

Handbook of Face Recognition

(The Second Edition)

Editors: *Stan Z. Li & Anil K. Jain*

Springer-Verlag

Use the template `dedic.tex` together with the Springer document class `SVMono` for monograph-type books or `SVMult` for contributed volumes to style a quotation or a dedication at the very beginning of your book in the Springer layout

Preface

Face recognition has a large number of applications, including security, person verification, Internet communication, and computer entertainment. Although research in automatic face recognition has been conducted since the 1960s, this problem is still largely unsolved. Recent years have seen significant progress in this area owing to advances in face modeling and analysis techniques. Systems have been developed for face detection and tracking, but reliable face recognition still offers a great challenge to computer vision and pattern recognition researchers.

There are several reasons for recent increased interest in face recognition, including rising public concern for security, the need for identity verification in the digital world, and the need for face analysis and modeling techniques in multimedia data management and computer entertainment. Recent advances in automated face analysis, pattern recognition, and machine learning have made it possible to develop automatic face recognition systems to address these applications.

This book was written based on two primary motivations. The first was the need for highly reliable, accurate face recognition algorithms and systems. The second was the recent research in image and object representation and matching that is of interest to face recognition researchers.

The book is intended for practitioners and students who plan to work in face recognition or who want to become familiar with the state-of-the-art in face recognition. It also provides references for scientists and engineers working in image processing, computer vision, biometrics and security, Internet communications, computer graphics, animation, and the computer game industry. The material fits the following categories: advanced tutorial, state-of-the-art survey, and guide to current technology.

The book consists of 16 chapters, covering all the subareas and major components necessary for designing operational face recognition systems. Each chapter focuses on a specific topic or system component, introduces background information, reviews up-to-date techniques, presents results, and points out challenges and future directions.

Chapter 1 introduces face recognition processing, including major components such as face detection, tracking, alignment, and feature extraction, and it points out the technical challenges of building a face recognition system. We emphasize the importance of subspace analysis and learning, not only providing an understanding of the challenges therein but also the most suc-

successful solutions available so far. In fact, most technical chapters represent subspace learning-based techniques for various steps in face recognition.

Chapter 2 reviews face detection techniques and describes effective statistical learning methods. In particular, AdaBoost-based learning methods are described because they often achieve practical and robust solutions. Techniques for dealing with nonfrontal face detection are discussed. Results are presented to compare boosting algorithms and other factors that affect face detection performance.

Chapters 3 and 4 discuss face modeling methods for face alignment. These chapters describe methods for localizing facial components (e.g., eyes, nose, mouth) and facial outlines and for aligning facial shape and texture with the input image. Input face images may be extracted from static images or video sequences, and parameters can be extracted from these input images to describe the shape and texture of a face. These results are based largely on advances in the use of active shape models and active appearance models.

Chapters 5 and 6 cover topics related to illumination and color. Chapter 5 describes recent advances in illumination modeling for faces. The illumination invariant facial feature representation is described; this representation improves the recognition performance under varying illumination and inspires further explorations of reliable face recognition solutions. Chapter 6 deals with facial skin color modeling, which is helpful when color is used for face detection and tracking.

Chapter 7 provides a tutorial on subspace modeling and learning-based dimension reduction methods, which are fundamental to many current face recognition techniques. Whereas the collection of all images constitutes high dimensional space, images of faces reside in a subspace of that space. Facial images of an individual are in a subspace of that subspace. It is of paramount importance to discover such subspaces so as to extract effective features and construct robust classifiers.

Chapter 8 addresses problems of face tracking and recognition from a video sequence of images. The purpose is to make use of temporal constraints present in the sequence to make tracking and recognition more reliable.

Chapters 9 and 10 present methods for pose and illumination normalization and extract effective facial features under such changes. Chapter 9 describes a model for extracting illumination invariants, which were previously presented in Chapter 5. Chapter 9 also presents a subregion method for dealing with variation in pose. Chapter 10 describes a recent innovation, called Morphable Models, for generative modeling and learning of face images under changes in illumination and pose in an analysis-by-synthesis framework. This approach results in algorithms that, in a sense, generalize the alignment algorithms described in Chapters 3 and 4 to the situation where the faces are subject to large changes in illumination and pose. In this work, the three-dimensional data of faces are used during the learning phase to train the model in addition to the normal intensity or texture images.

Chapters 11 and 12 provide methods for facial expression analysis and synthesis. The analysis part, Chapter 11, automatically analyzes and recognizes facial motions and facial feature changes from visual information. The synthesis part, Chapter 12, describes techniques on three-dimensional face modeling and animation, face lighting from a single image, and facial expression synthesis. These techniques can potentially be used for face recognition with varying poses, illuminations, and facial expressions. They can also be used for human computer interfaces.

Chapter 13 reviews 27 publicly available databases for face recognition, face detection, and facial expression analysis. These databases provide a common ground for development and evaluation of algorithms for faces under variations in identity, face pose, illumination, facial expression, age, occlusion and facial hair.

Chapter 14 introduces concepts and methods for face verification and identification performance evaluation. The chapter focuses on measures and protocols used in FERET and FRVT (face recognition vendor tests). Analysis of these tests identifies advances offered by state-of-the-art technologies for face recognition, as well as the limitations of these technologies.

Chapter 15 offers psychological and neural perspectives suggesting how face recognition might go on in the human brain. Combined findings suggest an image-based representation that encodes faces relative to a global average and evaluates deviations from the average as an indication of the unique properties of individual faces.

Chapter 16 describes various face recognition applications, including face identification, security, multimedia management, and human-computer interaction. The chapter also reviews many face recognition systems and discusses related issues in applications and business.

Acknowledgments

A number of people helped in making this book a reality. Vincent Hsu, Dirk Colbry, Xiaoguang Lu, Karthik Nandakumar, and Anoop Namboodiri of Michigan State University, and Shiguang Shan, Zhenan Sun, Chenghua Xu and Jiangwei Li of the Chinese Academy of Sciences helped proofread several of the chapters. We also thank Wayne Wheeler and Ann Kostant, editors at Springer, for their suggestions and for keeping us on schedule for the production of the book. This handbook project was done partly when Stan Li was with Microsoft Research Asia.

