Perception error modelling for autonomous driving

Andrea Piazzoni, Jim Cherian,
Martin Slavik, Justin Dauwels.
Cetran ERI@N – Interdisciplinary Graduate School
Nanyang Technological University

Ceci n’est pas de pluie…

- **Raindrop** was wrongly detected as a **car** by Fast R-CNN

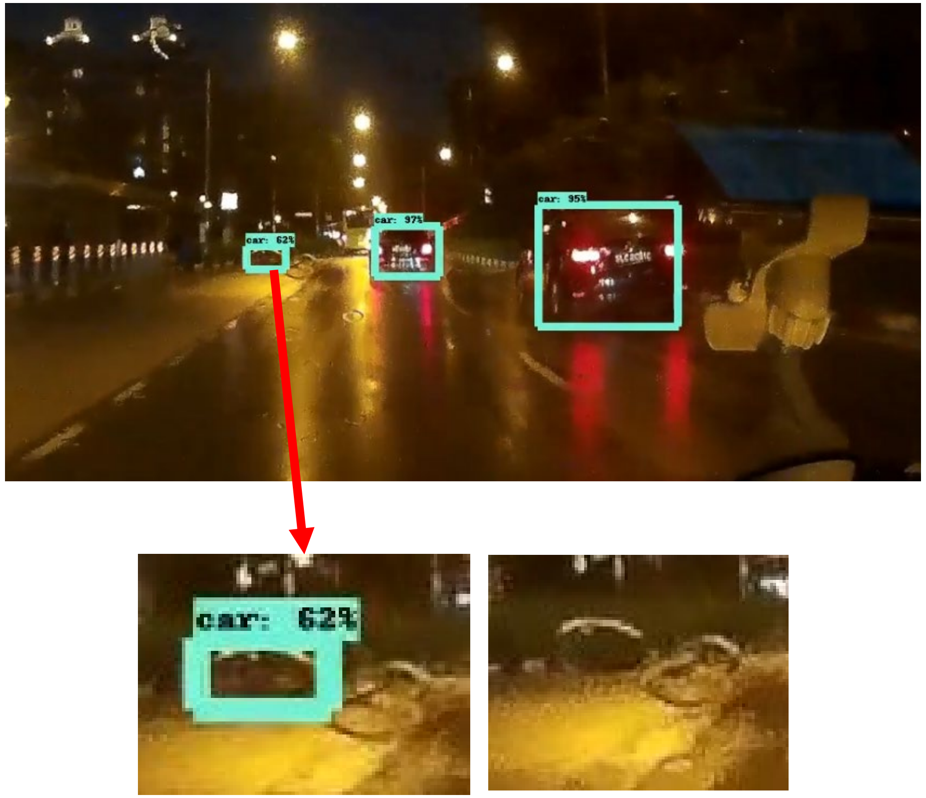

Fast R-CNN, Ross Girshick, 2015, https://arxiv.org/abs/1504.08083

Driverless Car Perception System

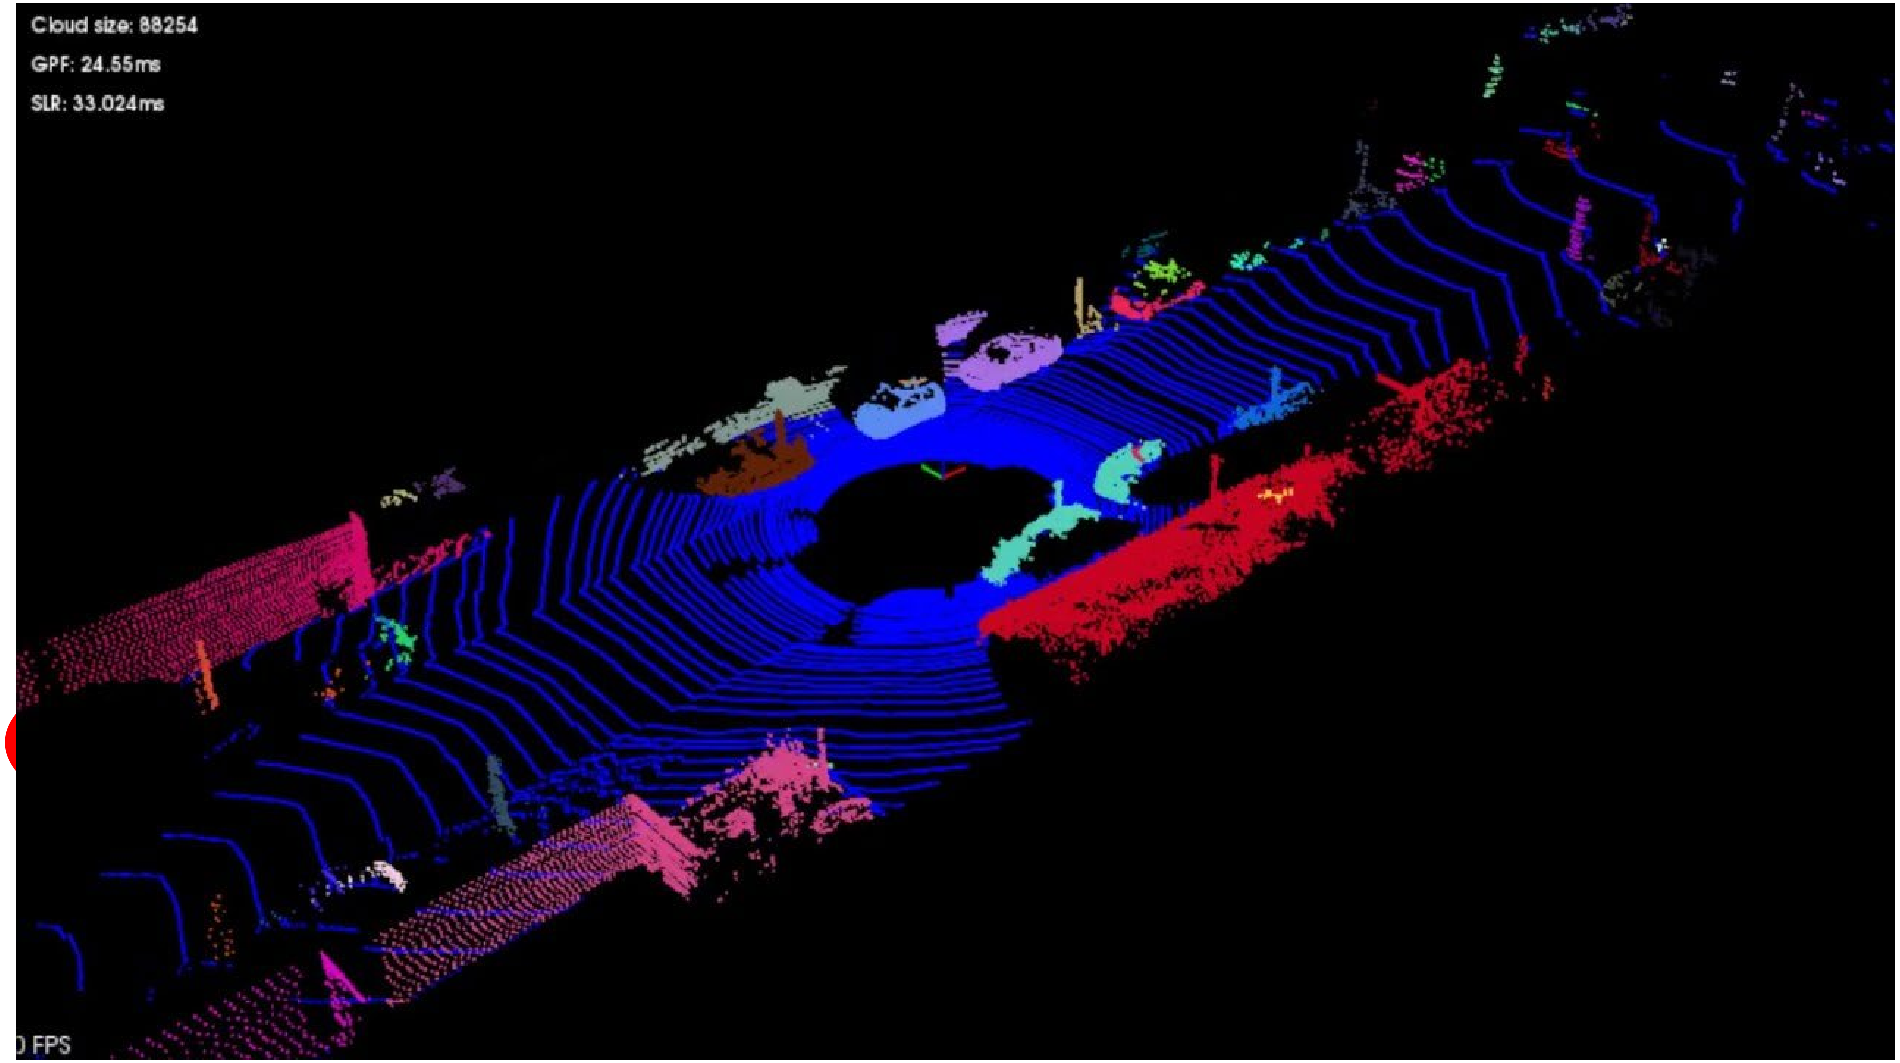

LIDAR = **Li**ght **D**etection **a**nd **R**anging

Driverless Car Perception System

What the Car Sees

The car’s sensors gather data on nearby objects, like their size and rate of speed. It categorizes the objects — as cyclists, pedestrians or other cars and objects — based on how they are likely to behave.

~~**LIDAR UNIT**~~

~~Constantly spinning, it uses laser beams to generate a 360-degree image of the car’s surroundings.~~

CAMERAS

Uses parallax from multiple images to find the distance to various objects. Cameras also detect traffic lights and signs, and help recognize moving objects like pedestrians and bicyclists.

RADAR SENSORS

Measure the distance from the car to obstacles.

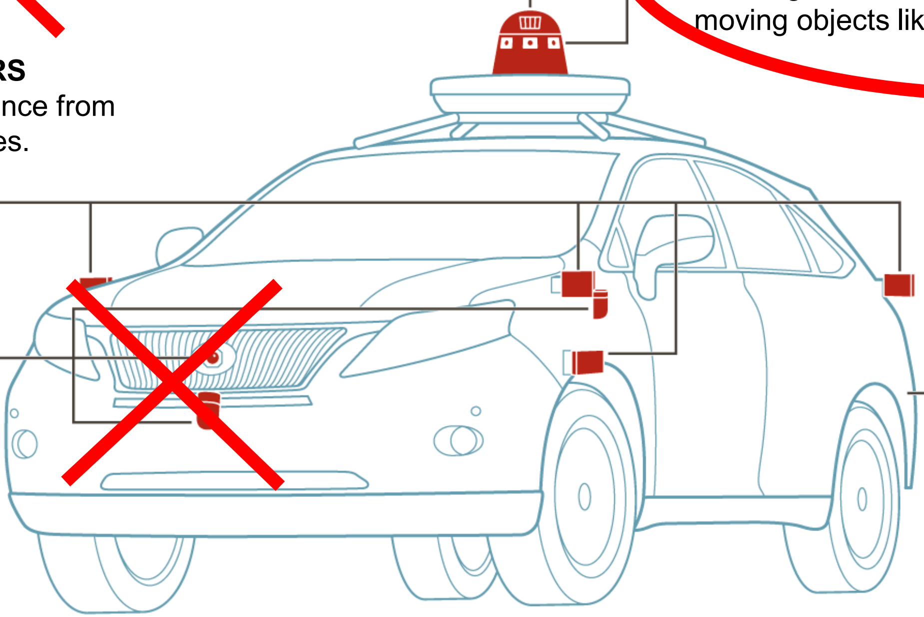

ADDITIONAL LIDAR UNITS

MAIN COMPUTER (LOCATED IN TRUNK)

Analyzes data from the sensors, and compares its stored maps to assess current conditions.

Perception via camera only can be vulnerable

THE DAILY NEWSLETTER
Sign up to our daily email newsletter

NewScientist

News Technology Space Physics Health Environment Mind Video | Tours Events Jobs

Tesla driver dies in first fatal autonomous car crash in US

TECHNOLOGY 1 July 2016

Source: https://www.newscientist.com/article/2095740-tesla-driver-dies-in-first-fatal-autonomous-car-crash-in-us/

Perception via camera only can be vulnerable

- The First **Driver Casualty** in Self-Driving Car Accident
  - 7th May 2016, Florida US
  - Tesla Model S, Autopilot mode
  - Driver passed away
  - The system didn't distinguish the white truck against the brightly lit sky, and failed to apply brakes.

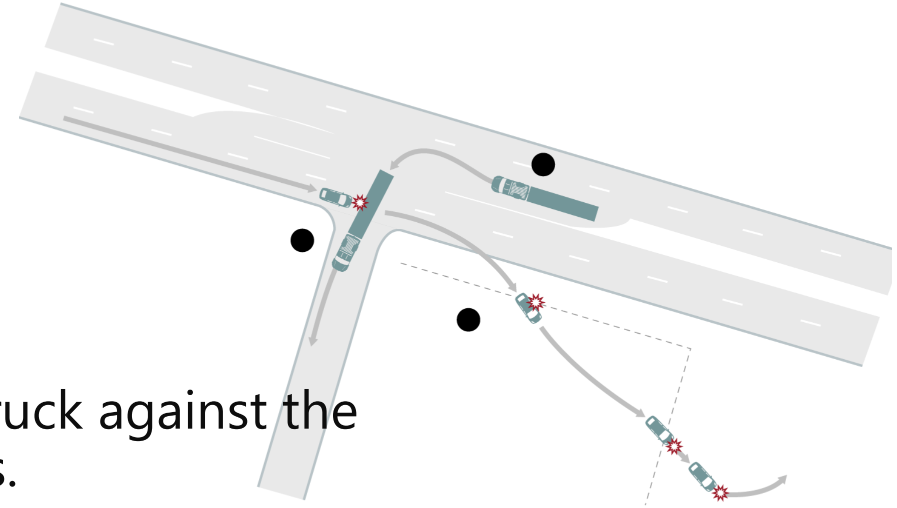

→ **Could have been prevented if the system had LIDAR sensors**

Source: https://www.newscientist.com/article/2095740-tesla-driver-dies-in-first-fatal-autonomous-car-crash-in-us/

1. Research Question

Can an autonomous vehicle make safe and robust decisions ***despite perception errors****?*

1. Contributions

3 areas:

- *Simulation Environment:* we propose a methodology to include perception errors in a simulation pipeline.

