

Framework for Objective Evaluation of Privacy Filters

Pavel Korshunov^a, Andrea Melle^b, Jean-Luc Dugelay^b, and Touradj Ebrahimi^a

^aMultimedia Signal Processing Group, EPFL, Switzerland;

^bMultimedia Department, EURECOM, France;

ABSTRACT

Extensive adoption of video surveillance, affecting many aspects of our daily lives, alarms the public about the increasing invasion into personal privacy. To address these concerns, many tools have been proposed for protection of personal privacy in image and video. However, little is understood regarding the effectiveness of such tools and especially their impact on the underlying surveillance tasks, leading to a tradeoff between the preservation of privacy offered by these tools and the intelligibility of activities under video surveillance. In this paper, we investigate this privacy-intelligibility tradeoff objectively by proposing an objective framework for evaluation of privacy filters. We apply the proposed framework on a use case where privacy of people is protected by obscuring faces, assuming an automated video surveillance system. We used several popular privacy protection filters, such as blurring, pixelization, and masking and applied them with varying strengths to people’s faces from different public datasets of video surveillance footage. Accuracy of face detection algorithm was used as a measure of intelligibility (a face should be detected to perform a surveillance task), and accuracy of face recognition algorithm as a measure of privacy (a specific person should not be identified). Under these conditions, after application of an ideal privacy protection tool, an obfuscated face would be visible as a face but would not be correctly identified by the recognition algorithm. The experiments demonstrate that, in general, an increase in strength of privacy filters under consideration leads to an increase in privacy (i.e., reduction in recognition accuracy) and to a decrease in intelligibility (i.e., reduction in detection accuracy). Masking also shows to be the most favorable filter across all tested datasets.

Keywords: Privacy protection, video surveillance, objective evaluations, privacy-intelligibility tradeoff

1. INTRODUCTION

The alarming rate at which video surveillance is being adopted has raised concerns among public and has motivated development of privacy protection tools. Typical techniques used for obscuring personal information in a video in order to preserve privacy include blurring and pixelization of sensitive regions or their masking. More advanced privacy protection techniques have also been developed recently, such as scrambling,¹ encryption of faces in video,² anonymization,³ geometrical warping,⁴ etc.

However, there is a noticeable lack of methods to assess the performance of privacy protection tools and their impact on the surveillance task. While many evaluation protocols and tools (most notably those developed as part of PETS* workshops and grand challenges) are available to test the robustness, accuracy, and efficiency of video analytics for surveillance, little attention has been devoted to privacy aspects.

Previously, we proposed a subjective evaluation methodology to analyze the tradeoff between the preservation of privacy offered by privacy protection filters and the intelligibility of activities under video surveillance.⁵ The proposed methodology performed well in both laboratory⁶ and crowdsourcing environments.⁷ These subjective evaluations demonstrated that the tradeoff between intelligibility and privacy can be quantified, and one can find a suitable balance between the two for a given surveillance objective and visual privacy filter. Although, subjective evaluation is undeniably the most accurate method of measuring subjective human perception (and privacy is a subjective notion), such evaluations are time consuming and impractical. Therefore, the development of objective metrics for measuring privacy and intelligibility is necessary for practical video surveillance systems.

emails: pavel.korshunov@epfl.ch, melle@eurecom.fr, dugelay@eurecom.fr, touradj.ebrahimi@epfl.ch

*IEEE International Workshop on Performance Evaluation of Tracking and Surveillance (PETS)

In this paper, we propose an objective framework for evaluation of visual privacy filters and for identifying the privacy-intelligibility tradeoff. Typically, privacy filter is a specific type of visual distortion which aims at obscuring personal visual details in video, leading to higher privacy protection, but retains their general appearance, thus keeping the same level of intelligibility. Hence, an ideal privacy filter should hide all possible personal details, resulting in the highest possible privacy protection, but keeps intact all visual information necessary for the performance of a surveillance task. Therefore, for evaluation of a given privacy protection tool, it would be intuitive to measure the amount of visual details as a metric of privacy and to measure an overall general shape as a metric of intelligibility. The relative change in the metric values would then constitute the privacy-intelligibility tradeoff.

To demonstrate this general framework, we consider the scenario where the face is a personal visual information that needs to be protected. In such a scenario, we can use face recognition as a privacy metric, since it requires fine details of faces to be visible to distinguish them from one another, and face detection as an intelligibility metric, since it only aims at detecting a general appearance of a face. Hence, face detection should continue to be possible after a privacy protection filter is applied, while facial personal features should be destroyed. Therefore, for evaluation of a given privacy filter, one should check that a protected face is detectable by a face detection algorithm but is not correctly identified by the recognition algorithm. Hence, the relative accuracy of a face detection algorithm would reflect the intelligibility level of the assessed visual privacy filter, while relative accuracy of recognition algorithm would reflect its level of privacy. By relative accuracy we understand the change in accuracy compared to the original undistorted image (e.g., if original face was detected by face detection and after privacy filter is applied, it is still detected, then the accuracy is 100%). As a practical application of this scenario, we can assume an automated people counting surveillance system (in airports or public transport), in which video face detection algorithm would be used for people counting and then served as a possible input to a tracking algorithm. Since facial recognition is not part of such systems, to protect privacy of people boarding a bus or walking in airports, faces should be obscured with privacy filters.

