

EVALUATION OF VISUAL PRIVACY FILTERS IMPACT ON VIDEO SURVEILLANCE INTELLIGIBILITY

*P. Korshunov*¹, *C. Araimo*², *F. De Simone*¹, *C. Velardo*², *J.-L. Dugelay*², and *T. Ebrahimi*¹

¹Multimedia Signal Processing Group, Institute of Electrical Engineering, EPFL, Switzerland, {pavel.korshunov, francesca.desimone, touradj.ebrahimi}@epfl.ch.

²Multimedia Department, EURECOM, France, {araimo, velardo, dugelay}@eurecom.fr.

ABSTRACT

Since privacy issues are becoming important with growth of the video surveillance, many tools are proposed for protection of personal privacy in the video. However, little is understood regarding the effectiveness of such tools and their effect on the underlying surveillance tasks. In this paper, we propose a subjective evaluation methodology that compares several popular privacy protection techniques applied to typical indoor surveillance video. We identify and analyze the tradeoff between the privacy preservation of these tools and the intelligibility of activities in the resulted surveillance video.

Index Terms— Privacy protection tools, video surveillance, subjective evaluation, methodology.

1. INTRODUCTION

The alarming rate, with which video surveillance is being adopted daily, has raised concerns of the public and demanded the development of privacy protection tools. Typical techniques that are used for obscuring personal information in the video in order to preserve privacy include blurring and pixelization of the sensitive video regions or covering them with a black box. More advanced privacy protections techniques have also been developed recently, such as scrambling [1] and anonymization [2].

However, there is a noticeable lack of methods for assessing the performance of privacy protection tools and their impact on the surveillance task. While many evaluation protocols and tools (most notable ones are developed as part of PETS workshops and datasets) are available for testing video analytics to robustly, efficiently, and accurately perform the surveillance task, little attention was paid to the privacy aspect of the surveillance. Therefore, a formal methodology for evaluation of the privacy protection filters is needed.

Since the typical end user of the privacy filters is human, the ground truth required for evaluation of the privacy protection filters performance is subjective. Therefore, in this

paper, we propose a subjective evaluation methodology. We focus on several typical use cases of benign and suspicious behavior in indoor video surveillance, and apply blurring, pixelization, and masking filters to obscure the privacy-sensitive regions. Then, we ask the human subjects to rate the resulted videos in terms of the degree of privacy preservation and the intelligibility of the surveillance events. The results of the evaluation allow us to identify the weaknesses of the existing privacy protection tools and provides a ground truth for the evaluation of future techniques.

2. USE CASES AND DATABASE

Privacy and surveillance are both heavily context dependent and therefore any evaluation methodology should take into account the issues relevant to the context, in which the task under evaluation 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. Individuals move in front of the camera, either behaving normally, or acting abnormally. The goal of the evaluation is to detect normal or abnormal behaviors in the scene from the video sequence, while various privacy protection filters have been applied to the latter, and, at the same time, assess the effectiveness of privacy protection applied.

Therefore, we have designed a specific dataset consisting of 9 different video sequences (the duration of 10 seconds each), representing different indoor video surveillance scenarios, such as a person walking towards and away from the camera (normal scenario), blinking into the camera (suspicious), and wearing sunglasses or scarf around the mouth (suspicious) to hide the personal identity.

To each video sequence in the dataset, a semi-automatic segmentation and tracking algorithm is applied in order to obtain a binary mask¹, identifying a foreground object of interest, which not only plays a certain role in the understanding

Authors are partners of EU NoE VideoSense <http://videosense.eu/>.

¹MIT annotation tool: <http://people.csail.mit.edu/ceiliu/motionAnnotation/>

of the specific situation under surveillance, but also may contain potentially privacy sensitive information. Different privacy protection filters are then applied to the extracted foreground objects. Blurring, pixelization, and masking (black foreground shape covering the ROI) privacy filters were selected to generate different versions for each video sequence resulting in 27 video sequences in total.

