Data Anonymization for Data Protection on Publicly Recorded Data

David Münch(✉), Ann-Kristin Grosselfinger, Erik Krempel, Marcus Hebel, and Michael Arens

Fraunhofer IOSB, Ettlingen, Germany
david.muench@iosb.fraunhofer.de

Abstract. Data protection in Germany has a long tradition (https://www.goethe.de/en/kul/med/20446236.html). For a long time, the German Federal Data Protection Act or Bundesdatenschutzgesetz (BDSG) was considered as one of the strictest. Since May 2017 the EU General Data Protection Regulation (GDPR) regulates data protection all over Europe and it strongly influenced by the German law. When recording data in public areas, the recordings may contain personal data, such as license plates or persons. According to the GDPR this processing of personal data has to fulfill certain requirements to be considered lawful. In this paper, we address recording visual data in public while abiding by the applicable laws. Towards this end, a formal data protection concept is developed for a mobile sensor platform. The core part of this data protection concept is the anonymization of personal data, which is implemented with state-of-the-art deep learning based methods achieving almost human-level performance. The methods are evaluated quantitatively and qualitatively on example data recorded with a real mobile sensor platform in an urban environment.

Keywords: Video data anonymization · Data protection

1 Introduction

Recording visual data in public while abiding by the applicable laws does require special attention concerning data protection. In this paper we present a formal data protection concept and a practical realization of that system with reliable results for handling the recorded data in a secure way. Thereby, a mobile sensor platform is used as underlying practical use case.

MODISSA (Mobile Distributed Situation Awareness) [4] is an experimental platform for evaluation and development of hard- and software in the context of automotive safety, security, and infrastructure applications. The basis of MODISSA is a Volkswagen Transporter T5, which has been equipped with a broad range of sensors on the car roof, see Fig. 1. The necessary hardware for raw data recording, real-time processing, and data visualization is installed inside the vehicle. The battery powered power supply and electronics are installed in the

D. Tzovaras et al. (Eds.): ICVS 2019, LNCS 11754, pp. 245–258, 2019.
https://doi.org/10.1007/978-3-030-34995-0_23

back and allow independent operation for several hours. The sensor equipment can be adapted to the requirements of the current task. To operate the sensors, several state-of-the-art computers are provided in the vehicle. A complete row of seats can be used by up to three people to operate the sensor system or to watch the results on the display devices (screens, virtual reality headsets).

The sensors are mounted on item aluminum profiles, which are mounted on roof racks. The current LiDAR sensor configuration consists of two Velodyne HDL-64E at the front and two Velodyne VLP-16 at the back. Additionally, in each corner of the car roof there are two Baumer VLG-20C.I color cameras which allow to generate a complete 360° panorama around the vehicle. On the middle of the sensor platform a gyro-stabilized pan-tilt unit is mounted and equipped with a thermal infrared camera, a grayscale camera, and a laser rangefinder.

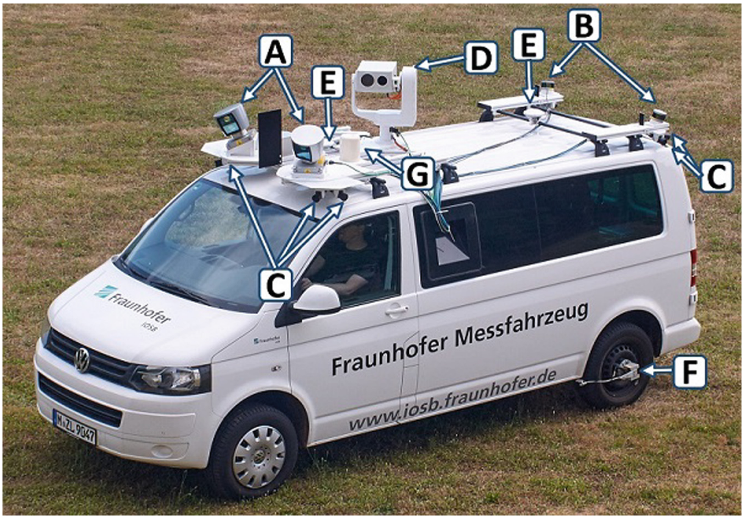

Fig. 1. The external components of the sensor system MODISSA. A: 2x Velodyne HDL-64E LiDAR, B: 2x Velodyne VLP-16 LiDAR, C: Panoramic camera setup 8x Baumer VLG-20C.I. 2 cameras per corner, D: Pan-tilt unit with Jenoptik IR-TCM 640 thermal infrared camera, JAI CM 200-MCL gray scale camera, and Jenoptik DLEM 20 laser rangefinder, E/F: Applanix POS LV V5 520 inertial navigation system with GNSS antennas (E), distance measuring indicator (DMI) (F), inertial measurement unit (IMU) and position computer (both not visible), G: External WiFi antenna. Figure and caption adapted from [4].