October 2004

Stan Z. Li
Anil K. Jain

Contents

Part I Face Representation and Modeling

Part II Face Recognition Technologies

3	Privacy protection and face recognition	5
	Andrew W. Senior and Sharath Pankanti	
3.1	Introduction	5
3.2	A model for visual privacy	9
3.3	The explosion of digital imagery	10
3.4	Technology for enabling privacy	14
3.5	Systems for face privacy protection	18
3.6	Guaranteeing visual privacy	19
3.7	Conclusions	21
	References	21

Part III Applications and Issues

Part IV Perceptual Principles

List of Contributors

Firstname Surname

ABC Institute, 123 Prime Street, Daisy Town, NA 01234, USA, e-mail: smith@smith.edu

Firstname Surname

XYZ Institute, Technical University, Albert-Schweitzer-Str. 34, 1000 Berlin, Germany, e-mail:
meier@tu.edu

Part I
Face Representation and Modeling

Part II
Face Recognition Technologies

Chapter 3

Privacy protection and face recognition

Andrew W. Senior and Sharath Pankanti

Abstract *In this chapter we describe the privacy issues surrounding the proliferation of digital imagery, in surveillance video, online photo-sharing, medical records and online navigable street imagery. We highlight the growing capacity for computer systems to process, recognize and index face images and outline some of the techniques that have been used to protect privacy while maintaining usefulness of the digital imagery.*

3.1 Introduction

Faced with the ubiquity of digital imagery — from personal cameras, cellphones, surveillance and television cameras — being used by governments, corporations and individuals; together with the new technologies for analysing and exploiting such images, it is important to ask what are the risks to privacy created or exacerbated and what protections are, or could be put, in place to protect individuals' privacy. Visual privacy has been of concern since the invention of photography, but the issues are becoming critical as digital imagery becomes widespread. At the same time, image processing, computer vision and cryptography techniques are, for the first time, able to deliver technological solutions to some visual privacy problems.

In this chapter, we first examine what is meant by privacy, and visual privacy in particular. In Section 3.2 we examine some of the factors that determine visual privacy, and in Section 3.3 we summarize particular domains in which visual privacy is important and Section 3.4 describes technologies for protecting privacy in images. Finally, Section 3.5 presents three systems that have been developed for applying privacy enhancing technologies to face images.

Andrew W. Senior
Google Research, New York, NY 10011, USA e-mail: andrewseniorATgoogle.com

Sharath Pankanti
IBM Research, Yorktown Heights, NY 10598, USA e-mail: sharatATus.ibm.com

3.1.1 What is privacy?

The problem of protecting privacy is ill-posed in the sense that privacy means different things to different people, and attitudes to its protection vary from the belief that this is a right and obligation, to an assumption that anyone demanding privacy must have something to hide [12]. Brin [11] argues that at some level privacy cannot be preserved and, in “The Transparent Society”, suggests that in the face of inevitable ubiquitous surveillance, our only choice is whether to leave this surveillance in the hands of the authorities or democratize access to the surveillance mechanisms and use these same tools to “watch the watchers” and so protect the populace against abuses of the tremendous power that the surveillance apparatus affords.

Danielson [19] views the ethics of video surveillance as “a continuously modifiable practice of social practice and agreement”. What is considered acceptable or intrusive in video privacy is a result of cultural attitudes (Danielson contrasts attitudes in the UK and Canada) but also technological capability. A report of the US General Accounting Office [49] quotes the 10th Circuit Court of Appeals decision to uphold the use of surveillance cameras on a public street without a warrant on grounds that “activity a person knowingly exposes to the public is not a subject of Fourth Amendment protection, and thus, is not constitutionally protected from observation.” However technology (with capabilities such as high zooms, automatic control, relentless monitoring, night vision and long term analysis) enables surveillance systems to record and analyze much more than we might naturally believe we are “exposing to the public”. It has been argued that the “chilling” effect of video surveillance is an infringement of US first amendment rights.

3.1.2 The risks to privacy from images

Just as it is difficult to define privacy, it is difficult to determine when privacy has been intruded upon. Equally, many people are happy to trade in their intangible privacy for negligible returns [] so it is hard to evaluate privacy and privacy intrusions.

In many applications where there are privacy concerns, it is hard to point to examples where there is material effect on a person “who has done nothing wrong”, yet the feeling of disquiet remains, perhaps because everyone has done something “wrong”, whether in the personal or legal sense (speeding, parking, jaywalking...) and few people wish to live in a society where all its laws are enforced absolutely rigidly, never mind arbitrarily, and there is always the possibility that a government to which we give such powers may begin to move towards authoritarianism and apply them towards ends that we do not endorse.

There is a continuum of privacy intrusion, and our comfort point on that continuum can easily be displaced, by a small incentive or a bout of media hype. The impact of being able to search for “embarrassing photos of person x ” on the internet may be small, but few people would argue that faces of patients should not be protected.

3.1.3 *Visual privacy vs. general data privacy*

In many legal systems, video privacy falls under the legislation dealing with general data privacy and thence data protection. In the European Union, for instance, this is covered by EU directive 95/46/EC which is enacted by member states in their own legislation and came into force in March 2000. In the United Kingdom, with perhaps the densest video surveillance, the relevant legislation is the 1998 Data Protection Act (DPA) which outlines the principles of data protection, saying that data must be:

- fairly and lawfully processed;
- processed for limited purposes;
- adequate, relevant and not excessive;
- accurate;
- not kept longer than necessary;
- processed in accordance with the data subject's rights;
- secure;
- not transferred to countries without adequate protection.

The act requires all CCTV systems to be registered with the Information Commissioner, extending the 1984 Data Protection act that only required registration of CCTV systems that involved "Automatic Processing" of the data. It further gives specific requirements on proper procedure in a CCTV system in order to protect privacy:

Users of CCTV systems must prevent unauthorized access to CCTV control rooms/areas; all visitors must be authorized and recorded in the visitors log and have signed the confidentiality proforma. Operators/staff must be trained in equipment use and tape management. They should also be fully aware of the Codes of Practice and Procedures for the system. The observation of the data by a third party is to be prevented e.g. no unauthorized staff must see the CCTV monitors.

It has been estimated [33] that 80% of CCTV systems in London's business district are not compliant with the DPA.

The act also guarantees the individual's right of access to information held about them, which extends to access to CCTV recordings of the individual, with protections on the privacy of other individuals who may have been recorded at the same time.¹

The European Convention on Human Rights guarantees the individual's right to privacy (see <http://www.crimereduction.gov.uk/cctv13.htm>) and further constrains the use of video surveillance, most explicitly constraining its use by public authorities.

The Swiss Federal Data Protection Commissioner has published these guidelines: [16]

¹ "The DPA supports the right of the individual to a copy of any personal data held about them. Therefore data controllers are obliged to provide a copy of the tape if the individual can prove that they are identifiable on the tape, and they provide enough detail to locate the image (e.g. 1 hour before/after the time they believe they were captured by CCTV, their location and what identifiable features to look for). They must submit an appropriate application to the Data Controller and pay a £10 fee. However, the request can be refused if there are additional data/images on the tape relating to a third party. These additional images must be blurred or pixelated out, if shown to a third party. A good example would be a car accident where one party is attempting to claim against another. The data controller is obliged to say no to a civil request to view the tape, as consideration must be given to the other party. A request by the police is a different matter though."