- *Perception evaluation:* we propose a novel holistic approach for the modeling and the evaluation of the capabilities of a perception subsystem.

- *Safety of AVs:* we aim to reach non-trivial conclusions regarding AV safety considering different settings and environmental conditions.

PEM: Perception Error Model for Virtual Testing of Autonomous Vehicles

Andrea Piazzoni, Jim Cherian, Justin Dauwels, *Senior Member, IEEE*, and Lap-Pui Chau, *Fellow, IEEE*

***Abstract*— Even though virtual testing of Autonomous Vehicles (AVs) has been well recognized as essential for safety assessment, AV simulators are still undergoing active development. One particular challenge is the problem of including the Sensing and Perception (S&P) subsystem into the virtual simulation loop in an efficient and effective manner. In this article, we define Perception Error Models (PEM), a virtual simulation component that can enable the analysis of the impact of perception errors on AV safety, without the need to model the sensors themselves. We propose a generalized data-driven procedure towards parametric modeling and evaluate it using Apollo, an open-source driving software, and nuScenes, a public AV dataset. Additionally, we implement PEMs in SVL, an open-source vehicle simulator. Furthermore, we demonstrate the usefulness of PEM-based virtual tests, by evaluating camera, LiDAR, and camera-LiDAR setups. Our virtual tests highlight limitations in the current evaluation metrics, and the proposed approach can help study the impact of perception errors on AV safety.**

***Index Terms*— Autonomous vehicles, computer vision, vehicle safety, simulation.**

driving behavior. For example, if the perception uncertainty increases, the AV could reduce its speed and adopt a more defensive driving, thus maintaining an adequate level of safety. Nevertheless, failures in obstacle detection may still lead to undesirable behavior such as collisions, emergency maneuvers, or traffic rules violations. For instance, the leading cause of a 2018 AV fatal accident was determined to be a perception error that was not adequately handled [4]. Thus, a deeper understanding of how perception errors affect the AV response is necessary for safety assurance.

This connection between perception errors and AV response can be explored via a holistic testing approach, both on a test track and in virtual environments. In this paper, we concentrate on virtual tests. Virtual testing of AVs by simulations offers a safe and convenient way to validate safety [5]. However, how to effectively include perception modules in the simulation pipeline is an open question. A common approach in the industry is to employ high-fidelity models that represent the

IEEE Trans ITS, 2023, accepted

2. Evaluation Metrics (Computer Vision)

Evaluation Metrics: from Computer Vision (CV)

- Object Detection: mAP (IoU) [1]
- Object Tracking: MOTP, MOTA [2]

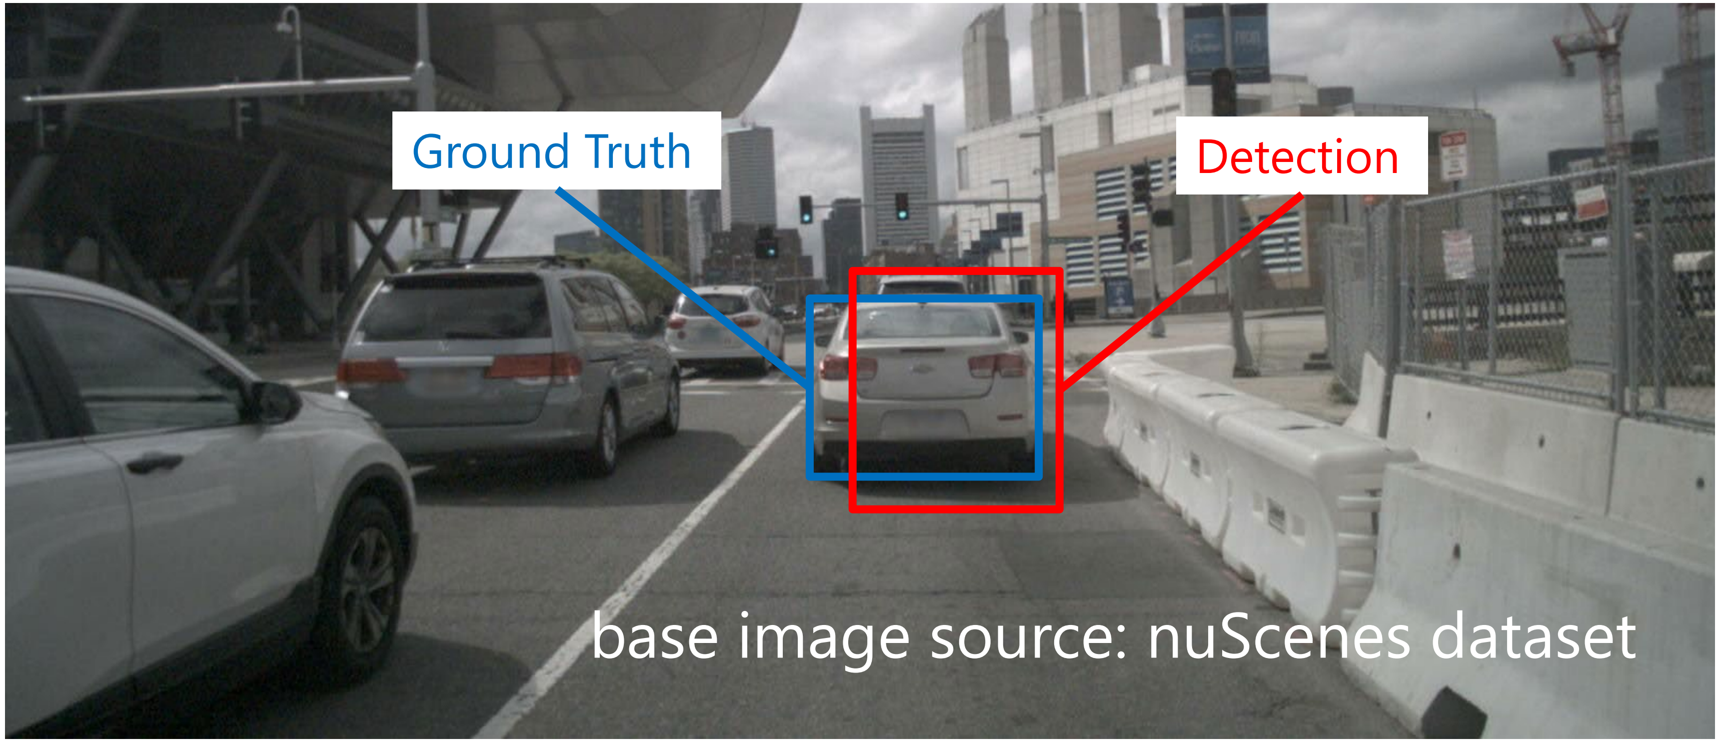

Intersection over Union

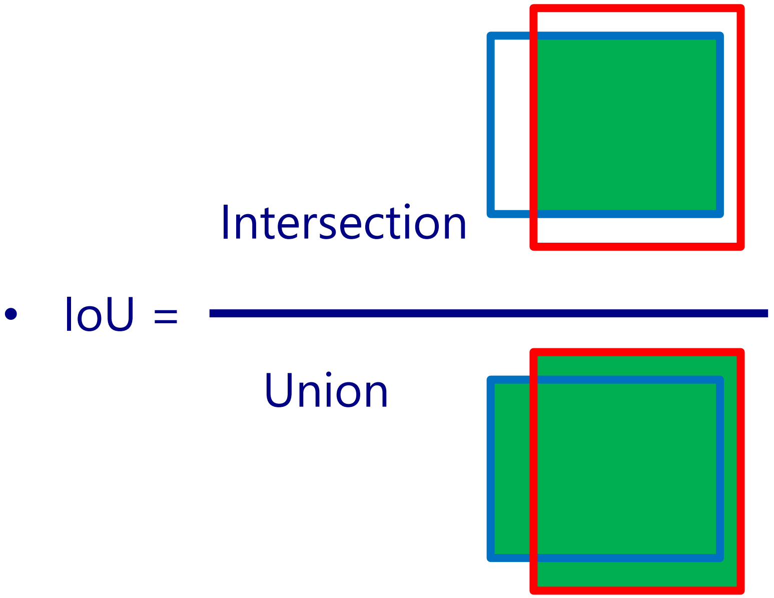

- IoU > threshold → True Positive

[1] Mark Everingham, Luc Van Gool, Christopher K I Williams, John Winn, Andrew Zisserman, M Everingham, L Van Gool, CKI Williams, J Winn, and A Zisserman. **The PASCAL Visual Object Classes (VOC) Challenge**. *Int J Comput Vis*, 88:303–338, 2010. doi: 10.1007/s11263-009-0275-4

[2] Keni Bernardin and Rainer Stiefelhagen**. Evaluating multiple object tracking performance: The CLEAR MOT metrics**. *Eurasip Journal on Image and Video Processing*, 2008, 2008. ISSN 16875176. doi: 10.1155/2008/246309.

2. Error Model

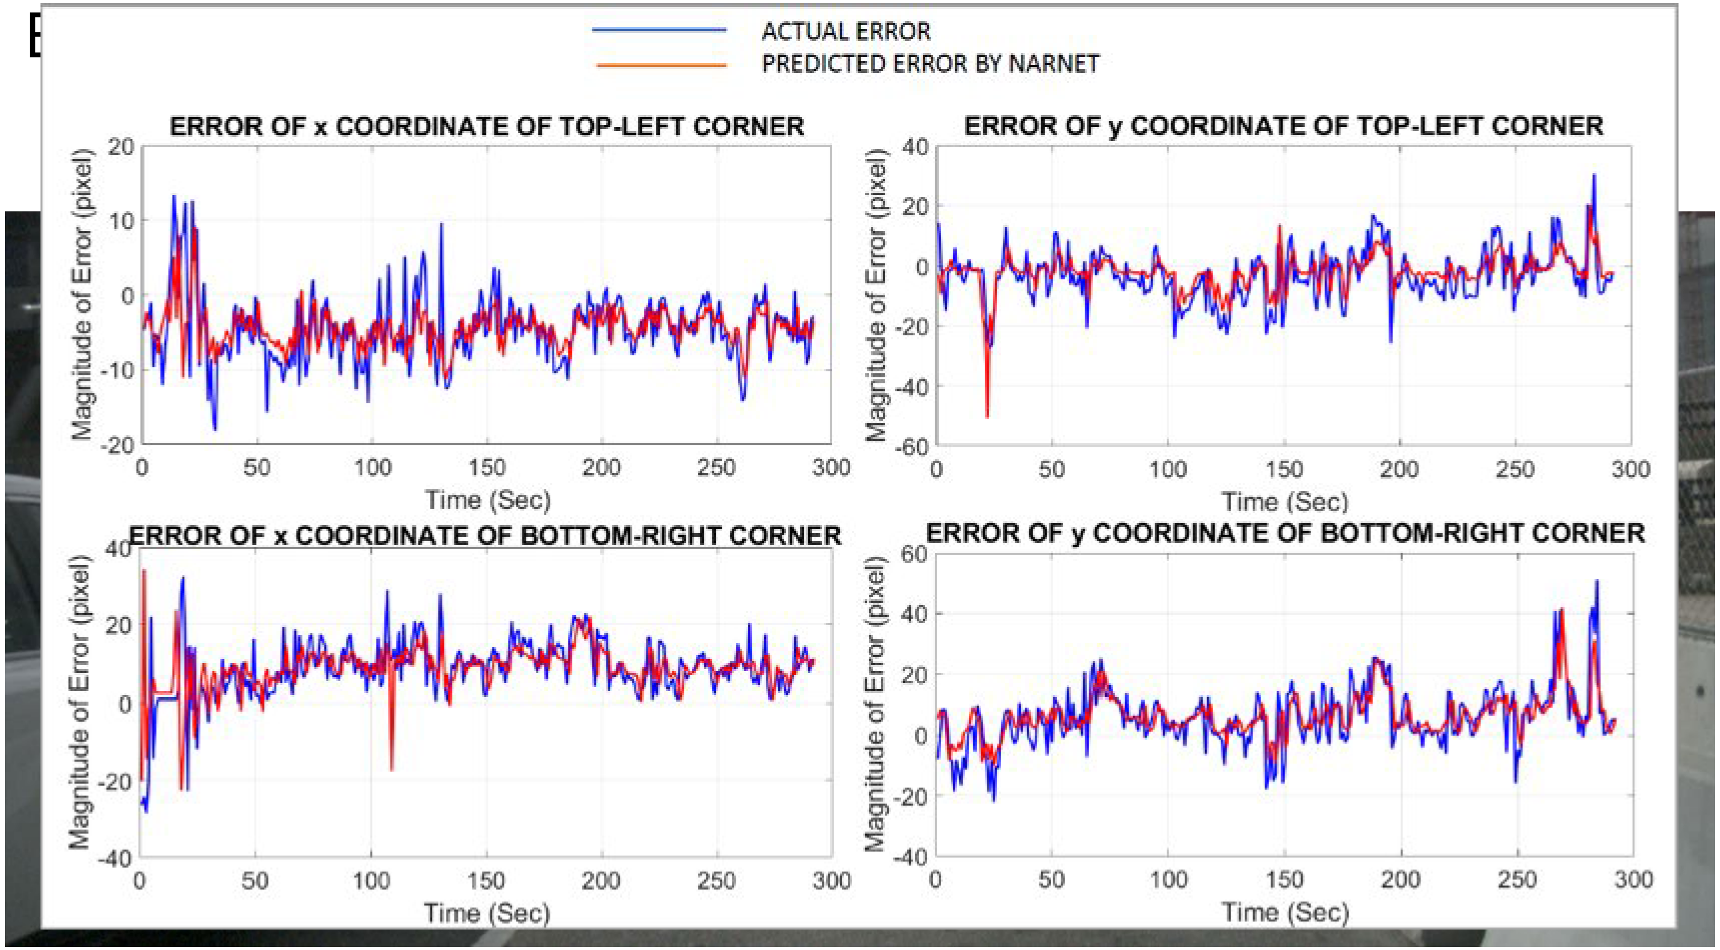

- Camera based
- 2D bounding boxes offsets
- Temporal Models

[3] Pallavi Mitra, Apratim Choudhury, Vimal Rau Aparow, Giridharan Kulandaivelu, and Justin Dauwels. **Towards Modeling of Perception Errors in Autonomous Vehicles**. In *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, 2018. ISBN 9781728103235. doi: 10.1109/ITSC.2018.8570015.