2 Data Protection on Publicly Recorded Data

Having the exemplary mobile sensor platform MODISSA described in the paragraph above, we are technically ready to start driving and recording data. Before this can be done, we have to identify applicable law and decide how to implement legal requirements. Processing of personal data in the EU is always regulated by the GDPR which in Germany is further qualified in the BDSG. Depending on the concrete scenario further area-specific regulations may exist.

As we operate with a mobile sensor platform in the public and record data there, it is inevitable that the recorded data contains personal data, such as license plates or persons. A legal opinion [9] and the critical feedback of the federal states data protection commissioner [3] come both to the conclusion that processing is possible when certain regulations and requirements are met.

By default processing of personal data in the EU is prohibited unless at least one of six different reasons legitimate it (Art. 6 GDPR). In short these possible reasons are: (a) data subject has given consent; (b) processing is necessary to fulfill a contract with the data subject; (c) processing is required by obligations of the processor; (d) processing is necessary in order to protect the vital interests of the data subject; (e) processing is necessary for the performance of a task carried out in the public interest; (f) legitimate interest of the processor. In the case of (f) legitimate interest, the data processor has to ensure, that his interest in data processing is not overridden by the interests or fundamental rights and freedoms of the data subject. Therefore, a high level of data protection and data security is mandatory. A data protection concept based on three core principles was developed to achieve this:

- **Avoidance.** Only collect as much data as necessary.
- **Security.** Protection against unauthorized access to the recorded data.
- **Privacy preserving.** Anonymization of personal data.

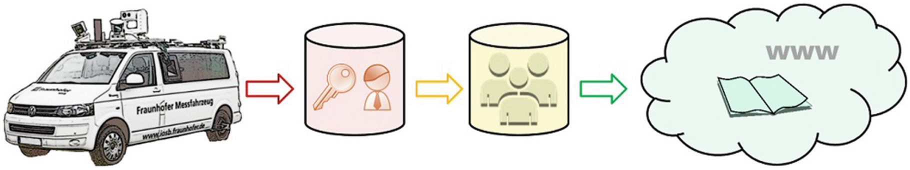

Fig. 2. Different zones for the recorded data: The red zone handles temporary stored raw data. In the yellow zone we have mostly anonymized data for research purpose with limited access only to people needing the data. In the green zone are few selected and manually checked anonymized data for print and online publication. (Color figure online)

Further we define three different environments where data is processed. In each environment appropriate measures enforce the three core principles. Figure 2 illustrates the three different environments:

- In the **red zone**, we handle temporary stored raw data. It has the highest level of access restrictions. Every access, i.e. using algorithms to anonymize the data and transform it in the yellow zone, is logged and restricted to selected persons. Additionally, physical access to the red zone is restricted.

- In the **yellow zone** we have algorithmically anonymized data for research purpose with limited access by people needing the data for their research.
- In the **green zone** we have selected and manually audited anonymized data for print and online publication.

How to deal with the three core principles in the red zone? **Avoidance:** The need for data recording in public space must exist. The reasons to record are clearly justified and documented in advance in a measurement plan. It is justified in the measurement plan which sensors are needed to achieve the research goals. Data recording is limited to these sensors. As part of the achievement of the research objectives a short time period and a route with a certainly low risk of the occurrence of sensitive data is selected. As a consequence, the responsible person for the data recording is generating a measurement plan and a measurement report.

Security: Access to the mobile sensor platform is limited to a few persons. Every ride on public ground is recorded in the logbook. The group of people involved in data recording is reduced to a necessary minimum. These individuals are aware of their responsibilities for the confidential treatment of sensitive information. Storing of data is minimized and access to this data is restricted at all times by technical measures. As a consequence the mobile sensor platform is never left unattended in public or accessible to third persons. By default, the raw data can be encrypted with AES-256. At the end of a recording, the responsible person has to transfer data stored on the mobile sensor platform to a separate and secured computer and has to delete all recorded data on the mobile sensor platform. All operations are logged in the measurement report.

Privacy preserving: For the planned scientific investigations (described in the measurement plan) personal data is removed from the recorded data. This includes faces and license plates. As a consequence anonymization of faces and license plates takes place automatically on a dedicated and secured computer. After anonymization the recorded and anonymized data is allowed to enter the yellow zone.

Detecting the objects of interest and then blurring them is the most common way in ensuring privacy in video data [1,6,7,15,16,19]. Another possibility instead of blurring sensitive data out, is the total removal of sensitive data, such as removing a whole person and inpainting background instead [20]. Recent works mainly address a street view scenario. In our case, we differ from previous approaches in developing a formal data protection concept abiding the law, gaining almost human performance, and providing a fast running overall system.

In the following we give a detailed description of the process of automatically anonymizing faces and license plates.

