Face Recognition Technology:

Security versus Privacy

KEVIN W. BOWYER

*They that can give up essential liberty to obtain
a little temporary safety deserve neither liberty nor safety.*

— Benjamin Franklin, *Historical Review of Pennsylvania,* 1759.

*Benjamin Franklin said that anyone who gives up
essential liberties to preserve freedom is a fool, but maybe
he didn't conceive of nuclear war and dirty bombs.*

— Songwriter Neil Young, discussing 9-11 memorial song "Let's Roll" [1]

Video surveillance and face recognition systems have become the subject of increased interest and controversy after the September 11 terrorist attacks on the United States. In favor of face recognition technology, there is the lure of a powerful tool to aid national security. On the negative side, there are fears of an Orwellian invasion of privacy. Given the ongoing nature of the controversy, and the fact that face recognition systems represent leading-edge and rapidly changing technology, face recognition technology is currently a major issue in the area of social impact of technology. This article analyzes the interplay of technical and social issues involved in the widespread application of video surveillance for person identification.

Right to Privacy

The United States Constitution declares a level of protection for the rights of individual citizens against oppression by their government that has made America unique. One right that has become a firmly entrenched American value, even though it is not explicitly enumerated in the Constitution, is the right to privacy. This is evidenced by the fact that phrases like "Orwellian nightmare" and "Big Brother" have passed into such common use that they have meaning even for those who have never read George Orwell's book [2].

The terrorist attacks of September 11, 2001, immediately focused attention on the field of *biometrics*, the study of ways to recognize people and verify identity. Face recognition is one type of biometric technology. Face recognition was highly controversial when used at the Super Bowl in early 2001. But the idea that video surveillance and face recognition systems might have alerted authorities to the terrorists that boarded airliners on September 11 has advanced public acceptance of the technology and motivated a number of airports to evaluate such systems. One idealized counter-terrorism scenario was outlined as follows in a *Business Week* article that appeared shortly after September 11 ([3, p. 39]):

"Khalid Al-Midhar came to the attention of federal law enforcement about a year ago. As the Saudi Arabian strolled into a meeting with some of Osama bin Laden's lieutenants at a hotel in Kuala Lumpur

 | 9

in December 1999, he was video-taped by a Malaysian surveillance team. The tape was turned over to U.S. intelligence officials and, after several months, Al-Midhar's name was put on the Immigration and Naturalization Service's "watch list" of potential terrorists. … The videotape of Al-Midhar also could have been helpful. Using biometric profiling, it would have been possible to make a precise digital map of his face. This data could have been hooked up to airport surveillance cameras. When the cameras captured Al-Midhar, an alarm would have sounded, allowing cops to take him into custody." (Fig. 1.)

This "dream scenario" outlined in the *Business Week* article has great public appeal. The average citizen is familiar with the mug-shot style images of the terrorists that have appeared in the media. The average citizen is also aware at some level that video surveillance is an everyday presence in activities such as banking, shopping, purchasing gas, driving through major intersections, and entering public and private buildings. There is great inherent appeal in the idea that future terrorist attacks could be prevented by high-tech video surveillance. A poll taken just after the 9-11 attacks asked for responses regarding "some increased powers of investigation that law enforcement agencies might use when dealing with people suspected of terrorist activity, which would also affect our civil liberties" [4]. This poll found that 86% of those responding were in favor of "use of facial recognition technology to scan for suspected terrorists at various locations and public events." Also, 63% were in favor of "expanded camera surveillance on streets and in public places." A follow-up poll taken six months later found that support for

use of face recognition technology fell by only a small amount, to 81% [5]. Support for other forms of increased surveillance generally dropped by larger amounts (Fig. 2).

If face recognition technology worked well enough, then it could become a valuable tool in fighting terrorism. However, it would also be a tool with great potential for enabling government invasion of privacy. Our society is currently in the midst of making important decisions about investment in, and deployment of, various forms of

> The terrorist attacks of September 11, 2001, immediately focused attention on the field of biometrics, the study of ways to recognize people and verify identity.

biometric technology. The debate is already engaged, with organizations such as the American Civil Liberties (ACLU) and Electronic Frontier Foundation (EFF) on one side, and government, law enforcement, and some corporations on the other side.

It is important that all citizens are able to think critically about the technical, social, and ethical issues that are involved in the decisions that confront our society. With this article, we hope to achieve the following: 1) outline a realistic understanding of the current state of the art in face recognition technology, 2) develop an understanding of fundamental technical tradeoffs inherent in such technology, 3) become familiar with some basic vocabulary used in discussing the performance of recognition systems, 4) be able to analyze the appropriateness of suggested analogies to the deployment of face recognition systems, 5) be able to assess the potential for misuse or abuse of such technology, and

6) identify issues to be dealt with in responsible deployment of such technology.

How Does A *"Face Recognition System"* Work?

The term *biometrics* refers to methods of identifying people based on some aspect of their biology. Fingerprints are a familiar example. However, fingerprints generally require the active participation of the person to be recognized. A technique such as face recognition could, at least in principle, be used to recognize people "passively," without their knowledge or cooperation. While there are numerous possible biometric approaches (ear image [25], gait analysis [26], ...), this paper focuses on face recognition technology of the type currently being considered for use in public spaces such as airports and tourist districts (Fig. 3).

The basic operation of a face recognition system is fairly simple. First, there is a camera that views some space — for example, the boarding area in an airport. Instead of a person continuously monitoring the video, the goal is to have a computer monitor the video and alert a human operator if an "interesting person" is in view. This presumes that there is a list of known "interesting persons." This is the *watch list*, or *gallery*. Enrolling a person in the gallery requires that you have a picture of the person.