2. Issues of Evaluation Metrics

I1: Temporal considerations: short vs. long non detection intervals.

I2: Overlap Sensitivity: how sensitive is DP to spatial errors?

I3: Relevance of the objects: which ones are actively affecting the decision taken?

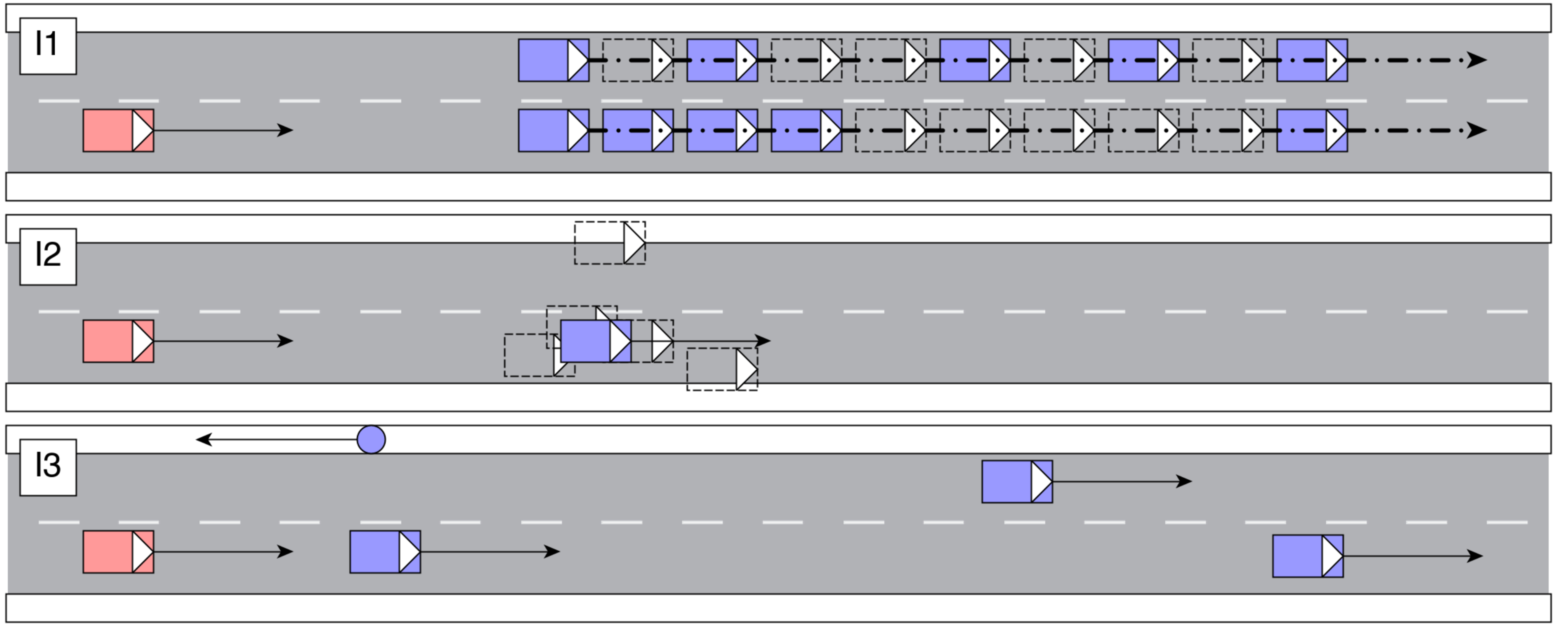

Remark: Current Evaluation Metrics do not provide insights on the kinds of errors.

2. Object detection over time

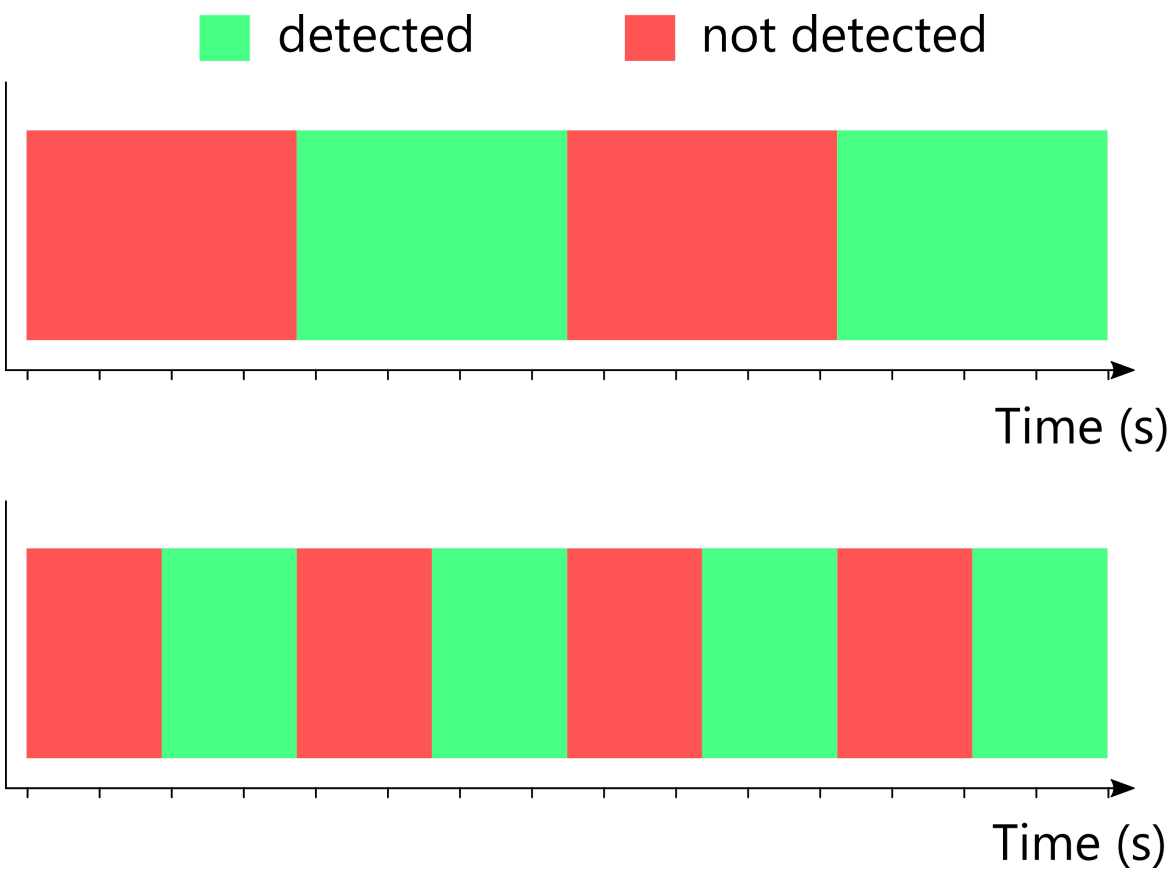

50% detection rate in both cases, but very different risks!

3. Proposed Approach

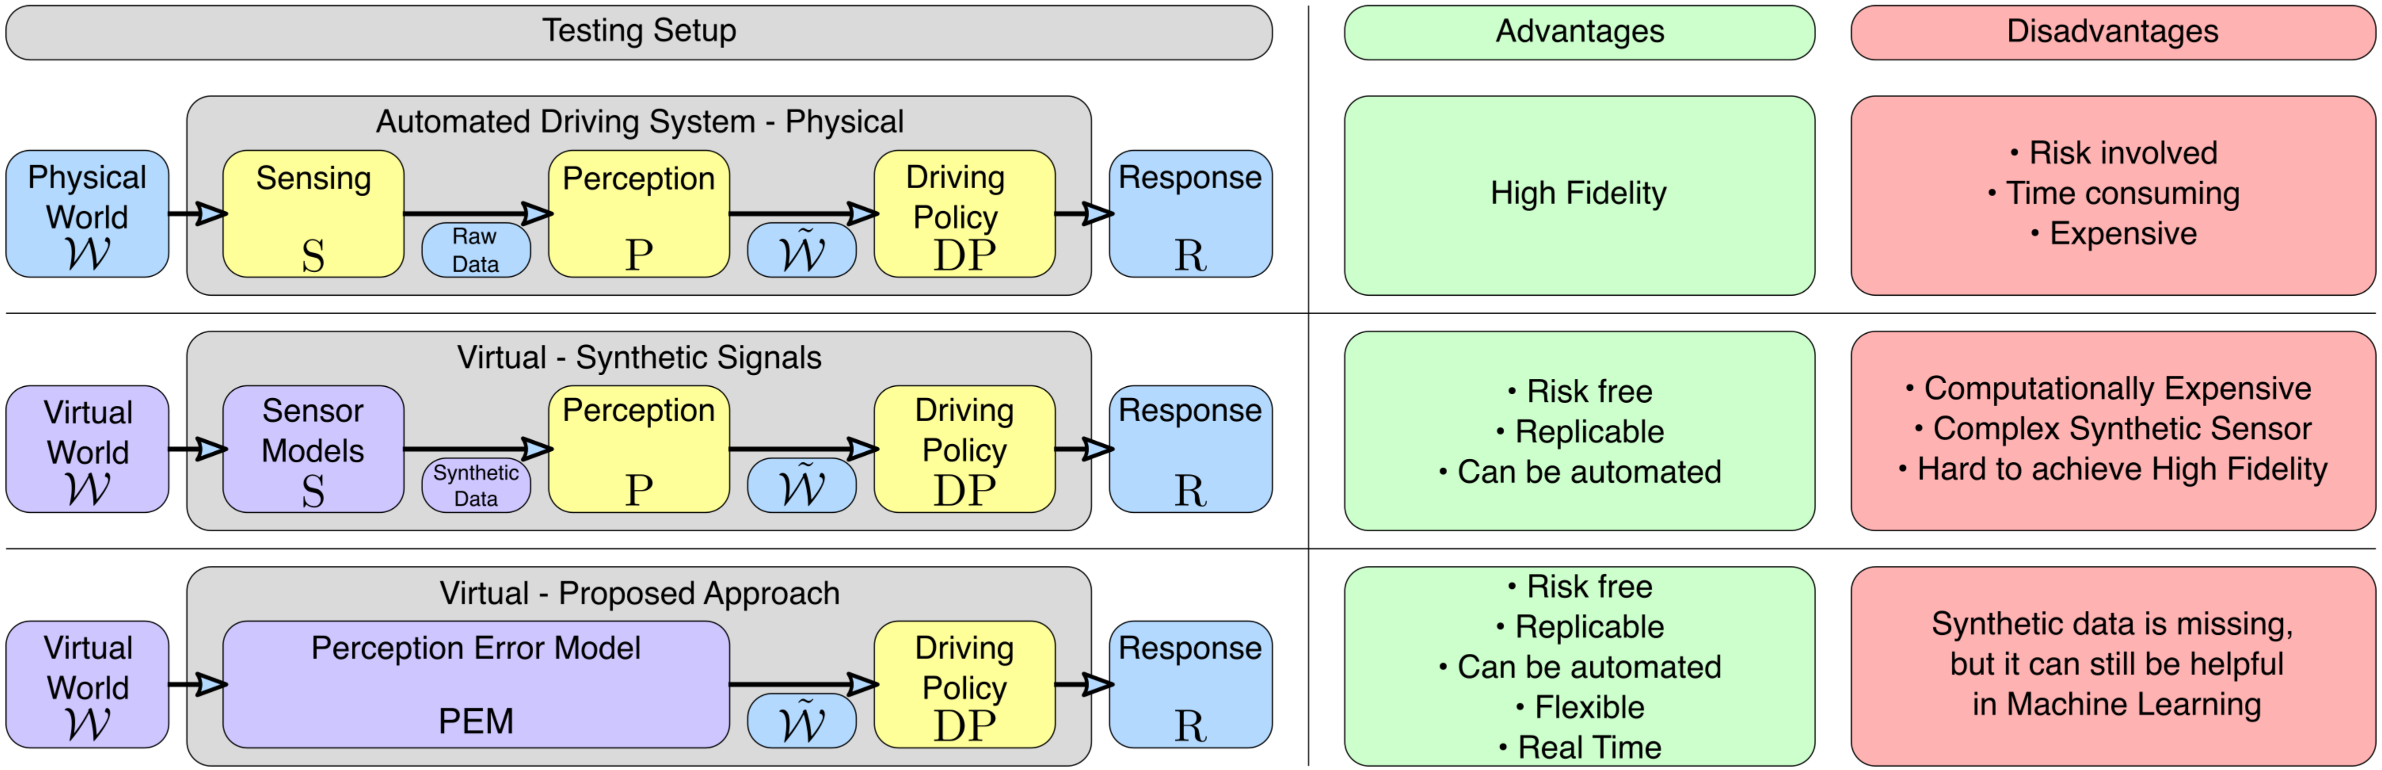

3. Dual Nature of PEMs

- **Explanatory** Model
  - Describe, and explain, the kinds of errors that can occur in the S&P subsystem.

$$\tilde{\mathcal{W}} = \mathcal{W} + \mathcal{E}$$

- **Generative** model:
  - Include a PEM in a simulation pipeline to replace the S&P subsystem.