3 Data Anonymization

The data anonymization methods address people and license plates. As license plates can be assigned to the vehicle owner, they are subject of the personal right of individuals [3,9]. Thus, they need to be anonymized when recording and

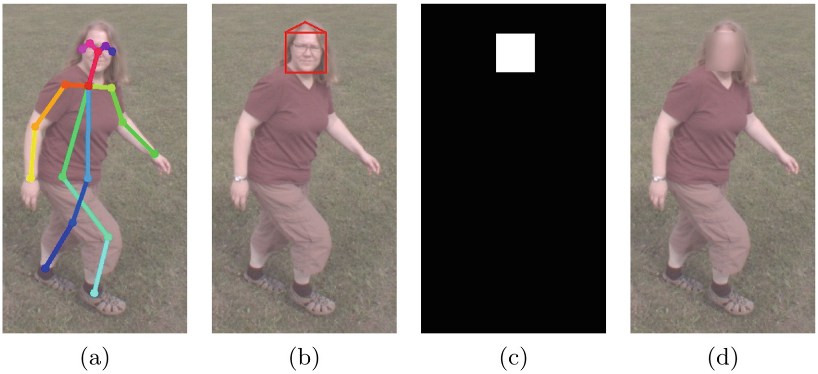

Fig. 3. After having estimated the facial regions from up to six upper body parts from OpenPose (a) the facial region can be inspected in detail (b) which is used as local mask (c) to anonymize, resulting in an image with blurred face (d). Figure from [8].

processing road scenes [3]. Our implemented processing pipeline is twofold: an automated license plate detection and localization system and all faces of persons in the images need to be anonymized. The source data from MODISSA contains images from cameras mounted on eight different positions around the vehicle, c.f. Fig. 1C. These images are of size 1624 × 1228 pixels and are stored in a raw Bayer pattern. In addition, images from cameras mounted on a pan-tilt-unit also need to be anonymized, c.f. Fig. 1D.

3.1 Face Anonymization

It has shown, that a fast and robust face detection including the whole person as contextual knowledge yields practical good results to gather the facial regions of the persons to be anonymized.

Towards this end, OpenPose [5,21], a real-time multi-person system is applied to estimate the keypoints of the persons and their facial region. From up to six positions of upper body parts the face region of the person is estimated. A subset of the facial parts nose, neck, leftEye, rightEye, leftEar, and rightEar, if detected above a given confidence threshold, contribute their weighted centroids to the calculation of a face center. The final result is the average of all passed face position suggestions. The size of the face region is determined from the distances of contributing nodes which are applied with a constant factor [8]. See Fig. 3 for a detailed overview.

3.2 License Plate Anonymization

Automated license plate recognition (ALPR) is well addressed in research, as it is a building block for many real applications, such as automatic toll collection, road traffic monitoring, or intelligent transportation systems. Many proposed systems are hierarchically designed and use a license plate detector in the first stage [2,10,11,13]. License plate detection is a difficult task with high variations in the appearance of license plates, see Fig. 4. The recordings in residential areas does not only contain preceding and following vehicles, but also cars on the roadside or in parking lots. Hence, the variation of orientations, scales, and positions of recorded vehicles and their license plates is large.

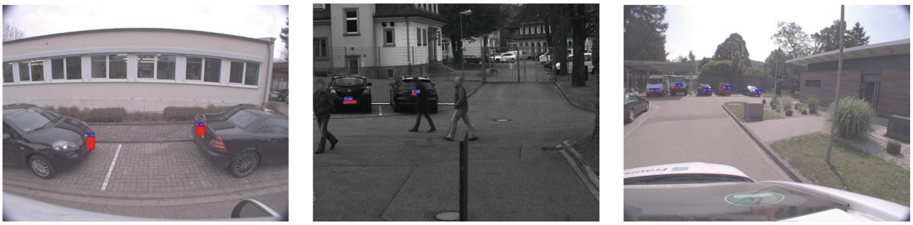

(a) (b) (c)

Fig. 4. Example images for the license plate anonymization. Challenge: High rotation angle (a), concealed parts (b), small plate sizes (c). Figure from [17].

For training and validation [12] an annotated data set of German license plates would be optimal, but is currently not available. Thus, for creation of an own dataset the following four approaches are taken: First, synthetic images with number plates of the same appearance as German license plates are generated. Secondly, annotated datasets of non-German license plates are used. Thirdly, images of German and European license plates are gathered and annotated. Lastly, the former dataset is expanded by using rectified early test results and data augmentation methods.

