THE HONG KONG UNIVERSITY OF SCIENCE AND TECHNOLOGY

PHD QUALIFYING EXAMINATION SURVEY

A Survey on Visual Privacy in Ubiquitous Computing

Student:

Jiayu SHU

Supervisor:

Dr. Pan HUI

Abstract

Recent advances in camera technology have contributed to extensive use of cameras everywhere. Together with the developments in camera-related applications and online social networking/media sites, our daily life have been dramatically changed. While people are experiencing the benefits of pervasive cameras, concerns on visual privacy invasion are inevitably raised. The potential to aggregate massive visual data from multiple sources, and the possibility of inferring private information using recognition techniques, result in negative reception to the increased amount of cameras from the public. As a result, a number of technical solutions have been proposed to protect visual privacy. In this survey, we introduce concepts about visual privacy and formalize three violations of visual privacy. Then we summarize findings from attitudes studies and investigate reasons behind growing visual privacy concerns in ubiquitous computing environment. Based on a general workflow and four key challenges we have identified, we classify latest visual privacy protection systems/frameworks according to the scenario they apply to. Next, we introduce and compare privacy requirement expression, association, and protection methods, which leverage computer vision techniques, communication technologies, and cryptography algorithms. Finally, we discuss challenges and opportunities of visual privacy that guide future research directions.

Contents

List of Figures

List of Tables

1 Introduction

Since the first camera phone being sold in 2000, people have witnessed the proliferation of built-in cameras on various mobile and wearable devices. The developments in cameras these years enable cameras with smaller size and higher resolution to be equipped on most of the mobile and wearable devices as shown in Figure 1. It is estimated that by 2018, over a third of the world's population is projected to own a smartphone, which is almost 2.53 billion smartphone users in the world [51]. On the other hand, revenue from wearable device sales are forecast to amount to around 38.84 billion U.S. dollars by 2018 [52]. More wearable devices will equip with built-in cameras, with more powerful characteristics such as automatic photo taken by a wink or a voice command. In addition to wearable cameras and smart glasses, smart watches and even smart contact lens are starting to embed cameras [74].

As a result, a large number of camera-related applications emerge in the mobile and wearable application markets. Camera not only helps record memorable moments, but has also been playing a more significant role in perceiving surroundings in the physical world. For example, mobile Augmented Reality (MAR) applications rely on cameras to sense the environment and then overlay digital information onto the real world. Lifelogging and continuous sensing systems capture audio-visual and other sensory data, with applications ranging from personal archival, journalism, medicine, to law enforcement. This trend of leveraging mobile and wearable cameras for more functionalities will keep growing, which in turn promotes the camera usage.

Meanwhile, online social networking and media sites are becoming extremely popular these years. These online communication channels are dedicated to community-based input, interaction, content-sharing and collaboration. People can post their personal information on a variety of social networking sites, including text, image, and video. For example, every 60 seconds on Facebook, there are 136000 photos uploaded, 510000 comments posted, and 293000 statuses updated [78]. Instagram, a photo and video sharing platform, has about 95 million photos uploaded per day [7]. Youtube, another one of the most popular social media sites, has 300 hours of videos uploaded every minute, and almost 5 billion videos watched every single day [27].

However, the developments of the camera technology, camera-related applications market, and current online environment, together present a significant threat to visual

Figure 1: Ubiquitous cameras. From left to right: Google Glass, LifeLogger, Narrative Clip 2, and HTC RE Camera.

privacy. People raise privacy concerns mainly because they are not aware of photograph action in the vicinity most of the time, therefore do not know what is captured. Cameras, especially those on wearable devices, have the capability of "always-on" and the feature of "non-overt act", which differs from traditional hand-held cameras. Moreover, people being captured cannot control where the image or video will appear. The media data with geo-tags may end up anywhere online, being viewed and commented by any Internet user. What makes it worse is that, huge multimedia data collected online can be used to infer personal or sensitive information. It is said "A picture is worth a thousand words," therefore, an image or a video can disclosure much more information than people may have realized.

In consequence, people's attitudes towards increasing amount of devices with built-in cameras, especially those wearables, are not completely positive. A representative example is Google Glass, which has been questioned by US Congressional Bi-Partisan Caucus and Data Protection Commissioners around the world, concerning privacy risks to the public, as well as to its users [1, 2]. They raised questions such as "How does Google plan to prevent Google Glass from unintentionally collecting data about non-users without consent?" and "Are product life-cycle guidelines and frameworks, such as privacy by Design, being implemented in connection with its design and commercialization?" Besides, there have been multiple reports of people being attacked for wearing Google Glass. For example, a woman named Sarah was attacked at a bar for wearing Google Glass, being suspected of recording others.

In fact, the society has taken action to address visual privacy issues caused by unauthorized or unnoticed visual information collection and sharing, by both legal and technical means. There are prohibition signs of camera use in some places, in order to remind people of turning cameras off. Devices with recording capability are banned in certain situations [18]. Some countries even have rules that sound or visual cues must be made

to show the recording is in action. Besides, users of the device can directly deny the permission of camera use from applications, to prevent possible visual privacy leaks. A more sophisticated way based on it is controlling applications' access to raw camera data, exposing part of the visual information [53], or only high-level objects, such as a skeleton or a face to applications [39].

Nevertheless, only efforts on recorder's side are far from enough to establish an efficient and effective visual privacy protection ecosystem. Recorders may not notice prohibition signs or forget to turning off background shooting. It is more often the case that recorders do not know the privacy intentions of people around or the privacy requirements of the objects and places, therefore, fail to protect visual information. As a result, the trend has been gradually shifting to involving the subjects being recorded in the loop, giving the privacy control back to them. For example, people can report photos and videos shared online that they believe to be in violation of their privacy rights, which is adopted by some mainstream social networking sites like Facebook [4]. People can also wear colored hats [59], markers such as QR code and tags [57, 10, 17], and show specific gestures [42, 63], to express their unwillingness of being photographed and help locate where they are. In addition to explicitly expressing their privacy requests, people can upload their privacy requirements to a cloud server, which will be responsible for visual privacy protection whenever a photo is captured at the same location [80]. People can also broadcast their privacy choices to be received by recorders using their smartphones [6]. Similarly, cameras detect WiFi access points or Bluetooth signals will get informed of privacy policies specified at certain places [57]. A more straightforward solution prevents unwanted recording at places by actively seeking cameras in the environment and directing a pulsing light at their lens to distort any imagery the camera records [50].

In this survey, we aim to provide an overview of visual privacy in ubiquitous computing, including possible invasion of privacy and latest technical solutions. The rest of the survey is organized as follows: in Section 2 we introduce the concept of visual privacy, and provide examples of visual privacy intrusion in the real world; in Section 3 we present studies on people's visual privacy concerns, and discuss reasons behind growing visual privacy concerns; in Section 4 we identify four key challenges that any visual privacy protection system should deal with, and classify existing technical solutions based on their scenarios; in Section 5 we describe how privacy requirements are expressed and associated using available technologies; in Section 6 we introduce commonly used pro-

tection methods and two way of protection enforcement; in Section 7 we summarize challenges and opportunities for visual privacy protection in ubiquitous computing environment; and finally, in Section 8 we conclude the paper and discuss future research work.

2 Visual Privacy Issues

In June 2013, Edward Snowden and the US National Security Agency (NSA) were pushed to the center of the storm, because of a global violation of data privacy. According to Snowden's leaks, NSA collected telephone records from tens of millions of Americans. NSA also accessed and collected data through back doors into 9 US Internet companies, including Facebook, Google, Microsoft and Yahoo, under a surveillance program called Prism [66]. The information from the Internet and phone use was then sifted and analyzed by the British and US intelligence agencies.

This incident has changed how people view their privacy. Not only does it reduce people's trust in the governments, but also raises their awareness of privacy protection. When asking people what they think about privacy, most of them people are likely to come up with massive data breaches, online social networks, wearable technologies, etc. But what does the privacy we are talking about actually mean? What are the differences regarding privacy when comparing the present with the past situations? What does the violation of privacy mean to us? And how will our daily life be affected?

In this section, we first introduce the terminology about privacy. We then give real life privacy invasion examples, which prove that privacy issues are not out of thin air, and such violation may happen on any of us.