When private individuals use video cameras, for example to protect individuals or prevent material damage, this is subject to the federal law of 19th June 1992 on data protection (DPL; SR 235.1) when the images filmed show identified or identifiable individuals. This applies irrespective of whether the images are stored or not. The processing of the images—such as acquisition, release, immediate or subsequent viewing or archiving—must comply with the general principles of data protection.

3.1.3.1 Why images are different

A big difference between ordinary data privacy and video data privacy is the amorphous nature of the latter, and the difficulty in processing it automatically to extract useful information. A video clip can convey negligible amounts of information (e.g. there is nobody in the street at 4 a.m.) or may contain very detailed and specific information (about times, a person's appearance, actions). Privacy is hard to define, even for explicit textual information such as name, address and social security number fields in a database, knowledge of which can be used for identity theft, fraud and the mining of copious information about the individual from other databases. It becomes much harder to assess the privacy-invasion that might result from the unstructured but potentially very rich information that could be harvested from surveillance video. A simple video of a person passing in front of a surveillance camera by itself affords little power over the individual, except in a few rare circumstances (such as proving or invalidating an alibi).

There are already strong restrictions on the use of microphones for surveillance because of the presumption of privacy of conversations, but video has been less restricted because there is an expectation of being observed when entering a public space. The UK DPA exempts from controls data where, "The information contained in the personal data has been made public as a result of steps deliberately taken by the data subject." While the act of walking along the street could be construed as deliberate steps to make one's visual appearance public, we have seen that the DPA does provide privacy safeguards for CCTV.

Until recently the unmanageability of video has limited its potential for abuse. It takes time to review video to find "interesting" excerpts, and the storage requirements have added to privacy reasons to ensure that recordings are retained for only short periods of time. Long term storage, and detailed analysis have been reserved for situations with strong economic or forensic motivation.

However, the advent of sophisticated computer algorithms to automate the extraction of data from video, mean that video is becoming as easy to mine as a queryable, machine-readable database. The data mined from an omnivident surveillance network will have a potential power that can only be guessed at today. Even in a liberal democracy and with many checks and balances, the potential for abuse is large. One possible threat is increasingly arbitrary justice—laws which are rarely enforced (such as speeding, or drug possession) end up being applied selectively and unfairly. The potential for this expands as the state becomes more able to monitor every individual's every action, though automation can (as with speed cameras) take some of the arbitrariness out of the system.

3.2 A model for visual privacy

3.2.1 *Absolute and relative Identification*

A major distinction that we [47] have drawn for privacy in surveillance systems, that significantly correlates with how likely they are to intrude on privacy, is the level of anonymity they afford. We distinguish three types of system: *Anonymous*, *Relative ID* and *Absolute ID*:

- **Anonymous** A typical CCTV system without computer augmentation is anonymous—it knows nothing about the individuals that are recorded onto the tape or presented on the monitors. While open to abuse by individuals watching the video, it does not facilitate that abuse in a systematic way.
- **Absolute ID** These systems have some method of identifying the individuals observed (such as face recognition or a badge swipe correlated with the video) and associating them with a personal record in a database. Such systems require some kind of enrollment process [9] to register the person in the database and link the personal information (such as name, social security number) with the identifying characteristic (face image or badge number), though the enrollment can happen without the knowledge or consent of the subject.
- **Relative ID** These systems can recognize people they have seen before, but have no enrollment step. Such systems can be used to collect statistics about people’s comings and goings, but do not know any individual information. A relative ID system may use weaker methods of identification (such as clothing colours) to collect short term statistics as people pass from one camera to another, but be unable to recognize people over periods of time longer than a day, or use face recognition without any external label.

Clearly, anonymity protects the individual’s privacy. An absolute ID system might, for instance be made to “Give a report on the movements of Joe Bloggs at the end of each day”. A relative ID system with a “strong identifier” can easily be converted retrospectively into an Absolute ID with a manual enrollment, but extracting Relative or Absolute ID from an Anonymous system would require storing and reprocessing the data.

3.2.2 *Threats to privacy — faces in video surveillance*

The goal of privacy protection is to prevent access to information that intrudes on an individual’s privacy, but specifying exactly what information is sensitive is difficult. We limit ourselves to considering the actions of people (including driving vehicles) in front of the camera, though certainly other visual information (e.g. documents, the presence of an object) can compromise privacy in certain circumstances. In particular we focus on preventing the identification of individuals – the major threat to individual privacy. In Table 3.1 we list a number of factors that play a role in the threat to privacy posed by an automatic video surveillance system.

The location of the system is certainly one factor. In high security environments, no privacy protection may be necessary, but in one’s home no level of video obfuscation may be considered acceptable. The person with access to the information also determines the level of privacy

intrusiveness, as shown by [30]. A person from any of the categories of Table 3.1 may be familiar with an individual observed by the system, increasing the risk of information being viewed as sensitive, but an unfamiliar person is still subject to voyeurism and prejudiced treatment. In each category the availability of each type of data must be limited as far as possible, consistent with the person's need to access information. The person seen by the camera also plays a role, being observed with some kind of informed consent (e.g. an employee); with active consent and the carrying of a privacy token [10, 45]; passively as a member of the public; or indeed as an intruder.

In preventing privacy breaches from a surveillance system, we must review the information that can be leaked, the access points to that information within the system, and the availability to different groups of people. Raw video contains much privacy-intrusive information, but much effort is required to get to that information. A keyframe may convey much less information, but if well-chosen presents that information more succinctly. An index with powerful search facilities can easily direct a user to a particular clip of video. The power to intrude on privacy is greatly enhanced if the system has the capability to identify individuals (Sec. 3.2.1). In principle, all the information stored in a networked digital system is vulnerable to hacking, but preventing such breaches is a matter of conventional information and physical security — in particular all privacy-sensitive information should always be encrypted when stored and transmitted, and strict access controls should be in place to limit access to such information, with the possibility of audit trails to know who accessed what data under what circumstances.

Scenario	Observer	Familiarity	Role of subject
High security	Law enforcement	Familiar	Member of general public
Low security e.g. workplace	System managers System operators	Unfamiliar	Employee Wearer of privacy tag
Public space	Authorized accessors		Intruder
Private space	Public Hackers Person observed		

Effort	Data type	Tools
Passive	Raw video	Summary
Opportunistic	Redacted video	Video review
Deliberate	Raw extracted data	Freeze-frame
Sophisticated.	Anonymized data Linked to an identity	Search

Table 3.1 Factors affecting video surveillance privacy protection.

