mobileye
An Intel Company

Mobileye "Under the Hood"

CES 2021

Prof. Amnon Shashua

Business Pillars

intel | mobileye An Intel Company | moovit

ADAS

L1- L2++ Advance Driver Assistance Systems

REM™ Mapping

HD maps for AV, ADAS and smart cities

Full Stack Self-Driving System

For consumer AV + Robotaxi MaaS

Our ADAS Business in 2020

19.3M

EyeQs shipped in 2020

10% YoY growth despite Covid-19 unprecedented challenges

49 RUNNING PROGRAMS

Globally across 28 OEMs

37 NEW DESIGN WINS

36.2M

Total lifetime volume of the new design wins

EyeQ Shipped (M units):

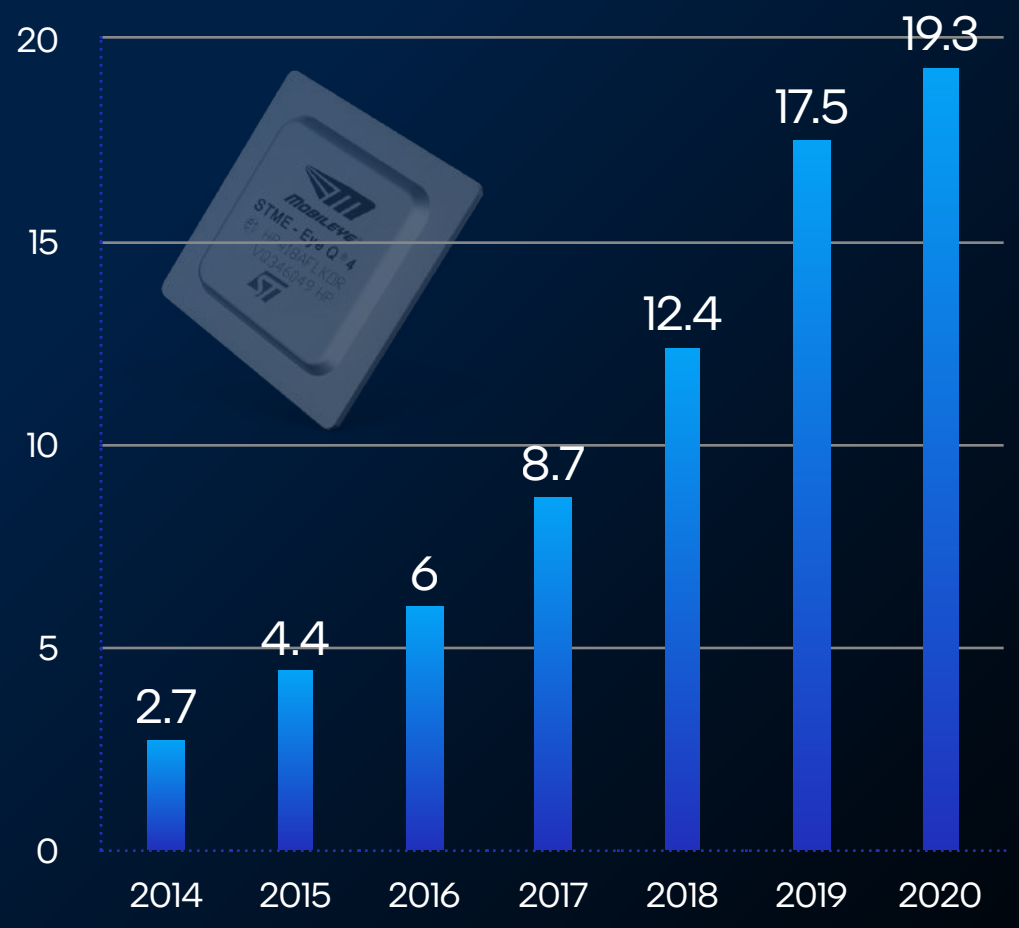

Product Portfolio

Moving beyond silicon

Silicon

EyeQ SoC + embedded ADAS SW

Silicon + PCB

Si+PCB integrated solution for L2++ (Supervision™)

Full stack SDS

Extension of the Si+PCB business line

PTOs, TNCs, and OEMs are potential customers

Including active sensors

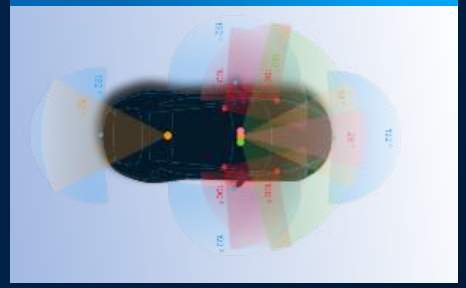

Full stack SDV VaaS

Based on AV-ready platform supplied by OEM partner

Potential customers

- Mobileye's MaaS
- PTOs
- TNCs

+Mobility intelligence layer

Customer-facing MaaS service

Full stack service

Owning the end customer

Moovit assets integrated in the service layers

Generating data from all 5 pillars to fuel our data offering for smart cities and infrastructure

The Trinity of Mobileye's Approach

The ADAS ←→ AV divide NOT range of capability, but MTBF

TRUE REDUNDANCY™

- AV is a system composed of independent subsystems; each is fully handsfree capable
- One of the subsystems is ADAS- we call it **SuperVision™**

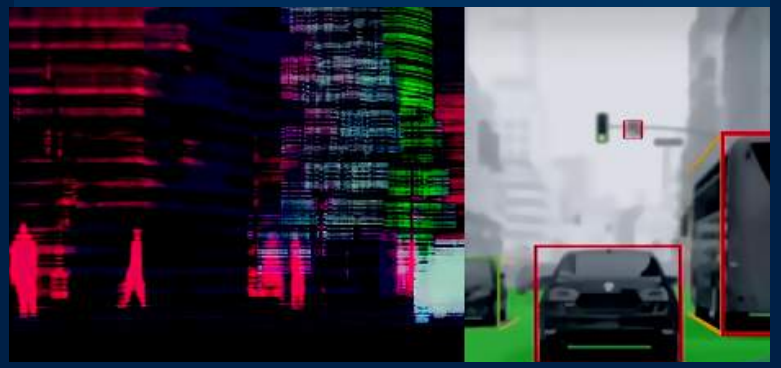

REM™- enabled scalability AV-Map Key to High MTBF

REM™

- To make this useful, geo scalability at low cost is imperative
- Crowdsource data collection followed by auto AV-map creation in the cloud
- Byproduct: data-driven business

RSS formal safety model Decision-making governance:

Responsibility-Sensitive Safety

- High MTBF is NOT sufficient for guaranteeing safety
- Need to guarantee that AV will not have "lapse of judgment"
- Standardizing human judgement (IEEE P2846)

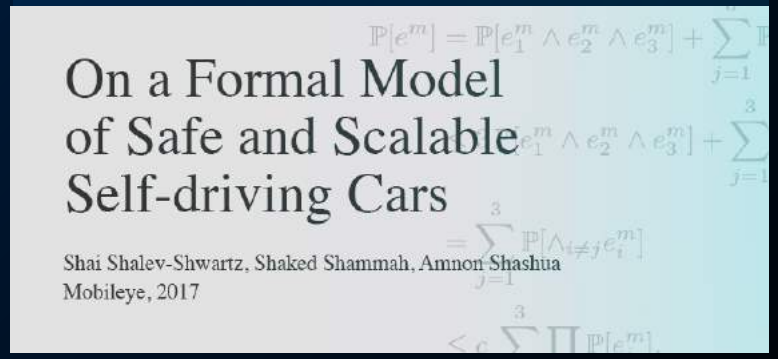

The Camera-only Subsystem

COMPUTE:

End-to-end operation based on 2x EyeQ5 running multiple independent computer vision engines for “algorithmic redundancy”

HW setup in our test vehicle

SENSOR SETUP:

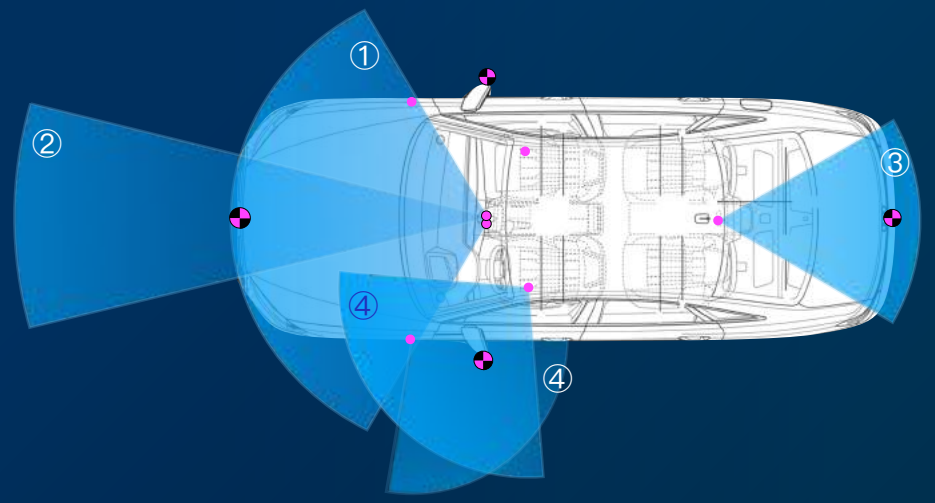

CAMERA

PARKING CAMERA

① Main	120°
② Narrow	28°
③ Rear	60°
④ Side x4	100°
⑤ Parking x4	192°

Water + air external cleaning solution

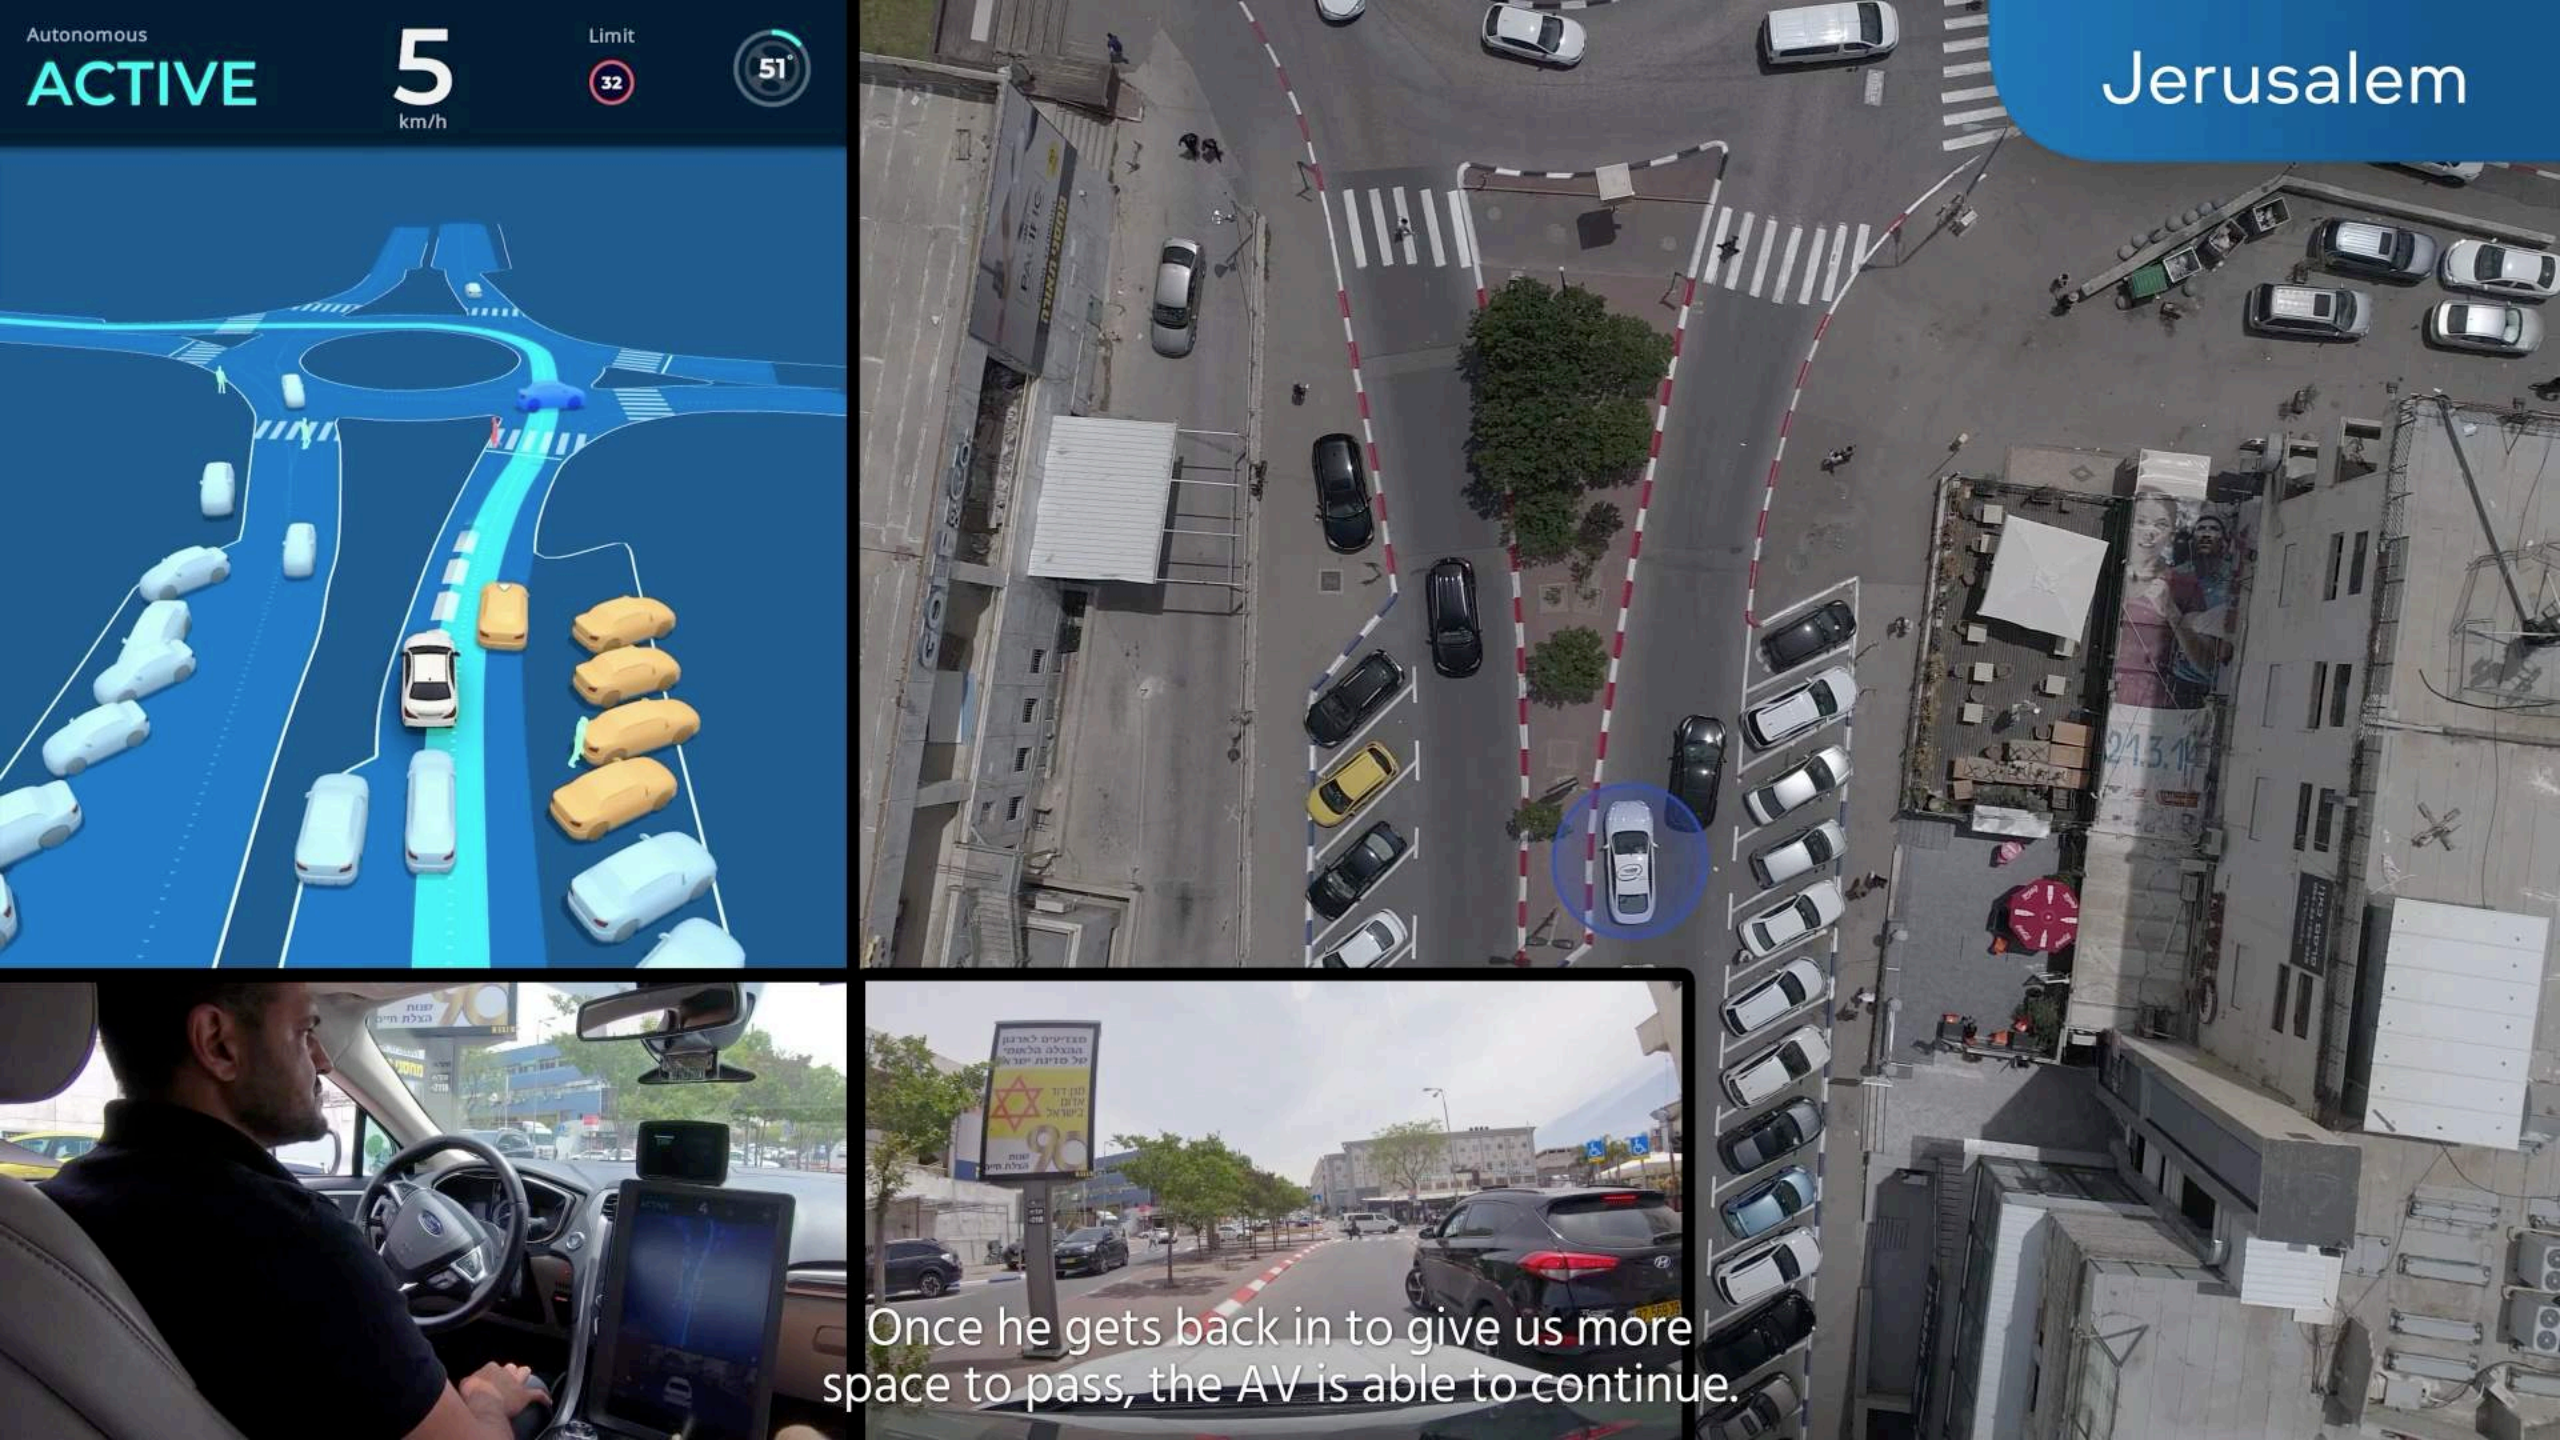
Autonomous
ACTIVE
5
km/h
Limit
32
51
Jerusalem
Once he gets back in to give us more
space to pass, the AV is able to continue.