The first step in the computer processing of the video is to detect when a person's face comes into view. The development of algorithms for finding a face in an image has been an active research area in its own right [6]. Once a face is detected, the face image is cropped from the video to be used as a *probe* into the *gallery* to check for possible matches. The face image is preprocessed to account for factors

such as image size and illumination, and to detect particular features of the face. The data from the probe image is then matched against the entries in the gallery.

There has been an enormous amount of research on data structures and algorithms for use in matching face images [7], [27]. In general, the matching algorithm will produce a similarity measure for the match of the probe image to each of the gallery images. A threshold can be set so that a match is reported to the operator only when the similarity measure between the probe and a gallery image exceeds the threshold. If more than one gallery image generates an above-threshold match, those images can be ranked according to the similarity measure. This threshold value can be raised or lowered to adjust the sensitivity of the system.

A reported match, or "alarm," typically alerts the system operator and displays the matched probe and gallery images. Incorrect matches may be dismissed by the operator based on differences in age, height, weight, gender, or other factors differing between the persons in the probe and gallery images. If the reported match is plausible, then the operator must decide what steps to take to confirm the identity of the person in the probe image. Thus the use of a face recognition system is generally conceived as a method of allocating scarce human resources to verifying the identity of the people that appear most similar to someone on the watch list.

There are several important elements of this scenario that should be clearly understood. First, the system can only recognize persons whose images have been enrolled in the gallery. If a terrorist is known by name and reputation, but no picture of the terrorist is available, then the face recognition system is useless. Another point is that the system must be able to acquire face images of reasonable quality to use as probes. If lighting conditions are poor, viewing angle is extreme, or

Khalid Al-Midhar
Terrorist on American Airlines Flight #77 that crashed into the Pentagon building in Washington, DC.

Mohamed Atta
Terrorist on American Airlines Flight #11 that crashed into the North Tower of the World Trade Center in New York.
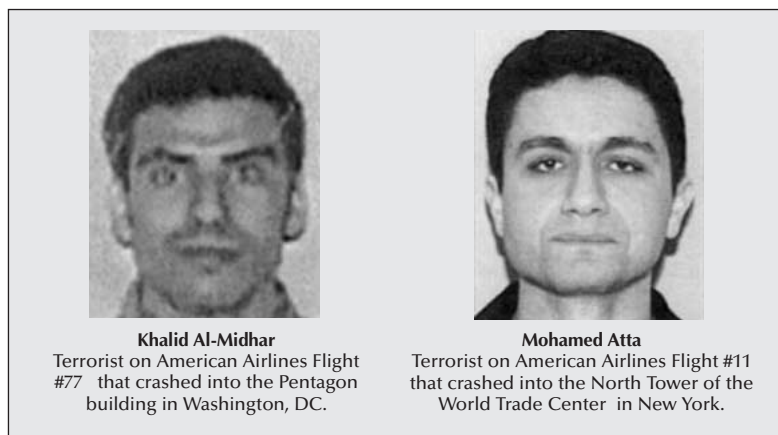
Fig. 1. Face images released by the FBI for two of the 9-11 terrorists (Images are available at www.fbi.gov/pressrel/penttbom/penttbomb.htm.)

people take measures to alter their appearance from that in the gallery image, then the performance of the face recognition system will generally suffer. A third point is that the system has a sensitivity threshold that must be set appropriately. Setting the threshold too low will result in too many *false positives* – an innocent person is subjected to some amount of scrutiny based on resemblance to a person on the watch list. Setting the threshold too high will result in too many *false negatives* – a terrorist is not recognized due to differences in appearance between the gallery and probe images of the terrorist. The diagram in Fig. 4 explains these possible outcomes of a recognition system's decision.

While the system threshold can be adjusted to be either more or less sensitive, it cannot be set to simultaneously give both fewer false positives and fewer false negatives. For any given current state of the technology, the system operator is inevitably faced with the choice of trading occurences of one type of error against occurences of the other type of error. Being able to achieve both fewer false positives and fewer false negatives requires new and better technology. Thus, if face recognition appears to be a potential solution in a given application, but the available tradeoff between true positives and false positives is not acceptable, then the answer is to develop improved technology.

One way of summarizing the potentional performance of this type of recognition system is a *Cumulative Match Characteristic* curve, or CMC curve. The CMC curve illustrates, in a certain way, the tradeoff of true positive versus false positive results. The Y axis is the *true positive rate*. In our current discussion, this is the fraction of the time that,

Fig. 2. Americans are familiar with widespread video surveillance. Surveillance cameras are prominent at banks, stores, gas stations, toll booths, road intersections, many public and private buildings, and other locations.
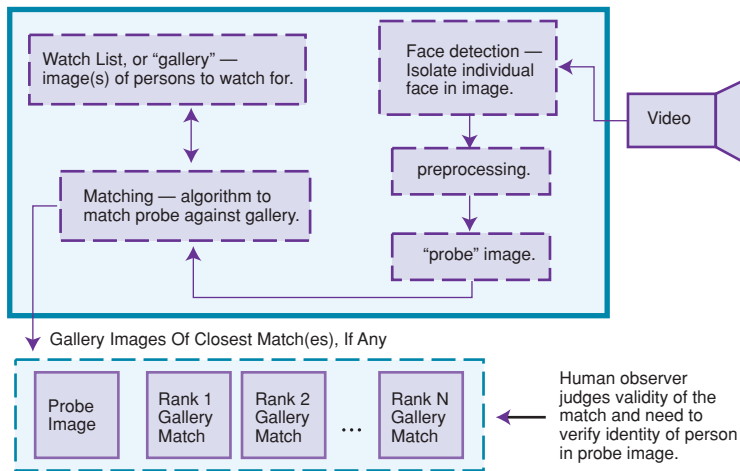
Fig. 3. Conceptual organization of a face recognition surveillance system. Important points are that the system can only recognize someone enrolled in the watch list, and that the candidate matches are screened by a human operator.

when someone on the watch list appears in the surveillance image, the system signals an alarm to the operator. The X axis is the cumulative rank, which we can think of as the maximum number of images that the system is allowed to report when giving an alarm for a given probe. If the system is allowed to report a larger number of possible matches, the true positive rate generally increases. However, the number of people that are incorrectly suggested for consideration, the false positive rate, also generally increases. And, of course, the workload on the human operator increases.