3.3 The explosion of digital imagery

In this section, we review some of the domains in which the privacy of face images is important.

3.3.1 Video surveillance

CCTV deployment is undoubtedly expanding rapidly. In 2003, McCahill and Norris [33] estimated that there were more than 4 million CCTV cameras in operation in the UK. At the time, most such CCTV systems were rarely monitored and of poor quality, installed as a deterrent without much regard for practical use. Automatic processing of surveillance video, however, is bringing a new era of CCTV with constant monitoring, recording and indexing of all video signals.

Many groups around the world [26, 50, 29, 34, 8, ?] are developing software tools to automate and facilitate the task of “watching” and understanding surveillance videos. These systems also have the potential for gathering much richer information about the people being observed, as well as beginning to make judgments about their actions and behaviours, as well as aggregating this data across days, or even lifetimes. It is these systems that magnify the potential for video surveillance, taking it from an expensive, labour-intensive operation with patchy coverage and poor recall, to an efficient, automated system that observes everything in front of any of its cameras, and allows all that data to be reviewed instantly, and mined in new ways: tracking a particular person throughout the day; showing what happens at a particular time of day over a long period; looking for people or vehicles who return to a location, or reappear at related locations.

Some CCTV systems have already publicly deployed face recognition software which has the potential for identifying, and thus tracking, people as effectively as cars are recognized today (see below). Currently face recognition technology is limited to operate on relatively small databases or under good conditions with compliant subjects [39]. Further algorithms bring the potential to automatically track individuals across multiple cameras, with tireless uninterrupted monitoring, across visible and non-visible wavelengths. Such computer systems may in future be able to process many thousands of video streams — whether from cameras installed for this purpose by a single body, or preinstalled private CCTV systems, access to which is subpoenaed or coerced — resulting in blanket, *omnivalent* surveillance networks. Such techniques can also be applied to public webcams [51] which provide continuous feed of digital images from many public locations.

Yu *et al.* [57], in work supported by the US Department of Justice, describe one potential future direction for higher-level learning based on face recognition. They show how automatically captured location tracks and face images from fixed and steerable cameras can be used to learn graphs of social networks in groups of people, particularly targeted at identifying gangs and their leaders in prisons.

3.3.1.1 Camera-based sensors

While surveillance has driven the widespread deployment of cameras, low cost sensors and more sophisticated algorithms are enabling many other applications that involve the installation of cameras that will see people, but in which it is not the images themselves that are of interest, but rather the data extracted from them. These range from today’s traffic cameras and cameras that anticipate drownings in swimming pools [1] to “human aware” buildings that adjust heating, lighting [32], elevators and telephones according to the locations and activities

of people, as well as controlling physical access and assisting with speech recognition by lip-reading [41]. Many future devices and systems will have cameras installed because they are a low-cost sensor that “sees the world as humans see it”. While the purpose of these sensors is often to merely detect a single piece of information, such as the number of people at a check-out line [53], the same hardware could equally be used for surveillance and face recognition. It is impossible for the subjects of the observation to know what is happening to the data once it has left the sensor, so without suitable oversight these devices are a potential and perceived privacy intrusion.

3.3.1.2 Ambient video connections

Some of the earliest work on image-based privacy relates to the use of video for ambient awareness in media spaces, particularly for the use of video in awareness of co-workers in a remote location. Here, a worker may choose to be in constant video feed to provide a sense of co-presence. However, in times when there is no explicit face-to-face conversation the worker may wish to reveal only general information, such as presence or general location without revealing specific details that would be visible in a full-resolution video. Such a privacy protection system that uses model-based face obscuration is described in Section 3.5.2.

3.3.2 Medical Images

Medical images are also proliferating, with the advances in medical science and the lowering cost of imaging devices. Much attention has been paid to the electronic patient record and its privacy implications. The ability to copy and transmit sensitive patient records in electronic form as well as access them remotely, together with the increasing richness of the records has led to stricter controls on medical record privacy, such as the HIPAA [2] regulations in the USA. These regulate medical records as a whole, but there is a specific area of interest to the face recognition community, which is photographs of patients that show their faces. Face images may be an important component of a patient record for such areas as oral and maxillofacial surgery, dentistry, dermatology and plastic and reconstructive surgery. It is important to protect the patient from exposure of the data both through unauthorized access and use for teaching or research material. It is essential in the latter case to remove identifying information while preserving the usefulness and accuracy for the intended purpose. De-identifying faces (Section 3.5.2 is an important technique here.

3.3.3 Online photograph sharing

Photo-sharing has recently become one of the most popular activities on the internet, featuring in social networking sites like Facebook, and special photo storage and sharing sites like Flickr

or photobucket.com. Billions of such photographs are stored by such services². As traffic has grown, the affordances for labelling have become more sophisticated. Text tagging has evolved to labelled bounding boxes and to the automatic face recognition found in Picasa and Windows Live Photo Gallery. Now the task of labelling a photo album has been made much easier by software which allows the user to name a person in one picture and then propagate that label to other similar photos of the same person. These new labels can be confirmed or corrected and the face model is improved accordingly, so a large photo collection can be iteratively labelled with relatively little manual intervention.

Companies such as PolarRose are seeking to apply these techniques to social network sites, and companies such as Google are applying face recognition technologies to label photographs and videos [44] of celebrities on the web. As recognition technology improves and the quantity of labeled data increases, it seems that it is only a matter of time before all photos of you on the internet can be tagged as such and searchable.

Google goggles has the potential to carry out face recognition on photos snapped anywhere with a smart phone, but privacy concerns have prevented it from being made available, according to a spokesman. "We do have the relevant facial recognition technology at our disposal.... But we haven't implemented this on Google Goggles because we want to consider the privacy implications and how this feature might be added responsibly." [3].

3.3.4 Street View

Online services such as Google's Street View, Bing Streetside, Mapjack and Everyscape present systematically captured street-level imagery on an unprecedented scale, allowing users to explore distant places through intuitive user interfaces in their computer browser. Currently only one (undated) instant is visible on these sites, but future storage and capture enhancements may change that and ultimately blur the distinction between street view and web cam imagery. The extent of their coverage, the high image quality and the easy access have aroused concern over the effect of the imagery on privacy. Individuals are concerned about the possibility of their presence in a particular location being publicly visible, about their property being easily examinable without their knowledge and scouted by burglars. In Japan privacy concerns led to Street View imagery being recaptured with the car-mounted cameras lowered by 40cm so that the service would not present imagery taken over people's garden walls, and there has been considerable opposition to the service on privacy grounds in Switzerland and Germany [42]. Mechanisms are provided for individuals to request that particular images are not made public, but the ubiquity of faces and license plates in the imagery, and the general unease that these elicit, required an automated solution to attempt to automatically obscure all the faces and licence plates. The automatic system that Google deployed to blur faces and license plates is described in Section 3.5.1.

