Toward Preserving User Privacy in Collected Visual Data of Autonomous Cars

Extended Abstract

Arghavan Hosseinzadeh da Silva
Fraunhofer IESE
Kaiserslautern, Germany
arghavan.hosseinzadeh@iese.fraunhofer.de

Masoud Naderpour
Fraunhofer IESE
Kaiserslautern, Germany
masoud.naderpour@iese.fraunhofer.de

ABSTRACT

In this paper, we motivate the idea of anonymizing textual data in images captured by autonomous cars that could potentially impair the location privacy of the passengers. More precisely, we focus on text instances that seem irrelevant to location at first glance but turn out to be revealing. In order to recognize privacy-sensitive texts, we take advantage of recent advancements in deep learning methods for visual perception as well as place search APIs. We evaluate our approach on real-world test data and report on a preliminary performance assessment.

KEYWORDS

Road scene images, privacy preserving, text recognition, text detection

1 INTRODUCTION

On-board sensors in autonomous cars capture a massive amount of road scene data, which can be used as an input to machine learning and deep learning techniques to enhance the perception system of autonomous cars. Researchers use the collected data for advanced analysis to improve scene understanding and decision making. This exploitation is crucial for the advancement of secure autonomous driving; however, it raises open questions about the data protection aspects of the collected data.

A recent user study [3] has shown that more than 50 percent of participants consider "capturing images" and "continuous analysis" to happen very likely in autonomous cars. Moreover, 85 percent of respondents are uncomfortable with the tracking of their vehicles, although one in three expects this to happen. Furthermore, data are subject to GDPR [4] if they contain personal information (e.g., location data). Likewise, the processing of data is subject to GDPR if it leads to profiling, that is, to any form of automated processing to predict aspects concerning the economic situation, health, personal preferences, interests, behavior, location or movements of this natural person, or similar.

In general, a captured image may threaten user privacy in various direct and indirect ways. A sequence of images shows the user's mobility traces as well as the exact locations where that user has been during a certain period of time. For example, the GPS coordinate records or the road signs directly indicate the location of the user at the time of recording. In addition, the names of shops, institutes, railway stations, restaurants and so on, indirectly—although with high accuracy—reveal the location of the captured image. In addition, names of the stores, institutes, railway stations, restaurants and so on, indirectly, although, with high accuracy, reveal the location of the captured picture. Li et al. [5] consider the location semantics and the frequency of visit as two important characteristics to measure the privacy relevance of a location. Moreover, they observe that users with similar traces have similar personal profiles. These observations justify the need to anonymize those place names that are considered privacy-sensitive in the captured images. In this paper, we focus on machine learning approaches for detecting and recognizing such texts that indirectly reveal user location data. We then use MYDATA [1], which is a Data Usage Control solution to enforce dynamic security and privacy policies within a system, in order to anonymize (e.g., to blur) the perceived privacy-sensitive texts in the captured images.

2 LOCATION-SENSITIVE TEXT DETECTION

We set our goal to recognize the privacy-sensitive texts in road scene images and to anonymize them accordingly. There are many conventional approaches to extract texts from images such as invoices, bank statements, computerized receipts and so on. However, due to the more complex nature of road scene images, these solutions are unsuitable. More recently, methods based on Deep Learning have been used to tackle the text recognition problem, since these methods have shown a better performance in various visual perception tasks compared to traditional techniques [6].

The text recognition task comprises the following two steps: **i.** Text detection: The presence of text blocks is predicted using convolutional neural networks in the form of bounding boxes. A bounding box includes the coordinates of a quadrilateral that tightly encloses a block of text and locates its exact position in a given image. Road scene images typically contain arbitrarily oriented text with various shapes and fonts. To identify as many text blocks as possible, we use the text detection method proposed by Liao et al. [7], namely Textboxes++, since it performs competently in the detection of both vertical and horizontal texts. **ii.** Text recognition: Various deep learning methods have been proposed for the text recognition subtask. In this step, we use a recognition method based on the work of Shi et al. [9], namely CRNN to recognize the actual word(s) inside the bounding boxes.

