CSC2125: Safety and Certification of Autonomous Vehicles

Lecture 1: Autonomous Driving System (ADS)

http://www.cs.toronto.edu/~chechik/courses19/csc2125

Goal of the Course

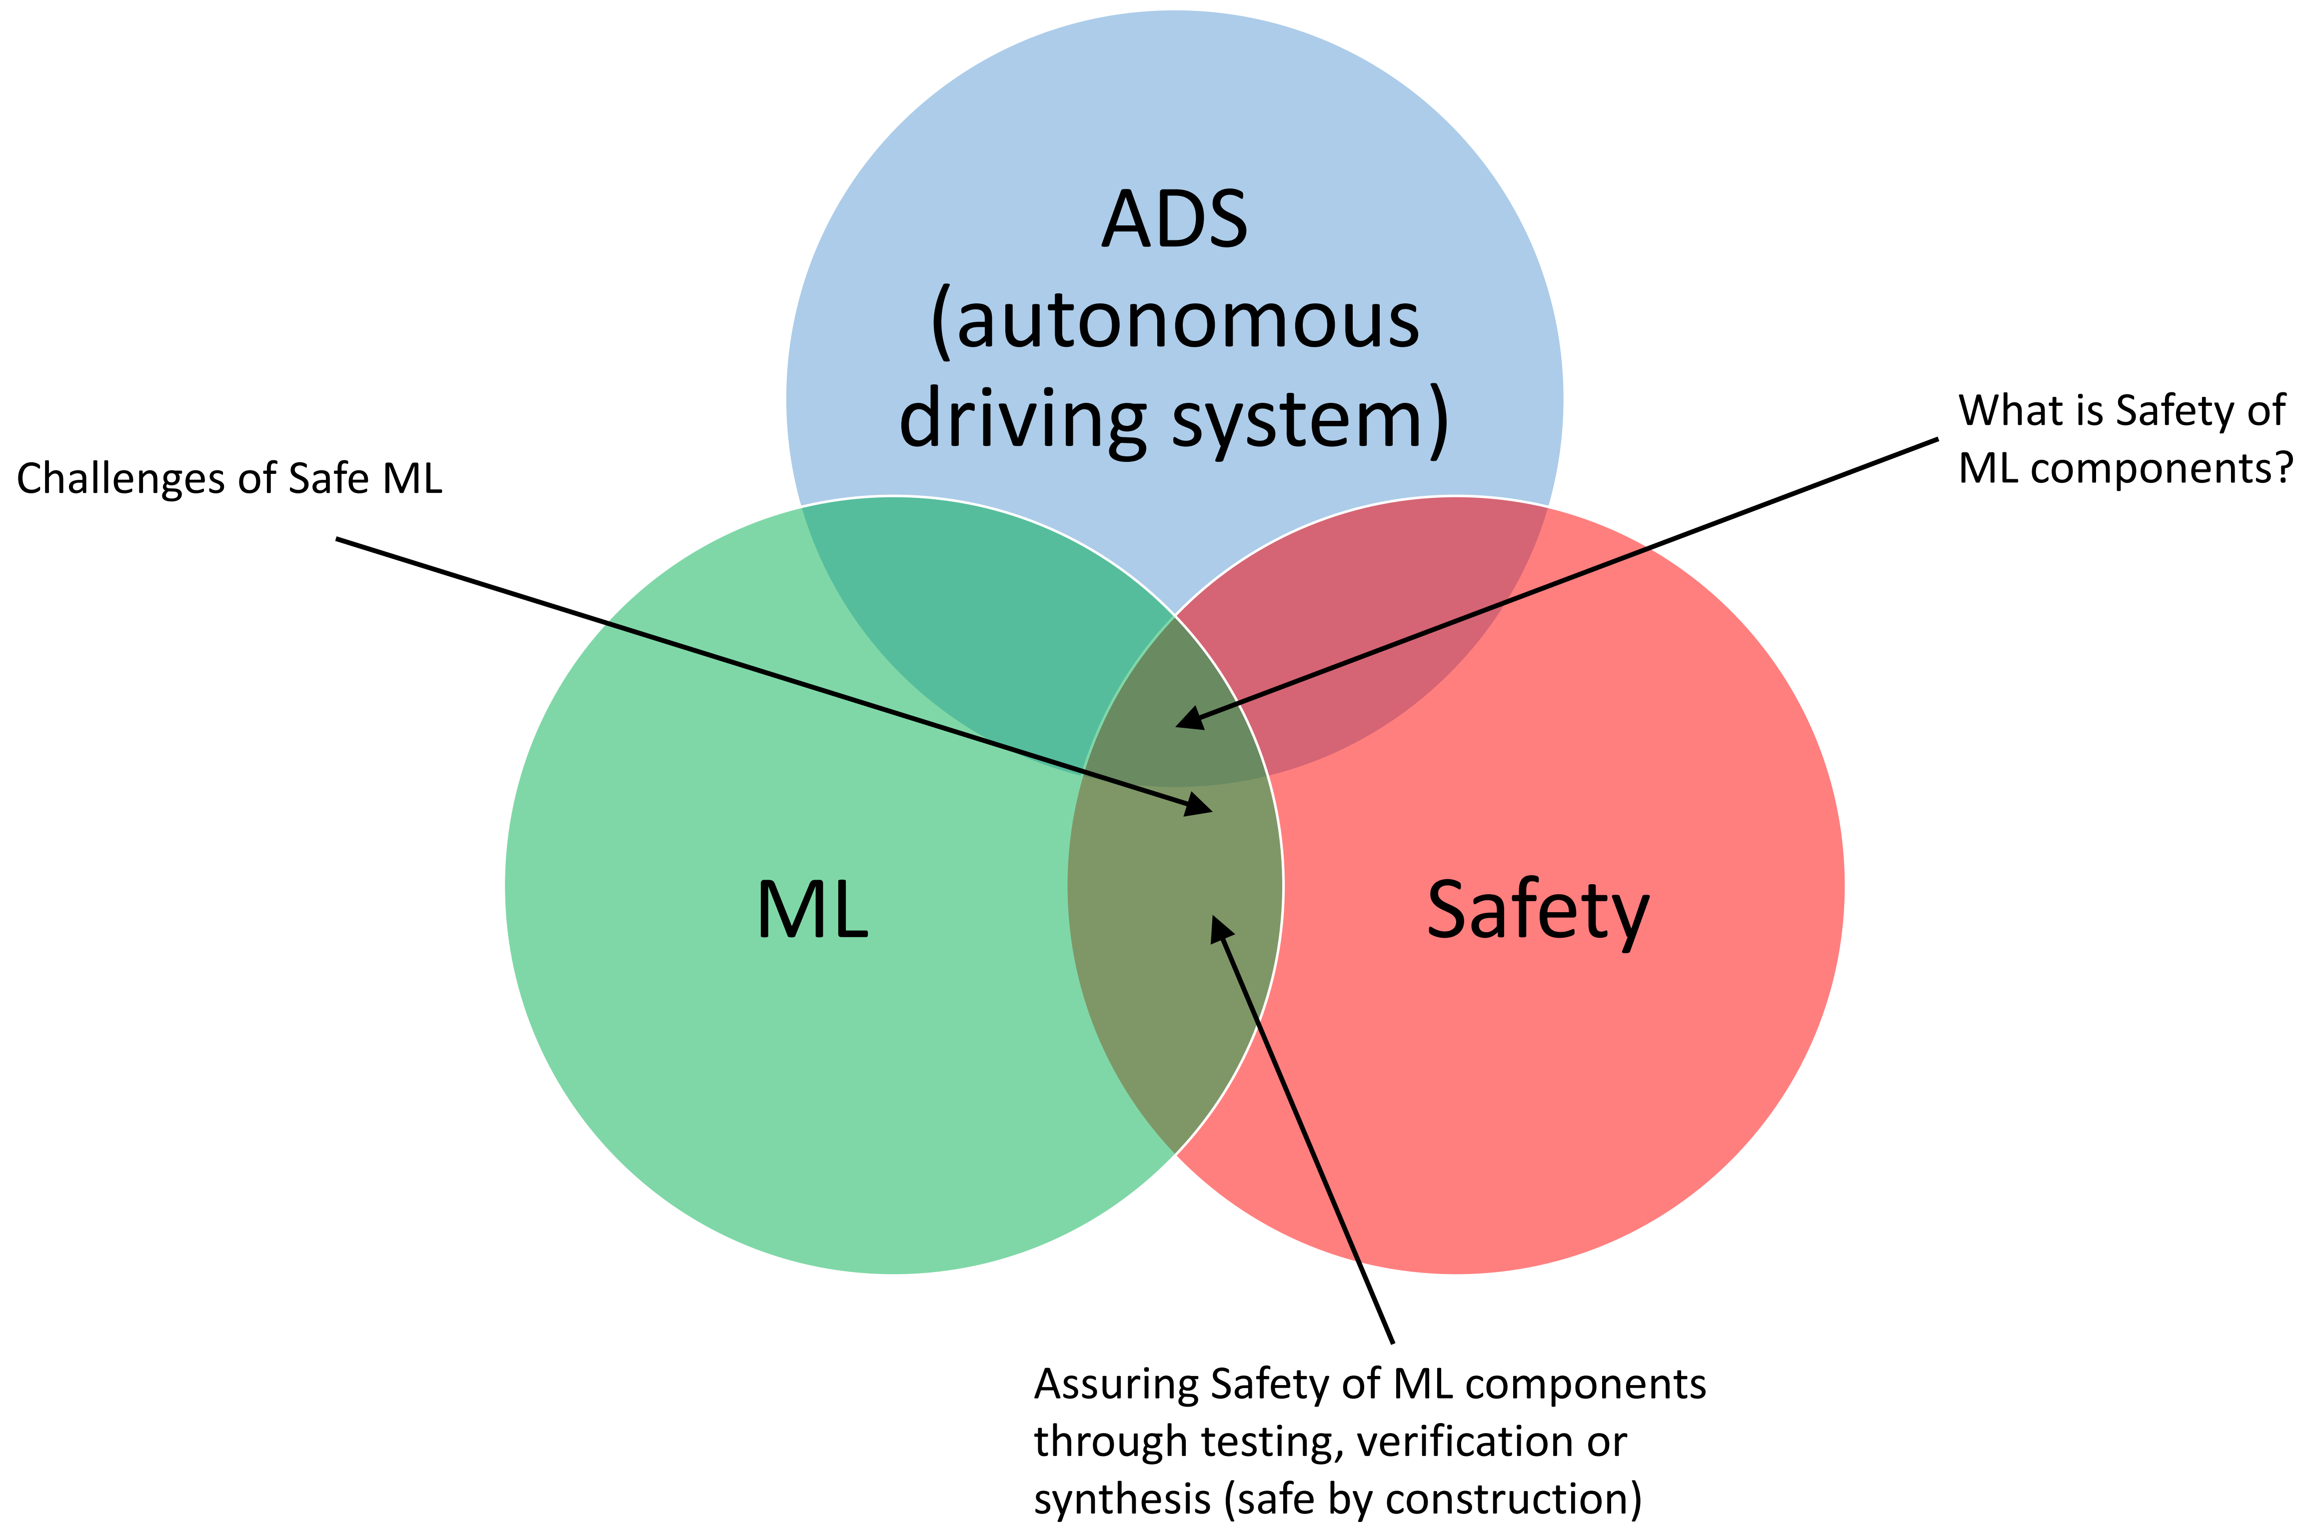

The Dream of Self-Driving

Lecture plan

- Levels of automation of ADS
- Current big players in ADS
- Functional reference architecture for ADS
  - And a look at some sensing technology
- A1 vs A2 autonomy
- A look at several ADS accidents
- Safety assurance of ADS
- ADS challenges

SAE J3016 Levels of Automation

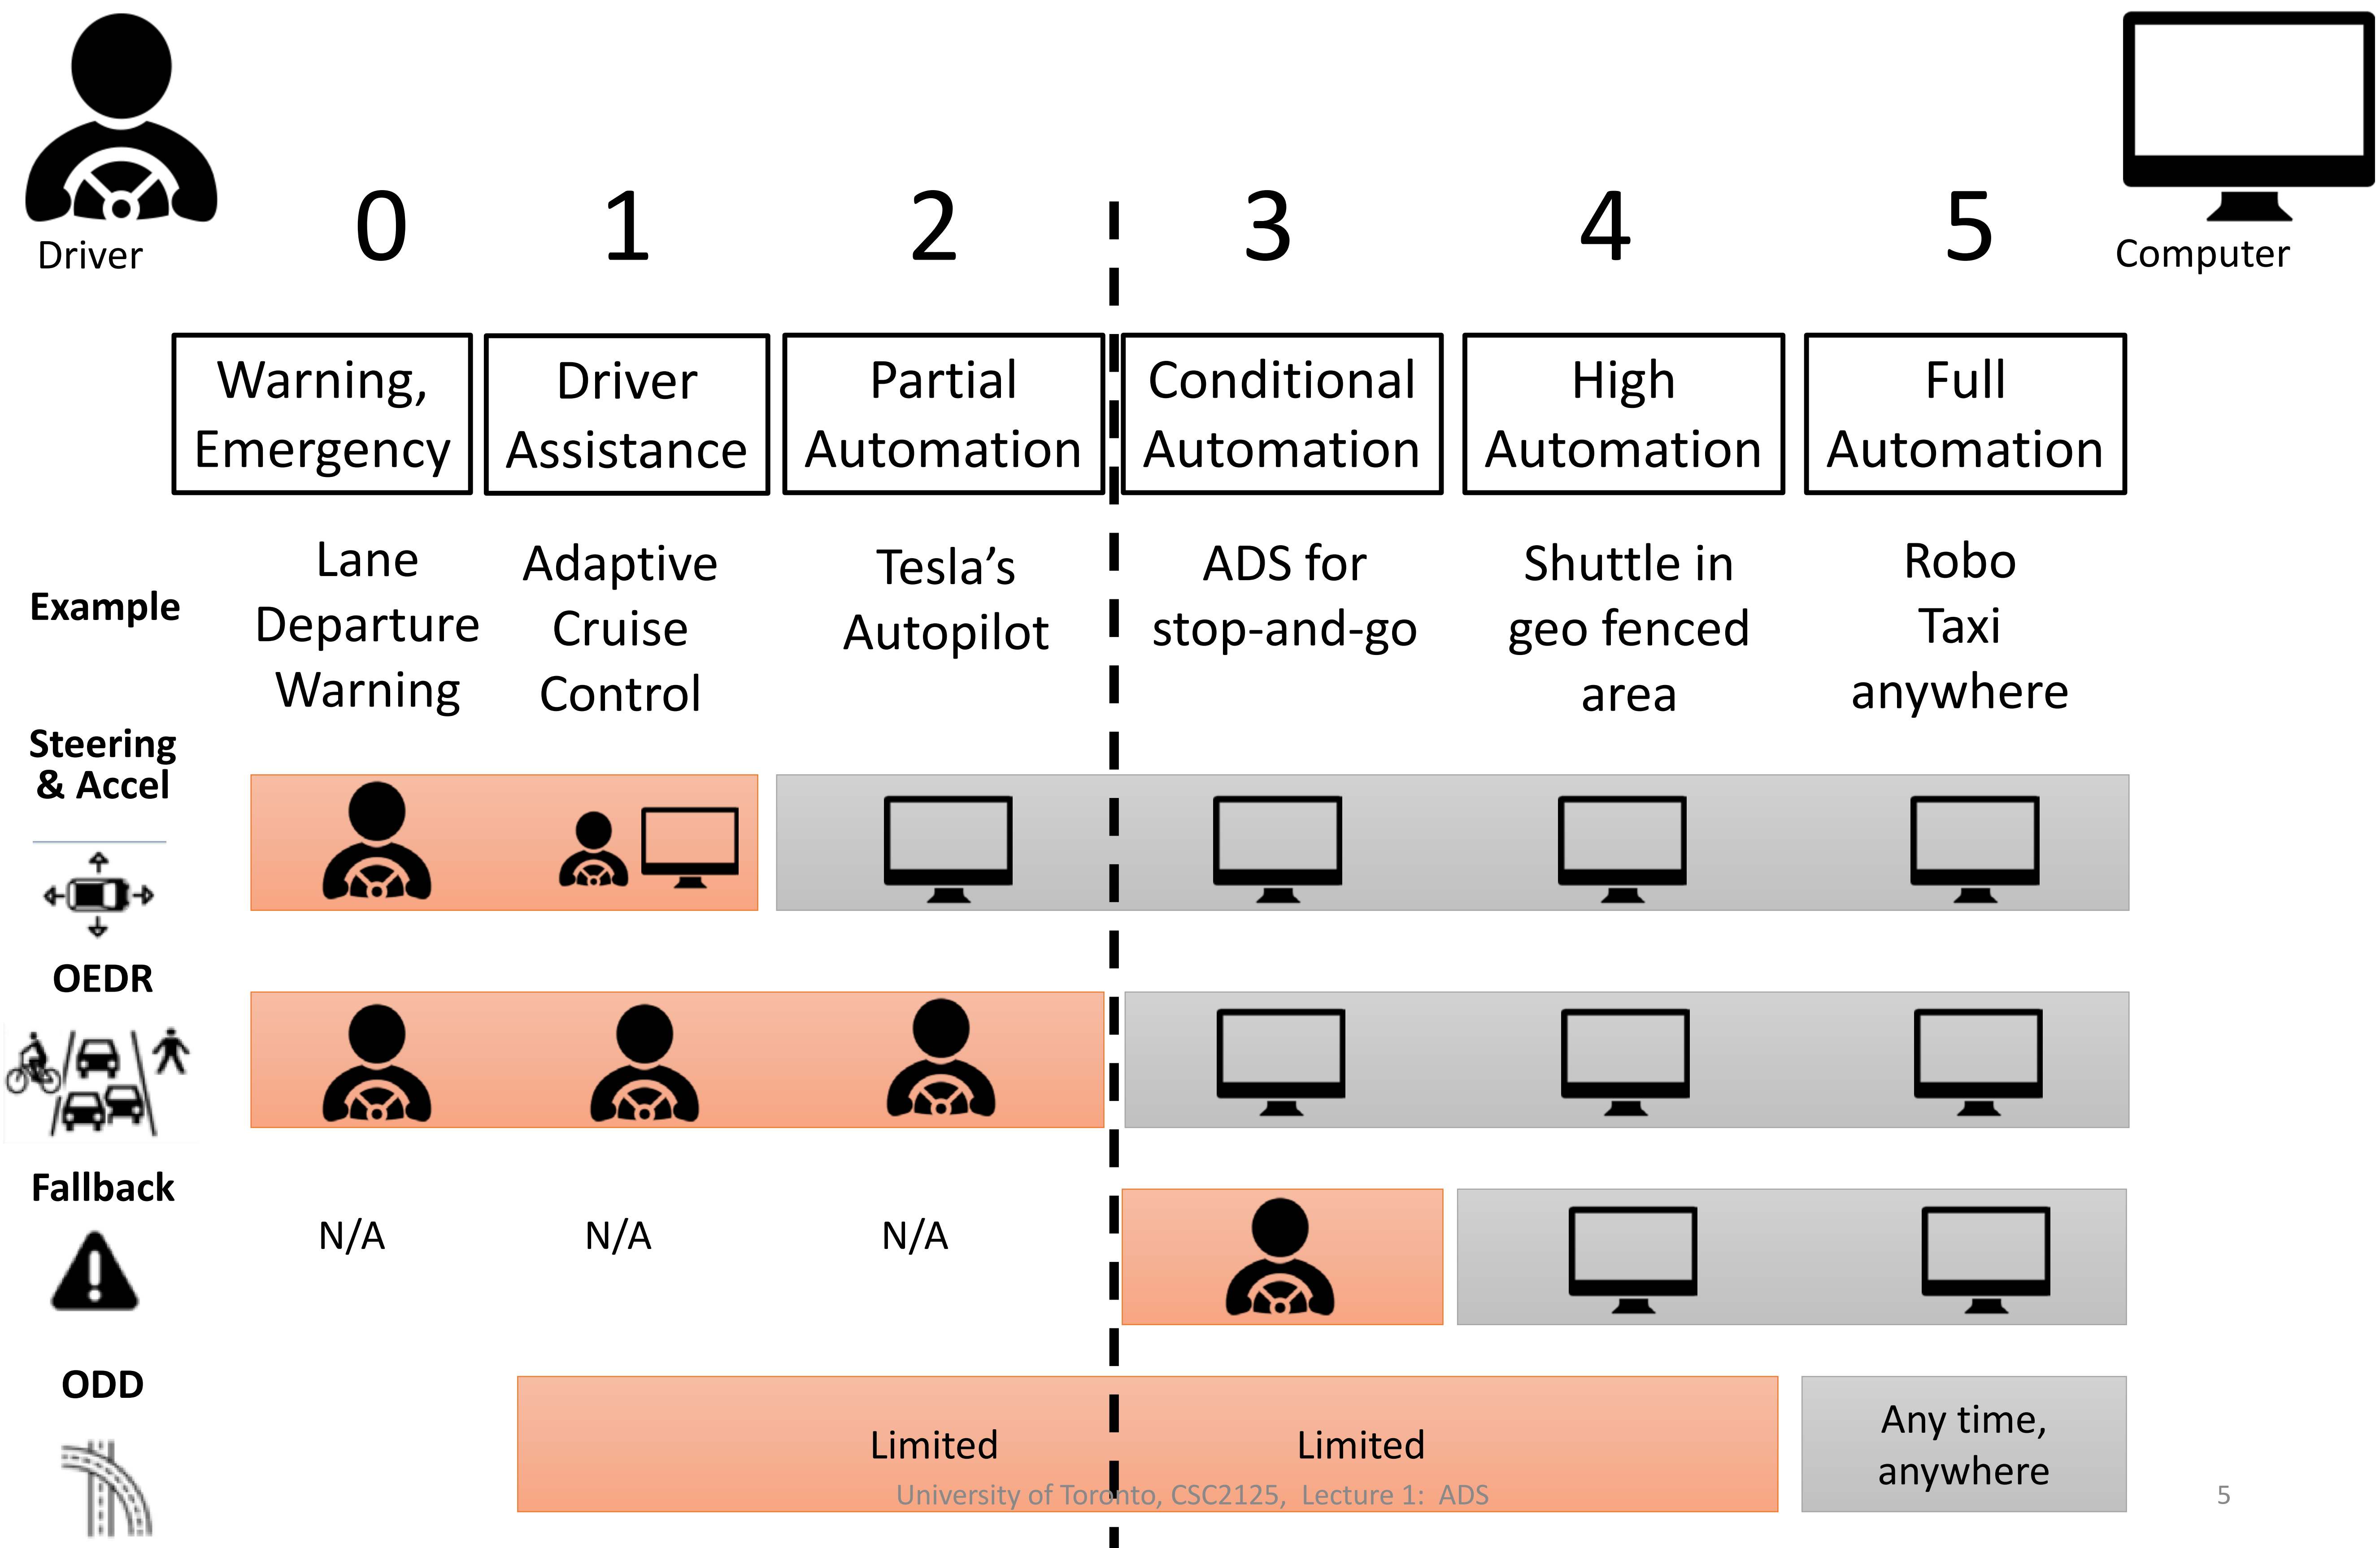

SAE J3016 Levels of Automation

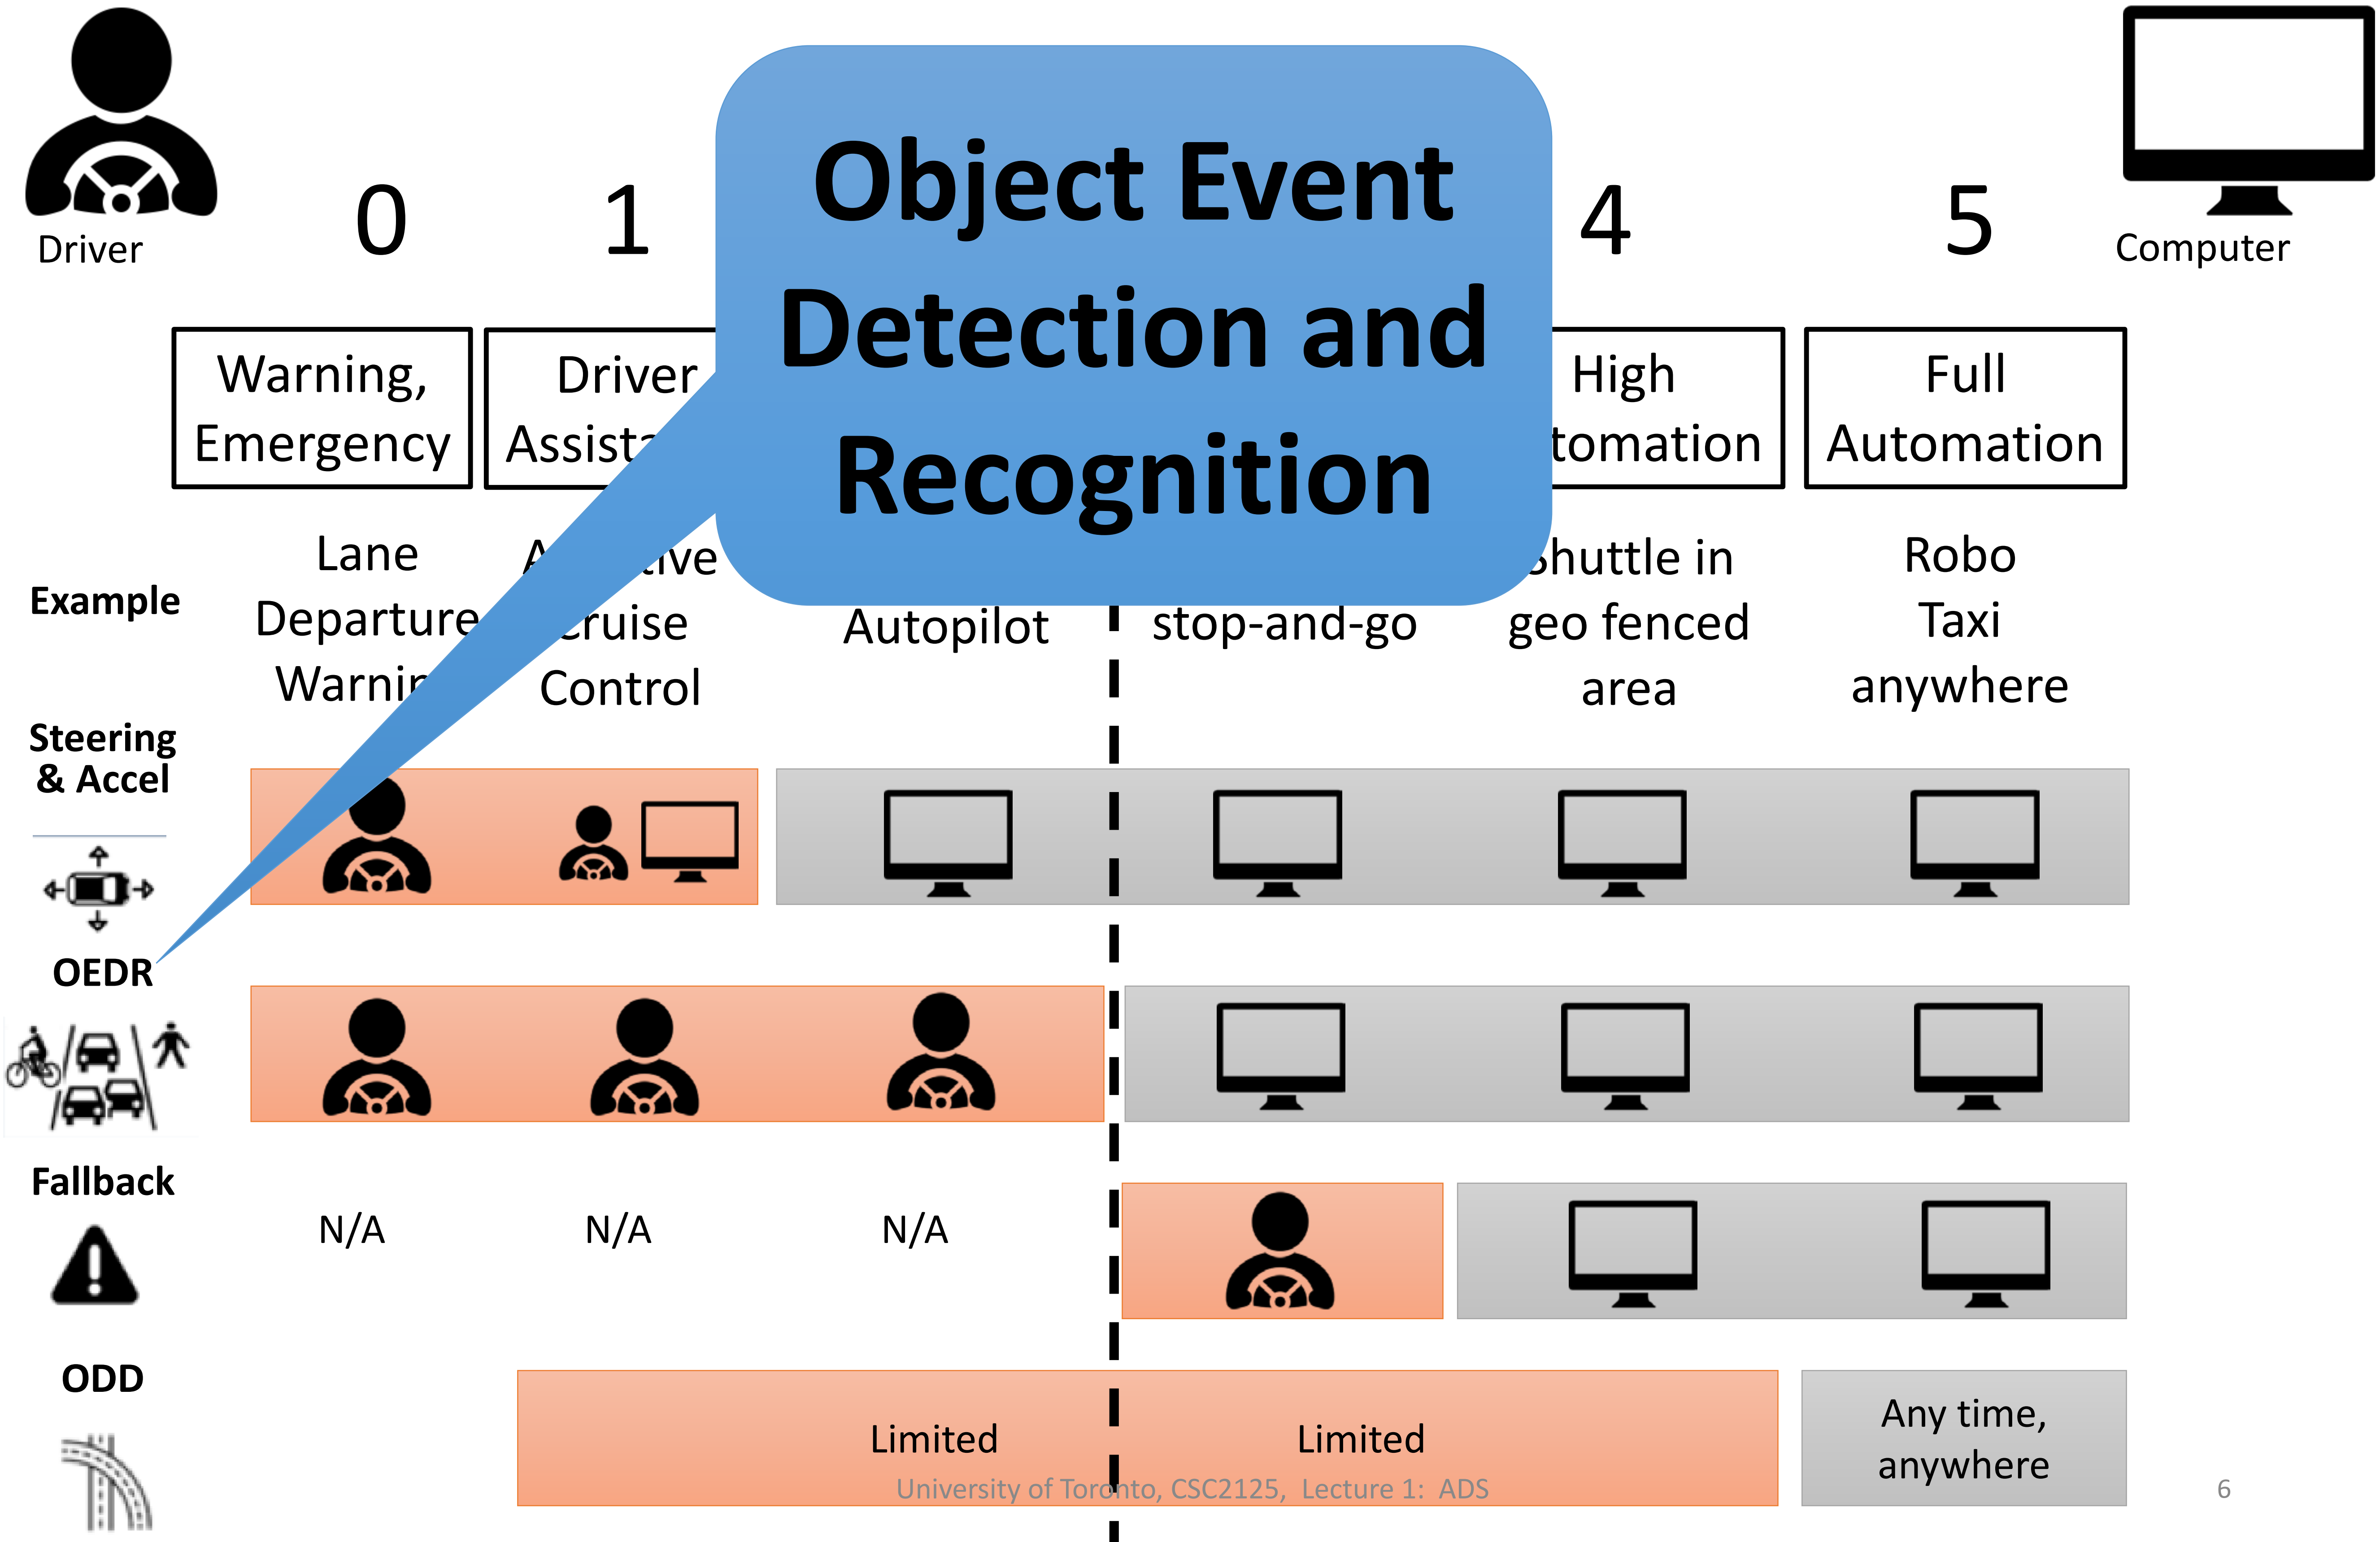

SAE J3016 Levels of Automation vs. Operational Design Domain (ODD)

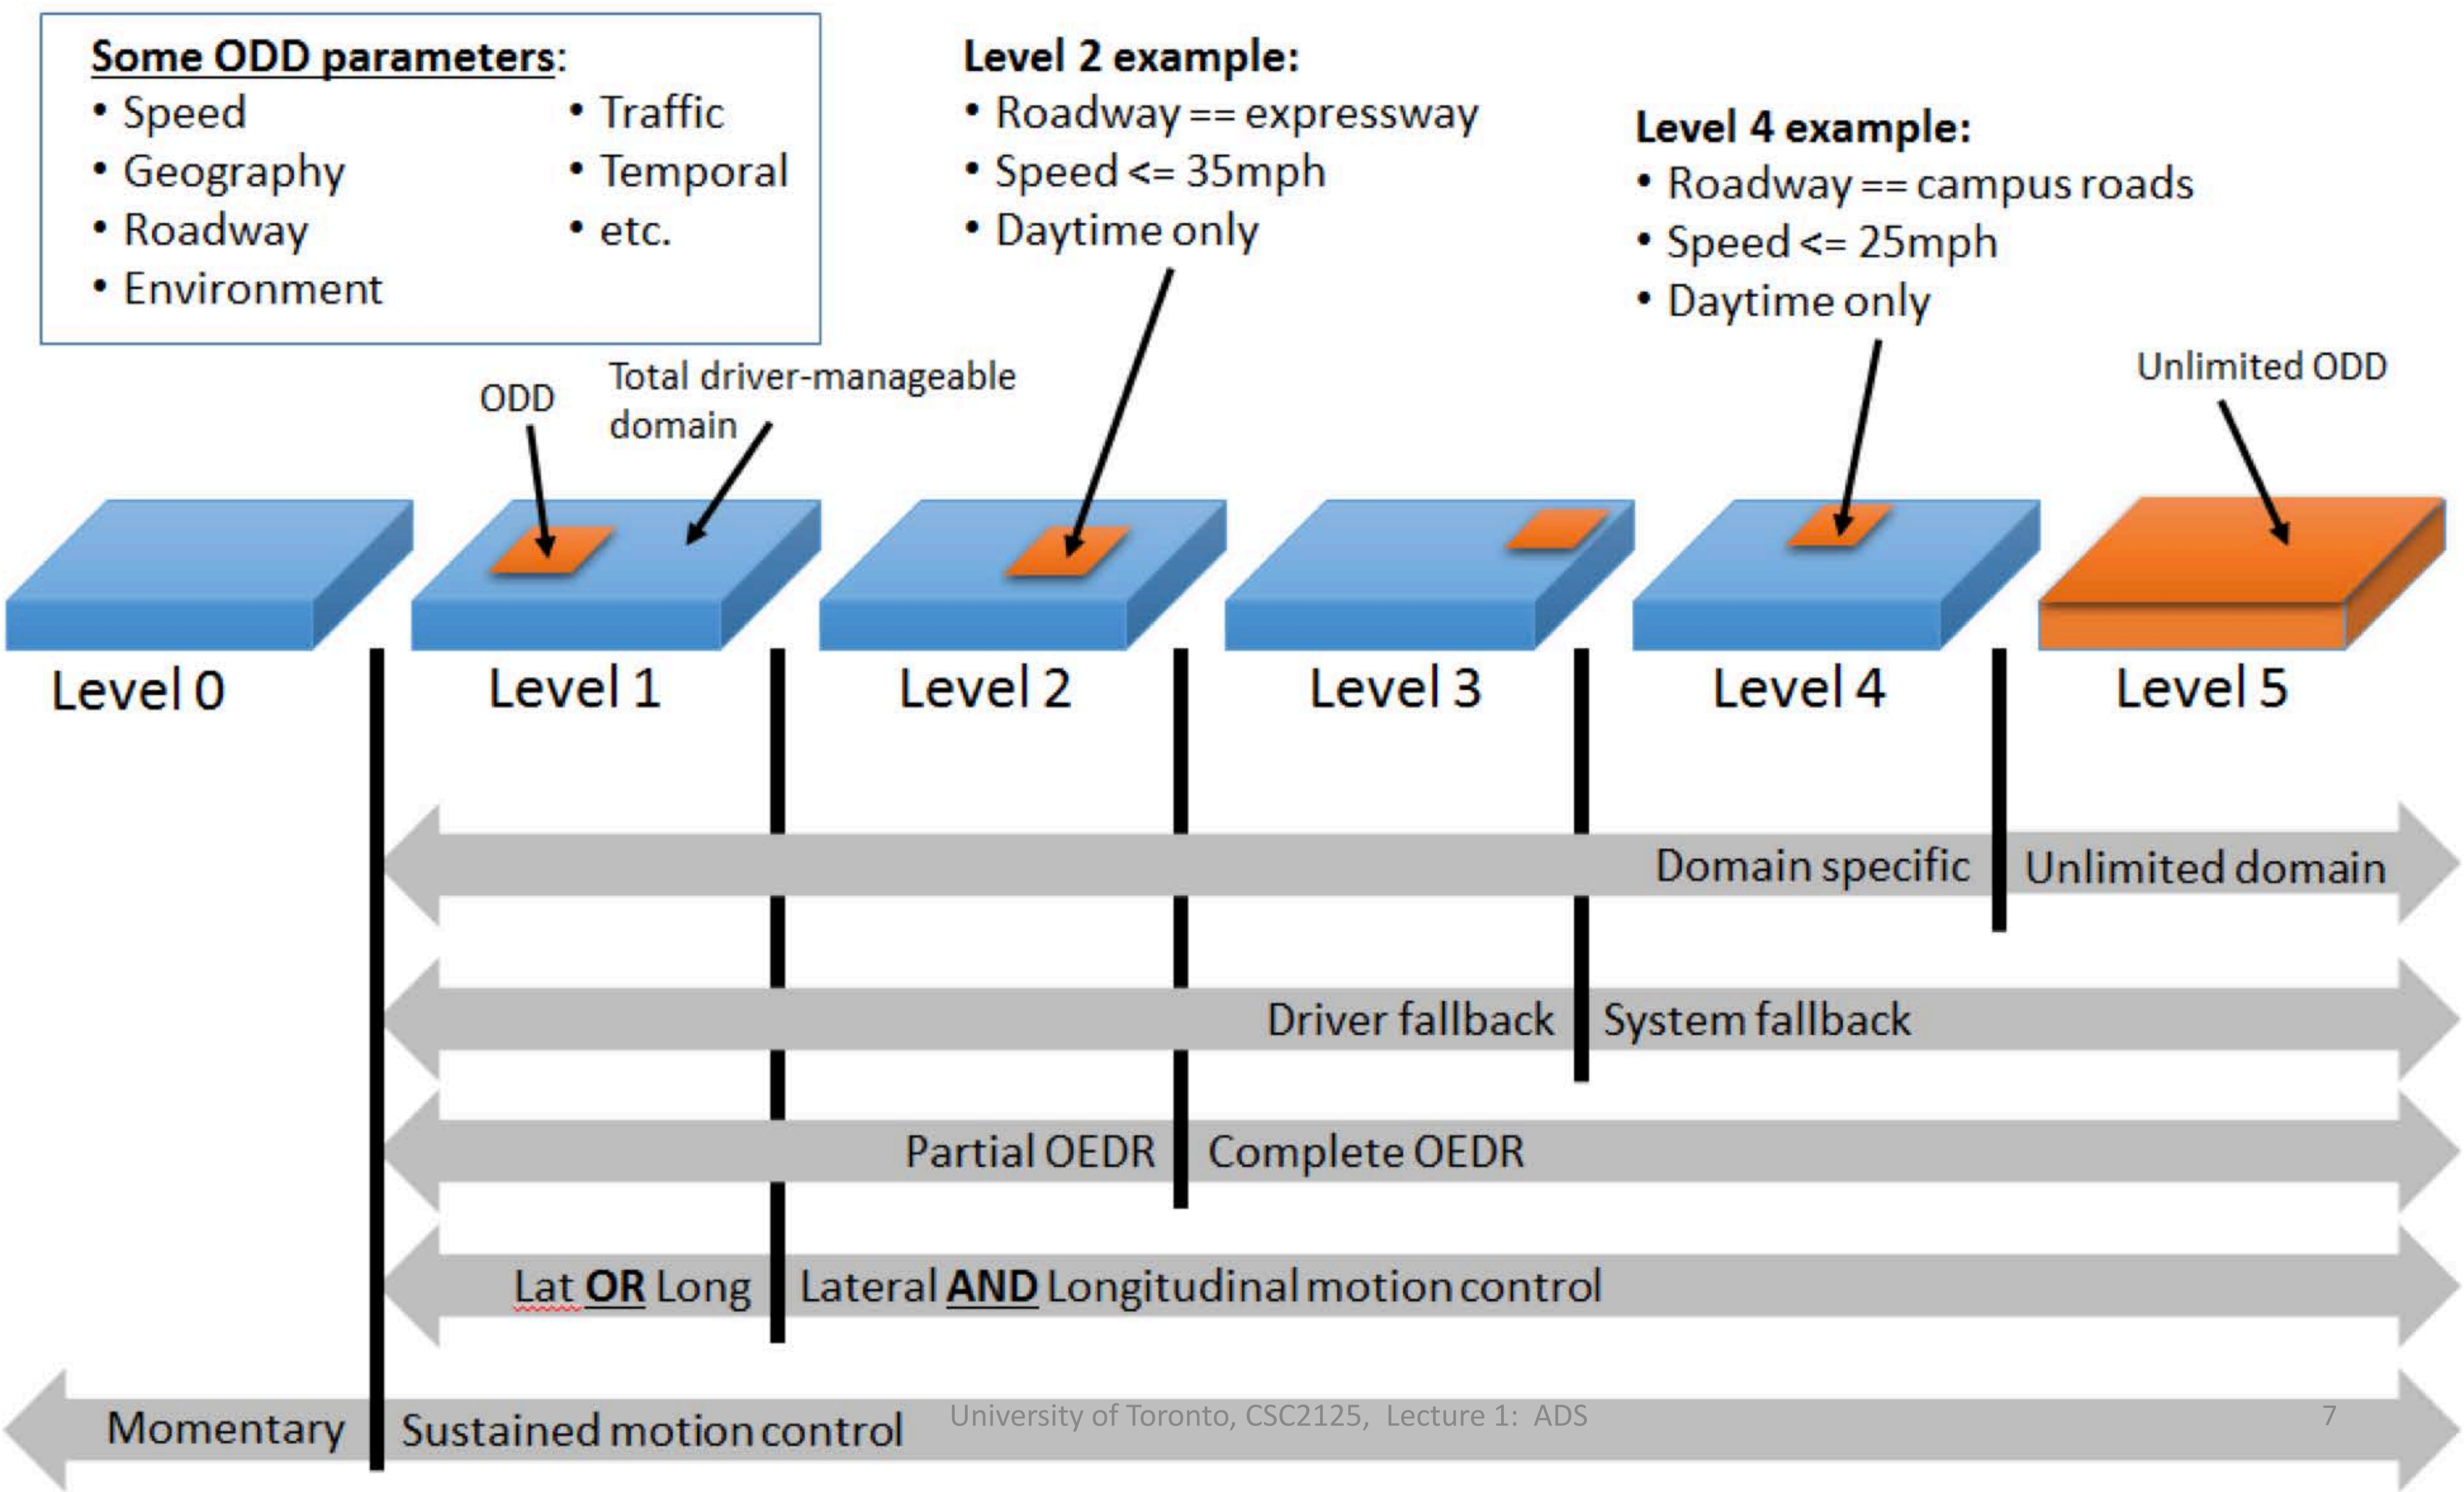

Beyond Traditional Levels: Two types of AI

- **Starting point:**
  - All cars are manually controlled until the AI system shows itself to be available and is elected to be turned on by the human.
- **A1:** Human-Centered Autonomy
  - **Definition:** AI is not fully responsible
  - Feature axis:
    - Where/how often is it "available"? (traffic, highway, sensor-based, etc.)
    - How many seconds for take-over? (0, 1, 10, etc)
    - Teleoperation support
- **A2:** Full Autonomy
  - **Definition:** AI is fully responsible
  - Notes:
    - No teleoperation
    - No 10-second rule: It's allowed to ask for human help, but not guaranteed to ever receive it.
    - Arrive to a **safe** destination or safe harbor.
    - Allow the human to take over **when they choose to**.

Beyond Traditional Levels: Two types of AI

L0 →
- **Starting point:**
  - All cars are manually controlled until the AI system shows itself to be available and is elected to be turned on by the human.

L1, L2, L3 →
- **A1:** Human-Centered Autonomy
  - **Definition:** AI is not fully responsible

L4, L5 →
- **A2:** Full Autonomy
  - **Definition:** AI is fully responsible

Example Players

Waymo

Notable:

- April 2017: Exits testing**:** first rider in Phoenix
- November 2017: 4 million miles driven autonomously
- December 2017: No safety driver in Phoenix

Uber

Notable:

- December 2017: 2 million miles driven autonomously

Tesla

Notable:

- Sep 2014: Released Autopilot
- Oct 2016: Started Autopilot 2 from scratch.
- Jan 2018: ~1 billion miles driven in Autopilot
- Jan 2018: ~300,000 Autopilot equipped vehicles

Audi A8 (released end of 2018)

- Thorsten Leonhardt, head of Automated Driving, Audio:
  "When the function is operated as intended, if the customer turns the traffic jam pilot on and uses it as intended, and the car was in control at the time of the accident, the driver goes to his insurance company and the insurance company will compensate the victims of the accident and in the aftermath they come to us and we have to pay them," he said.

Notable Progress

- Full autonomy (A2)
  - Waymo
  - Uber
  - GM Cruise
  - nuTonomy
  - OptimusRide
  - Zenuity
  - Voyage
  - …
- Human-centered autonomy (A1)
  - Tesla Autopilot - Model S/3/X
  - Volvo PilotAssist - S90/XC90/XC60/V90
  - Audi Traffic Jam Assist - A8
  - Mercedes-Benz Drive Pilot Assist - E-Class
  - Cadillac Super Cruise - CT6
  - Comma.ai openpilot
  - …

Paths to Autonomous Future

A1:
Human-Centered Autonomy

- **Localization and Mapping:** Where am I?
- **Scene Understanding:** Where/who/what/why of everyone else?
- **Movement Planning:** How do I get from A to B?
- **Human-Robot Interaction:** What is the physical and mental state of the driver?
- **Communicate:** How to I convey intent to the driver and to the world?