The CMC curve concept becomes important when evaluating and comparing the performance of face recognition systems. Improved technology would result in a better CMC curve, one that would run more toward the upper left corner of the plot as it is drawn in Fig. 5. In general, it makes no sense to compare the performance of systems based only on a true positive rate or false positive rate, since either number can be improved at the expense of making the other worse. To be meaningful, a comparison of two systems would need to take into account the CMC curve for each system, and the CMC curves would need to be obtained from the same experimental data.

Our discussion here primarily considers a recognition scenario as opposed to a verification scenario. That is, we assume that face recognition is being used to recognize the identity of an unknown person as matching someone on the watch list. In a verification scenario, face recognition is used to confirm that a given person is who they claim to be. The verification scenario is in a sense an easier problem, since the task is 'only' to verify a specific claimed identity, rather than to match against the entire gallery. The verification scenario is relevant to applications such as controlling access to a secure building. A person might approach the entry point and claim to be Jane Doe, someone who is on the list of persons allowed to enter. The face recognition system would then acquire the person's image, compare it to a known image of Jane Doe, and allow entry if the similarity is close enough. In this type of scenario, system performance would be described by a *Receiver Operating Characteristic* curve, or ROC curve. The ROC curve is effectively a special case of the CMC curve [30].

How Well Does Face Recognition Technology Work?

There are at least two senses in which one can discuss whether face recognition technology "works" – a system-technical sense, and an application-behavioral sense. The system-technical sense is concerned with the CMC curve and how many true and false positives the system produces in some test environment over some time period. The application-behavioral sense is concerned with how behavior is changed by the fact that the system is put into application. For example, say that a face recognition system is installed in an airport and it is known that it has only a 50% chance of correctly recognizing someone on the watch list when they appear in the airport. One might be tempted to think of this technical performance as a failure. However, a 50% chance of being identified might be sufficient to cause a terrorist to avoid the airport. If so, then one might be tempted to think of the system as a success in an application sense.

One often-repeated claim related to the effectiveness of video surveillance systems with face recognition

Fig. 4. Possible outcomes of face recognition for each person considered.

capability is that the introduction of such a system in London caused a major decrease in the crime rate [8]. News articles reported a 20% to 40% drop in crime, with the variation in the reported figures apparently arising from reporting statistics for different categories of crime, and/or adjusting for the increase in crime rate elsewhere in London over the same time period. Those who argue against the use of face recognition technology raise various objections to the reports of this experience.

One objection is that there is some inherent variability in reported crime statistics, and so perhaps the reported decrease was only a random occurence. Another objection is that the face recognition system was apparently not directly responsible for any arrests of criminals during this time. A related objection is that if the crime rate was reduced, perhaps the crimes were simply displaced to neighboring areas that were not using face recognition technology.

It seems plausible that this could be at least partially true. However, citizens in the area where the technology is deployed may still feel that it "works" for them, and this only begs the question of what would be the effect if the technology were more widely deployed.

One practical evaluation of face recognition technology was carried out at the Palm Beach International Airport. A face recognition system evaluated there captured about 10 000 face images per day, of about 5000 persons, during four weeks of testing. Of 958 images captured of volunteer subjects in the gallery, 455 were successfully matched, for a recogntion rate of approximately forty-seven percent. The report obtained from the Palm Beach County Department of Airports states that – "the false alarm rate was approximately 0.4% of total face captures, or about 2-3 false alarms per hour" [10], see Appendix A. A news article related to ACLU publicity about the test stated this
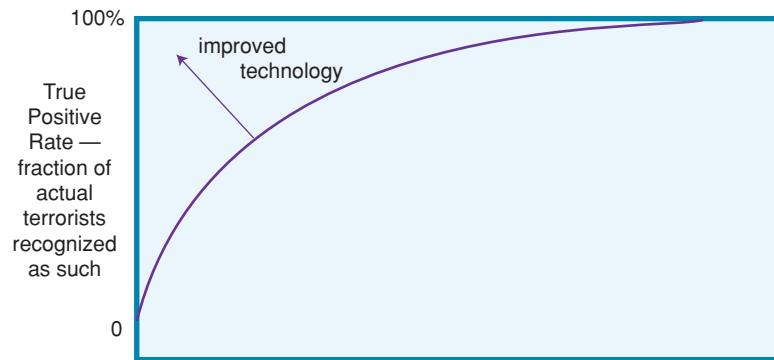

Fig. 5. Conceptual depiction of a cumulative match characteristic curve.

same information a little differently – "more than 1000 false alarms over the four weeks of testing" [9]. Problems were noted due to subjects wearing eyeglasses, lighting glare, and getting images that were not direct frontal views of the face.

An ACLU representative was quoted as saying that the results of the Palm Beach airport evaluation showed that "face recognition is a disaster" [9]. However, some observers may not agree with this assessment. A fifty percent chance of detection on a single face image may be sufficient to ward off some terrorists or criminals. Also, to the degree that the failures are due to psuedo-random effects such as lighting glare and viewing direction, acquiring images at multiple independent checkpoints could increase the chances of detecting a subject on a given trip through the airport. In other words, the chances of evading detection at both of two different checkpoint images might be as low as one in four, at three different checkpoints as low as one in eight, and so on. Also, researchers and companies in face recognition technology argue that the technology is still in very early stages of development. The systems are improving rapidly [22], and tests like the one at Palm Beach International Airport provide feedback to guide further research and development. A good sense of the current state of the art for commercial face recognition technology can be found in the 2002

Face Recognition Vendor Test [22].