As in several related works, the plate detection is performed in two stages: First, a YOLOv3 CNN detects and localizes vehicles and second, another YOLOv3 CNN detects number plates inside the vehicle regions, see Fig. 5. Even

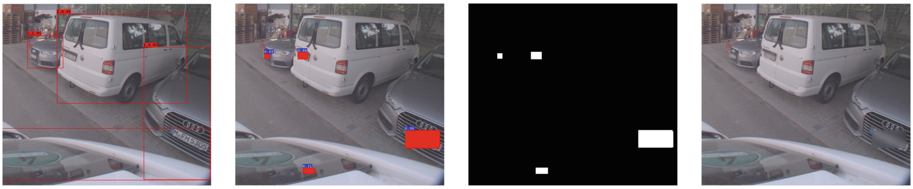

(a) (b) (c) (d)

Fig. 5. License plate detection pipeline. After having localized the cars (a) the license plate region can be detected (b) which is used as local mask (c) to anonymize, resulting in an image with blurred license plates (d).

though the same model is used for both steps, different weights and hyperparameters are set. The main benefit of a two stage detector is the reduction of false positive detections. Obviously, license plates are practically always mounted on vehicles and negligibly rare found isolated in other areas of the image. If the search area is restricted to image parts including vehicles, all false positive detections that are not on a vehicle are avoided. Furthermore, the precision is improved due to a smaller search area.

Properties that make using YOLOv3 a promising approach are the ability YOLOv3 outperforms other object detectors like SSDs in terms of average precision for small objects together with low runtime.

4 Anonymization of Image Regions

A markup language annotation file is used to store the estimated sensitive parts of faces and license plates to be anonymized. An external annotation tool can be used for manual refinement, if necessary. It is possible to add boxes for missing detections, eliminate false positives, or adjust the box boundaries. From the annotation file data, image masks are generated, which are used to determine regions where the image will be changed. Most of the image data to be anonymized is stored as raw Bayer pattern. A modified version of a Gaussian blur is applied to the image as 2D-filter [8].

Anonymization with Gaussian blur can be deanonymized under certain circumstances, see [14]. Since our regions to anonymize cover only few pixels, blur is applied with high size and σ and the blur regions are cropped from wholly blurred images—not phased out—our method can stay for the moment, but alternatives should be kept in mind.

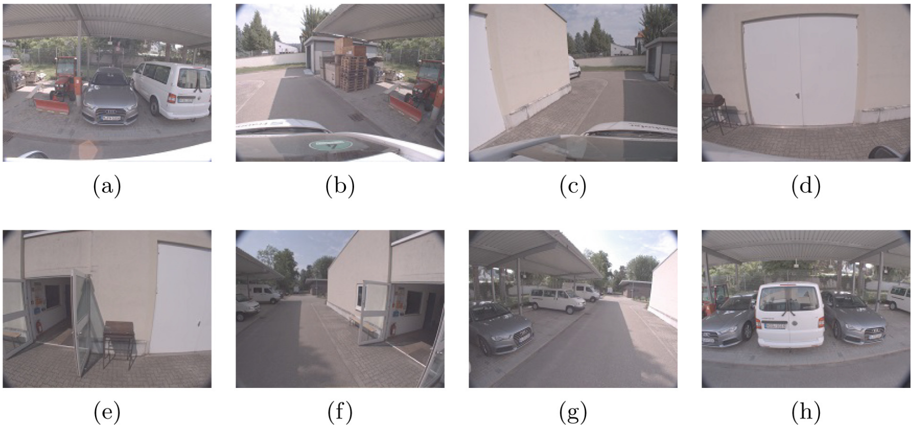
(a) (b) (c) (d)

(e) (f) (g) (h)

Fig. 6. Examples of MODISSA test dataset. All images are taken at the same time but with different cameras. Figure from [17].

5 Datasets

Three datasets are chosen for testing: A dataset from openALPR[1]. This dataset contains three subsets with Brazilian, European, and US license plates on vehicles. There is one vehicle per image and the perspective variations are little. This dataset is used to compare the performance to other license plate detectors.

For testing the face anonymization a face test dataset with 736 frames was recorded, also using the eight panorama cameras, see Fig. 1C. In this test dataset a person is moving around the MODISSA sensor platform and is visible in up to three cameras at the same time. Ground truth, consisting of 1053 boxes and 91 don't-care boxes is provided. The don't-care areas take into account that the back of a head does not need to be anonymized, but it is no matter if it is anonymized.

For evaluation of the performance for the actual overall anonymization process a third dataset is generated. Representative recorded data from MODISSA is used for evaluation. The dataset consists of 396 fully annotated frames from the eight panorama cameras, see Fig. 1C. A total amount of 1215 faces and 7933 license plates are present in the data. Figure 6 depicts some examples from the MODISSA test dataset.

6 Evaluation

Intersection over Union (IoU) with IoU ≥ 0.5 is a commonly used measure to evaluate detections. In our cases we observed that this measure does neither fit well for face detection nor for license plate detection. The reason is, that there are on the one side regions containing eyes, ears, nose, and mouth which should be anonymized and on the other side there are regions such as hair and forehead which do not care. Anonymized regions should not cover the whole image but can be bigger than IoU ≥ 0.5, if the needed (annotated) part is covered. To address this issue we introduce as additional measure Intersection over Area (IoA), the percentage the ground truth is covered with the detection.

