Picturing algorithmic surveillance: the politics of facial recognition systems

Introna, Lucas; Wood, David

Veröffentlichungsversion / Published Version
Zeitschriftenartikel / journal article

Surveillance & Society
CCTV Special (eds. Norris, McCahill and Wood)
2(2/3): 177-198
http://www.surveillance-and-society.org/cctv.htm

Picturing Algorithmic Surveillance: The Politics of Facial Recognition Systems

Lucas D. Introna[1] and David Wood[2]

Abstract

This paper opens up for scrutiny the politics of algorithmic surveillance through an examination of Facial Recognition Systems (FRS's) in video surveillance, showing that seemingly mundane design decisions may have important political consequences that ought to be subject to scrutiny. It first focuses on the politics of technology and algorithmic surveillance systems in particular: considering the broad politics of technology; the nature of algorithmic surveillance and biometrics, claiming that software algorithms are a particularly important domain of techno-politics; and finally considering both the growth of algorithmic biometric surveillance and the potential problems with such systems. Secondly, it gives an account of FRS's, the algorithms upon which they are based, and the biases embedded therein. In the third part, the ways in which these biases may manifest itself in real world implementation of FRS's are outlined. Finally, some policy suggestions for the future development of FRS's are made; it is noted that the most common critiques of such systems are based on notions of privacy which seem increasingly at odds with the world of automated systems.

Introduction: the circulation of faces

In a post 9/11 world security has become a big question for those feeling vulnerable. As in so many instances in social history the answer to this vulnerability is sought in a sort of certainty rooted in surveillance (Lyon, 1994, 2001, 2002; Dandeker, 1990). It is argued that through surveillance and early detection the problem can be solved. Security can be secured.

Surveillance is a powerful technology for social control, however, when surveillance becomes digitised then a there is a "step change in power, intensity and scope" (Graham and Wood, 2003). Digitisation permits the use of software algorithms (mathematical instructions –see Section 2) for automated identification of human biometrics (a bodily trace - see Section 2) With a biometric it is very difficult, if not impossible, for any individual to disassociate oneself (or

[1] Centre for the Study of Technology and Organisation, Lancaster University Management School, mailto:l.introna@lancaster.ac.uk

[2] Global Urban Research Unit (GURU), School of Architecture Planning and Landscape, University of Newcastle. mailto:d.f.j.wood@ncl.ac.uk

 ISSN: 1477-7487

be alienated) from one's biometric – in a sense you are your biometric (Van der Ploeg, 2002). Thus, if there is a match between your body and your biometric certainty over identity can be established. However effective surveillance also needs to be subtle, to be insinuated into the context of everyday life. Indeed in the security world the perfect unobtrusive biometric is considered the 'holy grail'. It is therefore not surprising that facial recognition system (FRS) have become a prime focus for the security establishment (Kopel & Krause, 2003). Not only are they relatively inexpensive, and supposedly effective, they require no involvement from their targets. Unlike other biometrics facial recognition can operate anonymously in the background. The targets do not need to surrender their face image, as they would their fingerprint, or their iris scan. A face can be captured and (de)coded without the consent or participation from those being targeted.

However, this 'captured' face image is only of use if it can be matched with an identifier. It requires a database of face images with associated identities. Unlike fingerprints or DNA samples, which are only collected when there is a reasonable level of suspicion of a crime, face images are routinely collected in society by a variety of institutions, such as when we apply for a driving licence, or a passport, or a library card, etc. It is the most common biometric in use by humans to identify other humans. Indeed, in any western society, if one would somehow cover, or be seen to attempt to disguise one's face, then there is almost an immediate assumption of guilt. One could almost say that there is an implicit common agreement to reveal our faces to others as a condition for ongoing social order. Thus, we tend to reveal our face to others and they to us. However, it seems that such an agreement only operates in a local and situated manner, as part of the social relationships we control. We would find it unacceptable if a stranger would photograph our face for no apparent reason. On the other hand we don't find it unacceptable to surrender our faces for the regulation of privileges—as long as we are in control of its use and circulation. In most cases it is used in the moment of authentication (by means of visual comparison) and then forgotten. However, what happens if our faces are collected anonymously, encoded, and start to circulate in an invisible network, even if it is for seemingly mundane reasons? When our face "becomes a bar code", in the words of Agre (2003). This seems not what we have in mind when we reveal our faces? It seems to us that the presence of FRS's may indeed be changing the implied relationships we assume when facing others. This concern becomes even more acute if we start to 'unpack' these algorithms to discover that to the algorithms 'all faces are not equal'. It is our contention that FRS's is a very powerful and ambiguous technology for social control. As such it requires much more scrutiny than it has had up to now.

The purpose of this paper is to open up for scrutiny the politics of facial recognition technology and its use in 'smart' CCTV. We aim to show that seemingly mundane design decisions may have important political consequences that ought to be subject to scrutiny.

This paper is one step in that direction. It is structured as follows: the first section will focus on the politics of technology and algorithmic surveillance systems in particular, considering first the broad politics of technology, then explaining the nature of algorithmic surveillance and biometrics, claiming that software algorithms are a particularly important domain of techno-politics; and finally considering both the growth of algorithmic biometric surveillance and the

potential problems with such systems. In the second section, we will give an account of FRS's, the algorithms they are based upon, and the biases embedded therein. In the third part, we will discuss the ways in which these biases may manifest itself in real world implementation of FRS's. Finally, we will make some policy suggestion for the future development of FRS's; it should be noted that the most common critiques of such systems are based on notions of privacy which seem increasingly at odds with the world of automated systems.

1. The Politics of Technology

The Micro-Politics of the Artefact

Technology is political (Winner, 1980). By this we mean that technology, by its very design, includes certain interests and excludes others. It is mostly an implicit politics, part of a mundane process of trying to solve practical problems. For example the ATM bank machine assumes a particular person in front of it. It assumes a person that is able to see the screen, read it, remember and enter a PIN code, etc. It is not difficult to imagine a whole section of society that does not conform with this assumption. If you are blind, in a wheelchair, have problem remembering, or unable to enter a PIN, because of disability, then your interest in accessing your account can be excluded by the ATM design. This exclusion of interests may not be obvious to the designers of ATMs as they may see their task as trying simply to solve a basic problem of making banking transactions more efficient for the 'average' customer doing average transactions. And they are mostly right — but if they are not, then their biases can become profoundly stubborn. These systems often seem like devices for surveillance and social control (in the sense of Foucault's *dispositif panoptique*), but as Lianos (2001, 2003) has recently pointed out, they are not designed with the monitoring and control of the human subject directly in mind, rather this is a potential (or secondary) function of systems for ensuring flow. Nevertheless the binary effects are in some senses quite irreversible. Where does the excluded go to appeal when they are faced with a stubborn and mute object such as an ATM? Maybe they can work around it, by going into the branch for example. This may be possible. However, this exclusion becomes all the more significant because of the political economic context in which these *dispositifs* exists and which they help to transform, for example if banks start to close branches or charge for an over-the-counter transaction (as is happening). Thus, as the micro-politics of the ATM becomes tied to, and multiplied through other exclusionary practices, what seems to be a rather trivial injustice soon may multiply into what may seem to be as an coherent and intentional strategy of exclusion (Introna and Nissenbaum, 2000). Yet there is often nobody there that 'authored' it as such (Foucault, 1975; Kafka, 1925). This paper will show how such an 'unauthored' strategy may be emerging in facial recognition technology.