In the end, those on either side of the debate can find interpretations of the performance data to support their argument. Face recognition technology does not yet work as well as desired in an technical sense [10], but the technology is improving [22], and reports such as those from London indicate that it does potentially produce some useful effect in application.

Another way of evaluating how well face recognition techology works is to look at how it has been adopted by private commercial users. One supplier of face recognition technology claims to have installed fifty systems in casinos, where the technology "is used by surveillance operators to identify cheaters and other casino undesirables, as well as casino VIPs" [29]. Casino owners are presumably evaluating the technology on a cost-effectiveness basis, and judging that it will allow them to increase their overall profit.

However, the context for deciding to deploy face recognition technology in such an application may involve important considerations that are not present in public anti-terrorism applications. For the purposes of a casino, the ability to automatically recognize "VIP's" – people who can be motivated to gamble big and lose big and may enjoy being recogonized – may be at least as important as the ability to recognize "undesirables." The

counter-terrorism scenario contains nothing analogous to the casino VIP. Thus the fact that face recognition technology is deployed in one application area does not automatically mean that it is appropriate for other application areas.

Viewpoints and Analogies

A number of media people, public officials, and special interest groups have gone on record with various opinions and arguments regarding the deployment of face recognition technology. A critical-thinking analysis of some examples can help to sharpen one's ability to discern the essence of the issues.

An article in *Time* magazine just after face recognition was used at Super Bowl XXXV provides an example of unrestrained belief in the technology [13] – "The beauty of the system is that it is disguise-proof. You can grow a beard and put on sunglasses, and FaceTrac will still pick you out of a crowd."

As any careful observer might suspect, the idea that any face recognition technology is "disguise-proof" is wildly optimistic! While some technologies may handle particular types of disguises better than others, no known current face recognition technology is "disguise-proof" in any general sense of the term. The results of the Face Recognition Vendor Test clearly show that there is much research and development still to be done in the area of face recognition [22]. It is also worth noting that researchers are beginning to look at ways of fooling face recognition systems [28].

Again in the context of face recognition used at Super Bowl XXXV, congressman Ed Markey (D., Mass.) was quoted as saying [13] – "It's chilling, the notion that 100 000 people were subject to video surveillance and had their identities checked by the government." This quote illustrates a subtle misunderstanding of the technology. If we hear that police "checked someone's identity," then we gener-

ally understand that police required the person to prove their identity, perhaps by examining the person's driver's license or taking their fingerprints. In using face recognition technology, the only people who can be recognized are those in the watch list, and the only people who might have to have their identities confirmed are those who look like someone on the watch list.

What really happened at the Super Bowl was that the system suggested that 99 980 of the attendees had no close resemblance to anyone on the watch list and that about 20 of the people were a possible match to someone on the list. This is not exactly the same as everyone having their identities checked. The system did not in any sense create a list of the identities of the 99 980 people.

A rather more obvious example of confusion about the technology occurred in a New York Times article following the Super Bowl. The article reported that [14]:

"A woman in Texas who saw the image claimed the man in the picture was wanted for crimes. She called the Tampa police, who questioned the man, a construction worker. It was the wrong person ... The system … is not 100 percent accurate."

The problem here is that the reported incident is a case of mistaken identity by a human, and not by the face recognition system! It is the viewer in Texas who suggested an incorrect identity for the person in the picture.

Another example of misunderstanding the technology occurs in this quote from a New York Times article [31]: "a computer glitch could match the face of an innocent person with the digital image of a criminal." The problem here is that the occurrence of such a false positive is not a "glitch," but instead it is an inescapable element of using the technology. To say that it is a "glitch" implies that it is something

that can be eliminated in a debugged version of the system.

The possibility of false positives cannot ever be completely engineered out of the system. In fact, as already mentioned, for any given state of the art of the technology, in order to be more certain of catching the terrorists, a higher rate of false positives has to be accepted. The public should never be encouraged to think that perfect performance – all terrorists caught and no innocent person inconvenienced – is something actually achievable.

A variety of analogies have been used in attempts to characterize the application of face recognition technology. The "police lineup" analogy was used by the ACLU in their protest over the use of face recognition at the Super Bowl [11]: "We do not believe that the public understands or accepts that they will be subjected to a computerized police lineup as a condition of admission." The term "police lineup" brings to mind a scenario of being called to the police station to be viewed along with a handful of other people for possible identification by a witness to a crime. There is a suggestion of a certain level of inconvenience involved, and also a bit of an implied accusation. It is doubtful that people who have walked through an area under video surveillance feel that they have participated in a police lineup.

A different analogy for the application of face recognition systems appears in the following quote [15]: "Police are enthusiastic about the system, saying it is no different from an officer standing on a street corner with a criminal's photograph in hand and checking out a passing crowd. They say it can reduce bias based on race or dress." The idea that a policeman might carry along a photo of a wanted person and scan crowds to try to find the person seems well accepted. And this is essentially the task that is being automated by the use of face recognition systems. Thus this seems like the most appro-

priate of the analogies mentioned here. However, there is the important difference that the use of face recognition technology makes it feasible to do this type of scanning on a scale that is impossible in practical terms using human observers.

Does Face Recognition Surveillance in Public Spaces Invade Privacy?

