# Privacy in Camera Networks: a Technical Perspective

*Marci Lenore Meingast*
*Sameer Pai*
*Stephen Wicker*
*S. Shankar Sastry*

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2007-94
http://www.eecs.berkeley.edu/Pubs/TechRpts/2007/EECS-2007-94.html

July 30, 2007

# Privacy in Camera Networks: a Technical Perspective *

Marci Meingast†, Sameer Pai ‡, Stephen Wicker‡, and Shankar Sastry†
†Dept. of Electrical Engineering and Computer Science
University of California, Berkeley, CA
{marci,sastry}@eecs.berkeley.edu
‡ Dept. of Electrical and Computer Engineering
Cornell University, Ithaca, NY
{sameer.pai@cornell.edu, wicker@ece.cornell.edu}

## Abstract

*The prevalence of camera networks in public places has increased substantially over the last ten years. This is in part due to improved vision algorithms and the development of simple, less expensive cameras. With this increased prevalence, a number of privacy concerns have been raised. Particular concern has focused on the level of detailed information that image data provides and the real time operation of most camera networks. In response to these concerns, policies and best practices have been suggested. In this paper we look at current deployment policies and suggest technical solutions that enhance privacy awareness. We propose the use of system design measures, such as validation codes and online notices. We also look at different levels of data abstraction that can be performed using computer vision techniques, and characterize the information these methods provide about the scene, allowing practitioners to determine which data must be used for to support the network mission as well as the information that can be withheld. We provide experimental results on these measures and suggest open areas of research.*

## 1 Introduction

Camera networks are a common type of high-bandwidth network of sensors. These systems consists of a number of cameras, which have some sensing, communication and processing capabilities. The cameras observe a designated area. The cameras communicate their image data to a central receiving location or base station, and possibly to other cameras in the network, to perform some specified tasks. In the last decade, particularly since the events of September 11, 2001, the placement of these camera networks in public spaces, also referred to as Public Video Surveillance (PVS) systems, has increased. Closed circuit televisions systems (CCTV), have grown in recent years in Britain, with over 4 million cameras being used [1]. In major U.S. cities, like Chicago and New York, cameras networks are also gaining ground with thousands of cameras already installed around the cities for surveillance [2, 3]. Examples of current applications for these networks include crime detection and traffic monitoring [4–7].

Camera networks provide certain benefits. For example, the footage from the London tube stations were used to help identify the attackers from the 2005 bombings [8, 9]. However, camera networks also raise a number of privacy concerns. The detailed level of information in an image, the remote nature of the cameras, the inability in some cases for the networks to be avoided by the public, and the ease in which image data can be transferred are just some of the aspects which make camera networks a privacy a concern. It is important to consider these issues when designing and deploying a camera network. Consideration of these issues is necessary in order to gain the benefits of the systems while limiting the privacy risks.

Many privacy issues surrounding camera networks in pubic spaces have been looked at by policy and social experts. Social factors, such as the expectations people have in the place where the network will be mounted, who has the right to see the image data, and who controls the network have been discussed. However, little has been done in the technical domain towards easing privacy concerns for these networks. Measures to protect privacy should not be, and many times cannot be, designed as a separate module to be added on top of the original network. Rather, technical privacy preservations measures should be integrated into the design of the components of camera networks. Policies can help set the regulations for camera networks while technical innovations can help uphold such regulations while allowing the

*This work was supported in part by TRUST (The Team for Research in Ubiquitous Secure Technology), which receives support NSF grant CCF-0424422, and the following organizations: Cisco, ESCHER, HP, IBM, Intel, Microsoft, ORNL, Qualcomm, Pirelli, Sun and Symantec. The second author is also supported by a National Science Foundation IGERT Fellowship

user to obtaining beneficial functionality from these PVS systems.

In this paper we explore how privacy concerns addressed by policies found in existing laws and guidelines can be supported through technical mechanisms. In section 2, we provide some basic background on camera networks and examples of current PVS systems. In section 3, we discuss privacy concerns that have arisen from PVS systems and the current policies and regulation for these camera networks. In Section 4, we examine some technical solutions that can be put in place and provide experimental results. Measures, such as validation codes and notice sites to address specific policy recommendations are shown. Different levels of data abstraction and other computer vision methods that can help in privacy preservation are also discussed. We conclude in Section 6 with recommendations and further areas of research.

# 2 Background on Camera Networks

At a general level, a camera network consists of multiple vido cameras with communication links in some predefined topology working to perform some specified function. Data from the camera network is fed back to a base station and then appropriate action is taken. For example, if the base station has an interface or console that shows the images from the cameras in real time, and people are actively monitoring the interface, then they can respond to what is being shown. If there is no active monitoring of the network, then the base station may just store the image data for later use or perform calculations and send the information back to the cameras as a means of control.

There are many ways to setup a camera network based on 1) hardware choice (including the type of cameras, the base station hardware, and other processing units) 2) wireless vs. wired communication 3) the level of distributed computing vs. centralized computing in the network and 4) the algorithms running on the system. One example of a camera network is the one setup in London. Many of these cameras are wired through coaxial cable to a base station where processing occrus and are fixed cameras without pan-tilt-zoom capabilities. Another example, is the SensEye network [10]. In this network, low resolution cameras are connected to motes which then communicate with higher resolution cameras connected to laptops. The communication occurs if further processing is needed. All cameras are fixed and have no pan-tilt-zoom capabilities and all communication is done over a wireless medium in a distributed fashion. In New Orleans, over 200 wireless IP cameras have been setup to monitor housing projects, cruise terminals and the French Quarter. In 2006, New York deployed wireless video cameras for monitoring as well [11]. In Los Angeles's MacArthur Park surveillance system employs CyberDome DayNite 25X cameras equipped with removable infrared cut filters, 25x optical zoom lenses, 12x digital zoom capabilities, and 360-degree PTZ (pan tilt zoom) capabilities. The system also includes GE Storesafe DVRs (digital video recorders) and GE KTD-405 keypad controllers. DVRs are capable of saving CCTV data to a PC for up to ten weeks of recording. Some even feature motion detection technology, which means that the footage is only recorded when motion triggers the surveillance system [12].

