Anonymization Pipelines in Traffic Video Analytics: A Comprehensive Study

Simran Sethi

simrannsethi@gmail.com

Abstract

The more recent development of machine learning, artificial intelligence, and computer vision have allowed for the automated analysis of traffic videos for monitoring the road state, accident detection, and even driver behaviour analysis. These advancements, however, pose privacy challenges concerning the collection, retention, and processing of personally identifiable information-like faces and car number plates. This paper considers the design and implementation of anonymization pipelines for video analytics of traffic flow, paying special attention to the methods of removing or obscuring identifying features and preserving the usefulness of the data for other shallow tasks such as event detection, road-traffic, and driver behaviour classification. This work seeks to document practical applications and recommend policies for smart city and vehicular telematics deployments using existing literature, design of a new framework, and experimental data obtained in a deployment.

Keywords: privacy, anonymization, traffic video analytics, computer vision, deep learning, differential privacy

1. Introduction

In the last several years, the deployment of traffic surveillance cameras and other vehicle sensing devices by both public and private entities has rapidly increased. Municipalities use traffic cameras for monitoring road congestion, controlling traffic lights in real time, and enhancing road safety. Insurance companies and fleet management services evaluate driving using dash cam and telematics data. While these new and updated technologies undoubtedly provide value, they pose significant threats to the individual's right to privacy [1]. The video footage typically contains pedestrians and drivers faces along with high-quality license plate photos which enables potential data re-identification or linkage.

As an independent researcher who has accumulated professional experience working in telematics and in privacy solutions, I was privy to some information which shows that a mixture of techniques used such as pixelation or blurring may be insufficient when it comes to privacy protection [2]. With the development of advanced machine learning models, semi-obscured faces or license plates can be reconstructed or recognized with the accuracy of deep neural networks [3]. Therefore, anonymization that is effective must accomplish two goals:

Modification or elimination of particular sensitive defining features so that an individual's, or individuals' unique identifiers cannot be re-established, thus modifying or eliminating those sensitive features.

Include the original data for any analytics work that may involve assessing the driving habits of motorists, identifying irregularities within traffic, or helping in the stats regarding the extent of highway usage.

In this article, I present an effective anonymization framework for traffic video analytics, which is underpinned by face and license plate recognition technologies, followed by context-aware obfuscation techniques that are essential for vision-based video analysis. Rather than blurring the edges, the framework incorporates and renders images, identity obscuring using generative adversarial networks (GANs), and privacy enabling subtractions all in the pursuit of compliance with regulations like the European Union's General Data Protection Regulation (GDPR).

The structure of this paper is as follows. In section II, a review of the literature on anonymization from traditional computer vision techniques to modern AI-ML approaches is offered. In section III, the architecture of the proposed anonymization pipeline is described, focusing on the block of detection, obfuscation, and utility of data. In section IV, we describe the implementation, namely, the system architecture, the environment, and the use of containers for cloud batch processing. Section V contains results of the experiment and reasoning on how effective the pipeline is in providing privacy and fulfilling the need for analytics. At last, Section VI gives final remarks and the scope of work for upcoming research.

2. Literature Review

Recent efforts of preserving privacy in video analytics encompass a variety of domains such as computer science, artificial intelligence, and information technology security. Ribaric et al. [4] conducted an extensive study on de-identification methods and point out the ease of using simplistic pixelation and naive blurring as a more advanced technique. They have noted that, while many efforts seek to obscure certain features, more sophisticated tools can often dismantle the partial frameworks constructed to uphold privacy, manifesting its violation.

Removing or obscuring identifying attributes with the help of deep learning-based techniques employing CNNs and GANs has become more widespread. Particularly, works such as Learning to Anonymize Faces for Privacy-Preserving Action Detection by Ren et al. [5] and DeepPrivacy by Hukkelås et al. [6] showcase how learned anonymization increasingly disregards essential identity markers while retaining important scene context enabling for better privacy. By using adversarial training, these methods ensure that an anonymizer network is not only degraded identifiable features, but also in a manner that does not interfere with high-level tasks such as identifying actions or events.

