even if each I frame is encoded. While we have not implemented an MPEG version, the performance of the full-frame software encrypted version only reduced frame rate by < 10%, and it is expected the "privacy" regions would be a much smaller fraction of the image but the detection and tracking will probably add about as much computation. In Mpeg, the "author" fields can be used for the encoding of the Public key and a data track or VBI data could be used to encode the region and encoded keys. To provide enhanced confidence in the resulting process, we strongly recommend a

cryptographic check sum of frames and of subsequences of the original be computed and also stored within the comments/XML.

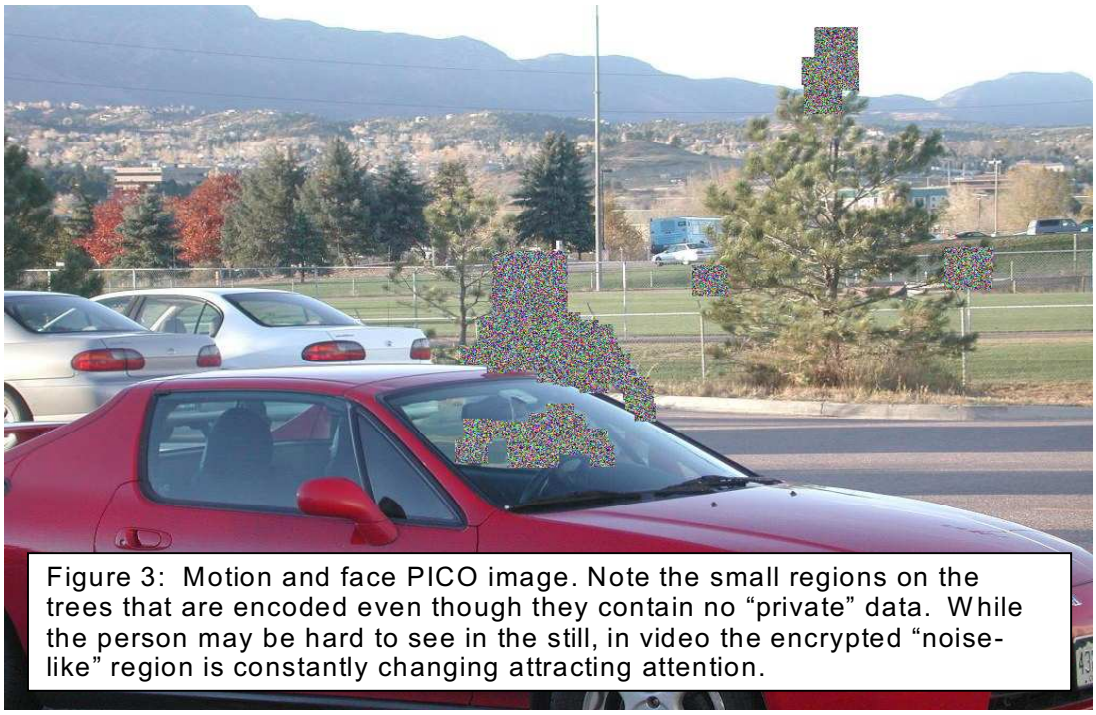

Figure 3: Motion and face PICO image. Note the small regions on the trees that are encoded even though they contain no "private" data. While the person may be hard to see in the still, in video the encrypted "noise-like" region is constantly changing attracting attention.

The resulting transformed images/video, with encrypted sub-regions, is suitable for display using standard display tools while successfully maintaining the privacy of individuals. Policy would determine when to release the decryption parameters and a specialized tool would then reconstruct the original data. It is important to note, however, that any non-lossless transcoding (decompression followed by recompression) would destroy the invertible nature of the data, but may be visibly indistinguishable. This might be used to improve overall transmission rates, where the server would store the original PICO data, but then provide a recompressed version for real-time viewing.

While the initial examples presented were simply protecting privacy by encrypting the face, this would not be sufficient for cameras placed in more sensitive areas such as dressing rooms (a major issue for retail theft) or in bathrooms, which are a issue for school security as well as airport and other critical facilities. For these settings we propose a more complex detection based on movement and/or skin-tone, e.g. using tracking technology such as [Boult-03] or [Senior-et-al-05]. Once the regions even slightly different from the background are detected, they can be protected by PICO's in place cryptographic obscuration. Some might question what value it would have if the whole person were obscured, but there is a tremendous amount of information security professionals can get from body language and behaviors. For physical security, e.g. airports, things like left-baggage detection, running, or loitering could still be computed during the processing.

