INSPIRE: Instance-level Privacy-preserving Transformation for Vehicular Camera Videos

Zhouyu Li, **Ruozhou Yu**, Anupam Das, Shaohu Zhang, Huayue Gu, Xiaojian Wang, Fangtong Zhou, Aafaq Sabir, Dilawer Ahmed, Ahsan Zafar

North Carolina State University

Outlines

Background and Motivation

Threat Model

Framework Design and Implementation

Performance Evaluation

Discussions, Future Work and Conclusions

Trending and Special Attributes of Vehicular Cameras

vehicular cameras are more and more popular

straits RESEARCH

Vehicle Camera Market Size is projected to reach USD 17.68 billion by 2030, growing at a CAGR of 10%: Straits Research

Four special attributes:

Ubiquity	High mobility	Privately-owned data	Lack of interface
↓	↓	↓	↓
Unavoidable	Large coverage	Less regulation	No opt-out method

C. Bloom, J. Tan, J. Ramjohn, and L. Bauer, “Self-Driving Cars and Data Collection: Privacy Perceptions of Networked Autonomous Vehicles,” p. 21.

Privacy Concerns of Vehicular Cameras

- Bystanders' feelings for vehicular camera video usages
- **Strong discomfort** for recognizing, identifying and tracking individuals/vehicles

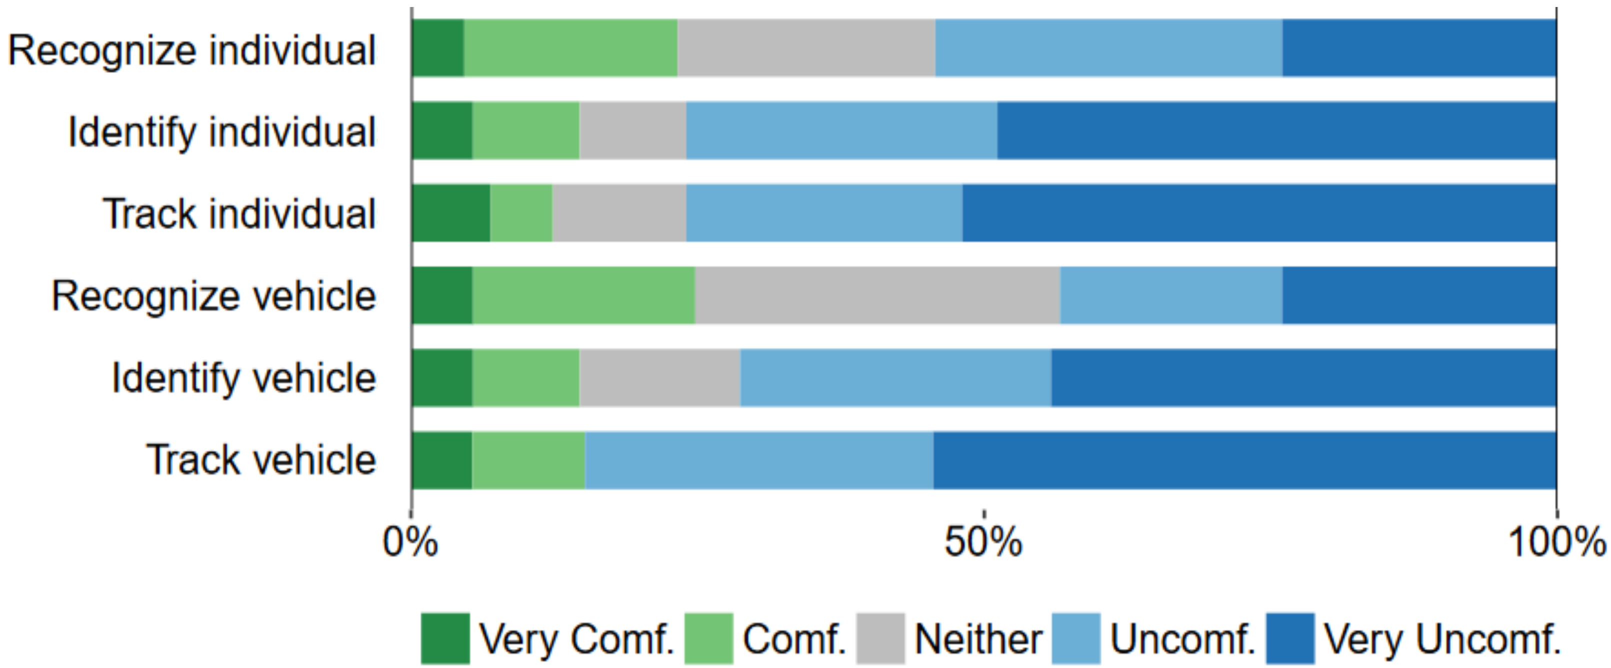

C. Bloom, J. Tan, J. Ramjohn, and L. Bauer, "Self-Driving Cars and Data Collection: Privacy Perceptions of Networked Autonomous Vehicles," p. 21.

Privacy Concerns of Vehicular Cameras

- Videos shared for different purposes.
- Attackers can launch attacks like **location inference attacks**.

Evidence grounding

Trip sharing

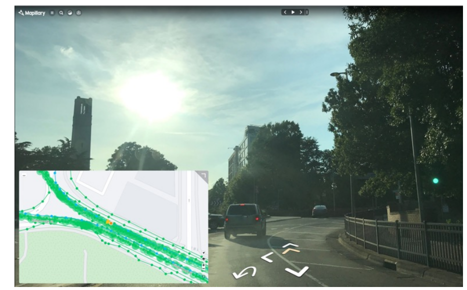

Street view building

Victim near the Triumphal Arch

Victim near the Eiffel Tower

Leaked location and trajectory

NC STATE UNIVERSITY

Z. Xiong, W. Li, Q. Han, and Z. Cai, “Privacy-Preserving Auto-Driving: A GAN-Based Approach to Protect Vehicular Camera Data,” in 2019 IEEE International Conference on Data Mining (ICDM), Beijing, China, Nov. 2019, pp. 668–677. doi: 10.1109/ICDM.2019.00077.