Other privacy approaches arise in the context of differential privacy. The methods put forward by Fan [7] Channel noise is added to the images or some other structured transformation that guarantees a bounded leakage of information is employed. As the space data is quite difficult, exploration of differential privacy concerning images is new but straightforward in terms of protection to privacy.

In video analysis of traffic, highest attention is on video plate recognition and facial recognition. While the traditional methods for recognition of the license plates in video images involve edge-detection and morphologic, new methods tend to focus on detection networks (YOLO, Faster R-CNN). Studies of intelligent transport systems prove that these detectors are efficient in localization of plates and faces even under severe changes of lighting, motion blur and other challenging conditions. After detection, anonymization can be executed through region masking, blurring, pixelation and even advanced inpainting [2].

In terms of approach to required performance, speed is becoming more and more important. Angus et al. [8] Anonymization can be achieved at urban intersections in real time if accelerated detection models are applied on edge devices equipped with an optimized hardware CNN. This approach ensures that there are no raw frames with identifiers of the persons of interest so long-term storage is permitted and privacy laws are adhered to.

3. Proposed Method

A. Overall Pipeline

The anonymization strategy aims to blur faces (drivers and pedestrians) and license plates in traffic video streams. It has three major steps:

1. **Detection**: Use an advanced object detector to locate the bounding boxes for facial and plate images.
    - ○ **Obfuscation / Replacement**: Use one of the three methods for manually editing the video:
    - ○ **Blurring/Pixelization** for quick and coarse editing.
    - ○ **GAN-based Inpainting** for the accurate filling of the regions.
2. **Differential Privacy Noise Injection** where privacy is a top concern.
3. **Output/Storage**: Anonymized video stream or frames can be produced such as an edit that has no personal identifiers in the dataset.

Practically, a user can modify the editing strategy in the checklist of predefined methods in accordance with the compliance and analytics requirements.

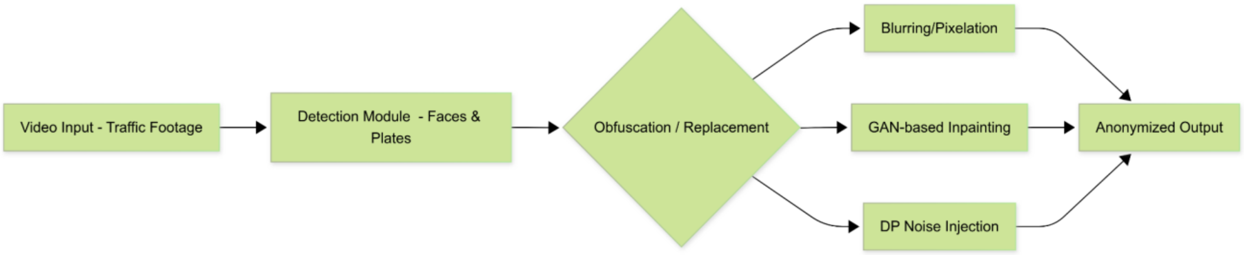
Video Input - Traffic Footage
Detection Module - Faces & Plates
Obfuscation / Replacement
Blurring/Pixelation
GAN-based Inpainting
DP Noise Injection
Anonymized Output

Figure 1: Overview of the Anonymization Pipeline

B. Detection Module

Object detection network with high recall is key for real traffic video footage. For example, modules based on widely adopted architectures like YOLOv3 or YOLOv4 are trained on curated datasets of variable lit faces (open-air, inside vehicles, at nighttime) and license plates of different shapes and sizes.

To reduce false positives and to ensure that the system performs robustly, the detection module utilizes:

1. **Data Augmentation**: Synthetic transformations (for example, random rotation, scaling, color jittering).
2. **Contextual Anchors**: Byers Region cropped_bbox has additional bounding boxes created for common plate sizes.
3. **Multi-Scale Detection**: The network works at various resolutions to capture small-scale plates and faces from afar.