In detail, IoU measures how good detection D and ground truth G match and IoA measures how good ground truth G is covered from detection D.

$$\text{IoU} = \frac{area(D \cap G)}{area(D \cup G)}, \qquad \text{IoA} = \frac{area(D \cap G)}{area(G)}$$

IoU threshold is applied to license plate detection evaluation with IoU ≥ 0.5, and IoU ≥ 0.3 respectively. On face detection evaluation IoU ≥ 0.3 is applied together with IoA ≥ 0.7.

6.1 Face Anonymization

Applying IoU ≥ 0.5 to the face test dataset, we gain weak performance; instead according to the paragraph above using the evaluation metrics IoU ≥ 0.3 and

[1] https://www.openalpr.com/.

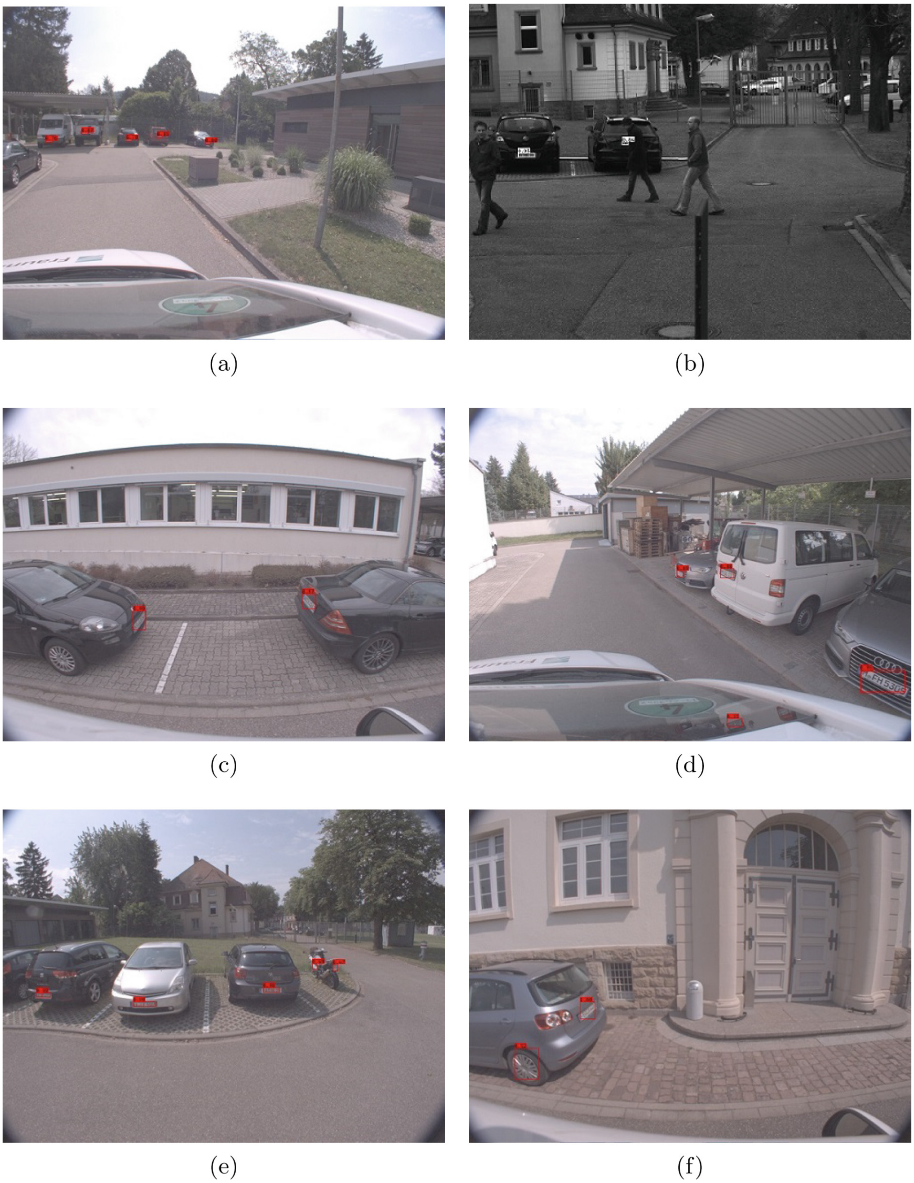
(a)

(b)

(c)

(d)

(e)

(f)

Fig. 7. Representative test results. High confidences for difficult scenes: Small plate sizes (a), concealed parts (b), high rotation angle (c-e), mirror images (d) and doubled lined plate design (e). False positives for bright surface (e) and wheel rim (f).

IoA ≥ 0.7 gives a much more realistic impression of the performance on real data.

With a confidence threshold of 0.575 the face test dataset evaluation achieves the highest precision value 0.998 from 1023 true positives, 2 false positives, 11 used don't-care areas, and 30 false negatives. With confidence threshold 0.32 we get the highest amount of 1039 true positives besides 11 false positives, 43 used don't-care areas and 14 false negatives, from which two are below the IoU threshold and 12 are not detected at all. All false negatives are next to the image border. In a further step the anonymization could be applied to the stiched panorama image instead of the eight single images. Then, the false negatives at the image borders are no longer expected. Thus, we argue, that false negatives at the image border (28 out of 30) can be neglected.