2.1 What is Visual Privacy

According to the Merriam-Webster dictionary, ***privacy*** is the freedom from unauthorized intrusion. To be more specific, privacy is the freedom from interference, the state of being alone, and the right to keep personal matters and relationships secret. As a subclass of privacy, ***information privacy***, also called ***data privacy***, is the right to have control over how your personal information is collected and used [3].

In this way, we can define ***visual privacy*** as the right to have control over how personal visual information is collected and used. Here, the personal visual information is regarded as information that can reveal personal information in images or videos. For example, visual information such as face, clothes, silhouette, buildings, can be used to infer the identity of a person, the relationship between the person and other objects. Figure 2 illustrates the relationship of privacy, information privacy, and visual privacy.

Based on this, we define ***violation of visual privacy*** as the act of collecting, sharing,

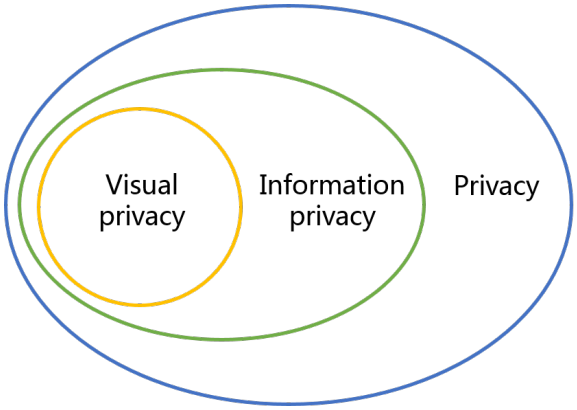

Figure 2: The classification of privacy.

and disseminating others' visual information for unauthorized use. We further formalize three categories of visual privacy invasion:

- **Uninformed photography:** The act of photographing or filming individuals or things without getting permissions (except those enforced by laws).
- **Hacking:** The act of accessing to personal or private visual data with no authorization.
- **Analysis with ulterior motives:** The act of analyzing visual data to get personal or sensitive information for unauthorized use.

The first and second category underline the unauthorized collection of visual information, and the third one emphasizes the use of visual information. It should also be noted that though normal daily photography is not intended to violate others' privacy, it may still lead to privacy infringement due to its subsequence.

2.2 Invasion of Privacy in the Real World

People's privacy concerns about visual information in the present era is not unwarranted. We present examples of visual privacy violations in the real world.

Uniformed Photography

Concerns on violation of privacy have been discussed since Google launched its Street View project in 2007. A couple in Pittsburgh sued Google because they found their home was clearly visible on the map, causing them "mental suffering" and diluting their home

value. Google Street View is banned in India, as security agencies and defence department object to the collection of data by Google's cars, mainly due to the security of sensitive defence installations. A mother found Google Street View published naked images of her two kinds outside their house playing in a padding pool. These stories are only the tip of the iceberg. Images with sensitive information are likely to appear on Google Maps. Maybe one day, your friends will tell you that you are found hanging out with a "secret" friend, which makes you quite unease and feel invaded.

Now in order to protect privacy and anonymity, Google Street View has blurred human faces and car license plates on images they capture, though it has not achieved 100% accuracy [28]. More importantly, blurring only the faces of people cannot solve the privacy issues. People can still be recognized by acquaintances according to their clothes or silhouettes. And in addition to car plates and faces, there are other personal and sensitive information can be discovered from images.

Hacking

Hacking may happen locally on the device that take images and record videos. Mobile malware can allow cyber criminals to intercept messages, monitor calls, and steal personal information. Last year the group FireEye discovered 11 malware apps being used on iPhones that gathered users' sensitive information and send it to a remote server, including text messages, Skype calls, contacts and photos. In addition to malware apps that are especially designed to launch a cyber-attack, apps installed on devices usually require permissions to access user information. For example, camera access permissions allow the app to use the camera at any time. SD card access permissions allow the app to read, and modify or delete the contents of your SD card, such as photos from a user's photo library.

However, it is difficult to judge the potential damage to a smartphone user that could be caused by access to any particular piece of personal or phone-collected information, although the scary fact is that nearly 30% of all free mobile apps capture and sell your contacts, text messages, Web browsing histories, and photos. It is more often the case that users are not be aware of that their personal information is collected by apps. But such privacy violation usually can be told from its consequences. For instance, after a malicious app is installed, it will collect information on your smartphone and upload to the ad server. Then you will see ad banners regardless of what you are doing on

the smartphone. Victims have even reported seeing ads pop up when staying on the Android home screen. In addition to ads, with opportunistic use of camera and other sensors on the smartphone, it is possible to construct three dimensional models of indoor environments [68]. Remote burglars can then study the environment for bad purposes.

Hacking can also happen on the cloud, or remote server. In 2014, hackers obtained the images from Apple's cloud services iCloud. They posted a collection of almost 500 private pictures of many celebrities, mostly women, with many containing nudity. The image were later disseminated by others on websites. Such leak is a massive invasion of privacy, and Apple took additional steps to protect the privacy and security of iCloud users very quickly. These days, many cameras installed at home and connected to the network for life-logging in China are reported to be hacked. Its consequences are too horrible to imagine, as the most sensitive and personal information is at the risk of being exposed.

Analysis with ulterior motives

Images can videos can provide much more information than plain text. With data gathered from various places, additional sensitive information can be inferred, such as age, address, working experience, social relationships, and even sexual orientation, credit scores. Researchers have investigated the feasibility of using publicly available online social network data to identity individuals both online and offline [5]. With off-the-shelf face recognition technology, they *1)* re-identified profiles on a US dating site with images from online social network profiles; *2)* identified students strolling on campus using publicly available images from Facebook; and *3)* inferred personal and sensitive information including interests, demographic information, and Social Security Numbers (SSN). In another work, researchers have developed a technique that can determine who is dating whom, given a large number of pictures shared on a social network [62].

3 People's Privacy Concerns

In this section, we present people's attitudes on visual privacy issues, from the perspectives of both users and bystanders. Then we discuss reasons that increase people's privacy and cause visual privacy risks.

3.1 Attitude Studies

Theoretically, visual privacy risks exist whenever individuals' identifiable information or other sensitive visual information is collected, stored, disseminated, in the form of images or videos, and finally interpreted or analyzed for other purposes.

In order to figure out people's perception of pervasive video recording, whether they have visual privacy concerns, to which degree, and to understand people's attitudes on visual privacy issues concerning ubiquitous cameras, especially built-in cameras on wearable devices (e.g., Google Glass), researchers have conducted some surveys and empirical studies hat investigate diverse aspects of visual privacy, in the scenario of pervasive video-recording such as closed-circuit television (CCTV) [48], daily image capture [6], AR applications [20], and life-logging systems [12, 36, 35]. Their findings related to ubiquitous computing can be summarized as follows:

- Generally, people are concerned about the amount of personal data being collected about them [48].
- People tend to be more restrictive when being captured in venues such as beaches, gyms, and hospitals [6].
- People are less comfortable when they are captured with strangers in a social situation, and when images are shared online [6].
- A number of factors affect people's feelings towards being recorded, including where they are, what they are doing, and how they feel about the recorder [20].
- The identities of people appearing in a photo, the context of the situation, and the appearance of the people determines if a photo is sensitive [20].
- People are highly concerned about their data being accessed by the "wrong" people, and being used for unauthorized purposes [48].

- People wearing life-logging cameras actually care about the privacy of bystanders, and actively try to delete or not share photos of them [36].

3.2 Reasons

The above findings are not surprising. It is understandable that people have concerned about their visual privacy for a long time since the emergence of CCTVs, digital cameras, and camera phones. Some of them even get used to being recorded when they are in public. But what makes people more serious about visual privacy today is the convergence of various technologies: camera technologies, online social networks, and recognition techniques.

First, camera technologies, in both hardware and software, enable cameras to be equipped on most of mobile and wearable devices. Cameras with smaller size, higher resolution, and more compact structure appear on more wearable devices these years. In addition to wearable cameras designed especially for taking images and recording videos, wearable devices such as smart glasses make applications like mobile AR more natural and user-friendly, which also promotes the development of wearable camera technology. Specifically, these hand-free wearable devices have the characteristics of always-on and non-overt act when taking images or recording videos. The situation is different from digital cameras or smartphones, where there are conspicuous actions that indicate the camera use. As a result, people are more likely to be captured without awareness in public places.