Thus, the politics of technology is more than the politics of this or that artefact. Rather these artefacts function as nodes, or links, in a dynamic socio-technical network, or collective, kept in place by a multiplicity of artefacts, agreements, alliances, conventions, translations, procedures, threats, and so forth: in short by relationships of power and discipline (Callon 1986, 1991). Some are stable, even irreversible; some are dynamic and fragile. Analytically we can isolate and describe these networks (see Law 1991, for examples). However, as we survey the landscape of networks we cannot locate, in any obvious manner, where they begin nor where

they end. Indeed we cannot with any degree of certainty separate the purely social from the purely technical, cause from effect, designer from user, winners from losers, and so on.

In these complex and dynamic socio-technical networks ATMs, doors, locks, keys, cameras, algorithms, etc. — function as political 'locations' where values and interests are negotiated and ultimately 'inscribed' into the very materiality of the things themselves—thereby rendering these values and interests more or less permanent (Akrich, 1992; Latour, 1991). Through these inscriptions, which may be more or less successful, those that encounter and use these inscribed artefacts become, wittingly or unwittingly, enrolled into particular programmes, or scripts for action. Neither the artefacts nor those that draw upon them simply except these inscriptions and enrolments as inevitable or unavoidable. In the flow of everyday life artefacts often get lost, break down, and need to be maintained. Furthermore, those that draw upon them use them in unintended ways, ignoring or deliberately 'misreading' the script the objects may endeavour to impose. Nevertheless, to the degree that these enrolments are successful, the consequences of such enrolments can result in more or less profound political 'ideologies' that ought to be scrutinised. We would claim that the politics of artefacts is much more mundane and much more powerful than most other politics, yet it often evades our scrutiny. It is with this in mind that we can introduce the politics of algorithmic surveillance

2. Algorithmic surveillance

What is an Algorithm?

The word 'algorithm' derives from the hugely influential 9th Century Muslim mathematician, Muhammed ibn Musa al-Khwarizmi, who produced the first extant text on algebra - a term which also originates with him. 12th Century Christian scholars used al-Khwarizmi's name, latinized as 'algorismus' to differentiate his method of calculation from commonly used methods like the abacus or counting tables[3].

An algorithm is simply a mathematical, or logical, term for a set of instructions. Algorithms can be divided into trivial and non-trivial types, the former being sets of instructions that are only applicable to a specific situation, or a task that needs to further explanation, the latter being instructions that will provide answers given any compatible input. Texts on algorithmics often give the example of a recipe as a useful metaphor for understanding the concept, though in fact this is slightly inaccurate: the recipe is more like software (see below).

Algorithms form the basis of modern mathematics and most importantly here, the foundation of computing. However in themselves algorithms are not accessible to computers, they need to be translated into a form that computers have been programmed to understand. This process, known as coding (or hacking) produces software. Software is essentially composed of many coded algorithms linked together to produce a desired output from the hardware. In the metaphor mentioned above, the software is therefore the recipe. Computer hardware is not in itself usually algorithmic, rather it is composed of many physical switches (however small) which

[3] For more on the history of algorithms, see Chabert *et al.* (1999).

have two positions, on/off, 1/0 etc. These switches then respond to instructions from the software, once it has been translated (assembled) into binary machine code[4].

Algorithmic Surveillance

The term 'algorithmic surveillance' was coined by Norris and Armstrong (1999) in their pioneering book, *The Maximum Surveillance Society*. It is in literal terms surveillance that makes use of automatic step-by-step instructions. However it is used specifically to refer to surveillance technologies that make use of computer systems to provide more than the raw data observed. This can range from systems that classify and store simple data, through more complex systems that compare the captured data to other data and provide matches, to systems that attempt to predict events based on the captured data.

Thus many surveillance technologies currently in use have algorithmic aspects, but not all. A city-centre CCTV system that provides images that are watched and analysed by guards or police is not algorithmic. If such a system contains a computer which compares the faces of people captured by the cameras which those of known offenders, then it is. If typists enter in the health details of a patient into the Health Service computer, then it is algorithmic to a limited extent in that software determines the extent of the information that can be entered, however it becomes what is usually understood as algorithmic surveillance when, for example, a program is installed which compares the patient records against signs of particular disease risk-factors, and defines or categorises patients automatically.

Algorithmic Surveillance in Practice

There are now many algorithmic surveillance systems which watch over almost all aspects of this planet and beyond (if one includes systems like the Hubble deep-space telescope and the Cassini project). Many of these systems monitor the non-human (water and electricity flow etc.) and are thus left largely unconsidered by social researchers; although there was a temporary wave of concern prior to the year 2000 with the fear of the so-called Millenium Bug that apparently had the potential to cause many of these systems to fail or malfunction. There are many algorithmic systems which have a hybrid monitoring function, for example the recordings of cash withdrawals from ATM machines, credit and debit card transactions and purchases in stores to which we refered above.

There are also systems which algorithmically record data and sort about things, but things which are related to human beings for example, car number plate recognition. Again these systems do indirectly monitor people, but there is no necessary correlation between a particular human and a number plate although this is quite likely and in some cases legally restricted. Systems like movement recognition can be used both for nonhuman and inhuman things and human beings depending on the circumstances and details of the technology used. Examples of these again are often about flow management for example, the Prismatica/Cromatica movement-recognition developed for the London Underground to ensure the movement of passengers is efficient and safe. The original system as it turned out in operation had the unintended consequence of being able to detect potential suicides – as it was observed that they remained relatively motionless on

[4] For more on algorithms and computing, see Harel (1992).

the platform for longer periods than most before jumping (Norris 2003) – but again this an indirect consequence of the behaviour of human beings observed through a system of flow management. These systems are extremely interesting because, whatever their intention, they do transform the context of social interaction in quite fundamental ways creating what Lianos and Douglas (2000) call Automated Socio-Technical Environments (ASTEs).

Recent years have however seen the largely experimental introduction of automated systems for the direct monitoring of human beings based on physical traits unique to the individual. These 'biometric' identification systems include: gait recognition; fingerprint and palmprint recognition; facial recognition; and iris recognition. Each has its own technical merits and drawbacks and each is suitable for different uses in varying physical environments. The oldest of these are hand-geometry recognition systems[5] which internally have remained largely unchanged since the 1970s, and which still work very well in environments where access is restricted to a relatively small database of people.

The Growth of Algorithmic Surveillance

Before the attacks of September 11th 2001, the biometrics industry was expanding steadily but not spectacularly and was also facing increasing opposition from civil rights and privacy groups. In the immediate aftermath of the attacks, particularly in the USA, there was a general assumption that rights arguments would lose out during what one of us has elsewhere characterised as a period of 'surveillance surge' (Wood, Konvitz and Ball 2003) wherein those with an interest in new surveillance technologies promote them to a polity shocked enough by events not to consider their efficiency, effectiveness or wider implications as carefully as they might normally do.

Zureik (2004)[6] shows that within a few weeks of the terrorist attacks almost 17 bills were introduced in the United States Congress, including measures "to tighten immigration, visa, and naturalization procedures, allow tax benefits to companies that use biometrics, and check employee background at border and maritime check points."