² More than 3 billion photos a day are uploaded to facebook [4]

3.3.5 Institutional databases

Increasingly in recent years, governments and corporations have sought to harness Information Technology to improve efficiency in their provision of services, to prevent fraud and to ensure the security of citizens. Such developments have involved collecting more information and making that information more readily available to searching and through links between databases. Silos of information, collected for an authorized process are readily accepted for the benefits they bring, but the public becomes more uneasy as such databases succumb to “function creep”, being used for purposes not originally intended, especially when several such databases are linked together to enable searches across multiple domains. Plans for Australian identity cards were rejected because of just such fears and there was a significant backlash when retired Admiral John Poindexter conceived the “Total Information Awareness” (TIA) project [40] which aimed to gather and mine large quantities of data, of all kinds, and use these to detect and track criminals and terrorists. The Orwellian potential for such a project raised an outcry that resulted in the project being renamed the *Terrorist Information Awareness* project, an epithet calculated to stifle objection in post-September 11th America.

Naturally faces are an important part of such electronic databases allowing the verification of identity for such purposes as border control and driver licensing, but registered faces provide a link between definitive, exploitable identification information such as name, address, social security number, bank accounts, immigration status, criminal record and medical history and the mass of images of individuals that is building up from other channels like surveillance and photo-sharing.³ Many authors, from Bentham [7] to the present have expressed concern about the potential for state oppression by the exercise of extensive monitoring and the projection that such monitoring is pervasive if unknowable.

The widespread use of electronic records and their portability has led to numerous cases of records being leaked or lost, and their potential value for identity theft has made them a target for theft and hacking, from within as well as outside the controlling institution. This inadvertent exposure is a major reason for strong automatic privacy protection controls such as encryption, tight access control and image redaction even in databases where normal use would not lead to privacy protection.

3.4 Technology for enabling privacy

In recent years a number of technological solutions have been proposed for the general problem of privacy protection in images and video, and for face privacy protection in particular. In this section we review the principal methods which are intervention, redaction, access control and provably secret processing.

³ Consider the case of the British fraudster John Darwin who faked his own death but was identified in a photograph on a real estate web site after buying property [54]

3.4.1 *Intervention*

Patel et al. [38] have proposed a system that prevents unauthorized photography by detecting cameras using their retroreflective properties. In their detection system, a bright infra-red source is located near a camera. If the lens of another camera is pointed toward the detector, a strong retroreflection is seen in the image, which can easily be detected automatically. When a camera is detected, a light is flashed towards it using a digital projector, spoiling any images that it may record. This unusual approach, dubbed an “anti-paparazzi” device exploits computer vision to create a privacy-protection solution where no control can be exerted over the use of the images once recorded. As well as privacy protection, the system is envisaged for copyright protection, for instance to prevent recording of new release films in cinemas.

3.4.2 *Visual privacy by redaction*

Most recent work on visual privacy protection has focused on systems that modify, or redact, visual images to mask out privacy sensitive information. Such systems typically use computer vision technology to determine privacy sensitive regions of the image, for instance tracking moving people in a video [], or detecting faces [22] in still or moving images. Given a region of interest, it is then obscured in some way to prevent subsequent viewers or algorithms from extracting privacy sensitive information. Obscuration methods that are commonly used include blurring, masking, pixellating [27] scrambling [20] or permuting pixels [13]. Recent work has investigated the limitations of some of these, for instance Gross *et al.* [24] show that pixellation and blurring may not be strong enough to defeat face recognition systems. Neustaedter *et al.* [35] have also found global blurring and other obscuration techniques to be unable to supply simultaneously both sufficient privacy and adequate information for always-on home video conferencing. Koshimizu *et al.* [31] have explored the acceptability of different obscuration and re-rendering techniques for video surveillance.

Stronger masking with greater changes to the image may have the limitation of reducing the usability of the video for its intended purpose, but re-rendering [47] may alleviate this by showing computer generated images to convey important information hidden by the redaction process. One example of this would be to obscure a person’s face in an image with a computer generated face — hiding the identity yet preserving the gaze direction and expression. Two extensions of this using face modelling are described in section 3.5.2.

One important aspect of redaction systems is reversibility. It may be desirable for some purposes to permanently destroy the privacy-intrusive information, but for others it may be desirable or necessary, perhaps for evidential reasons, to be able to reconstruct the original video.

In a privacy protecting surveillance system there are several design choices that must be made, involving the flow of protected and unprotected data; data protection methods; access control; and auditing. Following our model [48], this chapter is focused on privacy protection by the redaction that automatic processing determines regions in the image that may contain sensitive information, and that privacy protection is implemented by one of the methods mentioned above. We remain agnostic to the method — given a region of an image, any of these

methods can easily be applied and a particular system may well allow several methods for different circumstances.

When redacted information is to be presented to the user, one fundamental question is what redaction is necessary and when the redaction is to occur. In some scenarios (such as the PrivacyCam [48]) the system may need to guarantee that redaction happens at the earliest stage, and that unredacted data is never accessible. Only one level of redaction at a time is possible when such a system is a drop-in replacement for an analogue camera. However for a general surveillance system, it may be necessary to present both redacted and unredacted data to end users according to the task and their rights, and to allow different types and extents of redaction according to the circumstances.

In a distributed surveillance system, such as that shown in Figure ??, there are three principal locations through which the data must pass: the video processor, database and browser (or end-user application), each of which can be a reasonable choice for the redaction to take place:

Browser: Here the unredacted data is delivered to the client and client software carries out the redaction and presents the redacted information to the user. This scenario means that redacted data does not need to be stored and transmitted but metadata for redaction does need to be transferred with the raw data. Since the browser is the part of a system most exposed to attack, transmitting the unredacted data there is not secure.

Content management: The content management system can redact the information when requested for viewing, which will minimise storage requirements and allow complete flexibility, but involve additional processing (with the same keyframe perhaps being redacted multiple times), latency and imposes image modification requirements on the database system.

Video analytics: The video analytics system has access to the richest information about the video activity and content, and thus can have the finest control over the redaction, but committing at encoding time allows for no post-hoc flexibility. In the other two scenarios for instance, a set of people and objects could be chosen and obscured on-the-fly. Sending redacted and raw frames to the database imposes bandwidth and storage requirements.