Using the above scenario, we evaluate three commonly used privacy protection filters: blurring, pixelization, and masking. We apply these filters with different degrees of strength. For blurring, different degrees of strength are achieved by varying the size of a Gaussian filter, for pixelization, by varying the size of a block, and for masking – the level of opacity. We use widely popular OpenCV implementation of Viola-Jones face detection algorithm⁸ as an intelligibility metric. Since there exists several state-of-art face recognition algorithms, we consider three of them as candidates to be a privacy metric: based on Principal Component Analysis (PCA),⁹ based on Linear Discriminant Analysis (LDA),¹⁰ and based on local features (LBP).¹¹ Their implementations are also available in OpenCV library. Since face detection and recognition algorithms are especially sensitive to sizes of face visual regions and environmental conditions,¹² we used three different publicly available datasets: SCFace surveillance dataset,¹³ ChokePoint surveillance dataset,¹⁴ and FERET dataset¹⁵ for evaluation of face recognition algorithms.

For each evaluated privacy protection filter, the relative changes in detection and recognition accuracies constitute the privacy-intelligibility tradeoff. The filter can be considered to have a good tradeoff property if it does not change the intelligibility (accuracy of face detection is the same as for the original undistorted image) but it increases privacy to maximum (accuracy of face recognition decreases to zero).

2. BACKGROUND AND RELATED WORK

The objective evaluation of several primitive privacy filters was first performed by Newton *et al.*,¹⁶ where the authors demonstrated that such filters cannot adequately protect from the successful face recognition, because recognition algorithms are robust. The robustness of face recognition and detection algorithms to primitive distortions is also reported in Ref. 12. In the work by Dufaux *et al.*,¹⁷ a framework is defined to evaluate the performance of face recognition algorithms applied to images altered by various obfuscation methods, based on the Face Identification Evaluation System (FIES). Experiments using the FERET database showed the ineffectiveness of naïve face obfuscation techniques such as pixelization and blurring. The authors argue that more sophisticated scrambling techniques are more effective in foiling face recognition.

Since privacy is a subjective notion, Korshunov *et al.*⁶ argued that the evaluation should be done subjectively. Therefore, the authors define a subjective methodology for evaluation of privacy protection tools and propose a

subjective evaluation protocol, focusing on two important aspects: (i) how much of the privacy is protected by such a tool and (ii) how much it impacts the efficiency of the underlying surveillance task (intelligibility). The pixelization filter shows the best performance in terms of balancing between privacy protection and allowing sufficient intelligibility. Masking filter, instead, demonstrates the highest privacy protection with low incorrectness and high uncertainty, which can be suitable for the higher security surveillance applications.

For the crowdsourcing based evaluation,⁷ an online application VideoRate has been built which uses Facebook ID for login and a more reliable user authentication. Compared to laboratory based evaluations, this approach provides some flexibility to users, as one could stop participation in the experiments at any time, as well as better simulates the real-time scenario, since the users evaluate videos in different lighting conditions and using monitors with different resolutions and color settings. The call to participate was disseminated in the subjective test using the VideoRate application via such social networks as Facebook, Twitter, and LinkedIn, as well as various research mailing lists. With an estimated outreach to more than 1500, some 120 among them used the application and submitted subjective scores. These subjective results were then compared with the results from a similar evaluation conducted by a conventional approach in a designated research test laboratory at EPFL. The results demonstrate a high correlation with only minor differences favoring the crowdsourcing method, which means that it can be considered as a reliable and effective approach for subjective evaluation of visual privacy filters.

Our proposed objective evaluation framework extends the ideas of earlier works^{16,17} into a privacy-intelligibility tradeoff, demonstrated via subjective experiments by Korshunov *et al.*^{6,7} Hence, our framework, by considering both privacy and intelligibility components and using objective metrics to measure them, provides an effective practical tool for evaluation of privacy filters in video surveillance systems that heavily rely on video analysis algorithms.

There are also a few system extensions that incorporate privacy protection into the existing surveillance frameworks and systems.

The security of the privacy information in surveillance system is addressed in Ref. 18 by removing privacy-sensitive information from the video sequence. A perceptually-based compressed-domain watermarking technique is then used to securely embed this data in the video stream. Similarly, a secure reversible data hiding technique is introduced in Ref. 19 for privacy data embedding. Some privacy data management is also proposed to allow individual users to control access to their private data.

The authors of Ref. 20 focus on the problem of secure streaming of different surveillance video to different client devices with different display sizes. The authors propose a secure device registration inside the surveillance system to protect the private information of different users of the system. The system is domain-based, where different devices of the same user can freely share data among themselves, with one device selected as a domain controller that manages all other devices and stores the global relevant information. The system assumes a dedicated surveillance system manager that manages video streaming rules for different devices and provides them with a web interface. A dedicated licensing server is also proposed to manage users of the system and their usage rights. This kind of solution is not really a privacy respecting surveillance system, since it does not take into account the privacy of the surveyed persons at all, but it shows that secure management of the video surveillance streams, devices, and their users is a very important aspect in video surveillance.

Policy framework specifying whether a person should be shown in a surveillance video is proposed in Ref. 21. A set of control policies is developed using XML specification that describes who should be shown and who are not in the surveillance footage depending on the validity of their RFID tags in the surveillance area.