Current Countermeasures

- **Dashcam Cleaner**: blur faces and license plates
- **SecGAN**: blur the whole video

Dashcam Cleaner

Use pre-defined sensitive attributes

A. Nodari, M. Vanetti, and I. Gallo, “Digital privacy: Replacing pedestrians from Google Street View images,” p. 5.

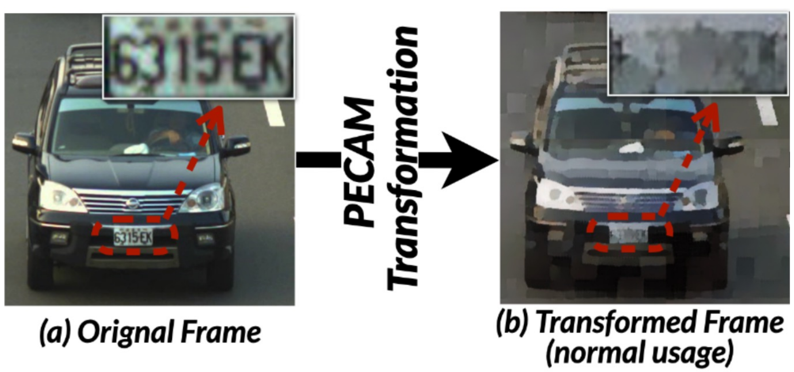

SecGAN

Also blur non-sensitive details

R. Uittenbogaard, C. Sebastian, J. Vijverberg, B. Boom, D. M. Gavrila, and P. H. N. de With, “Privacy Protection in Street-View Panoramas Using Depth and Multi-View Imagery,” in 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Long Beach, CA, USA, Jun. 2019, pp. 10573–10582. doi: 10.1109/CVPR.2019.01083.

INSPIRE Overview

- Replace protected instances with AI-synthesized non-existent counterparts

Original

Transformed

Outlines

Background and Motivation

Threat Model

Framework Design and Implementation

Performance Evaluation

Discussions, Future Work and Conclusions

Re-identification (Re-ID) Attack

- Re-identification (Re-ID) attack: finding the same instances across different images with deep-learning models.

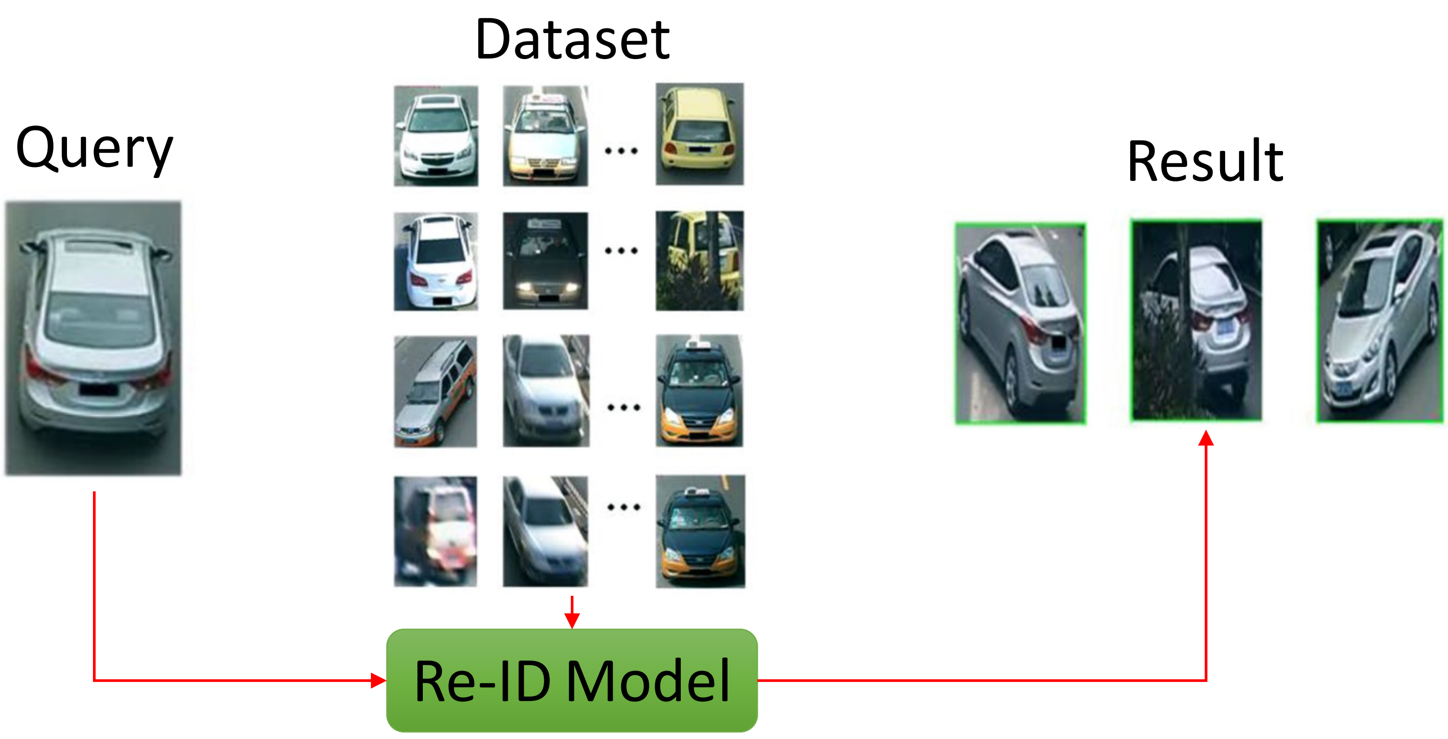

Threat Model

- INSPIRE as a Software plugin on Car's On-Board Unit or Mobile devices.
- Video contents are in a trusted environment before transformation, and exposed to attackers after transformation.