He concludes that:

> the combination of public fear, lobbying efforts of the industry, and linkages between political and economic interests, have catapulted the industry to centre stage in the fight against terrorism – an industry that until September 11 was a marginal player in the security field (*ibid.*)

On the Silent Politics of the Software Algorithm

Having argued that technology is political, and introduced algorithmic surveillance systems, we now want to claim that the politics of information technology (in the form of software algorithms) is, in a sense, of a different order (Graham and Wood, 2003). We want to contend that

[5] Manual fingerprinting is of course far older, but not initially automated, and even now remains only partially so in most countries.

[6] Quotations taken from a pre-publication proof (no page numbers).

scrutinising information technology is particularly problematic since information technology, in particular algorithms, is what we would term a *silent* technology as opposed to a *salient* technology (Introna, 1998). Obviously we do not see this distinction as a dichotomy but rather as a continuum. As an attempt to draw this distinction some aspects are highlighted in Table 1 below.

***Silent* technology is:**	***Salient* technology is:**
Embedded / hidden	On the 'surface' /conspicuous
Passive operation (limited user involvement)	Active operation (fair user involvement)
Application flexibility (open ended)	Application stability (firm)
Obscure (form/operation/outcome)	Transparent (form/operation/outcome)
Mobile (*soft*-ware)	Located (*hard*-ware)

Table 1: *Silent versus Salient Technology*

Facial recognition algorithms in 'smart' CCTV is a particularly good example of a silent technology. The facial recognition capability can be imbedded into existing CCTV networks, making its operation impossible to detect. Furthermore, it is entirely passive in its operation. It requires no participation or consent from its targets—it is "non-intrusive, contact-free process" (Woodward *et al.*, 2003: 7). Its application is flexible. It can as easily be used by a supermarket to monitor potential shoplifters (as was proposed and later abandoned, by the Borders bookstore), by casinos to track potential fraudsters, by law enforcement to monitor spectators at a Super Bowl match (as was done in Tampa, Florida), or used for identifying 'terrorists' at airports (as is currently in operation at various US airports). However, most important is the obscurity of its operation.

This obscurity is due to two factors. First, most of the software algorithms at the heart of facial recognition systems are propriety software objects. Thus, it is very difficult to get access to them for inspection and scrutiny. More specifically, even if you can go through the code line by line, it is impossible to inspect that code *in operation*, as it becomes implemented through multiple layers of translation for its execution. At the most basic level we have electric currents flowing through silicon chips, at the highest level we have programme instructions, yet it is almost impossible to trace the connection between these as it is being executed. Thus, it is virtually impossible to know if the code you inspected is the code being executed, when executed. In short, *software algorithms are operationally obscure*. Second, most of the algorithms in facial recognition are based on very sophisticated statistical methods that only a handful of experts can interpret and understand. Indeed it seems that even they have been surprised by the behaviour of their algorithms (Philips *et al.*, 2003). Thus, for most ordinary members of society facial recognition systems are somewhat exotic and obscure 'black boxes'. After all they do well what we find difficult to do — identify faces. This obscurity together with their obvious sophistication may give them a legitimacy beyond that which they deserve. In moments of uncertainty they may be taken as more authoritative than the humans involved — this could have important implications as we will argue and show below.

How then can we scrutinise these software algorithms? The only recourse we have to evaluating these algorithms is to look at their performance under controlled conditions as was done in the Facial Recognition Vendor Test (FRVT) of 2000 and 2002. These evaluations will be the core of our analysis below. This is not entirely satisfactory as these tests are mostly focused on the evaluation of the efficiency and effectiveness of the systems and not focused on the discovery of biases as such. In addition such tests are carried out through a highly interested and biased set of parties: a transatlantic but US-military and intelligence dominated collective led by the US Department of Defense (DoD).

It is our argument that the silent nature of information technology makes it difficult for society to scrutinise it. Furthermore, this inability to scrutinise creates unprecedented opportunities for this silent and 'invisible' micro politics to become pervasive (Graham and Wood, 2003). Thus, we tend to have extensive community consultation and impact studies when we build a new motorway. However, we tend not to do this when we install CCTV in public places or when we install FRSs in public spaces such as airports, shopping malls, etc. To put is simply: most informed people understand the cost (economic, personal, social, environmental) of a motorway, however they do not understand the cost of FRSs in 'smart' CCTV. This paper is an attempt to make this cost more visible.

3. The politics of Facial Recognition Systems

Getting a digital face: the facial recognition system

Figure 1 below depicts the typical way that a facial recognition system (FRS) system can be made operational.

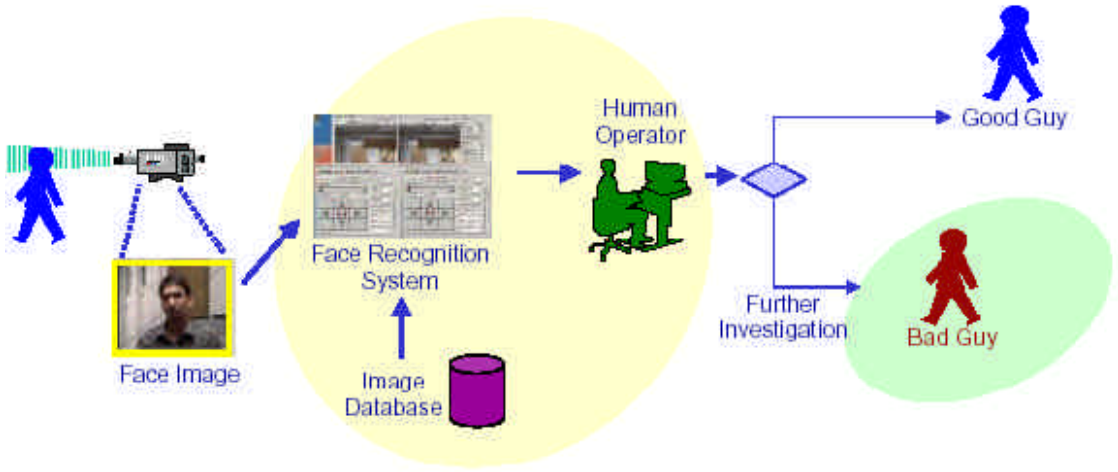

Figure 1: ***Overview of FRS***
(Source: FRVT, 2002)

The first step is the capturing of a face image. This would normally be done using a still or video camera. As such it can be incorporated into existing 'passive' CCTV systems. However,

locating a face image in the field of vision is not a trivial matter. The effectiveness of the whole system is dependent on the quality of the captured face image. The face image is passed to the recognition software for recognition (identification or verification). This would normally involve a number of steps such as normalising the face image and then creating a 'template' of 'print' to be compared to those in the database. If there is a 'match' then an alarm would solicit an operator's attention to verify the match and initiate the appropriate action. The match can either be a true match which would lead to investigative action or it might be a 'false positive' which means the recognition algorithm made a mistake and the alarm would be cancelled. Each element of the system can be located at different locations within a network, making it easy for a single operator to respond to a variety of systems.

For our analysis we want to concentrate on steps two and three of the system. We want to scrutinise the FR algorithms, the image database (also called the gallery) and the operators. At each of these points important decisions are made which may have an important political implication.