A camera, called MPEG-7 camera for its processing capabilities and support of MPEG-7 standard, is proposed in Ref. 22 that offers privacy protection. Its privacy protection is based on an object-oriented representation of the scene. The system re-renders a modified video based on the end-user access control authorizations. During re-rendering, areas of the image are blanked out or replaced by computer graphics. The relevant information in the scene is therefore preserved, but privacy-sensitive details are not transmitted.

Based on similar ideas as above, the system introduced in Ref. 23 relies on computer vision to analyze the video content and to automatically extract its components. Different users can selectively get access to these components, depending on their access-control rights. More specifically, the system renders a different version

of the video where privacy-sensitive objects have been hidden. This is achieved while information required to fulfill the surveillance task is preserved. The paper also describes a PrivacyCam, after the same idea as MPEG-7 camera, with built-in privacy protection tools, which directly outputs video streams with privacy-sensitive information removed.

The authors of Ref. 24 propose a real-time policy-based framework for privacy protection in video surveillance. The authors consider granular privacy filters with their hierarchical composition, as well as the set of verifiable policies using a standardized privacy specification language. The framework includes distributed live video database (LVDBMS) that can be used to search live video streams and perform actions of interests on them via its own query language. An administrator can associate privacy filters, defined using privacy specification language (PSL), with camera, query, user group, or view. When there is a combination of privacy filters for a given object, the filter with the highest priority (causing the most corruption to the video) is applied. For evaluation of the proposed framework, the authors adopted both their own recordings of university campus and road traffic scenarios and videos taken from the CAVIAR database. The authors evaluated the effectiveness of their framework in terms of managing different queries with different privacy filters (only blurring and masking filters were considered).

The focus of the above frameworks is on developing a smart and flexible system of policies and user control mechanisms that would allow manageable access to privacy sensitive regions of surveillance video. Our framework significantly differ from such policy-based frameworks since we provide the means to evaluate and to choose an appropriate privacy protection tool suitable for a given surveillance system. After such privacy protection tools are deployed in the surveillance system, policy-based frameworks can follow to help managing user access.

3. EVALUATION FRAMEWORK

In this section, we discuss the objective evaluation framework for privacy protection filters, which is based on measurement and assessment of the privacy-intelligibility tradeoff. Typically, a privacy filter is a specific type of visual distortion with the aim to obscure personal details in video (leading to higher privacy protection) while retaining the general appearance of a person (keeping the same level of intelligibility). Hence, an ideal privacy filter would hide all possible personal details, resulting in the highest possible privacy protection, but it would keep intact all visual information necessary for the performance of a surveillance task. Therefore, to evaluate privacy protection tools, we propose to measure the amount of visual details as a metric of privacy and to measure an overall general shape, pattern, or behavior as metrics of intelligibility. Changes in these measurements would then constitute the privacy-intelligibility tradeoff.

To demonstrate this general framework, we assume the face to be a personal visual information that should be protected. Since we focus on objective evaluation, a possible application scenario where such evaluation would be useful is an automated surveillance system that mostly relies on video analysis algorithms. An example of such system could be an automatic people counting system on public transport. In such systems, all analysis and scene understanding operations, such as face recognition or face detection, will be performed by a software system in an automatic fashion, without the intervention of a human operator. Faces naturally carry significant privacy sensitive information, since they allow associating images of subjects with their identity. At the same time, faces carry intelligibility information as well: being able to detect faces in a scene is often the first step towards higher-level intelligibility tasks, such as detecting whether a subject is wearing a mask, sunglasses or other elements which could potentially signal a safety threat.

To protect the privacy of faces, a number of filters are commonly used both in video surveillance, broadcasting and printed media applications. Most of such filters belong to the category of naïve filters, i.e., basic image processing techniques exploited to corrupt the original image from the human visual quality point of view. Little work has been done so far to prove whether such filters are also able to protect privacy when automatic algorithms are applied instead.

We selected three common privacy protection filters: blurring with a Gaussian kernel, pixelization by averaging the privacy sensitive ROI in blocks of dimension b , and masking with a black rectangle. For each filter, we can vary its strength, to obtain different levels of image data corruption: the blur filter can be tuned by varying the size of the kernel. For pixelization, we can choose the size of the averaging block, obtaining an effect which is




Figure 1: Sample images with privacy protection filters applied. (a)(f)(k) originals. (b)(c)(d)(e) blur at 5, 11, 17 and 23 kernel size respectively. (g)(h)(i)(j) pixelization with 4, 10, 16 and 22 averaging block size respectively. (l)(m)(n)(o) masking with 0.1, 0.4, 0.7 and 1.0 opacity respectively.

equivalent to downsampling the image at different resolutions. Masking is typically applied by just replacing the privacy sensitive ROI with a black rectangle, but we can assume that changing the opacity of the black rectangle leads to different levels of privacy protection. Figure 1 shows a number of examples of the selected privacy filters applied to the same subject image with different degrees of strength.

In recent years, there has been a growing research trend proposing new privacy protection filters, specifically in the context of face identity protection, but most methods focus on issues such as reversibility, encryption, or application to particular video compression standards. As a result, the visual quality of such advanced filters does not differ significantly from the naïve ones. Moreover, such advanced filters are currently not being used in video surveillance systems, which are still in their infancy when comes to privacy protection and, at the moment, mostly focus on implementation of privacy protection policies and user management rather than privacy protection filters.

