

DAI at the MediaEval 2013 Visual Privacy Task: Representing People with Foreground Edges

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@tu-berlin.de

ABSTRACT

In this paper, we present a method for removing identity-related information from image sequences for the privacy protection of individuals. The face, despite being an important feature to identify a person, is not the only body part that needs to be obscured. Therefore, we propose to replace the whole body of individuals by their silhouette defined by moving edges.

1. INTRODUCTION

The MediaEval 2013 Visual Privacy Task [1] 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 2013 Visual Privacy Task, our privacy filter is evaluated using the Privacy Evaluation Video Dataset (PEViD) [2].

In order to prevent the misuse of video surveillance systems, visual privacy filters are being developed to remove identity-related information from a video stream. A human operator or an automatic analysis system needs to be able to track persons and their actions in order to detect anomalies. Any other information such as identity, skin color, ethnicity and gender can be misused (e.g., abuse or discrimination) [4]. In this context, our privacy filter aims not only at obscuring facial identity, but also protecting other identity revealing features such as accessories and clothing. The goal of our approach is to prevent possible abuse and discrimination by overlaying a background image with silhouettes.

2. PROPOSED METHOD

The proposed privacy filter is an adaptation of the foreground privacy filter with stored background proposed by O’Gormans [3]. The approach presented in [3] is based on two observations [3]: 1) motion edge detection is more robust to lighting changes than intensity-based segmentation methods, and 2) video privatization can often be accomplished by obscuring edge regions only. Instead of using a stored background, we initialize the background using the first frame and update it using every new frame. Using the annotation of the dataset, only pixels that are not labeled as




Figure 1: The silhouette of a person dropping a bag.

person are updated. With this scheme, we can display scene changes (e.g., moving cars) that do not correspond to individuals and therefore, are not subject to privacy protection. The drawback of this approach is that the first frame of a video might already contain an individual. Instead of further filtering as in [3], we directly use the foreground edges as silhouettes. The rationale behind this is to achieve better intelligibility. Every pixel that is considered as a foreground edge and is within a person’s bounding box, is set to a particular color, namely green. The whole body annotation of individuals provided in the dataset helps to restrict interfering false positives edges to edges around individuals.

3. EVALUATION RESULTS

The paper by O’Gormans [3] uses various parameters for the foreground detection, in particular a threshold on Sobel horizontal and vertical gradient results ($T1$) and the value of the exponential moving average constant α which basically controls foreground pixel classification change to background. We adapted the teaching of this paper by respectively setting them to 60 and 0.5 (for details, the reader is referred to [3]). The privacy filter has been evaluated using objective and subjective measures [1]. The objective and subjective evaluation results and their comparison to the average score of all 9 teams participating in the MediaEval 2013 Visual Privacy Task are given in Table 1 and Table 2, respectively.

The objective intelligibility score (in Table 1) is far below average. This is an expected result, as the objective intelligi-




Figure 2: The silhouette of two persons fighting.

bility score is measured using an automatic human detector which classifies our silhouette representation as non-humans. A privacy filter could provide tracking information on a side-channel to compensate for this problem.

Our objective privacy score is above average for the same reason. The score is based on a face detection algorithm which detects natural faces. Due to the silhouette representation of individuals, the face detection algorithm is expected to find no faces in filtered image sequences.

Table 1: Objective evaluation results

	Our Method	Average Score
Intelligibility	0.313	0.502
Privacy	0.706	0.665
Appropriateness	0.435	0.561

In the subjective evaluation (Table 2), the intelligibility score is above average. This shows that users were able to track individuals and their actions, by only seeing the silhouette. The below average privacy score in the user study might suggest that the silhouettes still contain information related to the identity of individuals. In some cases, accessories, clothing and/or hair style can still be recognized by the users. Although our method has better appropriateness score in the subjective evaluation, the appropriateness scores are still below average in both objective and subjective evaluations. This is likely due to the visual artifacts produced by the imperfect foreground edge segmentation. The edges that belong to the background have a green color, when they are close to the bounding box of a person (see Figure 2). This reduces the appropriateness score of our method.

Table 2: Subjective evaluation results

	Our Method	Average Score
Intelligibility	0.678	0.656
Privacy	0.670	0.684
Appropriateness	0.464	0.492

The objective and subjective evaluations for our method and the average results of all 9 teams participating in the MediaEval 2013 Visual Privacy Task are summarized in Figure 3 and Figure 4, respectively.

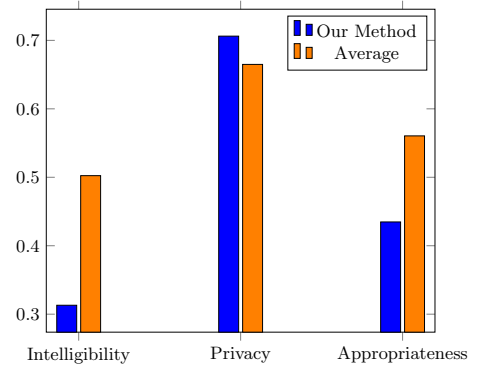


Figure 3: Objective evaluation




Figure 4: Subjective evaluation

4. CONCLUSIONS AND FUTURE WORK

In this paper, we propose a privacy filter that replaces the whole body by a silhouette. The user study shows that this filter is able to provide privacy while maintaining intelligibility. Future work needs to be done to improve foreground segmentation, and thus, to reduce artifacts produced by the imperfect foreground segmentation.

5. ACKNOWLEDGMENTS

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

6. REFERENCES

- [1] A. Badii, M. Einig, and T. Piatrik. Overview of the mediaeval 2013 visual privacy task. In *MediaEval 2013 Workshop*, Barcelona, Spain, October 18-19 2013.
- [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.
- [3] L. O'Gorman. Video privacy filters with tolerance to segmentation errors for video conferencing and surveillance. In *Pattern Recognition (ICPR), Int. Conf. on*, pages 1835–1838. IEEE, 2012.
- [4] T. Piatrik, V. Fernandez, and E. Izquierdo. The privacy challenges of in-depth video analytics. In *Multimedia Signal Processing (MMSP), 2012 IEEE 14th International Workshop on*, pages 383–386. IEEE, 2012.