On the Semantic AI Security in CPS: *The Case of Autonomous Driving*

Qi Alfred Chen

Assistant Professor, UC Irvine

UCIrvine

AS²Guard
Autonomous & Smart Systems Guard Research Group

A bit about myself & my group

- Assistant Professor of Computer Science, UC Irvine (2018 -)
  - Ph.D., University of Michigan
- Group: **AS²Guard** (<u>A</u>utonomous & <u>S</u>mart <u>S</u>ystems **<u>Guard</u>**)
- Expertise: **AI/Systems/Network Security**, mainly in **mobile/CPS/IoT**

AS²Guard

<u>A</u>utonomous & <u>S</u>mart <u>S</u>ystems <u>**Guard**</u> Research Group

Impact: Demo & vulnerability report

CVE cve.mitre.org

NDSS'16

Euro S&P'17

US-CERT UNITED STATES COMPUTER EMERGENCY READINESS TEAM

IEEE S&P'16

Android

Usenix Sec'14

NDSS'16

Apple

Comcast

Microsoft

WAYMO

apollo

TOYOTA

GM

DAIMLER

BOSCH ...

DEPARTMENT OF TRANSPORTATION UNITED STATES OF AMERICA

Linux

CCS'15

CCS'17

Usenix Sec'20

NDSS'18

My research so far in mobile/CPS/IoT security

- **CPS AI Security**
  - **Autonomous Driving (AD)** [ACM CCS’19, Usenix Security’20 (a), ’20 (b), ’21, IEEE S&P’21, NDSS’22, CVPR’22, ICLR’20]
  - **Intelligent transportation** [NDSS’18, TRB’18,’19,’20, ITS’21]
- **Network Security**
  - **Connected Vehicle (CV)** [Usenix Security’21]
  - **Automotive IoT** [Usenix Security’20, NDSS’20]
  - **Network protocol** [ACM CCS’15,’18, IEEE S&P’16]
- **UI (User Interface) Security**
  - **Smartphone** [Usenix Security’14, MobiSys’19]
- **Access Control / Policy Enforcement**
  - **Smartphone** [NDSS’16]
  - **Smart home** [NDSS’17]
- **Side Channel**
  - **Smartphone** [Usenix Security’14]
  - **Network** [ACM CCS’15]

Most recent focus (2018-): CPS AI security

- **CPS AI Security**
  - **Autonomous Driving (AD)** [ACM CCS'19, Usenix Security'20 (a), '20 (b), '21, IEEE S&P'21, NDSS'22, CVPR'22, ICLR'20]
  - **Intelligent transportation** [NDSS'18, TRB'18,'19,'20, ITS'21]

- Relatively new area:
  - AI security: Since 2013 [Szegedy et al., "Intriguing properties of neural networks"]
  - AI penetration in real-world CPS (e.g., since ~2015 in automotive industry)

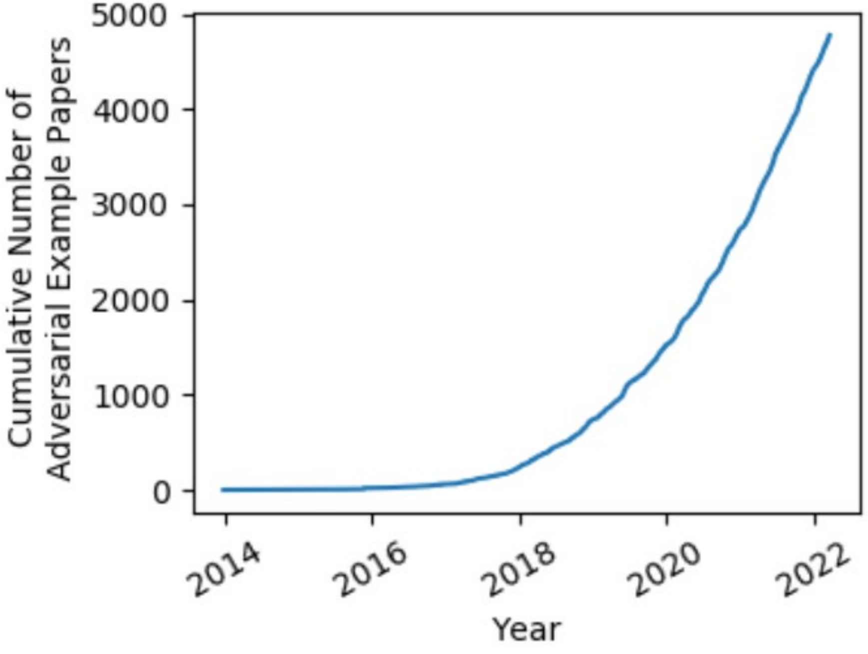

(*Image credit: Nicholas Carlini)

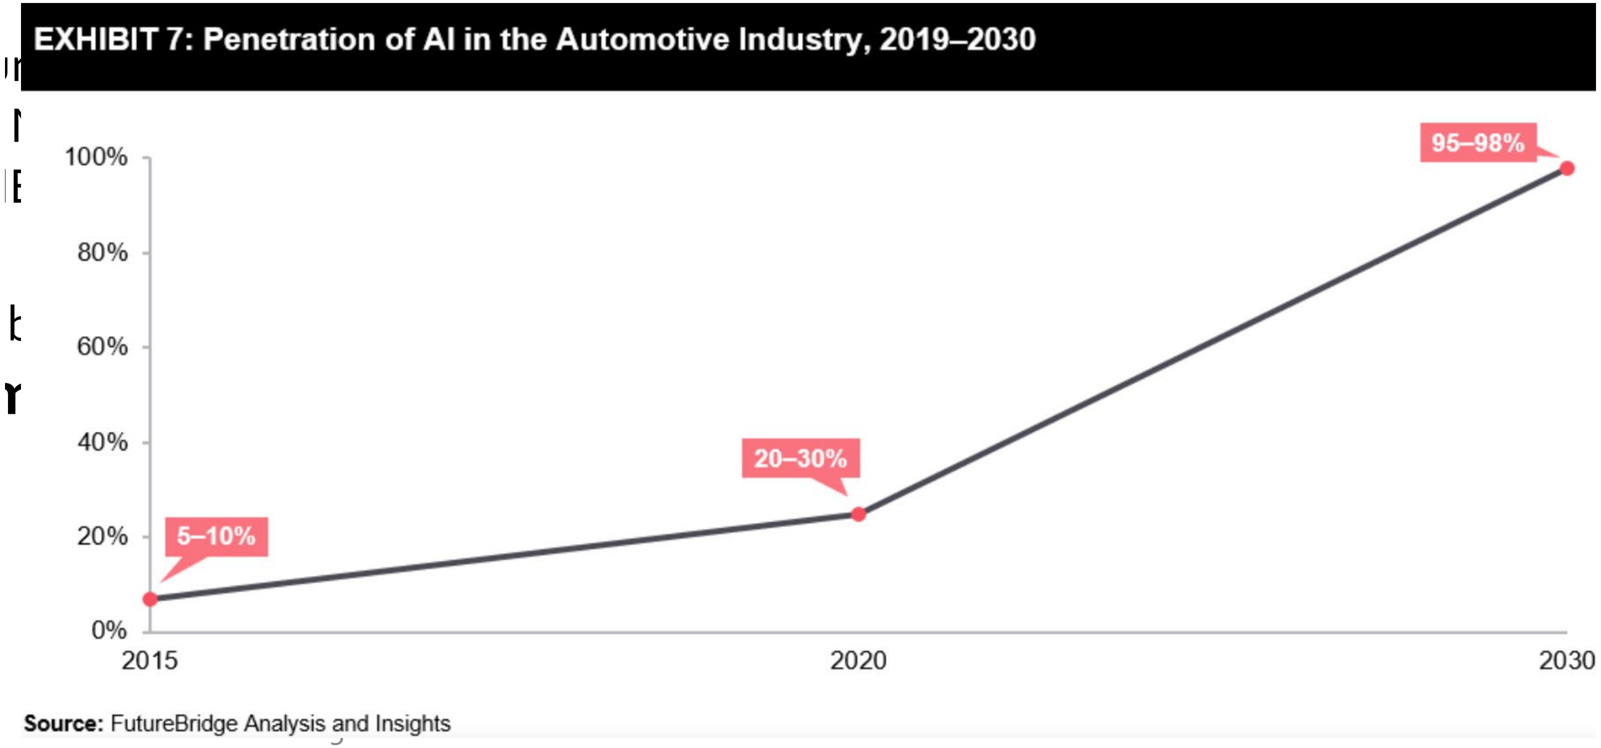

More recently, various kinds of AI-enabled autonomous systems coming into real life

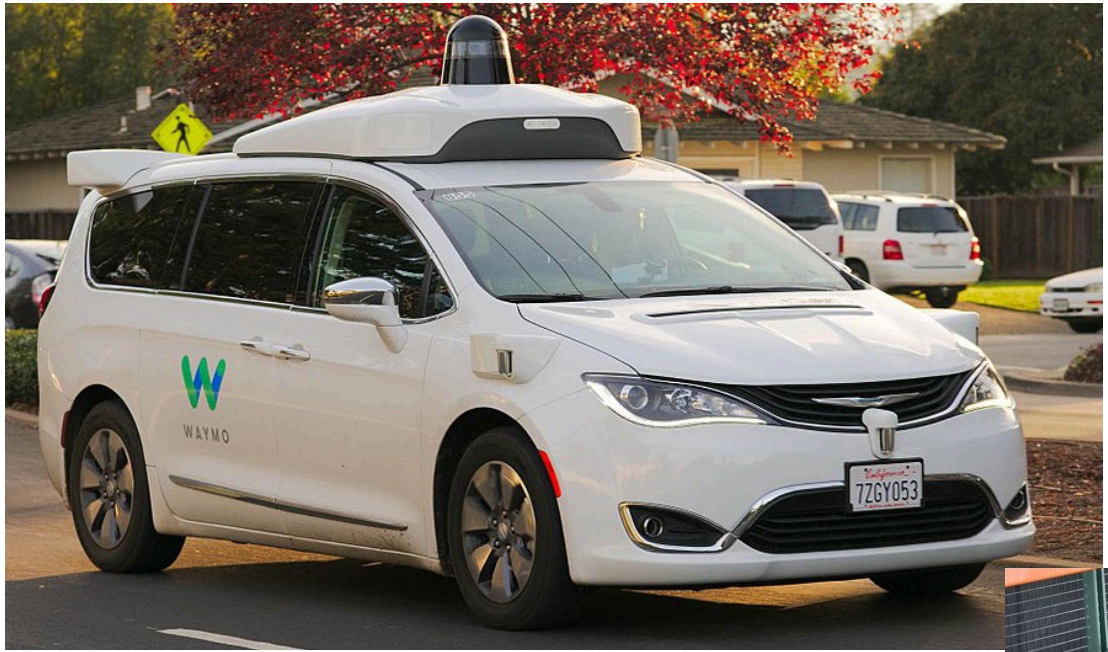

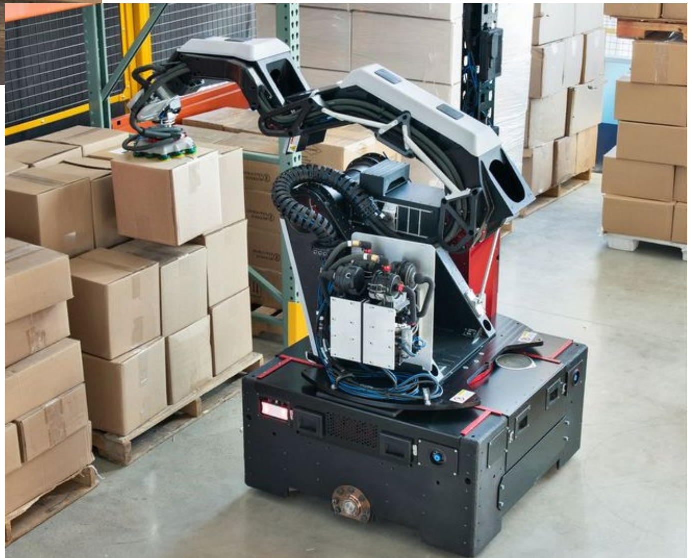

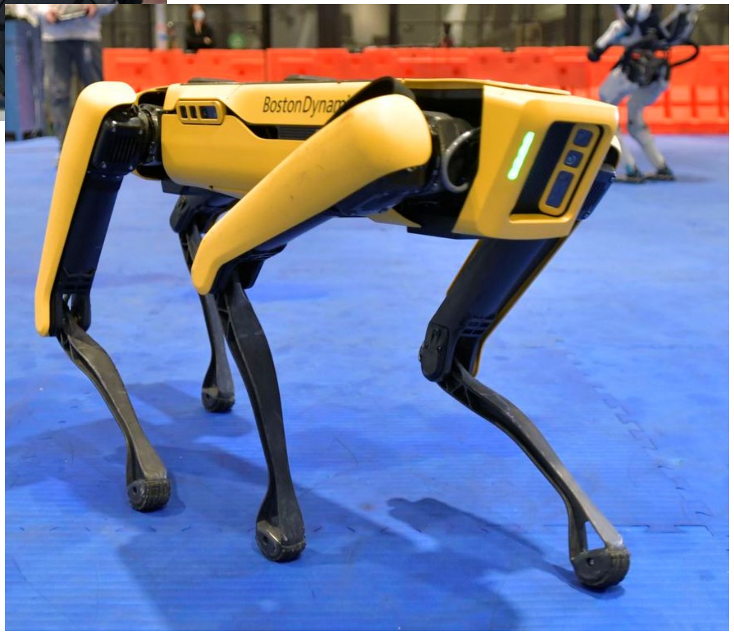

Highly desired to study their security

- In charge of highly safety-critical decision-making in the physical world
  → Security problems can have ***unprecedentedly high impacts on public safety & society*** (e.g., fatal crashes)

An Uber self-driving car hit & killed a woman crossing street in Arizona since it cannot classify her as a pedestrian. [1] [2]

Fatal crash of a Tesla model X w/ Autopilot on in 2018 at California [3]. From the 2016 crash that killed a Florida driver, >20 Autopilot-related crashes have occurred [4].

[1] https://www.nbcnews.com/tech/tech-news/self-driving-uber-car-hit-killed-woman-did-not-recognize-n1079281
[2] https://www.theverge.com/2019/11/6/20951385/uber-self-driving-crash-death-reason-ntsb-dcouments
[3] https://www.bbc.com/news/world-us-canada-43604440
[4] https://www.nytimes.com/2021/03/23/business/teslas-autopilot-safety-investigations.html

Highly desired to study their security

- In charge of highly safety-critical decision-making in the physical world
  → Security problems can have ***unprecedentedly high impacts on public safety & society*** (e.g., fatal crashes)

- Domain-specific system components that may come with **new security properties**
- To meaningfully affect the AI-enabled autonomous decision-making (e.g., driving), face new research challenges **as a "*semantic* AI security" problem**
  - Proposed by us recently[1], generalized from "semantic adv deep learning" by Seshia et al. in 2018[2]

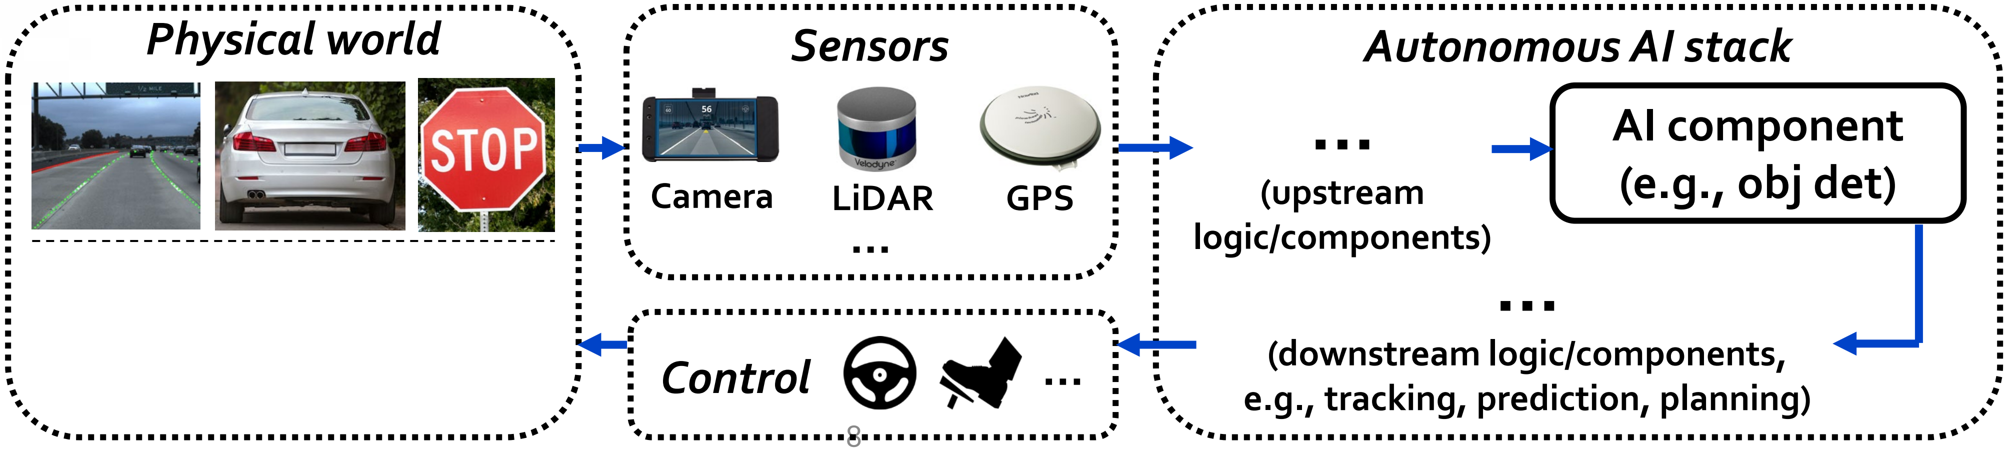

[1] Shen et al. @ arXiv 2203.05314, 2022 [2] Seshia et al. @ IEEE Design & Test'20

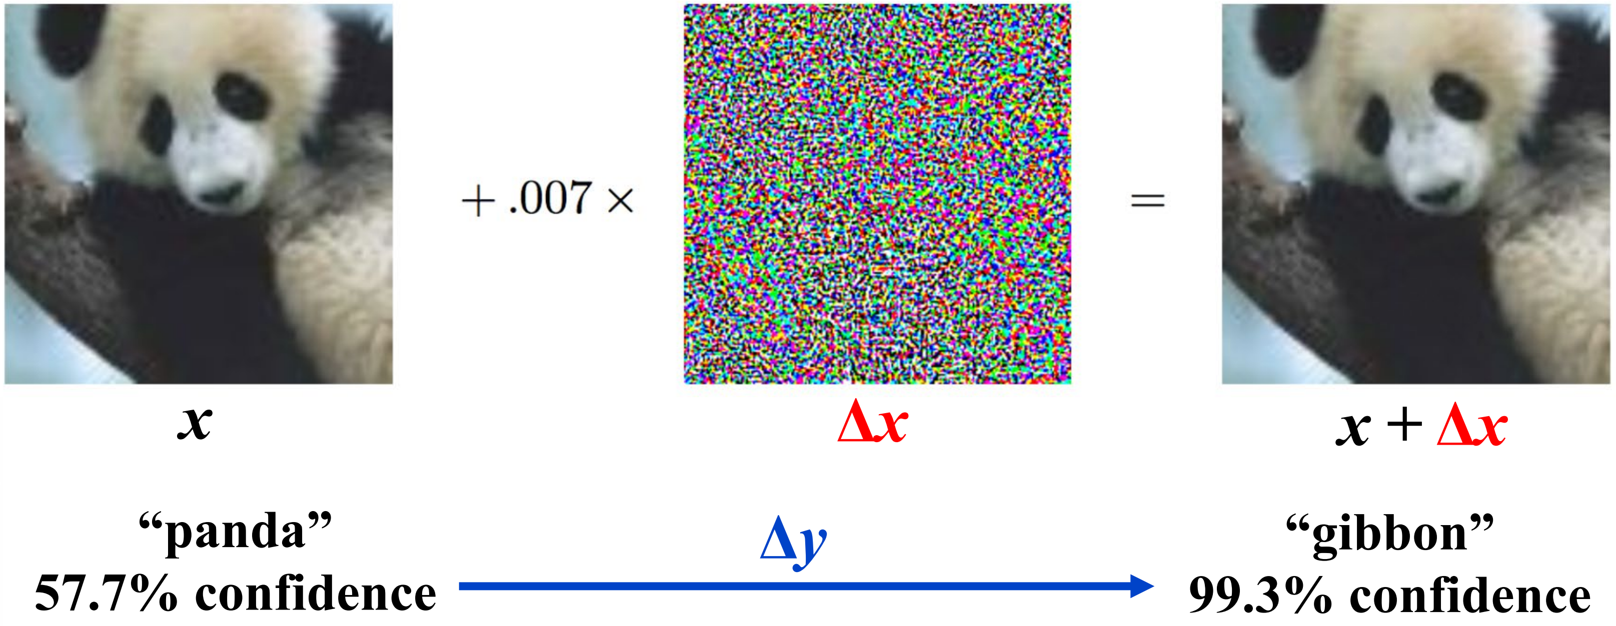

AI security papers

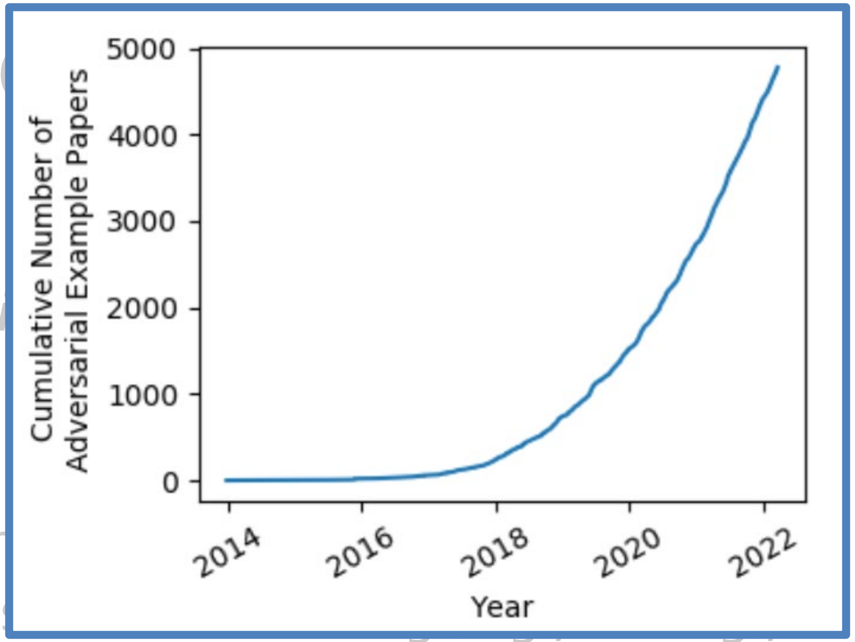

face new research challenges **as a "*semantic* AI security" problem**

- Proposed by us recently[1], generalized from "semantic adv deep learning" by Seshia et al. in 2018[2]

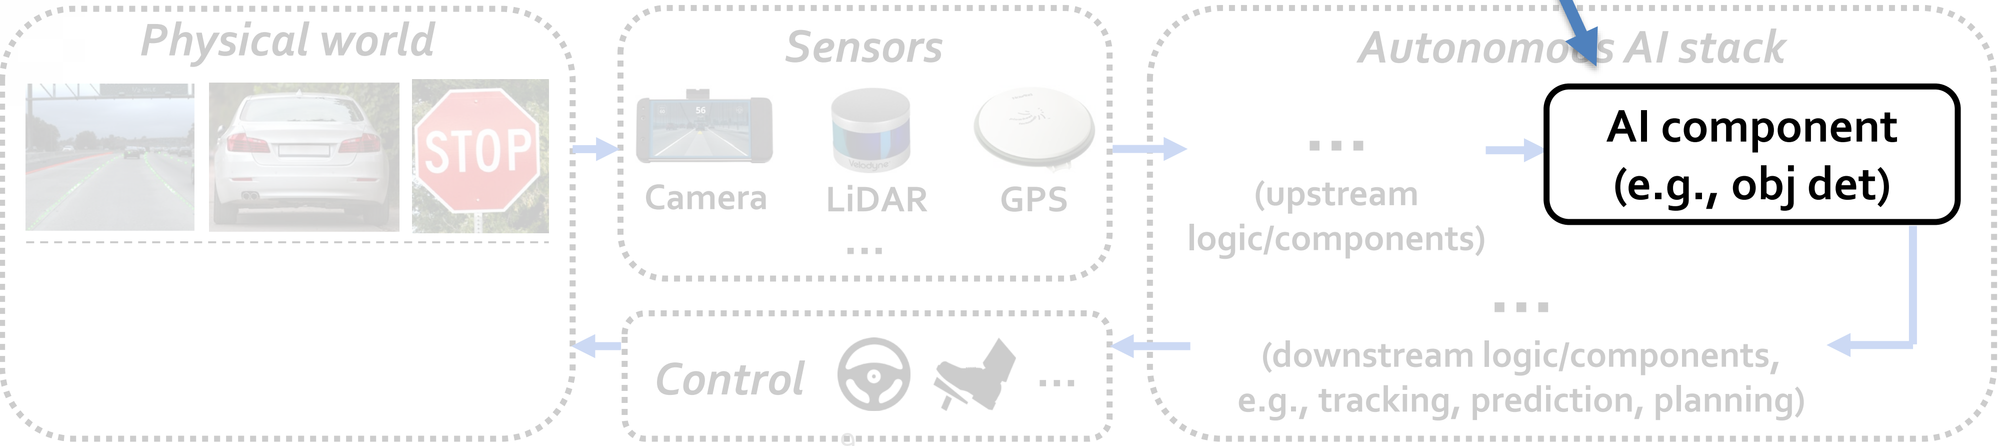

[1] Shen et al. @ arXiv 2203.05314, 2022

[2] Seshia et al. @ IEEE Design & Test'20

Highly desired to study their security

IMPORTANT

- In charge of highly safety-critical decision-making in the physical world
→ Security problems can have ***unprecedentedly high impacts on public safety & society*** (e.g., fatal crashes)

System-level attack input spaces (e.g., add stickers, laser shooting)
→ those at AI component level (e.g., image pixel changes)

- *Fundamentally challenging due to **inverse feature-mapping problem**[3]*

...at may come with **new security properties**
...utonomous decision-making (e.g., driving),
face new researc... ...as a "semantic AI security" problem

- Proposed by us rece... ...alized f...
- Need to further addre... ...*al seman*...

From AI component-level attack effects (e.g., misdetected objects)
→ those at CPS system level (e.g., vehicle collisions)

- *Challenging due to the **high end-to-end system-level complexity in CPS** & **dynamics from closed-loop control**[2,4,5]*

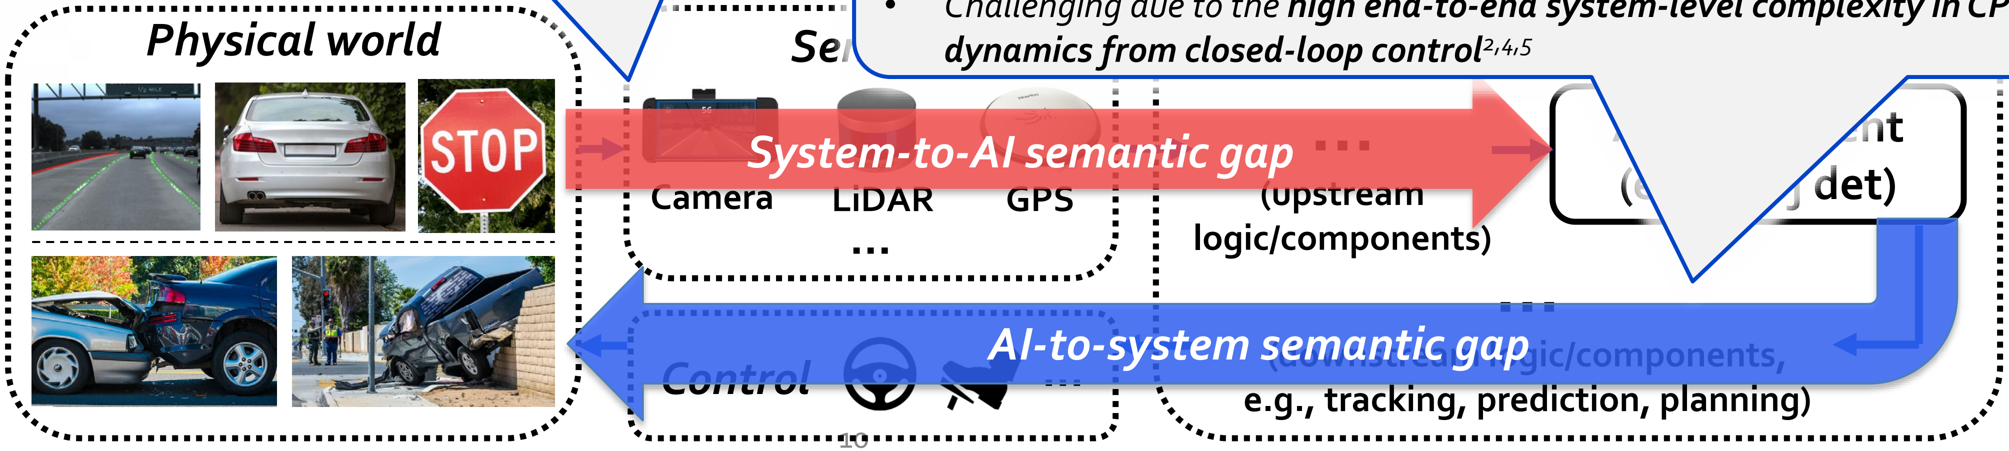

[1] Shen et al. @ arXiv 2203.05314, 2022 [2] Seshia et al. @ IEEE Design & Test'20 [3] Pierazzi et al. @ IEEE S&P'20 [4] Jia et al. @ ICLR'20 [5] Sato et al. @ Usenix Security'21

My recent focus (2018-): Automotive & transportation domain

Autonomous Driving (AD)

V2X-based Intelligent Transp.

TESLA
WAYMO
apollo
ZOOX
pony.ai
tusimple
GM
TOYOTA
Qualcomm
Ford
ARGO AI
Aurora
U.S. Department of Transportation

My recent focus (2018-): Automotive & transportation domain

Autonomous Driving (AD)

- **Fastest growing** AI-enabled autonomous system in industry today
- **Highly safety-critical**
  - ***Heavy***, ***fast-moving***, *& operate in* ***public*** *spaces*
- **Highly complex (to get right)**
  - Need to handle broad range of ***weather***, ***lighting***, ***road*** & ***traffic conditions***, while being ***safe*** & complying to ***traffic rule***

TESLA
WAYMO
apollo
ZOOX
pony.ai
tusimple
GM
TOYOTA
Qualcomm
Ford
ARGO AI
Aurora
U.S. Department of Transportation

Background: Autonomous Driving (AD) technology

- Equip vehicles with various types of sensors to enable self driving

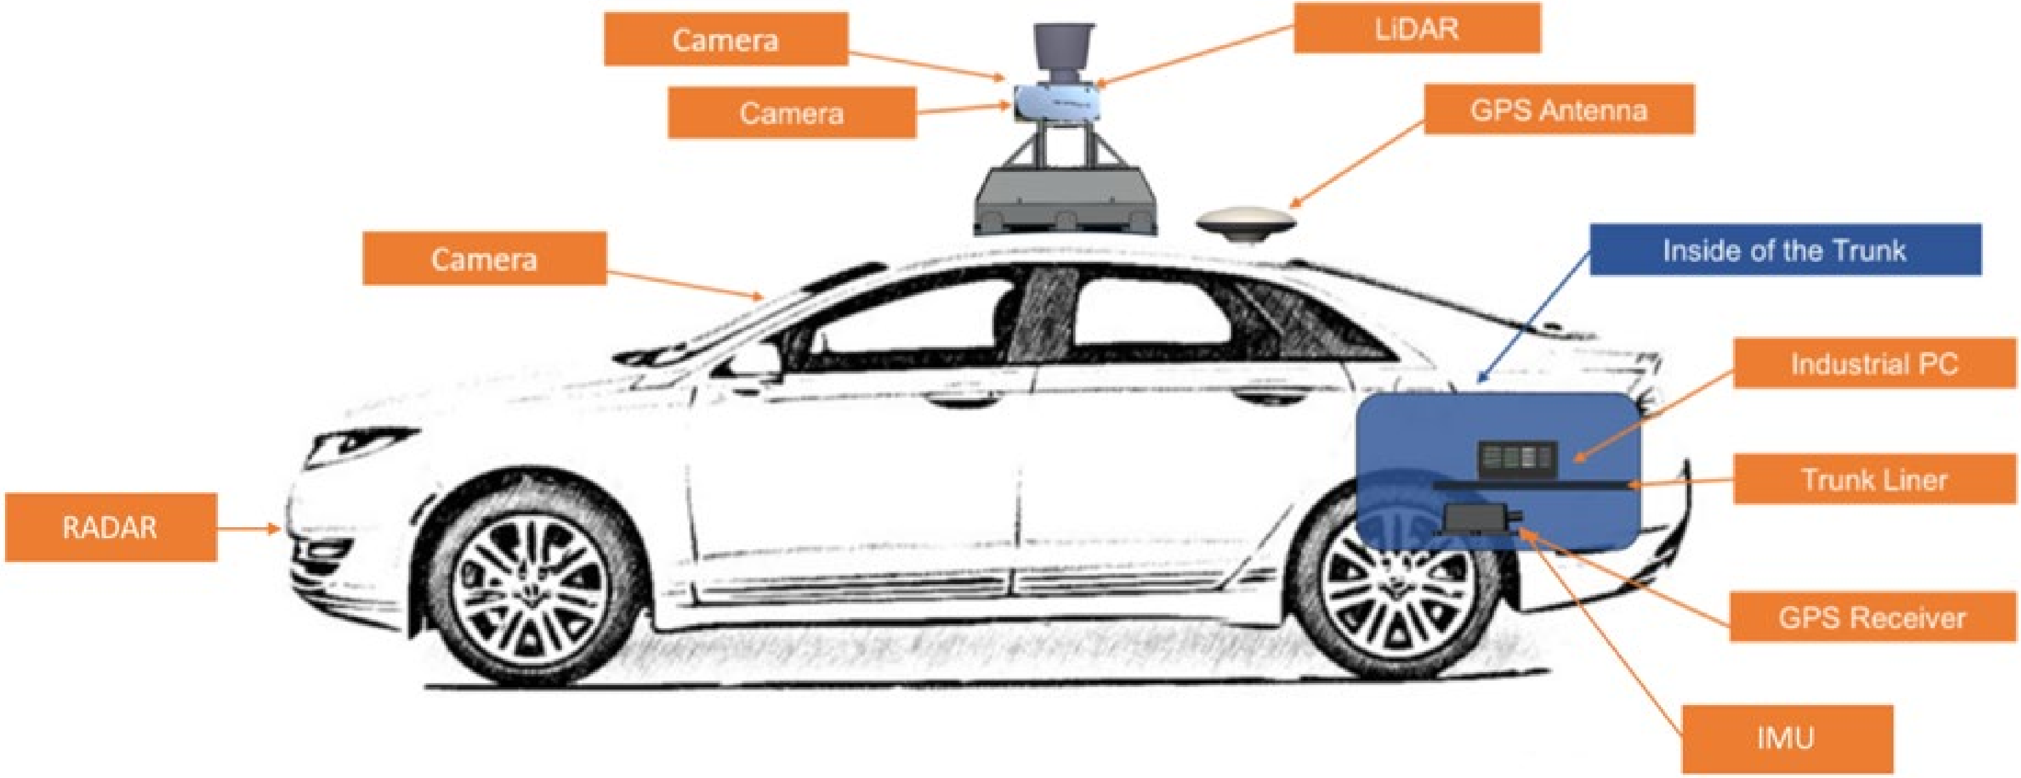

(*Image source: https://github.com/ApolloAuto/apollo)

Background: System architecture of industry-grade AD

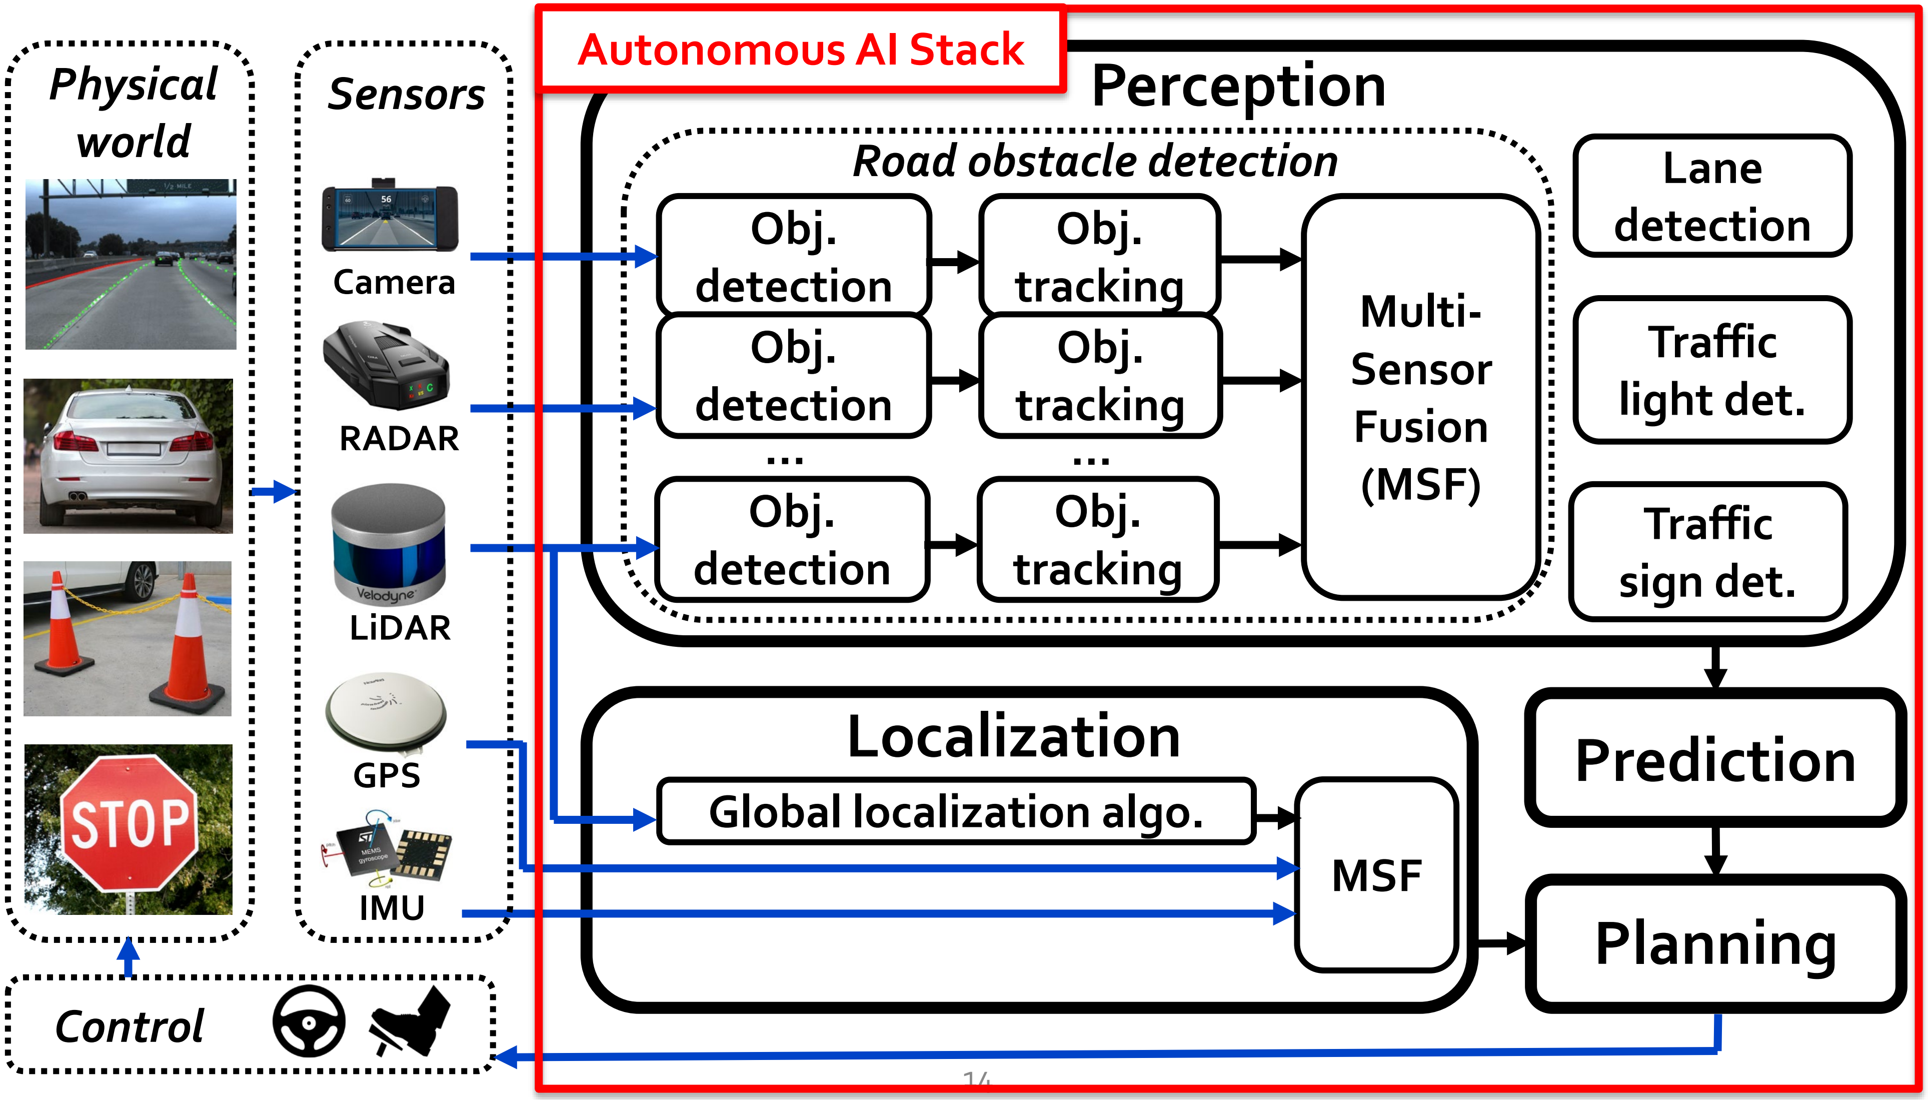

General & fundamental attack surface #1: Sensor attack

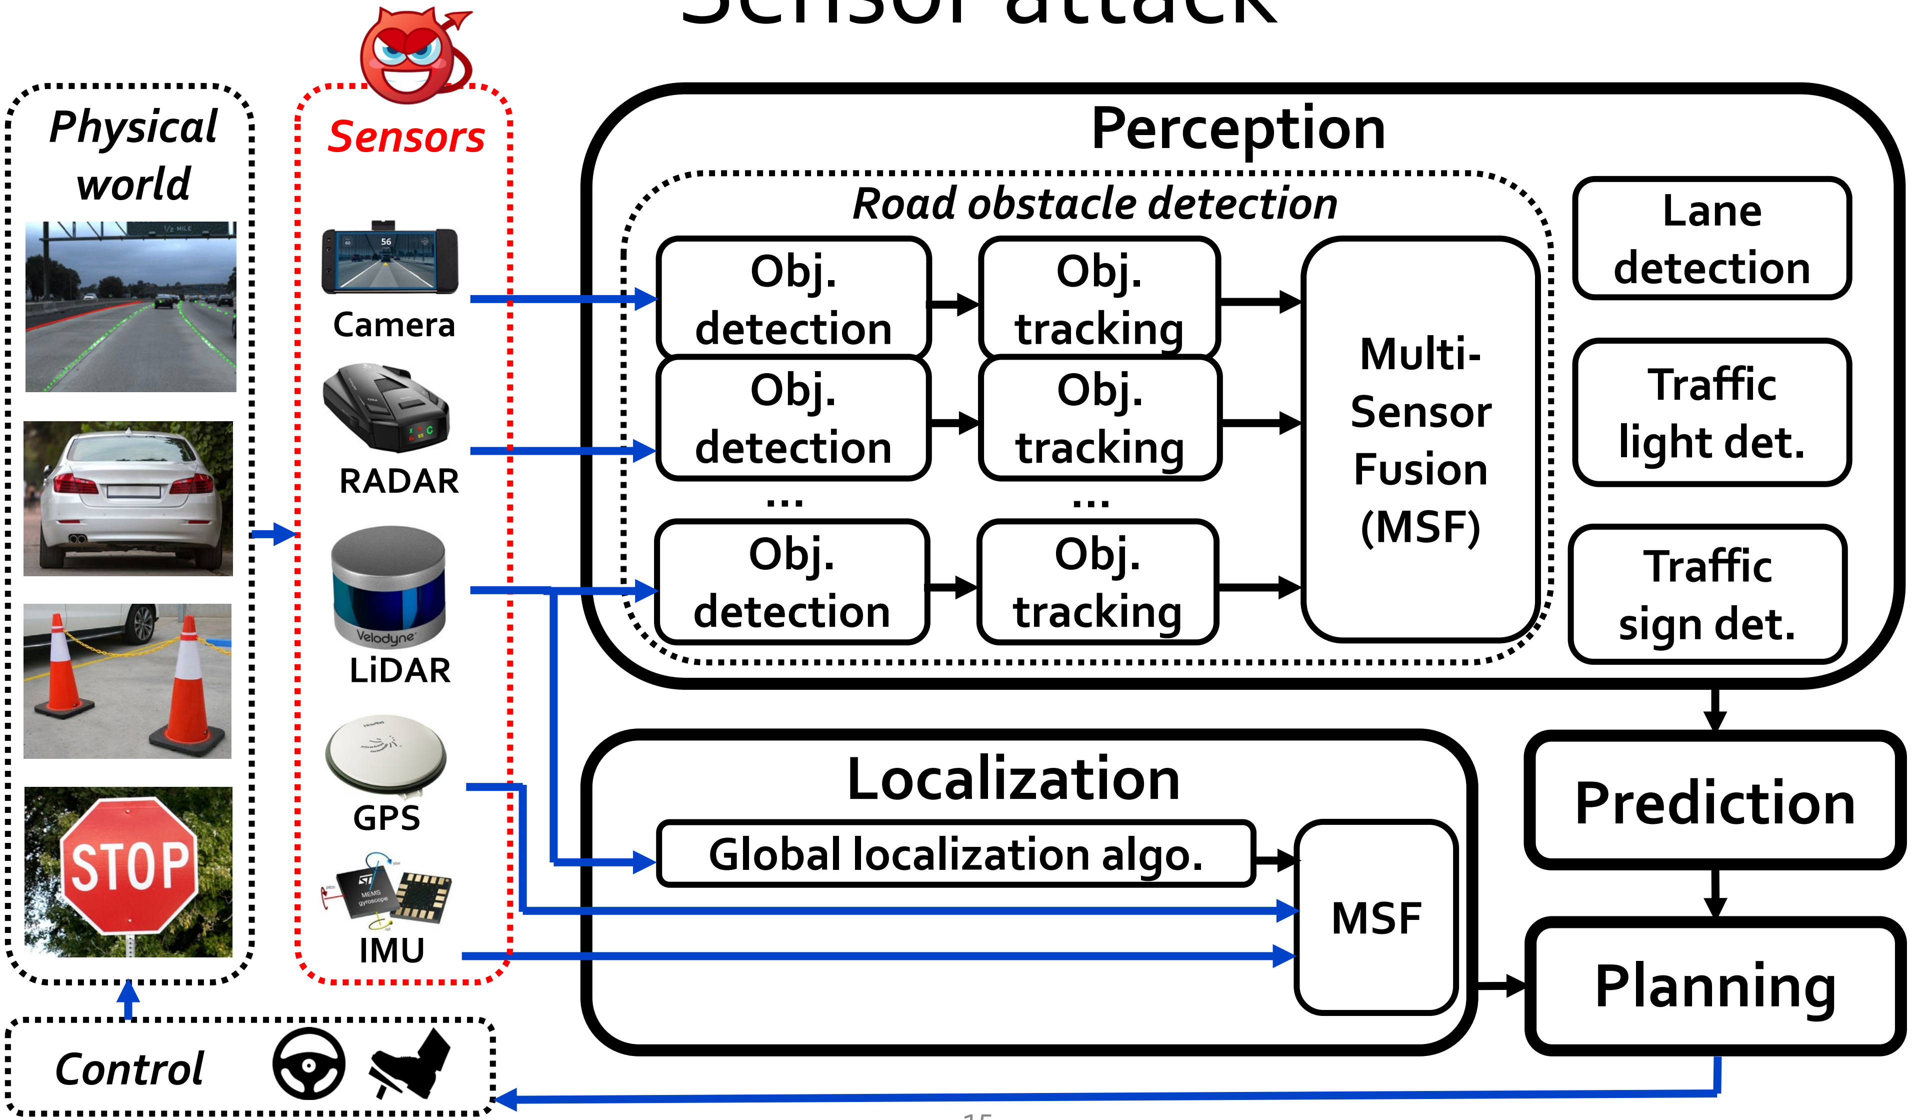

General & fundamental attack surface #1: Sensor attack

Spoofing/jamming attacks have been discovered on ***all popular sensor types used in AD systems***