In the context of the described scenario, we propose face detection as a measure of intelligibility, and the failure of automatic face recognition as a measure of privacy protection. More precisely, we adopt the percentage of true positives achieved by the face detection algorithm, since we apply filters only inside face regions, and therefore we are not interested in the behavior of the algorithm elsewhere (i.e., false positives). Similarly, for face recognition we adopt rank 1 of the Cumulative Matching Characteristic (CMS), which represents the probability of a subject being recognized among the list of all subjects in the training gallery. As a reminder, we recall that in this work we are not evaluating the performance of a specific detection or recognition algorithm, but we are rather interested in the relative variations of such performance between original and privacy protected images.

As face detection we adopt Viola-Jones object detection framework, described in Ref. 8, while for recognition we test on three state-of-art algorithms, two of them based on global descriptors: Principal Component Analysis (PCA)⁹ and Linear Discriminant Analysis (LDA),¹⁰ and one based on local features (LBP).¹¹ The similarity measure adopted for PCA and LDA is based on Cosine Distance, while for LBP the histograms are compared with the Chi Square test. All these algorithms are provided as open source implementations in the computer vision library OpenCV.

3.1 Experimental procedure

Since we consider a fully automated system, it is correct to assume that the privacy filter is also applied by an automatic procedure. If face detection on the original image fails, then no privacy filter is applied for this particular image, which will considerably influence the final recognition scores at a later stage. Based on this consideration, we first apply face detection to the original image, and then, conditionally, proceed with applying filters and running face recognition on faces detected with face detection. First, we apply all selected privacy filters with different degrees of strength. Then, the detection on protected image and recognition on an original and protected image are performed, and results are recorded.

3.2 Datasets

	Resolution (pixels)	Pixels Per Face	Image Format
FERET	256 × 384	50	tif (uncompressed)
SCFace - Far	75 × 100	8	jpeg (compressed)
SCFace - Mid	108 × 144	18	jpeg (compressed)
SCFace - Close	168 × 224	27	jpeg (compressed)
ChokePoint	800 × 600	50 (best) - 20 (worst)	jpeg (compressed)

Table 1: Image quality data for each dataset used in the experiments. Pixels per face refers to the distance between eyes in pixels.

To make our evaluations practical, it is important to evaluate the selected privacy protection filters under different environmental conditions and on surveillance video with different resolutions and various face sizes. Therefore, we used three different publically available video surveillance datasets: SCFace dataset,¹³ ChokePoint surveillance dataset,¹⁴ and FERET dataset.¹⁵

Each dataset provides images which are heterogeneous in terms of resolution and quality. Table 1 briefly summarizes the image quality characteristics of each dataset and partitions used, in terms of provided image resolution and average pixels per face (ppf) measure, which is taken as the number of pixels between eyes. Resolution alone is not sufficient to describe image quality for face recognition, since subjects are acquired at different unknown distances from the camera in each dataset or partition.

The Facial Recognition Technology (FERET) dataset¹⁵ consists of a corpus of 14'051 gray scale images of persons with clearly visible faces and different facial expressions, under various environmental conditions, including illumination, orientation, appearance and age variations. As recommended by the authors of FERET dataset, we selected the frontal faces subset A as the training gallery for face recognition, and the frontal faces subset B as a probe for all evaluation tests. For recognition purposes, the original set of 1'196 subjects was reduced to a randomly selected gallery of 100 subjects. While face detection is applied on the original

non-cropped images, which include part of the background and body information as well, face recognition is performed after aligning and cropping the face images according to ground truth for eye position.

The Surveillance Cameras Face (SCface)¹³ database is a set of static images of human faces. Images were recorded in uncontrolled indoor environment using five video surveillance cameras of various qualities. Database contains 4'160 static images of 130 subjects. Each face in the dataset is acquired with three different distances, far, mid and close, which results in faces to vary in sizes greatly, thus, simulating a typical surveillance scenario when a person is moving towards or away from the camera at different distances. We conducted our experiments separately for each of the three subsets corresponding to these three different distances: far, mid, and close.

ChokePoint¹⁴ is a dataset designed for experiments in person identification and verification under real-world surveillance conditions. An array of three cameras placed above several portals capture subjects walking through. Faces in such sets will have variations in terms of illumination conditions, pose, sharpness, as well as misalignment due to automatic face localization and detection. The dataset consists of 64'204 labeled face images, essentially video frames, of 25 subjects.

4. EVALUATION RESULTS

Figure 2 demonstrates detection (graphs in the left column) and recognition (graphs in the right column) results for blurring, pixelization, and masking filters for FERET dataset. For recognition graphs, the results from the three different recognition algorithms are shown in the same plot together. For every graph in the figure, vertical axis corresponds to the accuracy of detection or recognition respectively, and horizontal axis corresponds to the different degrees of strength for a privacy filter. For the reference, accuracy for the original unprotected images is indicated by a horizontal straight dashed line.