In the following years, PVS systems may employ wireless GSM (global system for mobile communication) or 3G (third generation) connections. These wireless systems are completely portable and can deliver live images to the Internet. IP-based solutions with LAN/WAN delivery options are also being looked at despite some lags in the technology's image quality and frame rate.

Despite the differences in the system setups, similar privacy concerns apply to all these networks. What information the cameras are gathering, who has access to this information, and how the data is getting used are are just some of the issues regarding privacy. Given these concerns, policies for these networks have appeared in multiple areas.

# 3 Policy and Regulations

Specific policy/guidelines for PVS systems and the collection of images have been developed in response to privacy issues. Issues such as equal protection, anti-discrimination, and anonymity need to be considered in the design and operation of a camera network. Currently, Canada [13], Great Britain [14], Sweden [15], and Belgium all have national policies regarding PVS systems. In the United States there is model legislation proposed by The Constitution Project [16] for camera networks and Privacy Impact Assessments must be conducted pursuant to the egovt act of 2002 [17] when systems are developed or deployed by the federal government. Here we give an overview of the key features of these policies/guidelines.

**Assessment of Necessity:** This requires that the need for a camera network be assessed before one is installed. An evaluation must be done of whether there is a real problem to address that cannot be handled with some less invasive means. In Canada, this requires a full privacy impact assessment must be done. In Sweden, the required license for a PVS system, obtained from the Provincial Administrative Board (lnsstyrelsen), can be obtained after this assessment. An additional constraint from the model legislation by The Constitution Project states that any proposed camera network for surveillance must be tested for a trial period of 120 days. Then, an assessment should be done on how well the camera network is fulfilling the purpose it was designed for and whether privacy being protected.

**Public Input:** When a PVS system is proposed, consultations with the public and relevant stakeholders, including representatives of the communities that will be affected, must take place. Soliciting of public comment in areas where the cameras will be placed is important. Soliciting public comment from others, for example, an independent organization, is important to do as well.

**Notice/Signage:** The public should be informed about the PVS system with clearly written signs in the area which the surveillance takes place. Signs should list information such as who is responsible for the surveillance, including the party accountable for compliance with privacy principles, and contact information for questions related to the system. In some cases, such as in the Swedish policy, there are some exceptions for this signage in particular security-sensitive areas. These include bank premises, post offices, or shops. In these cases, the installation of camera surveillance requires only official notification but not to the public.

**Limitation of Use/Restriction on Use:** A PVS system should be designed and operated so that the privacy intrusion it creates is no greater than absolutely necessary to achieve the systems goals. There are no specific ways listed in the policies on how to do this, but suggestions are given. For example, limited use of video surveillance (e.g., for limited periods of the day) should be preferred over always-on surveillance if it will achieve substantially the same result. If cameras are adjustable by the operators, this should be restricted so that operators cannot manipulate them to overlook spaces which are not intended to be covered by the surveillance scheme. It is also suggested by the model legislation that there should be restricted use of automatic identification, automatic tracking, and pan, tilt and zoom features on cameras, though to what extent is not explicitly stated.

**Minimal Data Collection:** The minimal amount of data necessary from a PVS system should be collected with restrictions on its use and disclosure. However, this is a broad suggestion with little detail for implementation.

**Storage Limitations:** Images can only be stored for a limited time and once that retention period has expired, the images should be erased. If the images are retained for evidential purposes, they should be retained in a secure place to which access is controlled. Again, this is a broad suggestions with little detail for implementation in any of the policies.

**Audits:** All PVS systems are subject to frequent audit by an independent party to evaluate the effectiveness of the system and to identify unintended negative consequences. Audits should ensure compliance with the policy governing the system. These include ensuring that only pertinent information is collected, that the system is used only for its intended purpose, and that privacy protections in the system are respected. However, no specific auditing party is designated in these policies.

**Access/Disclosure of Data:** Access of the image data should be limited to only those who demonstrate a need. When access to or disclosure of the images is allowed, then it should be documented with the date, parties involved, and other relevant information. Furthering this idea, the British policy specifically states that if images will be disclosed to the media, the images of individuals will need to be disguised or blurred so that they are not readily identifiable. The Canadian policy also states that people whose images are recorded should be able to request access to any personal information that was collected about them. Severing certain information in a recording (including technological blurring or blocking of the identities of others) may be necessary to allow a specific individual access to his/her information.

# 4 Technical Measures

In this section we discuss ways in which camera surveillance guidelines and privacy preservation can be aided by technology. We make technological suggestions for the video surveillance technology, the systems corresponding to their use, the oversight of this technology and these systems, and efficient conduction of the public impact assessments.

## 4.1 Public Input

To accomplish such a large information dissemination and query task, we suggest the use of the Internet and a dedicated server with a database component. The database will function to store information about the system(s) in need of public input as well as the input itself. Relevant questions and comments will be posted in the form of an online survey. Access to the online survey (e.g. implemented as a style sheet) will be given to individuals after an identification procedure (e.g online authentication) to identify those for which the PVS systems will have an immediate impact. If a targeted audience is desired for input, then access will be restricted to this audience and the identification procedure will accordingly limit who gains access. For example, if the target audience are those people who live in the geographic area where the cameras will be placed, then identifying codes could be mailed to this audience and used in the identification procedure online. If a non-targeted audience is desired, then more general identification procedure could be used. For instance, a person who wishes to comment on the system may just go the the site and get an identification code, via email, which would allow them access the survey.

The responses to the survey will be transmitted through the Internet and stored in the database on the server. It will still be necessary to make the written documentation of the PVS systems available, such as flyers and signage, and hard copies of the survey available so that those without Internet

access will also be able to comment. However, the online survey will alleviate some of the cost and time needed for public comment. It can also allow for comments from a larger and possibly more diverse population since the Internet can make it easier for those not in the geographic location of the the system to add their input.

## 4.2 Assessment of Necessity

In the Constitution Project, a trial period for testing the camera network is required. This is a detailed restriction for the assessment of necessity that other policies have on PVS systems. [1]

With the validation code system, the reviewer/auditor of the PVS system provides owners/users with validation codes when approval for the system is granted. These codes are tied to a record for the user in a reviewer/auditor-side database which also stores a list of approved validation codes for the system. The codes enable the camera network control console and the camera network to be operated for a 120 day period in agreement with the suggestions in the model legislation. After this period or before a valid code is entered, access to a video surveillance console, used to operate the camera network, is restricted. Valid codes entered, along with a timestamp of the date and time of code entry, are stored into the database within the user's record. Access to the video surveillance console is only granted while the 120 day duration is not exceeded. In the next section, we give a novel design of this concept and show experimental results. We discuss the general protocol and tools necessary make this service available for all PVS systems.