Double redaction: Perhaps the most flexible and secure method is double redaction, in which privacy protection is applied at the earliest possible stage, and privacy-protected data flows through the system by default. Separate encrypted streams containing the private data can be transmitted in parallel to the database and to authorized end users, allowing the inversion of the privacy protection in controlled circumstances. The operating point of the detection system can even be changed continuously at display time according to a user's rights, to obscure all possible detections, or only those above a certain confidence. Figure 3.1 shows an example of such a double redaction scheme, with two levels of redaction.

Several authors [28, 37] have adopted such a double redaction system and have explored options for embedding or hiding additional data streams in the redacted video data, for instance Zhang *et al.* store the information as a watermark, Carillo *et al.* [13] use a reversible cryptographic pixel permutation process to obscure information in a manner that can be reversed, given the right key, in a manner that is robust to compression and transcoding of video. Li *et al.* transform sensitive data using the Discrete Wavelet Transform, preserving only low frequency information and hide the encrypted high-frequency information in the JPEG image.




Fig. 3.1 Double redaction: Video is decomposed into information streams which protect privacy and are recombined when needed by a sufficiently authorized user. Information is not duplicated and sensitive information is only transmitted when authorized. Optionally an unmodified (but encrypted) copy of the video may need to be securely stored to meet requirements for evidence.

3.4.3 Cryptographically secure processing

A recent development in visual image privacy protection is in the development of cryptographically secure methods for processing images. These methods establish protocols by which two parties can collaborate to process images without risk of privacy intrusion. In particular if one party owns images and another party has, for instance a face detection algorithm algorithms such as “Blind Vision”[5] allow certain algorithmic operations to be carried out by the second party on the first party’s data without the data itself being made available.

Such systems have been applied to the problem of face detection [6] and more recently to face recognition [21] and iris recognition [56]. While such provably secure techniques do introduce overhead, progress is being made in finding more efficient implementations [43].

3.4.4 Privacy policies and tokens

An important aspect of privacy protection systems is the policies for determining what data needs to be obscured for which users. Determining privacy policies is a complex area, made more so when detection of the privacy sensitive information is not reliable. Brassil [10] and Wickramasuriya *et al.* [55] explore the use of devices that can be used to claim privacy protection in public cameras, and Schiff *et al.* [45] use hats or jackets to designate individuals whose faces are to be redacted, or preserved from redaction.

3.5 Systems for face privacy protection

In this section we describe three approaches that have been specifically designed for privacy protection of face images.

3.5.1 *Google Street View*

As mentioned in Section 3.3.4, Google's Street View and similar sites present particular privacy challenges with their vast coverage of street-level public imagery. Frome *et al.* [22] describe the system that they developed to address this problem. They highlight both the scale of the problem and the challenging nature of their "in the wild" data. They use a standard sliding-window face detector that classifies each square region of each image as face or non-face. The detector is applied with two operating points, and the results are combined with a number of other features (including face colour and a 3D position estimate) using a neural network to determine a final classification of the region as face or non-face. All face detections are blurred in the final imagery that is served in Google Maps. A similar system is used to detect and blur license plates.

They describe the criteria used for choosing the redaction method, that it should be: (1) irreversible; (2) visually acceptable for both true faces and false positives; (3) makes it clear to the public that redaction has taken place, a requirement that precludes the use of re-rendered face images. To meet these requirements, the authors choose to redact the faces with Gaussian blurring and the addition of noise. The system could also be cast as an active learning problem since over time the public take-down mechanism can provide locations of missed faces which can be used to improve the system in future.

3.5.2 *De-identifying face images*

Coutaz *et al.* [18], described a system for preserving privacy in the CoMedi mediaspace which gives remote participants a sense of co-presence. The system offered shadowing and resolution lowering redaction methods [27] for privacy protection but also used Eigen space filtering for face redaction. In this technique an eigenface representation is trained using a training set of face images, and faces detected in the mediaspace video are projected into that eigenspace before re-rendering. This effectively constrains the rendered face to conform to variations seen in the training set, and obscuring other kinds of appearance differences. While this can protect privacy in some ways, such as hiding socially incorrect gestures or expressions, it is also shown to have limitations. Since the choice of the correct model and the corresponding training set is crucial. Using a mismatched model may change the identity, pose, or expression of the face.

In several papers Sweeney and collaborators [36, 23, 25, 24] have described a series of algorithms for de-identifying faces that extend this eigenface approach, tackling the problem of identity hiding. They use the Active Appearance Model [17], face representation which normalizes for pose and facial expression. Their algorithms create a weighted average of faces from

different people, such that the resulting face is not identifiable. This de-identification is termed k -same in that it results in a face whose chance of being correctly identified is no more than $\frac{1}{k}$. In their more recent work [24] they use a multifactor decomposition in their face representations that reduces blending artefacts and allows the facial expression to be preserved while hiding the face identity. They also consider the application of this in a medical video database, showing patients responses to pain, in which facial expression, but not identity is important.

3.5.3 *Blind face recognition*

As described in Section 3.4.3, a new field of research is cryptographically provable privacy-preserving signal processing, or “Blind vision”. Recent work has applied this to face detection and recognition algorithms. Erkin *et al.* [21] describe a secure implementation of an eigenface face recognition algorithm [52]. Their system performs the operations of projecting a face image onto the eigenvectors of “face space” and calculating the distances to each of the enrolled faces, without the querying party, Alice, having to reveal the query image, nor the owner of the face recognizer, Bob, having to reveal the enrolled faces. Such a Secure Multiparty Computation can be very laborious and time consuming, with a single recognition taking 10-20s, though speed-ups have been proposed [43].

3.6 Guaranteeing visual privacy

Video information processing systems are error prone. Perfect performance can not be guaranteed, even under fairly benign operating conditions, and the system makes two types of errors when separating video into streams: missed detection (of an event or object) and false alarm (triggering when the event or object is not present). We can trade these errors off against one another, choosing an *operating point* with high sensitivity that has few missed detections, but many false alarms, or one with low sensitivity that has few false alarms, but more often fails to detect real events when they occur.

The problems of imperfect video processing capability can be minimized by selecting the appropriate system operating point. The *costs* of missed detection and false alarm are significantly different, and differ in privacy protection from those for a surveillance system. Given the sensitive nature of the information, it is likely that a single missed detection may reveal personal information over extended periods of time. For example, failing to detect, and thus obscure, a face in a single frame of video could allow identity information to be displayed and thus compromise the anonymity of days of aggregated track information associated with the supposedly anonymous individual. On the other hand, an occasional false alarm (e.g. obscuring something that isn't a face) may have a limited impact on the effectiveness of the installation. The operating point can be part of the access-control structure—greater authorization allows the reduction of the false alarm rate at a higher risk of compromising privacy. Additional measures such as limiting access to freeze-frame or data export functions can also reduce the risks associated with occasional failures in the system.