The top two graphs (Figures 2a and 2b) illustrate the detection and recognition results for the blurring privacy filter with a kernel varying from 0 to a maximum of 23. From Figure 2a it can be noted that blurring filter has little effect on the detection accuracy, even when extreme kernel sizes are used. Therefore, this type of filter offers an intelligibility level comparable to the original image. On the other hand, recognition accuracy (see Figure 2b) is significantly decreased by the blurring filter. Especially, the accuracy of local features-based LBP algorithm dramatically drops from 0.9 to 0.35 even for a kernel size of 5. Global features-based face recognition methods, such as PCA and LDA, are slightly more robust to blurring, showing steady decrease from kernel sizes larger than 6. Figure 2b also shows a slight increase in accuracy for small blurring kernel, such as 3 or 5. This phenomenon is probably due to the fact that a small amount of blurring can effectively be regarded as denoising, which often has a positive effect on face recognition.

In evaluation of pixelization filter (Figures 2c and 2d), the size of pixelization blocks varies from 0 to a maximum of 22 pixels. We can observe that pixelization has little effect on face detection (the same as for blurring), with only extreme sizes of the pixelization blocks (larger than 16 pixels) degrading the accuracy significantly. For recognition (see Figure 2d), the trend is similar to blurring, but the decrease in accuracy is more rapid and, unlike blurring, with no gain in performance for small block sizes. This sudden drop in performance is even more evident for the LBP-based face recognition algorithm, since its accuracy becomes almost 0 for pixelization filter's block sizes larger than 4.

While pixelization and blurring show similar results, masking is different as it is evident from Figures 2e and 2f. Various opacity levels do not affect the detection at all except for the extreme value of 1.0, when a black box covers the whole face. Such results can be expected from the Viola-Jones object detector, since it relies on features computed as the differences between neighboring pixels. Adding or subtracting the same amount to the gray level of all pixels will not affect much the values such features. Naturally, when we replace the privacy sensitive image ROI with a black rectangle, the detection drops immediately to 0. This result proves that full masking, a widely used filter in video surveillance application, is incapable of preserving any intelligibility information. Recognition accuracy (see Figure 2f) for the masked images show a slower decreasing trend compared to blurring and pixelization. We have to apply masking with a minimum of 0.5 opacity to see some significant loss in recognition performance. Again, adding a small constant value to each pixel gray level results in a non significant alteration of the source data.

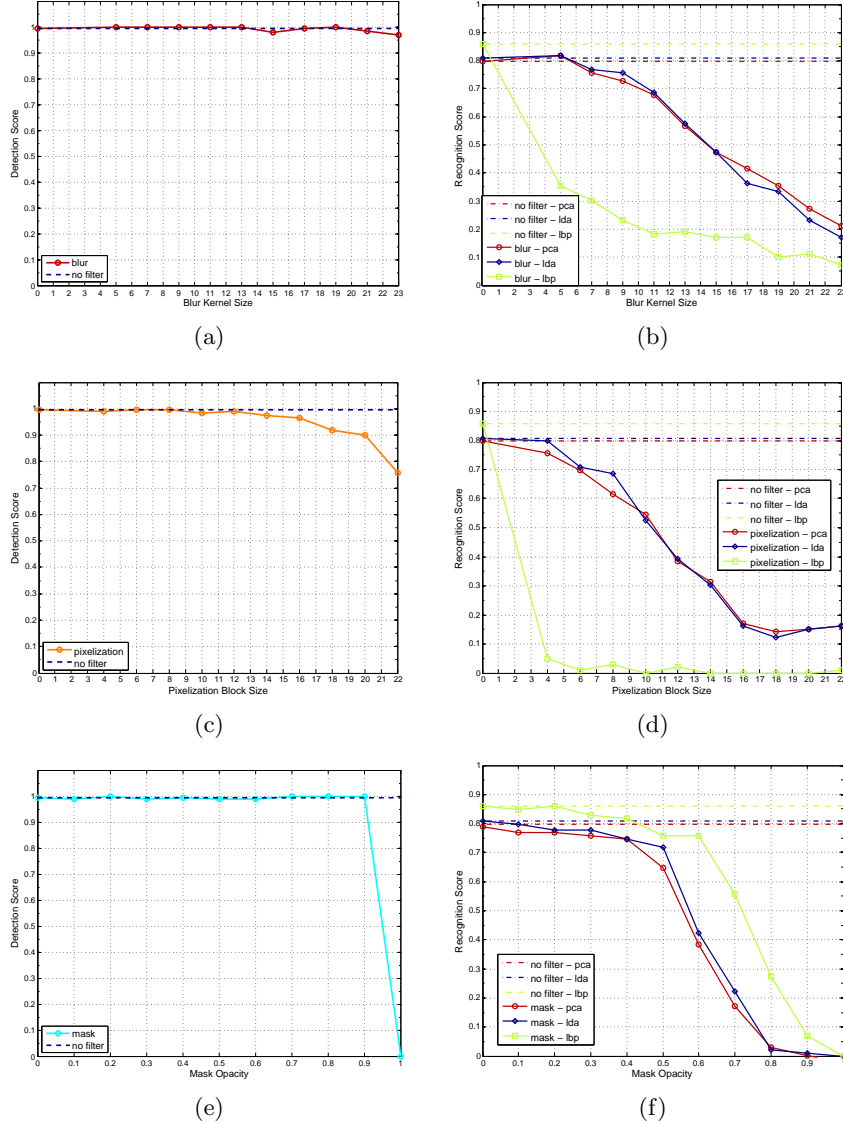


Figure 2: Detection and recognition results for varying filter strengths applied to FERET dataset. (a) and (b) correspond to blur filter, (c) and (d) to pixelization filter, and (e) and (f) to masking filter.