Blue Text: Easier

Red Text: Harder

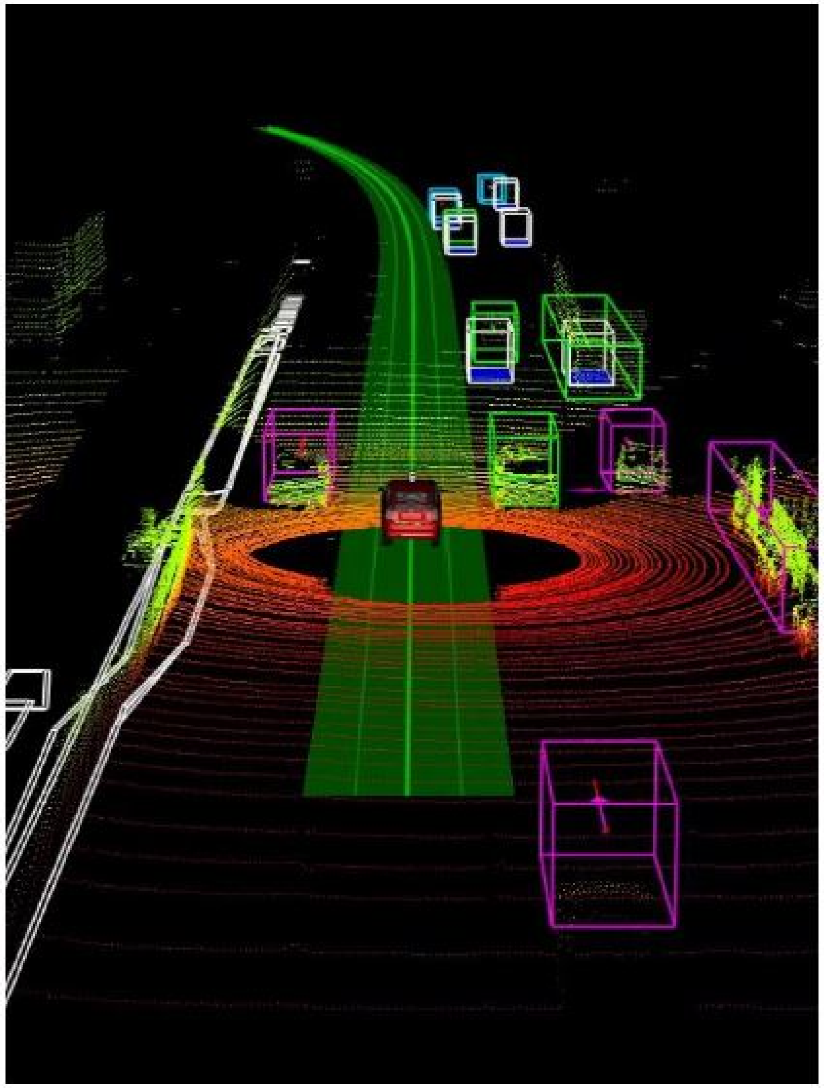

A2:
Full Autonomy

- **Localization and Mapping:** Where am I?
- **Scene Understanding:** Where/who/what/why of everyone else?
- **Movement Planning:** How do I get from A to B?
- **Human-Robot Interaction:** What is the physical and mental state of the driver?
- **Communicate:** How to I convey intent to the driver and to the world?

What does it take to drive a car?

1. Perception

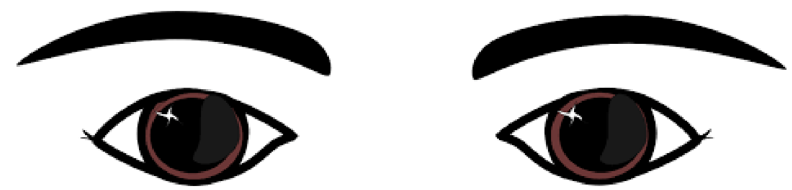

2. Decision making

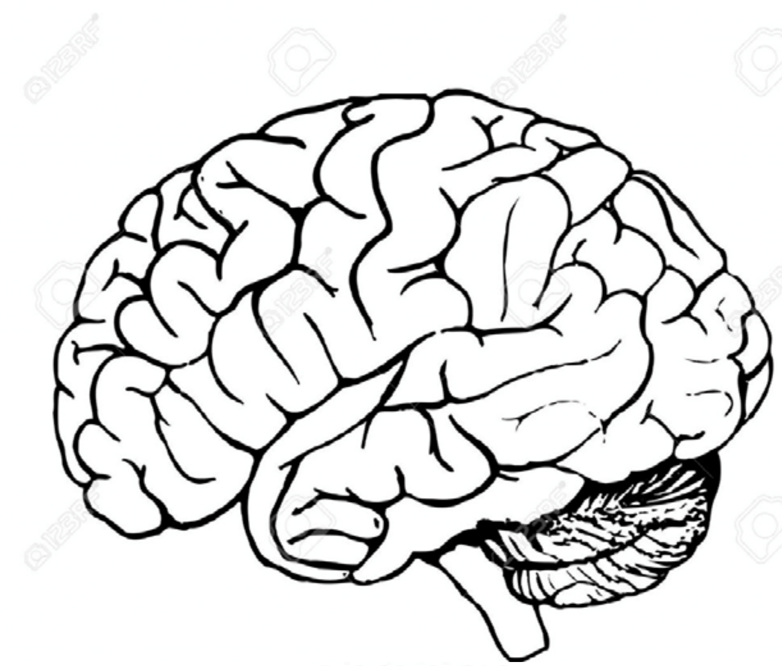

3. Control

Self-Driving Car Tasks

- Localization and Mapping – Where am I
- Scene Understanding – Where is Everyone Else?
- Movement Planning – How to get from Point A to Point B
- Driver State – What is the Driver Up to?
  - Essential if driver is part of the loop!
- Safety Monitoring