Second, online social networking and media sites serve as platforms that gather massive data. The appearance of online social networks have dramatically changed how people connect with others. It has become a habit for people to share their status, thoughts, and events with others, on a large number of websites. Therefore, a huge number of images and videos are available online. People on the Internet can collect a lot of data without much effort.

Finally, perception and understanding techniques can be used as tools for analyzing visual information. Advanced recognition techniques can link an image to a specific place, individual, or an event, thus makes searchable what was not considered searchable before. For example, face recognition technique can associate an image with a person. As a result, other information such as online profiles or images belonging to this person can be found further.

In summary, developments in the camera technology and its prospectives promote the camera usage, thus increasing the possibility that personal visual information will be collected without awareness. Later, images and videos being shared on social networks allow other people to collect data easily. The perception and understanding techniques then serve as tools for deeper exploitation of sensitive information. Whether being realized by people or not, it is various techniques together that pose threats to people's visual privacy. The technologies challenge people's expectations of privacy and anonymity in both the physical and digital world.

According to people's visual privacy concerns and the risks the privacy disclosure poses, we can therefore infer that the society is eager for effective visual privacy protection mechanisms to be provided and enforced.

4 Visual Privacy Protection Scenario

While visual privacy is a legal issue, visual privacy protection is a technical one essentially. To meet people's demands of visual privacy in ubiquitous computing environments, privacy and security, mobile and ubiquitous computing, image processing communities have proposed some methods to protect visual privacy.

Figure 3 depicts the main steps from cameras capturing visual information, to enforcing protection policies on images and videos. In general, any visual privacy protection framework should address the following key challenges: *1)* who or what should be protected, or how can the object that seeks protection express their privacy requirements; *2)* how to find the objects associated with the requirements on the image or video; *3)* how to protect privacy while keeping perceptual utility; and *4)* how and where to enforce privacy protection (e.g., before the data is accessed by others). Only if these challenges are addressed, can personal and sensitive information be protected.

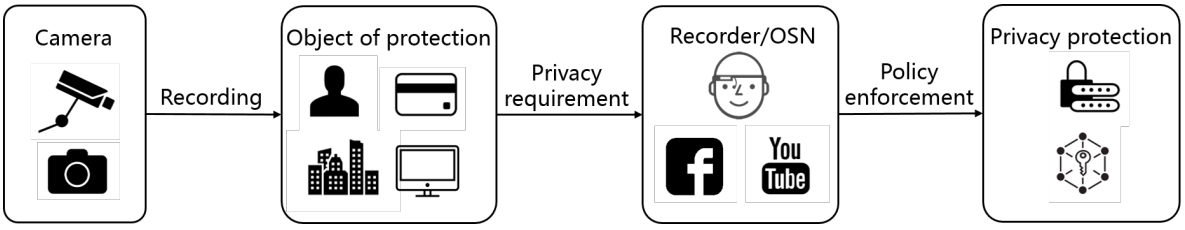

Figure 3: The flow of visual information collection and privacy protection.

However, it is impossible to find a perfect solution that satisfies everyone in all scenarios. People's understandings of privacy are different. Even one individual's expectations of privacy may be different in various scenarios. Moreover, as cameras in various scenarios usually capture or record different contents, the object that requires protection may be different, thus the protection methods vary according to the scenario.

Therefore, in this section, we classify representative visual privacy protection systems and frameworks based on scenarios of privacy issues in which these methods aim to solve. In general, visual privacy protection frameworks and systems are mainly trying to address privacy issues in three scenarios: 1) video surveillance systems, 2) traditional camera recording, and 3) perceptual applications. We summarize existing visual privacy protection systems and frameworks in Table 3.

4.1 Video Surveillance Systems

The amount of surveillance cameras has increased dramatically in recent years, especially in urban areas, to monitor surroundings and deter crimes. With pervasive CCTV cameras, individuals are observed in streets, subway stations, shopping malls, and office buildings.

However, the ubiquity of surveillance cameras, linked with the potential to aggregate information over thousands of cameras and many other networked information sources, such as security and police databases, and the power to automatically analyze and understand the videos using advanced techniques, have driven fears about the loss of privacy from the public. The privacy concerns are even worse for people who work in surveillance areas, as they are being monitored continuously. Thus visual privacy first gained attention due to the large amount of closed-circuit television (CCTV) cameras [14, 61].

The contradiction here is the trade-off between privacy and security, where the privacy issues are mainly concerned with the presence of the person and what the person is doing. In this sense, the video surveillance system can be "smarter", by providing multiple levels of access to the video data based on viewers' authorization levels [81, 60, 17]. For example, the general public can only see the part of or the modified video content, while special authorities such as the police, can observe all the information happened in the monitored space.

On the other hand, video surveillance systems can be more "privacy-enhanced", by allowing people to asserting privacy proactively and performing appropriate protection actions, before disseminating the captured videos to untrusted parties for whatever purposes [11, 59]. For instance, people can carry a privacy-enhancing device or wear a special marker, so that surveillance systems can know individuals' privacy requirements.

4.2 Personal Camera Recording

Personal camera recording refers to the most ordinary and traditional camera use. Devices including digital cameras, smartphones, wearable cameras, smart glasses, and even smart contact lens, all have the capability of taking pictures and recording videos. With increased popularity of online social networking and media sites, people usually put their captured photos and recorded videos online, sharing with others.

While people may not intend to capture bystanders or violate their privacy, the pri-

vacy risk here is that the "secondary use" of media data is unexpected, and the efforts to gather large amount of data online is extremely low. More importantly, people are not aware that they are being recorded most of the time. It differs from surveillance cameras in public places, which are usually expected by the public. As a result, more privacy concerns are raised in recent years, resulting from the explosion of personal cameras.

Though we cannot stop the existence of cameras, methods that directly prevent cameras from recording anything by first finding nearby cameras can work [70]. It does not require any effort from recorders' sides, and the protection is very straightforward. However, it can only work in some situations, as an additional device is needed to detect nearby cameras all the time.

On the other hand, it is believed that new social norms will finally evolve with pervasive personal cameras. Recorders are expected to respect bystanders' privacy preferences, as they may also be bystanders captured by others sometimes. Besides, the reason that recorders do not want to get involved into possible legal issues caused by photo or video they captured and disseminated, will encourage them to follow the social norms if a convenient privacy-. Based on this assumption, a number of methods that rely on collaboration between both recorders and bystanders have been proposed in recent years. For example, any privacy stakeholder (e.g., any person that is present when the recording is made) can decide if the recording is an invasion of privacy. No event will be recorded without the consent of all persons present, and no recording will be released without the consent of all persons present [34]. This method provides thorough but coarse privacy protection, as everything captured is considered as private. In a more practical way, recorders will protect bystanders' privacy, according to the privacy preferences proactively expressed from bystanders [10, 80, 6, 63]. These methods release privacy protection burdens on recorder's side by involving bystanders in the loop, which are more efficient and effective, though they are still far from real deployment.

4.3 Perceptual Applications

Perceptual applications sense the environment, sometimes interact with users via cameras and other sensors, in order to get voice command, gesture input, body movement, etc. Applications such as mobile AR, continuous sensing, and life-logging systems belong to this category.

Perceptual applications differ from personal camera recording in that these applica-

tions usually run at homes or public areas continuously, while users may even forget that applications are running. Moreover, information will be over-collected, which means the data collected is more than what is necessary. The over-collected data are highly likely to contain sensitive information, such as credit card numbers, license plates, computer monitors, etc. that accidentally end up in their field of vision [12, 5, 68, 56]. As a result, the privacy issues in perceptual applications are not only related to bystanders that may be captured, but also the users of the application.

To address privacy issues concerned with users, there are operating system-level protection layers proposed for untrusted perceptual applications for trusted devices. Applications can only have access to high level information instead of raw video feeds, which is sufficient for their functionalities, such as the skeleton or the face region [40, 39]. Or users can define the secure regions that applications can have access to [53]. These approaches are plausible for some applications. But before it is becomes possible to be deployed, fine-grained application access permissions are supposed to be proposed and followed by application developers. Based on the permission asked from applications, users then can know how to protect their privacy.

Besides, as life-logging systems are likely to capture some sensitive content in certain places, methods are proposed to assist cameras in deciding if images should be protected or cameras should be turned off [16, 67, 42]. The principles of these methods are similar to those proposed for personal camera recording, therefore they are faced with same challenges.