The most fundamental argument against government use of face recognition technology in public spaces is that it is a violation of the constitutional right to privacy. This core objection was advanced by the ACLU in the context of the use of face recognition at the Super Bowl [11]: "... this activity raises concerns about the Fourth Amendment right of all citizens to be free of unreasonable searches and seizures." However, essentially all legal commentators agree that use of face recognition systems in public spaces cannot be considered a "search" for constitutional purposes. Woodard's analysis of the issue seems careful and representative [12]:

"Under current law, however, the type of facial recognition used at the Super Bowl would almost certainly be constitutional. The Supreme Court has explained that government action constitutes a search when it invades a person's reasonable expectation or privacy. But the court has also found that a person does not have a reasonable expectation of privacy with regard to physical characteristics that are constantly exposed to the public, such as one's facial features, voice, and handwriting. So although the Fourth Amendment requires that a search conducted by government actors be 'reasonable,' which generally means that there must be some degree of suspicion that the person to be searched is engaged in wrongdoing, the scan of spectators' facial characteristics at the Super Bowl did not constitute a search."

However, this interpretation of the right to privacy was formulated before it was technically conceivable that a system could automatically match the face of every person entering a public space against a gallery of images of people wanted by authorities. Thus, some observers may argue that the scale of operations made possible by computerized face recognition technology should result in a change to the Supreme Court's traditional interpretation of the right to privacy.

Another concern is simply that citizens should be notified when they enter a public space where video surveillance is being used. The idea is apparently that people could then make an informed choice of whether or not to subject themselves to surveillance. Of course, if all airports install face recognition systems then there may be little practical "choice" left for some travelers. However, given the level of screening already in place for passengers boarding an airplane, posing for a picture for a face recognition system would seem to be a rather minimal added inconvenience.

There is also a policy question regarding the decision to add a person's image to the watch list. It seems clear that if there is an arrest warrant for a person, then the person's image could be put in the watch list. But what if a person is only wanted for questioning? Or what if authorities simply want to keep track of where a particular person goes and who they meet with? Should approval from a judge be required in order to place a person's image in the watch list, perhaps in a process similar to the requirement for a telephone wiretap? Should a permanent record be kept of when a person's image was entered into the watch list and at whose request? Clearly some level of official procedure is required in order to guard against abuse. If no approval is required for an individual law enforcement officer to enter images into the watch list, it is easy to imagine scenarios for abuse. A suspicious husband or wife might be tempted to want to keep tabs on his or her spouse. A person running for office might want to know whom an opponent is visiting. Such abuses of power are of course not new with face recognition technology, and so existing rules for other types of surveillance technology may suggest analogies. For example, there are well-developed rules for authorizing, using, and monitoring wiretaps [21], and this may provide a good starting point for developing rules for face recognition technology.

Another potential concern is the wholesale archiving of images for possible later use. One company advertises that the only probe images that are stored by the system are those that register a potential match to someone in the gallery. However, there is no real technical limitation to archiving all face images captured by the system, only the cost consideration for the volume of storage required. The archive of images could then be "mined" in the future to find out if a certain person had been at the checkpoint in the past.

Some observers have noted that "function creep" inevitably occurs after the installation of new technology. For instance, individuals can have an identifying card attached to their car that allows

> # There is great inherent appeal in the idea that future terrorist attacks could be prevented by high-tech video surveillance.

them to pay automatically at toll booths. The original purpose is simply to provide greater convenience to individuals and to generally speed the traffic at the toll booth. But if the record of toll payments made by an individual would later prove useful in another context, for example in a court case to prove or disprove who had been where at what time, then courts would typically order that the information be provided. If all images acquired by a face recognition system are archived, then the temptation toward function creep becomes strong. Again, this is not a new problem, but one that has been noted with earlier technologies and can be expected to arise in similar ways with new technologies.

For all of these reasons and others, the ACLU is one of the most vocal critics of face recognition technology. One press release emphasizes the threat to privacy, as illustrated in quotes such as: "There is an alarming potential for misuse of all of these systems" and "We are extremely troubled by this unprecedented expansion in high-tech surveillance in the United States" [23]. Another press release emphasizes the limitations of the technology, as illustrated in: "it is abundantly clear that the security benefits of such an approach would be minimal to nonexistent, for a simple reason: the technology doesn't work" and "Anyone who claims that facial recognition technology is an effective law enforcement tool is probably working for one of the companies trying to sell it to the government" [23].

Of course, there is a competing logic to these viewpoints. If the technology does not work, it can't be a real threat to privacy. And if it is a real threat to privacy, then it would seem to have at least some potential in law enforcement. It seems that it cannot be simultaneously both a technology that does not work and one that presents a serious threat to privacy.

Conceivably, it could be a tech-nology that does not yet work well enough to be of use in fighting terrorism, but that eventually will develop sufficiently to become useful in fighting terrorism and so then also then also a potential threat to privacy. But this only returns to the question of whether and how to trade off privacy versus security, and does not provide any easy answer to the question.

An issue that generally increases the intensity of all other concerns is that of networked recognition systems and databases. As initially envisioned, a face recognition system provides a way of detecting when a particular person enters a particular surveillance area. Every major airport and every major public building might have its own independent surveillance system. But what if all of these systems were networked together? For example, a person might be on the watch list for crossing the Canadian border into the United States. The person could be detected on crossing the border, and the person's image forwarded to the watch list for the expressway toll booths. The person could then be detected again as they take the exit to the airport, and their image forwarded to the watch list of the airport system, and so on. Each of these events can also be reported back to a central location. The result is a continuous monitoring of the person's whereabouts. While we might think this is good for a non-citizen who is on a list of suspected terrorists, we would see it as a invasion of privacy if done for a person who has simply been a vocal critic of current goverment policies. And if done on a large scale without due cause, it begins to approach the level of Orwell's "Big Brother."

Of course, the potential for monitoring a person's location and actions through a combination of existing technologies such as credit card records, cell phones, and global positioning systems in automobiles may already exceed the poten-tial invasion of privacy that could come from mis-use of face recognition technology. It may be that, to some extent, face recognition technology is evaluated more emotionally simply because it uses images of a person's face.