$$\mathcal{W} \rightarrow \tilde{\mathcal{W}}$$

3. Experimental Setup: Architecture

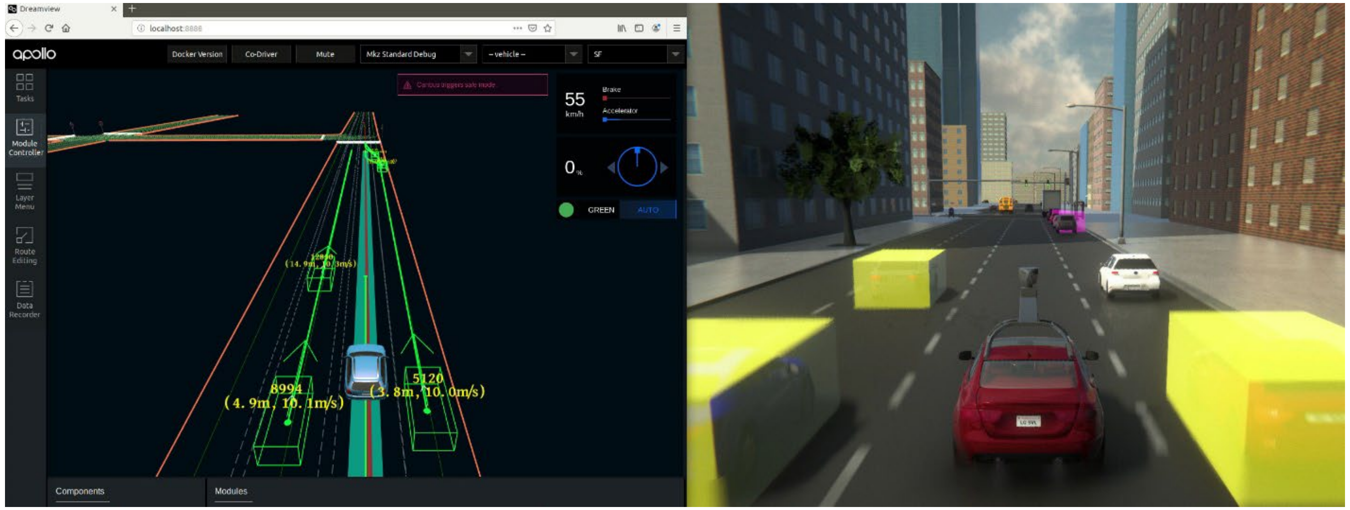

- a virtual vehicle driven by Apollo (ADS) in a virtual environment modeled in LGSVL simulator
- replaced Apollo's Perception with a virtual sensor (implementing PEM)
  - communicates over CyberRT
- Virtual sensor to
  - detect ground truth,
  - generate the error applying the PEM,
  - publish the Object Map
- automated test execution framework (Python-based)

4. Experimental Setup: Test Cases

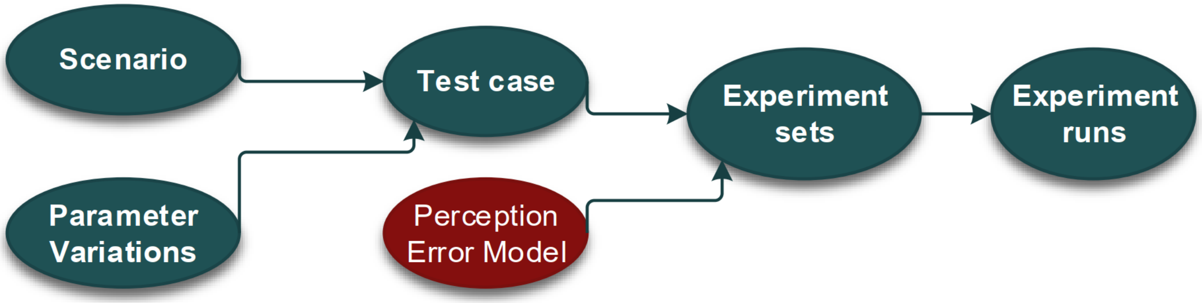

TC1-3: false negative errors modeled by means of Markov chains

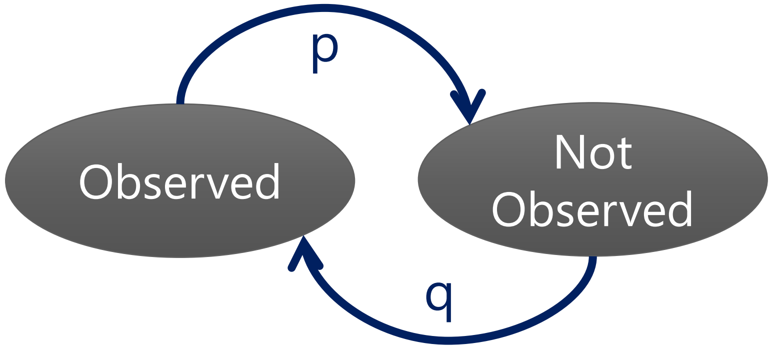

TC1-3

TC4-5

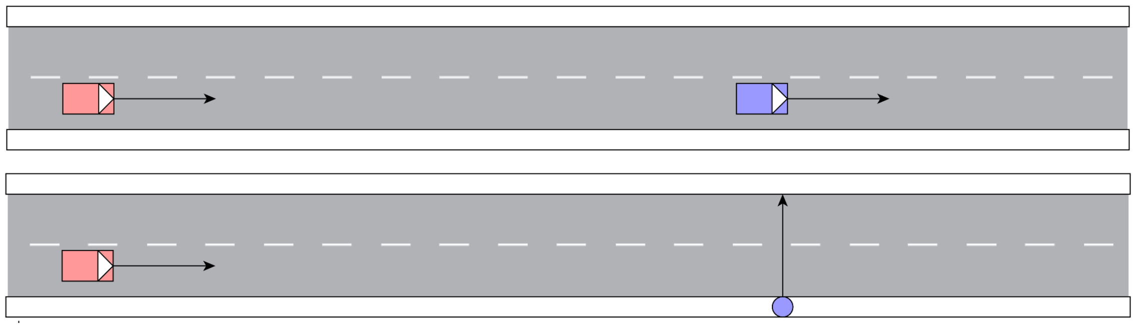

TC4-5a: Gaussian White Noise in polar coordinates (relative position):

- multiplicative noise on radius d as $\sigma_d \in [0\%, 12\%]$;
- additive noise on azimuth θ as $\sigma_\theta \in [0°, 1.5°]$

TC5b: Uber accident: perfect detection, but with tracking loss probability $p_{tl} \in [0, 1]$ for the previously detected obstacle.