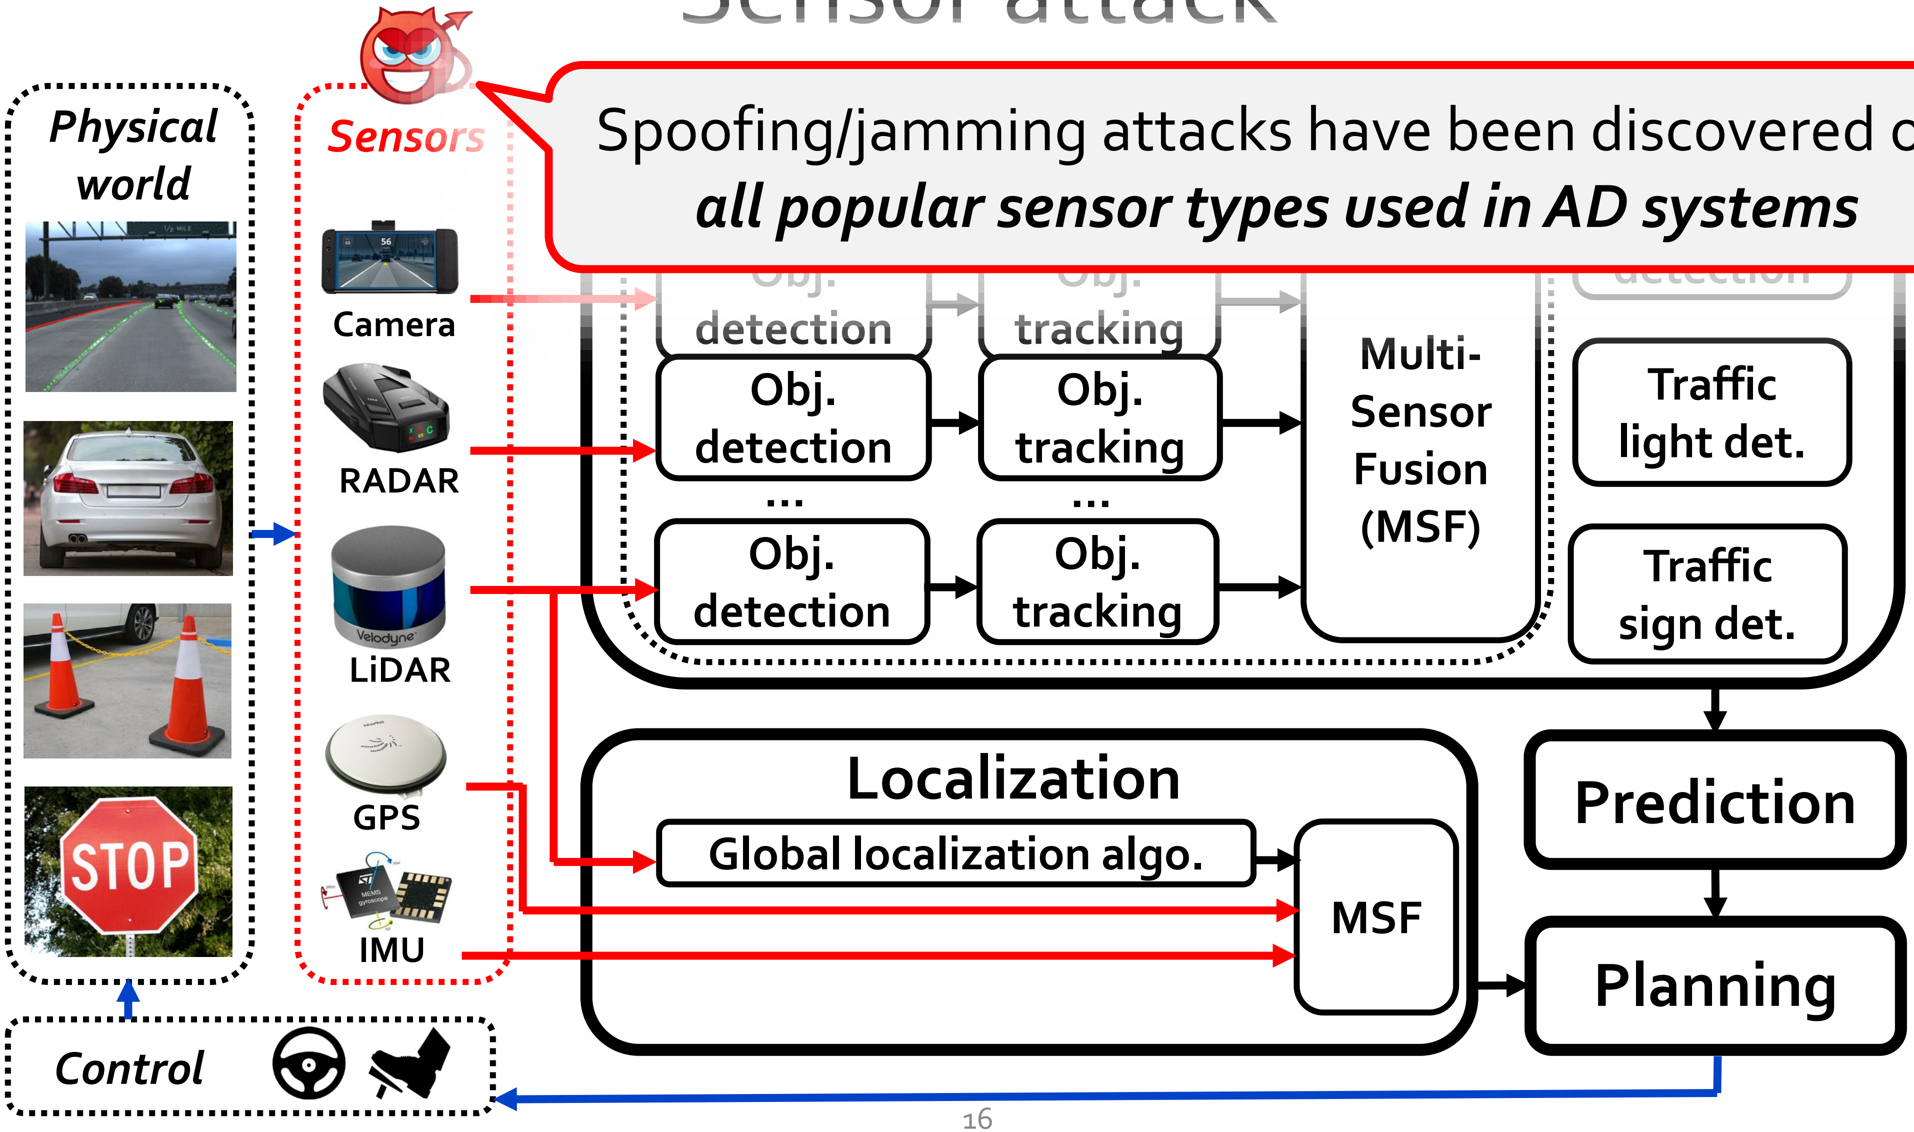

General & fundamental attack surface #2: Physical-world attack

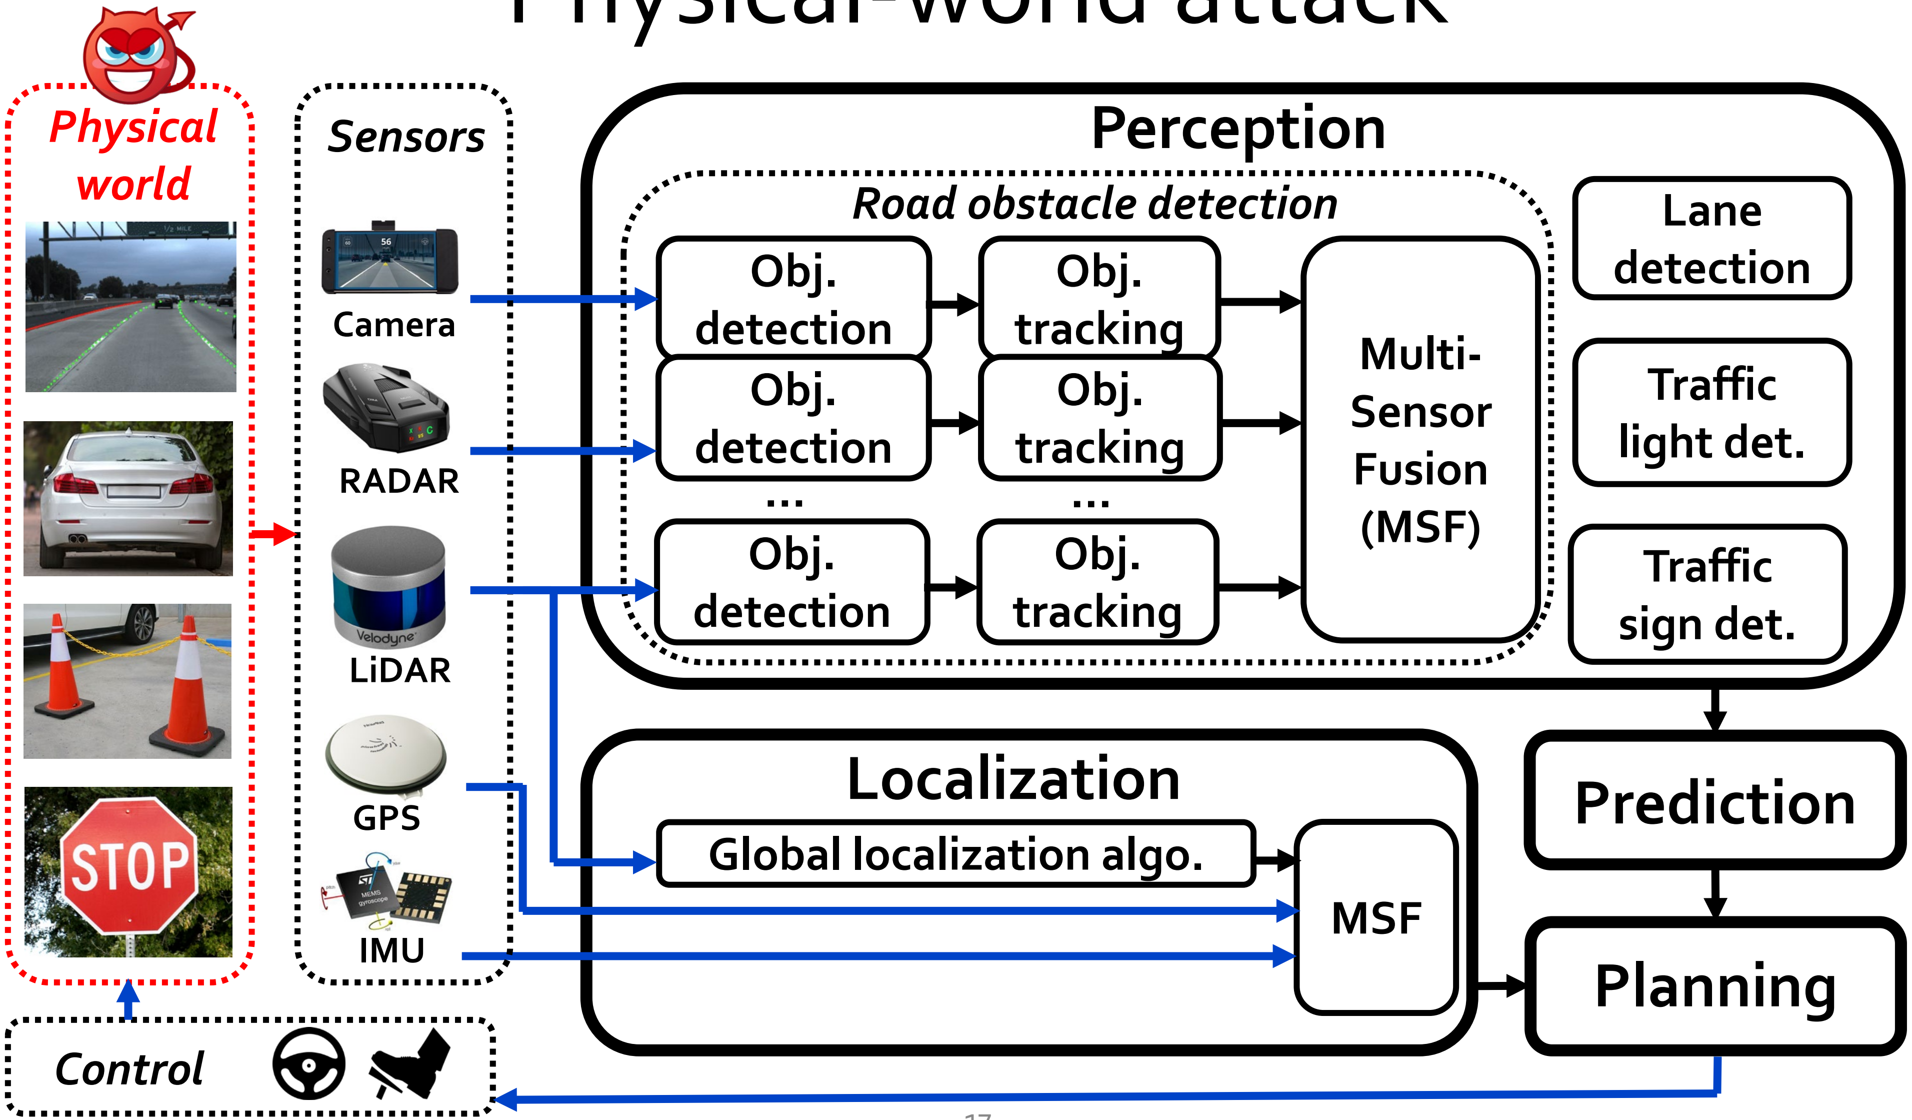

General & fundamental attack surface #2: Physical-world attack

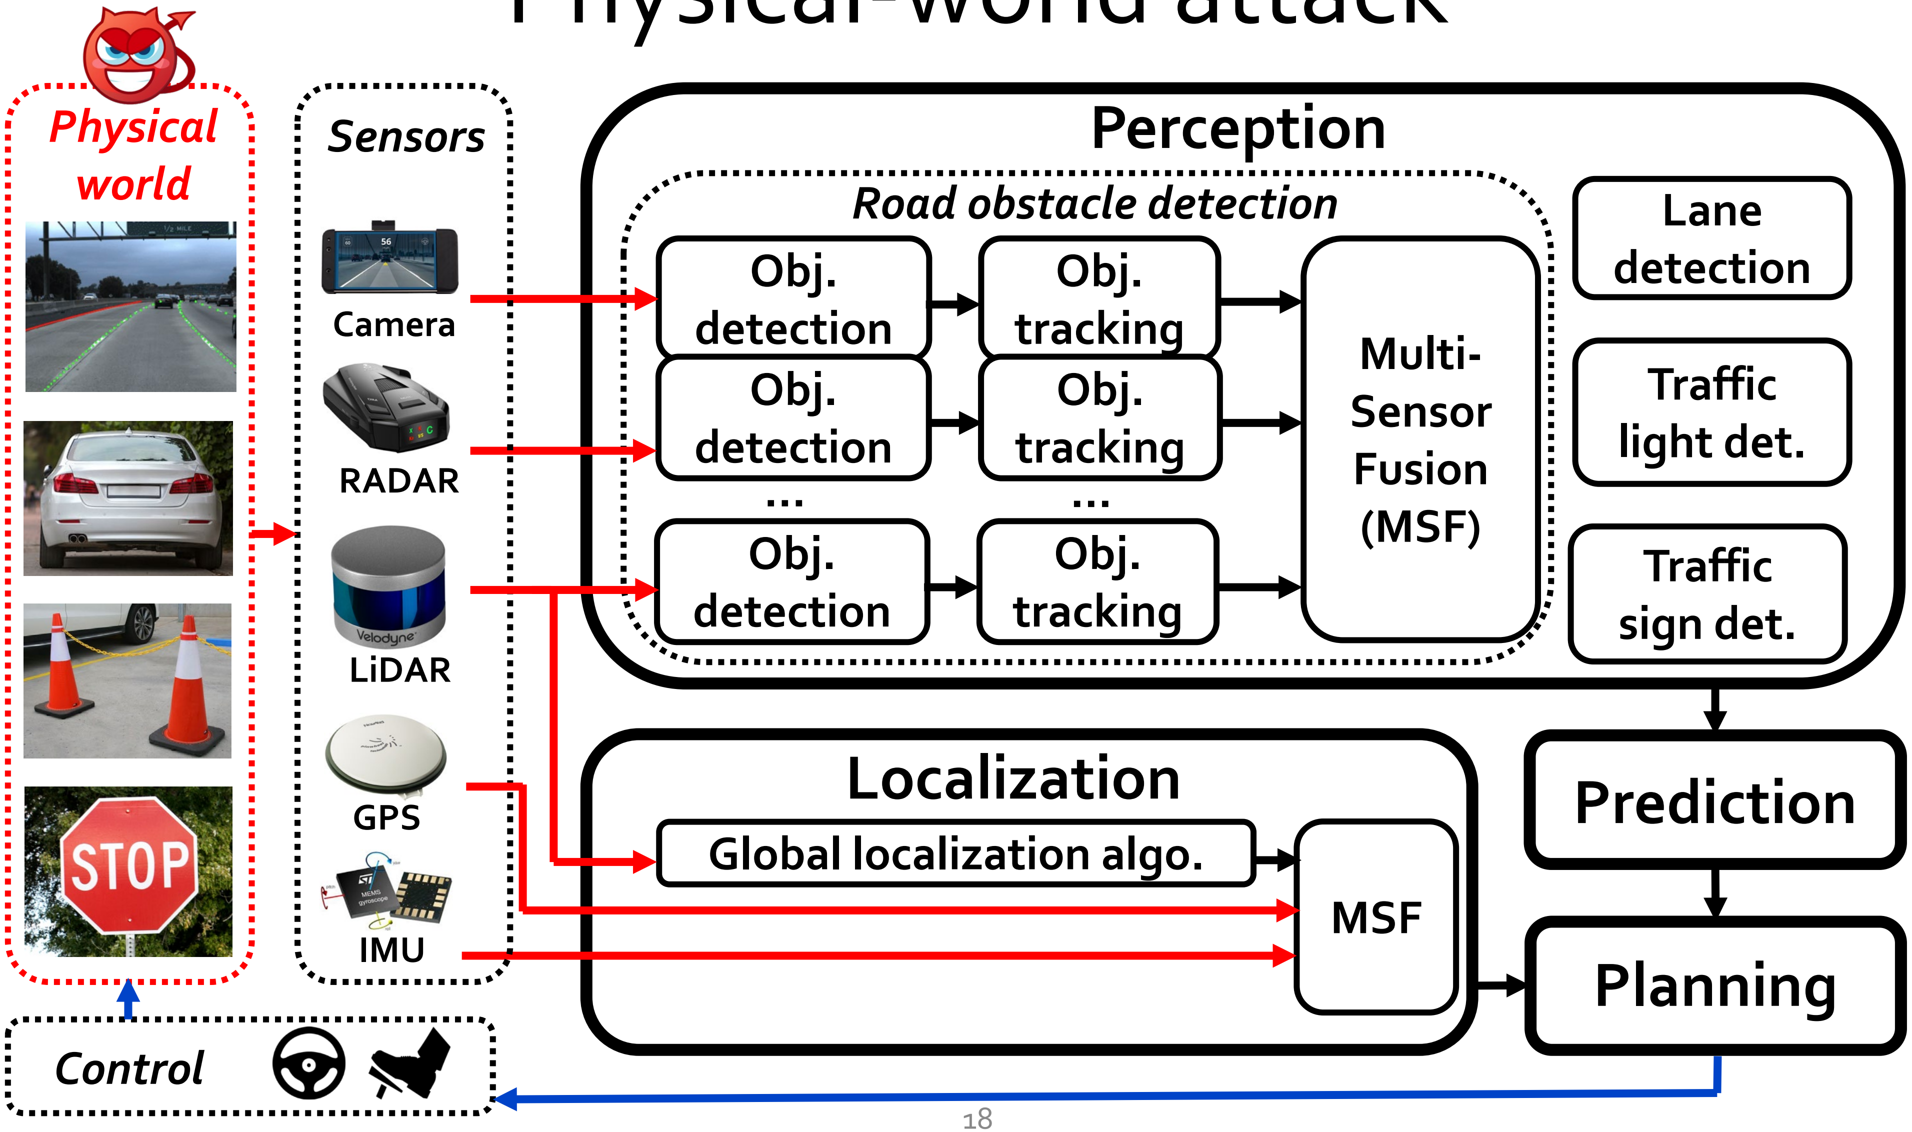

Both are considered in my research

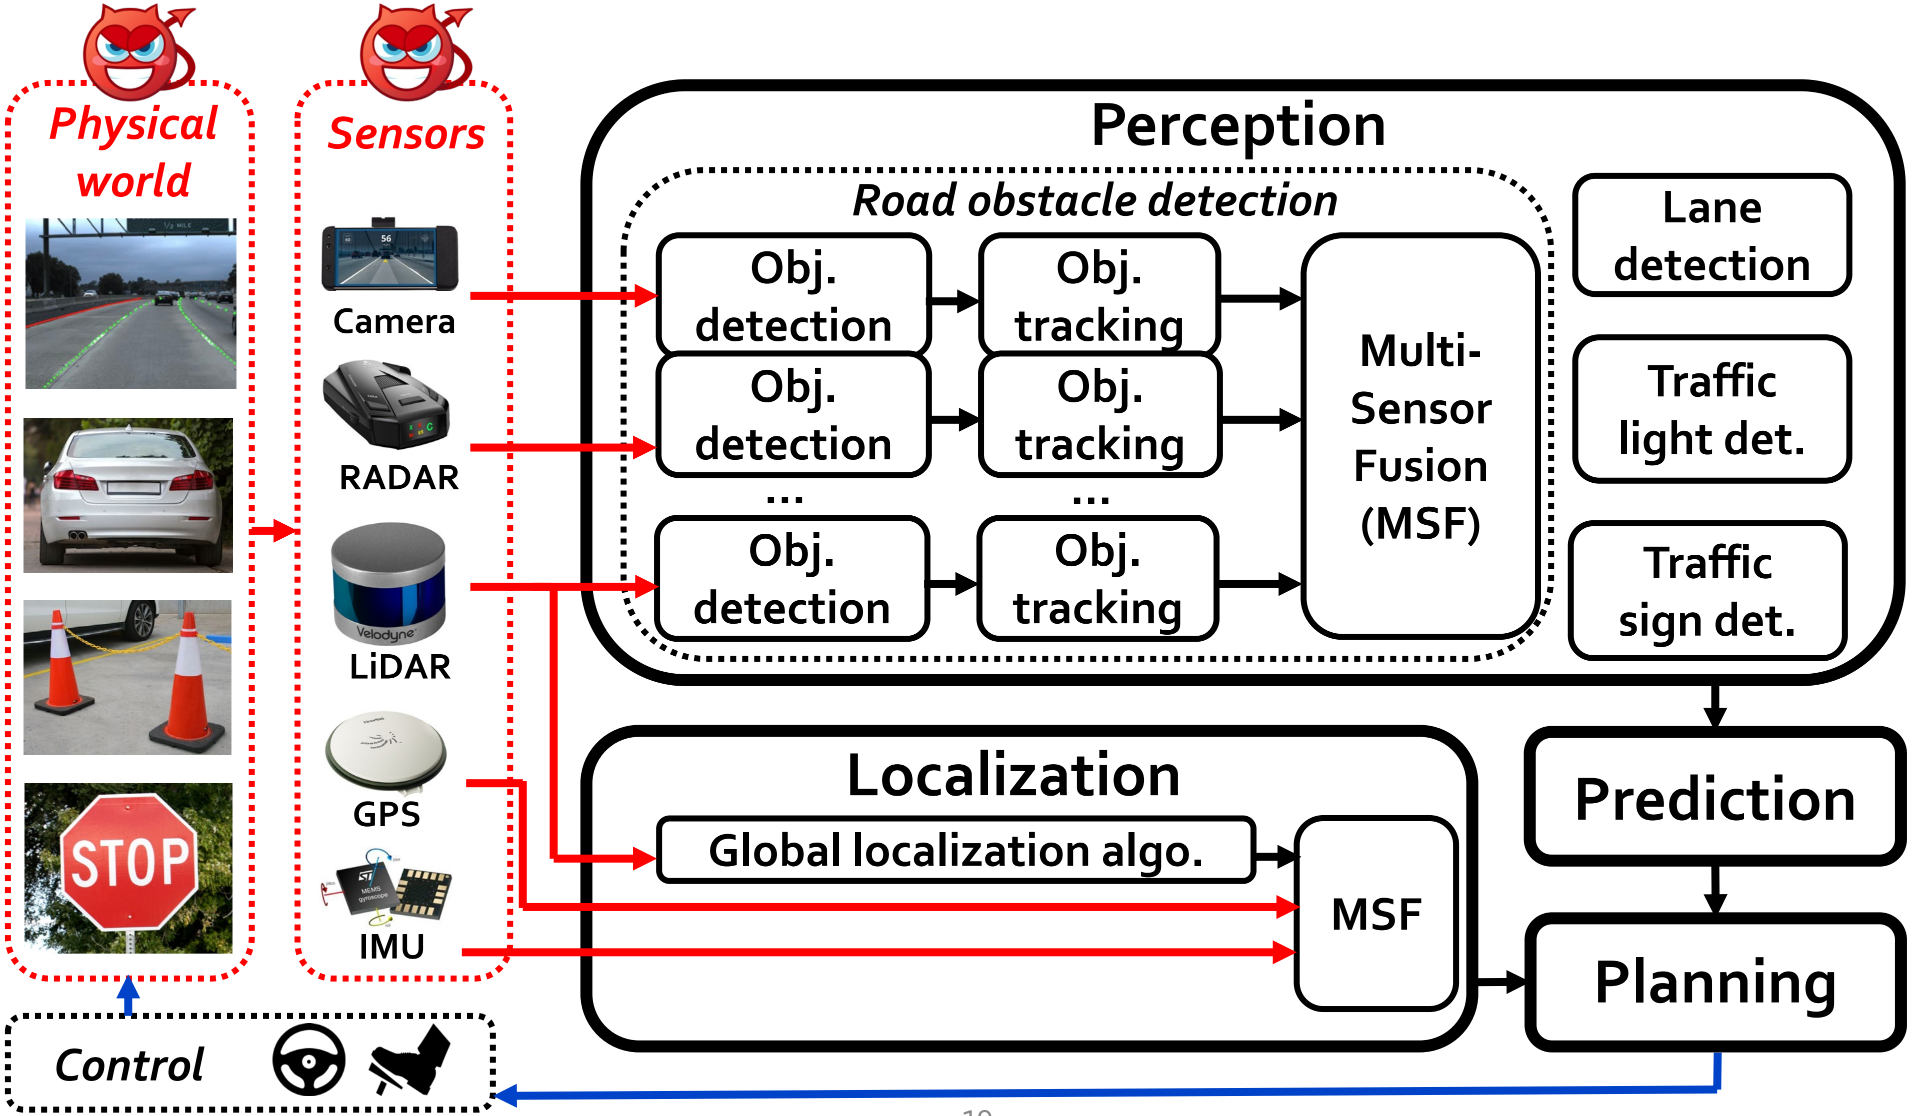

Black Hat'15, DEF CON'16

- Target: Sensors in production AD systems (e.g., Tesla)
- Attack vector: Sensor spoofing/jamming
- Impact: Make road obstacle disappear, or spoof fake ones

(a) Normal. (b) Spoofed. (c) Jammed.

(a) Drive gear. (b) Autopilot. (c) Jammed.

[Yan et al. @ DEF CON 2016]

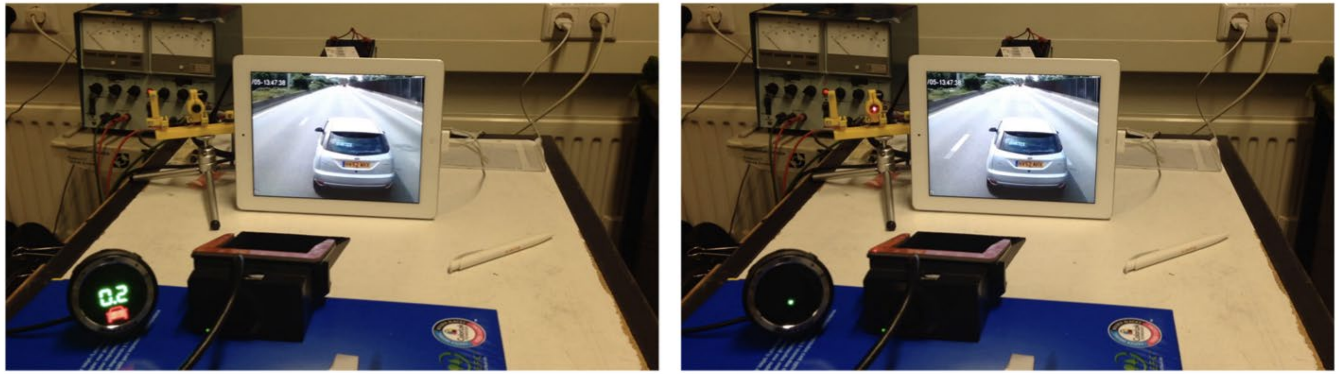

(a) Laser off, normal behavior of MobilEye C2-270

(b) Laser on, MobilEye C2-270 does not detect vehicle ahead

[Petit et al. @ Black Hat 2015]

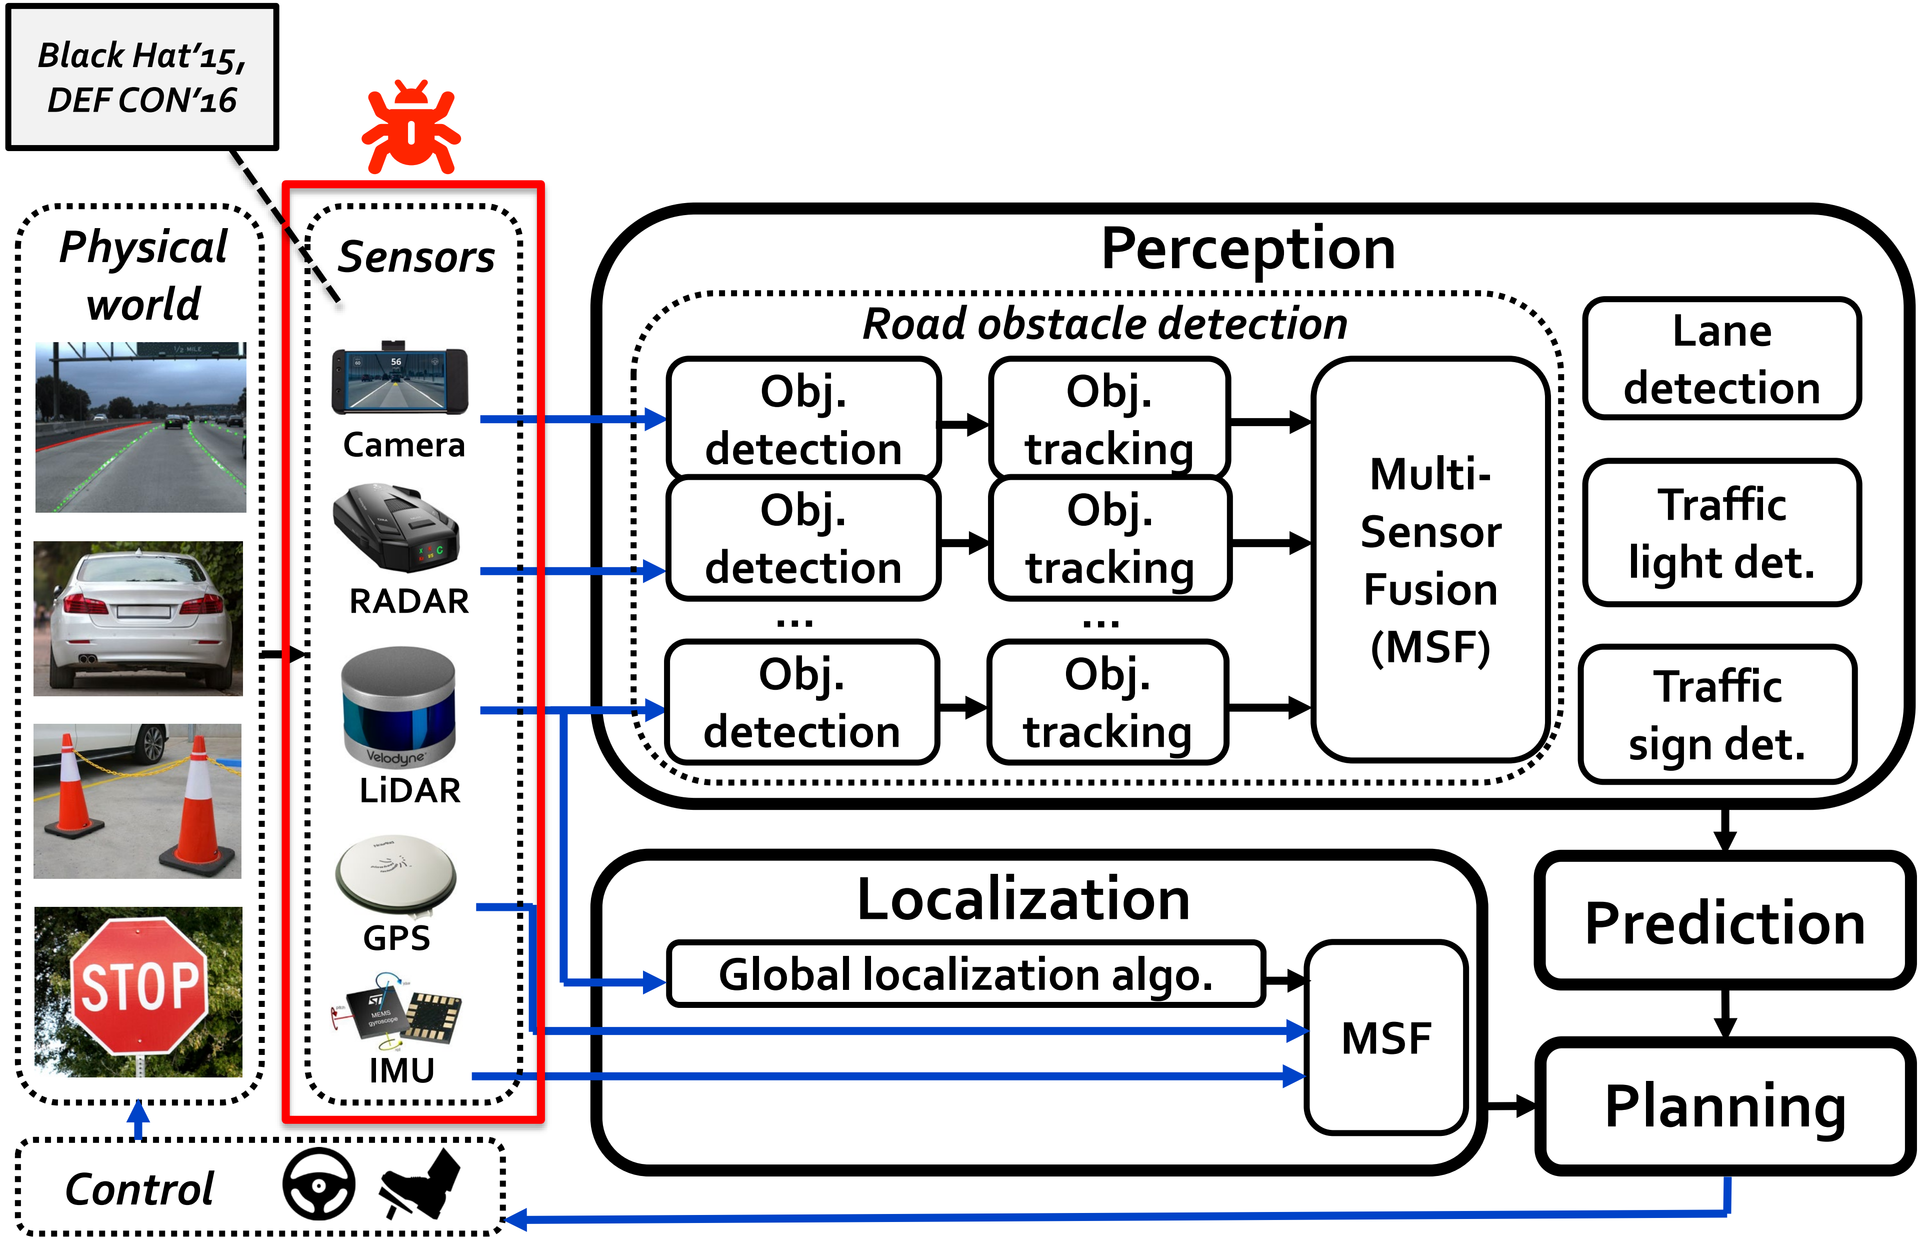
Black Hat'15, DEF CON'16
Physical world
STOP
Sensors
Camera
RADAR
LiDAR
Velodyne
GPS
IMU
Perception
Road obstacle detection
Obj. detection
Obj. tracking
Obj. detection
Obj. tracking
...
...
Obj. detection
Obj. tracking
Multi-Sensor Fusion (MSF)
Lane detection
Traffic light det.
Traffic sign det.
Localization
Global localization algo.
MSF
Prediction
Planning
Control

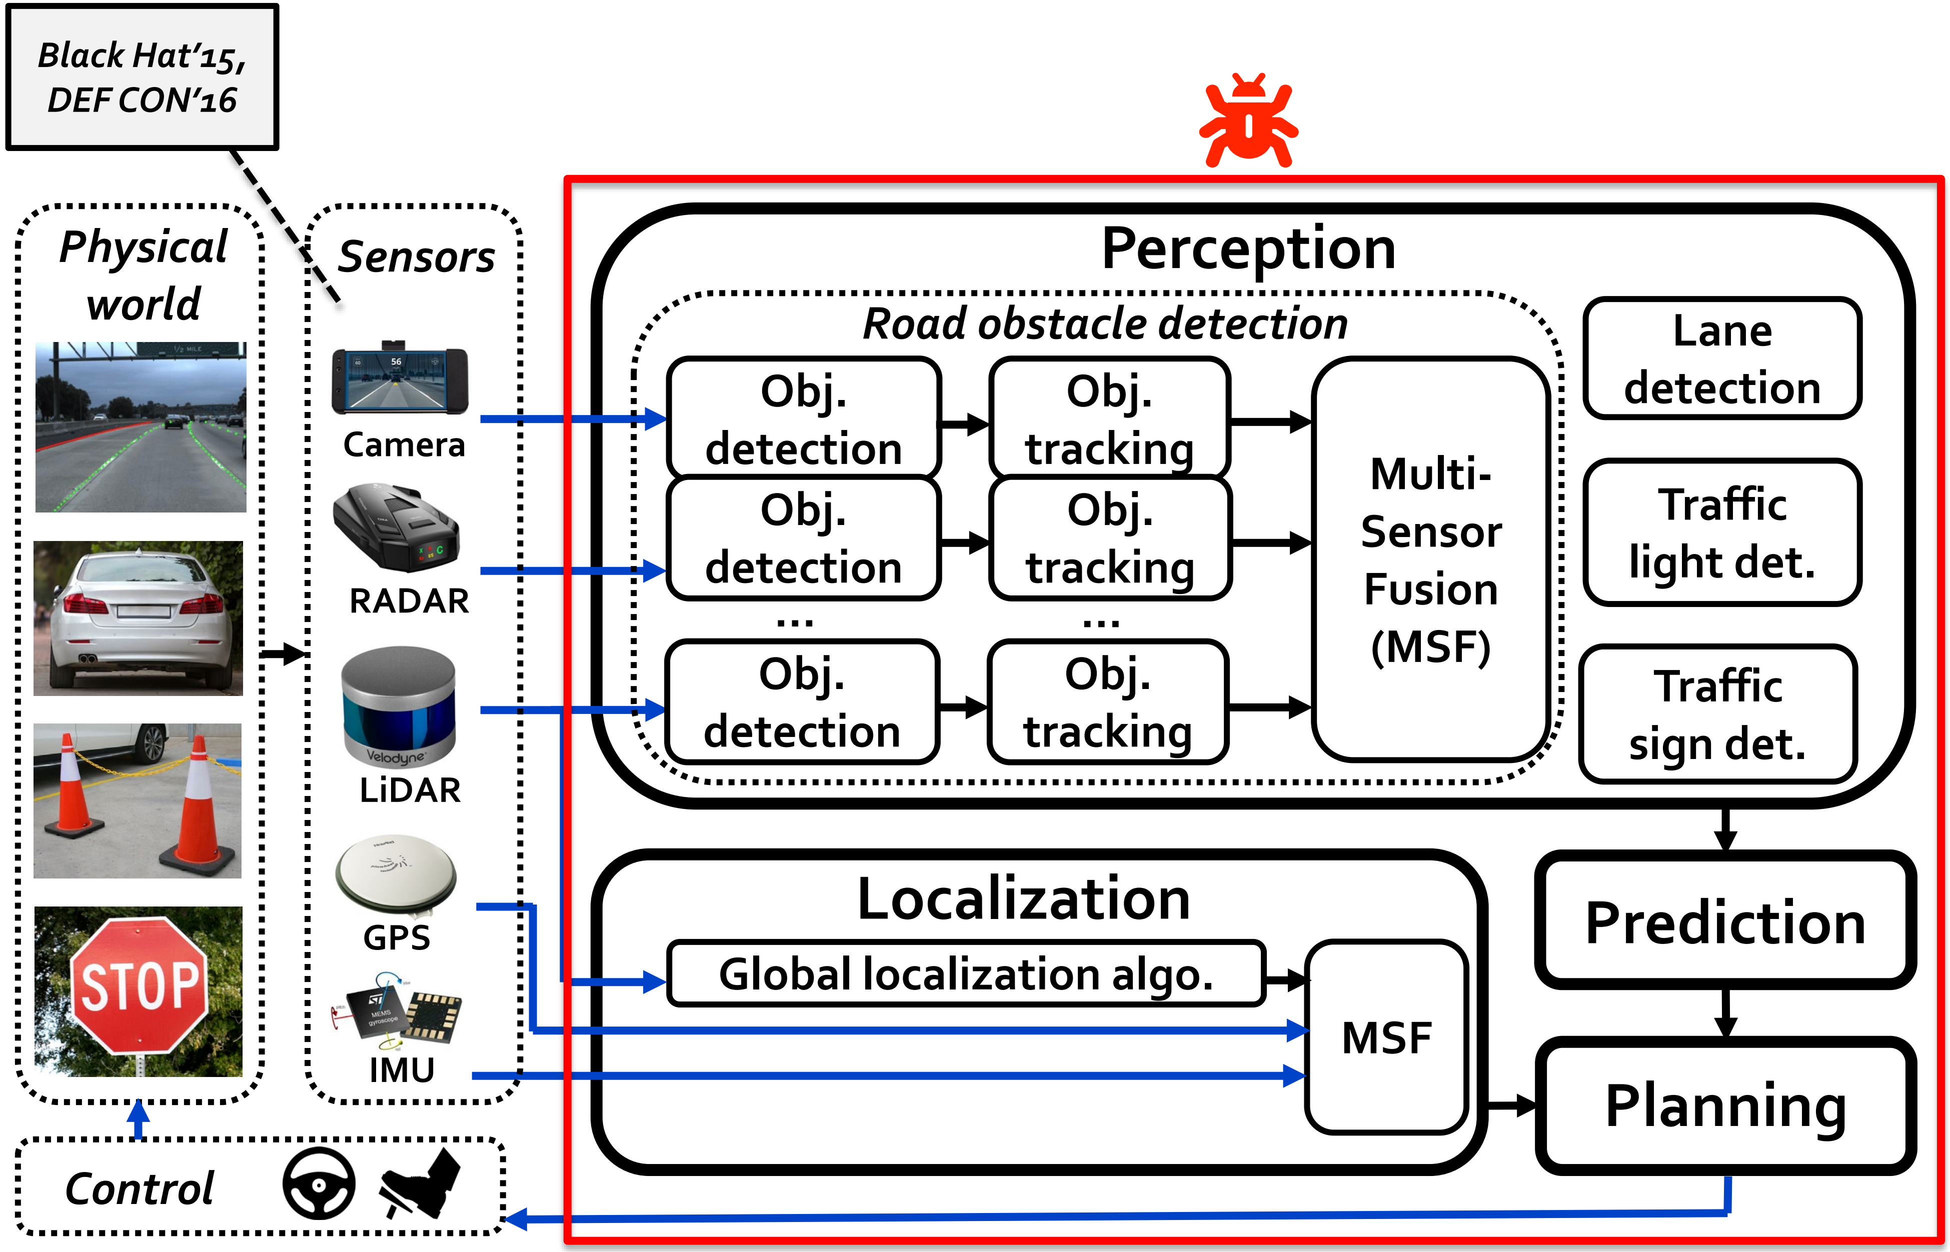

Black Hat'15, DEF CON'16
Physical world
STOP
Sensors
Camera
RADAR
LiDAR
Velodyne
GPS
IMU
Control
Perception
Road obstacle detection
Obj. detection
Obj. tracking
Obj. detection
Obj. tracking
...
...
Obj. detection
Obj. tracking
Multi-Sensor Fusion (MSF)
Lane detection
Traffic light det.
Traffic sign det.
Localization
Global localization algo.
MSF
Prediction
Planning