3. EVALUATION METHODOLOGY AND RESULTS

A total number of $N = 36$ subjects were asked to view a subset of the database described above. An important issue to resolve was the memory effect during viewing, when observation of a video could potentially affect the evaluation score of the next video. For instance, observation of a blurred video could provide information otherwise invisible in a masked version of the same video. Therefore, to avoid the memory effect in the assessment, subjects were shown the contents of each video sequence only once, but with different filter applied. Hence, subjects were equally divided into three separate sessions designated as A, B, and C, with each session containing 9 different sequences (27 in total). Every session also contained an equal number of blurred, pixelated, and masked video. This arrangement insures that every subject has a balanced overview of the used privacy filters helping to avoid bias in the results.

Each session took about 5 minutes, during which test subjects assessed the video sequences. Each video sequence was displayed to a subject after a short message informing that the start of scoring for that sequence was imminent. Test subjects then respond to the questions about gender and race of the person in the video, what he/she is wearing such as glasses, scarf, and sunglasses, and whether the person blinks into the camera. Subjects were provided with 25 seconds to respond by ticking the corresponding checkboxes in the scoring sheet. They were instructed that they should give a definitive answers (such as “Yes” or “No”) only if they 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.

Given the context dependent nature of privacy and intelligibility, in the surveillance scenario under consideration, the questions related to gender, race, and whether a person wears glasses (personal item that can be used to identify someone) were assumed to be relevant to privacy. The questions about scarf around the face, sunglasses, and blinking are considered as relevant to intelligibility.

For each privacy filter, the aggregated results are presented in Figure 1 demonstrating the tradeoff between privacy preservation and intelligibility. The privacy scores were computed as follows. If an observer correctly answers to the privacy related question, the privacy value is considered to be 0, since the privacy was not protected in this case. Incorrect

answer or no answer (option “I don’t know”) yields 1. Then, the average privacy score, with 0 corresponding to no privacy protection and 1 to full protection, of all three privacy related questions across all test subjects was computed for each type of filter and each video sequence. The computation of intelligibility scores was done similarly, but three intelligibility related questions were taken into account instead.

Figure 1 demonstrates that blurring filter yields the highest intelligibility while providing the lowest privacy protection. 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 correct. A surprising result shows pixelization filter demonstrating high privacy protection while still yielding high degree of intelligibility.

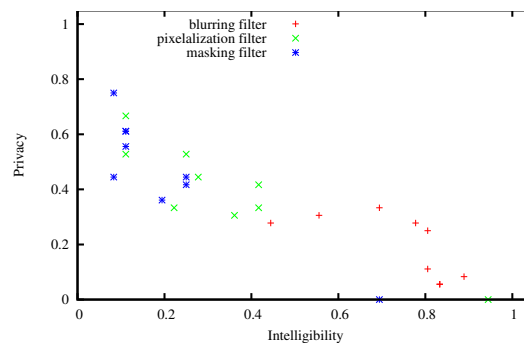


Fig. 1: Intelligibility vs. privacy for different filters.

4. CONCLUSION AND FUTURE WORK

This paper is a work in progress aiming to develop an extensive methodology for evaluation of privacy protection tools for video surveillance. In the proposed evaluation protocol, we focus on the two important aspects: (i) how much of the privacy is protected by such tool and (ii) how much it degrades the quality of the underlying surveillance task. We are extending the set of evaluation questions to identify other tradeoff in the privacy protection task. We have also created a dataset for evaluation of the privacy protection tools that includes masks of the foreground objects and results of several filtering tools. The complete dataset will be available for download and use in research.

5. REFERENCES

- [1] F. Dufaux and T. Ebrahimi, “Video surveillance using JPEG 2000,” in *proc. SPIE Applications of Digital Image Processing XXVII*, Denver, CO, Aug 2004, vol. 5588, pp. 268–275.
- [2] 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.