FastCodeAI

Autonomy at scale: where are we headed?

Arjun Jain, Founder @ FastCode.AI

25th July 2023

Automotive: Use Cases and Customers

User Comfort and Assist
(Inside the Vehicle)

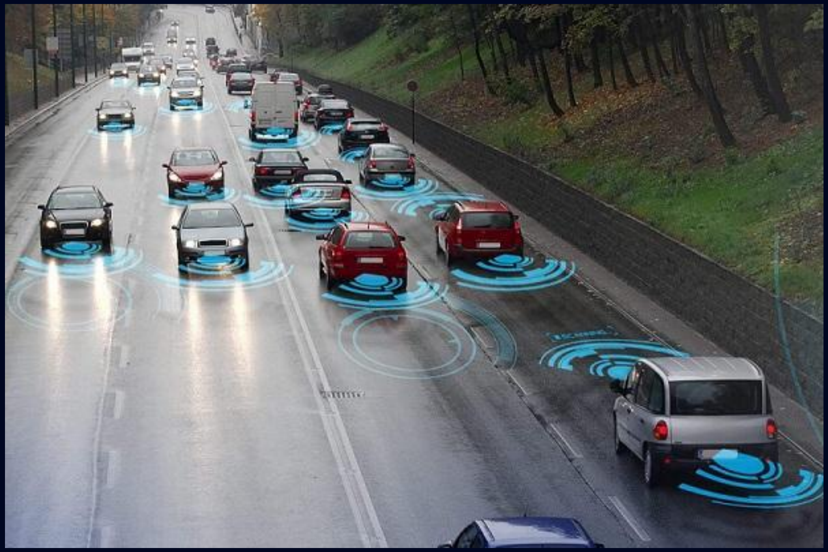

Automated Driving
(Outside the Vehicle)

Mercedes Benz MBUX (Paris Auto Show 18)

AUTOMOTIVE 2021

Method and system for triggering an event in a vehicle

Embodiments of present disclosure relates to method for identifying a hand pose in a vehicle, and a system for performing an event in a vehicle. A hand image for a hand in the vehicle, is extracted from a vehicle image of the vehicle for identification. Plurality of contextual images of the hand image is obtained based on the single point. A hand pose associated with the hand is identified based on the plurality of contextual features using a classifier model.

Download PDF

AUTOMOTIVE 2021

System and method for deployment of airbag based on head pose estimation

Advanced airbag deployment control system designed for vehicular use incorporating an image sensor unit that captures real-time images of a vehicle occupant, focusing particularly on head localization. Utilizing a processing unit equipped with Long Short Term Memory (LSTM) neural network architecture, the system analyzes images to determine and predict the future position and orientation of the occupant's head. The system dynamically adjusts the direction in which the airbag flap is removed and the airbag's inflation pressure, ensuring optimal safety by adapting the deployment to the predicted head position at the moment of impact.

Download PDF

AUTOMOTIVE 2019

Method for Identifying a Hand Pose in a Vehicle

A method for activating vehicle functions via hand gestures, utilizing a 3D Convolutional Neural Network (3D-CNN) and Gated Recurrent Unit (GRU) to analyze video frames and extract spatio-temporal features. A prediction module simultaneously assesses the gesture's progression and classifies it, employing predefined models to determine the gesture's type and its completion rate. Upon accurate detection and classification, a corresponding event is triggered within the vehicle.

View Source

All | Automotive | Healthcare | Finance | Retail | Oil & Gas

Gesture Recognition based User Experience

Enhance interaction between driver and car by detecting driver hand gesture for superior level driver experience with Automotive OEM.

Read More

Vulnerable Roadside User Protection System

Accurate, low footprint detection with vulnerable roadside users (VRUs) and road signs covered in the Automotive Tier 1 autonomous driving initiative.

Read More

Federated Learning on the Edge

Revolutionizing privacy and efficiency in data processing with Federated Learning (FL) on edge devices, enabling real-time, secure analytics in several key sectors.

Read More

Autonomous Driving: Pedestrians and VRUs

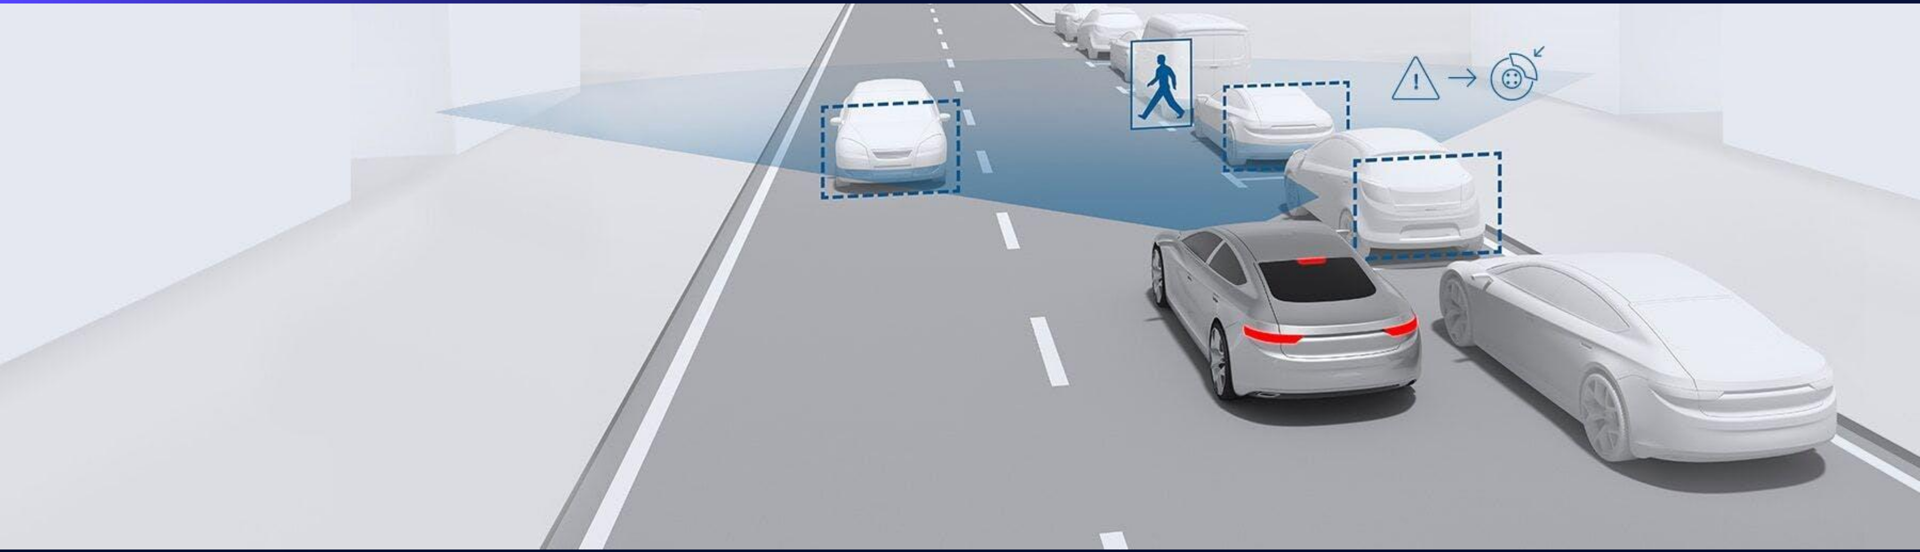

Impact

01. **Research Collaboration:** Co-authored the research paper "VRU Pose-SSD: Multiperson Pose Estimation for Automated Driving" with Bosch, Mercedes, and the Indian Institute of Science.

02. **NCAP Compliance:** The predictive pedestrian protection system meets NCAP requirements for automatic emergency braking.

03. **Potential of Automatic Emergency Braking:** Shows significant potential in preventing or mitigating frontal collisions with pedestrians at speeds up to 60 km/h, significantly reducing injury risks, avoiding or mitigating half of the accidents with cyclists resulting in personal injury in Germany, and reducing up to 30% of relevant pedestrian accidents.

History of Autonomous Cars

- Eureka Project **PROMETHEUS** Europe between 1987-1995
- VITA II by Daimler-Benz, succeeded in automatically driving on highways

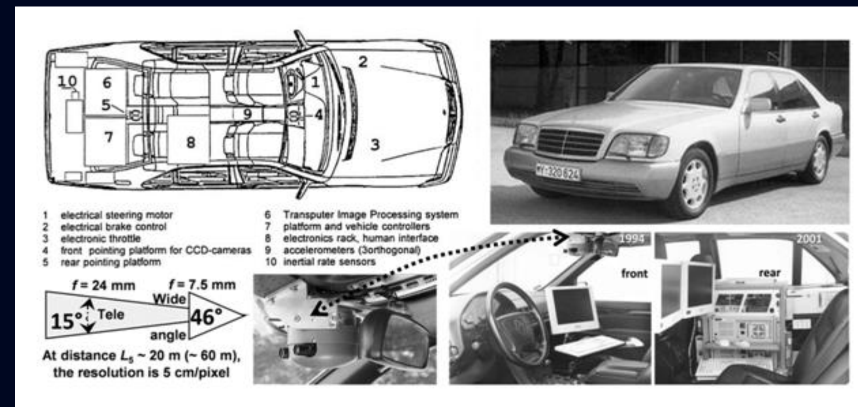

- DARPA Grand Challenge in 2004 - all participants FAILED to finish the 150-mile off-road track.

History of Autonomous Cars – cont.

- Another similar DARPA Grand Challenge was held in 2005. This time five teams managed to complete the off-road track without any human interference. Velodyne supplier to all teams.
- DARPA Urban Challenge held in 2007, test environment that was modelled after a typical urban scene. Six teams managed to complete the event.

Tartan Racing – Pittsburgh, PA
1st Place

Stanford Racing – Stanford, CA
2nd Place


Victor Tango – Blacksburg, VA
3rd Place

Nobody won. Nobody finished. Yet the first DARPA race left an indelible mark

Nobody won, but participants fondly recall the race that started the push toward the modern-day autonomous vehicle industry.

March 10, 2024 12:00 AM

PETE BIGELOW

AUTOMOTIVE NEWS VIDEO - 03/12/2024: DARPA Grand Challenge anniversary

The DARPA challenges were three foundational events which catalyzed and set the stage for the modern-day self…

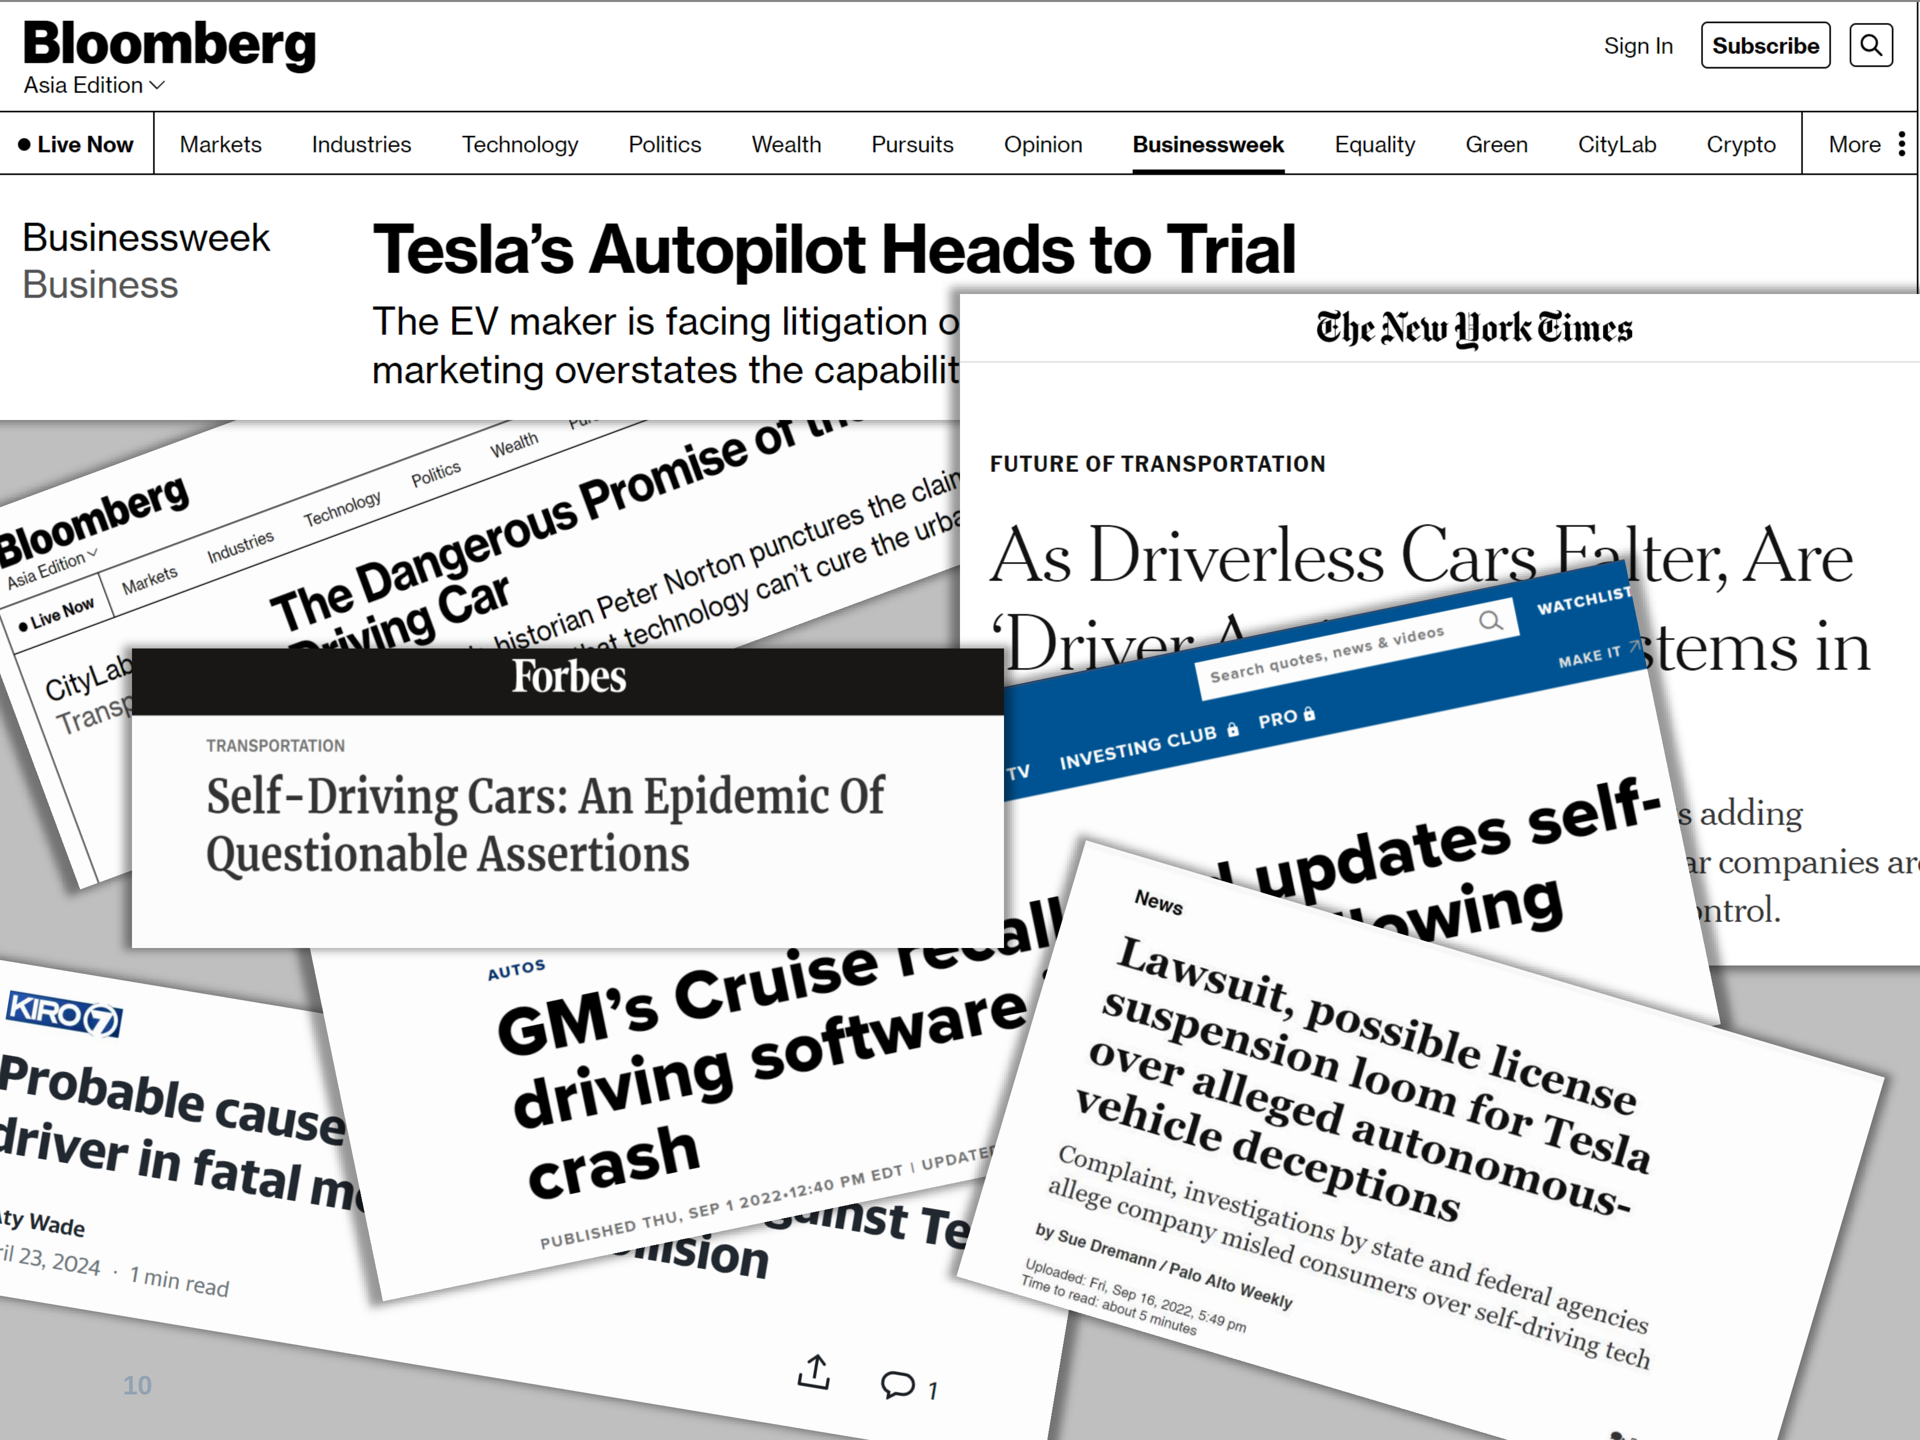
Bloomberg
Asia Edition
Sign In
Subscribe
Live Now | Markets | Industries | Technology | Politics | Wealth | Pursuits | Opinion | Businessweek | Equality | Green | CityLab | Crypto | More
Businessweek
Business
Tesla’s Autopilot Heads to Trial
The EV maker is facing litigation o
marketing overstates the capabilit

Bloomberg
Asia Edition
Live Now | Markets | Industries | Technology | Politics | Wealth
CityLab
Transp
The Dangerous Promise of th
Driving Car
historian Peter Norton punctures the clai
that technology can’t cure the urba

The New York Times
FUTURE OF TRANSPORTATION
As Driverless Cars Falter, Are
‘Driver A
stems in
s adding
ar companies ar
ontrol.