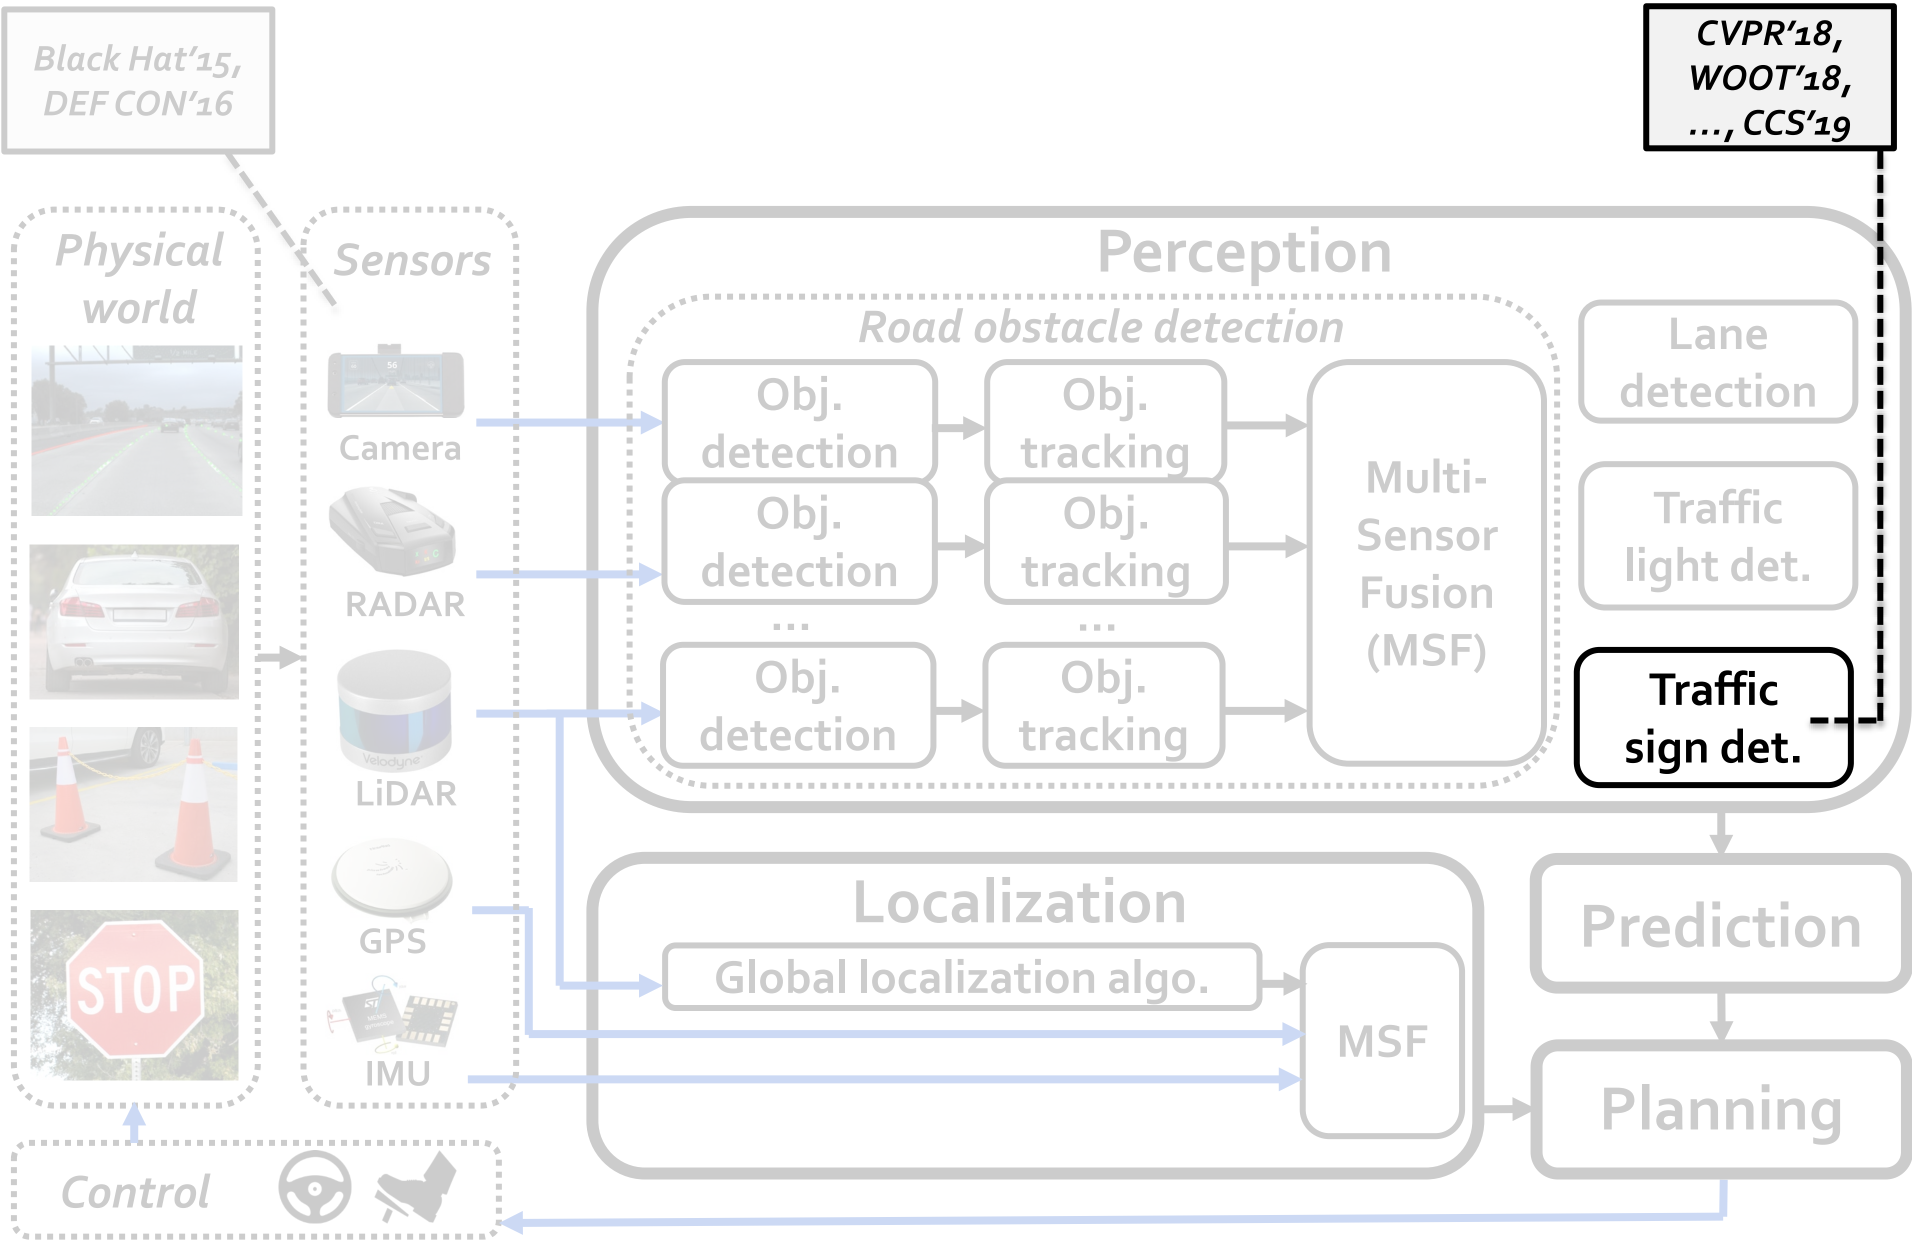
Black Hat'15, DEF CON'16
CVPR'18, WOOT'18, …, CCS'19
Physical world
STOP
Sensors
Camera
RADAR
LiDAR
Velodyne
GPS
IMU
Perception
Road obstacle detection
Obj. detection
Obj. tracking
Obj. detection
Obj. tracking
…
…
Obj. detection
Obj. tracking
Multi-Sensor Fusion (MSF)
Lane detection
Traffic light det.
Traffic sign det.
Localization
Global localization algo.
MSF
Prediction
Planning
Control

Black Hat'15, DEF CON'16

CVPR'18, WOOT'18, …, CCS'19

- Target: State-of-the-art (SOTA) camera object detection DNN
- Attack vector: Malicious stickers/posters
- Impact: Make a traffic sign disappear, misclassified, or spoof one

[Zhao et al. @ CCS'19]

Black Hat'15, DEF CON'16

CVPR'18, WOOT'18, …, CCS'19

- Target: State-of-the-art (SOTA) camera object detection DNN
- Attack vector: Malicious stickers/posters
- Impact: Make a traffic sign disappear, misclassified, or spoof one
  - Latest work shown to work on **commercial vehicle system**[1]

Control

[1] Jia et al. @ NDSS'22

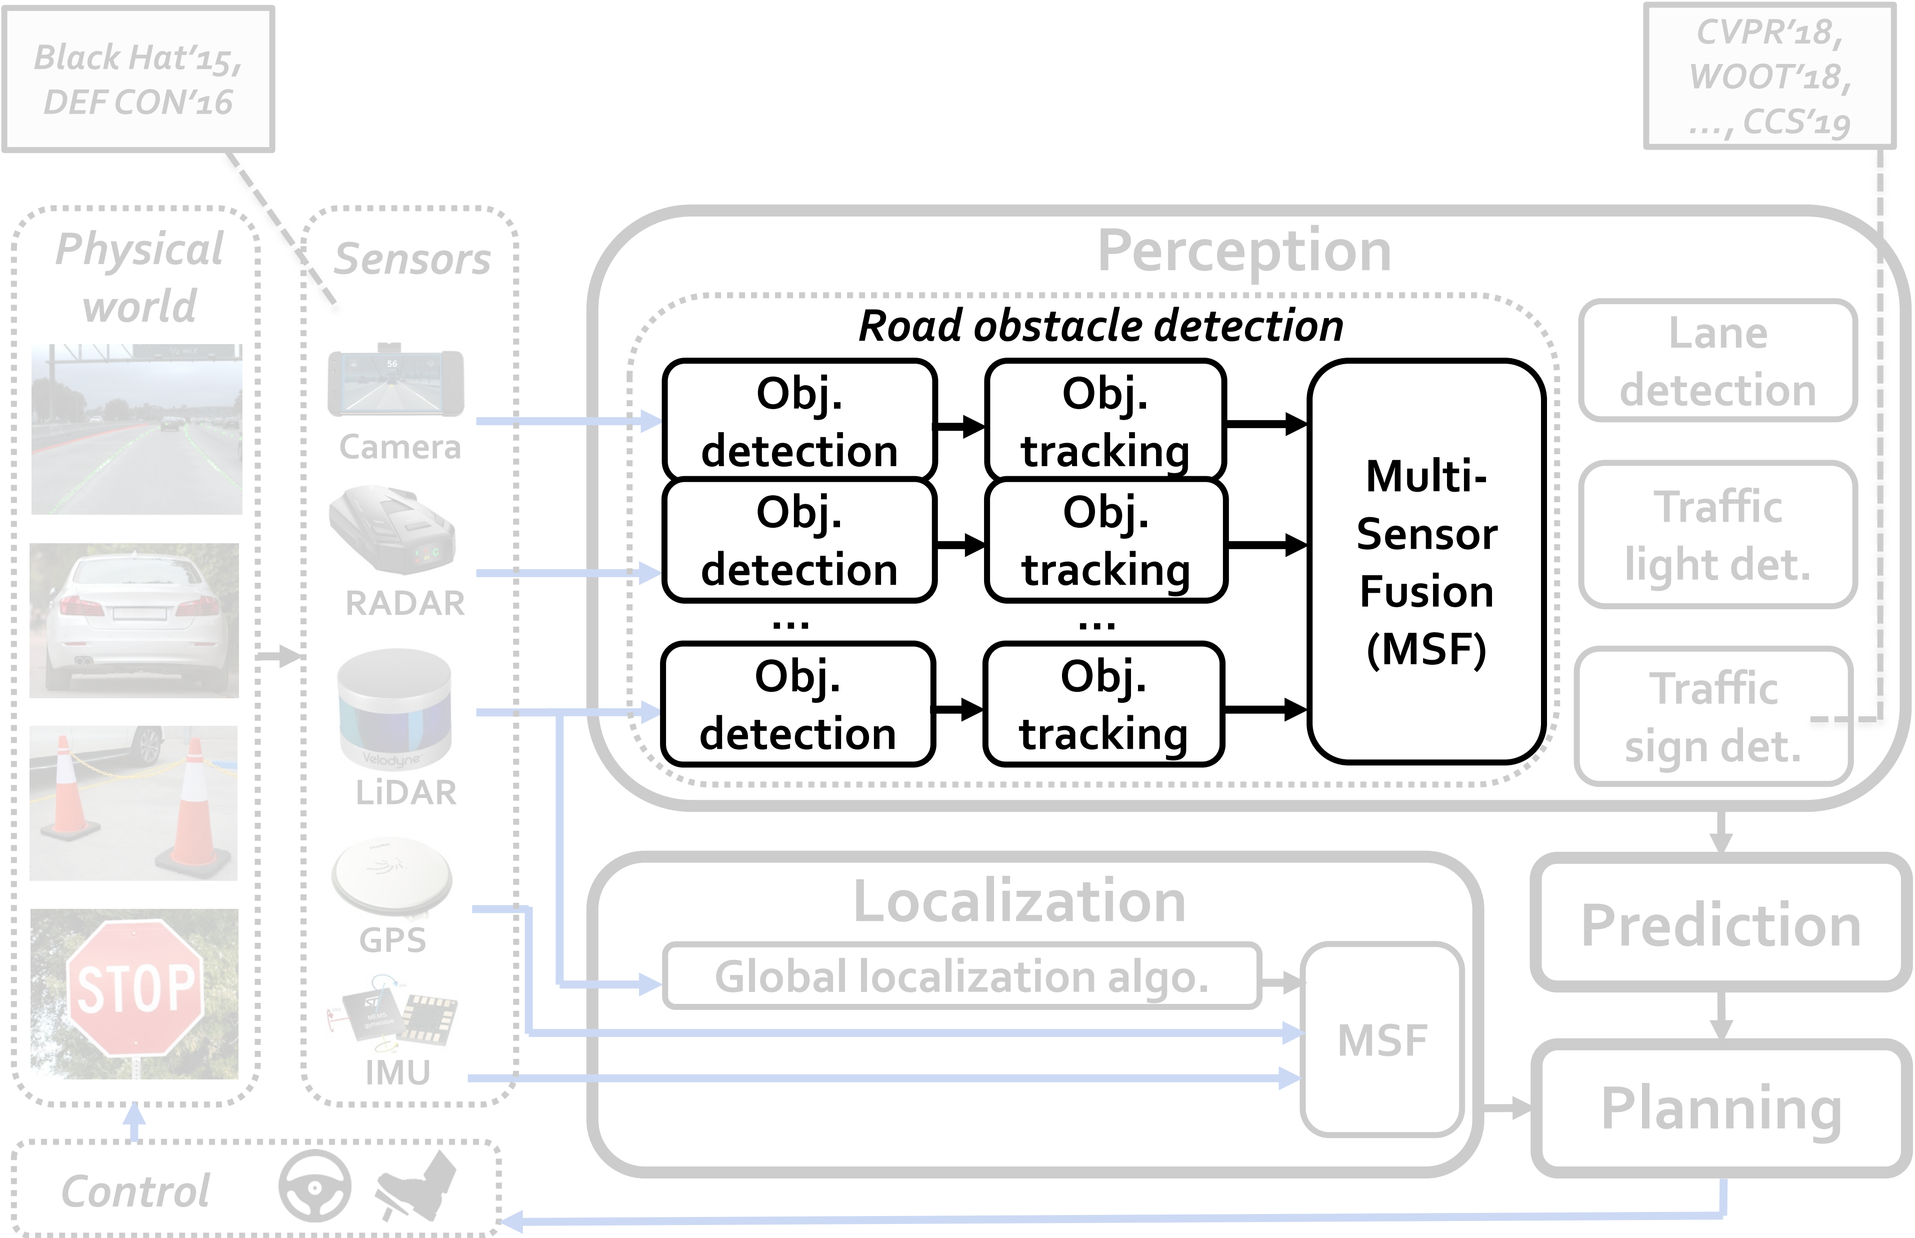

Black Hat'15, DEF CON'16
CVPR'18, WOOT'18, ..., CCS'19
Physical world
Sensors
STOP
Camera
RADAR
LiDAR
GPS
IMU
Perception
Road obstacle detection
Obj. detection
Obj. tracking
Obj. detection
Obj. tracking
...
...
Obj. detection
Obj. tracking
Multi-Sensor Fusion (MSF)
Lane detection
Traffic light det.
Traffic sign det.
Localization
Global localization algo.
MSF
Prediction
Planning
Control

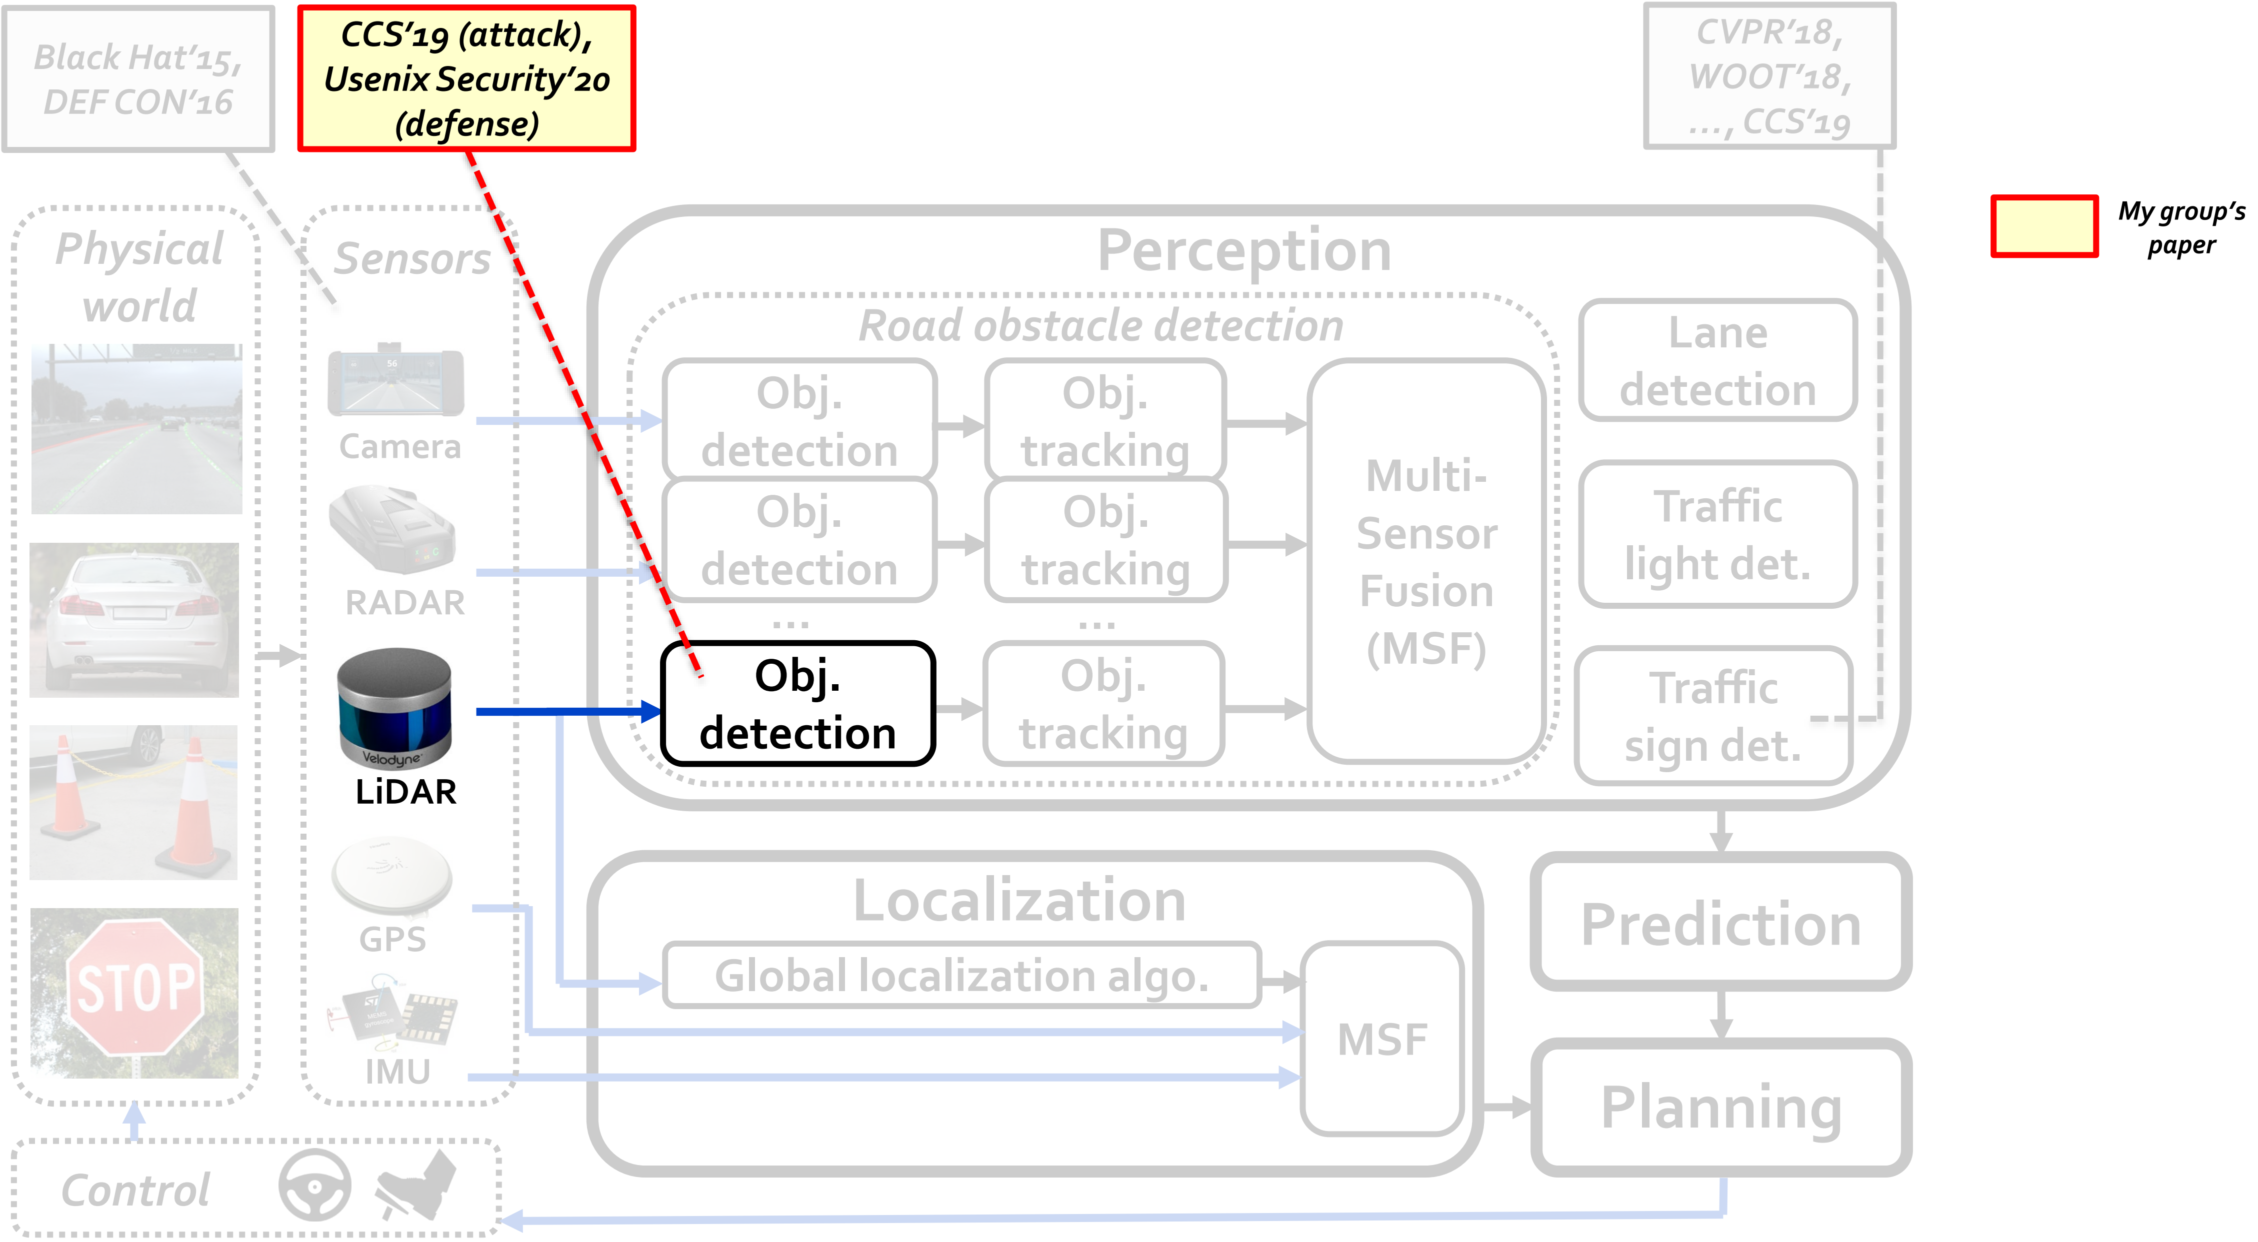
Black Hat'15,
DEF CON'16
CCS'19 (attack),
Usenix Security'20
(defense)
CVPR'18,
WOOT'18,
…, CCS'19
My group's paper
Physical world
STOP
Sensors
Camera
RADAR
LiDAR
GPS
IMU
Perception
Road obstacle detection
Obj. detection
Obj. tracking
Obj. detection
Obj. tracking
Obj. detection
Obj. tracking
Multi-Sensor Fusion (MSF)
Lane detection
Traffic light det.
Traffic sign det.
Localization
Global localization algo.
MSF
Prediction
Planning
Control

- ***First*** security analysis for **3D object detection**
- Attack vector: LiDAR spoofing

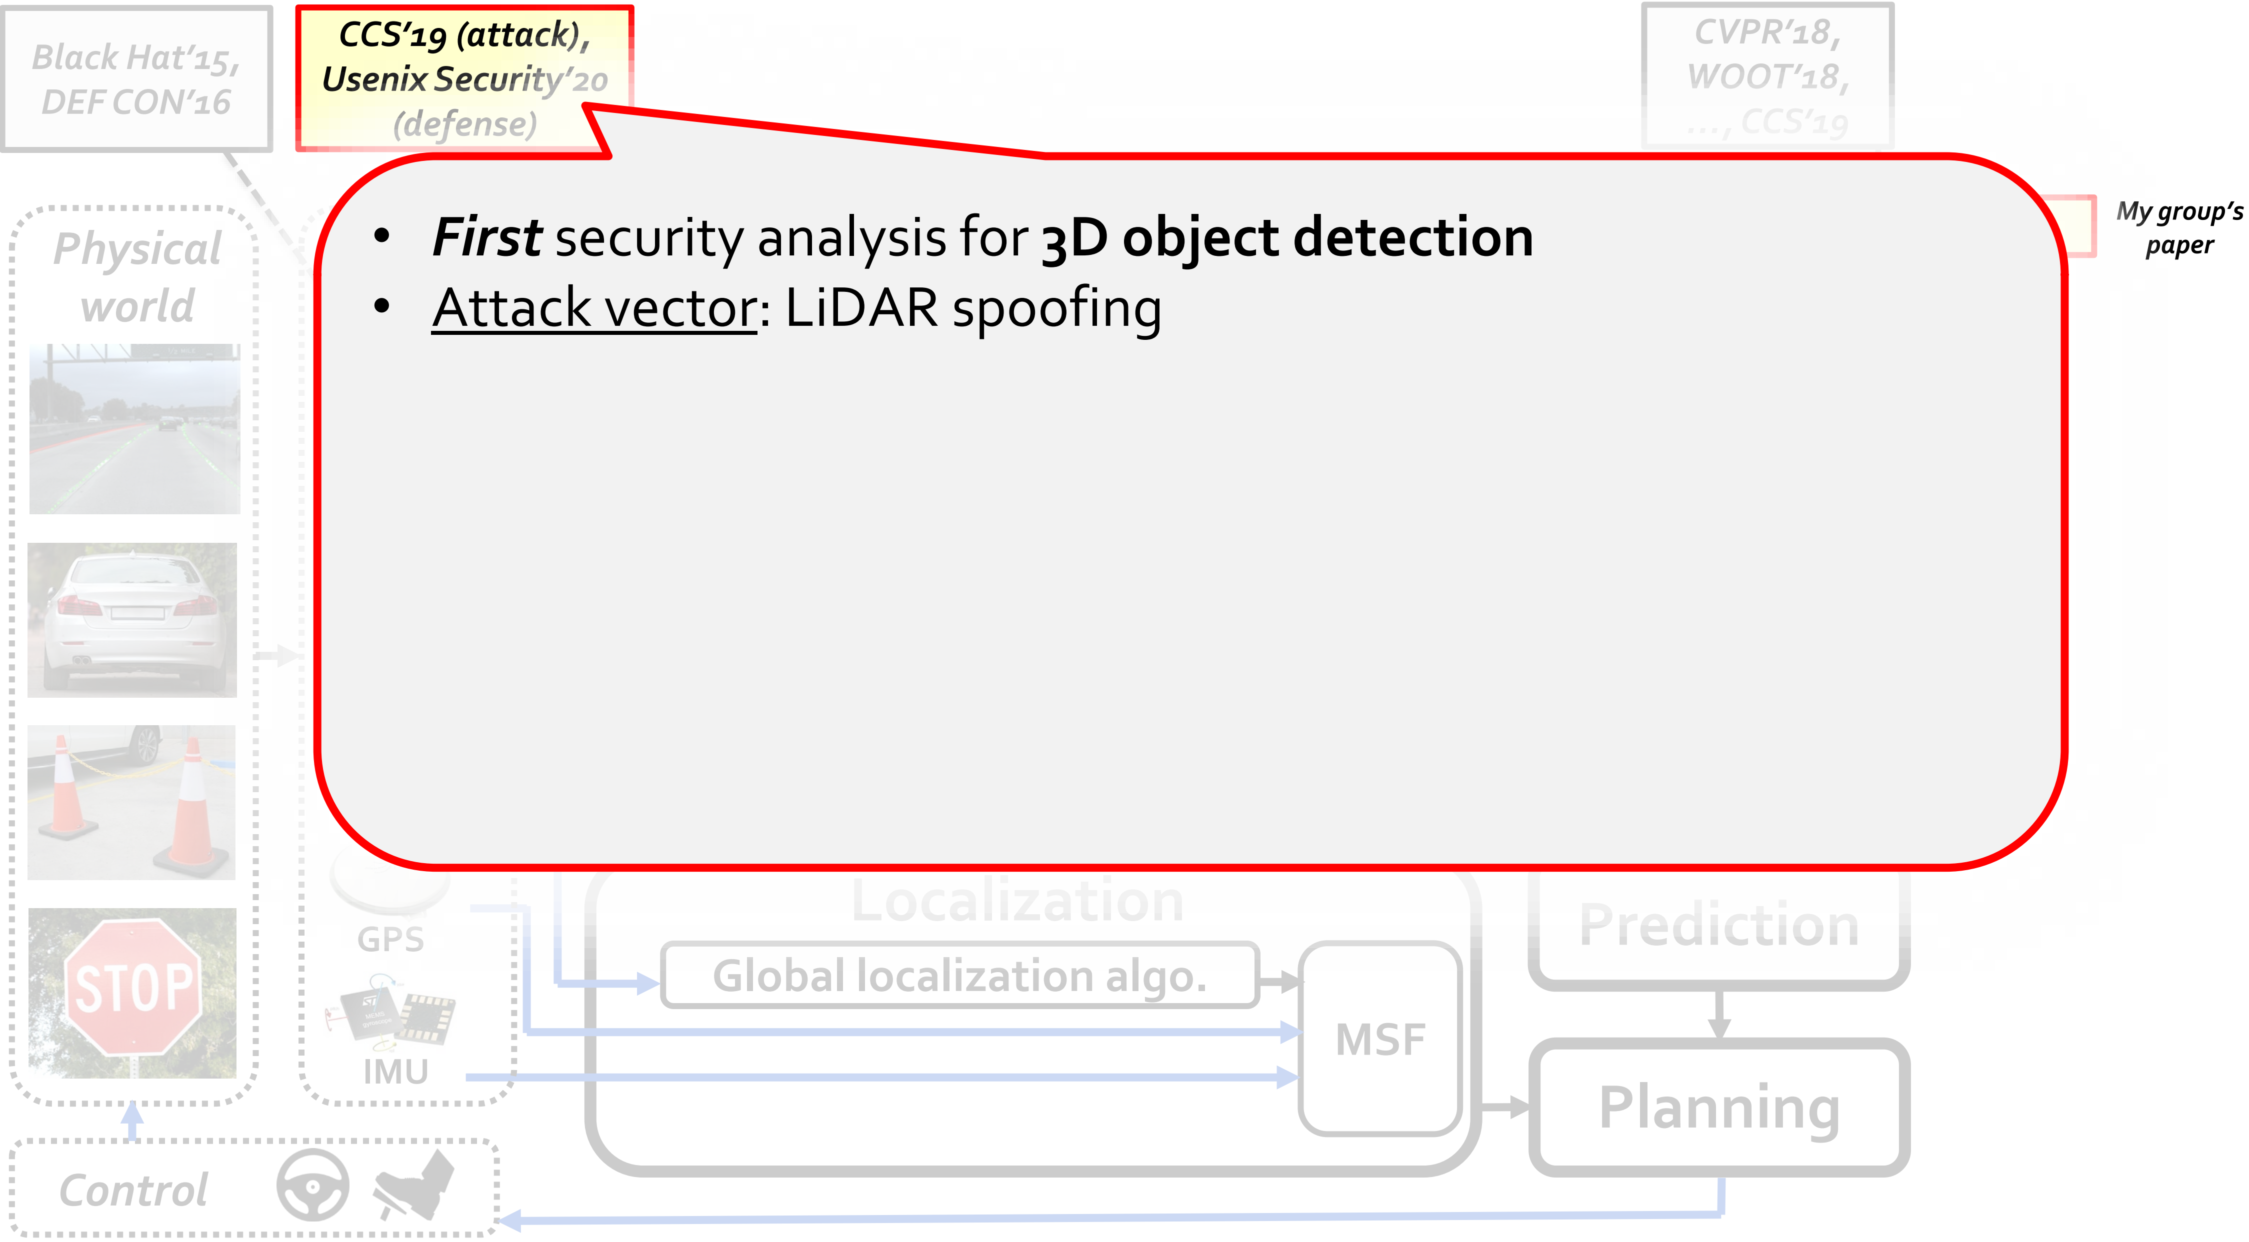

Background: LiDAR basics

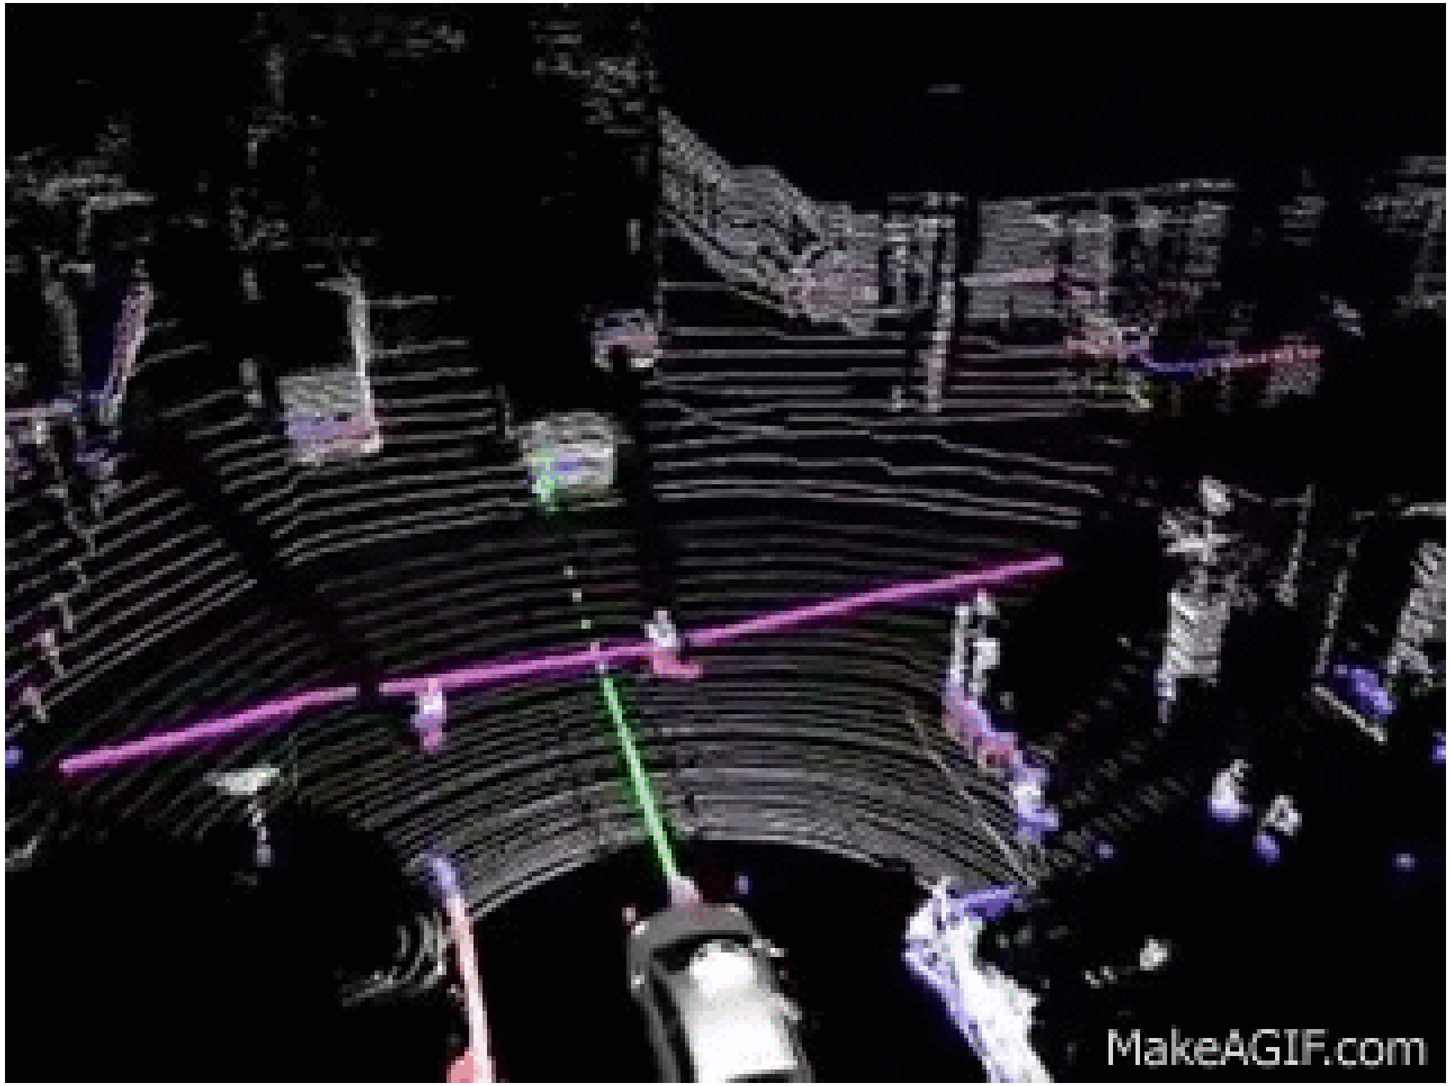

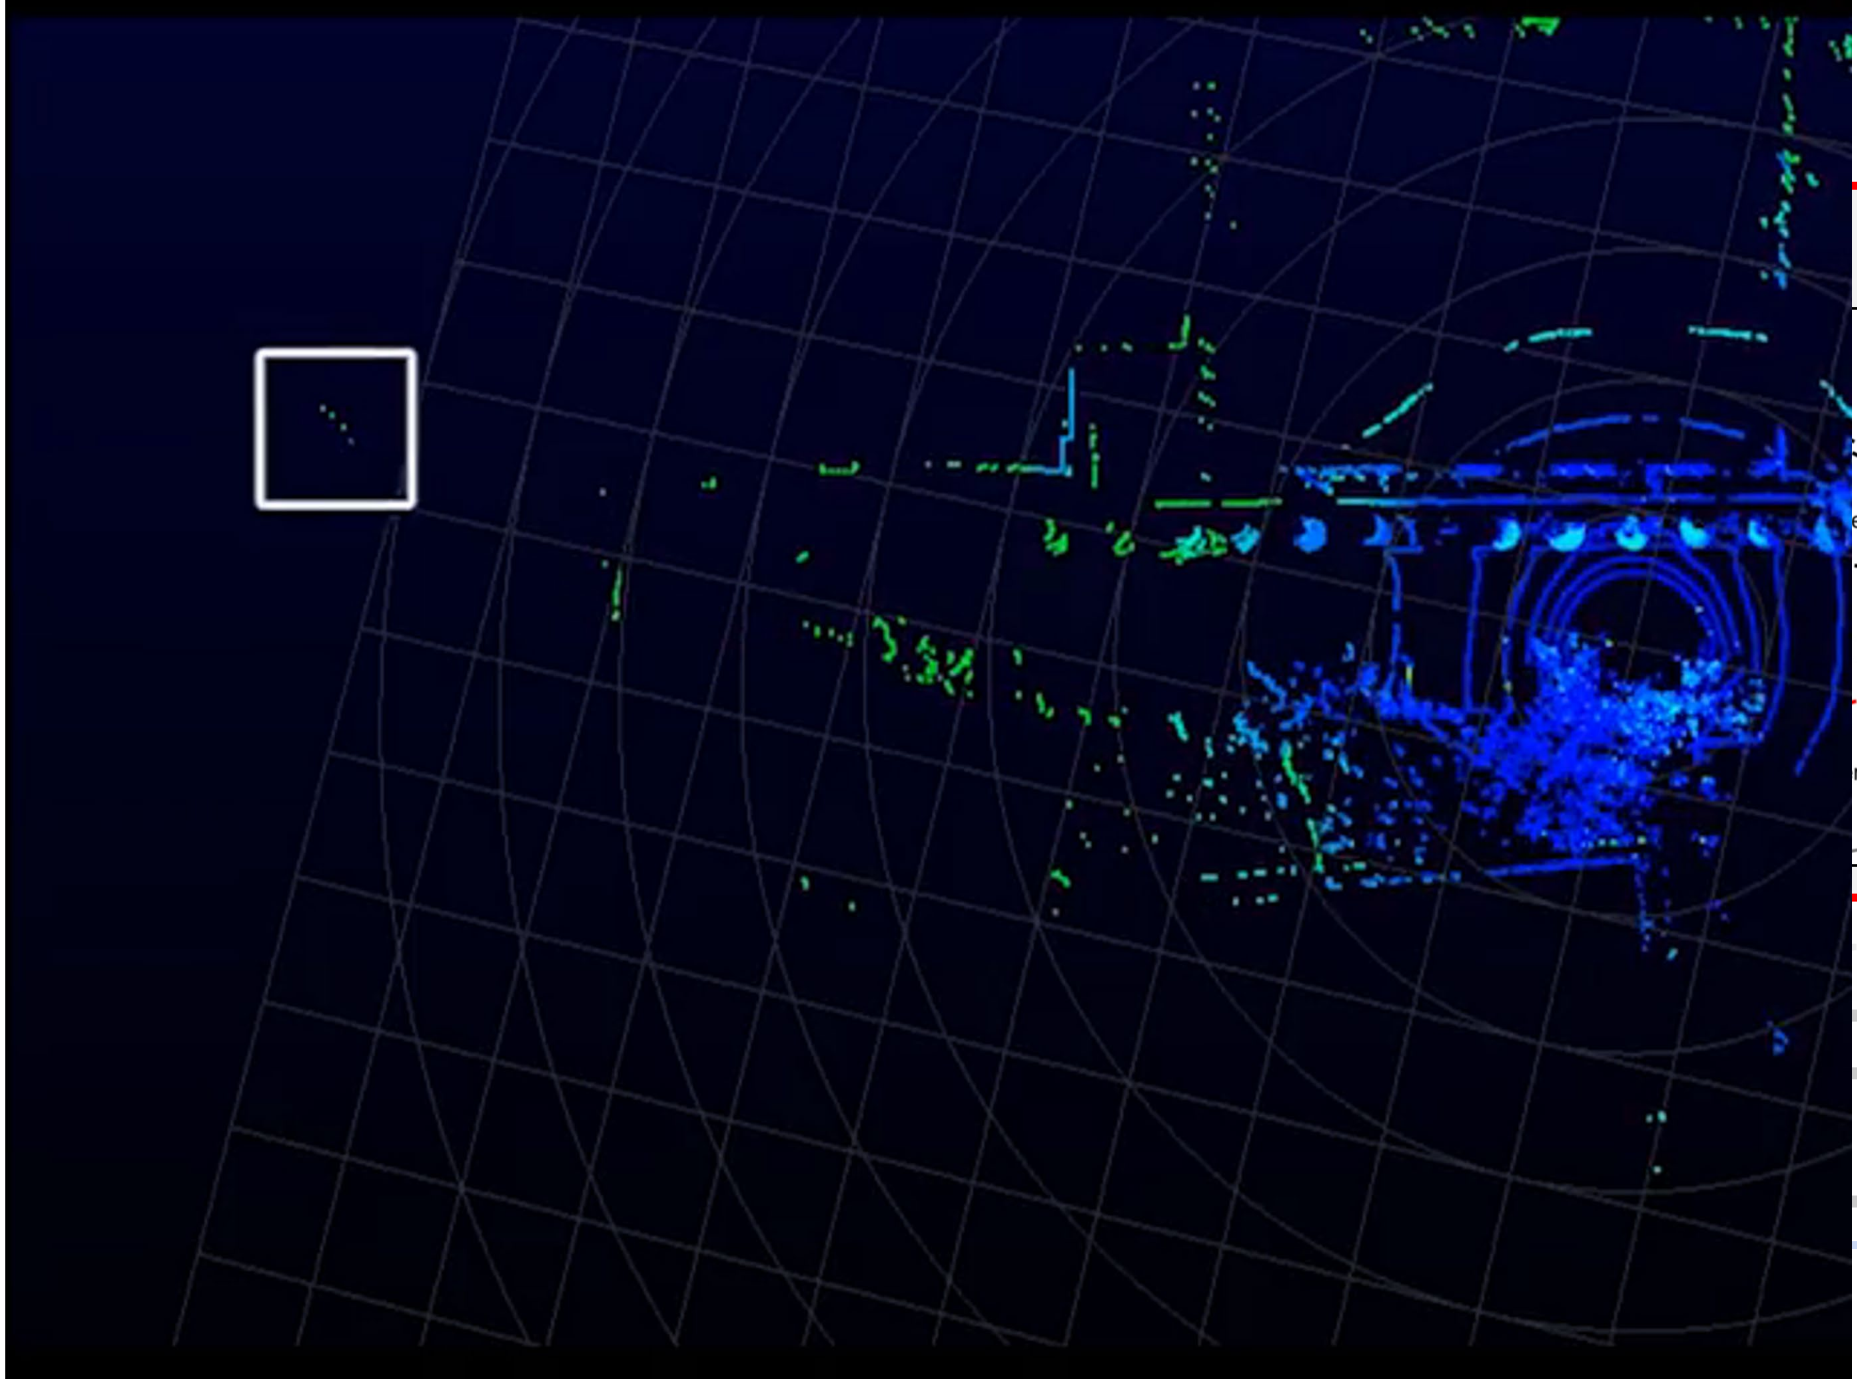

[Shin et al. @ CHES'17]

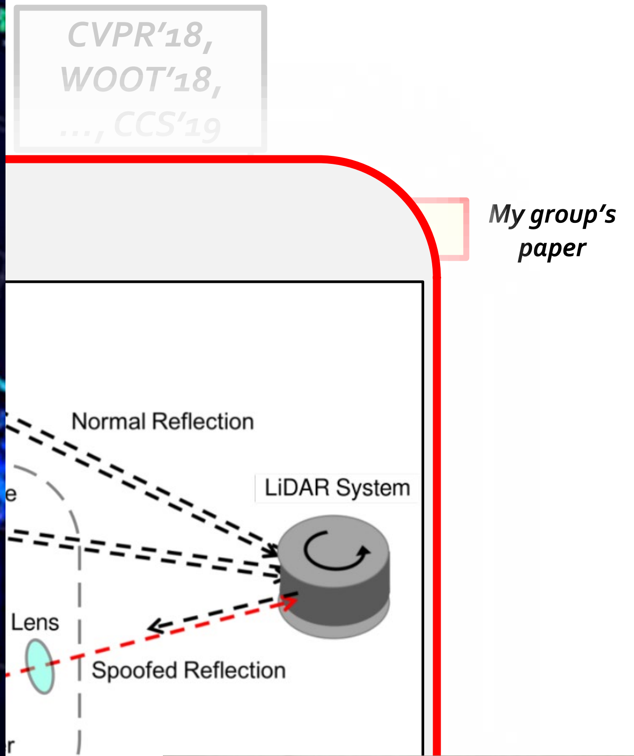

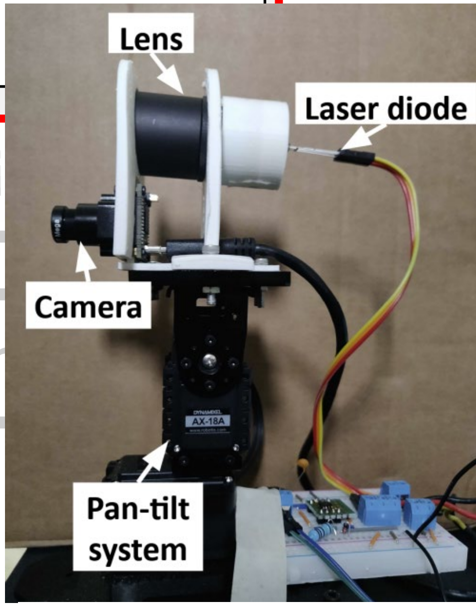

[Cao et al. @ AutoSec'21]

Black Hat'15, DEF CON'16

***CCS'19 (attack), Usenix Security'20** (defense)*

CVPR'18, WOOT'18, …, CCS'19

My group's paper

Physical world

- ***First*** security analysis for **3D object detection**
- Attack vector: LiDAR spoofing

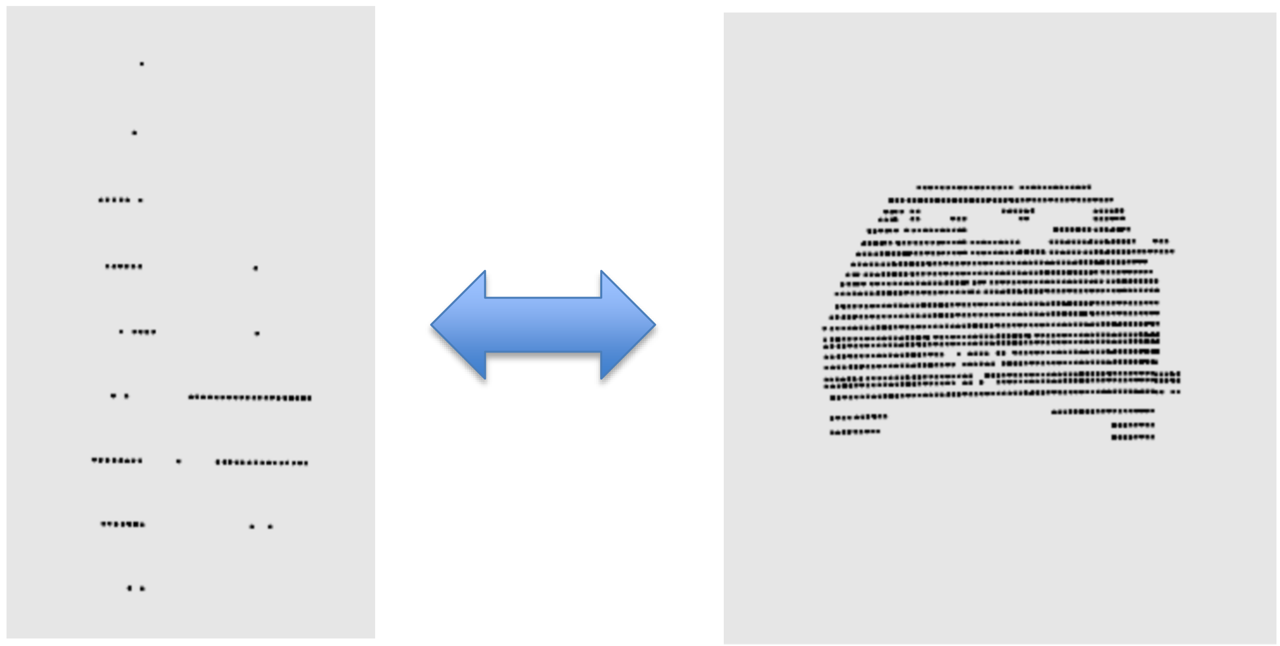

Control

ediction

anning

Blind spoofing is not enough: Tried various different angles & shapes, *cannot spoof fake obstacles at SOTA LiDAR detection model output at all*

Black Hat'15, DEF CON'16

CCS'19 (attack), Usenix Security'20 (defense)

CVPR'18, WOOT'18, …, CCS'19

My group's paper

Physical world

- ***First*** security analysis for **3D object detection**
- Attack vector: LiDAR spoofing
- Solution: Combine sensor spoofing with adversarial AI attack!