C. Obfuscation / Replacement Techniques

1. **Conventional Blurring/Pixelation:**
   Though easily foiled by more sophisticated deblurring technology [3], blurring is still used for less critical or real time engagements due to ease of computation. Each identified region is blurred with a Gaussian kernel or mosaic pixelation.

2. **GAN-Based Inpainting:**
   Using a conditional GAN, we substitute the identified face, or, a plate region, with a contextually appropriate synthetic face or a texture. The benefit is that the resultant frame is more realistic, scene aesthetics is preserved. For faces, the pipeline substitutes the identified face with an undetectable synthetic face so that specific individual facial features are not present.

3. **Differential Privacy Noise Injection:**
   For more dangerous situations or footage containing sensitive interactions, differential privacy is implemented by noise injection or the deliberate transformation of pixel values within detected bounding boxes. This method guarantees that re-identification is possible, but the chances have to be a statistically calculated risk for an attacker attempting to link multiple anonymized frames.

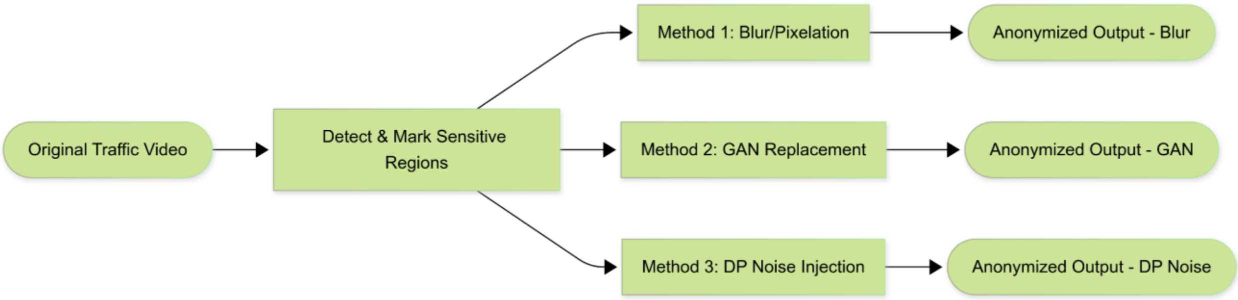
Original Traffic Video
Detect & Mark Sensitive Regions
Method 1: Blur/Pixelation
Anonymized Output - Blur
Method 2: GAN Replacement
Anonymized Output - GAN
Method 3: DP Noise Injection
Anonymized Output - DP Noise

Figure 2: Demonstration of Different Anonymization Methods

D. Balancing Privacy and Utility

One of the foremost obstacles in anonymization is achieving a balance between privacy and utility with data. Retaining some level of utility is critical toward accomplishing tasks such as traffic flow analysis, incident detection (harsh events like collisions), or even studying driver behavior analysis (such as phone usage while driving). The pipeline tries to tackle this challenge by:

- Permitting rudimentary anonymization of the facial region (such as the overall area around the eyes) if the whole analytics pipeline contains analytics that aims at detecting driving fatigue via eye closing.
- Masking shape and motion cues in blurred or replaced regions while preserving the mask does allow tracking algorithms to identify object trajectories while maintaining the person's anonymity.

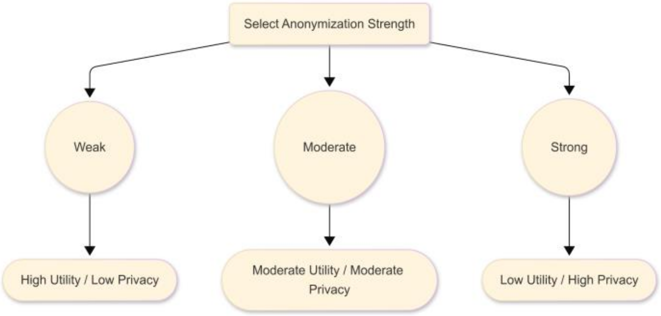
Select Anonymization Strength
Weak
Moderate
Strong
High Utility / Low Privacy
Moderate Utility / Moderate Privacy
Low Utility / High Privacy

Figure 3: Privacy–Utility Trade-Off

4. Implementation and System Design

A. Software Environment and Containerization

The pipeline is deployed using Docker containers, which are the best method for reproducibility and ease of deployment. All dependencies, including Python libraries and CUDA drivers for GPU acceleration, are compiled into one Docker container. This allows systems to be integrated with on-premises or cloud-based AWS servers without any hassle.