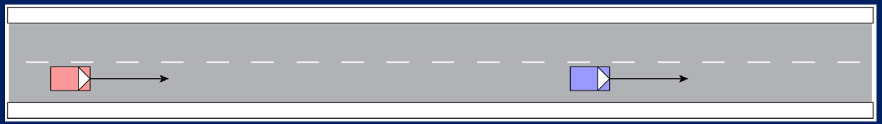

4. Results TC1-3: Detection Frequency

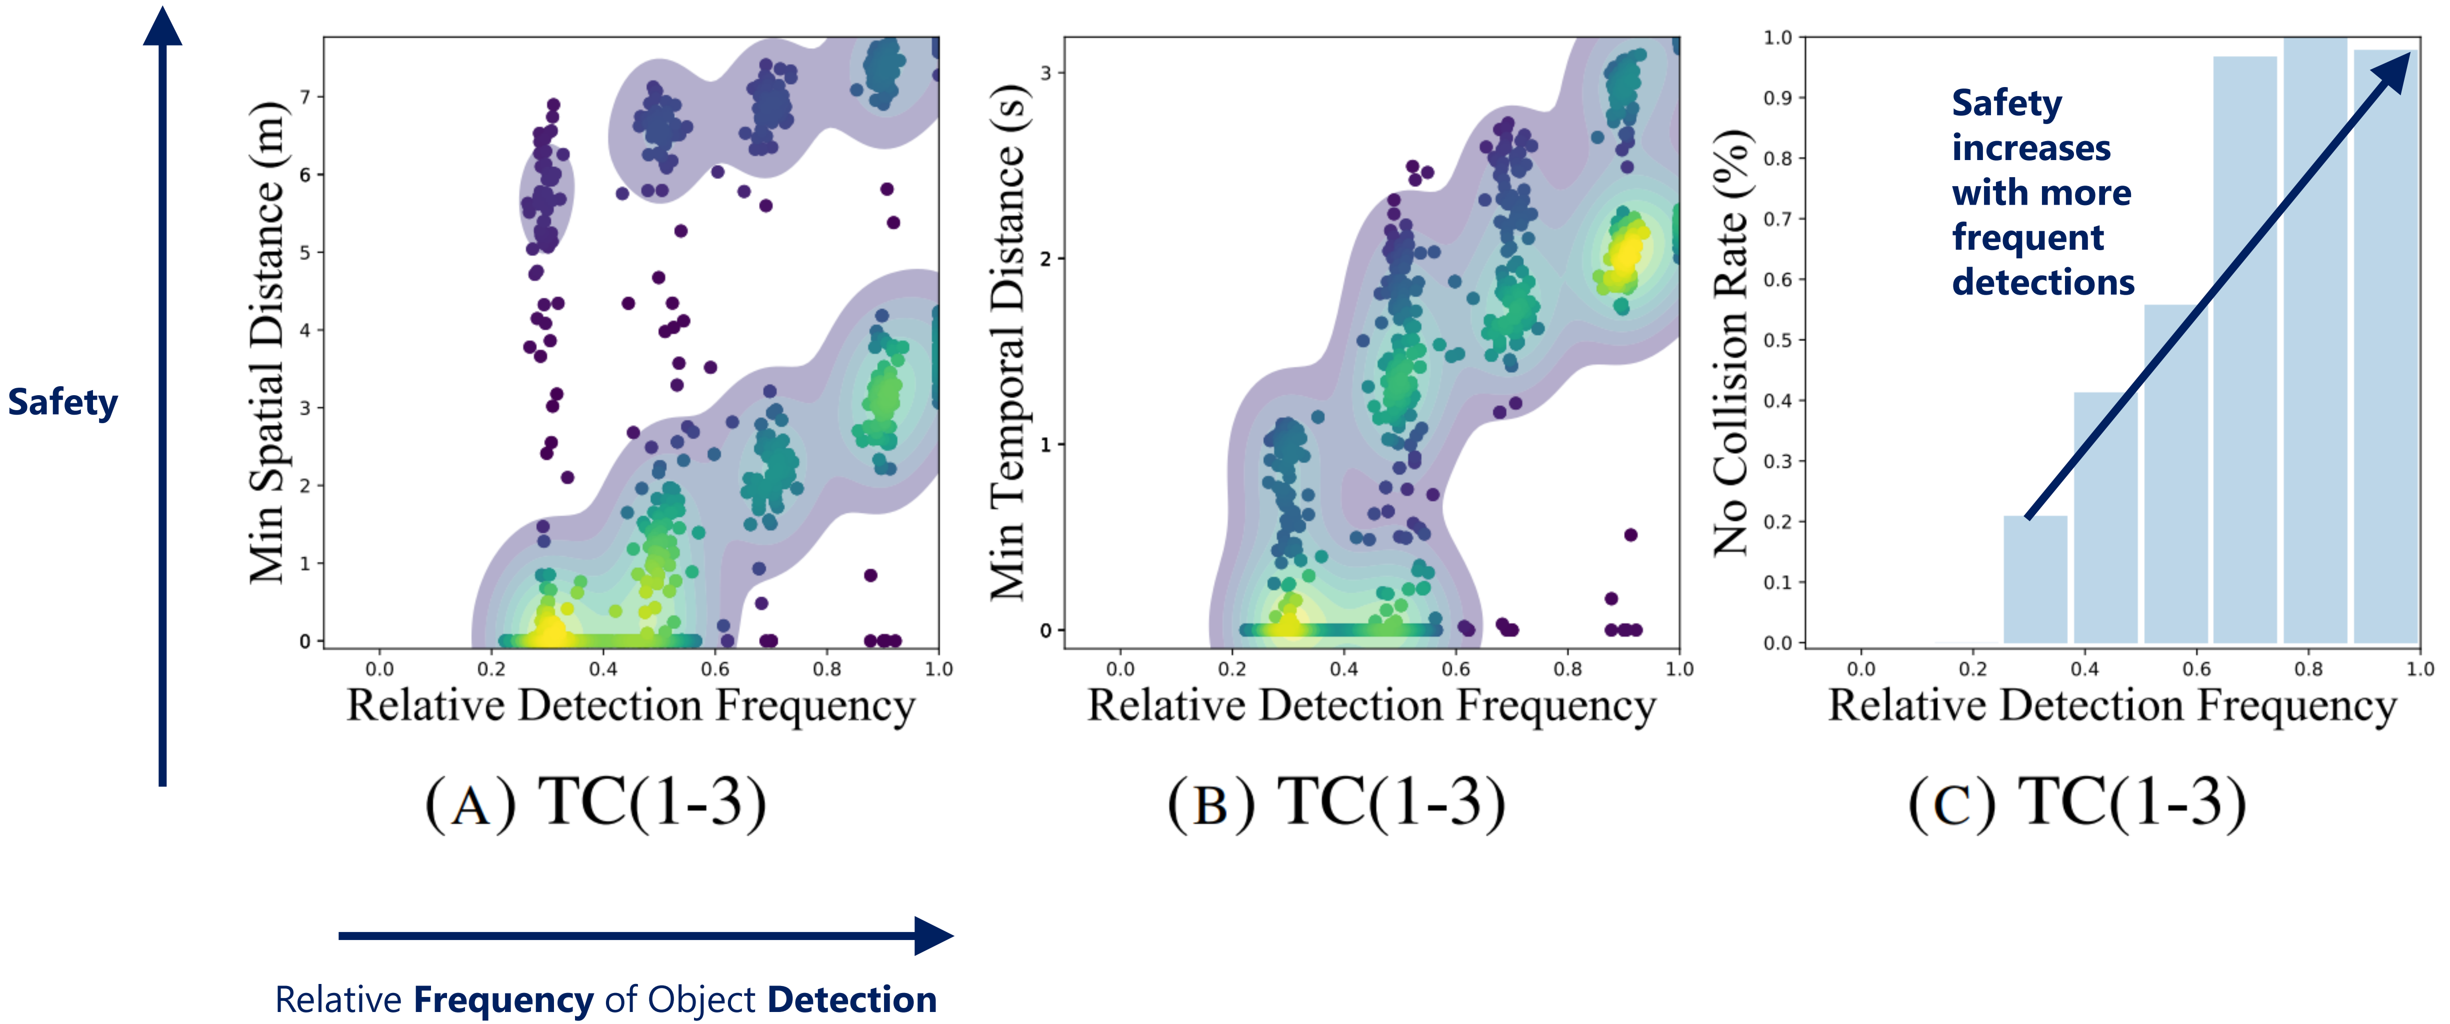

2. Object detection over time

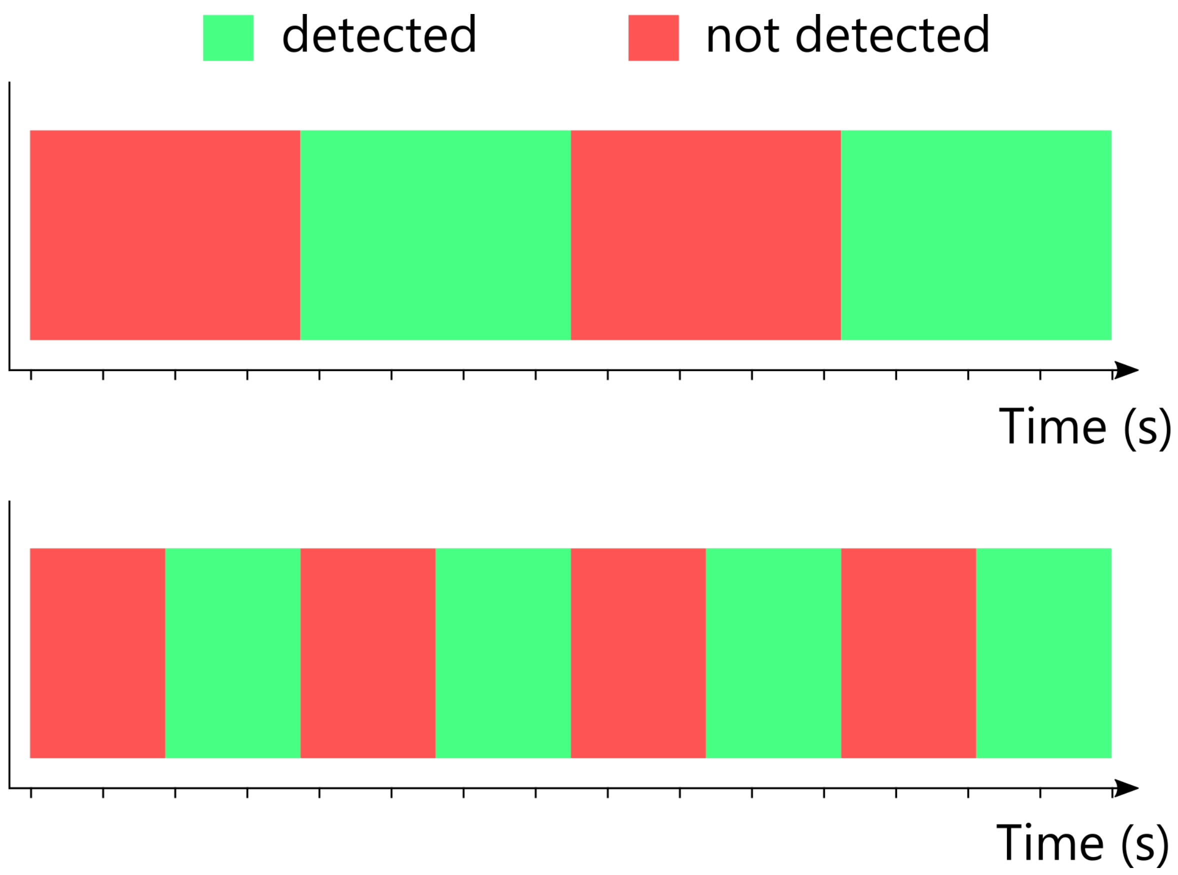

50% detection rate in both cases, but very different risks!

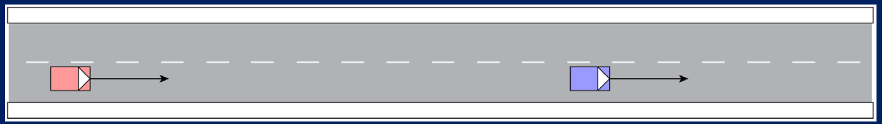

4. Results TC1-3: Non-detection interval length

Safety

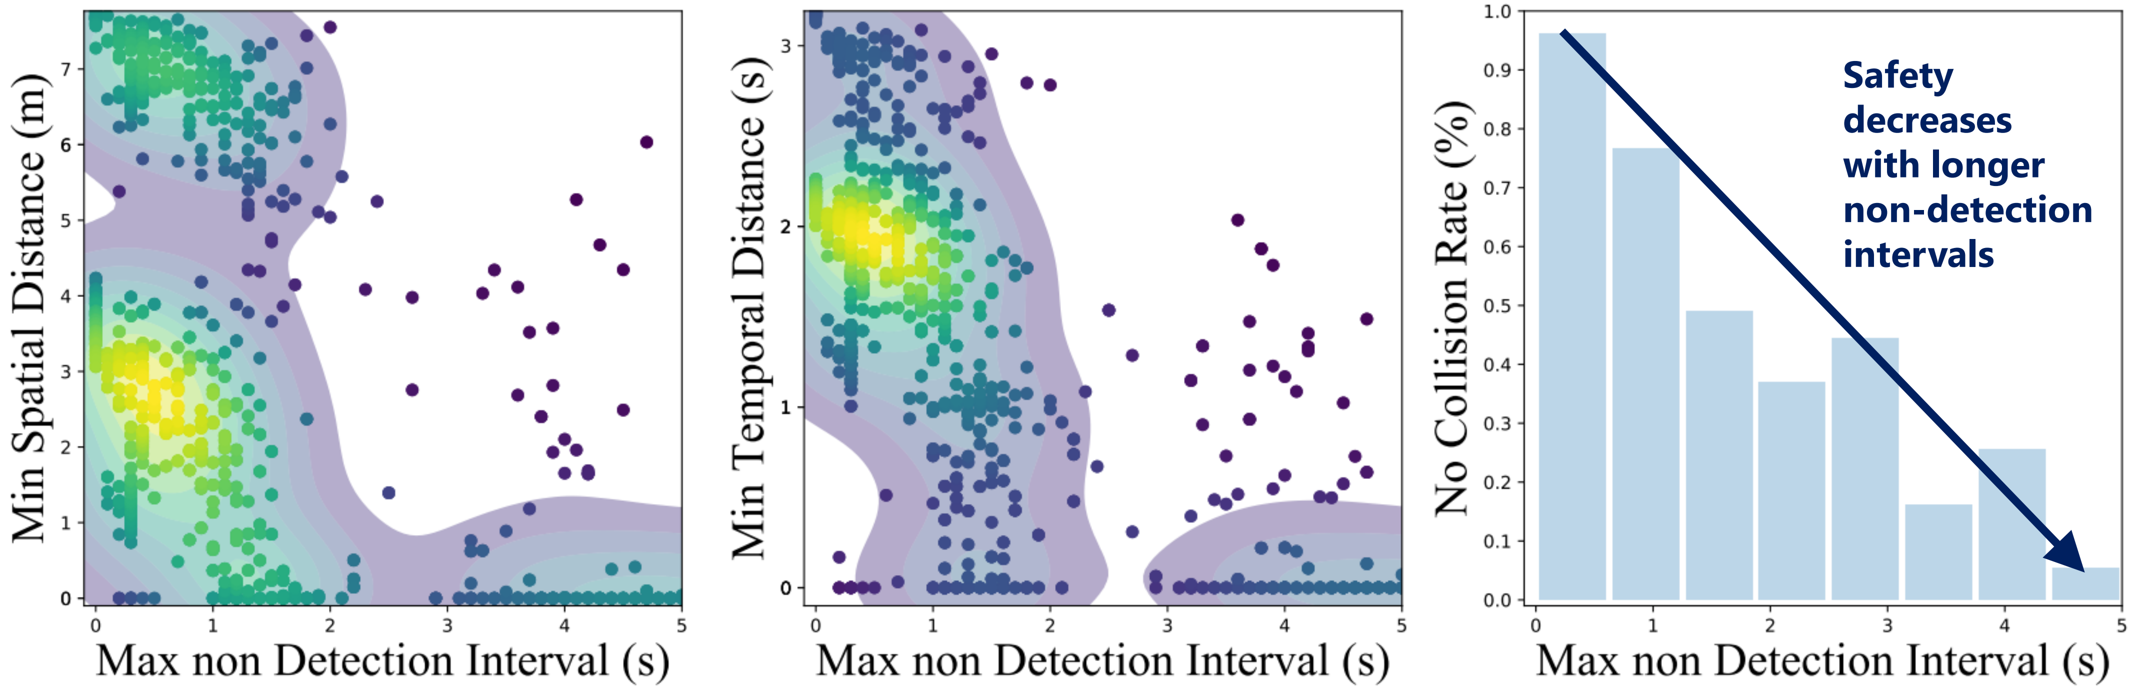

(D) TC(1-3) (E) TC(1-3) (F) TC(1-3)

Longest interval of
Non-detection of Objects

Longest non-detection interval is more relevant for safety

4. Results TC4-5

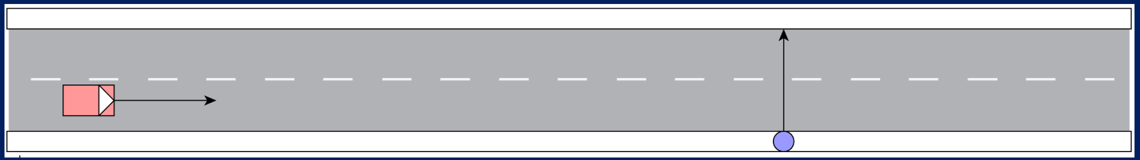

Safety

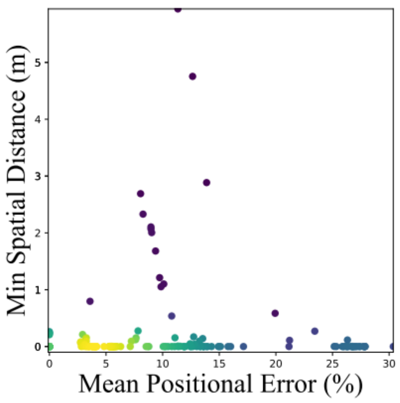

(G) TC4

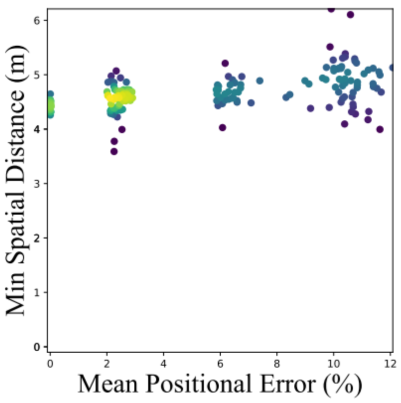

(H) TC5a

Safety

Safety decreases with increasing tracking loss

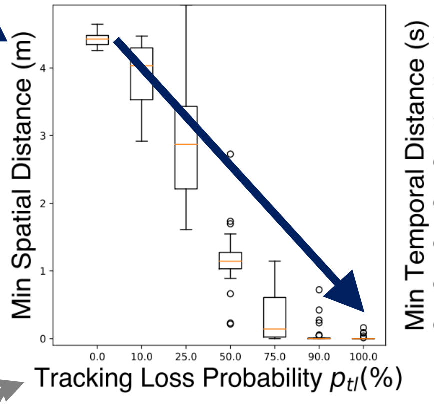

(I) TC5b

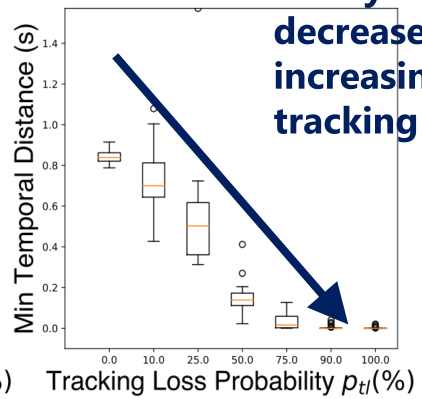

(J) TC5b

Uber accident

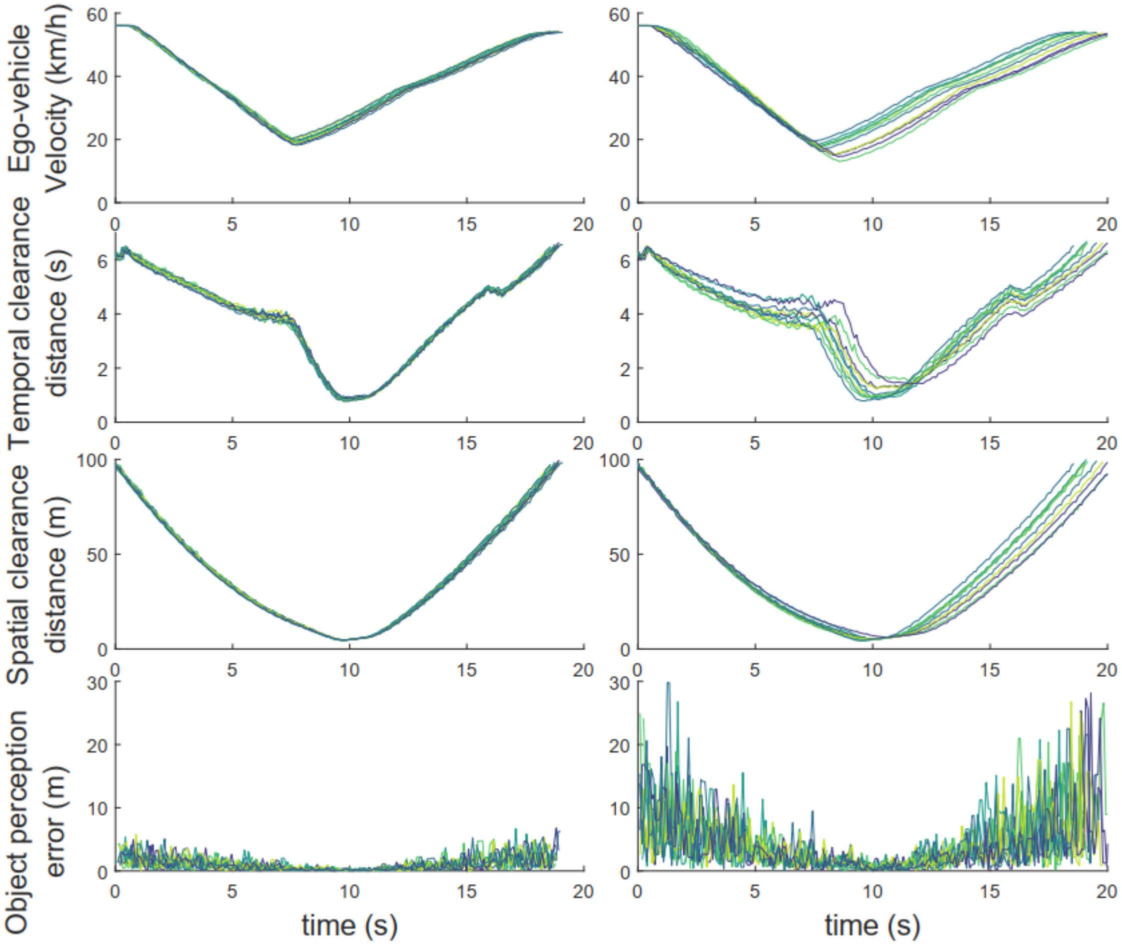

5. Data-driven models

NUSCENES by Motional

- In the previous study, the model was hand-crafted to inject controlled errors. In our latest study, we use the **nuScenes** dataset (sensor data and ground truth) and Apollo perception module.