For the other two datasets, the results show similar trends as for FERET dataset, hence, we omitted such graphs to save space. Instead, we present the detection and recognition results as relative changes in detection and recognition accuracies in Figures 3, 4 and 5 for FERET, SCFace, and ChokePoint datasets respectively, which more clearly demonstrate privacy-intelligibility tradeoff. Intelligibility loss and privacy gain for each type of privacy filter in these figures is computed as following,

$$\text{IntelligibilityLoss} = D_o - D_p \quad (1)$$

$$\text{PrivacyGain} = (1 - R_o) - (1 - R_p) \quad (2)$$

where D indicates the detection accuracy, R the recognition accuracy, o indicates a result obtained on original unprotected images and p a result obtained on privacy protected images. We recall that privacy is inversely proportional to the recognition rate. As a reference, in the figures, the original detection accuracy (for unprotected images) is shown in parenthesis next to the intelligibility loss, and the original recognition accuracy (for unprotected images) is shown next to each privacy gain.

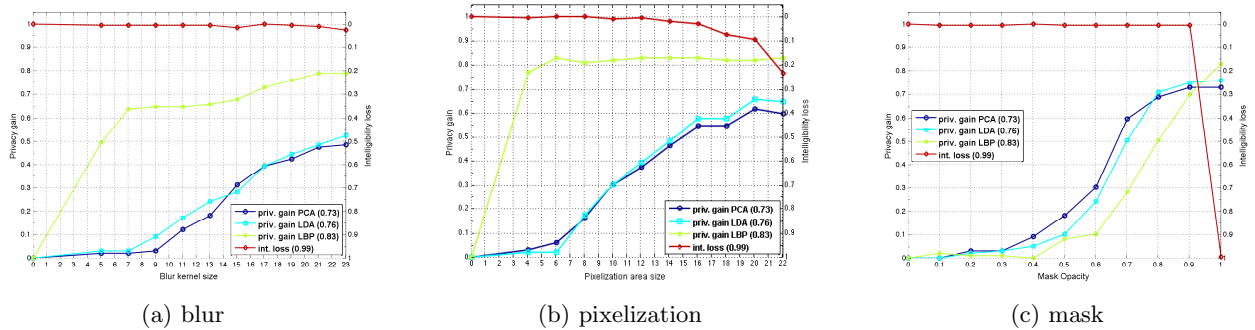


Figure 3: Intelligibility loss and privacy gains for FERET dataset. Values in parenthesis refer to original absolute detection accuracy for intelligibility loss and to original absolute recognition accuracies for privacy gains.


5. DISCUSSION

From privacy-intelligibility tradeoff results for FERET dataset in Figure 3, it is also evident that blurring and pixelization show similar behavior, letting us increase the level of privacy protection by making the filter stronger, while still keeping a very high intelligibility performance. In case of pixelization, for example, up to an 80% gain in privacy protection on LBP based recognition system and around 60% on PCA/LDA systems can be obtained, by only sacrificing about 10% of intelligibility. A pixelization block size equal to 20 can be used in practice. Pixelization appears to be slightly better than blurring at preserving privacy, but at the cost of affecting more the intelligibility for larger filter sizes.


The results in Figure 3c demonstrate that masking is effective only if applied with high but still partial opacity. When opacity is less than 0.5, it has little effect on recognition accuracy, leading to low privacy protection. When applied with full opacity, both recognition and detection fall to zero, offering a perfect privacy protection, but no intelligibility information whatsoever. However, if 0.9 level of opacity is chosen, detection accuracy remains as high as for smaller opacity level, while the drop in recognition accuracy is almost as significant as for full opacity. That means applying masking filter with opacity 0.9 leads to both high privacy protection and high intelligibility. This result suggests that partial opacity masking is the best filter to protect privacy in automatic video surveillance systems, which contradicts previous findings from subjective evaluations⁶ and confirms the observation that computer vision is different from human vision.¹²

Privacy-intelligibility tradeoff for SCFace dataset is presented in Figure 4 for different acquisition distances separately. This dataset is characterized by a very low resolution of the input images (see Table 1) and acquisition distances, which also reduce face sizes even further. For these reasons, Figure 4 shows extremely low recognition accuracy even for original images. It can imply that in a video surveillance system with video of such low resolutions, privacy is already protected, at least when no other advanced technology such as image super resolution and multi-frame integration has been implemented. This observation is confirmed by the results, since applying privacy filters does not lead to any gain in privacy on such low baseline. However, the detection and, therefore, intelligibility performance depends on the distance for SCFace dataset. For example, for the farthest camera, blurring with a small kernel of size 5 or pixelating with a block size of 2 already causes a 20% loss in intelligibility (see Figures 4a and 4b). Masking offers better intelligibility performance on far distances (see Figure 4c), allowing us to add up to 0.5 opacity without losing no more than 0.2 in the intelligibility. Performance on intelligibility progressively improves, as expected, when the distance between the faces and the camera is reduced. For example, we can blur with size of a kernel up to 7 on mid distances (see Figure 4d) and 11 on close distances (see Figure 4g), with a decrease no lower than 10% intelligibility. A general trend we observe, independently from the distance, is that pixelization causes the fastest drop in intelligibility, followed by blur and by masking, which offer a much smoother decrease.