Functional Reference Architecture

Source: Krzysztof Czarnecki, Waterloo

Functional Reference Architecture

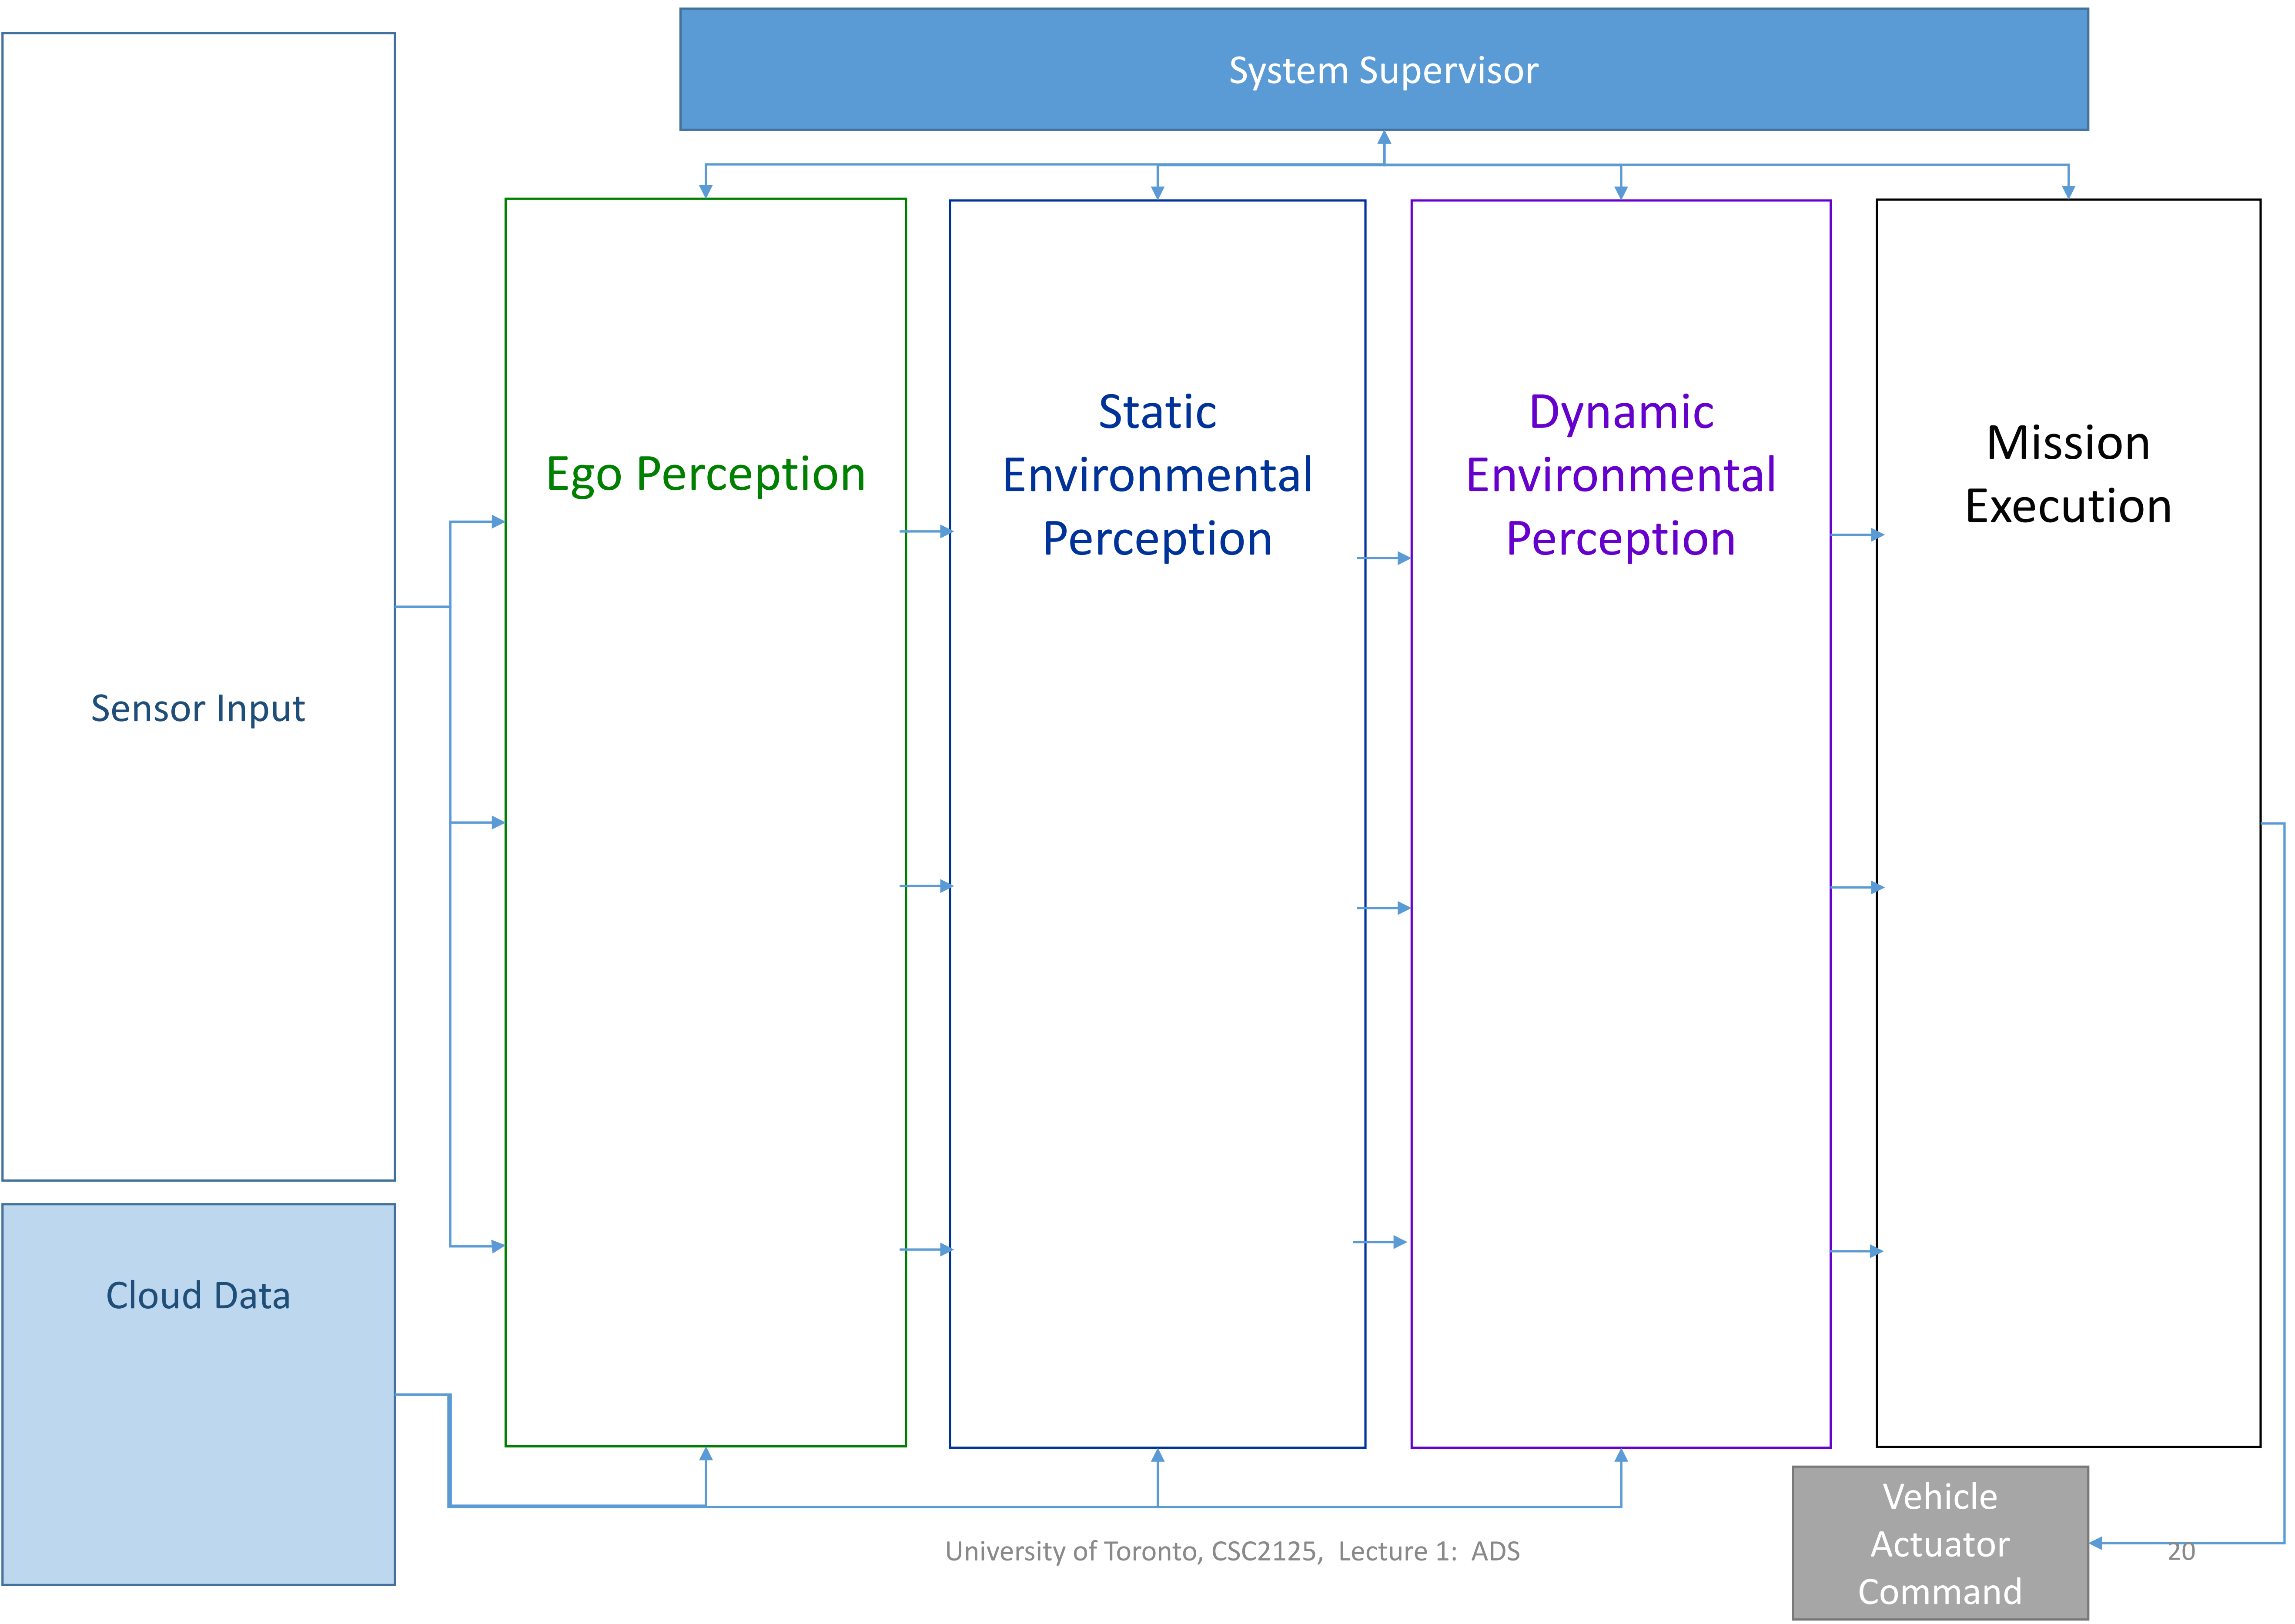

Functional Reference Architecture

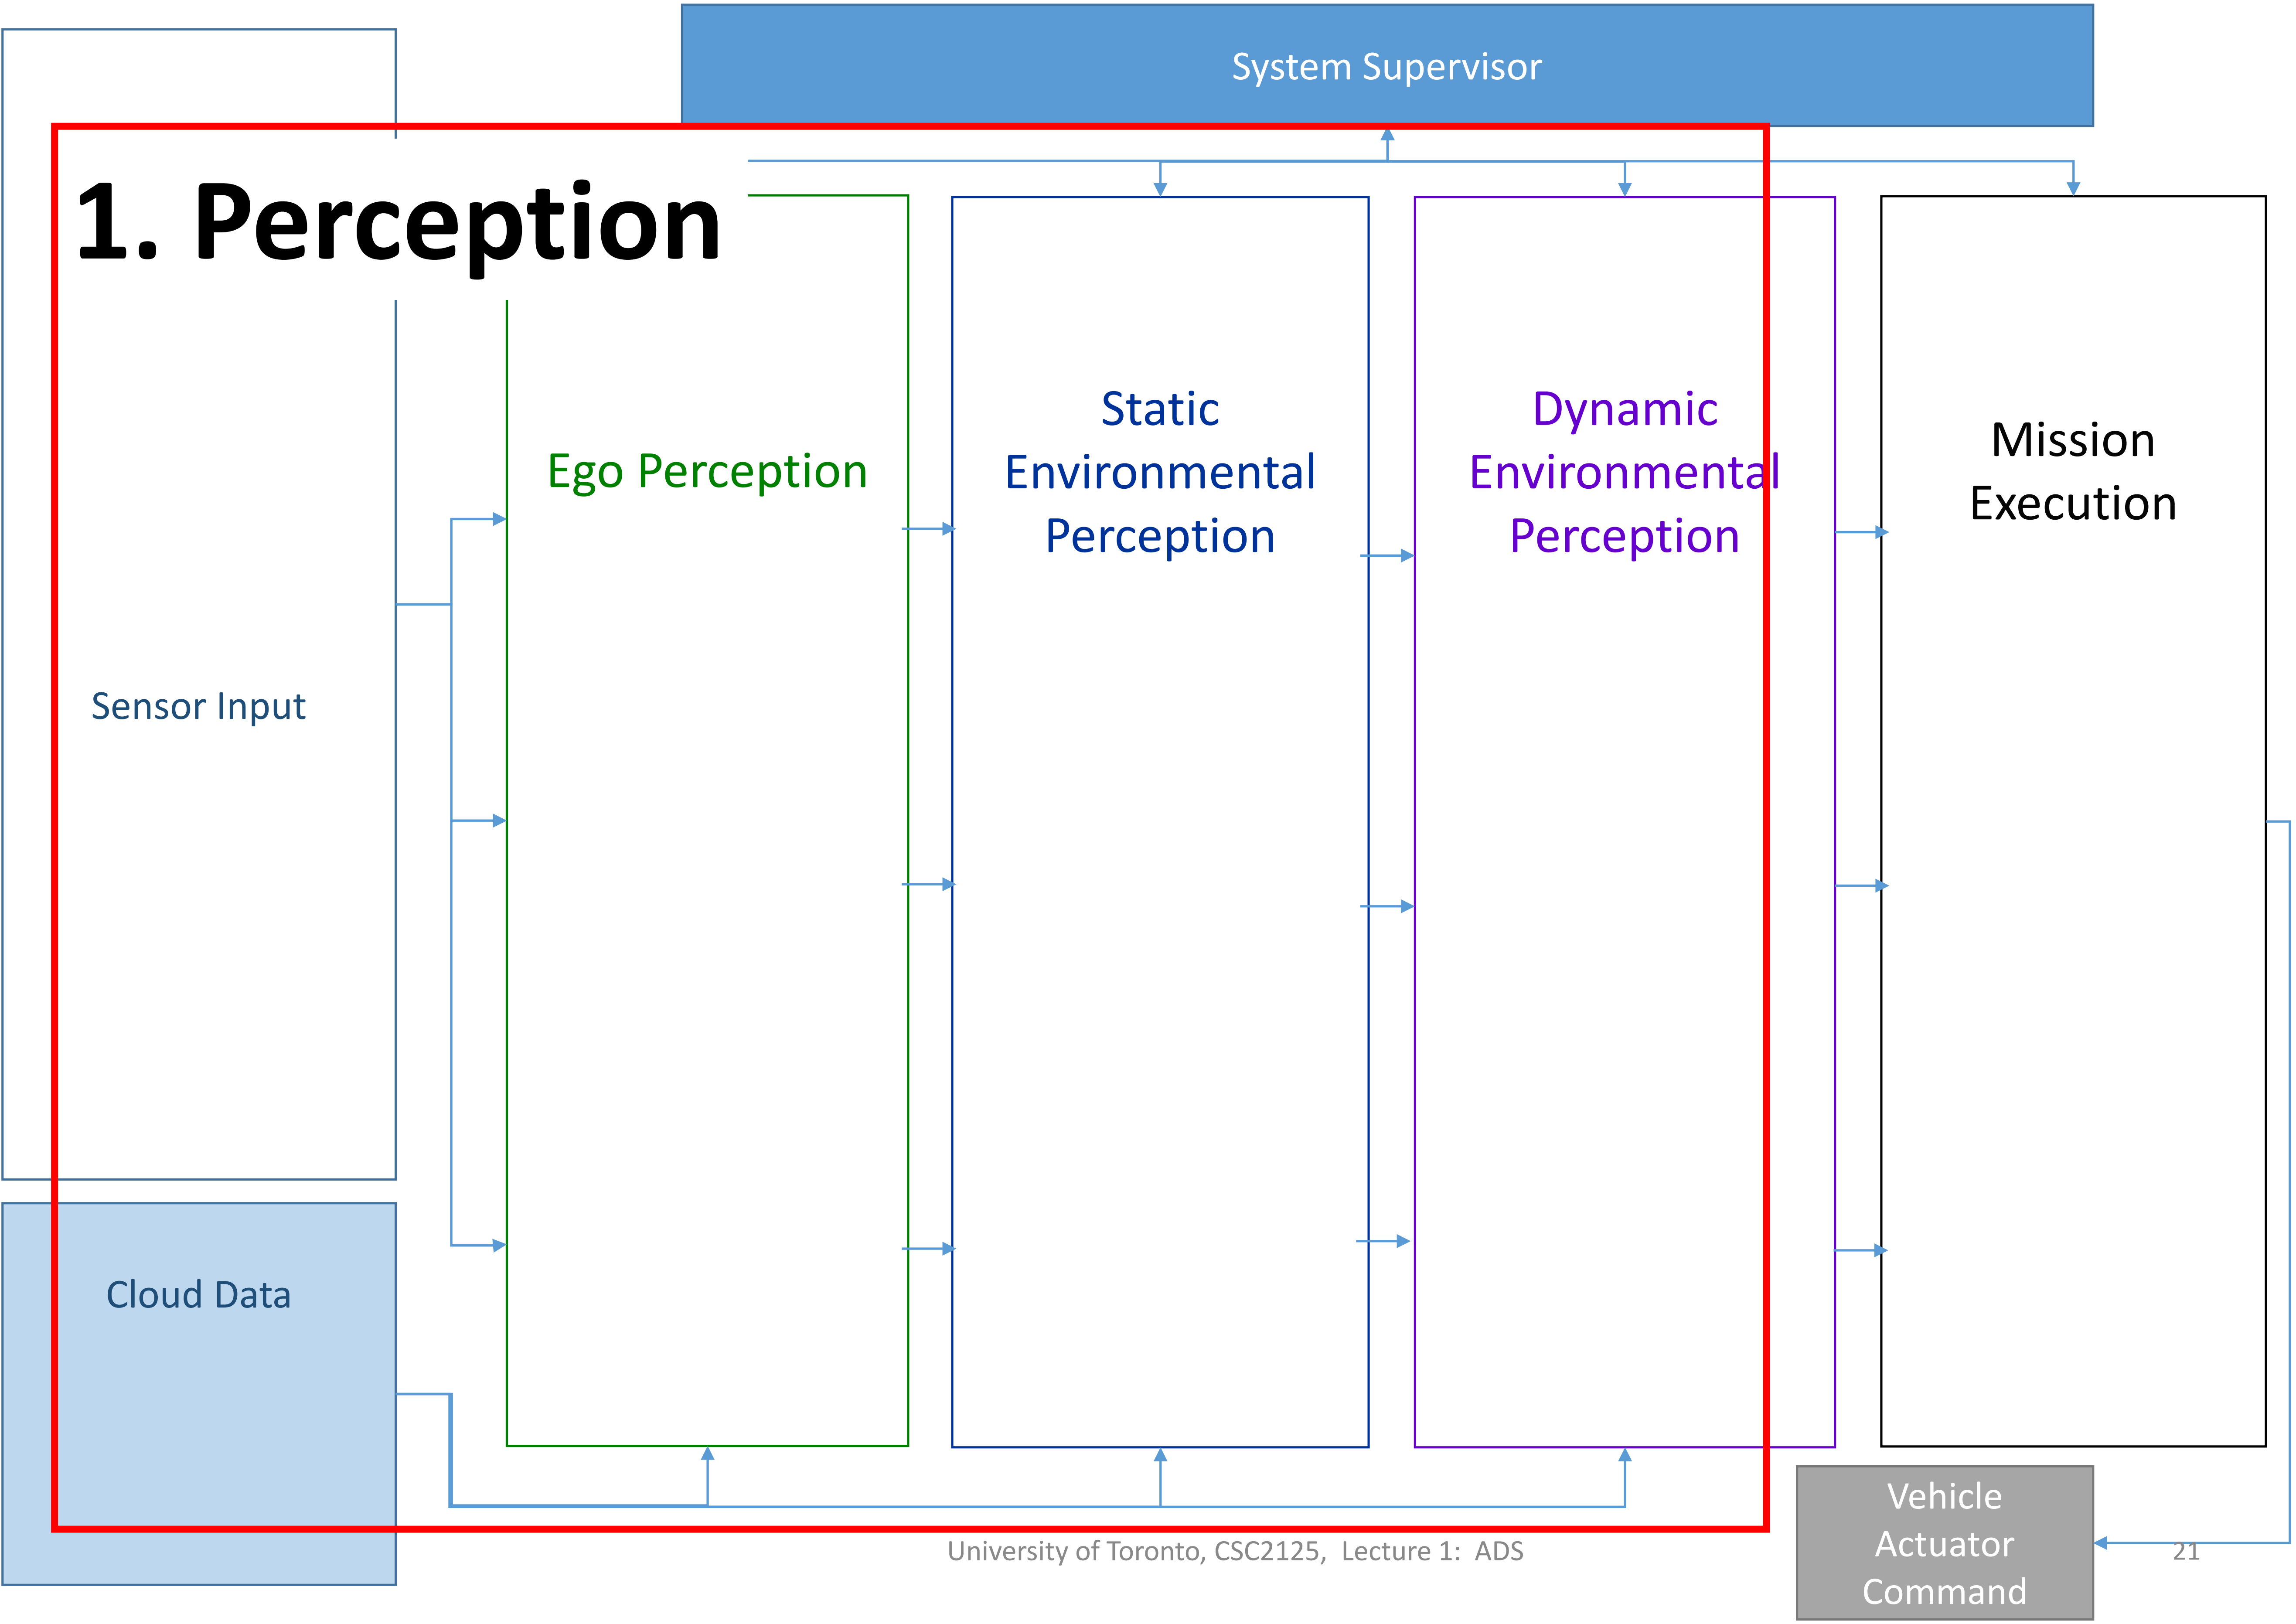

Functional Reference Architecture

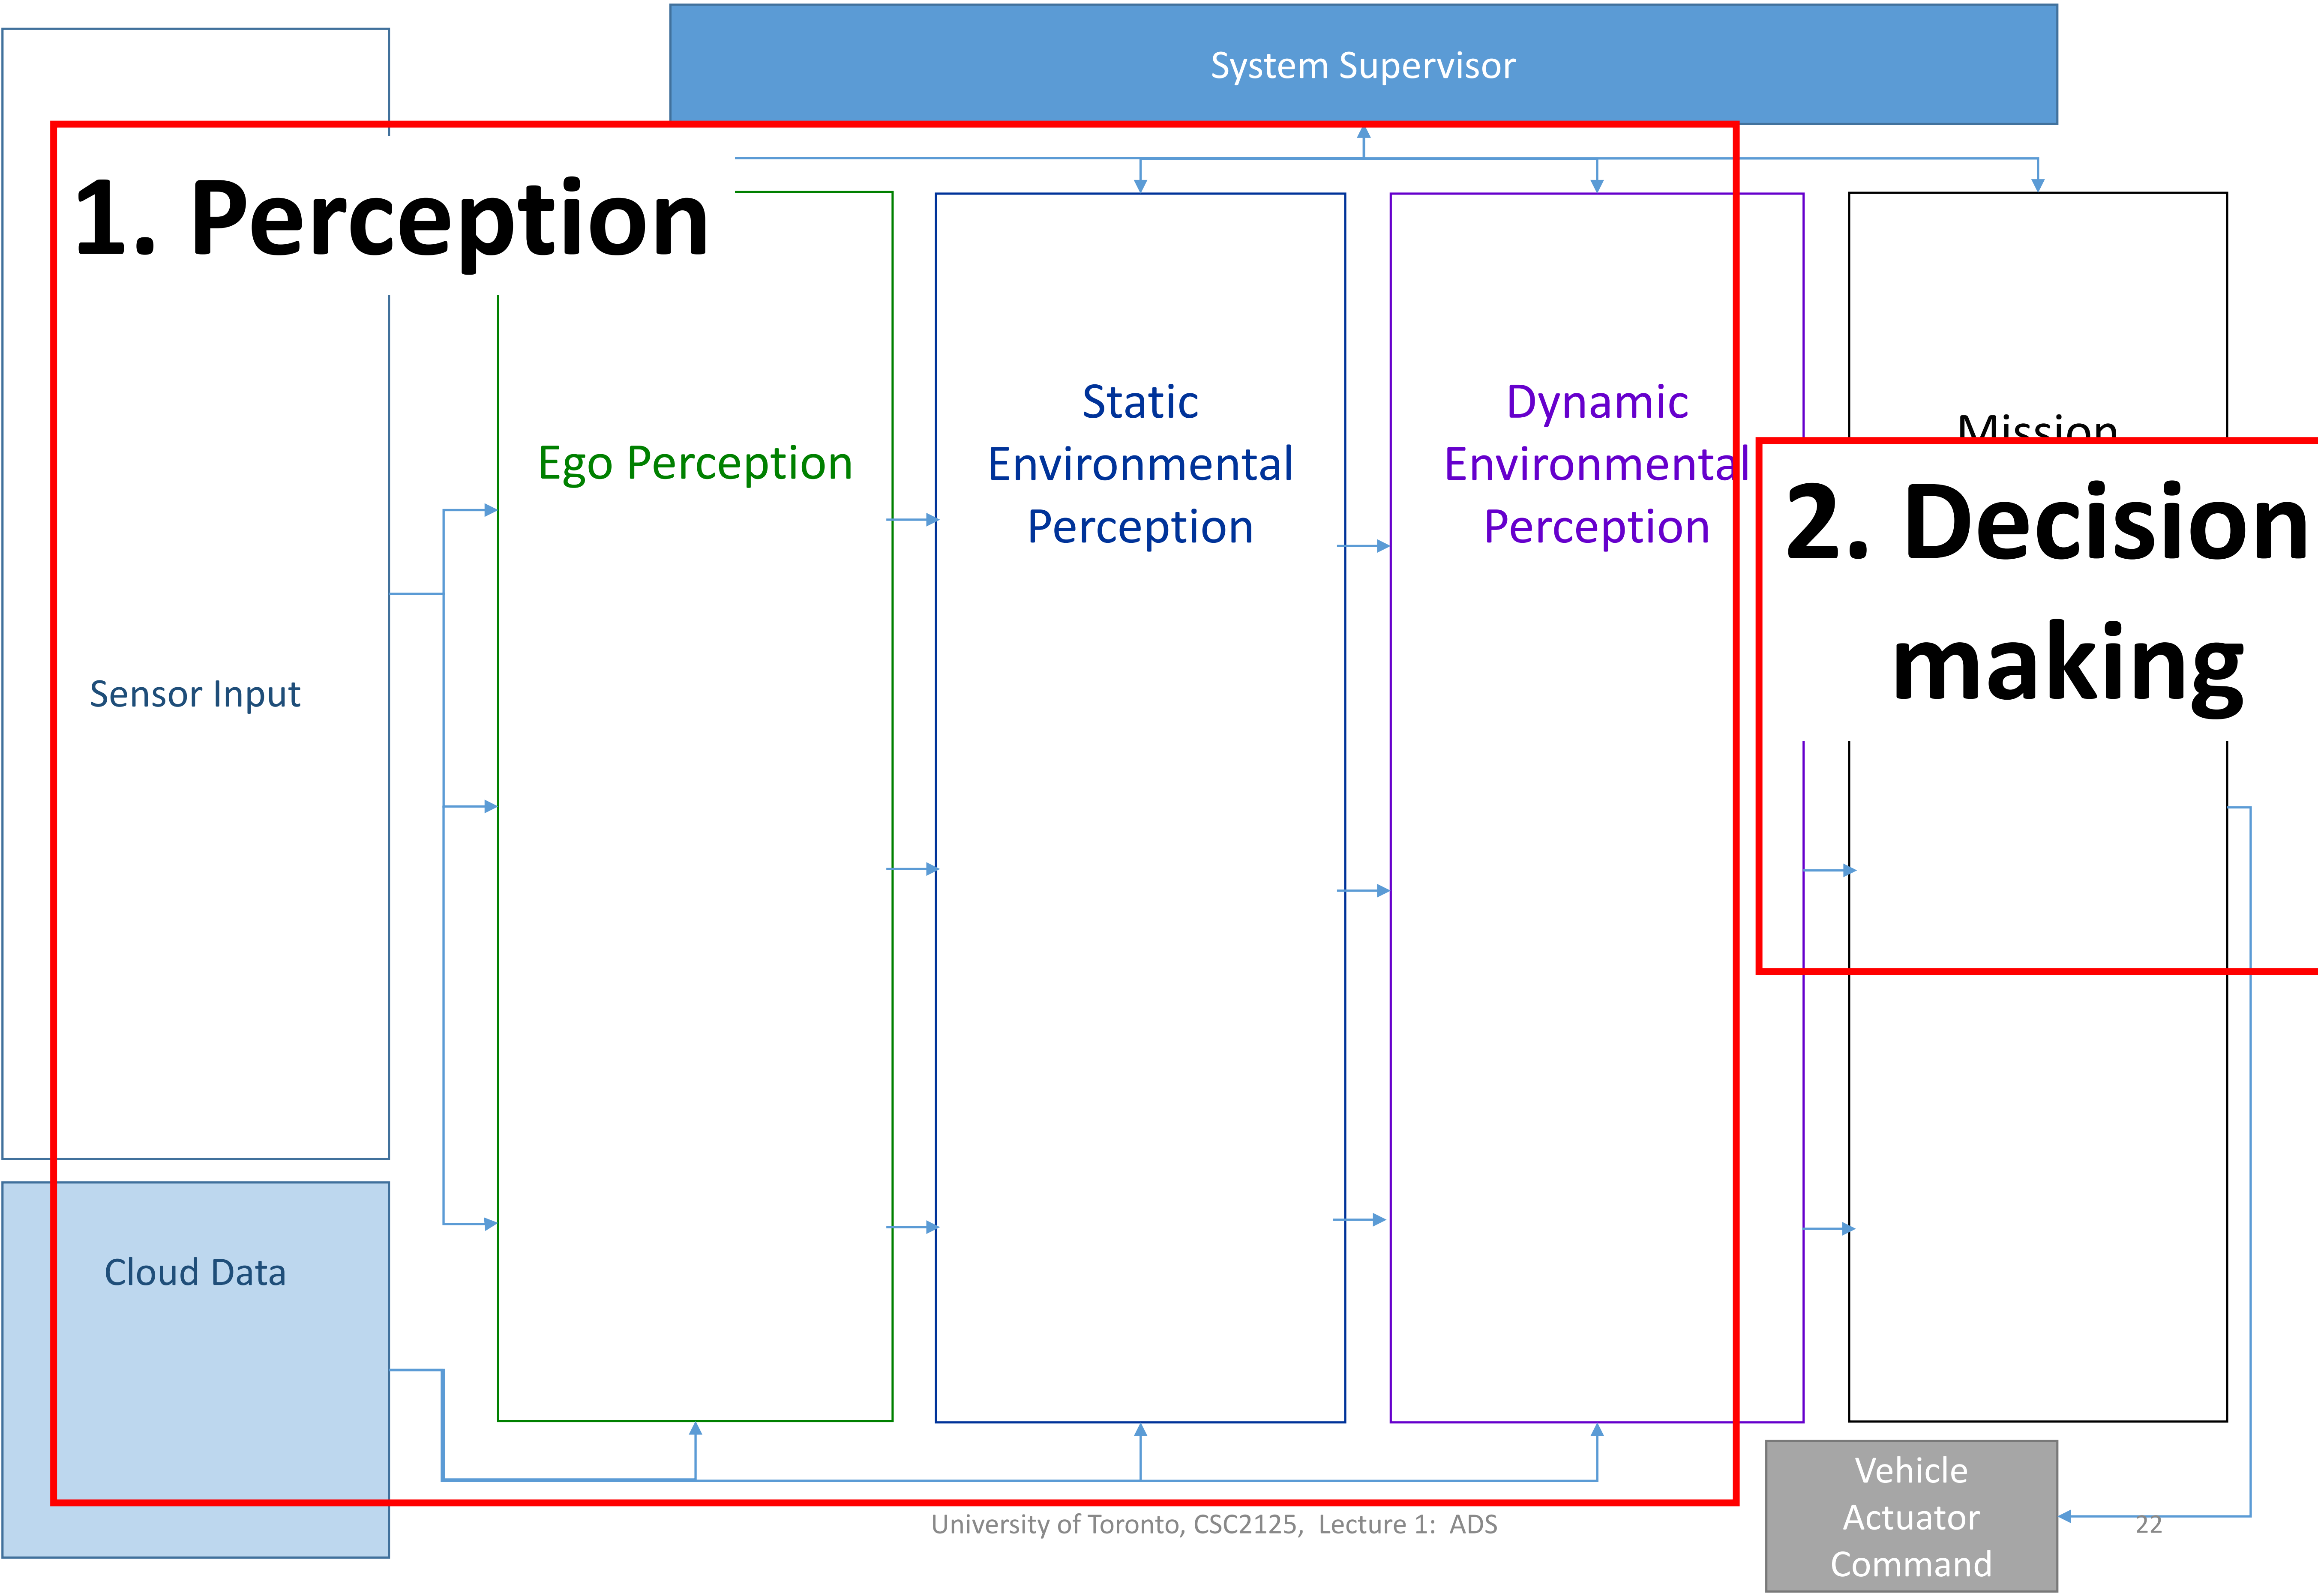

Functional Reference Architecture

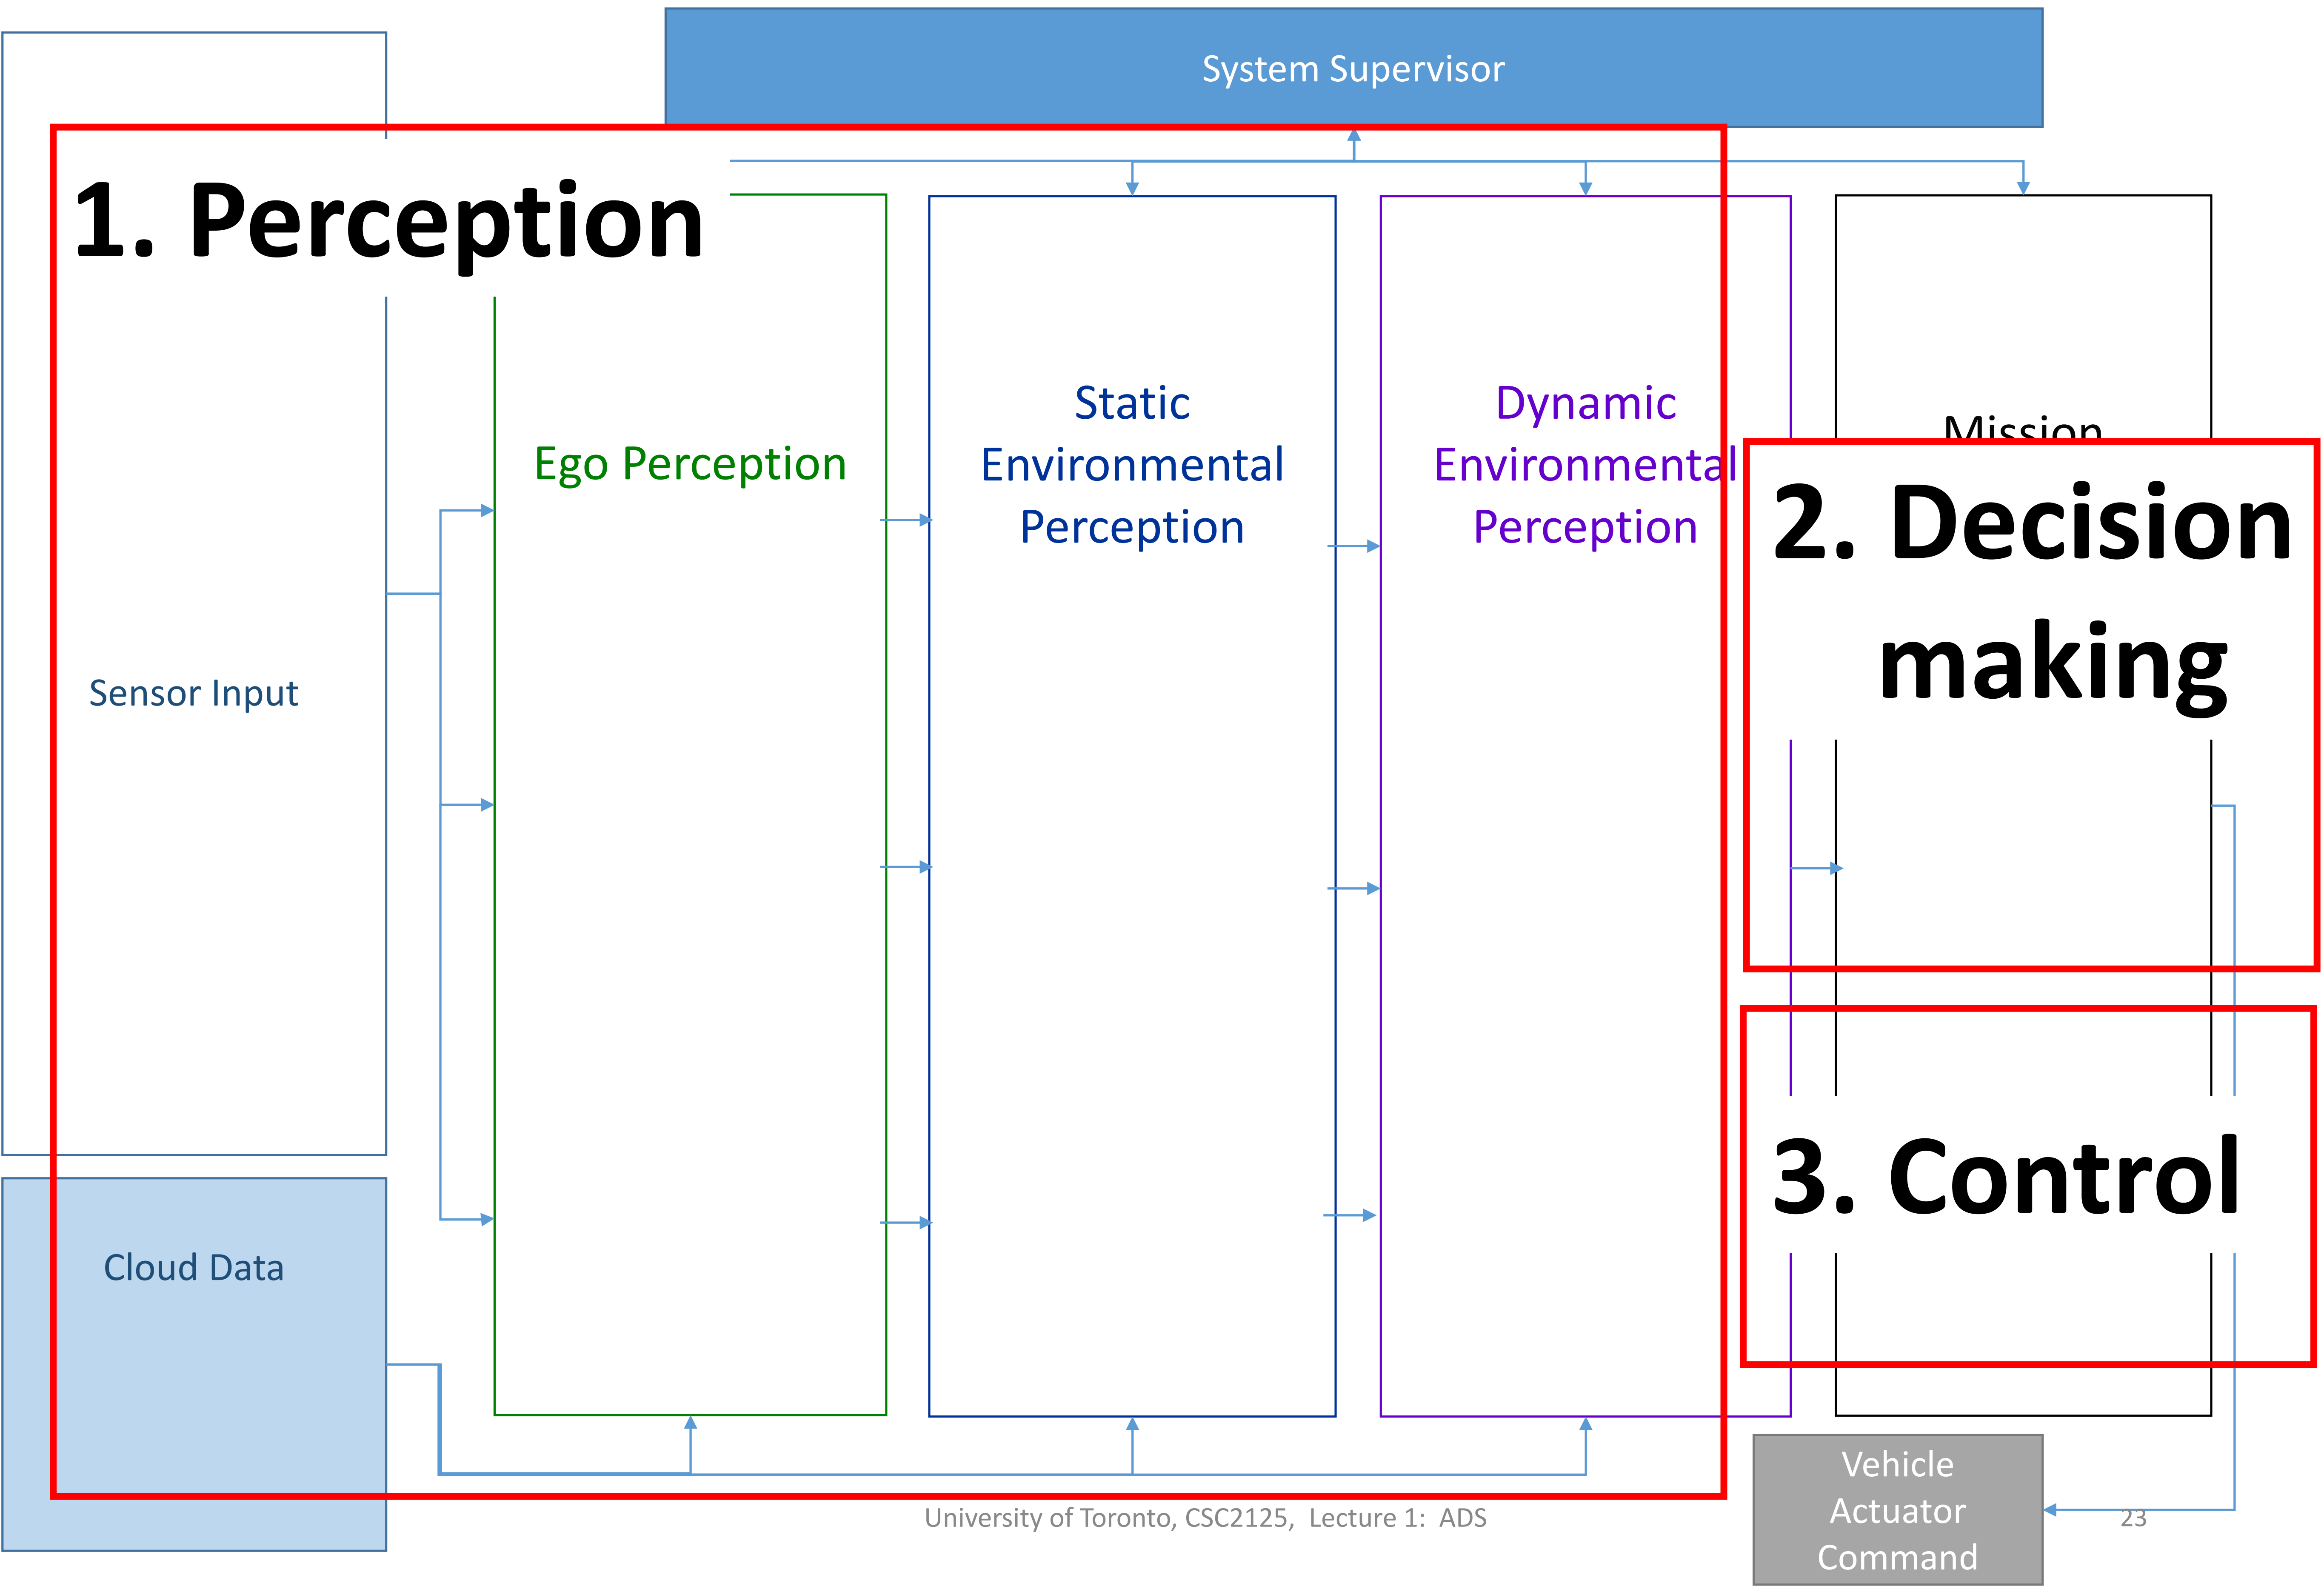

Functional Reference Architecture

System Supervisor

GPS

IMU

Wheel Odometry

Cameras

LIDAR

RADAR

Cloud Data

Maps

Traffic

Weather

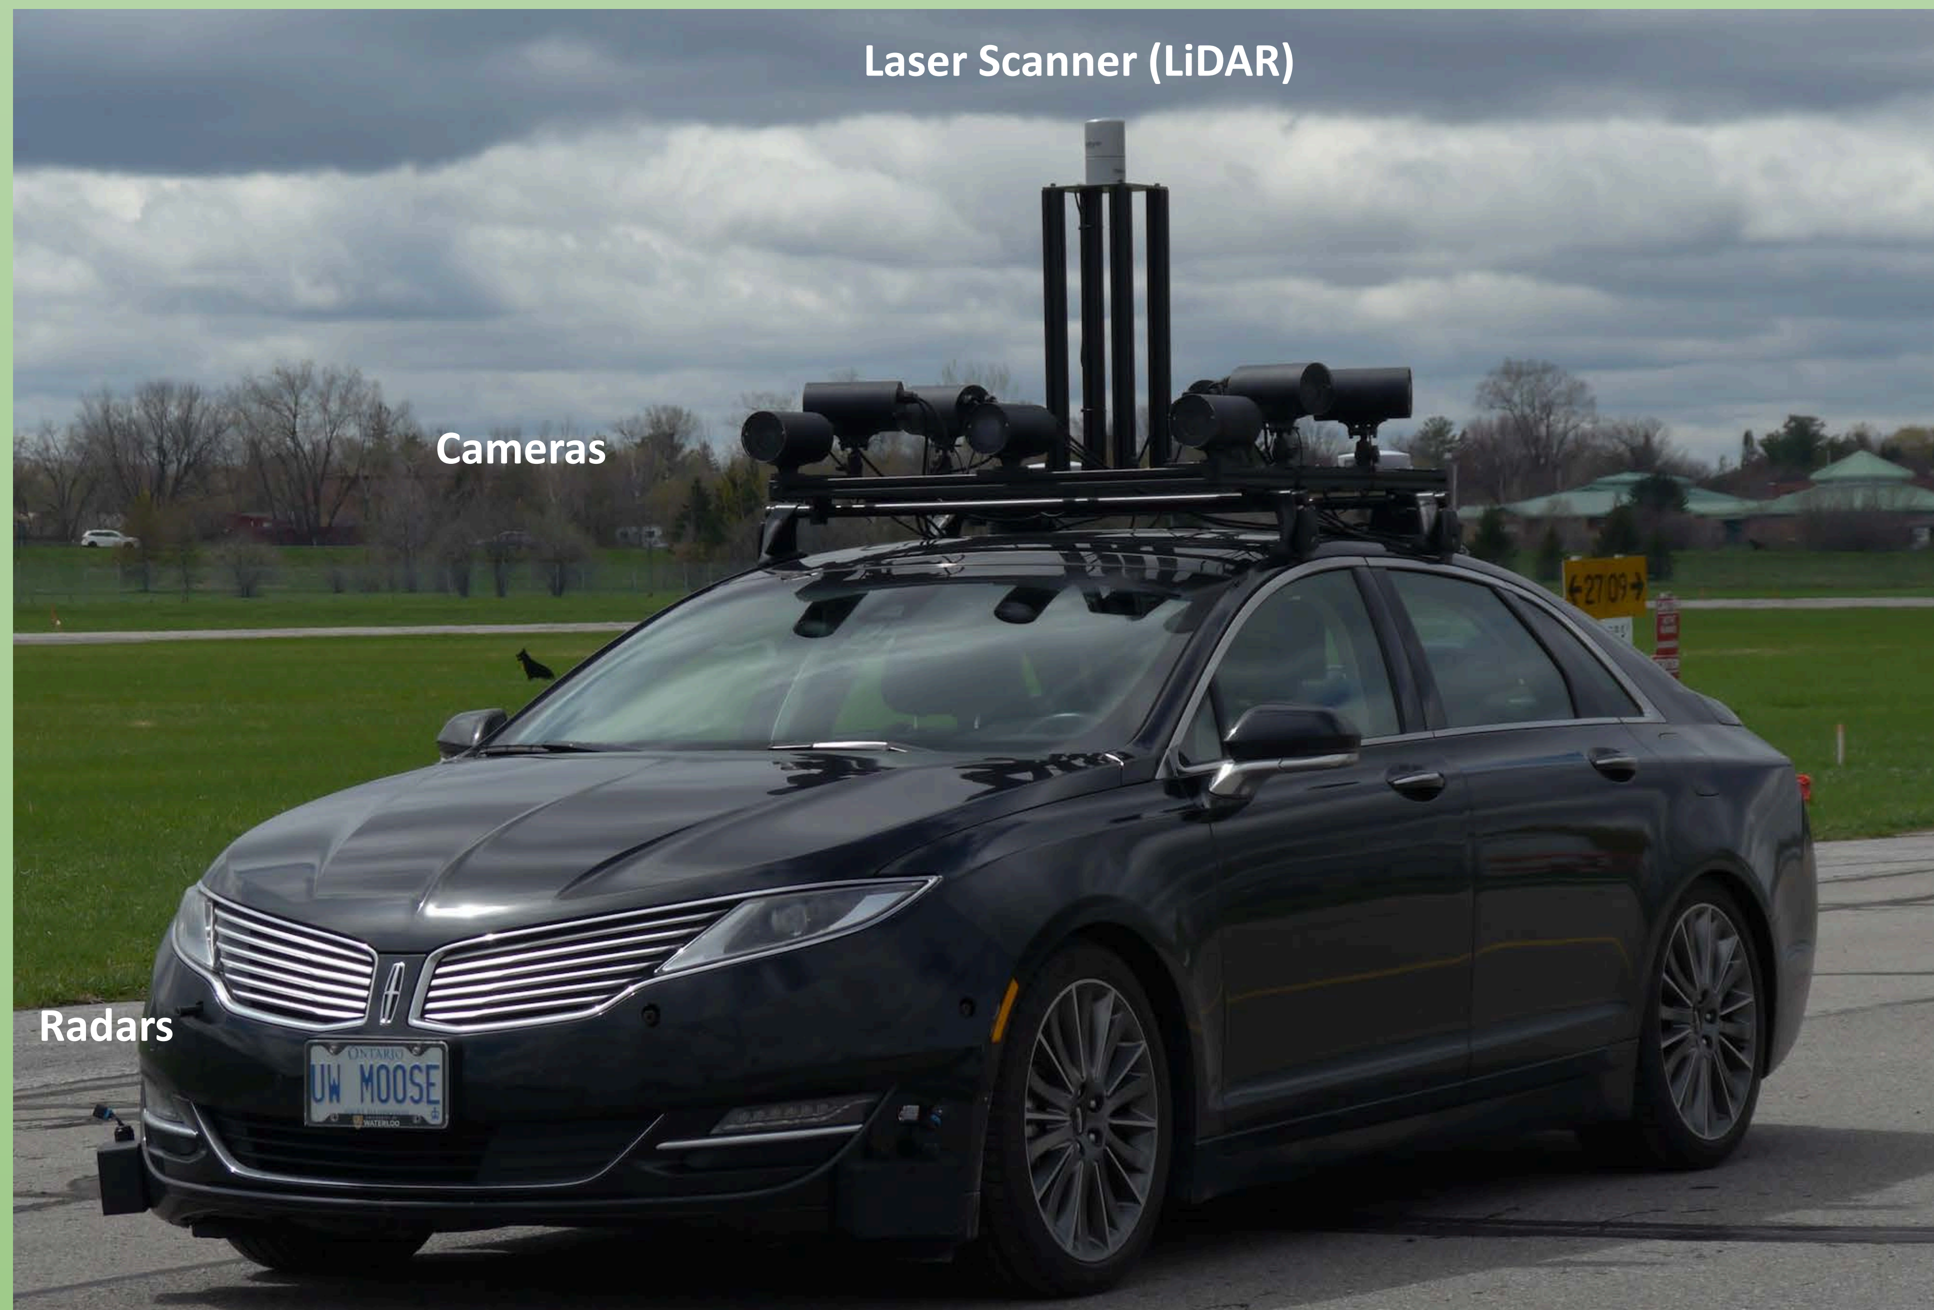

Vehicle Actuator Command

Functional Reference Architecture

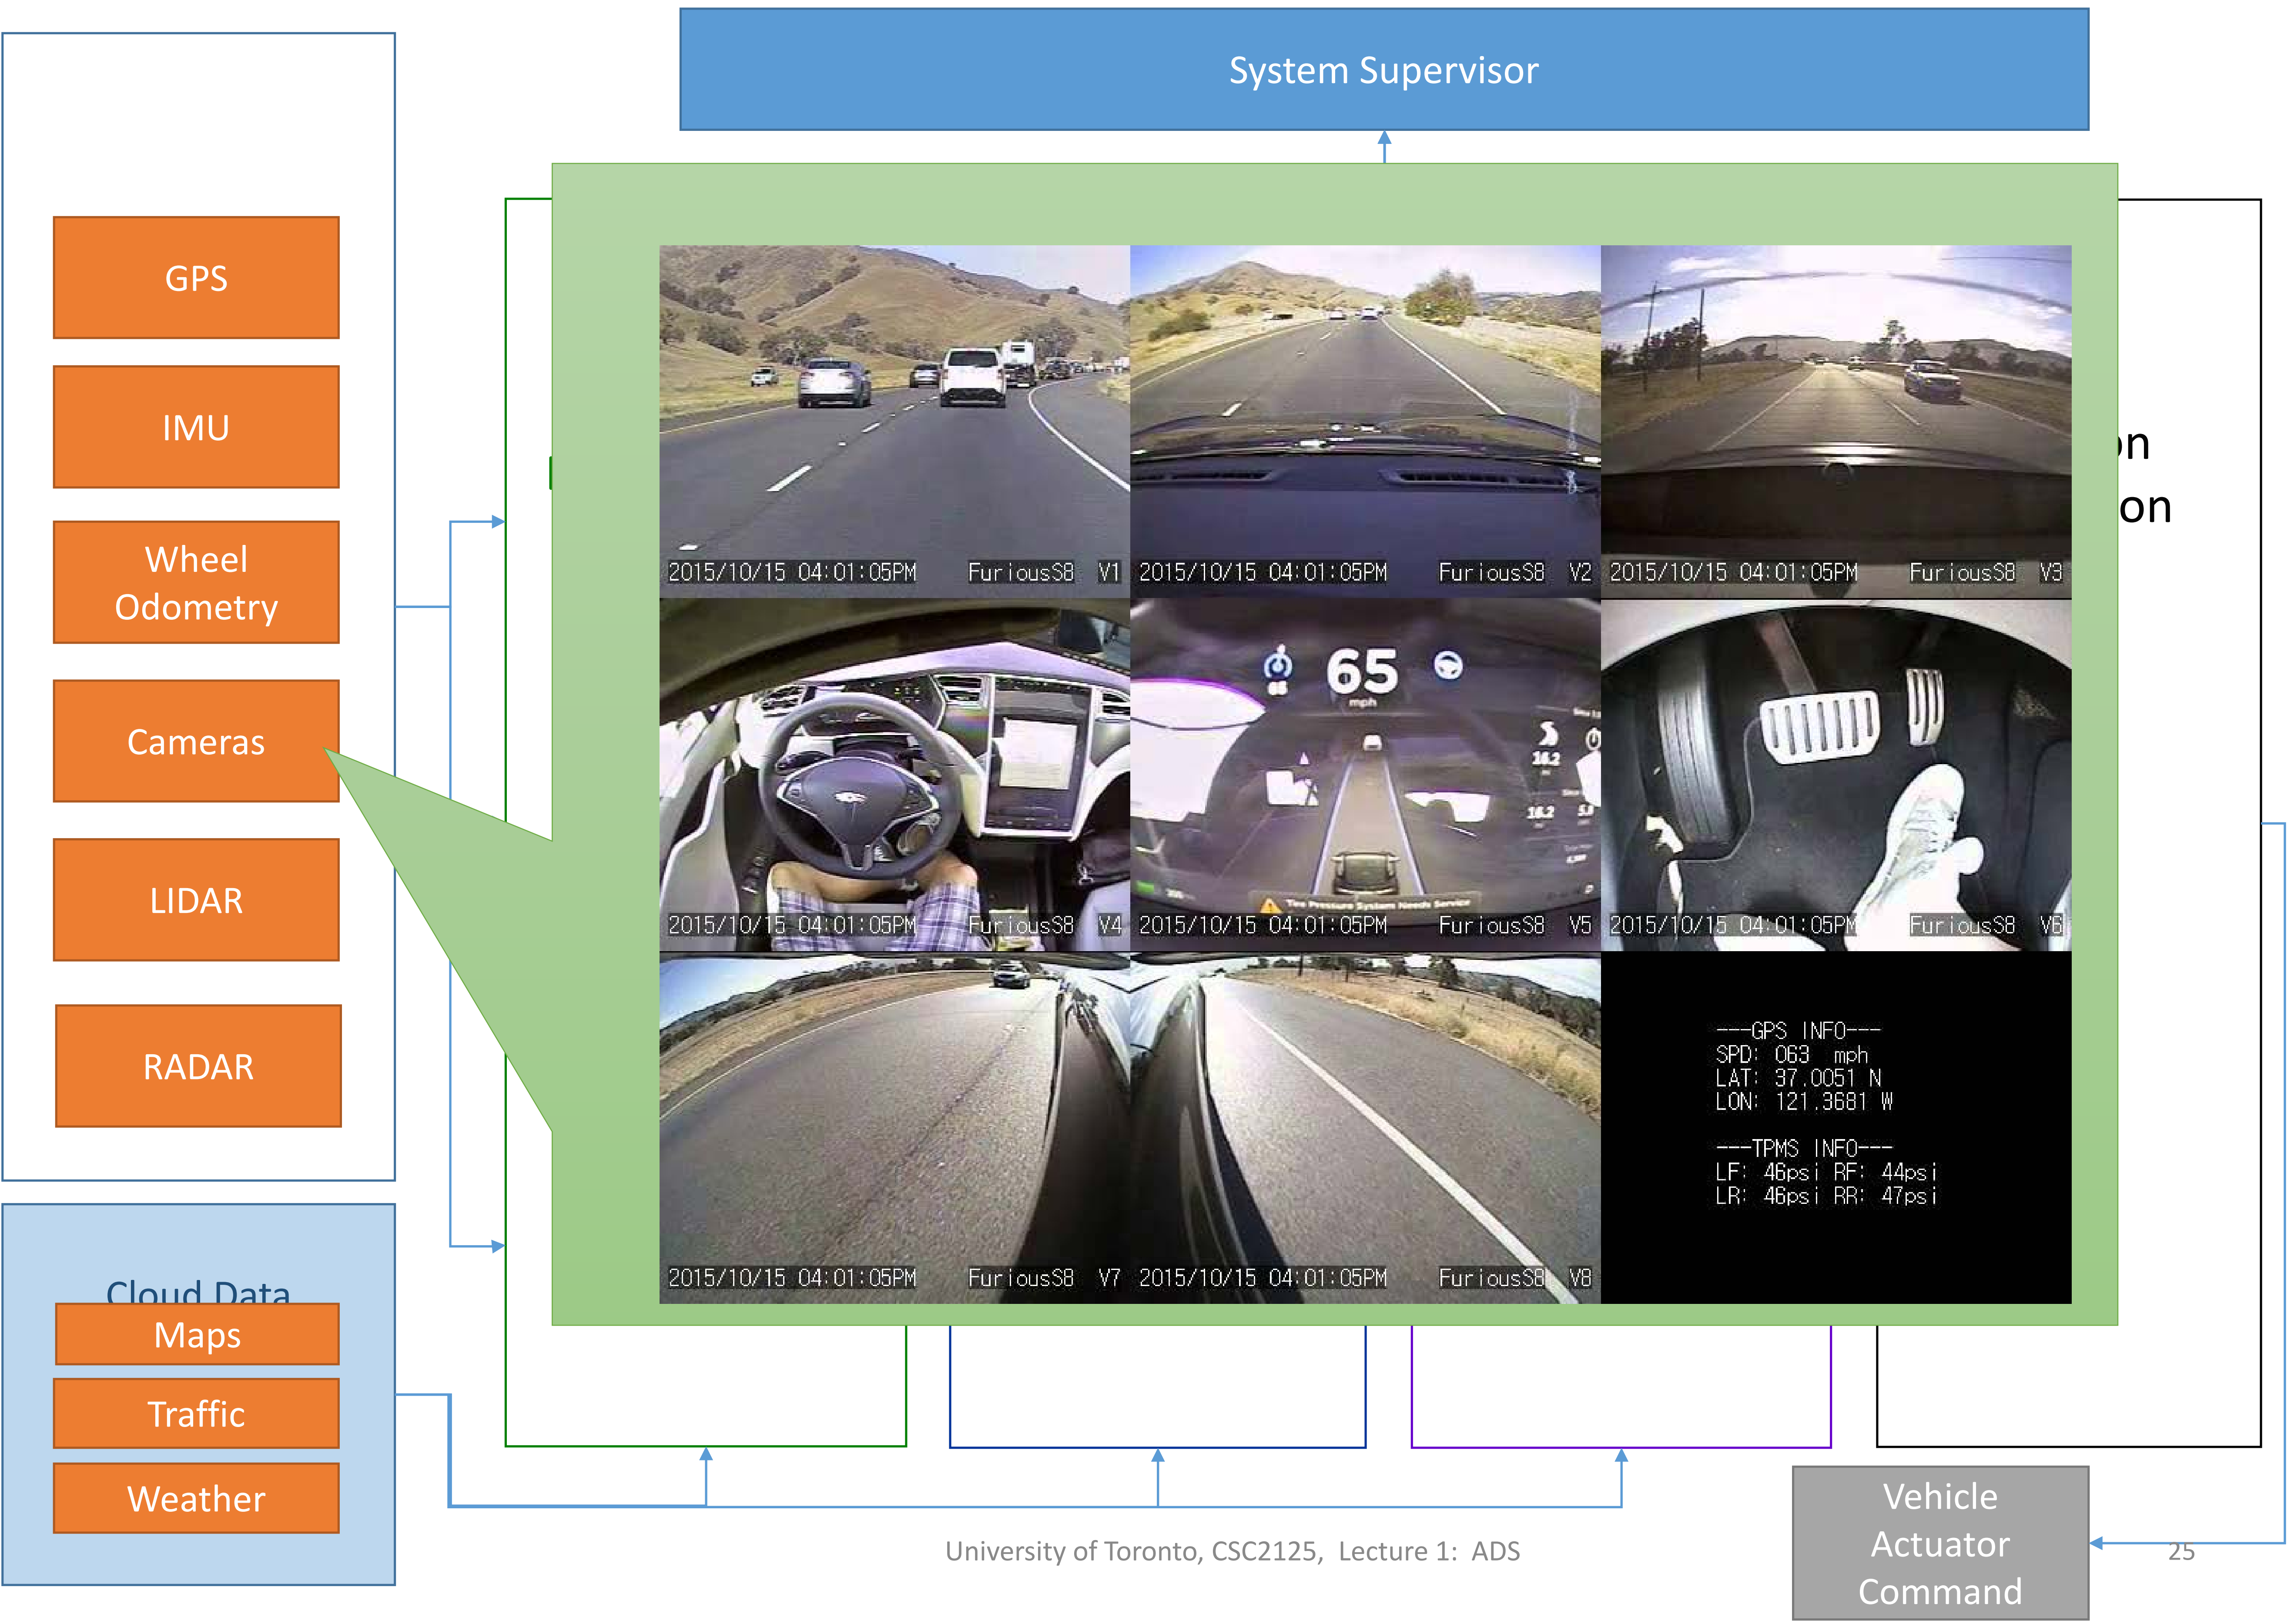

Functional Reference Architecture

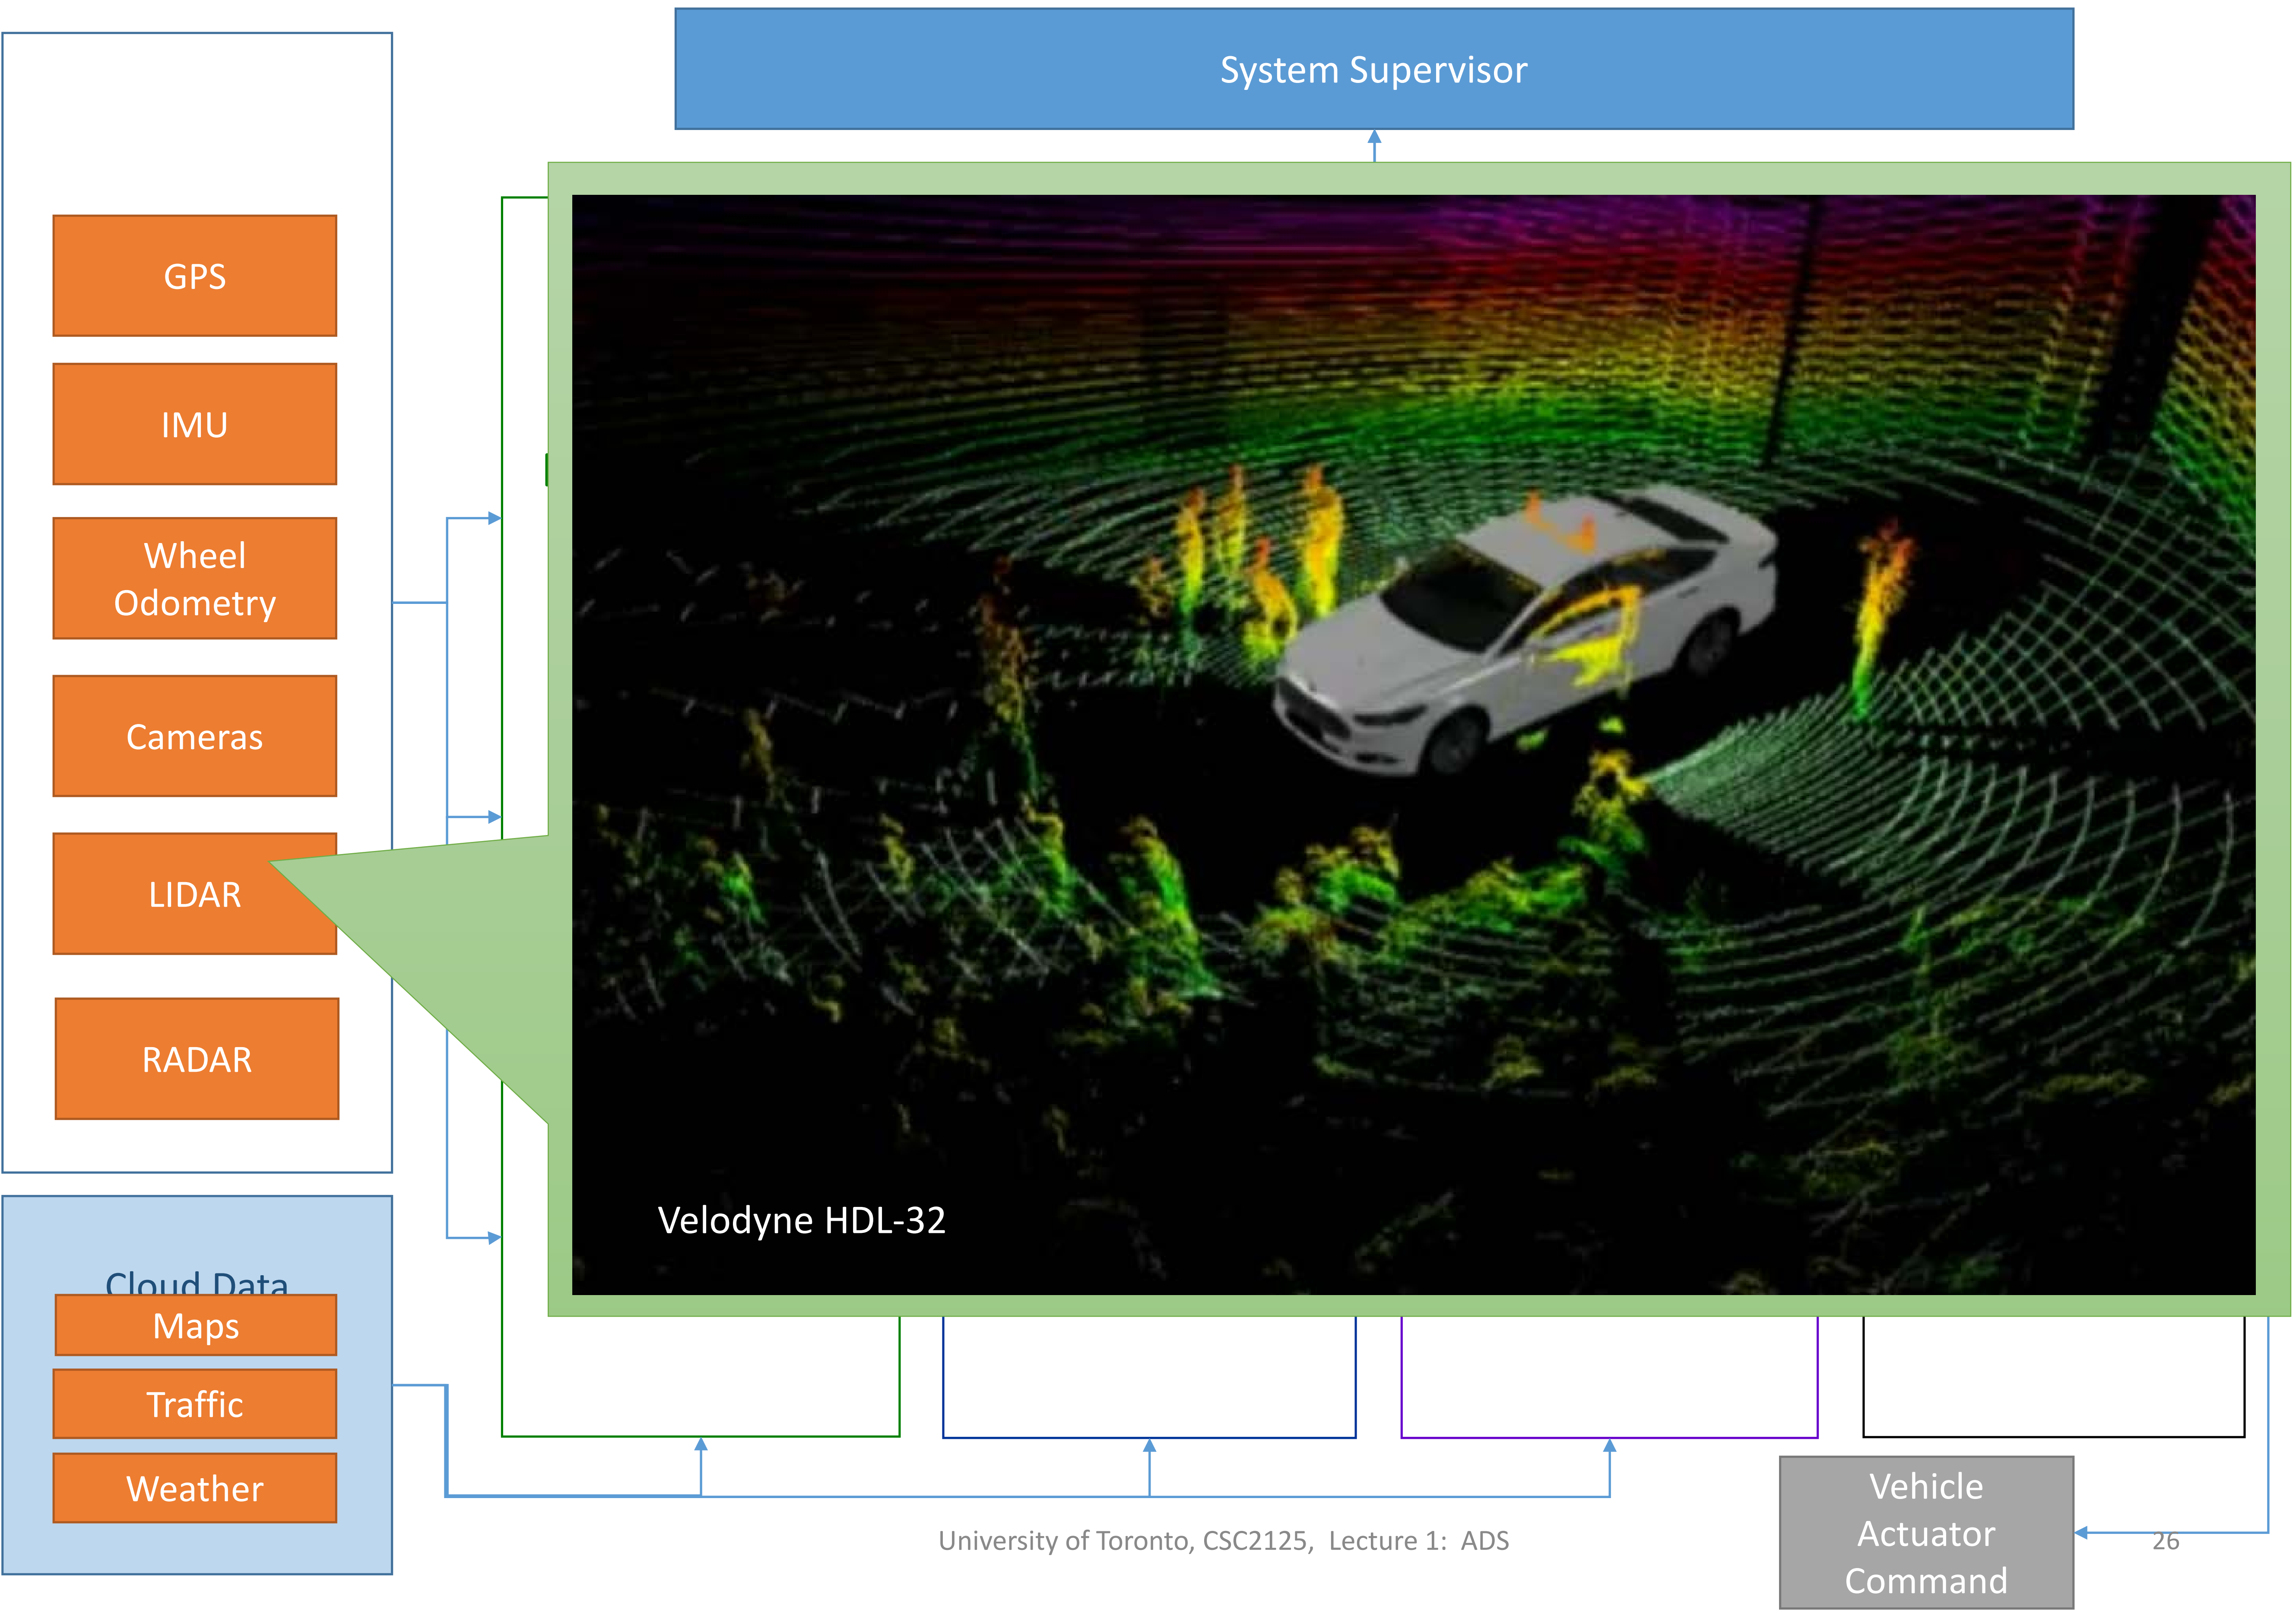

Functional Reference Architecture

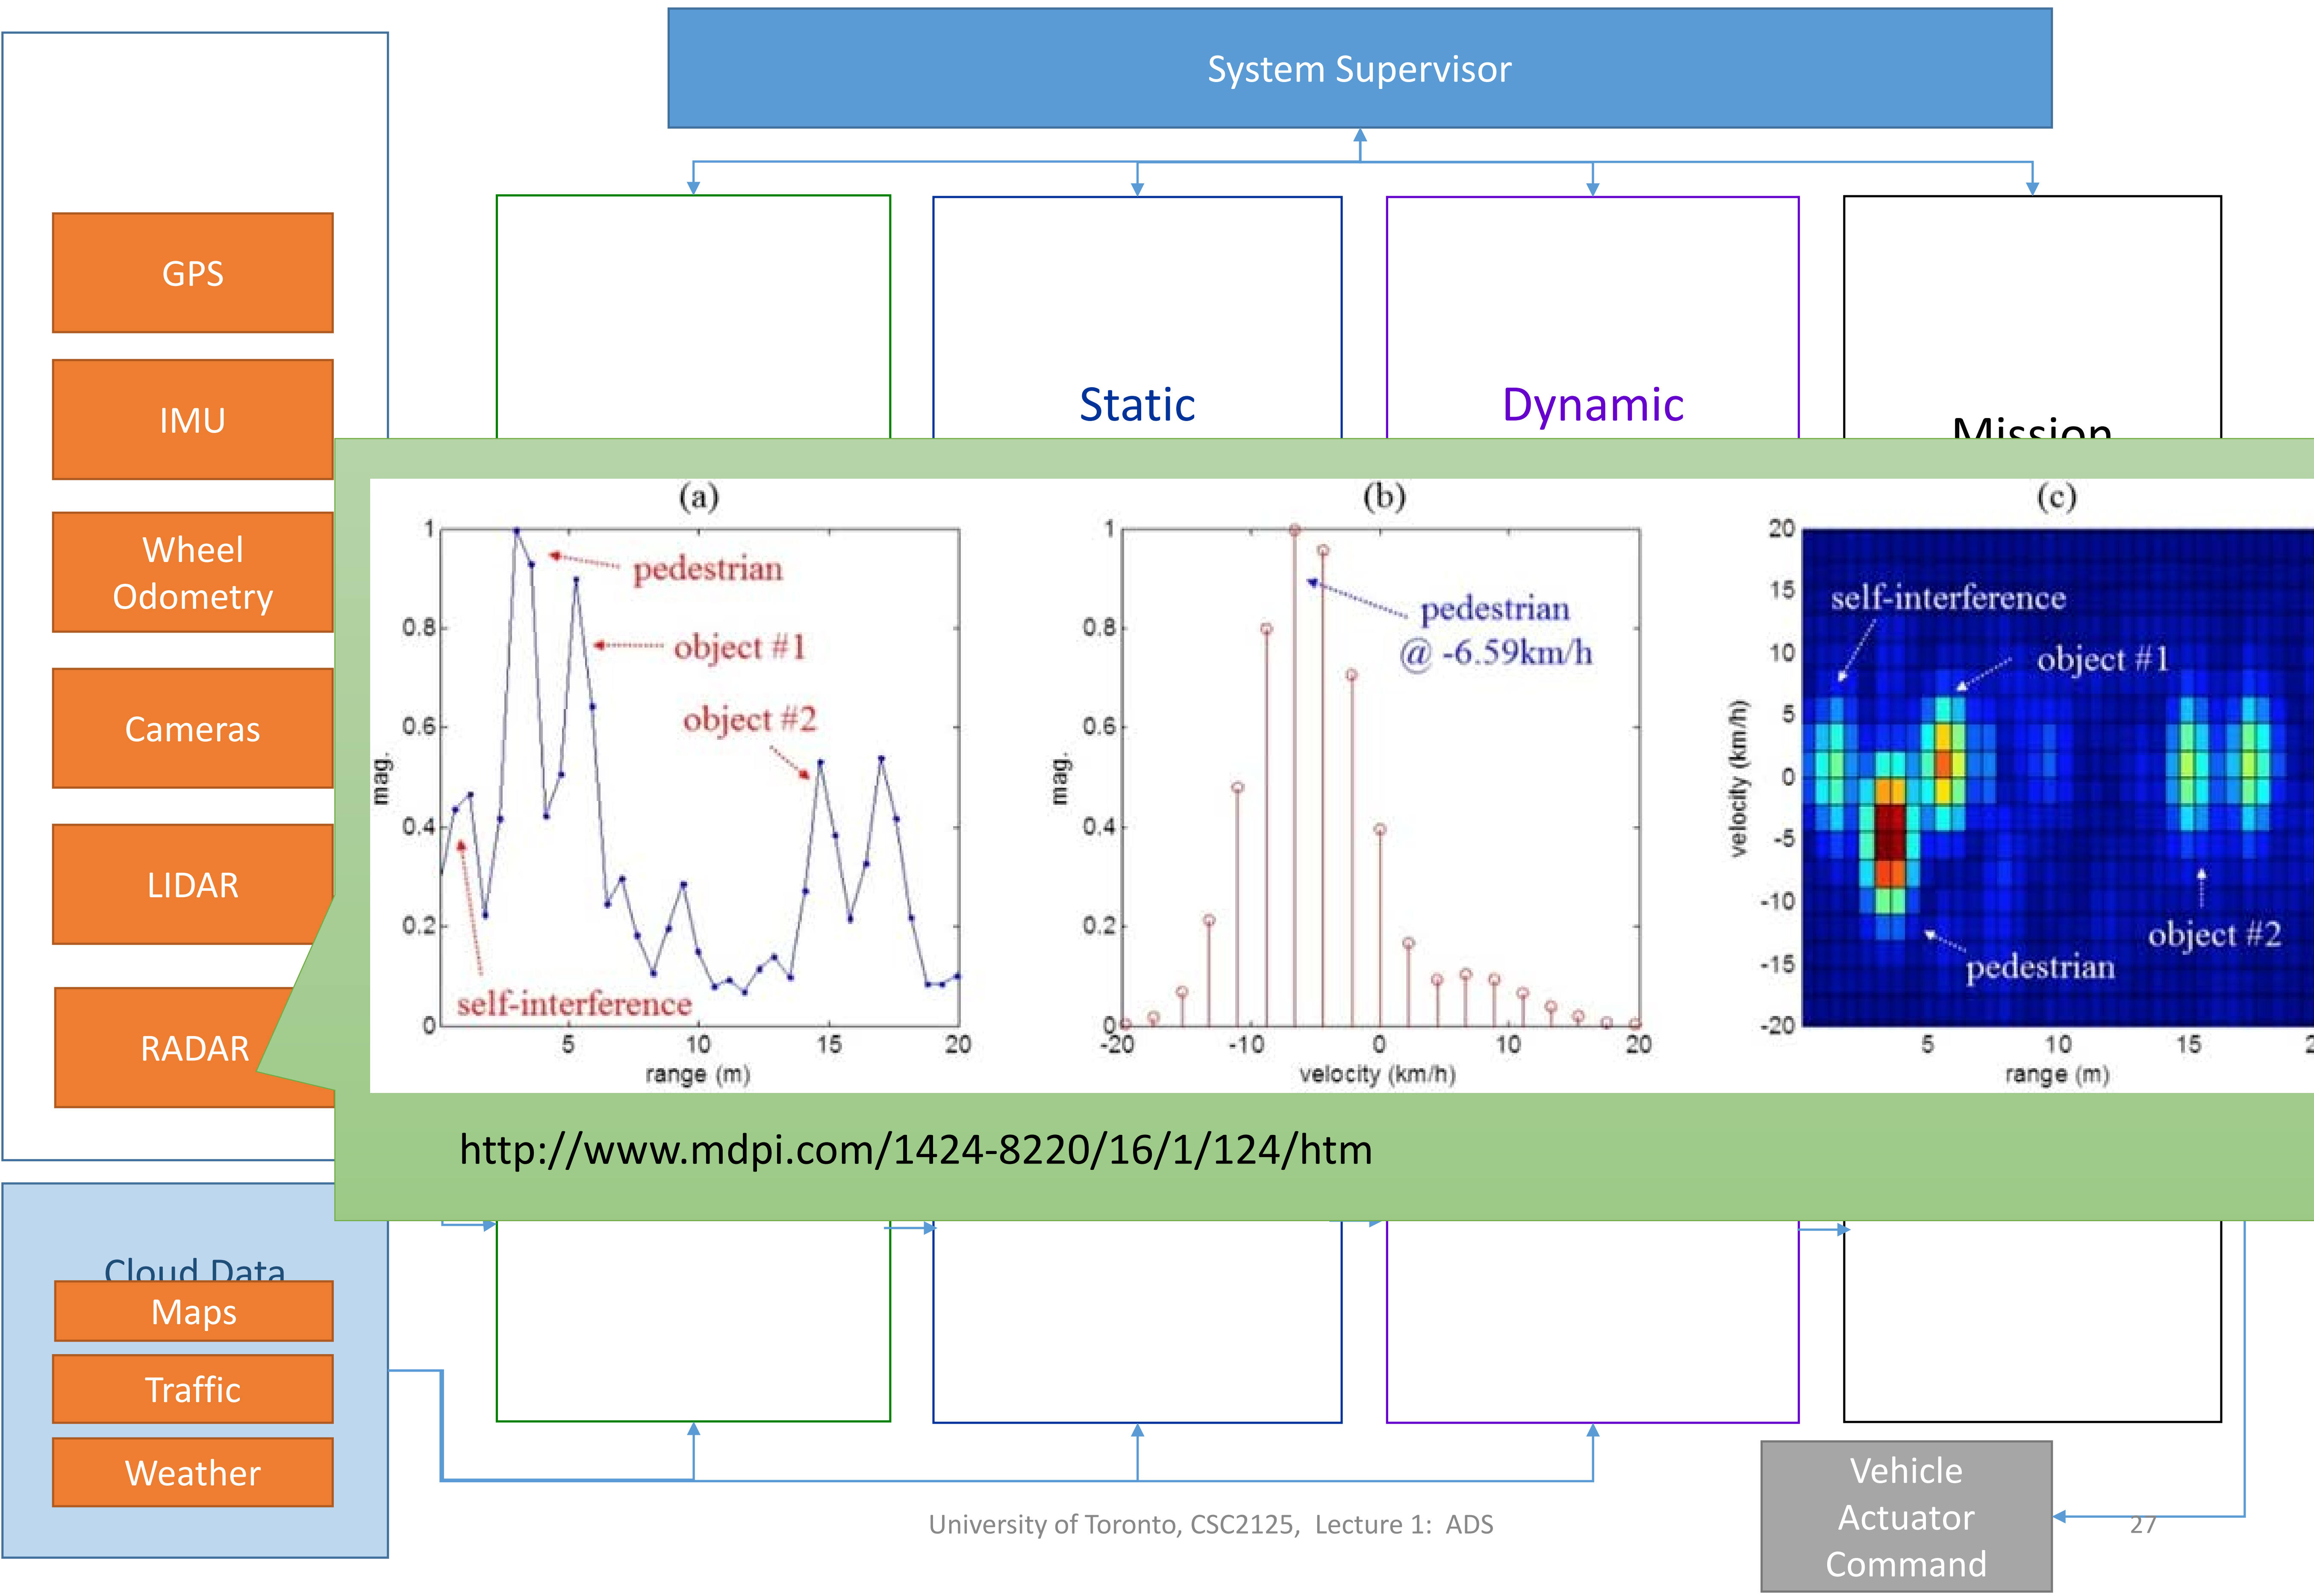

Radar

- Cheap
- Does well in extreme weather
- Low resolution
- Most used automotive sensor for object detection and tracking

LIDAR

- Expensive
- Extremely accurate depth information
- Resolution much higher than radar
- 360 degrees of visibility

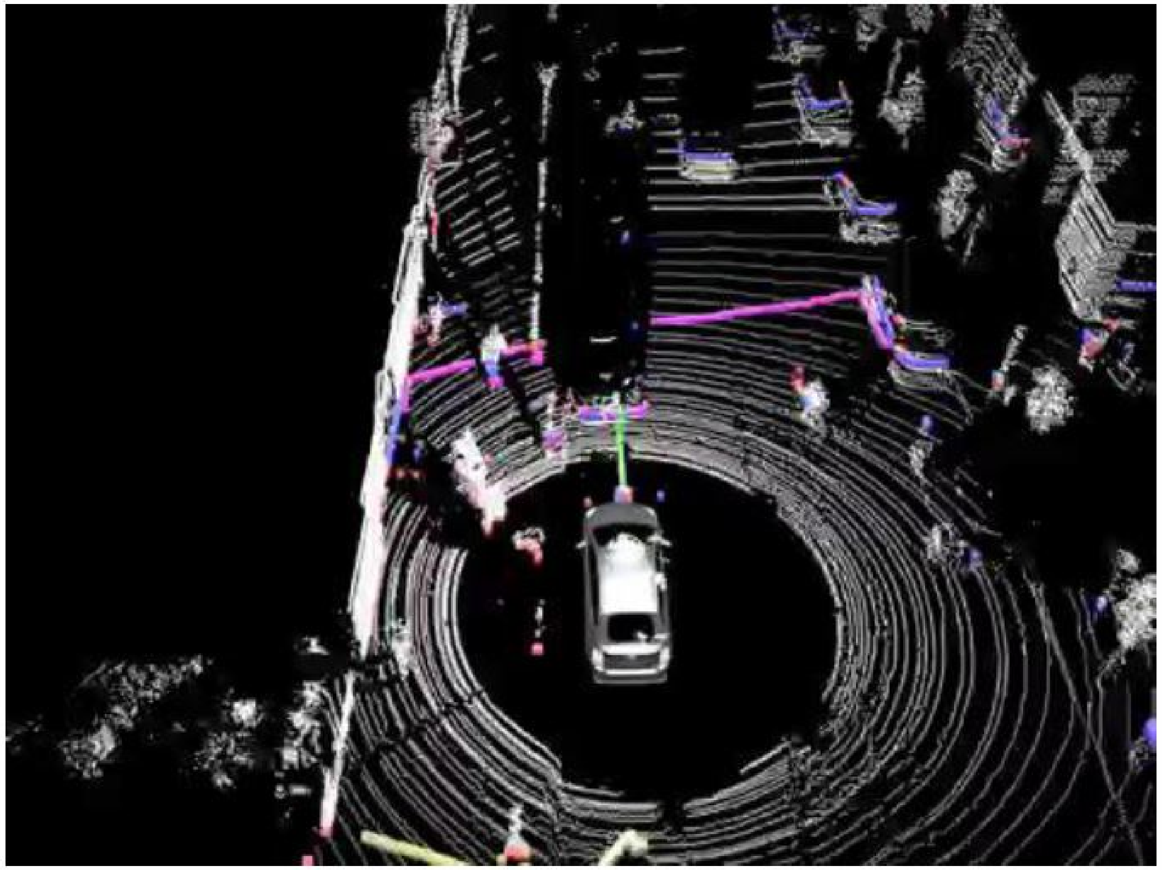

Camera

- Cheap
- Highest resolution
- Huge data = deep learning
- Human brains use similar sensor technology for driving
- Bad at depth estimation
- Not good in extreme weather

