Subjective Study of Privacy Filters in Video Surveillance

P. Korshunov [#1], C. Araimo [*2], F. De Simone [#3], C. Velardo [*4], J.-L. Dugelay [*5], and T. Ebrahimi [#6]

[#] *Multimedia Signal Processing Group – MMSPG,*
Institute of Electrical Engineering – IEL,
École Polytechnique Fédérale de Lausanne – EPFL,
CH-1015 Lausanne, Switzerland
[1] pavel.korshunov@epfl.ch
[3] francesca.desimone@epfl.ch
[6] touradj.ebrahimi@epfl.ch

[*] *Multimedia Department, EURECOM*
2229 Route des Crêtes, 06560 Valbonne, France,
[2] araimo@eurecom.fr
[4] velardo@eurecom.fr
[5] dugelay@eurecom.fr

***Abstract*—Extensive adoption of video surveillance, affecting many aspects of the daily life, alarms the concerned public about the increasing invasion into personal privacy. Therefore, to address privacy issues, 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. In this paper, we propose a subjective evaluation methodology to analyze the tradeoff between the preservation of privacy offered by these tools and the intelligibility of activities under video surveillance. As an example, the proposed method is used to compare several commonly employed privacy protection techniques, such as blurring, pixelization, and masking applied to indoor surveillance video. The results show that, for the test material under analysis, the pixelization filter provides the best performance in terms of balance between privacy protection and intelligibility.**

I. Introduction

The alarming rate at which video surveillance is being adopted has raised concerns among public and motivated development of privacy protection tools. Typical techniques (i.e. filters) 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], obscuring [3] and complete removal of the body silhouettes [4], anonymization [5], 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[1] workshops and grand challenges) are available to test the robustness, accuracy, and efficiency of video analytics for surveillance, little attention has been devoted to the privacy aspects. Therefore, a formal methodology for evaluation of the privacy protection filters is needed.

As the typical end user of privacy filters is a human subject, the ground truth required for evaluating their performance is also subjective. In this paper, we propose 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. We focus on several typical use cases of benign and suspicious behavior in indoor video surveillance, and apply commonly used privacy protection filters, such as blurring, pixelization, and masking to obscure the privacy-sensitive regions. Then, we ask human subjects to evaluate the resulting videos in terms of degree of privacy preservation and intelligibility of the surveillance events. The proposed evaluation method allows to identify the weaknesses of existing privacy protection tools and provide a reference for evaluation of other techniques.

The rest of the paper is organized as follows. In Section II, we describe the evaluation methodology with underlying dataset and evaluation protocol. In Section III, we focus on the evaluation criteria, and in Section IV, we discuss the results of the subjective evaluation. We conclude the paper with Section V.

II. Evaluation methodology

This section describes the evaluation methodology that is designed for effectiveness assessment of the various visual filters to protect privacy of individuals on one hand, and their

MMSP'12, September 17-19, 2012, Banff, Canada.
???-?-????-????-?/10/$??.?? *©2012 IEEE.*

[1] IEEE International Workshop on Performance Evaluation of Tracking and Surveillance (PETS)

impact on the intelligibility of the surveillance task on the other.

A. Use cases and underlying database

Privacy and surveillance are both heavily context dependent concepts. Therefore, any evaluation methodology should take into account the context in which the surveillance task is performed. In this paper, we focus on a simple use case, namely, a monitoring situation, without recording, where an observer (test subject) watches a video of an indoor scene under surveillance with a single standard definition camera. In the monitored scene, individuals move in front of the camera, either behaving normally, or acting abnormally.

To evaluate this use case, we have built a dataset consisting of 9 different video sequences with a duration of 10 seconds each. Different indoor video surveillance scenarios were considered, such as a person walking towards and away from the camera (normal scenario), blinking into the camera (suspicious), and wearing hat, sunglasses, or scarf around the mouth (suspicious) to hide personal identity. Table I provides a short description of each video sequence in the database.