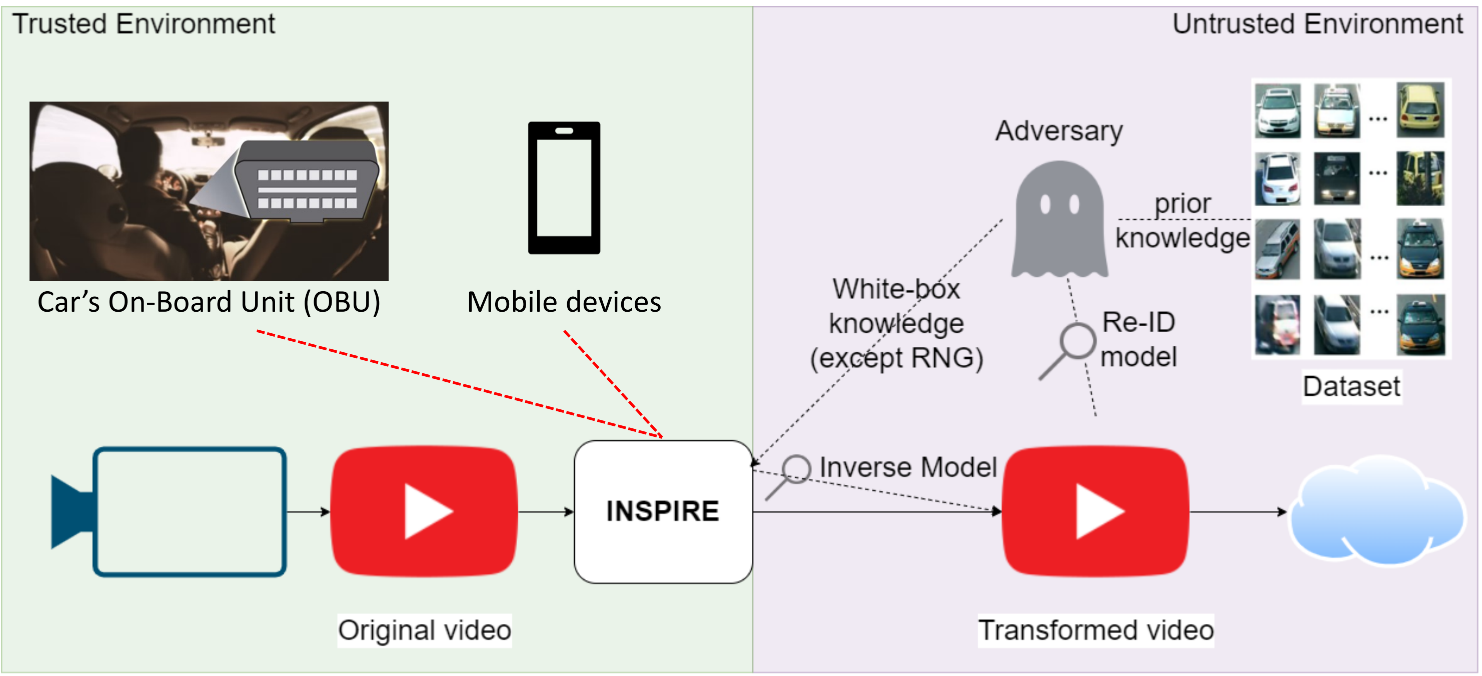

*RNG: random number generator

Threat Model

- Attackers have white-box access to the system.
- Attackers launch Re-ID attack and Model-inversion attack to transformed videos.