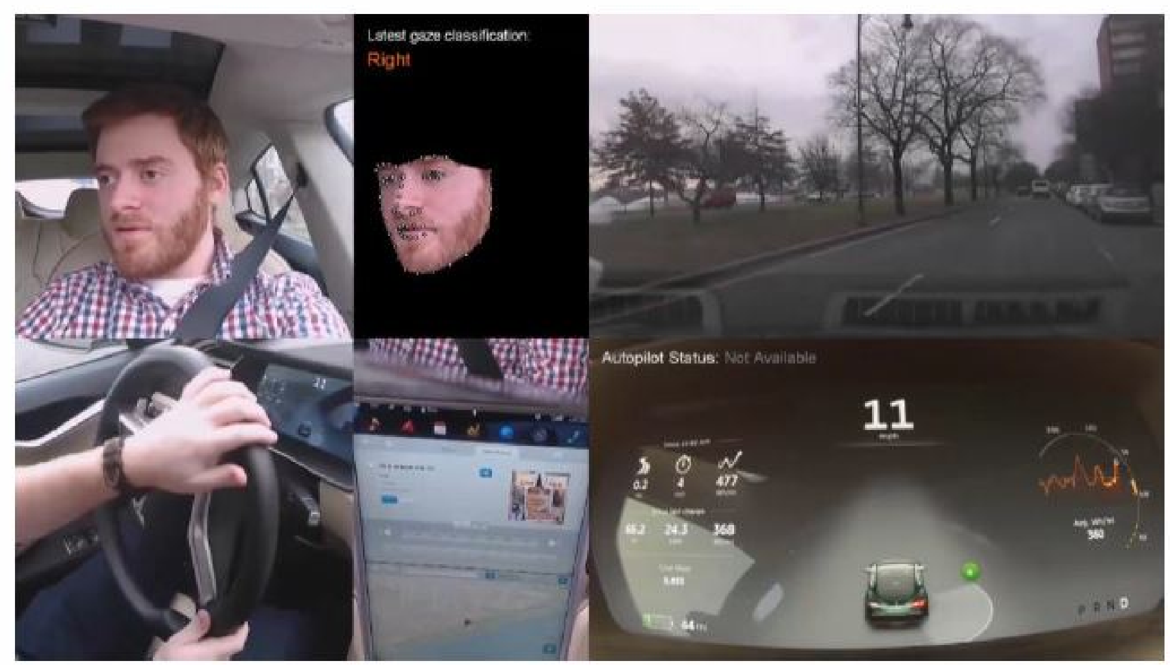

Comparisons

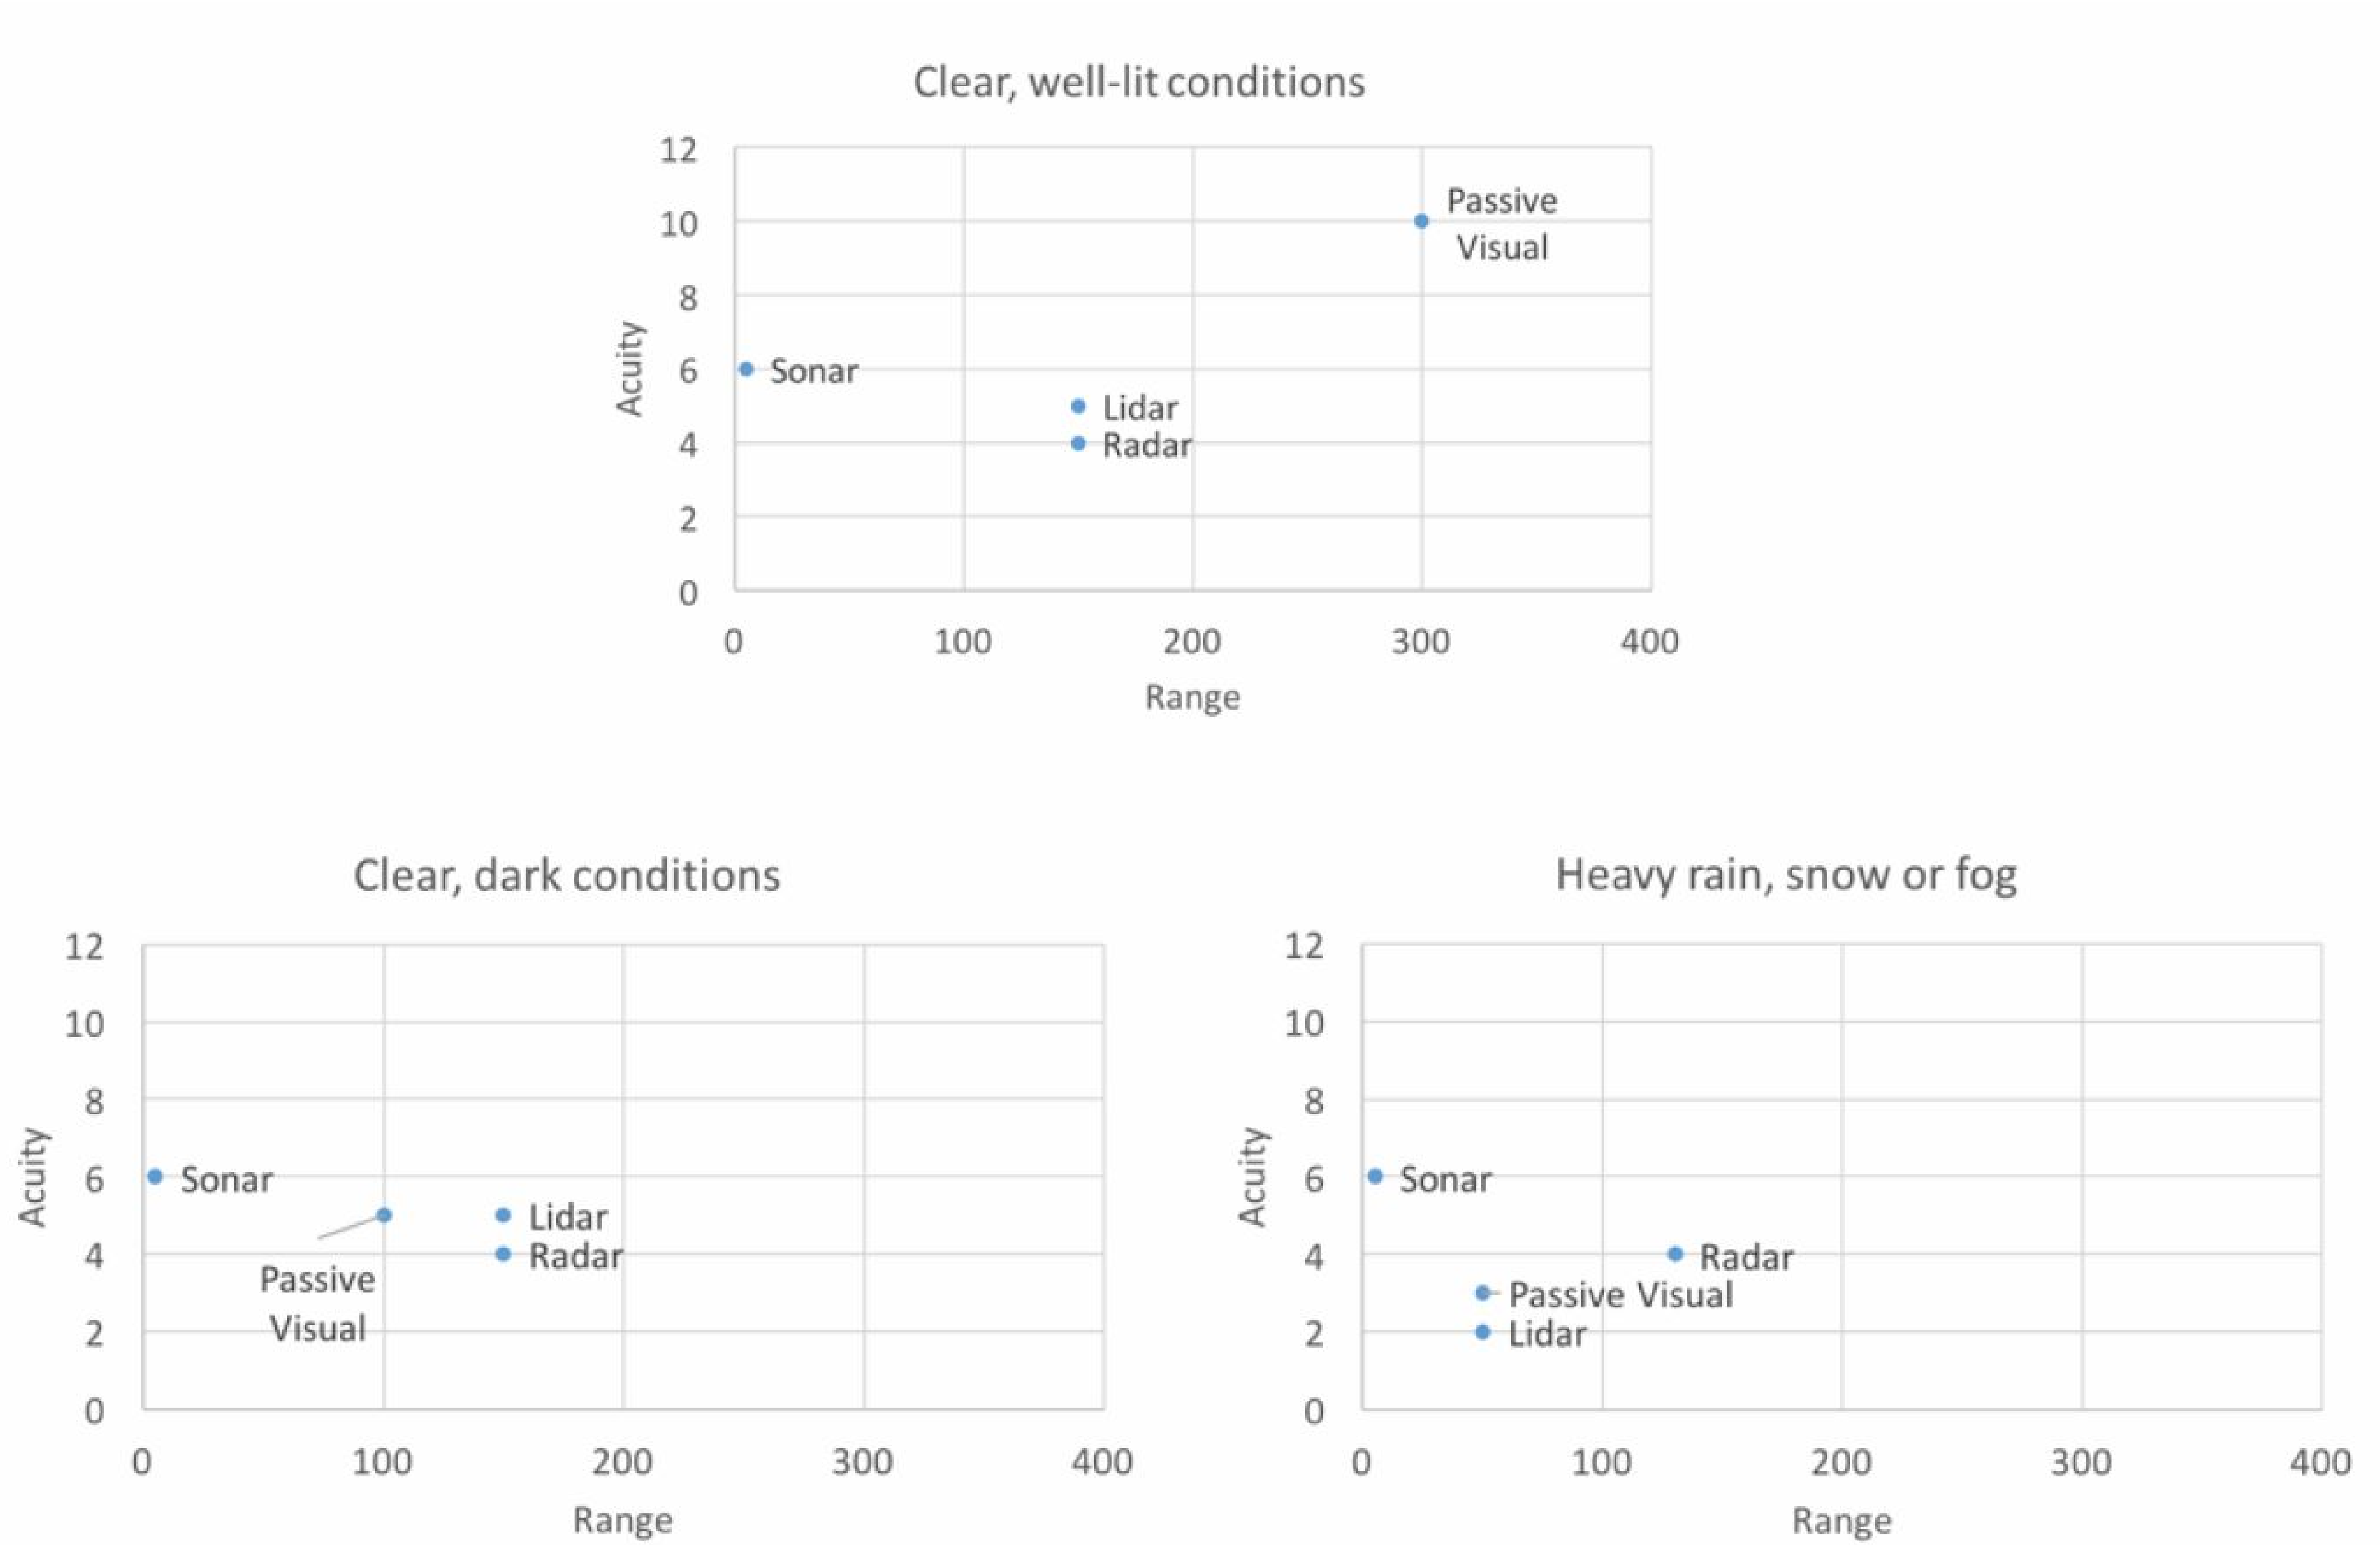

Source: https://cleantechnica.com/2016/07/29/tesla-google-disagree-lidar-right/

Lidar

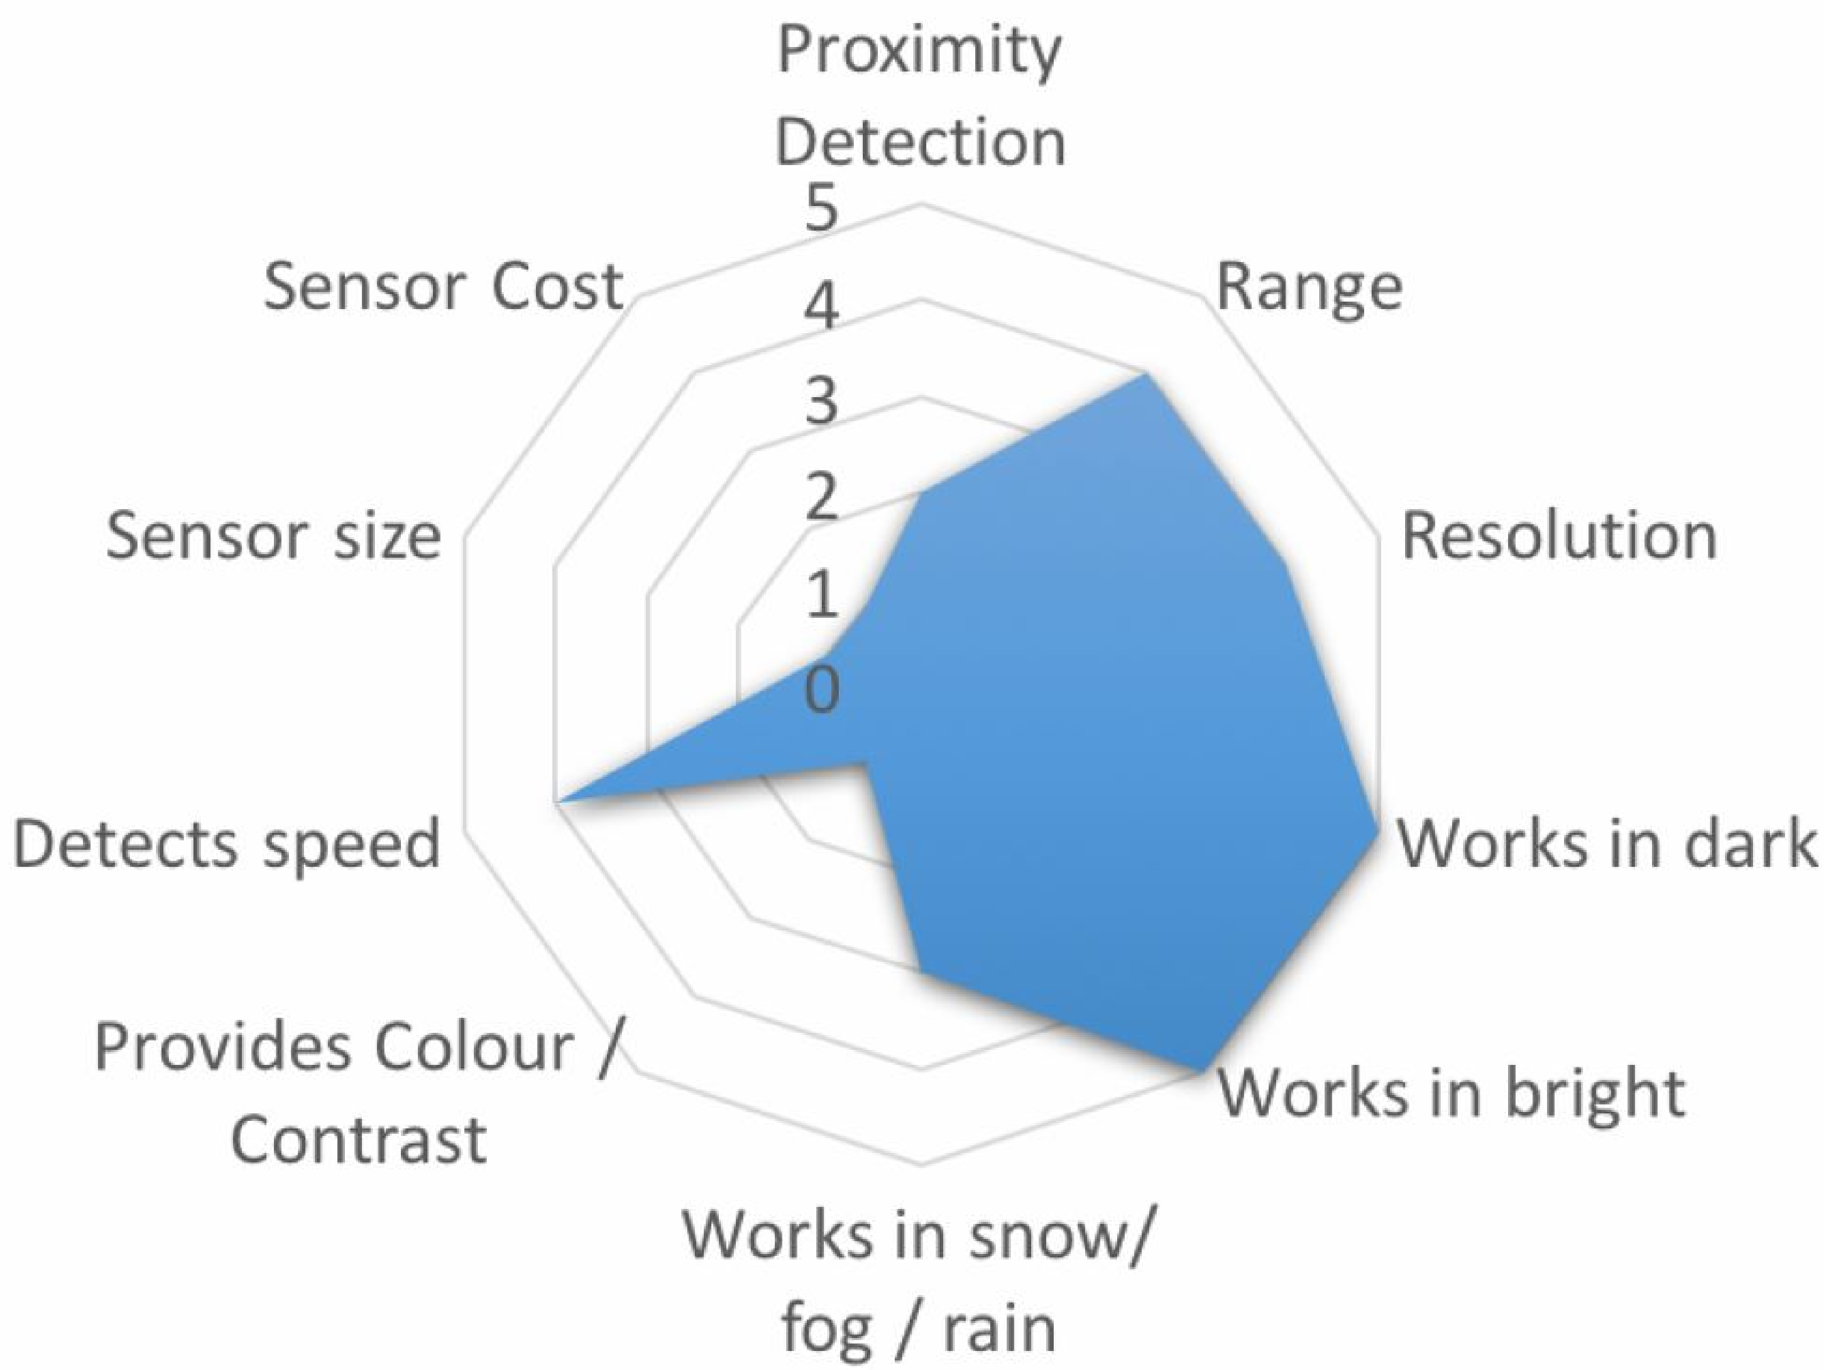
Proximity Detection
Range
Resolution
Works in dark
Works in bright
Works in snow/ fog / rain
Provides Colour / Contrast
Detects speed
Sensor size
Sensor Cost
5
4
3
2
1
0

Ultrasonic

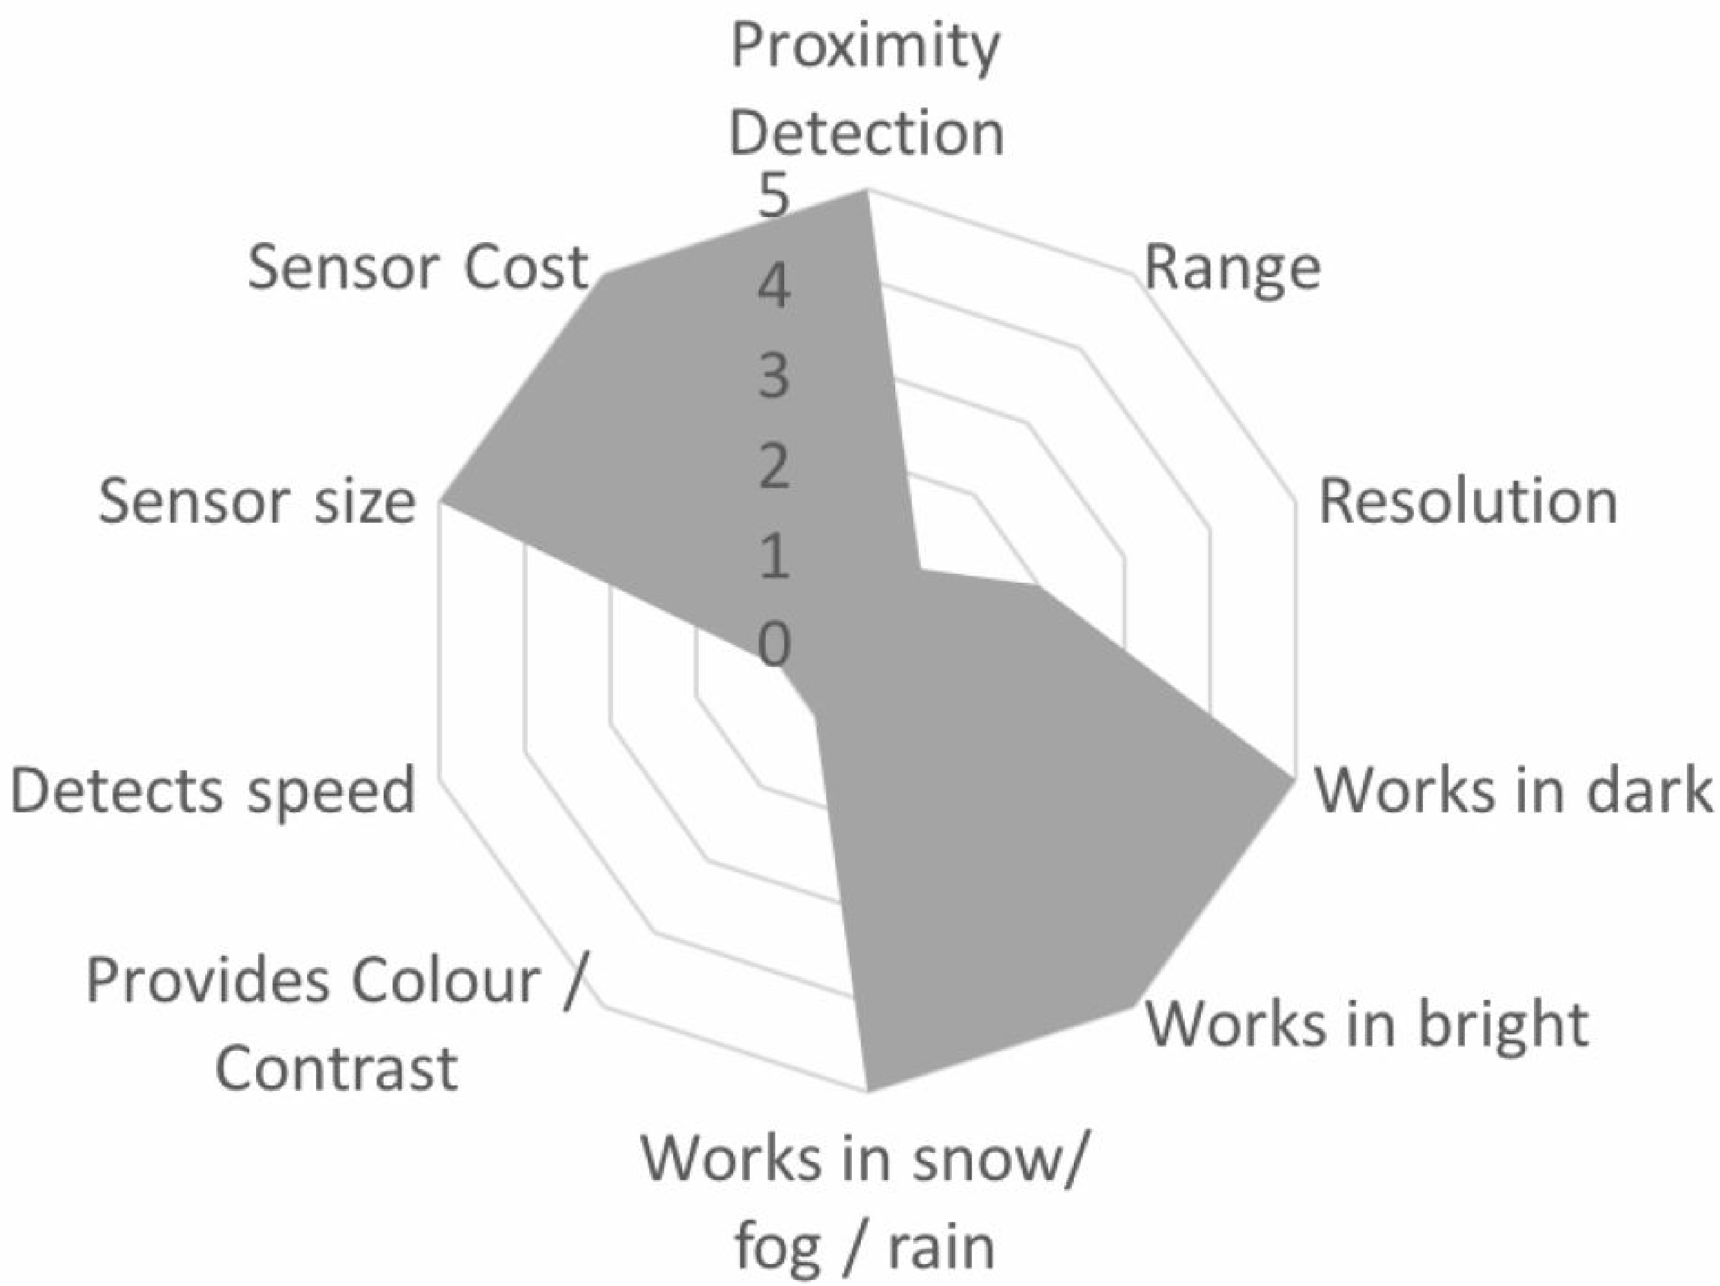

Proximity Detection
Range
Resolution
Works in dark
Works in bright
Works in snow/ fog / rain
Provides Colour / Contrast
Detects speed
Sensor size
Sensor Cost
5
4
3
2
1
0

Radar

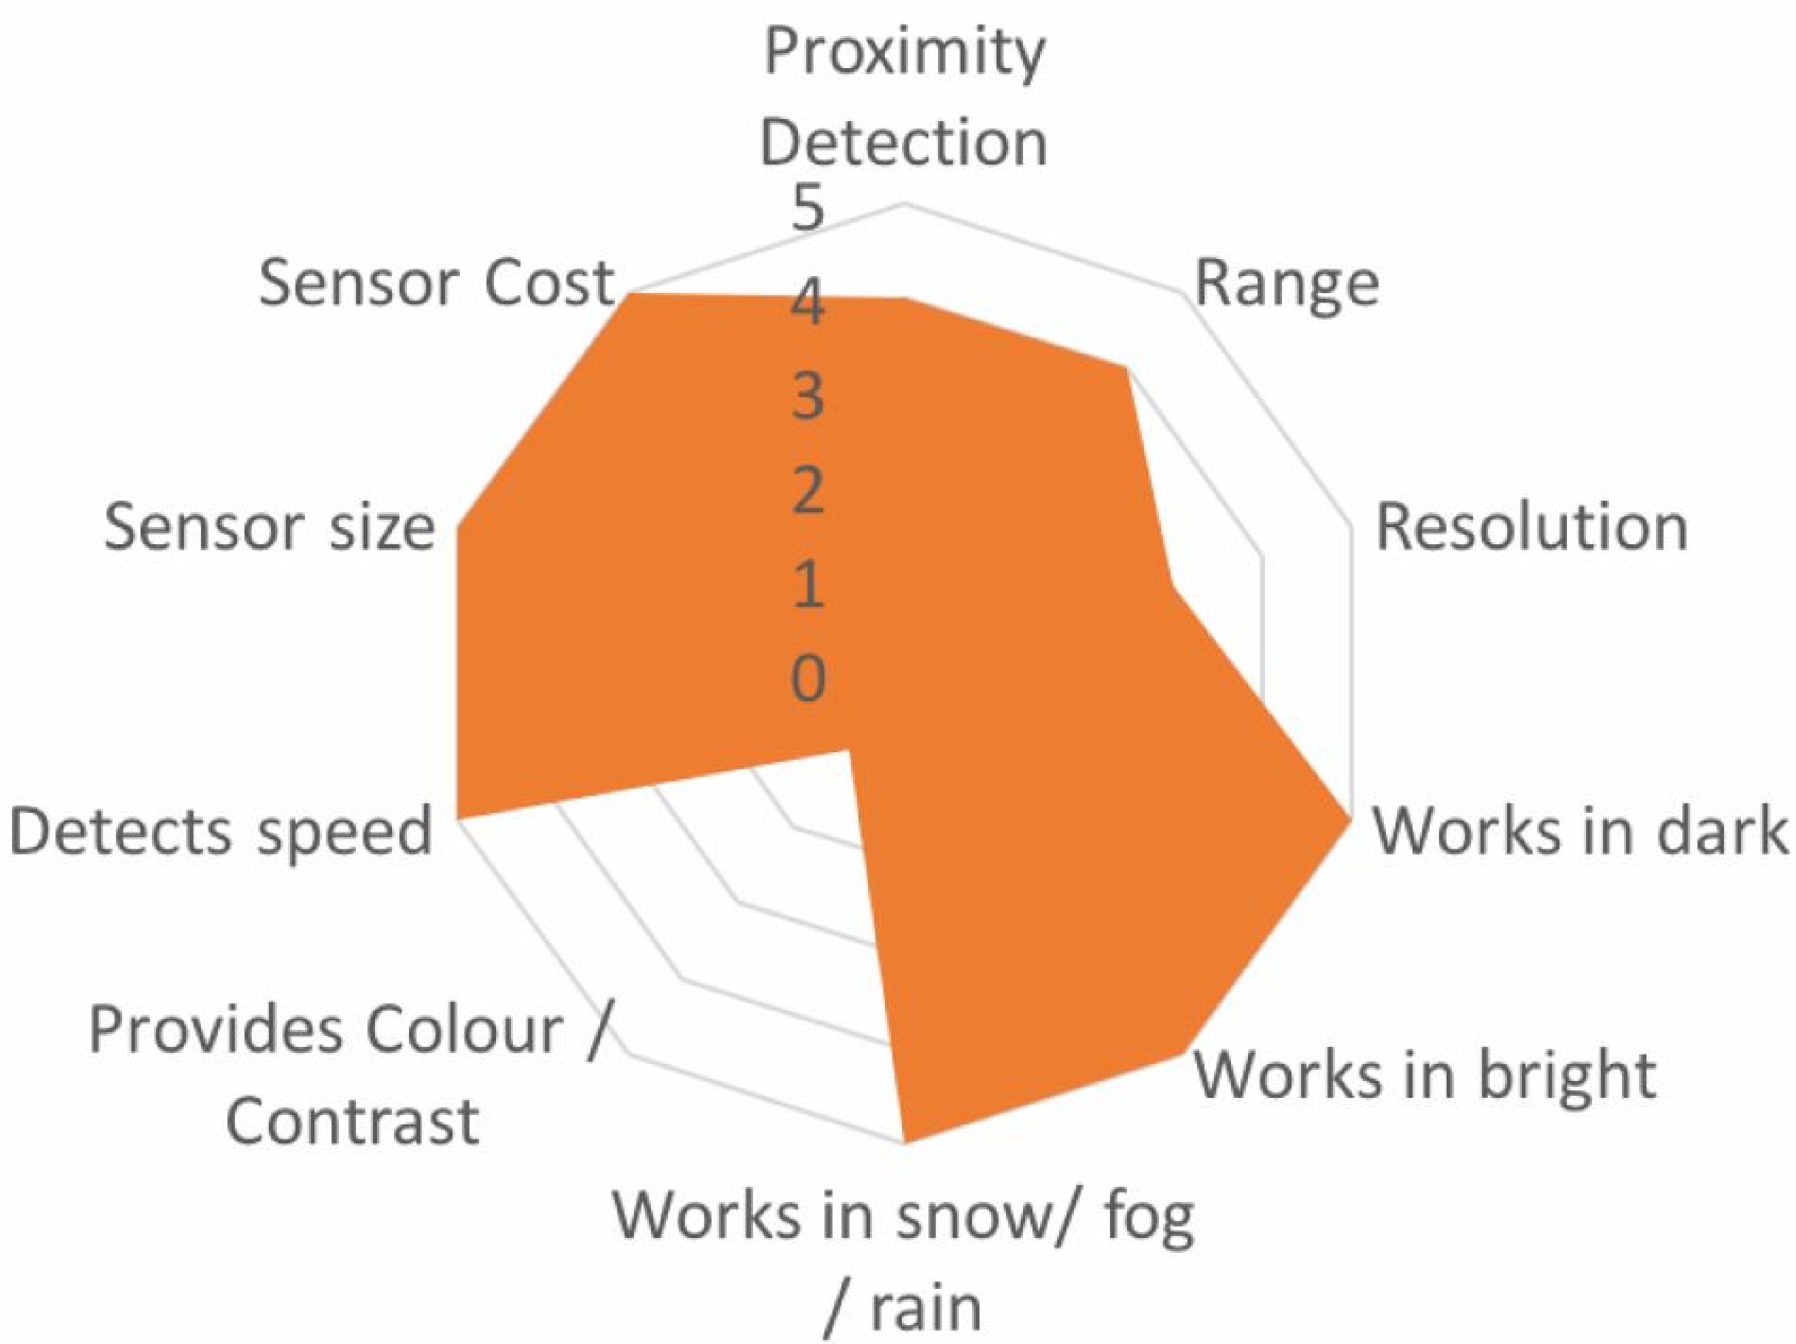
Proximity Detection
Range
Resolution
Works in dark
Works in bright
Works in snow/ fog / rain
Provides Colour / Contrast
Detects speed
Sensor size
Sensor Cost
5
4
3
2
1
0

Passive Visual

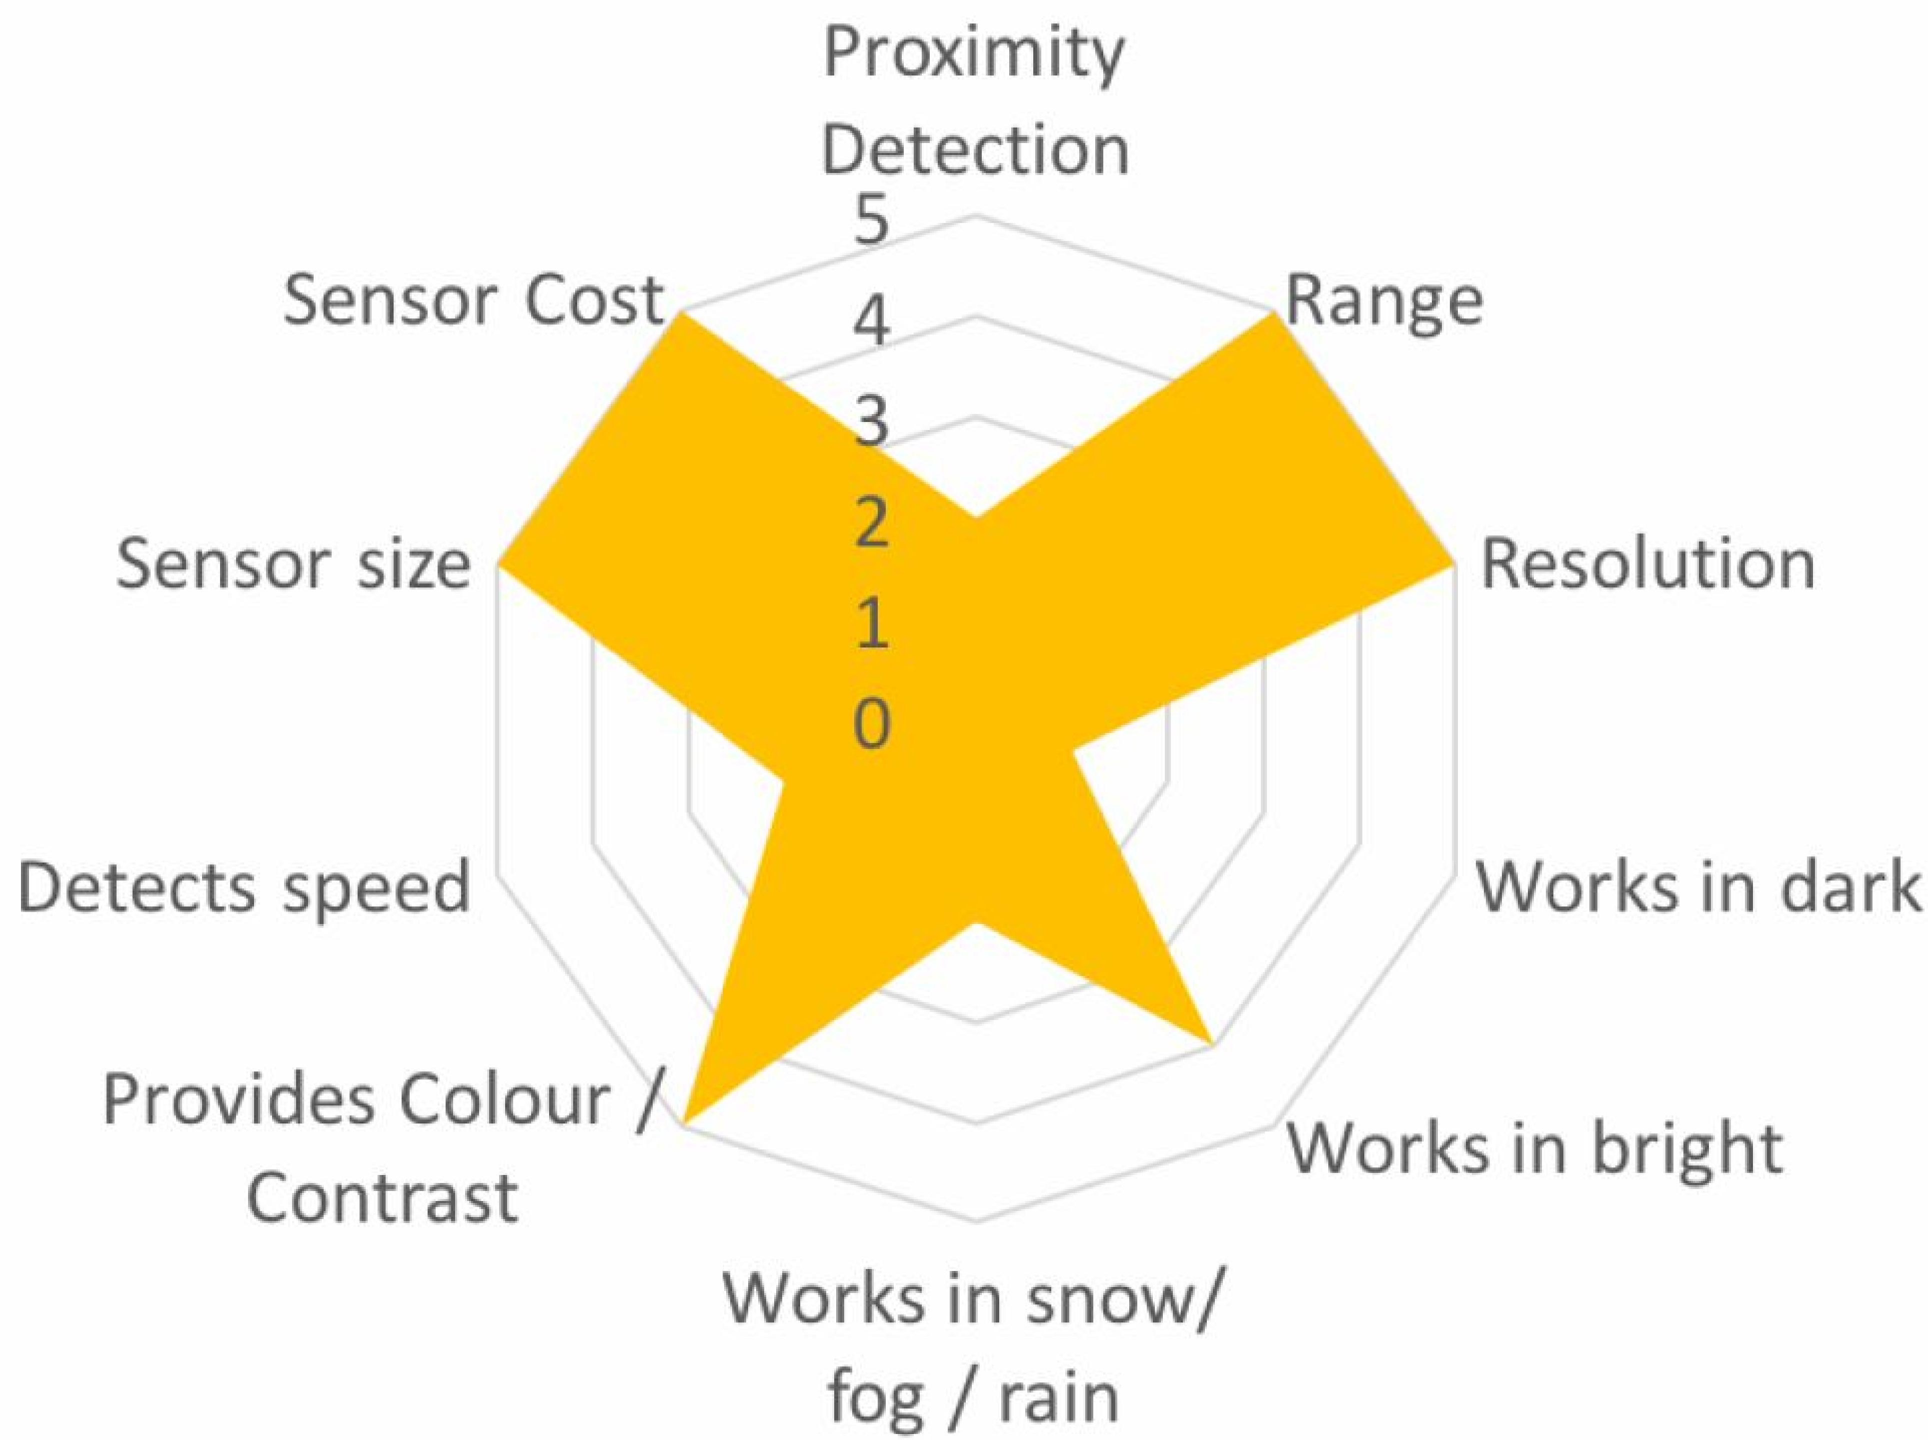

Sensor Fusion

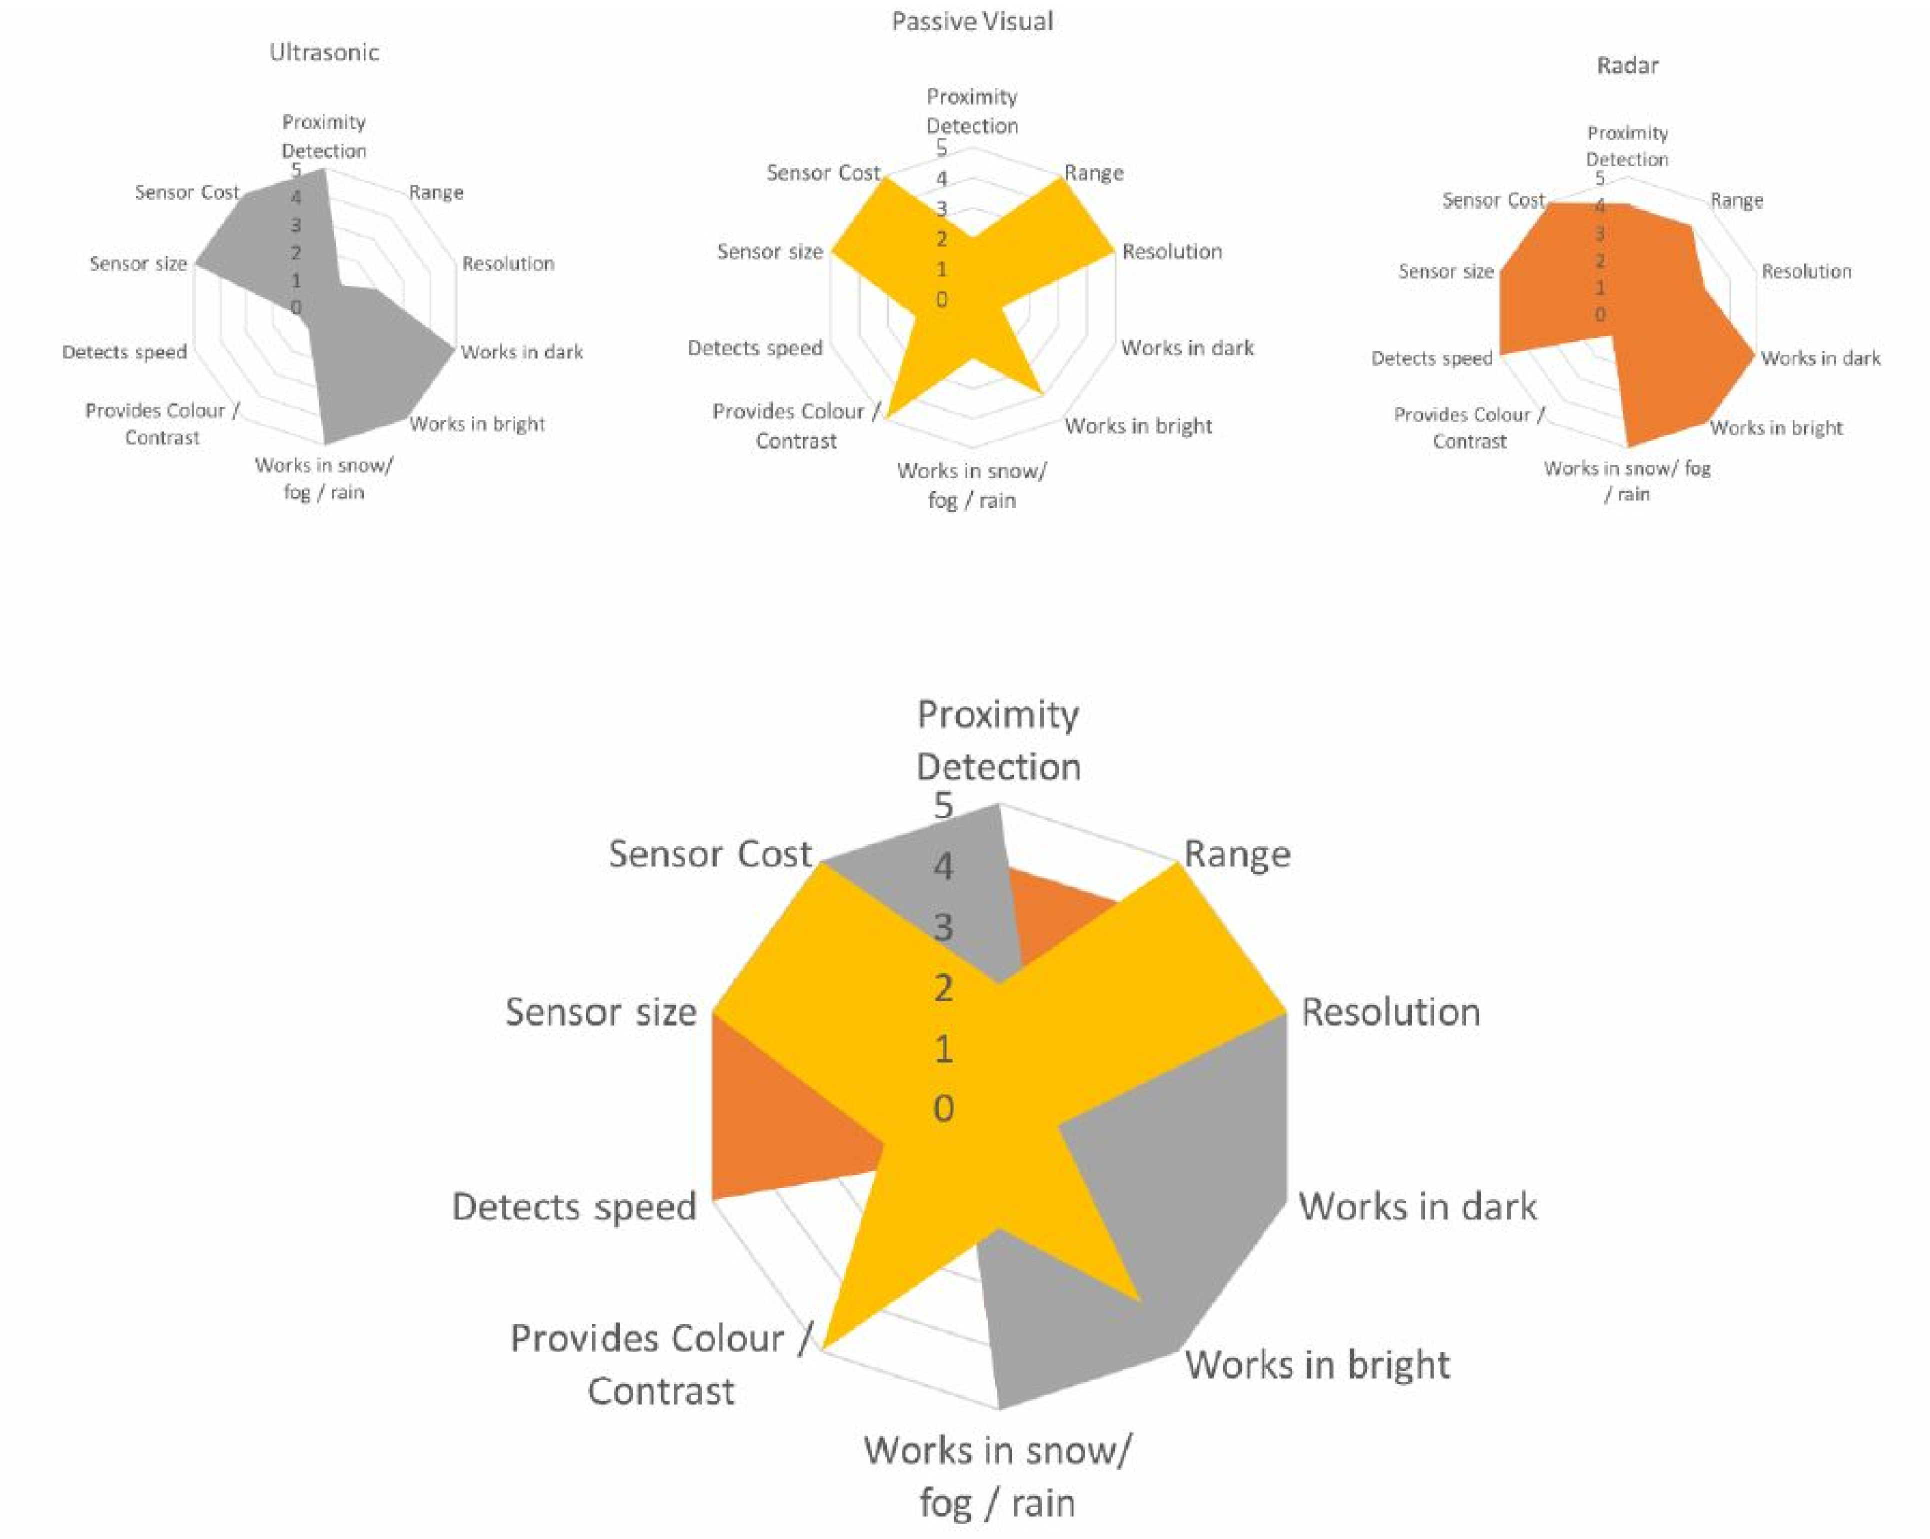

Future of Sensor Technology: Camera vs Lidar

- **Radar and Ultrasonic:**
  - Always there to help
- **Camera:**
  - Annotated driving data grows
  - Deep learning algorithms improve
- **LIDAR:**
  - Range increases
  - Cost drops (solid-state LIDAR)