Achieving the MTBF Goal for L2+ with Vidar

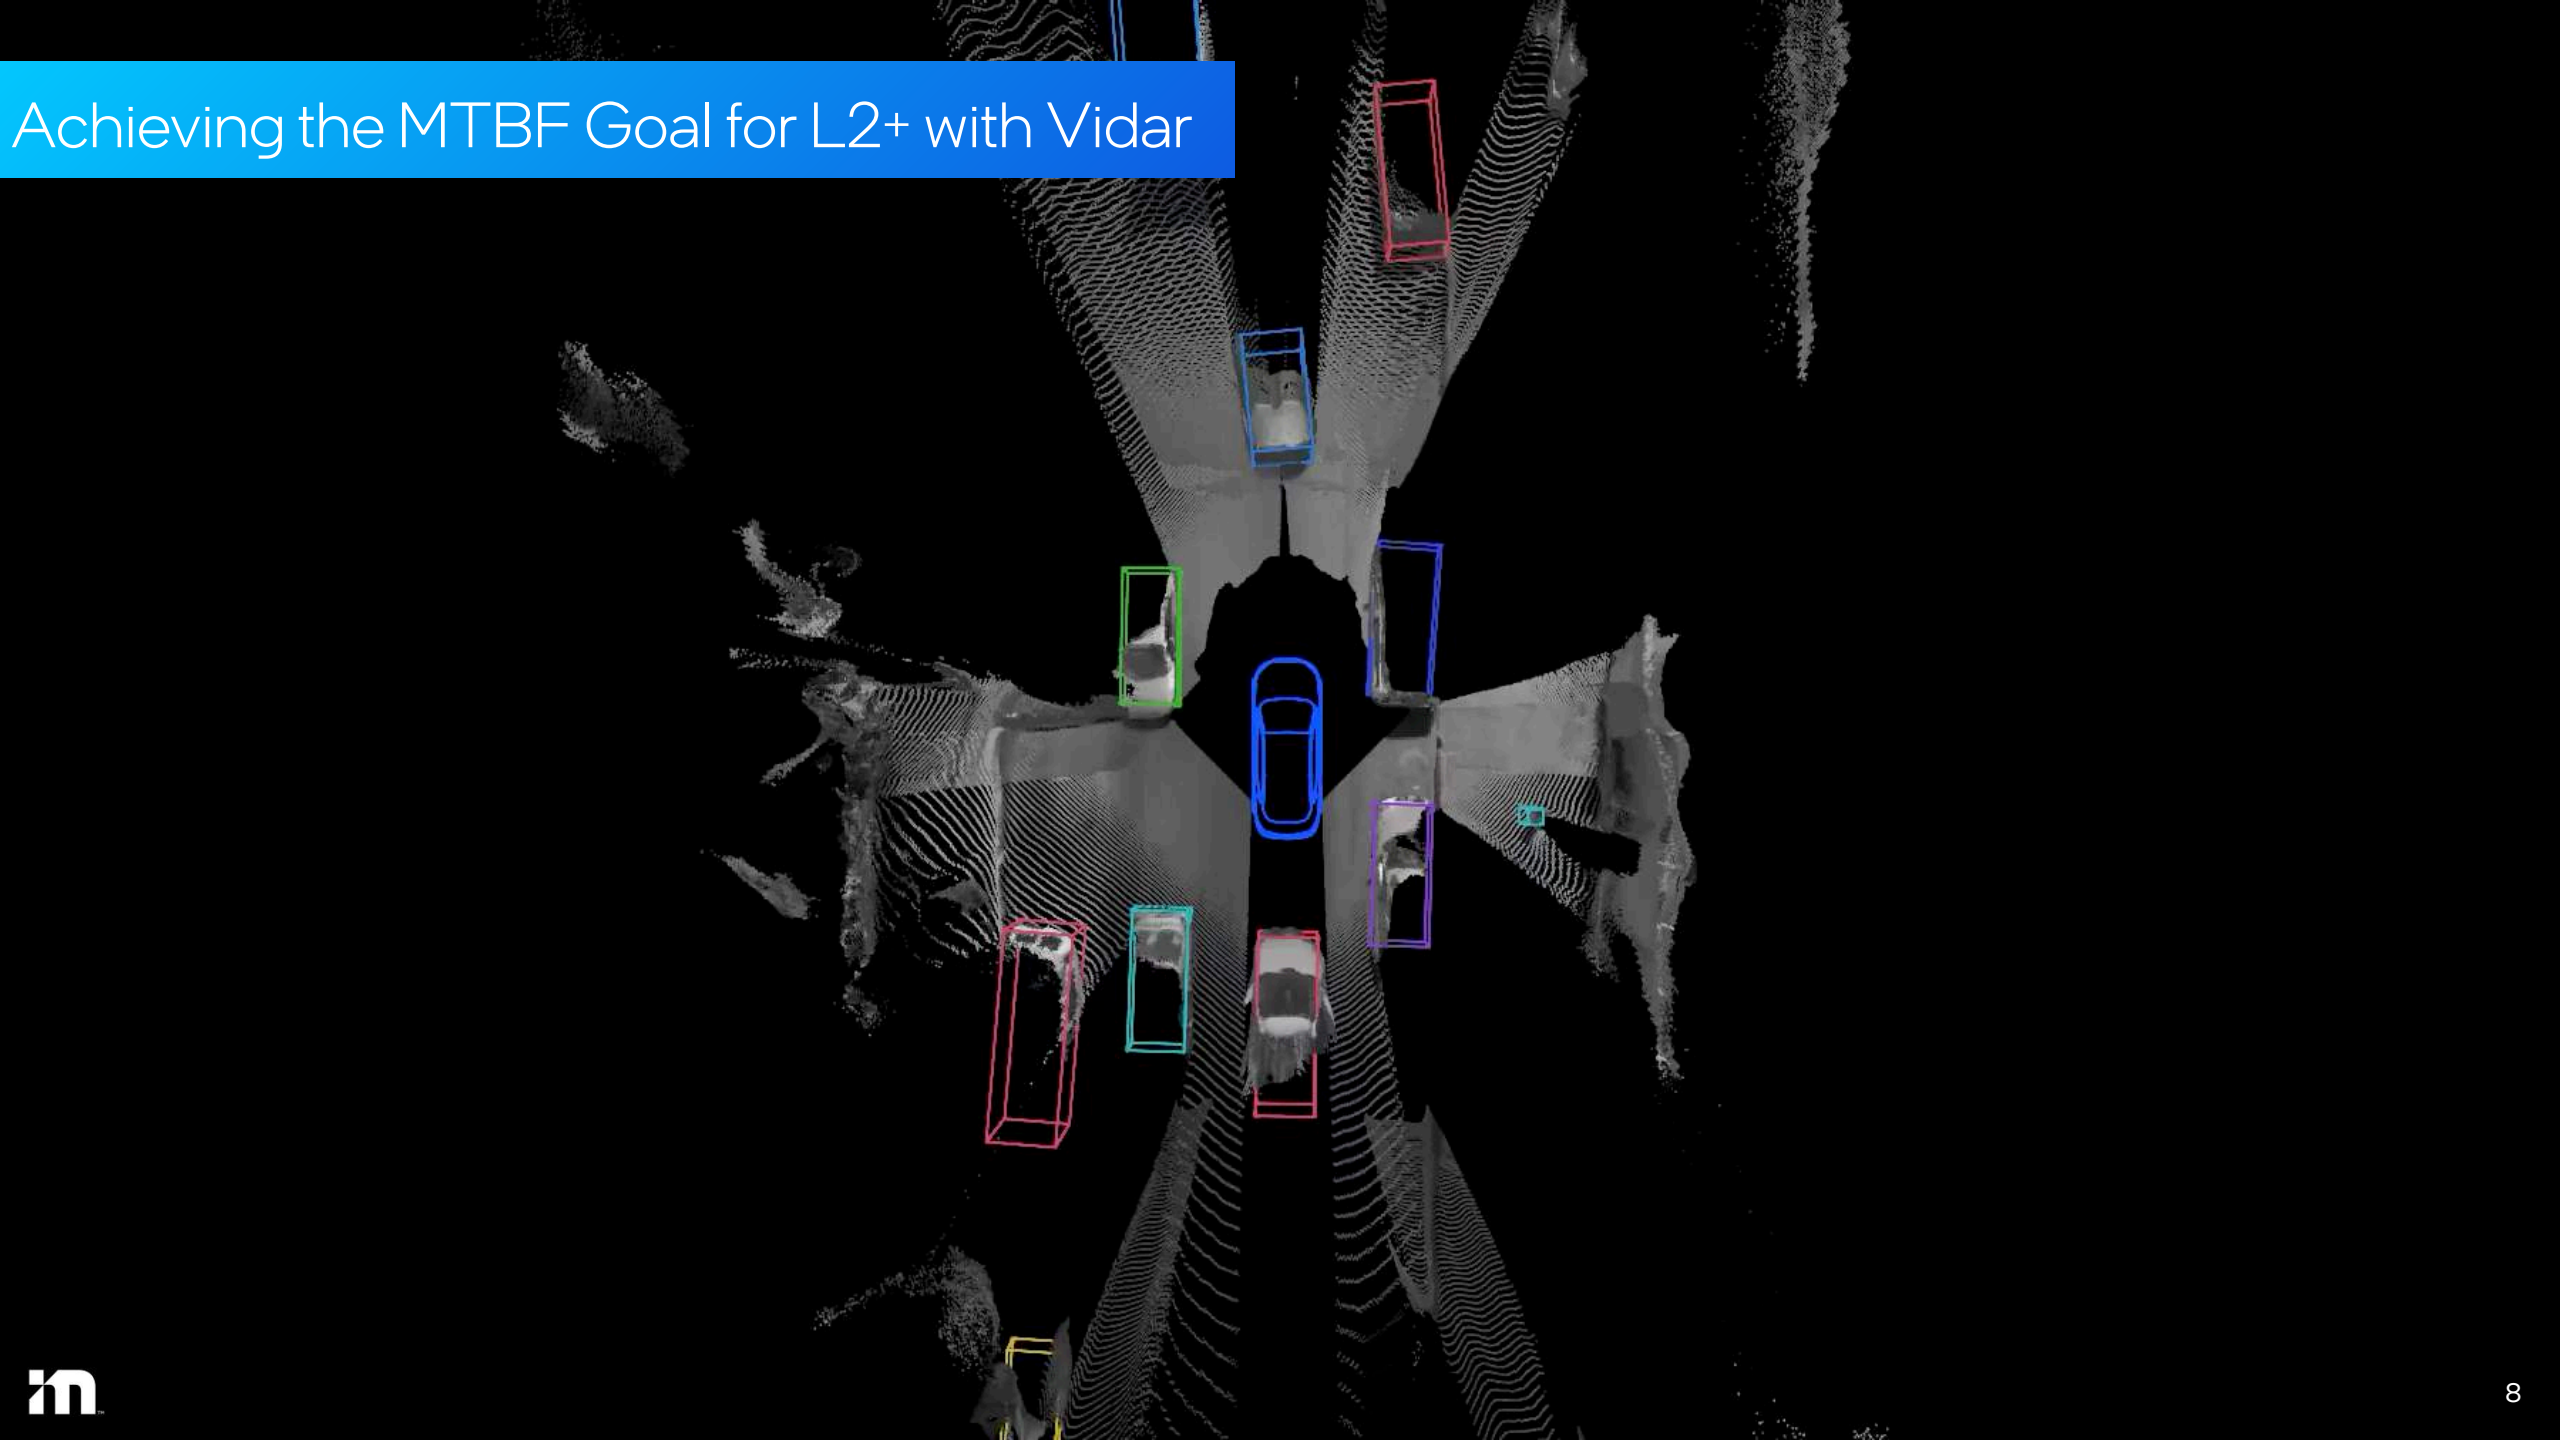

SuperVision™

Productizing AV Camera-only Subsystem to ADAS

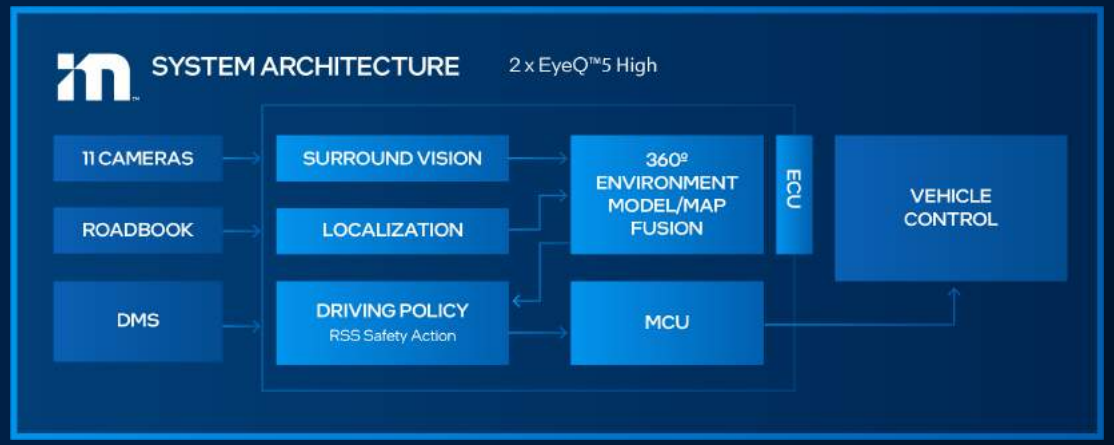

Value proposition:

+ Safety: "Vision Zero" next-gen ADAS
+ Comfort: market-leading L2+ experience

First customer:

+ Geely, Lynk & Co models
+ SoP: Q4, 2021

Content:

+ 2x EyeQ 5, PCB, perception, driving policy, REM HD-maps, control, parking

Capability:

+ Covering all road types: HW, rural, arterial, urban

Automotive World est. 1992

Home / News Releases / Automaker News Releases / Mobileye and Geely to offer most robust driver assistance features

Mobileye and Geely to offer most robust driver assistance features

Geely Auto Group, the
vehicle (EV) Zero Cond
conjunction with the Be

September 24, 2020

TC

Join Extra Crunch

Mobileye signs driver-assistance deal with Geely, one of China's largest privately held auto makers

Catherine Shu @cat

EET ASIA

Home News Tech Room Downloads Webinars Special Reports Videos

China's Geely to Feature Mobileye SuperVision for Scalable ADAS

Article By : Junko Yoshida

Expanding Footprint of **SuperVision™**

ISRAEL

Main development site in Jerusalem
Conducting tests across all of Israel

DETROIT

Ongoing testing and development

MUNICH

Already completed thousands of AV driving hours
Main hub for customer demonstrations

NEXT

TOKYO
SHANGHAI
PARIS
NYC*

*based on regulation

Achieving the MTBF Goal for L4

Building an End-to-End AV Capability Based on Radar+LiDAR as an Independent Subsystem

Why this is not a trivial problem

+ Classifying and tracking objects in the absence of texture
+ Segmenting and tracking objects in cluttered environments
+ LiDAR-based localization in a sparse semantic camera-generated map