- nuScenes:
  - 1000 scenes, ~20s each.
  - **1.1M bounding boxes** (ground truth)
  - ~19% pedestrians, ~57% vehicles, ~24% other obstacles (traffic cones, barriers, etc.)

- Two steps:
  - **Train** PEM on nuScenes dataset & Apollo's detections
  - With this PEM, **simulate** behavior of AV for test cases

5. nuScenes: Examples

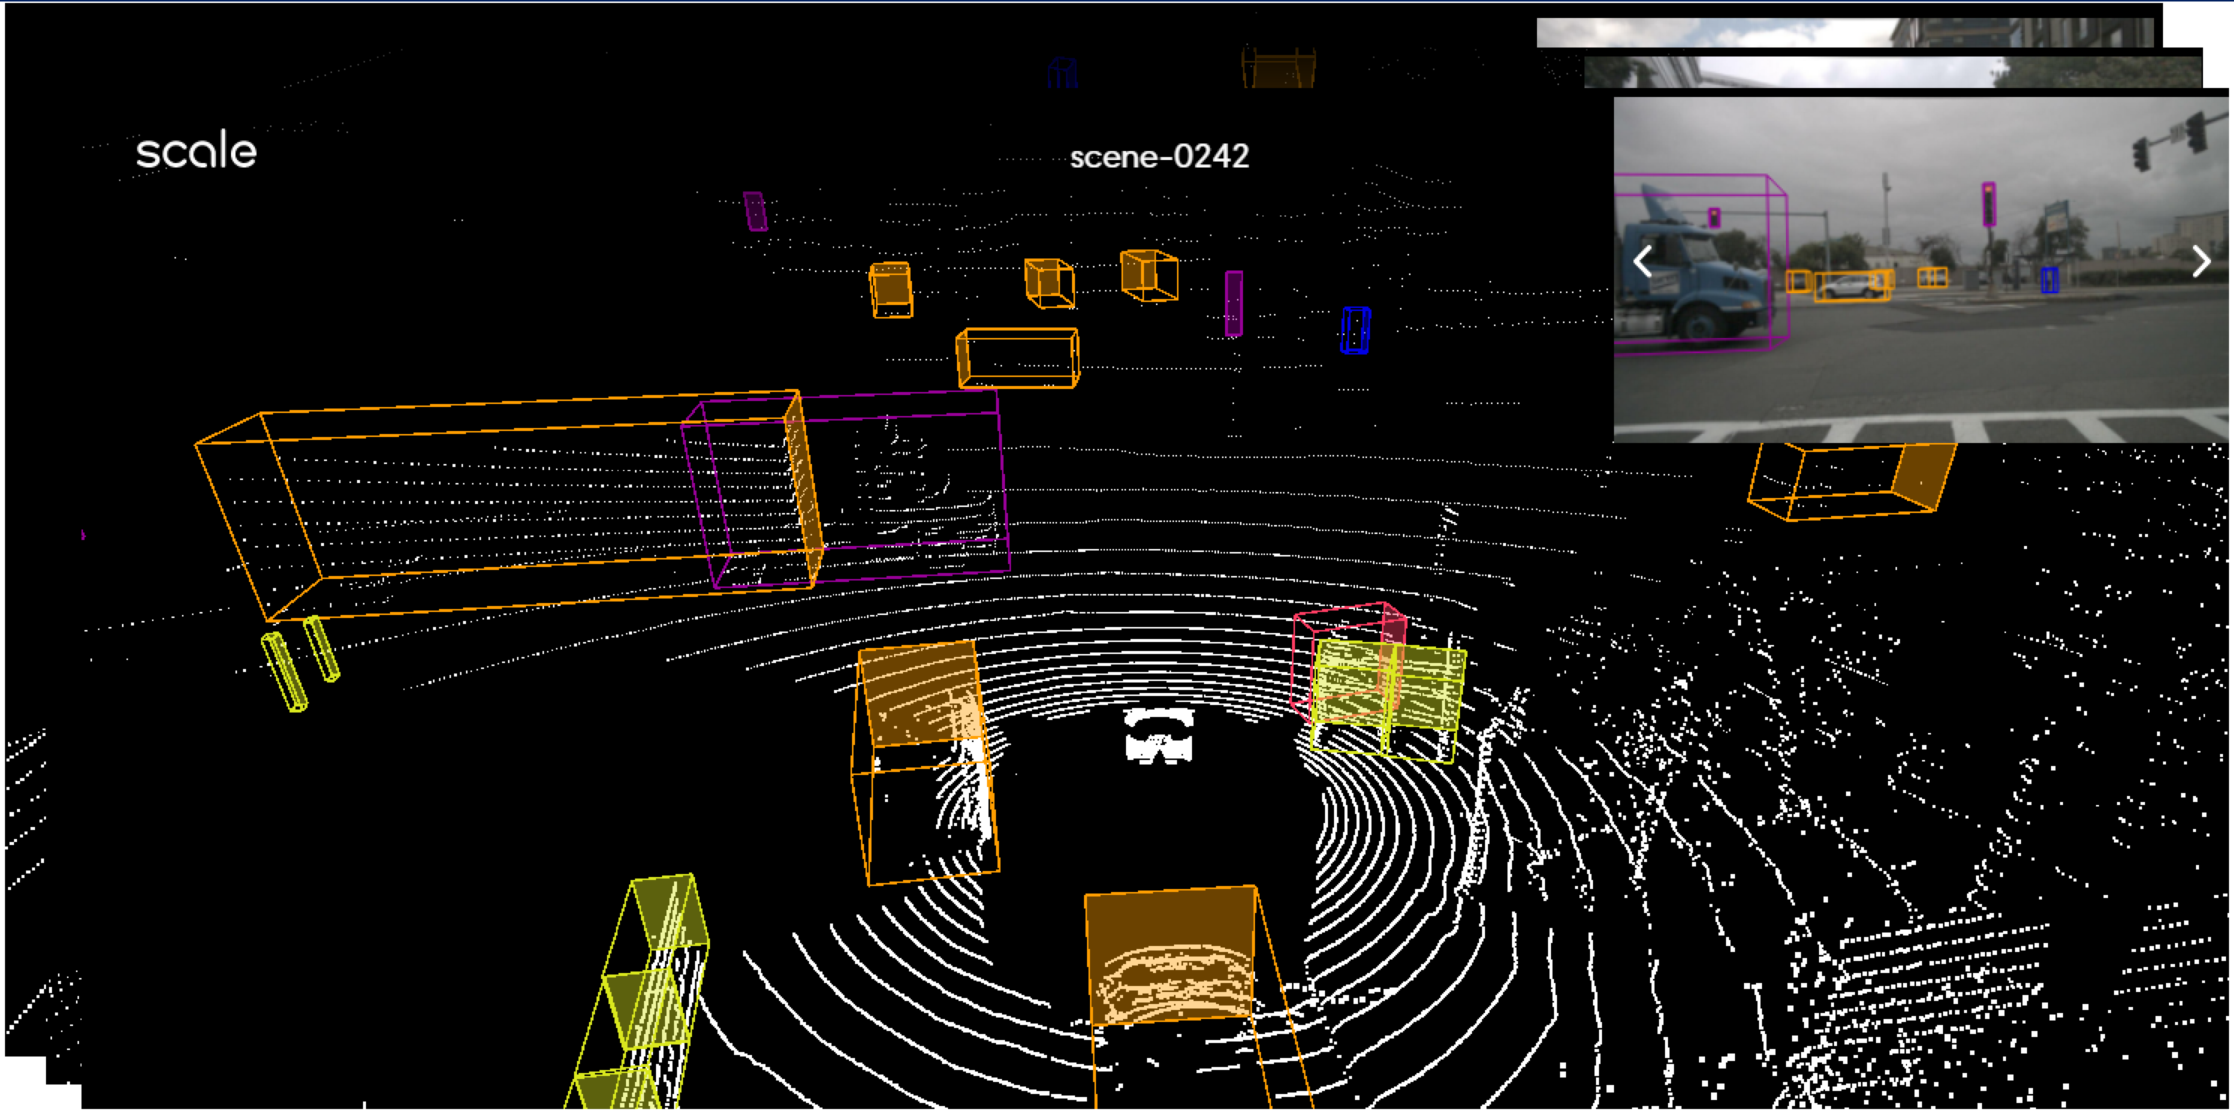

5. Training data for PEM: nuScenes and Apollo

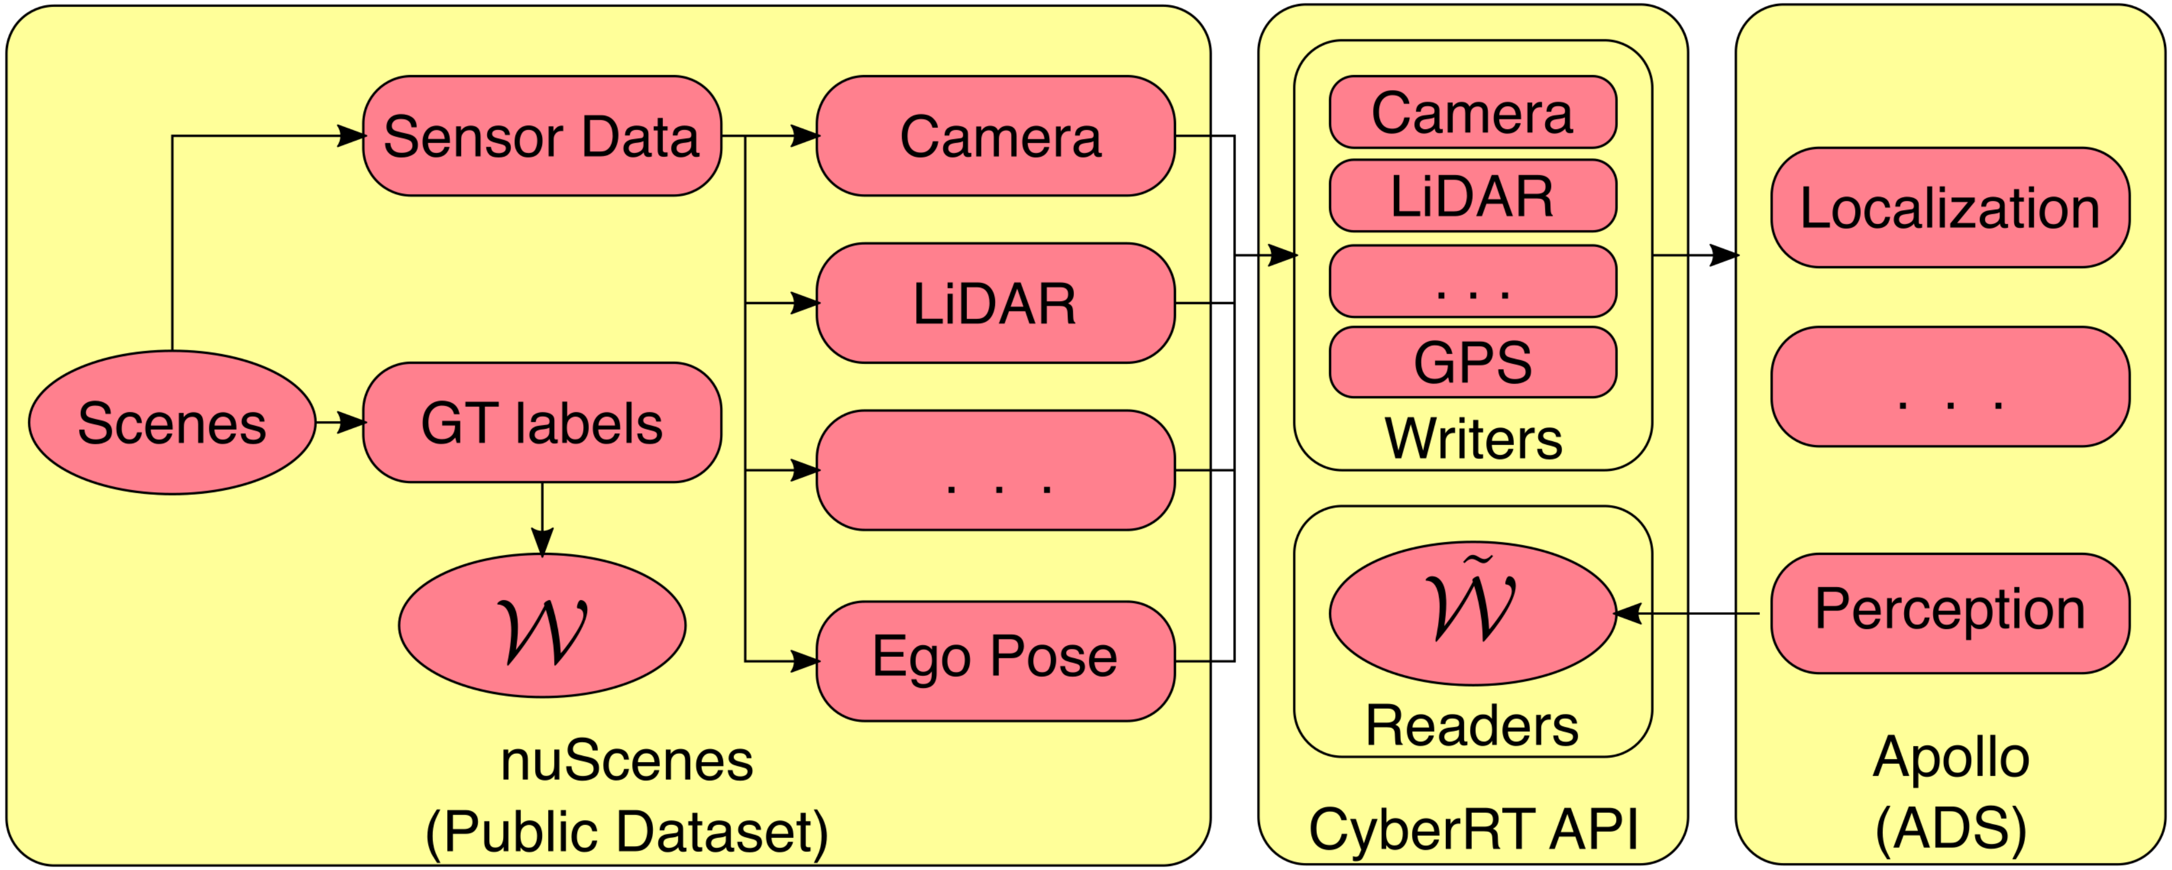

NUSCENES by Motional

Baidu 百度 | apollo

5. Three perception systems

ID	Sensor Configuration
CAM	Only frontal camera data
LID	Only Lidar data, i.e., only the point cloud
FUL	Full setup, combining camera, LiDAR, and RADAR data