## 4.3 Periodic Audit

Periodic audits require review of a PVS system in case of misuse or harm, alterations or change in purpose. In particular, the auditor is in charge conducting oversight of the system use. For on-site audits of the PVS systems, it is important to determine the position of cameras, the orientation, and if the cameras are turned on.

In order to aid the auditing process, unless the approved purpose of the system states otherwise, we recommend that a Light Emitting Diode (LED), or another visual marker, be put on each camera. The marker should be positioned in an orientation perpendicular to the direction of the camera lens, thus indicating the field of view of the camera. The LED will be used to indicate if the camera is turned on or not as well. The orientation of the LED will also assist the auditor in determining the location of the camera. The location, operation, orientation, and field of view of all the cameras in the system are all important in auditing a PVS system. This information verifies if the camera is being used as it should be and only observing the appropriate area and nothing further.

[1] In order to enforce the trial period. With a validation code system, the cameras are turned on only for a specified period of time based on the specific code used. These codes are created by a reviewer/auditing party which then distributes them to the owners of the PVS system. This limits the amount of time the auditor has to spend checking systems since the codes automatically force the cameras to turn off once the trial period has expired.

## 4.4 Limitation of Use/Restriction on Use

Policies suggest a limited use or restricted use of the camera network. This includes restrictions on use of recording, automatic identification, automatic tracking, and pan, tilt and zoom. These limitations could be done through an enforcement on hardware choice. However, there are also other methods and limitations to look at.

While digitization of data has made it easier to analyze images, allowing for better identification and tracking of objects, this is both a benefit and a drawback. As mentioned in [18], the future lies in digitization of the information received from a camera system and the further analysis of the data. However, what forms of analysis are done on the digitized data should be carefully looked at. In order to limit the use of the cameras, we suggest multiple levels of abstraction of the video data. As an example, in the next section we demonstrate how face detection, background subtraction and motion segmentation allow us to de-identify objects in the scene. For example, in background subtraction, upon detection of a foreground object, the object is covered with a mask and the camera image is sent out with the object is blocked out. Using analysis of digitized image data, information not pertinent to the task at hand, can be abstracted out leaving only the necessary information in the image. Therefore, the camera operators see only information necessary for what they are trying to detect using surveillance [19–21]. In this way, privacy of individuals being monitored can be maintained while still allowing functionality of the PVS sysem. Also, the large burden of monitoring all of the surveillance streams can be reduced when data from multiple cameras is combined through abstraction.

While policies discuss limiting automatic methods, there are some specific exceptions mentioned. For example, automatic identification and tracking pursuant to a federal counterterrorism watch list are deemed not prohibited. This implies that in certain situations it may become necessary to use automatic identification and tracking methods on digitized data. In these cases, it becomes important to hide as much identifying information as possible from those that are not targeted. Therefore, although information is digitized and processed by automatic identification and tracking algorithms in use by the PVS system, preserving the privacy of certain individuals to the greatest extent possible is still a

necessity.

This same idea could be used to notify the camera operators only in the event that individuals or targets, that must be monitored, appear in video captured by the PVS cameras. For example, in the case when there is a federal watch list, the software would automatically try to correlate individuals, captured by the camera, with the federal counterterrorism database containing images/names/IDs/licenses. If a match is found, only then would the software alert the PVS operator. Although false-positive identifications would not be completely eliminated, this system would allow for PVS operators not have undue access to the underlying data from the identification or tracking software. Furthermore this would allow for a strong filter that would allow only important information to be given to these operators. The same would go for automatic tracking methods. Images would be processed and only the tracks of those that show individuals that have a similarity to data from a watch list would be shown to operators. We discuss more about different levels of abstraction in the next section, demonstrating different methods and their tradeoffs with respect to privacy preservation.

### 4.5 Storage Limitations

A limit on how long data is stored is in many polices. A time duration based deletion of archived data can be employed through software. Although it is not necessarily an enforcement tool, when used by the PVS operator or the company which owns the PVS system, archived video will automatically be deleted after some period of time.

Furthermore, for the limitation of the duration of surveillance (i.e. 1-hour initial review limit), we suggest the use of fixed length media. If videotapes are used to store data, then the recording limit on the videotapes be at maximum 1-hour. Once the tape is done recording, it will be recorded over from the beginning [18]. If a computer is being used to record the video, then we suggest limiting the quota on the file to which the video data stream is recording. Although the file size would be dependent on parameters such as the quality of video taken and the compression amount, we recommend allowing a maximum file size such that it is only sufficient for the recording time limit (1-hour) after considering these parameters. Quotas can easily be set using most modern operating systems for the PC (UNIX, Windows, Linux, etc.). These types of mechanisms could also be verified by auditors in a fairly straightforward manner, if necessary.

### 4.6 Access/Disclosure of Data

In order to only allow the operator of the PVS system or the operating agency access to stored data, we suggest the use of proper levels of authentication prior to access. Access control mechanisms and role-based authentication should be used in order for an person to see image data. Data should be stored in an encrypted form such that only those with appropriate decryption keys can access the image data.

Furthermore, the creation and maintenance of PVS system control access logs should be done. For this we suggest regular monitoring of file access logs as well as logging and monitoring of logs which detail the person(s) who have used the PVS system controls at anytime. We note that file access logs can be created and viewed in most modern operating systems for the PC. Also, authentication logs seem sufficient to identify any individuals who have accessed the PVS system controls.

When data needs to be transferred to a third party, the operators of the PVS system should only provide the receiver with as much abstracted data as possible. Also, it is important that data be watermarked so when data is transferred to a third party, the source can be verified and the data authenticated.

## 5 Experiments

Here we present experimental results on technical solutions that have not been well explored in the literature in regards to privacy. We present the validation code protocol and discuss how it can be used in a real system to enforce policy requirements. We also look at different methods of data abstraction using computer vision techniques, and discuss the tradeoffs between these methods in regards to privacy preservation.