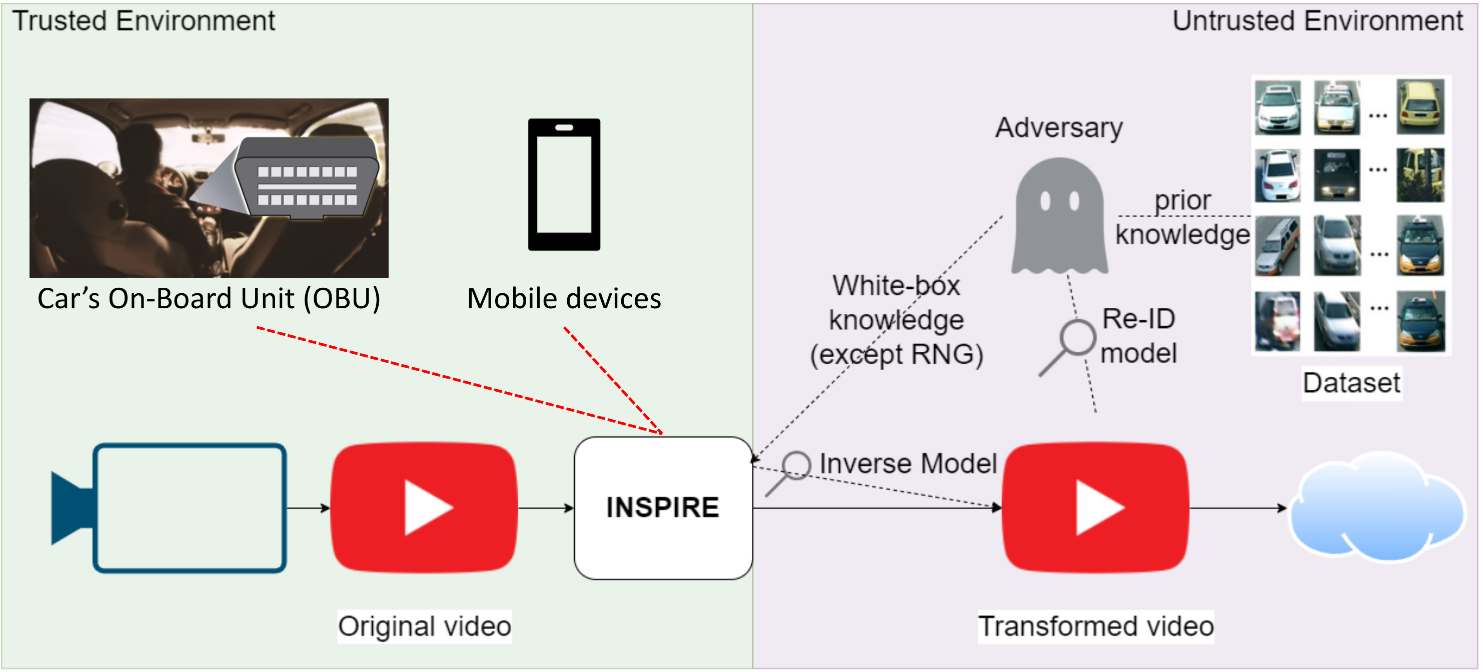

*RNG: random number generator

Outlines

Background and Motivation

Threat Model

Framework Design and Implementation

Performance Evaluation

Discussions, Future Work and Conclusions

Framework Design

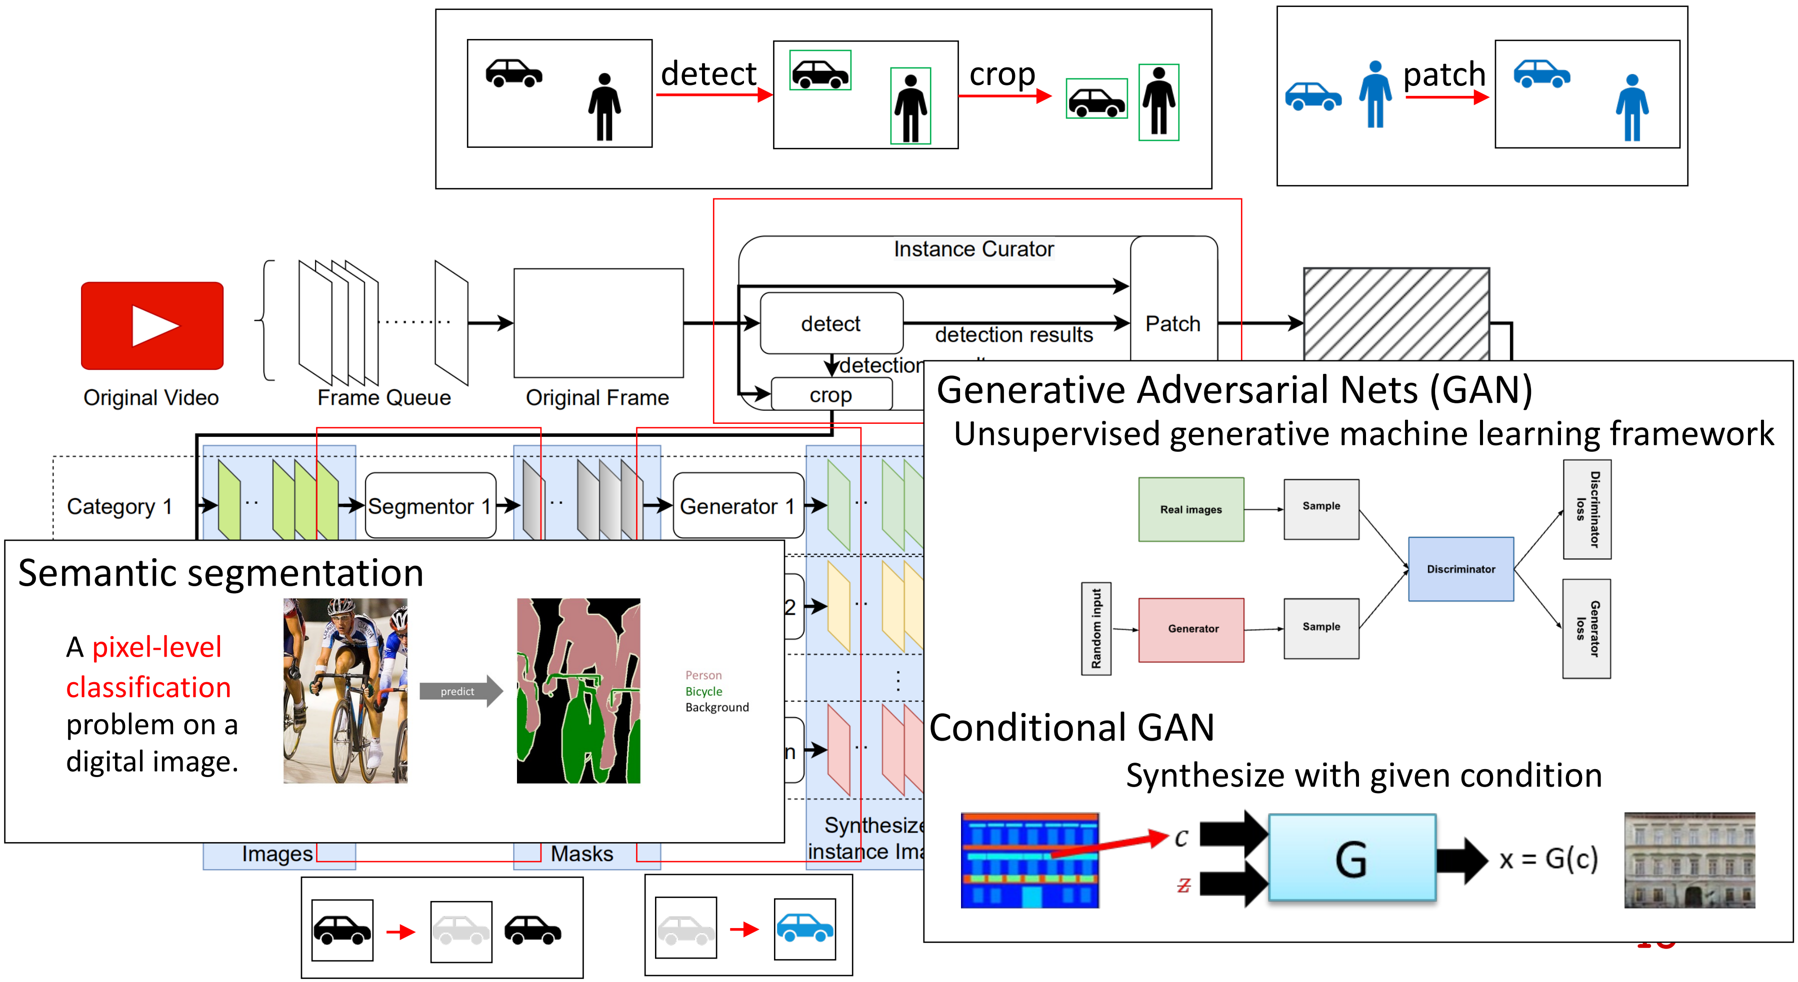

System Implementation

- An INSPIRE system protect people and cars in the vehicular video.