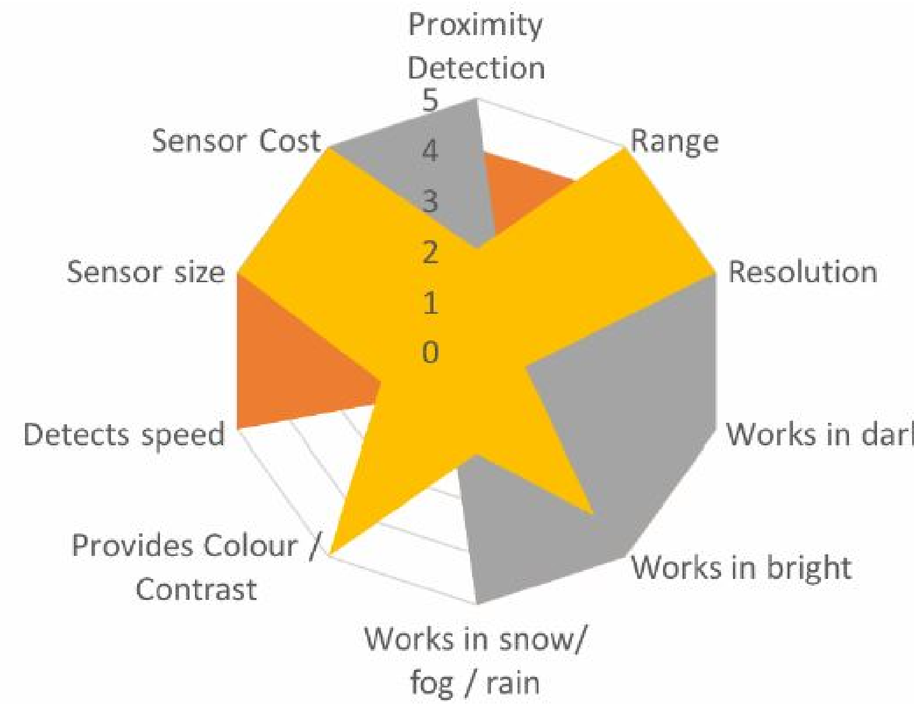

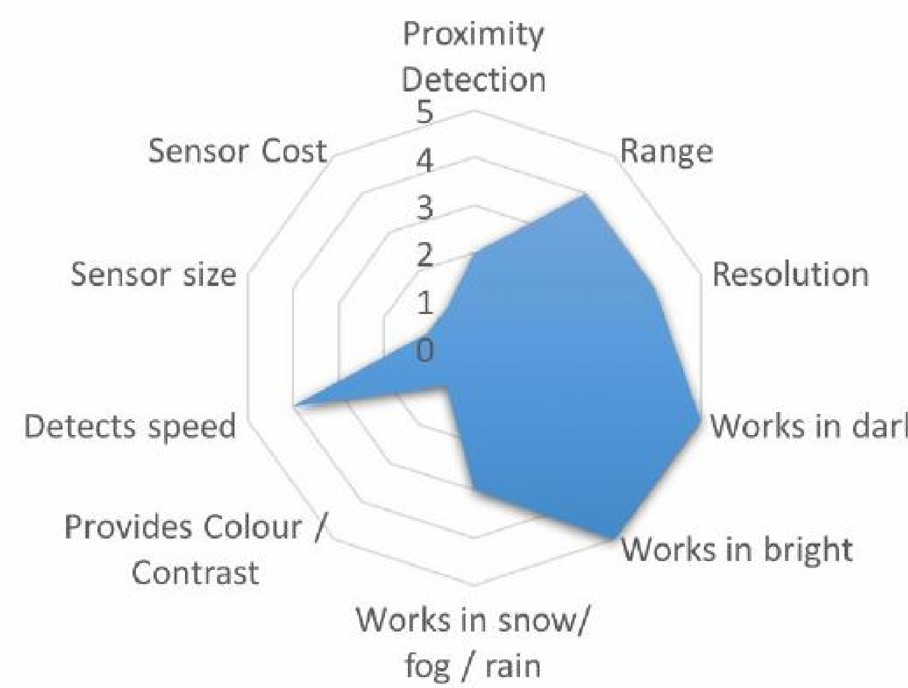

Functional Reference Architecture

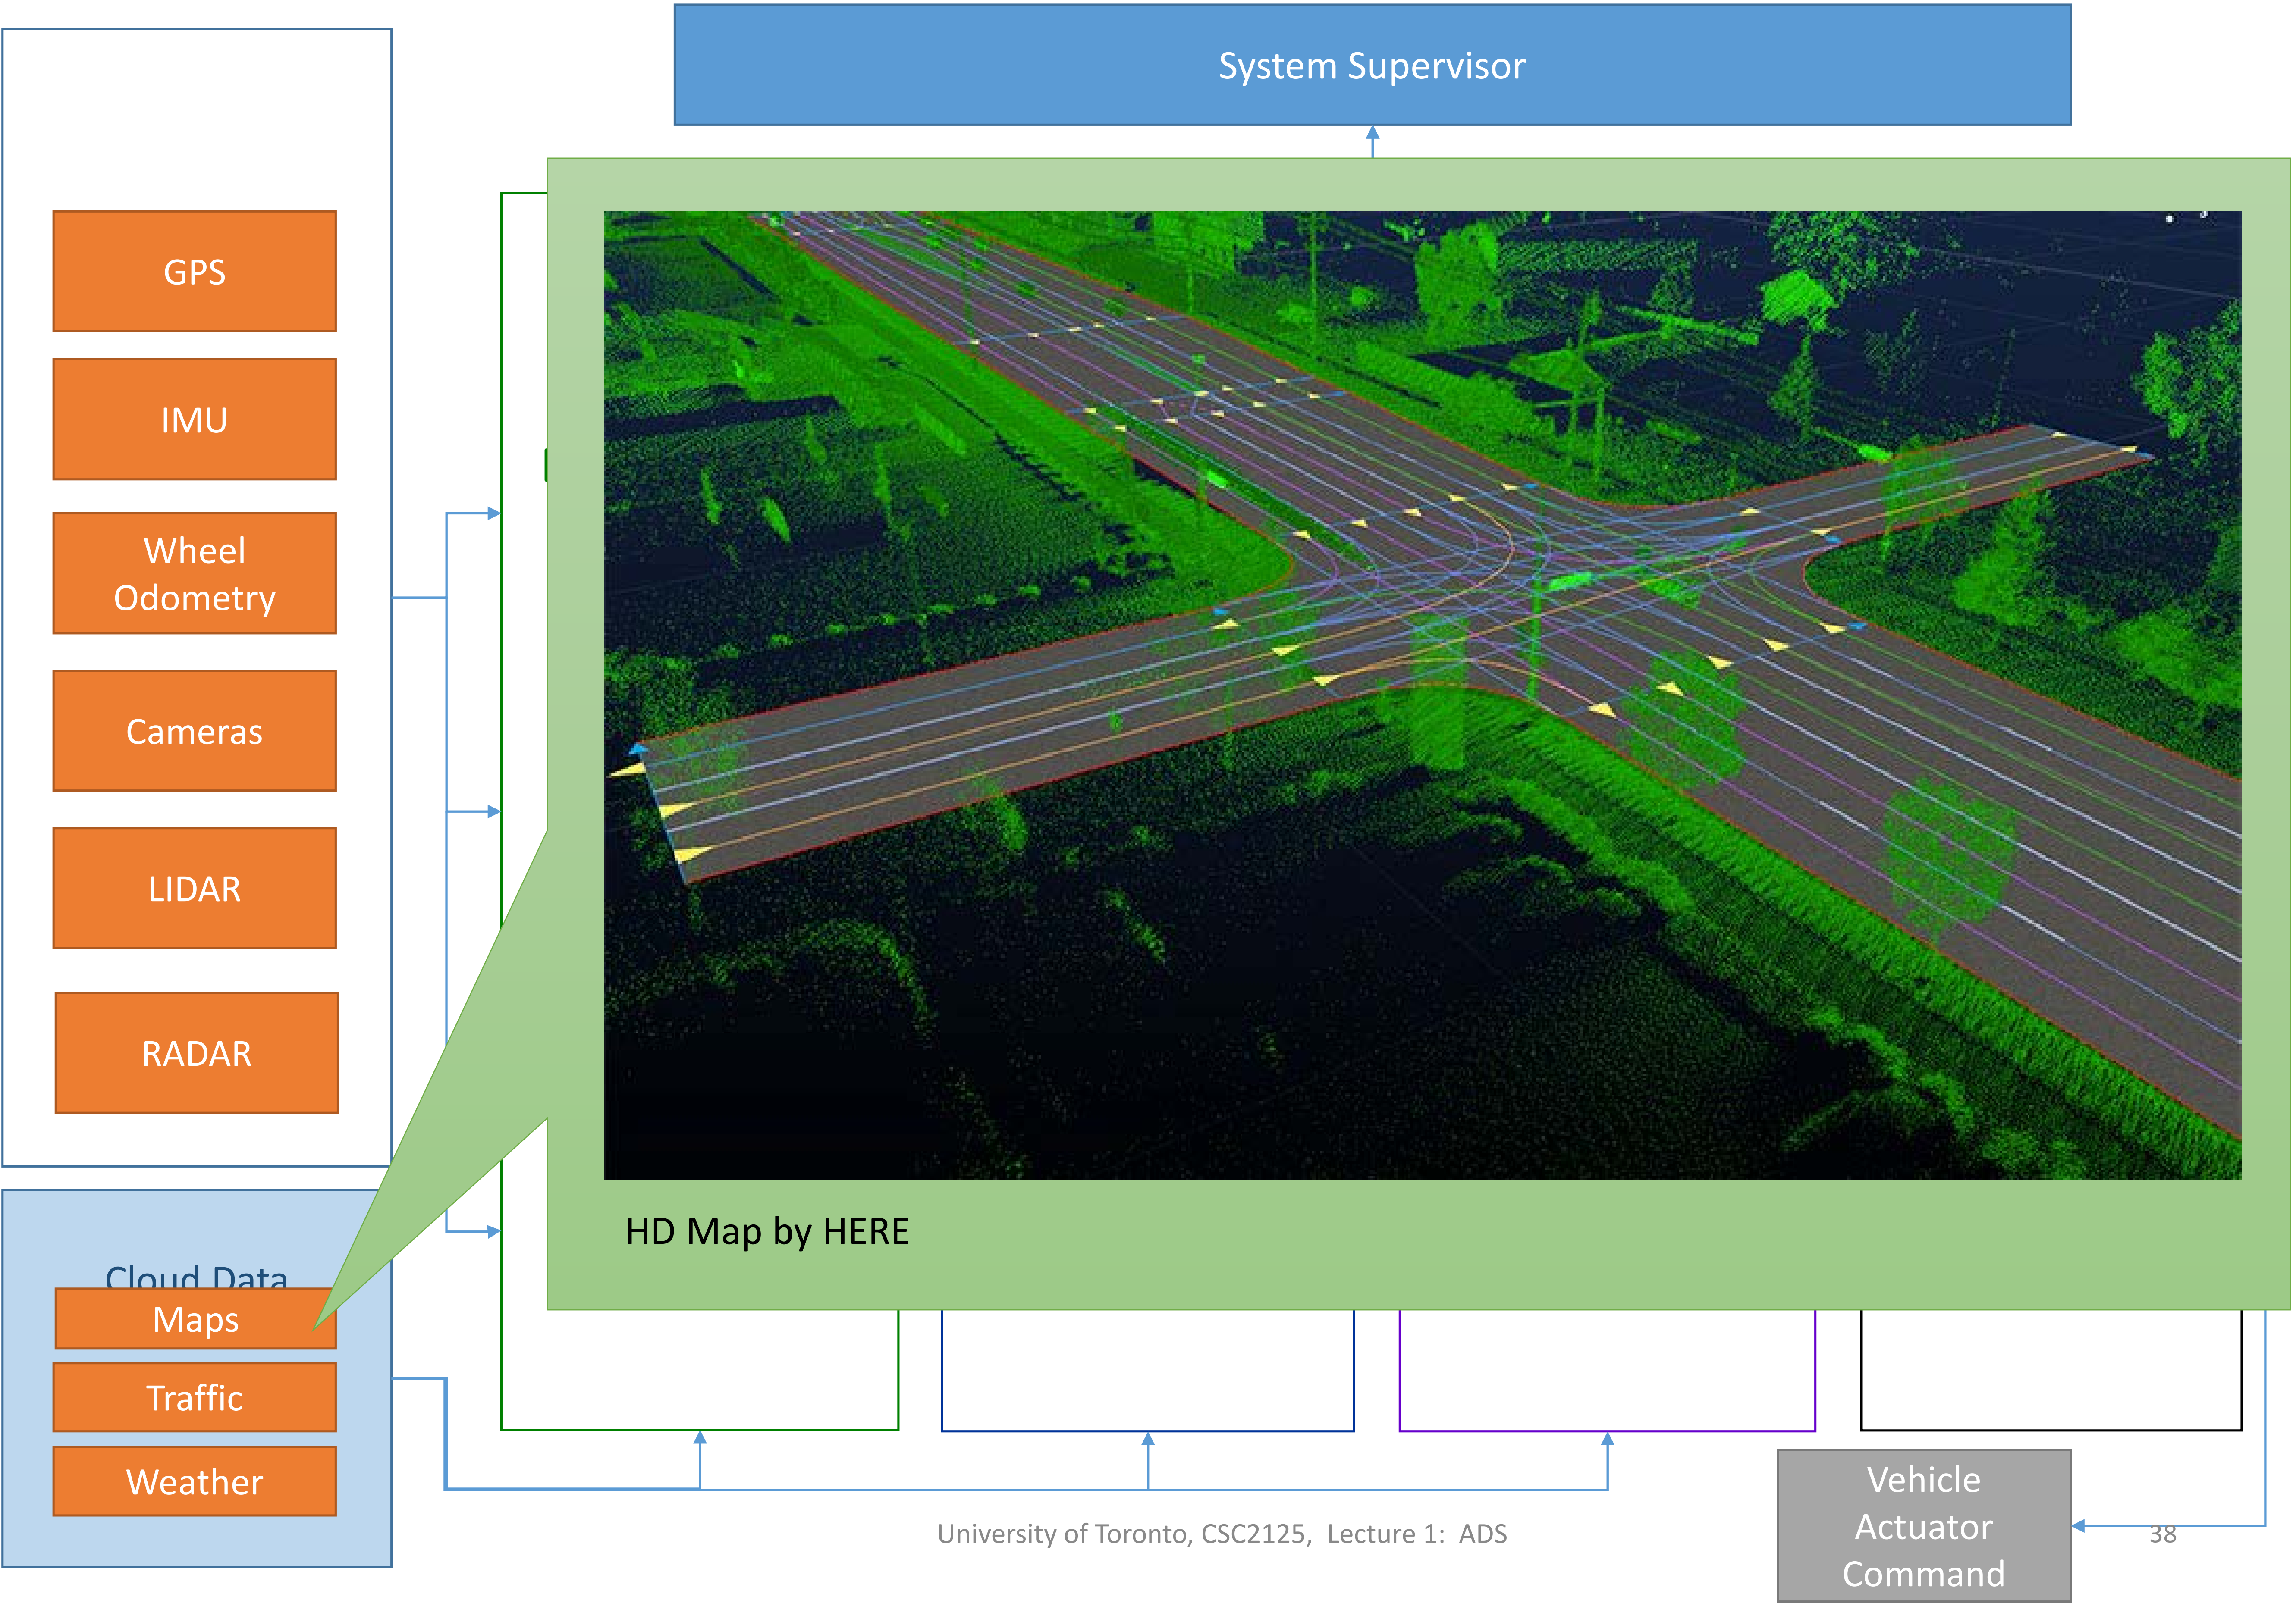

Functional Reference Architecture

System Supervisor

GPS

IMU

Wheel Odometry

Cameras

V2X

V2I

V2V

...