Simple "motion" detection may include encryption of irrelevant details, e.g. the moving tree in figure 3. This would be an issue for some areas (e.g. watching coastal waters) but techniques for detecting only salient motion are already commercially available and could be incorporated. Because we are not using the tracking/detection for alarming, such systems should be very conservative in their detection settings to ensure privacy is maintained. Further privacy improvements would be to use detections and interpolate between them for any frames where the detection was "lost".

More interesting issues involve hierarchal invertible obscuration, e.g. where face, skin and motion cues are used to encrypt data. While the face is generally critical to identity, and skin often important for privacy, all moving parts of the scene may impart some information about identity and hence may violate privacy. A multi-layer approach better addresses the different amounts of privacy invasion. The approach encrypts the face data with one key then encrypting all of the data where there is skin (including the face) with another and the areas with motion (including the face and skin) with third key. The resulting image would still allow visual detection of people and important things like a "left bag" or loitering around a car, but would not show any detail about the person. And when an alarm goes off, a local decision, by a second person, could release the motion key to show the areas encrypted as motion, leaving skin and faces encrypted.. With the "motion" key the close and items being carried could be seen, but the skin and "face" data is recoverable and with additional keys. The multi-layer is similar to the ideas of multiple data streams being produced in [Senior-et-al-05], except that the primary data encryption is done in place and does not require a significant new infrastructure to use it

3.0 Conclusions and future work

While encryption is ubiquitous across security, the key in this application was finding a balance between encryption of the data and maintaining some of the data in some unencrypted form so that it can still be used without decryption and without invading privacy. If we cannot find that balance, security will continue to dominate the near-term decision-making process and privacy will be lost.

In PICO, privacy through invertible cryptographic obscuration, the key issues were finding minimal privacy preserving regions and then embedding the encryption into image, audio or video formats. Multi-level encryption allows multiple layers of privacy enhancement within the media. To support decryption, it is critical that the data is preserved, so encryption must be applied after any lossy transforms are applied to the data. Keeping the data in the original media file format, with the key and region info in comments or optional fields, allows preexisting software to render the data supporting the traditional use of the non-private parts of the data. The "in place" encrypted data shows the location and activity of the subjects and hence can still be used for analysis.

The demonstrated cryptographically invertible obscuration is, once described, a simple concept that allows a new balance between security and privacy. The real key to ensuring that "privacy" is immutable would be its introduction into the sensors themselves so that there is no concern that "software" would be circumventing the policy. Until that time, software-based solutions for use with web cameras and digital video recorders or digital audio recorders will have to suffice. In either case, key management and strong general policies, e.g. multi-stage "surveillance" warrants, must be addressed. The basic concept is quite simple; the more difficult issues for future work are developing the infrastructure and codecs needed to support its use in the many different video/image file formats.

Acknowledgements

Arun Viswanathan of UCCS did the implementation of PICO as an MS project. The work was supported through Dr. Boult's El Pomar Chair funds.

References:

[Adams-Sasse-01] Adams, A. & Sasse, M. A (2001) "Privacy in multimedia communications: protecting users not just data" in Proceedings of IMH HCI'01. pp. 49-64

[Agre-03] Philip E. Agre (September 2003). "Your face is not a bar code: Arguments against automatic face recognition in public places." An Abridged version appeared in *Whole Earth* 106, Winter 2001, and pages 74-77. Extended version available (October 25, 2004) from http://polaris.gseis.ucla.edu/pagre/bar-code.html

[Ansley-93] Leslie Ansley, "Safety in Schools: It Just Keeps Getting Worse," USA Weekend magazine, August 13–15, 1993, pp. 4–6.

[Bellotti-96] V. Bellotti, "What You Don't Know Can Hurt You: Privacy in Collaborative Computing", in A. Sasse, R. J. Cunningham & R. Winder (eds.), People and Computers XI (Proceedings of HCI'96), Springer-Verlag, pp.241–61, 1996.