6.2 License Plate Anonymization

The number plate localization system is assessed using the openALPR benchmark dataset. Thereby, the overall license plate detection system is evaluated instead of the two initial stages separately. Using the standard IoU ≥ 0.5, an average precision (AP) of 85.36% is yielded. Reducing IoU ≥ 0.3, the AP is increased to 98.73%. Considering the fact that the bounding boxes for the test images use another norm than the training images, the smaller threshold is reasonable: While the complete number plate is labeled as ground truth for the benchmark dataset, only the smallest axes-aligned rectangle that includes all characters of the plates states the bounding box for the training images.

Figure 7 shows some example results applying the license plate detection to the MODISSA test dataset.

Even though false positive detections are reduced by the two stage approach, they occur occasionally. Since the distance between the confidence for false positive and false negative detections is large, this is not a problem in general. However, there are few exceptions for false negatives with higher confidences. Mostly, they occur in the following cases:

- periodic, in particular vertical structures, e.g. wheel rims, characters
- bright surfaces, e.g. white cars, overexposed areas
- rectangular objects, e.g. other plates, side mirrors, rear lights.

To improve the license plate detection system, mostly the second stage, i.e. the YOLOv3 for plate detection, should be considered. The training is limited by the number and quality of the training dataset. Thus, further data augmentation and/or suitable annotated datasets might enhance the precision and recall.

6.3 Overall Evaluation

For quantitative results of the MODISSA test dataset, see Fig. 8. Since the dataset was recorded for development and optimized to contain many samples of different challenging situations, the results are correspondingly. The sequences

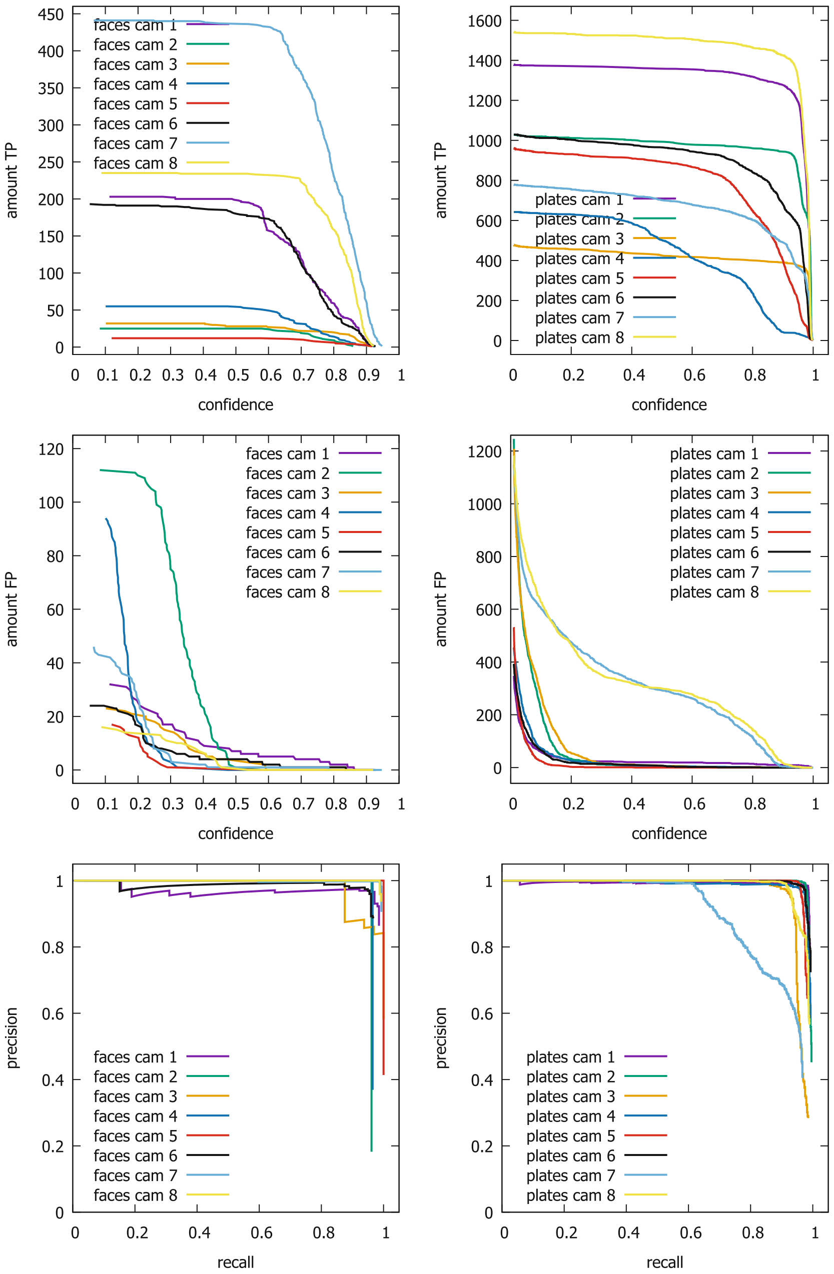

Fig. 8. Quantitative results from the MODISSA test dataset. Left side, results of face anonymization, right side plate anonymization. True positive (TP). False positive (FP).