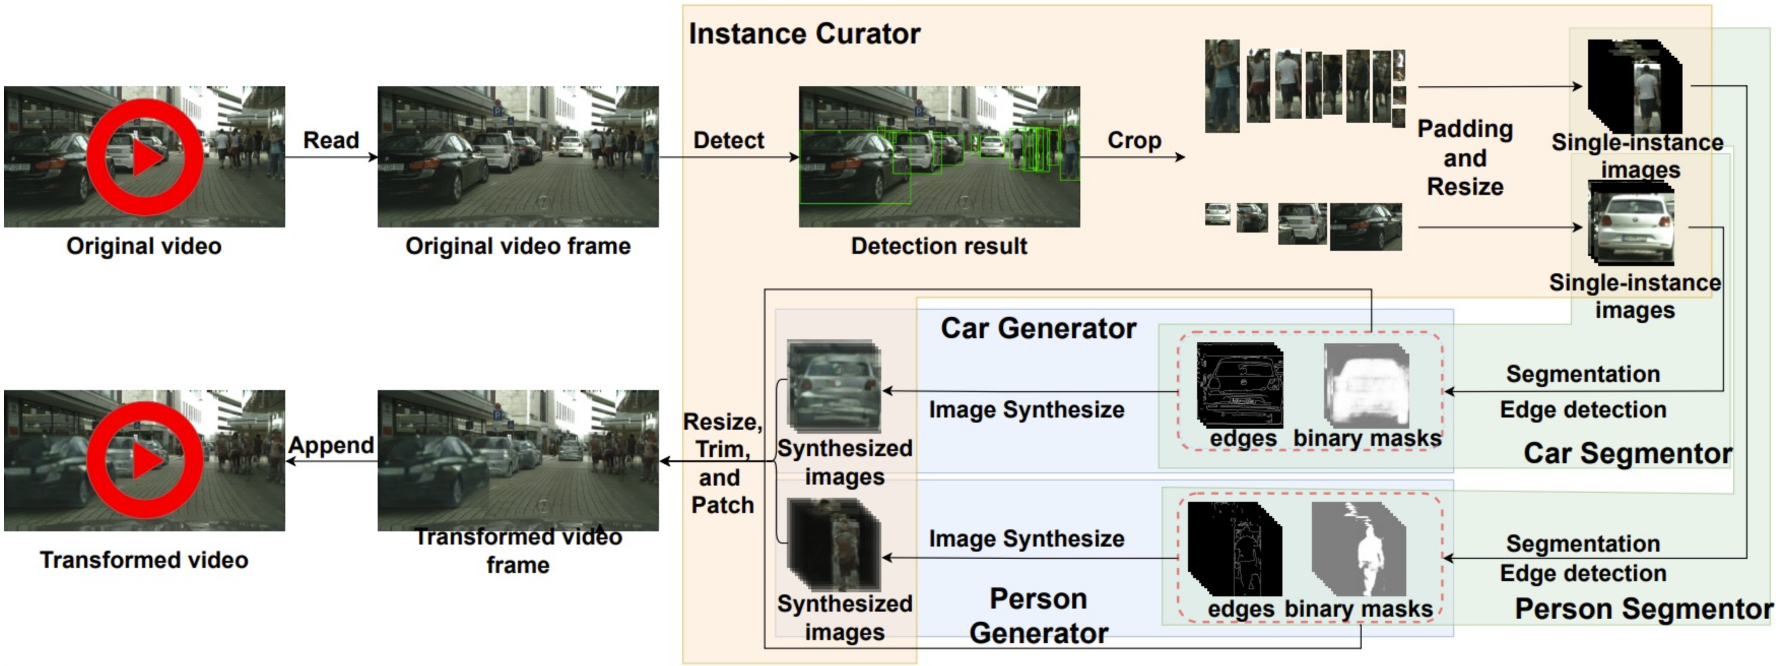

System Implementation

- Object Detection: YOLOv5
- Semantic Segmentation: U-Net
- Instance Synthesis: Pix2pixHD

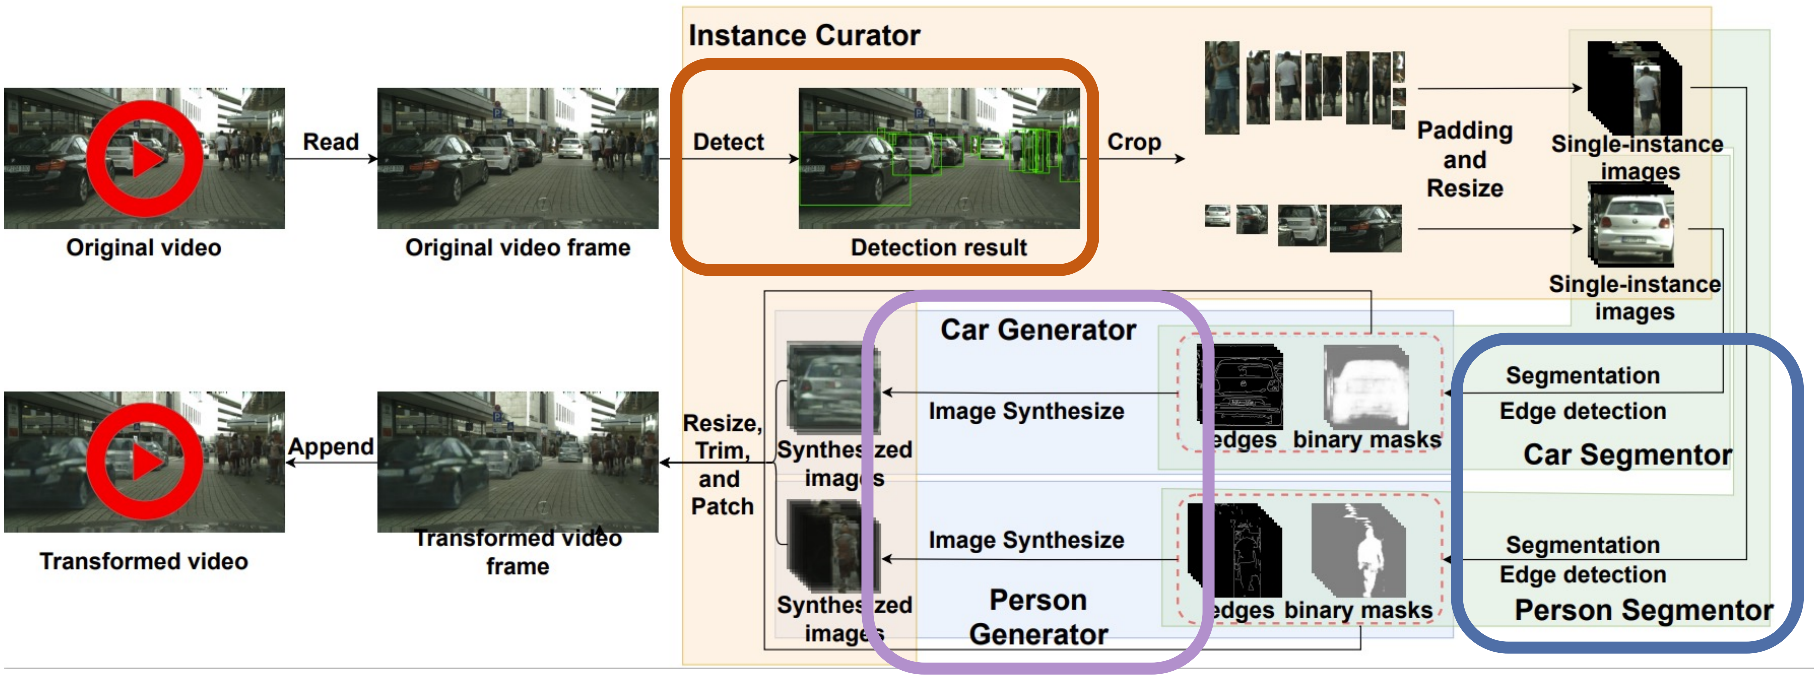

System Implementation

- **Challenge**: Contour alone cannot deal with overlapped instances.
- **Solution**: Use both contour and edge detection result for instance synthesis.

Original

Contour Only

With Edge Detection

System Implementation

- Challenge: How to control privacy leaked by edge information?
- Solution: Apply a Gaussian filter before edge detection.