Forbes
TRANSPORTATION
Self-Driving Cars: An Epidemic Of Questionable Assertions

Search quotes, news & videos
WATCHLIST
MAKE IT
TV | INVESTING CLUB | PRO
updates self-
owing

KIRO 7
Probable cause
driver in fatal m
ty Wade
il 23, 2024 · 1 min read

AUTOS
GM’s Cruise recall
driving software
crash
PUBLISHED THU, SEP 1 2022·12:40 PM EDT | UPDATE
ainst Te
ision

News
Lawsuit, possible license suspension loom for Tesla over alleged autonomous-vehicle deceptions
Complaint, investigations by state and federal agencies allege company misled consumers over self-driving tech
by Sue Dremann / Palo Alto Weekly
Uploaded: Fri, Sep 16, 2022, 5:49 pm
Time to read: about 5 minutes

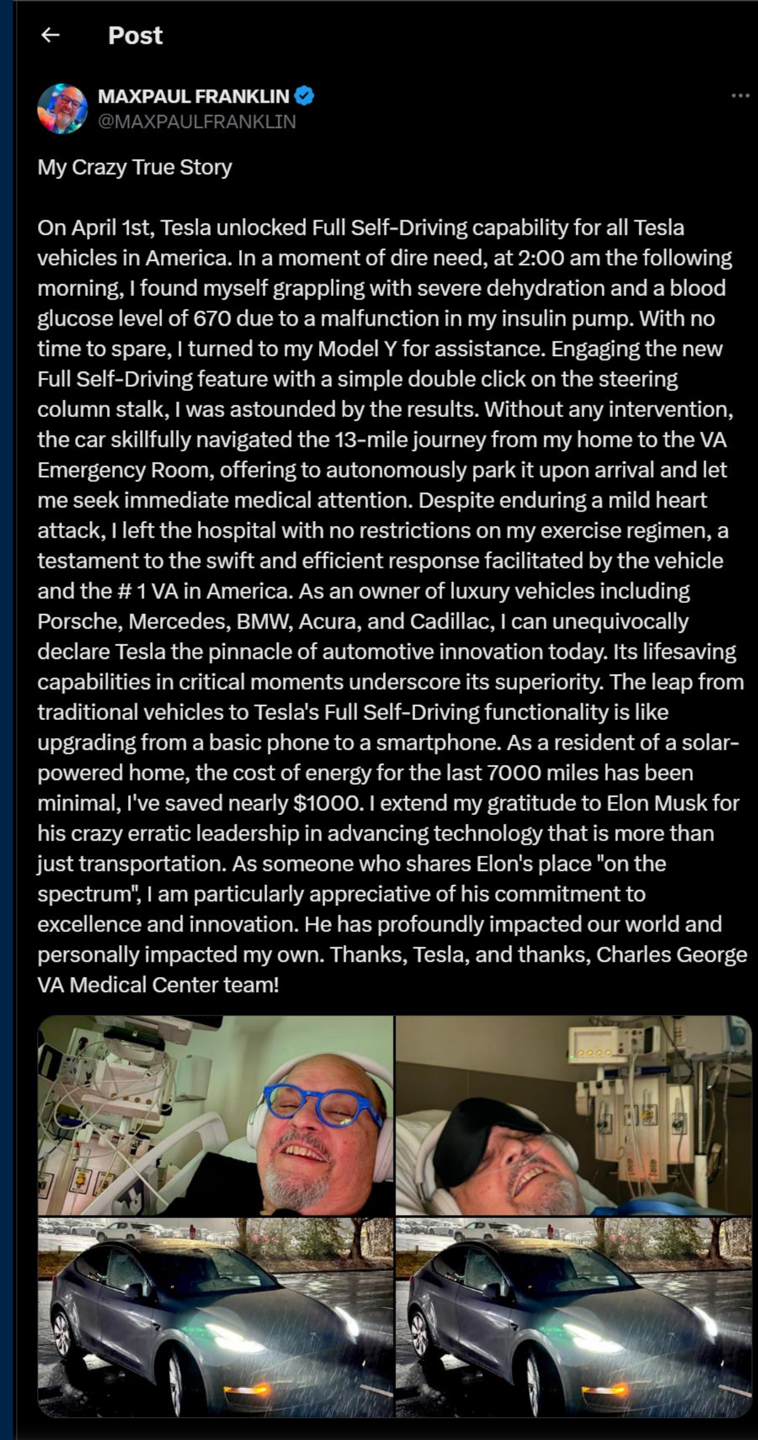
Post

MAXPAUL FRANKLIN
@MAXPAULFRANKLIN

My Crazy True Story

On April 1st, Tesla unlocked Full Self-Driving capability for all Tesla vehicles in America. In a moment of dire need, at 2:00 am the following morning, I found myself grappling with severe dehydration and a blood glucose level of 670 due to a malfunction in my insulin pump. With no time to spare, I turned to my Model Y for assistance. Engaging the new Full Self-Driving feature with a simple double click on the steering column stalk, I was astounded by the results. Without any intervention, the car skillfully navigated the 13-mile journey from my home to the VA Emergency Room, offering to autonomously park it upon arrival and let me seek immediate medical attention. Despite enduring a mild heart attack, I left the hospital with no restrictions on my exercise regimen, a testament to the swift and efficient response facilitated by the vehicle and the # 1 VA in America. As an owner of luxury vehicles including Porsche, Mercedes, BMW, Acura, and Cadillac, I can unequivocally declare Tesla the pinnacle of automotive innovation today. Its lifesaving capabilities in critical moments underscore its superiority. The leap from traditional vehicles to Tesla's Full Self-Driving functionality is like upgrading from a basic phone to a smartphone. As a resident of a solar-powered home, the cost of energy for the last 7000 miles has been minimal, I've saved nearly $1000. I extend my gratitude to Elon Musk for his crazy erratic leadership in advancing technology that is more than just transportation. As someone who shares Elon's place "on the spectrum", I am particularly appreciative of his commitment to excellence and innovation. He has profoundly impacted our world and personally impacted my own. Thanks, Tesla, and thanks, Charles George VA Medical Center team!

Genuine Progress getting Masked

- **Automated emergency braking is standard** on every new car as of September 2022 - 2016 agreement - automakers, I.I.H.S., National Highway Traffic Safety Administration
- Radar or camera-linked brakes have cut police-reported rear-end collisions by a **50%** (I.I.H.S.)

The New York Times

FUTURE OF TRANSPORTATION

As Driverless Cars Falter, Are 'Driver Assistance' Systems in Closer Reach?

With investigations and lawsuits over accidents adding skepticism toward fully driverless technology, car companies are betting on systems that take some, but not all, control.

- Automated pedestrian braking has reduced the number of car-human collisions by **30%** versus cars without the feature.
- And anti-lock brakes; cameras, radar and ultrasonic sensors to manage blind spot and lane departure monitors; and adaptive cruise control have become standard

End-to-end or Modular?

- And we can do autonomy algorithms either using:
  - End to End Systems

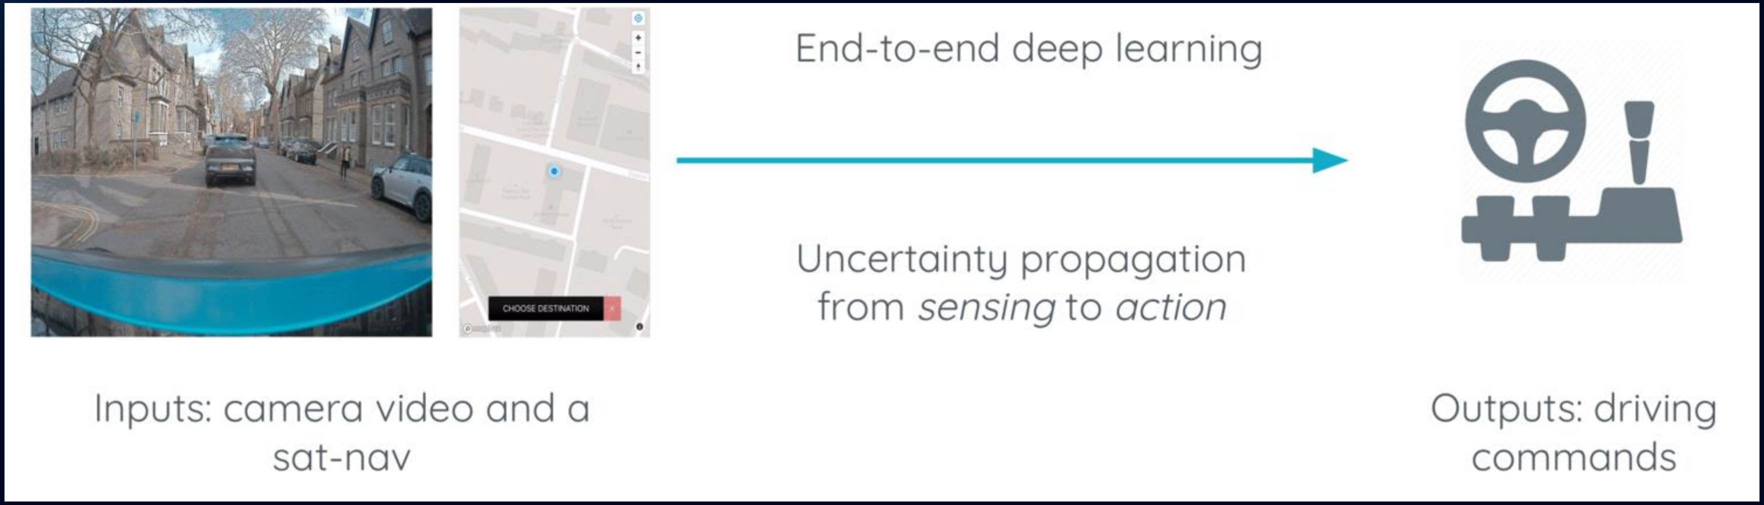

End-to-end or Modular?

- And we can do autonomy algorithms either using:
  - End to End Systems or
  - Modular Systems
    - Perception
    - Scene Representation and Localization
    - Prediction
    - Planning and decision making
    - Vehicle control

Perception

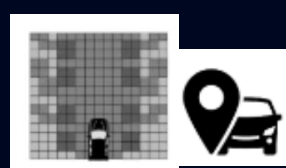

Scene Representation and Localization

Prediction

Planning and Decision Making

Vehicle Control

Perception

Perception

Scene Representation and Localization

Prediction

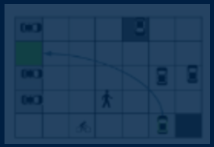

Planning and Decision Making

Vehicle Control

Perception: Autonomous Driving

perception

/pəˈsɛpʃ(ə)n/

noun

1. the ability to see, hear, or become aware of something through the senses.
   "the normal limits to human perception"

Perception: Autonomous Driving

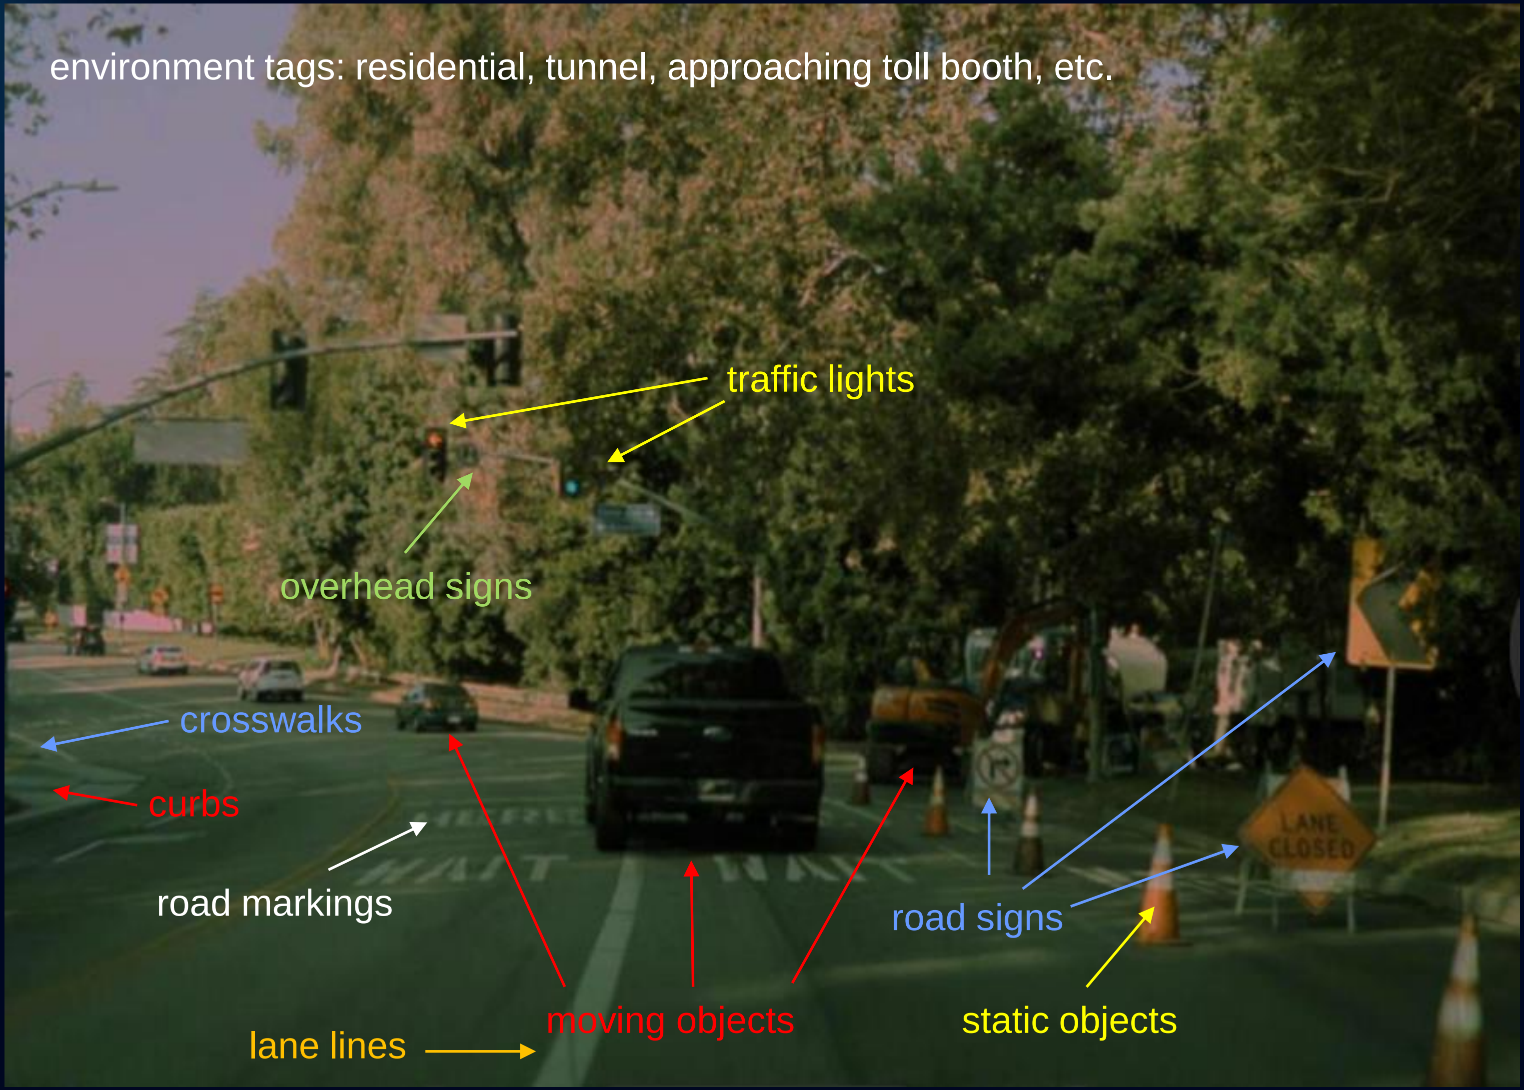

Perception: Long Tail

Each task has additional sub-task: e.g.: object types, vehicle classes, blinkers, brakes, parked, collision, etc.

Perception: Long Tail

Even a sub-task is (car detection) a very difficult problem! Note, 97% is not good enough, finally we need 99.9999%

The Long Tail Problem

- Want to capture behaviour corner cases, but those *by definition* do not much training data.