Policies For Industry and Government

Presently, there are essentially no legal guidelines for where and how face recognition technology can be used in public spaces. Local governments are making decisions on the deployment of face recognition technology on a case-by-case basis, sometimes without a clear understanding of what they are approving [16]. There are important questions to be answered in terms of the policies that are going to control how government makes use of this technology. How should the installation of a face recognition system for a particular public space be proposed, approved, and monitored? What technical performance requirements should there be for such systems? What sort of notice needs to be given about the use of the technology to persons entering the public space? What rules control whose image can be put in the gallery and how long it can remain there? What images acquired by the system can be kept and for how long? At least one company involved in face recognition technology has made initial proposals with regard to some of these issues. The Visionics web pages list a set of proposed "Privacy Protection Principles" (see Appendix B). However, there are many details still to be worked out before any comprehensive answers emerge.

Guidance from Codes of Ethics for the Computing Professions

Each of the codes of ethics relevant to the computing-oriented professions [21] indicates a general con-

cern for privacy. However, none of the codes can be looked to for direct and consistent guidance on the question of whether face recognition systems should be deployed in public spaces. For example, the standards of conduct of the Association of Information Technology Professionals (AITP) contains several statements potentially relevant to the debate over deployment of face recognition surveillance systems:

"In recognition of my obligation to society I shall ...

Protect the privacy and confidentiality of all information entrusted to me. ...

Use my skill and knowledge to inform the public in all areas of my expertise. ...

To the best of my ability, insure that the products of my work are used in a socially responsible way."

The statement about protecting the privacy of information entrusted to the computing professional can be understood to suggest that the developers of a surveillance system have a responsibility to build in protections for the security and privacy of the information in the system. The statement about informing the public can be understood to suggest that computing professionals should be active in the public debate concerning surveillance systems and privacy. Computing professionals should particularly make inputs to the debate based on their special expertise, striving to give a clear understanding of the potential and the limitations of the technology. The statement about insuring that technology is used in a socially responsible way can be understood to suggest that the system should be constructed so as to minimize the possibility of misuse.

The code of ethics for the Association for Computing Machinery (ACM) has, in its list of "general moral imperatives," some similar general statements about obligation to society and respect for privacy:

"As an ACM member I will...:

1.1 – Contribute to society and human well-being.

1.7 – Respect the privacy of others."

Also, as part of its "organizational leadership imperatives," the ACM code of ethics states:

"As an ACM member and an organizational leader, I will...

3.4 – Ensure that users and those who will be affected by a computing system have their needs clearly articulated during the assessment and design of requirements. Later the system must be validated to meet requirements."

These statements in the ACM code can be understood to support concerns similar to those in the AITP standards of conduct. The point that is a bit more explicit here is that the computing professional should, during the design and validation of the system, consider the needs of those who will be affected by the system.

The Software Engineering code of ethics contains the following statements:

"Software engineers shall act consistently with the public interest. In particular, software engineers shall, as appropriate: …

1.3 Approve software only if they have a well-founded belief that it is safe, meets specifications, passes appropriate tests, and does not diminish quality of life, diminish privacy or harm the environment. The ultimate effect of the work should be to the public good.

1.4 Disclose to appropriate persons or authorities any actual or potential danger to the user, the public, or the environment, that they reasonably believe to be associated with software or related documents.

1.5 Cooperate in efforts to address matters of grave public concern caused by software, its installation, maintenance, support or documentation."

These statements in the software engineering code of ethics reflect concerns similar to those in the AITP and ACM codes. Item 1.3 suggests that software engineers involved in the development of surveillance systems should approve the system only if they believe that it does not diminish privacy and that its ultimate effect is to the public good. But one likely scenario in this case is that a software engineer might believe both that the system does diminish privacy and yet also does contribute to the public good by affecting some increase in security. The code of ethics does not say how to balance these competing goals.

In the end, the codes of ethics indicate that 1) the computing professional should have a concern for privacy, 2) the computing professional should give informed input to public debate, and 3) the decision should ultimately be that which is judged to best contribute to the well-being of society. The codes provide a framework to guide decision-making, but do not directly suggest whether or not the proposed application is ethical. We are left with the value judgment of whether the cost in diminished privacy is offset by an increase in security.

Several of the codes of ethics also contain elements that warn against unrealistic statements about the potential of a technology. For example, the Software Engineering Code of Ethics contains the admonition, "Be fair and avoid deception in all statements, particularly public ones, concerning software." Similarly the International Biometric Industry Association advocates this principle [32], "Accountability in Marketing: Because truth is the key to industry credibility, members attest that their stated claims are accurate and can be independently verified by a competent authority." The CEO of one company in the biometrics industry gave the following quote that some would consider over-optimistic and others might consider as appropriately qualified:

"We know that three out of the nineteen were already known on watch lists and knowing how many checkpoints these people had to go through, we had a high probability to alert, intercept, these individuals maybe August 21st or 23rd when they crossed the Canadian border and we would have perhaps foiled the whole plot" [20].

Despite the presence of qualifiers such as "high probability" and "perhaps," it is this author's opinion that much of the general public would form unrealistic expectations of the technology based on hearing this type of quote.

Trading Liberty for Security

The full depth and meaning of Benjamin Franklin's warning about trading liberty for security is not always appreciated. He posed the tradeoff as one of giving up "*essential* liberty" in order to obtain a "*little temporary* safety." Thus we can expect that much of the disagreement in this area comes down to whether a particular liberty is judged as essential or inessential, and whether the increase in safety is judged to be little or much, and temporary or permanent. Perhaps this can be appreciated by considering a sort of complement to Franklin's statement:

"They that insist on keeping an inessential liberty at the cost of a large and permanent threat to safety ...."