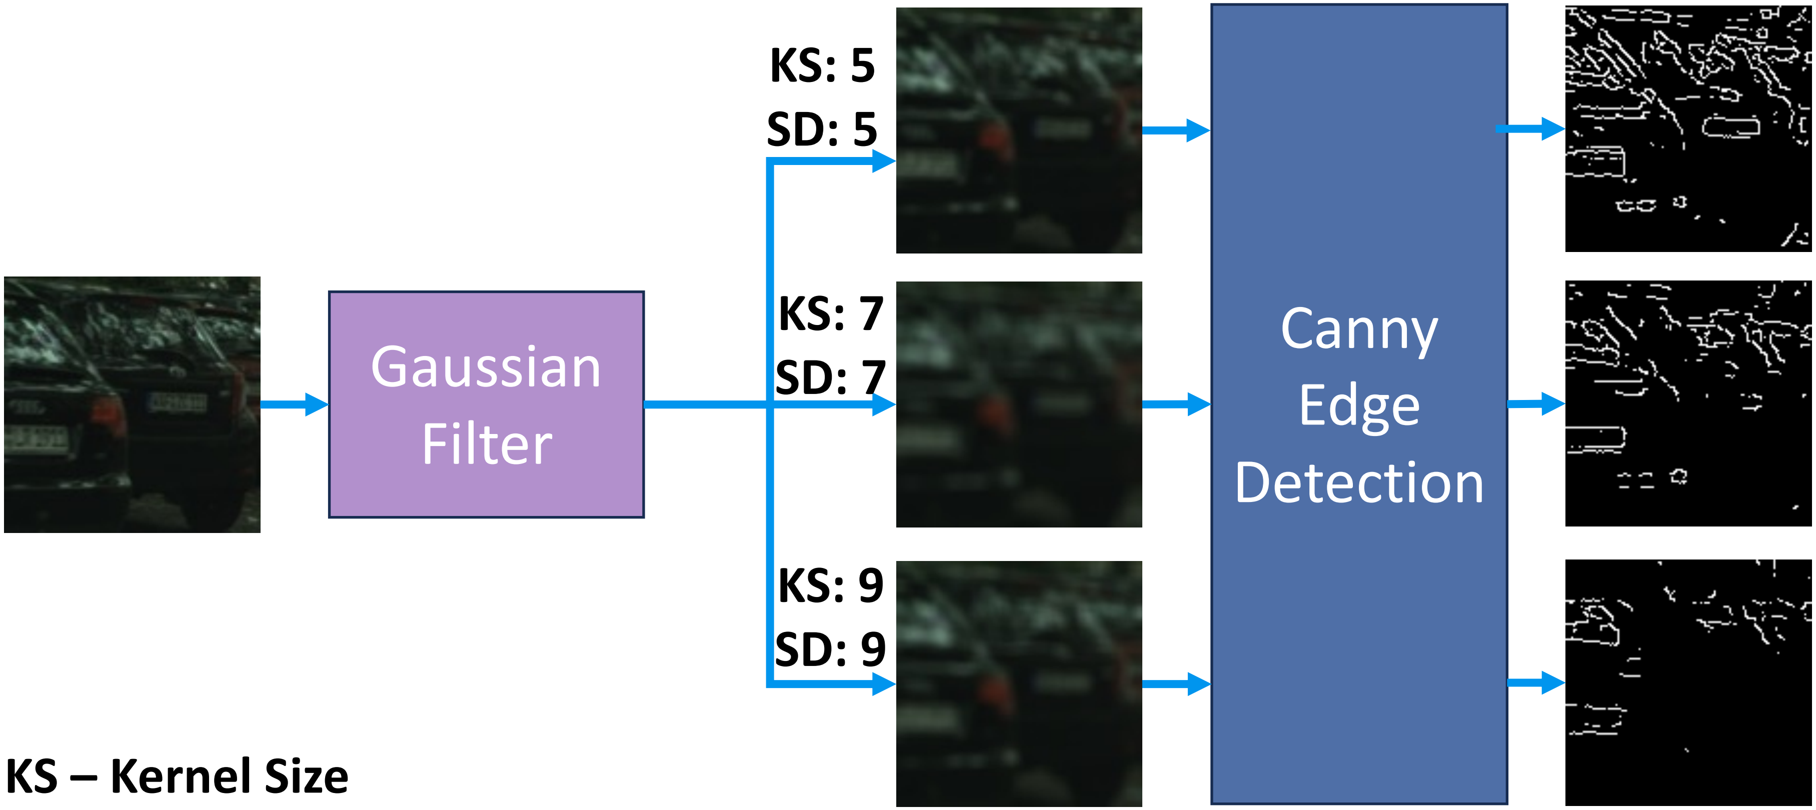

KS – Kernel Size
SD – Standard deviation

Outlines

Background and Motivation

Threat Model

Framework Design and Implementation

Performance Evaluation

Discussions, Future Work and Conclusions

Compared Systems

- INSPIRE: Replace instances with synthesized counterparts.
- SecGAN: Blur the whole video frame

Original

INSPIRE

SecGAN

Compared Systems

- Dashcam Cleaner: Blur faces and license plates.
- Bbox Blur: Blur instances according to their object detection bounding boxes with Gaussian filters.

Original

Dashcam Cleaner

Bbox Blur

Evaluation Settings – Privacy & Utility

- Settings
  - Privacy:
    - Re-ID attack
      - Image-wise thwarting rate
    - Model inversion attack
      - Train adversarial models
  - Utility:
    - Statistical counting
      - Accuracy
    - Object detection
      - mean average precision

TABLE II: Details about Re-ID datasets

Name	Query images	Gallery images	Gallery instances	Category	Real world
Cityscapes (person)	4924	4924	267	person	✓
Duck MTMC	2228	17661	1110	person	✓
Market-1501	3368	19732	752	person	✓
Cityscapes (car)	10450	10450	147	car	✓
VeRi	1678	11579	200	car	✓
VeRi-CARLA	424	3823	50	car	✗

TABLE III: Details about utility evaluation datasets.

Dataset Names		Number of videos	Average people per frame	Average cars per frame
Cityscapes		3	5.70	4.68
Accident	Positive	17	2.08	4.45
	Negative	31	2.60	4.82
BDD100K		54	0.95	4.04

Re-ID Attack: Influence of Gaussian filters

- Applying the Gaussian filter in INSPIRE can improve and stabilize the protection performance against Re-ID attacks.
- In INSPIRE, applying a Gaussian filter with small kernel size and SD is sufficient to thwart most Re-ID attacks.
- For INSPIRE and BBox Blur, improving the kernel size and SD of the Gaussian filter enhances the Re-ID thwarting rate.