First sensor-AI co-designed attack

- Call it "adversarial sensor attack"

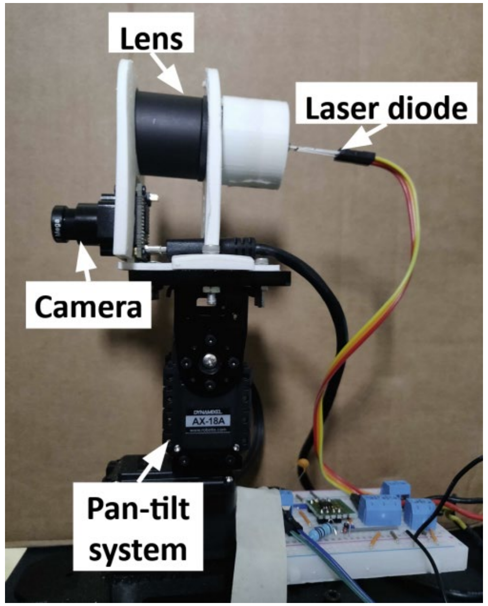

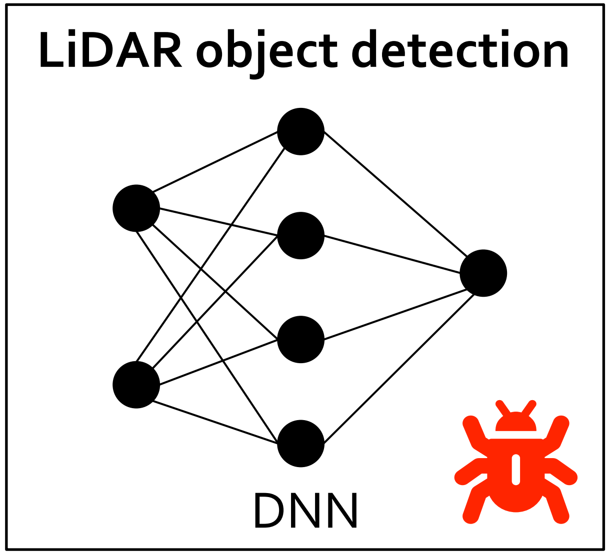

Control

Black Hat'15, DEF CON'16

CCS'19 (atta
Usenix Securi
(defense)

Physical world

- ***First***
- Atta
- Solu ... rsarial AI attack!
  - Optimize w/ **differentiable spoofing capability modelling & spatial transformation** of attack trace

Differentiable

$$x' = x \bigoplus t' = \begin{bmatrix} I^{x}_{cnt} + I^{t'}_{cnt} \\ (I^{x}_{avg\_h} \cdot I^{x}_{cnt} + I^{t'}_{avg\_h} \cdot I^{t'}_{cnt})/(I^{x}_{cnt} + I^{t'}_{cnt}) \\ \max(I^{x}_{max\_h}, I^{t'}_{max\_h}) \\ (I^{x}_{avg\_int} \cdot I^{x}_{cnt} + I^{t'}_{avg\_int} \cdot I^{t'}_{cnt})/(I^{x}_{cnt} + I^{t'}_{cnt}) \\ \sum I^{x}_{max\_int} \cdot \mathbf{1}\{I^{x}_{max\_h} = max\{I^{x}_{max\_h}, I^{t'}_{max\_h}\}\} \end{bmatrix}$$

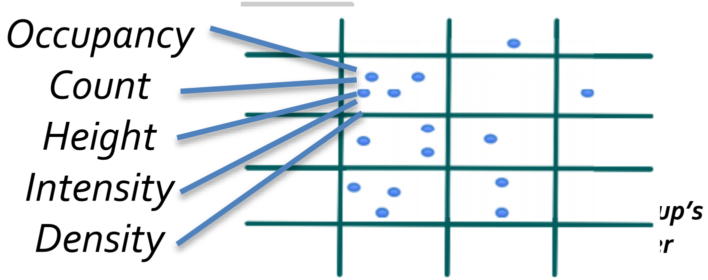

Agg. features

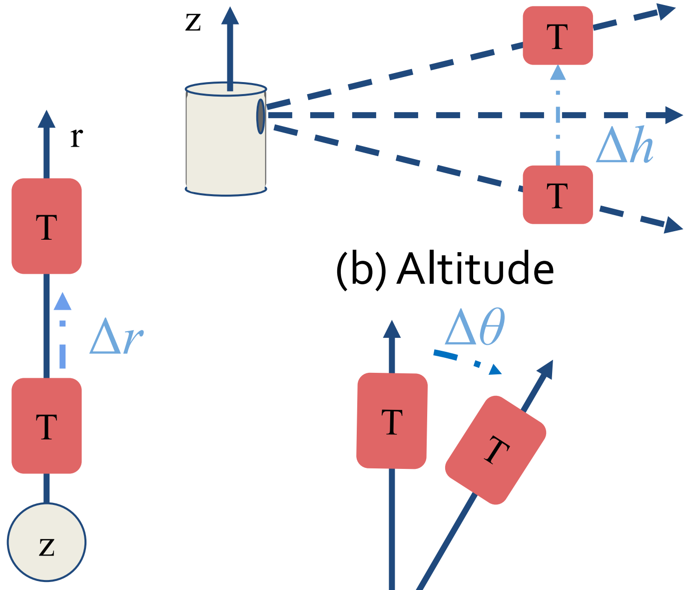

(a) Distance

(b) Altitude

(c) Azimuth

Differentiable

$$\begin{bmatrix} T'_{w_x} \\ T'_{w_y} \\ T'_{w_z} \\ 1 \end{bmatrix} = \begin{bmatrix} \cos\theta & -\sin\theta & 0 & \tau_x \\ \sin\theta & \cos\theta & 0 & 0 \\ 0 & 0 & s_h & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} T_{w_x} \\ T_{w_y} \\ T_{w_z} \\ 1 \end{bmatrix}$$

Spatial Transformation

$G_t\ (\theta,\ \tau_x,\ s_h;t)$

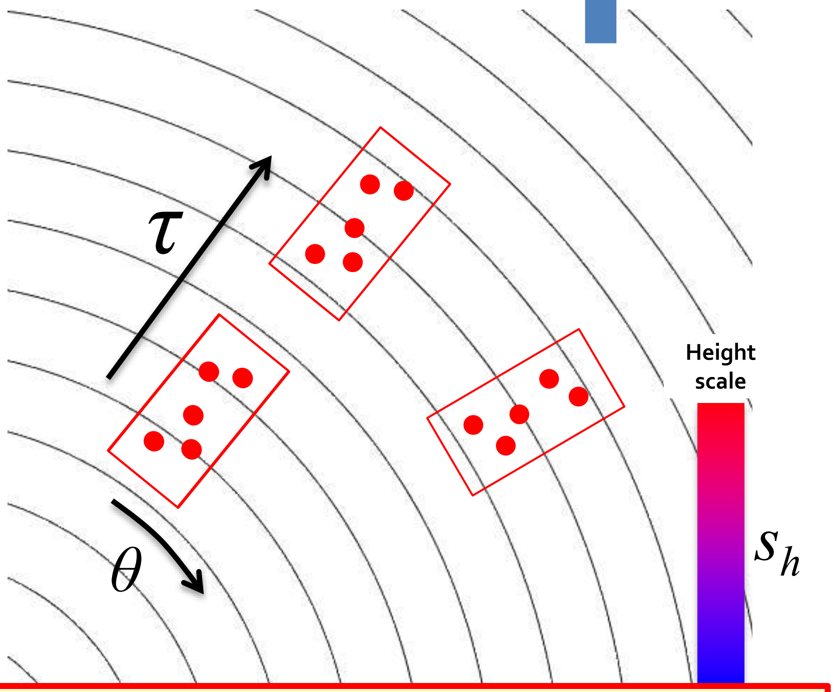

Addressing *System-to-AI semantic gap*: from attack perturbation capability at CPS system input space (i.e., LiDAR spoofing capability) to that at AI component input space (i.e., DNN model input)

Black Hat'15, DEF CON'16

CCS'19 (attack), Usenix Security'20 (defense)

CVPR'18, WOOT'18, …, CCS'19

My group's paper

Physical world

- ***First*** security analysis for **3D object detection**
- Attack vector: LiDAR spoofing
- Solution: Combine sensor spoofing with adversarial AI attack!
  - Optimize w/ **differentiable spoofing capability modelling** & **spatial transformation** of attack trace
  - **Global sampling** to avoid trapping at local minima due to hard perturbation constraints imposed by spoofing capability
  - ***0% → 75% success rate*** *in spoofing a near-front vehicle!*

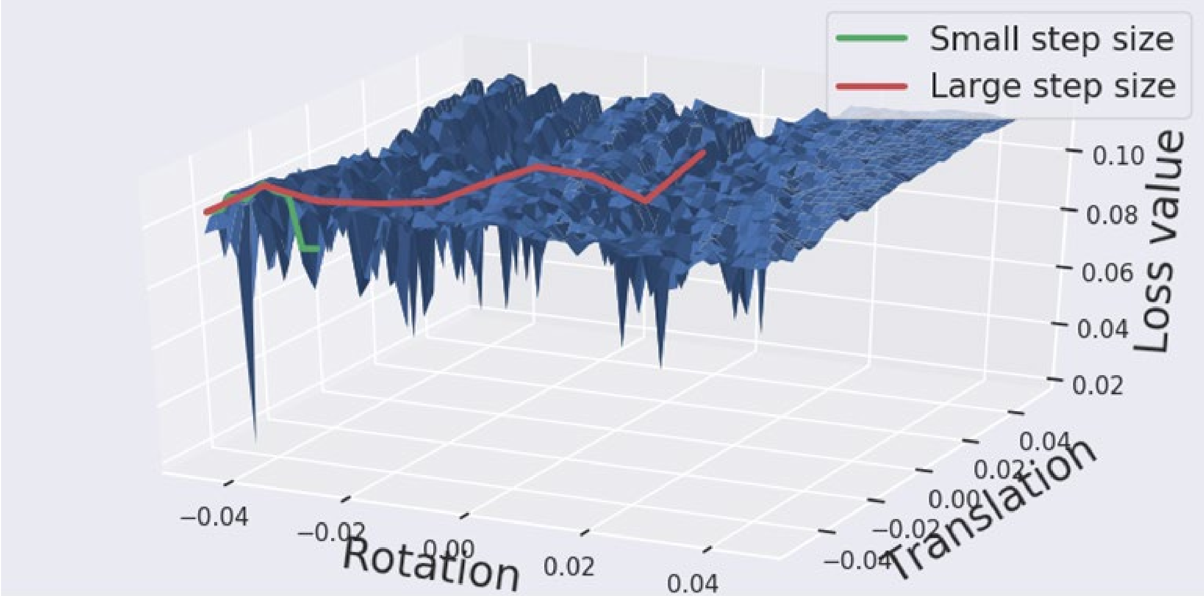

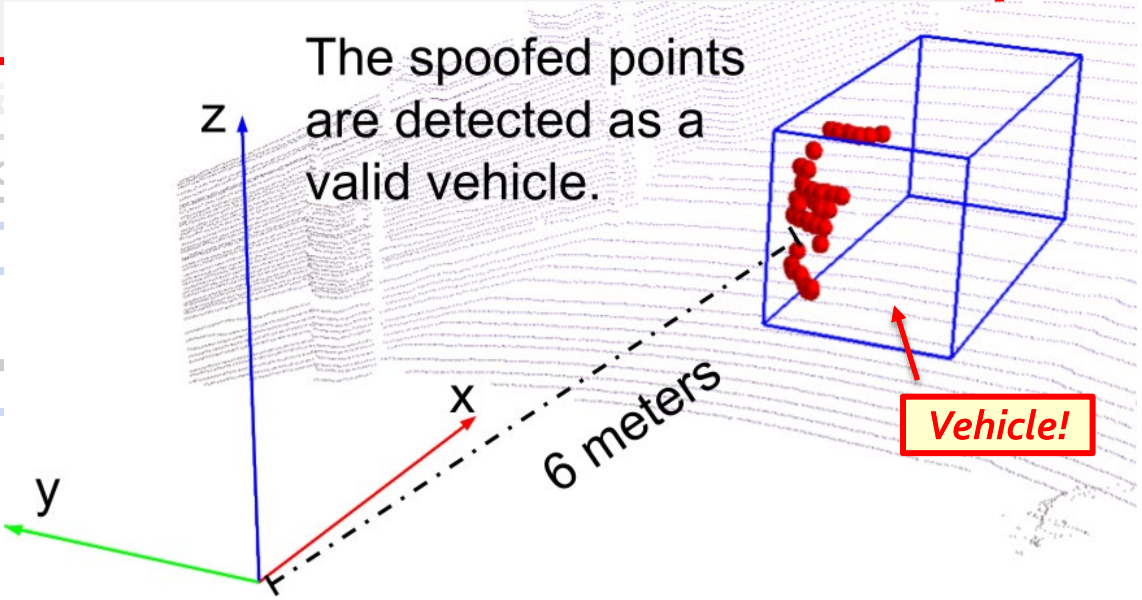

Black Hat'15, DEF CON'16

CCS'19 (attack), Usenix Security'20 (defense)

CVPR'18, WOOT'18, …, CCS'19

My group's paper

Physical world

- ***First*** security analysis for **3D object detection**
- Attack vector: LiDAR spoofing
- Solution: Combine sensor spoofing with adversarial AI attack!
  - Optimize w/ **differentiable spoofing capability modelling** & **spatial transformation** of attack trace
  - **Global sampling** to avoid trapping at local minima due to hard perturbation constraints imposed by spoofing capability
  - ***0% → 75% success rate** in spoofing a near-front vehicle!*
- Impact: Causing emergency brake or permanent stop

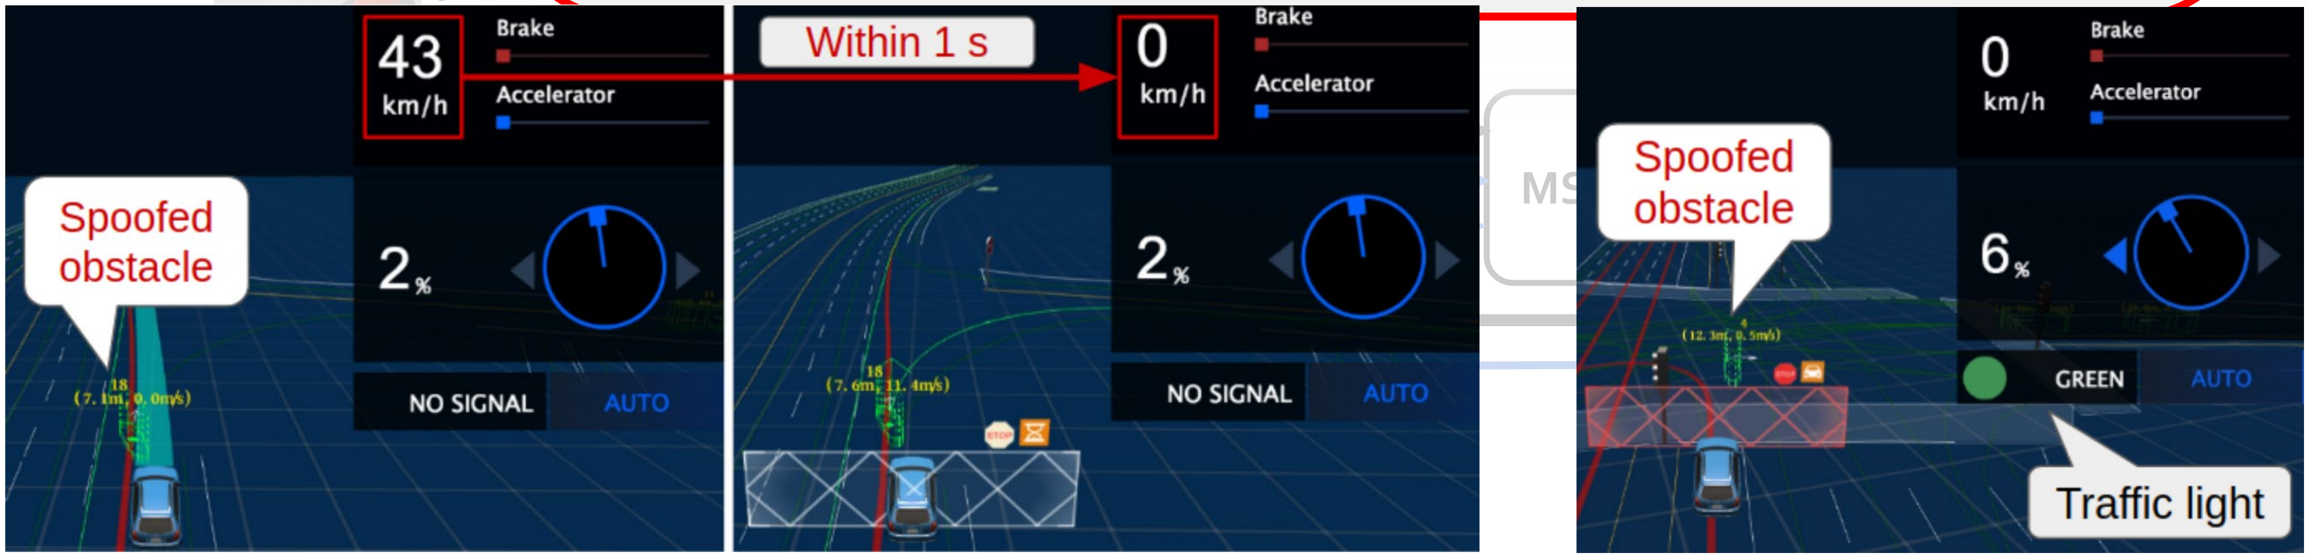

Black Hat'15, DEF CON'16

CCS'19 (attack), Usenix Security'20 (defense)

CVPR'18, WOOT'18, …, CCS'19

My group's paper

Physical world

- ***First*** security analysis for **3D object detection**
- Attack vector: LiDAR spoofing
- Solution: Combine sensor spoofing with adversarial AI attack!
    - Optimize w/ **differentiable spoofing capability modelling** & **spatial transformation** of attack trace
    - **Global sampling** to avoid trapping at local minima due to hard perturbation constraints imposed by spoofing capability
    - ***0% → 75% success rate** in spoofing a near-front vehicle!*
- Impact: Causi

Addressing *AI-to-System semantic gap*: from AI component-level errors (i.e., DNN output misdetection) to CPS system-level attack effect (i.e., emergency brake)

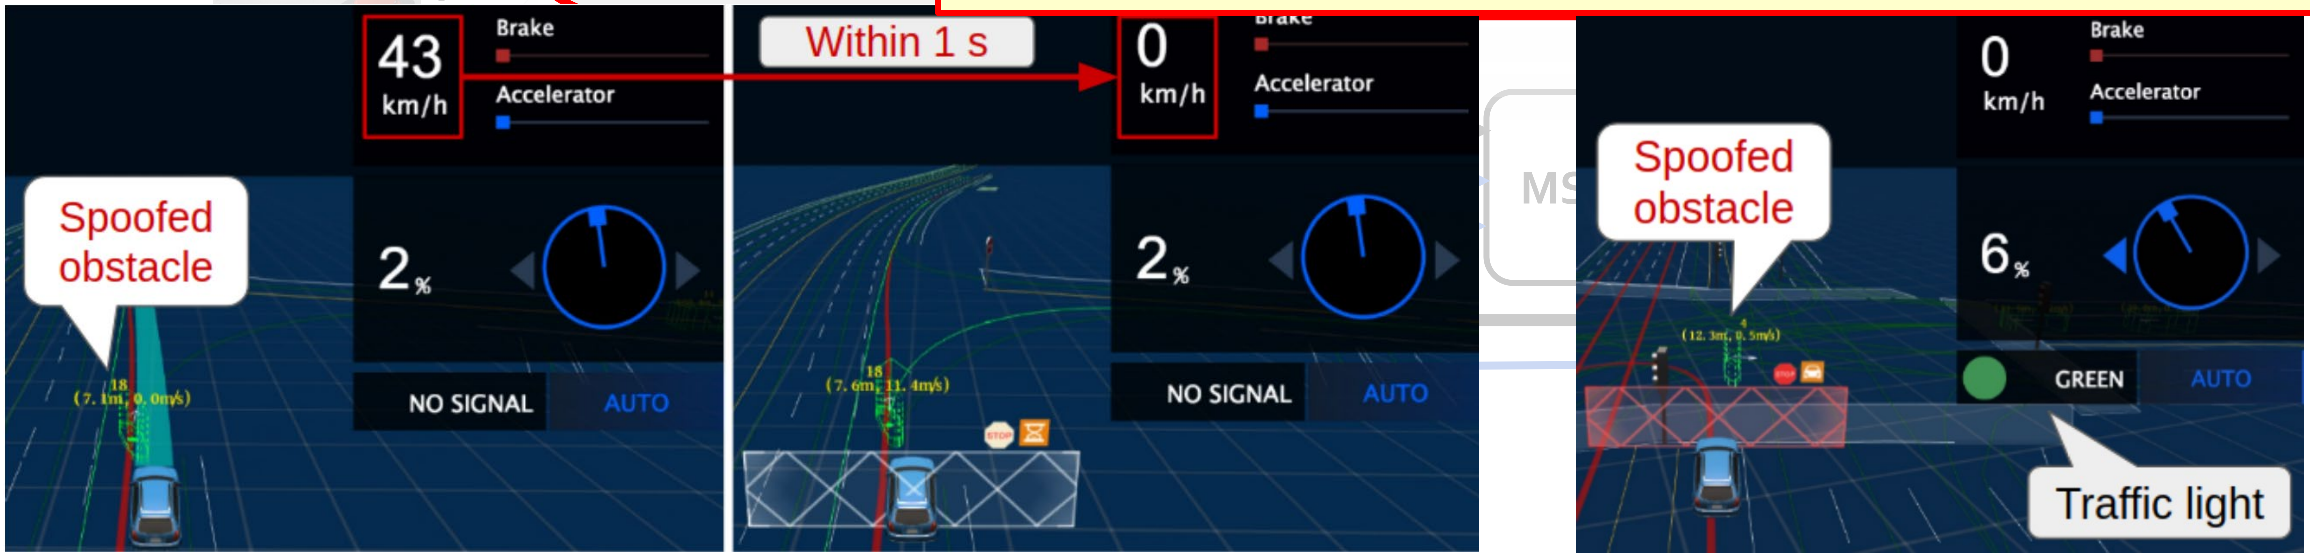

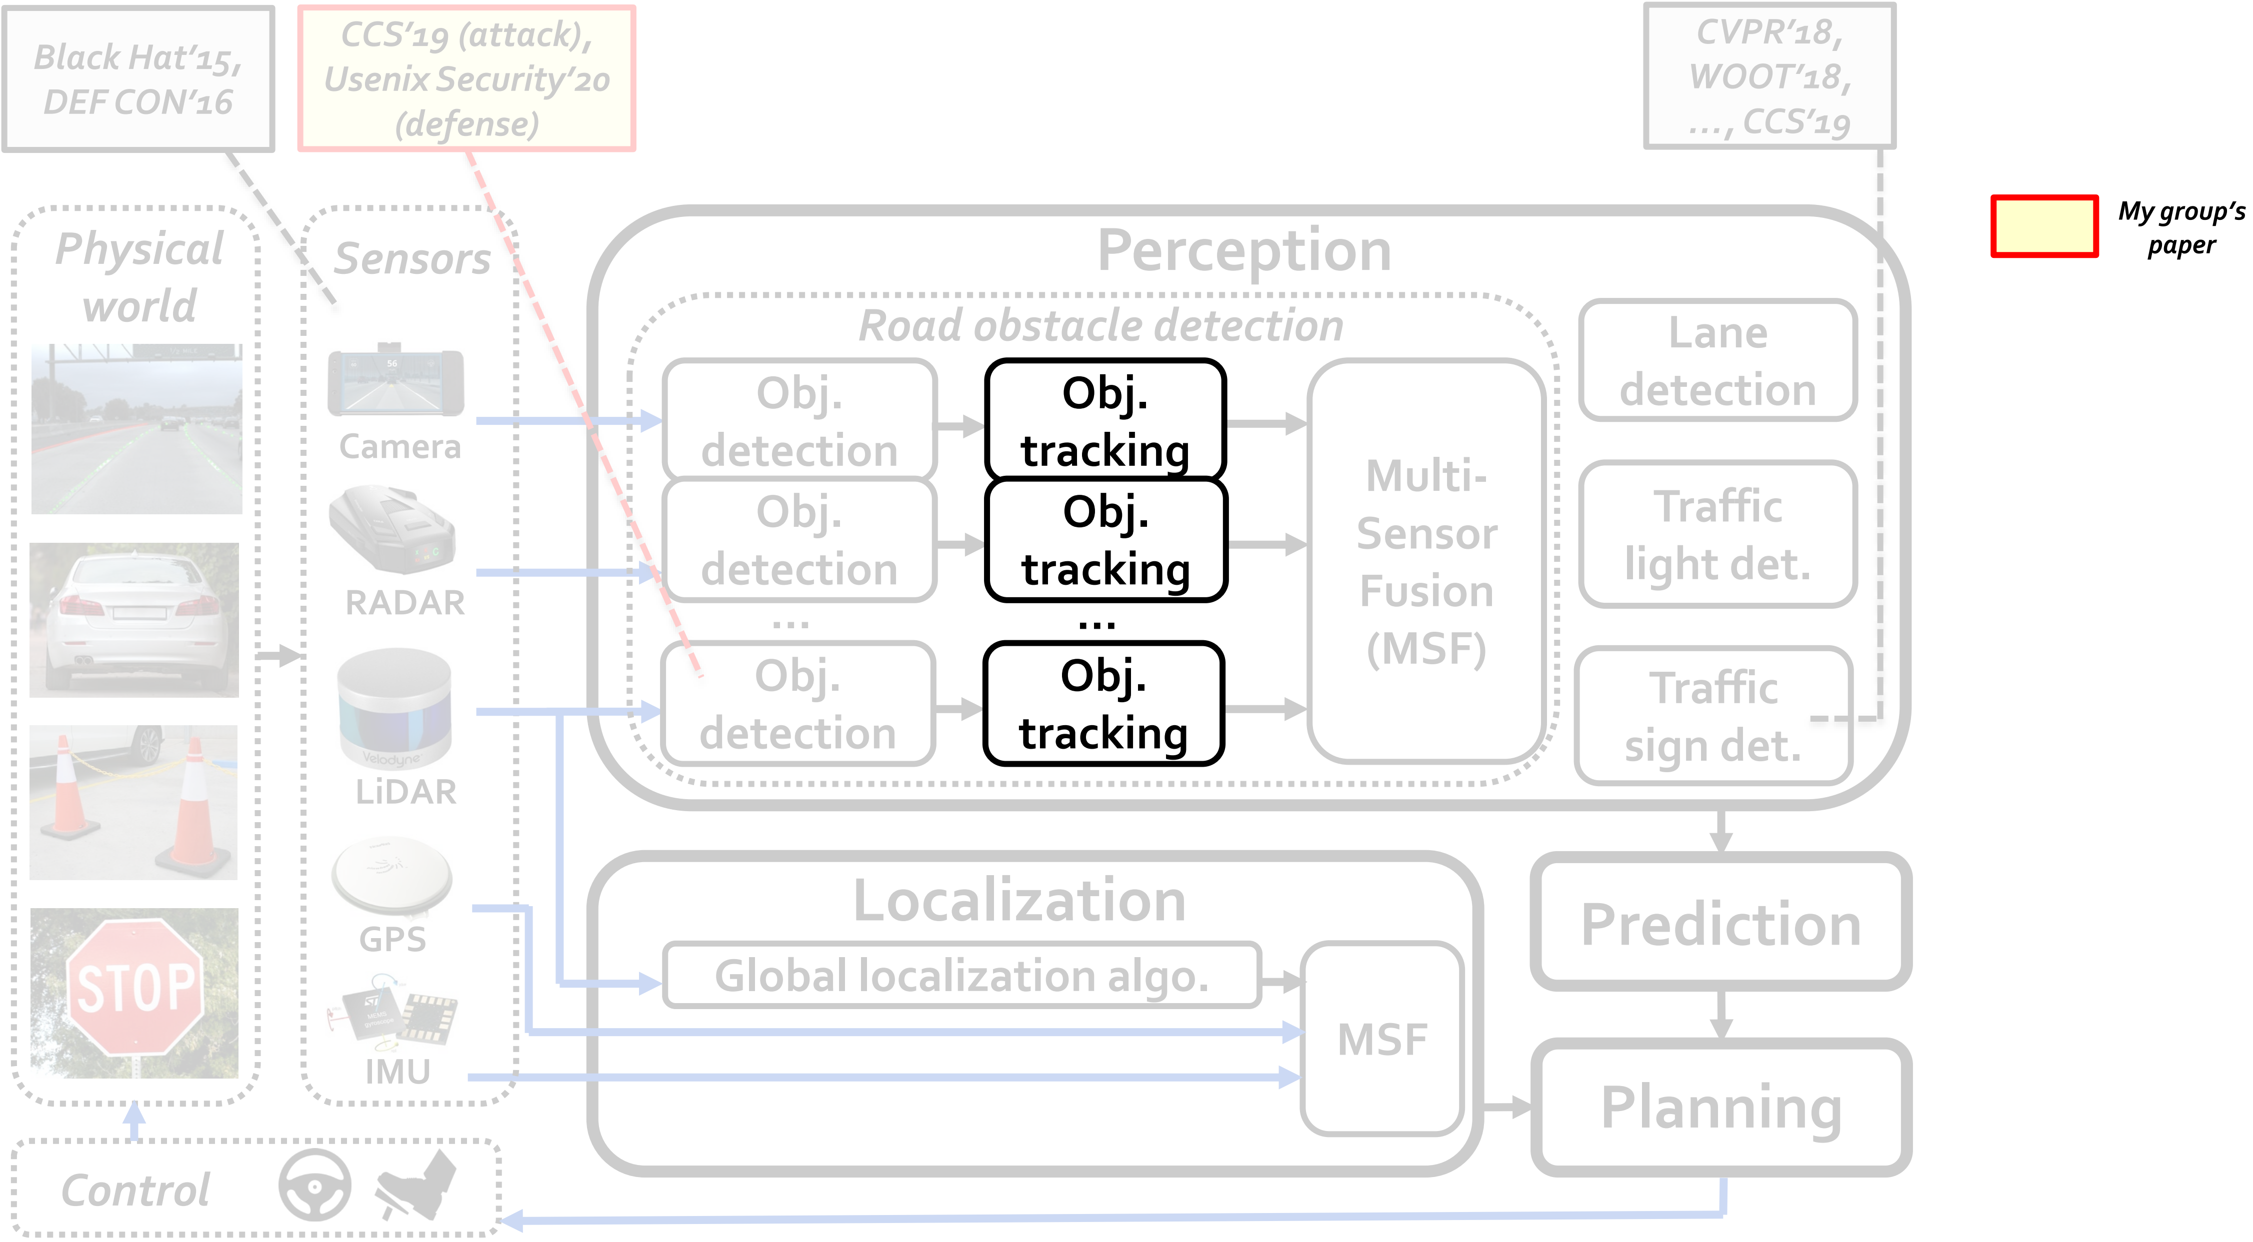
Black Hat'15, DEF CON'16
CCS'19 (attack), Usenix Security'20 (defense)
CVPR'18, WOOT'18, …, CCS'19
My group's paper
Physical world
STOP
Sensors
Camera
RADAR
LiDAR
GPS
IMU
Perception
Road obstacle detection
Obj. detection
Obj. tracking
Obj. detection
Obj. tracking
…
Obj. detection
Obj. tracking
Multi-Sensor Fusion (MSF)
Lane detection
Traffic light det.
Traffic sign det.
Localization
Global localization algo.
MSF
Prediction
Planning
Control

Black Hat'15, DEF CON'16

CCS'19 (attack), Usenix Security'20 (defense)

ICLR'20

- ***First*** security analysis for object tracking
- <u>Attack vector</u>: Stickers on the back of front car
- <u>Methodology</u>: Optimize bounding box position shifting
- <u>Impact</u>: Move a road-side object into the current lane, causing **emergency brake**; or move a front car away, causing a **crash**.

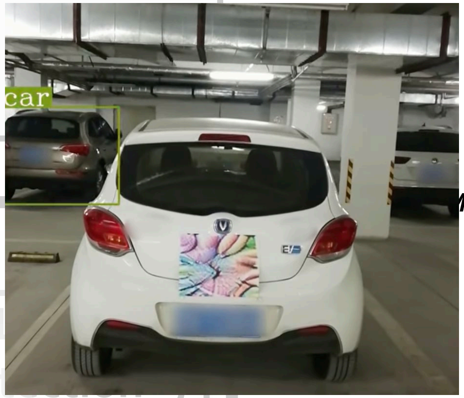

My group's paper

[Zhao et al. @ CCS'19]

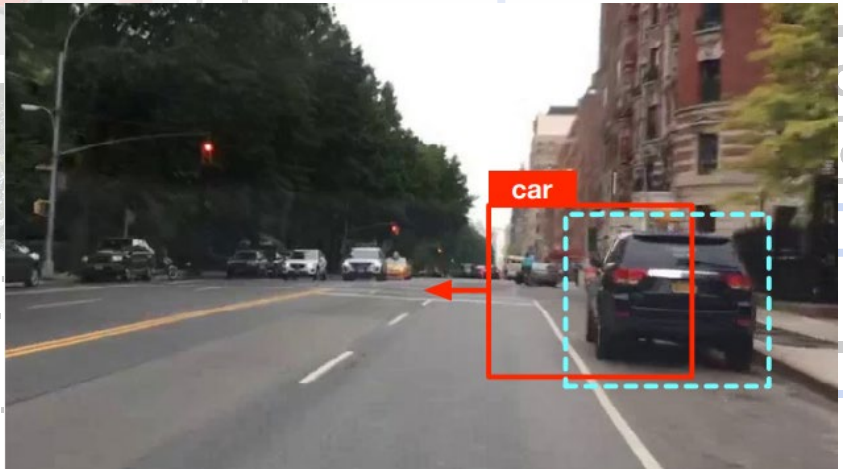

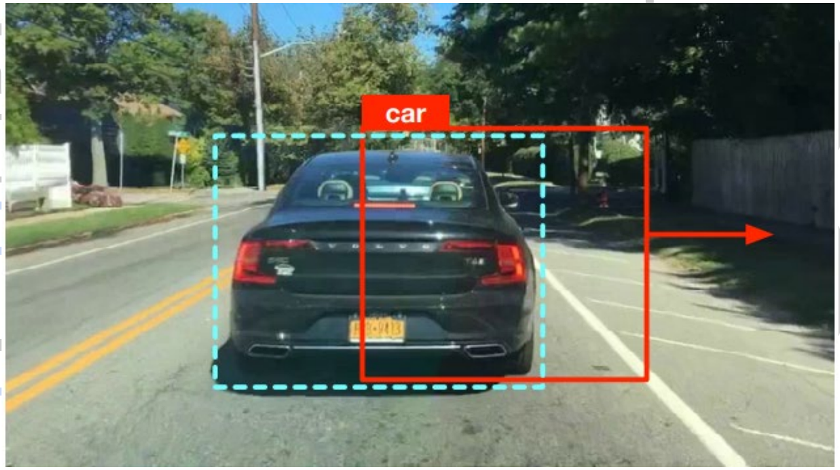

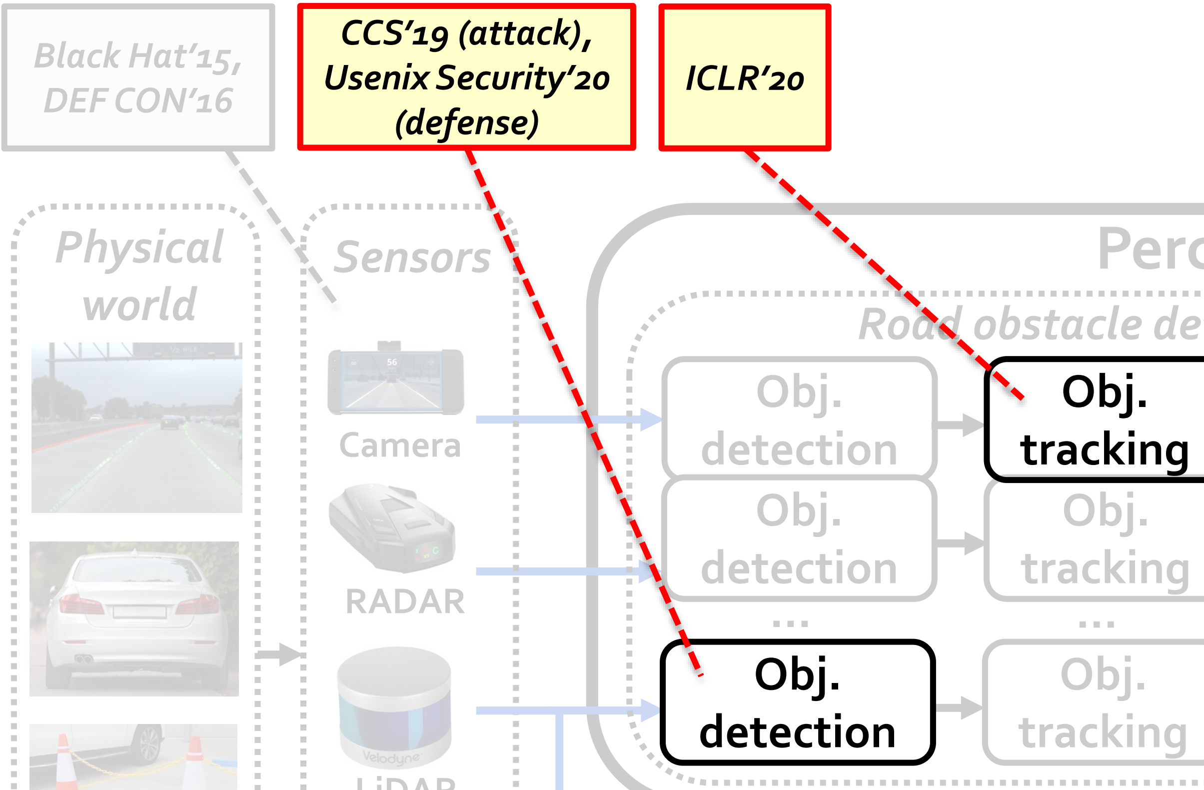

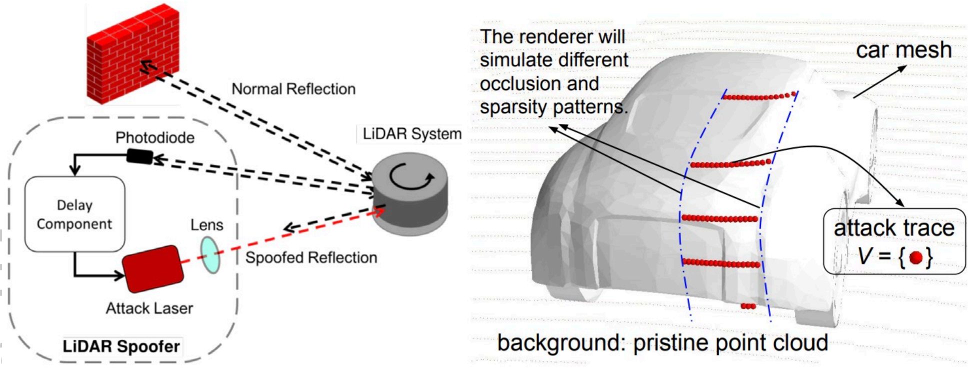

[Cao et al. @ CCS'19, Sun et al. @ USENIX Security'20]

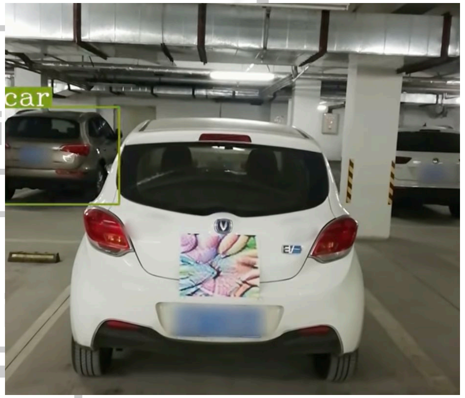

[Zhao et al. @ CCS'19]

[Nassi et al. @ CCS'20]

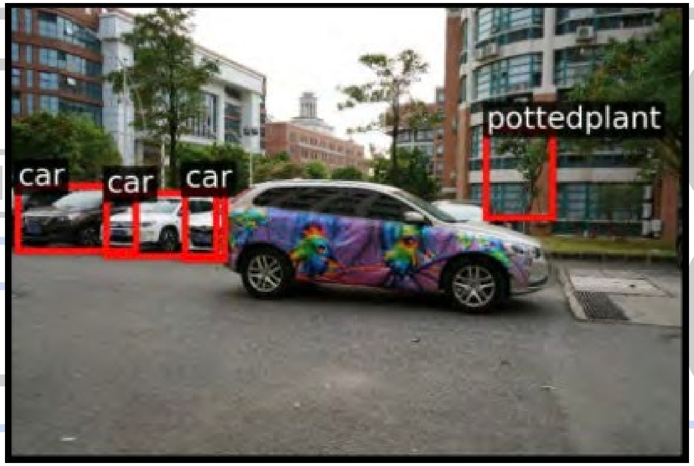

[Huang et al. @ CVPR'20]

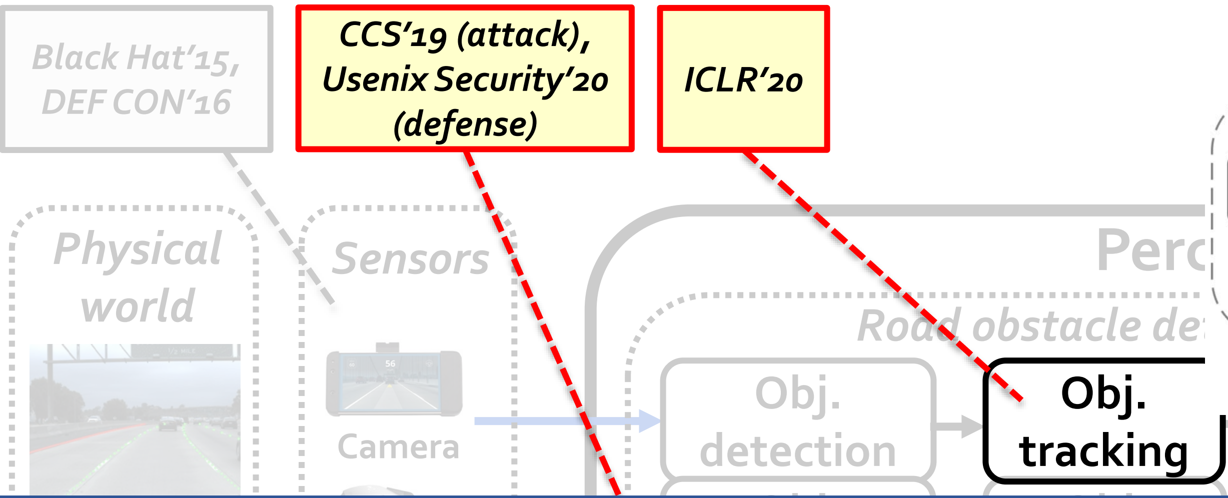

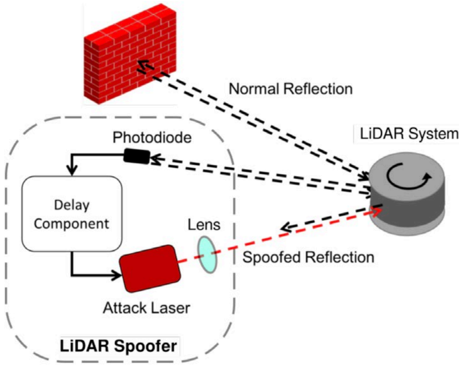

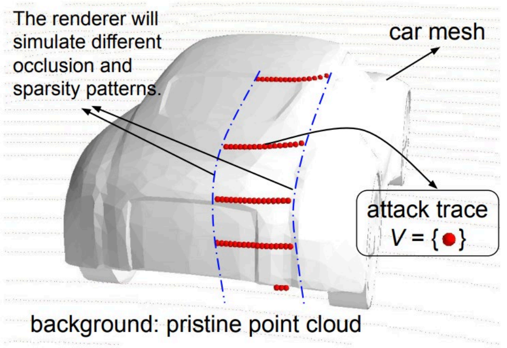

[Cao et al. @ CCS'19, Sun et al. @ USENIX Security'20]

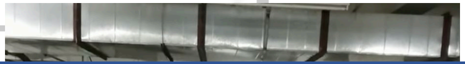

All limited to attacks on a single source of AD perception, e.g., camera- or LiDAR-based AD perception alone!

[Nassi et al. @ CCS'20]

[Huang et al. @ CVPR'20]

[Zhao et al. @ CCS'19]

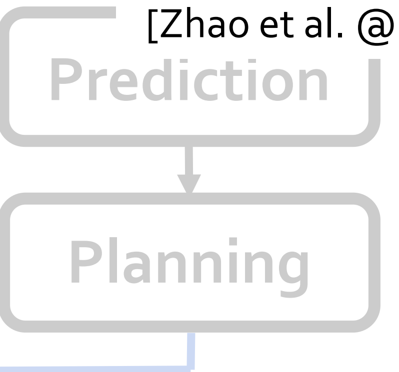

- **Generally adopted by production AD** to achieve overall high robustness & accuracy in practical settings
  - Typically *camera* + *LiDAR*, based on *DNN*
- Assuming not all perception sources are (or can be) attacked simultaneously, should generally be able to ***at least detect single-source attacks***

Basic security design assumption:
Believed to hold in general

Black Hat'15, DEF CON'16 | CCS'19 (attack), Usenix Security'20 (defense) | ICLR'20 | CVPR'18, WOOT'18, ..., CCS'19

My group's paper

Multi-Sensor Fusion (MSF)

Obj. detection → Obj. tracking → MSF

Traffic light det.

Traffic sign det.

LiDAR

DNN-based

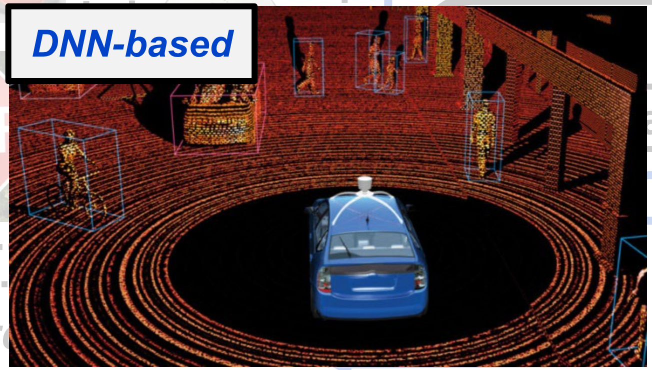

✓ *Depth & shape*

+

Fusion (DNN/rule-based)

DNN-based

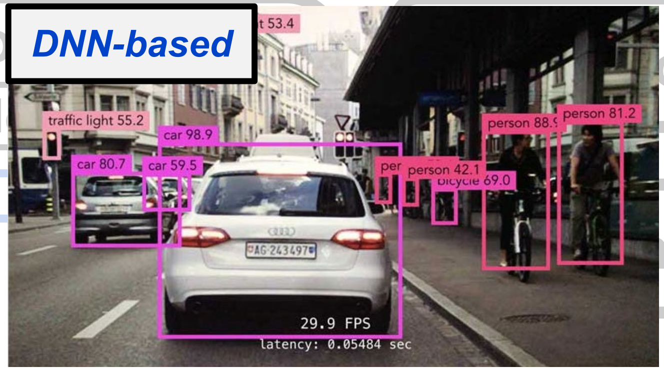

✓ *Texture (e.g., color)*

MSF: Widely recognized as a general defense strategy against existing attacks on AD perception

10.3.2 Sensor-Level Defenses. Several defenses could be adop against spoofing attacks on LiDAR sensors:

Detection techniques. Sensor fusion, which intelligently c bines data from several sensors to detect anomalies and impr performance, could be adopted against LiDAR spoofing attacks systems are often equipped with sensors beyond LiDAR. Came radars, and ultrasonic sensors provide additional information redundancy to detect and handle an attack on LiDAR.