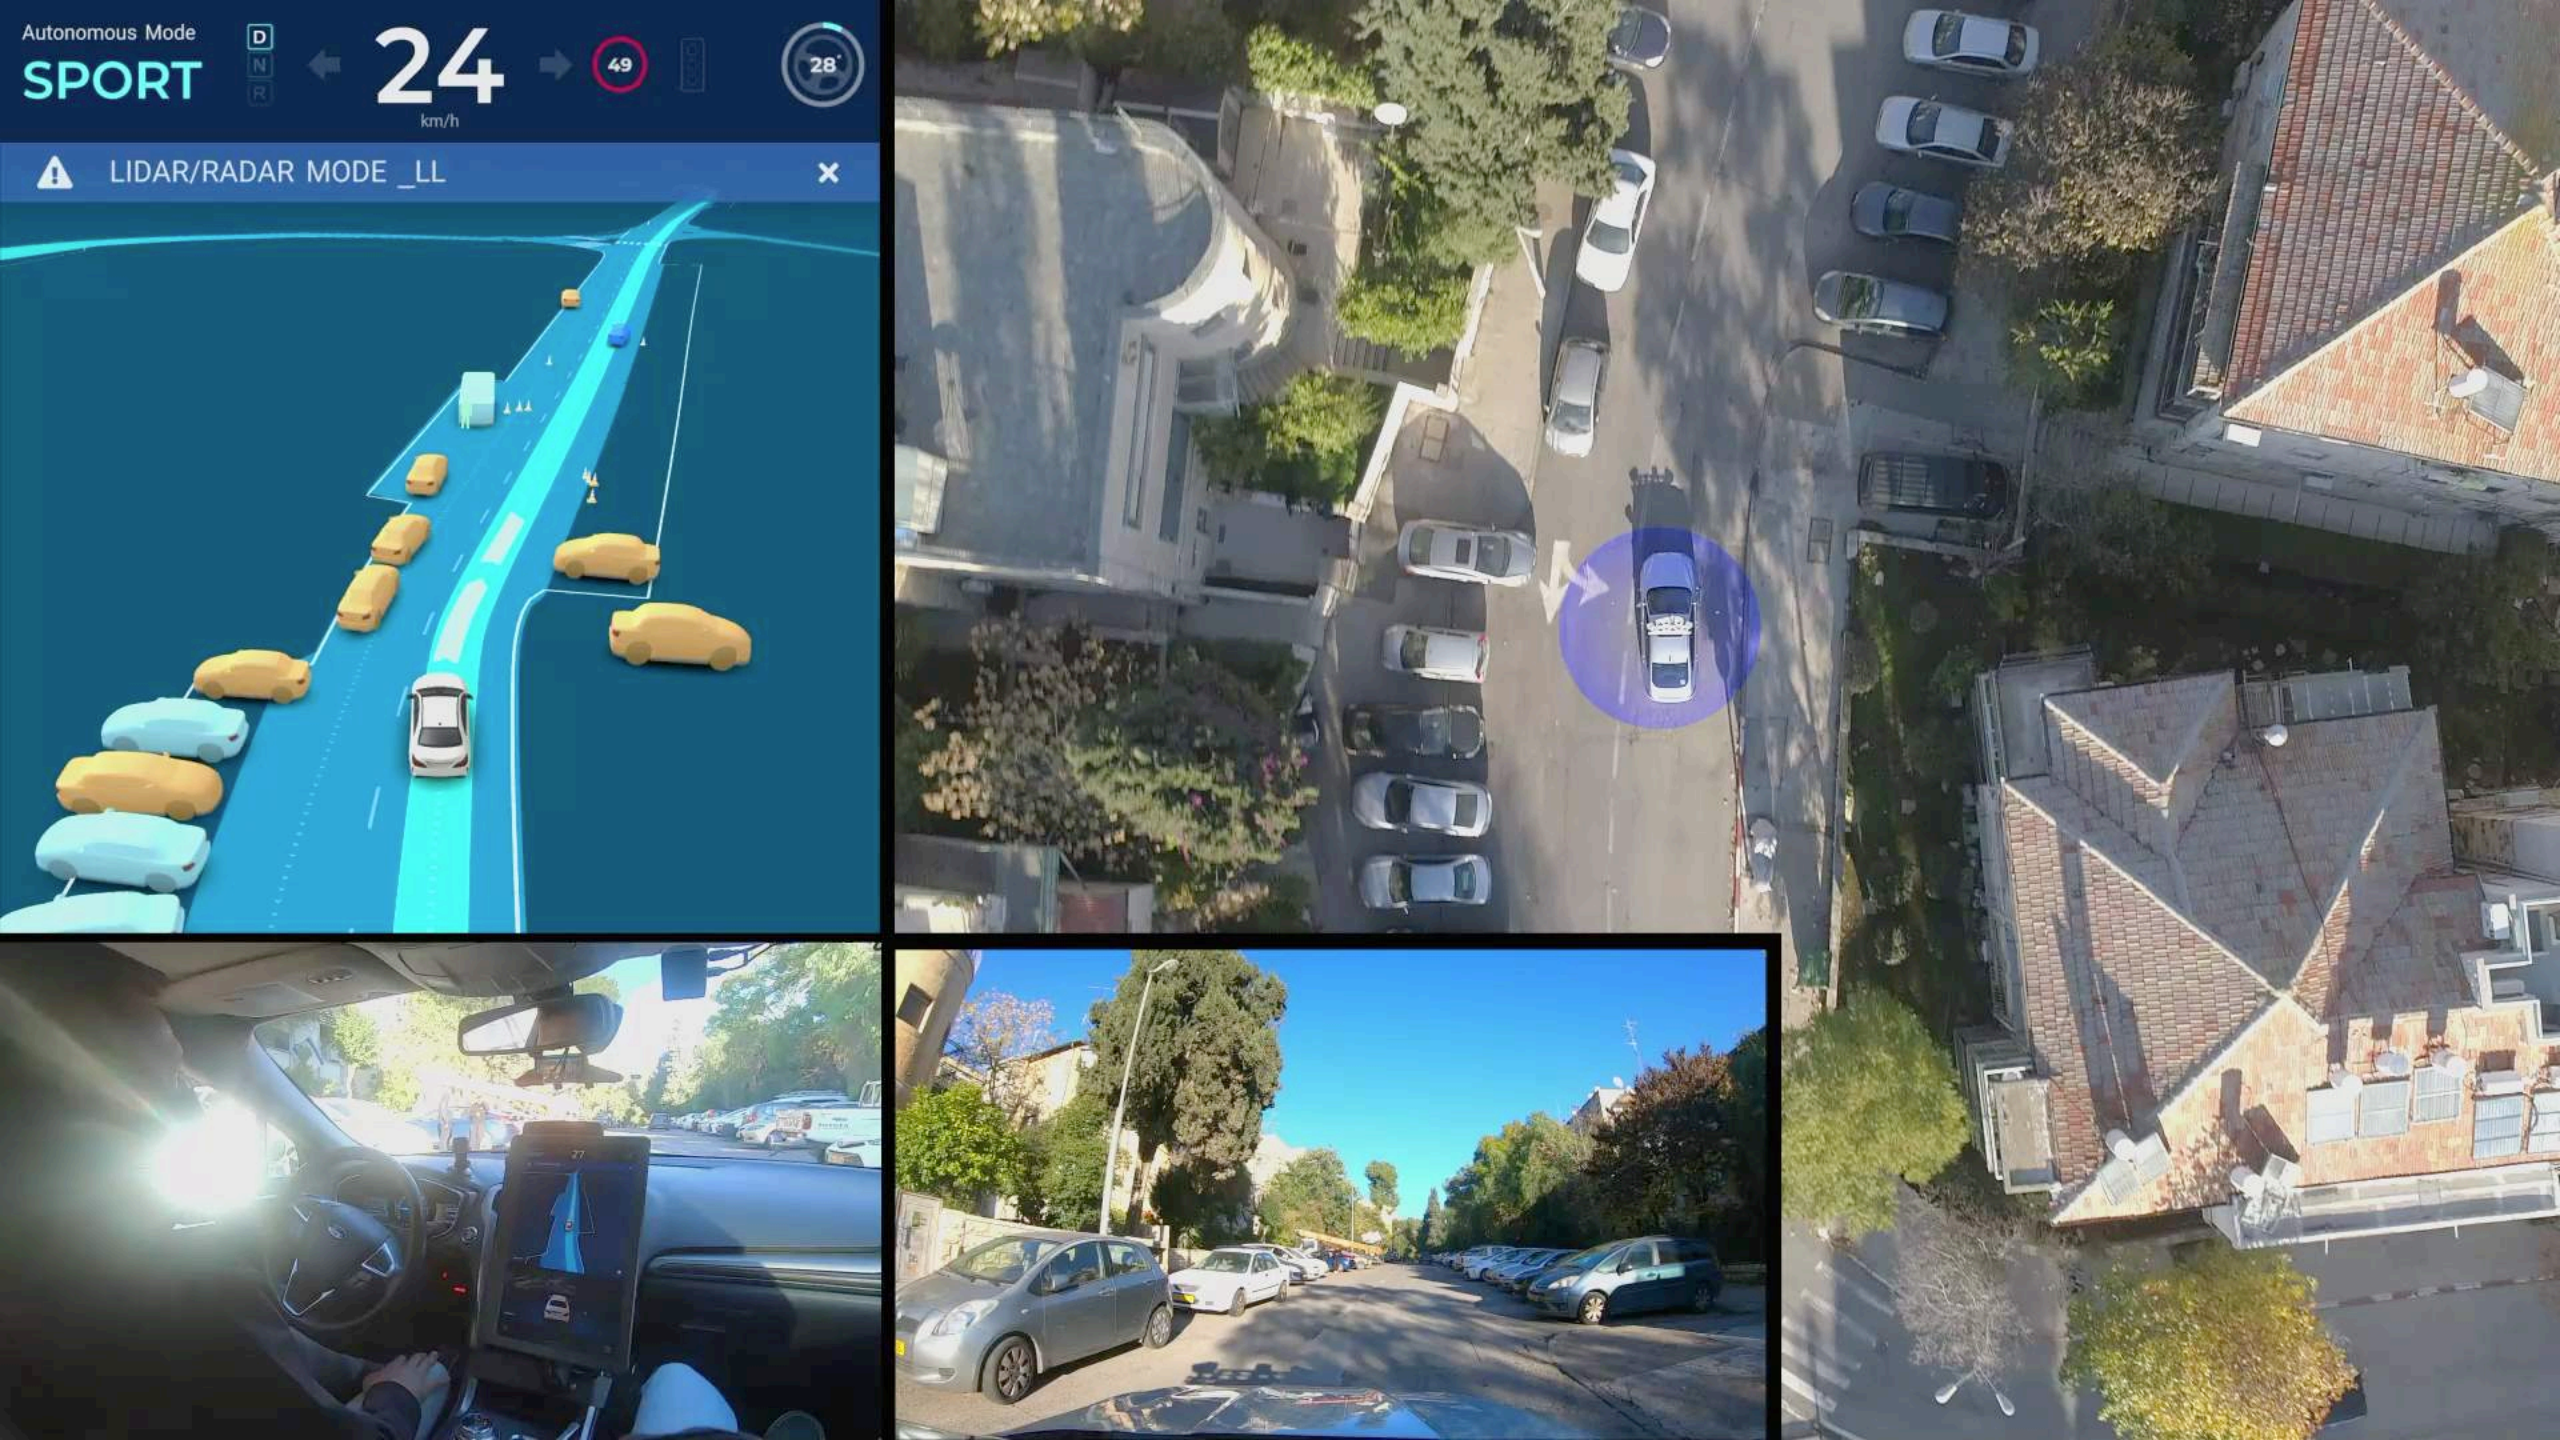
Autonomous Mode
SPORT
D
N
R
24
km/h
49
28
LIDAR/RADAR MODE _LL

REM™ Maps for ADAS and Autonomous Driving

The Motivation Behind High Resolution Maps

It is possible to detect and correctly interpret all roadway data, online:

- Drivable path
- Lane priory
- Path delimiters
- Lane-TFL association
- Crosswalks
- Stopping/ yield points

But,

Making this work at a very high MTBF is not realistic for current AI

Thus,

Prepare in advance all the above information

The Challenges

MAPS AT SCALE

Robotaxi can sustain lack of scalability (for now)

Consumer-AV needs to drive “everywhere”
Lack of scalability is a showstopper

Advanced ADAS needs to drive “everywhere”
Lack of HD maps yields low MTBF

UP-TO-DATENESS

Time-to-reflect-reality, ideally close to real time

ACCURACY

Host vehicle and other surrounding road users need to be positioned at a cm-level accuracy

The Common Approach for Building HD Maps

- Survey fleet with LiDARS, INS and Cameras
- Create a geometrical road model using the dense LiDAR & INS output
- Use manual / semi-automatic annotation to fine-tune semantic features
- Fuse external data layers to enrich map with missing information

Mobileye mapping vehicle (used for ground truth)

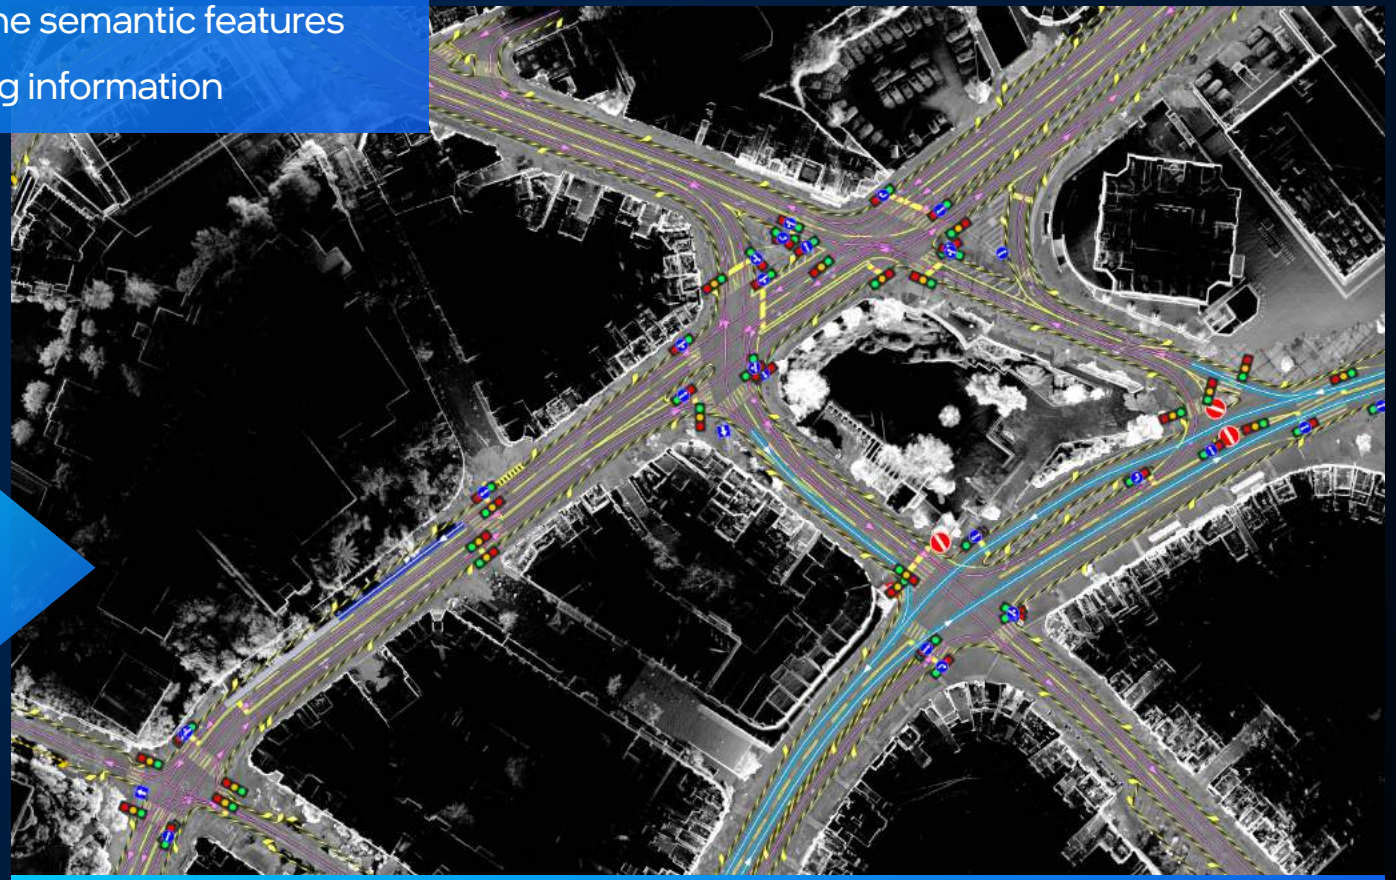

Mobileye "pure" HD mapping (used for ground truth)

What's Wrong with the Typical HD Map Approach?