The introduction of video surveillance and face recognition systems into public spaces presents our society with important decisions. We must decide: 1) when or whether a sophisticated high-tech application works well enough to be worth deploying, 2) which elements of privacy are essential and which are inessential, and 3) what level of increased safety can come through the introduction of this technology. Especially for those in computing-related professions, who may be tempted to focus only on the technology, it is important that the potential of this technology both in terms of increased security and in terms of potential abuse be properly understood.

Additional Information

For additional information, there are numerous information resources on the web related to this topic. The Biometrics Consortium provides access to the industry viewpoint (www.biometrics.org). The Electronic Frontier Foundation (www.eff.org) and Electronic Information Privacy center (EPIC) (www.epics.org/privacy/facerecognition/) provide information from the privacy-advocate viewpoint. Also, an article by Philip Agre attempts to counter essentially all known arguments in favor of face recognition technology (dlis.gseis. ucla.edu/~pagre/bar-code.html). Much of the academic research concerned with advancing face recognition technology appears in journals such as the *IEEE Transactions on Pattern Analysis and Machine Intelligence* and conferences such as *Face and Gesture Recognition* and the *International Conference on Audio- and Video-based Biometric Person Authentication*.

Teachers who wish to cover this topic, or other topics related to social impact of computing, in their classes may find it useful to browse the site www.cse.nd.edu/~kwb/nsf-ufe. This site contains a variety of educational materials related to teaching ethics and computing. It also contains a powerpoint file that can be used to lead a class discussion along the lines of the material presented in this paper.

Acknowledgments

The author would like to thank Cherie Ann Sherman of Ramapo College of New Jersey, Sudeep Sarkar of the University of South Florida, and Jonathon Phillips of the National Institute of Standards for reading and commenting on earlier drafts of this paper. The author would also like to thank Patrick Flynn of the University of Notre Dame for numerous helpful discussions in this area.

Author Information

The author is Schubmehl-Prein Department Chair, Department of Computer Science & Engineering, University of Notre Dame, South Bend, IN. Email: kwb@cse.nd.edu.

References

[1] E. Gunderson, "Neil Young is passionate about 9/11 anthem," *USA Today*, Apr. 11, 2002.
[2] G. Orwell, 1984. www-online-literature.com z/orwell/1984/
[3] "Special report: privacy in an age of terror," *Business Week*, Nov. 5, 2001, pp. 39-46.
[4] T. Humphrey, "Overwhelming public support for increasing surveillance powers …," Oct. 3, 2001, www.harrisinteractive.com/harris_poll/index.asp?PID=260
[5] T. Humphrey, "Overwhelming public support for increasing surveillance powers …," Apr. 3, 2001, www.harrisinteractive.com/harris_poll/index.asp?PID=293
[6] M.-H. Yang, D. Kriegman, and N. Ahuja, "Detecting faces in images: A survey," *IEEE Trans. Pattern Analysis & Machine Intelligence*, vol. 24, no. 1, pp. 34-58, Jan. 2002.
[7] R. Chellappa, C.L. Wilson, and S. Sirohey, "Human and machine recognition of faces – A survey," *Proc. IEEE*, vol. 83, no. 5, pp 705-740, May 1995.
[8] B. Lack, "Development of facial recognition technologies in CCTV systems," *The Source: Public Management J.*, Oct. 25, 1999, www.sourceuk.net/articles/a00624.html
[9] S. Bonisteel, "Face recognition technology fails again, ACLU claims," May 16, 2002, washingtonpost.com, www.newsbytes.com/news/02/176621.html
[10] www.aclu.org/issues/privacy/FaceRec_data.pdf.
[11] D. McCullagh, "Call it Super Bowl Face Scan I," Feb. 2, 2001, www.wired.com/news/print/0,1294,41571,00.html
[12] J.D. Woodard, "Super Bowl surveillance: Facing up to biometrics," /www.rand.org/publications/IP/IP209/
[13] L. Grossman, "Welcome to the snooper

bowl," *Time*, Feb 12, 2001.

[14] "Electronic surveillance: From 'Big Brother' fears to safety tool," *New York Times*, Dec. 6, 2001.

[15] J. Cienski, "Police cameras denounced as threat to privacy," July 12, 2001, *National Post Online*, ww.nationalpost.com

[16] D. Nguyen, "Council: So that's what we okayed," *St Pete Times Online*, July 6, 2001. www.sptimes.com/News/070601/TampaBay/Council__So_that_s_wh.shtml

[17] A.M. Burton, S. Wilson, M. Cowan, and V. Bruce, "Face recognition in poor quality video: Evidence from security surveillance," *Psychological Sci.*, vol. 10, no. 3, pp. 243-248, May 1999.

[18] A. Jain, R. Bolle, and S. Pankanti, *Biometrics: Personal Identification In A Networked Society*. Kluwer, 1999.

[19] R. Beveridge, "Evaluation of face recognition algorithms." www.colostate.edu/evalfacerec/

[20] "Morning Edition," National Public Radio, Feb. 25, 2002. Transcript from Burrelle's Information Services, Box 7, Livingston, NJ 07039.

[21] K. Bowyer, *Ethics and Computing: Living Responsibly In A Computerized World*, rev. ed. New York, NY: IEEE/Wiley, 2001.

[22] P.J. Phillips, P. Grother, R. Michaels, D.M. Blackburn, E. Tabassi, and J. Bone, "Face Recognition Vendor Test 2002: Evaluation Report," www.frvt.org

[23] "Proliferation of surveillance devices threatens privacy," archive.aclu.org/news/2001/n071101a.html.

[24] "ACLU opposes use of face recognition software in airports due to ineffectiveness and privacy concerns," archive.aclu.org/issues/privacy/FaceRec_Feature.html.

[25] K. Chang, K. Bowyer, S. Sarkar, and B. Victor, "Comparison and combination of ear and face images for appearance-based biometrics," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 25, no.9, pp.1160-1165, Sept. 2003.