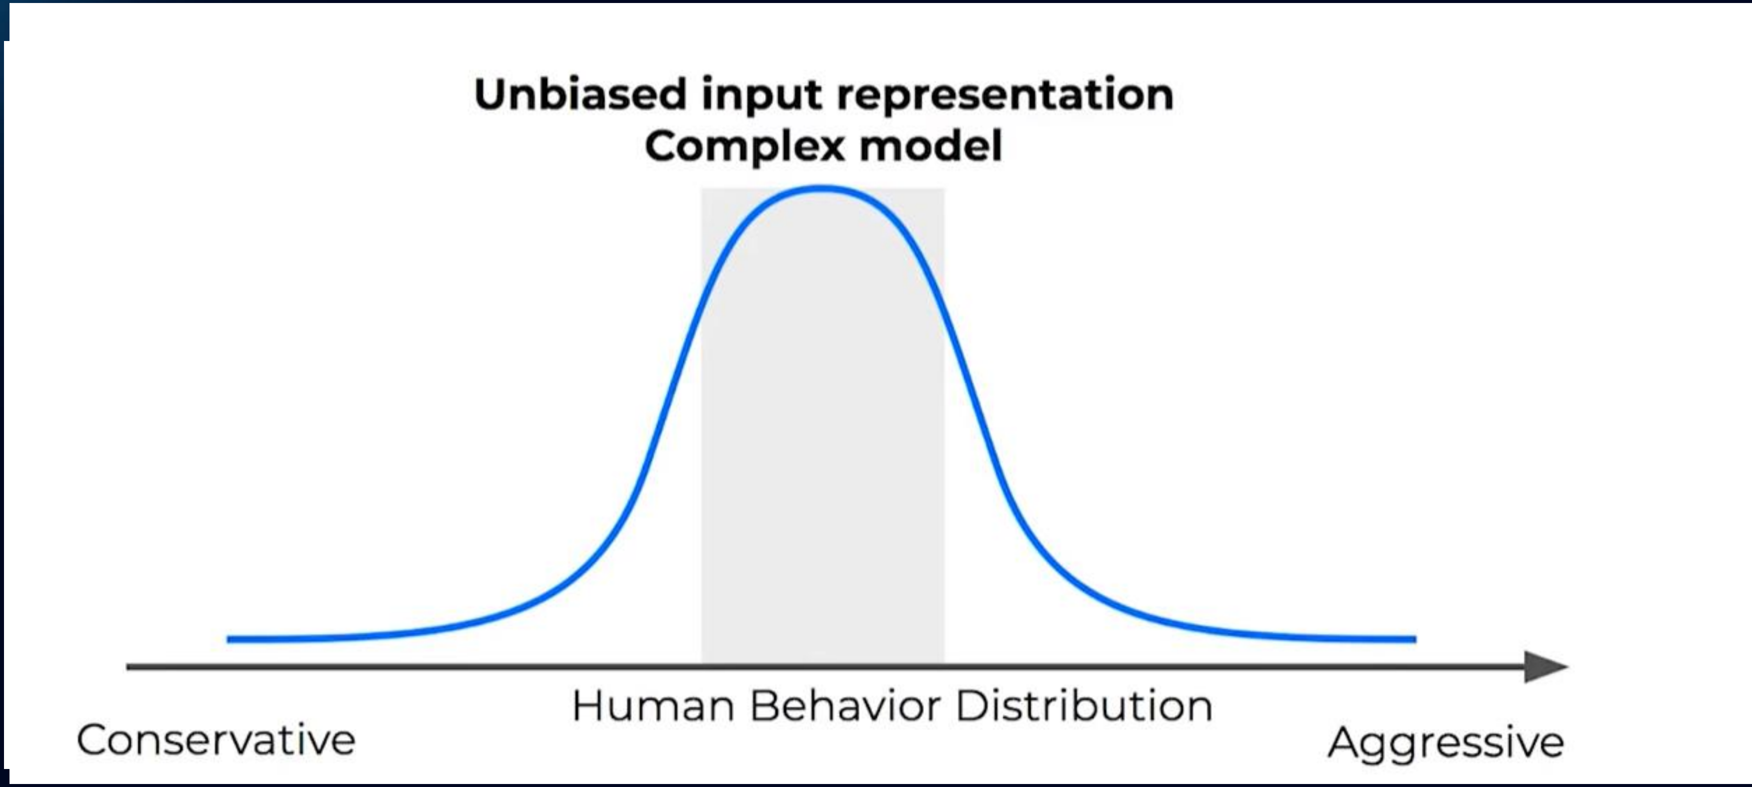

The Long Tail Problem – skydio 2+

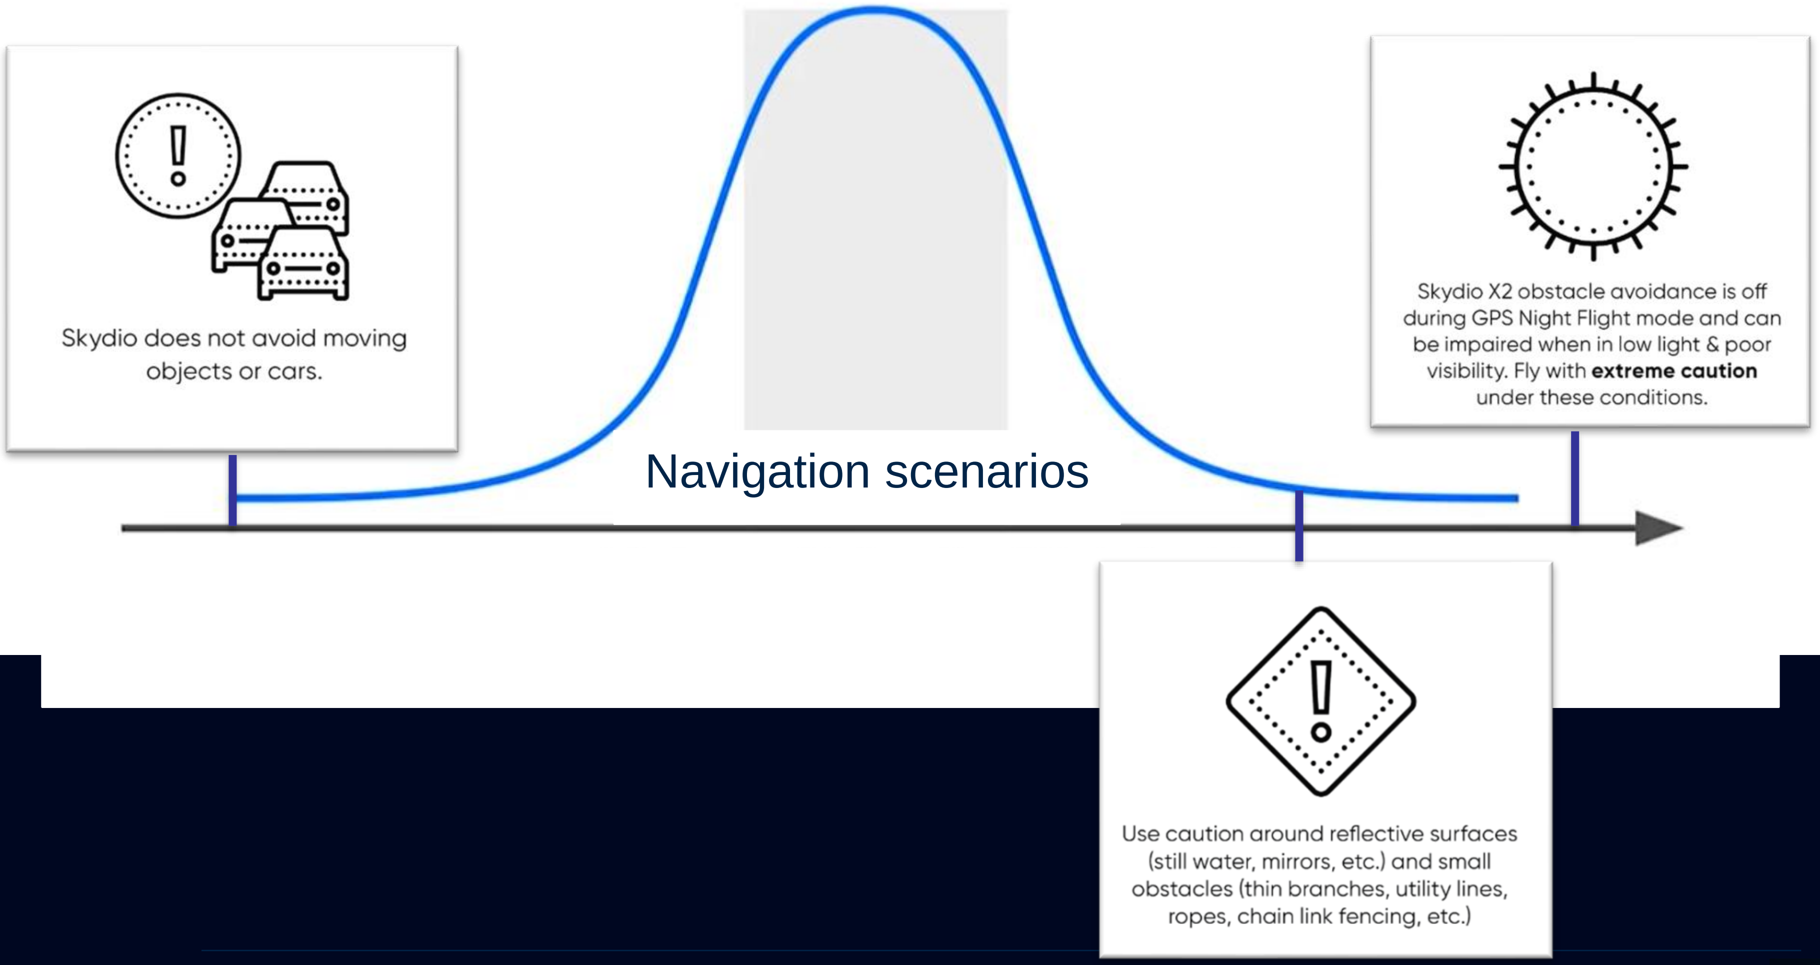

The Long Tail Problem

Another example – Traffic Sign Recognition

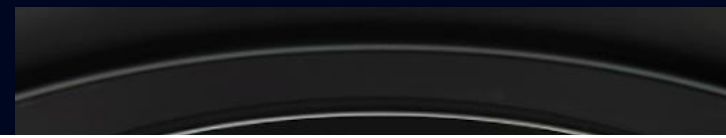

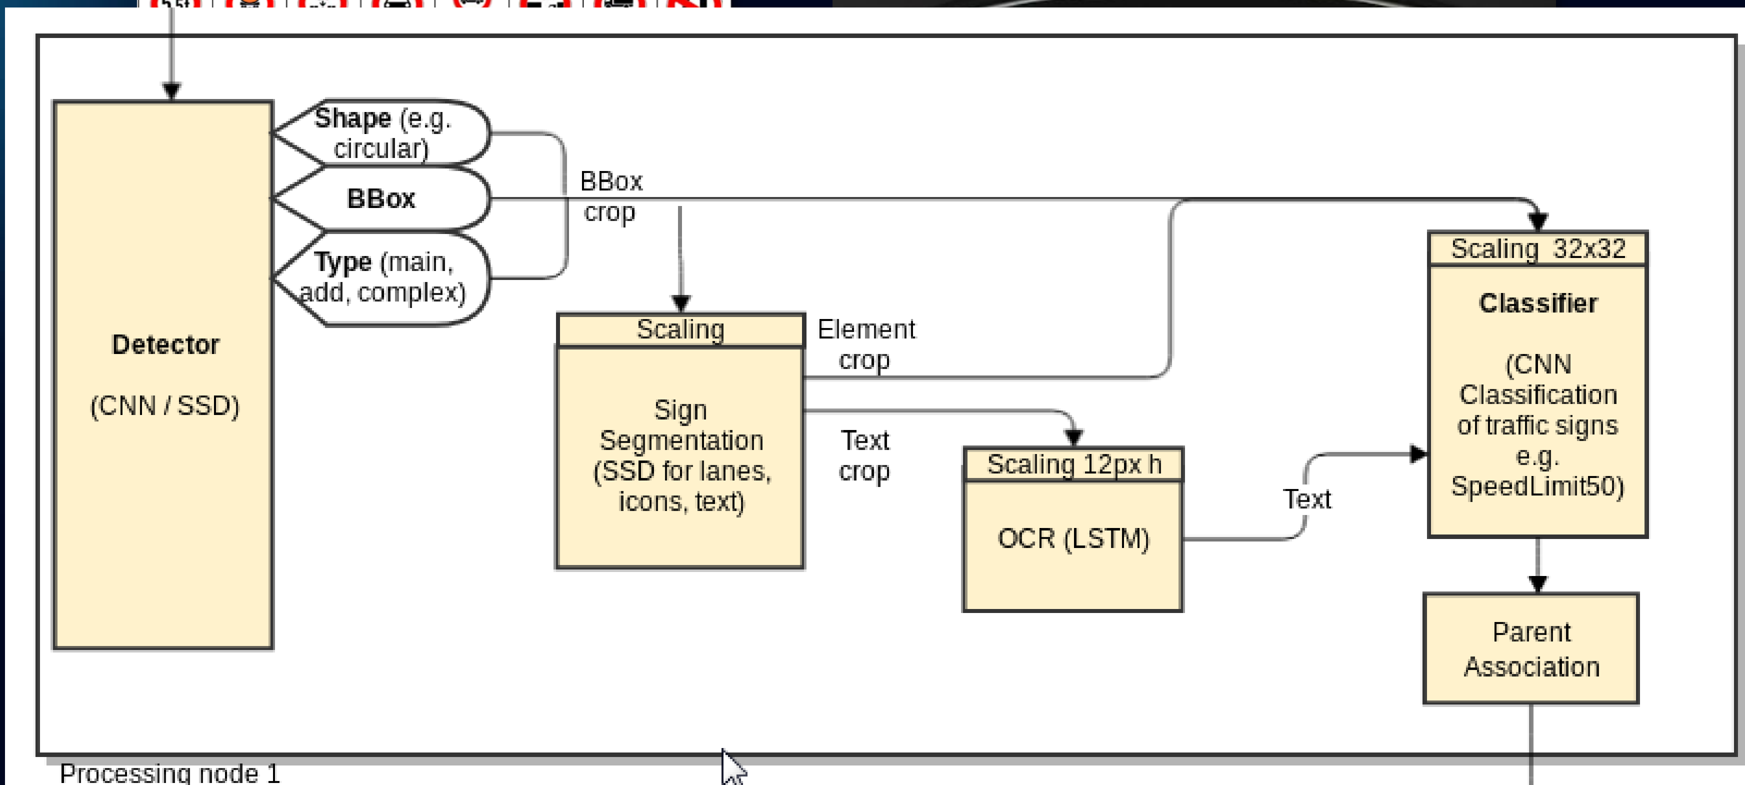

- Low latency, high accuracy system

The Long Tail Problem

Scene Representation and Localization

Localization: LIDAR

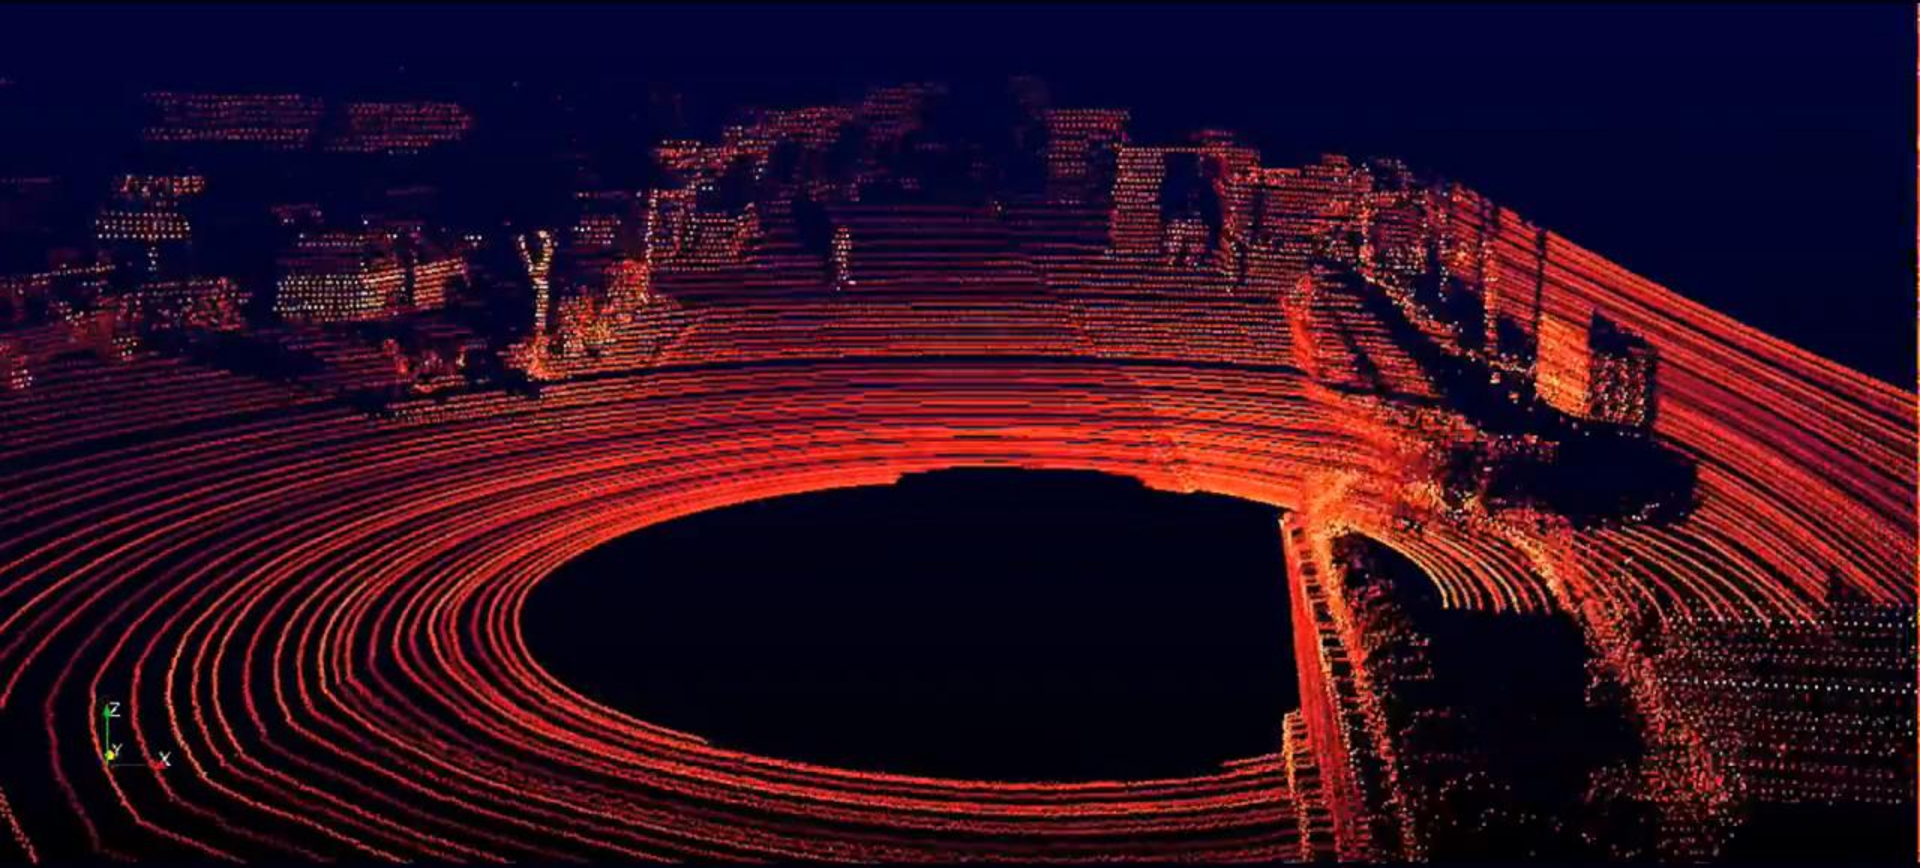

Localization: Only Vision

Scene Representation: Autonomous Driving

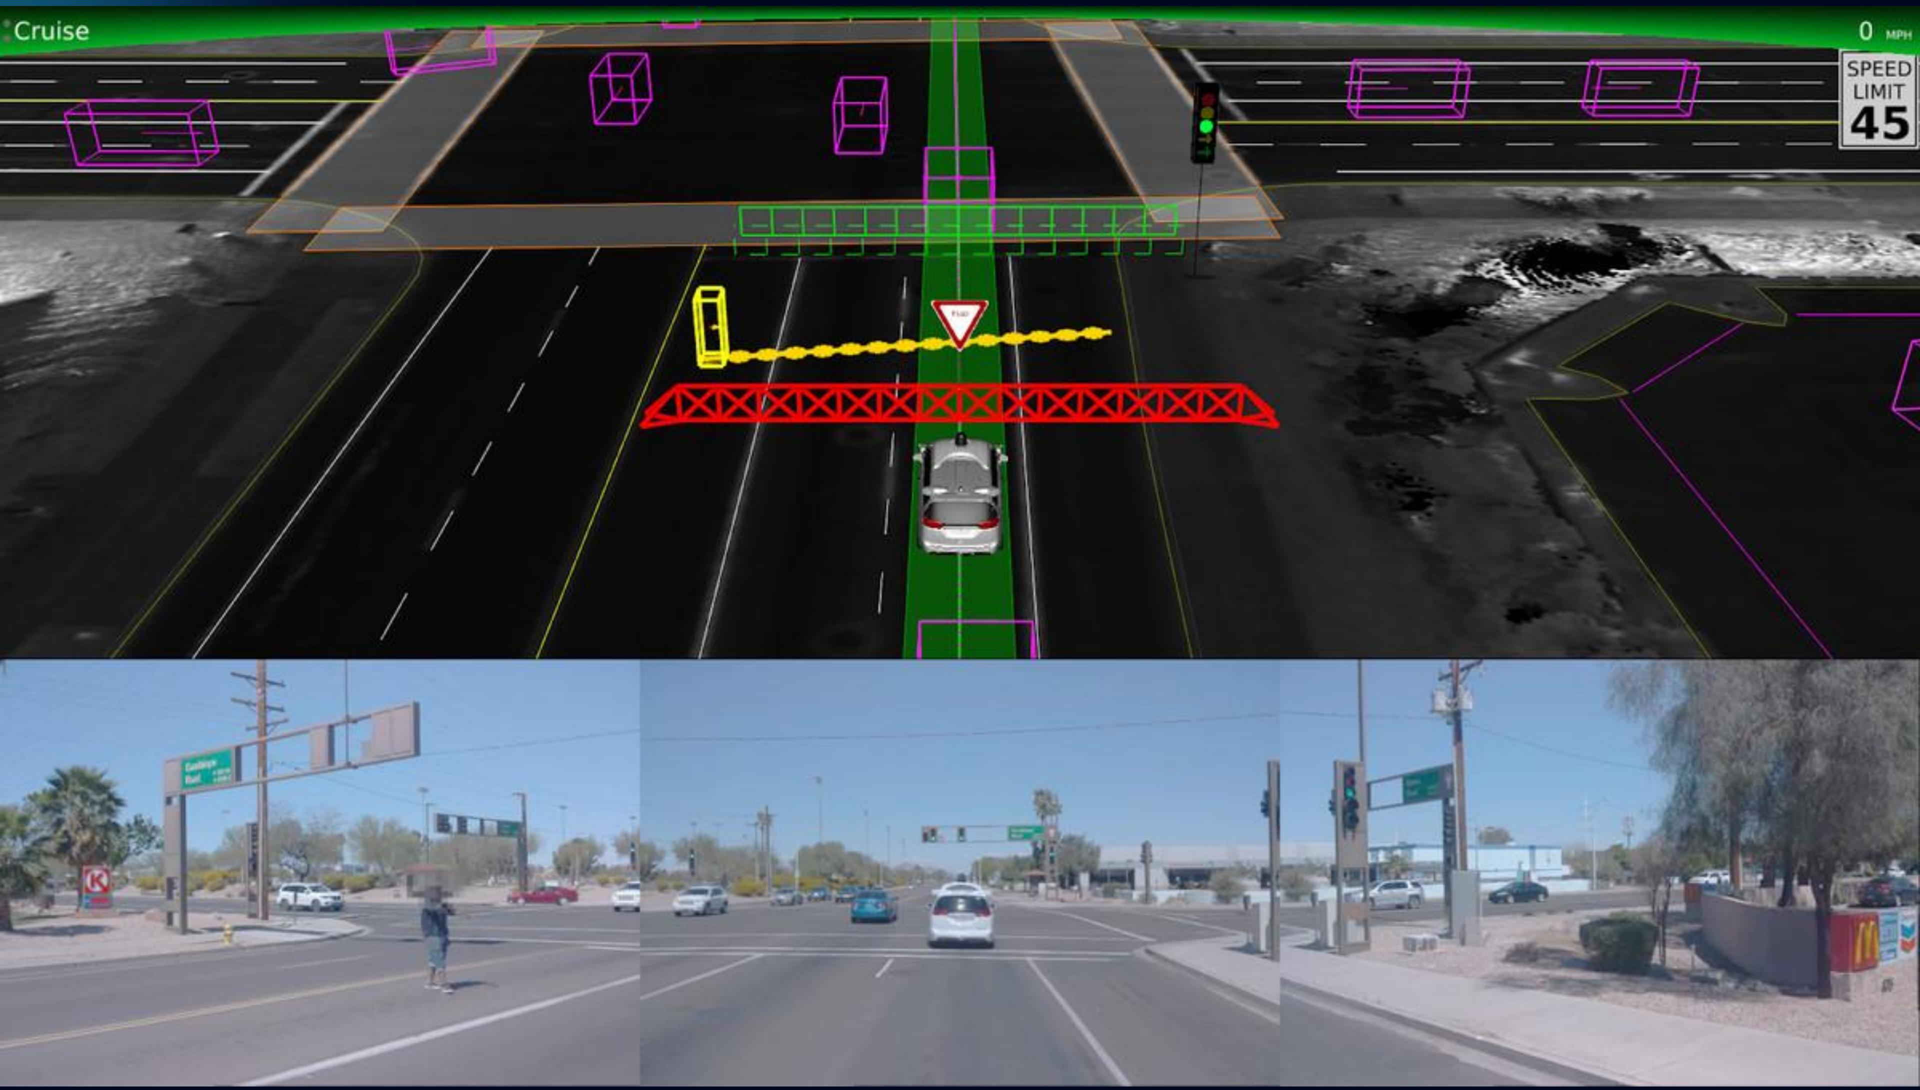

Scene Representation and Localization

Perception

Scene Representation and Localization

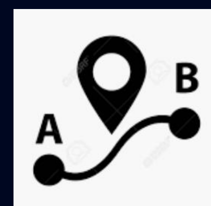

Prediction

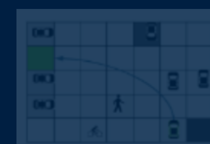

Planning and Decision Making

Vehicle Control

Prediction: Autonomous Driving

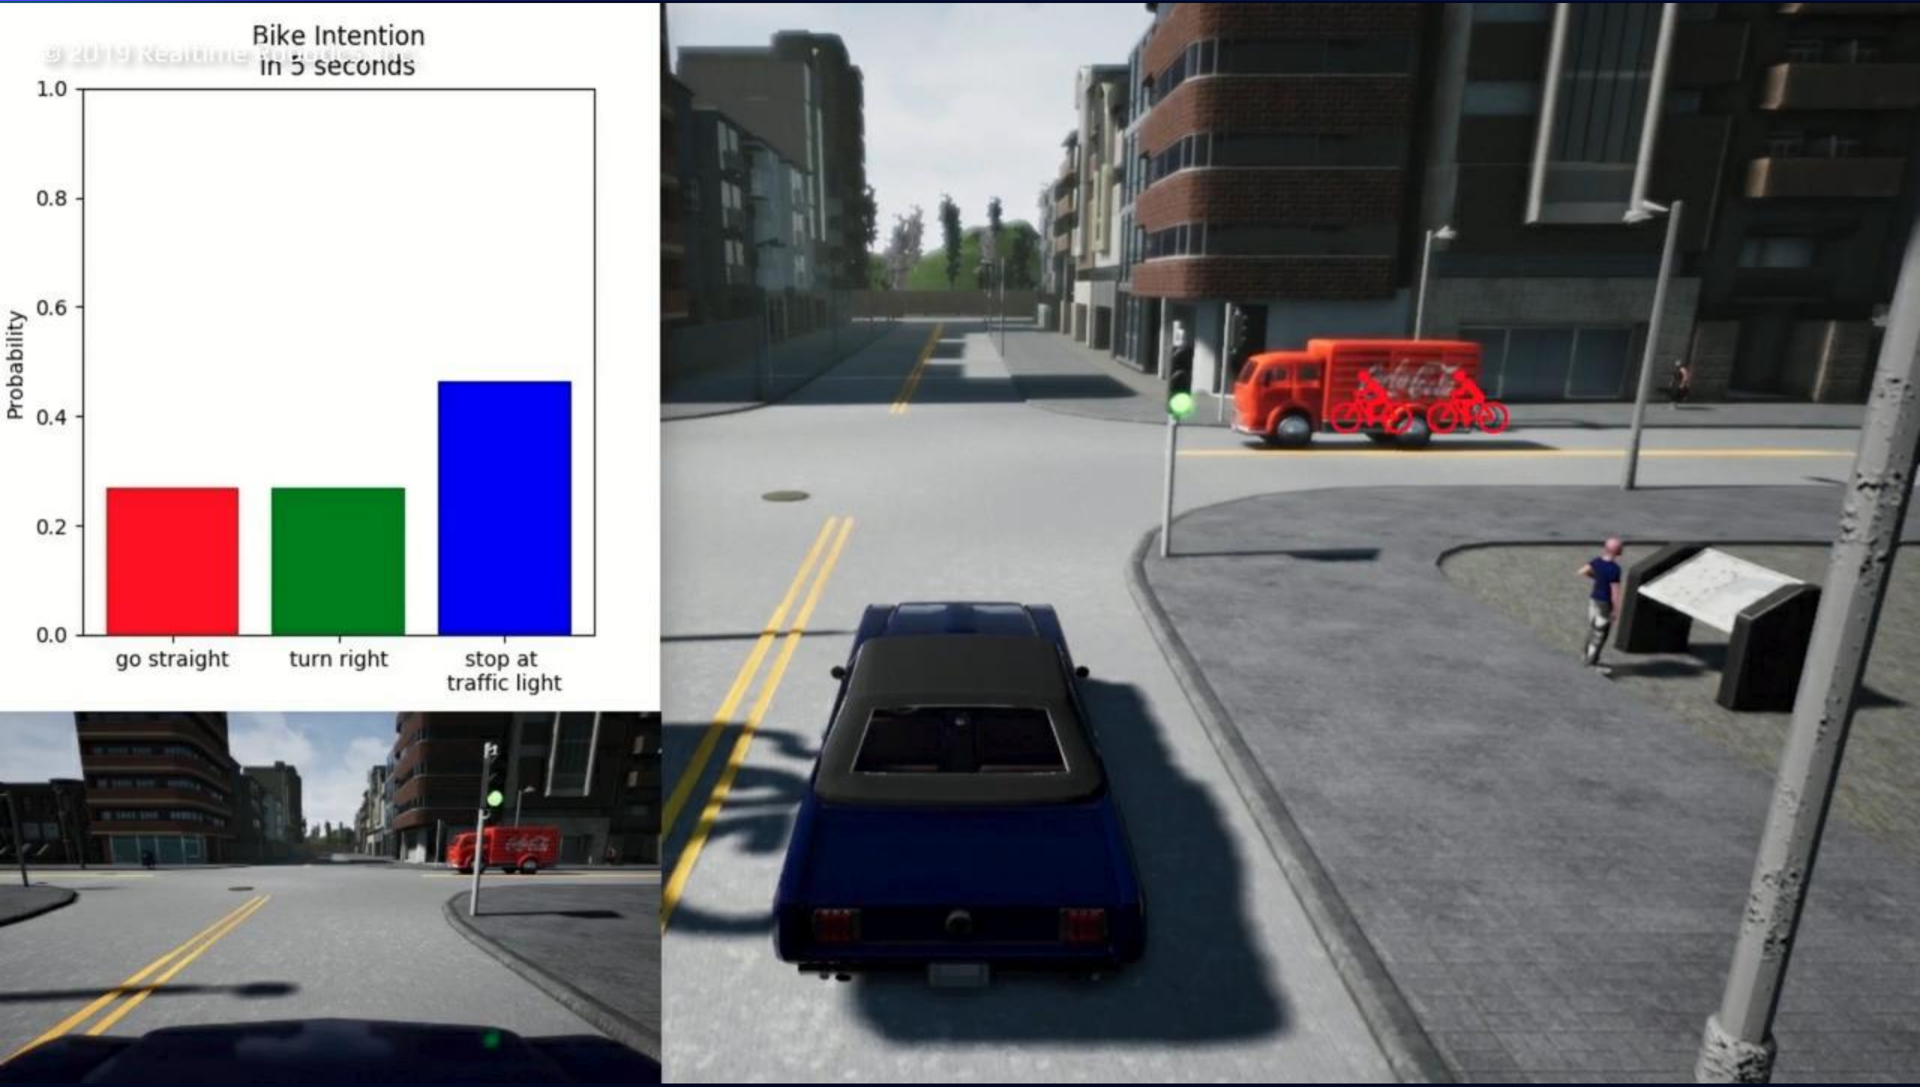

Prediction: Autonomous Driving

Cruise recalls its robotaxis after passenger injured in crash

By Matt McFarland, CNN Business
Updated 1927 GMT (0327 HKT) September 1, 2022

Washington, DC (CNN Business) — This week Cruise, which counts General Motors as its largest shareholder, became the first robotaxi operator to recall its vehicles, following a June crash involving "major" damage and minor passenger injuries.

The crash occurred after the Cruise robotaxi making a left turn stopped in the intersection, thinking that an oncoming vehicle would turn in front of it. But the oncoming vehicle instead drove straight, striking the Cruise vehicle. Both the San Francisco police department and National Highway Traffic Safety Administration launched investigations.

Cruise has said that the oncoming vehicle drove in the right-turn lane and was traveling at "approximately 40 mph" in a 25-mph lane before it exited the lane and proceeded forward. Cruise acknowledged in its recall filing that its robotaxi was not "sufficiently reactive." Cruise spokeswoman Hannah Lindow declined to say what the Cruise vehicle could have done differently, and declined to release video of the crash.

Prediction: In the Sky

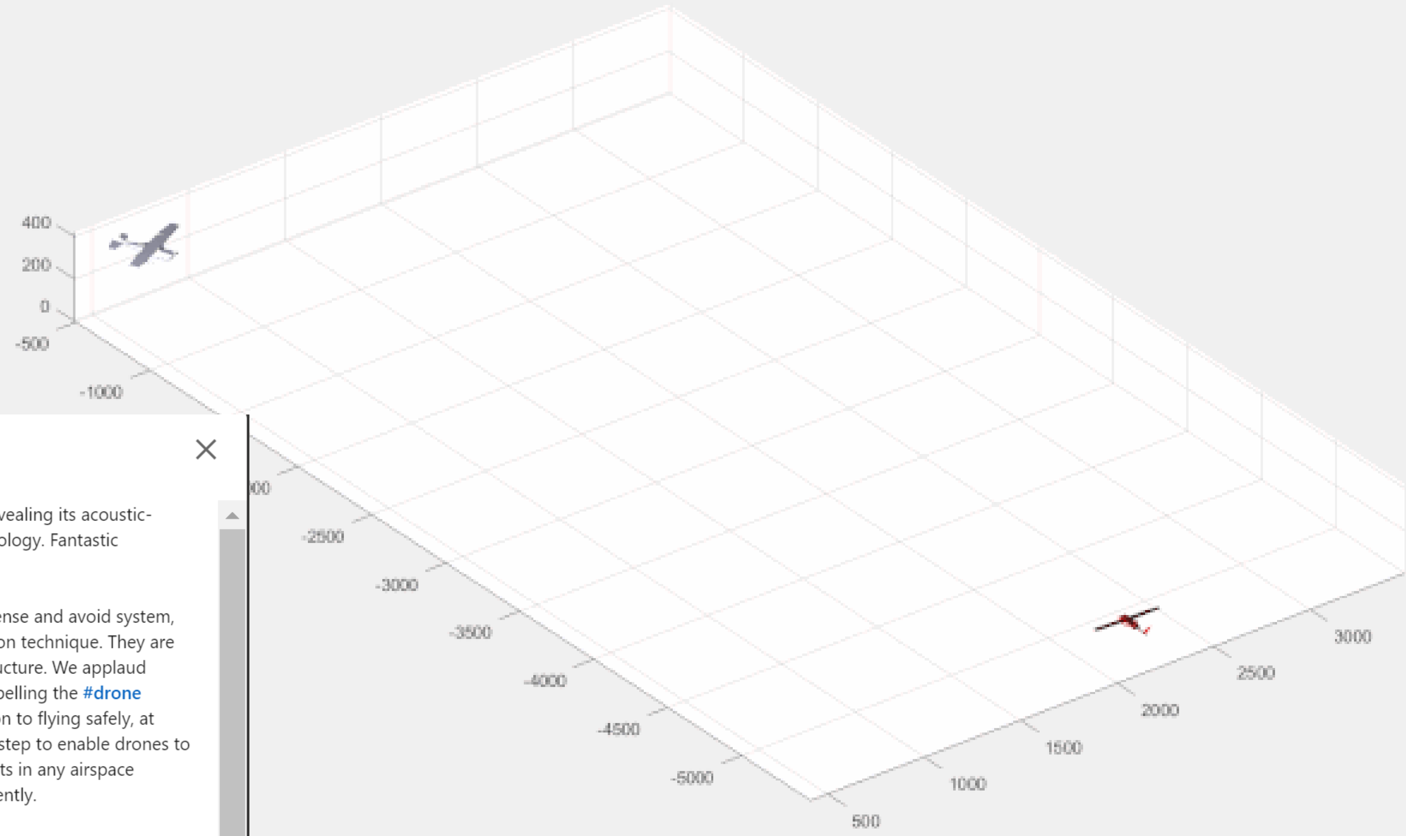

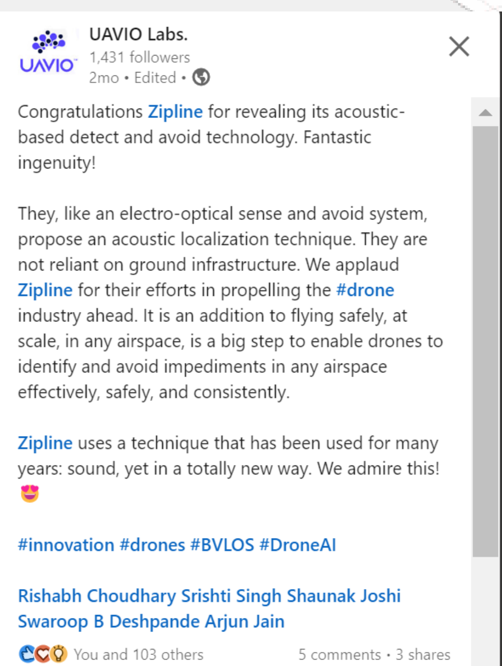

UAVIO Labs.
1,431 followers
2mo • Edited •

Congratulations **Zipline** for revealing its acoustic-based detect and avoid technology. Fantastic ingenuity!