[Cao et al. @ CCS'19]

5.2 Potential Countermeasures

Redundancy and Fusion: If a vehicle is equipped with multiple lidars having an overlapping field of view, the effect of saturating and spoofing can be mitigated to a certain extent. However, this directly increases the cost, and is not a definitive solution because attackers car Besides, it is also not easy to detect spo non-overlapped zones. Likewise, the fusic be an ultimate solution either. Radars [44 sors [44] have all been revealed to be vul spoofing.

[Shin et al. @ CHES'17]

In this work, we do not assume any particular sensing or actuation workflow to be trusted. However, we do assume that not all sensor readings can be corrupted simultaneously. Under the design where workflows run with isolation (see Section II-A), attacks or failures in a workflow can be constrained within. Admittedly, such cases could be possible in carefully crafted attacks. However, it is difficult for attackers. Firstly, for heterogeneous sensors, holding a vulnerability and a corresponding exploit which targets one sensing workflow is already costly [6], [9], not to mention corrupting all. Secondly, even if an attacker is capable of corrupting all sensors, the nch the attacks simultaneously to avoid at challenge to launch such coordinated arget sensing workflows [9].

[Guo et al. @ DSN'18]

As the system's autonomy increases, so does the concern about its security. In modern vehicles, a malicious attacker may deceive the controller into performing a dangerous action by altering the measurements of some sensors [1], [2]. D on the attacker's goal and capabilities, the consequer range from minor disturbances in performance to cra loss of human lives. Consequently, performing attack sensor fusion is essential for the safety of such syste

[Ivanov et al. @ DATE'14]

2.1 System Model and Current Approach

We consider a system with n sensors measuring the same physical variable. As mentioned above, we assume *abstract* sensors; therefore, each sensor provides the controller with an interval of all possible values. We assume the system queries all the sensors periodically such that a centralized estimator receives measurements from all sensors, and then performs attack detection/identification and *sensor fusion* (SF). We now explain the current approach to attack detection, referred to herein as a SF-based detector, before providing the improved version addressed in this paper.

[Park et al. @ ICCPS'15]

Black Hat'15, DEF CON'16

CCS'19 (attack), Usenix Security'20 (defense)

ICLR'20

CVPR'18, WOOT'18, …, CCS'19

My group's paper

- **Generally adopted by production AD** to achieve overall high robustness & accuracy in practical settings
  - Typically *camera* + *LiDAR*, based on *DNN*
- Assuming not all perception sources are (or can be) attacked simultaneously, should generally be able to ***at least detect single-source attacks***

Basic security design assumption:
Believed to hold in general

Obj. detection → Obj. tracking → Multi-Sensor Fusion (MSF)

Traffic light det.

Traffic sign det.

Research Question:

Can such basic security design assumption actually be broken, especially in practical AD settings?

✓ Depth & shape

✓ Texture (e.g., color)

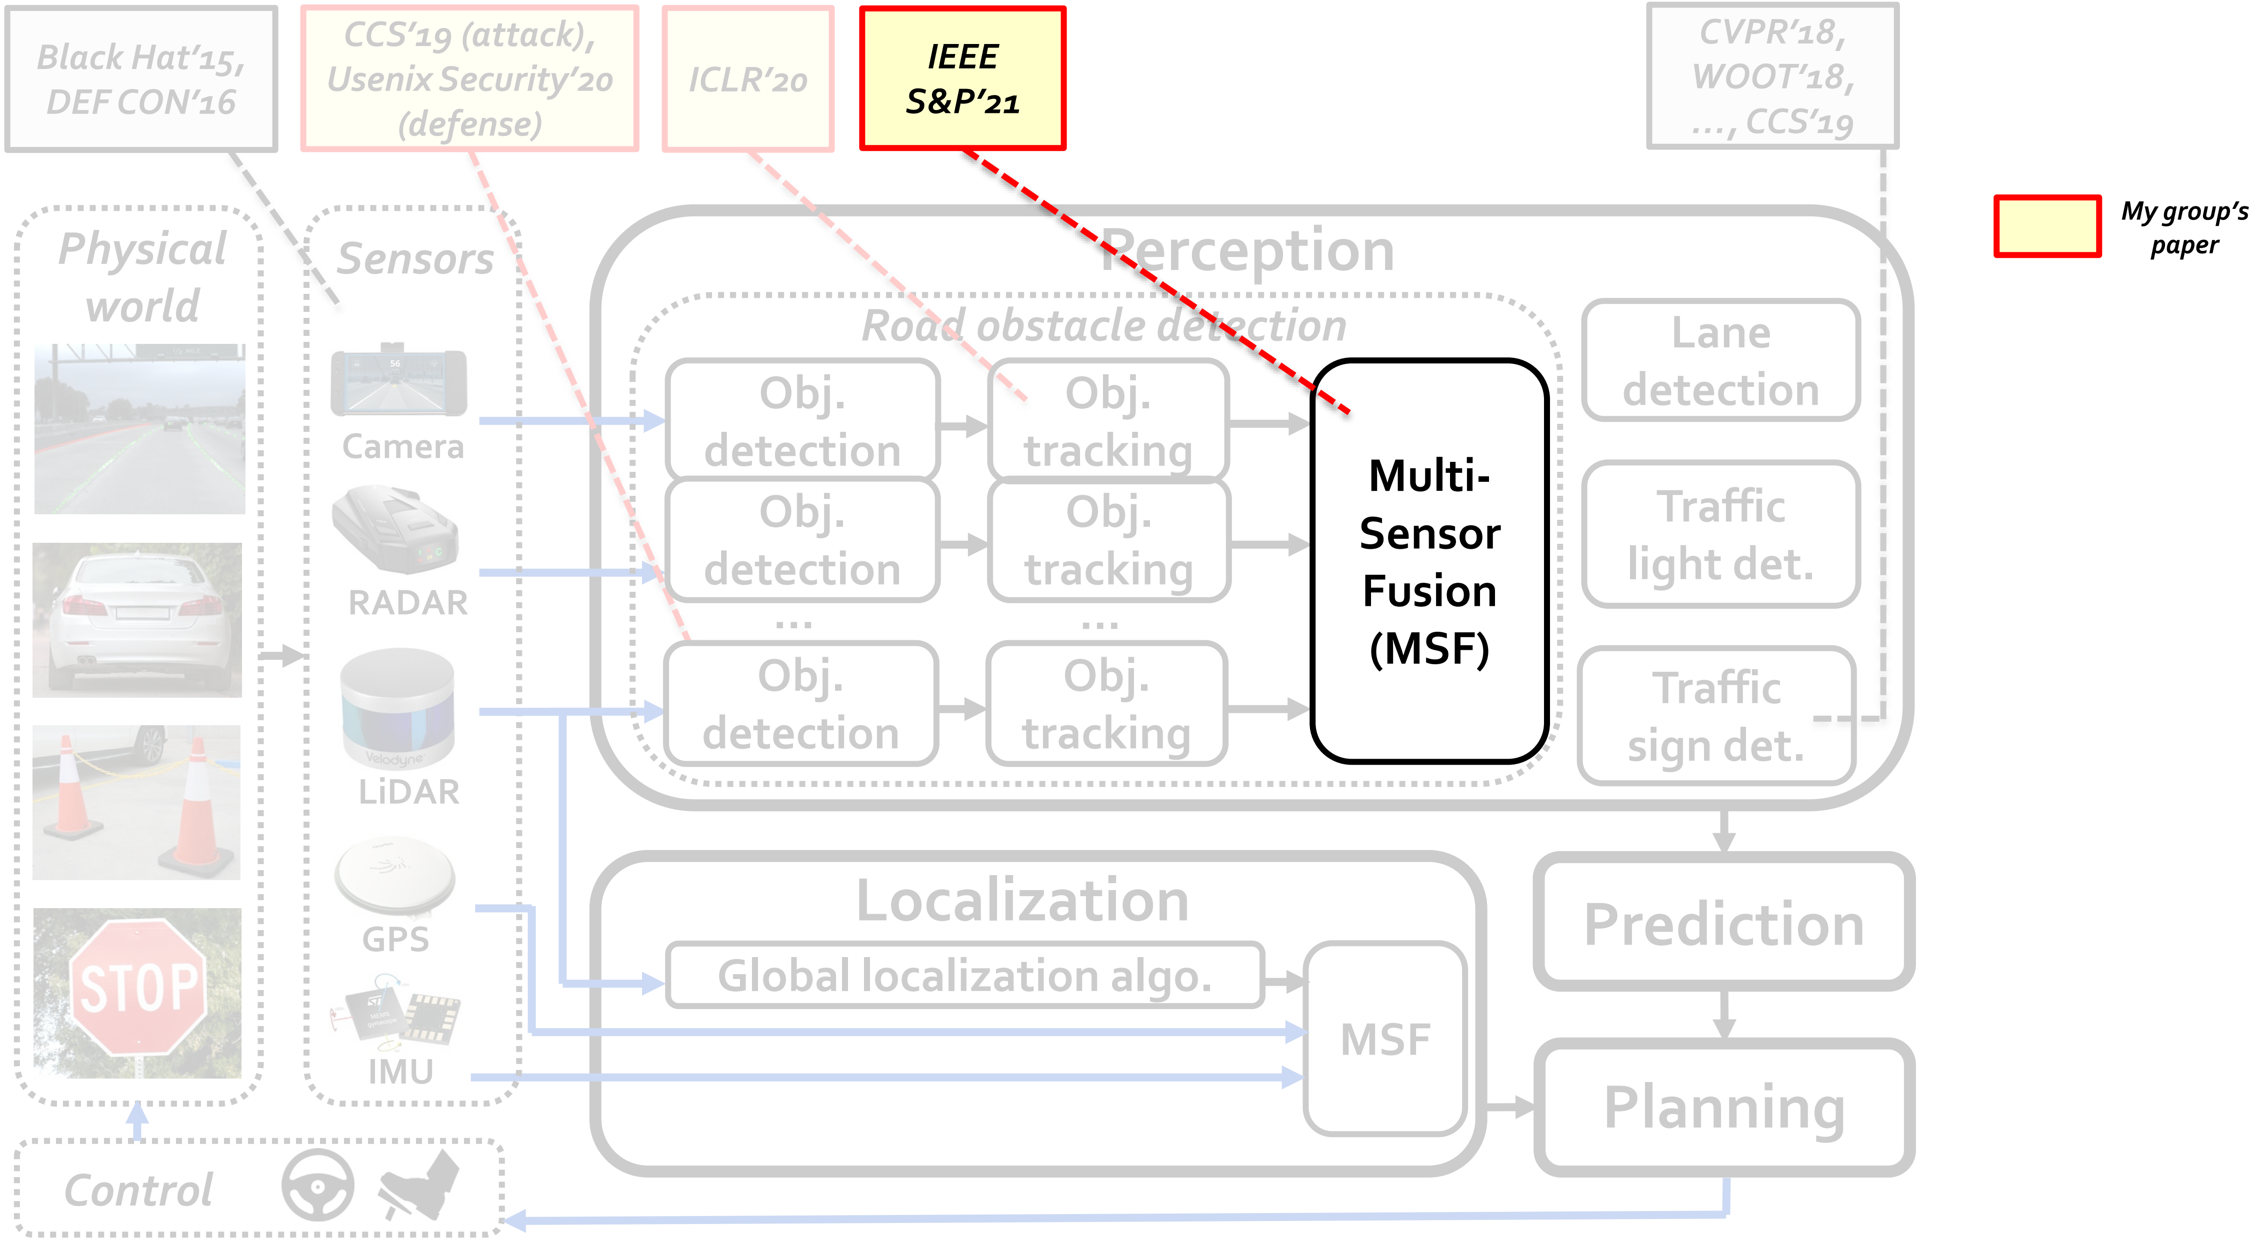
Black Hat'15, DEF CON'16
CCS'19 (attack), Usenix Security'20 (defense)
ICLR'20
IEEE S&P'21
CVPR'18, WOOT'18, …, CCS'19
My group's paper
Physical world
Sensors
Camera
RADAR
LiDAR
GPS
IMU
Perception
Road obstacle detection
Obj. detection
Obj. tracking
Multi-Sensor Fusion (MSF)
Lane detection
Traffic light det.
Traffic sign det.
Localization
Global localization algo.
MSF
Prediction
Planning
Control

CVPR'18, WOOT'18, …, CCS'19

My group's paper

- ***First*** study on security of MSF perception
- Directly challenge security design assumption: explore possibility of ***effectively*** & ***simultaneously*** attacking ***all* perception sources**
- <u>New attack vector</u>: Maliciously-shaped adversarial 3D object (e.g., traffic cone or rock) → can influence both camera pixels & LiDAR point cloud
  - Fool victim to fail in detecting front obstacle, thus **crash into it**
  - ***Physically-realizable*** (via 3D printing) & ***stealthy*** (by mimicking)

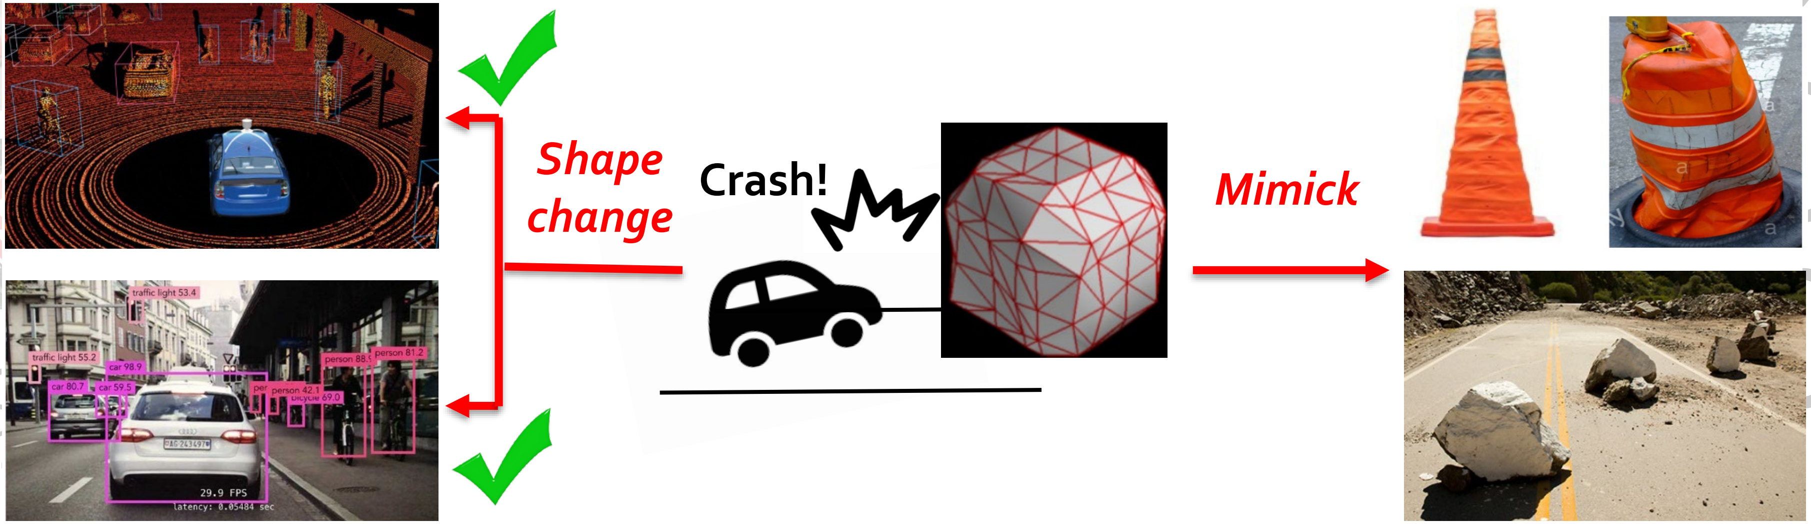

- ***First*** study on security of MSF perception
- Directly challenge security design assumption:
  & ***simultaneously*** attacking ***all* perception sou**
- New attack vector: Maliciously-shaped adversa
  or rock) → can influence both camera pixels & L
  - Fool victim to fail in detecting front obstacl
  - ***Physically-realizable*** (via 3D printing) & ***stealthy*** (by mimicking)
- New methodology: Customized differentiable rendering & new differentiable approx func designs for pre-processing (esp. cell-level aggregated feature calc)
  - **<10% → 100%** in attack success rate

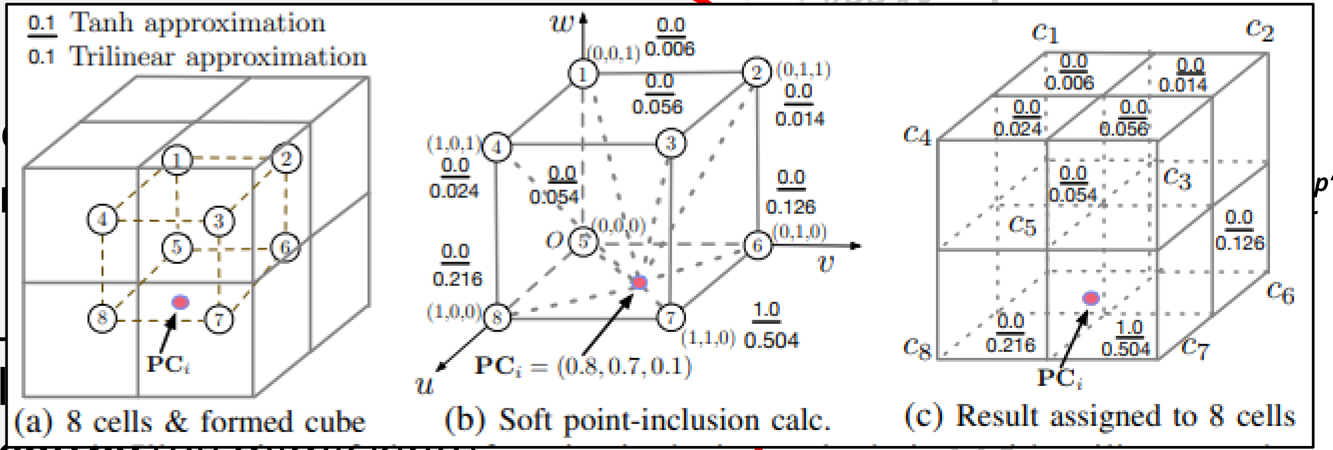

Addressing *System-to-AI semantic gap*: from attack perturbation capability at CPS system input space (i.e., 3D object shape changes) to that at AI component input space (i.e., camera pixel & LiDAR point cloud changes)

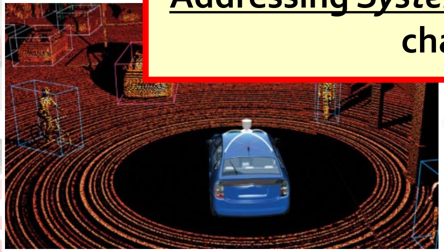

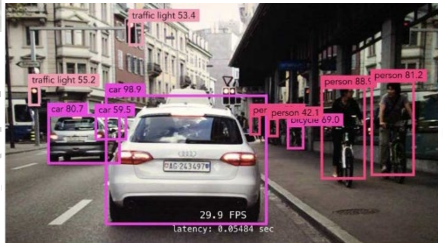

Shape change

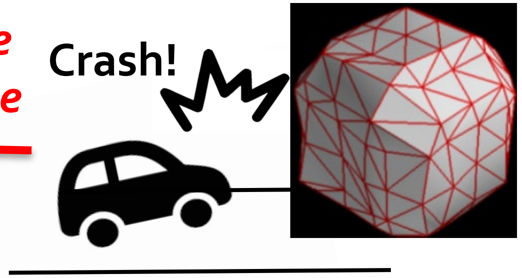

Mimick

Attack demos: Benign case

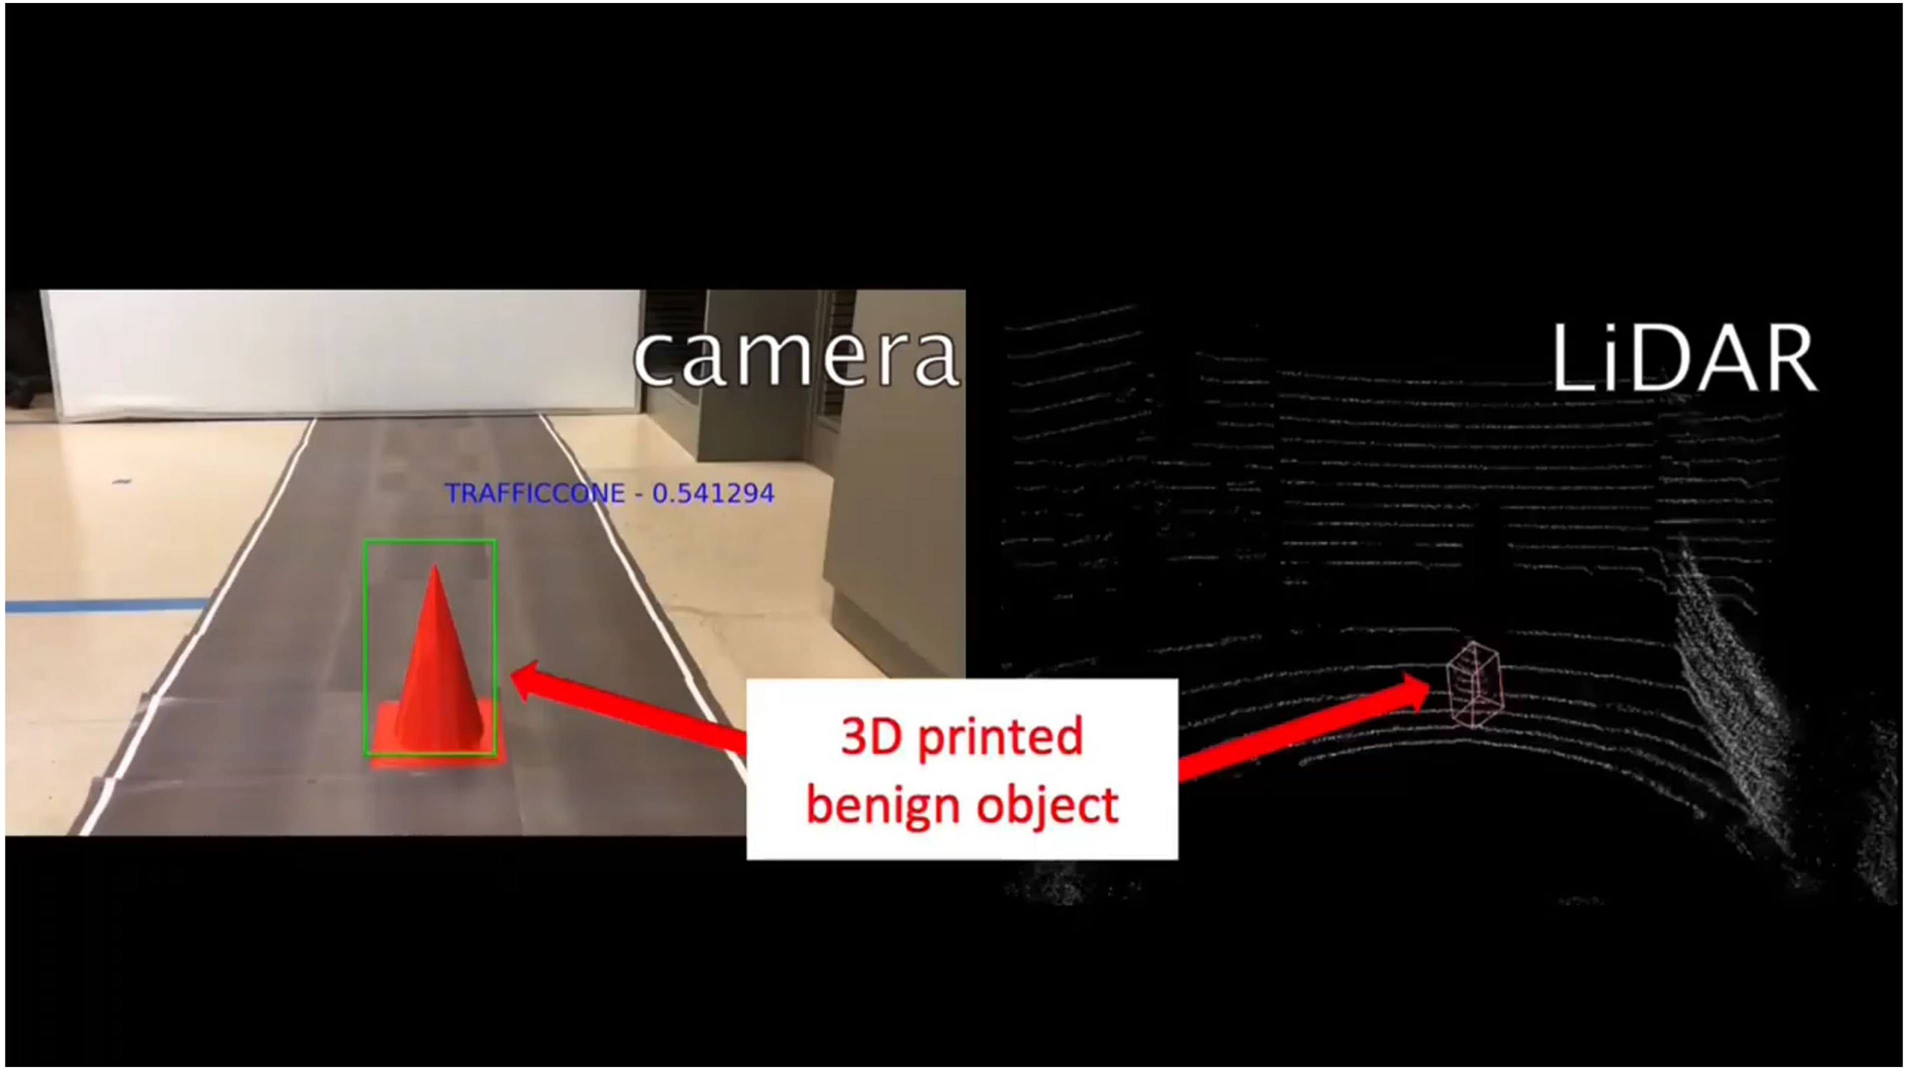

Attack demos: Adversarial case

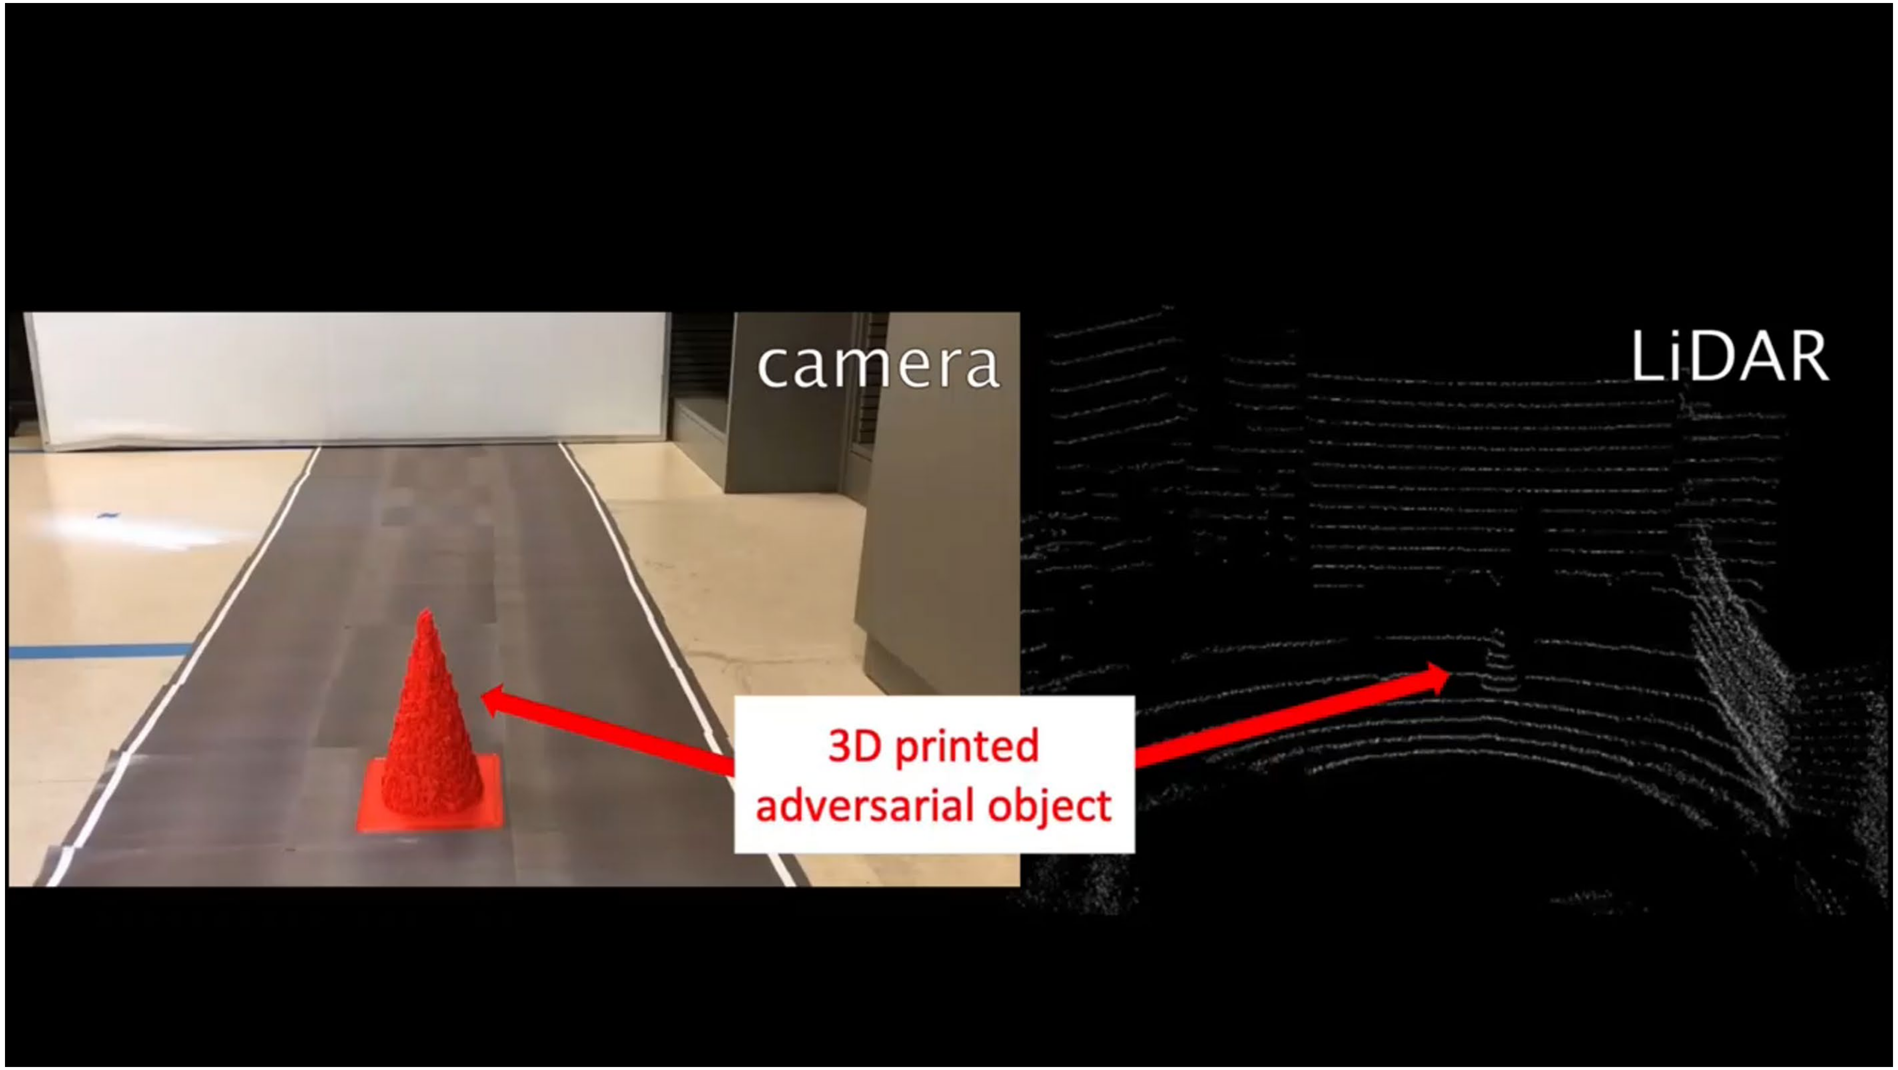

Attack demos

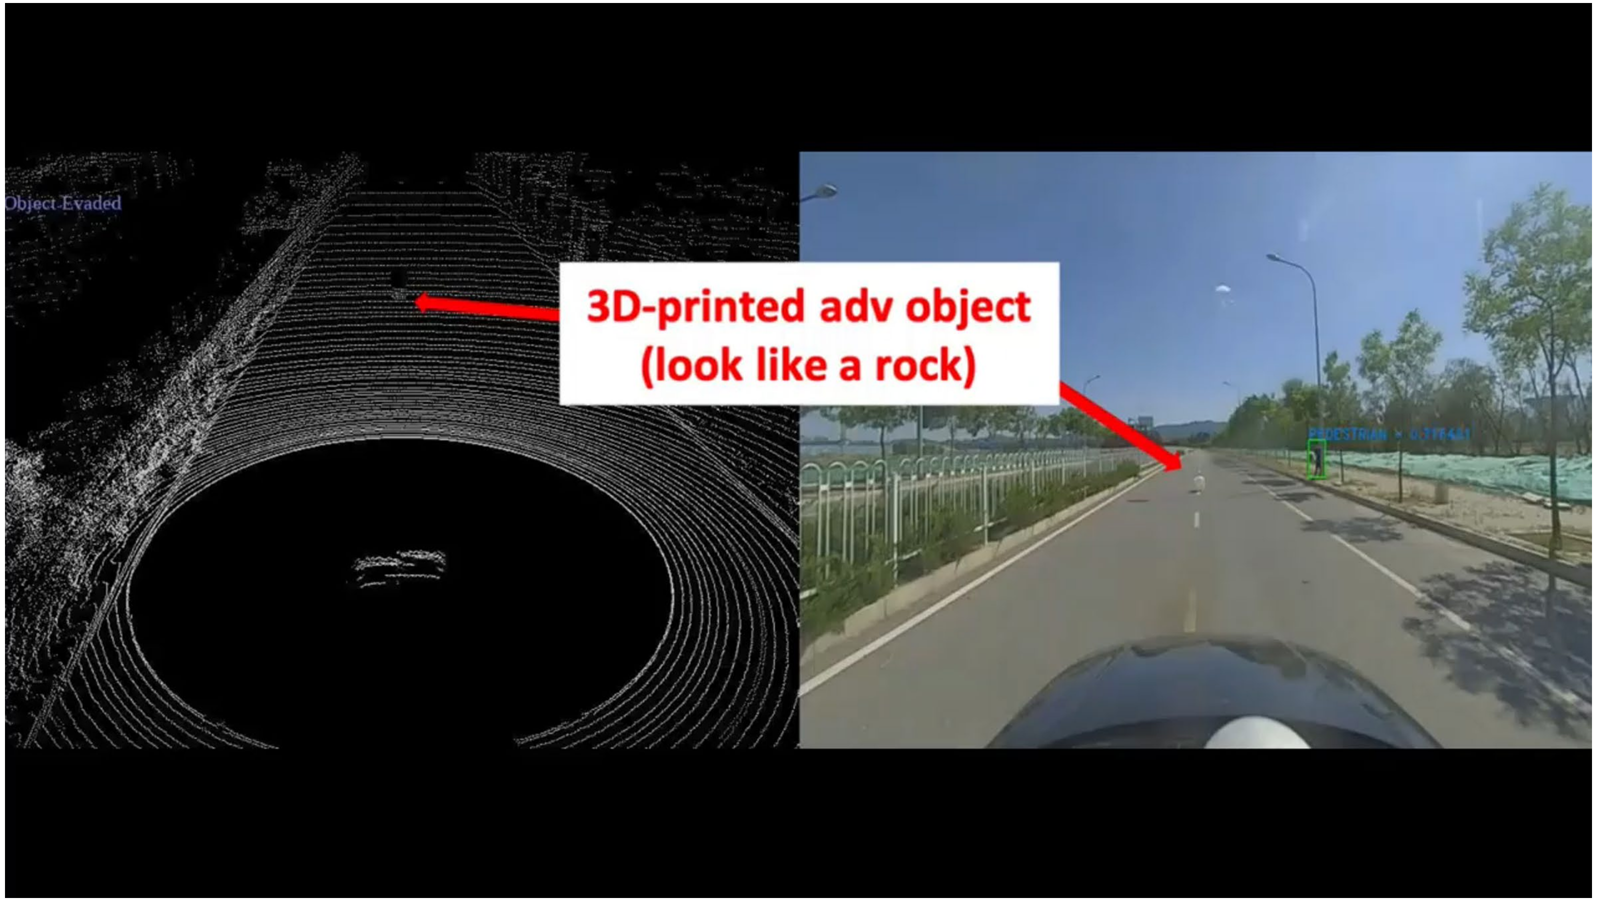

Demo website: https://sites.google.com/view/cav-sec/msf-adv

Black Hat'15, DEF CON'16

CCS'19 (attack), Usenix Security'20 (defense)

ICLR'20

IEEE S&P'21

NDSS'20 Best Poster, Usenix Security'21

CVPR'18, WOOT'18, …, CCS'19

My group's paper

- **One of the first** to study production lane detection DNN
- Attack vector: Malicious dirty road patterns

Lane detection

Traffic

Black Hat'15, DEF CON'16

CCS'19 (attack), Usenix Security'20 (defense)

ICLR'20

IEEE S&P'21

NDSS'20 Best Poster, Usenix Security'21

CVPR'18, WOOT'18, …, CCS'19

My group's paper

- **One of the first** to study production lane detection DNN
- Attack vector: Malicious dirty road patterns
- Method: Optimization-based method

Lane detection

Traffic light det.

Traffic

Black Hat'15, DEF CON'16

CCS'19 (attack), Usenix Security'20 (defense)

ICLR'20

IEEE S&P'21

NDSS'20 Best Poster, Usenix Security'21

CVPR'18, WOOT'18, …, CCS'19

My group's paper

Lane detection

Traffic light det.

Traffic

- **One of the first** to study production lane detection DNN
- Attack vector: Malicious dirty road patterns
- Method: Optimization-based method
- Impact: Cause a victim to ***drive out of the current lane boundaries within 1 sec***
  - *Far below normal driver reaction time (~2.5 sec)*

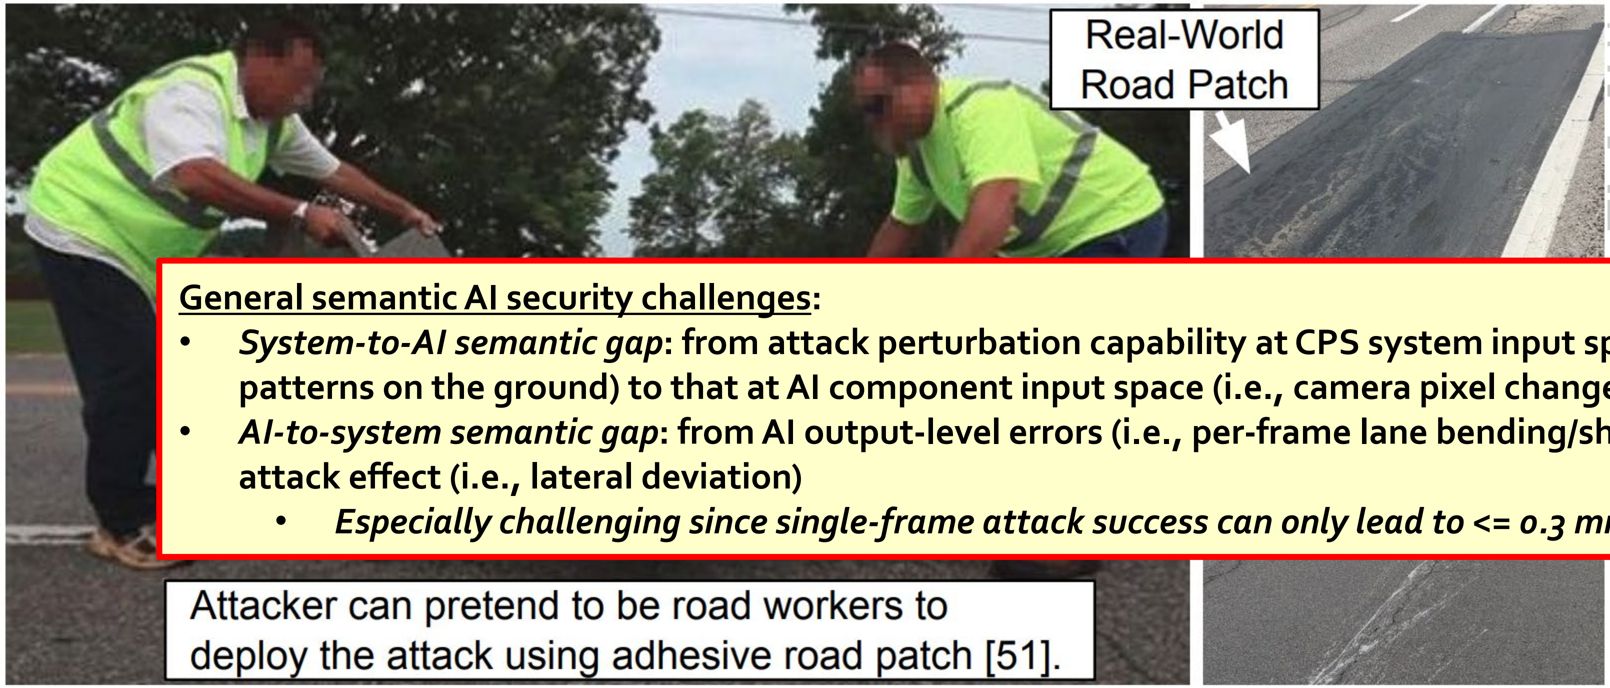

General semantic AI security challenges:
- *System-to-AI semantic gap*: from attack perturbation capability at CPS system input space (i.e., malicious dirty patterns on the ground) to that at AI component input space (i.e., camera pixel changes)
- *AI-to-system semantic gap*: from AI output-level errors (i.e., per-frame lane bending/shifting) to CPS system-level attack effect (i.e., lateral deviation)
  - *Especially challenging since single-frame attack success can only lead to <= 0.3 mm lateral dev. at 45 mph*

Demo: Dirty road patch attack on lane detection

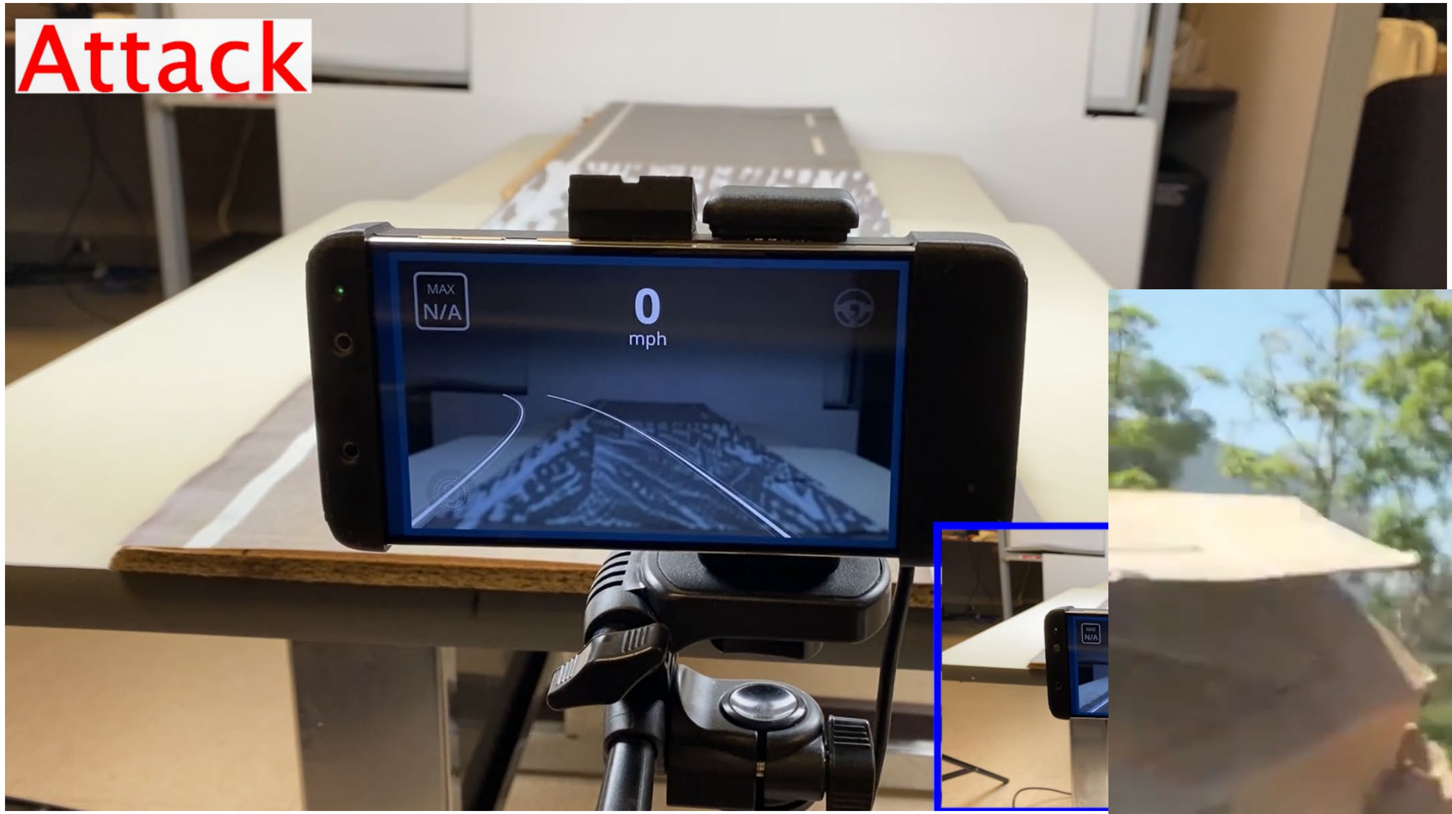

100% (10/10) crash rate for real vehicle w/ AEB

Demo website: https://sites.google.com/view/cav-sec/drp-attack/

Black Hat'15, DEF CON'16

CCS'19 (attack), Usenix Security'20 (defense)

ICLR'20

IEEE S&P'21

AutoSec 2021

NDSS'20 Best Poster, Usenix Security'21

CVPR'18, WOOT'18, …, CCS'19

My group's paper

Physical world

Control

- First security analysis of traffic light detection
- Target: Industry-grade AD traffic light detection pipeline
  - Specifically, the use of **ROI (Region-of-Interest)** to narrow down detection scope in raw camera input
- Attack vector: GPS spoofing
- Impact: Move right traffic light out of ROI, causing **DoS;** or move wrong traffic light into ROI, causing **red light running**
- Demo website: https://sites.google.com/view/roiattack