Facial Recognition Algorithms and Reduction

Research in software algorithms for facial recognition has been ongoing for the last 30 years or so (Gross *et al.*, 2001). However, advances in information technology and statistical methods have given impetus to this development with seemingly excellent recognition results and low error rates—at least in ideal laboratory conditions. It is possible to identify two main categories of algorithms according to Gross *et al.* (2001):

- *Image template algorithms*. These algorithms use a template-based method to calculate the correlation between a face and one or more standard templates to estimate the face identity. These standard templates tend to capture the global features of a gallery of face images. Thus, the individual face identity is the difference between (or deviation from) the general or 'standard' face. This is an intuitive approach since we as humans tend to look for distinctive features (or differences from the general) when we identify individuals. Some of the methods used are: Support Vector Machines (SVM), Principal Component Analysis (PCA), Neural Networks, Kernel Methods etc. The most commercially known template based algorithm is the MIT Bayesian Eigenface technique, which has been developed with the PCA method. During various tests conducted in 1996, its performance was consistently near the top compared to other available at the time.

- *Geometry feature-based algorithms*. These methods capture the local facial features and their geometric relationships. They often locate anchor points at key facial features (eyes, nose, mouth, etc), connect these points to form a net and then measure the distances and angles of the net to create a unique face 'print'. The most often cited of these is the technique known as Local Feature Analysis (LFA), which is used in the Identix (formerly Visionics) face recognition system called FaceIt. The LFA method, in contrast to the PCA technique, is less sensitive to variations in lighting, skin tone, eye glasses, facial expression, hair style, and individual's pose up to 35 degrees.

The commonality in both of these groups of techniques is the issue of *reduction*. In order to be

efficient in processing and storage the actual face image gets reduced to a numerical representation (as small as 84 bytes or 84 individual characters in the case of FaceIt). With this reduction certain information is disregarded (as incidental or irrelevant) at the expense of others. It is here that we need to focus our analysis. What are the consequences of the process of reduction? It would be best to understand this through some detailed study of the logic and operation of these algorithms in diverse settings with diverse databases. This has not yet being done (not even in the FRVT 2002, which has been the most comprehensive thus far). Nevertheless, with our limited knowledge we can make some logical conclusions and then see how these may play out in the FRVT 2002 evaluations. How will the reduction effect the performance of these algorithms?

- *Template based algorithms*. In these algorithms certain biases become built into the standard template. It obviously depends on the gallery used to create the standard template as well as the range of potential variations within a population. For example, because minorities tend to deviate the most from the standard template they might become easier to recognise.

- *Feature based algorithms*. These algorithms do not have an initial bias. However, because of the reduction the 'face prints' generated are in close proximity to each other. Thus, as the gallery database increases more and more face prints are generated in ever diminishing proximity, thereby making the discrimination required for the recognition task more difficult. Therefore the operation of the system deteriorates rapidly as the database increases (this is also true for template based algorithms). It also makes the system dependent on good quality face images. The implication of this is that the system will operate at its best with a small database and good quality face capture, such as an operator assisted face capture (reintroducing the operator bias). In addition to this it will tend to be better at identifying those that are more distinctive, or less similar, to those already in the database (such as minorities).

Thus, in both cases we would expect some form of bias to emerge as a result of the reduction. Is this conclusion borne out by the performance of these algorithms in the FRVT? Let us now consider the results of these evaluations.

The Evaluations: Reduction, Operation and Error

The most significant evaluation of FRSs happened with the Facial Recognition Vendor Tests of 2002 (Philips et al., 2003). These test were independent tests sponsored by a host of organizations such as Defense Advanced Research Projects Agency (DARPA), the Department of State and the Federal Bureau of Investigation. This evaluation followed in the footsteps of the earlier FRVT of 2000 and the FERET evaluations of 1994, 95 and 96. In the FRVT 2002 ten FRS vendors participated in the evaluations. The FRVT of 2002 were more significant than any of the previous evaluations because of:

- The use of a large database (37437 individuals)
- The use of a medium size database of outdoor and video images
- Some attention given to demographics

The large database (referred to as the HCInt data set) is a subset of a much larger database which was provided by the Visa Services Directorate, Bureau of Consular Affairs of the U.S. Department of State. The HCInt data set consisted of 121,589 images of 37,437 individuals with at least three images of each person. All individuals were from the Mexican non-immigrant visa archive . The images were typical visa application type photographs with a universally uniform background, all gathered in a consistent manner.

The medium size database consisted of a number outdoor and video images from various sources. *Figure 2* below gives an indication of the images in the database. The top row contains images taken indoors and the bottom contains outdoor images taken on the same day. Notice the quality of the outdoor images. The face is consistently located in the frame and similar in orientation to the indoor images.

Figure 2: **Indoor and outdoor images from the medium data base.**
(Source: FRVT, 2002)

For the identification task an image of an unknown person is provided to a system (assumed to be in the database). The system then compares the unknown image (called the probe image) to the database of known people. The results of this comparison are then presented by the system, to an operator, in a ranked listing of the top *n* 'candidates' (referred to as the 'rank', typically anywhere from 1 to 50). If the correct image is somewhere in the top *n*, then the system is considered to have performed the identification task correctly. *Figure 3* below indicates the performance at rank 1, 10 and 50 for the three top performers in the evaluation.

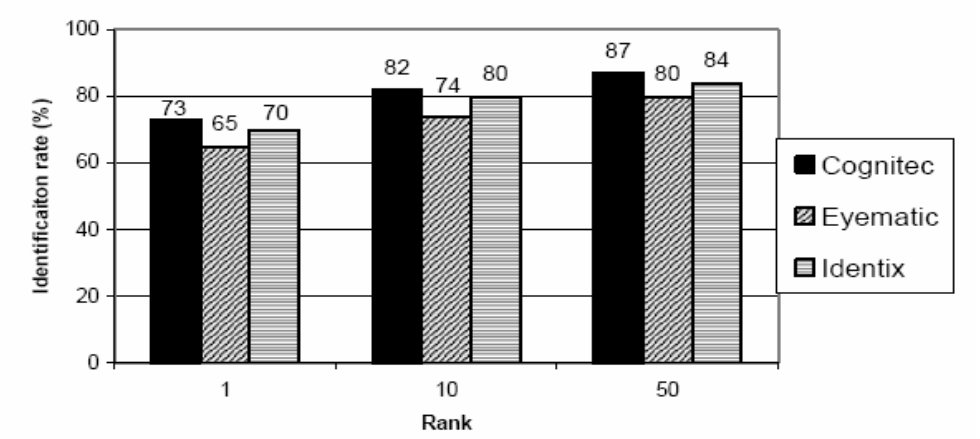

Figure 3: **Performance at rank 1, 10 and 50 for the three top performers in the evaluation** (from FRVT 2002, Overview and Summary, p.9.)

With the very good images from the large database (37,437 images) the identification performance of the best system at rank one is 73% at a false accept rate of 1%. There is a tradeoff between the recognition rates and the level of 'false accepts' (incorrect identification) one is prepared to accept, the false accept rate. If you are prepared to accept a higher false accept rate then the recognition performance can go up. However, this will give you more cases of false identification to deal with. This rate is normally a threshold parameter that can be set by the operators of the system.