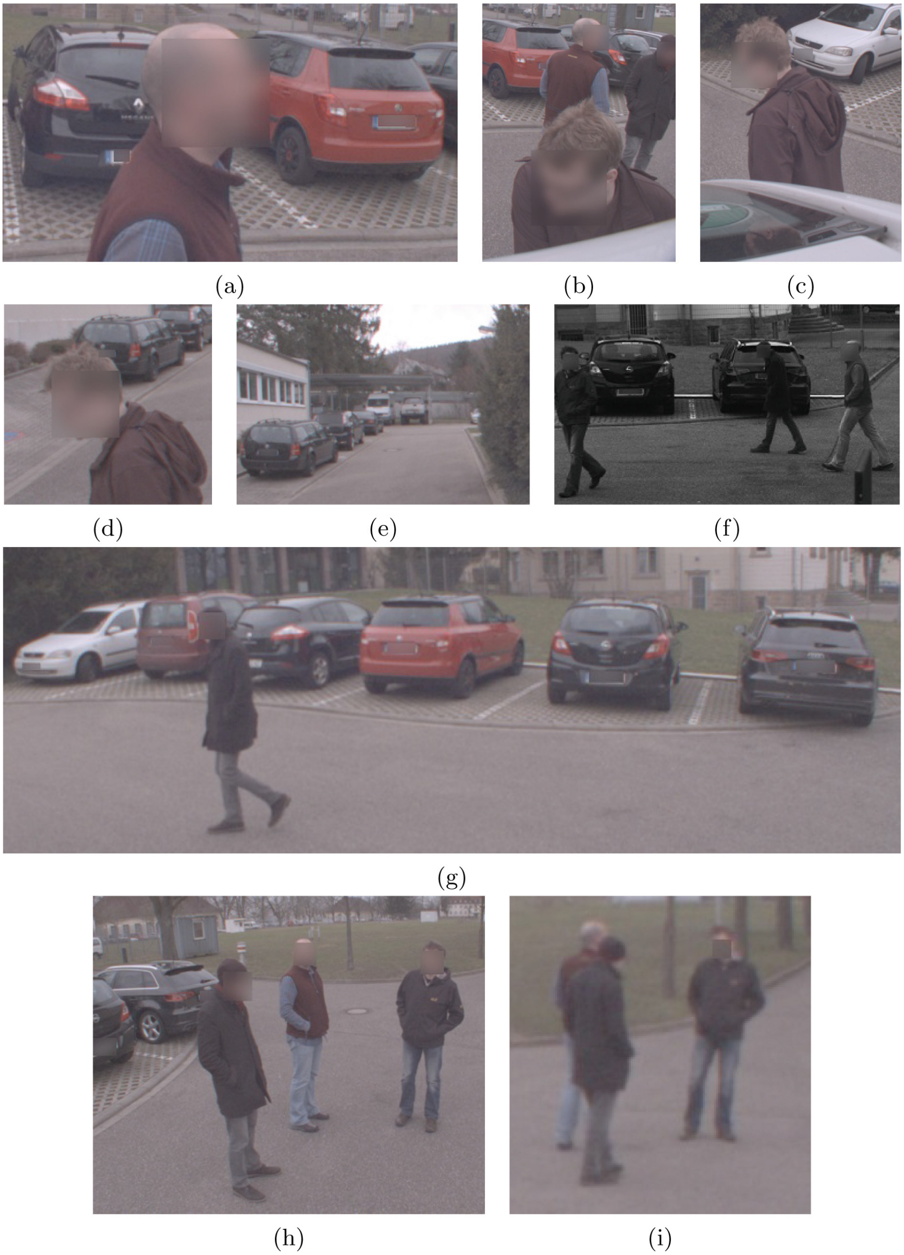

(a) (b) (c)

(d) (e) (f)

(g)

(h) (i)

Fig. 9. Qualitative final automatically anonymized results from the MODISSA test dataset.

of camera 1 to 8 contain a ground truth amount of 206, 26, 32, 57, 12, 199, 446, and 237 faces, as well as 1392, 1032, 485, 647, 978, 1036, 805, and 1558 plates. Besides, there are 53, 16, 9, 2, 4, 301, 148 and 26 don't care faces and 10, 120, 508, 10, 285, 416, 100 and 12 don't care plates annotated. Don't-care plates are too small to show details, that is they do not need to be anonymized, but it does not matter if they are anonymized. During the experiments 17, 0, 0, 0, 2, 32, 21, and 22 don't care faces and 444, 112, 182, 32, 329, 92, 286, and 345 don't care plates are used. Qualitative results are visualized in Fig. 9.