In order to train our detection model, we used an off-the-shelf action camera with an ultrawide lens to mimic the embedded cameras used in current autonomous vehicles. We collected a set of images in the urban areas of Kaiserslautern and Mannheim (a total of 440 images, including 56 images reserved for evaluation). Our training data also includes a set of images from ICDAR MLT 2017 dataset

ACM Computer Science in Cars Symposium, 2019, Kaiserslautern, Germany
.

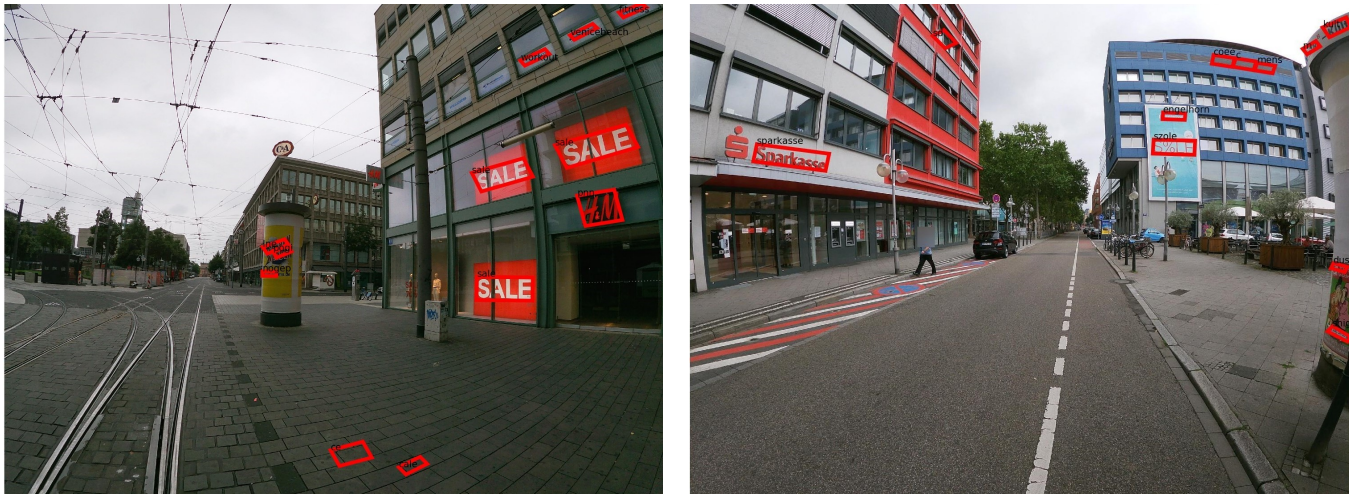

a b

Figure 1: These images are part of our test data. The combination of the places shown in the images (i.e., "H&M", "C&A" and "venicebeach" in Figure a and "Sparkasse" and "Engelhorn" in Figure b) reveal the exact locations of the taken pictures which are Kunststraße 6+7, 68161 Mannheim, Germany and E1 2, 68159 Mannheim, Germany, respectively.

[8] (1000 natural scene images). For text recognition, we use the pre-trained model provided by Shi, Bai and Yao in [9].

Anonymizing all detected text blocks could result in an excessively defaced or distorted image. Therefore, we lay an intermediary step between recognition and anonymization steps to filter out insensitive data. To this end, we use a place search API, Google Places API [2], to limit the candidate words to only those that return a specific address when queried against Places API. This API compensates to some extent the partially incorrectly recognized texts; it returns the expected address even if the input place name includes a few misspelled letters. As a final step, we anonymize the recognized privacy-sensitive texts, e.g. through redaction, blurring, or similar in the captured images.

3 EVALUATION

We evaluated our solution on a test dataset of 56 road scene images containing 155 privacy-sensitive texts. Figure 1 shows the result of text detection and recognition on two test images. Running the text detection and text recognition methods on these images, we correctly recognized "Sparkasse", "Engelhorn" and "venicebeach", as depicted in Figure 1. On the other hand, "C&A" and "H&M" were either not detected or not correctly recognized (i.e., even Google Places API was not able to return the desired address). Hence, these cases are considered as *false negatives* in our evaluation. In general, we observed that the selected (text detection & recognition) methods tend to find texts that have lighter background as well as the ones in which there are less spaces between the characters and where characters are not mixed with numbers and symbols.