What are the factors that that can detract from this 'ideal' performance? There might be many. The FRVT 2002 considered three:

- Indoor versus outdoor images
- The time delay between the database image and the probe image
- The size of the database

The identification performance drops dramatically when outdoor images are used—in spite of the fact that they can be judge as relatively good—as indicated above. One would not expect a typical video camera to get this quality of image all the time. For the best systems the recognition rate for faces captured outdoors (i.e. less than ideal circumstances) was only 50% at a false accept rate of 1%. Thus, as the report concluded: "face recognition from outdoor imagery remains a research challenge area." The main reason for this problem is that the algorithm cannot distinguish between the change in tone, at the pixel level, caused by a relatively dark shadow, versus such a change caused by a facial feature. As such it starts to code shadows as facial features. The impact of this on the identification may be severe if it happens to be in certain key areas of the face.

As one would expect, the identification performance also decreases as time laps increases between the acquisition of the database image and the newly captured probe image presented to a system. FRVT 2002 found that for the top systems, performance degraded at approximately 5% points per year. It is not unusual of the security establishment to have a relatively old

photograph of a suspect. Thus, a two year old photograph will take 10% off the identification performance. A study by the US National Institute of Standards and Technology found that two sets of mugshots taken 18 months apart produced a recognition rate of only 57% (Brooks, 2002). Gross et al (2001: 17) found an even more dramatic deterioration. In their evaluation the performance dropped by 20% in recognition rate for images just two weeks apart. Obviously these evaluations are not directly comparable. Nevertheless, there is a clear indication that there may be a significant deterioration when there is a time gap between the database image and the probe image.

What about the size of the database? For the best system, "the top-rank identification rate was 85% on a database of 800 people, 83% on a database of 1,600, and 73% on a database of 37,437. For every doubling of database size, performance decreases by two to three overall percentage points" (Philips *et al*., 2003: 21). What would this mean for extremely large databases? For example the UK fingerprint database consists of approximately 5.5 million records. If one had a similar size 'mugshot' database how will the algorithms perform in identifying a probe image in that database? If one takes the decrease to be 2.5% for every doubling of the database, and use 73% at 37,437 as the baseline, then one would expect the identification performance to be approximately 55% in ideal conditions and as low as 32% in less than ideal conditions.

To conclude this discussion we can imagine a very plausible scenario where we have a large database, less than ideal image due to factors such as variable illumination, outdoor conditions, poor camera angle, etc, and the probe image is relatively old, a year or two. Under these conditions the probability to be recognized is very low, unless one sets the false accept rate to a much higher level, which means than there is a risk that a high number of individual may be subjected to scrutiny for the sake of a few potential identifications. What will be the implications of this for practice? We will take up this point again below. Obviously we do know how these factors would act together and they are not necessarily cumulative. Nevertheless it seems reasonable to believe that there will be some interaction that would lead to some cumulative affect.

Such a conclusion can make sense of the Tampa Police Department case reported by ACLU (Stanley and Steinhardt, 2002) as well as the Palm Beach International Airport also reported by the ACLU. In the Tampa case the system was abandoned because of all the false positive alarms it generated. As far as it could be ascertained it did not make one single positive identification. In the Palm Beach Airport case the system achieved a mere 47% correct identifications of a group of 15 volunteers using a database of 250 images (Brooks, 2002). In Newham, UK, the police admitted that the FaceIt system had not made a single positive identification, in spite of working with a small database. One could argue that there might not have been the potential for a match to be made as none of the individual in the database actually appeared in the street. Nevertheless, the system could not 'spot' a *Guardian* journalist, placed in the database, that intentionally presented himself in the two zones covered by the system (Meek, 2002). These cases indicate the complexity of real world scenarios. We now want to move to the focal concern of this paper namely the question of biases in the algorithms themselves.

Reduction and Biased Code

The most surprising outcome – for those involved – of the FRVT 2002 is the realization that the algorithms displayed particular identification biases. First, recognition rates for males were higher than females. For the top systems, identification rates for males were 6% to 9% points higher than that of females. For the best system, identification performance on males was 78% and for females was 79%. Second, recognition rates for older people were higher than younger people. For 18 to 22 year olds the average identification rate for the top systems was 62%, and for 38 to 42 year olds was 74%. For every ten years increase in age, on average performance increases approximately 5% through age 63. Unfortunately they could not check race as the large data set consisted of mostly Mexican non-immigrant visa applicants. However, research by Givens *et al.* (2003), using PCA algorithms, has confirmed the biases in the FRVT 2002 (except for the gender bias) and also found a significant race bias. This was confirmed using balanced databases and controlling for other factors. They concluded that: "Asians are easier [to recognize] than whites, African-Americans are easier than whites, other race members are easier than whites, old people are easier than young people, other skin people are easier to recognize than clear skin people…" (8). Their results are indicated in Figure 4 below.

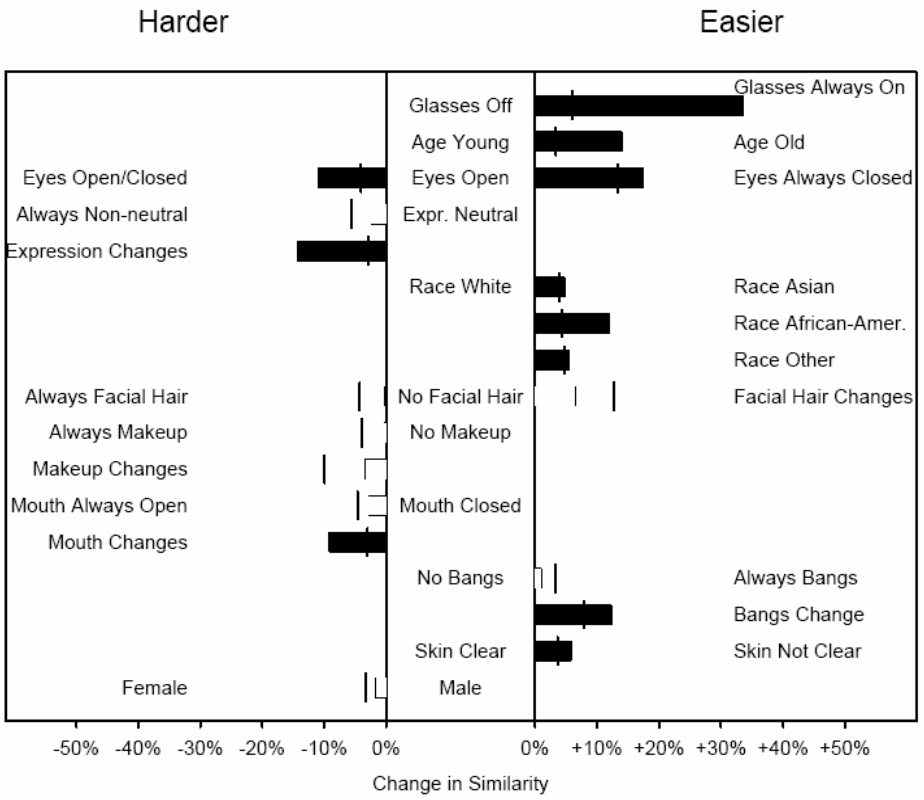

Figure 4: ***Factors making it harder or easier to correctly identify a probe image presented to a system***
(Source: Givens *et al.*, 2003)

These results were also found in another context by Furl, Phillips and O'Toole (2002) in their study of recognition performance by thirteen different algorithms. One can legitimacy ask

whether these differences, probably in the order of 5-10%, really makes a difference? Are they not rather trivial? We would argue that taken by themselves they may seem rather trivial. However, as we argued earlier on, it is when these trivial differences become incorporated into a network of practices that the may become extremely important. This is what we now want to explore: the politics of the digital face as it becomes imbedded in practices.