Cloud Data

Maps

Traffic

Weather

Vehicle to infrastructure (V2I) and vehicle to vehicle (V2V) communication

Vehicle Actuator Command

Functional Reference Architecture

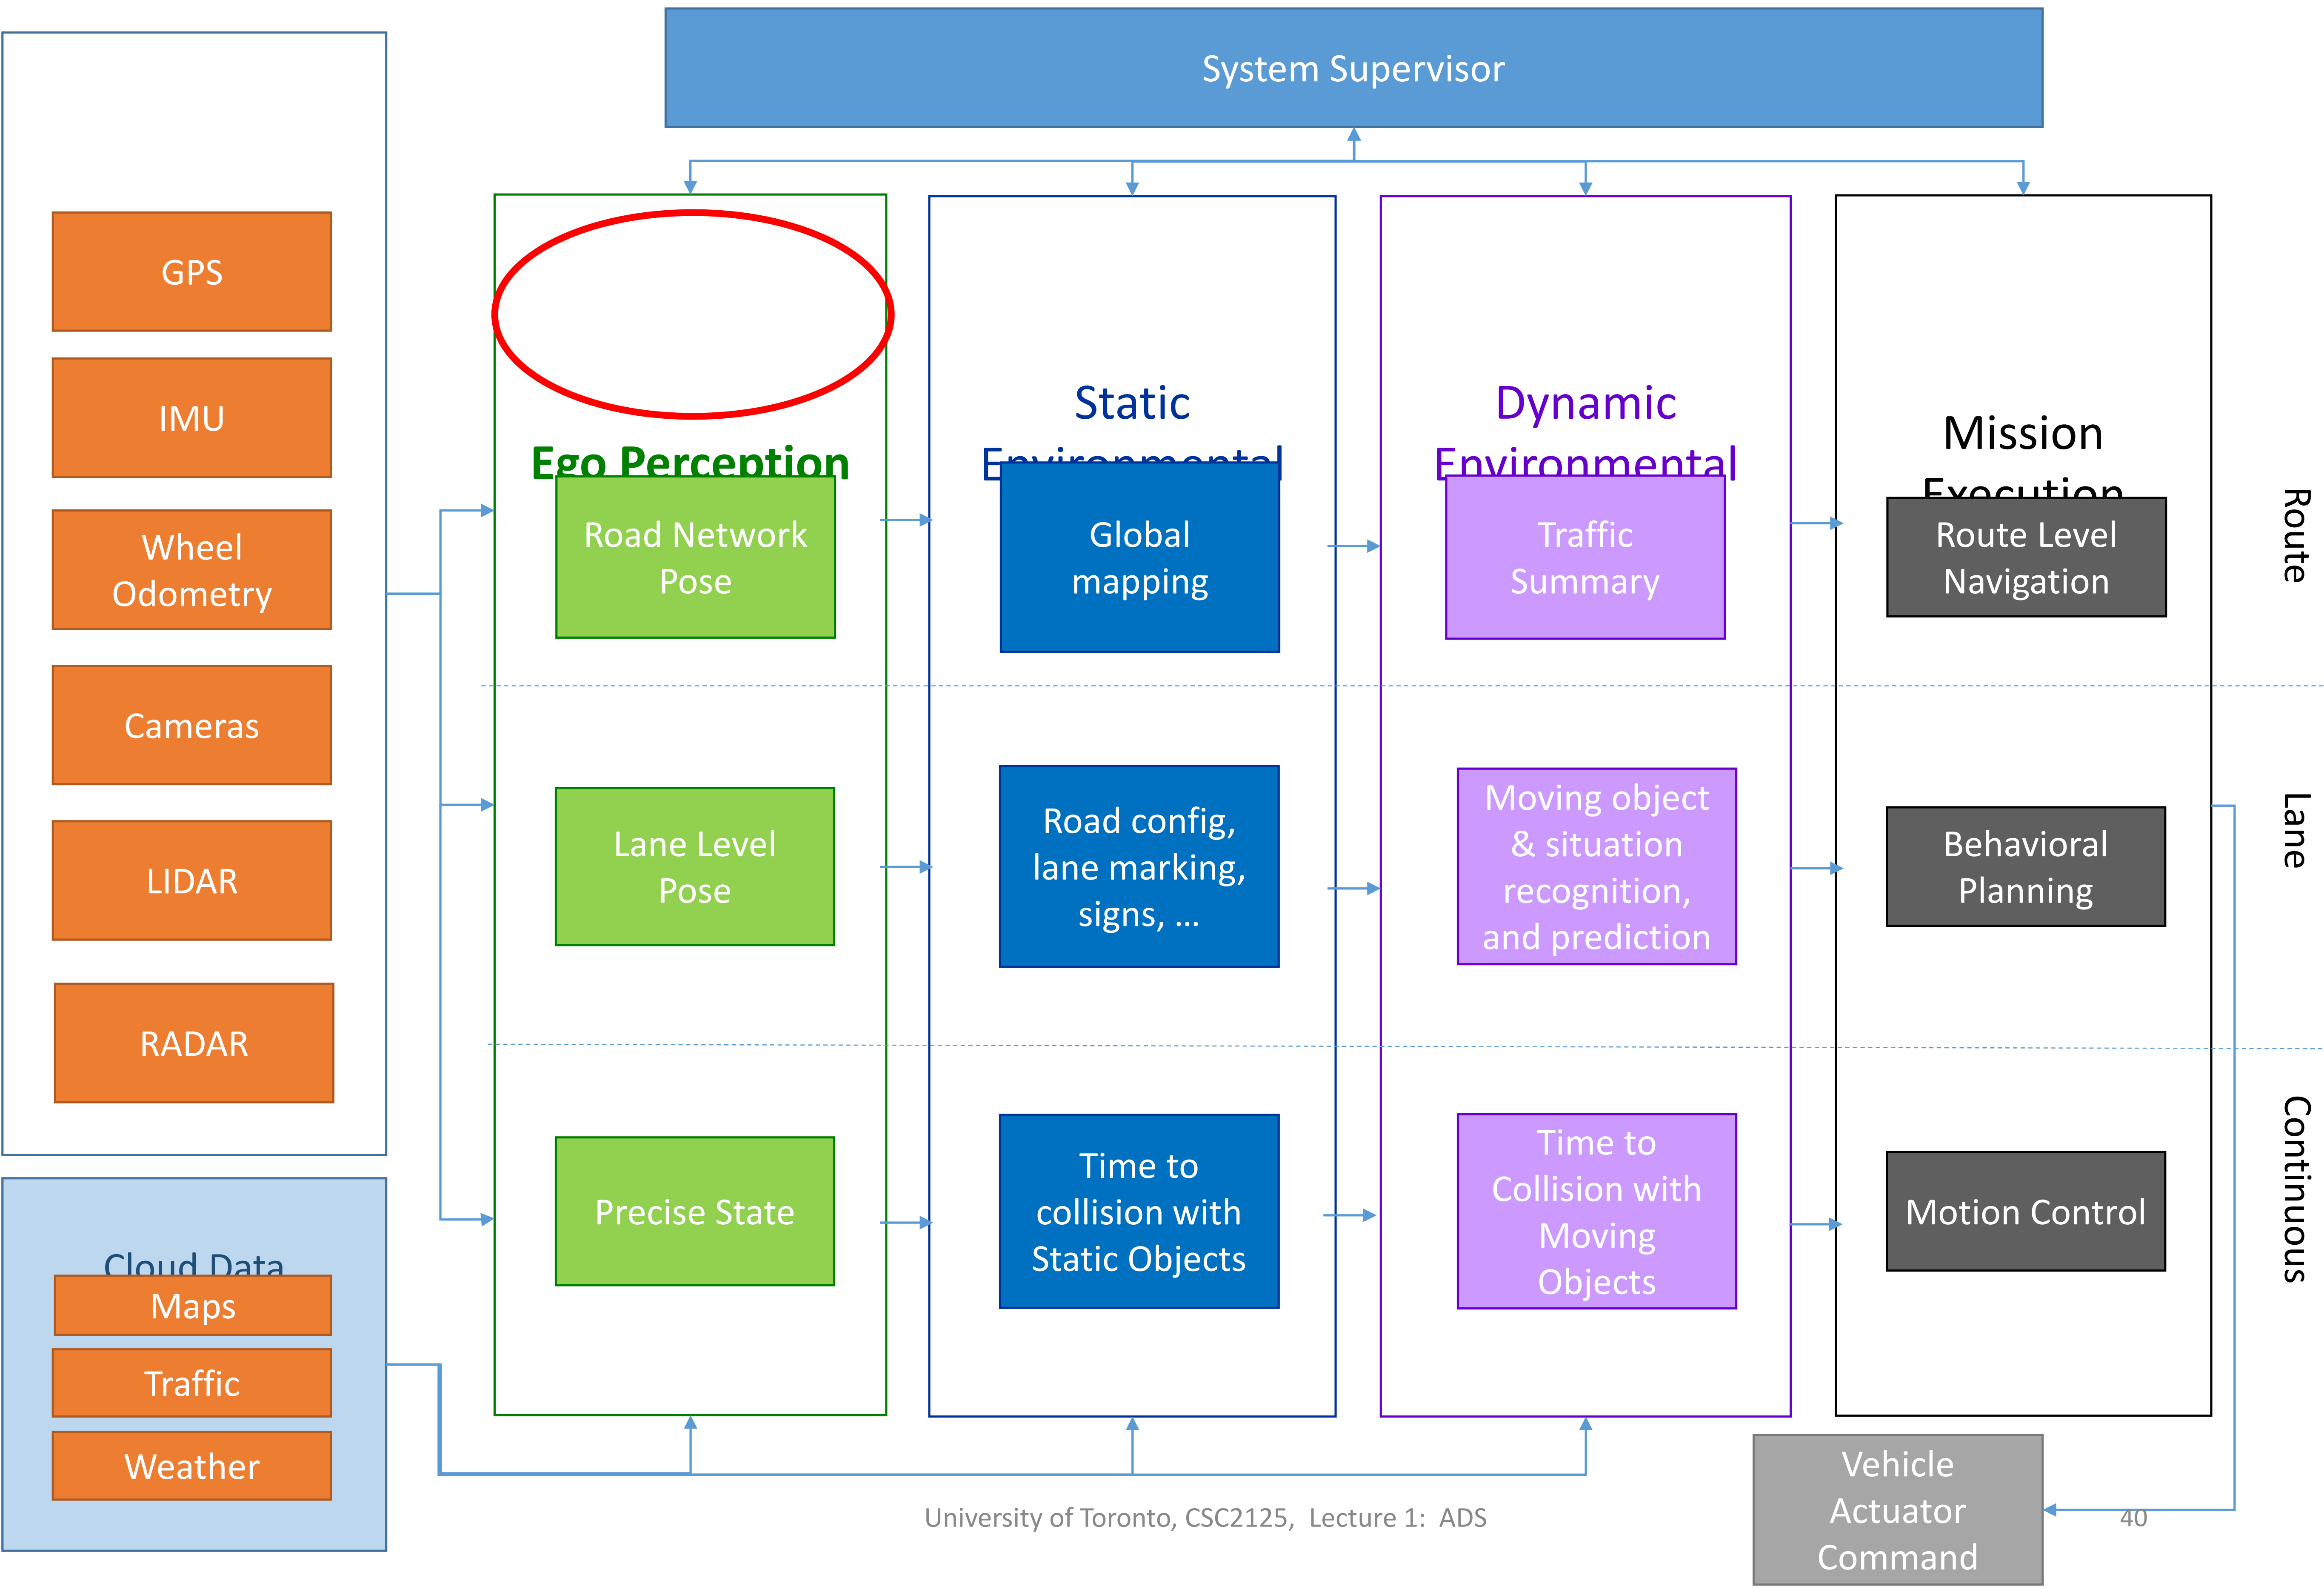

Functional Reference Architecture

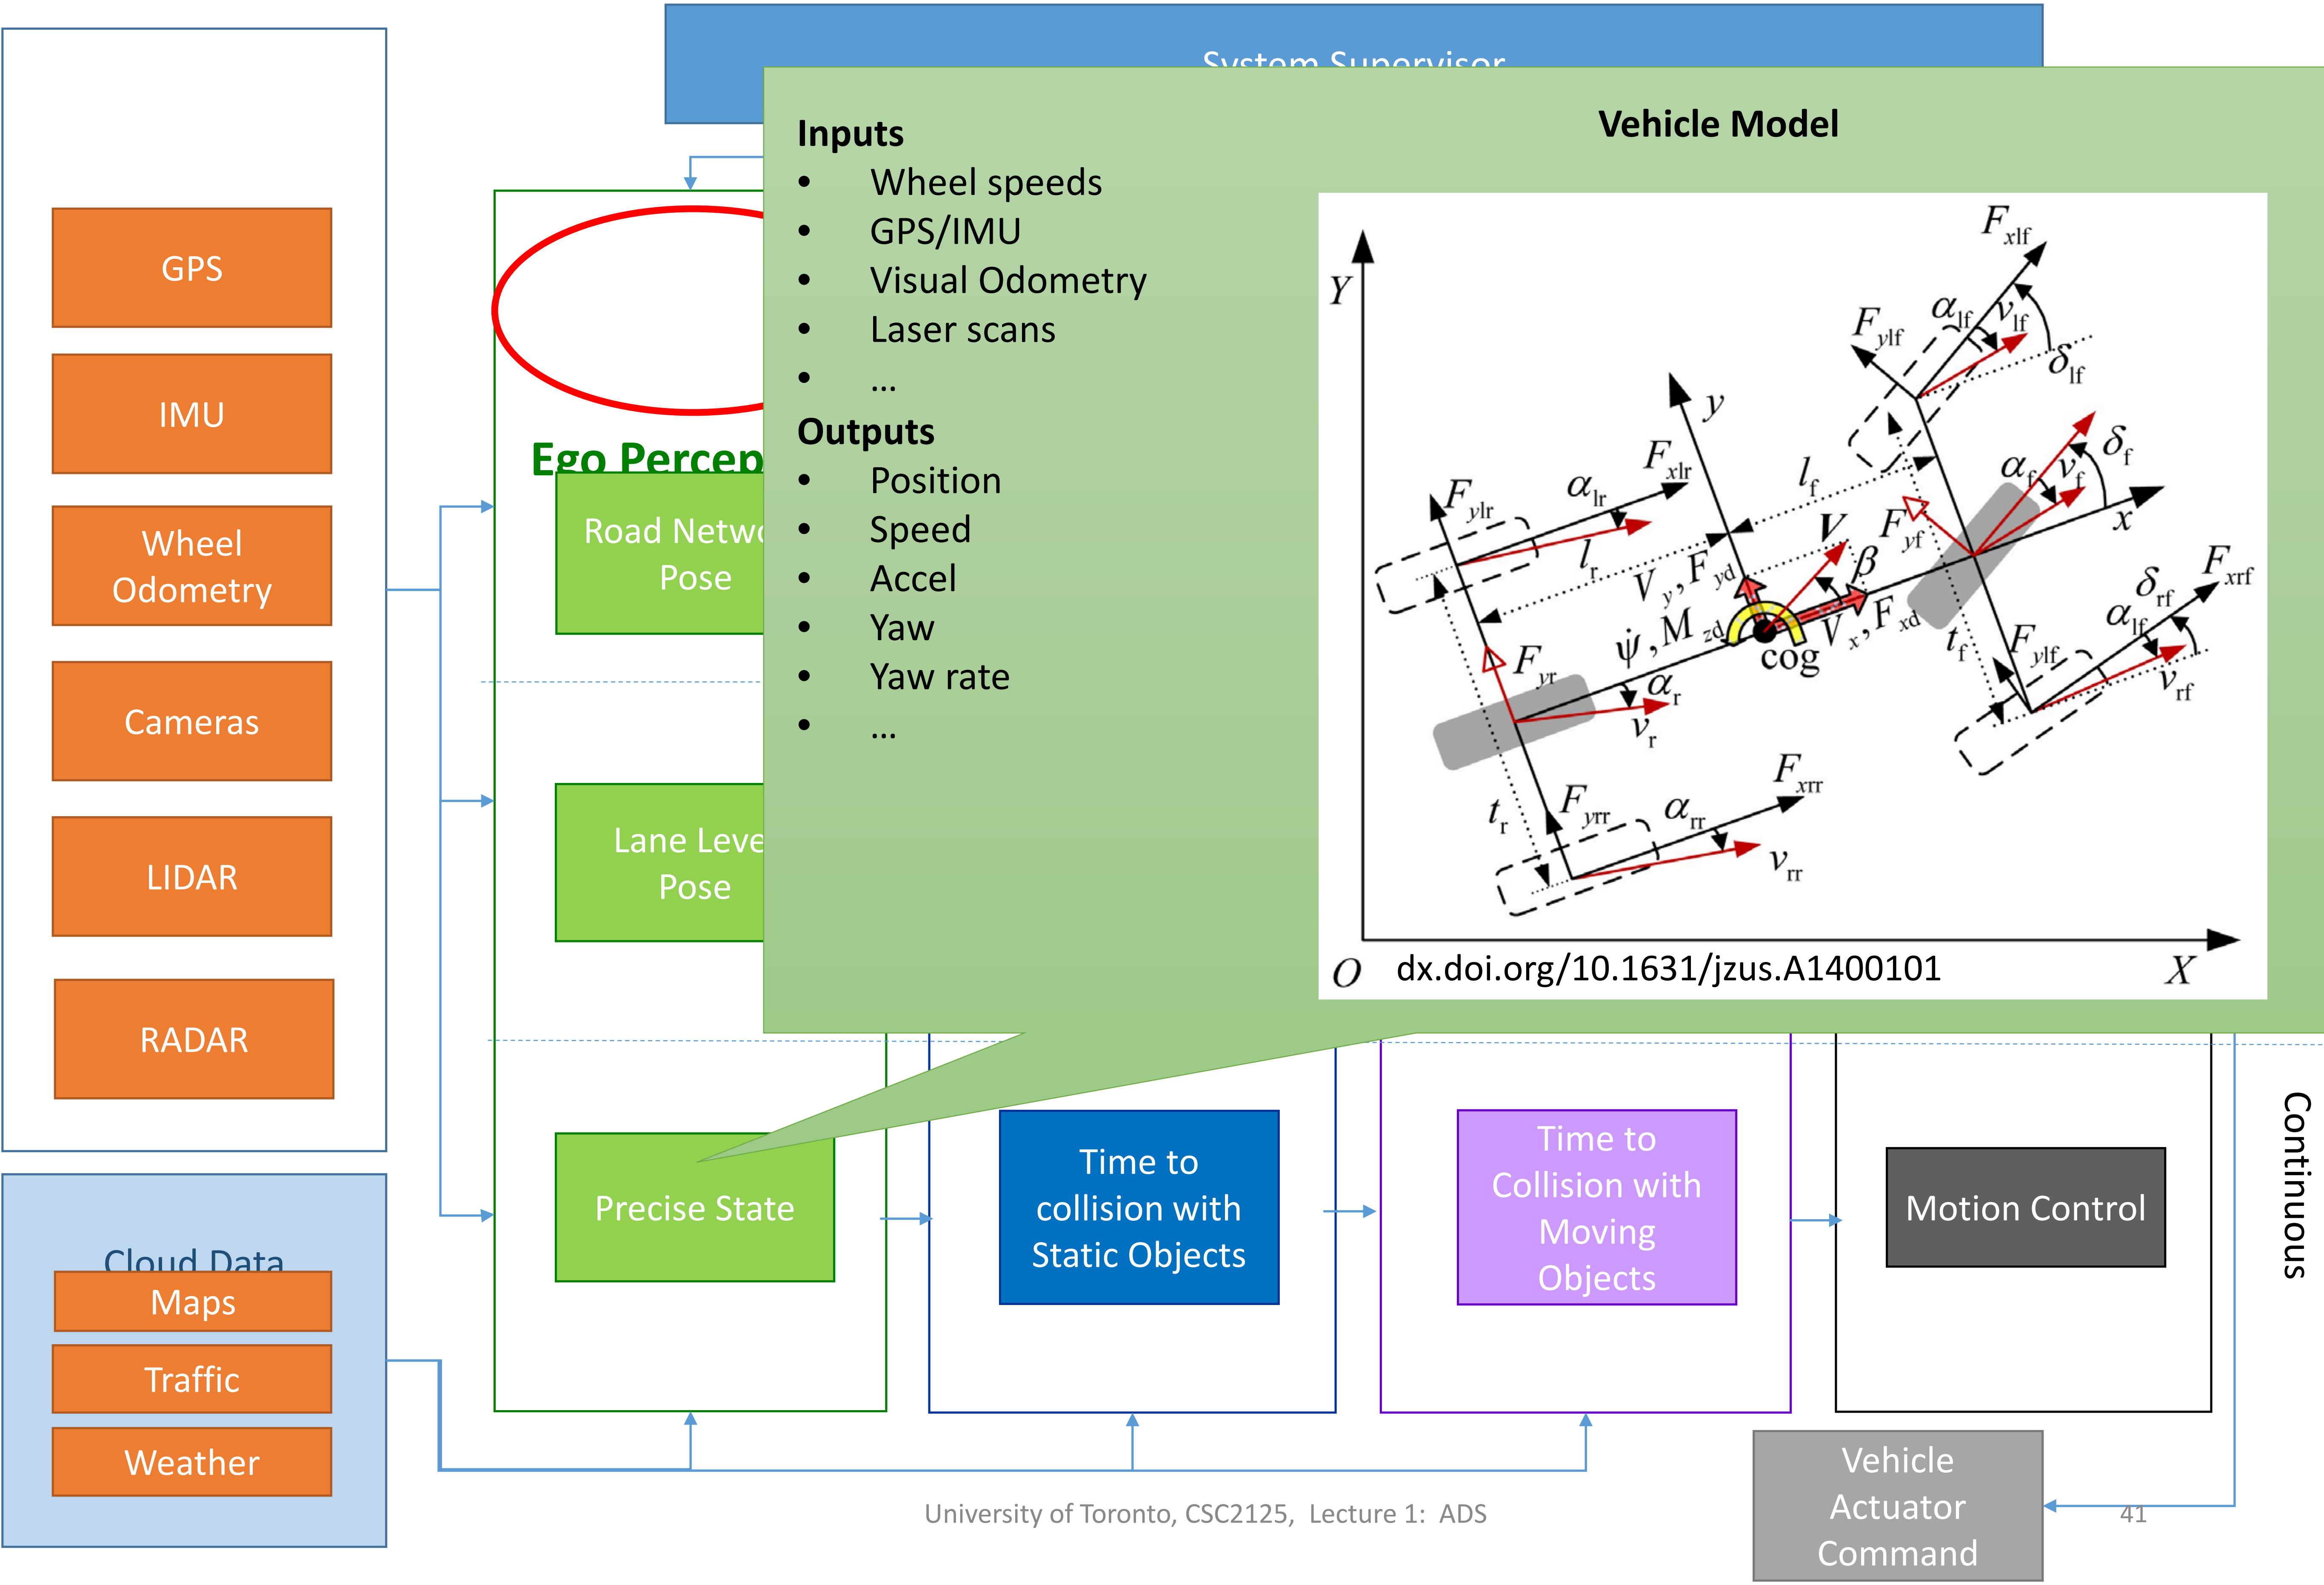

Vehicle Model

Inputs

- Wheel speeds
- GPS/IMU
- Visual Odometry
- Laser scans
- …

Outputs

- Position
- Speed
- Accel
- Yaw
- Yaw rate
- …

dx.doi.org/10.1631/jzus.A1400101

Functional Reference Architecture

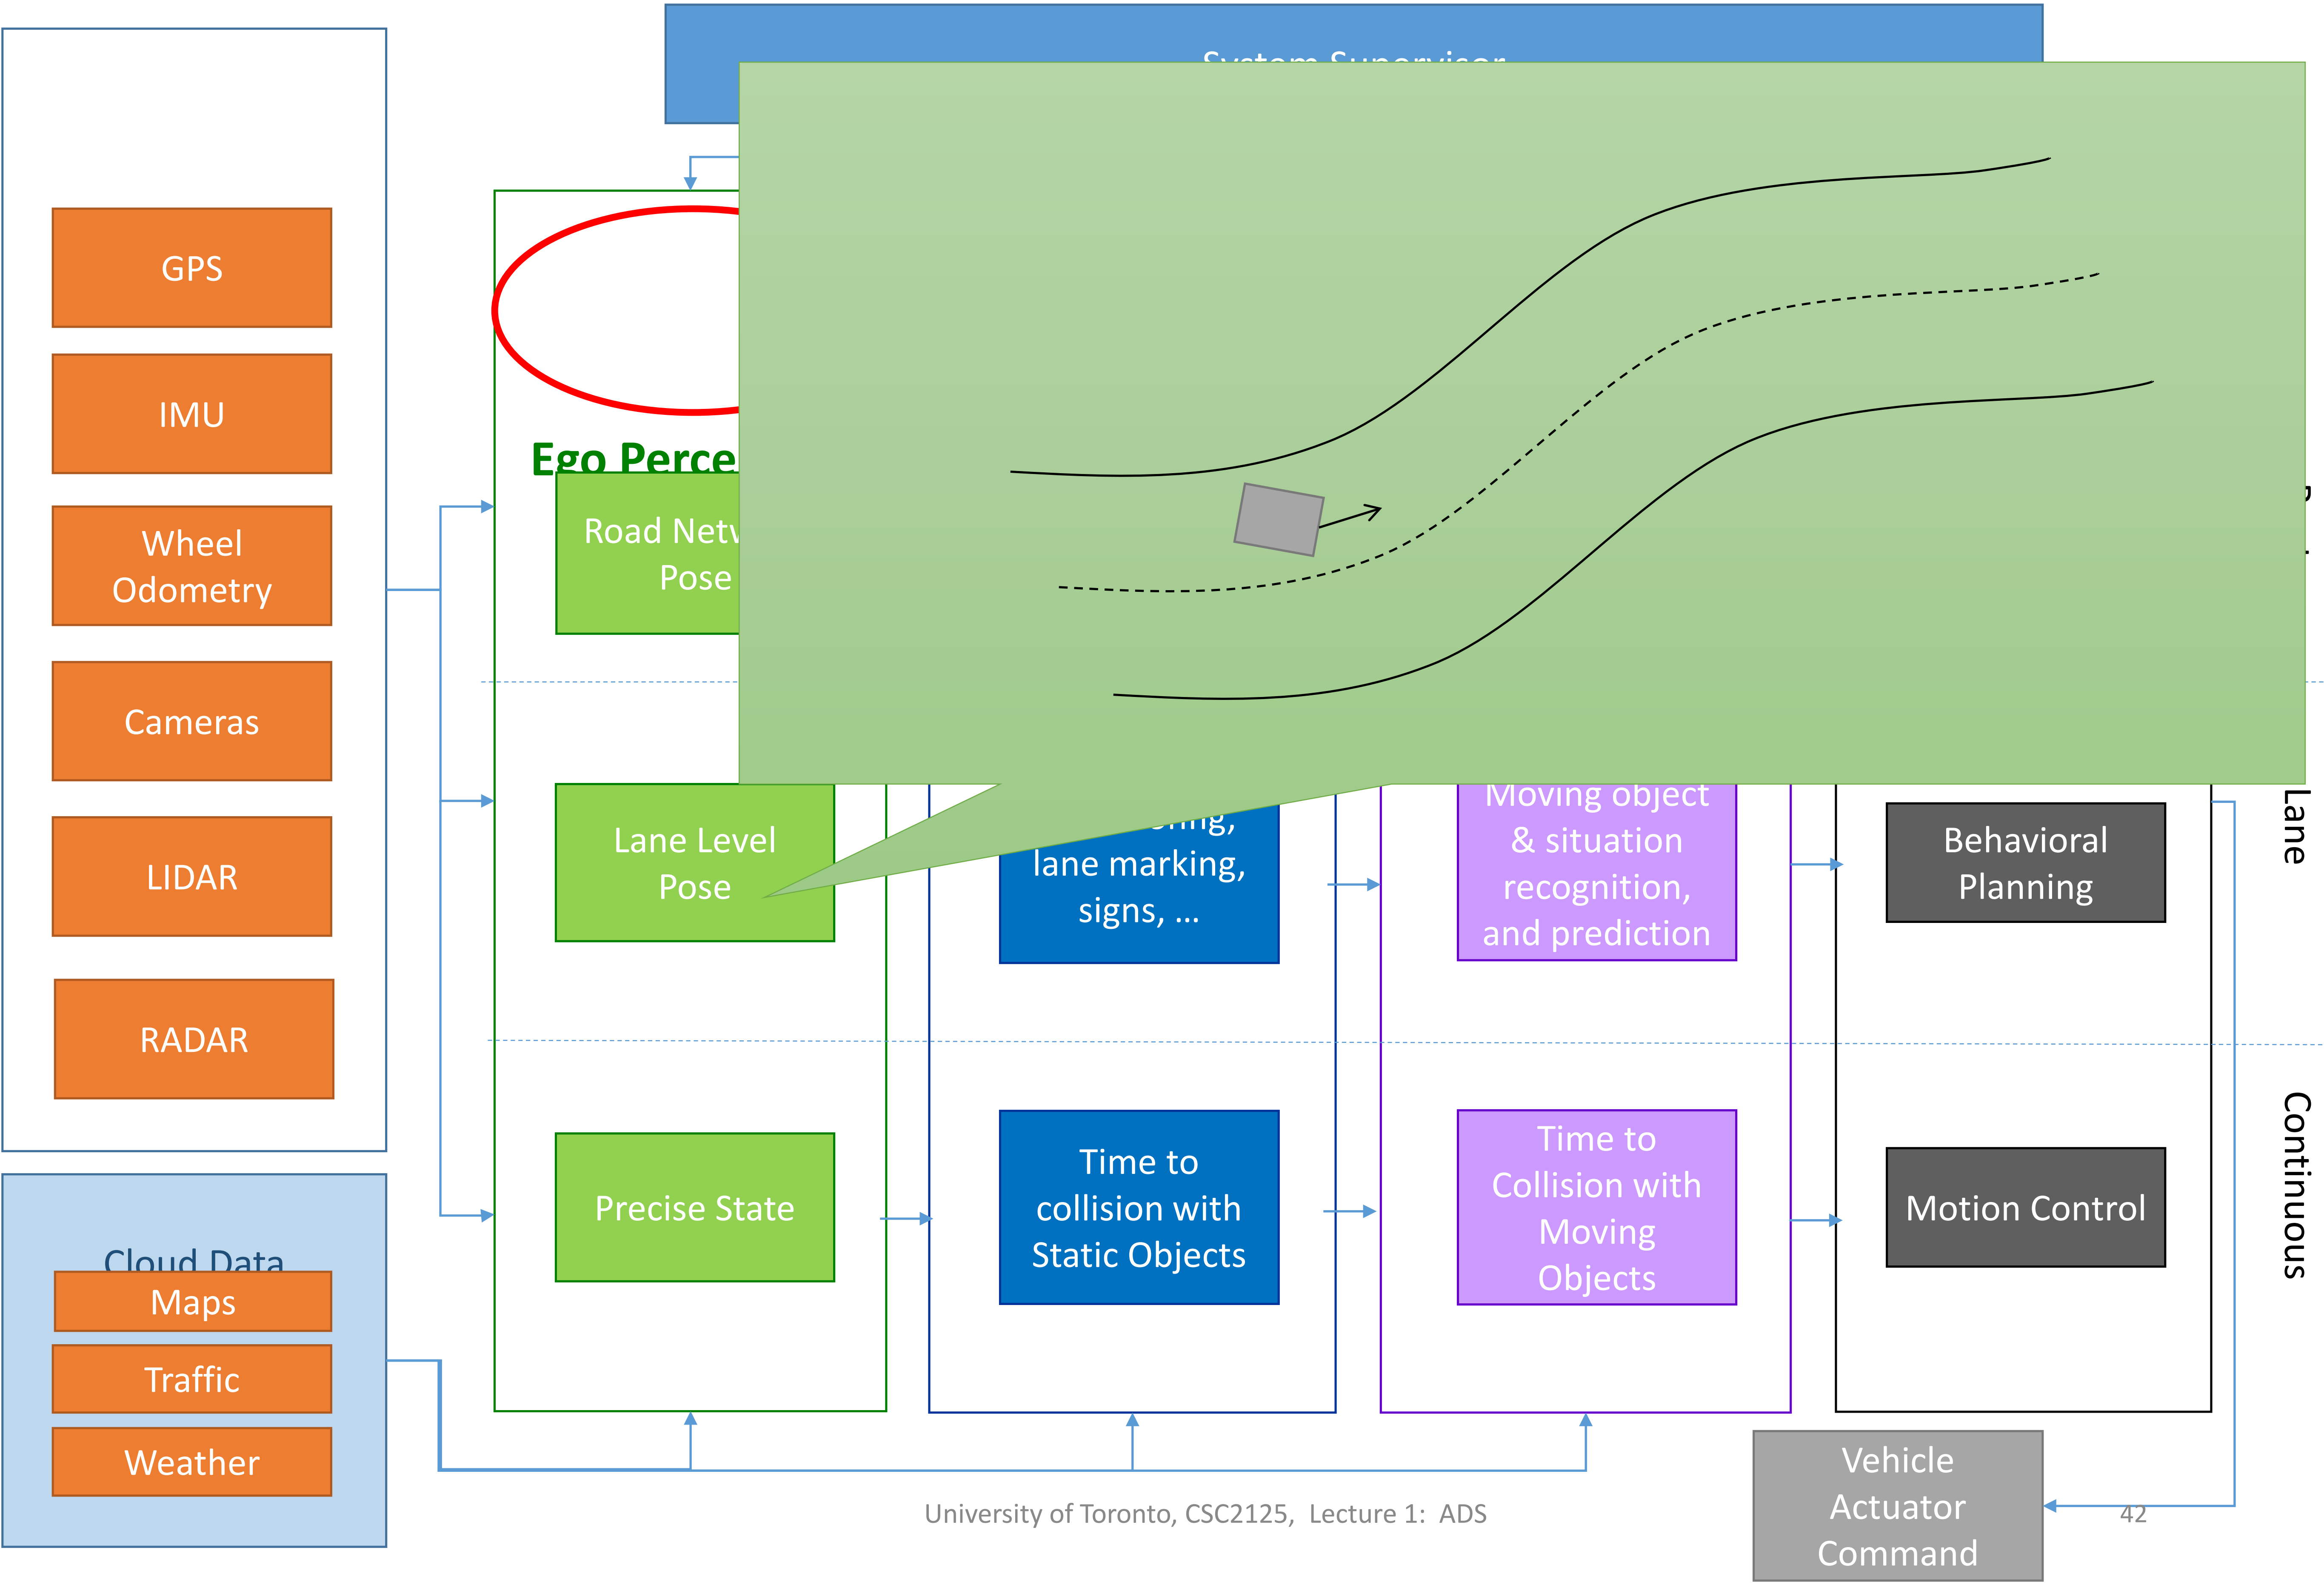

Functional Reference Architecture

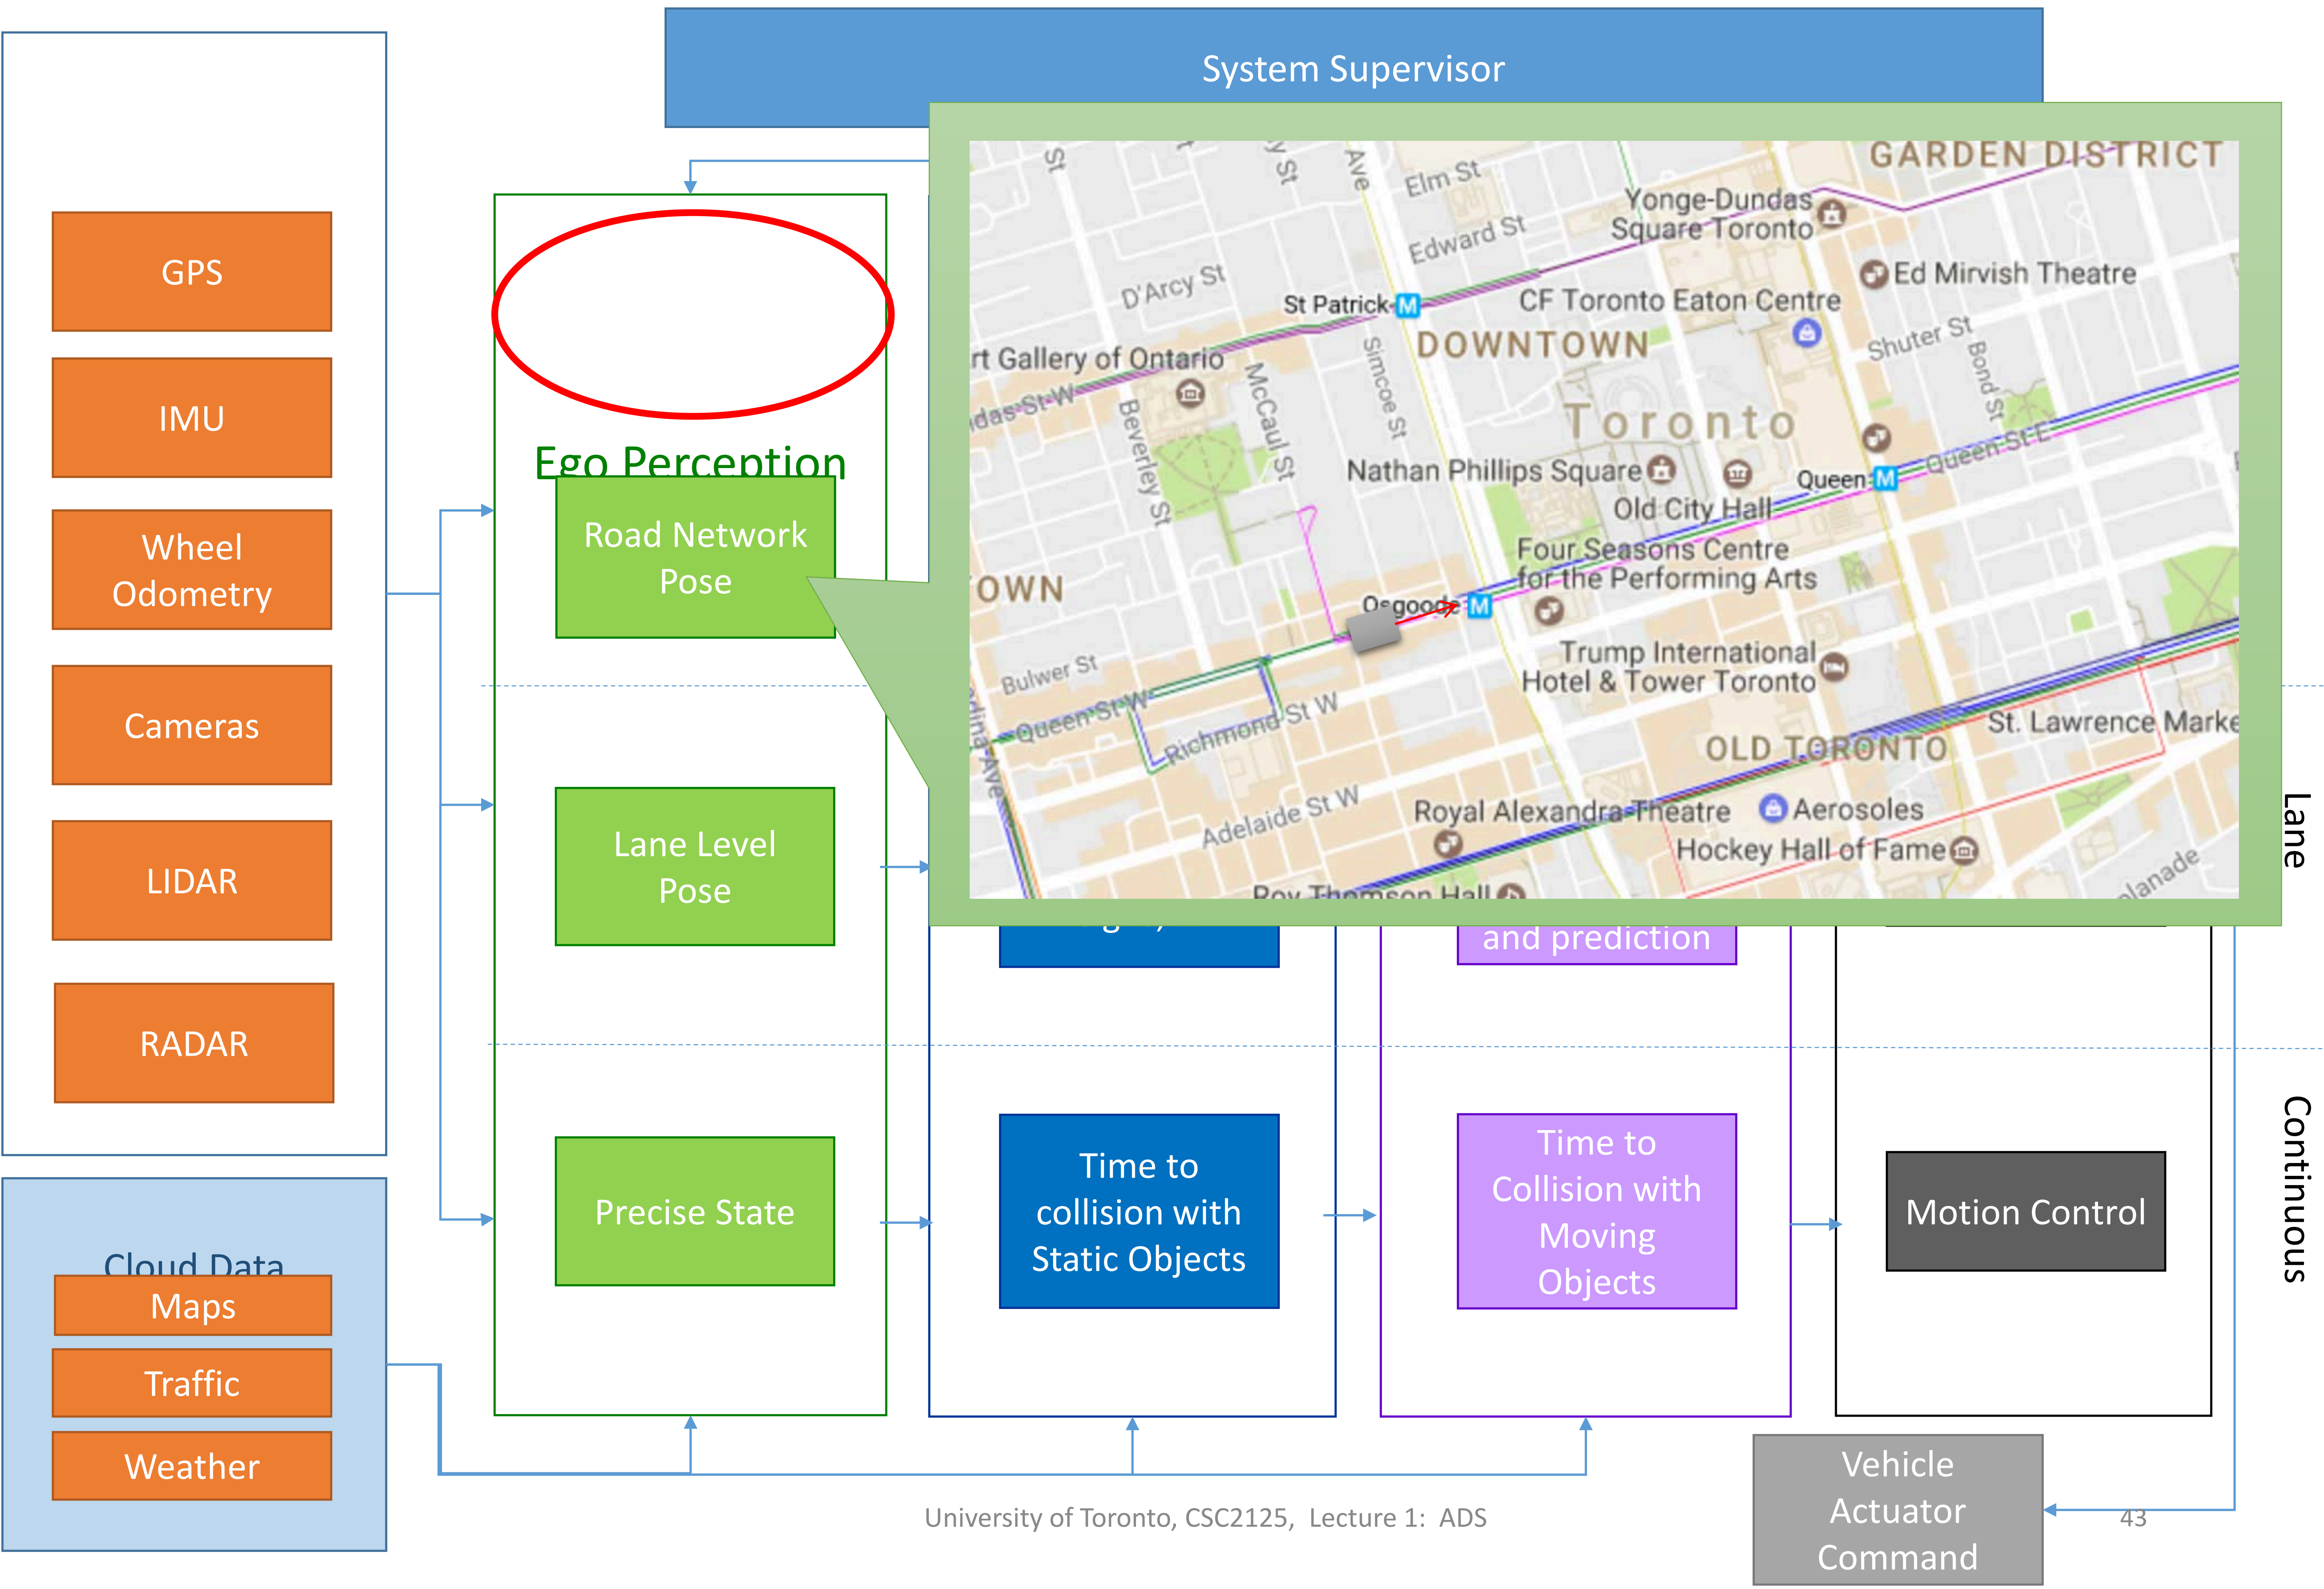

Functional Reference Architecture

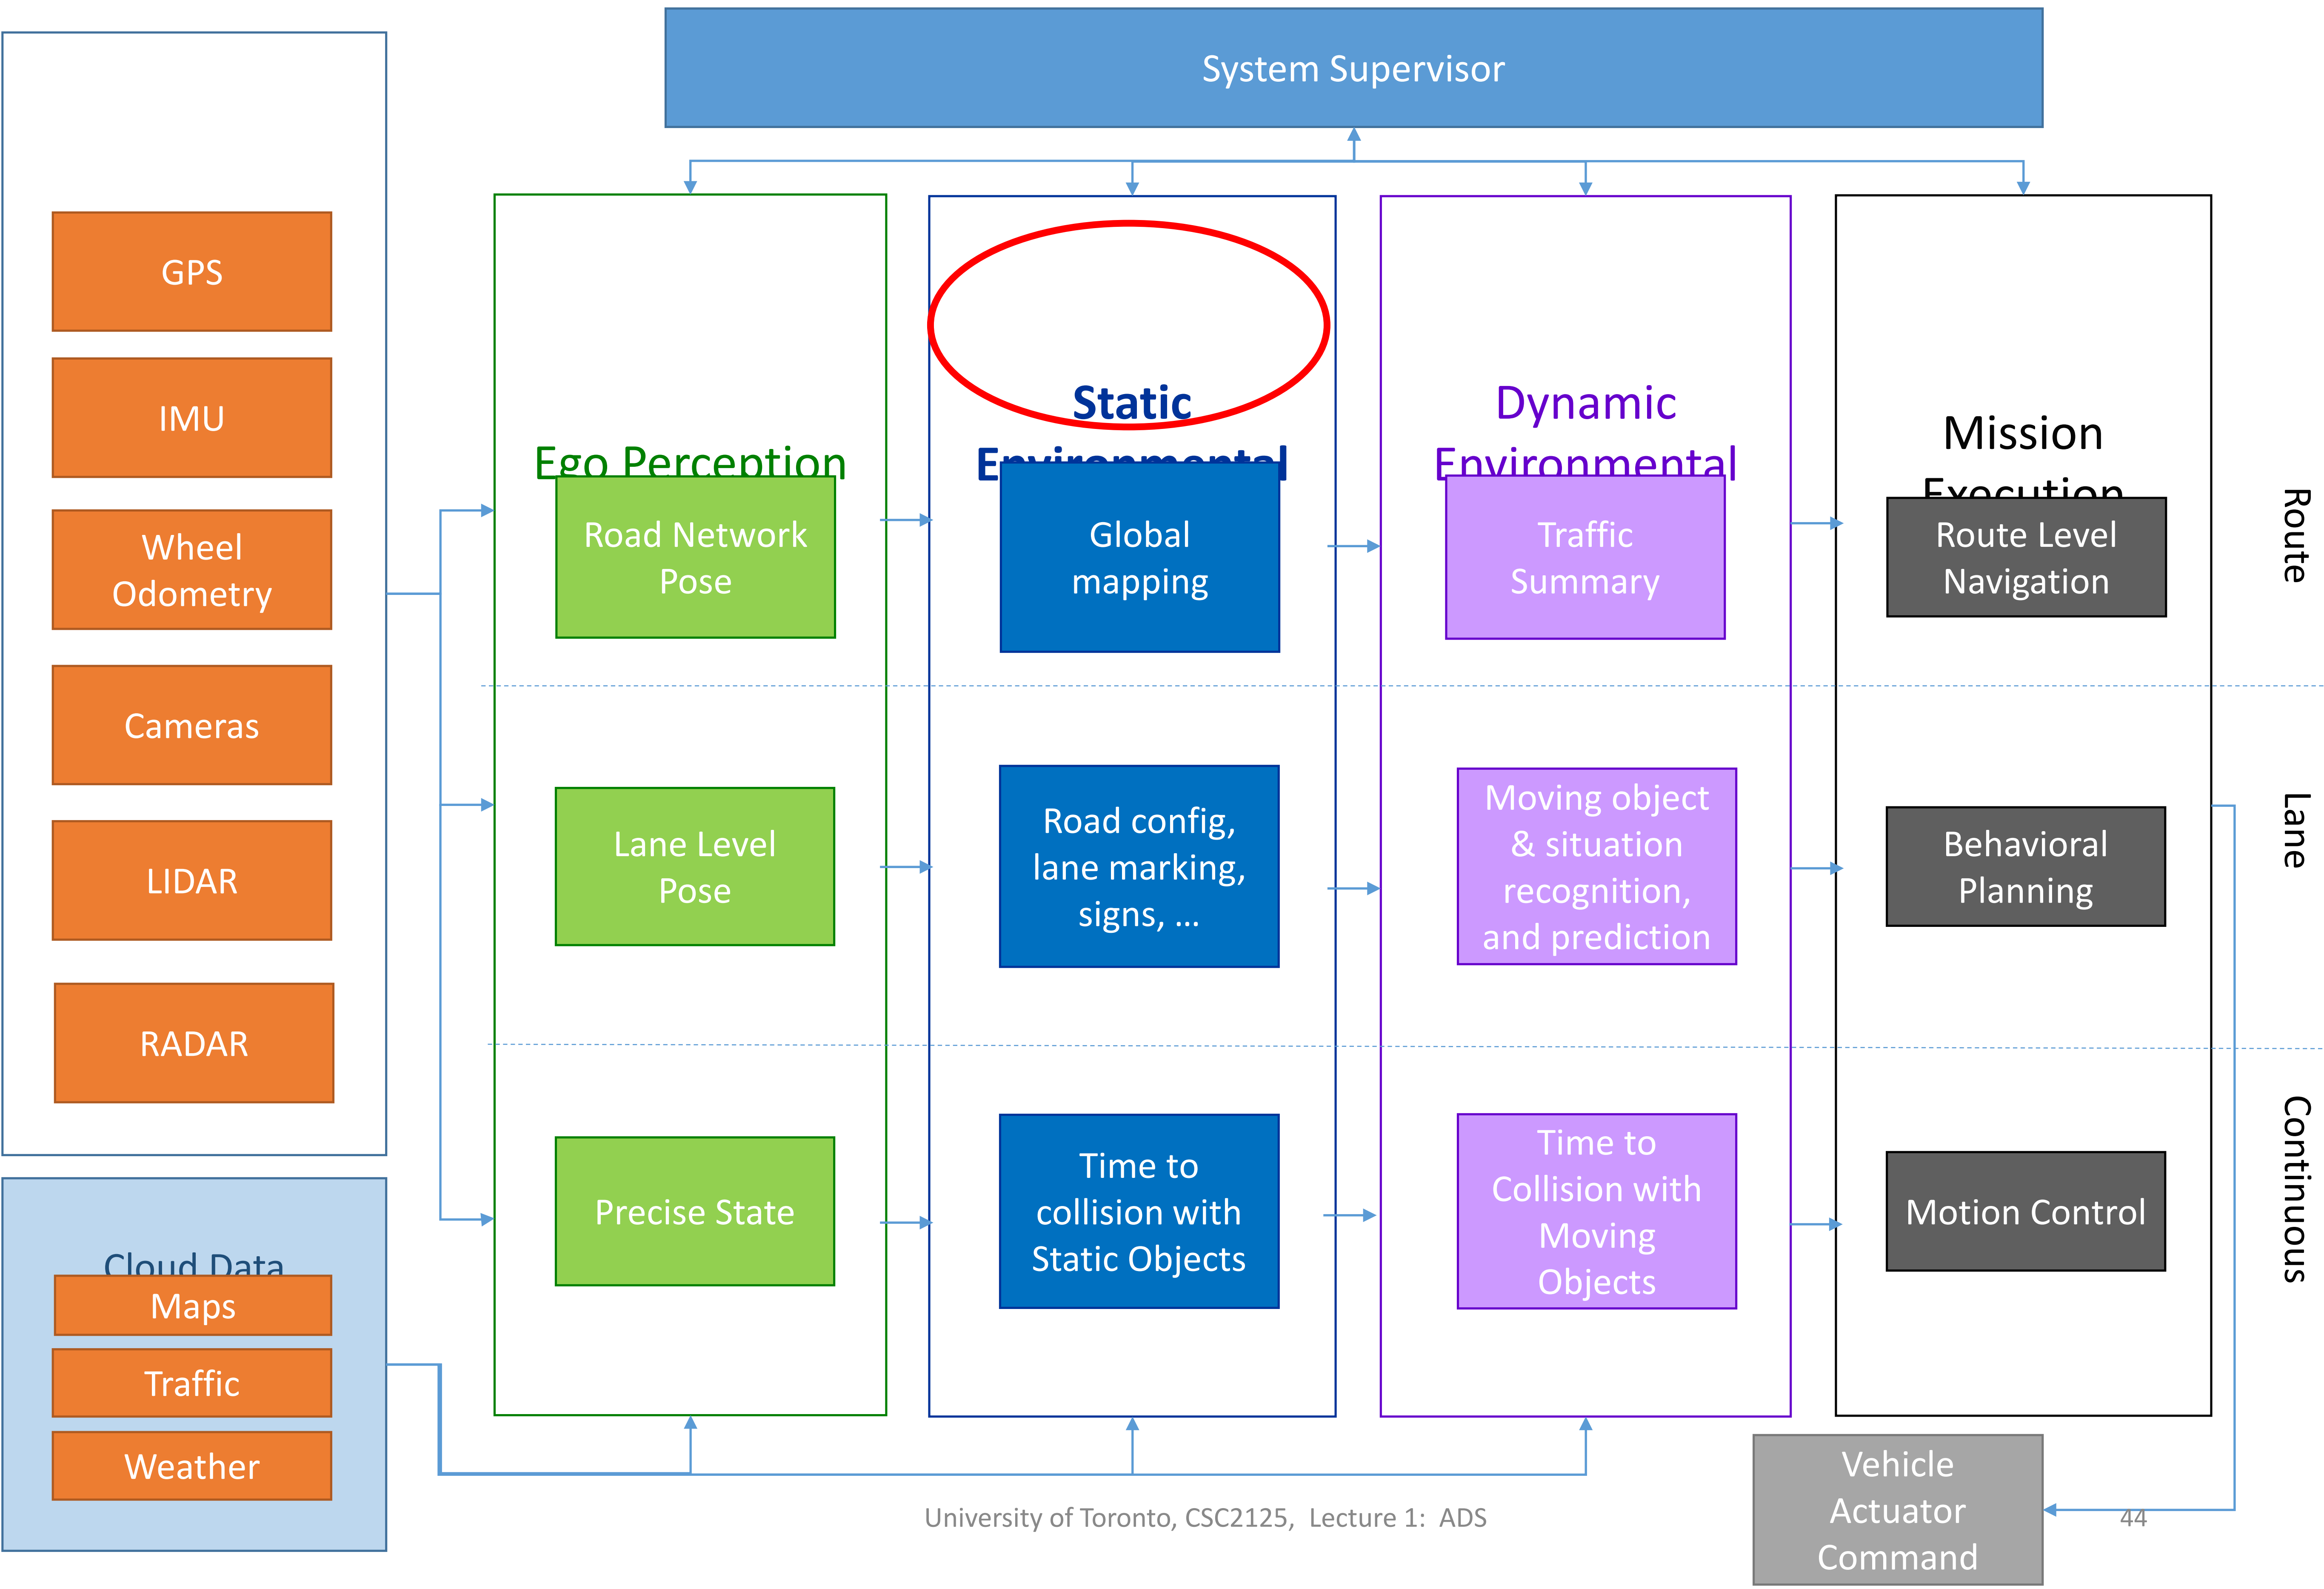

Functional Reference Architecture

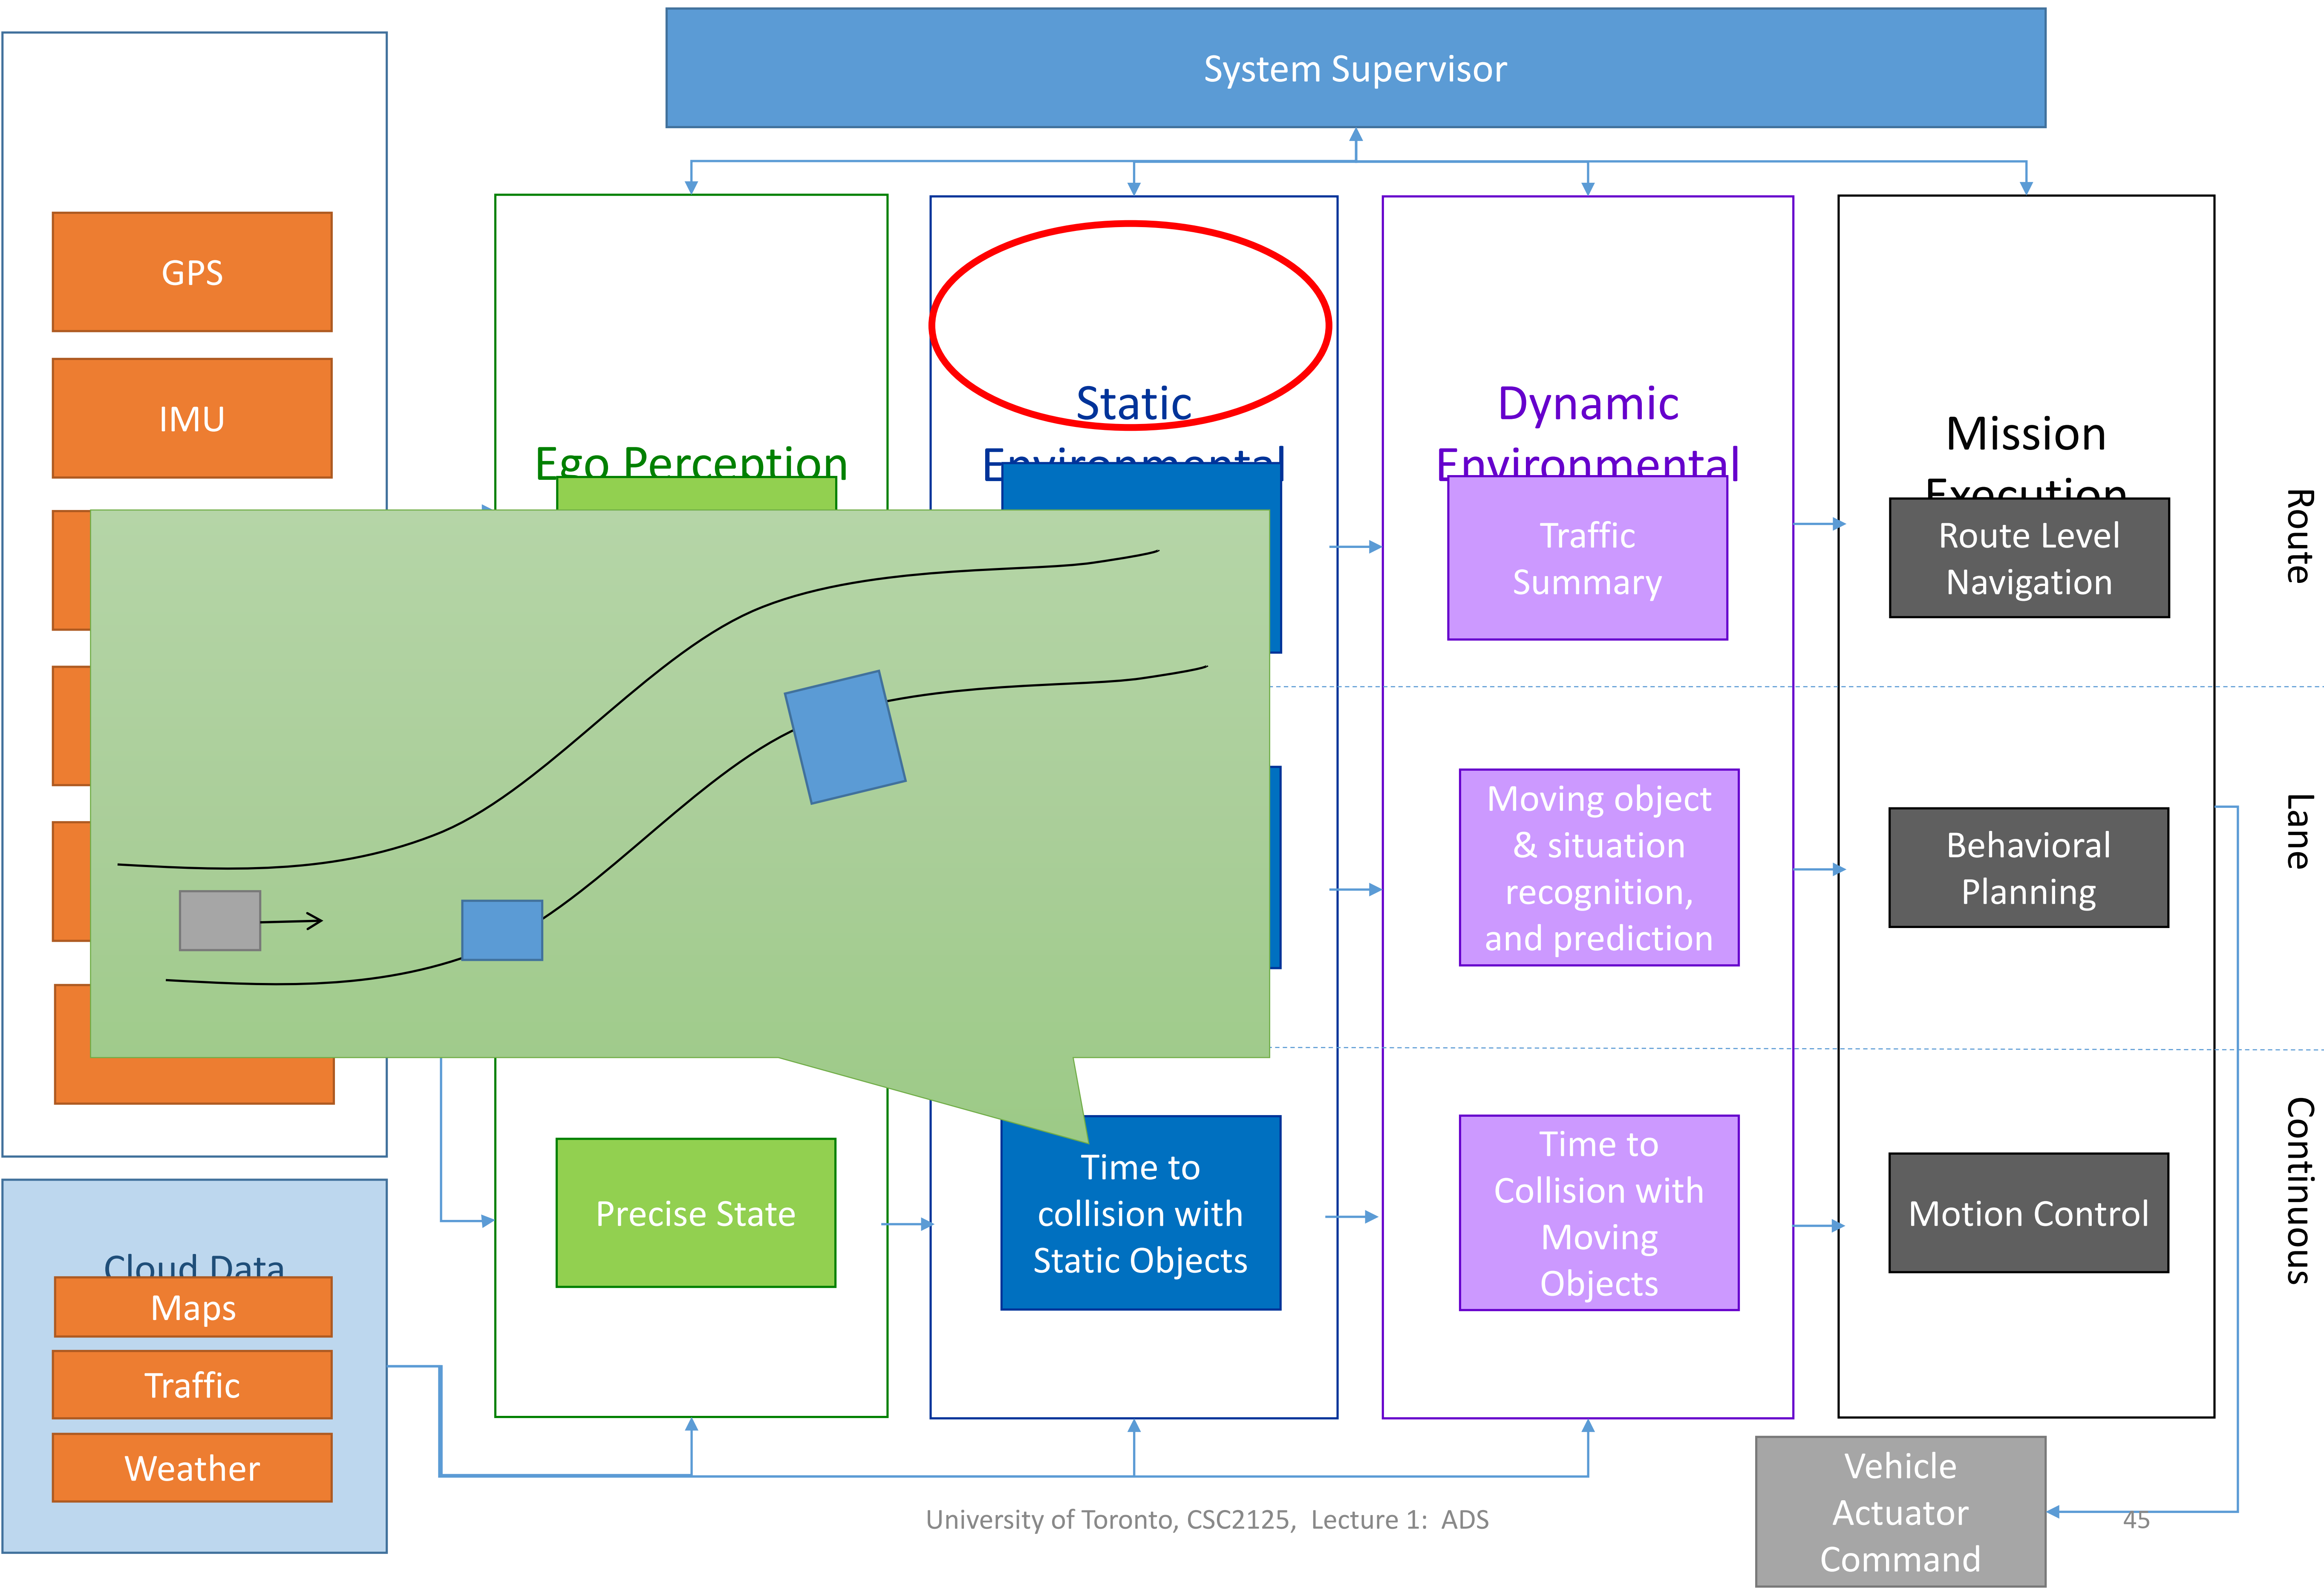

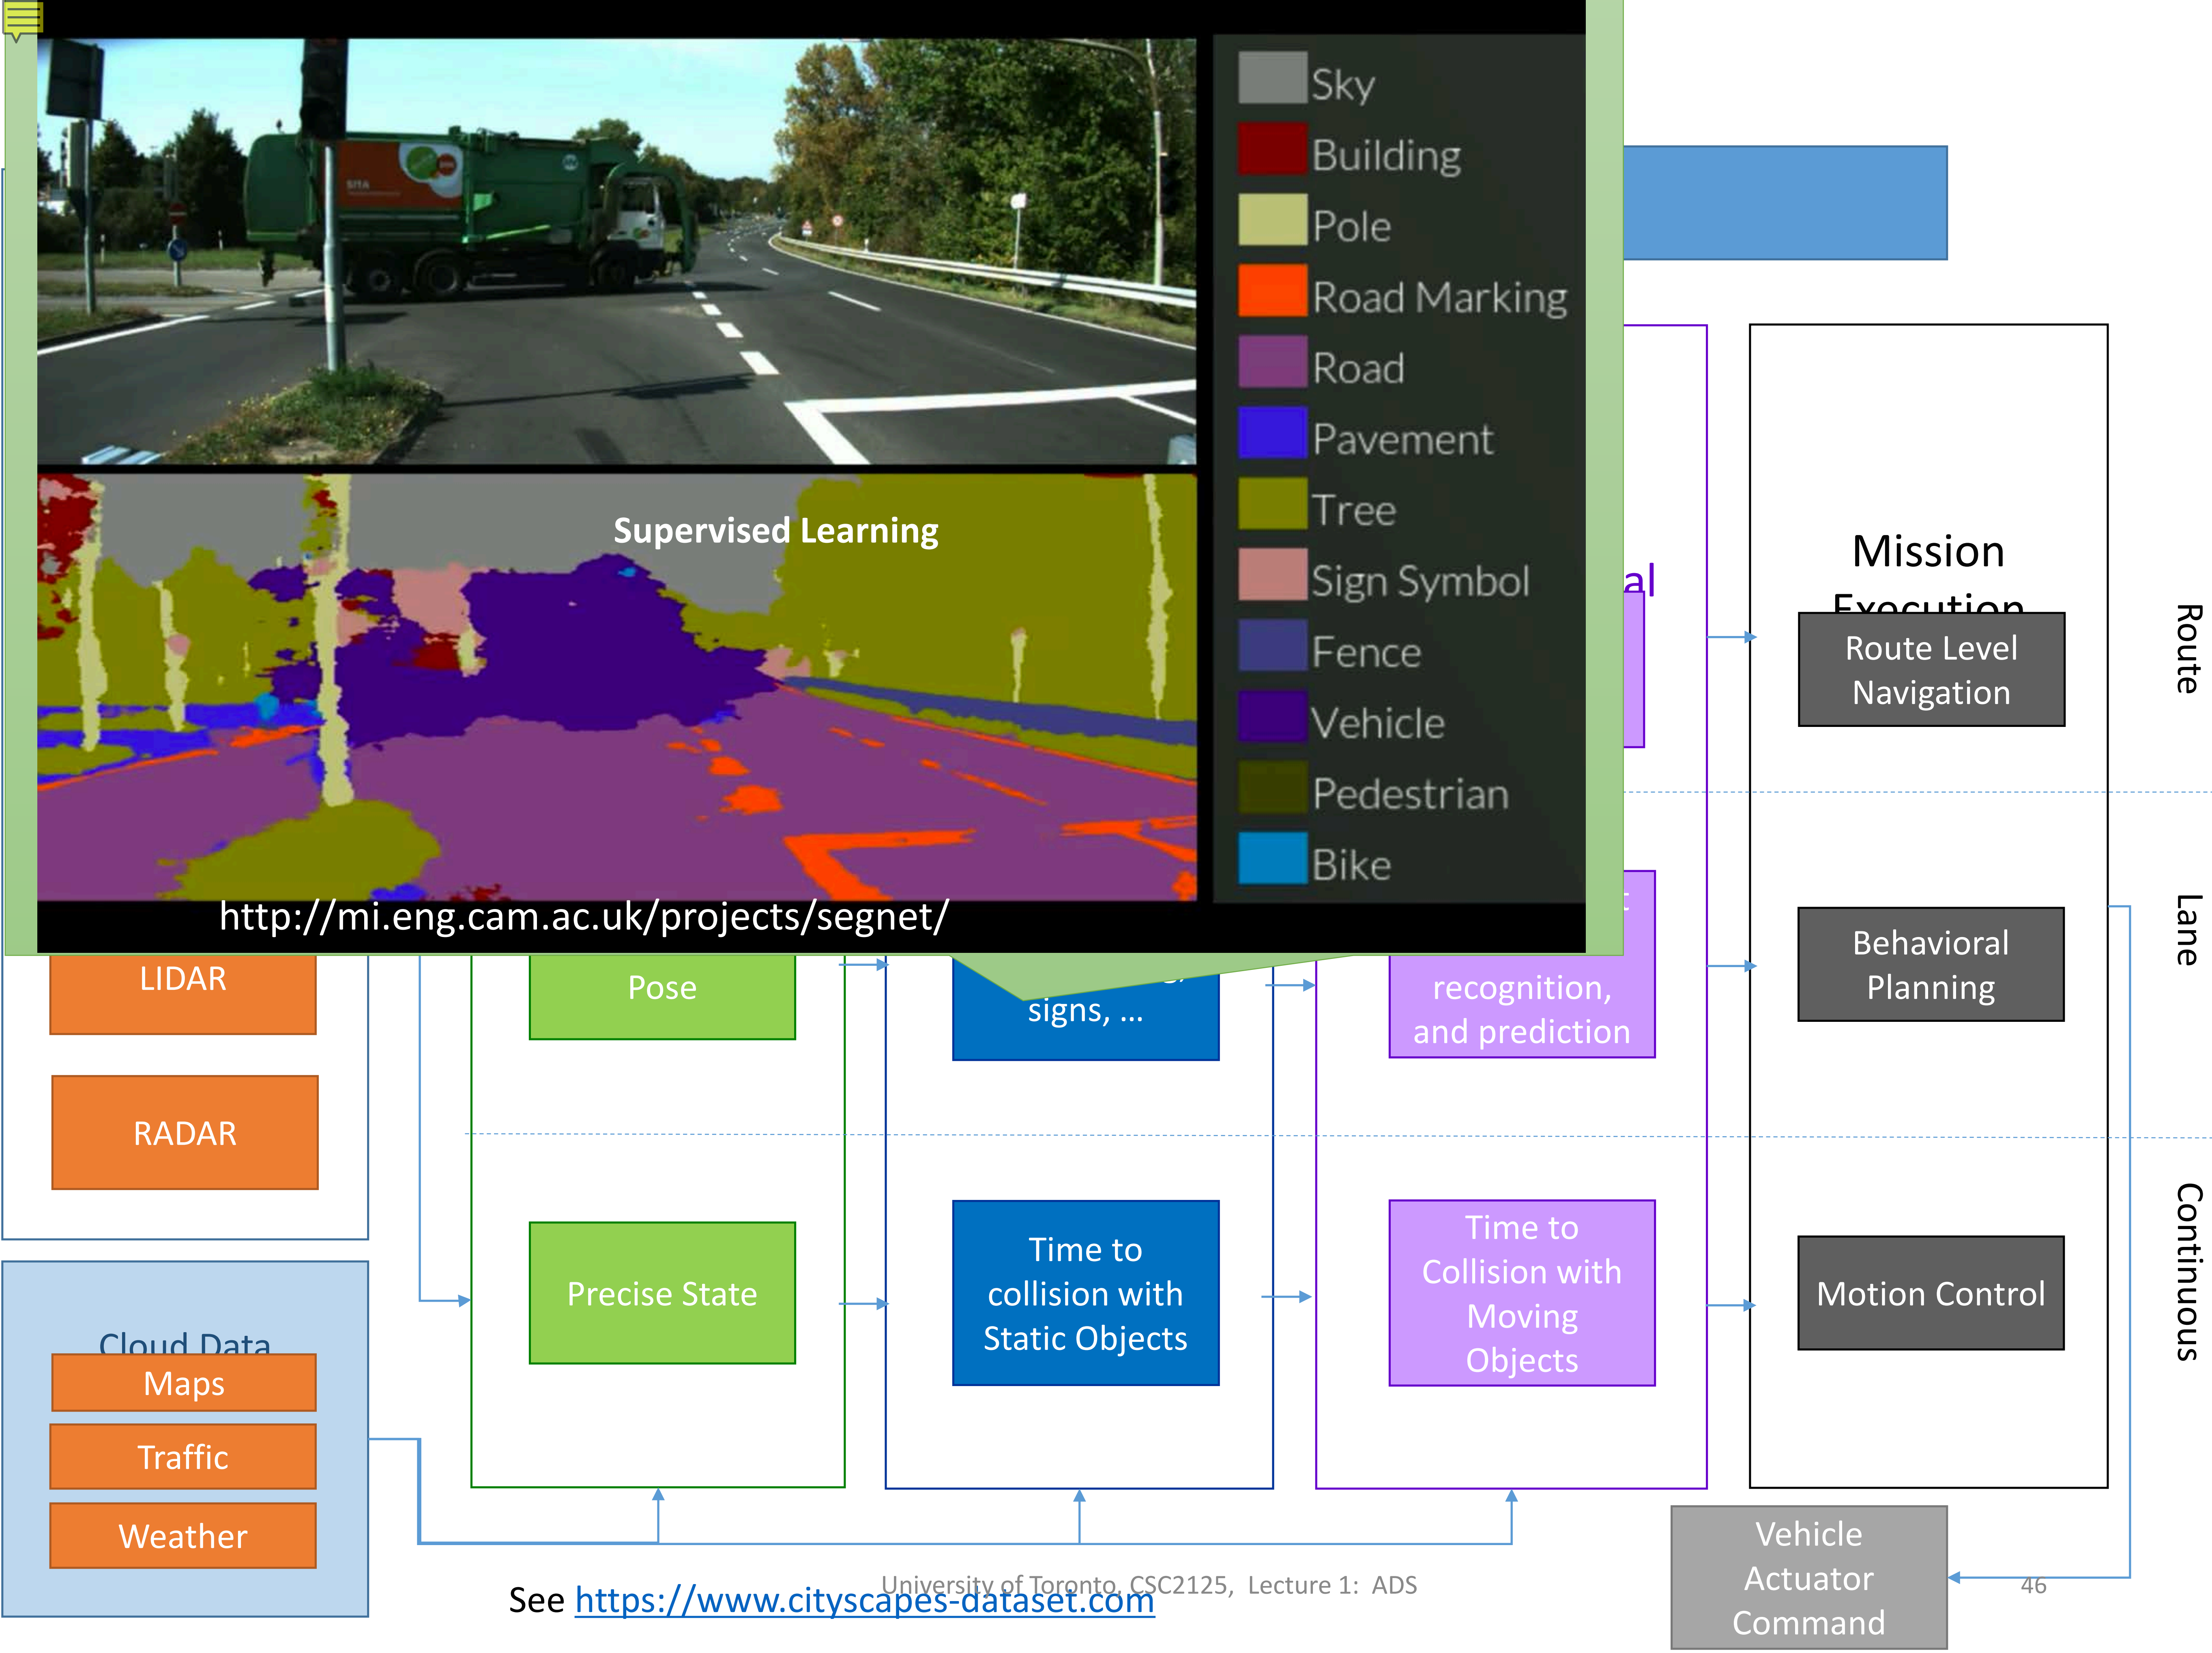

See https://www.cityscapes-dataset.com

Functional Reference Architecture

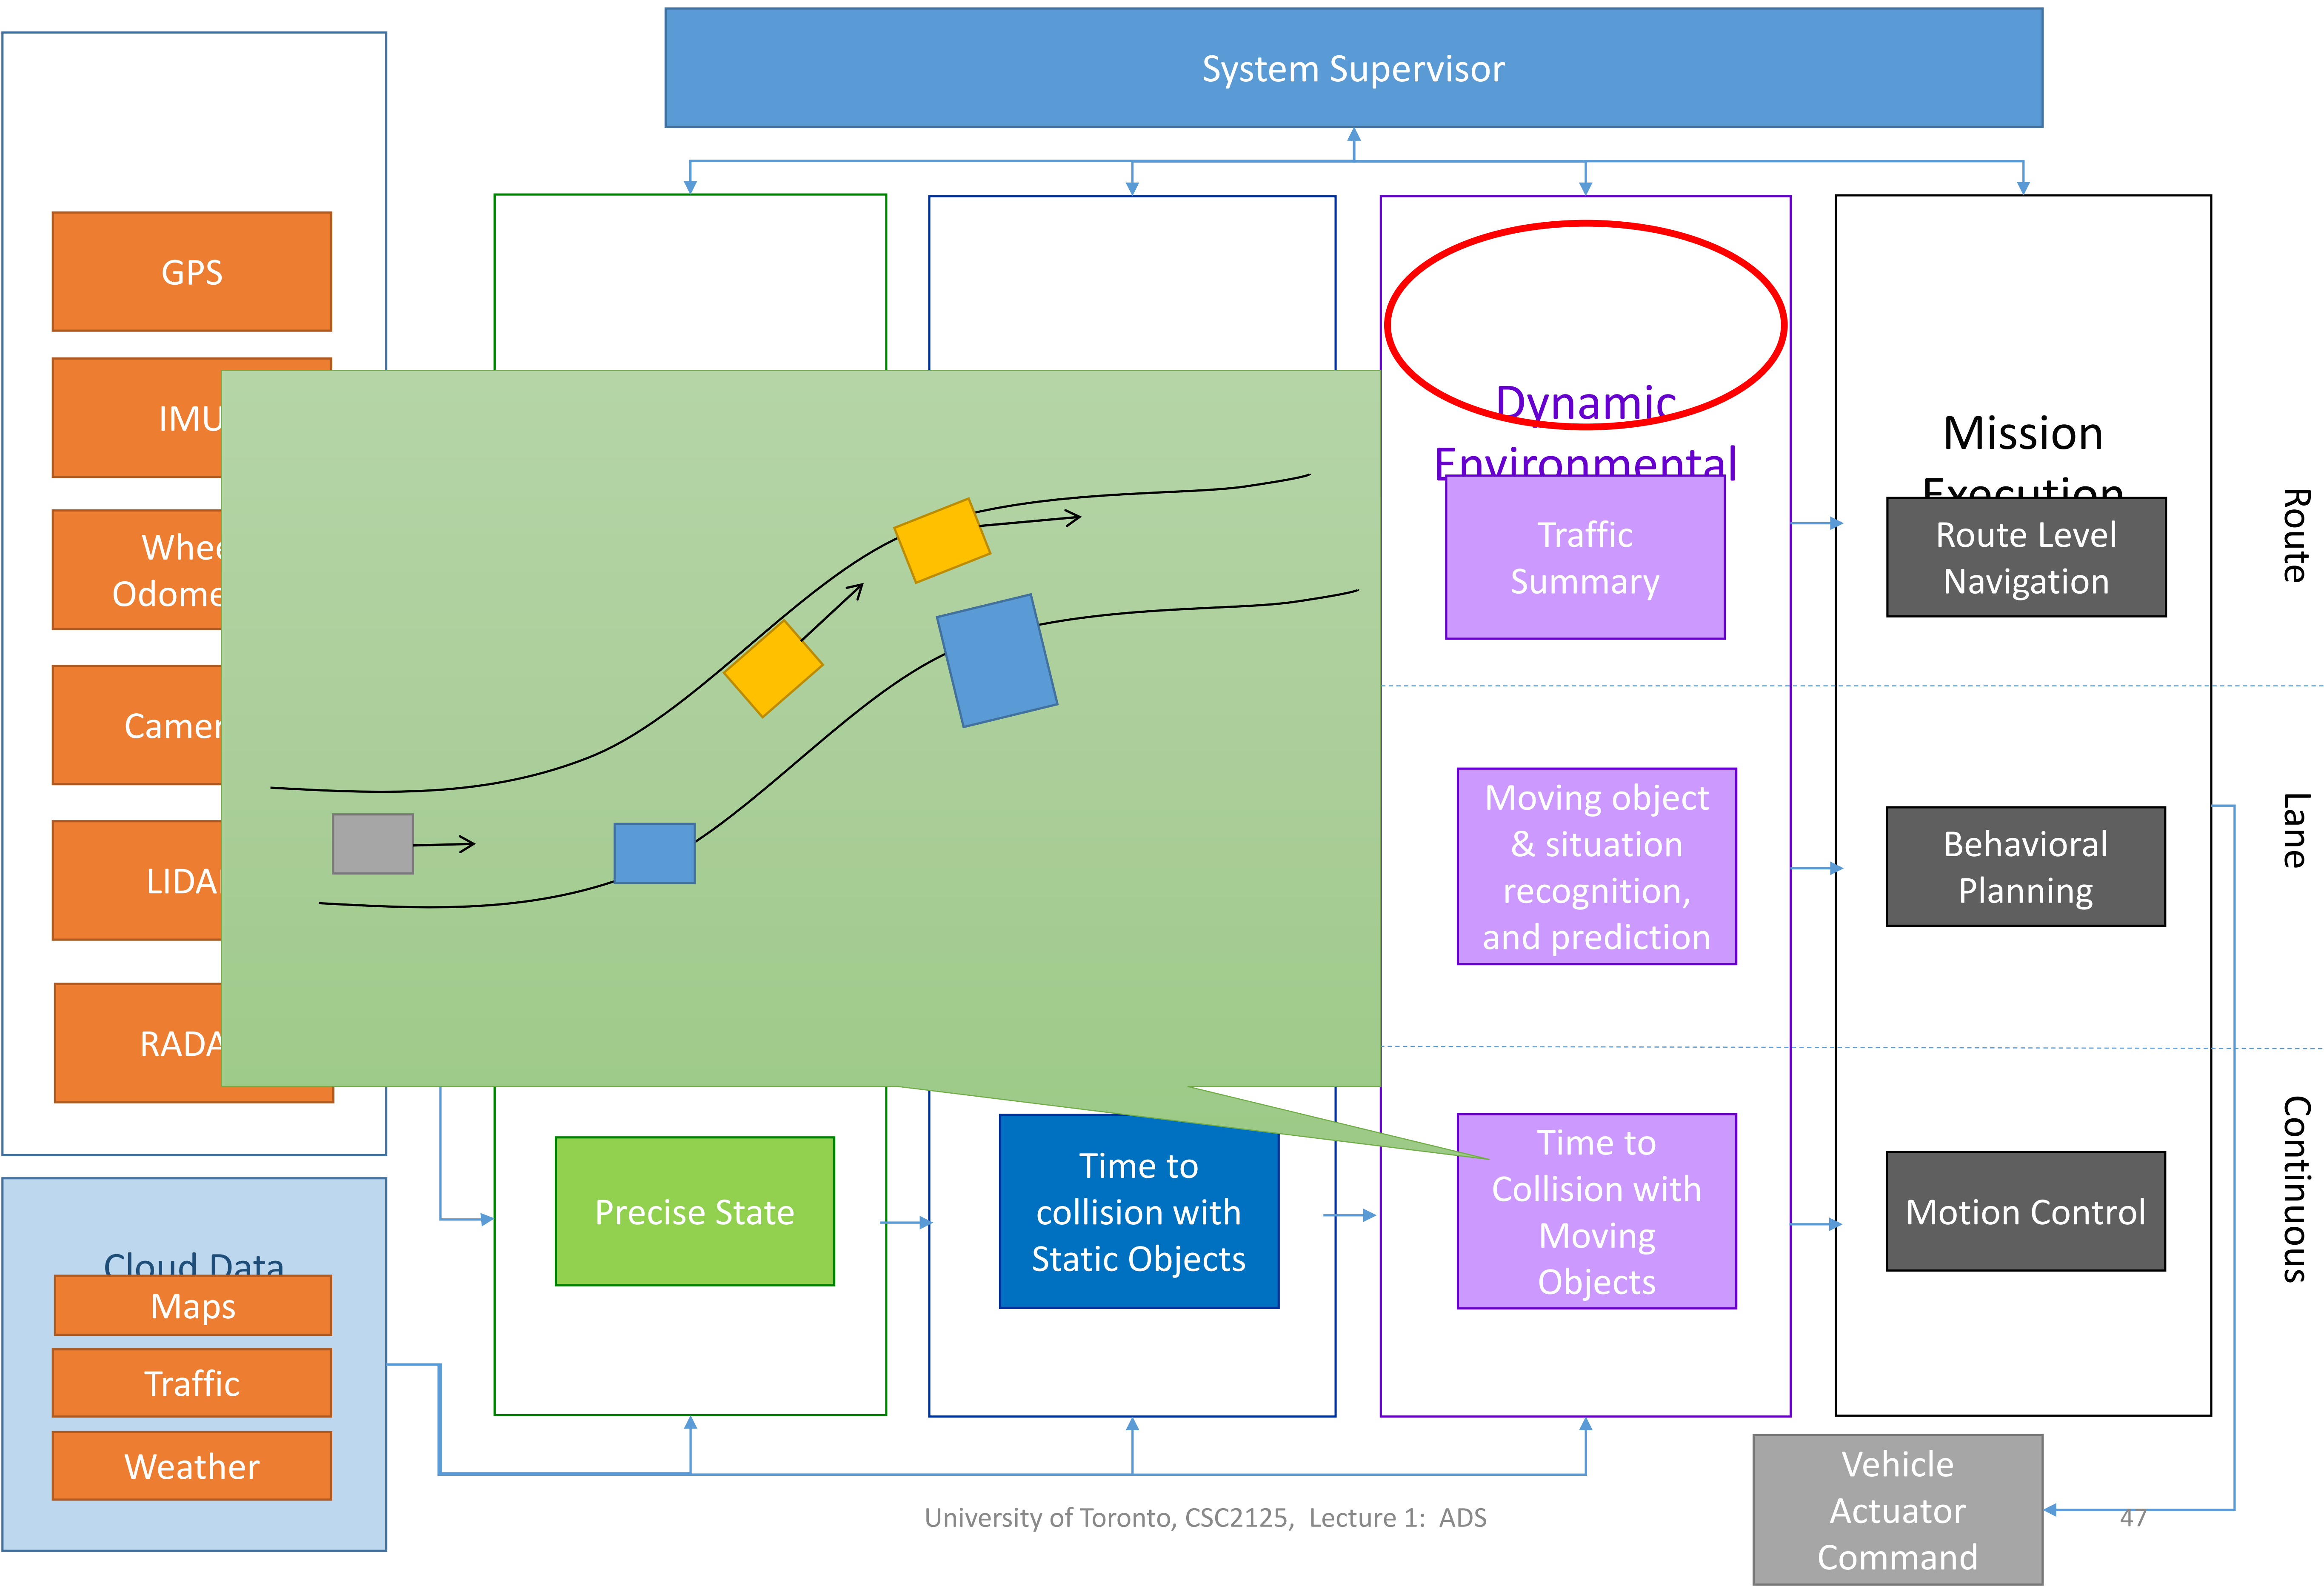

ture

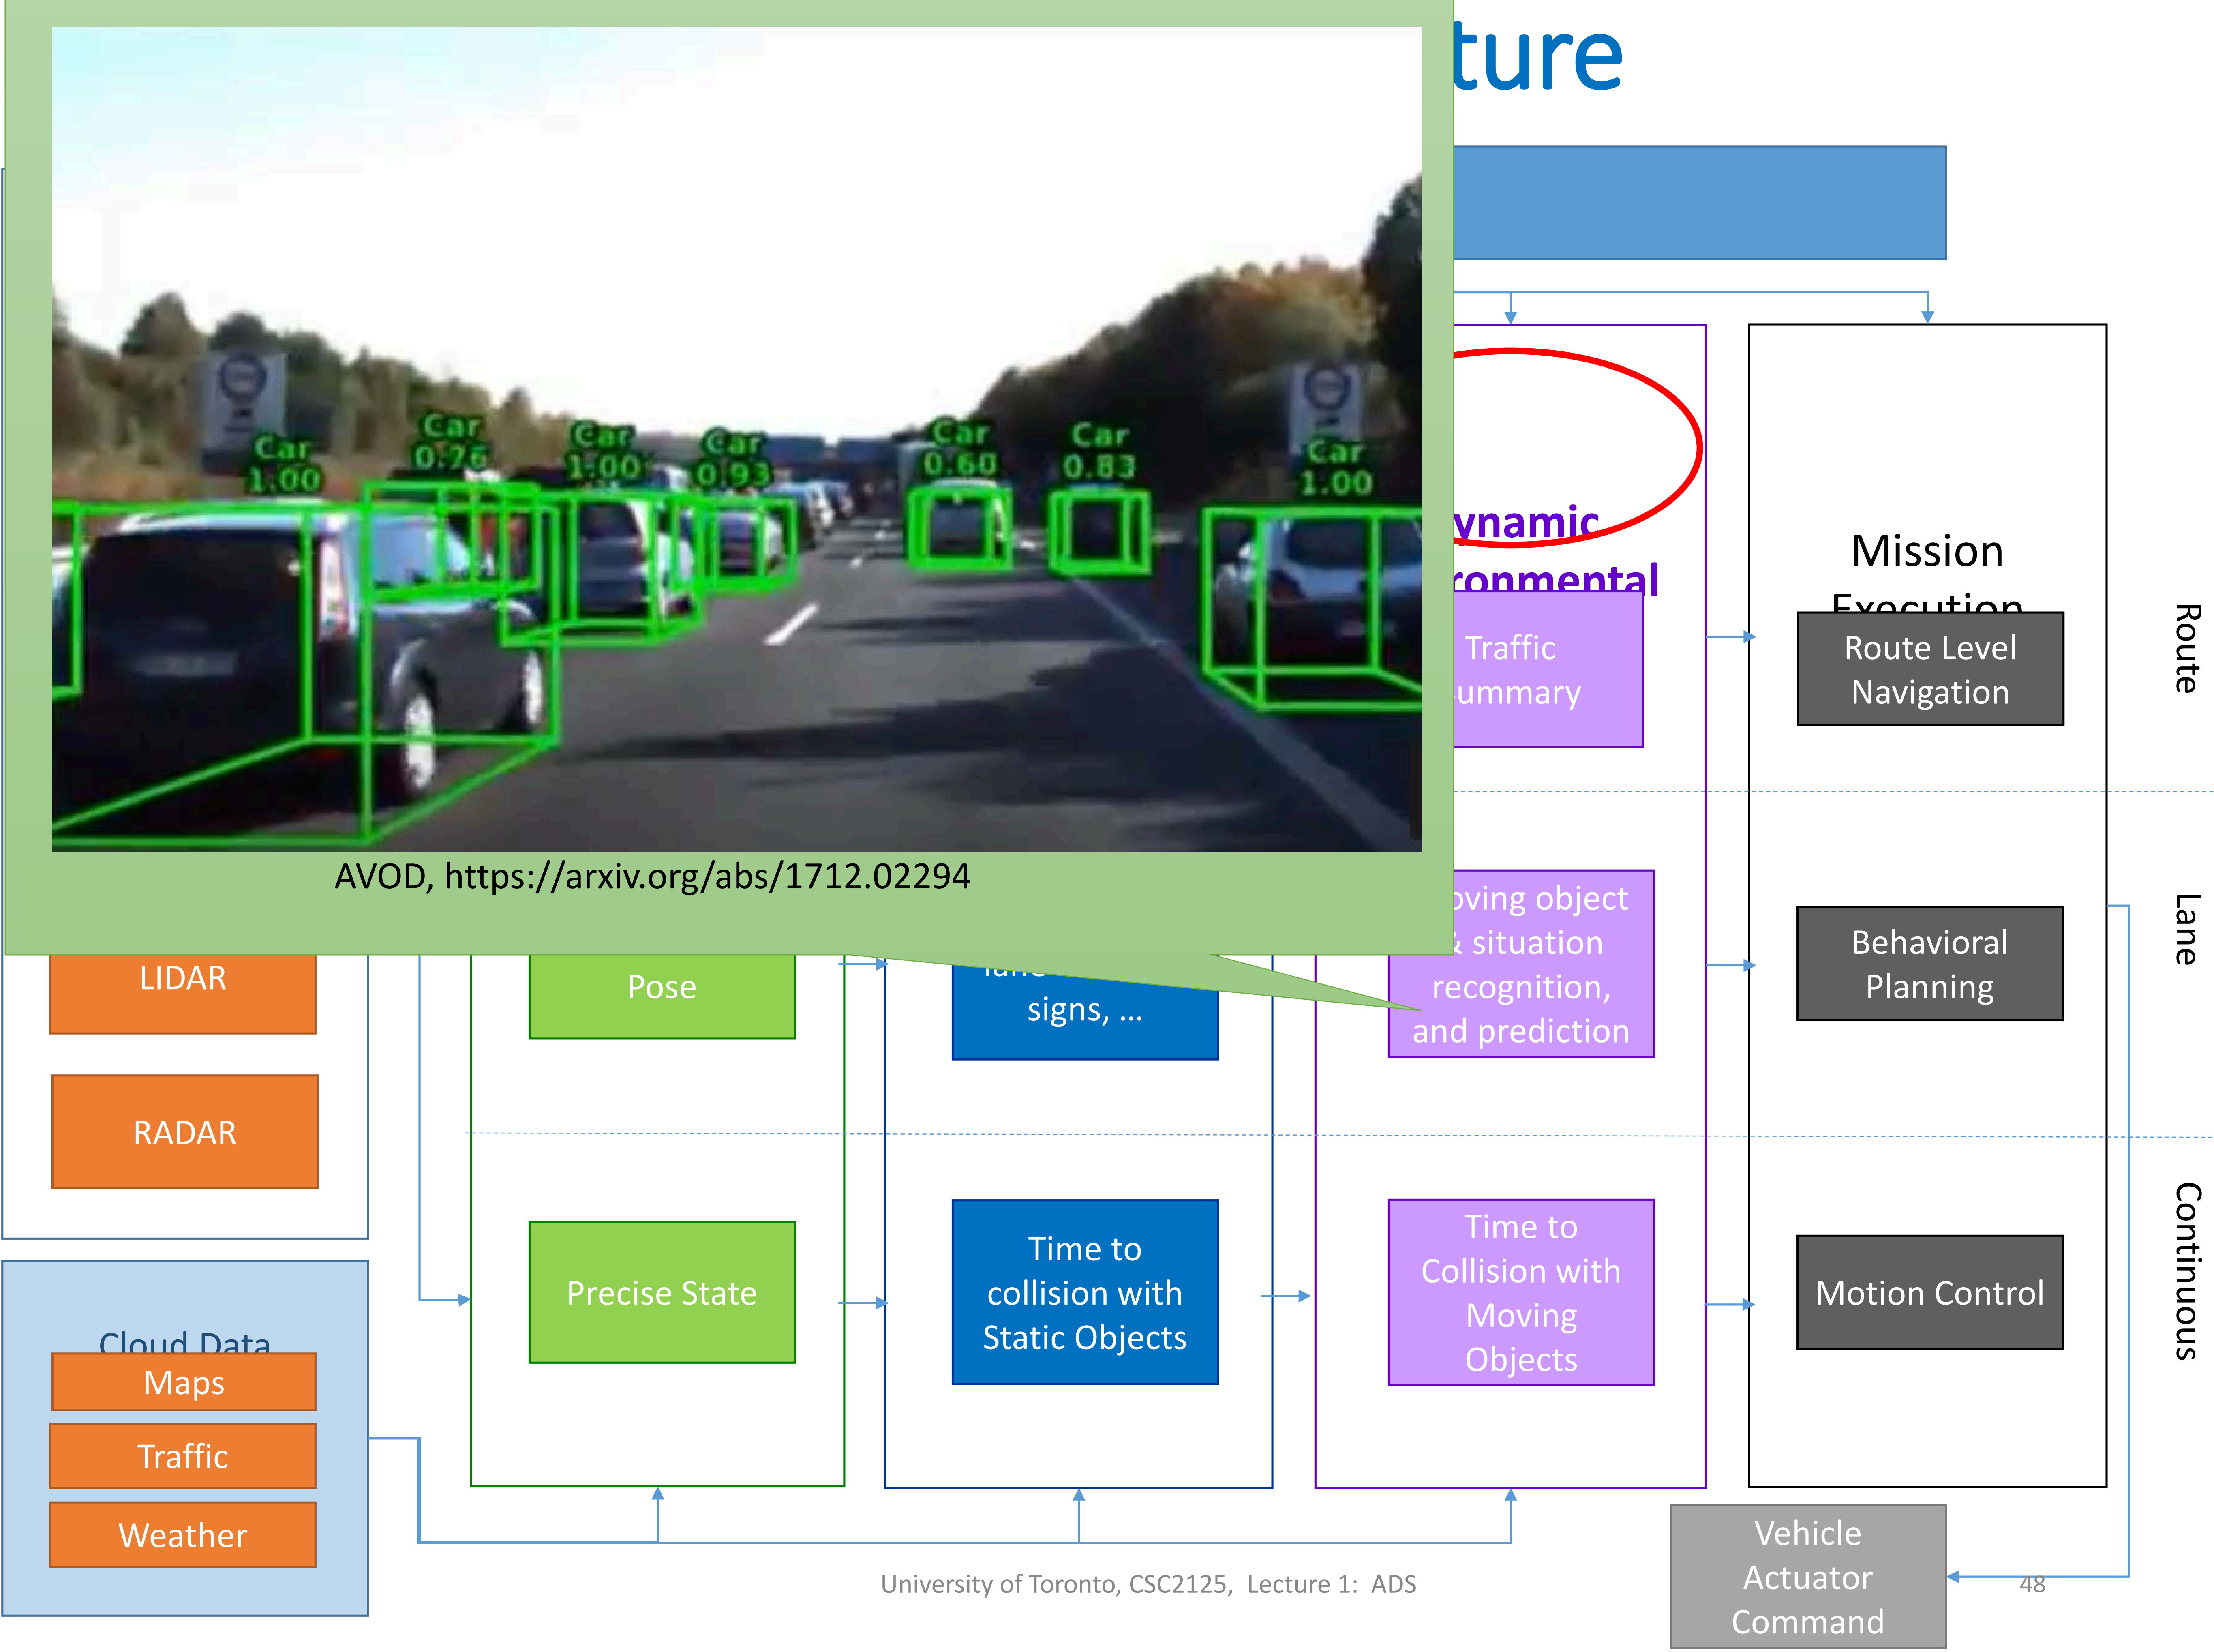

AVOD, https://arxiv.org/abs/1712.02294

Functional Reference Architecture

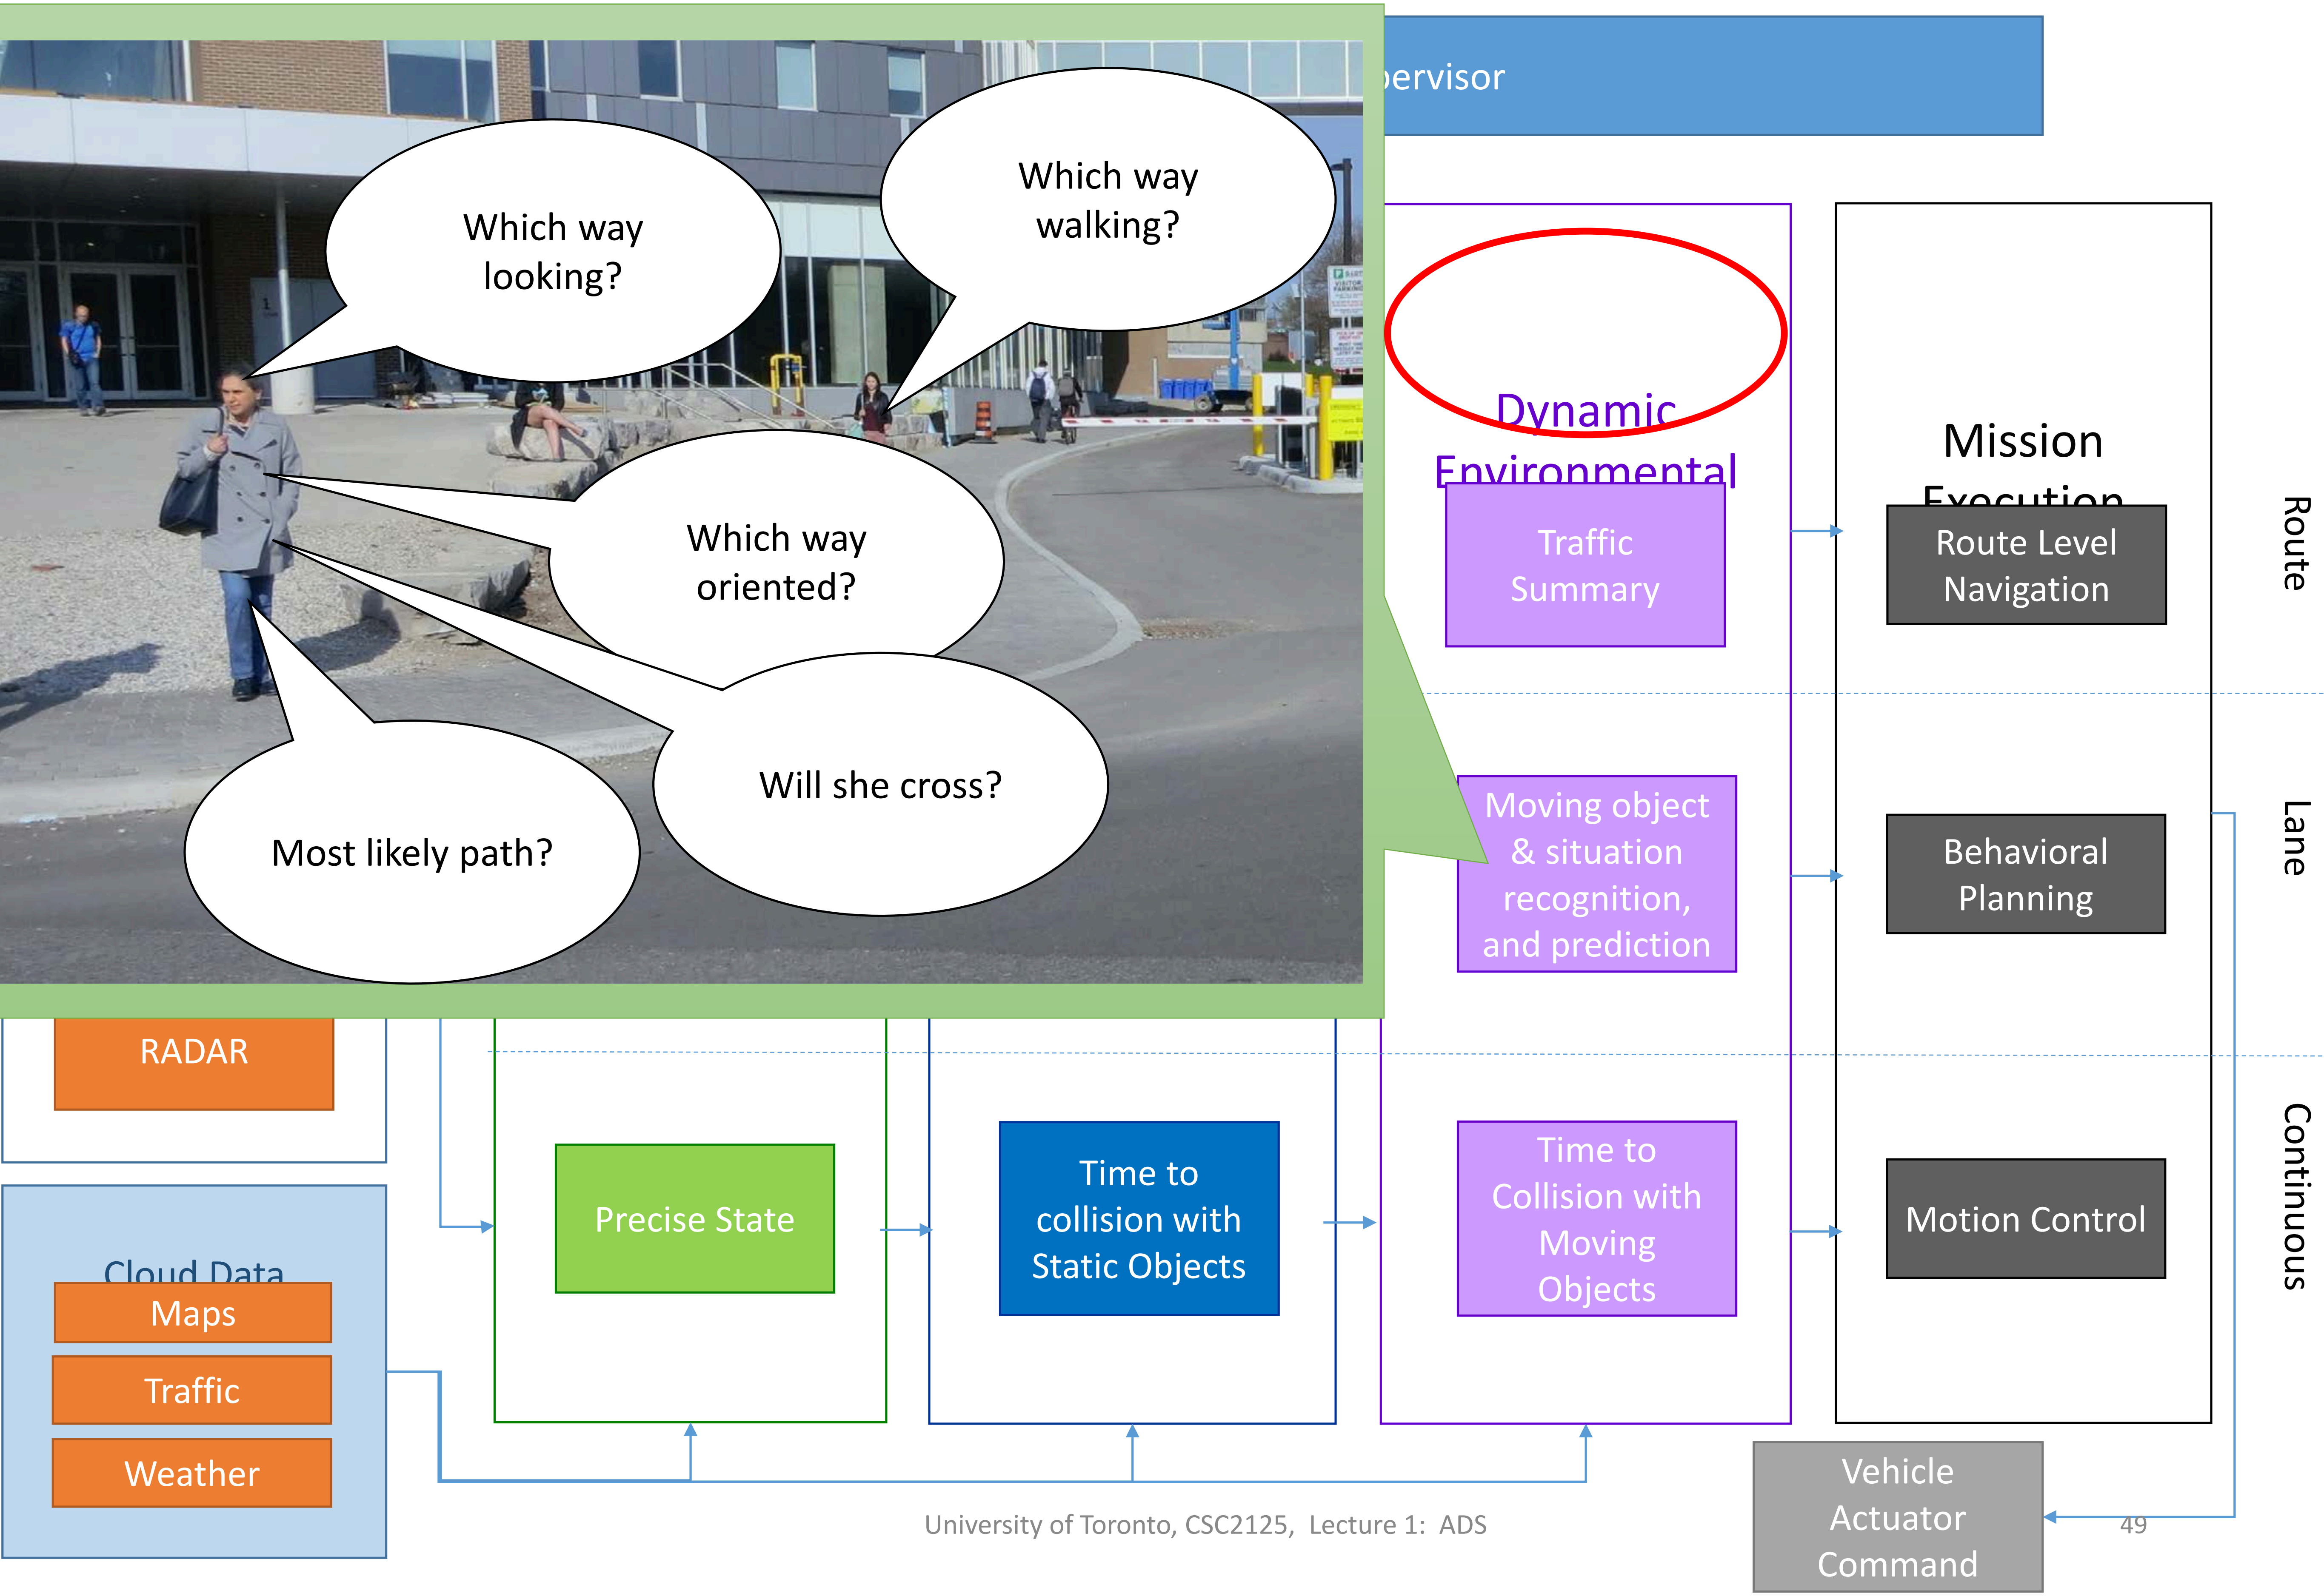

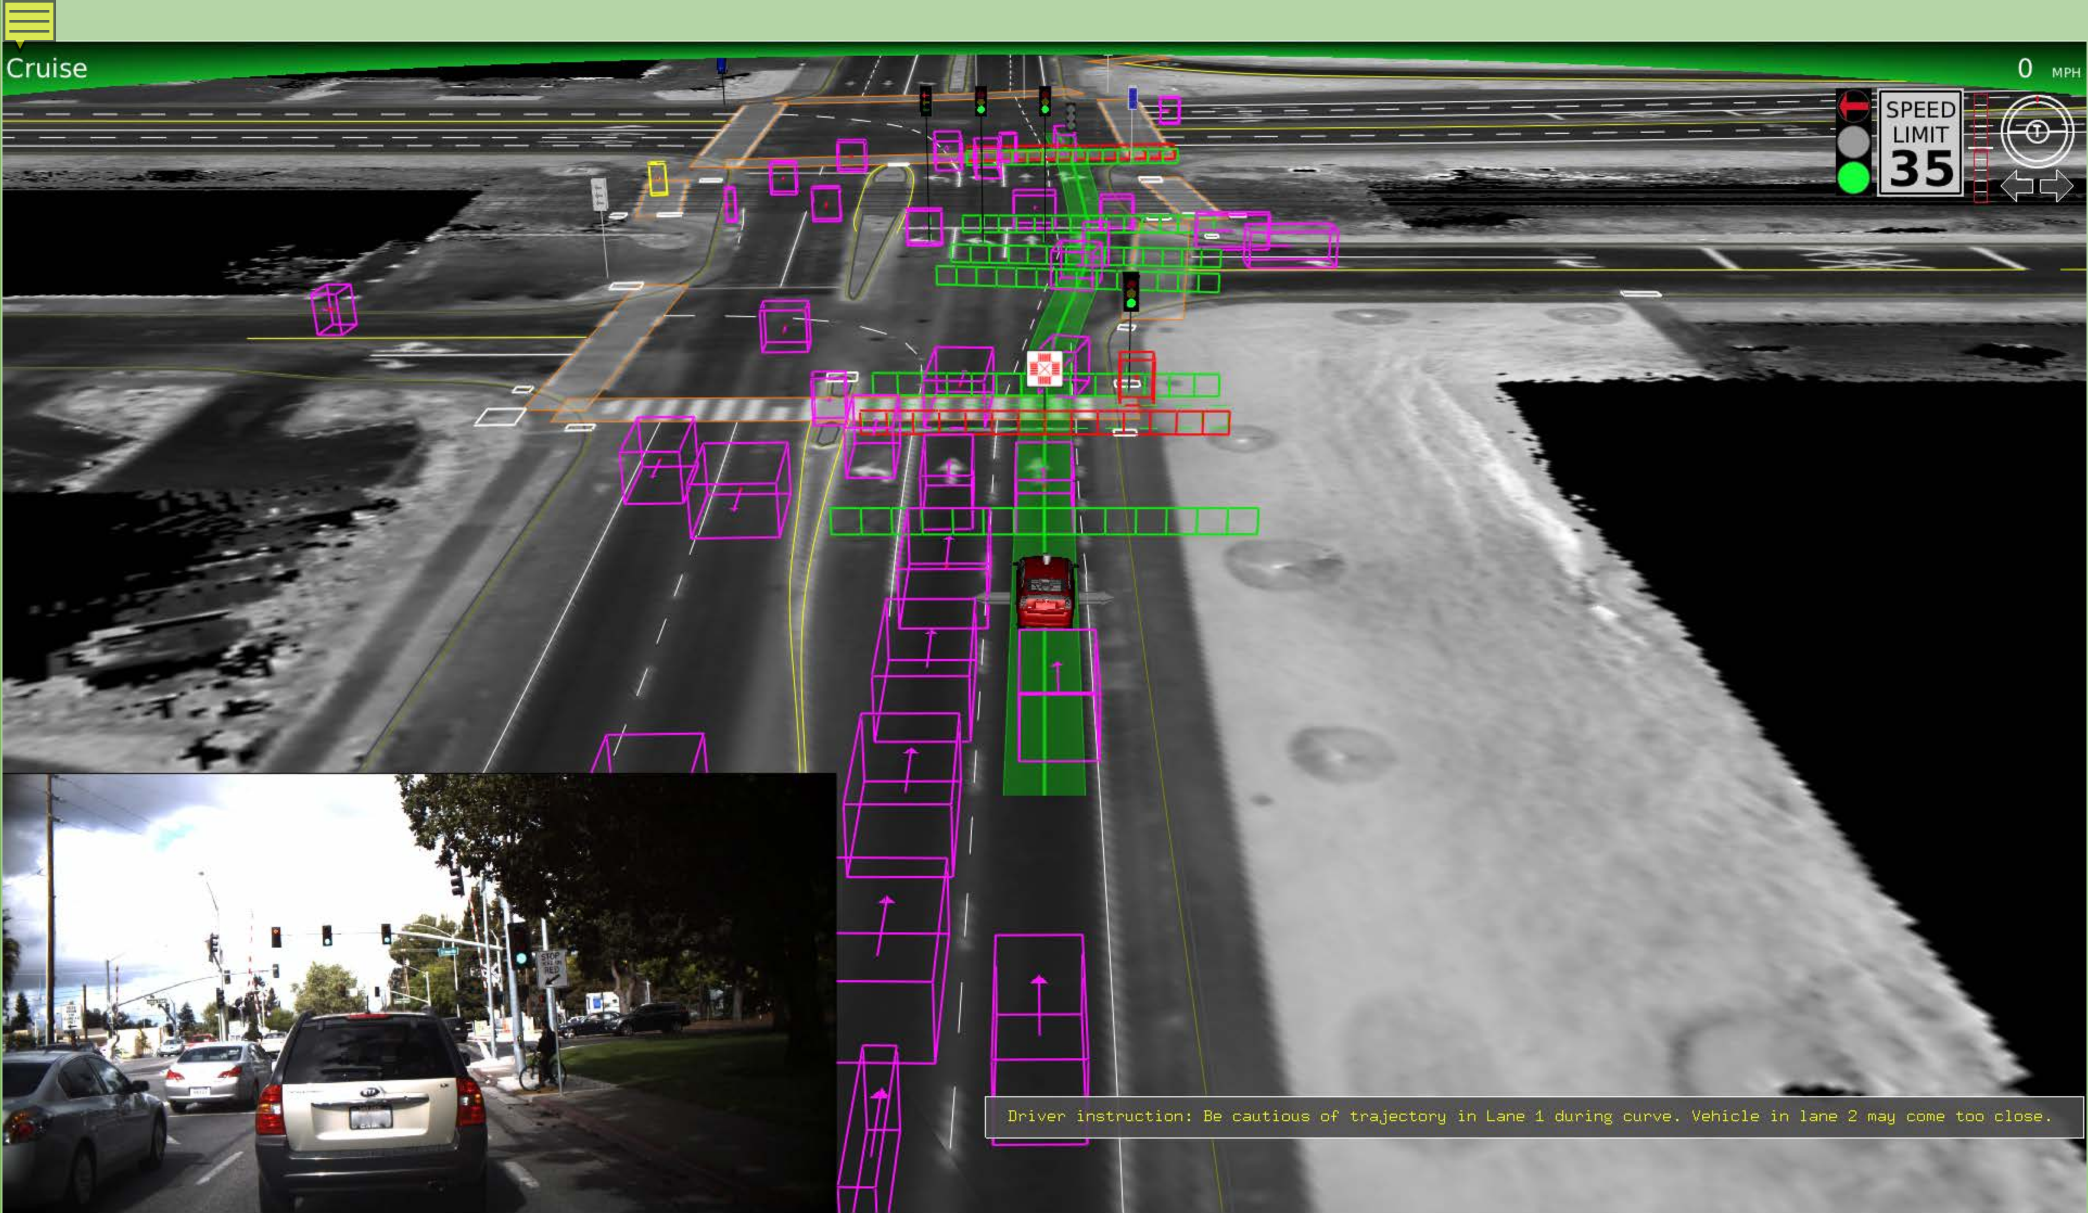

Google’s environment representation

Mission Execution

Route Level Navigation

Behavioral Planning

Motion Control

Route

Lane