They, like an electro-optical sense and avoid system, propose an acoustic localization technique. They are not reliant on ground infrastructure. We applaud **Zipline** for their efforts in propelling the **#drone** industry ahead. It is an addition to flying safely, at scale, in any airspace, is a big step to enable drones to identify and avoid impediments in any airspace effectively, safely, and consistently.

Zipline uses a technique that has been used for many years: sound, yet in a totally new way. We admire this! 😍

#innovation #drones #BVLOS #DroneAI

Rishabh Choudhary Srishti Singh Shaunak Joshi Swaroop B Deshpande Arjun Jain

You and 103 others 5 comments • 3 shares

Scene Representation and Localization

Motion Planning: Autonomous Driving

Scene Representation and Localization

Perception

Scene Representation and Localization

Prediction

Planning and Decision Making

Vehicle Control

Vehicle Control: Autonomous Driving

- Execute path from Motion Planning by giving relevant actuator commands (steering, acceleration, brake, etc.)
- However, tracking errors are generated due to inaccuracies in vehicle model (e.g. wheel slip during hard breaking)
- Two approaches to fix these errors:
  - **Classic control**: Feedback control uses the measured system response and actively compensates for any deviations
  - **Model predictive control**

Modular Vs End-to-End

Perception

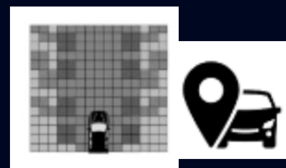

Scene Representation and Localization

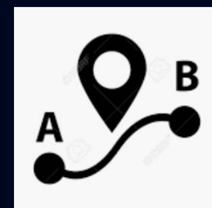

Prediction

Planning and Decision Making

Vehicle Control

- Modular approach helps divide tasks within a team, engineering is easier, more interpretable
- Hand coded rules (priors) easier to include in each module
- Each step of process performed in sequential modular silos
- Stack has to be manually tuned for each introduced change
- Limited information shared across modules (cascaded errors)
- Rely on “modules” for each problem, e.g. foliage

Modular Vs **End-to-End**

Perception

Neural Network

Vehicle Control

- Simple: Single AI System
- Trained for the end task
- Difficult to add priors and rules
- Lack of interpretability for validation and safety
- Require a lot of examples