4. The politics of the digital face

FRS's: Efficient, Effective and Neutral

Many security analysts see FRSs as the ideal biometric to deal with the new emerging security environment. They claim that it is efficient (FaceIt only requires a single 733 Mhz Pentium PC to run) and effective, often quoting close to 80% recognition rates from the FRVT 2002 evaluation while leaving out of the discussion issues of the quality of the images used in the FRVT, size of the database, the elapsed time between database image and probe image, etc. But most of all they claim that these systems "performs equally well on all races and both genders. Does not matter if population is homogeneous or heterogeneous in facial appearance" (Faceit technical specification[7]). This claim is not only made by the suppliers of FRSs such as Identix and Imagis Technologies. It is also echoed in various security forums: "Face recognition is completely oblivious to differences in appearance as a result of race or gender differences and is a highly robust Biometrics"[8] Even the critical scholar Gary Marx (1995: 238) argued that algorithmic surveillance provides the possibility of eliminating discrimination. The question is not whether these claims are correct or not. One could argue that in a certain sense they are correct. The significance of these claims is the way they *frame* the technology. It presents the technology itself as neutral and unproblematic. More than this it presents the technology as a solution to the problem of terrorism. Atick of Identix claimed, in the wake of the 9/11 attacks, that with FaceIt the US has the "ability to turn all of these cameras around the country into a national shield" (O'Harrow, 2001). He might argue that in the face of terrorism 'minor' injustices (biases in the algorithms) and loss of privacy is a small price to pay for security. This may be so, although we would disagree.

Nevertheless, our main concern is that these arguments present the technical artifacts in isolation with disregard to the socio-technical networks within which they will become imbedded. As argued above, it is not just the micro-politics of the artifact that is the issue. It is how these become multiplied and magnified as they become tied to other practices that is of significance. We need to understand the 'network effects', as it were, of the micro-politics of artifacts. This is especially so for silent digital technology. There is every reason to believe that the silent and non-invasiveness of FRSs make it highly desirable as a biometric for digital surveillance. It is therefore important that this technology becomes scrutinized for its potential in the socio-technical network of digital surveillance. Thus, not just as isolated objects as was done in the FRVTs but in its multiplicity of implementations and practices. We would claim it is here where

---

7 http://www.identix.com/newsroom/news_biometrics_face_acc.html

8 http://www.ats-computers.com/biometrics/face.html
http://www.biocom.tv/BIOMETRICS_types.htm

the seemingly trivial biases may become very important as they become incorporated into actual practices.

FRS's in Practice: Alarms, Biases and Suspects

There is an urgent need for an in-depth study of FRSs in practice (as has been done with CCTV by Norris and Armstrong (1999) and others). However, since we currently only have a limited number of systems in operation and due to the sensitivity of these implementations it is unlikely that we would be able do so in the near future. Thus, in the face of this limitation, we propose to outline what we consider to be a highly probable scenario of how these digital biases may become incorporated into other practices that would render these seemingly trivial biases significant.

Based on the FRVT of 2002 we know that, although FRSs have the capability to achieve a 70-85% accuracy rate, this is only in ideal circumstances. The system's performance degrades significantly in an uncontrolled 'face-in-the-crowd' environment, with a large database, and where there is an elapsed time between the database image and the probe image. This would seem to us to be a usual rather than an unusual situation. What will happen if the system's performance degrades under these rather usual conditions?

We would propose that two possibilities are most likely. First, it is possible that the operators will become so used to false positives that they will start to treat all alarms as false positives thereby rendering the system useless. Alternatively, they may deal with it by increasing the identification threshold (requesting the system to reduce the number of false positives). This will obviously also increase the false negatives, thereby raising all sorts of questions about the value of the system into question. However, more important to us, with an increased threshold small differences in identifiability (the biases outlined above) will mean that those that are easier to identify by the algorithms (African-Americans, Asians, dark skinned persons and older people) will have a greater probability of triggering the alarm. If the alarm is an actual positive recognition then one could argue that nothing is lost. However, it also means that these groups would be subjected to a higher probability of scrutiny as false positives, i.e. mistaken identity. Moreover we would propose that this scrutiny will be more intense as it would be based on the assumption that the system is working at a higher level and therefore would be more accurate. In such a case existing biases, against the usual suspects (such as minorities), will tend to come into play (Norris and Armstrong, 1999). The operators may even override their own judgements as they may think that the system under such high conditions of operation must 'see something' that they do not. This is highly likely as humans are not generally very good a facial recognition in pressurised situations as was indicated in a study by Kemp *et al.* (1997). Thus, under these conditions the bias group (African-Americans, Asians, dark skinned persons and older people) may be subjected to disproportionate scrutiny, thereby creating a new type of 'digital divide' (Jupp in Graham and Wood, 2003: 234).

How likely is this scenario? We believe it to be more likely than we presume. We have only the following anecdotal evidence reported in the *Discover Magazine* of an installation at the Fresno Yosemite International Airport to suggest:

> "[The system] generates about one false positive for every 750 passengers scanned, says Pelco vice president Ron Cadle. Shortly after the system was installed, a man *who looked as if he might be from the Middle East* set the system off. "The gentleman was detained by the FBI, and he ended up spending the night," says Cadle. "We put him up in a hotel, and he caught his flight the next day."
>
> (Garpinkle, 2002: 19 – *emphasis added*)

To produce only one false positive per 700 passengers the system had to operate with a very restricted false positive rate, thereby suggesting that an alarm must 'mean something'. Notice that one of the false positives was a man supposedly from 'Middle Eastern' origin. The individual was detained and questioned by the FBI because "looked as if he might be from the Middle East" in spite of the fact that he was obviously a false positive. There could be many explanations for this action. Nevertheless, it is likely that they may have decided to detain him 'just in case' the system saw something they did not see. This case clearly demonstrates the scenario we outline above. Our analysis has demonstrated that seemingly trivial differences in recognition rates, within the algorithm, can indeed have important political implications for some when it becomes incorporated into a whole set of socio-technical surveillance practices.

One might imagine that in an environment where there is an acute sense of vulnerability it would not be unreasonable to store these false positives in a database 'just in case'. These false positive may then become targets for further scrutiny. Why? Just because they have features that make them more distinctive. We are not saying that this will happen. We are merely trying to indicate how seemingly trivial 'technical issues' can add up to political ideologies at the expense of some for the sake of others. This is the issue of the politics of FRSs. This is particularly dangerous politics in the case of silent technologies such as FRSs.