TABLE I: Description of the video sequences used for the evaluation

Seq. 1	White male, sunglasses, walks away and towards the camera
Seq. 2	White female, walks towards the camera, blinks three times
Seq. 3	Asian male, glasses, walks in from the right side, blinks three times to the camera
Seq. 4	White male, walks toward the camera, blinks three times
Seq. 5	Asian female, walks towards the camera, blinks three times
Seq. 6	White female, walks toward the camera, blinks three times
Seq. 7	Asian male, glasses, walks toward the camera, blinks three times
Seq. 8	White female, walks toward the camera
Seq. 9	White female, wears scarf around her face, walks toward the camera

To each video sequence in the dataset, a semi-automatic segmentation and tracking algorithm was applied in order to obtain a binary mask[2], identifying a foreground object of interest, which not only plays a certain role in the understanding of the situation under surveillance, but also may contain potentially privacy sensitive information. Different privacy protection filters were then applied to the extracted foreground objects. Blurring, pixelization, and masking (black foreground shape covering the region of interest) privacy filters were selected (see examples in Figure 1) to generate different versions for each video sequence. Thus, a total of 27 processed video sequences were produced and used in the subjective evaluation, as described in the next section.

B. Evaluation protocol

The goal of the subjective evaluation was to assess whether the detection of the normal or abnormal behaviors in the scene was possible, while various privacy protection filters were applied. At the same time, the effectiveness of privacy protection was assessed, as the identities of the individuals in the sequences might have been hidden. Particularly, each subject was asked to watch a video sequence and then answer to questions presented in Table II.

TABLE II: Questions asked during the assessment

1. What is the gender of the person?	☐ Female	☐ Male	☐ I don't know
2. What is the race of the person?	☐ White	☐ Asian	☐ I don't know
3. Does the person wear glasses?	☐ Yes	☐ No	☐ I don't know
4. Does the person wear sunglasses?	☐ Yes	☐ No	☐ I don't know
5. Does the person wear a scarf?	☐ Yes	☐ No	☐ I don't know
6. Does the person blink into the camera?	☐ Yes	☐ No	☐ I don't know

An important issue to resolve was the memory effect during viewing, when observation of a video could potentially affect the evaluation of a following video. In our case, the main concern came from interactions between different versions of the same video, since different details could be visible in different video of the same scene obfuscated by different privacy filters. For instance, observation of a blurred video could provide information otherwise invisible in the pixelated version of the same video. Consequently, if the former precedes the latter, the memory effect could affect the evaluation of the latter.

To avoid such memory effect in the assessment of the privacy protection and the task performed in video surveillance, each subject was shown each of the 9 contents only once. To insure that, the 27 processed video sequences were divided into three separate sessions designated as A, B, and C, with each session containing 9 sequences including all the different contents. Furthermore, every session contained an equal number of blurred, pixelated, and masked video. Table III illustrates how video sequences were divided into the these three sessions.

Each session lasted about 5 minutes and was attended by a different group of 12 subjects, thus, overall 36 subjects took part in the evaluation. In such an arrangement, every subject had a balanced overview of the used privacy filters, which helps avoiding bias in the results. The subjects were naive viewers of mixed gender (almost equally distributed) and various nationalities. Subjects' age was in range from middle twenties up to late forties.

TABLE III: Arrangement order of the filtered video sequences into evaluation sessions A, B, and C

Seq.	Blurring	Pixelization	Masking
Seq. 1	A	C	B
Seq. 2	B	A	C
Seq. 3	C	B	A
Seq. 4	A	C	B
Seq. 5	B	A	C
Seq. 6	C	B	A
Seq. 7	A	C	B
Seq. 8	B	A	C
Seq. 9	C	B	A

[2]MIT annotation tool: http://people.csail.mit.edu/celiu/motionAnnotation/

(a) Original, no filters

(b) Blurring

(c) Pixelization

(d) Masking

Fig. 1: An example of video sequence (Seq. 9 from Table I) with privacy filters applied

Each video sequence was displayed after a short message informing the subject that the evaluation for that sequence was imminent. Subjects were given 25 seconds to respond the questions in Table II by ticking the corresponding check boxes in the scoring sheet. They were instructed to give definitive answers (such as "Yes" or "No") only if reasonably certain about the answer, and answer "I don't know" in all other cases. The same procedure was repeated for each video sequence until the end of the session when a message informed the test subjects that the session was over. Figure 2 displays the photo demonstrating the test lab and how the subjective tests were performed.