5. Detection Probability: CAM

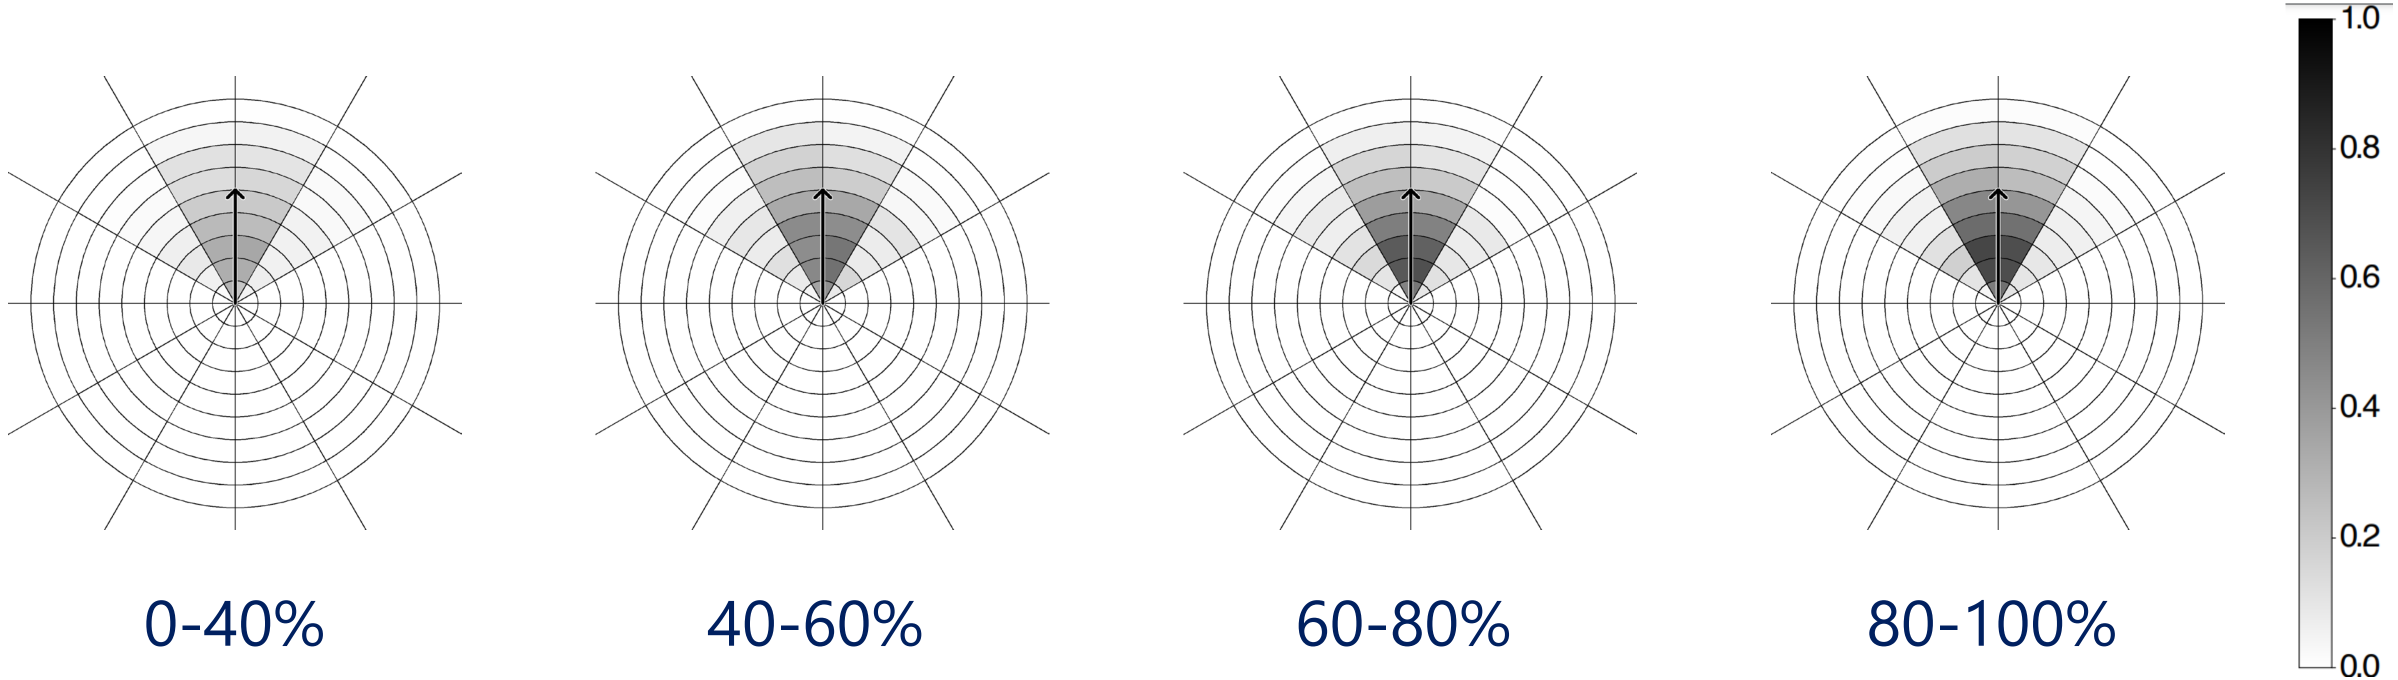

- Camera Field of View and detection range
- Detection Probability increases with object visibility

5. Detection Probability: LID

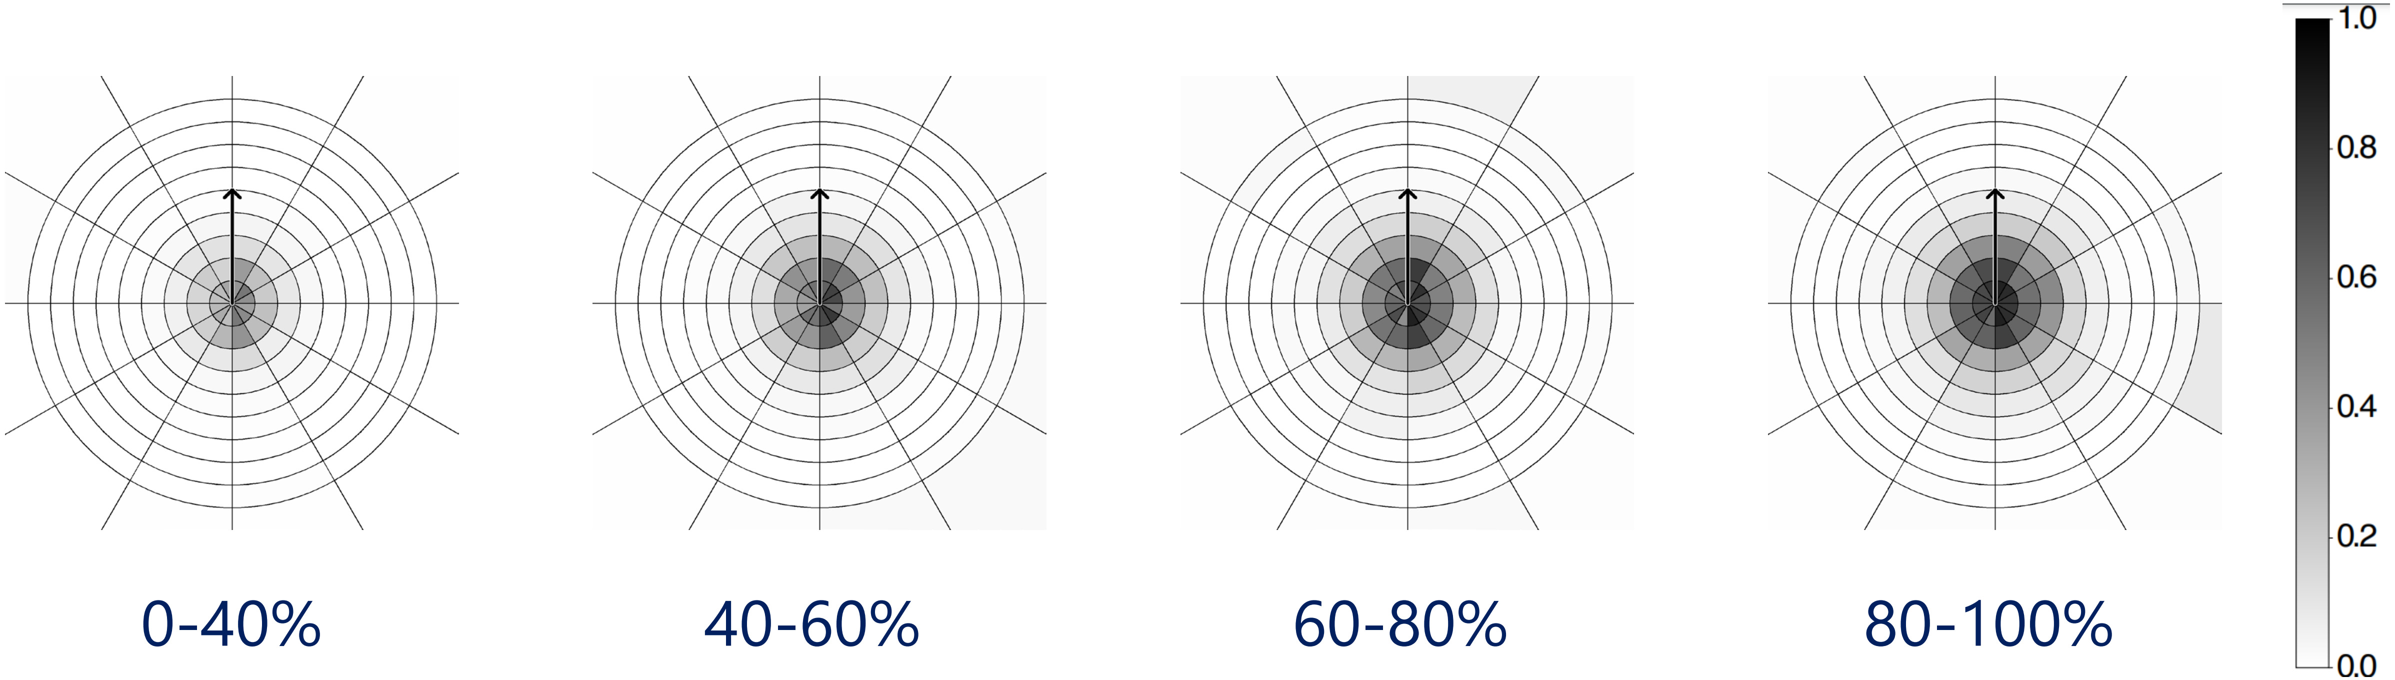

0-40% 40-60% 60-80% 80-100%

- LiDAR Field of View and detection range
- Detection Probability increases with object visibility