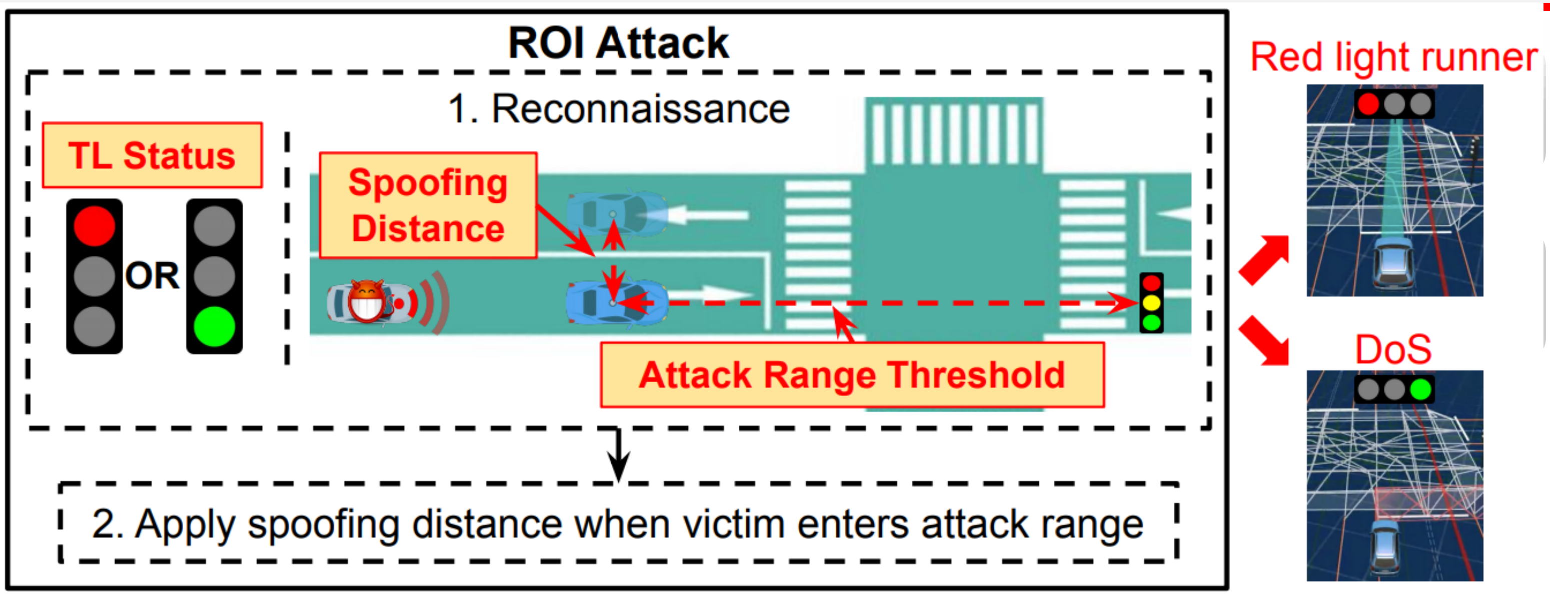

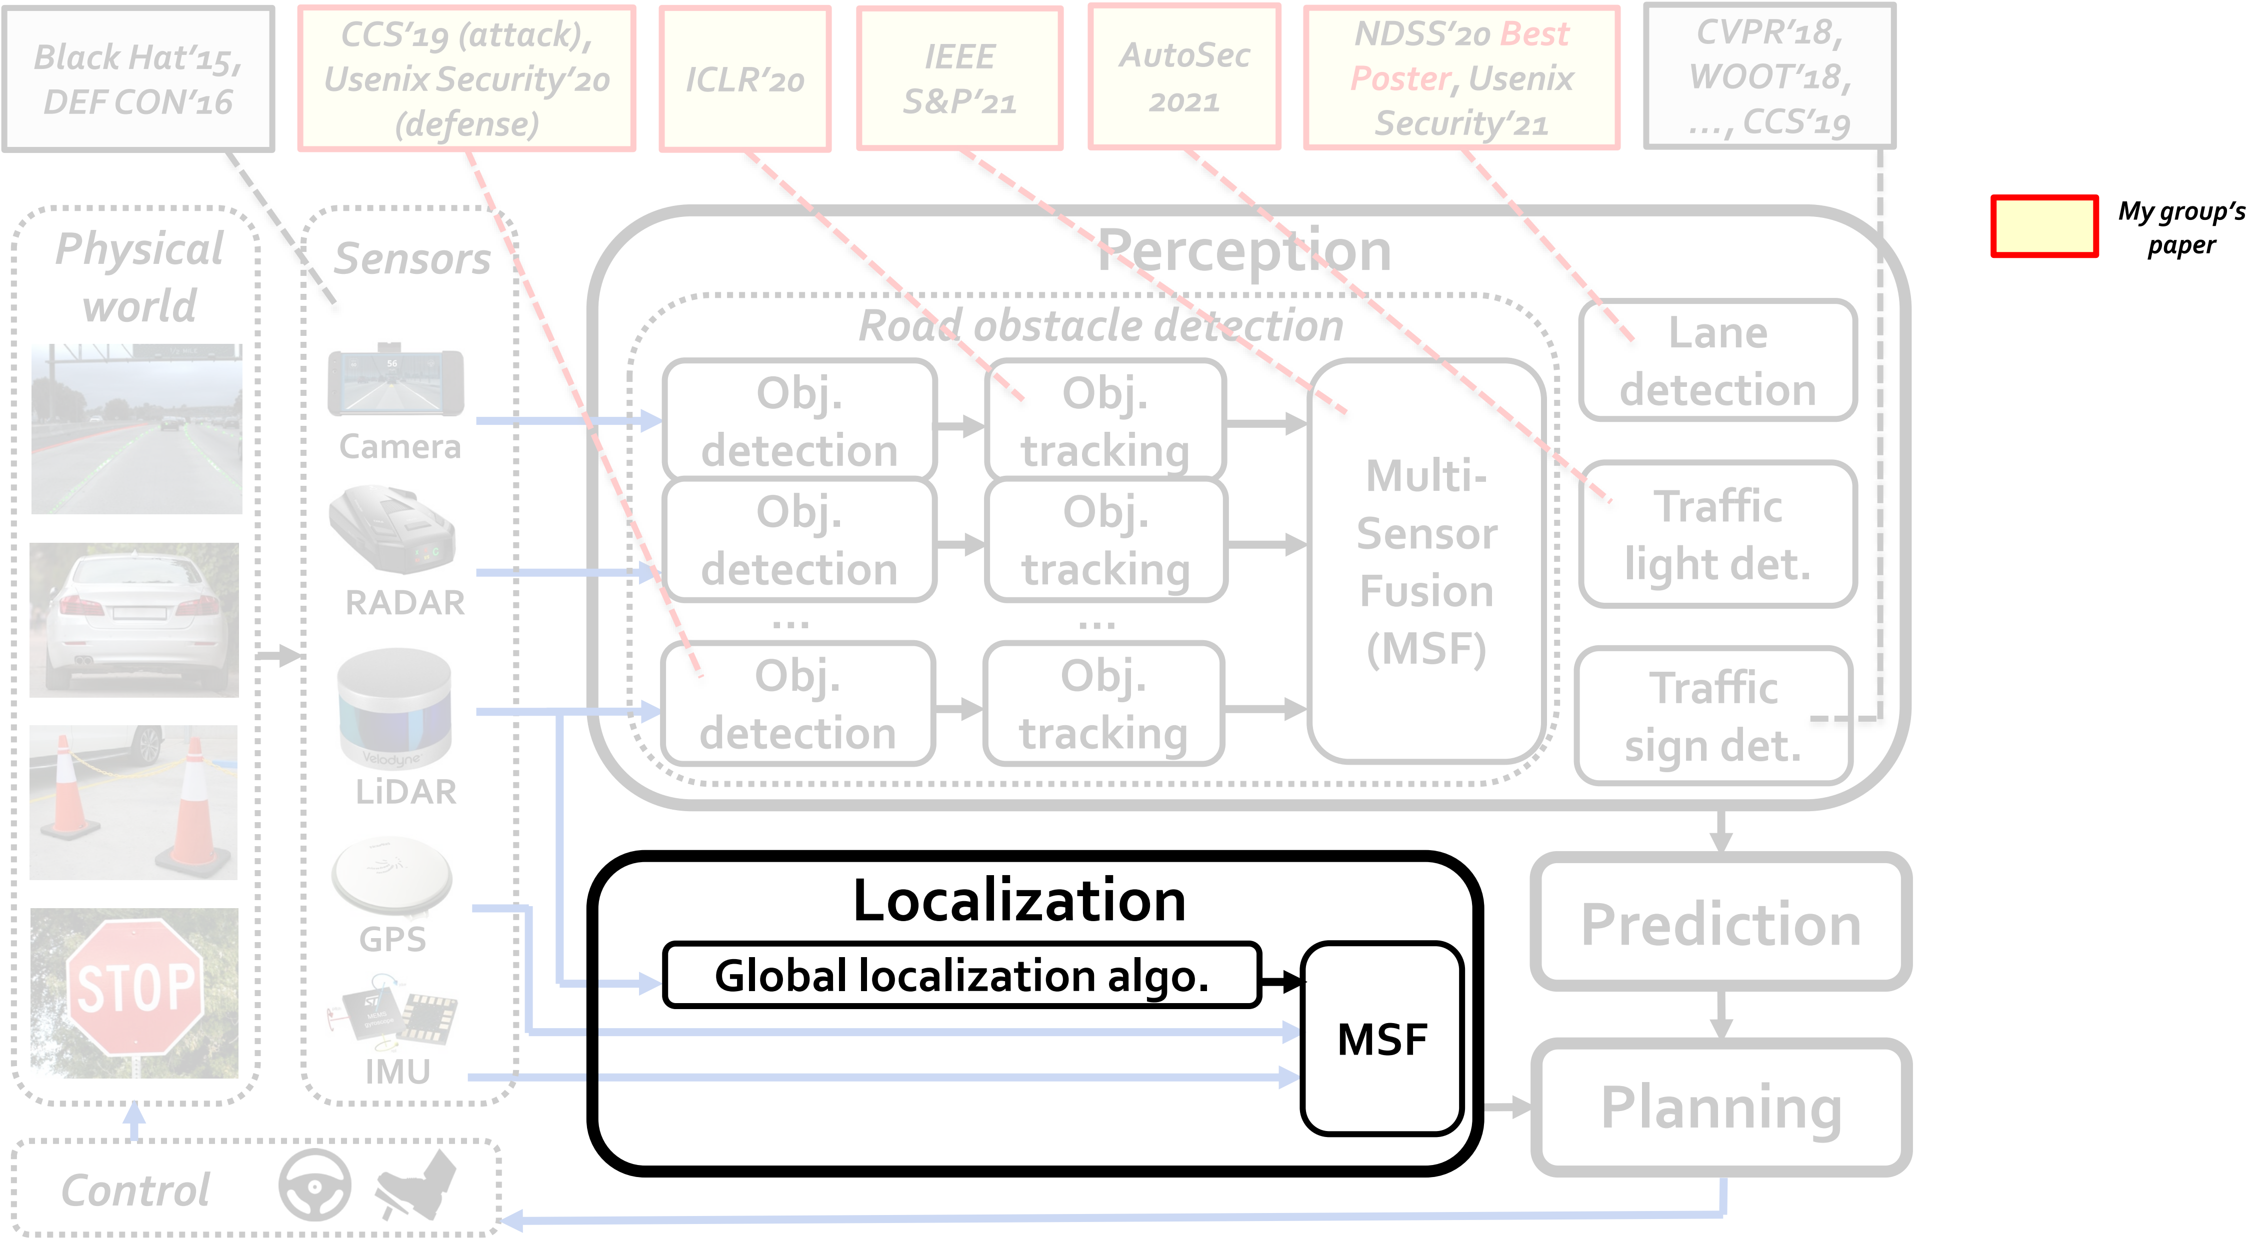
Black Hat'15, DEF CON'16
CCS'19 (attack), Usenix Security'20 (defense)
ICLR'20
IEEE S&P'21
AutoSec 2021
NDSS'20 Best Poster, Usenix Security'21
CVPR'18, WOOT'18, ..., CCS'19
My group's paper
Physical world
STOP
Sensors
Camera
RADAR
LiDAR
GPS
IMU
Perception
Road obstacle detection
Obj. detection
Obj. tracking
Obj. detection
Obj. tracking
Obj. detection
Obj. tracking
Multi-Sensor Fusion (MSF)
Lane detection
Traffic light det.
Traffic sign det.
Localization
Global localization algo.
MSF
Prediction
Planning
Control

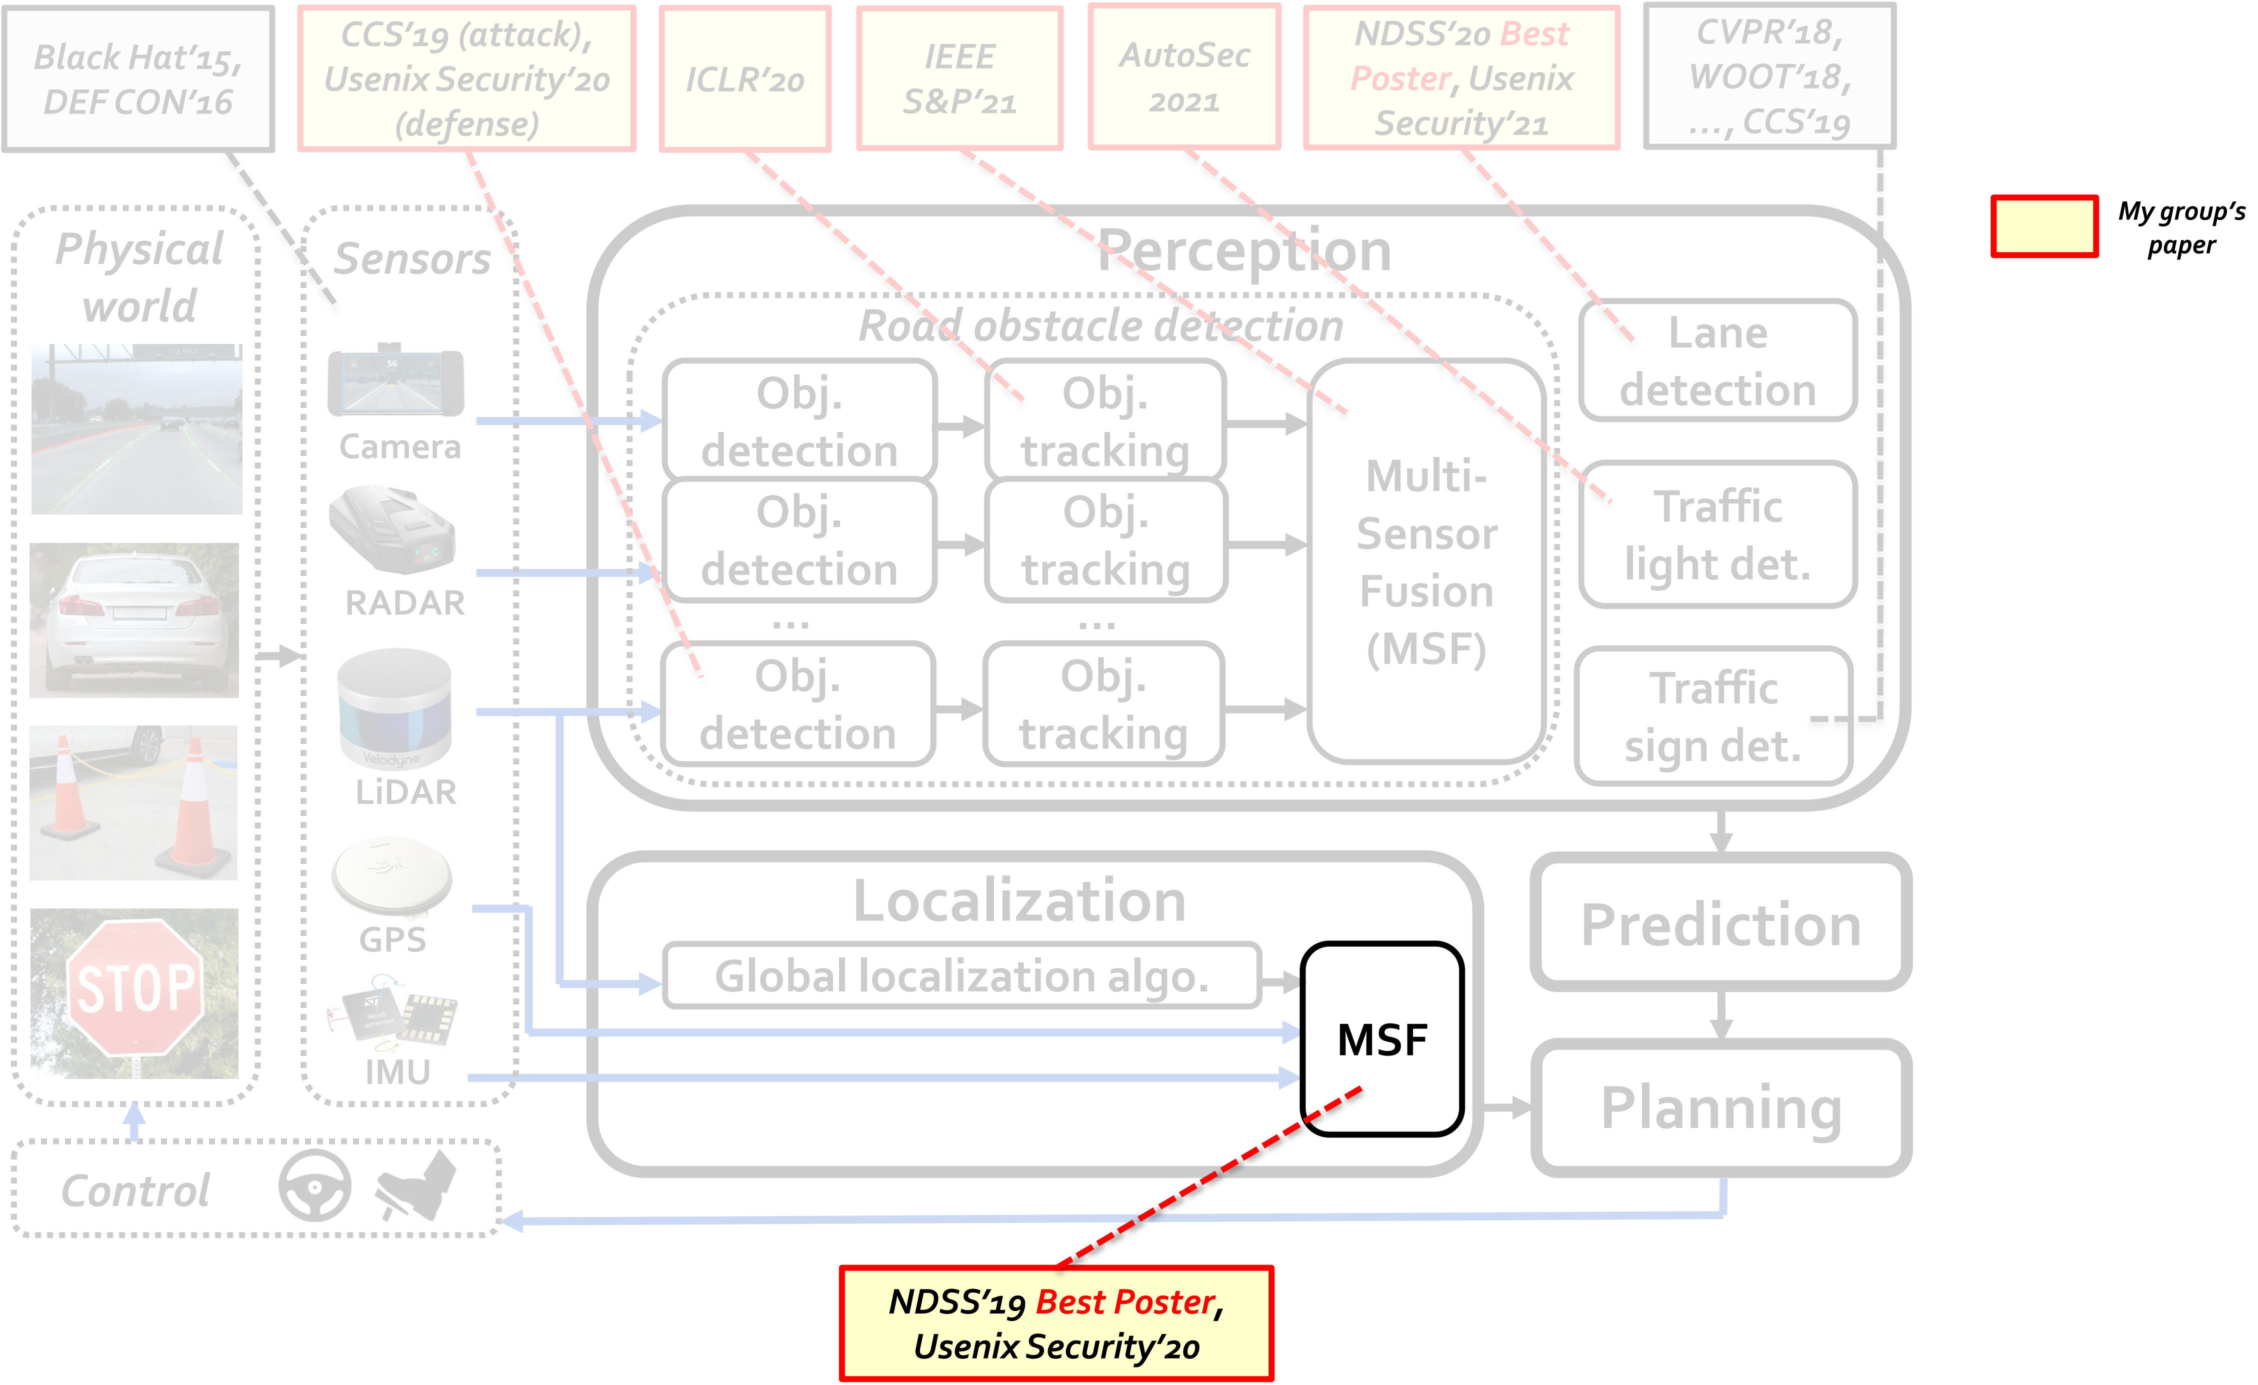

Black Hat'15, DEF CON'16
CCS'19 (attack), Usenix Security'20 (defense)
ICLR'20
IEEE S&P'21
AutoSec 2021
NDSS'20 Best Poster, Usenix Security'21
CVPR'18, WOOT'18, …, CCS'19
My group's paper
Physical world
Sensors
Camera
RADAR
LiDAR
GPS
IMU
Perception
Road obstacle detection
Obj. detection
Obj. tracking
Multi-Sensor Fusion (MSF)
Lane detection
Traffic light det.
Traffic sign det.
Localization
Global localization algo.
MSF
Prediction
Planning
Control
NDSS'19 Best Poster, Usenix Security'20

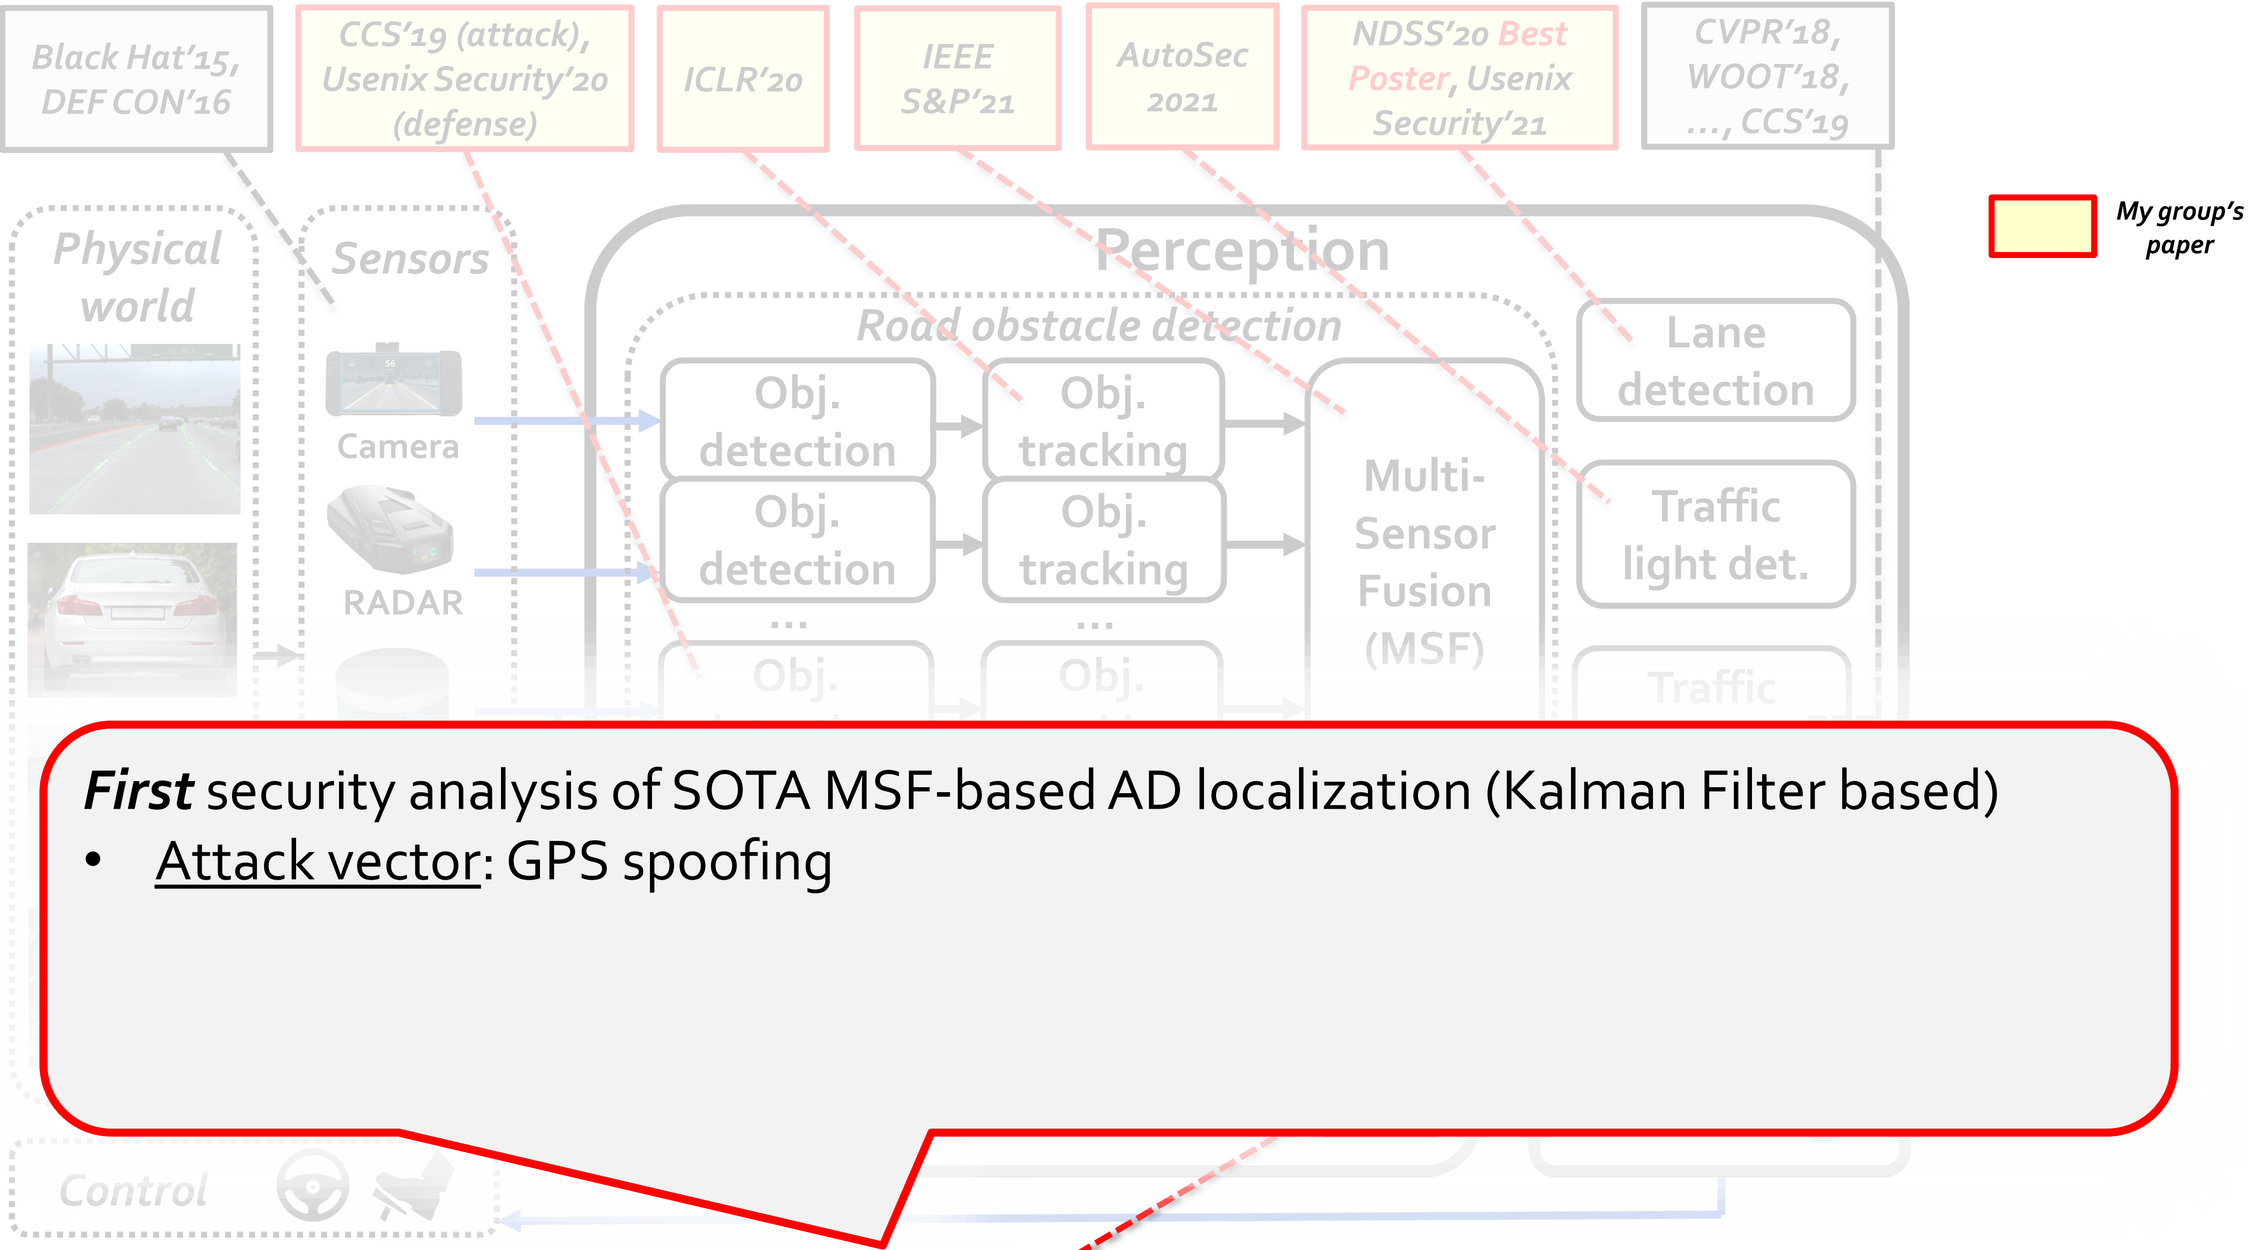

First security analysis of SOTA MSF-based AD localization (Kalman Filter based)

- Attack vector: GPS spoofing

NDSS'19 Best Poster, Usenix Security'20

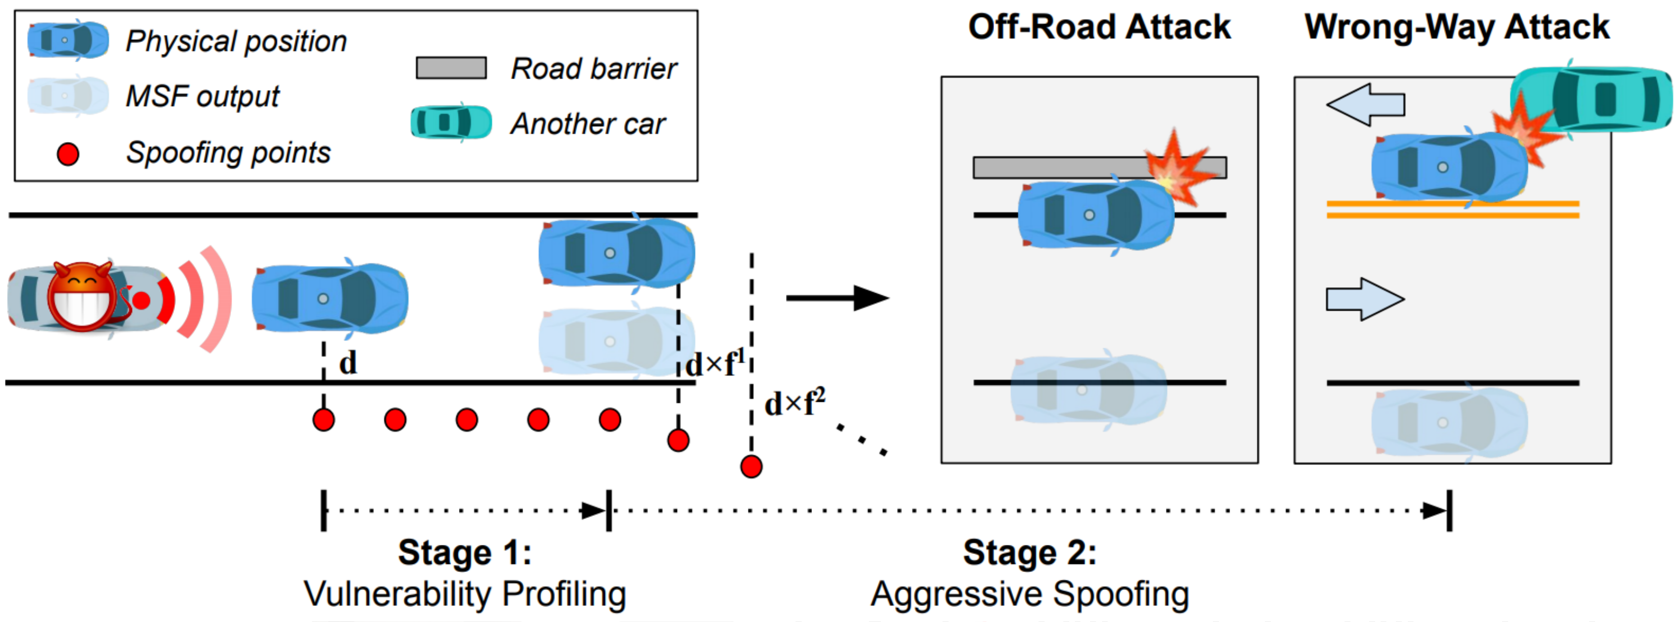

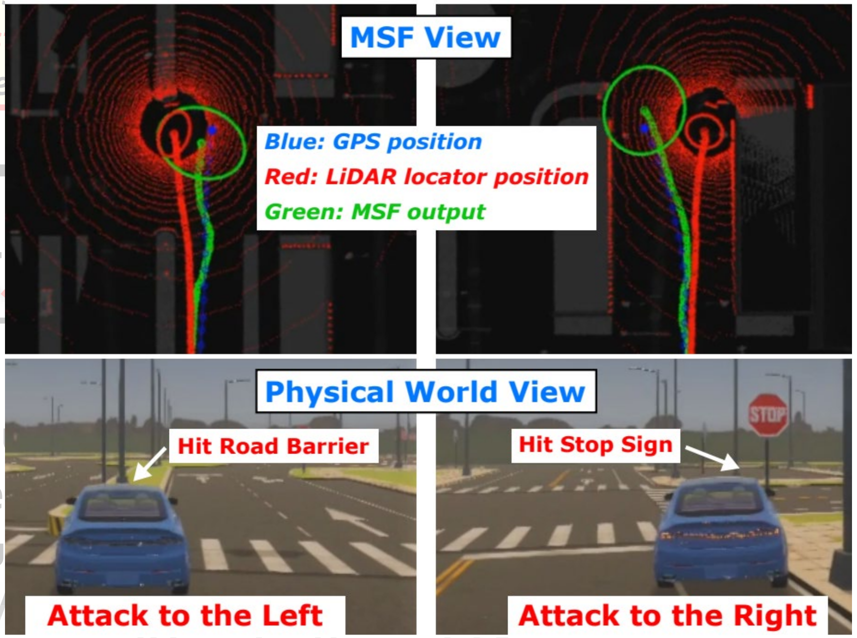

First security analysis of SOTA MSF-based AD localization (Kalman Filter based)

- Attack vector: GPS spoofing
- Impact: If tailgate for **2 min**, ***almost always (97% chance)*** can find an opportunity to break sensor fusion, and cause a victim to ***drive off road or to the wrong way***
- Demo website: https://sites.google.com/view/cav-sec/fusionripper

NDSS'19 Best Poster,
Usenix Security'20

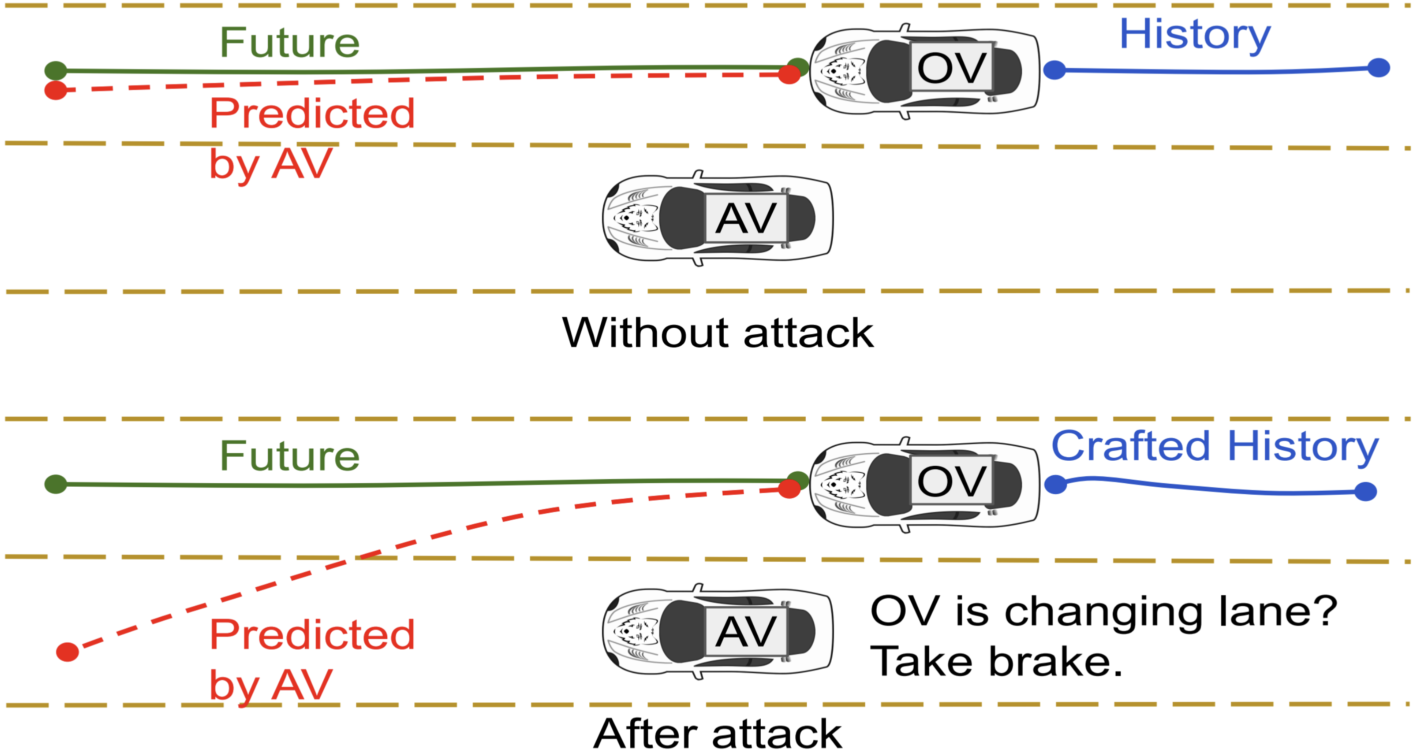

First security analysis of AD trajectory prediction DNN

- Attack vector: Malicious driving trajectory
- Impact: Spoof a ***cut-in*** trajectory prediction, causing **emergency brake**; or spoof a ***drive-away*** prediction, causing **crash**

CVPR'22

My group's paper

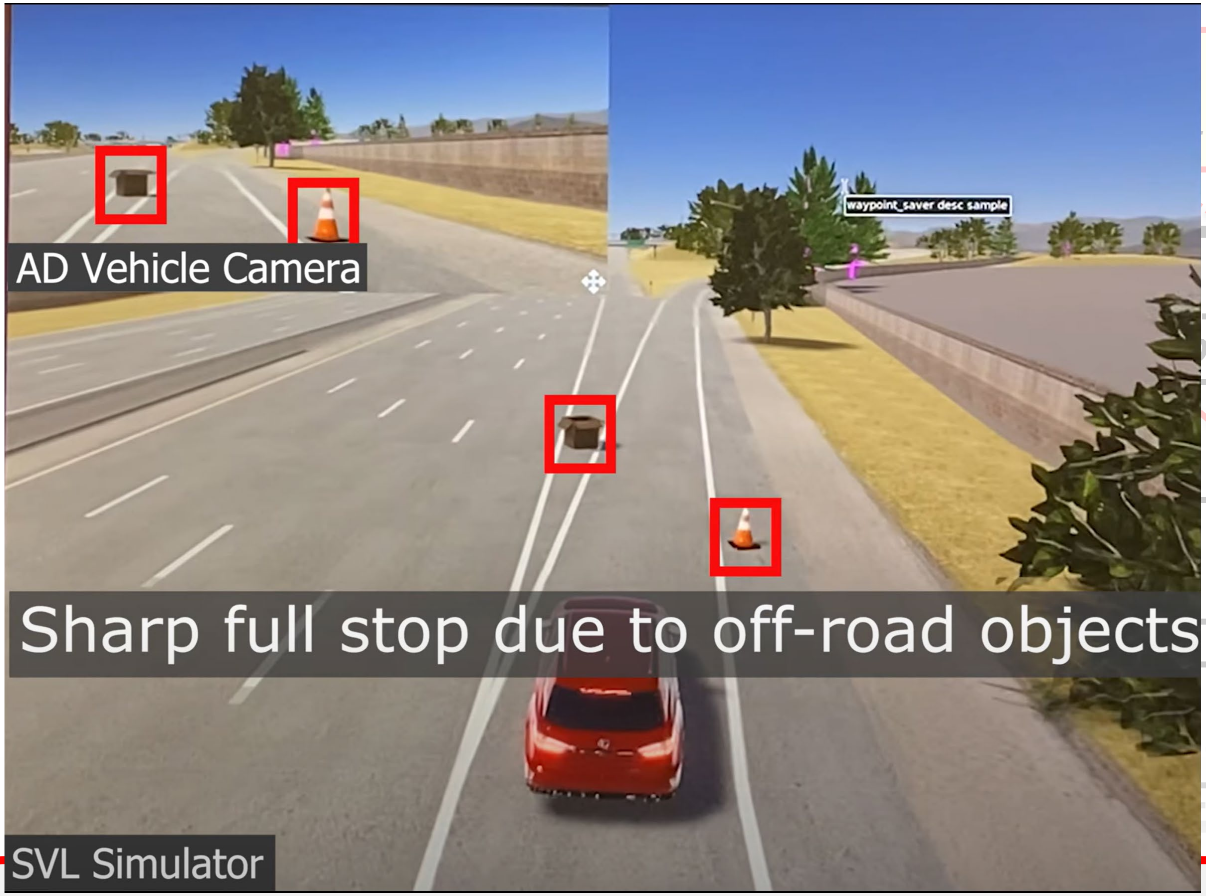

AutoSec 2021

NDSS'20 Best Poster, Usenix

CVPR'18, WOOT'18,

My group's paper

First security analysis of AD behavior planning (*program-based*)

- Attack vector: Common road objects (e.g., road-side cardboard boxes, parked bikes, etc.)
- Methodology: Domain-customized evolutionary testing
- Impact: Unnecessary sharp braking, stopping, giving up mission-critical driving decisions, etc.