The results for ChokePoint dataset in Figure 5 show similar trends as for FERET and SCFace datasets with results lying somewhere in between for those two datasets.




(a) blur - far distance




(b) pixelization - far distance




(c) mask - far distance




(d) blur - mid distance




(e) pixelization - mid distance




(f) mask - mid distance



(g) blur - close distance



(h) pixelization - close distance



(i) mask - close distance

Figure 4: Intelligibility loss and privacy gains for SCFace dataset. Values in parenthesis refer to original absolute detection accuracy for intelligibility loss and to original absolute recognition accuracies for privacy gains.

By analyzing the performance of each face recognition algorithm (see Figures 3, 4 and 5), we can observe that LBP is more sensitive to filtering than recognition methods based on global features (PCA and LDA). This result can be explained by the fact that Local Binary Pattern relies on local contrast measure when computing its features, which are strongly affected by pixelization and blurring filtering, since these filters can be considered as local-based distortions. Since, local feature based methods for face recognition are becoming increasingly popular in practice, due to their robustness to occlusion and variation in pose and illumination, then, a higher level of privacy protection can be obtained even when using a weaker privacy protection filter.

From Figures 3, 4 and 5, we can also note that performance of both recognition and detection algorithms are significantly affected by the original resolution of images (tested datasets have very different original resolutions, see Table 1), or from another perspective, by the distance between the subjects and the camera (see results in Figure 4 corresponding to different distances of SCFace dataset). Therefore, original video resolution should be taken into account when designing privacy protection filters for a practical video surveillance system.

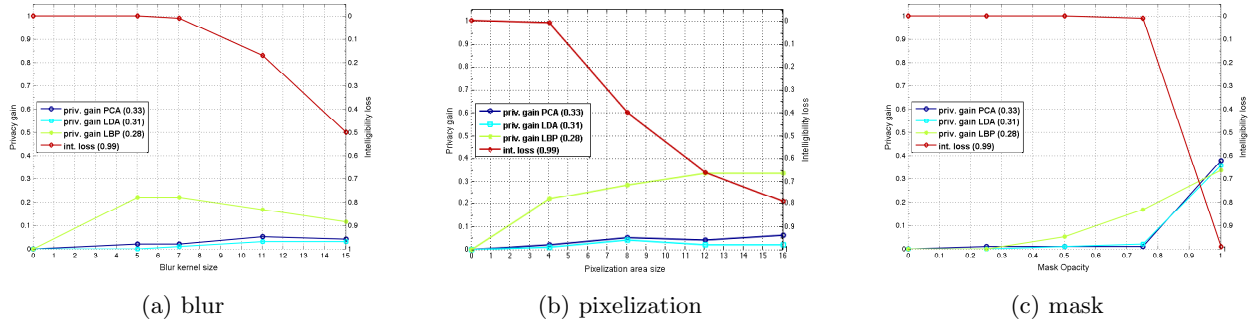


Figure 5: Intelligibility loss and privacy gains for ChokePoint dataset. Values in parenthesis refer to original absolute detection accuracy for intelligibility loss and to original absolute recognition accuracies for privacy gains.

6. CONCLUSION AND FUTURE WORK

In this paper, we have proposed an objective framework for evaluation of privacy protection filters by assessing the privacy-intelligibility tradeoff. We demonstrated the practical application of this framework using examples of face privacy protection that can be applicable in people counting automatic video surveillance systems. We compared three typical privacy protection filters blurring, pixelization, and masking by varying their degrees of strength. By using face detection as a metric for intelligibility and face recognition as a metric for privacy, we investigated on different datasets the privacy-intelligibility tradeoff. The results showed that resolution of the original face (alternatively, distance of a person from the camera) significantly affects the absolute accuracies of detection and recognition, however, the corresponding relative values show consistent trend. These evaluations suggest that masking filter with opacity 0.9 leads to the highest intelligibility (i.e., high detection accuracy) and the highest privacy protection (i.e., low recognition accuracy). This result contradicts previous findings from subjective evaluations, confirming an *a priori* observation that human perception differs from machine perception in video analytics.

The considered simplest filters: blurring, pixelization, and masking, though popular, do not satisfy requirements of desirable privacy filters in current surveillance systems. Therefore, as future work, the evaluations should also include more advanced privacy protection filters, such as scrambling or encryption based filters. Objective privacy evaluation framework should also be demonstrated in other surveillance scenarios such as object tracking or event detection.

ACKNOWLEDGMENTS

This work was conducted in the framework of the EC funded Network of Excellence VideoSense.