### 5.1 Validation Codes

We offer a proof of concept design of a validation code system for use in conjunction with a PVS systems. The system, prototyped using Microsoft Visual Studio and was tested on a system of cameras, is detailed in figure 1. In this system, the User is a client, who in our case is also the PVS system operator. The Application Server can either reside on the User side, the reviewer/auditor side or with a trusted third party. Finally, the Validation Database and Service resides with the reviewer/auditor. Each subfigure of figure 1 shows a particular use of the system:

Figure 1(a) shows a scenario in which the User is first authenticated by an Application Server. Next, the user enters a new legitimate validation code, which has not yet been used, into a validation form provided by the Application Server. The Application Server verifies the validation code against a list of legitimate authentication codes residing on the reviewer/auditor side Validation Database. Furthermore the Application Server stores the validation code and a timestamp of when the code was entered into the User record

on the reviewer/auditor side Validation Database. The Application Server then allows the user to access the Camera Operation Console.

In figure 1(b), the user is first authenticated by an Application Server. Upon authentication, the Application Server uses the Validation Database Service and determines, according to the the user record, that the user had previously entered a validation code for which the time period of PVS system access has not expired. The Application Server then allows the user to access the Camera Operation Console.

Finally, in figure 1(c), the user is, once again, first authenticated by an Application Server. Next, the user either enters an illegitimate validation code or a validation code which has already been used, and for which the time has expired, into the validation form presented by the Application Server. The Application Server verifies that either the code does not exist or that the code has already been used using the Validation Database Service. This information can be checked by examining the list of legitimate authentication codes and the user record residing on the reviewer/auditor side Validation Database. The Application Server then presents the user with the validation form again.

In general, to implement validation code system for use with camera network deployments, we recommend that the Camera Operation Console be made standard and served using an Application Server.

## 5.2 Data Abstraction for Images

Data abstraction can be used to protect the privacy of those being monitored by a camera network. There are many different ways information from image can be abstracted. Here we look at three different approaches, discuss their tradeoffs, and what is required of the PVS system in order to use them. While these approaches have been used for other purposes, such as tracking, there is little literature exploring how they can be use to anonymize image data for privacy preservation.

### 5.2.1 Track/Position Extraction

Noting the position or trajectory of an object in the scene being observed by the PVS system is one form of data abstraction. With this approach, no images are directly relayed to a PVS operator, only the tracks and positions of objects in the scene. This could be displayed, as an overlay on top of a 3D map of the environment the cameras are observing.

In order to do this kind of abstraction, there are certain requirements that must be upheld in the PVS system. First,there must be an overlap of field of views for at least every pair cameras so that image information can be combined to do 3D estimates. Second, the cameras in the system must be well calibrated, with both intrinsic and extrinsic parameters, so that measurements from a given pair of cameras can be used to get a relatively accurate estimate of 3D points. Third, there must be some way to notice that a similar item is seen in at least two cameras so that only information from matching objects in images is used to do 3D estimation.

We demonstrate this method of abstraction using 8 cameras, used as 4 stereo pairs, observe the moving object from all sides. These cameras have already been calibrated for extrinsic and intrinsic parameters. The moving object, which in this case is a truck, is partitioned from the images in each camera using background subtraction methods as shown in 2. The stereo pair from each group of cameras is used to find 3D information about the object by finding matching points on the object in each camera and using this to project to a 3D point. Next, the 3D information from each camera group is combined together to give absolute 3D coordinates. We take the centroid of the object, based on image estimates, and plot its trajectory; this shows the path of the object. This method of abstraction can be used to see how objects are moving throughout a scene, without actually showing any images. The method protects the privacy of the objects and also is an easier form of data to transmit than whole images. Good calibration is a necessity for this method and stereo pairs for 3D reconstruction should be used. Some results from this experiment can be see in Figure 2. The method accurately traces the path of the observed object.

### 5.2.2 Background Subtraction

Background subtraction can be used on images to detect objects in the scene which differ enough from what is considered the background scene. These objects are considered in the foreground of the scene. In order to preserve privacy, these foreground objects can then be de-identified by placing a mask over them and then transmitting the new image. The PVS operator would see how the masked object interacted with the environment, but would not see the masked object. Thus, this would provide a level of privacy. While the motions of the object through the environment would be seen by the PVS operator, the identity would not be exposed. Here we explore how background subtraction could be used for privacy preservation and what issues exist with it. We show results four sequences and how well background subtraction does at hiding the identifying information of objects needing privacy protection.

While there are unimodal distribution approaches for background subtraction [22–24], these really only work in static background scenarios. These methods cannot cope with multi-valued background distributions. As such, they will be prone to errors whenever those situations arise. For example, lighting changes or a change of background content causes persistent errors in the background calculation.

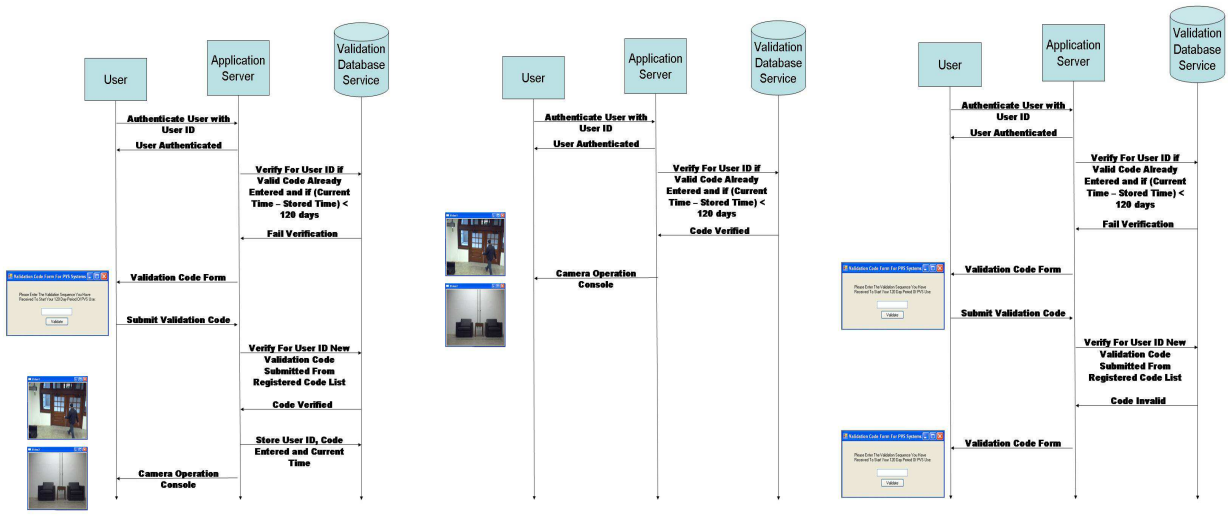


(a) A validation code sequence is entered for the first time.