Other areas remain problematic.In a legal-technical review of the technology based largely on a previous FRVT, Michael Bromby (2002) claims that facial recognition offers a significant improvement over human identification, which can be extremely limited. However he is also careful to argue that because of shortcomings, it can only be considered as supplementary to other human and technological forms of recognition. In fact the flaws in Facial Recognition are similar to those of human identification: the problems of environmental conditions, angle of view, the gradual decrease in ability to recognise as the number of individual faces to choose from increases, the inability to deal with aging faces, and so on. Psychological research has shown that human beings can only definitively recognise a limited number of people, and whilst this has sometimes been attributed to evolutionary factors in human neural development, Facial Recognition's troubles with larger numbers may indicate that human facial types are simply more limited in their basic variety than we often assume. The problem is that the simplest technical solution – increasing the number of variables – leads to increasing complexity (and therefore more problems) and also potentially makes the systems prey to less fundamental changes in faces: hair, colour etc.

More complex solutions, for example, the use of predictive algorithms (as for example in the Cromatica system operating on the London Underground) and neural networks (connectionism)

to allow heuristics (learning) in recognition systems, or 'best guess' and fuzzy logic systems, take the system further away from the (at least relative) certainty about identity demanded by the state and other users and more towards simulation or reconstruction. Surveillance of course has always had a strong connection to simulation (see: Bogard, 1996; Graham, 1998), but it seems that with heuristic systems and even with simpler technologies analysed by FRVT 2002 like normalisaton and three-dimensional morphing (where several images are converted into a virtual model of the head), we are moving further along the spectrum towards outright simulation which raises questions as to what exactly software-driven surveillance systems are 'seeing'. The complex modeling technologies also make the problem of enrolment – the difficulties of obtaining the necessary images for the database – still more problematic.

A vital point is that the human actants involved in facial recognition collectives need to be aware of these limitations. And it is not always evident that they are. A paper by Clive Reedman, former Director of the UK's Police Information Technology Organisation (PITO) posted onto the Biometrics discussion list, refers to "the success of a CCTV/Facial Recognition implementation in London's East End" (Reedman 2002: 7), when we have seen that success is hardly how the implementation of the FaceIT system in Newham can be characterised. Indeed the same paper lists quite clearly the theoretical attractions of facial recognition to law enforcement:

> manually attempting to find a 'face in the crowd', or identify a suspect from pictures of known offenders is a notoriously difficult task, as well as a very costly one in terms of police time. Just watch a single video monitor in a local council's control room for hours on end waiting for a particular individual to appear for a second or two and you will soon realise the concept of 'face blindness'. See the success of a television programme such as the BBC's Crimewatch, which relies heavily on the fact that images can be shared amongst a national audience and you will soon grasp that the chances of identifying an individual increase dramatically the wider that audience is. (Reedman, 2002: 6)

However if there is any 'law' in the history of technology it is that technologies are rarely used in ways that their inventors intended. Mitch Gray (2003) in a recent piece asks 'Will we recognise the facial recognition society?' This is the right question with the wrong focus. There will not be a facial recognition society, only partly because societies are never entirely defined by particular technologies even in the era of ASTEs, but mainly because the limits of facial recognition may mean that it will only ever be of limited use on its own. The use of facial recognition in the way in which Reedman describes it above is far from the techno-optimistic PR of Biometrics companies, rather is seen more mundanely as an efficient labour-saving device. The Tampa experience suggests however that notions of efficiency are also questionable. Gray concentrates on the development of micro-expression recognition technologies, but given the problems of facial recognition at much more basic stages, it seems that this is some distance away.

None of this means that progress will slow the development and improvement of FRS's. There is no doubt in our minds that facial biometrics will remain a very important part of the future security infrastructure. Kopel and Krause (2003) report that: "As of June 2001 the Departments

of Justice and Defence had given about $21.3 million and $24.7 million, respectively, to the research and development of FRSs." Its perceived efficiency, ease of implementation and invisible nature make FRS the ideal biometric technology for the foreseeable future.

Conclusions

This paper has shown simply that at present FRS's do not work in the way that is claimed by their manufacturers, and that their use is at present highly limited in conjunction with open-street CCTV. However the current enthusiasm for FRS's means that FRS's require careful scrutiny and regulation. We can not remain naïve about such a powerful technology. On the other hand, facial recognition can also be seen as merely a temporarily fashionable biometric amongst many other surveillance technologies currently being advanced, and it is not the success or failure of particular technologies that is driving the development of surveillance, but a move towards integration of multiple technologies and new ways of managing information flow from these multiple technologies.

This has important implications for theorising CCTV in that it emphasises that surveillance at root is founded on sorting and categorisation not on vision (see Lyon 2004). Theorising visibility or vision cannot therefore provide in itself a general understanding of 'surveillance'. Further it emphasises the necessity of a socio-technical approach which is able to integrate technological development into a more general understanding of what 'society' means, rather than seeing it as something external to core social variables.

An important part of our democratic society is our supposed equality before the law. Unfortunately, as many studies of police practices have shown, this is not always the case (Marx, 1988). These studies show that humans carry their biases into their workplace. That it is not possible to simply exclude these prejudices from your workplace behaviour if they are already part of your social make-up. Nevertheless, we are mostly aware of this. As such, one can always legitimately appeal for further consideration and scrutiny. Indeed we have developed a variety of mechanisms and procedures to scrutinise the behaviour of law enforcement officers. However, when it comes to technology we mostly assume it to be neutral and value free. Thus, we tend not to subject artefacts to the same level of scrutiny.

It is our view that this social and technical distinction with respect to technology is inappropriate for two reasons: (a) technical artefact already embody values in its design, it is 'society made durable' as suggested by Latour, and (b) technical artefact never act in isolation but become imbedded into a socio-technical network in which the micro-politics of the artefact can become multiplied and sized upon in many unexpected ways (Introna and Nissenbaum, 2000). We believe we have given an illustration of this as exemplified in FRS. What would the policy implication of such an analysis be. There are many. We will highlight the following:

- A need for more detailed studies of FR algorithms with a particular emphasis on biases. We need to understand why these biases emerge and what we ought to do to eliminate or limit them.

- A need for more detailed studies of actual implementations of FRSs. What are the appropriate ways to imbed this technology into a larger security infrastructure
- The development of an appropriate legal framework to prevent the misuse of the technology (especially as private installations increase). There is no doubt that this technology will also contribute to 'surveillance creep' as argued by Marx (1988: 2).
- A very strong legal framework that prohibit or control the circulation of individuals facial biometric ('face prints') without due process.

Obviously more in-depth study of actual installations of FRS's and a continuous and careful watch on technological development in the biometrics industry and its political economy are required. Nevertheless, we believe we have demonstrated that there are many aspects of this silent technology that still needs to be scrutinised. It is not feasible to remain naïve to the politics of such a powerful technology.

References

Akrich, M. (1992) 'The De-scription of Technical Objects.' *Shaping Technology/Building Society*. Ed. W.E. Bijker and J Law. Cambridege: MIT Press, 205-224.

Agre, P.E. (2003), 'Your Face is Not a Bar Code: Arguments against face recognition in public places', http://dlis.gseis.ucla.edu/pagre [accessed 1 September 2004]

Bogard W. (1996) *The Simulation of Surveillance: Hypercontrol in Telematic Societies* Cambridge: Cambridge University Press.

Bromby, M. (2002) 'To be Taken at Face Value? Computerised identification', *Information and Technology Law Journal Journal*, 11(1): 63-73.