Prediction

Planning

NDSS'22

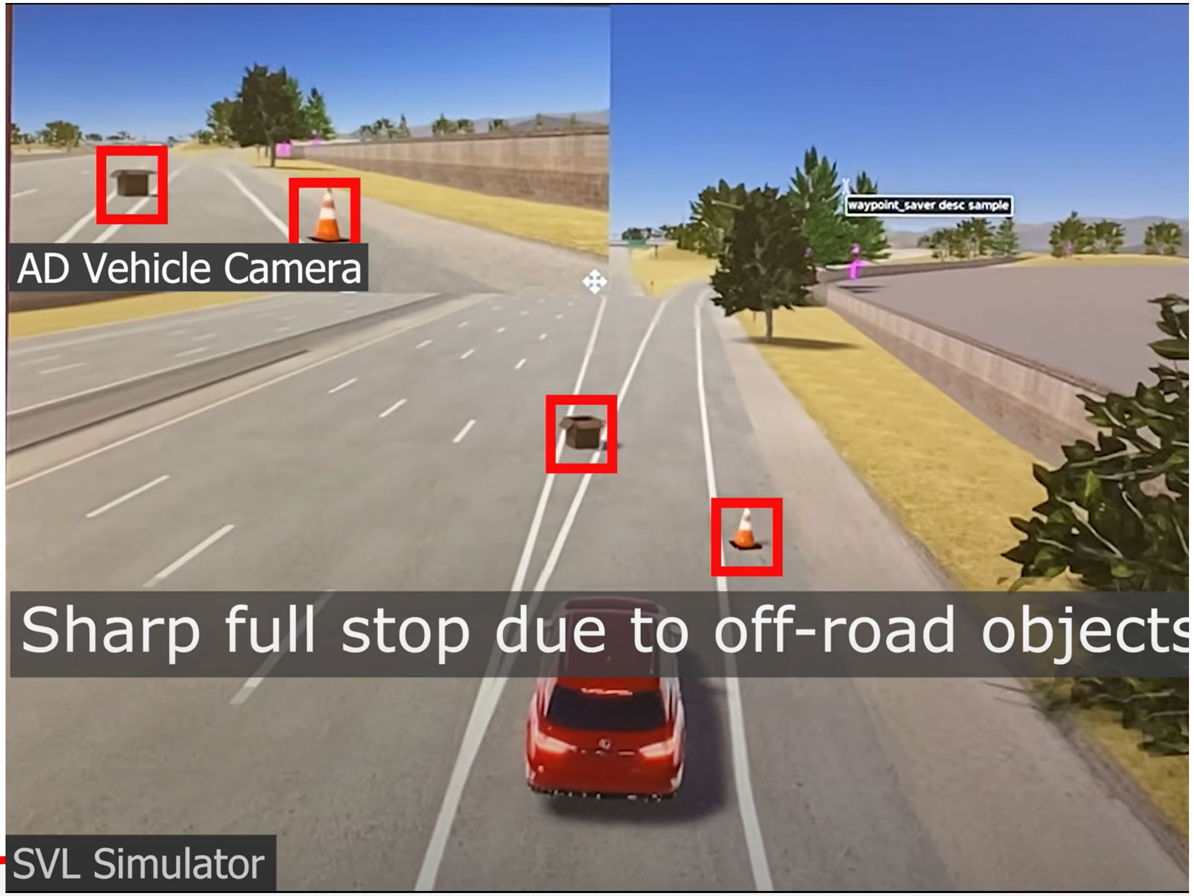

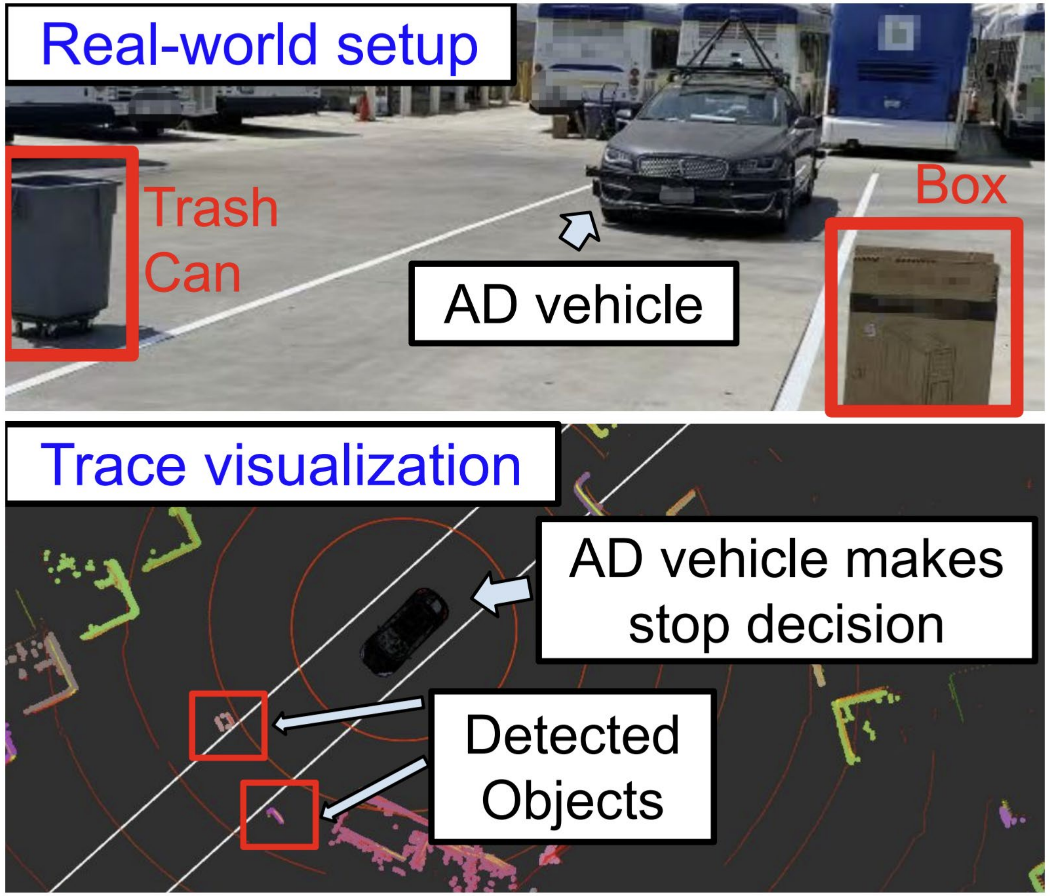

My group's paper

First security analysis of AD behavior planning (*program-based*)

- Attack vector: Common road objects (e.g., road-side cardboard boxes, parked bikes, etc.)
- Methodology: Domain-customized evolutionary testing
- Impact: Unnecessary sharp braking, stopping, giving up mission-critical driving decisions, etc.
- Demo website: https://sites.google.com/view/cav-sec/planfuzz

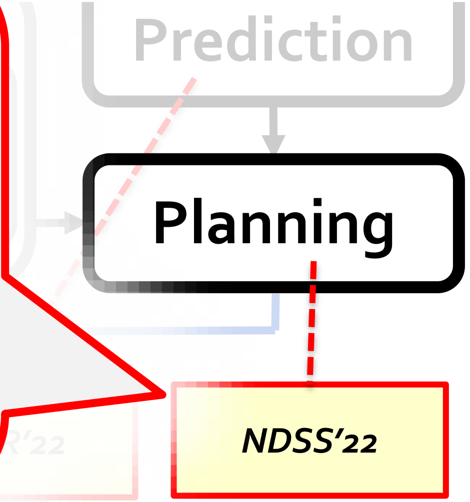

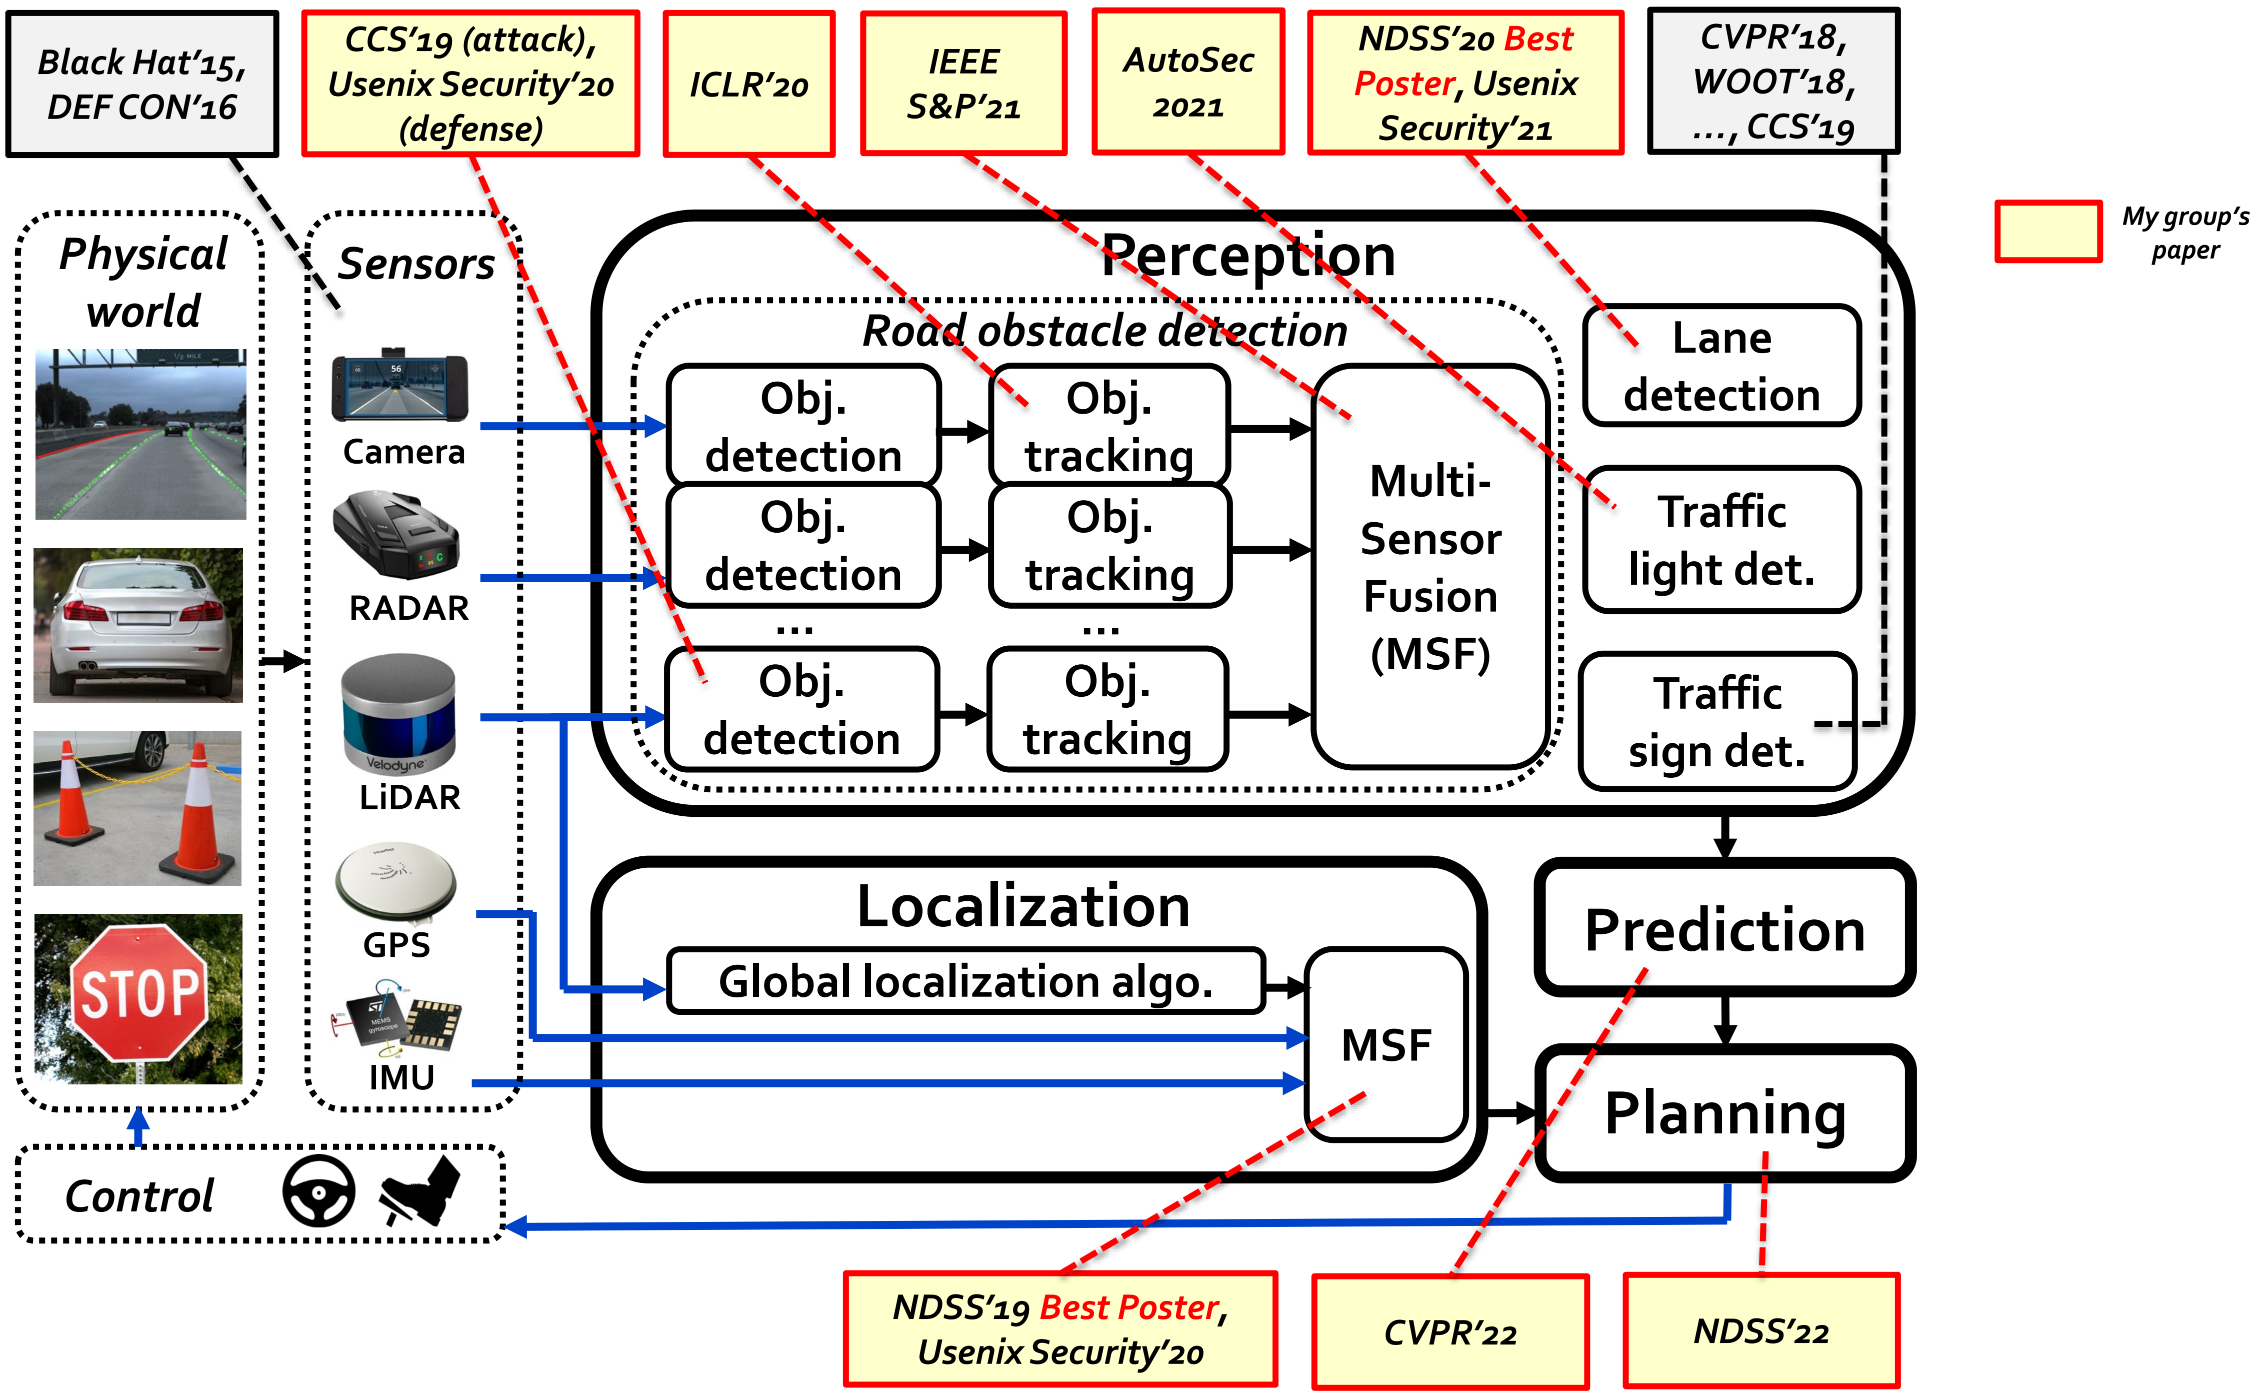

Black Hat'15, DEF CON'16
CCS'19 (attack), Usenix Security'20 (defense)
ICLR'20
IEEE S&P'21
AutoSec 2021
NDSS'20 Best Poster, Usenix Security'21
CVPR'18, WOOT'18, …, CCS'19
My group's paper
Physical world
Sensors
Camera
RADAR
LiDAR
GPS
IMU
Perception
Road obstacle detection
Obj. detection
Obj. tracking
Multi-Sensor Fusion (MSF)
Lane detection
Traffic light det.
Traffic sign det.
Localization
Global localization algo.
MSF
Prediction
Planning
Control
NDSS'19 Best Poster, Usenix Security'20
CVPR'22
NDSS'22

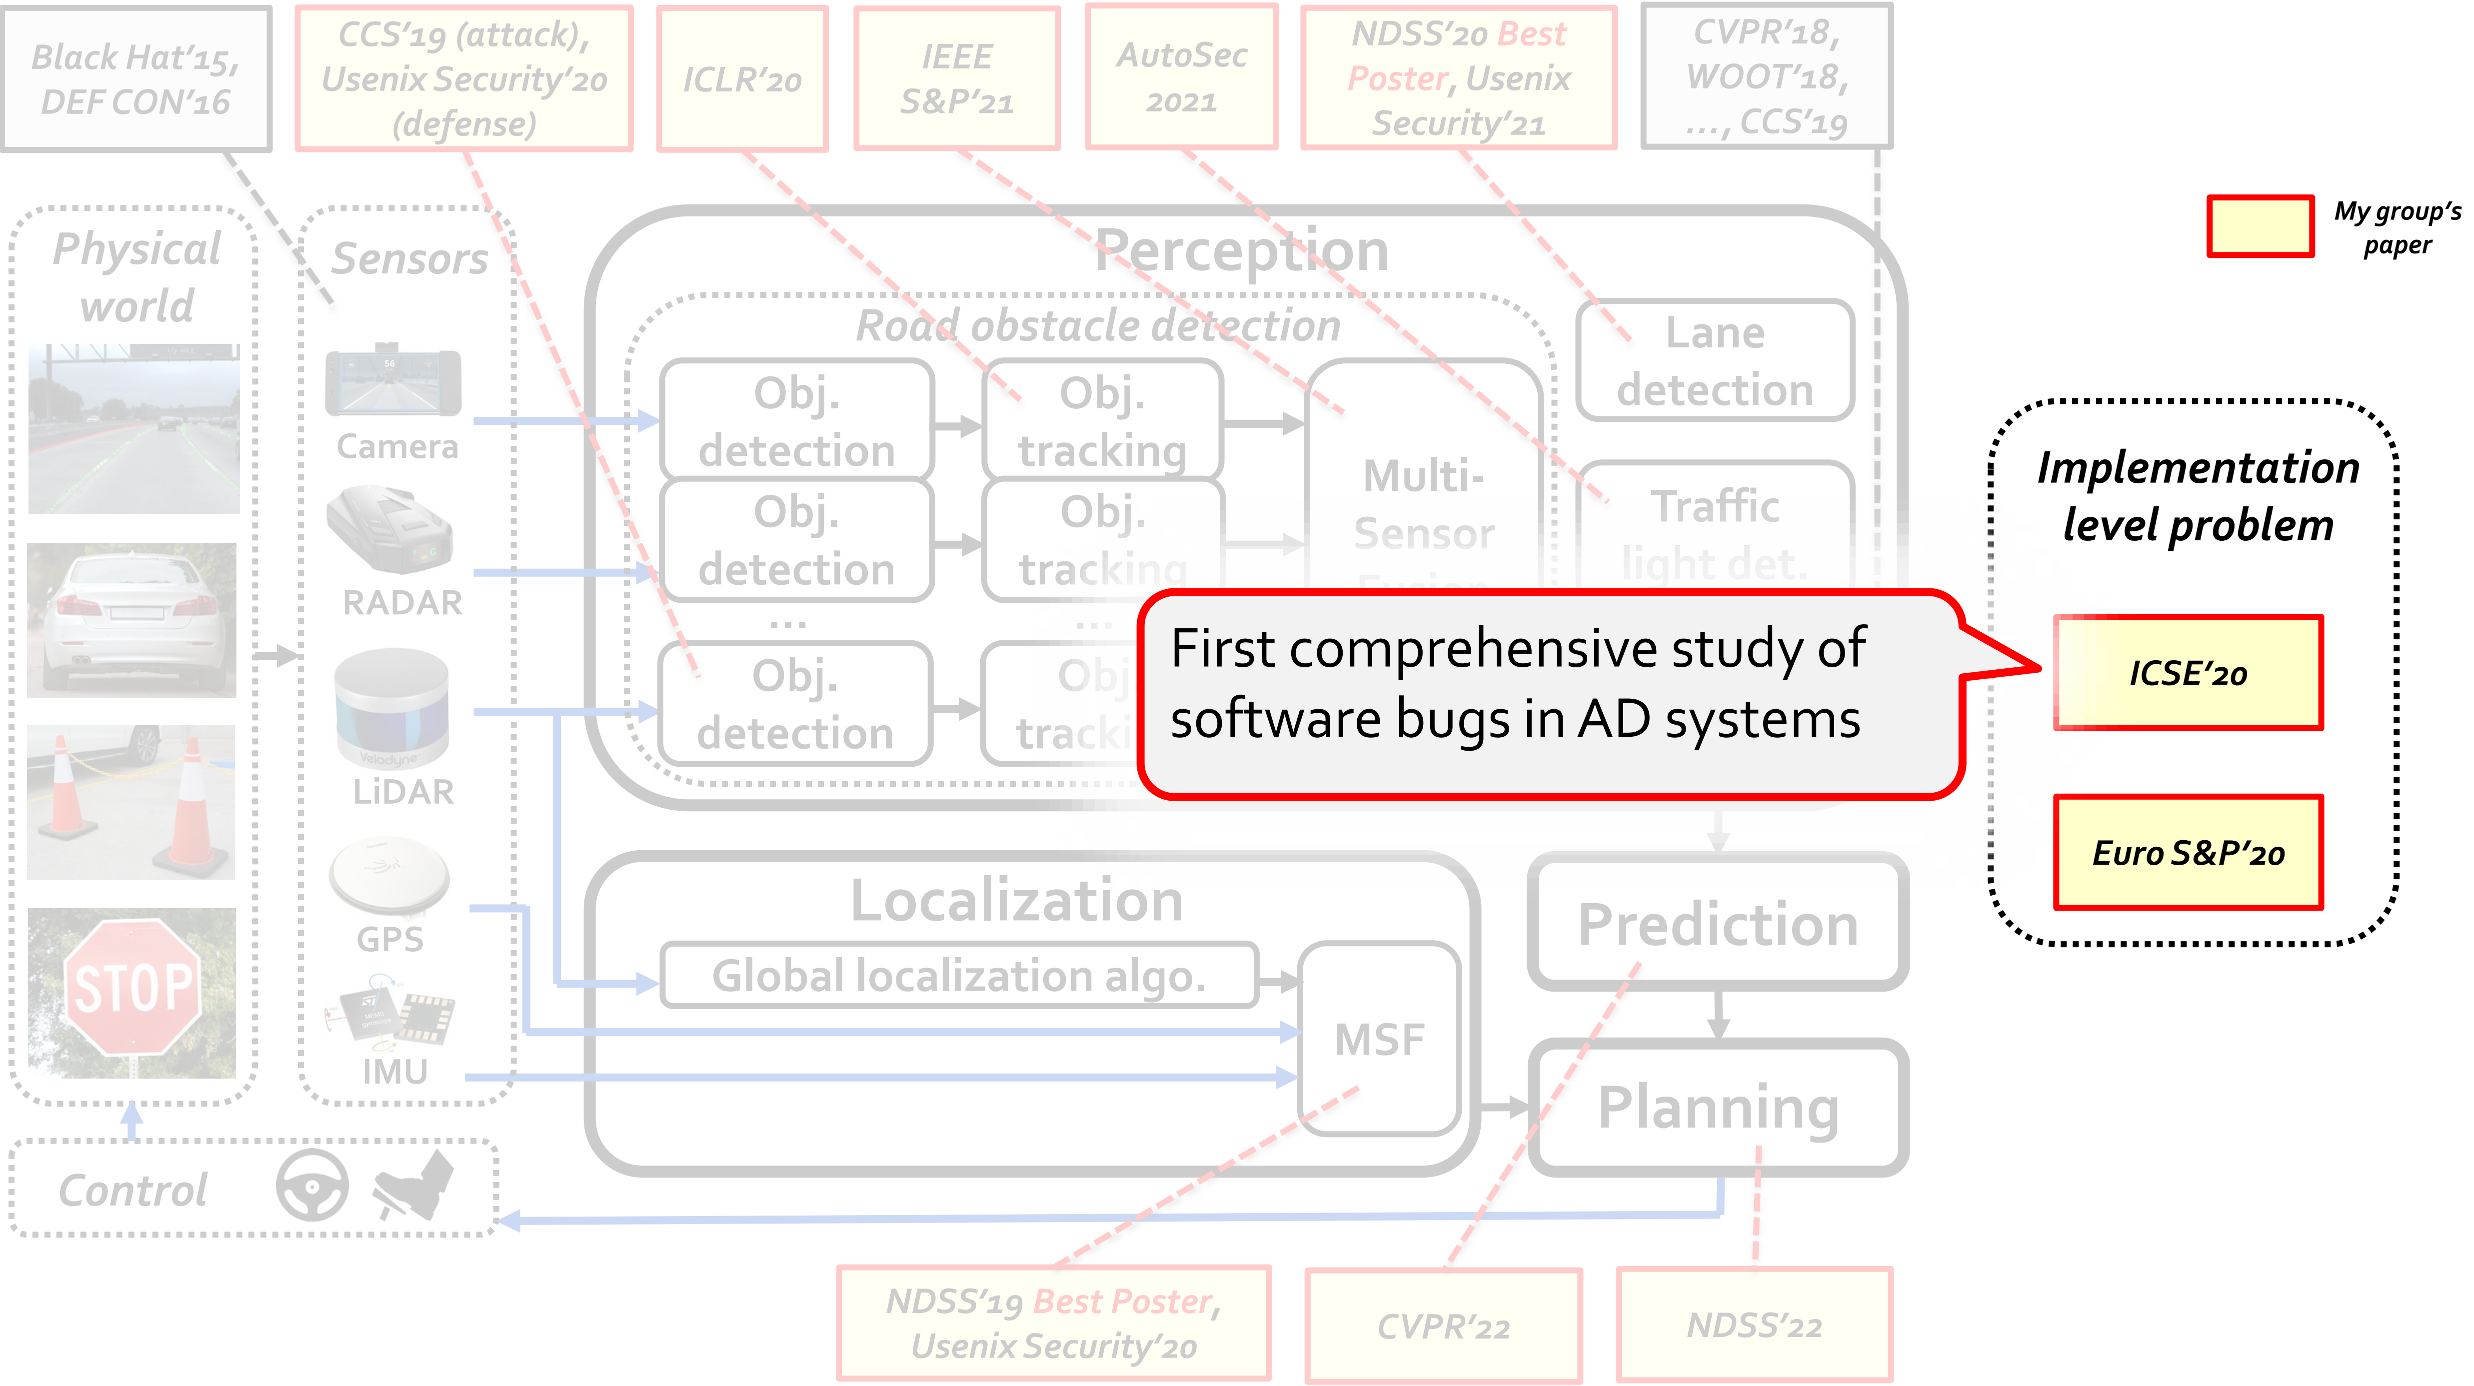

Black Hat'15, DEF CON'16
CCS'19 (attack), Usenix Security'20 (defense)
ICLR'20
IEEE S&P'21
AutoSec 2021
NDSS'20 Best Poster, Usenix Security'21
CVPR'18, WOOT'18, …, CCS'19
My group's paper
Physical world
Sensors
Camera
RADAR
LiDAR
GPS
IMU
Perception
Road obstacle detection
Obj. detection
Obj. tracking
Multi-Sensor Fusion
Lane detection
Traffic light det.
Localization
Global localization algo.
MSF
Prediction
Planning
Control
First comprehensive study of software bugs in AD systems
Implementation level problem
ICSE'20
Euro S&P'20
NDSS'19 Best Poster, Usenix Security'20
CVPR'22
NDSS'22

Black Hat'15,

CCS'19 (attack), Usenix Security'20

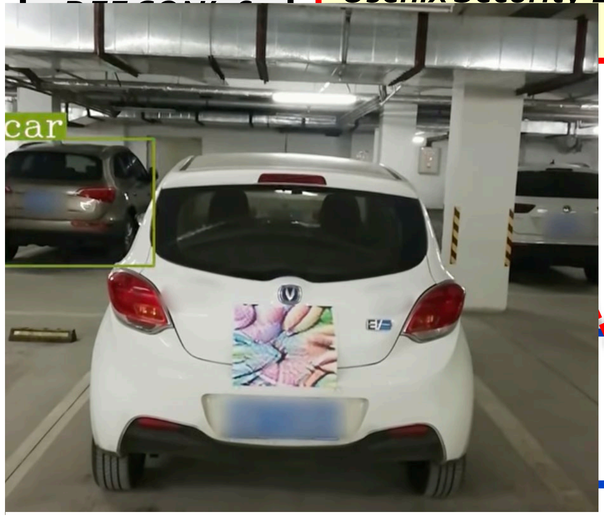

[Zhao et al. @ CCS'19]

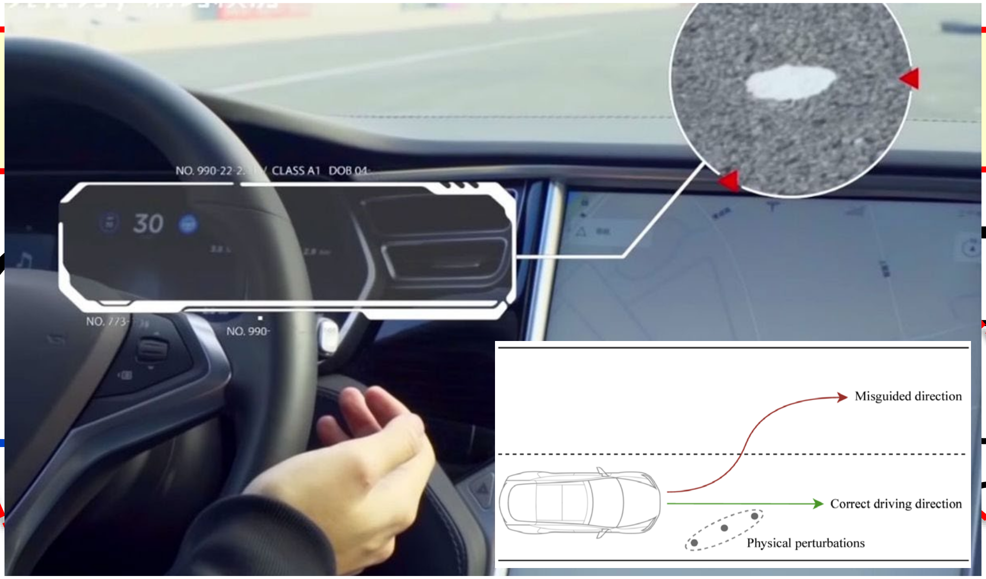

[Jing et al. @ Usenix Security'21]

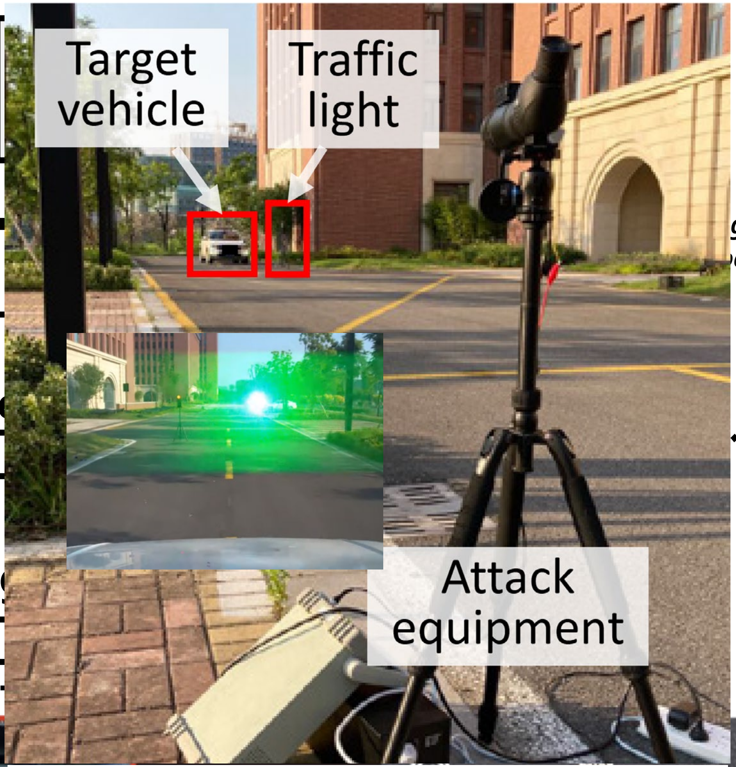

[Yan et al. @ Usenix Security'22]

group's paper

[Nassi et al. @ CCS'20]

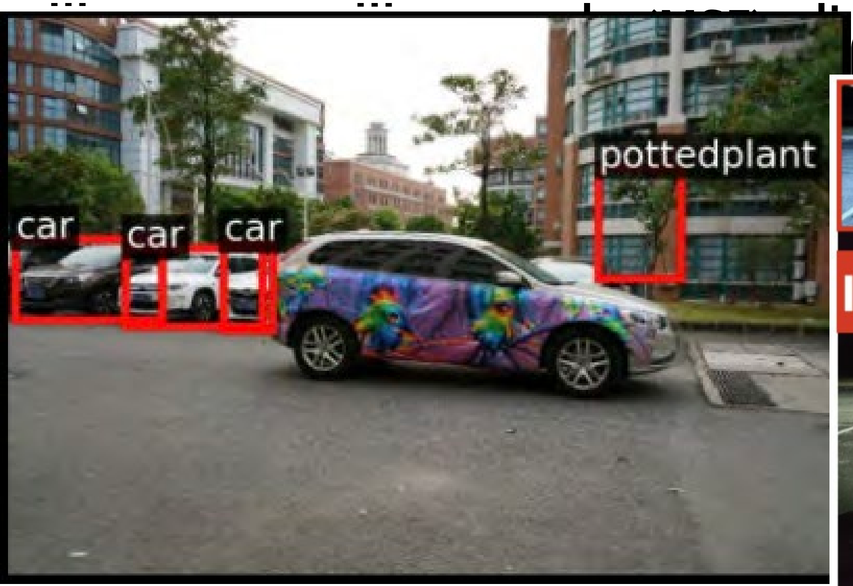

[Huang et al. @ CVPR'20]

NDSS'19 Best Poster, Usenix Security'20

CVPR'22

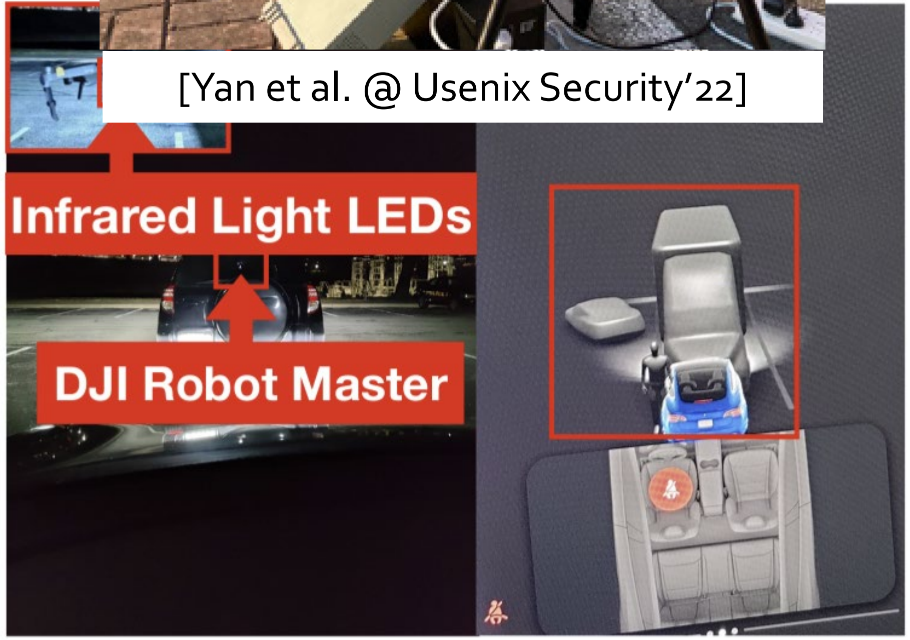

[Wang et al. @ CCS'21]

Black Hat'15,

CCS'19 (attack), Usenix Security'20

AD AI security papers

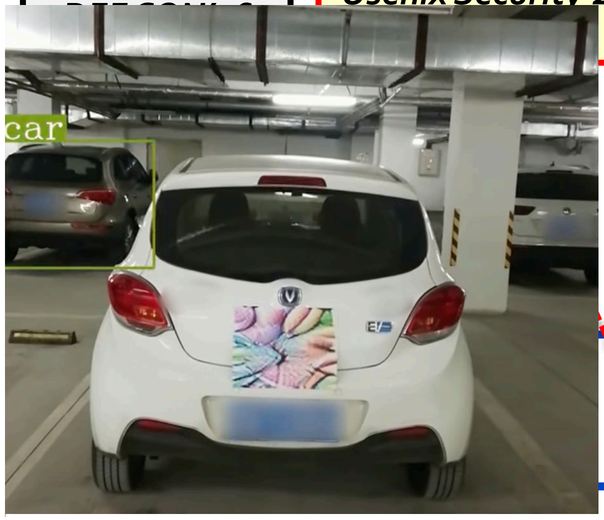

[Zhao et al. @ CCS'19]

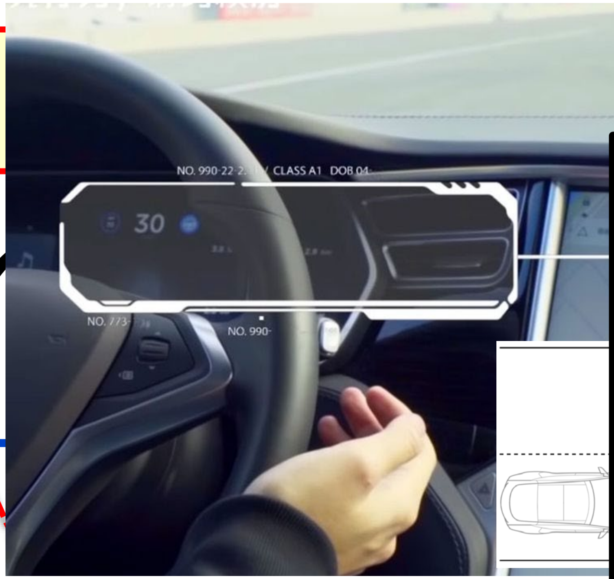

[Jing et al. @ Usenix Se

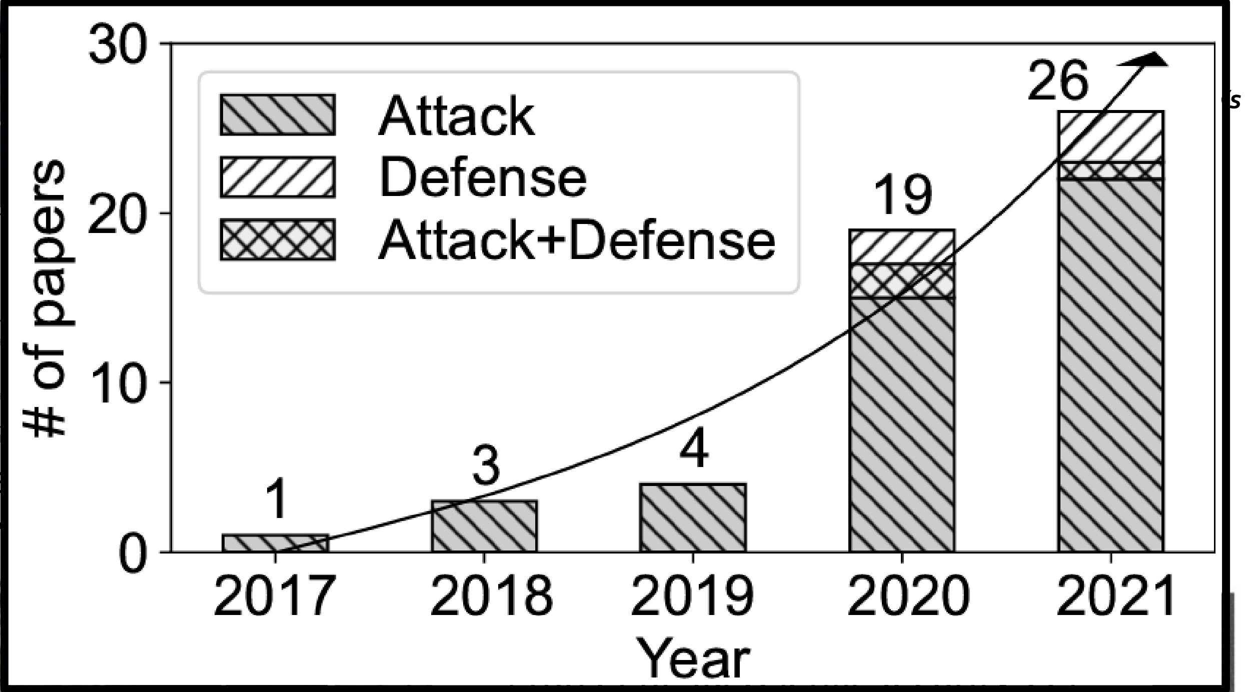

(*Systematization of Knowledge (SoK) effort from my group)

[Nassi et al. @ CCS'20]

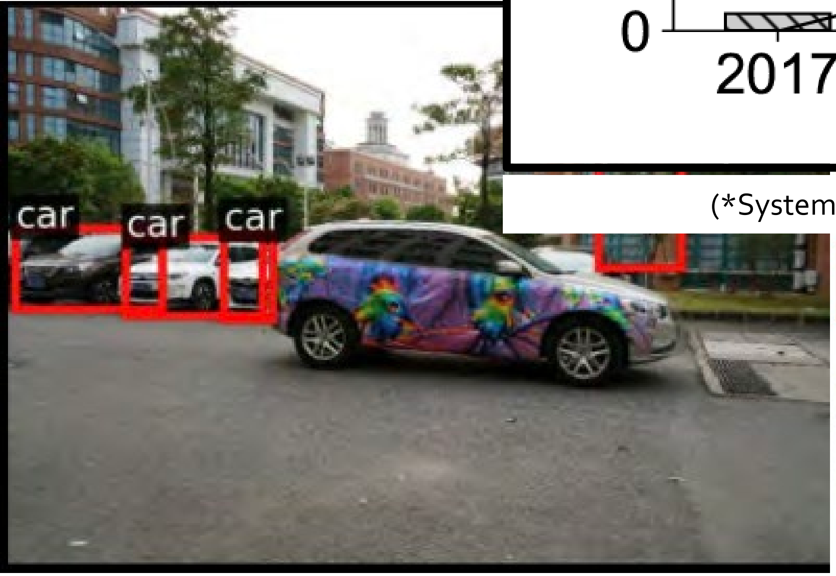

[Huang et al. @ CVPR'20]

NDSS'19 Best Poster, Usenix Security'20

CVPR'22

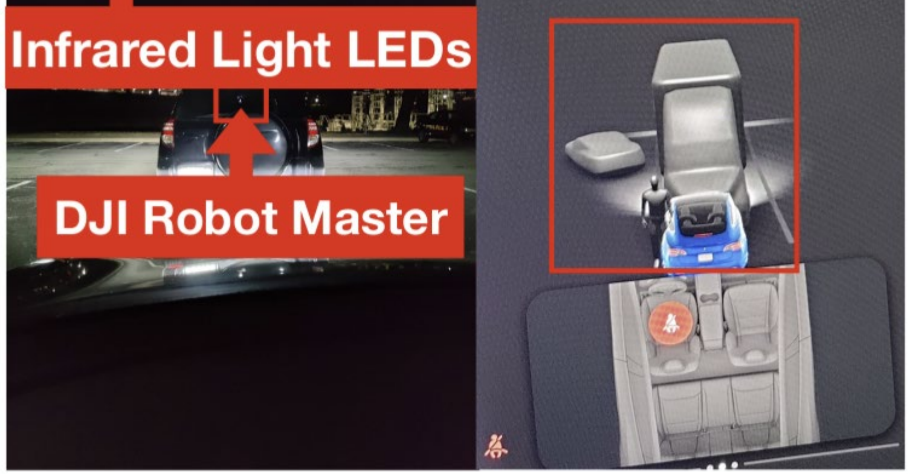

[Wang et al. @ CCS'21]

NDSS SYMPOSIUM NDSS About Previous Events

NDSS 2022 Attend Program Submissions Committees Sponsorship NDSS 2023

Automotive and Autonomous Vehicle Security (AutoSec) Workshop 2022

Note: All times are in PDT (UTC-7) and all sessions

Best Demo Award Voting (end at 4:40pm): https://

Future of AutoSec Voting (always open for your in

https://www.surveymonkey.com/r/9Q7JJMH

Proceedings Frontmatter

Sunday April 24

9:00 am - 9:10 am	Welcome to AutoSec 2022 and
9:10 am - 10:10 am	Keynote #1: Prof. Dongyan Xu Conte Professor of Computer S Purdue University)

Keynote #1

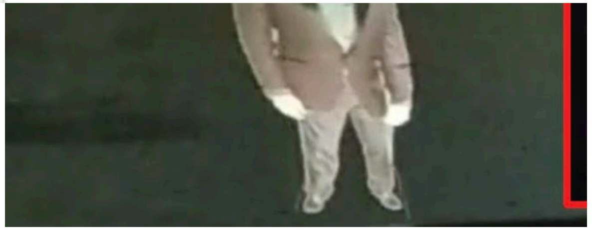

[Nassi et al. @ CCS'

AutoSec2022@NDSS
@autosec_conf

First-ever AutoSec PC meeting just occurred!! >18 PCs attended & looooots of paper debating and even new ideas on how to run the workshop in the future --- what a healthy community 🤩 ! Paper decisions will come out tomorrow. Stay tuned! #autosec22 @NDSSSymposium

AutoSec2022@NDSS @autosec_conf · Jan 13
Wow, another year of a record number of submissions #autosec22 @NDSSSymposium ! 32 regular/short/wip+ 17 demo submissions, which are 23%+70% more than last year!! Looks like the community is growing crazily 🤩 Now the review process begins… Good luck to all authors!

AD AI security papers

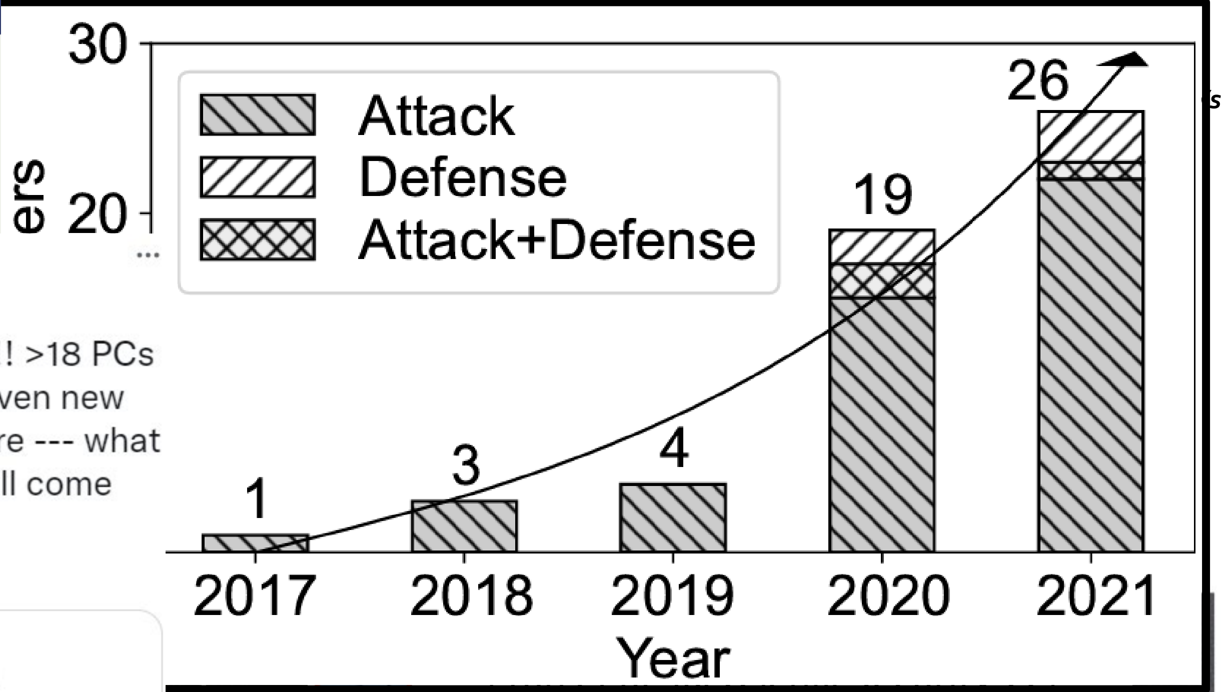

(*Systematization of Knowledge (SoK) effort from my group)

VPR'22

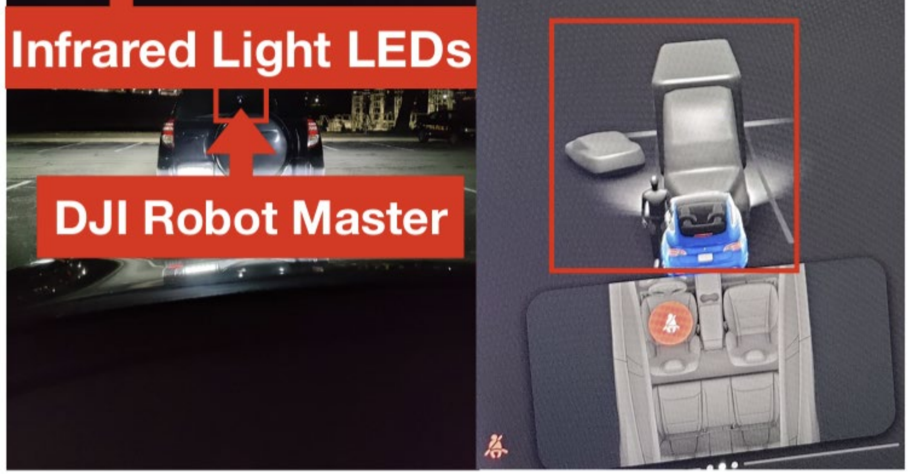

[Wang et al. @ CCS'21]

NDSS About Previous Events

NDSS 2022 Attend Program Submissions Committees Sponsorship NDSS 2023

Automotive and Autonomous Vehicle Security (AutoSec) Workshop 2022

Note: All times are in PDT (UTC-7) and all sessions

Best Demo Award Voting (end at 4:40pm): https://

Future of AutoSec Voting (always open for your in

https://www.surveymonkey.com/r/9Q7JJMH

Proceedings Frontmatter

Sunday April 24	
9:00 am - 9:10 am	Welcome to AutoSec 2022 and
9:10 am - 10:10 am	Keynote #1: Prof. Dongyan Xu Conte Professor of Computer S Purdue University)

Keynote #1

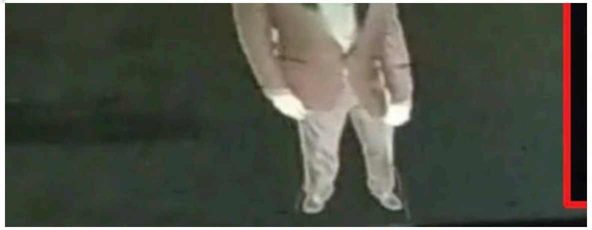

[Nassi et al. @ CCS'

AutoSec2022@NDSS
@autosec_conf

First-ever AutoSec PC meeting just occurred!! >18 PCs attended & loooooots of paper debating and even new ideas on how to run the workshop in the future --- what a healthy community 🤩 ! Paper decisions will come out tomorrow. Stay tuned! #autosec22 @NDSSSymposium

AutoSec2022@NDSS @autosec_conf · Jan 13
Wow, another year of a record number of submissions #autosec22 @NDSSSymposium ! 32 regular/short/wip+ 17 demo submissions, which are 23%+70% more than last year!! Looks like the community is growing crazily 🤩 Now the review process begins... Good luck to all authors!