(b) A valid validation code sequence has already been entered at during computer system startup.

(c) The processor and radio specification of different sensor.

Figure 1: Validation Code System

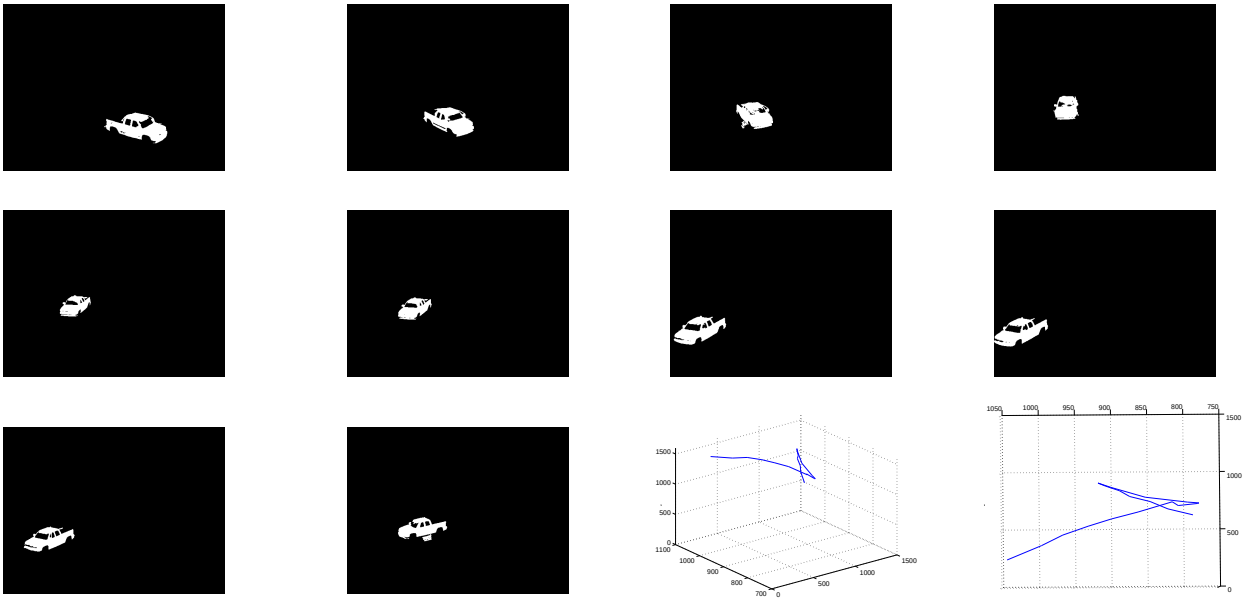

Figure 2: Track Abstraction

However, if such errors connect into relatively small blobs, they can be removed from the classified image by an adequate size filter. This type of method will work very well when there are no dynamics in the background and only the dynamics in the foreground.

The methods with a background model based on a single scalar value can guarantee adaptation to slow illumination changes, hut cannot cope with multi-valued background distributions. As such, they will be prone to errors whenever those situations arise. However, if such errors connect into relatively small blobs, they can be removed from the classified image by an adequate size filter.

As a result of the limits of unimodal approaches, many adaptive background subtraction methods that have been developed. A good overview of the different types of these approaches is given in [25]. Sophisticated adaptation meth-

ods are required to handle issues such as changes in illumination and changes of background content. Many of these adaptive methods are not too difficult to implement and run in real time and could be employed on PVS system components.

When considering the use of adaptive background subtraction methods for privacy preservation there are some aspects common to all methods that can create issues.

- **Fading to Background**: The rate of adaption will have an affect on how well the algorithm works on detecting foreground objects. If a foreground object is moving and then stops and remains still for a while, the adaption rate will make the object fade into the background. It will no longer be considered a foreground object. An example of this is shown in Figure 3.

- **Ghosts**: Ghosts are a set of connected image points detected as a foreground object due to motion, but not corresponding to any real moving objects in the scene. This can happen when moving objects are stopped for a long time and become part of the background, as shown in Figure 4. When these objects start again, a ghost is detected in the area where they were stopped. This is again affected by the rate of background update.

- **Similarity with Background**: If an object is too similar in color or intensity to what is considered background, then these objects will not be detected as foreground. When just part of a foreground object is too similar with the background model, then this portion of the object will be categorized as background, splitting the foreground object apart into multiple smaller objects. These multiple objects are too small to be considered a larger, single object and thus, become part of the background. An example of this happening is shown in Figure 5

- **Shadows**: Shadows are often detected as foreground objects. While many methods can minimize this and separate out what is a true foreground object from background, there may still be some shadow detected as a foreground object. This can creates a problem in some cases as parts of the scene that need to be observed will get hidden due to the masks put over the shadow detected as foreground. We would like the moving object detection to not classify shadows as belonging to foreground objects. Since the appearance and geometrical properties of the object can be distorted, many subsequent tasks such as object classification and the assessment of moving object position (normally considered to be the shape centroid) can be affected. Moreover, the probability of object under segmentation (where more than one object is detected as a single object) increases due to connectivity via shadows between different objects.

- **Unnecessary Objects**: The methods proposed detect any object that differs enough from the background. For example, in a scene with humans walking and moving cars as well, both the cars and the humans will be detected as foreground. Thus, if masks are placed over foreground objects, the cars will be masked. This might create a problem if too much of the scene is covered by masks. Essentially, the mask could cover objects that need privacy preservation, such as the humans, as well as other objects that might not need privacy. Figure 6 shows how cars are detected as foreground. In some cases, post-processing methods, such as human model detectors or face detectors, as discussed in the next section, could be used to prune out the foreground objects not needing privacy preservation.