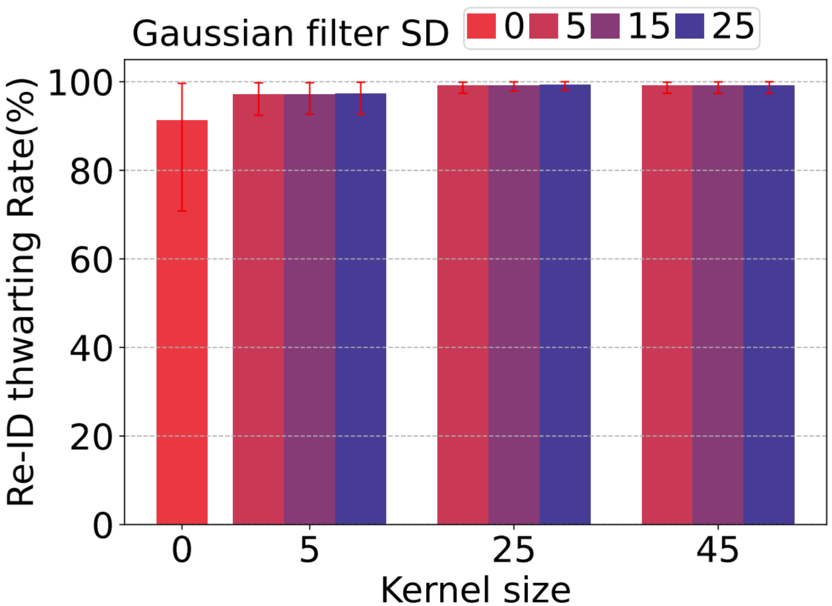

INSPIRE

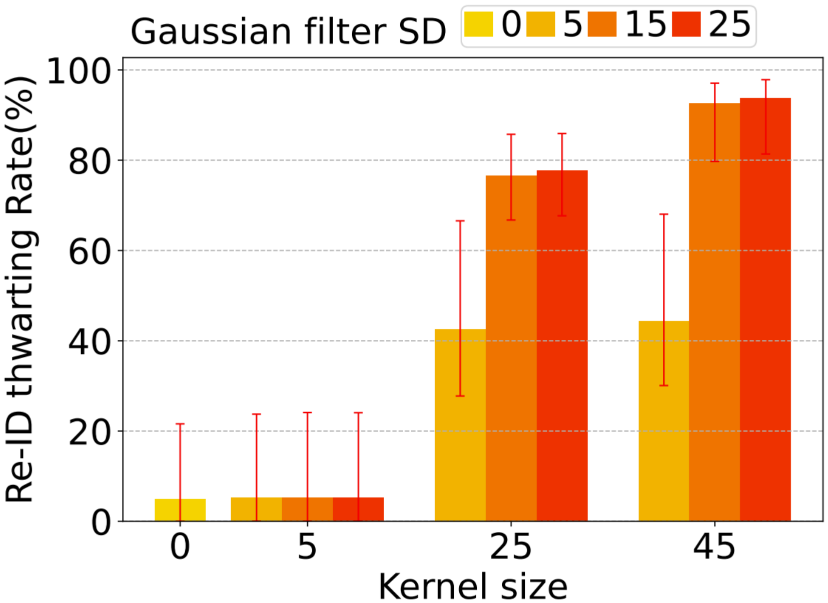

Bbox Blur

Re-ID Attack: System-wise Comparison

- In practice, INSPIRE with Gaussian filter can effectively thwart Re-ID attacks for its protected instances.
- Attribute-level and frame-level obfuscation cannot thwart Re-ID attacks with state-of-the-art deep learning models

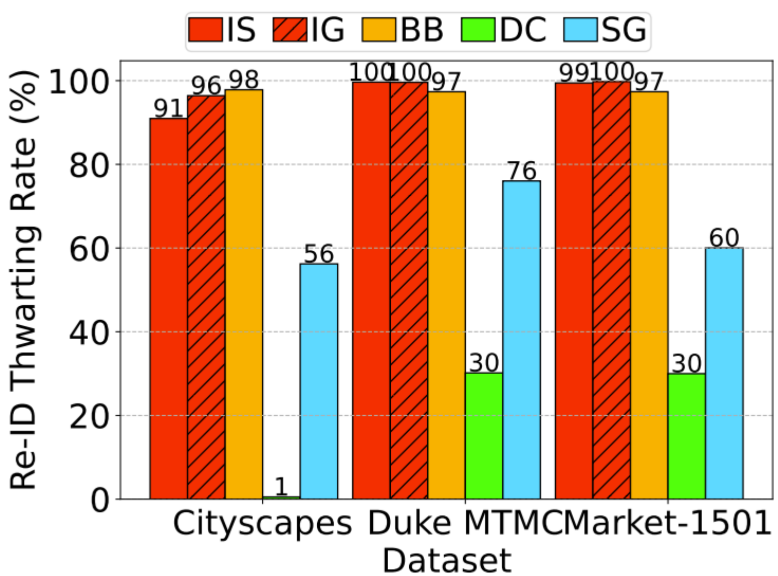

(a) Person Re-ID thwarting rates

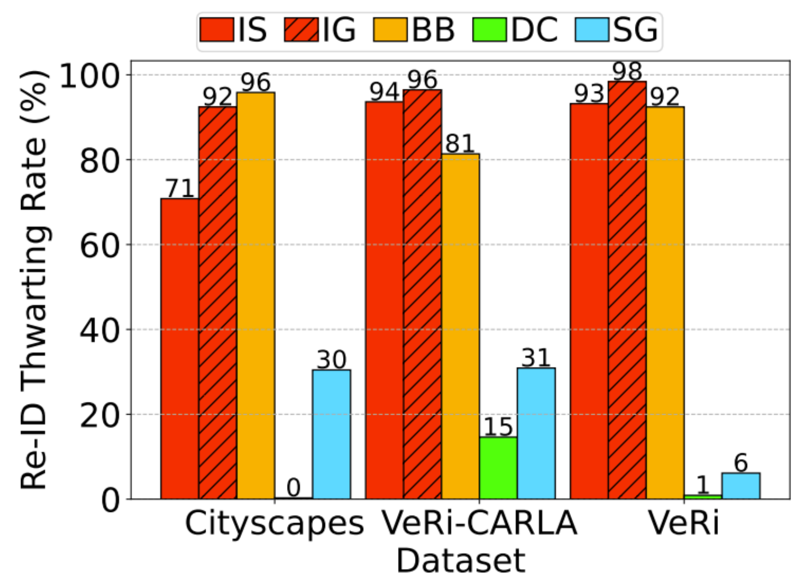

(b) Car Re-ID thwarting rates

IS: INSPIRE; IG: INSPIRE with Gaussian Filter (KS: 5, SD: 5); BB: BBox Blur; DC: Dashcam Cleaner; SG: SecGAN.

Model Inversion Attack

- Inverse model: Pix2pixHD, tries to restore original images from transformed images.
- Collected 9948 transformed-original image pairs for training.
- Trained and applied adversarial models to SecGAN and INSPIRE.
- INSPIRE can thwart model inversion attacks by design.

Original | SecGAN | SecGAN inverse | INSPIRE | INSPIRE inverse

Utility of Transformed Videos

- Dashcam Cleaner maintains best utility (however, no privacy aginst Re-ID attacks).
- INSPIRE performs better than Bbox Blur and SecGAN, and preserves higher utility on the Cityscapes datasets.