End-to-end Learning - Wayve

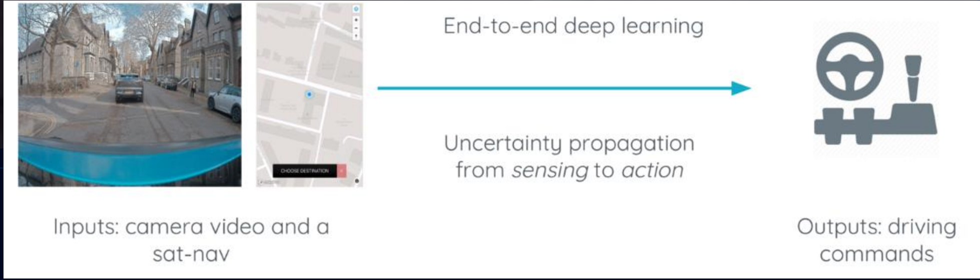

End-to-end Learning - Wayve

- And we can do autonomy algorithms either using:
  - End to End Systems

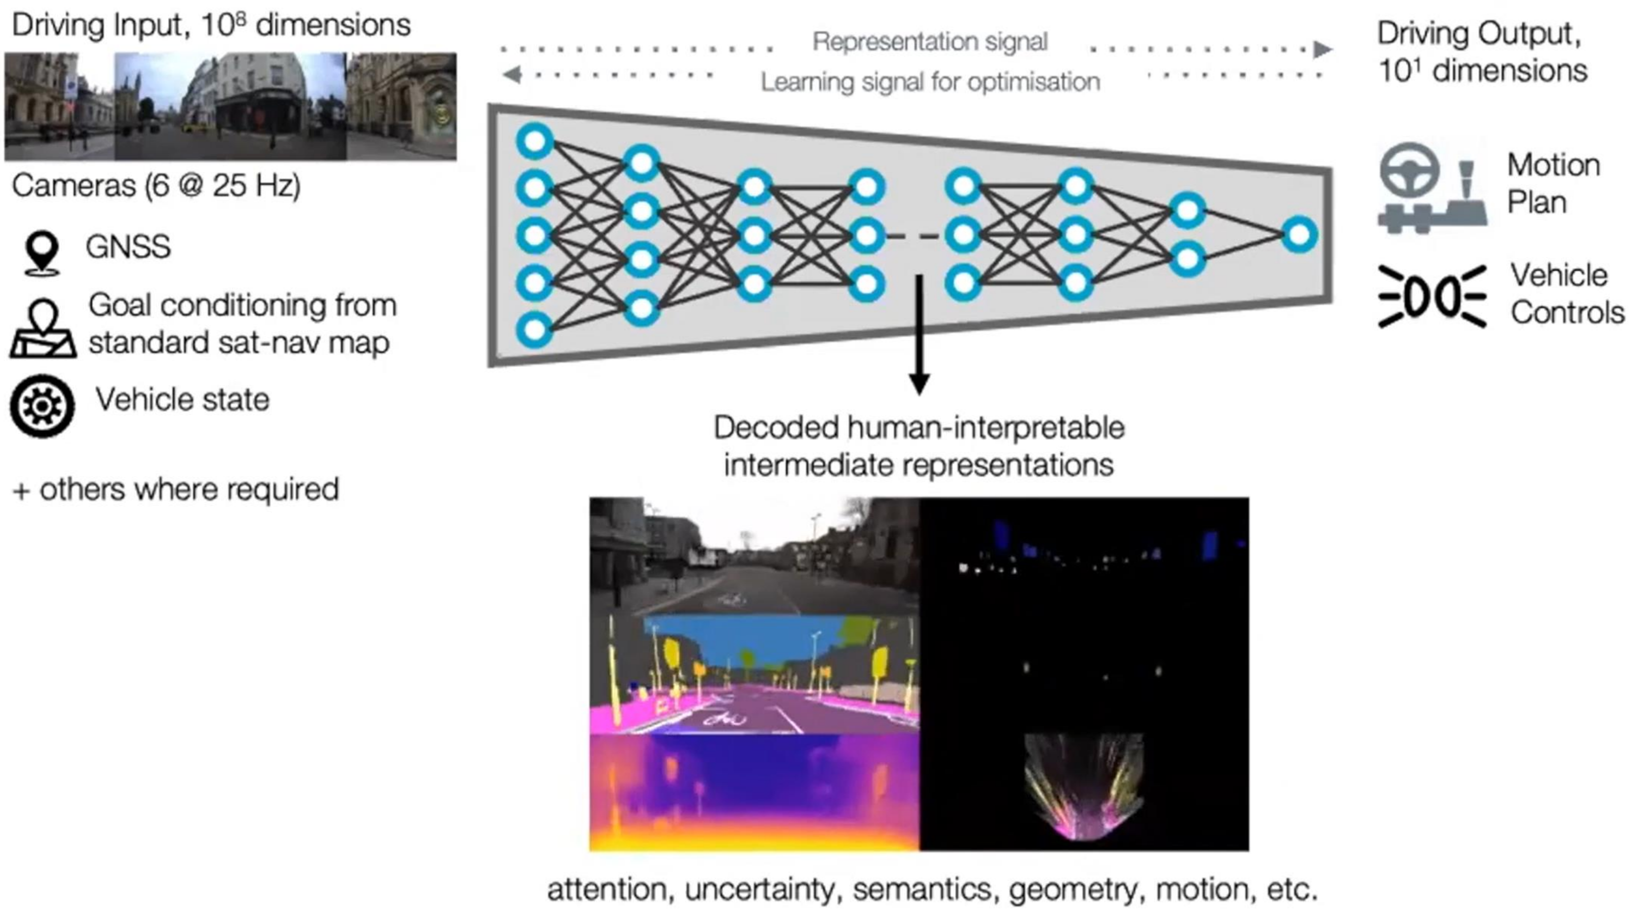

Emergent Behaviour - Wayve

- ❑ End-to-end Transformer Architecture for driving model
- ❑ Attention changes from Traffic lights to road when it turns Green

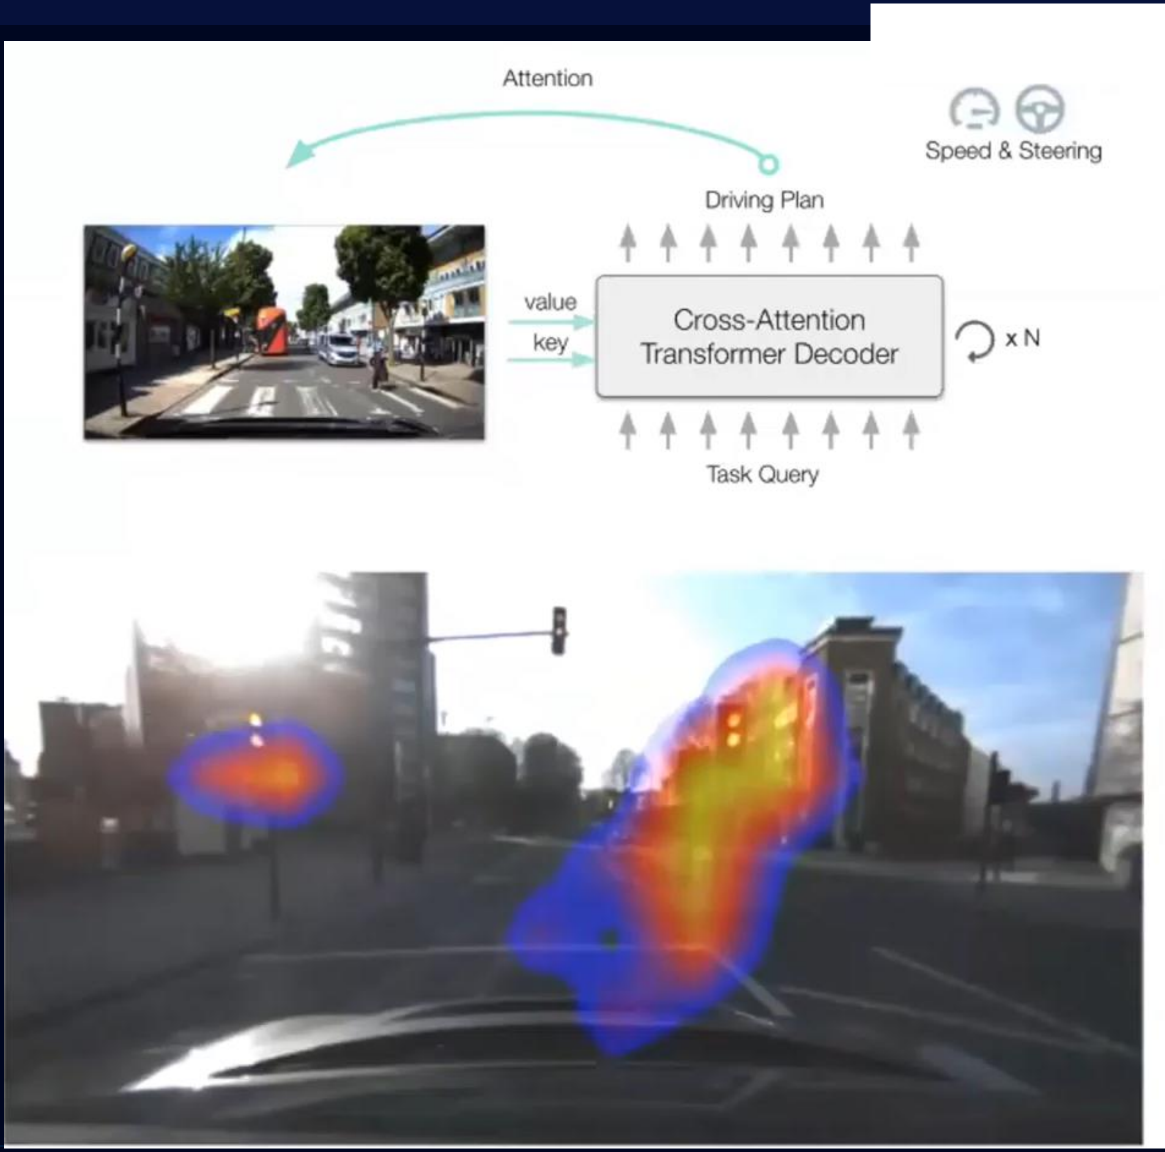

Hybrid: Joint Perception and Prediction

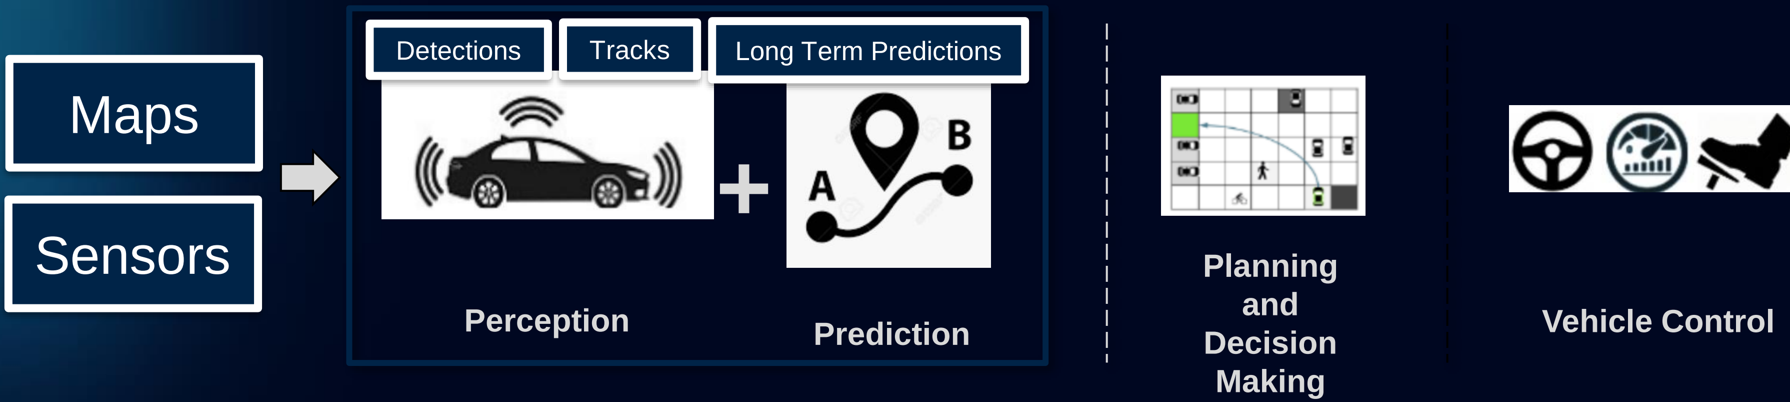

- Mix of both Modular and end-to-end

Hybrid: Joint Perception and Prediction

Maps

Sensors

Detections | Tracks | Long Term Predictions

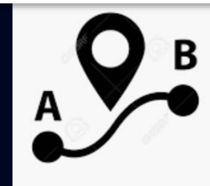

Perception

Prediction

Planning and Decision Making

Vehicle Control

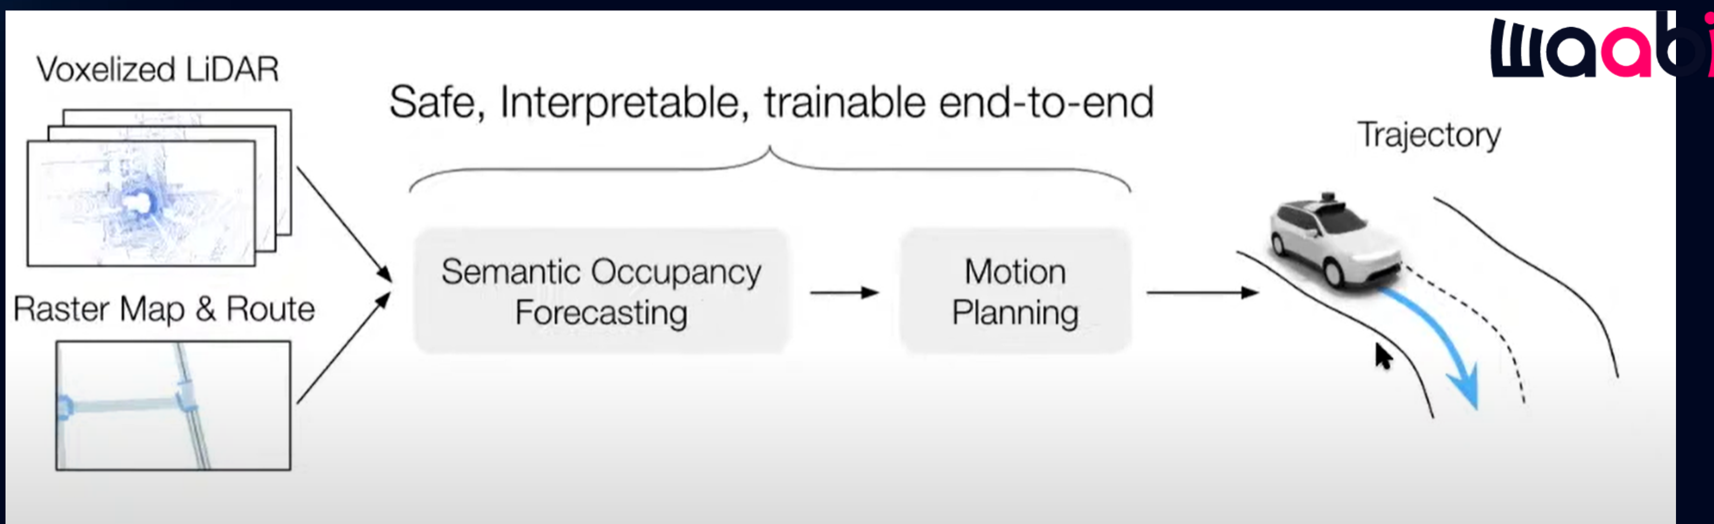

J. Phillips, J. Martinez, I. A. BÃ¢rsan, S. Casas, A. Sadat and **R. Urtasun**
Deep Multi-Task Learning for Joint Localization, Perception, and Prediction
*In Conference on Computer Vision and Pattern Recognition (**CVPR**), Virtual, June 2021*

S. Casas, A. Sadat and **R. Urtasun**
MP3: A Unified Model to Map, Perceive, Predict and Plan **(oral)**
*In Conference on Computer Vision and Pattern Recognition (**CVPR**), Virtual, June 2021*