Even with perfect detection, anonymity cannot be guaranteed. While face recognition is the most salient identifier in video, a number of other biometrics such as face, gait or ear shape; and weak identifiers (height, pace length, skin colour, clothing colour) can still be preserved after face redaction. Contextual information alone may be enough to uniquely identify a person even when all identifying characteristics are obscured in the video. Obscuring biometrics and weak identifiers will nevertheless reduce the potential for privacy intrusion. These privacy-protection algorithms, even when operating imperfectly, will serve the purpose of making it harder, if not impossible, to run automatic algorithms to extract privacy-intrusive information, and making abuses by human operators more difficult or costly.

3.6.1 Increasing public acceptance

The techniques described in this chapter can be considered as an optional addition to a system that displays images — one that will cost more and risk impinging on the usefulness of the system, while the privacy-protection benefits accrue to different stakeholders. We must then ask the question of why anybody would accept this extra burden, and how the tragedy of the commons can be overcome. A principal reason is likely to be through legislation. In the future, it may be required by law that CCTV systems impose privacy protection of the form that we describe. Existing legislation in some jurisdictions may already require the deployment of these techniques in domains such as surveillance as soon as they become feasible and commercially available. Without legislation, it may still be that companies and institutions deploying image-based systems choose, or are pressured (by the public, shareholders or customers), to “do the right thing” and include privacy-protecting technology in their systems. Liability for infringement of privacy may encourage such a movement.

Even when privacy protection methods are mandated, compliance and enforcement are still open to question, particularly in private systems such as medical images and surveillance. McCahill and Norris [33] estimated that nearly 80% of CCTV systems in London’s business space did not comply with current data protection legislation, which specifies privacy protection controls such as preventing unauthorized people from viewing CCTV monitors. Legislating public access to surveillance systems as proposed by Brin [11] is one solution, but that still begs the question — are there any additional video feeds that are not available for public scrutiny? A potential solution is certification and registration of systems, perhaps along the lines of the TRUSTe system that evolved for internet privacy. Vendors of video systems might invite certification of their privacy-protection system by some independent body. For purpose-built devices with a dedicated camera sensor (like PrivacyCam [48]) this would suffice. Individual surveillance installations could also be certified for compliance with installation and operating procedures, with a certification of the privacy protection offered by the surveillance site prominently displayed on the equipment and CCTV advisory notices. Such notices might include a site (or even camera) identification number and the URL or SMS number of the surveillance privacy registrar where the site can be looked up to confirm the certification of the surveillance system. Consumer complaints would invoke investigations by the registrar, and conscientious companies could invite voluntary inspections.

3.7 Conclusions

There is an explosion in the dissemination of images, in domains from photo-sharing to surveillance and medical imaging, with a corresponding increase in the potential for privacy intrusive uses of those images. Thus far, controls on the privacy intrusions that these technologies bring have been very limited and primarily legislative. We have examined how images in different domains can contain sensitive information particularly images of faces that allow individuals to be identified, and have described ways in which that information can be obscured by redaction, based on computer vision techniques to identify the regions of interest, and image processing techniques to carry out the redaction in a secure, invertible manner. Finally we have described three particular systems that have been used to apply privacy preserving techniques to face images.

References

1. Poseidon. <http://www.poseidon-tech.com/>.
2. The health insurance portability and accountability act (hipaa) privacy and security rules, 1996.
3. Privacy fears force search giant to block facial recognition application on Google goggles. *The Daily Mail Online*, Dec. 2009.
4. Facebook press room, 21 April 2010.
5. S. Avidan and M. Butman. Blind vision. In *European Conference on Computer Vision*, volume 3953, pages 1–13, 2006.
6. S. Avidan and M. Butman. Efficient methods for privacy preserving face detection. In *NIPS*, pages 57–64, 2006.
7. Jeremy Bentham. *Panopticon Letters*. London, 1787. <http://cartome.org/panopticon2.htm>.
8. J. Black and T. Ellis. Multi camera image tracking. In *International Workshop on Performance Evaluation of Tracking and Surveillance*, 2001.
9. R.M. Bolle, J.H. Connell, S. Pankanti, N.K. Ratha, and A.W. Senior. *Guide to Biometrics: Selection and Use*. Springer-Verlag, New York, 2003.
10. T. Boult, R.J. Micheals, X. Gao, and M. Eckmann. Into the woods: Visual surveillance of non-cooperative and camouflaged targets in complex outdoor settings. *Proceedings of the IEEE*, 89(10):1382–1402, October 2001.
11. J. Brassil. Using mobile communications to assert privacy from video surveillance. In *19th IEEE International Parallel and Distributed Processing Symposium*, page 290a, 2005.
12. David Brin. *The Transparent Society: Will Technology Force Us to Choose Between Privacy and Freedom*. Perseus Publishing, 1999.
13. Michael Caloyannides. Society cannot function without privacy. *IEEE Security and Privacy magazine*, May/June 2003.
14. P. Carrillo, H. Kalva, and S. Magliveras. Compression independent reversible encryption for privacy in video surveillance. In *EURASIP Journal on Information Security* [15].
15. D. Chen, Y. Chang, R. Yan, and J. Yang. Protecting personal identification in video. In Senior [46].
16. S. S. Cheung, A. Senior, and D. Kundur, editors. *EURASIP Journal on Information Security Special issue on Enhancing Privacy Protection in Multimedia Systems*. Hindawi, 2010.
17. Swiss Federal Data Protection Commissioner. *Leaflet on video surveillance by private individuals*. 3003 Bern, January 2003.
18. T. Cootes, G. Edwards, and C. Taylor. Active appearance models. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 23(6):492–7, 2001.