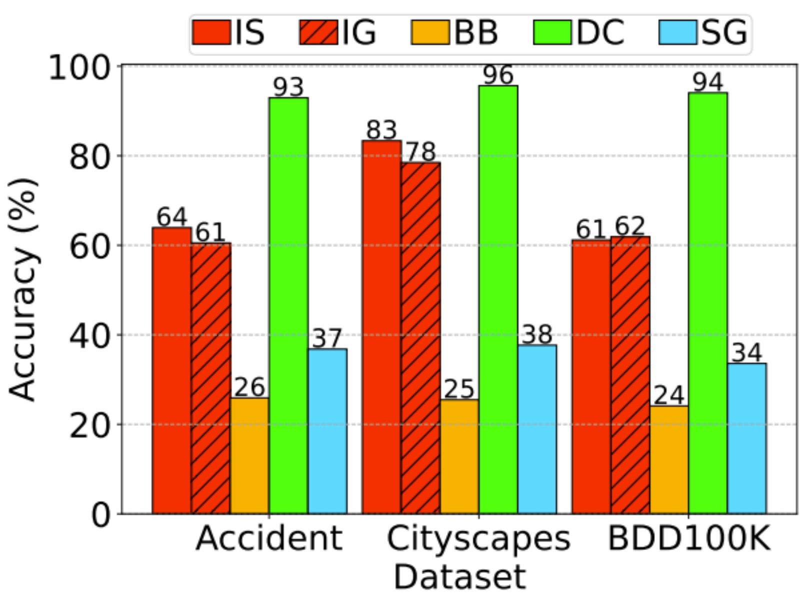

(a) Counting accuracy

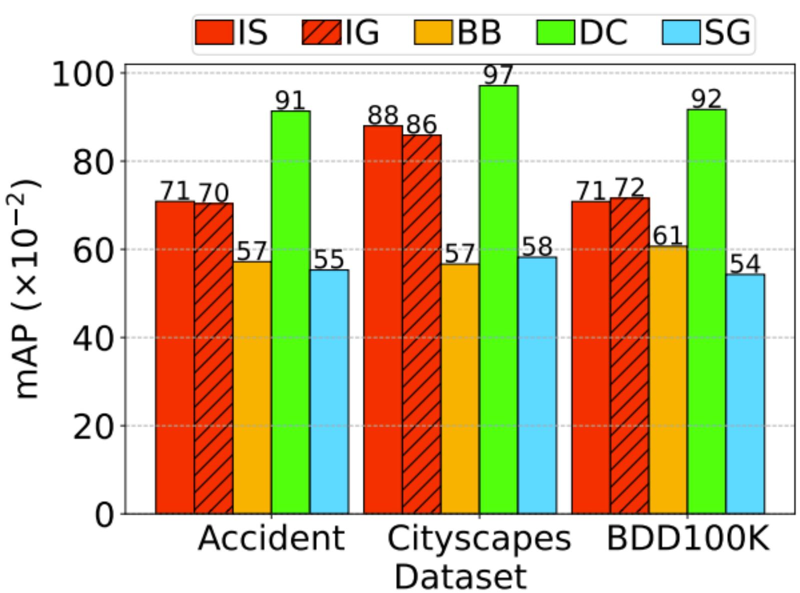

(b) Detection mAP

IS: INSPIRE; IG: INSPIRE with Gaussian Filter (KS: 5, SD: 5); BB: BBox Blur; DC: Dashcam Cleaner; SG: SecGAN.

Privacy-utility Trade-off

- Metrics:
  - Utility metric: Object detection mAP.
  - Privacy metric: Re-ID thwarting rate.
- Utility-privacy product
  - $Object\ detection\ mAP \times Re - ID\ thwarting\ rate$

INSPIRE achieves the best privacy-utility trade-off among compared systems.

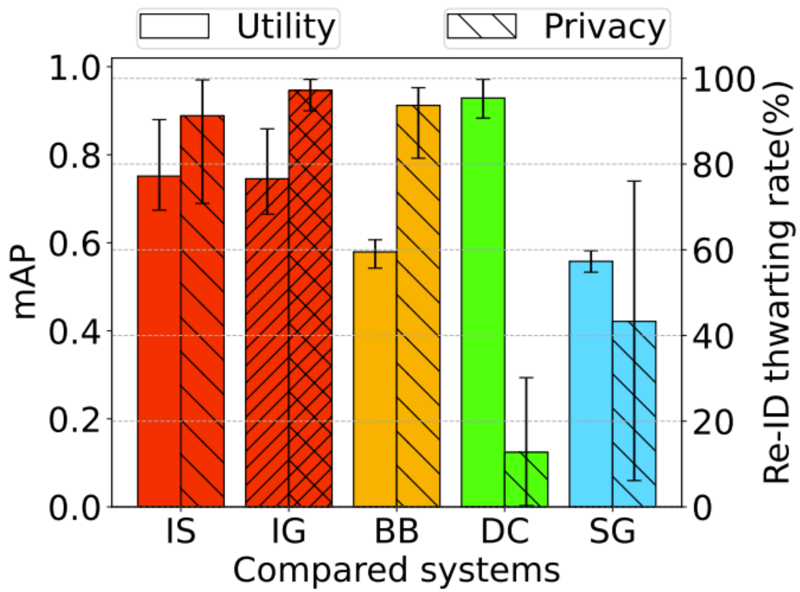

(a) Utility-privacy trade-off.

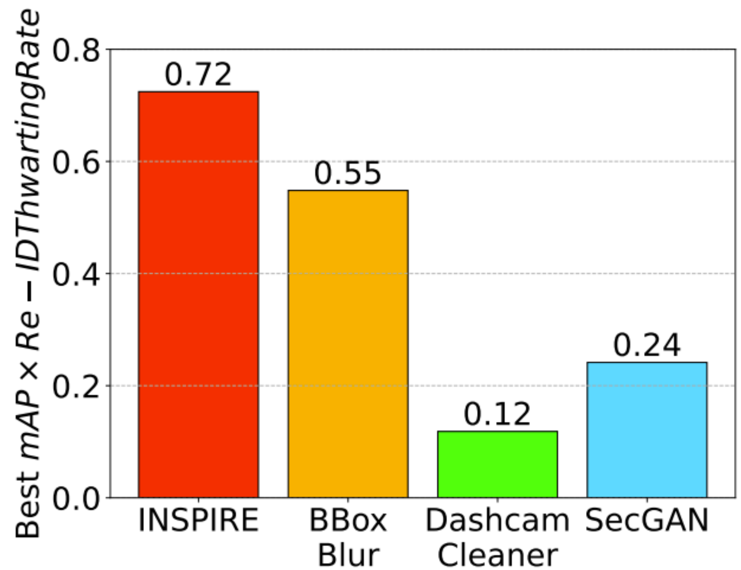

(b) Utility-privacy product.

Outlines

Background and Motivation

Threat Model

Framework Design and Implementation

Performance Evaluation

Discussions, Future Work and Conclusions

Other Perspectives, Conclusions

- INSPIRE achived
  - Instance-level privacy protection
    on
    Highly dynamic vehicular videos

    Replace instances with AI-synthesized ones

- What could be improved
  - Better object detection and segmentation
  - Better synthesized instances

    Use latest models
    (e.g. YOLOv8 & Diffusion)

  - Usability for computational constraint devices
    - Privacy protected by the image segmentation
    - Image synthesis is computational heavy

    Transplant to Mobile edge computing framework

  - Better visual effects
    - Currently only for machine analysis.

    Apply Object Tracking Algorithms (e.g. DeepSORT)

- **Conclusions:** Instance-level Privacy Protection on Vehicular Camera Videos

This research was supported in part by NSF grants 2007391 and 2045539. The information reported here does not reflect the position or the policy of the funding agency.

Thank you very much!

Q&A?