Classical vs Modern ML Systems

The Factory: Data Engine

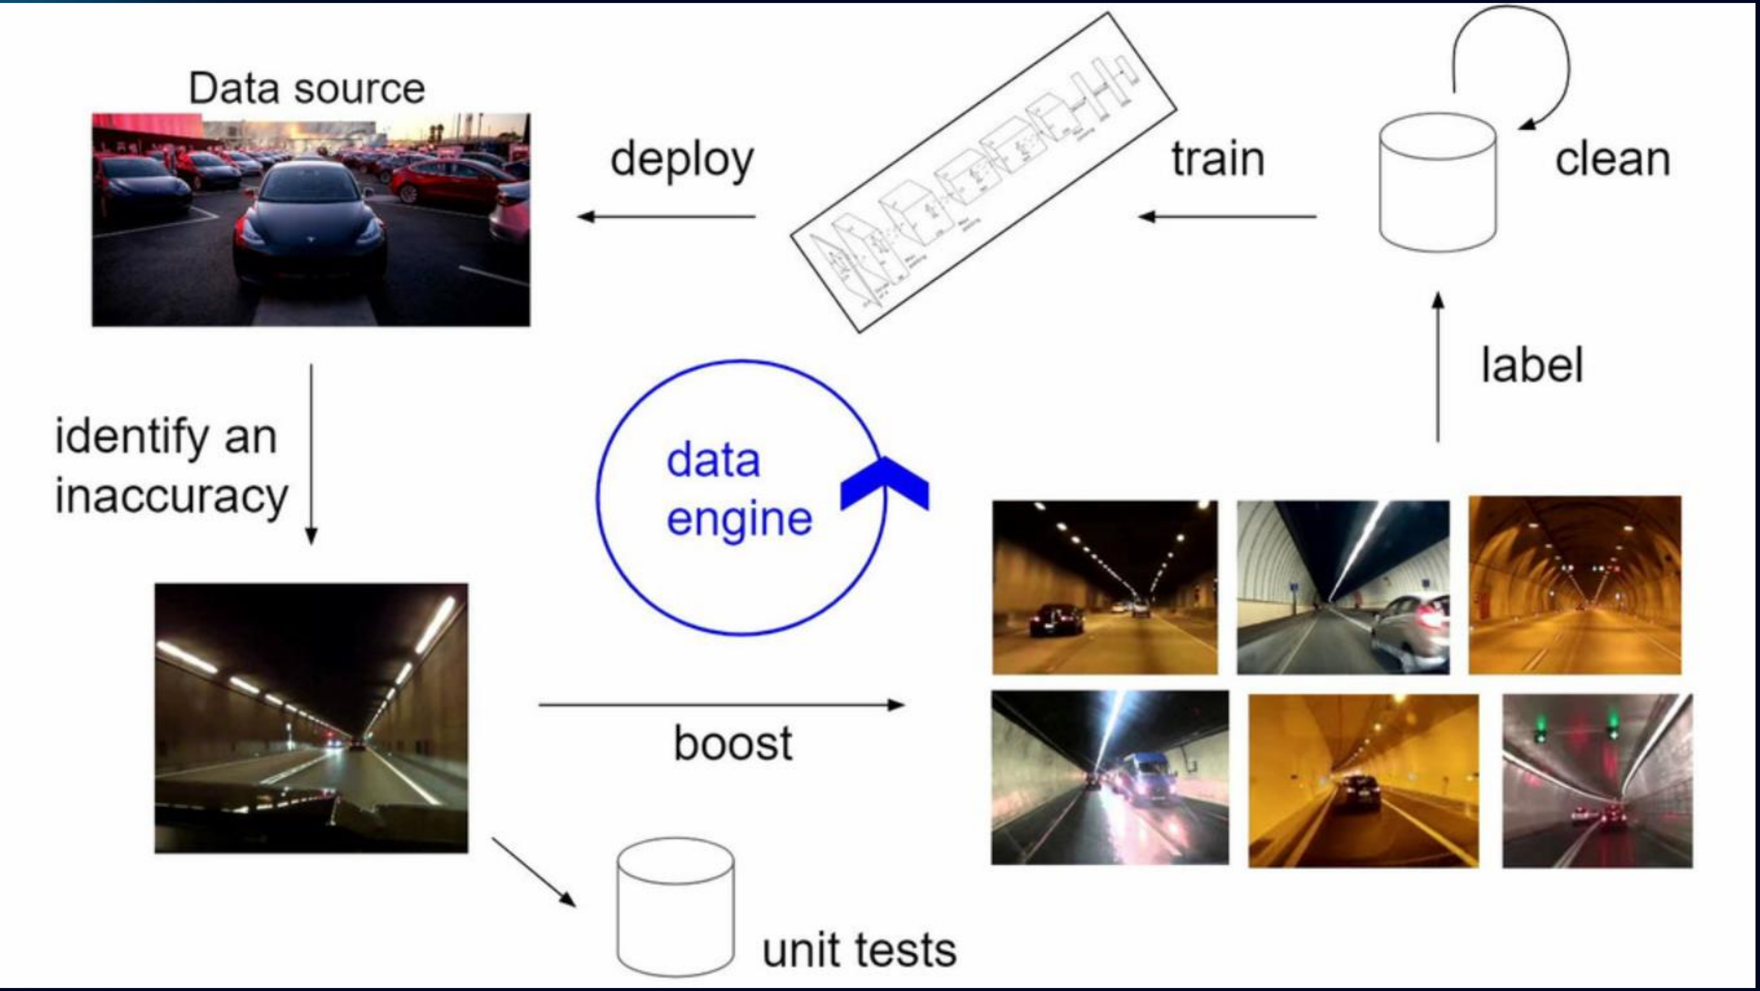

Sensors: Complementary and Redundant

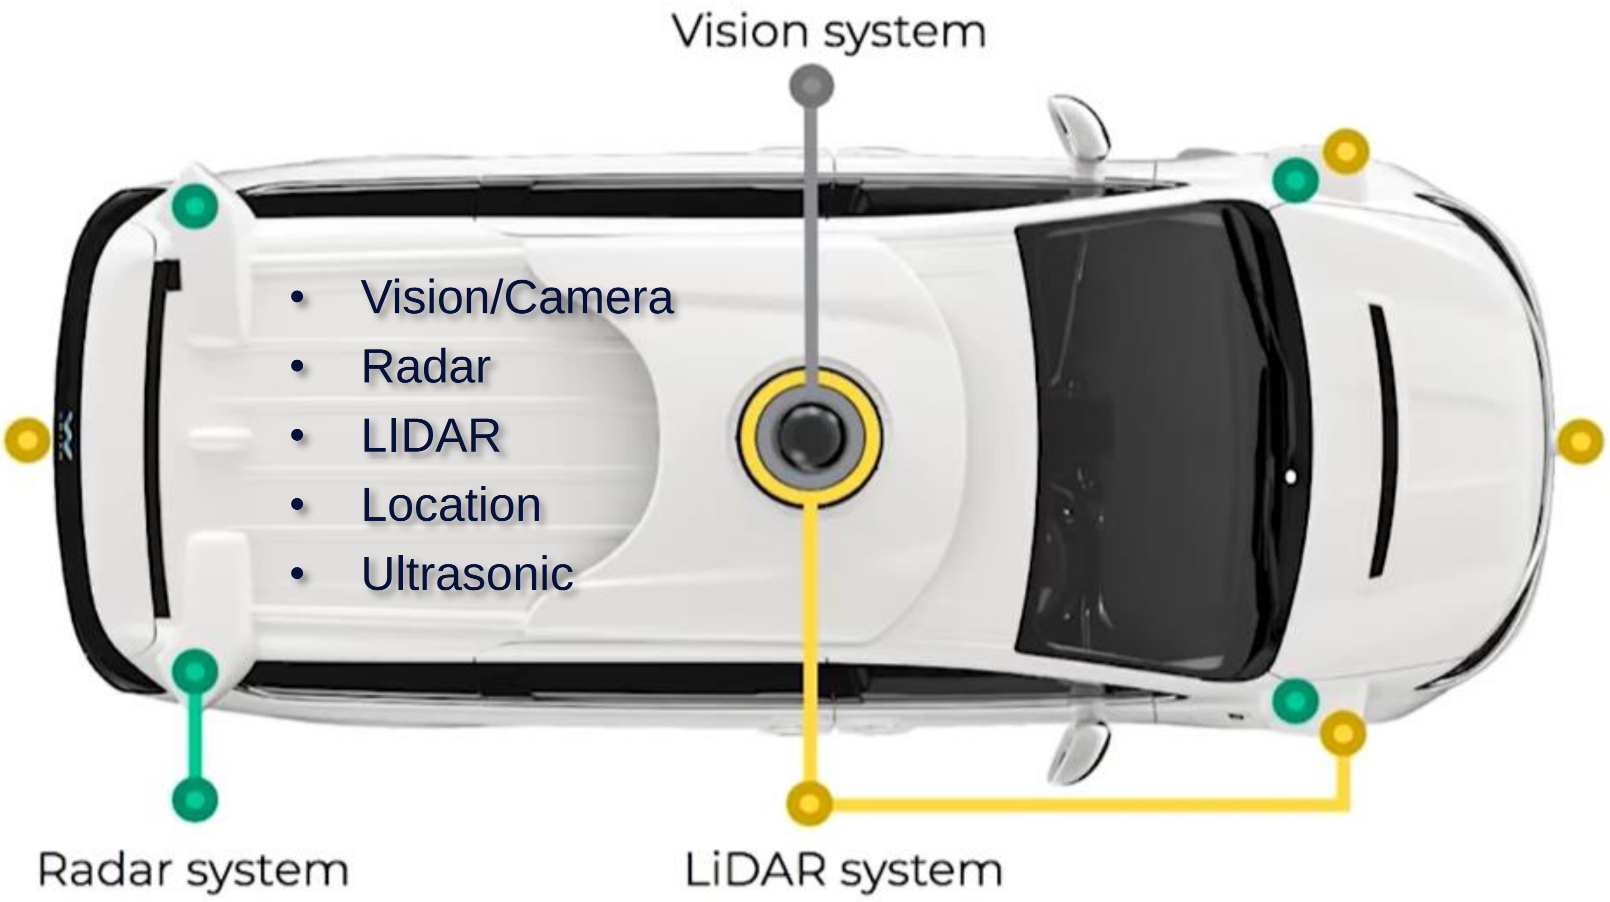

- Vision/Camera
- Radar
- LIDAR
- Location
- Ultrasonic

Boeing 737 MAX AOA sensor did not have proper redundancy!

Sensors: Complementary and Redundant

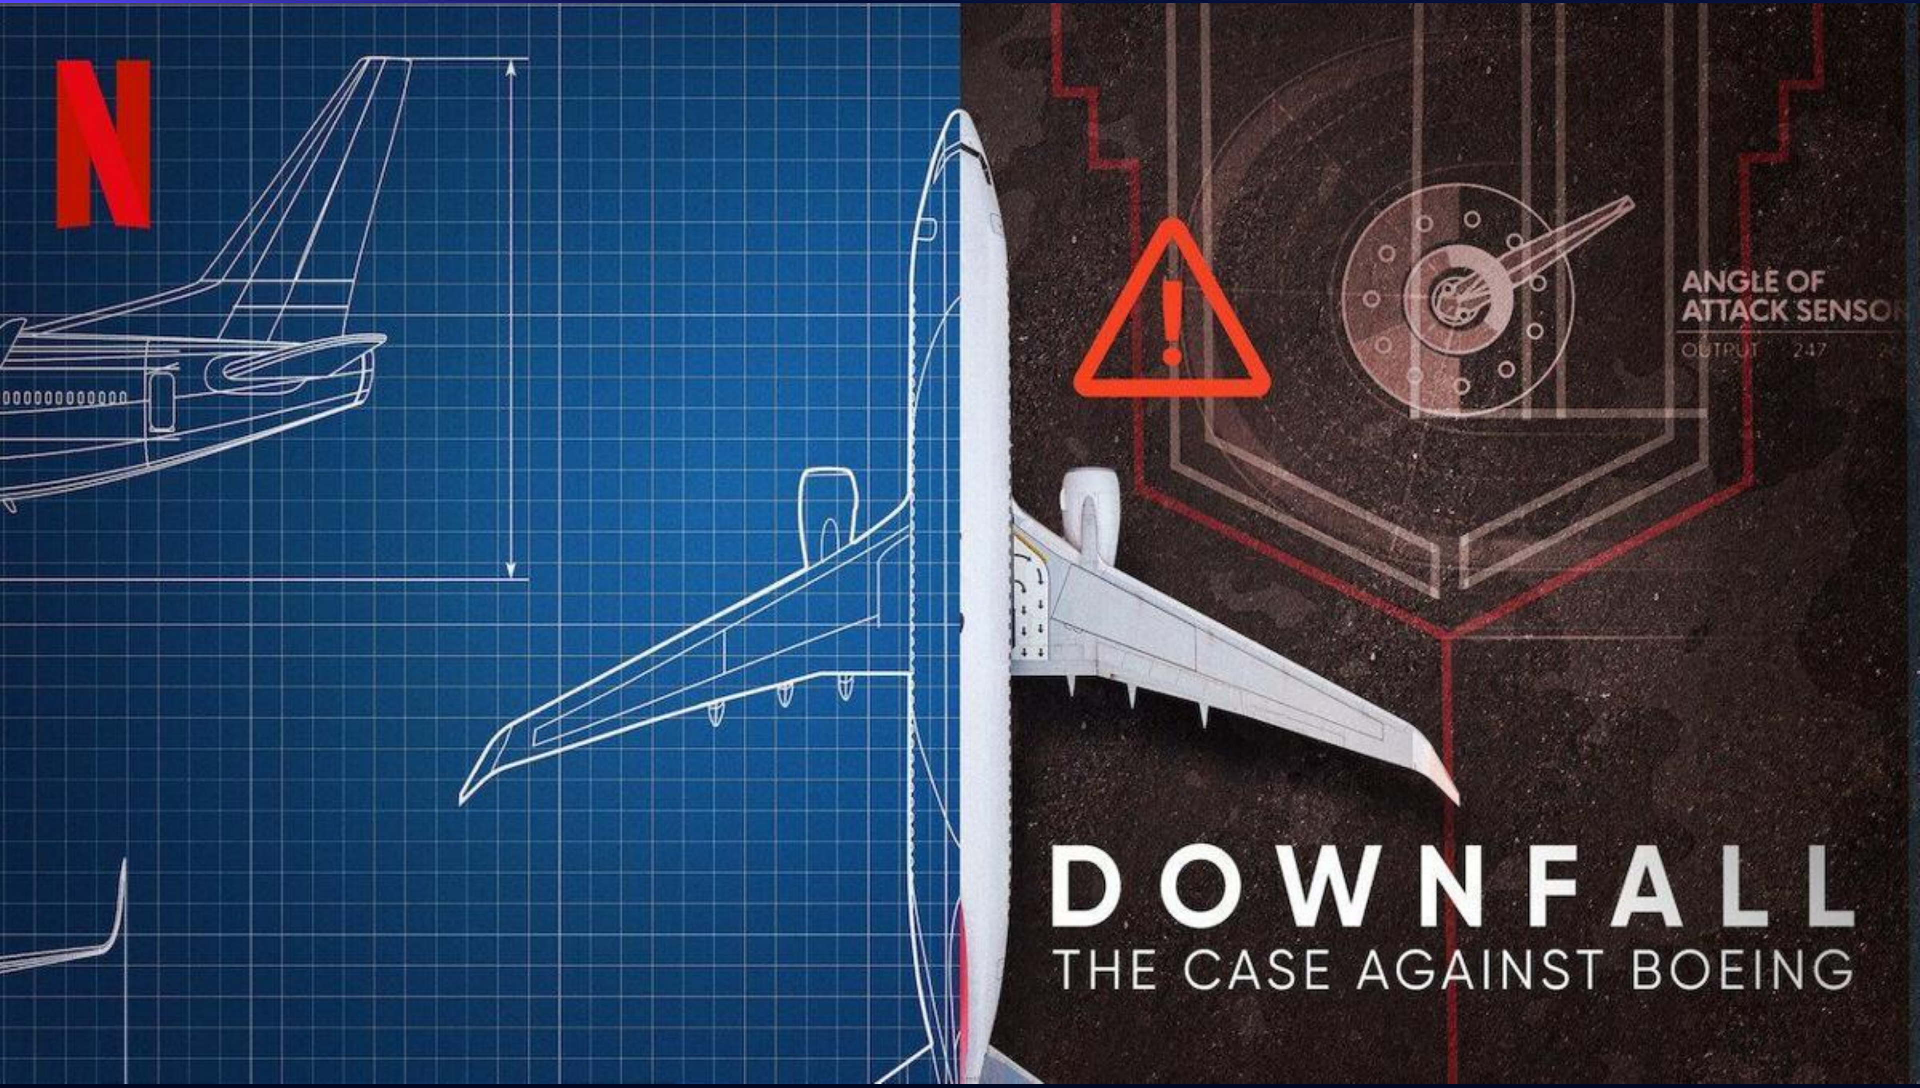

Boeing 737 Max

Federal Aviation Administration

Use of Single Angle of Attack (AOA) Sensor In the original design, erroneous data from a single AOA sensor activated MCAS and subsequently caused airplane nose-down trim of the horizontal stabilizer.

In the new design, Boeing eliminated MCAS reliance on a single AOA sensor signal by using both AOA sensor inputs and through flight-control law changes that include safeguards against failed or erroneous AOA indications. The updated FCC software with revised flight-control laws uses inputs from both AOA sensors to activate MCAS. This is in contrast to the original MCAS design, which relied on data from only one sensor at a time, and allowed repeated MCAS activation as a result of input from a single AOA sensor. The updated FCC software compares the inputs from the two sensors to detect a failed AOA sensor**. If the difference between the AOA sensor inputs is above a calculated threshold, the FCC will disable the STS, including its MCAS function, for the remainder of that flight and provide a corresponding indication of such deactivation on the flight deck.**

Date: November 18, 2020

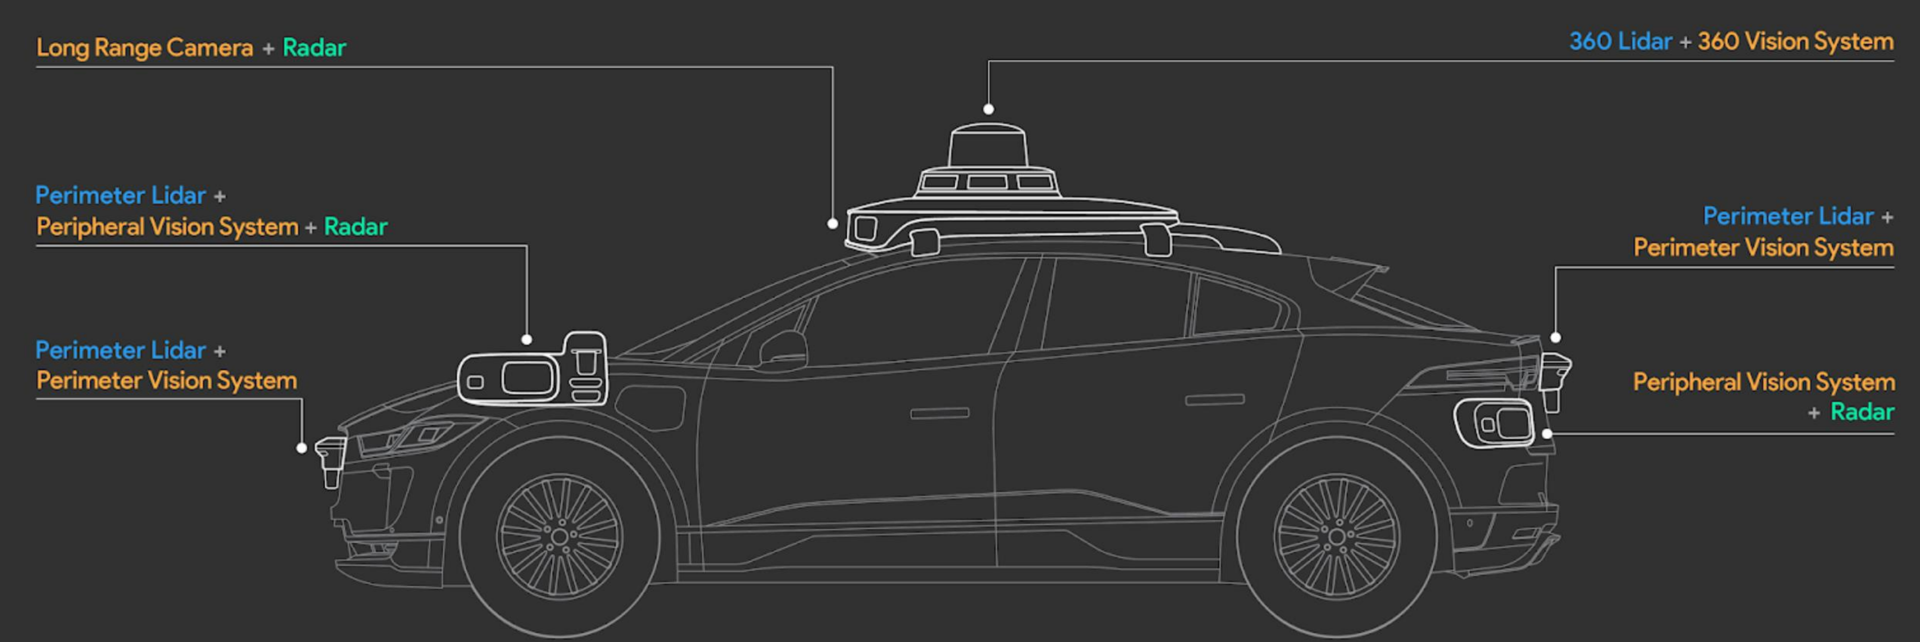
Long Range Camera + Radar
360 Lidar + 360 Vision System
Perimeter Lidar +
Peripheral Vision System + Radar
Perimeter Lidar +
Perimeter Vision System
Perimeter Lidar +
Perimeter Vision System
Perimeter Lidar +
Perimeter Vision System
Peripheral Vision System
+ Radar


Sensor-Setup Comparison

- **Tesla Model S**: Reliable environment perception only in standard driving use cases; most corner cases not manageable or with severe restrictions and no redundancy
- **Google “Koala”**: Very good near-field due to availability of LiDAR. Two different types of LiDAR
- **Uber XC 90**: Problems anticipated for the case of entering priority roads due to missing corner LiDARs at the front (they have to rely on radar only!)

Simulation is essential

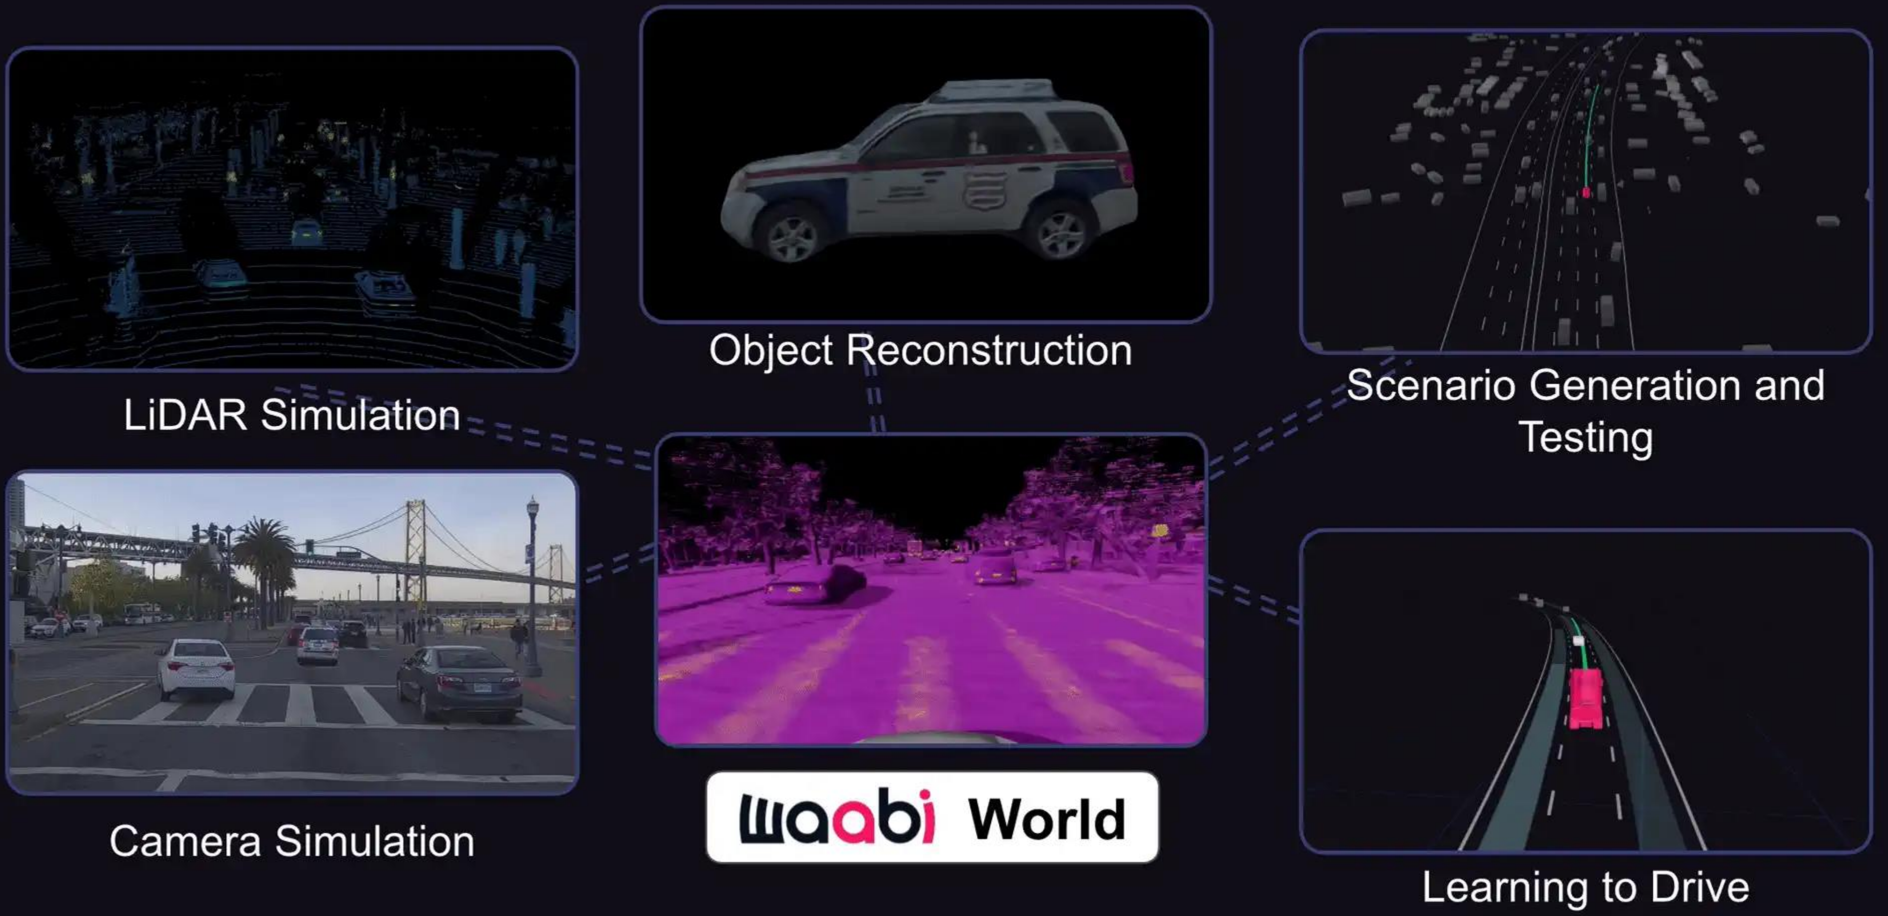
LiDAR Simulation
Object Reconstruction
Scenario Generation and Testing
Camera Simulation
Waabi World
Learning to Drive


World State
Realistic Sensor Simulation
Waabi Driver
Update world state
SDV-Updated World State
Reactive and Intelligent Actor Models
Update world state

Simulation – all the way to the Sensors!