AD AI security papers

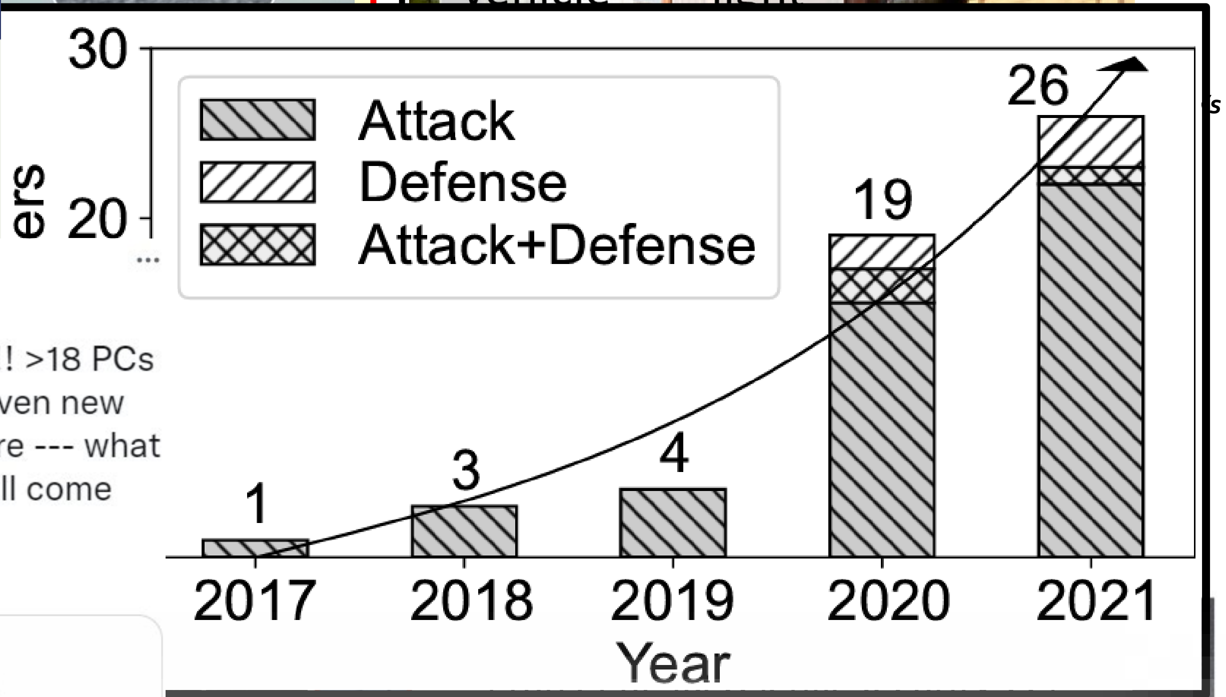

(*Systematization of Knowledge (SoK) effort from my group)

VPR'22

>5 years of growth in research space & community now!
→ *Time for some reflection?*

[Wang et al. @ CCS'21]

A reflection of the 5+ years of AD AI security research

- Conduct the **first Systematization of Knowledge (SoK) effort** on **semantic AI security** research in AD
  - Collect & analyze ***53 papers in past 5 years***, mainly from ***top-tier venues in security, CV (Computer Vision), ML (Machine Learning), AI, and robotics***

SoK: On the Semantic AI Security in Autonomous Driving

Junjie Shen, Ningfei Wang, Ziwen Wan, Yunpeng Luo, Takami Sato, Zhisheng Hu†, Xinyang Zhang†, Shengjian Guo†, Zhenyu Zhong†, Kang Li†, Ziming Zhao‡, Chunming Qiao‡, Qi Alfred Chen

{junjies1, ningfei.wang, ziwenw8, yunpel3, takamis, alfchen}@uci.edu,
†{zhishenghu, xinyangzhang, sjguo, edwardzhong, kangli01}@baidu.com, ‡{zimingzh, qiao}@buffalo.edu
UC Irvine, †Baidu Security, ‡University at Buffalo

Link: https://arxiv.org/abs/2203.05314

Our SoK effort

- **Taxonomization, status & trend analysis,** based on critical research aspects for security
  - E.g., attack/defense goal, attack vector, defense deployability, evaluation methodologies, etc.

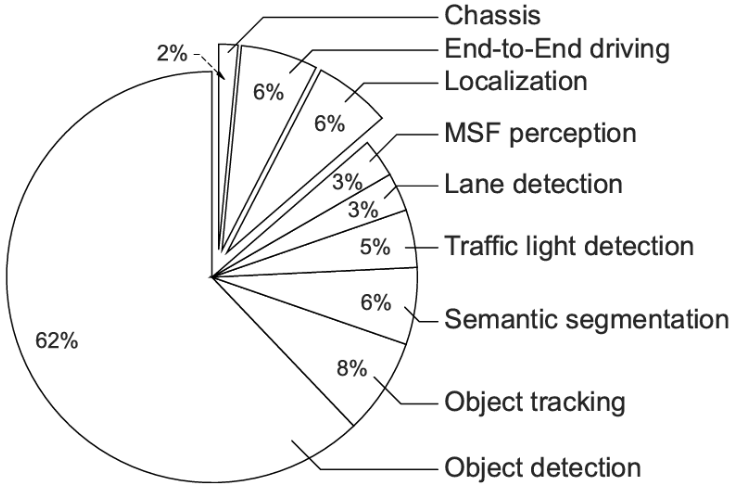

Figure 6: Distribution of (attack/defense) targeted AI components in semantic AD AI security papers.

Targeted AI component		Paper	Year	Field	Attack goal			Attack vector: Physical-layer: Phys. world			Physical-layer: Sensor attack						Cyber layer			Eval. level		
					Integrity	Confidentiality	Availability	Object texture	Object shape	Object position	GPS spoofing	LiDAR spoofing	Radar spoofing	Laser/IR light	Acoustic signal	Translucent patch	ML backdoor	Malware & s/w compromise	Attacker's knowledge	Component-level	System-level	Open source
Camera perception	Object detection	Lu et al. [54]	'17	V	✓			✓											○	✓		
		Eykholt et al. [18]	'18	S	✓			✓											○	✓		
		Chen et al. [37]	'18	M	✓			✓											○	✓		✓
		Zhao et al. [26]	'19	S	✓			✓											○	✓		
		Xiao et al. [55]	'19	V	✓			✓	✓										○	✓		
		Zhang et al. [56]	'19	M	✓			✓											◐	✓		✓
		Nassi et al. [57]	'20	S	✓			✓											○	✓	✓	
		Man et al. [58]	'20	S	✓									✓					○	✓		✓
		Hong et al. [59]	'20	S	✓													✓	○		✓	
		Huang et al. [60]	'20	V	✓			✓											○	✓		✓
		Wu et al. [61]	'20	V	✓			✓											○	✓		
		Xu et al. [62]	'20	V	✓			✓											○	✓		
		Hu et al. [63]	'20	V	✓			✓											◐	✓		
		Hamdi et al. [64]	'20	M	✓				✓										◐	✓		
		Ji et al. [65]	'21	S	✓										✓				◐	✓		
		Lovisotto et al. [66]	'21	S	✓			✓											●	✓		✓
		Wang et al. [67]	'21	S	✓									✓					◐	✓		
		Köhler et al. [68]	'21	S	✓									✓					●	✓		
		Wang et al. [69]	'21	S	✓			✓											●	✓		✓
		Zolfi et al. [70]	'21	V	✓											✓			○	✓		
		Wang et al. [71]	'21	V	✓			✓											◐	✓		✓
		Zhu et al. [72]	'21	M	✓									✓					○	✓		
	Semantic segmentation	Nakka et al. [73]	'20	V	✓			✓											○	✓		
		Nesti et al. [74]	'22	V	✓			✓											○	✓		
	Object tracking	Jha et al. [75]	'20	S	✓													✓	○		✓	
		Jia et al. [17]	'20	M	✓			✓											○	✓		✓
		Ding et al. [76]	'21	M	✓			✓											○	✓		
		Chen et al. [77]	'21	M	✓			✓											○	✓		
	Lane detection	Sato et al. [78]	'21	S	✓			✓											○	✓	✓	✓
		Jing et al. [79]	'21	S	✓			✓											◐	✓	✓	
	Traffic light detection	Wang et al. [67]	'21	S	✓									✓					○	✓		
		Tang et al. [80]	'21	S	✓						✓								○		✓	
LiDAR perception	Object detection	Cao et al. [19]	'19	S	✓							✓							○	✓	✓	
		Sun et al. [81]	'20	S	✓							✓							●	✓		
		Hong et al. [59]	'20	S	✓													✓	○		✓	
		Tu et al. [82]	'20	V	✓				✓										○	✓		
		Zhu et al. [83]	'21	S	✓					✓									◐	✓		
		Yang et al. [84]	'21	S	✓				✓										◐	✓	✓	
		Hau et al. [85]	'21	S	✓							✓							○	✓		
		Li et al. [86]	'21	V	✓						✓								○	✓		✓
		Zhu et al. [87]	'21	O	✓					✓									◐	✓		
	Semantic segmentation	Tsai et al. [88]	'20	M	✓				✓										○	✓		✓
		Zhu et al. [87]	'21	O	✓					✓									◐	✓		
RADAR perception	Obj. detection	Sun et al. [89]	'21	S	✓								✓						◐	✓	✓	
MSF perception		Cao et al. [38]	'21	S	✓				✓										○	✓	✓	✓
		Tu et al. [90]	'21	O	✓				✓										○	✓		
LiDAR localization		Luo et al. [91]	'20	S		✓												✓	○		✓	
MSF localization		Shen et al. [92]	'20	S	✓						✓								○	✓	✓	
Camera localization		Wang et al. [67]	'21	S	✓									✓					○	✓		
Chassis		Hong et al. [59]	'20	S		✓												✓	○		✓	
End-to-end driving		Liu et al. [93]	'18	S	✓			✓									✓		○	✓	✓	✓
		Kong et al. [94]	'20	V	✓			✓											○	✓		✓
		Hamdi et al. [64]	'20	M	✓				✓										◐	✓		
		Boloor et al. [95]	'20	O	✓			✓											●	✓	✓	✓

Field: S = Security, V = Computer Vision, M = ML/AI, O = Others, e.g., Robotics, arXiv;
Attacker's knowledge: ○ = white-box, ◐ = gray-box, ● = black-box

Table I. Overview of existing semantic AD AI attacks in our SoK scope (§II-C). (s/w = software)

Our SoK effort: Scientific gaps identification

- Most importantly, identify **6 most substantial scientific gaps**
  - Observed based on quantitative comparisons both ***vertically*** among existing AD AI security works and ***horizontally*** with security works from closely-related domains
  - <u>Scientific Gap 1</u>: **Evaluation**: General lack of system-level evaluation
    - Only 25.4% of existing works perform system-level evaluation
  - <u>Scientific Gap 2</u>: **Research goal**: General lack of defense solutions
    - Only 14.3% propose defenses
    - In comparison, much more balanced in drone security area (49% on defense)
  - <u>Scientific Gap 3</u>: **Attack vector**: Cyber-layer attack vectors under-explored
    - Only 11.1% assume cyber-layer attack vectors, e.g., malware, ML backdoors
  - <u>Scientific Gap 4</u>: **Attack target**: Downstream AI components under-explored
    - Limited study on prediction & planning
  - <u>Scientific Gap 5</u>: **Attack goal**: Attack goals other than "integrity" under-explored
    - Limited study on confidentiality & availability
  - <u>Scientific Gap 6</u>: **Community**: Substantial Lack of Open Sourcing
    - <20.6% (7/34) papers from security conferences release source code

Our SoK effort
(https://arxiv.org/abs/2203.05314)

Most critical gap: General lack of system-level evaluation

- Most importantly, identify **6 most substantial scientific gaps**
  - Observed based on quantitative comparisons both ***vertically*** among existing AD AI security works and ***horizontally*** with security works from closely-related domains
  - Scientific Gap 1: **Evaluation**: General lack of system-level evaluation
    - Only 25.4% of existing works perform system-level evaluation
  - Scientific Gap 2: **Research goal**: General lack of defense solutions
    - O... defenses

- Widely recognized that in autonomous system, AI component-level errors (e.g., obj det errors) **do not necessarily lead to system-level effect (e.g., collisions)**
  - Essentially the AI-to-system semantic gap mentioned earlier
- However, today **vast majority (74.6%)** of existing works ***did not perform any form of system-level evaluation***
  - I.e., eval w/ full-stack AD system & closed-loop control via simulation/real-vehicle setups
- Without it, may lead to ***meaningless*** attack/defense progress at the system level

<20.6% (7/34) papers from security conferences release source code

Our SoK effort
(https://arxiv.org/abs/2203.05314)

Demo: Necessity of system-level evaluation

- <u>Setup</u>: Existing STOP sign disappearance attack [1]
  - **Effective at component level**: **>** ***70% frame-level success rate*** to make STOP sign disappear (consistent success pattern w/ [1])
  - However, **failed at system level**: ***0% stop sign violation rate*** due to **object tracking**

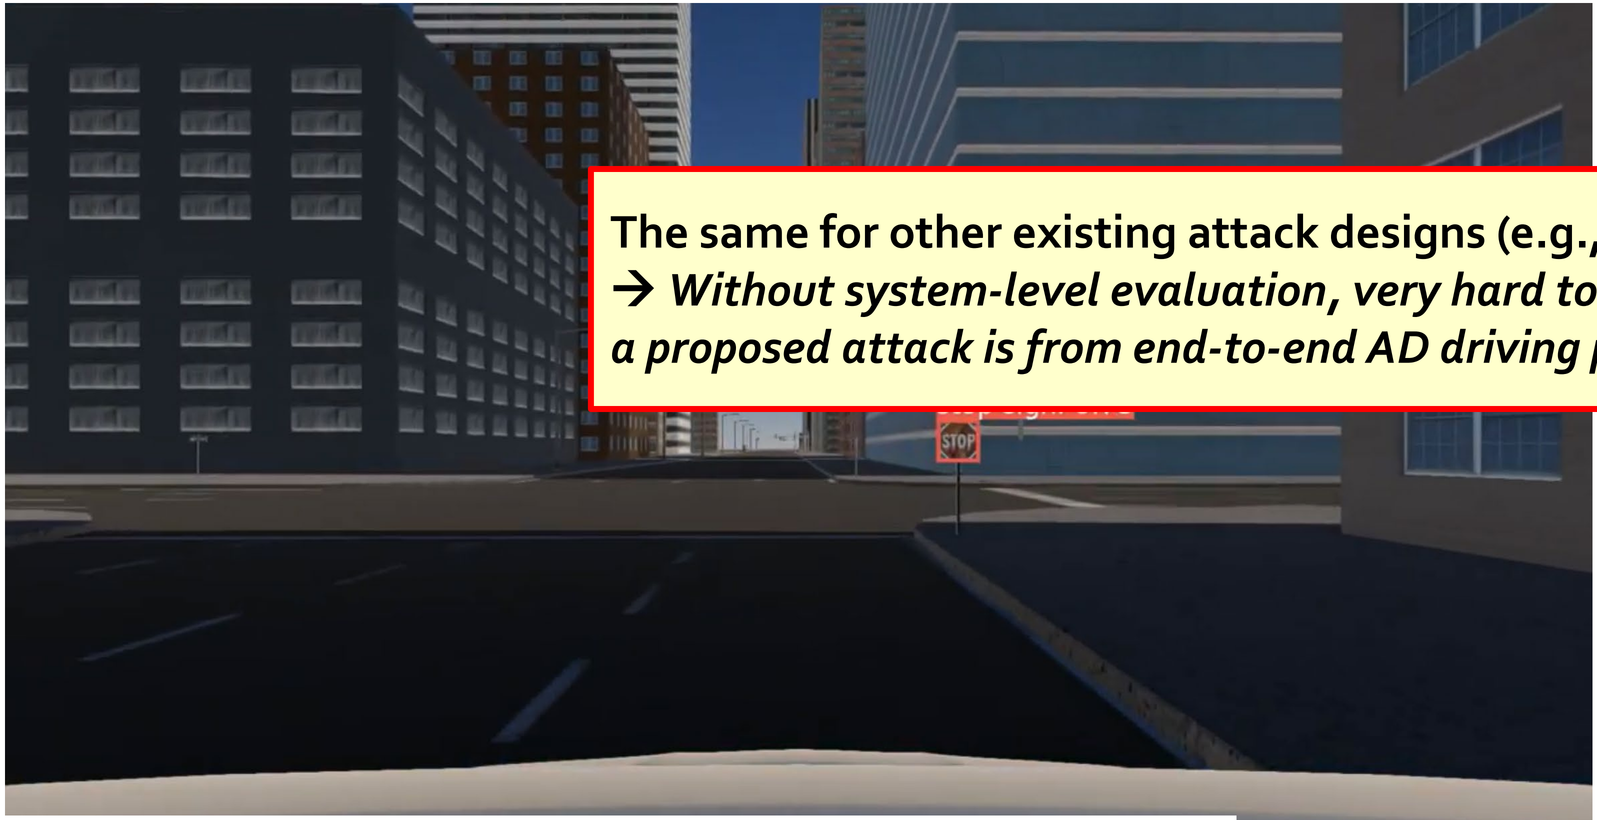

The same for other existing attack designs (e.g., [2] [3])
→ *Without system-level evaluation, very hard to know how meaningful a proposed attack is from end-to-end AD driving perspective!*

[1] Chen et al., "ShapeShifter: Robust Physical Adversarial Attack on Faster R-CNN Object Detector." ECML PKDD, 2018.
[2] Eykholt et al., "Physical Adversarial Examples for Object Detectors," WOOT, 2018
[3] Zhao et al., "Seeing isn't Believing: Towards More Robust Adversarial Attack Against Real World Object Detectors," ACM CCS, 2019

How to systematically address this?

- **Various challenges** to effectively fill this gap **at the *research community* level**
  - Real AD vehicle testing: Low affordability/accessibility, safety, flexibility, & reproducibility
  - Simulation-based testing: Still non-trivial engineering efforts to instrument simulation environment & engine for security testing
- A **community-level effort** can greatly help!
  - Collectively build a ***common system-level evaluation infrastructure***
  - <u>Benefits</u>:
    - (1) Avoid repeated engineering efforts in instrumenting the simulator/vehicle
    - (2) Improve result comparability due to the more unified evaluation setup, benefitting scientific advances

Our proposal: PASS (Platform for Autonomous driving Security and Safety)

- ***Open***, ***uniform*** & ***extensible*** system-driven evaluation platform

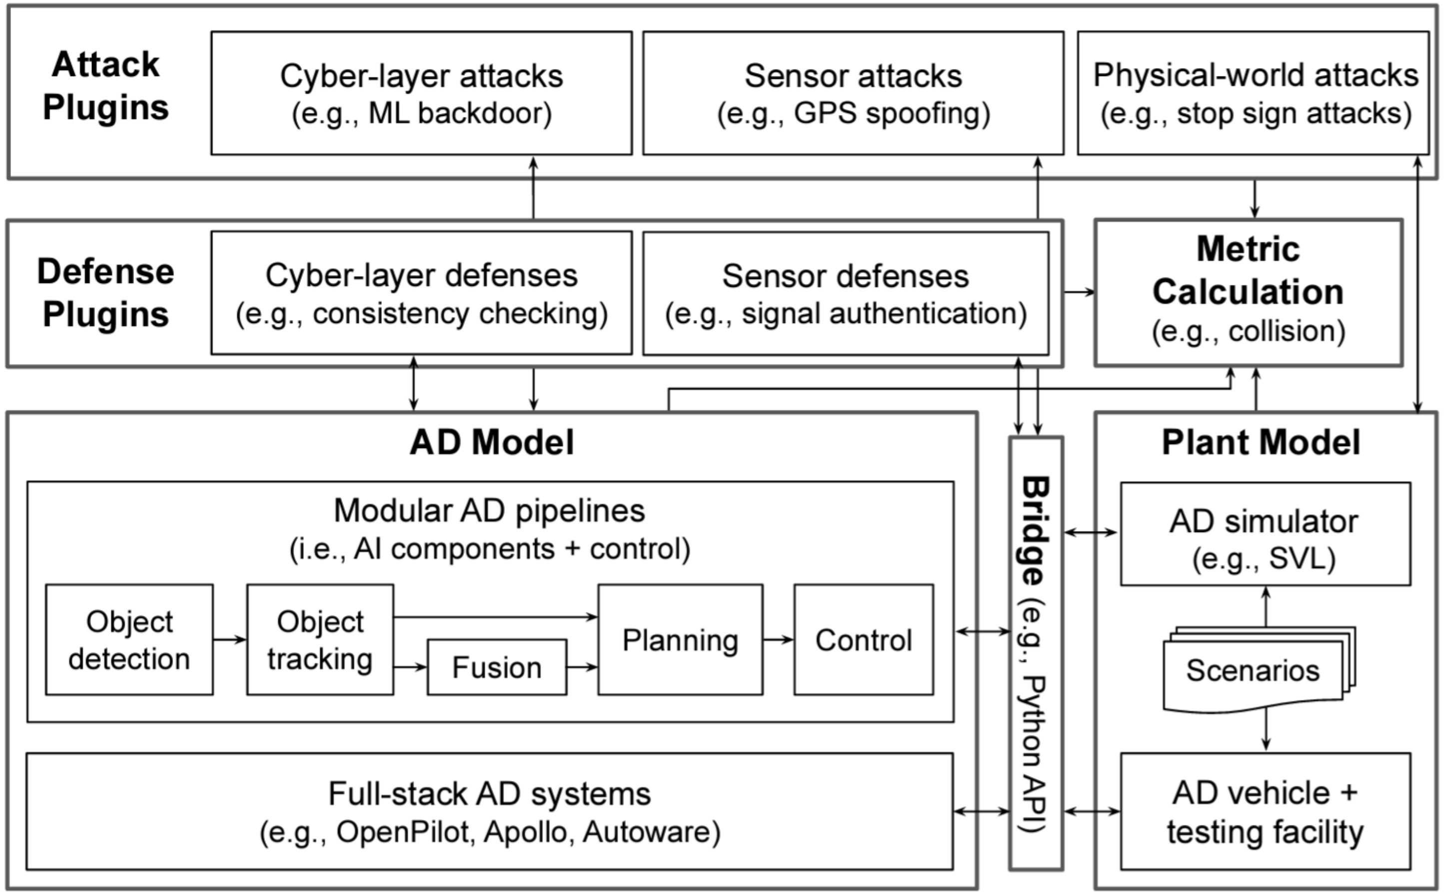

Our proposal: PASS (Platform for Autonomous driving Security and Safety)

- ***Open***, ***uniform*** & ***extensible*** system-driven evaluation platform

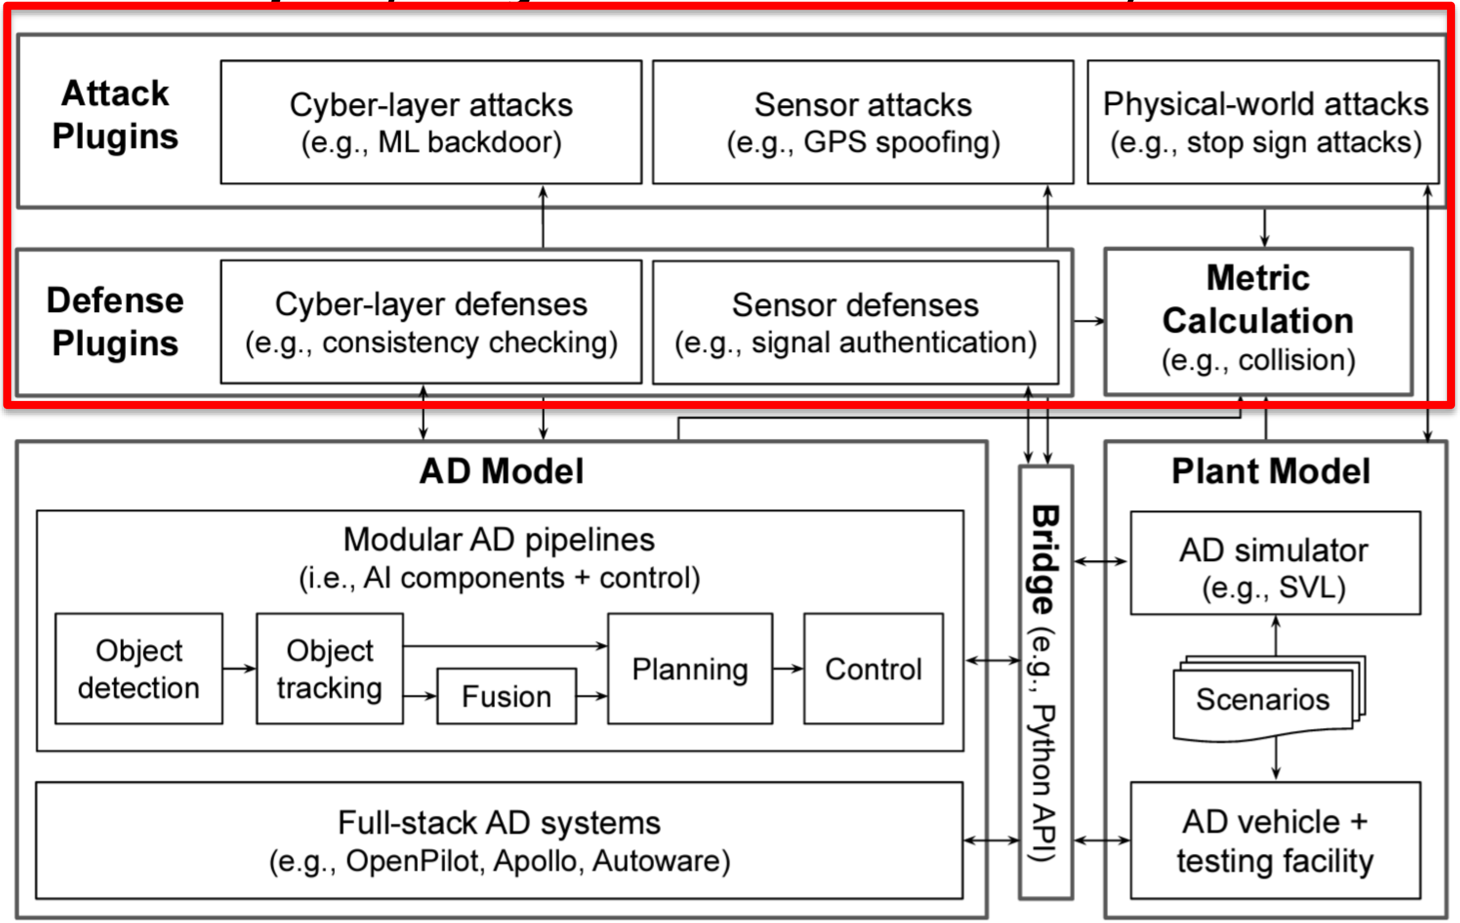

- **Attack/defense plugins** in Python APIs

```
# Physical-world attack API: adding a 2D patch or 3D object
#  - is_3d: True for 3D, False for 2D
#  - path: 2D image (.PNG) or 3D object (.FBX) file path
#  - transform: position, rotation, scaling vectors
add_physical_object(is_3d:bool,path:str,transform:{})->bool

# Sensor attack/defense API: register a callback function
# for modifying the sensor data at runtime
#  - sensor_type: GPS, LiDAR, IMU, Camera, etc.
#  - op: callback function defines how to modify sensor data
register_sensor_callback(sensor_type:enum,op:Callable)->bool

# Cyber layer attack/defense API: add a new component or
# replace an existing one with a new component
#  - comp_type: ObjectDetector, ObjectTracker, etc.
#  - new_comp: the Python code of the new component
add_component(comp_type:enum,new_comp:str)->bool
replace_component(comp_type:enum,new_comp:str)->bool
```

Our proposal: PASS (Platform for Autonomous driving Security and Safety)

- ***Open***, ***uniform*** & ***extensible*** system-driven evaluation platform

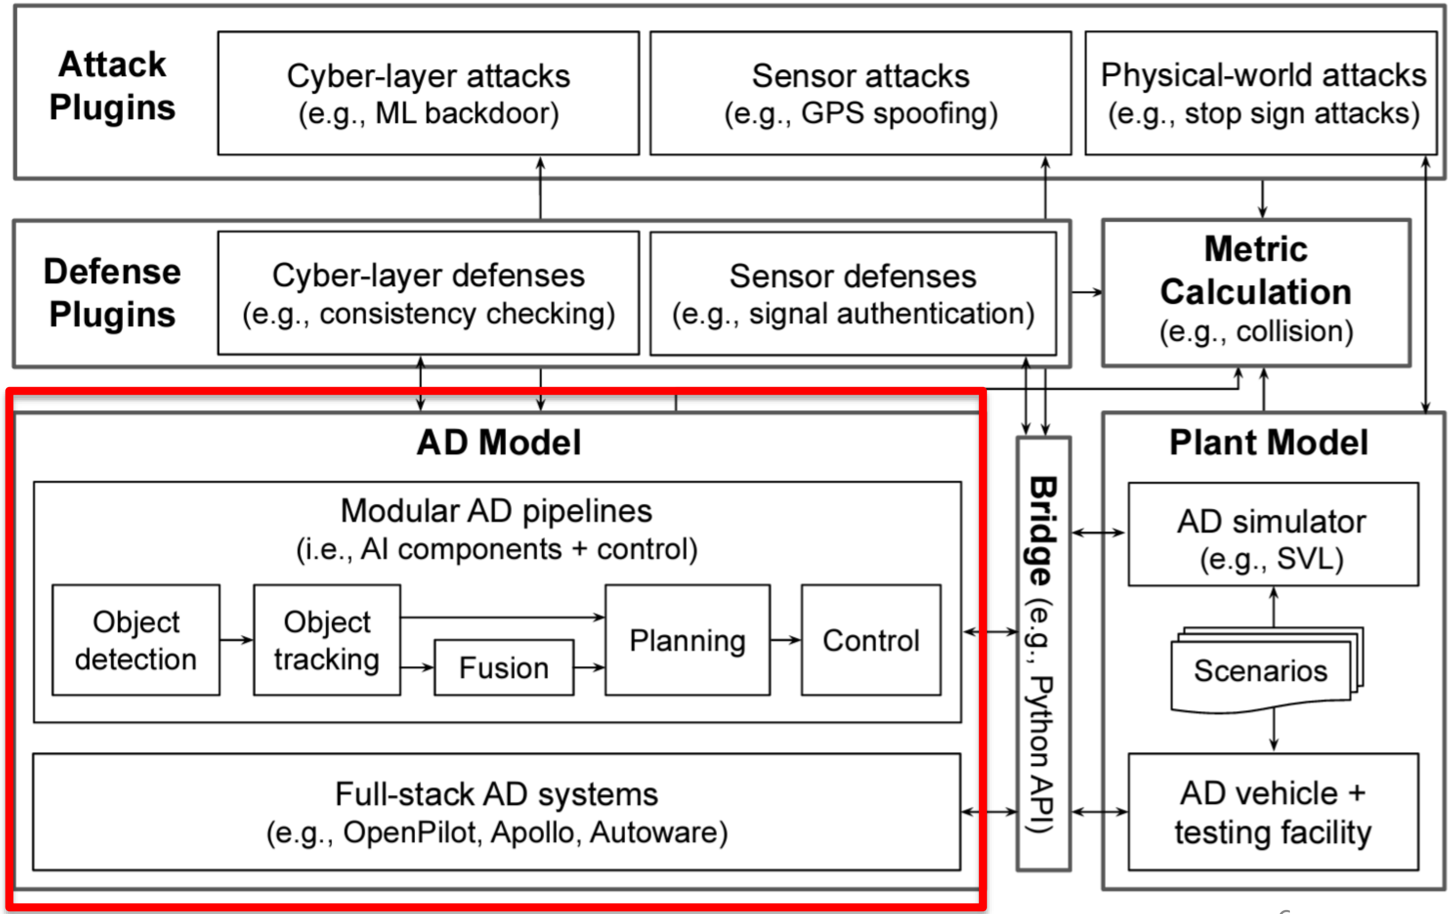

- **Attack/defense plugins** in Python APIs
- **Plug & play modular AD design**

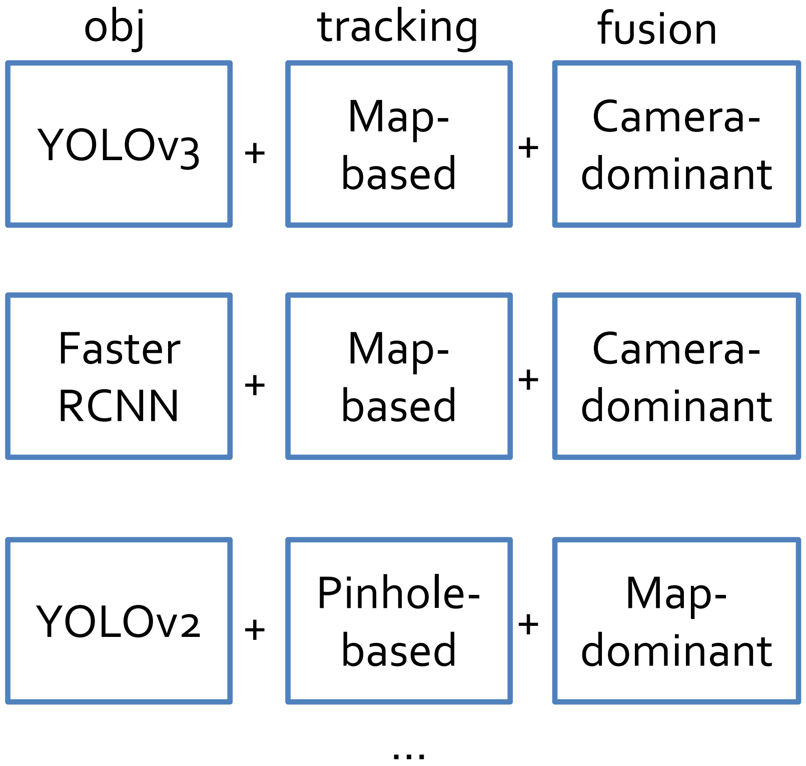

Our proposal: PASS (Platform for Autonomous driving Security and Safety)

- ***Open***, ***uniform*** & ***extensible*** system-driven evaluation platform

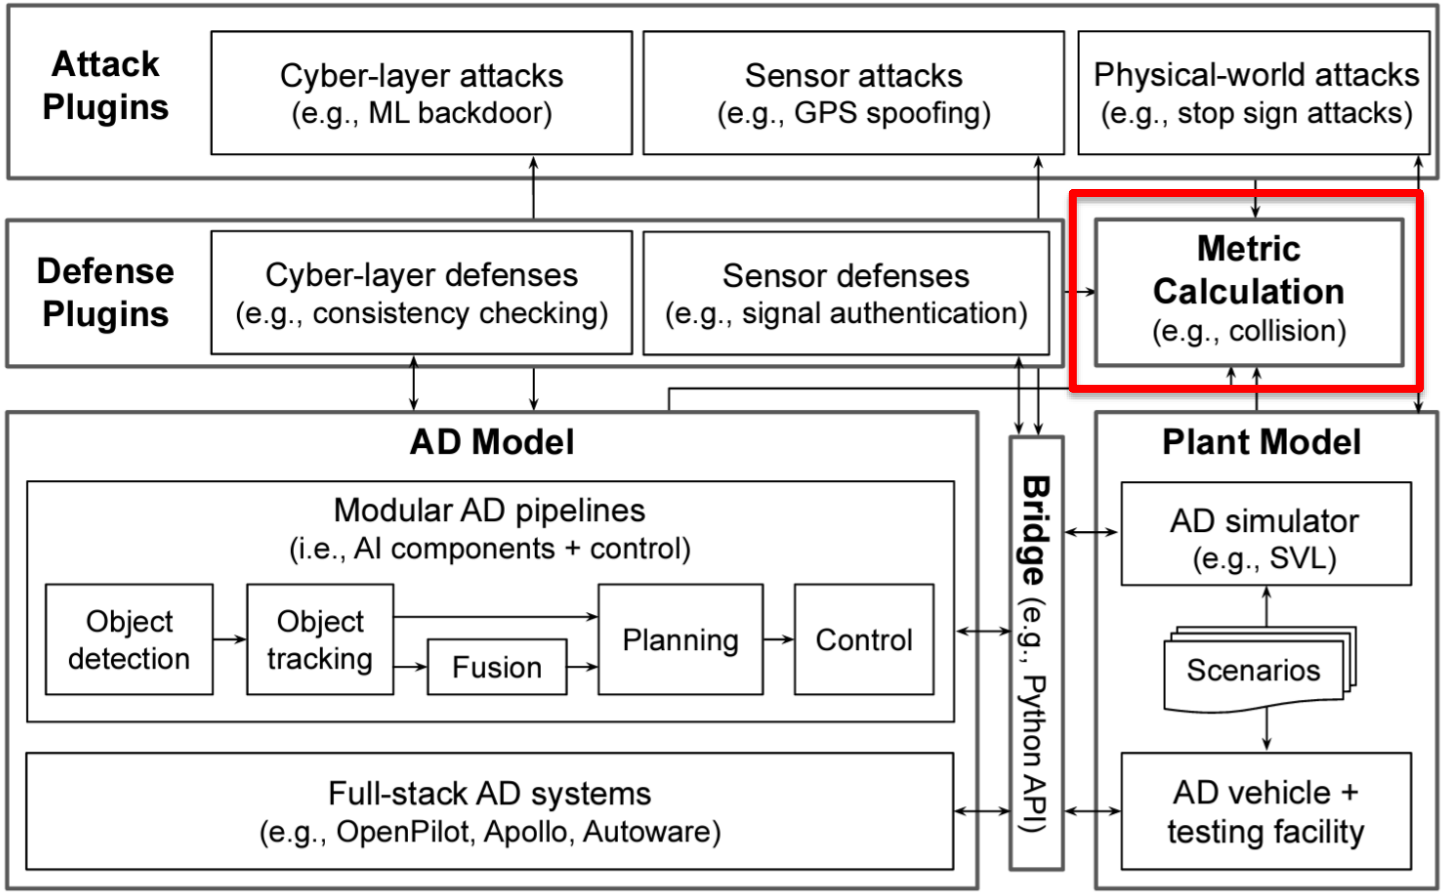

- **Attack/defense plugins** in Python APIs
- **Plug & play modular AD design**
- **Standardized system-level eval metrics**

Component level metrics				System level metrics	
Mis detection rate (MSR)			Best MSR over 50 consecutive frames	Stop sign violation rate	Avg. distance exceed stop line (m)
near-range (<10m)	mid-range (10-20m)	far-range (20-30m)			

Our proposal: PASS (Platform for Autonomous driving Security and Safety)

- ***Open***, ***uniform*** & ***extensible*** system-driven evaluation platform

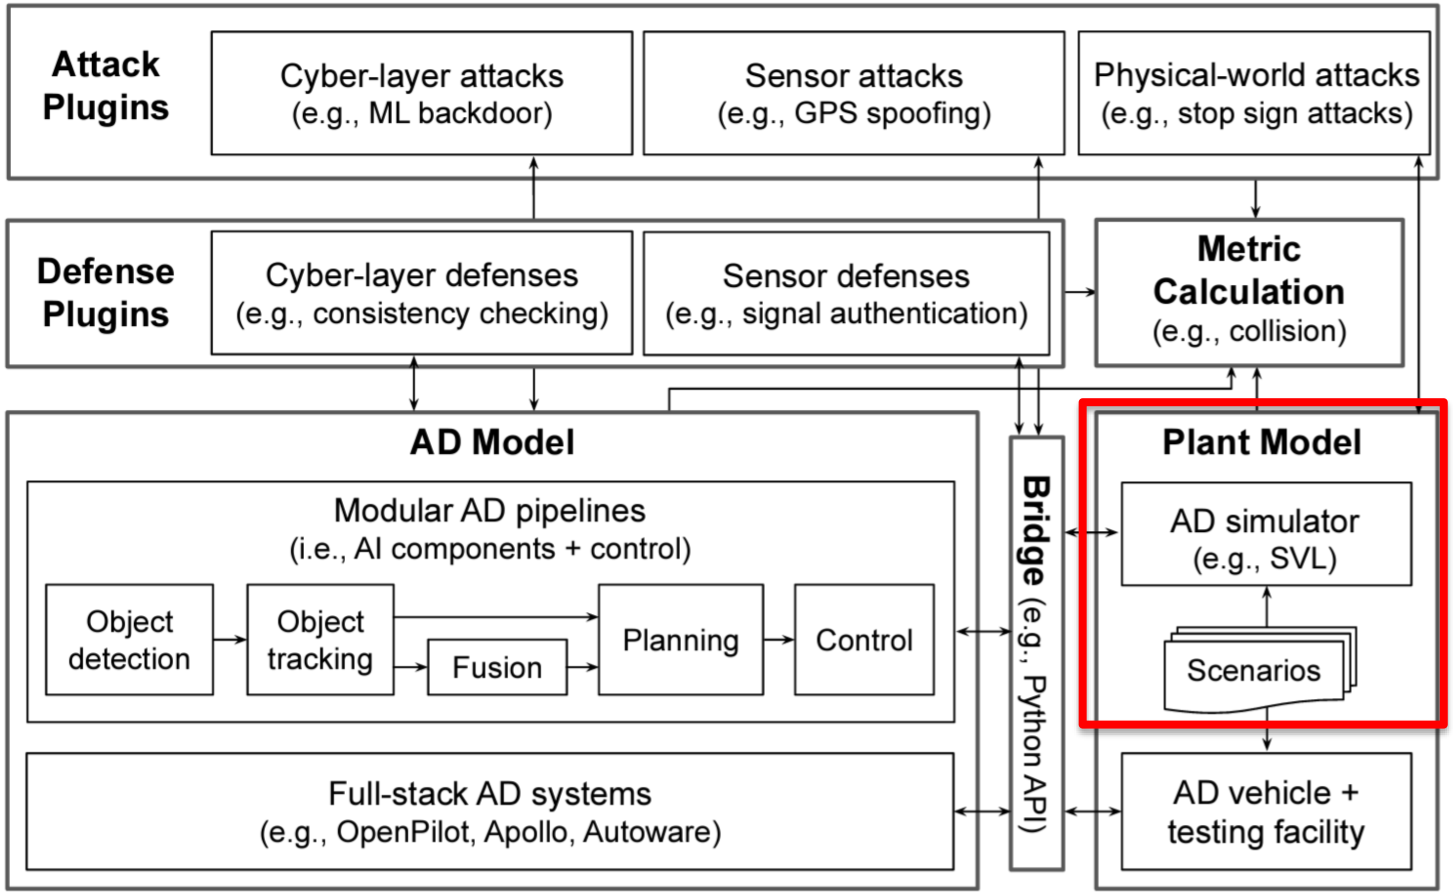

- **Attack/defense plugins** in Python APIs
- **Plug & play modular AD design**
- **Standardized system-level eval metrics**
- **Simulation-centric design** for affordability, accessibility, safety, flexibility & reproducibility

Our proposal: PASS (Platform for Autonomous driving Security and Safety)

- ***Open***, ***uniform*** & ***extensible*** system-driven evaluation platform

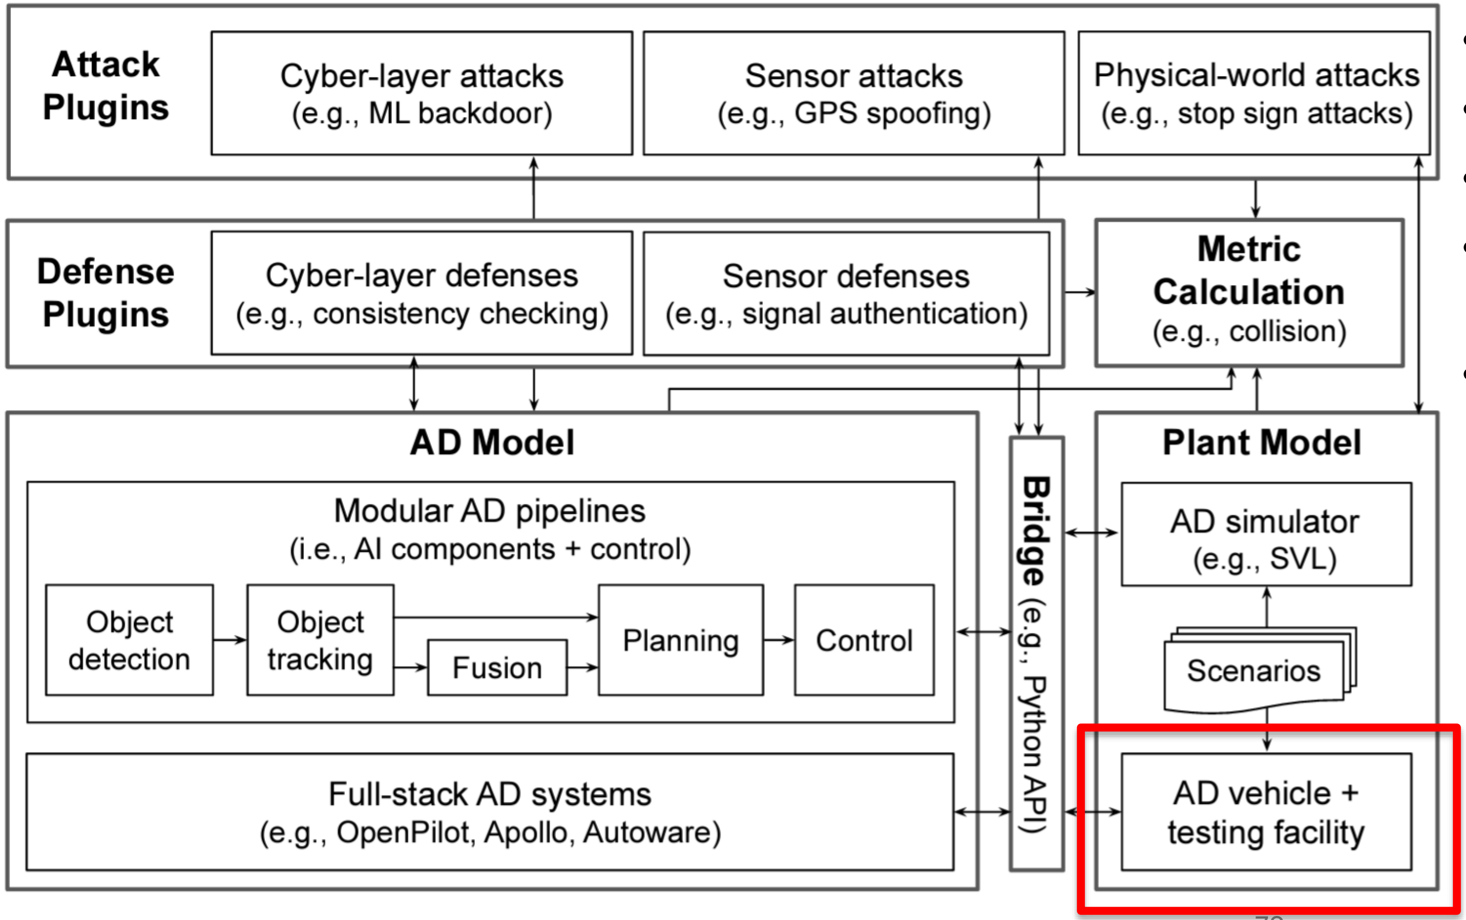

- **Attack/defense plugins** in Python APIs
- **Plug & play modular AD design**
- **Standardized system-level eval metrics**
- **Simulation-centric design** for affordability, accessibility, safety, flexibility & reproducibility
- **Test AD vehicles** for fidelity improvement

Available L4 AD vehicle & AD development chassis

In the process of solicitating community feedback!

- Do you think such a platform can be **useful/beneficial** to you (e.g., in *research*, *education*, *training*, and/or *outreaching*)?
- Any **features** you wish to *add/improve*?
- Any **concerns** you have regarding our current *design/vision*?
- Feel free to let us know your feedback anytime via ***the survey below*** or ***directly email me***!
  - Such info can also be found at the PASS website: https://sites.google.com/view/cav-sec/pass

Platform feedback Survey
(https://docs.google.com/forms/u/1/d/e/1FAIpQLSf94hAZMKCdW-L5uROGnFrmI7XUakxYNkSA9JZydPZUM4I5fg/viewform)

Our SoK effort
(https://arxiv.org/abs/2203.05314)

Conclusion

- **My group: *Actively developing*** research space on **autonomous system AI security**, currently most in AD & intelligent transportation
  - Collection of our efforts: https://sites.google.com/view/cav-sec
- ***Only the beginning*** of this research problem space
  - Now mostly on attack side, need more on ***defense & research infra*. sides**
  - To facilitate community building & broader impacts:
    - Co-found ***ACM/ISOC AutoSec (Automotive & Autonomous Vehicle Security) Workshop (2019 -)***, co-located **w/ *NDSS'21*** & ***'22***
    - Co-created ***DEF CON's first AutoDriving-themed hacking competition*** in 2021 (one of world's most famous hacker convention)
    - Served on **NIST focused group & panel on *AD AI test standards & metrics***

Sponsors:

NSF

Department of Transportation, United States of America

gm

TOYOTA

Qualcomm

AutoSec NDSS SYMPOSIUM

AutoDriving CTF

DEF CON 29

NIST National Institute of Standards and Technology

Conclusion

- **My group: *Actively developing*** research space on **autonomous system AI security**, currently most in AD & intelligent transportation
  - Collection of our efforts: **https://sites.google.com/view/cav-sec**
- ***Only the beginning*** of this research problem space
  - Now mostly on attack side, need more on ***defense & research infra.* sides**
  - To facilitate community building & broader impacts:
    - Co-found ***ACM/ISOC AutoSec (Automotive & Autonomous Vehicle Security) Workshop (2019 -)***, co-located **w/ *NDSS'21*** & ***'22***
    - Co-created ***DEF CON's first AutoDriving-themed hacking competition*** in 2021 (one of world's most famous hacker convention)
    - Served on **NIST focused group & panel on *AD AI test standards & metrics***
  - **Happy to chat more & form collaborations!**

Sponsors:

Contact

Alfred Chen (alfchen@uci.edu)

Homepage:https://www.ics.uci.edu/~alfchen/

AS²Guard Autonomous & Smart Systems Guard Research Group