Here we look at accuracy in terms of a couple items which are important when thinking about in a PVS system with privacy protection. If detected foreground objects are to be masked as a means of privacy protection, then false positives and false negatives for what is detected as foreground and what additional objects should also get privacy protection are important to look at. False negatives(FN) occur when an object which needs privacy is not detected as foreground. Thus, when the background scene is shown, the object's image also appears revealing the identity of the object. For our purposes, we define humans as the objects needing privacy. We further delineate the false negatives as the face getting shown, the body getting shown, and the whole body, including the head getting shown. There is different identity information that can be gathered given what parts of the object are shown. For example, while a face may get hidden, if the lower body of the object is shown, then a PVS operator might be able to determine the gender of the person based on the type of shoes. If high heels are shown, it is more likely that the person is a woman. If the face of the object is shown, this reveals more personal identifying information about the human, but if the body is not shown, a PVS operator might not know what action the person is taking. If the whole body is hidden, including the face, this is hides the most information from being learned about the object.

We also look at the false positive (FP) rate of areas that are labeled as foreground, but really are background portions of the scene, such as trees branches moving due to a breeze, ghosts and shadows. False positives represent areas getting masked that should not be getting masked, preventing areas of the scene from being shown to the PVS operator. While moving objects that do not necessarily need privacy protection, such as cars as dogs, also appear as fore-

ground objects, we do not consider there as false positives as we are not using any additional measures beyond background subtraction to differentiate between object classes. We will discuss more about human detection in the next section.

For our background subtraction method we use an version of the method in [26]. This method adapts to many conditions, such as shadows and ghosts, in a rapid amount of time relative to other algorithms. It is based off of the Gaussian Mixture Model statistical analysis on pixels. Both indoor and outdoor sequences were tested and at two different background adaption rates.Representative frames of the sequences are shown in Figures 7, 8, 9, and 10, and the results are shown in Tables 1 and 2.

From the results it can be seen that slowly adapting background models preserve privacy more than the fast adapting models. This especially works well in the indoor office setting where there is relatively constant lighting and a fixed background. However, in certain scenarios, such as the outdoor sequence, more false positive occur due the internal motions of fixed background objects, like trees, and lighting changes. In order to determine the best rate of adaption that would best tradeoff privacy protection with the goals of the PVS system, a trial period for the system could be used. This would allow time to learn what the most beneficial rate would be based on the number of false positives and false negatives. For example, if the PVS system is used for detecting loitering objects, then a fast adaption rate would be better. Foreground objects that were still for too long would fade into the background and this would show when an object was loitering. If just motion (traffic) patterns of objects are desired, but it is important to hide the identities of the objects, then a slower adaption rate could be used. While this would show ghosts, the ghost would be stationary and would not effect motion patterns.

#### 5.2.3 Face/Human Detection using Models

Finding humans using models is an on going area of research. Models are used to find humans or features of humans in a scene. Once found, a mask could be put over the humans to protect their privacy when the image is shown to the PVS operator. Here we specifically consider face detection as a means to protect the identities of individuals. Using face detection, faces of people would be masked out in the image, while the rest of the image would be transmitted as is. This would allow a PVS operator to observe the actions of individuals, while still providing a level of privacy to the individual by hiding the identifying features on their face.

A good listing of the state of the art in face recognition is given in [27]. While there are different methods for doing face detection, there are still many problems that all methods face. We tested multiple methods and found common issues that will effect privacy.

- **Sizing**:Faces have to be of a certain size and feature prominent enough for detectors to find it. While many methods say they can find faces larger than a 20 x 20 window, this is not a absolute threshold. As seen in Figure 11(a), this face, even though it is over around 300 pixels, can still not be found by any face detector. Thus, the identity of this individual would not be protected in a PVS system.

- **Profiles:** Faces in profiles, or turned to the side too much, are hard to detect. If they can be detected, they often lead to many false positive, as shown in Figure 11(b). While most methods are good at detecting large enough faces in profile, false positives will hide more information from the PVS operator. Privacy is still maintained, but the benefit of using face detectors to show more of the scene while maintaining privacy, may be lost.

- **Orientation**: Faces that are not close to upright in their location are hard for all face finder to detect. Thus, if a person tilts his/her head to the side too much or tilts his/her head down, the person's face will show. This can also happen when a camera is not positioned upright with respect to the scene, thus people's faces will often be rotated causing them to not be masked . This can be seen in Figure 11(c)

- **Additional Items on the Face**: Items worn on the face, such as sunglasses, and facial hair also cause issues for most facial detection algorithms. While sunglasses could help hide protect the identity of an individual, the remaining bits of the face might be identifying enough.

While face detection is still an active area of research, it is a promising method for helping to manage privacy concerns while allowing a PVS operator to see information about the actions of objects in the scene. We show a few results in Figure 12 which demonstrate this. Without second hand knowledge about objects, such as the color of the hair, or knowledge of what a person is wearing, it is hard to determine and individual's identity. However, the gender of a person or their relative age could be determined. For example in 12(c) it can be seen that there is a woman with a younger child.

One robust way of finding faces for privacy protection is by using a marker on the object (e.g. a hat on a human) [28]. However, this is not very generalizable since it might not be feasible in all setups and it is potentially intrusive for those desiring privacy protection. Instead, an inherent feature of

| False Negative Percentages | | | | | |
|---|---|---|---|---|---|
| Sequence | Rate | # Frames | FN face | FN body | FN whole body |
| Outdoor Walking then Sitting | fast | 542 | 71.4 | 66.05 | 64.21 |
| Outdoor Walking then Sitting | slow | 542 | 14.5 | 5.17 | 3.69 |
| Outdoor Sitting then walking | fast | 721 | 63.66 | 63.11 | 62.27 |
| Outdoor Sitting then walking | slow | 721 | 64.08 | 63.08 | 63.66 |
| Outdoor walking with cars | fast | 155 | 17.04 | 0 | 0 |
| Outdoor walking with cars | slow | 155 | 17.04 | 0 | 0 |
| Office Scene | fast | 77 | 44.16 | 6.49 | 5.19 |
| Office Scene | slow | 77 | 3.9 | 0 | 0 |