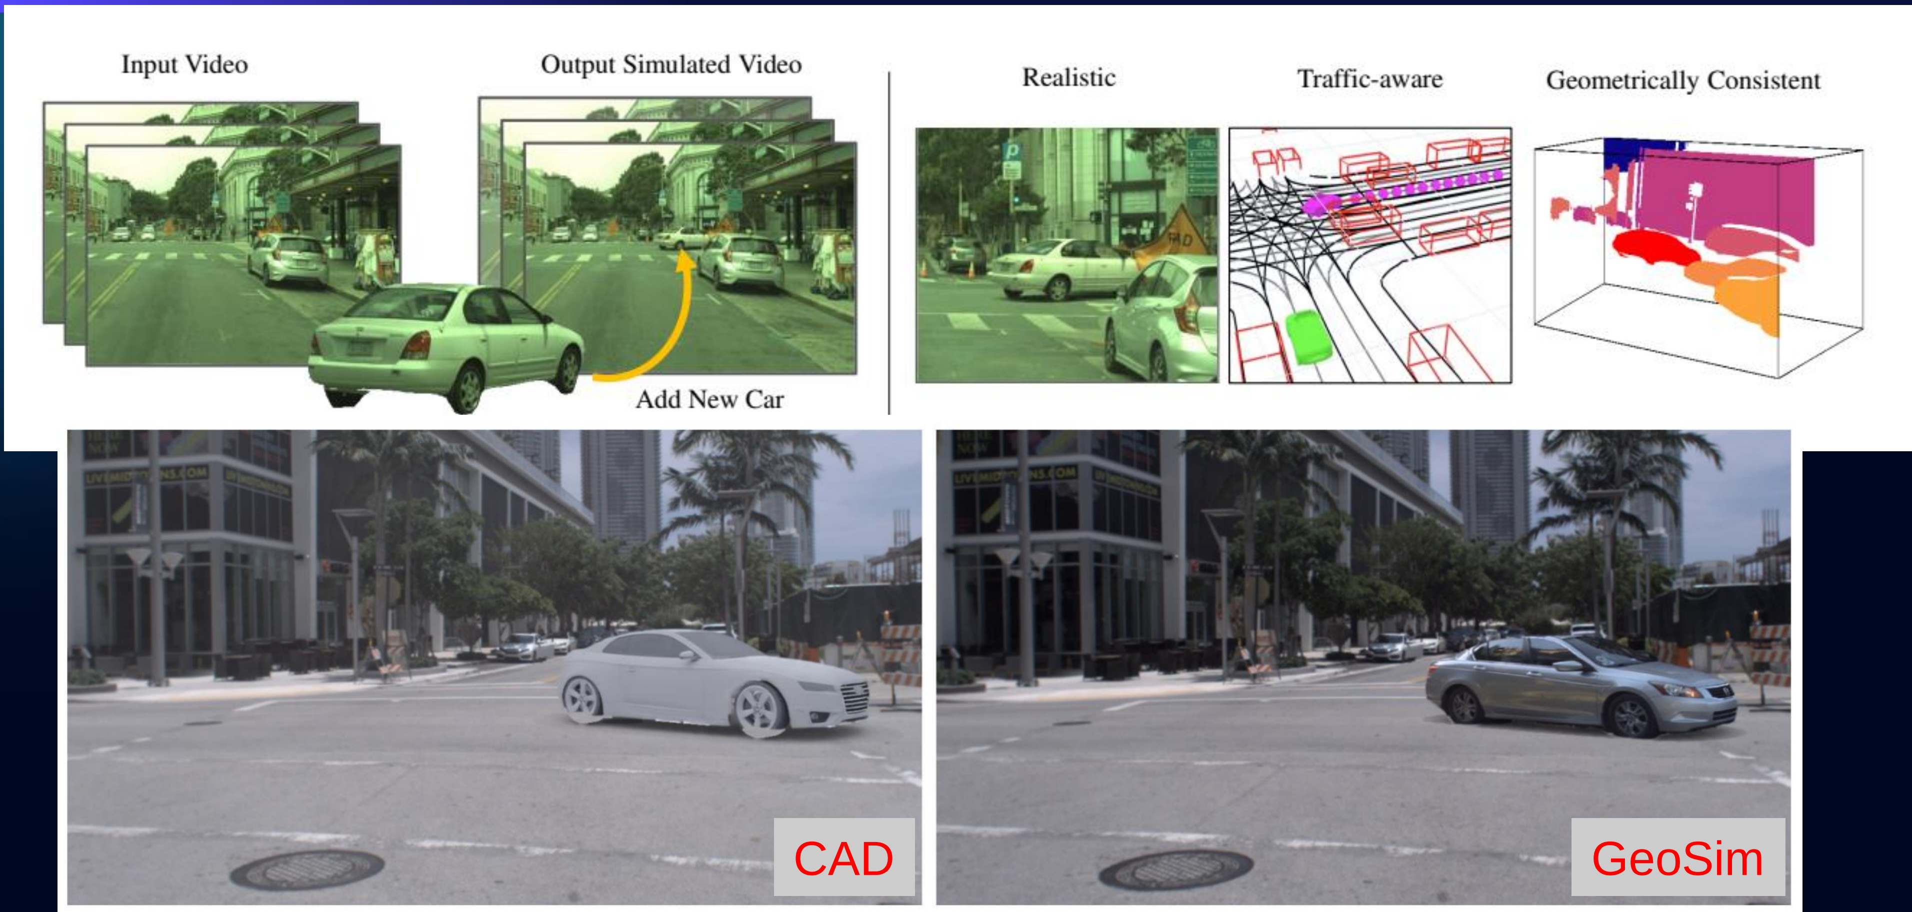

Y. Chen, F. Rong, S. Duggal, S. Wang, X. Yan, S. Manivasagam, S. Xue, E. Yumer **R. Urtasun**
GeoSim: Photorealistic Image Simulation with Geometry-Aware Composition **(oral)**
*In Conference on Computer Vision and Pattern Recognition (**CVPR**), Virtual, June 2021*

Simulation

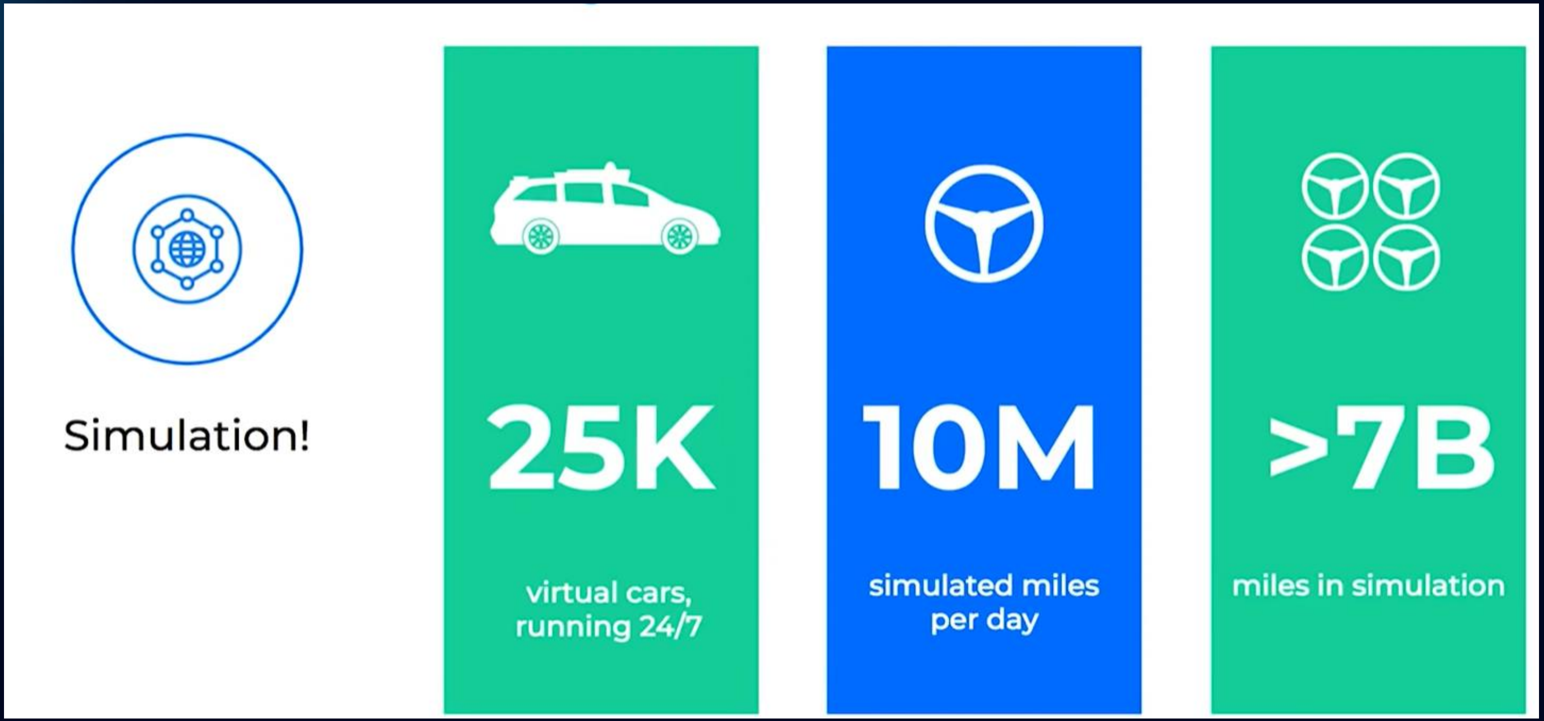

S. Tan, K. Wong, S. Wang, S. Manivasagam, M. Ren and **R. Urtasun**
SceneGen: Learning to Generate Realistic Traffic Scenes
*In Conference on Computer Vision and Pattern Recognition (**CVPR**), Virtual, June 2021*

J. Wang, A. Pun, J. Tu, S. Manivasagam, A. Sadat, S. Casas, M. Ren, **R. Urtasun**
AdvSim: Generating Safety-Critical Scenarios for Self-Driving Vehicles
*In Conference on Computer Vision and Pattern Recognition (**CVPR**), Virtual, June 2021*

Copilot4D: Fundamental Models for AV

Similar to how LLMs learn by predicting the next word in a sentence, Copilot4D learns by predicting how a machine will observe the world in the future. However, while LLMs learn from discrete tokens that represent words, LiDAR data is continuous in nature. To bridge this gap between language and the physical world, Copilot4D features a 3-stage architecture.

- First, a LiDAR tokenizer abstracts continuous sensor data into a set of discrete tokens, similar to words in language.
- Then, our foundation model forecasts how the world will evolve as a set of tokens, leveraging the recent breakthroughs in LLMs. Importantly, it takes into account how the future actions of the embodied AI agent will affect the world.
- Finally, a LiDAR renderer brings these tokens back to LiDAR point clouds, something robots can observe just like humans see through their eyes, enabling us to learn from raw sensor recordings without requiring human supervision.

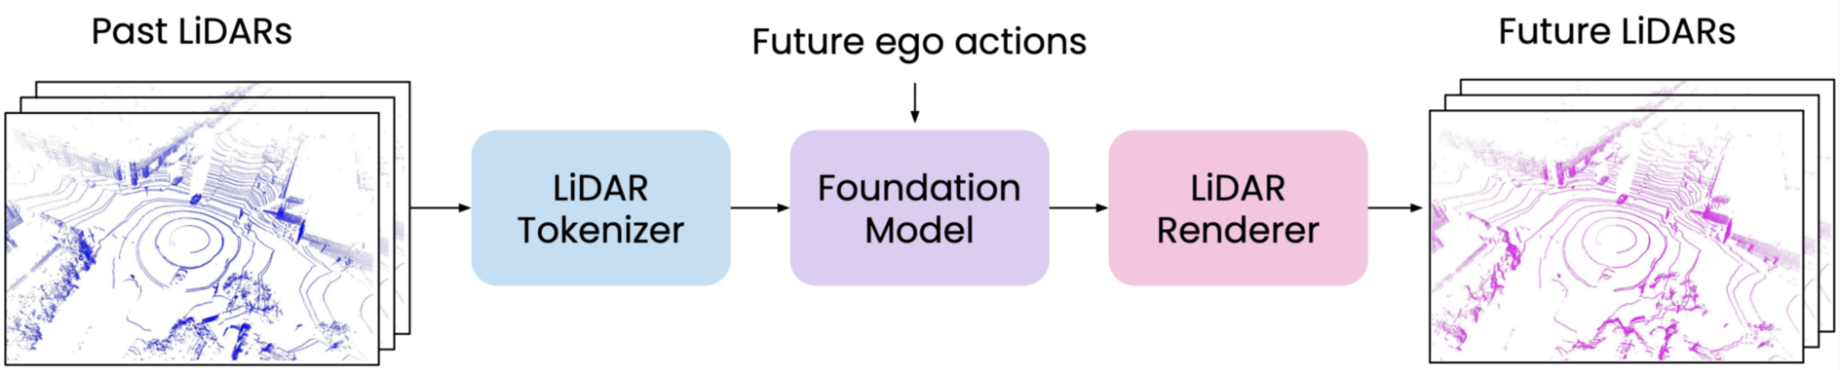

Copilot4D predicts future LiDAR point clouds from a history of past LiDAR observations, akin to how LLMs predict the next word given the preceding text. We design a 3 stage architecture that is able to exploit all the breakthroughs in LLMs to bring the first 4D foundation model.

Test Facility: Waymo

Castle, CA (91 acres)

What is a good mAP?

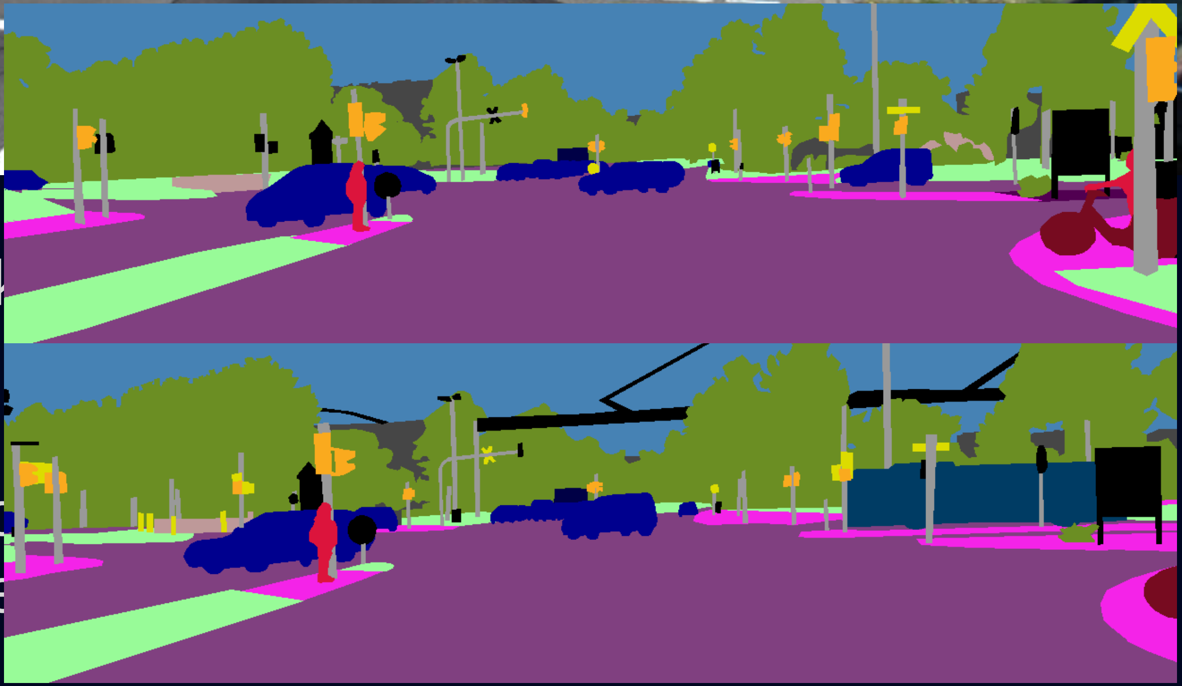

- Most metrics lacl
  on the road
- Consider **mAP** fc
  threshold for whe

Is recall of 96.6 for Cars enough?