5 Privacy Requirement Expression and Association

5.1 Requirement Expression

How can the object that seeks privacy protection express their privacy requirements is the first challenge that a visual privacy protection system should address. The simplest and the most basic requirement is whether the person or the object wants to be recorded. It is also possible to include more information, like how to protect the privacy if the object is captured in the image or video. If recorders or cameras can "see" or "hear" such requirements, visual privacy protection measures can be performed efficiently.

In this subsection, we present three adopted ways to express privacy requirement: 1) via vision channels, 2) via wireless communication channels, and 3) with the help of a server/cloud. Except the last one which involves a central server, the other two are in an ad-hoc way. Although each requirement expression manner has its limitations, they present possibilities of establishing a healthy visual privacy protection ecosystem by solving the first key challenge.

5.1.1 Visual Indicators

Visual indicators are visual clues, such as markers and tags, that encode privacy requirements for objects that require protection. Examples of visual indicators used are QR code [10, 57], colorful hints like hats [59], especially designed markers (e.g., dotted rectange with solid rectangle) [53], and hand gestures [42, 63] for individuals to express their unwillingness to be captured. Figure 4 shows the visual indicators proposed in different systems.

(a) QR code (b) hat (c) marker (d) gesture

Figure 4: Multiple visual indicators.

Visual indicators are technically feasible for deployment in real life, as existing object detection methods are able to detect them. However, indicators they are limited in their

specific scenarios. First, few people would like to wear them in public places. Moreover, others can explicitly tell people's privacy intentions according to conspicuous markers, which results in another privacy leakage. Finally, privacy concerns vary widely among individuals and change from time to time, following patterns which cannot be conveyed by static visual markers.

5.1.2 Wireless Communication

Objects can also express their privacy requirements by leveraging the short-range wireless capability available on devices, such as Blurtooth or WiFi. Using device installed in the building or in the room, the environment can notify cameras that the space does not allow photography [57]. With the device carried by individuals, people can inform nearby cameras that they do not want to be captured [34, 11, 57, 6]. Figure 5 from [34] illustrates the general idea, that devices discover and exchange short messages with other devices within the recording area.

Privacy requirement expression via wireless communication channels can make privacy intentions invisible to people, while cameras can still get informed. Compared with visual indicators, this method is more convenient and flexible. The limitation, however, is that an additional device with wireless communication capability is required, either to be carried with individuals, or installed in advance. Moreover, the camera should also be equipped with wireless communication module in order to detect wireless signals. Besides, protection is opportunistic as transmitted signal may not be received by cameras, especially in those Also, as messages transmitted in wireless channels can be received by any device in the vicinity, which may contain some sensitive information itself. How to prevent privacy leak in this stage should also be considered.

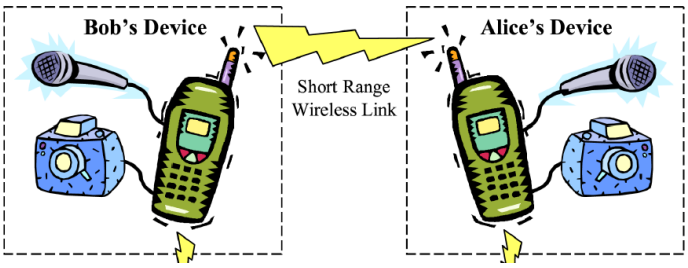

Figure 5: Two devices communicate to get privacy requirements and protect the content.

5.1.3 Remote Server/Cloud

In addition to ad hoc requirement expression, central privacy management is another option. In this way, a remote server or the cloud will get users' privacy preferences, and process images accordingly. In [80], a cloud as depicted in Figure 6 takes charge of location, communication, and computation services. A user creates his/her profile to express his/her privacy requirement (e.g., he/she can can ask recorders to make him/her invisible in the image once captured).

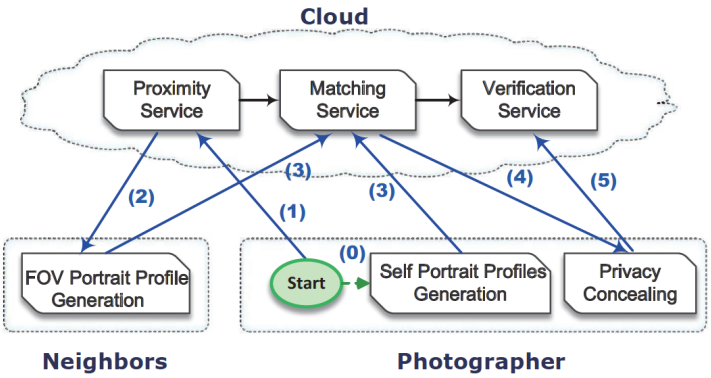

Figure 6: Privacy requirement expression and association with a cloud.

The benefits of involving a remote server lie in that privacy requirements can be changed at different locations and environments. The requirement is also transparent to recorders. Besides, for recorders, some computational tasks can be outsourced to the server, reducing the overhead on cameras. More importantly, with a central privacy management service provider, people can express more fine-grained privacy requirements. It is also possible to integrate the service into online social networks, so that images will be filtered before being shared online. The biggest limitation, however, is that Internet connection is required, in order to get processed image instantly. And another big concern is the security of the server: how can uses trust the server that their privacy profile stored in the server will not be leaked.

We conclude the advantages and limitations of different privacy requirement expression in Table 1. Although a perfect expression manner does not exist, these methods have chances to be adopted in different scenarios, since some of their limitations can be overcome with more sophisticated designs.

Table 1: Comparison of privacy requirement expression methods.

Method	Examples	Advantages	Limitations	Adopted in
Visual indicator	marker, tag, gesture	1. easy to be detected	1. static 2. intention leak 3. overhead on camera	[59, 10, 57, 53, 42, 63]
Wireless link	WiFi, BLE	1. dynamic 2. transparent	1. additional device 2. "always-on" discovery 3. opportunistic	[34, 11, 57, 6]
Server	cloud	1. dynamic 2. transparent	1. security 2. Internet connection	[80]

5.2 Requirement Association

Privacy requirement association aims to find the region of interest (i.e., regions that should be protected) in images or videos, with requirements either directly expressed by objects that seek privacy, or explicitly set by people in advance. In this subsection, we describe requirement association approaches divided into three categories: 1) vision-based detection, 2) sensor-based identification, and 3) feature-based matching.

5.2.1 Vision-Based Detection

Vision-based detection refers to using computer vision techniques to detect the object, track the object, and sometimes recognize the object. The object can be sensitive visual information itself, or visual indicators that express privacy requirements. For example, in [16, 40, 39], the goal is to detect face or other sensitive objects. In [59, 56, 10, 53, 6, 63], visual indicators are supposed to be detected.

To this end, face detection [71, 73, 28], individual recognition/identification [45, 30, 72, 43], object detection [9, 55, 54], and tracking [19] are commonly used in privacy protection systems. It can be imagined that with developments in computer vision, performance of systems that use computer vision techniques can be improved, both in terms of accuracy and efficiency.

5.2.2 Sensor-Based Identification

Sensor-based identification refers to methods that use information provided by sensors, or directly rely on sensors to identify objects that requires protection in the image or video. In Cloak, location information is provided by the GPS on devices carried by in-

dividuals [11]. In PriSurv, RFID-tags are carried by individuals so that RFID-readers in the environment can identify objects in original images according to received signal strength[17]. In Courteous Glass, a gesture shown by individuals can be detected using far-infrared imager [42].

These methods are preliminary attempts towards visual privacy protection with additional sensors in the real world. In the near future, we can expect more advanced sensor-based identification techniques to be developed, which will definitely benefit visual privacy protection.

5.2.3 Feature-Based Identification

Feature-based identification refers to methods that identify the object, usually the person that seeks privacy in the image or video, based on biometric features. The goal is to match the information provided by individuals with those captured. As a result, secure matching is used, especially when a third-party, such as a central server/cloud is involved to associate the privacy requirement with objects that request protection. Though the nature is still identifying object based on some visual clues, we put it into another category as visual clues are biometric features that point to the sensitive object directly. They are quite sensitive and private in this case. For example, biometric information like face features being accessed by others may also leak private information. Thus there is a need to protect biometric data, usually by encrypting or projecting the representations.