6.4 Runtime

Manually annotating 1000 faces takes about 120 min. Automatic anonymization takes between 0.6 and 1.1 s per camera and frame for license plate detection, depending on car density in the image. Face detection takes about 0.15 s per camera and frame. Applying the Gaussian blur takes about one minute for 1000 frames. The computation time of automatic data anonymization including manual post-treatment is about ten times faster than completely manually anonymizing.

7 Conclusion

In this paper, we addressed recording visual data in public while abiding by the applicable laws. As a consequence, a data protection concept is developed for a mobile sensor platform. The core part of this data protection concept is the anonymization of personal data. Face anonymization is implemented based on OpenPose [5,8,21] and license plate anonymization is realized as a hierarchical approach based on YOLOv3 [17,18]. The quantitative evaluation on example data recorded with a mobile sensor platform in an urban environment shows almost human-level performance and manual post-treatment data anonymization is reduced to a minimum.

References

1. Agrawal, P., Narayanan, P.: Person de-identification in videos. IEEE Trans. Circuits Syst. Video Technol. **21**(3), 299–310 (2011)
2. Arsenovic, M., Sladojevic, S., Anderla, A., Stefanovic, D.: Deep learning driven plates recognition system. In: 17th International Scientific Conference Industrial Systems (IS 2017) (2017)
3. Bayerisches Landesamt für Datenschutzaufsicht: Bundesdatenschutzgesetz (BDSG); Rechtsgutachten zum möglichen Erprobungseinsatz eines mit Laser- und Kameratechnik ausgestatteten Sensorfahrzeuges im Stadtgebiet von Ettlingen, December 2015
4. Borgmann, B., Schatz, V., Kieritz, H., Scherer-Klöckling, C., Hebel, M., Arens, M.: Data processing and recording using a versatile multi-sensor vehicle. ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci. **IV-1**, 21–28 (2018)

5. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2D pose estimation using part affinity fields. In: CVPR (2017)
6. Flores, A., Belongie, S.: Removing pedestrians from google street view images. In: CVPR - Workshops, June 2010
7. Frome, A., et al.: Large-scale privacy protection in Google street view. In: ICCV. IEEE (2009)
8. Grosselfinger, A.K., Münch, D., Arens, M.: An architecture for automatic multimodal video data anonymization to ensure data protection. SPIE Security + Defence (2019)
9. HÄRTING Rechtsanwälte: RECHTSGUTACHTEN zu Rechtsfragen zum möglichen Erprobungseinsatz eines mit Laser- und Kameratechnik ausgestatteten Sensorfahrzeuges im Stadtgebiet Ettlingen durch das Fraunhofer Institut für Optronik, Systemtechnik und Bildauswertung (IOSB), Abteilung Objekterkennung, June 2013
10. Jørgensen, H.: Automatic license plate recognition using deep learning techniques. Master's thesis, NTNU (2017)
11. Laroca, R., et al.: A robust real-time automatic license plate recognition based on the YOLO detector. arXiv preprint arXiv:1802.09567 (2018)
12. Li, H., Shen, C.: Reading car license plates using deep convolutional neural networks and LSTMs. arXiv preprint arXiv:1601.05610 (2016)
13. Masood, S.Z., Shu, G., Dehghan, A., Ortiz, E.G.: license plate detection and recognition using deeply learned convolutional neural networks. arXiv preprint arXiv:1703.07330 (2017)
14. McPherson, R., Shokri, R., Shmatikov, V.: Defeating image obfuscation with deep learning. arXiv preprint arXiv:1609.00408 (2016)
15. Nodari, A., Vanetti, M., Gallo, I.: Digital privacy: Replacing pedestrians from Google street view images. In: ICPR, November 2012
16. Padilla-López, J.R., Chaaraoui, A.A., Flórez-Revuelta, F.: Visual privacy protection methods: a survey. Expert Syst. Appl. **42**(9), 4177–4195 (2015)
17. Peter, R., Grosselfinger, A.K., Münch, D., Arens, M.: Automated license plate detection for image anonymization. SPIE Security + Defence (2019)
18. Redmon, J., Farhadi, A.: Yolov3 an incremental improvement. arXiv preprint arXiv:1804.02767 (2018)
19. Ribaric, S., Ariyaeeinia, A., Pavesic, N.: De-identification for privacy protection in multimedia content: A survey. Sig. Process. Image Commun. **47**, 131–151 (2016)
20. Uittenbogaard, R., Sebastian, C., Vijverberg, J., Boom, B., Gavrila, D.M., et al.: Privacy protection in street-view panoramas using depth and multi-view imagery. In: CVPR (2019)
21. Wei, S.E., Ramakrishna, V., Kanade, T., Sheikh, Y.: Convolutional pose machines. In: CVPR (2016)