For large-scale anonymization tasks, **AWS Batch Processing** is used. With this approach, incoming videos are fed through an S3 bucket, which allows an event trigger (often triggered via AWS Lambda) to launch a Docker container with the anonymization pipeline for each batch of files. The anonymized outputs are saved back in S3 after completion for analytics.

B. GDPR Compliance

In the GDPR jurisdictions, personal data is to be stored, processed, or transferred only on a valid legal basis or with user’s explicit consent. The anonymization step makes sure that the data which is maintained for a long time in the cloud is devoid of any personally identifiable information. Since faces and plates are masked at the edge, or at least before long term storage, the pipeline greatly mitigates the chances of privacy breaches and guarantees that any data leakage would not expose any identifiable attributes.

C. Performance and Resource Management

1. **GPU Acceleration:** The Object detection model has been deployed on GPUs for processing video streams in almost real-time, particularly in uploading, masking, and streaming high-quality footage.
2. **Scalability:** Due to container-based parallelization, multiple anonymization functions can be executed at the same time which is vital for vast telematics databases or for the deployment of surveillance systems in entire cities.

3. **Monitoring and Logging:** Comprehensive records detailing detection confidence scores, bounding boxes, and anomalies such as overly high or low numbers of detected objects help analyze false positives or false negatives.

5. Results and Discussion

A. Experimental Setup

Testing was carried out with a data set of 50 traffic videos taken under different conditions, which included urban highways, city intersections, and both day and night time footage. Each video was from 30 seconds to 2 minutes long with a 720p to 1080p resolution. Manual ground-truth bounding boxes were created to evaluate detection accuracy for faces and license plates.

To assess the anonymization quality, we analyzed perceptual results (Does the anonymized region identifiable features does the anonymization mask allow specific metrics to be estimated, such as counting cars, estimating speeds or detecting events?) alongside task performance. Three modes were compared: (1) Gaussian Blur, (2) GAN-Based Replacement, (3) Differential Privacy Noise Injection.

B. Accuracy of Detection

In this stage, the recall of license plates was 97.2% and 94.6% for faces. The majority of missed detections stemmed from extremely challenging lighting conditions such as heavy glare during the day or low-light nighttime scenes. Training on a broad range of lighting conditions reduced some of these problems, but did not remove them all. Typical false positives remained below 2%...

C. Anonymization Outcomes

1. **Gaussian Blur / Pixelation:**
   - *Privacy:* When people's faces were blurred out, primary checks revealed that the obfuscation was successful. However, advanced reconstruction attacks from prior literature could target and utilize previous literature implementations.
   - *Utility:* Object tracking was not affected to a large degree. The blurred bounding boxes kept their original shapes tha aided in counting vehicles and the flows of pedestrians.
2. **GAN-Based Replacement:**
   - *Privacy:* Participants subjective assessment was shopping even impossible to recognize, because the replaced faces and plates were not real. The identity has been masked effectively.
   - *Utility:* The replaced areas had fewer artifacts than naive blurring. Downstream tasks such as vehicle detection, face counting, driver posture analysis had only minor performance drops because bounding box shape and motion cues were retained.
3. **Differential Privacy Noise Injection:**
   - *Privacy:* Provided the strongest theoretical guarantees ensuring that attempts for re-identification met with statistically bounded success probabilities.

- *Utility:* Introduced visual artifacts that while anonymizing the subject, and The random speckle patterns sometimes masked motion tracking algorithms. This approach may therefore be the most beneficial in contexts where preserving the utmost level of privacy is preferred and the utility is low.