In [6], secure dot product [29] based on Paillier homomorphic encryption scheme [49], and secure threshold computation based on garbled circuits [58, 37] are used to recognize the face. Specifically, homomorphic encryption allows computations to be carried on cipertext, thus generating an encrypted result. After decrypting the result, it will match the result of operations performed on the plaintext. Garbled circuits allow two pairs with inputs x and y, respectively, to compute an arbitrary function $f(x, y)$ without disclosing inputs. In [80], an encryption-free privacy preserving vector distance protocol was proposed to conduct the portrait graph matching in a non-interactive manner against untrusted server. After recorders and bystanders obtain a same random number for transforming feature vectors, the cloud will compute the distance between vectors.

However, to achieve privacy-preserving matching by using methods like multi-party computation (SMC) or garble circuit that compute Euclidean distance between vectors usually require frequent online interactions among data owners. Moreover, their large

Table 2: Characteristics of privacy requirement association methods.

Method	Techniques	Features	Used in
Vision-based detection	CV, ML	+ straightforward - computation overhead	[16, 40, 39, 59, 10, 57, 53, 42, 63]
Sensor-based identification	RFID, GPS, FIR imager	+ flexible - additional sensor	[11, 17, 42]
Feature-based identification	secure matching	+ reliable - inefficient	[80, 6]

computation cost and ciphertext size make them unsuitable for mobile applications. But in order to gain the trust from people who are worried about visual privacy, privacy preserving at the remote cloud must be treated carefully.

In summary, involving a third-party to protect visual privacy seems to be a feasible and practical approach for pervasive personal cameras. The third-party can be responsible for privacy requirement management and association, bringing the gap between recorders and bystanders. Thus privacy-preserving biometric identification techniques such as those proposed in [25, 26] will play an important role here. How to make these algorithms more suitable for specific scenarios, such as large-scale users and mobile application platforms are worth exploring.

Table 2 summarizes the characteristics of different privacy requirement association methods. With advances in fields such as computer vision and cloud security, privacy requirement association will be more user-friendly and efficient to be adopted in the future.

6 Privacy Protection Method and Enforcement

6.1 Protection Method

As the third key challenge, how to protect the private information on the images and videos is also important. In some situations, pictures are taken for some purposes, therefore utility should still be kept. In this subsection, we classify privacy protection methods into four classes: 1) intervention, 2) data modification, 3) visual abstract, and 4) data encryption.

6.1.1 Intervention

Intervention prevents cameras from recording any visual information about the object. It is a more through way of protecting privacy from the source. For example, Figure 7 shows the design of a capture-resistant environment in [70, 50]. When a person takes a camera into the capture-resistant environment, the system locates any number of retro-reflective CCD or CMOS camera sensors within its field of view. A pulsing light is then directed at the lens, distorting any imagery the camera records.

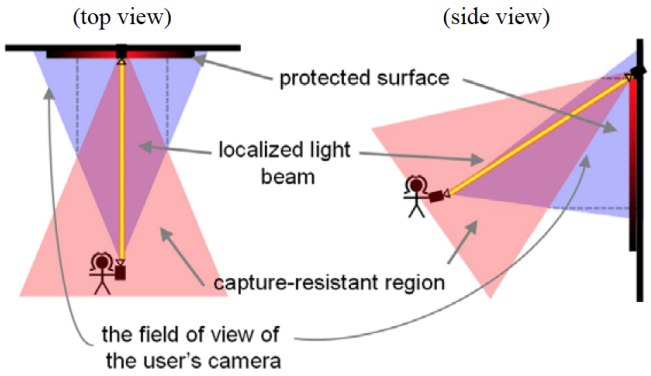

Figure 7: Prevent camera recording by designing a capture-resistant environment.

In addition to stopping cameras from recording by force, systems proposed in [42] and [57] asks cameras to turn off recording automatically. Slightly different from above methods, Yamada et al. designed an approach that prohibits computer vision algorithms from mining or interpreting media data. They made faces in captured images undetectable by using a device worn on the face that corrupts facial features with light absorbing and reflecting materials [76].

6.1.2 Data Modification

Assuming regions that contain private or sensitive information have been identified, data modification methods change the regions on images to prevent sensitive visual information from being viewed by others. The most common method is obscuration, which has been studied for years. Obscuration methods such as blurring [73, 28], masking [77, 59], pixelization [38], scrambling/distortion [21, 22, 23], and permuting pixels [15, 13, 10] are commonly used to remove or protect sensitive visual information.

However, recent work has shown limitations of some obscuration methods. Neustardter et al. found video blurring is unable to balance privacy with awareness for risky situations. Participants also suggested that other popular image masking techniques would be problematic as well [47]. Gross et al. showed that simple blurring or pixelization may not defeat face recognition system [32]. Dufaux et al. also showed the ineffectiveness of naive privacy protection techniques such as blurring and pixelization, and demonstrated the effectiveness of more sophisticated scrambling techniques [24].

The conclusion here is that more sophisticated method should be explored, to further meet people's privacy protection requirements, and make the media data acceptable even after removal of some visual information. Therefore, advanced techniques such as image inpainting can be applied, which restores missing or damaged areas in an image [33].

6.1.3 Visual Abstraction

Visual abstraction controls how much visual information can be accessed. By implementing visual abstraction, only part of information is presented, in the form of text or visual data, while hiding all the rest of the information.

Visual abstraction is useful in video surveillance systems. Figure 8 shows different video content rendering options proposed in [60]. Generally, the system policies might require partially or fully obscuring or statistically perturbing certain components of the extracted information, such as a subject's location, pose, activity, and so on. For example, a particular policy may offer statistical information, such as when the street is the most crowded. Another policy might require that all faces in the video be obscured, so that only gender information (but not identity, age, or expression) is available. Similarly, in [17], different levels of visual abstraction is provided. Details of the object can be disclosed, partially-hidden, or completely hidden.

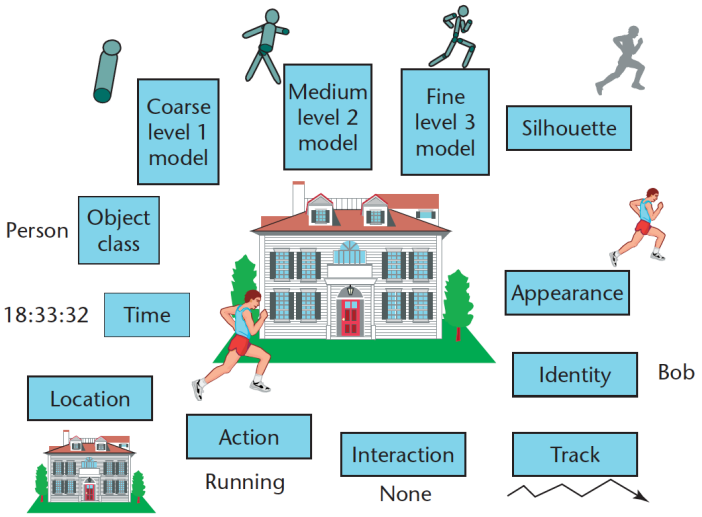

Figure 8: Different video rendering options.

Visual abstraction also works for application such as mobile AR, where coarse personal information is needed. To this end, Jana et al. designed methods that only expose higher-level objects, such as a skeleton or a face, to applications [39]. They also proposed and integrated a privacy protection framework with computer vision library OpenCV, which conveys only high-level information to applications, while hiding more specific and sensitive details [40].

6.1.4 Data Encryption

Another method is to prevent the whole image or video from unauthorized access by using cryptographic techniques. The raw data are then encoded or encrypted directly at the source.

The basic steps are as follows. Two parties (e.g., reviewer and video surveillance system) first exchange keys. The original image or video then can be encrypted using a key. Next, a person can use a decryption key to get the original image. Such data encryption mechanism can work for video surveillance systems, so that people with high authority can use their secret keys to access the video [81, 60]. It can also work for online social networks, where an image uploaded by a user can only be seen by another specific user [65]. When the image is a secret between multiple persons, anyone else who wants to access the raw data should get permissions from all privacy stakeholders [34].

6.2 Protection Enforcement

Privacy protection can be enforced immediately or at any moment before others can access the raw image or video. Therefore we divide protection enforcement into two classes according to when protection happens: in-situ and before dissemination.