The geometric layer is over specified

Cm-level accuracy in global coordinates is **NOT** needed to support AV

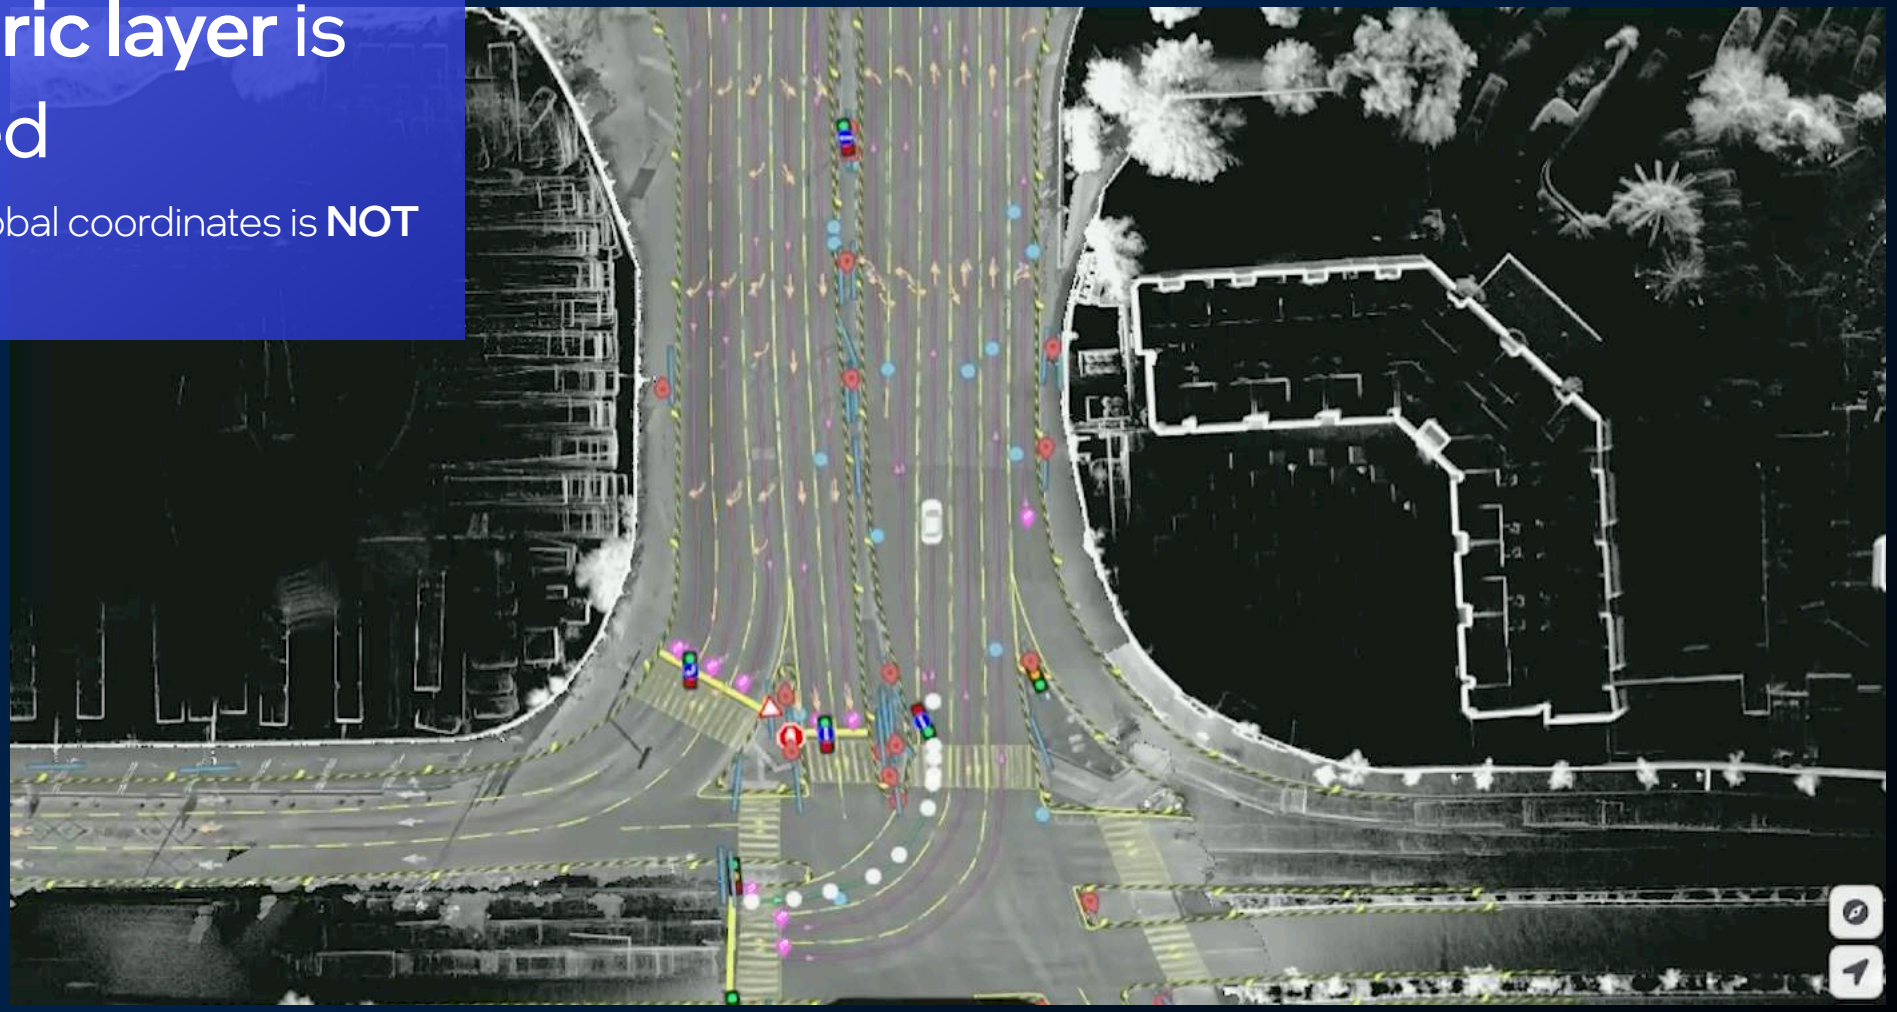

What's Wrong with the Typical HD Map Approach?

The richness of **the semantic layer** is very difficult to automate

The semantic layer gives an actionable meaning for each element in the map:

DRIVABLE PATHS

PRIORITY

TFL/CROSSWALKS TO LANE ASSOCIATION

STOPPING/ YIELD POINTS

COMMON SPEED

What's Wrong with the Typical HD Map Approach?

DRIVABLE PATHS

PRIORITY

TFL/CROSSWALKS TO LANE ASSOCIATION

STOPPING/ YIELD POINTS

COMMON SPEED

How can the drivable path be determined when there are no lane marks

3-way junction with no lane marks

What's Wrong with the Typical HD Map Approach?

DRIVABLE PATHS

PRIORITY

TFL/CROSSWALKS TO LANE ASSOCIATION

STOPPING/ YIELD POINTS

COMMON SPEED

Disambiguating traffic lights priority rules

What's Wrong with the Typical HD Map Approach?

DRIVABLE PATHS

PRIORITY

TFL/CROSSWALKS TO LANE ASSOCIATION

STOPPING/ YIELD POINTS

COMMON SPEED

Traffic lights structure and logic vary dramatically from country to country

What's Wrong with the Typical HD Map Approach?

DRIVABLE PATHS

PRIORITY

TFL/CROSSWALKS TO LANE ASSOCIATION

STOPPING/ YIELD POINTS

COMMON SPEED

Determining the optimal stopping point in terms of viewing range:

- Occlusions may happen from objects that are not part of the map (buildings, fences, poles, etc.)
- To calculate the optimal stopping point, a map with all stationary objects that may affect the viewing range is needed- **clearly not feasible**

What's Wrong with the Typical HD Map Approach?

DRIVABLE PATHS

PRIORITY

TFL/CROSSWALKS TO LANE ASSOCIATION

STOPPING/ YIELD POINTS

COMMON SPEED

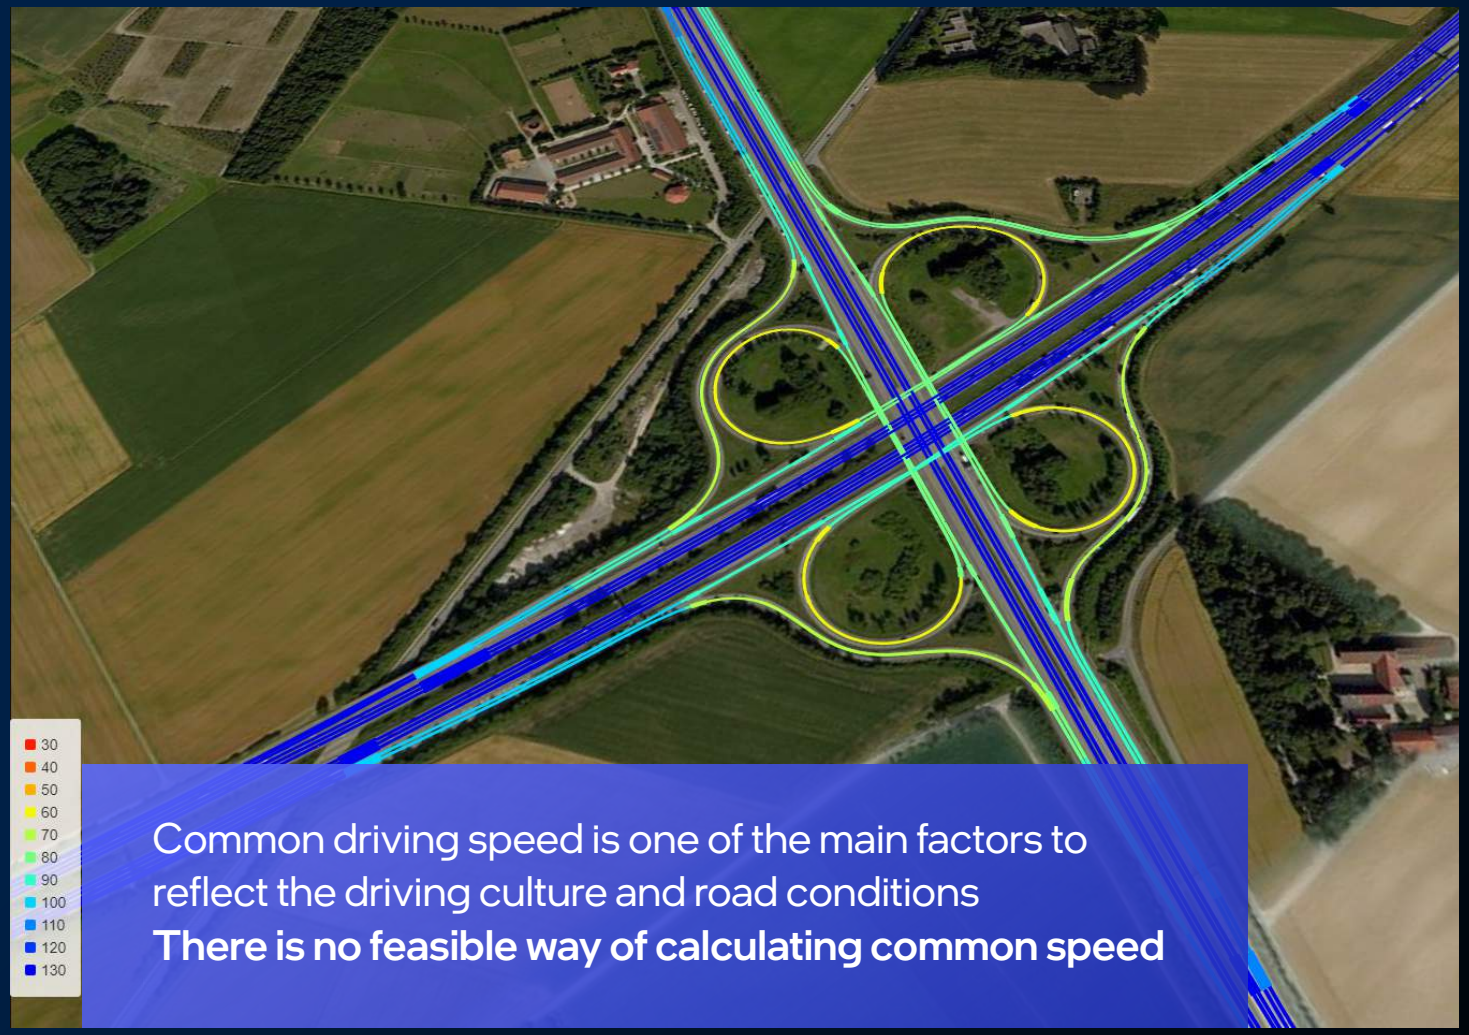

Common driving speed is one of the main factors to reflect the driving culture and road conditions

There is no feasible way of calculating common speed

Mobileye's Approach:

"AV Map", not "HD Map"

Designing a map that suits exactly what AVs need:

Scalability	Accuracy	Detailed Semantic features
Unlocking millions of "mapping agents" in every relevant region	Use novel state-of-the-art algorithms to achieve high accuracy levels where it matters	Use explicit attributes and "wisdom of the crowd" to generalize traffic rules and driving culture

REM process:

1. RSD Harvesting
2. Automatic Map Creation
3. Localization & AD

REM™ Under the Hood

Harvesting

Single drive reports everything it detected

Alignment

Creating a coherent statistical model in a global coordinates system by aligning random crowd drives accurately

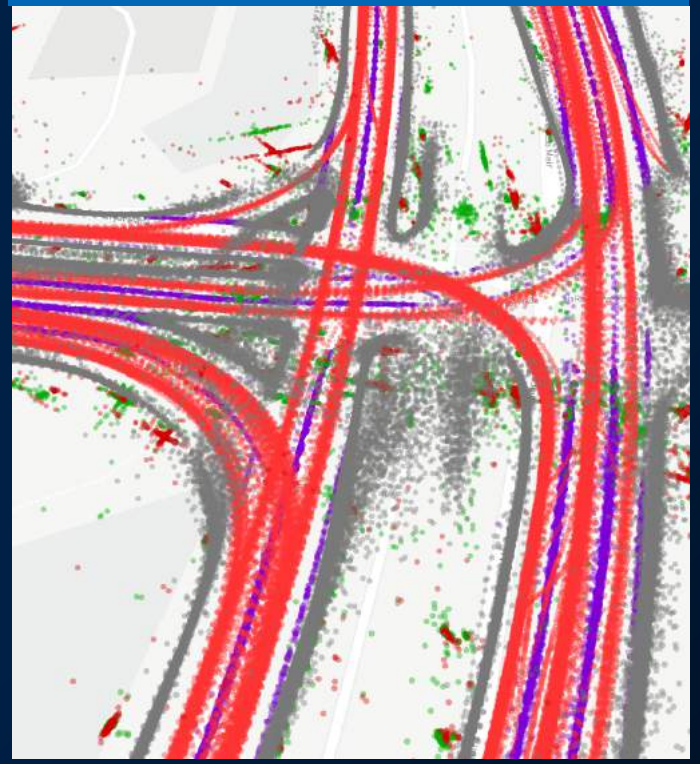

Modeling & Semantics

Creating a definitive, actionable modeling of all road elements with accurate semantic description

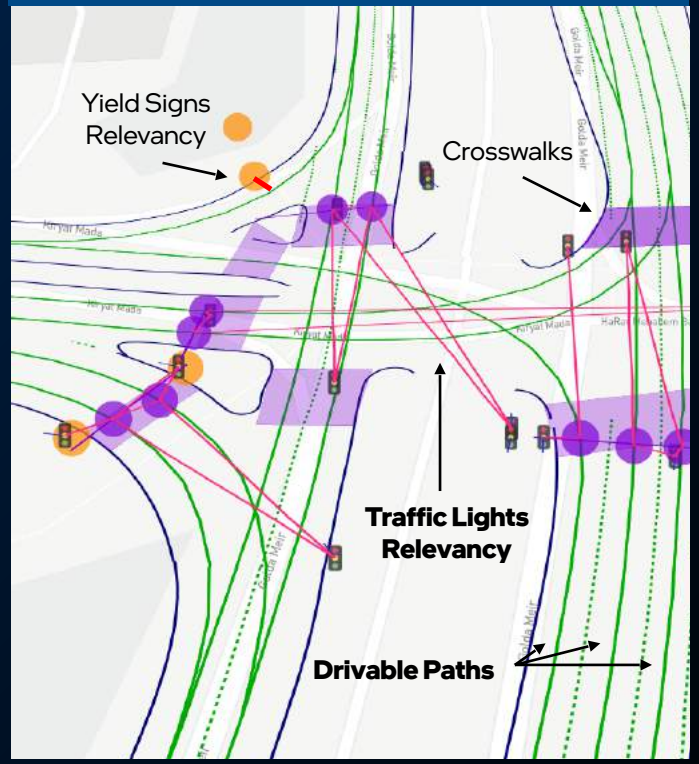

RSD Harvesting

Each Harvesting agent collects information on all the relevant road elements

Road Geometry

- Visual landmarks for localization
- Traffic signs & traffic lights
- Lane marks & road edges
- Crosswalks

Semantics

- Drivable paths
- TSR type
- Stopping points
- Speed

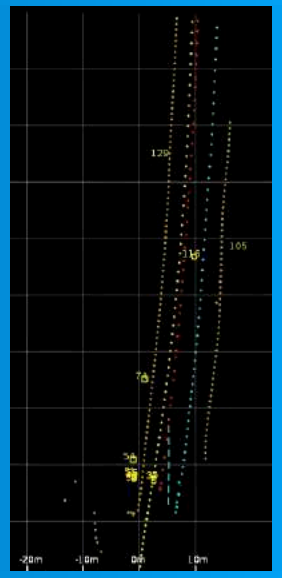

Information is sent as a compact “RSD” (Road Segment Data)
packet, with sparse, economic representation of these elements (10KB/ Km)
Full compliance with GDPR regulation

Aligning Drives

Aligning random crowdsourced drives to create a 3D model at a cm-level accuracy in the surrounding of the host vehicle. The task is to find the 6D pose (in the map coordinates) for each detected element in each RSD

Translate hundreds of random drives that traversed the same location in different driving patters into a coherent aligned model.

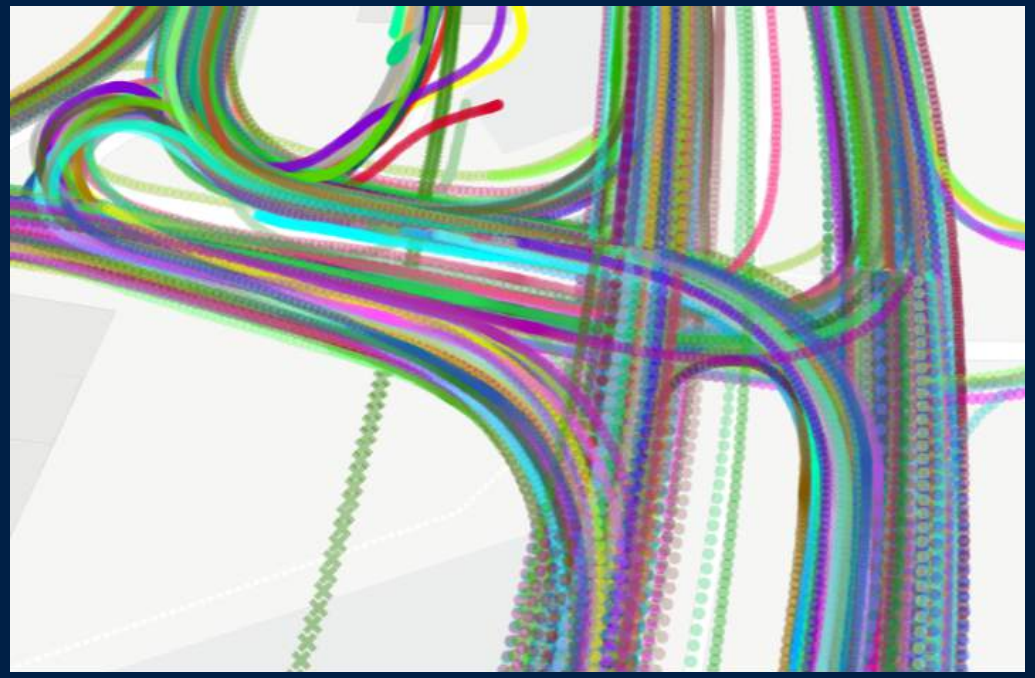

Aggregated GPS trajectories of ~300 drives that traversed a junction from multiple directions.