Brooks, M. (2002) 'Face-off' , *New Scientist*, 175 (2399), 9 July.

Callon, M. (1986) 'Some elements of a sociology of translation: domestication of the scallops and the fishermen of St Brieuc Bay.' In J. Law (ed.) *Power, Action and Belief*. London: Routledge & Kegan Paul, 196-233.

Callon, M. (1991) 'Techno-economic networks and Irreversibility', in J. Law (ed.) *The Sociology of Monsters: Essays on Power, Technology and Domination*. London: Routledge, 132-161.

Chabert, J. (ed.) (1999) *A History of Algorithms: From the Pebble to the Microchip*, Berlin: Springer-Verlag.

Dandeker, C. (1990), *Surveillance, Power and Modernity*, Oxford, UK: Polity Press.

Elmer, G. (2003) 'A diagram of panoptic surveillance' *New Media & Society* 5(2): 231-247.

Furl, N., J.P. Phillips, and A.J O'Toole, (2002) 'Face recognition algorithms and the other-race effect: computational mechanisms for a developmental contact hypothesis'. *Cognitive Science* 26: 797-815.

Facial Recognition Vendor Test 2002 (FRTV2002) http://www.frvt.org/FRVT2000/default.htm [accessed 1 September 2004]

Foucault, M. (1975) *Discipline and Punish, The Birth of the Prison*, London, UK: Penguin.

Garpinkle, S. (2002) 'Don't count on face-recognition technology to catch terrorists', *Discover*, 23(9): 17-20.

Givens, G., J.R. Beveridge, B.A. Draper and D. Bolme, (2003), *A Statistical Assessment of Subject Factors in the PCA Recognition of Human Faces.* http://www.cs.colostate.edu/evalfacerec/papers/csusacv03.pdf [accessed 1 September 2004]

Graham, S and D. Wood (2003) 'Digitizing surveillance: categorization, space and inequality', *Critical Social Policy*, 20(2), 227-248.

Gray, M. (2003) 'Urban surveillance and panopticism: will we recognise the facial recognition society?' *Surveillance & Society* 1(3): 314-330. http://www.surveillance-and-society.org/articles1(3)/facial.pdf

Gross, R., J. Shi and J.F. Cohn (2001) *Quo Vadis Face Recognition?* http://dagwood.vsam.ri.cmu.edu/ralph/Publications/QuoVadisFR.pdf [accessed 1 September 2004]

Harel, D. (1992) *Algorithmics: the Spirit of Computing*, Reading MA: Addison-Wesley.

Huber, P. and M.P. Mills (2002) 'How technology will defeat terrorism', *City Journal* 12(1). http://www.city-journal.org/html/12_1_how_tech.html [accessed 1 September 2004]

Introna, L.D. (1998) 'Oppression, Resistance and Information Technology: Some Thoughts on Design and Values', Paper presented at Design for Values: Ethical, Social and Political Dimensions of Information Technology, workshop sponsored by the NSF DIMACS held at Princeton University, USA, 27 February - 1 March, 1998.

Introna, L.D. and H. Nissenbaum (2000) 'The Internet as a democratic medium: why the politics of search engines matters', *Information Society* 16(3): 169-185.

Kafka, F. (1925) *The Trial*, London: Penguin Books.

Kemp, R., N. Towell and G. Pike (1997) 'When seeing should not be believing: photographs, credit cards and fraud'. *Applied Cognitive Psychology*, 11: 211-222.

Kopel, D and M. Krause (2003) 'Face the FactsFacial recognition technology's troubled past – and troubling future'. http://www.reason.com/0210/fe.dk.face.shtml [accessed 1 September 2004]

Latour, B. (1991) 'Technology is society made durable.' *A Sociology of Monsters: Essays on Power, Technology and Domination*. Ed. J. Law. London: Routledge, 103-131.

Law, John. (ed.) (1991) *The Sociology of Monsters: Essays on Power, Technology and Domination*. London: Routledge.

Lianos, M. (2001) *Le Nouveau Contrôle Social*, Paris: L'Harmattan.

Lianos, M. (2003) 'Social Control After Foucault' (trans. D. Wood and M. Lianos) *Surveillance & Society* 1(3): 412-430 http://www.surveillance-and-society.org/articles1(3)/AfterFoucault.pdf

Lyon, D. (1994) *The Electronic Eye: the Rise of the Surveillance Society*, Cambridge, UK: Polity Press.

Lyon, D. (2001) *Surveillance Society, Monitoring everyday life*, Milton Keynes, UK: Open University Press.

Lyon, D. (2002), *Surveillance After September 11, 2001*, http://www.fine.lett.hiroshima-u.ac.jp/lyon/lyon2.html [accessed 1 September 2004]

Lyon, D. (ed.) (2003) *Surveillance as Social Sorting: Privacy Risk and Digital Discrmination.* London: Routledge.

Marx, G.T. (1988) *Undercover: Police surveillance in America*. Berkeley: University of California Press.

Marx, G.T. (1995) 'The Engineering of Social Control: The search for the silver Bullet', in J. Hagen and R. Peterson (eds.) *Crime and Inequality*. Stanford, CA: Stanford University Press, 225-46.

Meek, J. (2002) 'Robo cop: Some of Britain's 2.5 million CCTV cameras are being hooked up to a facial recognition system designed to identify known criminals. But does it work?' *Guardian*, 13 June, 2002.

Norris, C. and G. Armstrong (1999), *The Maximum Surveillance Society, The Rise of CCTV*, Oxford: Berg.

O'Harrow Jr., R. (2001) 'Facial Recognition System Considered For U.S. Airports' *Washington Post*, Monday, 24 September, 2001: A14.

Phillips, P., P. Grother, R. Micheals, D.M. Blackburn, E. Tabassi and J.M. Bone, (2003), *Face Recognition Vendor Test 2002: Overview and Summary.* http://www.biometricsinstitute.org/bi/FaceRecognitionVendorTest2002.pdf [accessed 1 September 2004]

Reedman, C. (2002) 'Biometrics and Law Enforcement'. Paper posted to Yahoo Biometrics Mailing List.

Stanley, J. and B. Steinhardt (2002) *Drawing a Blank: the Failure of Facial Recognition in Tampa, Florida*. Washington DC: American Civil Liberties Union.

Wood, D., E. Konvitz and K. Ball (2003) 'The constant state of emergency: surveillance after 9/11', in K. Ball and F. Webster (eds.), *The Intensification of Surveillance: Crime, Terror and Warfare in the Information Era*, London: Pluto Press.

Van der Ploeg, I. (2002) 'Biometrics and the Body and Information: Normative issies of the socio-technical coding of the body', in D. Lyon (ed.) *Surveillance as Social Sorting*. London: Routledge, 57-73.

Winner, L. (1980) 'Do Artefacts Have Politics?' *Daedalus* 109: 121-136.

Woodward, J., C.Horn, J. Gatune, and A. Thomas, (2003) 'Biometrics: A Look at Facal Recognition', Documented Briefing prepared for the Virginia State Crime Commission. http://www.rand.org

Zureik, E. (with K. Hindle) (2004 forthcoming) 'Governance, Security and Technology: The Case of Biometrics', *Studies in Political Economy*, 73: 113-137.