3D Object Detection Results on KITTI Dataset

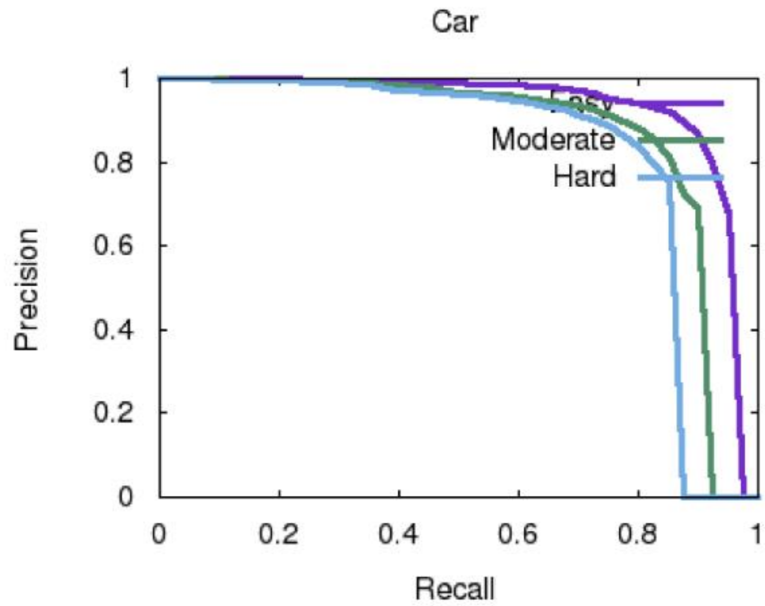

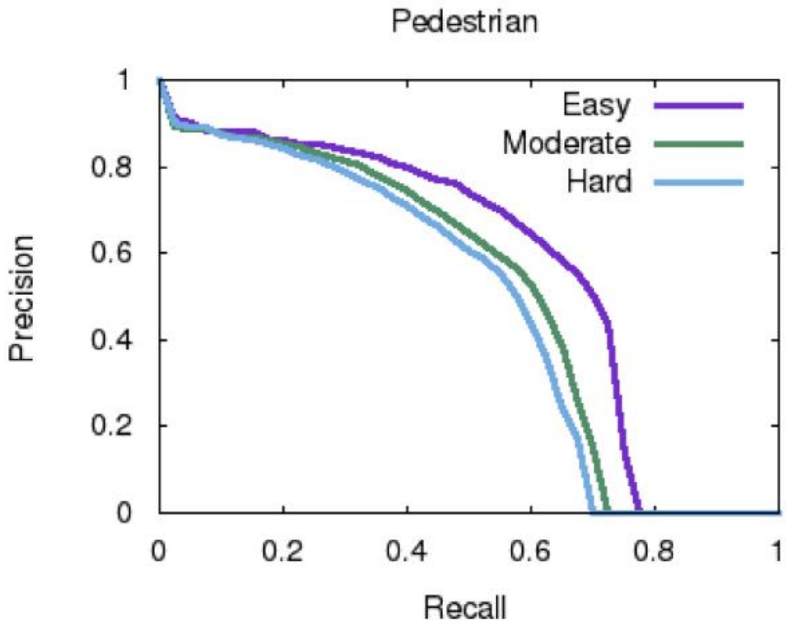

- Struggles even with the “easy” cases of pedestrians
- Hard to decide on the exact expectations on precision and recall in object detection

Unit Testing Perception Systems – example

Unit Testing Perception Systems – examples

Unit Testing Perception Systems – examples

Also extend to 3D (*presence* and *absence* polygons)

Compute Hardware

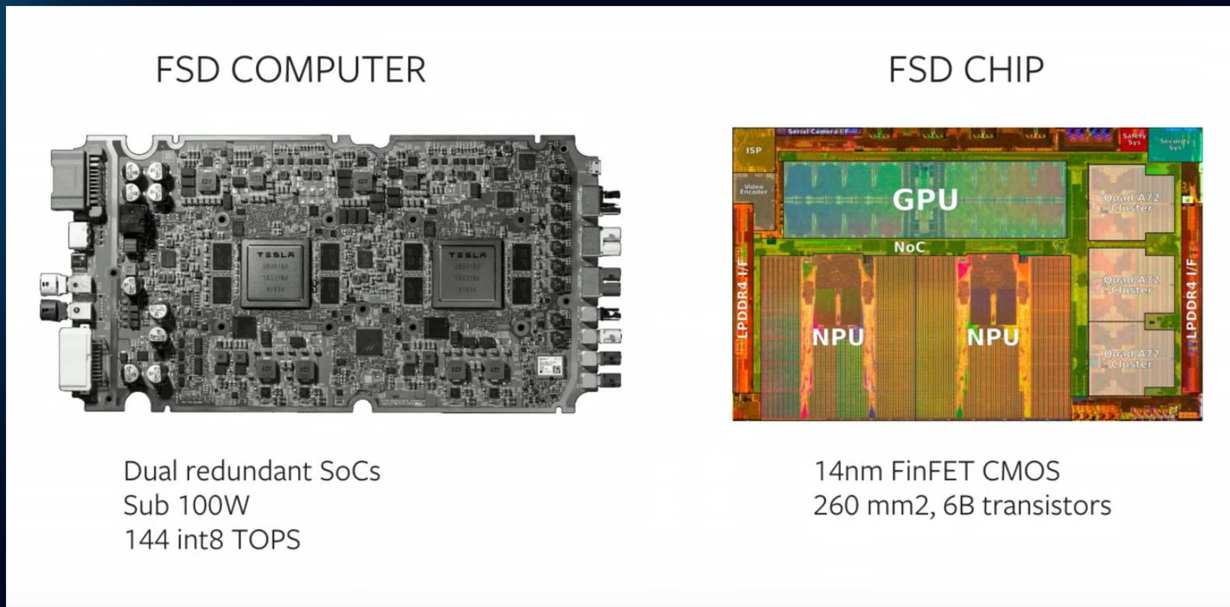

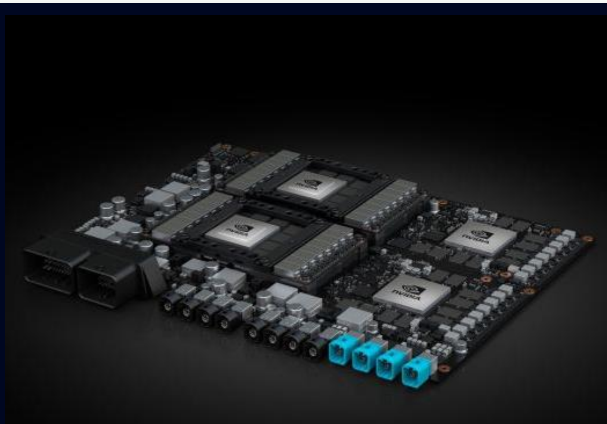

DRIVE PX PEGASUS

LEVEL 5: DRIVERLESS ROBOTAXIS

320 TOPS for AI Inferencing

ASIL D Functional Safety

Few Companies going Autonomous

Audi
Bai du 百度
BOSCH
DAIMLER
DELPHI
Ford
GM
SAIC 上汽集团 SAIC MOTOR
Google
HONDA
HYUNDAI
JAGUAR
LAND ROVER
Mercedes-Benz
MOBILEYE
NISSAN
RENAULT
SCANIA
TATA ELXSI engineering creativity
TESLA
TOYOTA
UBER
Volkswagen
VOLVO
YUTONG
ZOOX
www.cbinsights.com
Aurora
WAYVE
waabi
EINRIDE
ARGO AI
WAYMO
cruise
nuro
TESLA

When will we get truly Autonomous Cars?

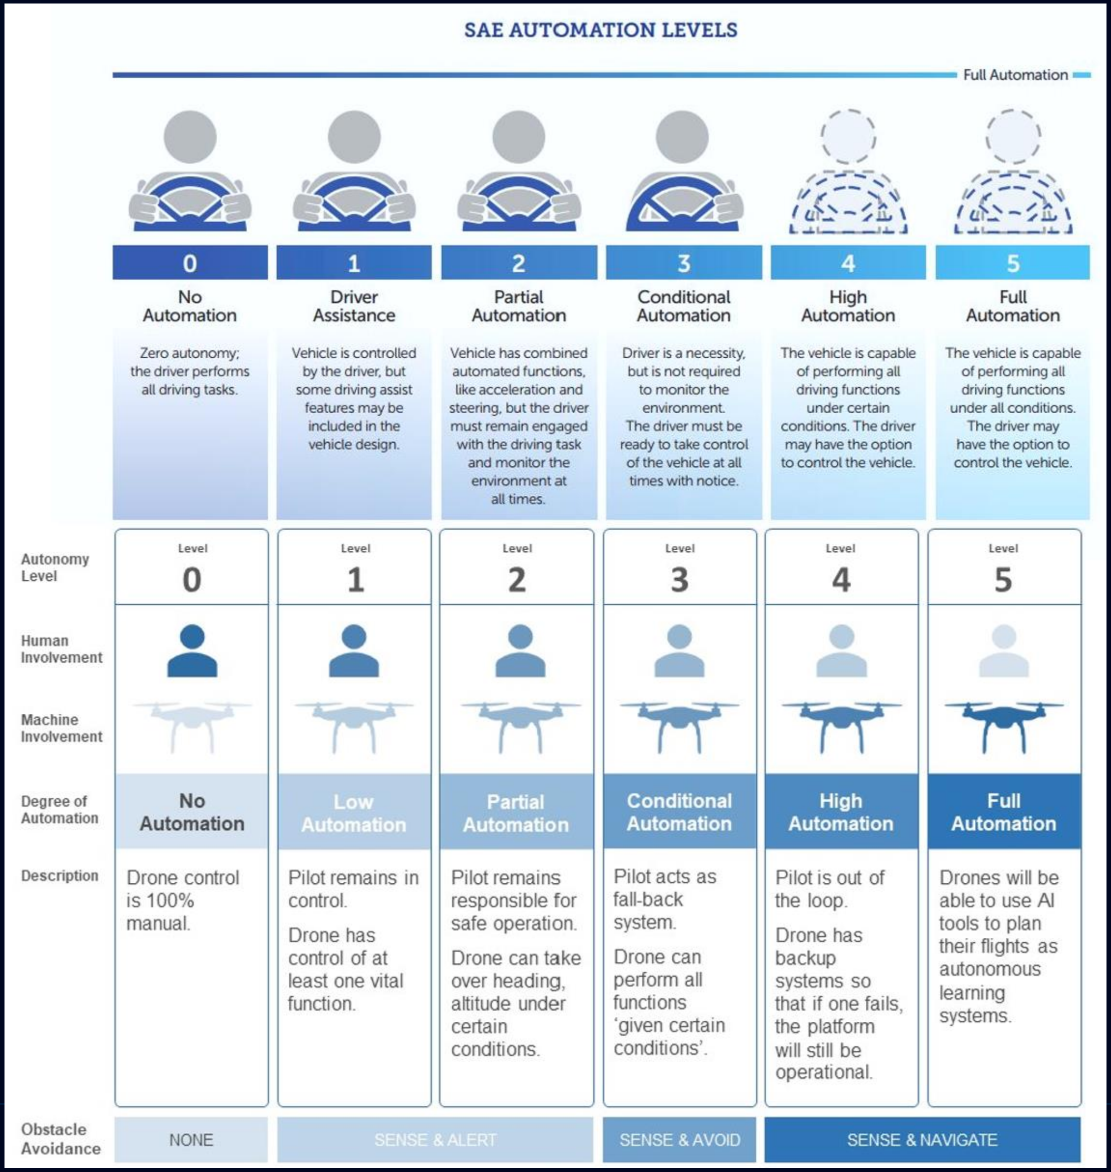

How far have we come?

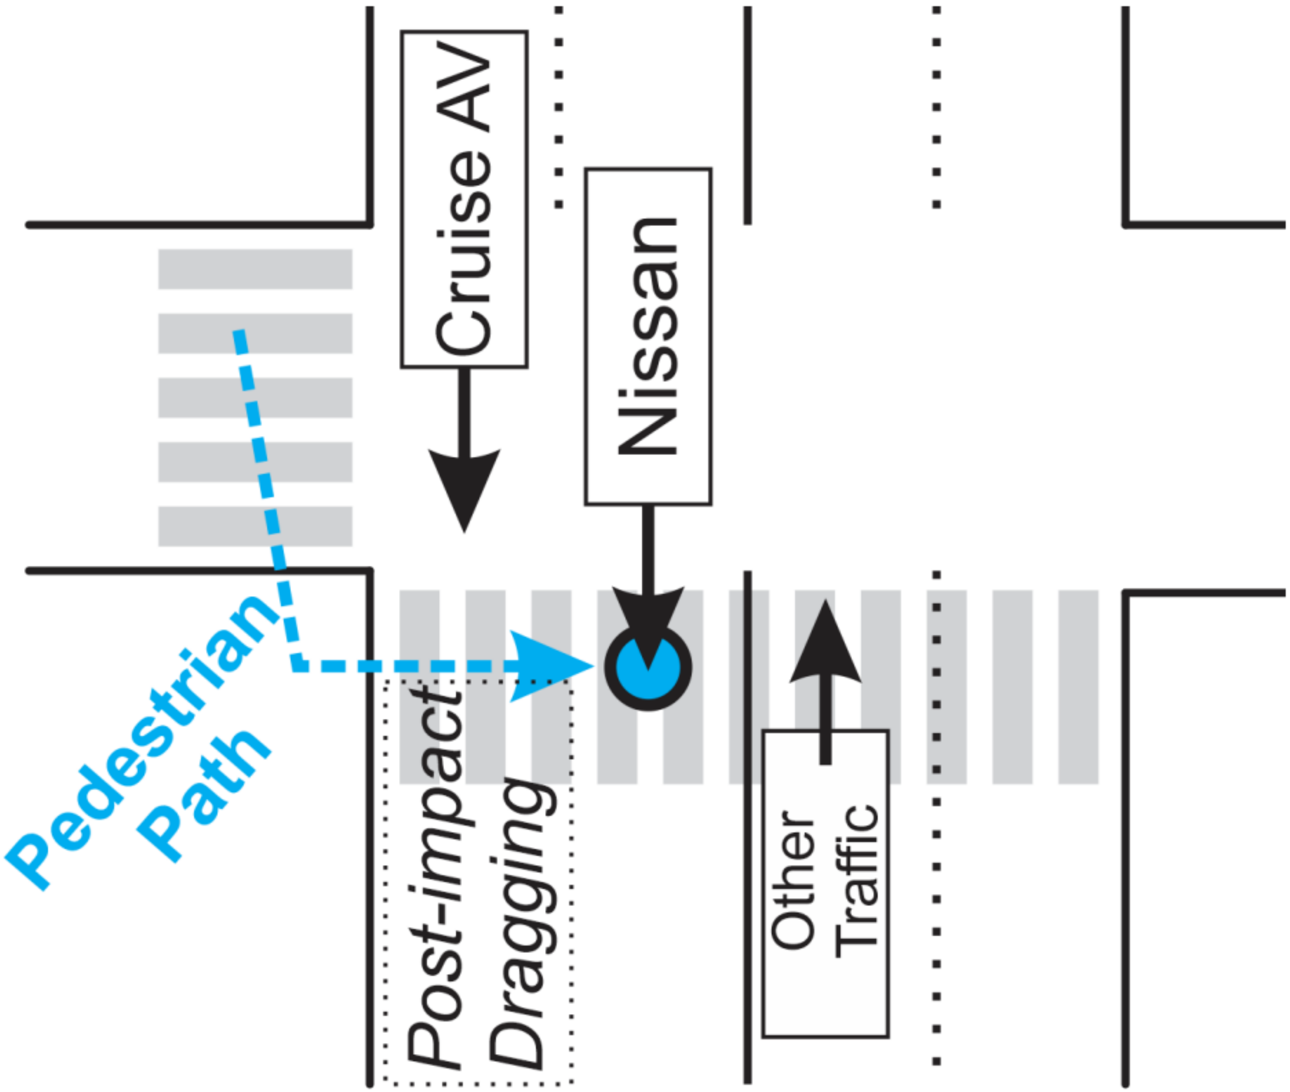

Figure 1. Simplified diagram of mishap; not to scale.

Cite as: arXiv:2402.06046 **[cs.RO]**
(or arXiv:2402.06046v2 **[cs.RO]** for this version)
https://doi.org/10.48550/arXiv.2402.06046


apollo
apollo

When will we get truly Autonomous Cars?

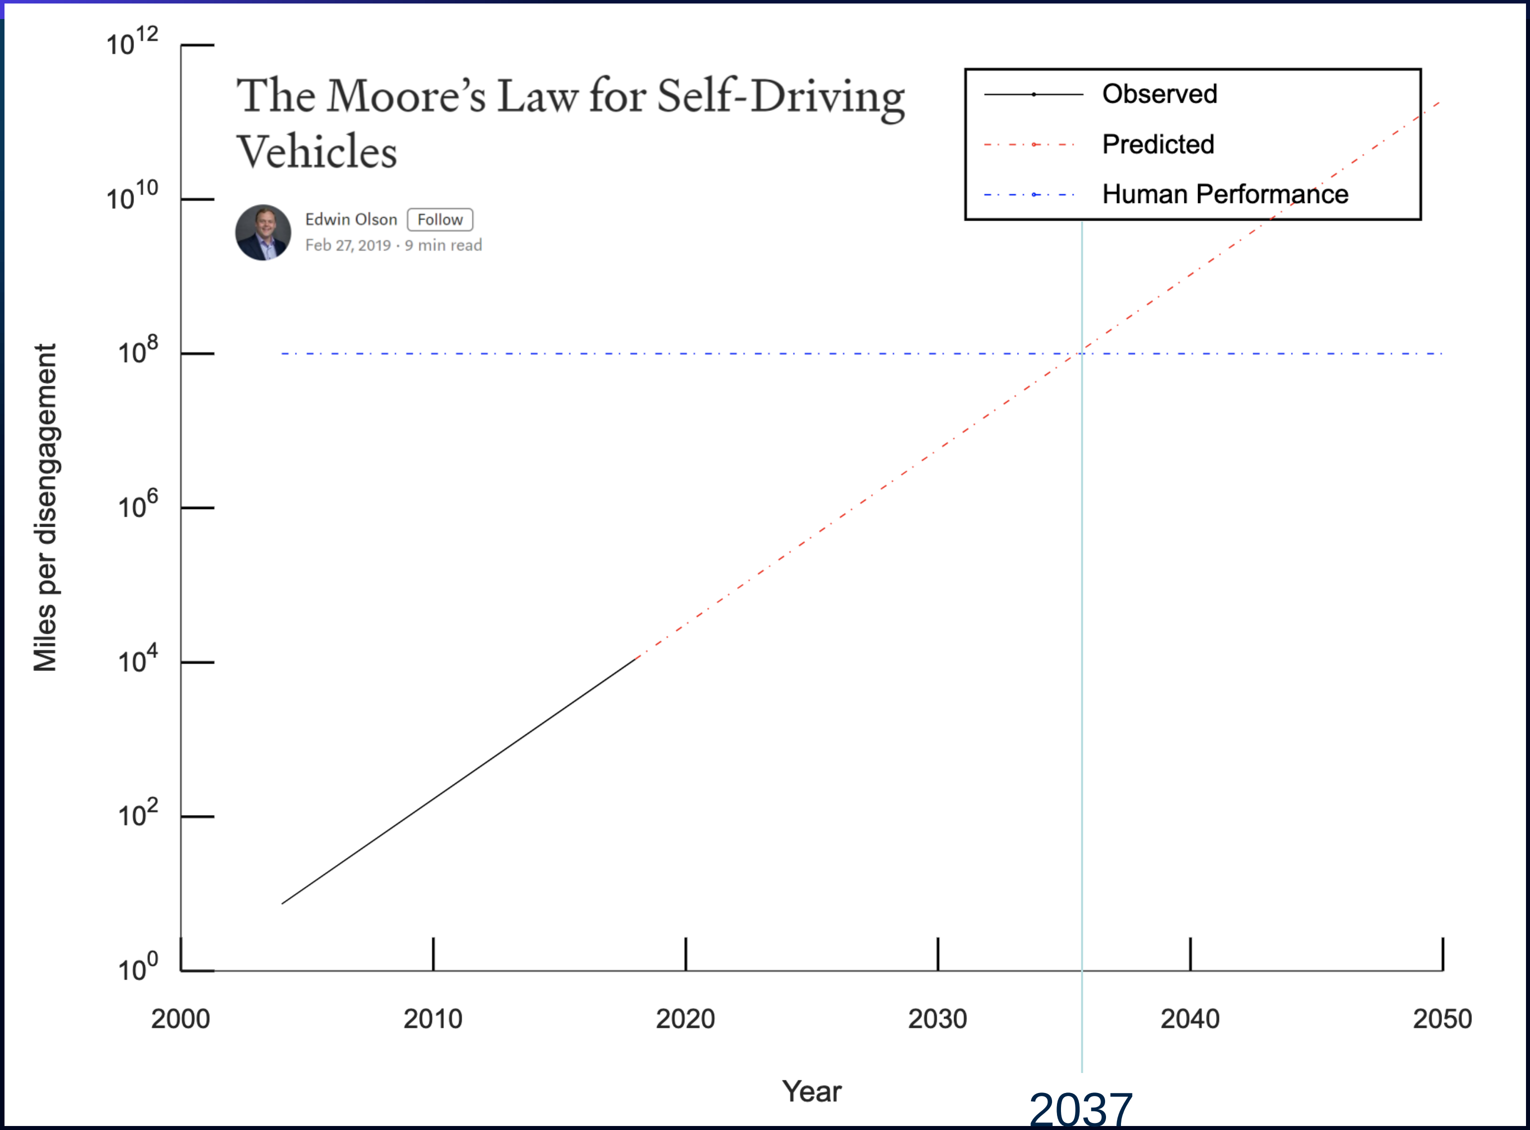

Case for India?

Tesla: Urban and Highway Full AutoPilot

FastCodeAI

Thank you!

http://arjunjain.co.in

arjun@fastcode.ai