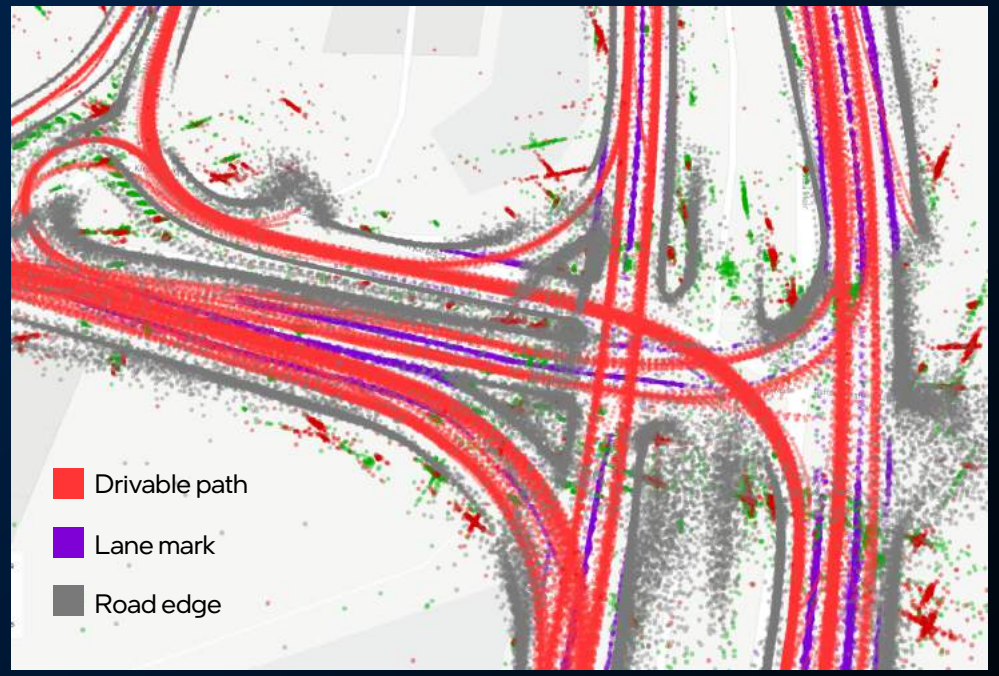

Aligned statistical model based on Mobileye's proprietary Alignment algorithm

Why this is hard?

Every class has its own noise model, thus impossible to implement naïve averaging methods

ALIGNMENT BASED ON GPS

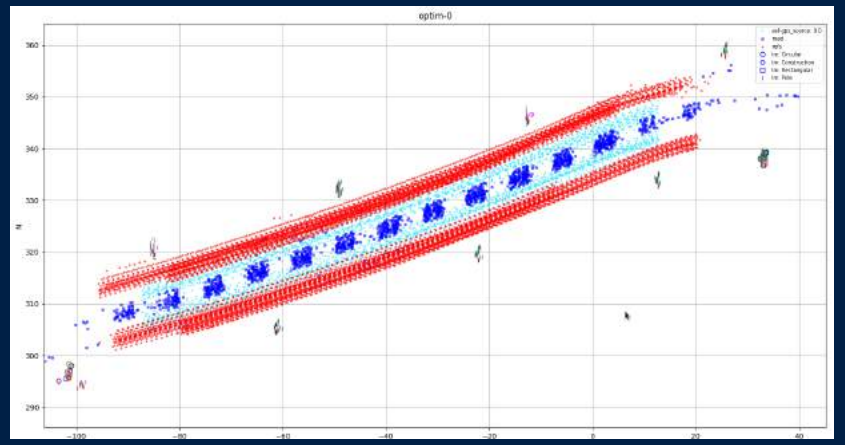

ZOOM IN

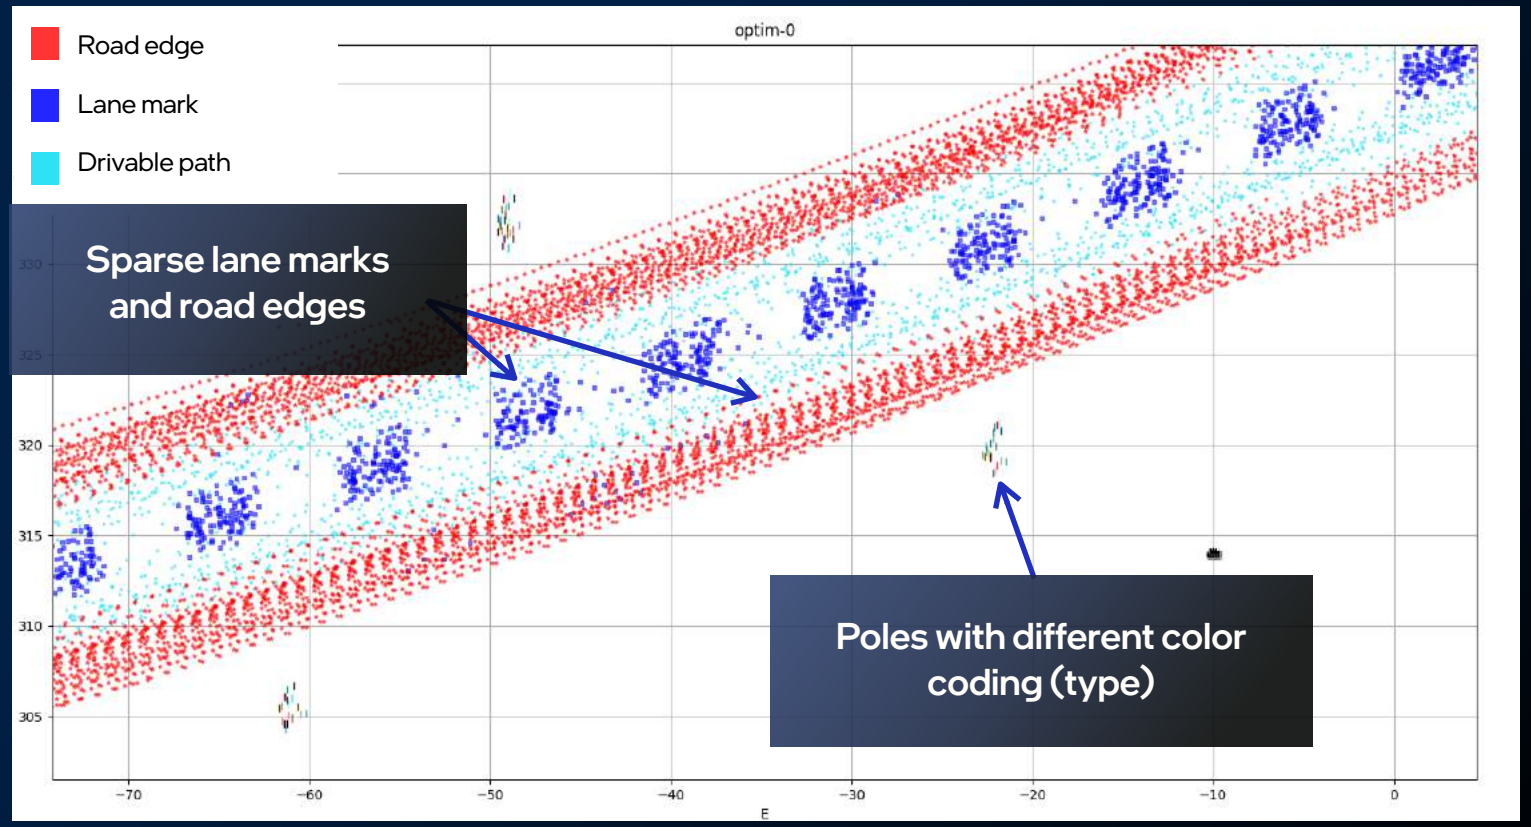

Creating an Extremely Accurate Road Geometry without any Priors

THE RESULT:

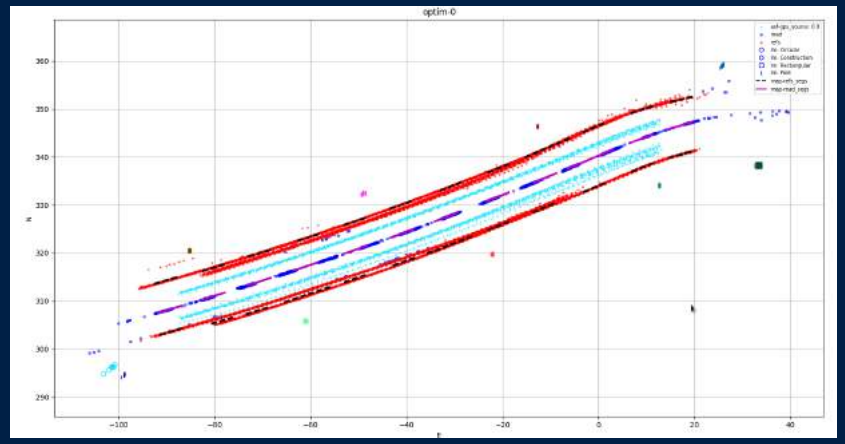

ZOOM IN

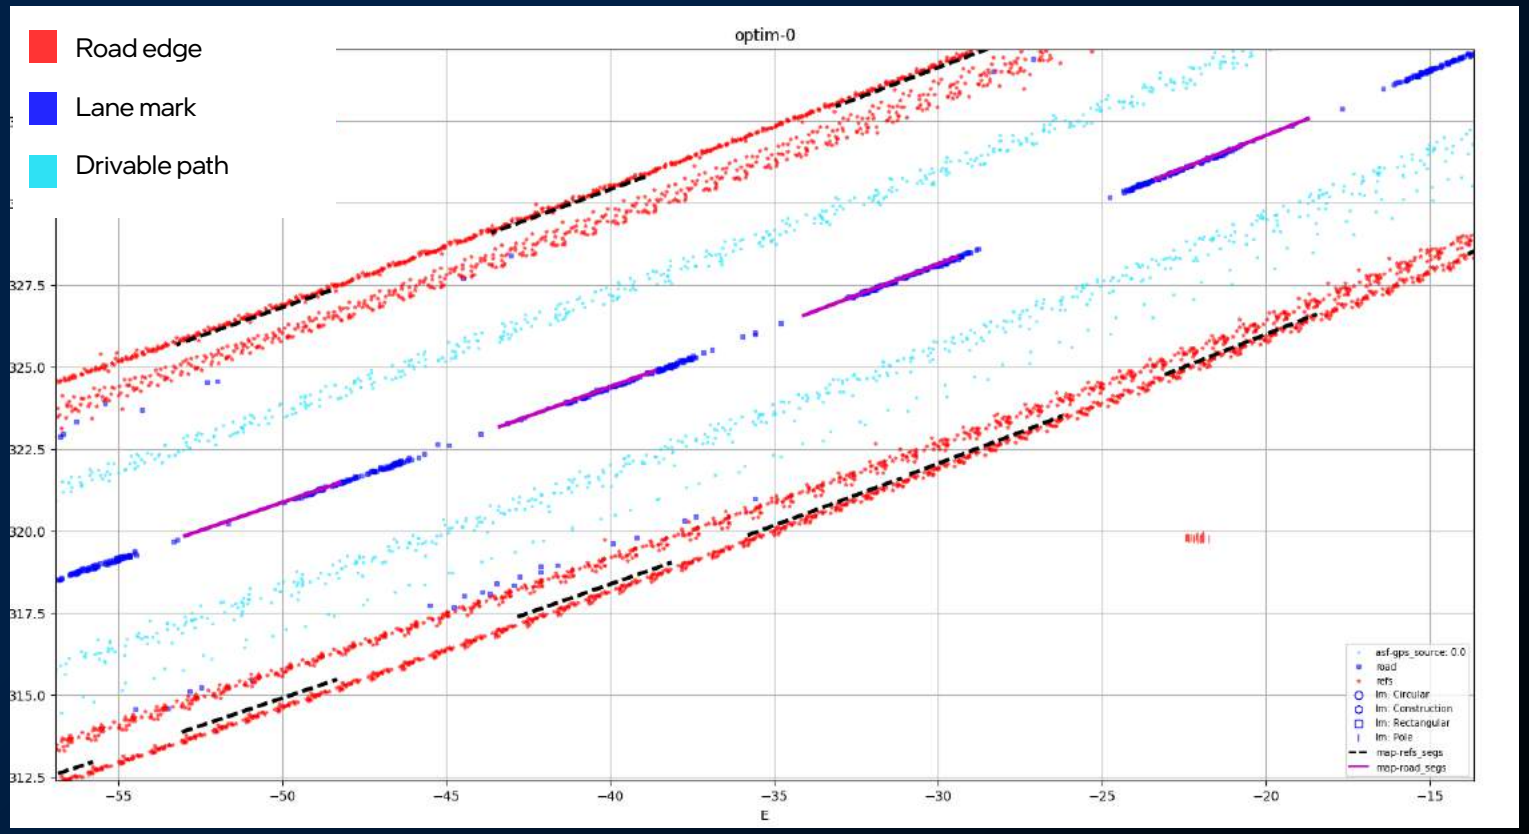

Modeling Process

DNN for modelling allows to compensate for noisy data and to increase robustness

Using “handmade” HD maps as a training set

Clustering and Spline fitting output

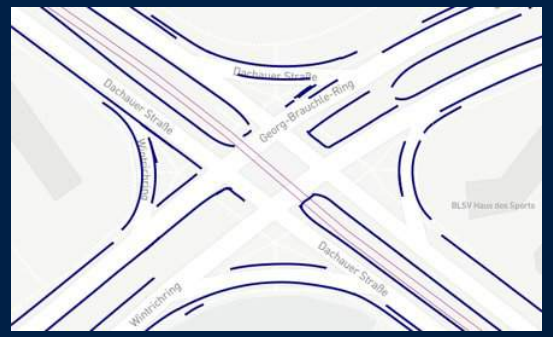

DNN output

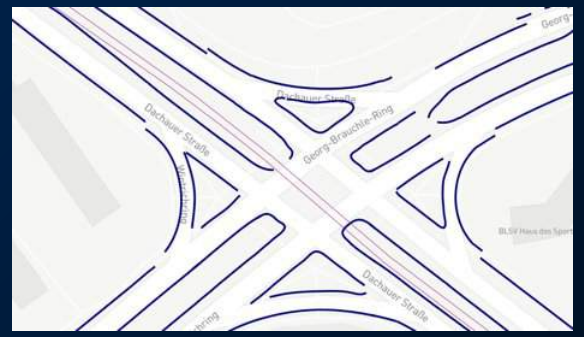

INPUT

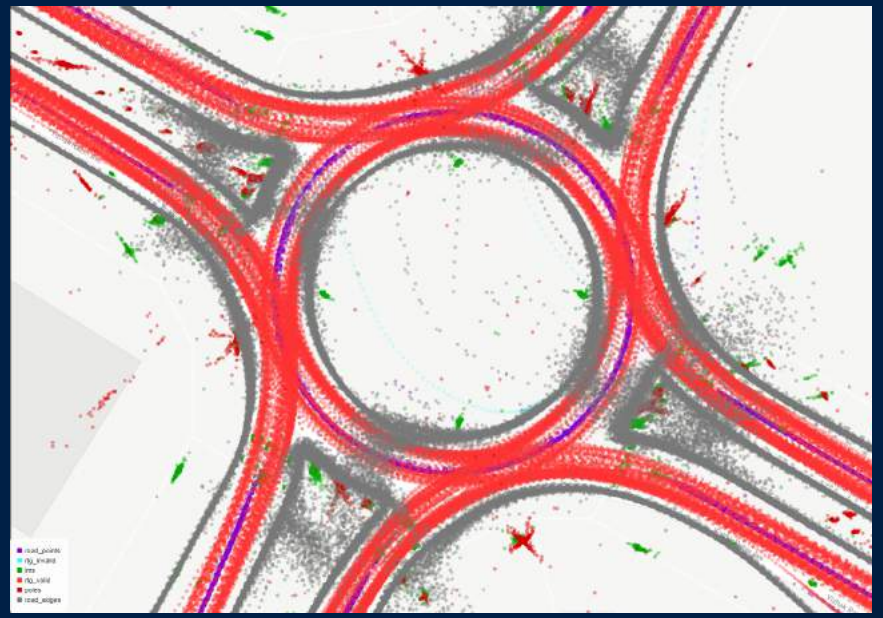

OUTPUT

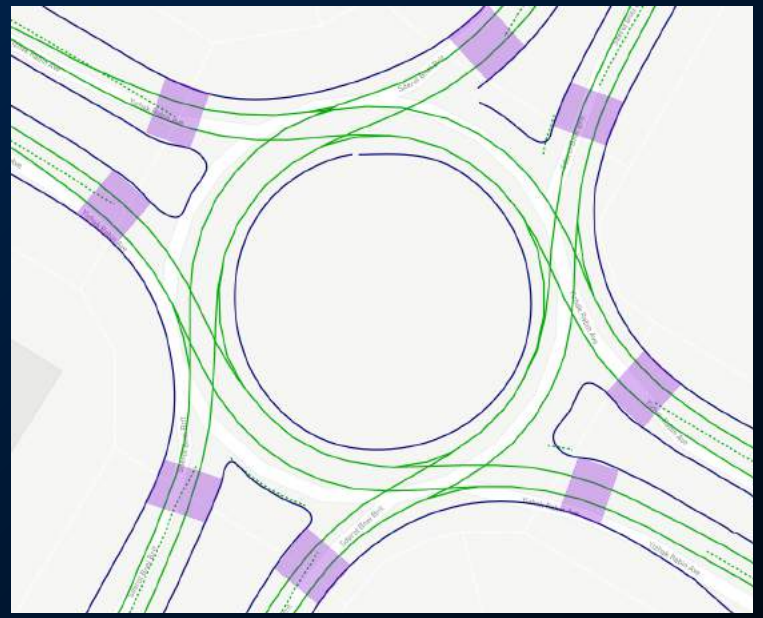

Why Crowd Sourcing is Perfect for Semantic Understanding

DRIVABLE PATH IN AREAS WITH NO LANE MARKS

- Determining where's the Drivable Path without relying on Lane Marks – imperative in urban and rural scenes

By evaluating drivable path samples from many drives (red dots), we can easily determine the actual definitive drivable path (green line)
Right-of-way can also be determined by evaluating which DP has more stopping points

DETERMINING ALL POSSIBLE DRIVABLE PATHS IN ADVANCE

- Crowd automatically determines where are the drivable paths, even in complex scenarios.
- Relying on crowd requires no external information / prior assumptions.

