Privacy Protection Filter Using Shape and Color Cues

Hajer Fradi
EURECOM, Campus Sophia Tech
450 route des Chappes
06410 BIOT
fradi@eurecom.fr

Yiqing Yan
EURECOM, Campus Sophia Tech
450 route des Chappes
06410 BIOT
yiqing.yan@eurecom.fr

Jean-Luc Dugelay
EURECOM, Campus Sophia Tech
450 route des Chappes
06410 BIOT
dugelay@eurecom.fr

ABSTRACT

The steady growth in the adoption of video surveillance systems emphasizes the need for privacy protection techniques. In this paper, we present a method inspired from image abstraction and non-photorealistic rendering fields for creating privacy protection filters. The effectiveness of the proposed filter has been demonstrated by assessing the intelligibility vs. privacy vs. pleasantness in a subjective evaluation framework using different videos.

1. INTRODUCTION

In recent years, a steady growth in the adoption of digital video surveillance systems for monitoring buildings and public spaces has been observed. These systems raised many concerns related to the privacy rights of the subjects being monitored [8, 2, 3]. At the same time, video analytic tools created additional problems, since algorithms such as face recognition or person re-identification can expose the identity of any individual that appears in the field of view of the camera.

Current surveillance systems either do not implement any mechanism for privacy protection, or they apply traditional techniques such as masking, Gaussian blurring, and pixelization. The lack of appropriate methods to detect regions of the image which contain privacy sensitive information and to evaluate the amount of privacy protection required in a specific scenario often causes failure in either minimizing the intrusion of the surveillance system or goes against the purpose of the surveillance itself.

Recently, much effort has been devoted on this field. One big challenge in defining privacy protection policies for video surveillance applications is the identification of the correct trade-off between intelligibility of the video, which should be adequate to the monitoring tasks, and privacy protection itself. Consequently, a number of recent studies have been conducted to propose more adequate systems for privacy protection. Among the proposed techniques, non-photorealistic rendering ones achieve artistic results. For instance, in [4], the authors propose to use segmentation in order to obtain a pixelizated result resembling to pixel art. However, this method applied to privacy preservation suffers from the drawbacks of the adopted pixelization filter.

In this paper, we propose a privacy protection filter inspired from [7]. As proposed in [7], the filter transforms the original privacy sensitive RoI to a simplified version, while preserving the general appearance. The algorithm is essentially based on a color quantization and a patch rendering step, to obfuscate fine details containing personal visual information. A binary masking step is involved in the process as well to refine the results. The goal behind this is to enable people and action detection tasks, while obfuscating identification details such as face and clothe traits.

2. PROPOSED APPROACH

The proposed algorithm essentially consists of applying a color-based segmentation in order to divide any region of interest (RoI) into different patches of different colors. Then, the RoI is abstracted by replacing the pixels belonging each patch by one single color. The privacy sensitive RoIs, are defined in our case by the bounding boxes provided by the annotation step. These regions can be further refined if additional foreground segmentation masks are available. An illustration of the proposed protection filter is shown in Figure1.

(a) Original frame

(b) Filtered frame

Figure 1: Illustration of the proposed protection filter estimation using color segmentation and binary masking

2.1 Color-based segmentation

To process videos, we apply a color-based segmentation algorithm in the bounding boxes defining the privacy sensitive regions of interest. This algorithm proceeds as follows: First, for each RoI, we define every pixel by its color in RGB space (r_i, g_i, b_i) and its spacial coordinates (x_i, y_i). So, every pixel could be presented as this vector $[r_i, g_i, b_i, \alpha x_i, \alpha y_i]$, where α is a coefficient used to adjust the balance between the spatial proximity and the color similarity of the resulting clusters. α is set according to the frame size and the number of clusters. After that, the pixels belonging to that RoI are divided into different groups using k-means clustering algorithm [6], which minimizes the euclidean distance between them. As known, k-means choose N points randomly as the assumed centroids at the beginning. Then, it computes the euclidean distance of every point to these N centroids. Similar to other pixels, the centroid is represented as $[r_{c_n}, g_{c_n}, b_{c_n}, \alpha x_{c_n}, \alpha y_{c_n}]$. By varying the number of clusters (N), different abstraction levels can be obtained, either globally or locally in certain regions.