6.2.1 In-situ

The in-situ enforcement includes situations that prevent cameras from recording anything, or process the visual data locally on the device, before any applications can have access to the original data. In the later case, image and videos can still be captured, but they will be processed and then used without any concern of privacy violation. Therefore, an ideal approach is to integrate privacy protection into the default camera subsystem. For example, when a wearable camera is taking pictures, the operating system will process the image instantly, before any third-party applications can access the image. Systems such as [53] implement the in-situ protection that works in real-time to prevent sensitive visual information being exposed.

However, considering the large amount of cameras on the market and in use now, it is not easy for device manufacturers to take actions. Moreover, it requires high computational power from devices due to heavy computer vision tasks, such as detection and recognition.

6.2.2 Dissemination

An alternative is to process the video as long as it is not available to the public, since the privacy issues are usually raised at sharing or dissemination moments.

As a result, visual privacy protection can be a service integrated in image and video sharing platforms as proposed in [10]. Recorders leave the visual privacy protection job to the service provider, which are responsible for enforcing protection actions according to privacy requirements from the object. For example, a life-logging camera records some videos and synchronizes to an archive application on the personal computer. The archive application then will process the video to respect people's privacy. A video is uploading to an online social network. The platform will check the video to make sure it does not violate others' visual privacy.

Table 3: Visual privacy protection systems and frameworks.

System	Scenario	Target	Privacy expression	Privacy association	Protection method	Protection enforcement
[60]	video surveillance	all	n/a	n/a	encryption & abstraction	in-situ
[81]	video surveillance	person	n/a	vision-based detection	masking	n/a
[11]	video surveillance	person	wireless signal	sensor-based identification	n/a	dissemination
[17]	video surveillance	person	wireless signal	sensor-based identification	abstraction	n/a
[59]	video surveillance	person	visual indicator	vision-based detection	masking	n/a
[34]	personal camera	all	wireless signal	n/a	data encryption	in-situ
[70]	personal camera	person	n/a	n/a	intervention	in-situ
[10]	personal camera	person	visual indicator	vision-based detection	permuting pixel	dissemination
[80]	personal camera	person	server/ cloud	feature-based identification	data modification	dissemination
[6]	personal camera	person	wireless signal	feature-based identification	data modification	n/a
[63]	personal camera	person	visual indicator	vision-based detection	data modification	n/a
[16]	perceptual application	person	n/a	vision-based detection	masking	n/a
[67]	perceptual application	place	n/a	n/a	n/a	n/a
[42]	perceptual application	person	visual indicator	sensor-based identification	intervention	in-situ
[40]	perceptual application	all	n/a	vision-based detection	abstraction	in-situ
[39]	perceptual application	all	n/a	vision-based detection	abstraction	in-situ
[56]	perceptual application	all	visual & wireless	vision-based & sensor-based	intervention & modification	in-situ
[53]	perceptual application	all	visual indicator	vision-based detection	intervention	in-situ

7 Challenges and Opportunities

In this section, we discuss challenges and opportunities in visual privacy protection, based on existing research works and our reflections after practice.

7.1 Challenges

There are different challenges in systems and frameworks specially designed for certain applications or scenarios. A perfect solution that applies to all objects and settings, and satisfies various privacy requirements does not exist. Some challenges we are faced with in multiple scenarios are presented as follows:

First, only in-situ privacy protection can prevent privacy invasion fundamentally. However, enforcing privacy-preserving mechanism for cameras requires huge efforts, such as designing an additional module or layer in the operating system of the camera device. Problems like what should be protected and how to protect should be addressed first. Besides, considering the large number of camera devices already in the market, it cannot solve existing problem in the near future.

Second, if protection enforcement will work before dissemination moments, we need to persuade recorders to comply with visual privacy protection policies. It usually includes asking recorders to install a specific camera application, which embeds protection mechanism in it. If recorders are not willing to use such applications due to whatever reasons, bystanders' privacy can never be protected.

Third, if privacy protection service will be provided by a third-party, the security of a server/cloud that takes charge of the privacy management should be guaranteed. For example, if an online social networking site now offers privacy protection settings, that images can be processed automatically before they are available to the public, users are supposed to first provide some privacy preference information. As a result, users may worry about the information leak in this step, which resists their willingness to use the service.

Finally, a system could not have real practical use, unless it is widely adopted (e.g., enforced by law) in the real world. Currently, most of the solutions are conceptual frameworks, which are based on assumptions and have not taken real situations into consideration, such as efficiency, people's acceptance. In fact, most of the methods are far away from real life deployment.

7.2 Opportunities

Though a number of challenges remains unsolved, which makes we are far away from a healthy visual privacy protection ecosystem, opportunities still exist that provide promising directions to improve current visual privacy protection approaches.

First, identifying risks caused by visual information that endanger people's privacy, both online or offline, will raise people's awareness of visual privacy, and further encourage them to use privacy protection service. It has been known by social psychologists for decades that the relationship between attitudes and behaviors is complex, if not weak. For example, though you are aware that applications installed on your smartphone may hack images and videos stored, you are probably still give applications permissions to access visual data. In other words, even though people claim they care about visual privacy, the concerns may not be strong enough to drive serious actions against privacy intrusion. Therefore, efforts can be put in identifying severe violations of visual privacy, in order to persuade people that protecting visual privacy is extremely important.

Second, exploring what is considered sensitive in fact is necessary. There are some researches along this direction. PriFir used low-power sensors (e.g., accelerometer, light sensor) embedded in smartphones and smart watches to understand user's preferences about certain scenarios being sensitive or private, and identify sensitive scenarios [75]. Spyromitros-Xioufis et al. developed personalized privacy classification models for images on online social networks by utilizing users' feedback [64]. Besides, there are many works on privacy-aware image classification, in the context of online image sharing [79, 69]. Furthermore, Badii et al. stated the needs for context identification [8]. If what is sensitive or private can be defined by different persons, users will be more willing to use privacy protection service.

Third, integrating visual privacy protection into small world social networks will be more convincing. As people upload images and videos to multiple online social media sites to share with others, these sites are the places that can guard media data before they are disseminated. In this context, more fine-grained privacy preserving actions can be performed.

Fourth, studying people's reception towards current visual privacy protection system, will indicate how people like the idea, and what should be done to improve the solution.

Finally, there are other topics related to preserving privacy when using the data. With

the rapidly expanding filed of machine learning, a large number of data are provided by individuals who wish to retain a degree of privacy. Privacy then is formalized via the notion of "*differential privacy*". It defines a probabilistic channel between the data and the outside world such that an observer of the output channel cannot infer reliably whether particular individuals have supplied data or not [41]. For example, it turns out that people are starting to feel uncomfortable about sending a lot of personal information to various services they use. As a result, differential privacy technology is used such as by Apple to obscure an individual's identity by adding noise to the data collected from users, to improve user experience [31].

8 Conclusion and Future Research

In this survey, we introduced the concept of visual privacy and privacy issues in ubiquitous computing environments. We presented people's privacy concerns, and investigated the reasons behind. Then we reviewed a number of visual privacy protection systems and frameworks proposed in three different ubiquitous computing scenarios. We classified these methods according to how privacy requirement is expressed, associated, and enforced, along the general workflow of privacy protection, followed by discussions about challenges and opportunities.

A useful visual privacy protection framework/system is a considerate high-level design, with a combination of techniques to address a series of problems and challenges. It is of tremendous value, but also faced with limitations in the real world. Our future work aims to address part of the challenges, in order to design a practical visual privacy protection framework, that can be applied in some real life scenarios.

Gabriel Garca Mrquez, a novelist, said in one of his books:"All human beings have three lives: public, private, and secret." We have our public life, which is what we willingly do and share with others in a wide range of social settings. We also have our private life, which we reluctantly give away in the hope that it is not fully revealed to the world or to those who should not see it. Then there is our secret life, which, for now, can only be found offline. To enjoy the public, private, and secret life requires efforts from everyone in the society. Otherwise, we can only enjoy the illusion of private life, and finally give up secret life.

9 References

[1] Congress's letter. http://marketingland.com/google-glass-44279, 2013.

[2] Data protection authorities' letter. https://www.priv.gc.ca/en/opc-news/news-and-announcements/2013/nr-c_130618/, 2013.

[3] Iapp. https://iapp.org/about/what-is-privacy/, 2017.