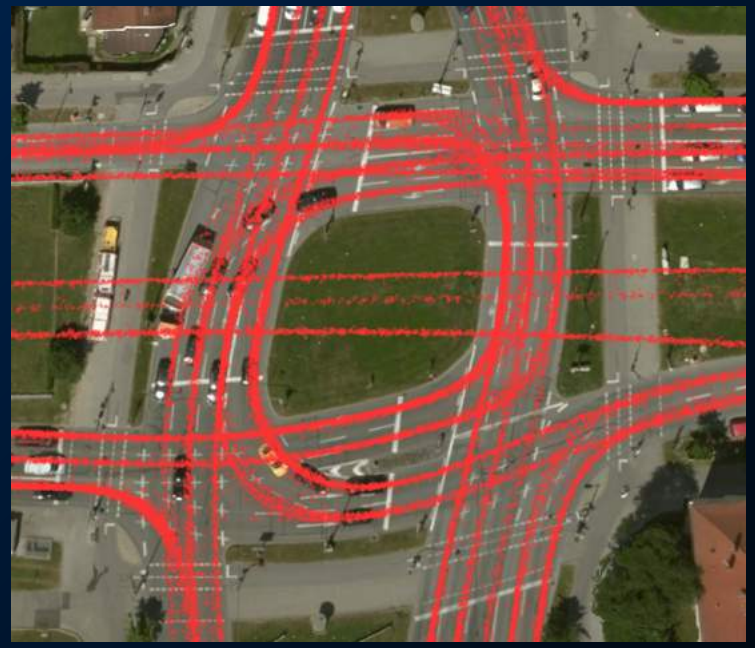

AV will know where are all the drivable paths before it enters this complex junction, which is essential for safe route planning and driving.

Why Crowd Sourcing is Perfect for Semantic Understanding

By using crowd observations we can associate instructive road elements with drivable paths in a generalized manner

TRAFFIC LIGHT ASSOCIATION

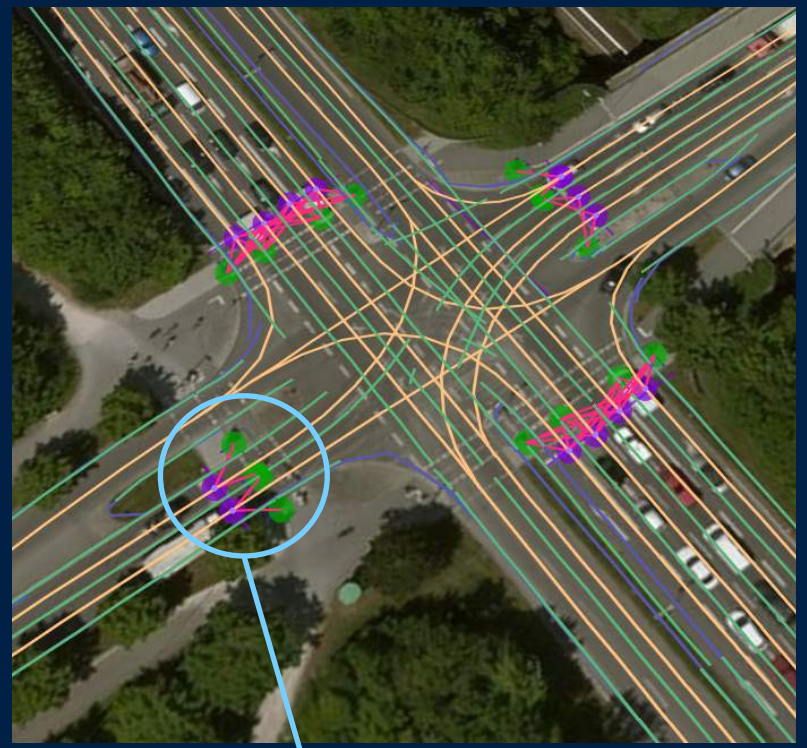

Different association to the two drivable paths.
(TFLs are not synchronized)

YIELD & STOP SIGN ASSOCIATION

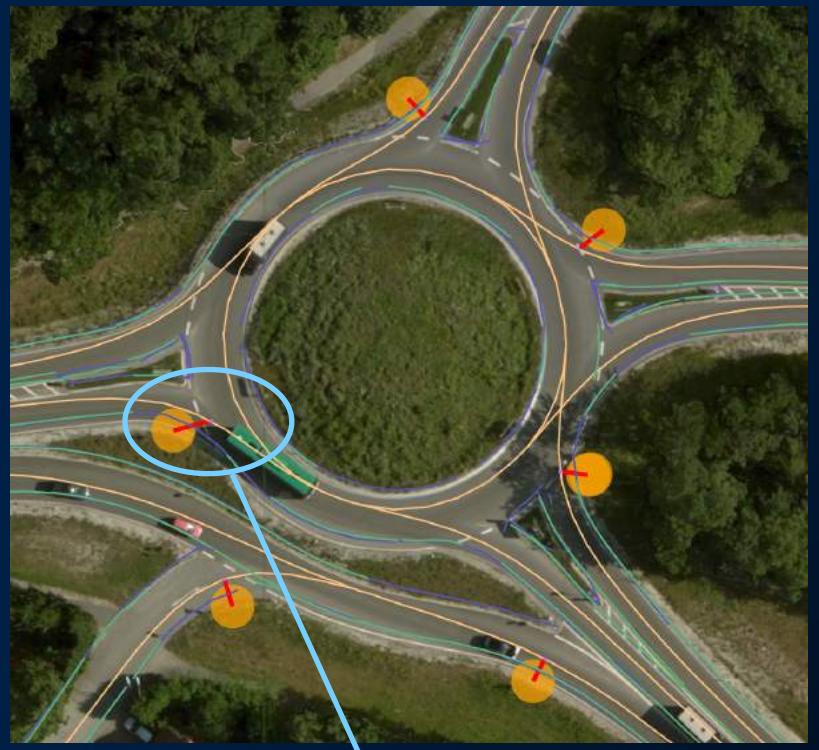

Each drivable path entering the roundabout has a yield sign associated to it, as the traffic rules instruct.

CROSSWALKS ASSOCIATION

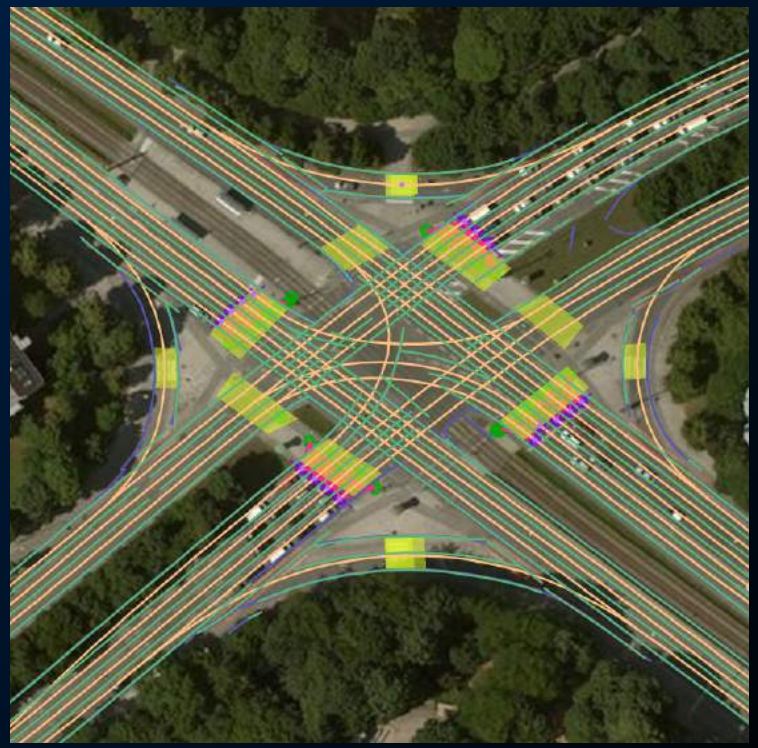

Why Crowd Sourcing is Perfect for Semantic Understanding

RIGHT OF WAY WHERE THERE ARE NO TRAFFIC SIGNS

- Relying strictly on traffic signs will result in safety-critical errors where these are not available.
- Relying on crowd stop & slow down points easily shows which drivable path has right of way in every occasion.

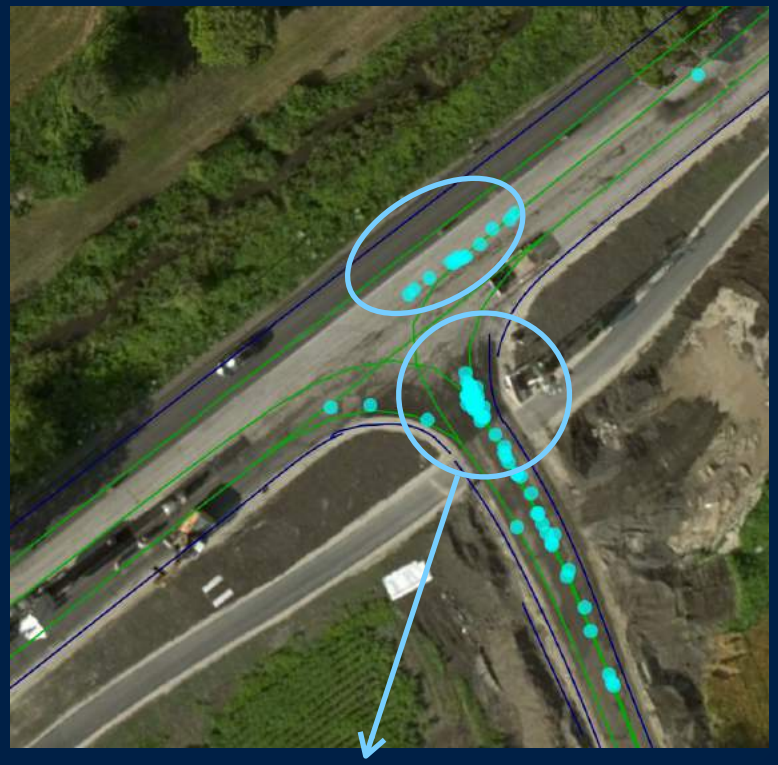

By evaluating in which drivable path there are more crowd stops , we can easily determine which direction has right of way.

WHERE TO STOP

- Crowd behavior provides a robust understanding of where the AV should yield for traffic,
- Crowd behavior generalizes all cases , without relying on complex models.

The map will indicate that the Yield Point will be where the majority of drivers yielded.

DETERMINING YIELD ON GREEN TRAFFIC LIGHTS

- Using crowd behavior to achieve a robust right-of-way determination in yield on green left turns, without relying on country specific traffic signs.

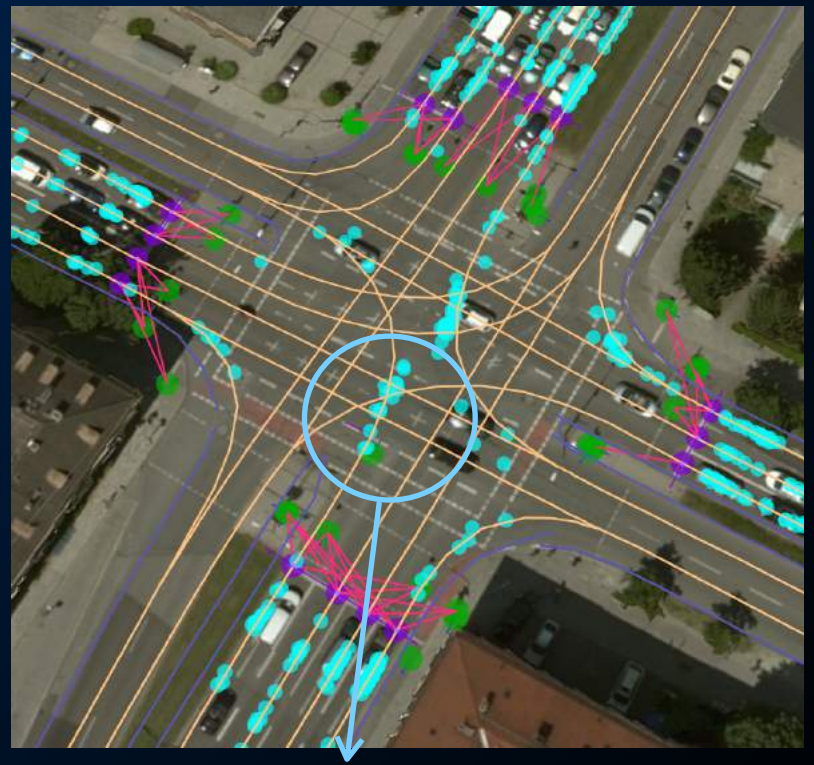

Multiple stop points inside a junction, indication for Yield On Green.

Why Crowd Sourcing is Perfect for Semantic Understanding

COMMON DRIVING SPEED

- Provides the target speed in which the AV should drive in when the road is clear.
- Understanding how the “average driver” drives helps the AV to blend in naturally, without interrupting traffic.
- Crowdsource is the only effective way to obtain this critical data

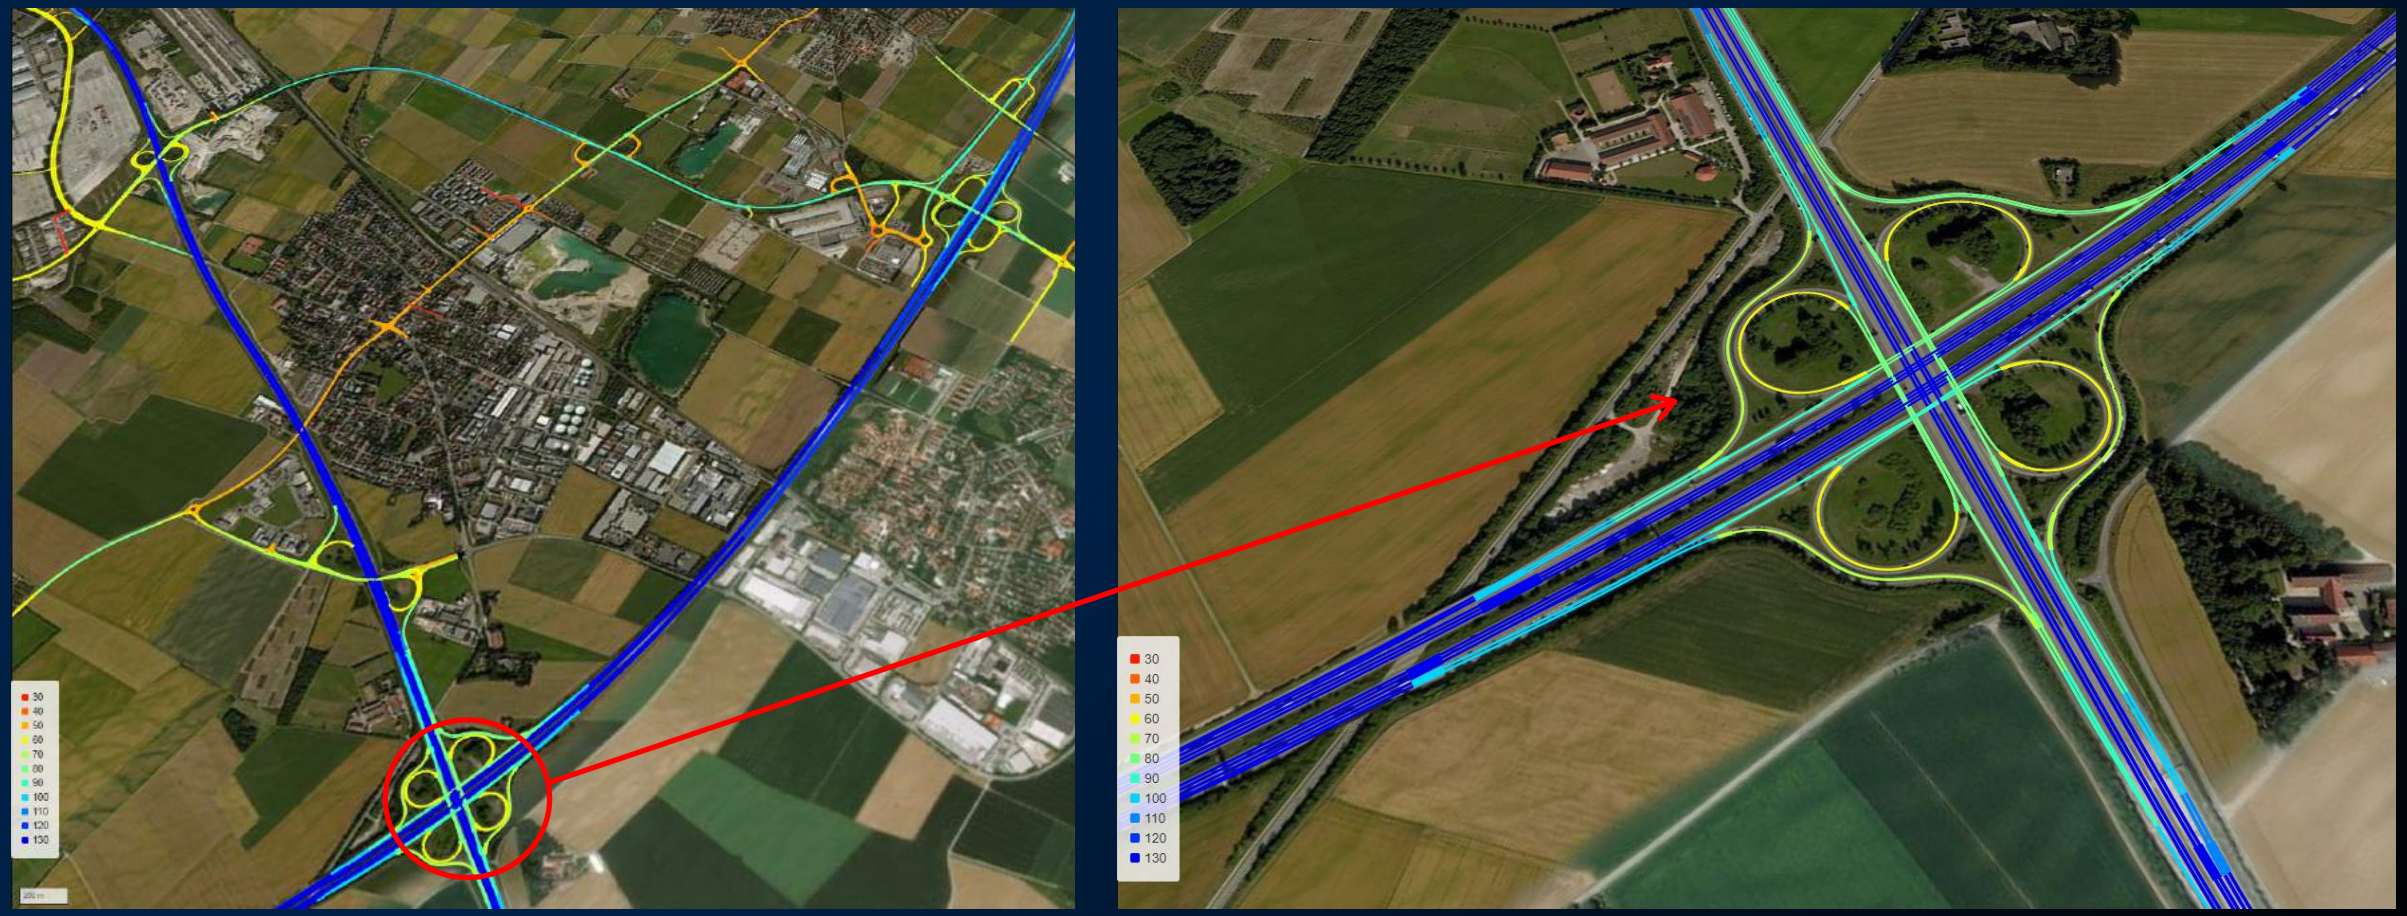

We are Building the World's Largest Crowdsourcing Fleet

Today

7.5B KMS of roads harvested globally

8M KMS of roads are covered daily

2024

1B KMS of roads covered Daily

Via 6 major OEM partners and retrofit solution for fleets

Active Sensor Development

The Motivation Behind LiDAR and Radar Development

2022 LiDAR/radar subsystem

- Best-in-class ToF LiDAR (Luminar) - 360° coverage
- Advanced stock radars- 360° coverage

