TUB-IRML at MediaEval 2014 Visual Privacy Task: Privacy Filtering through Blurring and Color Remapping

Dominique Maniry, Esra Acar, Sahin Albayrak
DAI Laboratory, Technische Universität Berlin
Ernst-Reuter-Platz 7, TEL 14, 10587 Berlin, Germany
dmaniry@cs.tu-berlin.de, esra.acar@tu-berlin.de, sahin.albayrak@dai-labor.de

ABSTRACT

This paper describes the participation of the TUB-IRML group to the MediaEval 2014 Visual Privacy task. We present a method for privacy protection of individuals in surveillance videos. In order to achieve this, our method obscures both shape and appearance of identity-related regions through blurring and color remapping. The intelligibility is preserved by displaying edges and anomalous events are hinted at by special colors. The experimental results obtained on surveillance videos show that our method considerably outperforms other participating teams in terms of privacy score. However, the drawback is that the results in terms of intelligibility are below average.

1. INTRODUCTION

The MediaEval 2014 Visual Privacy Task addresses the problem of privacy protection in video surveillance, which is gaining more and more importance due to concerns raised about the privacy of monitored individuals. Detailed description of the task, the dataset and the evaluation methodologies are given in the paper by Badii et al. [1]. As part of the MediaEval 2014 Visual Privacy Task, our privacy filter is evaluated using the Privacy Evaluation Video Dataset (PEViD) [2].

In the context of this task, we propose a simple but effective privacy filter which aims not only at obscuring facial identity, but also at protecting other identity revealing features such as accessories and clothing. This is achieved by obscuring both shape and appearance of identity-revealing regions in videos.

2. THE PROPOSED METHOD

The application of our privacy filter is a four-step process. First, we convert each frame into grayscale and apply a Gaussian blur to all privacy-related regions of the frame. The intensity of the blurring can be controlled using three different blur levels (obtained by varying the standard deviation of the Gaussian kernel) for regions labeled with *low*, *medium* and *high* privacy requirements.

As a second step, the pixel values are quantized to a given number of values (e.g., 8). These values are remapped to either a green or red color with the corresponding pixel intensity, so that the relation between light and dark regions remains same. The red color is used whenever an anomalous event (e.g., fighting, stealing or dropping a bag) happens. In other cases (i.e., non-anomalous), the individuals are shown in a green color. The aim of this red-green coloring is to enable human operators to focus on any event which requires particular attention. The second step removes, depending on the blur level and number of colors, most of the shape and appearance information that could potentially reveal a person's identity, gender or ethnicity, while preserving their movements and actions.

Figure 1: A walking person is shown as a green silhouette.

In the third step, the obscured image $\hat{I}(x, y)$ is blended back into the original frame $I(x, y)$ to create a smooth transition between obscured regions and the background. The blending mask $mask(x, y)$ is a binary image where annotated regions have a value of 1 and remaining regions have a value 0. The smoothing is achieved by applying a Gaussian blur to the blending mask. The result is:

$$result(x, y) = mask(x, y) \cdot \hat{I}(x, y) + (1 - mask(x, y)) \cdot I(x, y)$$

In the final step, we target a better intelligibility by including some shape information in the image. The obscured regions are overlaid with edges obtained with Canny Edge detection. Edges in regions with a *high* privacy requirement (i.e., faces) are discarded in order not to reveal identity through the edges of facial features. The remaining edges are emphasized using morphological dilation with a 3x3 circle as structuring element.

MediaEval 2014 Workshop, October 16-17, 2014, Barcelona, Spain

Figure 2: The silhouettes of two people fighting. The red color indicates that an anomalous event is happening.

Table 1: User study results for *Stream 1*, *2* and *3*.

	Stream 1	Stream 2	Stream 3
Our Intelligibility	73.4%	67.7%	60.7%
Median	74.9%	79.3%	69.6%
Our Privacy	69.4%	80.0%	78.7%
Median	50.2%	46.5%	41.7%
Our Pleasantness	22.7%	46.0%	52.6%
Median	24.8%	69.6%	59.7%

3. RESULTS AND DISCUSSION

Our submitted run was created using a constant blur level of 14 for all three privacy levels. The number of colors is 8. This choice of parameters favors privacy over intelligibility. The submissions of eight teams have been evaluated in a user study. The user study has been conducted with three different groups (i.e., streams). *Stream 1* represents 230 crowd-sourcing workers, *Stream 2* is 65 people working at *Thales* (mainly in Research&Development – R&D) and *Stream 3* has 59 participants from sectors including R&D, data protection and law enforcement from all around the world. The results for our method and the median across all 8 submissions can be seen in Table 3.

Among the participants, our proposed method achieved the highest privacy score. The privacy protection of our method still comes with a trade-off in intelligibility, as seen by the consistent below-average scores. We think that this could be improved by adding additional hints during and after anomalous events.

The appropriateness/pleasantness score is also consistently below average. One possible cause for this is that the privacy filter obscures the whole rectangular region around a person including a significant portion of the background. This could be improved with a pixel-wise foreground segmentation. However, this requires the foreground segmentation to be very accurate, since every false positive could potentially reveal identity-related information. Another unpleasant artifact is the blinking of the overlaid edges. When edge values oscillate around threshold values, the edges can become distracting. We think that adaptive thresholds or temporal smoothing should be explored as a future work.

The evaluations of *Stream 1*, *Stream 2* and *Stream 3* for our method and the other participating teams are summarized in Figure 3, Figure 4 and Figure 5, respectively.

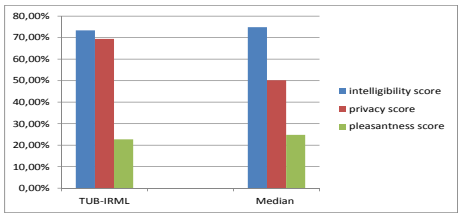

Figure 3: Stream 1 results.

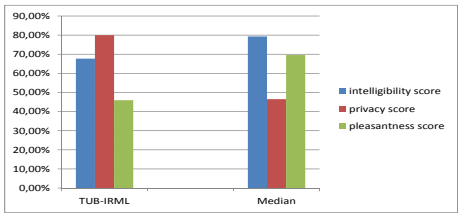

Figure 4: Stream 2 results.

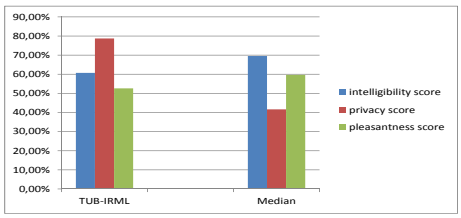

Figure 5: Stream 3 results.

4. CONCLUSIONS

In this paper, we proposed a privacy filter that obscures both shape and appearance of privacy-related regions. The user study has shown that our method is very effective at protecting privacy. As a future work, we plan to evaluate different parameters to find a suitable balance between privacy and intelligibility for different contexts. Another interesting future work would be to improve the appropriateness by reducing the obscured regions using a pixel-wise segmentation.

Acknowledgments

The research leading to these results has received funding from the European Community FP7 under grant agreement number 261743 (NoE VideoSense).

5. REFERENCES

[1] A. Badii, E. Touradj, C. Fedorczak, P. Korshunov, T. Piatrik, V. Eiselein, and A. Al-Obaidi. Overview of the mediaeval 2014 visual privacy task. In *MediaEval 2014 Workshop*, Barcelona, Spain, October 16-17 2014.

[2] P. Korshunov and T. Ebrahimi. PEViD: privacy evaluation video dataset. In *Applications of Digital Image Processing XXXVI*, San Diego, CA, August 25-29 2013.