[26] P.J. Phillips, S. Sarkar, I. Robledo, P. Grother, and K.W. Bowyer, "Baseline results for the challenge problem of Human ID using gait analysis," in *Proc. Face and Gesture Recognition 2002*, Washington, DC, pp.

137-142.

[27] W. Zhao, R. Chellappa, A. Rosenfeld, P.J. Phillips, "Face Recognition: A literature survey," *ACM Computing Surveys*, vol. 35, no. 4, pp. 399–458, Dec. 2003.

[28] J. Alexander and J. Smith, "Engineering privacy in public: Confounding face recognition," presented at Workshop on Privacy Enhancing Technologies (PET 2003), Dresden, Germany, Mar. 2003, www.cis.upenn.edu/~jalex/papers/pet2003.pdf

[29] "Viisage technology and biometrica systems achieve 50th face recognition installation at Mirage Resort, Las Vegas," Viisage Corporation press release, Mar. 29, 2000, www.viisage.com/March29_2000.htm

[30] P. Grother, R. Michaels, and P.J. Phillips, "Face Recognition Vendor Test 2002 performance metrics," presented at Int. Conf. on Audio- and Video-based Biometric Person Authentication, Surrey, U.K., June 2003.

[31] "Super bowl snooping," *New York Times*, Feb 4, 2001.

[32] www.ibia.org/principl.htm

Appendix A

Results of face recognition test at the Palm Beach International Airport.

Facial Recognition System Test (Phase I) Summary.

Palm Beach County, Department of Airports conducted a test of the Visionics "Argus" facial recognition sys-tem. The purpose of the test was to ascertain the effectiveness of this technology in an airport checkpoint environment.

Utilizing a test group of 15 airport employees and a database of 250 photographs, the system features that were tested included:

- Face capture rate.
- False alarm rate.
- The ability to successfully identify test group against database photographs.

The data collected and compared to the manufacturer's advertised specifications revealed the following:

- Input photographs populating the database need to be of a good quality to avoid false alarms and in-sure successful matches.
- Motion of test subject head has a significant effect on the system ability to both capture and alarm on test subject.
- There was a substantial loss in matching if test

subject had a pose 15 to 30 degrees (up / down, right / left) off of input camera focal point.

- Eyeglasses were problematic, glare from ambient light and tinted lenses diminished the system's effectiveness.
- System required approximately 250 lux of directional lighting to successfully capture faces and alarm on test subjects.
- Face capture rate was approximately 10,000 face captures per day. The actual traffic through the security checkpoint is approximately 5,000 ticketed passengers and airport employees. There were multiple face captures for each event.
- The false alarm rate was approximately 0.4% of total face captures. Or about 2-3 false alarms per hour.
- Of the 958 total combined attempts there were 455 successful matches, (47% successful rate).

Test conducted at Palm Beach International Airport Concourse C security checkpoint, March 11-th through April 15-th, 2002.

This document was obtained from the Palm Beach Count Department of Airports by the American Civil Liberties Union of Florida.

Appendix B

"Privacy protection principles" from the Visionics Corporation web page www.faceit.com/newsroom/biometrics/privacy.html (February 2, 2002).

Privacy Protection Principles

In recent weeks, much media attention had been given to a new tool for combating crime, installed for the first time in the United States in Ybor City, the entertainment district of Tampa, Florida. The system uses FaceIt® face recognition technology from Visionics to alert police to the presence of known criminals in crowds in a public place.

Visionics has pioneered facial recognition and continues to be the worldwide leader in advancing the state-of-the-art. The company is most qualified to understand the power of the technology and its ability to impact society. The massive reduction in crime in the Newham borough of London (up to 40%) and elsewhere is a testament to the benefits that society can reap from the use of facial recognition. At the same time, the company is cognizant that powerful technologies require responsible use.

In keeping with its belief that its leadership must extend beyond the technology itself, Visionics has takend an active role in articulating the most appropriate and ethical ways for society to benefit from its technology while ensuring there are no opportunities for abuse.

Thus far Visionics has formulated responsible use guidelines, secured their acceptance by those adopting its technology, been vigilant in ensuring compliance and, where possible the company has built technical measures to maintain control over the installations.

So far these measures have been effective. But without systematic oversight and enforcement, they may not be enough in the long run. Consequently, Visionics is calling for federal legislation in the United States to help transform these responsible use principles into responsible public policy.

Among the issues that must be addressed when examining the use of facial recognition as a tool for combating crime in public places.

- Public Knowledge: Guidelines that establish the proper ommunication mechanisms to the public (such as street signage and media alerts) and the circumstances under which exceptions could be made (e.g., matters of national security at airports and borders).

- Database: Face recognition requires a watch-list database. Specific guidelines must be established for database protocols such as: need, justification for inclusion and removal, valid duration of information, dissemination, review, disclosure and sharing. This should include very explicit guidelines that spell out who can be in the "watch-list" database (e.g., fugitive with warrants, criminals under court supervision, etc.). Technical measures should be in place to ensure control over database size and its integrity.

- No Match – No Memory: Guidelines to ensure that no audit trail is kept of faces that do not match a known criminal or person under active police inestigation. Non-matches should be purged instantly.

- Authorized Operation and Access: Technical and physical safeguards such as logon, encryption, control logs, and security to ensure that only authorized, trained individuals have access to the system and to the database.

- Enforcement and Penalty: Oversight procedures and penalties for violaton of the above principles should be formulated.

In the months to come, the company will work closely with federal legislators, privacy interest and industry groups to share its knowledge, experience and privacy protection principles pertaining to all applications of face recognition technology. The company will continue to promote an improved public understanding of the creation, use and control of facial recognition systems and will support public policy that ensures that deployment of facial recognition are conducted in a way that upholds all privacy rights.

© 2001 Visionics Corp.