5. Detection Probability: FUL

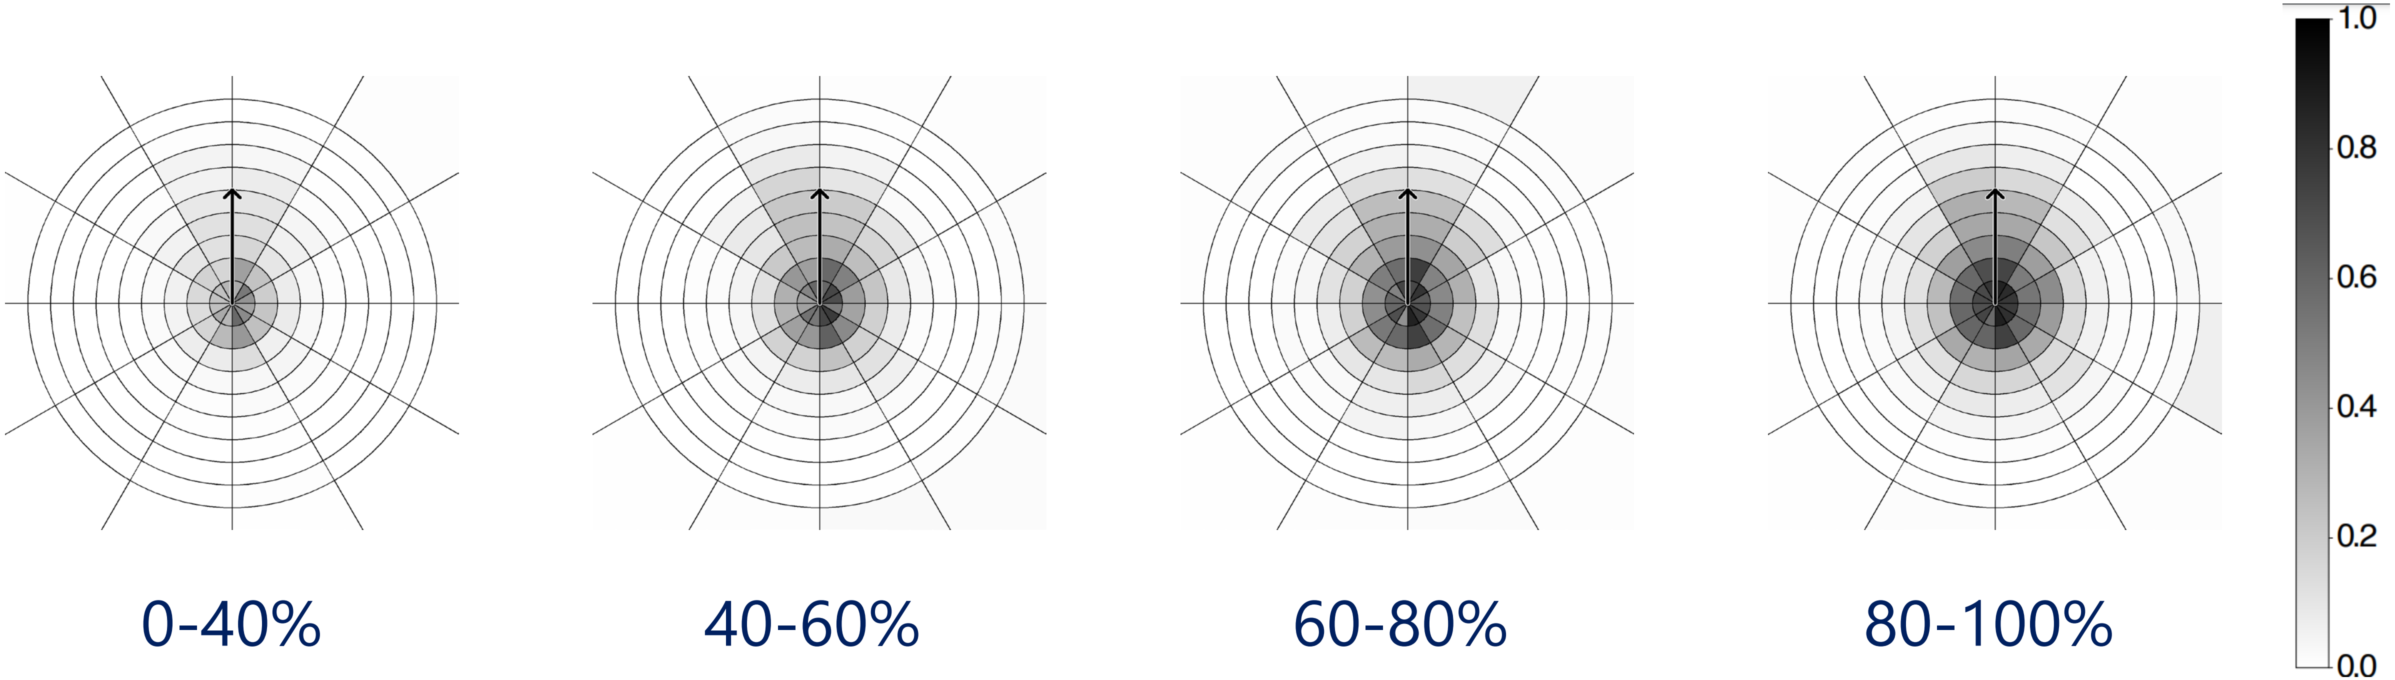

0-40% 40-60% 60-80% 80-100%

- Fusion Field of View and detection range
- Detection Probability increases with object visibility

5. Detection Probability: Frontal Cone

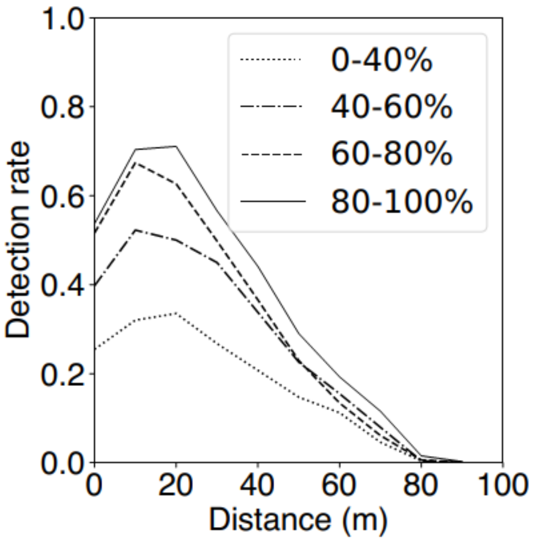

CAM

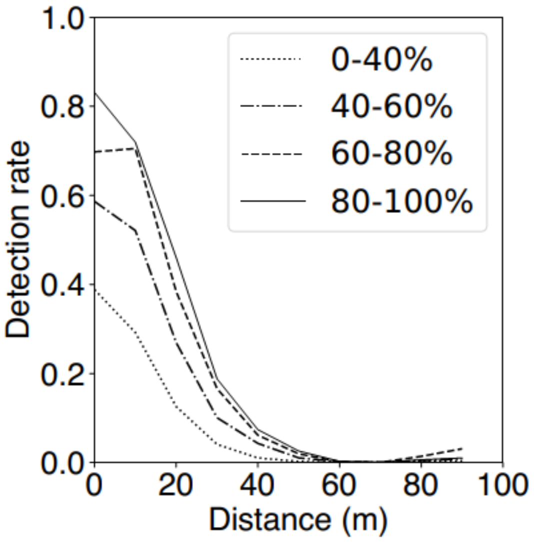

LID

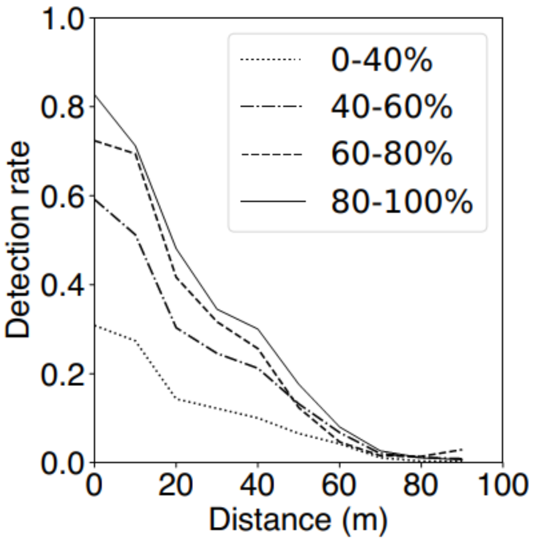

FUL

- Detection probability with CAM setup decays slower than LID
- FUL setup is better than LID, except for small distances

5. TC6: Jaywalking pedestrian

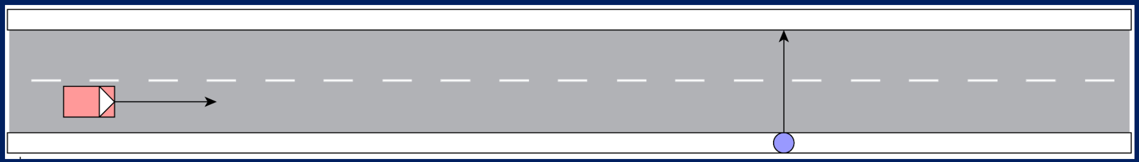

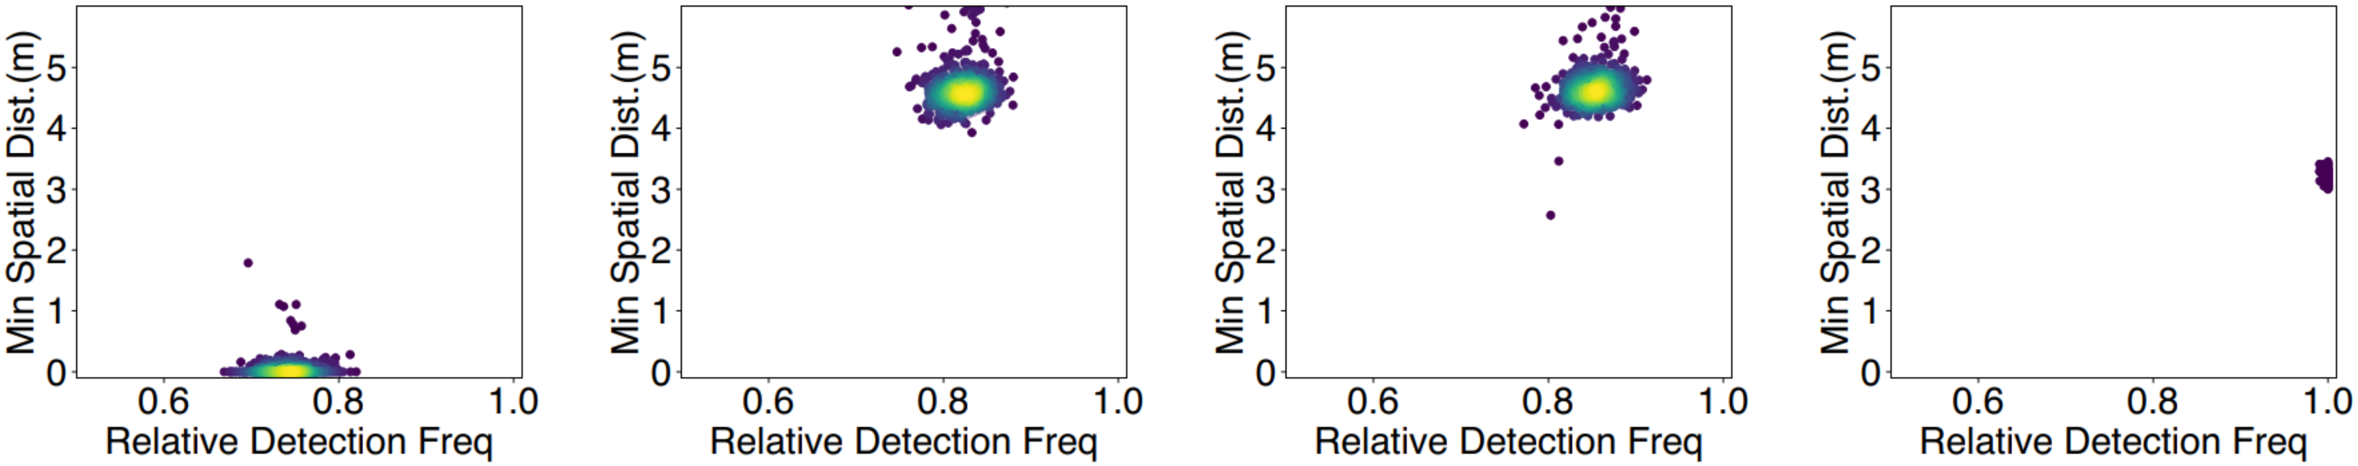

CAM LID FUL GT

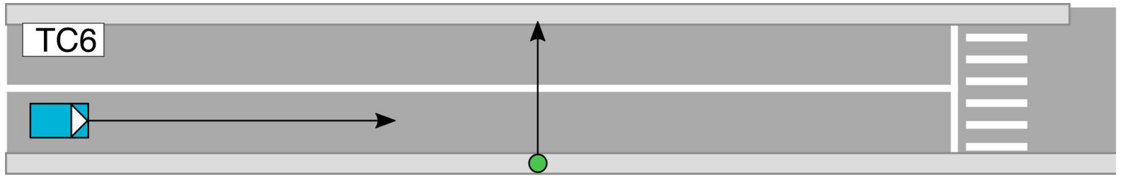

- Under CAM setup, the AV does not perform safely due to **localization errors**.
- LID and FUL both lead to safer behavior.
- Comparing LID, FUL, and GT, we observe how LID and FUL are more unstable than GT but maintain a larger distance from the pedestrian.

6. Conclusions

- Effects of misdetection errors on AV behavior
  - Perception Error Model (PEM)
  - Realistic simulations with errors generated by PEM
  - Risk assessment
- Future developments:
  - Physical HiL testing;
  - On-road AV trials.

6. Future Works

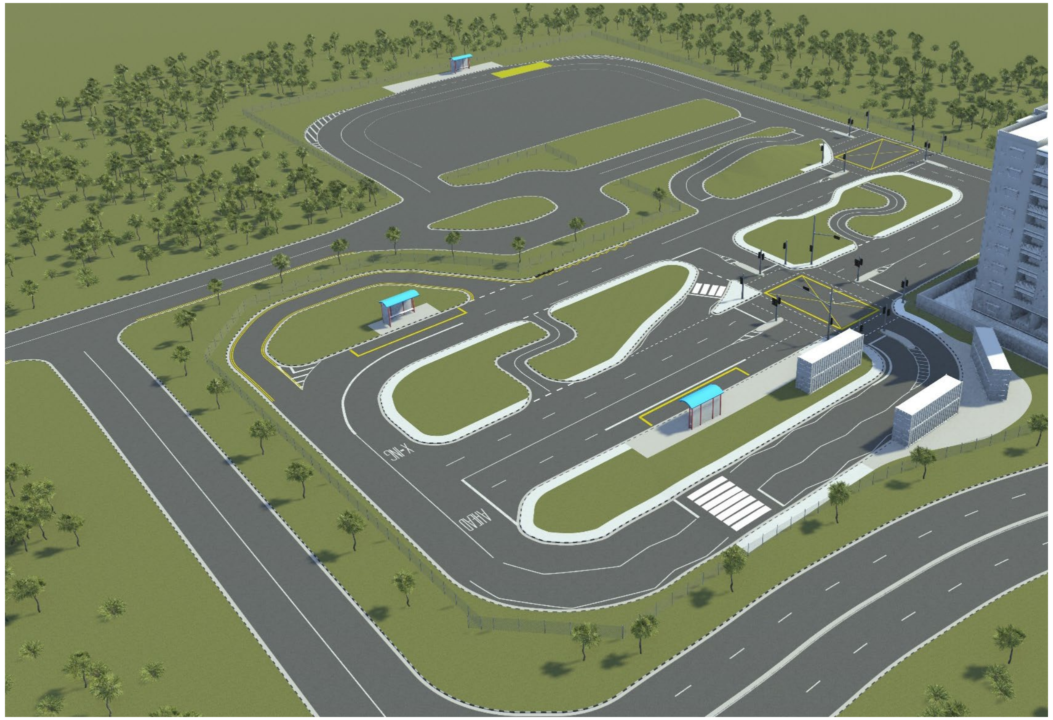

1. Bus stop with bay
2. Rain simulator
3. Slope
4. Signalled intersection
5. S-course
6. Signalled intersection
7. V2X communication
8. Charging station for vehicle and AutOnomous VehicLe MonItoring and EValuation SystEm (OLIVE)
9. Urban canyon
10. Pedestrian crossing
11. Crank course
12. Bus stop
13. Flash flood area

CETRAN

Thank you for the attention

Q&A

Perception error modelling for autonomous driving

Andrea Piazzoni, Jim Cherian,
Martin Slavik, Justin Dauwels.
Cetran ERI@N – Interdisciplinary Graduate School
Nanyang Technological University