Need both to build a sensing state

Current radar-lidar subsystem configuration

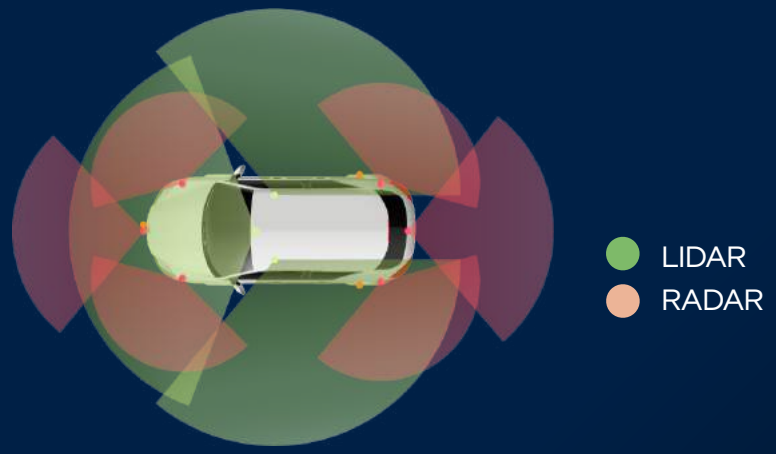

2025 LiDAR/radar subsystem

- Front sector- 3-way redundancy ⇒ ODD case to L5
- Remaining FoV- 2-way redundancy of cameras + radars ⇒ massive cost reduction

The enablers:

"Drive by" LiDAR capabilities

- Solving range limitations, interferences ,and velocity

"Drive by" radar capabilities

- Solving angular res., dynamic range, and side lobes effect

Next gen radar-lidar subsystem: 360° radar cocoon + 1 front facing LiDAR

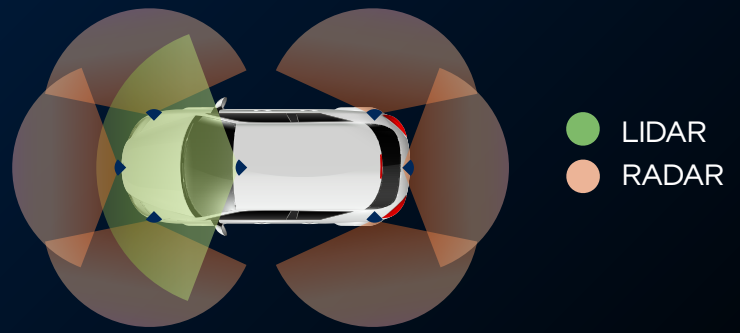

SW-defined Imaging Radar Development

The Goal

1. Increasing angular and vertical resolution
2. Assuring extremely high probability of detection (Pd) by:
   - Reducing mutual interaction between echoes from different objects
   - High dynamic range
   - Very low side lobes levels (SLL)

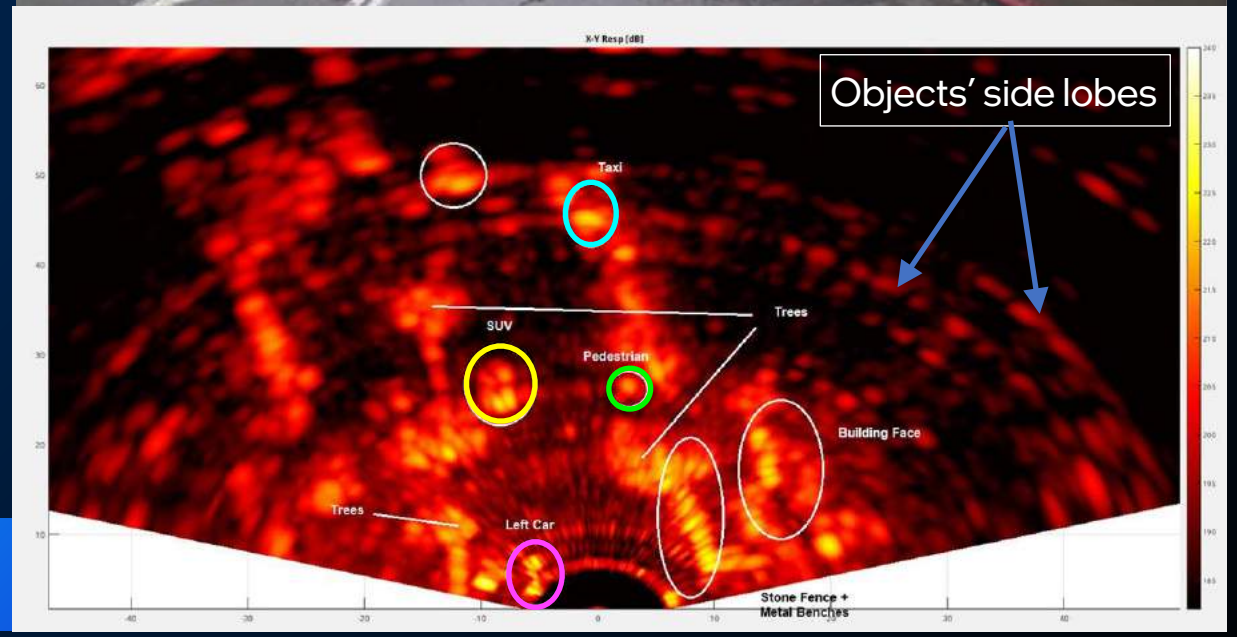

Inherent radar limitations must be overcome to usefully process data from raw output

The Required Capabilities

Moving from (industry standard):	**To:**	**To achieve:**
Simple MIMO radar with 192 virtual channels (12x16)	Massive MIMO radar with 2,304 virtual channels (48x48)	High angular and vertical res. (0.5°x2°) enabling accurate lane assignment and vertical discrimination
25dBc azimuth side lobe levels (SLL)	40dBc SLL	Detection of weak targets in the presence of strong targets at similar range and velocity
60dB dynamic range	100dB dynamic range	Detection of weak far targets in the presence of close strong targets

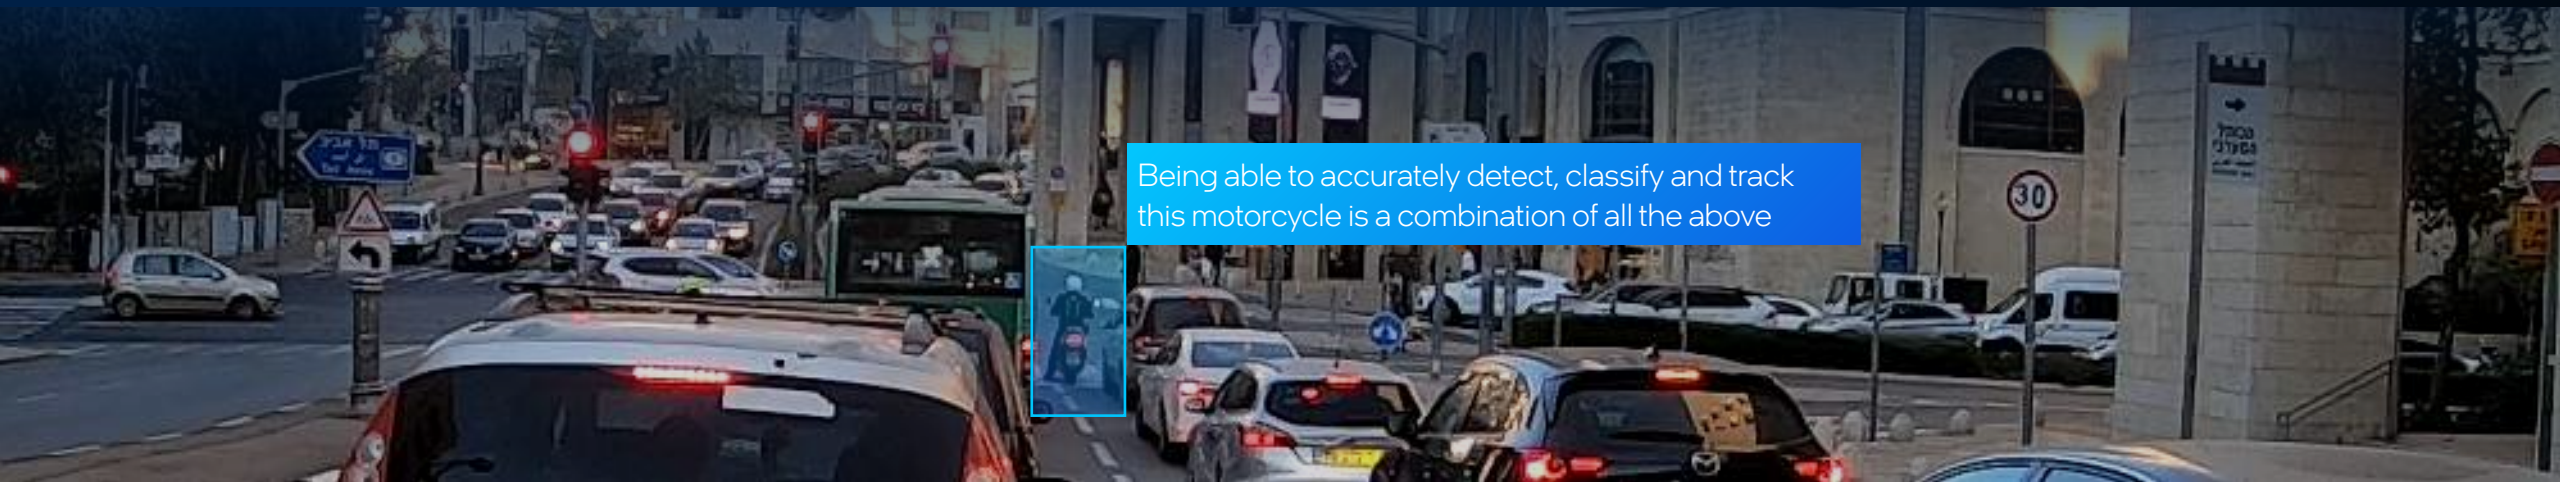

The Solution: SW-defined Imaging Radar

Paradigm shift in architecture to enable the leap in performance:

Running on 11 TOPS SoC where usually, the processing of 2,304 channels would require up to 100 TOPS in a non-optimized solution

- Increasing the number of virtual channels creates a computational load existing radar processing SoC cannot handle
- Through sophisticated algorithms, computational load can be reduced to run on 11 TOPS SoC

Signal processing is shifted in full to the digital domain for superior accuracy

- Best in class RF and high fidelity sampling of FoV in 1 GHz/ 11 bits (industry standard 20 Mhz/ 8 bits)
- Applying advanced digital signal processing enabling detecting up to 500K detections per sec.

Implications:

Detecting motorcycles beyond 200m
Detecting cars up to 350m

Detecting remote low RCS targets in presence of strong close targets

Contour detection
500K PPS

Hazard cue, e.g. detecting a rimless tire @ 130m

CAM_FRONT

R-frame

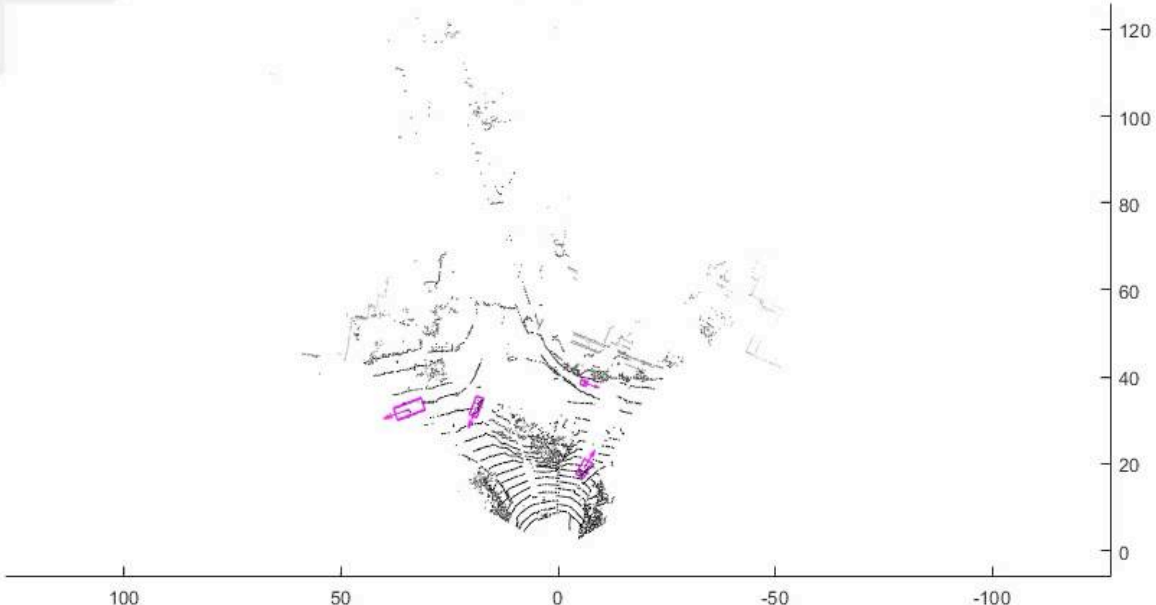

Detecting two closed pedestrians (low RCS) behind a vehicles (high RCS)

Detecting a pedestrian (low RCS) behind parked vehicles (high RCS)

Stable tire detection at 140m

FMCW LiDAR Development

FMCW LiDAR

Frequency-Modulated Continuous Wave (FMCW) LiDAR is conducting a coherent detection and uses Doppler effect to measure both range and closing speed while ToF are focused on short bursts and range measurement

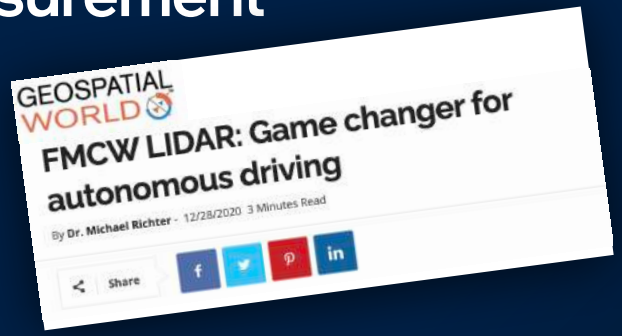
GEOSPATIAL WORLD

FMCW LIDAR: Game changer for autonomous driving

By Dr. Michael Richter - 12/28/2020 3 Minutes Read

Share

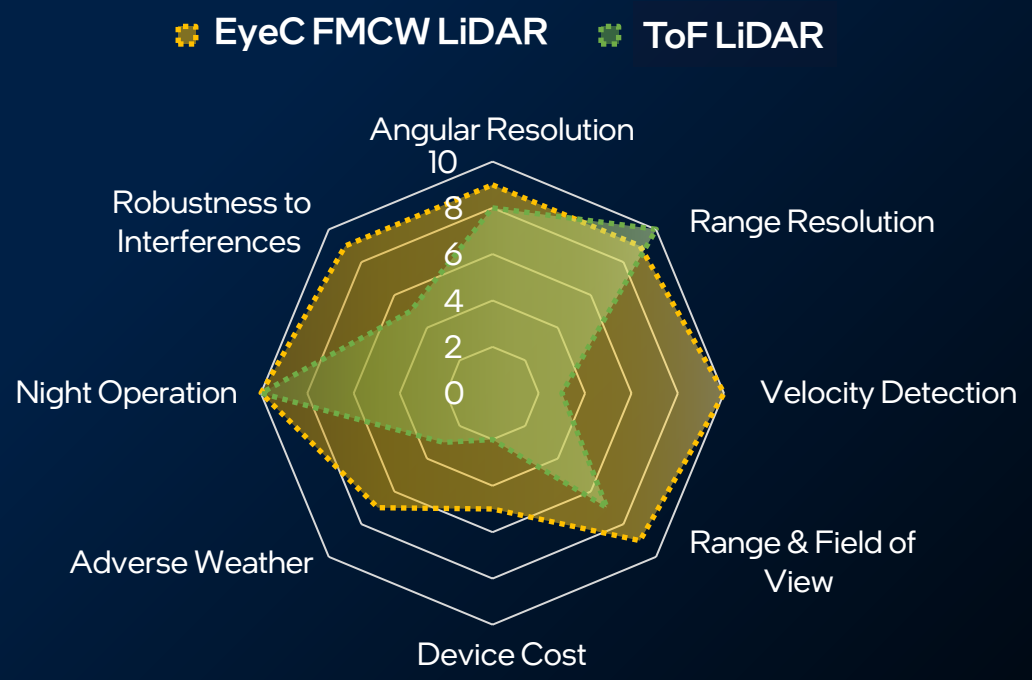

The Required Capabilities

Moving from ToF LiDARs	To FMCW	To achieve:
▪ 3D sampling (range, azimuth, elevation)	 ▪ 4D (+Doppler for relative velocity measurement)	▪ Instant heading measurement ▪ Point-levels TTC ▪ Reliable clustering and tracking
▪ 200m max. range ▪ $1/R^2$ decay	 ▪ 300m max. Range ▪ 1/R decay	▪ Enhanced ODD ▪ Higher effective dynamic range
▪ Sensitivity to interferences	 ▪ High immunity to interference through coherent detection	▪ Avoid impairments from sun, other LiDARs, retroreflectors
	 Maintaining high res. sampling ▪ 2M PPS ▪ 600 pts per degree^2	▪ Achieve accurate environmental perception

Examples

Direct Doppler velocity measurement

- Measurement of TTC at a point level
- Reliable clustering and object tracking
- Instant heading calculation