19. J. Coutaz, F. Bérard, E. Carraux, W. Astier, and J.L. Crowley. CoMedi: Using computer vision to support awareness and privacy in mediaspaces. In *CHI*, pages 13–14. ACM Press, 1999.
20. P. Danielson. Video surveillance for the rest of us: Proliferation, privacy, and ethics education. In *International Symposium on Technology and Society*, pages 162–167, 6–8 June 2002.
21. F. Dufaux and T. Ebrahimi. Scrambling for video surveillance with privacy. In *Proceedings of Computer Vision and Pattern Recognition*, page 160, June 2006.
22. Z. Erkin, M. Franz, J. Guajardo, S. Katzenbeisser, I Lagendijk, and T. Toft. Privacy-preserving face recognition. In *Privacy Enhancing Technologies Symposium, 2009*.
23. A. Frome, G. Cheung, A. Abdulkader, M. Zennaro, B. Wu, A. Bissacco, H. Adam, H. Neven, and L. Vincent. Large-scale privacy protection in Google Street View. In *Proceedings of Computer Vision and Pattern Recognition, 2009*.
24. R. Gross, E. Airolidi, B. Malin, and L. Sweeney. Integrating utility into face de-identification. In *Workshop on Privacy Enhancing Technologies*. CMU, 2005.
25. R. Gross, L. Sweeney, J. Cohn, F. de la Torre, and S. Baker. Face de-identification. In Senior [46].
26. R. Gross, L. Sweeney, F. de la Torre, and S. Baker. Model-based face de-identification. In *Workshop on Privacy Research in Vision*. IEEE, 2006.
27. A. Hampapur, L. Brown, J. Connell, A. Ekin, M. Lu, H. Merkl, S. Pankanti, A. Senior, and Y.L. Tian. Multi-scale tracking for smart video surveillance. *IEEE Transactions on Signal Processing*, 2005.
28. S. Hudson and I. Smith. Techniques for addressing fundamental privacy and distribution tradeoffs in awareness support systems. In *CSCW*, pages 248–257, 1996.
29. I. Ito and H. Kiya. One-time key based phase scrambling for phase-only correlation between visually protected images. In *EURASIP Journal on Information Security* [15].
30. S. Khan and M. Shah. Tracking people in presence of occlusion. In *Asian Conference on Computer Vision, 2000*.
31. T. Koshimizu, T. Toriyama, and N. Babaguchi. Factors on the sense of privacy in video surveillance. In *Proceedings of the 3rd ACM workshop on Continuous archival and retrieval of personal experiences*, pages 35–44. ACM, 2006.
32. T. Koshimizu, I. Umata, T. Toriyama, and N. Babaguchi. Psychological study for designing privacy protected video surveillance system: PriSurv. In Senior [46].
33. A.J. Lipton, J.I.W. Clark, B. Thompson, G. Myers, Z. Zhang, S. Titus, and P. Venetianer. The intelligent vision sensor: Turning video into information. In *Advanced Video and Signal-based Surveillance*. IEEE, Sept. 2007.
34. Mike McCahill and Clive Norris. *CCTV*. Perpetuity Press, 2003.
35. S. McKenna, J.S. Jabri, Z. Duran, and H. Wechsler. Tracking interacting people. In *International Conference on Face and Gesture Recognition*, pages 348–53, March 2000.
36. C. Neustaedter, S. Greenberg, and M. Boyle. Blur filtration fails to preserve privacy for home-based video conferencing. In *ACM Transactions on Computer Human Interactions*, 2006.
37. E. Newton, L. Sweeney, and B. Malin. Preserving privacy by de-identifying facial images. Technical Report CMU-CS-03-119, Carnegie Mellon University, School of Computer Science, Pittsburgh, 2003.
38. J. K. Paruchuri, S. S. Cheung, and M. W. Hail. Video data hiding for managing privacy information in surveillance systems. In *EURASIP Journal on Information Security* [15].
39. S.N. Patel, J.W. Summet, and K.N. Truong. Blindspot: Creating capture-resistant spaces. In Senior [46], pages 185–201.
40. P.J. Phillips, W.T. Scruggs, A.J. O’Toole, P.J. Flynn, K.W. Bowyer, C.L. Schott, and M. Sharpe. FRVT 2006 and ICE 2006 large-scale results. Technical Report NISTIR 7408, NIST, Gaithersburg, MD 20899, March 2006.
41. J. Poindexter. Overview of the information awareness office, August 2002.
42. G. Potamianos, C. Neti, G. Gravier, A. Garg, and A.W. Senior. Recent advances in the automatic recognition of audiovisual speech. *Proceedings of the IEEE*, 2003.
43. Associated Press. Swiss official demands shutdown of Google Street View. *New York Times*, 2009.
44. A.R. Sadeghi, T. Schneider, and I. Wehrenberg. Efficient privacy-preserving face recognition. In *12th International Conference on Information Security and Cryptology, 2010*.

45. M.E. Sargin, H. Aradhye, P. Moreno, and M. Zhao. Audiovisual celebrity recognition in unconstrained web videos. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, Apr. 2009.
46. J. Schiff, M. Meingast, D.K. Mulligan, S. Sastry, and K. Goldberg. Respectful cameras: Detecting visual markers in real-time to address privacy concerns. In Senior [46].
47. A.W. Senior, editor. *Protecting Privacy in Video Surveillance*. Springer, 2009.
48. A.W. Senior, S. Pankanti, A. Hampapur, L. Brown, Y.-L. Tian, and A. Ekin. Blinkering surveillance: Enabling video privacy through computer vision. Technical Report RC22886, IBM T.J.Watson Research Center, NY 10598, August 2003.
49. A.W. Senior, S. Pankanti, A. Hampapur, L. Brown, Y.-L. Tian, and A. Ekin. Enabling video privacy through computer vision. *IEEE Security and Privacy*, 3(5):50–7, May/June 2004.
50. R. Stana. Video surveillance. Technical Report GAO–03–748, United States General Accounting Office, June 2003.
51. C. Stauffer and W. E. L. Grimson. Learning patterns of activity using real-time tracking. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 22(8):747–757, August 2000.
52. L. Sweeney and R. Gross. Mining images in publicly-available cameras for homeland security. In *AAAI Spring Symposium on AI technologies for Homeland Security*, 2005.
53. M. Turk and A. Pentland. Eigenfaces for recognition. *Journal of Cognitive Neuro Science*, 3(1):71–86, 1991.
54. P. Venetianer, Z. Zhang, A. Scanlon, Y. Hu, and A. Lipton. Video verification of point of sale transactions. In *AVSS*, 2007.
55. M. Weaver. Canoe mystery man arrested for fraud. *Guardian*, Dec. 2007.
56. J. Wickramasuriya, M. Alhazzazi, M. Datt, S. Mehrotra, and N. Venkatasubramanian. Privacy-protecting video surveillance. In *SPIE International Symposium on Electronic Imaging*, 2005.
57. S. Ye, Y. Luo, J. Zhao, and S.-C.S. Cheung. Anonymous biometric access control. In *EURASIP Journal on Information Security* [15].
58. T. Yu, S.-N. Lim, K. Patwardhan, and N. Krahnstoever. Monitoring, recognizing and discovering social networks. In *Proceedings of Computer Vision and Pattern Recognition*, 2009.

Part III
Applications and Issues

Part IV
Perceptual Principles