Table 1: False Negatives

| False Positives Percentages | | | | | |
|---|---|---|---|---|---|
| Sequence | Rate | # Frames | FP Ghost | FP Shadow | FP Background Bits |
| Outdoor Walking then Sitting | fast | 542 | 0 | 0 | 0 |
| Outdoor Walking then Sitting | slow | 542 | 0 | 0 | 0 |
| Outdoor Sitting then walking | fast | 721 | 1.11 | 0.37 | 0 |
| Outdoor Sitting then walking | slow | 721 | 38.97 | 0.37 | 0 |
| Outdoor walking with cars | fast | 155 | 0 | 0 | 3.23 |
| Outdoor walking with cars | slow | 155 | 17.04 | 0 | 16.13 |
| Office Scene | fast | 77 | 0 | 0 | 3.90 |
| Office Scene | slow | 77 | 0 | 0 | 3.90 |

Table 2: False Positives

the face or head should be used. We have shown that current face detectors have limitations that have to be resolved so that they are more robust. An interesting direction of research would be to look at the curvature of the head [29] and how this curvature changes based on the facial orientation.

# 6 Conclusion

We have shown that privacy is an important consideration in the deployment of camera networks. Policies and best practices have already started taking into account this concern. We have examined representative examples of these policies and explored the key points necessary to uphold privacy. We have shown technical methods that can be used to support these policies. In particular, we examined specific alterations we can make in camera network systems as well as examining image specific methods for privacy protection based on abstracting data. Finally, we discussed open areas of research that still remain in this domain.

# References


[1] M. McCahill and C. Norris, "From cameras to control rooms: the mediation of the image by cctv operatives," *CCTV and Social Control:The politics and practice of video surviellance-European and global perspectives*, 2004.

[2] M. Anderson, "Picture this: Aldermen caught on camera," *Chicago Sun-Times*, Jan. ,14 2006.

[3] NYCLU, "NYCLU report documents rapid proliferation of video surveillance cameras, calls for public oversight to prevent abuses," December 13, 2006. [Online]. Available: http://www.nyclu.org/whoswatching_pr_121306.html

[4] S. Kamijo, Y. Matsushita, K. Ikeuchi, and M. Sakauchi, "Traffic monitoring and accident detection at intersections," *IEEE Trans. Intelligent Transportation Systems*, vol. 1, no. 2, pp. 108–118, June 2000.

[5] O. Masoud, N. P. Papanikolopoulos, and E. Kwon, "The use of computer vision in monitoring weaving sections," *IEEE Trans. Intelligent Transportation Systems*, vol. 2, no. 1, pp. 18–25, March 2001.

[6] J. Tai, S. Tseng, C. Lin, and K. Song, "Real-time image tracking for automatic traffic monitoring and enforcement applications," *Journal of Image and Vision Computing*, vol. 22, no. 6, pp. 485–501, June 2004.

[7] R. Cucchiara, M. Piccardi, and P. Mello, "Image analysis and rule-based reasoning for a traffic monitoring system," *EEE Trans. Intelligent Transportation Systems*, vol. 1, no. 2, pp. 119–130, June 2000.

[8] K. Sullivan and K. Adam, "Jury sees tape from london bomb scare," *Washington Post*, Jan. 16 2007.

[9] E. Pape, "More watchful eyes on the continent," *Newsweek:International Edition*, Feb. 20 2006.

[10] P. Kulkarni, D. Ganesan, and P. Shenoy, "The case for multi-tier camera sensor networks," *Proceedings of the Fifteenth International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*, 2005. [Online]. Available: http://sensors.cs.umass.edu/projects/senseye/

[11] M. T. Moore, "Cities opening more video surveillance eyes," *USA Today*, July 18, 2005.

[12] L. A. P. Department, "Returning the alvarado corridor/macarthur park to the community," *Press Release*, March 10, 2004.

[13] Office of the Privacy Commissioner of Canada, "OPC guidelines for the use of video surveillance of public places by police and law enforcement authorities." [Online]. Available: http://www.privcom.gc.ca/information/guide/vs_060301_e.asp

[14] I. Commissioner, "CCTV code of practice," July 2000. [Online]. Available: http://www.ico.gov.uk/upload/documents/library/data_protection/detailed_specialist_guides/cctv_code_of_practice.pdf

[15] "Public camera surveillance act," 1998.

[16] The Constitution Project, "Model legislation to implement guidelines for public video surveillance: A guide to protecting communities and preserving civil liberties."

[17] "E-government act of 2002," Pub. L. No. 107-347, December 17 2002, §B.1.a.i. SEC. 208. PRIVACY PROVISIONS. See also, M-03-22 OMB Guidance for Implementing the Privacy Provisions of the E-Government Act of 2002, September 26, 2003.

[18] C. Norris, "Closed circuit television: a review of its development and its implications for privacy," *Paper Prepared for Department of Home Land Security Data Privacy and Integrity Advisory Committee Quarterly Meeting*, 2006.

[19] B.A. Boghossian and S.A. Velastin, "Motion-based machine vision techniques for the management of large crowds," *Proceedings of IEEE Electronics, Circuits and Systems*, 1999.