[Berger-00] A. M. Berger (May 2000). "Privacy mode for acquisition cameras and camcorders", US Patent 6,067,399, Sony Corporation, May 23, 2000

[Davies-96] S. Davies, *Big Brother: Britain's web of surveillance and the new technological order*. Pan Books, London, p. 183. (1996)

[Dieteman-Time-71] Dietemann v. Time, Inc., 449 F.2d 245 (9th Cir. 1971) (California Law).

[Fay-98] S. Fay, "Tough on Crime, Tough on Civil Liberties: Some Negative Aspects of Britain's Wholesale Adoption of CCTV Surveillance During the 1990s", Int. Review of Law, Computers & Technology, Vol. 12, Num. 2, Pages: 315 – 347, July 1998.

[Flaherty-98] D.H. Flaherty, "Video surveillance by public bodies: A discussion, Investigation report, Investigation", P98-012, Information and Privacy Commissioner for British Columbia 1998 http://www.oipcbc.org/investigations/reports/invrpt12.html

[Honess-Chaman-92], T. Honess, E. and Charman; "Closed Circuit Television in public places", Crime Prevention Unit paper no. 35 London HMSO. 1992.

[Johnson-Hartcort-75] Johnson v. Hartcourt, Brace, Jovanovich, Inc., 43 Cal. App. 3d 880, 885 (2d Dist. 1975).

[Krause-01] Mike Krause. "The Expanding Surveillance State: Why Colorado should scrap the plan to map every driver's face and should ban facial recognition in public places," Independence Institute, Issue Paper, Number 8-2001, Oct. 2001. Retrieved 10/22/ 2004 from http://i2i.org/articles/8-2001.PDF

[Lee-et-al-97] A. Lee, A. Girgensohn, and K Schlueter, "NYNEX Portholes: Initial User Reactions and Redesign Implications, in S. C. Hayne & W. Prinz (eds.), Proc. of International ACM SIGGROUP Conf. on Supporting Group Work, Group'97, ACM Press, pp.385–94, 1997.

[Li-et-al-96] Y. Li, Z. Chen, S. Tan, and R. Campbell. "Security enhanced MPEG player". In Proc. of IEEE First Int. Workshop on Multimedia Software Development (MMSD'96), March 1996

[McCray-State-90] McCray v. State, 581 A.2d 45 (Ct. App. Md. 1990). (California Law).

[Newton-et-al-05] E. Newton, L. Sweeney, and B. Malin. Preserving Privacy by De-identifying Facial Images. IEEE Transactions on Knowledge and Data Engineering, IEEE TKDE, February 2005.

[Pawson-Tilley-94] R. Pawson and N Tilley, "What Works in Evaluation Research?" British Journal of Criminology, 34(3), 291-306, 1994.

[Priv-05] Privacy International, ND "CCTV FAQ", Website address: (accessed Nov 10 2005). http://www.privacy.org/pi/issues/cctv/cctv_faq.html

[RCFP-05] Reporters Committee for Freedom of the Press, http://www.rcfp.org/taping/ accessed on 11/10/05.

[Senior-et-al-05] Andrew Senior, Sharath Pankanti, Arun Hampapur, Lisa Brown, Ying-Li Tian, Ahmet Ekin, "Blinkering Surveillance: Enabling Video Privacy trough Computer Vision", IEEE Security & Privacy, volume 3, (no 3), pages 50-57 in 2005

[Smith-Hudson-95] I. Smith and S. Hudson, "Low Disturbance Audio for Awareness and Privacy in Media Space Applications", in R. Heller (ed.), Proceedings of Multimedia'95, ACM Press, pp.91–97, 1995.

[Thieme-03] Michael Thieme, International Biometrics Group, Presentation at *13th Annual Conference on Computers, Freedom & Privacy*, NYC, NY April 2003

[US389] Katz v. United States, 389 U.S. 347 (1967)

[US468] United States v. Karo, 468 U.S. 705-755 (1984).

[Wada-et-al-01]Jyoji Wada, Koji Kaiyama, Ken Ikoma, Haruo Kogane. "Monitor camera system and method of displaying picture from monitor camera thereof," European Patent, EP 1 081 955 A2, Matsushita Electric Industrial Co. Ltd., April 2001