Fig. 2: A photo of one of the subjective test sessions

III. Evaluation criteria

Given the context dependent nature of privacy and intelligibility, in the surveillance scenario under consideration, the first three questions from the Table II were assumed to be relevant to privacy and the last three questions to intelligibility. Information about gender, somatic traits, and glasses (first three questions) are privacy related. These characteristics do not carry anything unusual, given the surveillance scenario, while they can be used to identify people in the indoor environment, and they can be discriminated against based on

these features. On the other hand, blinking three times into the camera, which looks like sort of a code (at least, it's an unusual behavior), sunglasses worn indoor (possibly for hiding eyes), and scarf around the face (to hide the identity) are considered unusual and alarming, since either of these characteristics are not typical for the indoor environment. These unusual features therefore are set as related to intelligibility and should be visible to the observers.

Therefore, the following criteria were used for understanding how well a given filter protects privacy. If an observer correctly answers the privacy related question for a given video sequence and privacy protection filter, the privacy is not protected in this case. Incorrect answer or no answer (option "I don't know") would mean that the privacy is preserved. For intelligibility, on the other hand, a correct answer to the corresponding question would mean that surveillance task can be performed successfully, while incorrect or uncertain answers would lead to the failure of recognizing an important unusual event.

Such tradeoff between privacy and intelligibility can be used to compare different privacy protection techniques and understand how these techniques perform, given various video contents.

If an observer correctly answers to the privacy related question, the privacy value is 0, since the privacy was not protected in this case. Incorrect answer or no answer (option "I don't know") yield 1. Then, the average privacy score of all three privacy related questions across all test subjects was computed for each type of filter and each video sequence.

IV. Evaluation results

For each privacy filter, the aggregated results are illustrated on a two dimensional space in Figure 3, with the amount of privacy preservation and the degree of recognition of activities under surveillance (i.e., intelligibility), as vertical and horizontal axes respectively. The privacy score is ranging from 0 (no privacy protection) to 1 (fully protected), which is the average of the scores of the privacy related questions from the test subjects, as described in the previous section. Each point in the figure corresponds to a different video sequence and a different privacy protection filter. Points corresponding to a privacy filter are marked with distinguishing point-style.

The best privacy preserving filter would be a blacked out camera with no video feed, but, in such case, there would be no surveillance possible and intelligibility would be zero. Therefore, a usable privacy protection filter should have a balance between privacy and intelligibility. In an ideal situation, the evaluation scores for such filter would lie in the top right corner of the tradeoff graph, having the highest values of privacy and intelligibility. In practice, however, it means that a filter with points close to the 45° line provides the best balance between privacy and intelligibility.

Figure 3, with evaluation scores of the typical privacy filters, demonstrates that blurring filter yields the highest intelligibility while providing the lowest privacy protection. Not surprisingly, the masking filter shows the highest privacy protection, while having the lowest intelligibility, since a person from the video sequence is replaced with the black boundary. However, the highest privacy score for the masking filter is still below 0.8, which means that at least 20% of the answers to the privacy questions were still correct. By looking into details, we noticed that the largest number of correct answers for masking filter is to the gender question, which is because people can recognize gender by the shape of a person in the video. Therefore, the shape of persons' masks should be distorted to hide the actual shape of the person. A surprising result shows pixelization filter demonstrating high privacy protection while still yielding high degree of the activities recognition, which makes it the filter with the best balance of privacy and intelligibility.

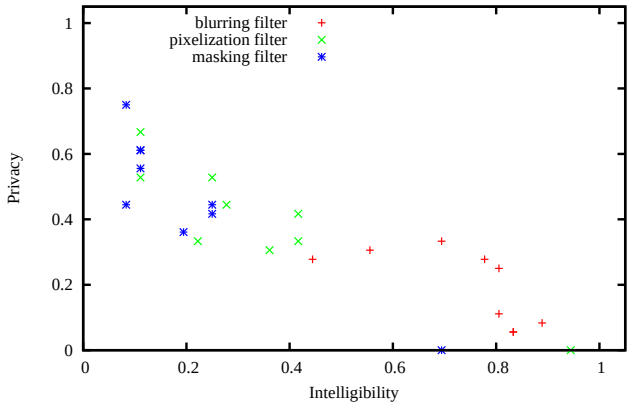

Fig. 3: Intelligibility vs. privacy for different filters

It can be noted in Figure 3 that one video sequence demonstrates an odd results for every filter, having the smallest value of privacy and a significantly high intelligibility (the points of each different color with the lowest privacy scores). In this video, the face of the person walking from a distance was left visible (unprotected by a filter) just for a couple of frames, which immediately rendered privacy protection filters useless. This video sequence indicates that even a slight inaccuracy or inconsistency in the way the filters are applied can lead to the complete loss of the privacy protection effort.