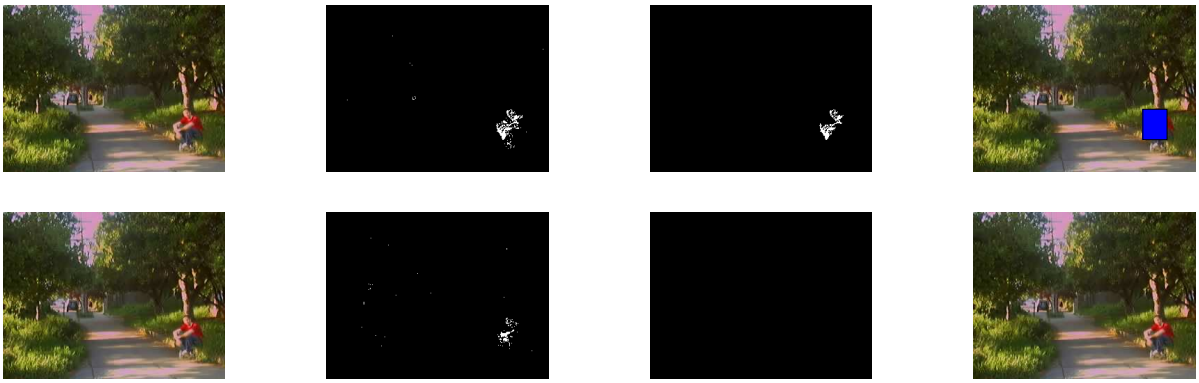

Figure 3: Object fading into the background

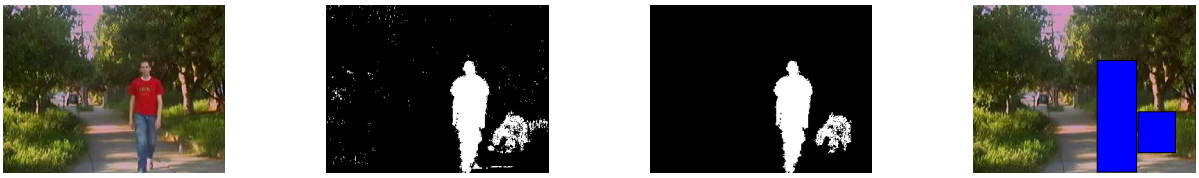

Figure 4: Ghost

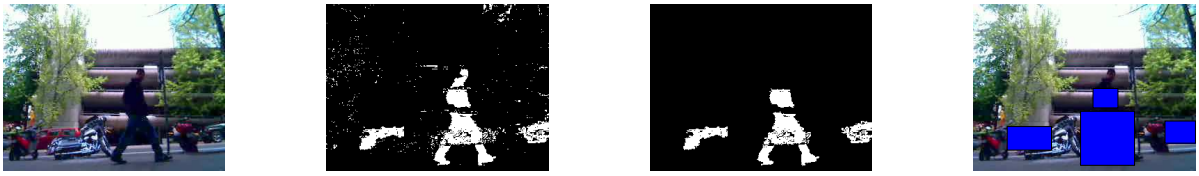

Figure 5: Similarity with background

Figure 6: Unneccesary objects

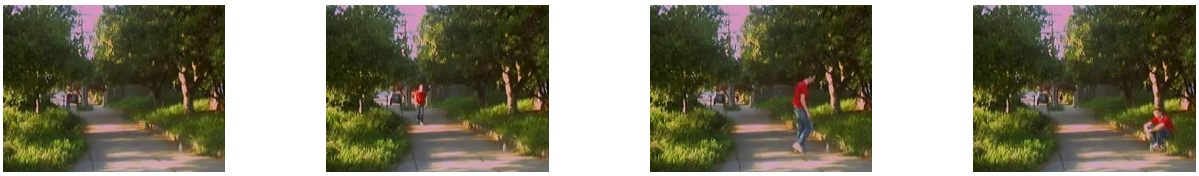

Figure 7: Outdoor walking then sitting sequence

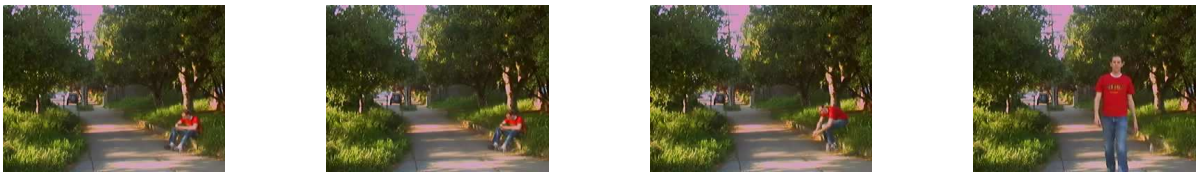

Figure 8: Outdoor sitting then walking sequence

Figure 9: Outdoor walking with cars sequence

Figure 10: Office scene sequence

(a) (b) (c)

Figure 11: Face detection results

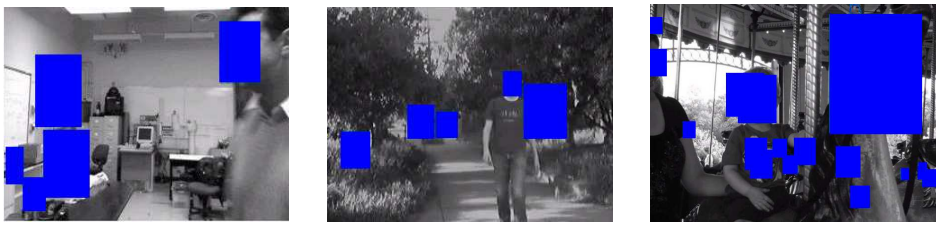

(a) (b) (c)

Figure 12: Face finding to hide identity

[20] B. Boghossian and S. Velastin, "Real-time motion detection of crowds in video signals," *High Performance Architectures for Real-Time Image Processing*, 1998.

[21] W. Chau, O. Au, and T. Chong, "Motion-based machine vision techniques for the management of large crowds," *Proceedngs IEEE Conference on Multemedia and Expo*, 2004.

[22] C. Wren, A. Azarbayejani, T. Darrell, and A. Pentland, "Pfinder: Real-time tracking of the human body," *Transactions on PAMI*, vol. 19, no. 7, pp. 780–785, 1997.

[23] I. Haritaoglu, D. Harwood, and L. Davis, "Real-timesurveillance of people and their activities," *IEEE Trans. on PAMI*, 2000.

[24] T. Horprasert, D. Harwood, and L. Davis, "Statistical approach for real-time robust background subtraction and shadow detection," *IEEE Frame-Rate ApplicationsWorkshop*, 1999.

[25] M. Piccardi, "Background subtraction techniques: a review," *IEEE International Conference on Systems, Man and Cybernetics*, 2004.

[26] Z.Zivkovic and F. der Heijden, "Efficient adaptive density estimation per image pixel for the task of background subtraction," *Pattern Recognition Letters*, vol. 27, no. 7, pp. 773–780.

[27] [Online]. Available: http://www.face-rec.org/new-papers/

[28] [Online]. Available: http://www.cs.berkeley.edu/~jschiff/RespectfulCameras/index.html

[29] B. Wu and R. Nevatia, "Detection of multiple, partially occluded humans in a single image by bayesian combination of edgelet part detectors," in *International Conference On Computer Vision*, 2005.