[4] Image Privacy Rights – Facebook. https://www.facebook.com/help/428478523862899, 2017.

[5] Alessandro Acquisti, R Gross, and F Stutzman. Privacy in the age of augmented reality. *Proc. National Academy of Sciences*, 2011.

[6] Paarijaat Aditya, Rijurekha Sen, Peter Druschel, Seong Joon Oh, Rodrigo Benenson, Mario Fritz, Bernt Schiele, Bobby Bhattacharjee, and Tong Tong Wu. I-pic: A platform for privacy-compliant image capture. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services, MobiSys*, volume 16, 2016.

[7] Salman Aslam. Instagram by the numbers: Stats, demographics & fun facts. https://www.omnicoreagency.com/instagram-statistics/, 2017.

[8] Atta Badii, Mathieu Einig, Marco Tiemann, Daniel Thiemert, and Chattun Lallah. Visual context identification for privacy-respecting video analytics. In *Multimedia Signal Processing (MMSP), 2012 IEEE 14th International Workshop on*, pages 366–371. IEEE, 2012.

[9] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. *Computer vision–ECCV 2006*, pages 404–417, 2006.

[10] Cheng Bo, Guobin Shen, Jie Liu, Xiang-Yang Li, YongGuang Zhang, and Feng Zhao. Privacy. tag: Privacy concern expressed and respected. In *Proceedings of the 12th ACM Conference on Embedded Network Sensor Systems*, pages 163–176. ACM, 2014.

[11] Jack Brassil. Using mobile communications to assert privacy from video surveillance. In *Parallel and Distributed Processing Symposium, 2005. Proceedings. 19th IEEE International*, pages 8–pp. IEEE, 2005.

[12] Kelly E Caine, Arthur D Fisk, and Wendy A Rogers. Benefits and privacy concerns of a home equipped with a visual sensing system: A perspective from older adults. In *Proceedings of the human factors and ergonomics society annual meeting*, volume 50, pages 180–184. Sage Publications Sage CA: Los Angeles, CA, 2006.

[13] Paula Carrillo, Hari Kalva, and Spyros Magliveras. Compression independent reversible encryption for privacy in video surveillance. *EURASIP Journal on Information Security*, 2009(1):429581, 2010.

[14] Andrea Cavailaro. Privacy in video surveillance [in the spotlight]. *IEEE Signal Processing Magazine*, 24(2):168–166, 2007.

[15] Ankur Chattopadhyay and Terrance E Boult. Privacycam: a privacy preserving camera using uclinux on the blackfin dsp. In *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*, pages 1–8. IEEE, 2007.

[16] Jayashri Chaudhari, S Cheung Sen-ching, and M Vijay Venkatesh. Privacy protection for life-log video. In *Signal Processing Applications for Public Security and Forensics, 2007. SAFE'07. IEEE Workshop on*, pages 1–5. IEEE, 2007.

[17] Kenta Chinomi, Naoko Nitta, Yoshimichi Ito, and Noboru Babaguchi. Prisurv: privacy protected video surveillance system using adaptive visual abstraction. In *International Conference on Multimedia Modeling*, pages 144–154. Springer, 2008.

[18] Albert Costill. Top 10 places that have banned google glass. `https://www.searchenginejournal.com/top-10-places-that-have-banned-google-glass/66585/`, 2013.

[19] Martin Danelljan, Gustav Häger, Fahad Khan, and Michael Felsberg. Accurate scale estimation for robust visual tracking. In *British Machine Vision Conference, Nottingham, September 1-5, 2014*. BMVA Press, 2014.

[20] Tamara Denning, Zakariya Dehlawi, and Tadayoshi Kohno. In situ with bystanders of augmented reality glasses: Perspectives on recording and privacy-mediating technologies. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 2377–2386. ACM, 2014.

[21] Frederic Dufaux and Touradj Ebrahimi. Scrambling for video surveillance with privacy. In *Computer Vision and Pattern Recognition Workshop, 2006. CVPRW'06. Conference on*, pages 160–160. IEEE, 2006.

[22] Frederic Dufaux and Touradj Ebrahimi. H. 264/avc video scrambling for privacy protection. In *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*, pages 1688–1691. IEEE, 2008.

[23] Frederic Dufaux and Touradj Ebrahimi. Scrambling for privacy protection in video surveillance systems. *IEEE Transactions on Circuits and Systems for Video Technology*, 18(8):1168–1174, 2008.

[24] Frédéric Dufaux and Touradj Ebrahimi. A framework for the validation of privacy protection solutions in video surveillance. In *Multimedia and Expo (ICME), 2010 IEEE International Conference on*, pages 66–71. IEEE, 2010.

[25] Zekeriya Erkin, Martin Franz, Jorge Guajardo, Stefan Katzenbeisser, Inald Lagendijk, and Tomas Toft. Privacy-preserving face recognition. In *International Symposium on Privacy Enhancing Technologies Symposium*, pages 235–253. Springer, 2009.

[26] David Evans, Yan Huang, Jonathan Katz, and Lior Malka. Efficient privacy-preserving biometric identification. In *Proceedings of the 17th conference Network and Distributed System Security Symposium, NDSS*, 2011.

[27] Fortunelords. Youtube Statistics. `https://fortunelords.com/youtube-statistics/`, 2017.

[28] Andrea Frome, German Cheung, Ahmad Abdulkader, Marco Zennaro, Bo Wu, Alessandro Bissacco, Hartwig Adam, Hartmut Neven, and Luc Vincent. Large-scale privacy protection in google street view. In *Computer Vision, 2009 IEEE 12th International Conference on*, pages 2373–2380. IEEE, 2009.

[29] Bart Goethals, Sven Laur, Helger Lipmaa, and Taneli Mielikäinen. On private scalar product computation for privacy-preserving data mining. In *International Conference on Information Security and Cryptology*, pages 104–120. Springer, 2004.

[30] Shaogang Gong, Marco Cristani, Shuicheng Yan, and Chen Change Loy. *Person re-identification*, volume 1. Springer, 2014.

[31] Matthew Green. What is differential privacy? https://blog.cryptographyengineering.com/2016/06/15/what-is-differential-privacy/, 2016.

[32] Ralph Gross, Latanya Sweeney, Jeffrey Cohn, Fernando De la Torre, and Simon Baker. Face de-identification. In *Protecting Privacy in Video Surveillance*, pages 129–146. Springer, 2009.

[33] Christine Guillemot and Olivier Le Meur. Image inpainting: Overview and recent advances. *IEEE signal processing magazine*, 31(1):127–144, 2014.

[34] J Alex Halderman, Brent Waters, and Edward W Felten. Privacy management for portable recording devices. In *Proceedings of the 2004 ACM workshop on Privacy in the electronic society*, pages 16–24. ACM, 2004.

[35] Roberto Hoyle, Robert Templeman, Denise Anthony, David Crandall, and Apu Kapadia. Sensitive lifelogs: A privacy analysis of photos from wearable cameras. In *Proceedings of the 33rd Annual ACM conference on human factors in computing systems*, pages 1645–1648. ACM, 2015.

[36] Roberto Hoyle, Robert Templeman, Steven Armes, Denise Anthony, David Crandall, and Apu Kapadia. Privacy behaviors of lifeloggers using wearable cameras. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 571–582. ACM, 2014.

[37] Yan Huang, David Evans, Jonathan Katz, and Lior Malka. Faster secure two-party computation using garbled circuits. In *USENIX Security Symposium*, volume 201, 2011.

[38] Scott E Hudson and Ian Smith. Techniques for addressing fundamental privacy and disruption tradeoffs in awareness support systems. In *Proceedings of the 1996 ACM conference on Computer supported cooperative work*, pages 248–257. ACM, 1996.

[39] Suman Jana, David Molnar, Alexander Moshchuk, Alan Dunn, Benjamin Livshits, Helen J Wang, and Eyal Ofek. Enabling fine-grained permissions for augmented reality applications with recognizers. In *Presented as part of the 22nd USENIX Security Symposium (USENIX Security 13)*, pages 415–430, 2013.

[40] Suman Jana, Arvind Narayanan, and Vitaly Shmatikov. A scanner darkly: Protecting user privacy from perceptual applications. In *Security and Privacy (SP), 2013 IEEE Symposium on*, pages 349–363. IEEE, 2013.