Table 1 shows the evaluation result on the test data. The text detection method shows a satisfactory performance (detection recall of 72 %) yet without any optimization. To better demonstrate the effectiveness of the solution, we do not only use the prevalent metrics in computer vision literature, like recall, but also some new metrics which we believe are more relevant and critical in our case. We define a "privacy-sensitive text block" in a given image as a piece of information that returns a specific address when queried against the Google Places API. Hence, the privacy-sensitive recall factor shows explicitly the detected share of all existing privacy-sensitive information in our test dataset (118 out of 155). By calling the Google Places API for all detected texts, we could avoid anonymizing almost all 160 false positive detected texts and therefore, prevent distorting images. However, we could only anonymize 63 text blocks out of 155 detected privacy-sensitive text blocks due to low performance of the recognition method.

Table 1: Performance of Text Detection and Recognition Methods on Test Data

Detection recall	Privacy-sensitive recall	Anonymized privacy-sensitive texts
72%	76%	40%

4 CONCLUSION

The road scene images captured by autonomous cars contain various privacy-sensitive data. In this paper, we evaluated a machine learning approach to recognize and anonymize textual data in such images. Our preliminary analysis shows that the idea of anonymizing the privacy-sensitive text in road scene images could be promising. In future work, we will focus on recognition models that are more suitable for a road scene setup. In parallel, the detection model could be enhanced to decrease the false negative rate. We note that we have assumed that a service similar to Google Places API exists that supports us to recognize privacy-sensitive texts.

ACKNOWLEDGMENTS

This project has received funding from the ECSEL Joint Undertaking (JU) under grant agreement No 783119. The JU receives support from the European Union's Horizon 2020 research and innovation program and Netherlands, Austria, Belgium, Czech Republic, Germany, Spain, Finland, France, Hungary, Italy, Poland, Portugal, Romania, Sweden, United Kingdom, and Tunisia.

REFERENCES

[1] 2018. MYDATA Control Technologies. https://www.mydata-control.de/de/. Accessed: 2019-08-06.
[2] 2019. Google Places API. https://developers.google.com/places/web-service/search. Accessed: 2019-08-06.
[3] Cara Bloom, Joshua Tan, Javed Ramjohn, and Lujo Bauer. 2017. Self-driving cars and data collection: Privacy perceptions of networked autonomous vehicles. In *Thirteenth Symposium on Usable Privacy and Security ({SOUPS} 2017)*. 357–375.
[4] European Parliament. 2016. "General Data Protection Regulation," EU/2016/679.
[5] Huaxin Li, Haojin Zhu, Suguo Du, Xiaohui Liang, and Xuemin Sherman Shen. 2016. Privacy leakage of location sharing in mobile social networks: Attacks and defense. *IEEE Transactions on Dependable and Secure Computing* 15, 4 (2016), 646–660.
[6] Shangbang Long, Xin He, and Cong Yao. 2018. Scene Text Detection and Recognition: The Deep Learning Era. *CoRR* abs/1811.04256 (2018). arXiv:1811.04256 http://arxiv.org/abs/1811.04256
[7] Baoguang Shi Minghui Liao and Xiang Bai. 2018. TextBoxes++: A Single-Shot Oriented Scene Text Detector. *IEEE Transactions on Image Processing* 27, 8 (2018), 3676–3690. https://doi.org/10.1109/TIP.2018.2825107
[8] Robust Reading Competition. 2017. ICDAR2017 Competition on Multi-lingual scene text detection and script identification. https://rrc.cvc.uab.es/?ch=8. Accessed: 2019-08-07.
[9] Baoguang Shi, Xiang Bai, and Cong Yao. 2017. An End-to-End Trainable Neural Network for Image-Based Sequence Recognition and Its Application to Scene Text Recognition. *IEEE TPAMI* 39, 11 (2017), 2298–2304.