Figure 4 demonstrates the effect of different privacy protection filters on the uncertainty and incorrectness in the answers of the test subjects. Uncertainty axis reflects the average normalized amount of "I don't know" answers, which were given to both privacy and intelligibility questions. Incorrectness is computed as normalized average of wrong answers, when subjects were certain but wrong. Each point on the graph corresponds to one video sequence distorted by one privacy protection filter. This figure shows masking filter yielding the largest uncertainty, while blurring filter results in the largest false positive (incorrectness). Such unbalance indicates that blurring filter is less applicable in the surveillance scenarios with little tolerance for false positives. The masking filter on the other hand would be better in a typical surveillance system when some uncertainty can be tolerated, i.e., an uncertain observation can be checked via other means, such as an

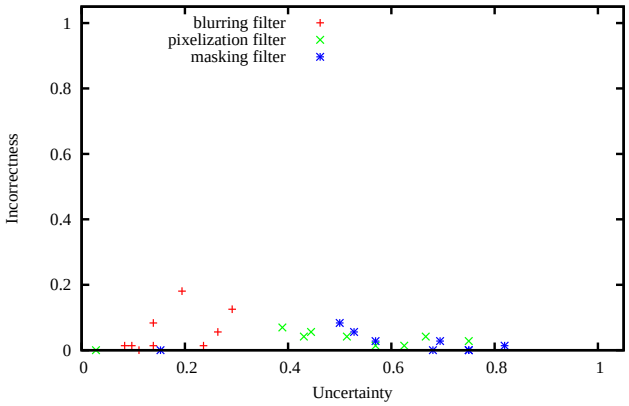

Fig. 4: Uncertainty vs. incorrectness for different filters

additional security check, but false positive is required to be low.

V. Conclusion and Future Work

This paper defines a methodology for evaluation of privacy protection tools for video surveillance. In the proposed evaluation protocol, we focus on two important aspects: (i) how much of the privacy is protected by such 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 high intelligibility. Masking filter, on the other hand, demonstrates the highest privacy protection with the lowest false positives. Future work includes extending the set of evaluation questions to identify other tradeoffs in privacy protection. The effect of applying protection tools with different levels of strength also need to be evaluated. We also plan to extend the dataset to include both more content and several additional privacy filtering tools. The complete dataset used in this paper will be made available to public.

Acknowledgment

The authors would like to thank Zdenek Svachula for help in generating some of the masks used in this work and with the score sheets. This work was conducted in the framework of the EC funded Network of Excellence VideoSense.

References

[1] F. Dufaux and T. Ebrahimi, "Video surveillance using JPEG 2000," in *proc. SPIE Applications of Digital Image Processing XXVII*, vol. 5588, Denver, CO, Aug 2004, pp. 268–275.

[2] T. E. Boult, "PICO: Privacy through invertible cryptographic obscuration," in *IEEE Workshop on Computer Vision for Interactive and Intelligent Environments*, Lexington, KY, Nov 2005, pp. 27–38.

[3] S.-C. S. Cheung, M. V. Venkatesh, J. K. Paruchuri, J. Zhao, and T. Nguyen, *Protecting privacy in video surveillance*. Springer-Verlag, 2009, ch. Protecting and Managing Privacy Information in Video Surveillance Systems, pp. 115–128.

[4] J. Wickramasuriya, M. Datt, S. Mehrotra, and N. Venkatasubramanian, "Privacy protecting data collection in media spaces," in *Proceedings of the 12th annual ACM international conference on Multimedia (ACMM'04)*, New York, NY, USA, Oct 2004, pp. 48–55.

[5] C. Velardo, C. Araimo, and J.-L. Dugelay, "Synthetic and privacy-preserving visualization of video sensor network outputs," in *5th ACM/IEEE International Conference on Distributed Smart Cameras (ICDSC'11)*, Ghent, Belgium, Aug 2011, pp. 1–5.