[41] Michael I Jordan and Tom M Mitchell. Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245):255–260, 2015.

[42] Jaeyeon Jung and Matthai Philipose. Courteous glass. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*, pages 1307–1312. ACM, 2014.

[43] Erik Learned-Miller, Gary B Huang, Aruni RoyChowdhury, Haoxiang Li, and Gang Hua. Labeled faces in the wild: A survey. In *Advances in Face Detection and Facial Image Analysis*, pages 189–248. Springer, 2016.

[44] Ming Li, Shucheng Yu, Yao Zheng, Kui Ren, and Wenjing Lou. Scalable and secure sharing of personal health records in cloud computing using attribute-based encryption. *IEEE transactions on parallel and distributed systems*, 24(1):131–143, 2013.

[45] Ju Man and Bir Bhanu. Individual recognition using gait energy image. *IEEE transactions on pattern analysis and machine intelligence*, 28(2):316–322, 2006.

[46] Shibnath Mukherjee, Zhiyuan Chen, and Aryya Gangopadhyay. A privacy-preserving technique for euclidean distance-based mining algorithms using fourier-related transforms. *The VLDB JournalThe International Journal on Very Large Data Bases*, 15(4):293–315, 2006.

[47] Carman Neustaedter, Saul Greenberg, and Michael Boyle. Blur filtration fails to preserve privacy for home-based video conferencing. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 13(1):1–36, 2006.

[48] David H Nguyen, Aurora Bedford, Alexander Gerard Bretana, and Gillian R Hayes. Situating the concern for information privacy through an empirical study of responses to video recording. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 3207–3216. ACM, 2011.

[49] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 223–238. Springer, 1999.

[50] Shwetak N Patel, Jay W Summet, and Khai N Truong. Blindspot: Creating capture-resistant spaces. In *Protecting Privacy in Video Surveillance*, pages 185–201. Springer, 2009.

[51] The Statistics Portal. Smartphones - Statistics & Facts. `https://www.statista.com/topics/840/smartphones/`, 2017.

[52] The Statistics Portal. Wearable device revenue worldwide 2015-2021. `https://www.statista.com/statistics/610447/wearable-device-revenue-worldwide/`, 2017.

[53] Nisarg Raval, Animesh Srivastava, Ali Razeen, Kiron Lebeck, Ashwin Machanavajjhala, and Landon P Cox. What you mark is what apps see. In *ACM International Conference on Mobile Systems, Applications, and Services (Mobisys)*, 2016.

[54] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 779–788, 2016.

[55] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.

[56] Franziska Roesner, Tadayoshi Kohno, and David Molnar. Security and privacy for augmented reality systems. *Communications of the ACM*, 57(4):88–96, 2014.

[57] Franziska Roesner, David Molnar, Alexander Moshchuk, Tadayoshi Kohno, and Helen J Wang. World-driven access control for continuous sensing. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 1169–1181. ACM, 2014.

[58] Ahmad-Reza Sadeghi, Thomas Schneider, and Immo Wehrenberg. Efficient privacy-preserving face recognition. In *International Conference on Information Security and Cryptology*, pages 229–244. Springer, 2009.

[59] Jeremy Schiff, Marci Meingast, Deirdre K Mulligan, Shankar Sastry, and Ken Goldberg. Respectful cameras: Detecting visual markers in real-time to address privacy concerns. In *Protecting Privacy in Video Surveillance*, pages 65–89. Springer, 2009.

[60] Andrew Senior, Sharath Pankanti, Arun Hampapur, Lisa Brown, Ying-Li Tian, Ahmet Ekin, Jonathan Connell, Chiao Fe Shu, and Max Lu. Enabling video privacy through computer vision. *IEEE Security & Privacy*, 3(3):50–57, 2005.

[61] Andrew Senior and Andrew W Senior. *Protecting privacy in video surveillance*, volume 1. Springer, 2009.

[62] Yan Shoshitaishvili, Christopher Kruegel, and Giovanni Vigna. Portrait of a privacy invasion. *Proceedings on Privacy Enhancing Technologies*, 2015(1):41–60, 2015.

[63] Jiayu Shu, Rui Zheng, and Pan Hui. Demo: Interactive visual privacy control with gestures. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services Companion*, pages 120–120. ACM, 2016.

[64] Eleftherios Spyromitros-Xioufis, Symeon Papadopoulos, Adrian Popescu, and Yiannis Kompatsiaris. Personalized privacy-aware image classification. In *Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval*, pages 71–78. ACM, 2016.

[65] Weiwei Sun, Jiantao Zhou, Ran Lyu, and Shuyuan Zhu. Processing-aware privacy-preserving photo sharing over online social networks. In *Proceedings of the 2016 ACM on Multimedia Conference*, pages 581–585. ACM, 2016.

[66] Paul Szoldra. This is everything edward snowden revealed in one year of unprecedented top-secret leaks. `http://www.businessinsider.com/snowden-leaks-timeline-2016-9`, 2016.

[67] Robert Templeman, Mohammed Korayem, David J Crandall, and Apu Kapadia. Placeavoider: Steering first-person cameras away from sensitive spaces. In *NDSS*, 2014.

[68] Robert Templeman, Zahid Rahman, David Crandall, and Apu Kapadia. Placeraider: Virtual theft in physical spaces with smartphones. *arXiv preprint arXiv:1209.5982*, 2012.

[69] Ashwini Kishore Tonge and Cornelia Caragea. Image privacy prediction using deep features. In *AAAI*, pages 4266–4267, 2016.

[70] Khai N Truong, Shwetak N Patel, Jay W Summet, and Gregory D Abowd. Preventing camera recording by designing a capture-resistant environment. In *International Conference on Ubiquitous Computing*, pages 73–86. Springer, 2005.

[71] Paul Viola and Michael J Jones. Robust real-time face detection. *International journal of computer vision*, 57(2):137–154, 2004.

[72] He Wang, Xuan Bao, Romit Roy Choudhury, and Srihari Nelakuditi. Visually fingerprinting humans without face recognition. In *Proceedings of the 13th Annual International Conference on Mobile Systems, Applications, and Services*, pages 345–358. ACM, 2015.

[73] Jian-Gang Wang, Andy Suwandy, and Wei-Yun Yau. Face obscuration in a video sequence by integrating kernel-based mean-shift and active contour. In *Control, Automation, Robotics and Vision, 2008. ICARCV 2008. 10th International Conference on*, pages 2314–2318. IEEE, 2008.

[74] Raymond Wong. Samsung patents smart contact lenses with a built-in camera. `http://mashable.com/2016/04/05/samsung-smart-contact-lenses-patent/\#jCGVCY0GYaqp`, 2016.

[75] Muchen Wu, Parth H Pathak, and Prasant Mohapatra. Enabling privacy-preserving first-person cameras using low-power sensors. In *Sensing, Communication, and Networking (SECON), 2015 12th Annual IEEE International Conference on*, pages 444–452. IEEE, 2015.

[76] Takayuki Yamada, Seiichi Gohshi, and Isao Echizen. Privacy visor: Method based on light absorbing and reflecting properties for preventing face image detection. In *Systems, Man, and Cybernetics (SMC), 2013 IEEE International Conference on*, pages 1572–1577. IEEE, 2013.

[77] Xiaoyi Yu, Kenta Chinomi, Takashi Koshimizu, Naoko Nitta, Yoshimichi Ito, and Noboru Babaguchi. Privacy protecting visual processing for secure video surveillance. In *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*, pages 1672–1675. IEEE, 2008.

[78] Zephoria. The top 20 valuable facebook statistics. `https://zephoria.com/top-15-valuable-facebook-statistics/`, 2017.

[79] Sergej Zerr, Stefan Siersdorfer, Jonathon Hare, and Elena Demidova. Privacy-aware image classification and search. In *Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval*, pages 35–44. ACM, 2012.

[80] Lan Zhang, Kebin Liu, Xiang-Yang Li, Cihang Liu, Xuan Ding, and Yunhao Liu. Privacy-friendly photo capturing and sharing system. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, pages 524–534. ACM, 2016.

[81] Wei Zhang, SS Cheung, and Minghua Chen. Hiding privacy information in video surveillance system. In *Image Processing, 2005. ICIP 2005. IEEE International Conference on*, volume 3, pages II–868. IEEE, 2005.