Once the clustering process is finished, the pixels in the same patch are replaced by one color from this patch. This color is chosen to be the color of the centroid of every patch.

2.2 Binary masking

Since our primary goal is to keep the main information to enable detecting persons and actions while obscuring the sensitive data to protect personal privacy, our actual results obtained by applying k-means clustering in RoI have to be further processed. At this stage, we have as input the results of the clustering and the patch rendering steps applied on the whole RoI, which is in our case the bounding box of the annotation. These results are refined using the foreground masks. The final abstracted RoI is computed as:

$$I_{out} = I_{in} \cap [S(L, Fg) \geq T] \oplus I \cap [S(L, Fg) < T] \qquad (1)$$

where I_{out}, I_{in}, and I denote the final abstracted RoI, the abstracted RoI resulting from the previous step, and the original RoI, respectively. Also, L is the segmentation label map, Fg is the foreground mask, and S is a support operator that counts the number of foreground pixels of each patch label.

By this way, each patch is either fully rendered abstracted or fully rendered original. Finally, the refined and filtered RoIs are added to the original frame by replacing the original ROIs. This step of binary masking is important to make the results visually more appealing and to avoid filtering non-sensitive regions.

3. EXPERIMENTAL RESULTS

We tested our proposed approach on different video sequences from PEViD dataset [5]. The evaluation has been performed according to MediaEval 2014 Visual Privacy Task guidelines, more details can be found in [1].

In Table 1, we report our results, compared to the median score of all participants to the challenge. These results are the average of three subjective evaluations.

It is important to mention that our proposed filter achieved good results for both intelligibility and pleasantness, however the privacy protection level was low. This limitation of privacy protection can be mainly explained by the fact that our approach did not enforce maximum protection for sensitive regions such as faces. We would expect higher protection level if this point has been taken into account.

Evaluation	Score(%)	Median(%)
Intelligibility	79.35	74.57
Privacy	29.13	46.12
Pleasantness	59.02	51.34

Table 1: Our results compared to the median score of all participants to MediaEval privacy task

4. CONCLUSIONS

In this paper, we present our approach for privacy protection filters based on color segmentation and binary masking. The resulting abstracted (filtered) image has the advantage of resembling to the original image in the general shape and color appearance, while destroying fine details.

5. ACKNOWLEDGMENTS

This work has been conducted within the framework of the EC funded Network of Excellence VideoSense.

6. REFERENCES

[1] A. Badii, T. Ebrahimi, C. Fedorczak, P. Korshunov, T. Piatrik, V. Eiselein, and A. Al-Obaid. Overview of the mediaeval 2014 visual privacy task. October 2014.

[2] F. Dufaux and T. Ebrahimi. A framework for the validation of privacy protection solutions in video surveillance. In *ICME*, pages 66–71. IEEE, 2010.

[3] H. Fradi, A. Melle, and J.-L. Dugelay. Contextualized privacy filters in video surveillance using crowd density maps. In *ISM 2013, IEEE International Symposium on Multimedia, 9-11 December 2013, Anaheim, CA, USA*, Anaheim, ÉTATS-UNIS, 12 2013.

[4] T. Gerstner, D. DeCarlo, M. Alexa, A. Finkelstein, Y. Gingold, and A. Nealen. Pixelated image abstraction. In *NPAR 2012, Proceedings of the 10th International Symposium on Non-photorealistic Animation and Rendering*, June 2012.

[5] P. Korshunov and T. Ebrahimi. PEViD: privacy evaluation video dataset. In *Proceedings of SPIE Volume 8856*, volume 8856 of *Proceedings of SPIE*, Bellingham, 2013. Spie-Int Soc Optical Engineering.

[6] S. Lloyd. Least squares quantization in pcm. *IEEE Trans. Inf. Theor.*, 28(2):129–137, Sept. 2006.

[7] A. Melle and J.-L. Dugelay. Shape and color-aware privacy protection. Number EURECOM+4152, Barcelona, SPAIN, 10 2013.

[8] H.-M. Moon and S. B. Pan. Implementation of the privacy protection in video surveillance system. In *SSIRI*, pages 291–292. IEEE Computer Society, 2009.