REFERENCES

- [1] Dufaux, F. and Ebrahimi, T., “Video surveillance using JPEG 2000,” in *[proc. SPIE Applications of Digital Image Processing XXVII]*, **5588**, 268–275 (Aug 2004).
- [2] Boulton, T. E., “PICO: Privacy through invertible cryptographic obscuration,” in *[IEEE Workshop on Computer Vision for Interactive and Intelligent Environments]*, 27–38 (Nov 2005).
- [3] Velardo, C., Araimo, C., and Dugelay, J.-L., “Synthetic and privacy-preserving visualization of video sensor network outputs,” in *[5th ACM/IEEE International Conference on Distributed Smart Cameras (ICDSC’11)]*, 1–5 (Aug 2011).
- [4] Korshunov, P. and Ebrahimi, T., “Using warping for privacy protection in video surveillance,” in *[18th International Conference on Digital Signal Processing (DSP)]*, *DSP’13* (June 2013).
- [5] Korshunov, P., Araimo, C., De Simone, F., Velardo, C., Dugelay, J., and Ebrahimi, T., “Evaluation of visual privacy filters impact on video surveillance intelligibility,” in *[2012 Fourth International Workshop on Quality of Multimedia Experience (QoMEX)]*, 150–151 (July 2012).

- [6] Korshunov, P., Araimo, C., De Simone, F., Velardo, C., Dugelay, J., and Ebrahimi, T., “Subjective study of privacy filters in video surveillance,” in [*2012 IEEE 14th International Workshop on Multimedia Signal Processing (MMSP)*], 378–382 (Sept. 2012).
- [7] Korshunov, P., Cai, S., and Ebrahimi, T., “Crowdsourcing approach for evaluation of privacy filters in video surveillance,” in [*Proceedings of the ACM Multimedia 2012 Workshop on Crowdsourcing for Multimedia*], *CrowdMM’12*, 35–40 (Oct. 2012).
- [8] Viola, P. and Jones, M. J., “Robust real-time face detection,” *Int. J. Comput. Vision* **57**, 137–154 (May 2004).
- [9] Turk, M. and Pentland, A., “Face recognition using eigenfaces,” in [*Computer Vision and Pattern Recognition, 1991. Proceedings CVPR ’91., IEEE Computer Society Conference on*], 586–591 (1991).
- [10] Belhumeur, P., Hespanha, J., and Kriegman, D., “Eigenfaces vs. fisherfaces: recognition using class specific linear projection,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **19**(7), 711–720 (1997).
- [11] Ahonen, T., Hadid, A., and Pietikainen, M., “Face description with local binary patterns: Application to face recognition,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **28**(12), 2037–2041 (2006).
- [12] Korshunov, P. and Ooi, W. T., “Video quality for face detection, recognition, and tracking,” *ACM Trans. Multimedia Comput. Commun. Appl.* **7**, 14:1–14:21 (Sept. 2011).
- [13] Grgic, M., Delac, K., and Grgic, S., “Scface — surveillance cameras face database,” *Multimedia Tools Appl.* **51**, 863–879 (Feb. 2011).
- [14] Wong, Y., Chen, S., Mau, S., Sanderson, C., and Lovell, B. C., “Patch-based probabilistic image quality assessment for face selection and improved video-based face recognition,” in [*Computer Vision and Pattern Recognition Workshops (CVPRW)*], (2011).
- [15] Phillips, P. J., Moon, H., Rizvi, S. A., and Rauss, P. J., “The feret evaluation methodology for face-recognition algorithms,” *IEEE Trans. Pattern Anal. Mach. Intell.* **22**, 1090–1104 (Oct. 2000).
- [16] Newton, E., Sweeney, L., and Malin, B., “Preserving privacy by de-identifying face images,” *IEEE Trans. on Knowledge and Data Engineering* **17**, 232–243 (Feb 2005).
- [17] Dufaux, F. and Ebrahimi, T., “A framework for the validation of privacy protection solutions in video surveillance,” in [*in Proceedings of IEEE International Conference on Multimedia & Expo (ICME 2010)*], (July 2010).
- [18] Zhang, W., Cheung, S., and Chen, M., “Hiding privacy information in video surveillance system,” in [*in Proc. IEEE International Conference on Image Processing*], (Sep 2005).
- [19] Cheung, S. S., Paruchuri, J. K., and Nguyen, T. P., “Managing privacy data in pervasive camera networks,” in [*in Proc. IEEE International Conference on Image Processing*], (Oct 2008).
- [20] Park, S.-W., Han, J. W., and Shin, S.-U., “Secure service mechanism of video surveillance system based on h.264/svc,” in [*International Conference on Information Technology and Multimedia (ICIM 2011)*], 1–4 (nov. 2011).
- [21] Wickramasuriya, J., Datt, M., Mehrotra, S., and Venkatasubramanian, N., “Privacy protecting data collection in media spaces,” in [*Proceedings of the 12th annual ACM international conference on Multimedia*], 48–55 (Oct. 2004).
- [22] Ebrahimi, T., Abdeljaoued, Y., i Ventura, R. F., and Escoda, O. D., “MPEG-7 camera,” in [*Proceedings of International Conference on Image Processing (ICIP 2001)*], **3**, 600–603 (2001).
- [23] Senior, A., Pankanti, S., Hampapur, A., Brown, L., Tian, Y.-L., Ekin, A., Connell, J., Shu, C., and Lu, M., “Enabling video privacy through computer vision,” *IEEE Security and Privacy* **3**, 50–57 (May 2005).
- [24] Aved, A. J. and Hua, K. A., “A general framework for managing and processing live video data with privacy protection,” *Multimedia Syst.* **18**(2), 123–143 (2012).