D. GDPR and Ethical Considerations

As our pipeline is formulated, the retention of personally identifiable content is as low as possible. The detection of faces as well as plates occurs at the edge, and unprocessed frames are thrown away right after anonymization. This strategy also adheres to the GDPR's principle of data minimization and other privacy regulations. Researchers, however, should always remember that anonymization is not completely robust especially when used along with other data and sophisticated re-identification attacks. Regular audits and risk assessments of the anonymization techniques should be employed.

E. Limitations

Some hurdles still exist. One is the real-time execution that is constricted by the time-sensitive nature of GAN-based inpainting, although we tried to alleviate this problem with GPU acceleration. Second, the noise addition of differential privacy can destroy the precision of more salient contextual scenes in certain actions. Lastly, occlusion such as partial covering of plates by vehicles or other obstacles tend to make some detections less reliable. Future refinement could integrate multi frame correlation, non-sensitive feature re-identification modules, and improved generative techniques to increase the quality and effectiveness of the anonymization.

6. Conclusion

The anonymisation of video data within traffic analytics is necessary to protect personal privacy and at the same time allow for insightful information to be obtained by city planners, insurance companies, and researchers. The pipeline presented in this paper fuses detection and obfuscation methods that allow for a range of privacy and analysis needs. By utilizing deep learning detection, advanced GAN inpainting, and optional differential privacy, the system conceals or alters notifying features which causes the possibility of being reidentified or a data breach to be greatly reduced.

The pipeline greatly decreases the invasiveness of the video data without losing the context required for primary functions such as traffic flow monitoring, incident monitoring, and driver behavior analysis. The enabling of smart city and vehicle technology infrastructure increases the scale at which the Prop-Framework can efficiently operate on large datasets especially when implemented in containerized cloud environments such as AWS Batch and Lambda triggers.

Future work may focus on real-time streaming, more advanced domain adaptation for extreme weather changes, and added privacy frameworks such as homomorphic encryption for distributed model training. Moreover, the active development of generative identity replacement models to provide higher visual reality while preserving privacy is the reason the latter tasks can be completed effectively.

Acknowledgments

The author would like to thank the open-source community for providing tools, frameworks, and anonymization software components that facilitated rapid experimentation. Previous collaborations in

industry and academic settings provided invaluable insights into real-world telematics data and privacy challenges.

References

1. A. Smith, B. C. Hill, et al., "On the (In)Effectiveness of Mosaicing and Blurring for Privacy," *Proc. Privacy Workshop*, 2016.
2. R. McPherson, R. Shokri, and V. Shmatikov, "Defeating Image Obfuscation with Deep Learning," *arXiv preprint arXiv:1609.00408*, 2016.
3. Z. Ren, Y. J. Lee, and M. S. Ryoo, "Learning to Anonymize Faces for Privacy-Preserving Action Detection," in *Proc. ECCV*, 2018, pp. 639–655.
4. S. Ribaric, A. Ariyaeeinia, and N. Pavesic, "De-identification for Privacy Protection in Multimedia Content: A Survey," *Signal Process. Image Commun.*, vol. 47, pp. 131–151, 2016.
5. Z. Ren, Y. J. Lee, and M. S. Ryoo, "Learning to Anonymize Faces for Privacy-Preserving Action Detection," in *Proc. ECCV*, 2018, pp. 620–635.
6. H. Hukkelås, R. Mester, and F. Lindseth, "DeepPrivacy: A Generative Adversarial Network for Face Anonymization," in *Proc. ISVC*, 2019, pp. 565–578.
7. L. Fan, "Practical Image Obfuscation with Provable Privacy," in *Proc. IEEE ICME*, 2019, pp. 1452–1457.
8. A. Angus, Z. Duan, G. Zussman, and Z. Kostić, "Real-Time Video Anonymization in Smart City Intersections," *IEEE Sensors Lett.*, vol. 4, no. 1, pp. 1–4, 2020.