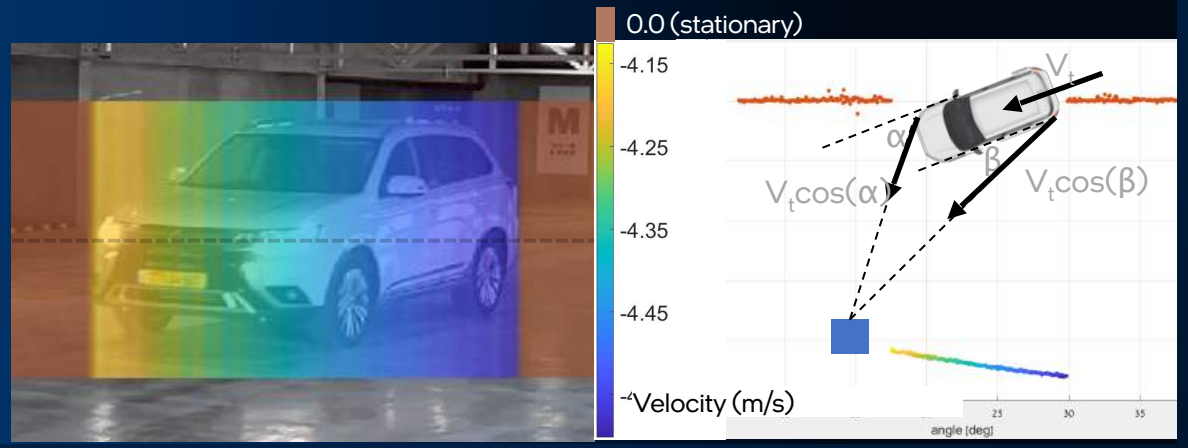

High Pd at longer ranges

FMCW LiDAR can transmit more energy over time (continuous waves while assuring Eye Safety)

Resolution: 0.05° x 0.05° demonstrating over 250 points (4D) on object @ 184m

Better immunity to interferences

thanks to coherent detection

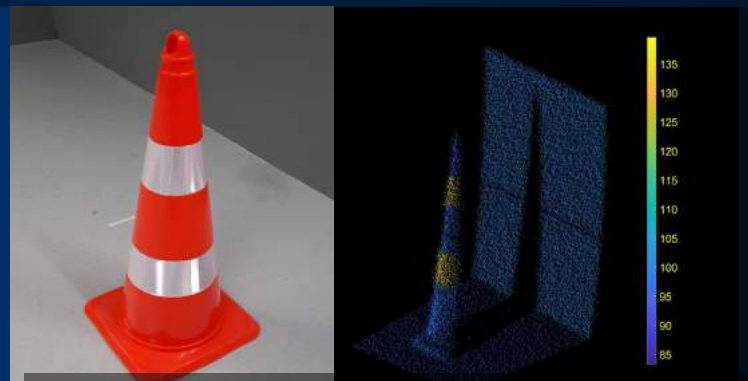

Detecting a cone with 600% reflectivity @ 10m with no blur

Harnessing Intel's Si Photonics Leadership to FMCW Sensor Development

Intel owns the only FAB in the world capable of integrating Active and Passives Components, enabling highly compressed and large scale PIC with the required on-chip interconnects and at high volume manufacturing

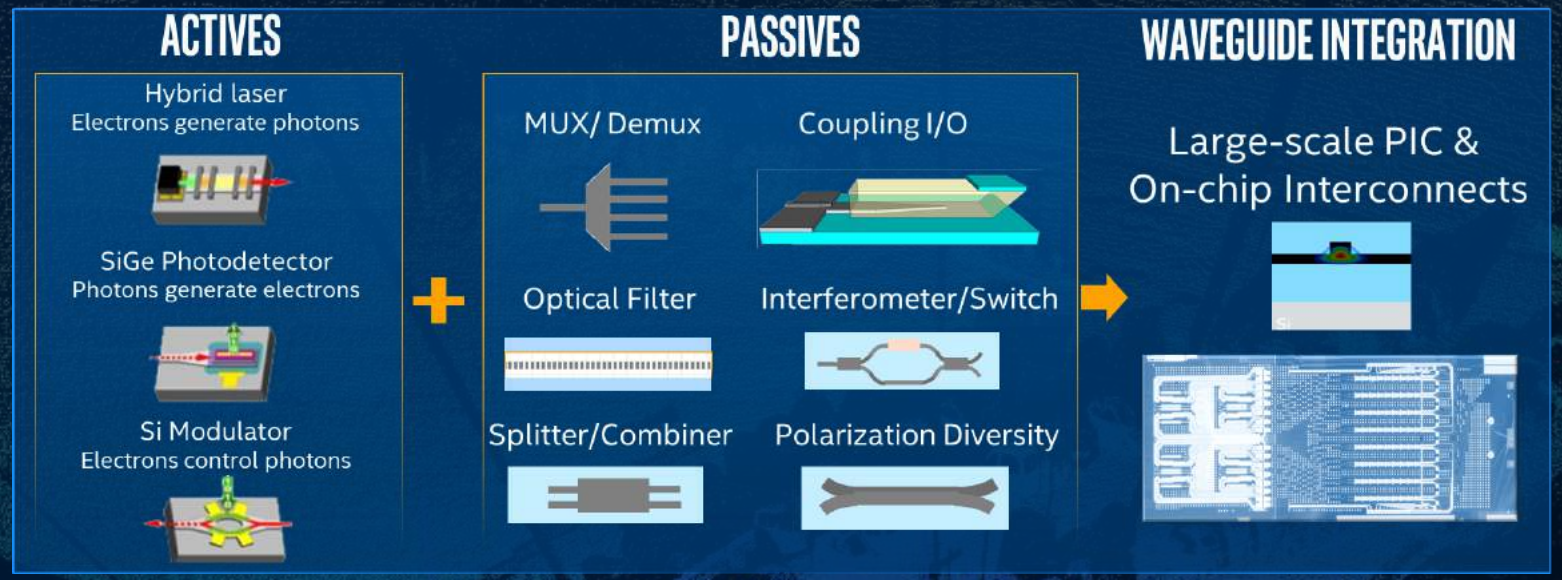

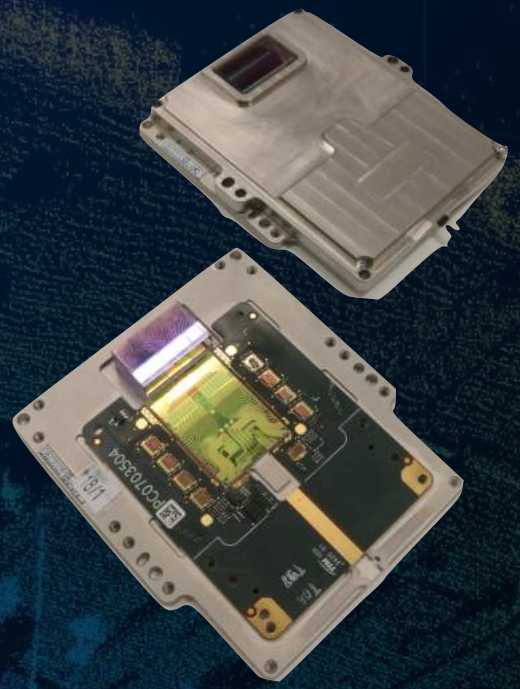

Optical head enabling 184 vertical lines per scan

RSS-based Driving Policy

What is RSS?

A formal model for safety, that provides mathematical formalization for the AV to never cause an accident

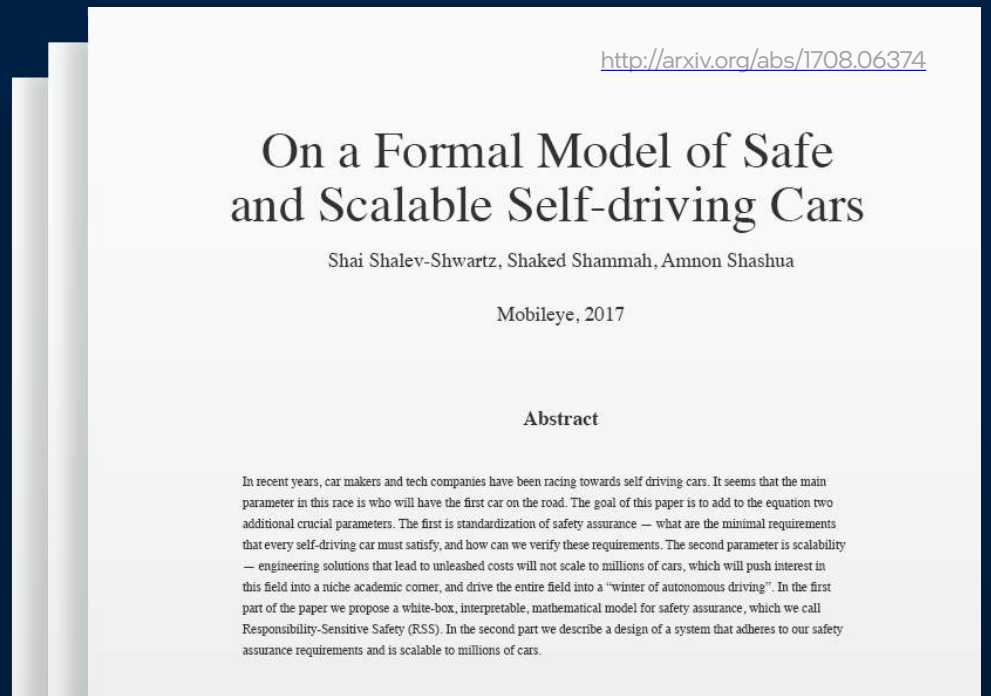

http://arxiv.org/abs/1708.06374

On a Formal Model of Safe and Scalable Self-driving Cars

Shai Shalev-Shwartz, Shaked Shammah, Amnon Shashua

Mobileye, 2017

Abstract

In recent years, car makers and tech companies have been racing towards self driving cars. It seems that the main parameter in this race is who will have the first car on the road. The goal of this paper is to add to the equation two additional crucial parameters. The first is standardization of safety assurance — what are the minimal requirements that every self-driving car must satisfy, and how can we verify these requirements. The second parameter is scalability — engineering solutions that lead to unleashed costs will not scale to millions of cars, which will push interest in this field into a niche academic corner, and drive the entire field into a "winter of autonomous driving". In the first part of the paper we propose a white-box, interpretable, mathematical model for safety assurance, which we call Responsibility-Sensitive Safety (RSS). In the second part we describe a design of a system that adheres to our safety assurance requirements and is scalable to millions of cars.

The Method

01 Defining reasonable boundaries on the behavior of other road users

02 Within the boundaries specified by RSS, one must always assume the worst-case behavior of other agents

03 The boundaries capture the common sense of reasonable assumptions that human drivers make

04 Any action beyond the defined boundaries is not reasonable to assume

RSS Principles

RSS formalizes five common sense rules apply to all driving situations

RSS

01	Do not hit someone from behind
02	Do not cut-in recklessly
03	Right-of-way is given, not taken
04	Be careful of areas with limited visibility
05	If you can avoid an accident without causing another one, you must do so

Industry Standardization Efforts

IEEE
Advancing Technology for Humanity

IEEE WORKGROUP TO DEFINE A FORMAL MODEL FOR AV SAFETY CHAIRED BY INTEL-MOBILEYE

- Workgroup consists of 30 leading industry players
- 1st Project: Standardizing assumptions (i.e. parameters from RSS) that must be used in AV Decision Making
- This standard will provide governments the framework they need to set the acceptable balance between safety and usefulness

ISO

ISO/TR 4804:2020

SAFETY AND CYBERSECURITY FOR AUTOMATED DRIVING SYSTEMS

- World's first ISO Technical Report defining a Safe-By-Design Automated Driving System
- RSS featured as a key element to implementing a safe Driving Policy
- Recently approved for promotion to an ISO Technical Specification

Government Efforts

U.S. Department of Transportation

ADVANCED NOTICE OF PROPOSED RULE MAKING: FRAMEWORK FOR AV SYSTEM SAFETY

- US DOT Seeking public comment on the development of a framework for Automated Driving System (ADS) Safety
- RSS recognized as a "Notable Effort Under Consideration" as an Engineering Measure for Safety
- RSS cited as an approach that could require AV's "**to be programmed to drive defensively in a risk-minimizing manner in any scenario within their ODD"**

Law Commission
Reforming the law

THIRD AND FINAL CONSULTATION, PROPOSING A REGULATORY FRAMEWORK FOR AV'S IN UK

- Builds on three years of work including two previous consultations
- RSS proposed as a way to define "how safe is safe enough" by defining a "does not cause a fault"
- RSS featured as a way to define "road craft" - a safety envelope around the AV defined by safe distances

Mobility-as-a-Service

Mobility Supply is Shaping in Two Main Streams

PRESENT

PTOs
Public Transport Operators
E.g. RATP, Keolis, Transdev, MTA, etc.

FUTURE

APTO
Automated Public Transport Operators

PRESENT

TNC
Transportation Network Companies
E.g. Uber, Lyft, Grab, Didi, etc.

FUTURE

ATNC
Automated Transportation Network Companies

E.g. Waymo, Cruise, Argo, etc.

Asset Light
E.g. Uber, Lyft, Grab, Didi, etc.

Mobility Supply is Shaping in Two Main Streams

PRESENT

PTOs
Public Transport Operators
E.g. RATP, Keolis, Transdev, MTA, etc.

TNC
Transportation Network Companies
E.g. Uber, Lyft, Grab, Didi, etc.

PRESENT

VaaS/RaaS

RaaS

MaaS

mobileye
An Intel Company

moovit

Through a diversified XaaS strategy Mobileye-Moovit will serve as the enabler to drive PTOs and TNCs to the autonomous revolution

FUTURE

APTO
Automated Public Transport Operators

ATNC
Automated Transportation Network Companies

E.g. Waymo, Cruise, Argo, etc.

Asset Light
E.g. Uber, Lyft, Grab, Didi, etc.

FUTURE

Harnessing the World's leading Mobility platform to power our robotaxi service

moovit

World's most popular Urban Mobility app

- +950 Million users worldwide
- Service in 3400 Cities, 112 Countries

Largest and most accurate transit data repository

- Local Editors Community of over 700,000 volunteers
- Used also by Microsoft, Apple, Uber to power their offering

Most Advanced Multi-Modal Journey Planner of alll Mobility Forms

Mobile Payments and Ticketing

- Plan, Pay and Ride - all in one app to offer a holistic experience

Transit On-Demand:

- End-to-end platform to plan and power Fixed & Dynamic Routes
- Human driven and Autonomous vehicle services
- Dozens of existing customers

Partnerships with hundreds of transit agencies, transit operators and cities

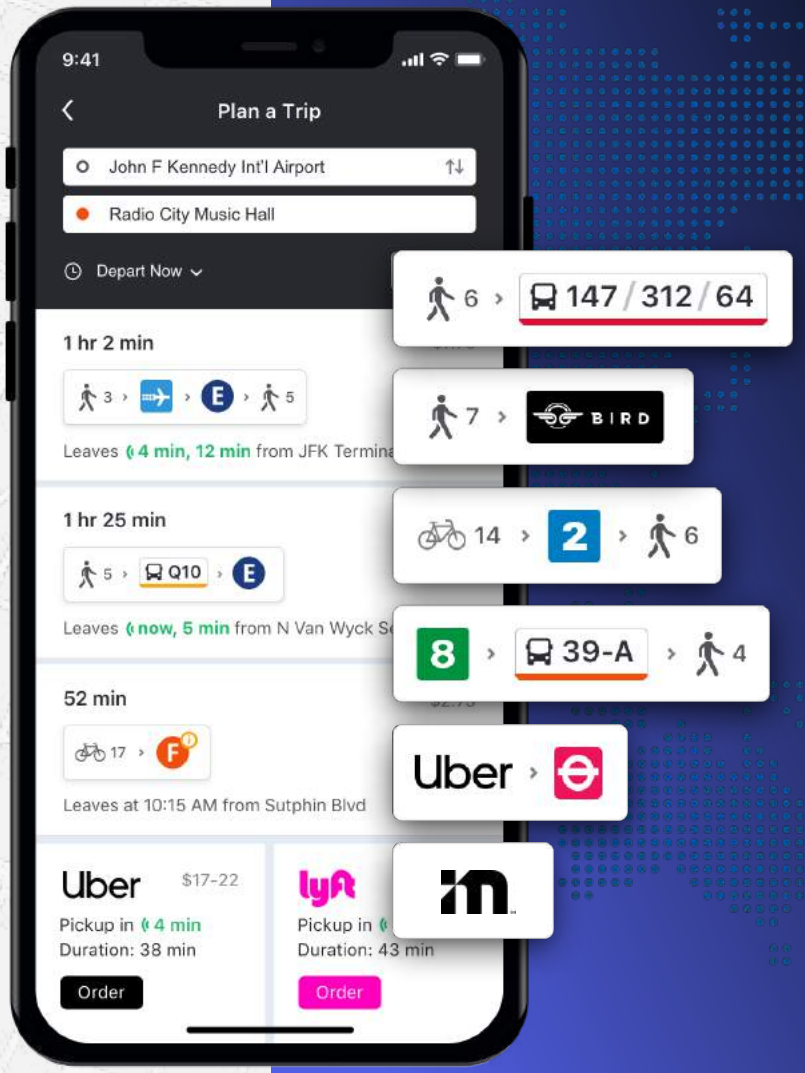

Mobileye-Moovit Driverless MaaS

MaaS Layer 5 moovit	UX & UI
MaaS Layer 4 moovit	Mobility Intelligence and demand analysis
MaaS Layer 3 mobileye | moovit	Tele Ops Fleet Ops & Control Center
MaaS Layer 2 mobileye	Self-Driving Vehicles
MaaS Layer 1 mobileye	Self-Driving System

The Value of Moovit in our playbook:

- Building an E2E service platform (layers 3-5) using Moovit's assets
- Pre-acquired public transit User base
- Global PTO relationship network
- Pre-planning and optimization based on demand entries, patterns and insights
- Solving the mixed-fleet issue

Mobility-as-a-Service Global Footprint

First MaaS deployment in Israel

- Israel as a sandbox- JV with VW
- Early riders program to start 2021
- Commercial driverless pilot to start in 2022 upon regulatory approval

Collaboration with WILLER Group

- Targeting commercial launch in Osaka in 2023
- Future expansion to Singapore and southeast Asia

MaaS deployment in France

- Deals with the two largest EU PTOs
- Testing in Paris to start next month
- Future expansion to other EU countries

MaaS deployment in Daegu city

- Collaboration with the city of Daegu
- Starting AV Testing by mid 2021
- Promoting AV regulation

intel | mobileye An Intel Company | moovit

mobileye
An Intel Company

Thank you!