Continuous

RADAR

Cloud Data

Maps

Traffic

Weather

Precise State

Time to collision with Static Objects

Time to Collision with Moving Objects

Vehicle Actuator Command

Functional Reference Architecture

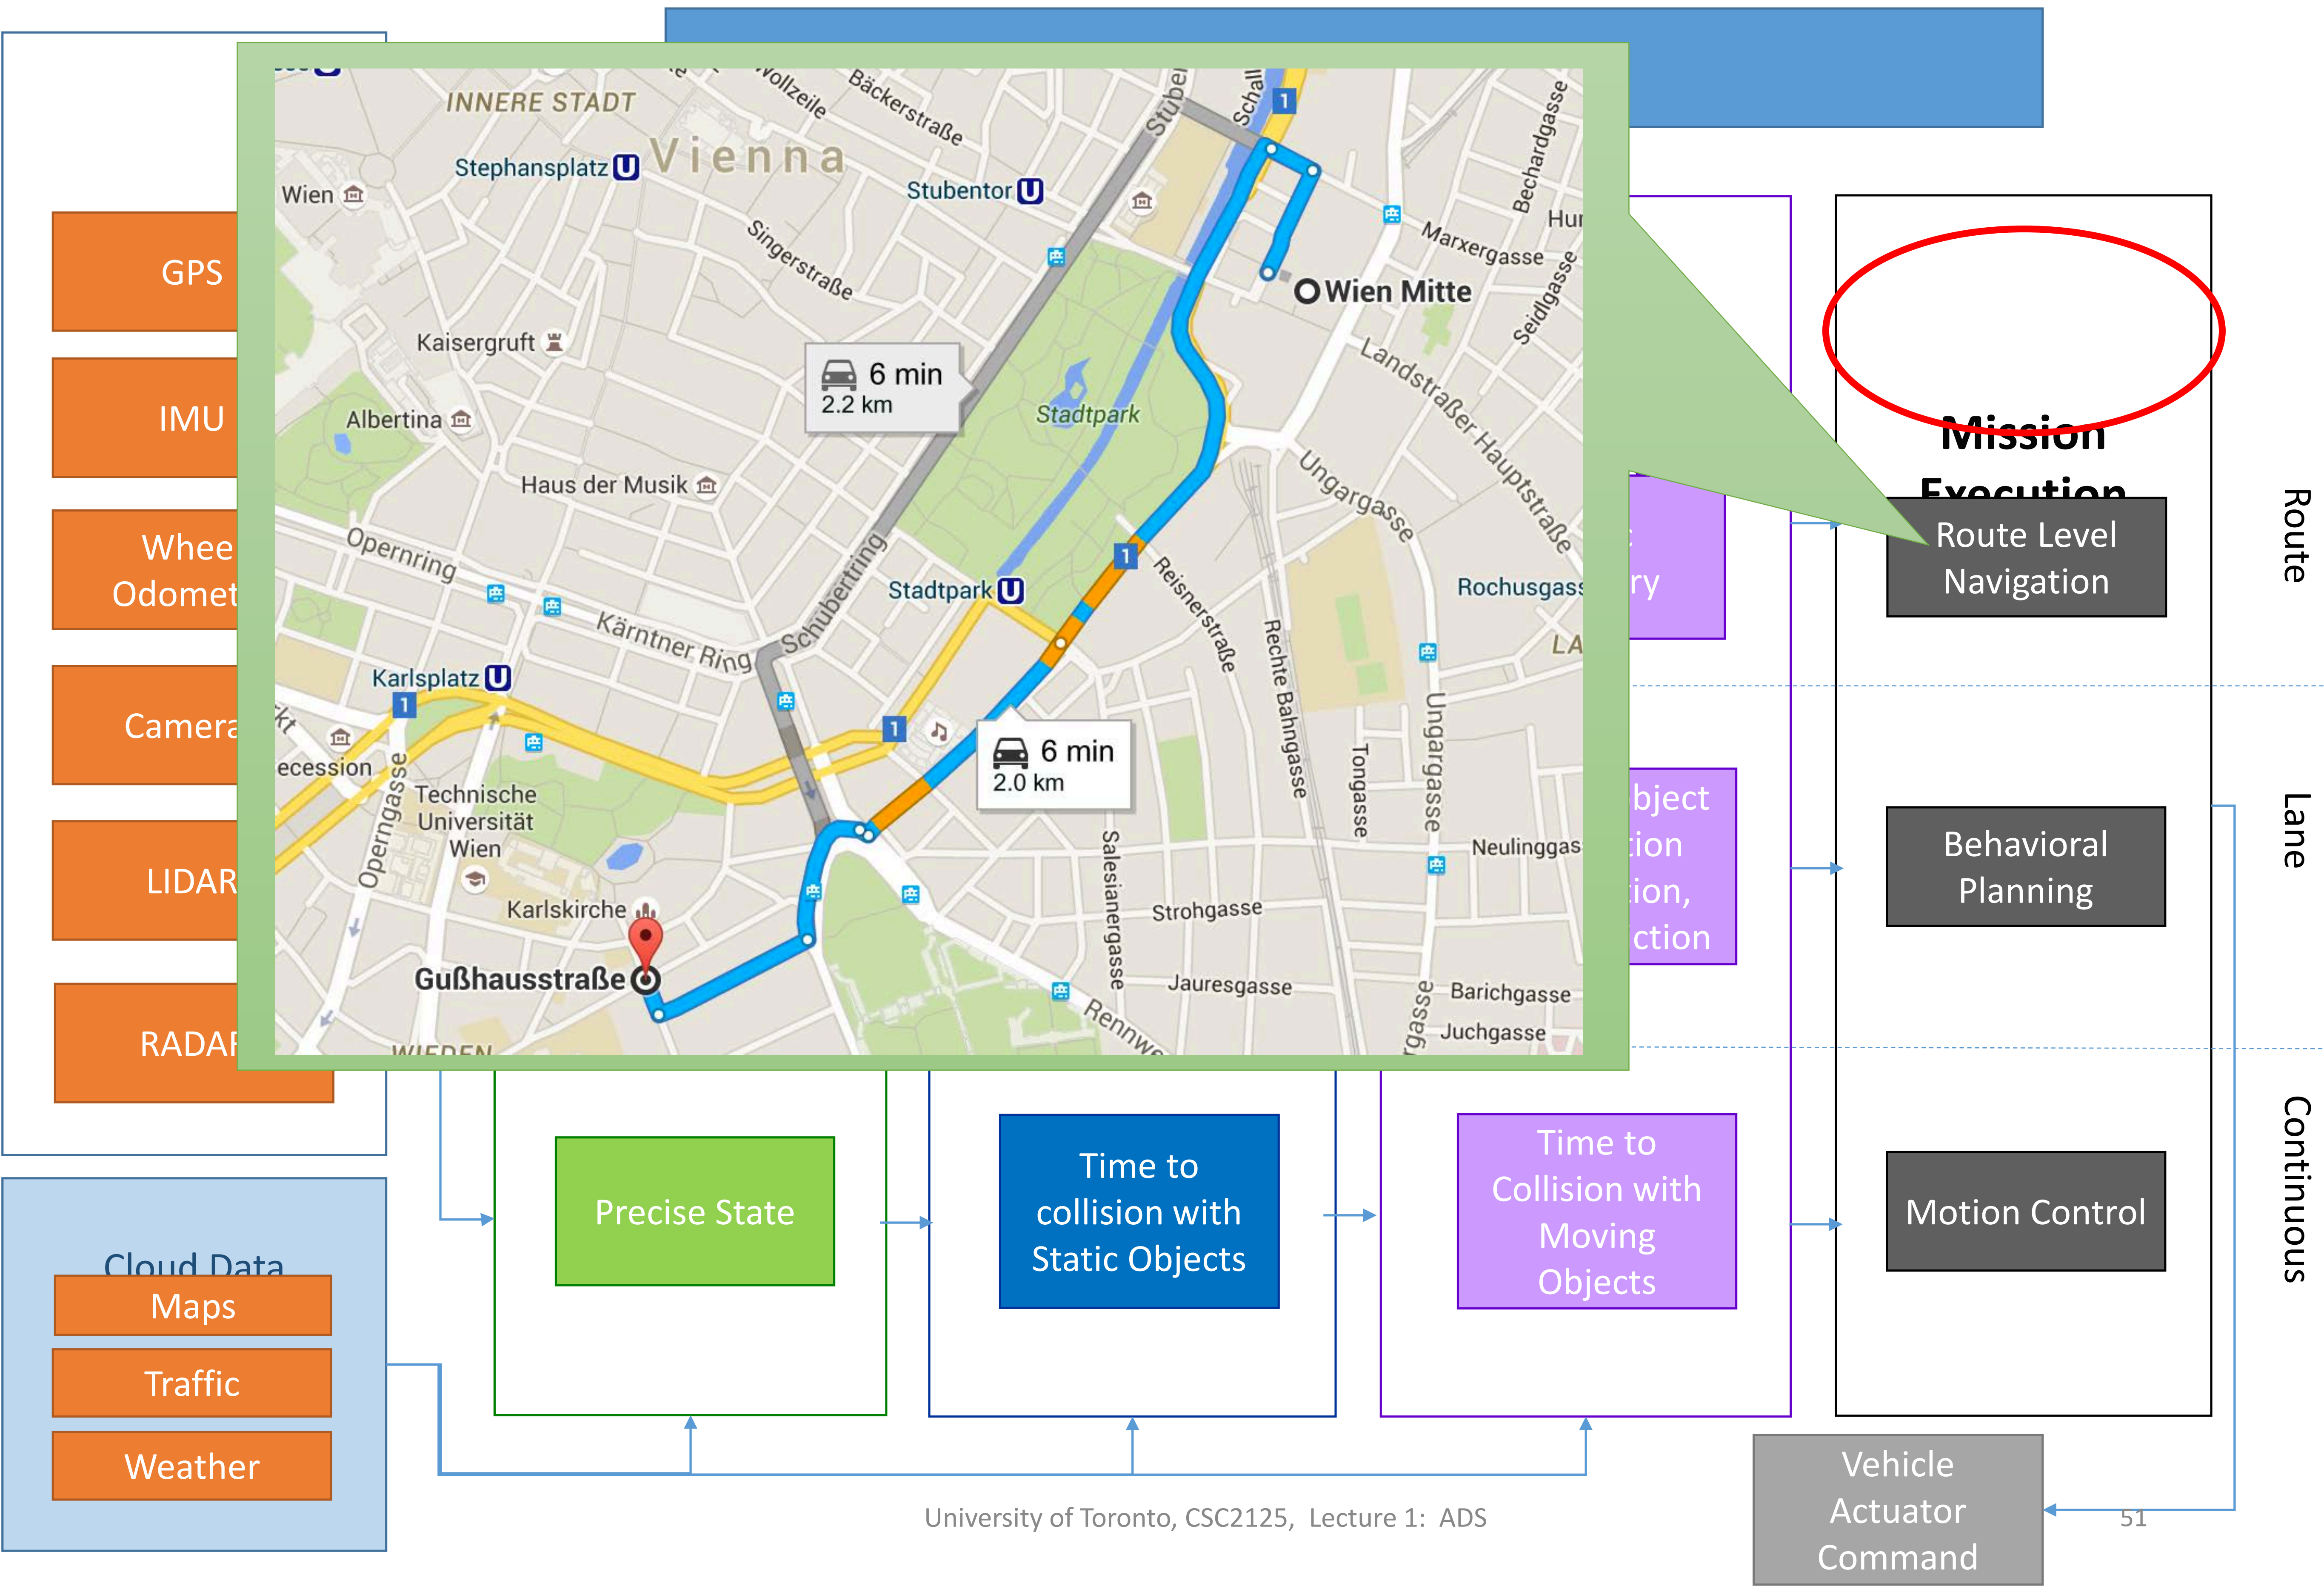

Functional Reference Architecture

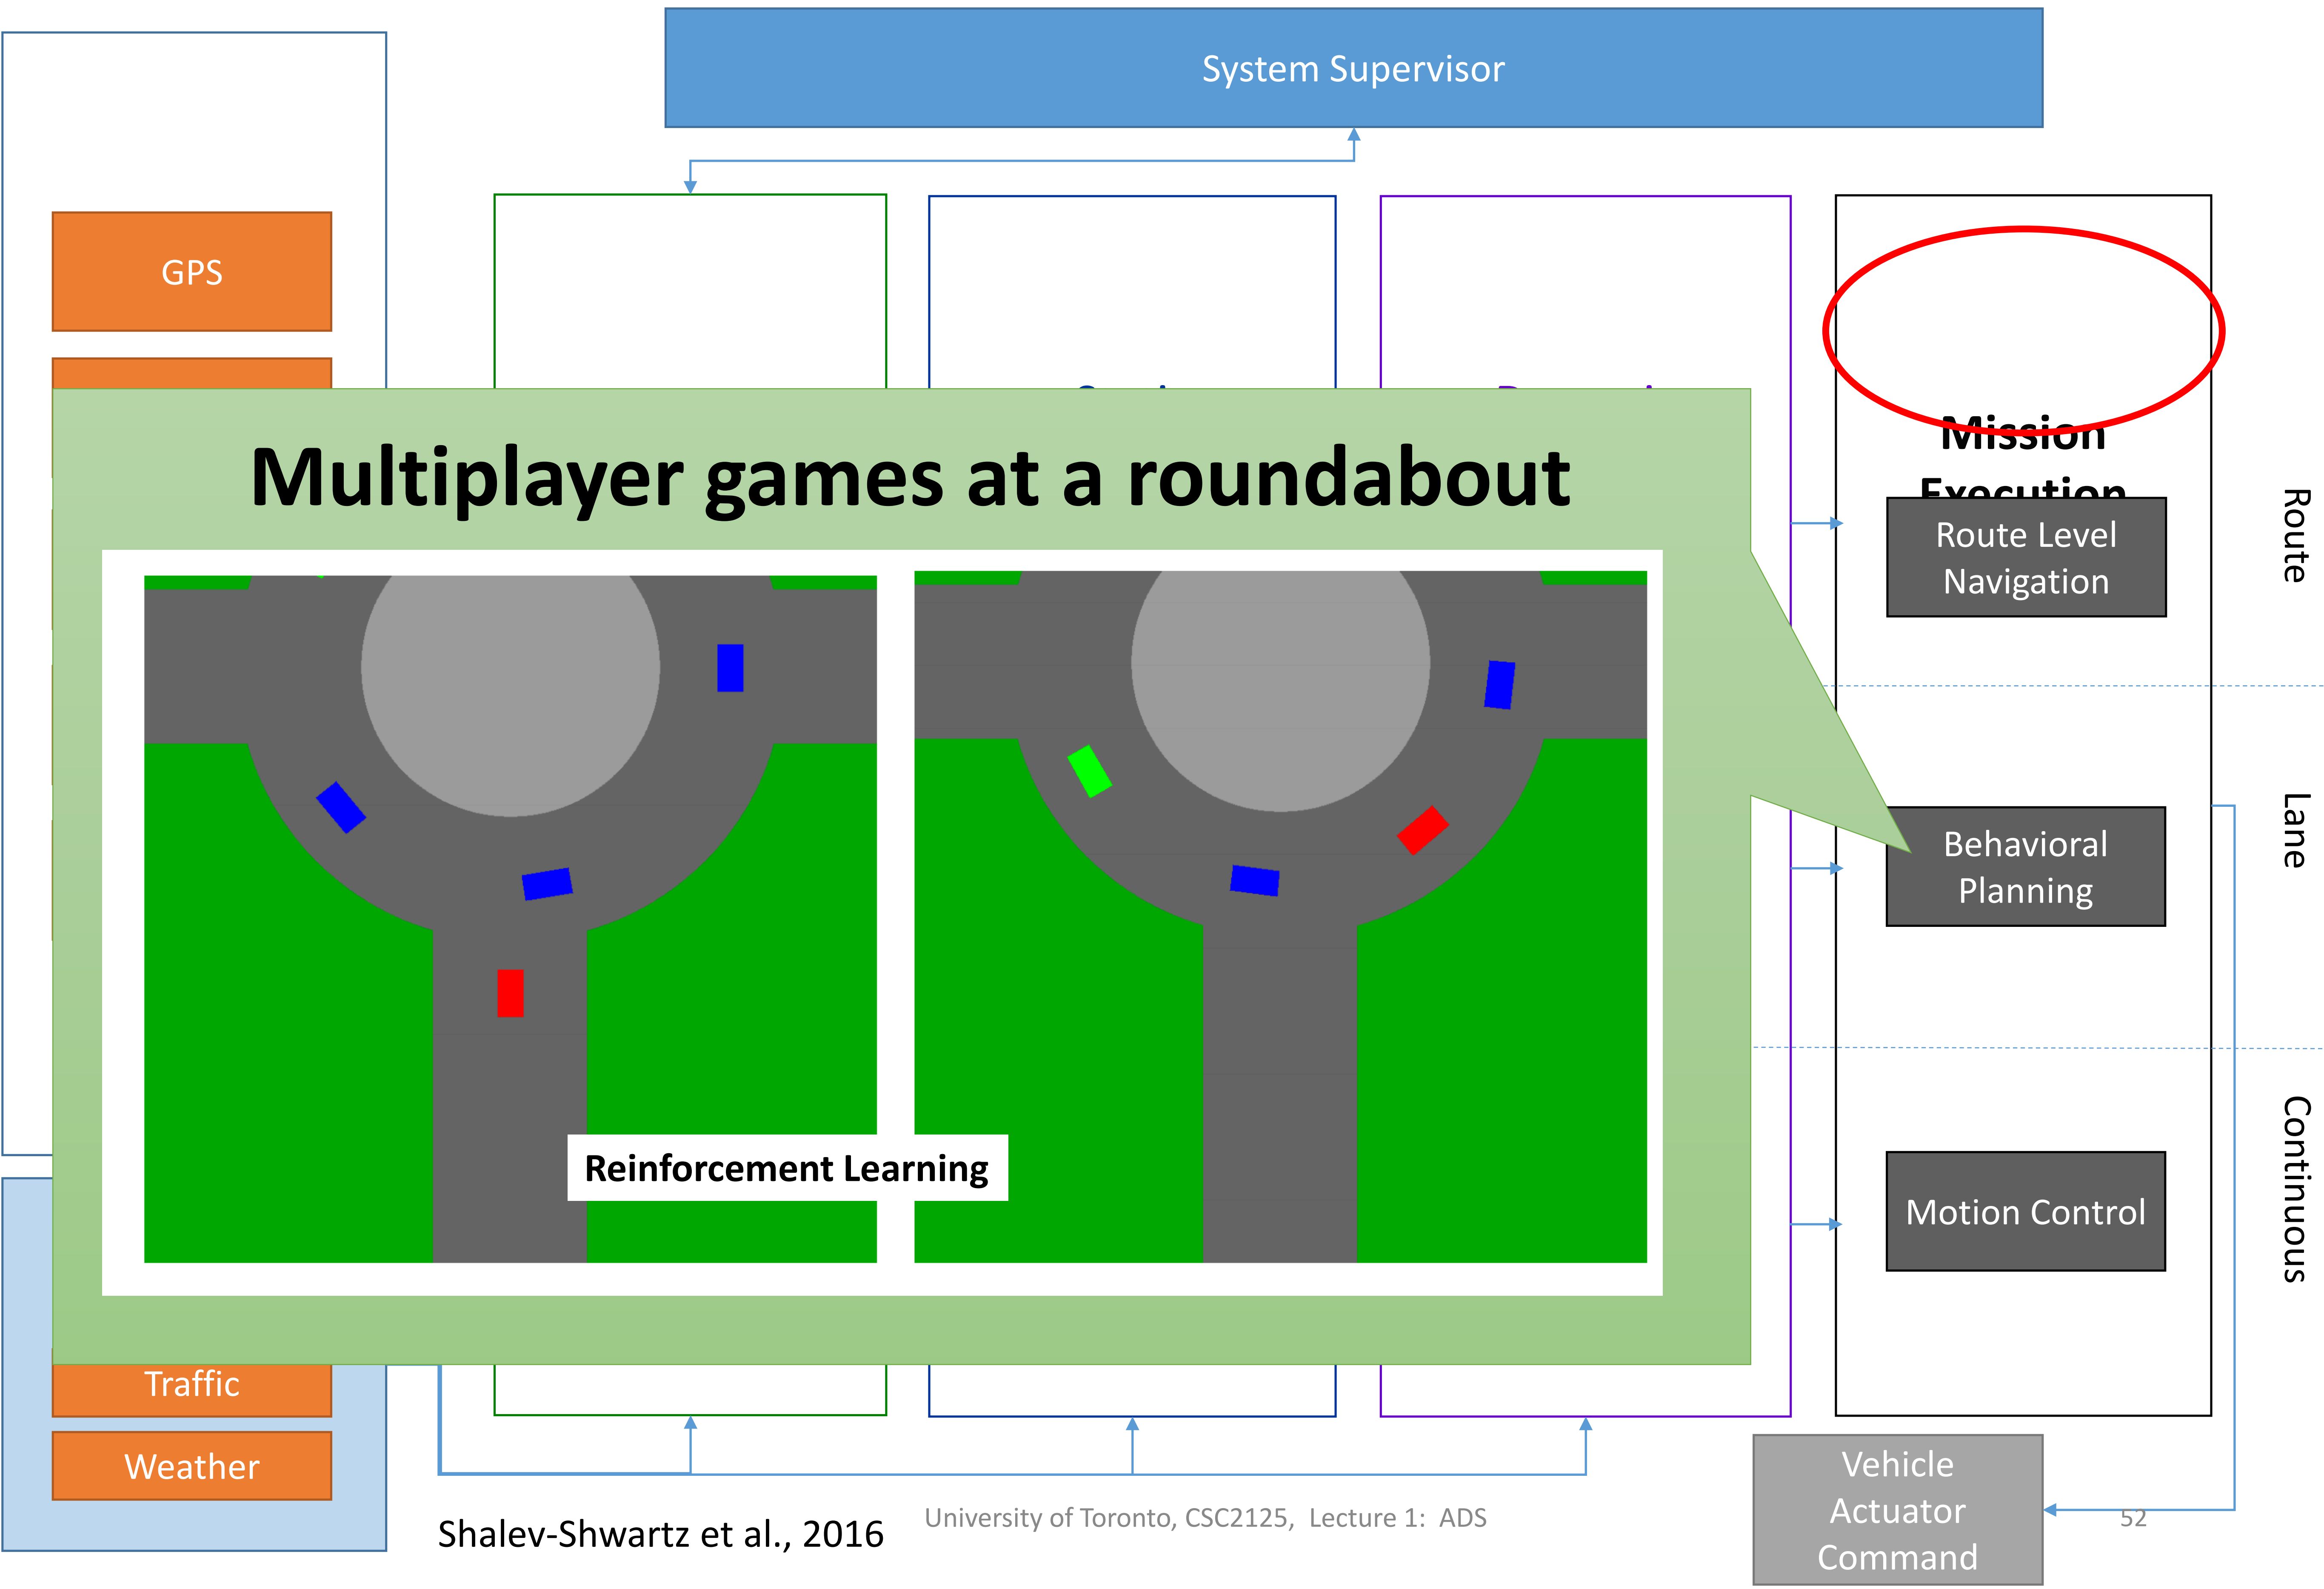

Shalev-Shwartz et al., 2016

Functional Reference Architecture

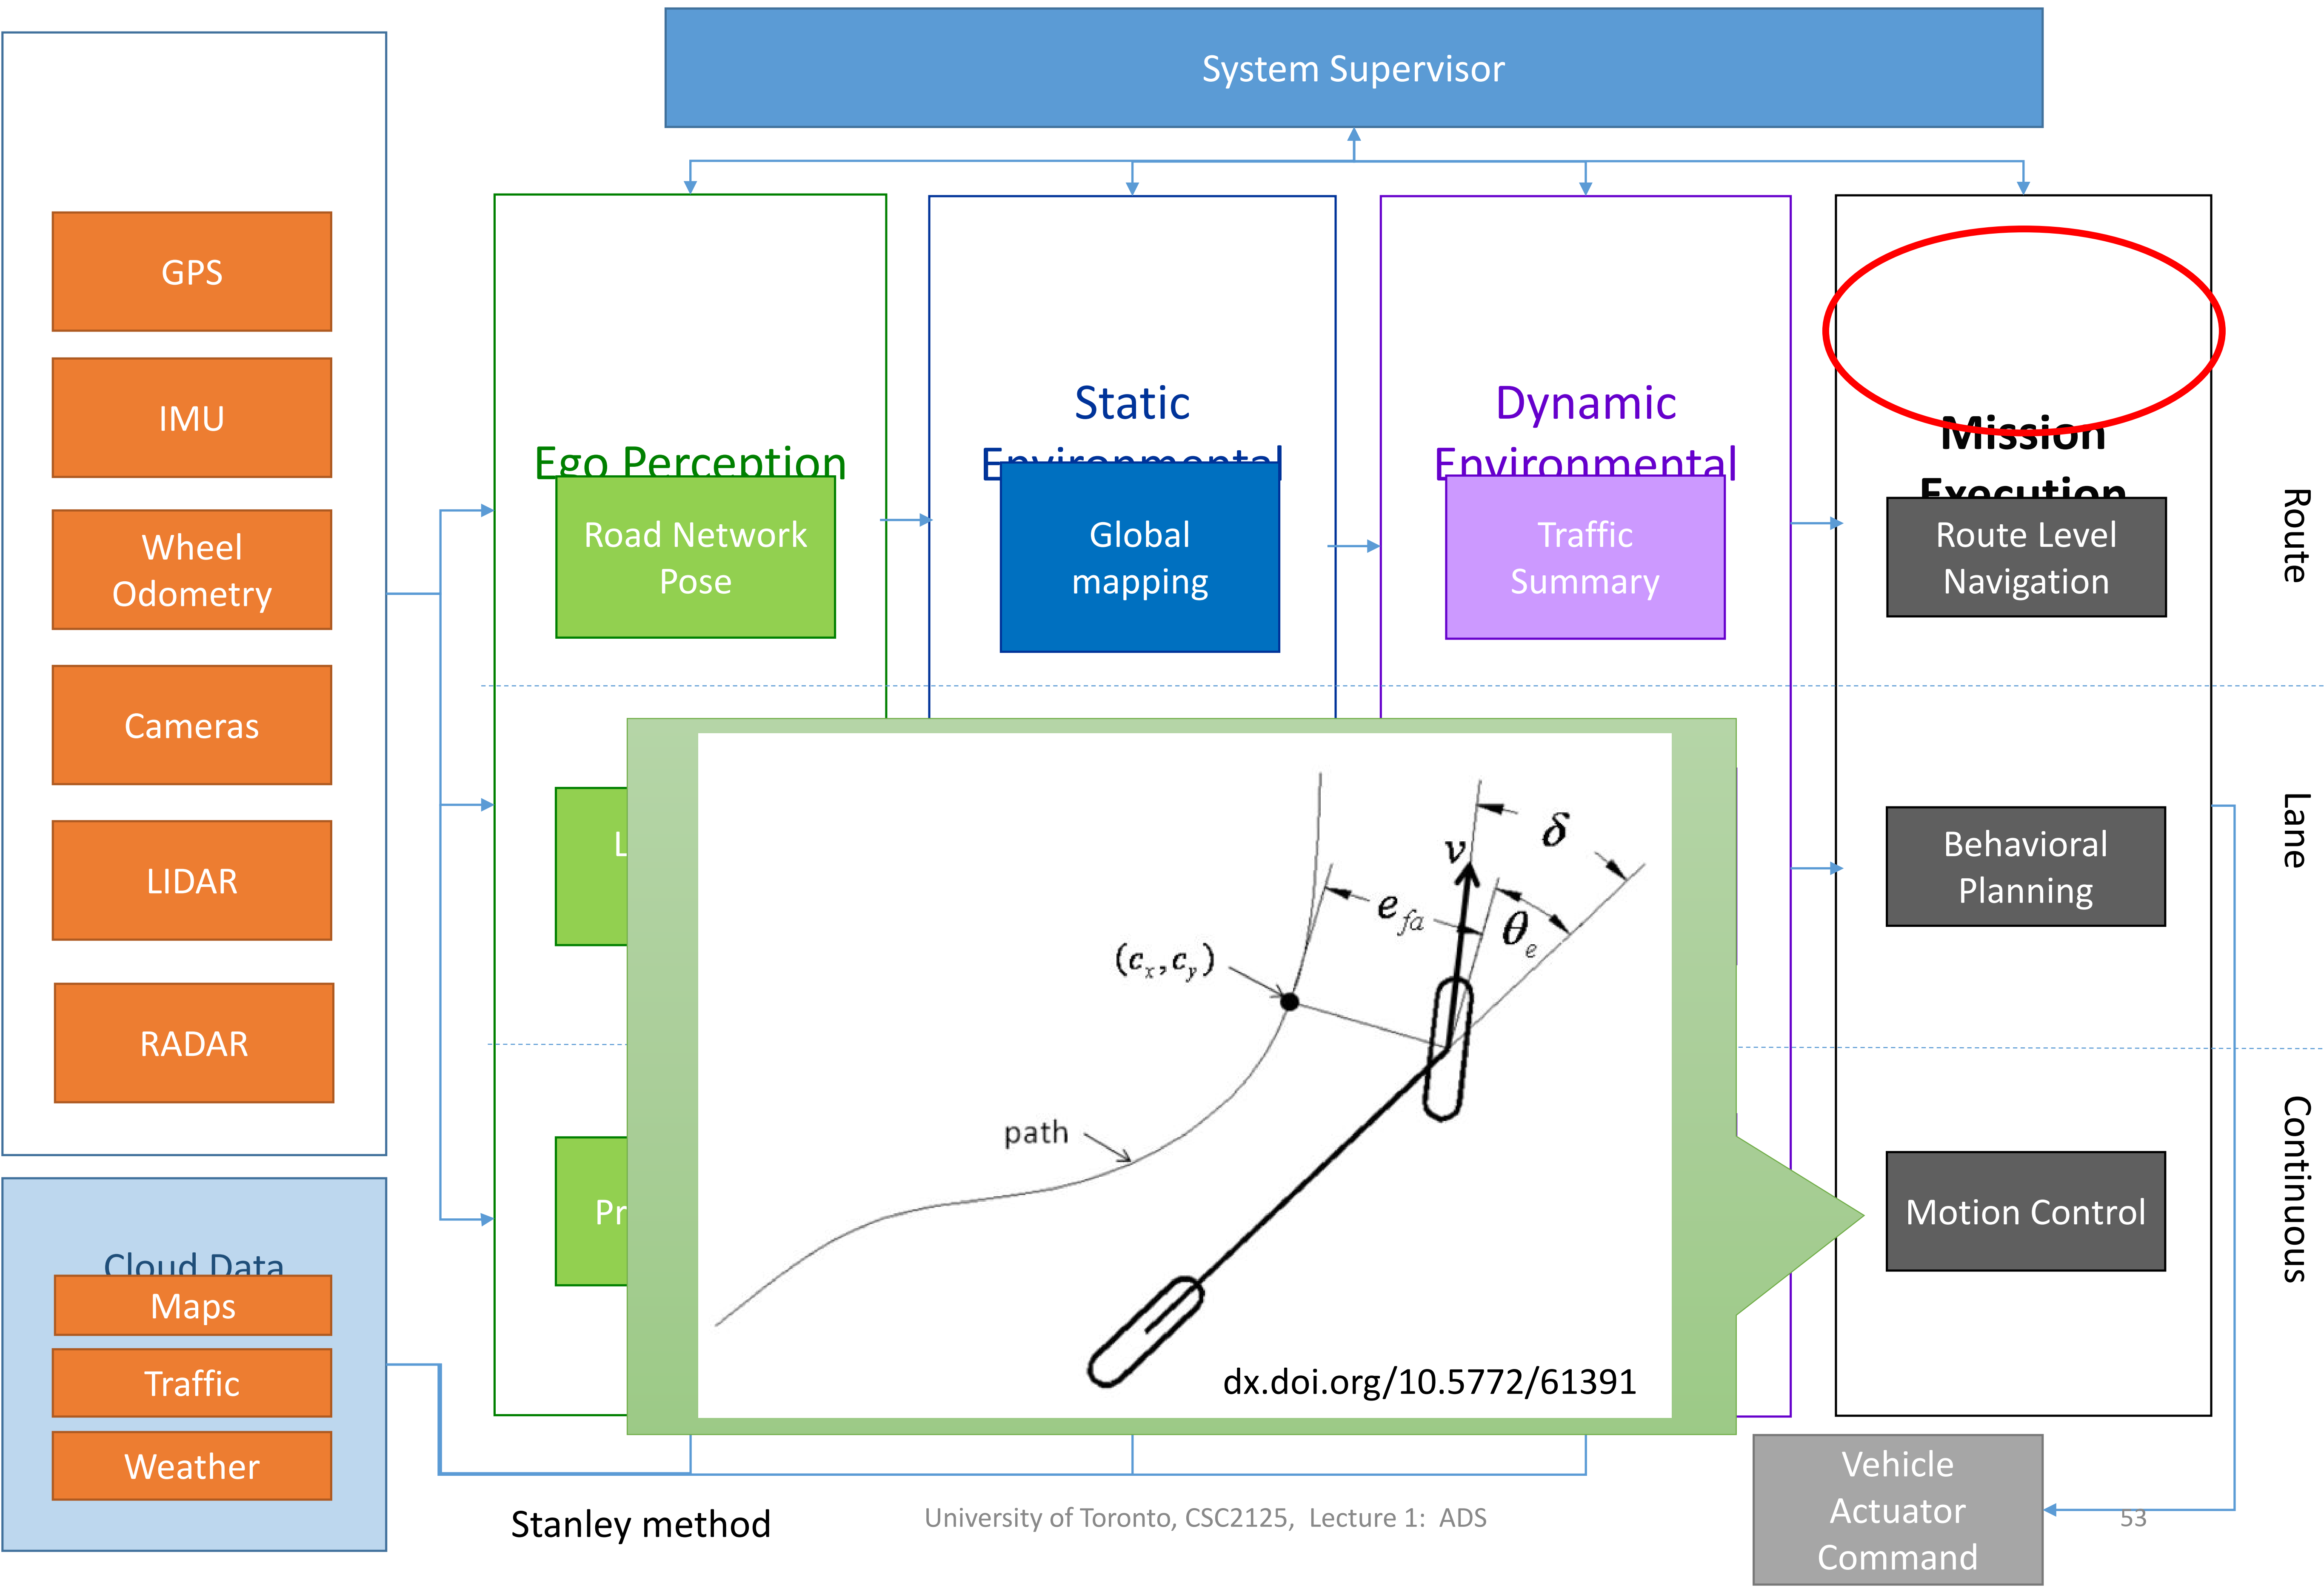

Stanley method

Functional Reference Architecture

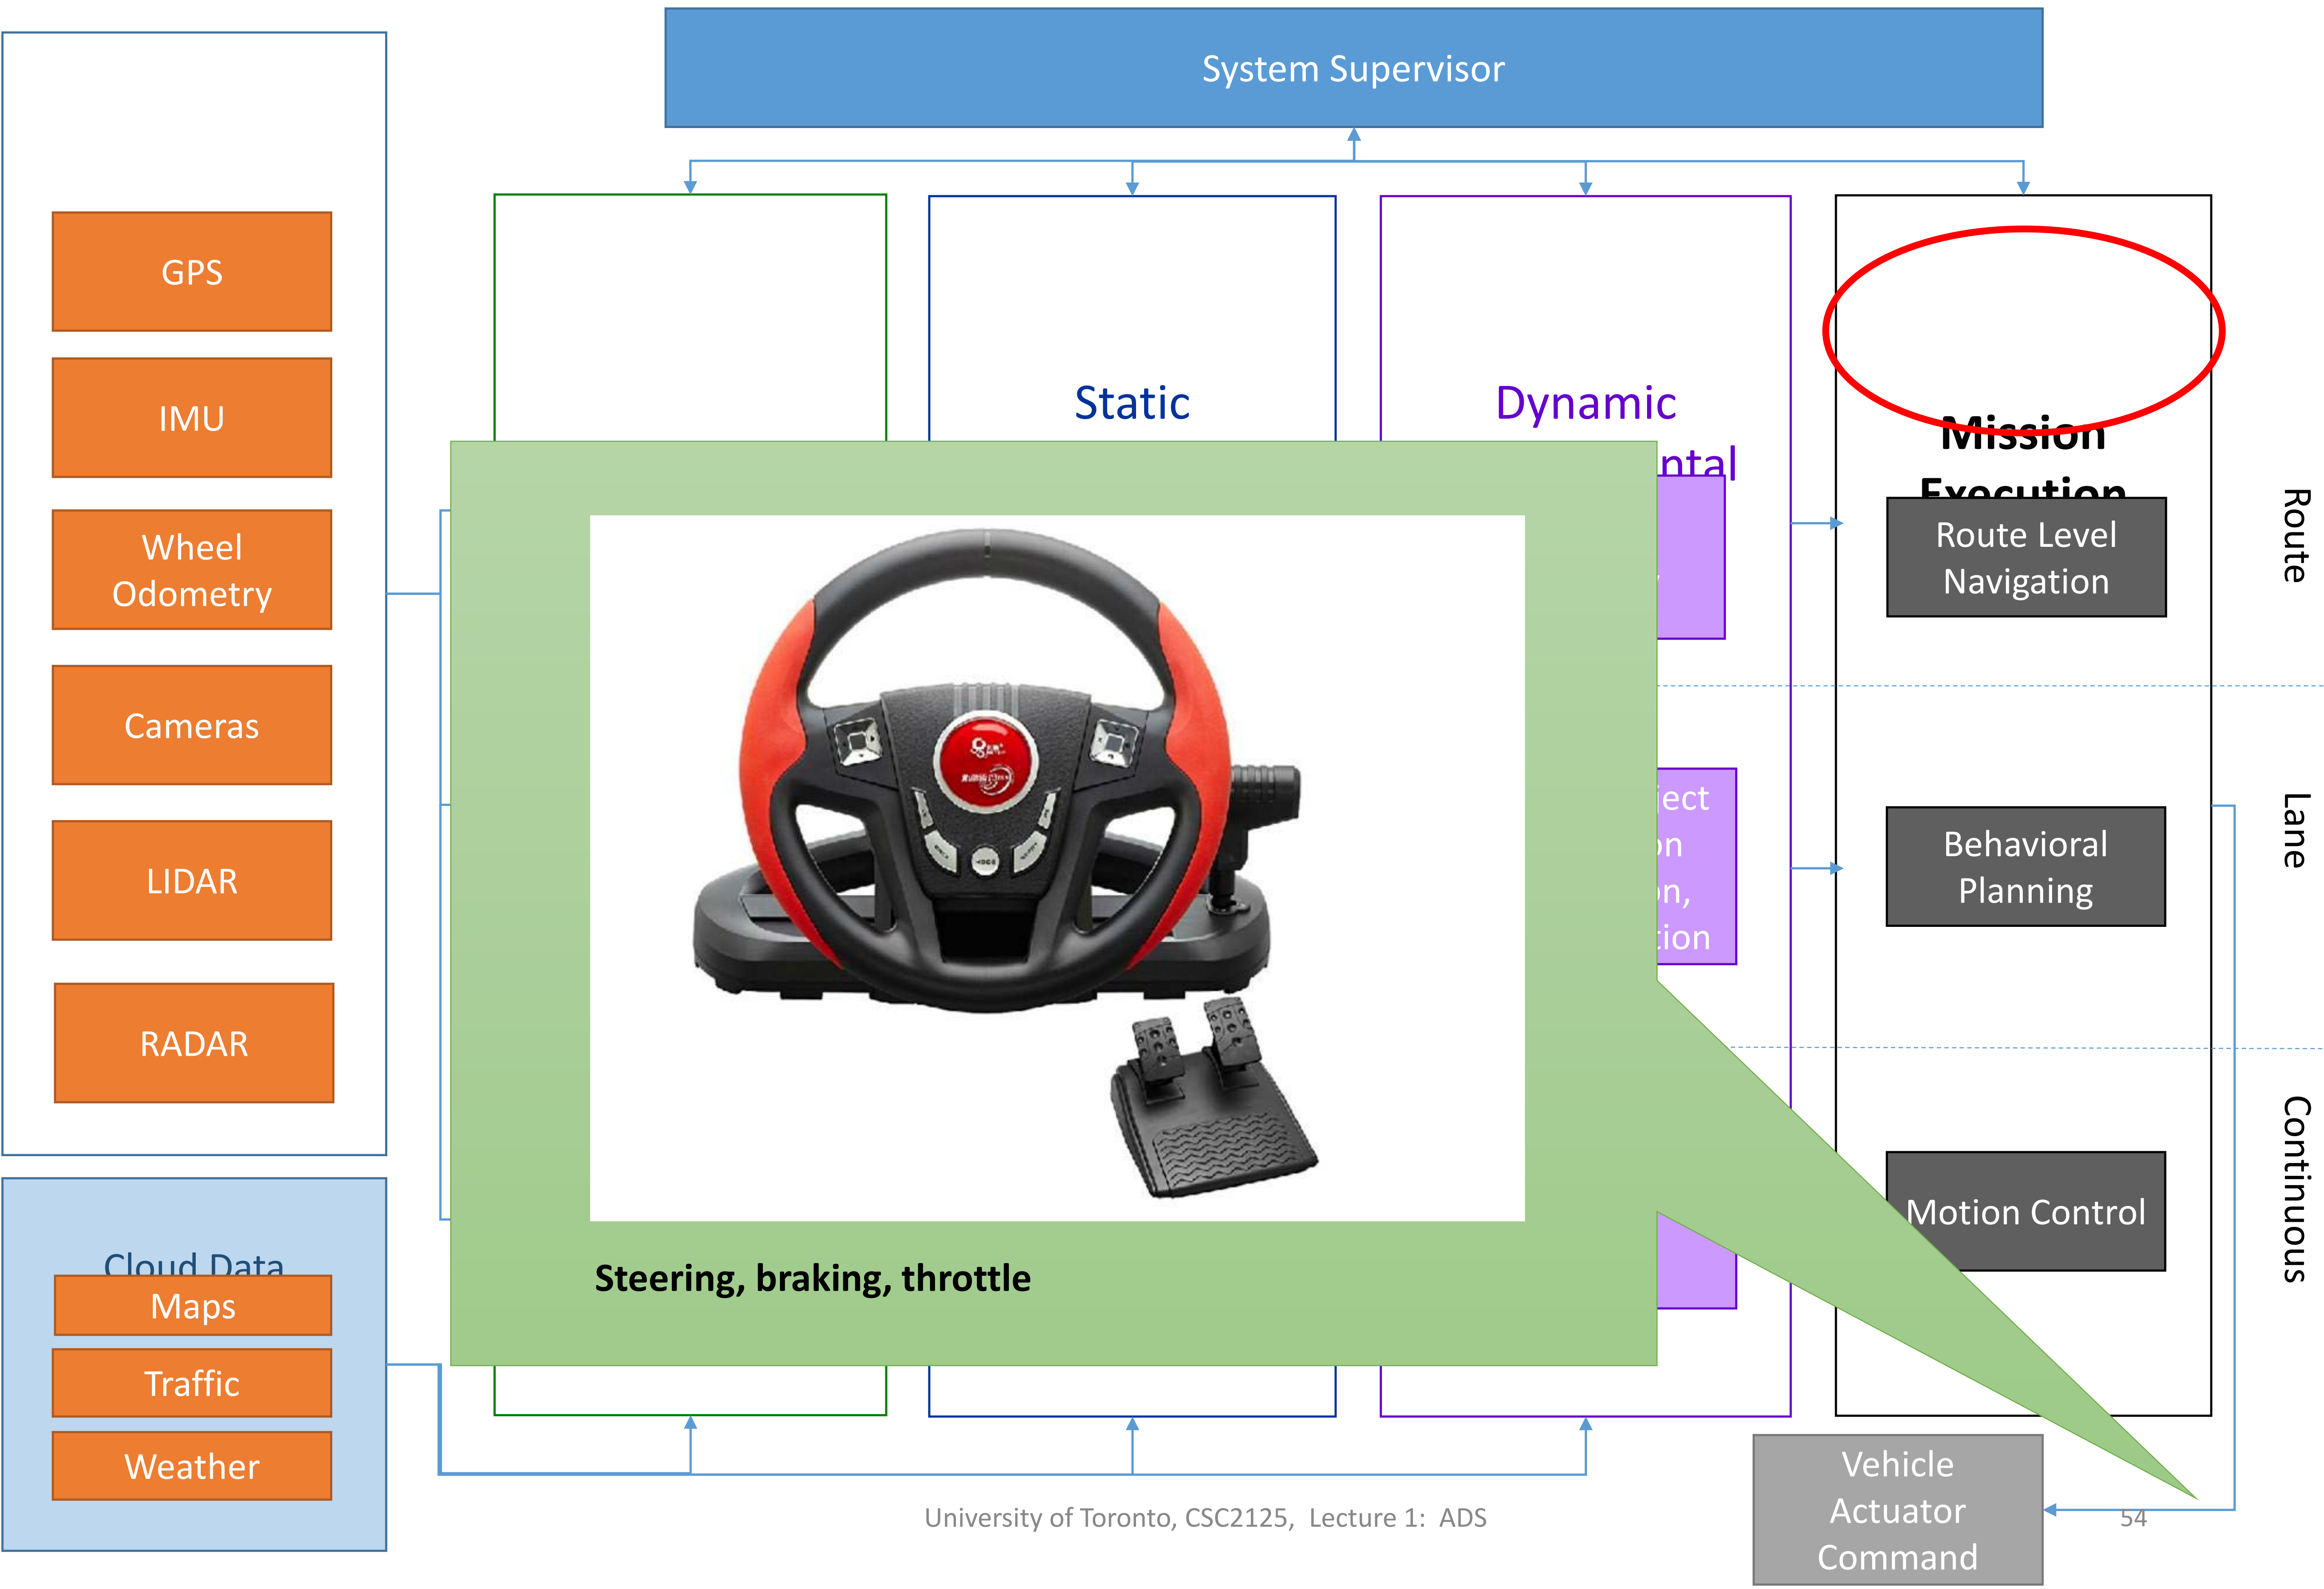

Functional Reference Architecture

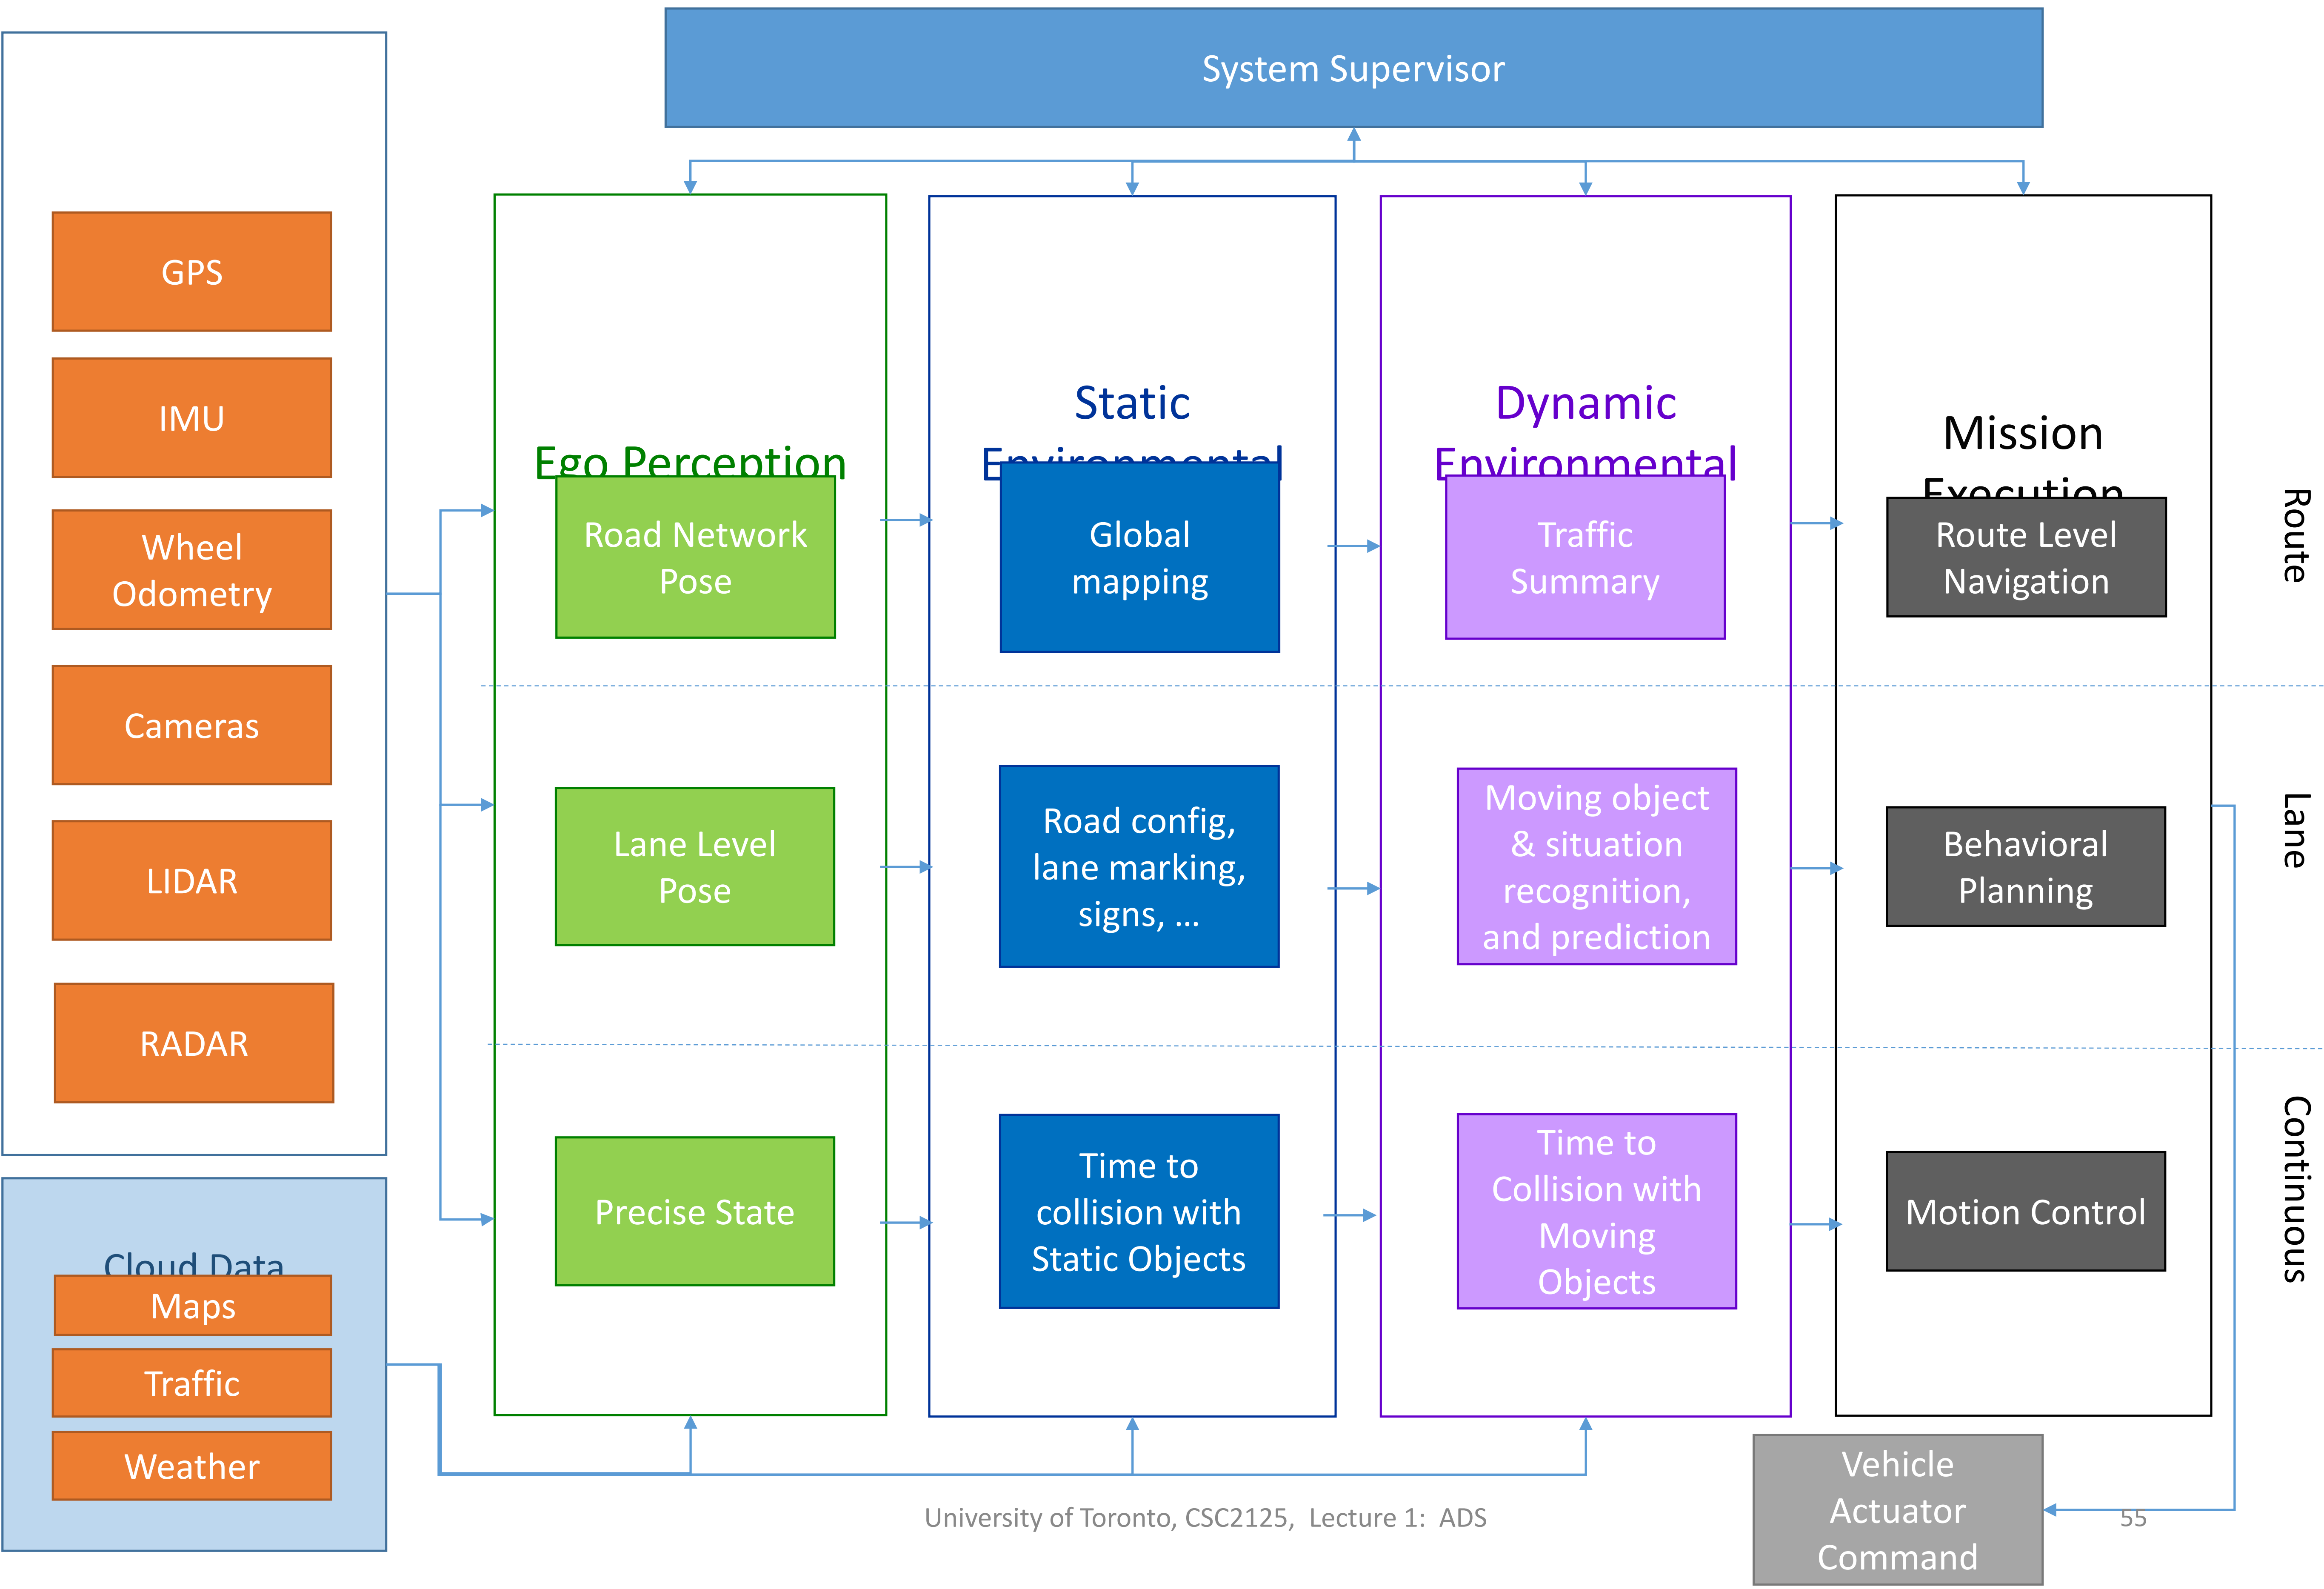

Traffic Data

Naturalistic driving

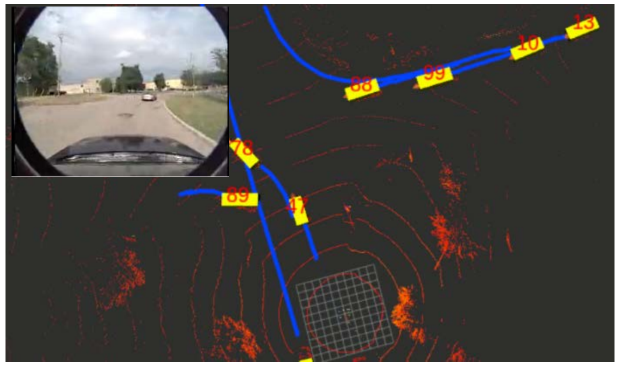

AV sensors & perception

Infrastructure mounted

Birds-eye view

A1 vs A2 Autonomy

- **Starting point:**
  - All cars are manually controlled until the AI system shows itself to be available and is elected to be turned on by the human.
- **A1:** Human-Centered Autonomy
  - **Definition:** AI is not fully responsible
  - Feature axis:
    - Where/how often is it "available"? (traffic, highway, sensor-based, etc.)
    - How many seconds for take-over? (0, 1, 10, etc)
    - Teleoperation support
- **A2:** Full Autonomy
  - **Definition:** AI is fully responsible
  - Notes:
    - No teleoperation
    - No 10-second rule: It's allowed to ask for human help, but not guaranteed to ever receive it.
    - Arrive to a **safe** destination or safe harbor.
    - Allow the human to take over **when they choose to**.

A1 vs A2 Autonomy

L0 →
- **Starting point:**
  - All cars are manually controlled until the AI system shows itself to be available and is elected to be turned on by the human.

L1, L2, L3 →
- **A1:** Human-Centered Autonomy
  - **Definition:** AI is not fully responsible

L4, L5 →
- **A2:** Full Autonomy
  - **Definition:** AI is fully responsible

Human-Centric Approach to AI (also see Safety)

90% ← No — Human Needed — Yes → 10%

Solve the perception-control problem where **possible**:

And where **not possible**: involve the human

Perception / control (via Deep-Learning)

Effective human-robot interaction

Paths to Autonomous Future

A1:

Human-Centered Autonomy

- **Localization and Mapping:**
  Where am I?
- **Scene Understanding:**
  Where/who/what/why of everyone else?
- **Movement Planning:**
  How do I get from A to B?
- **Human-Robot Interaction:**
  What is the physical and mental state of the driver?
- **Communicate:**
  How to I convey intent to the driver and to the world?

Blue Text: Easier

Red Text: Harder

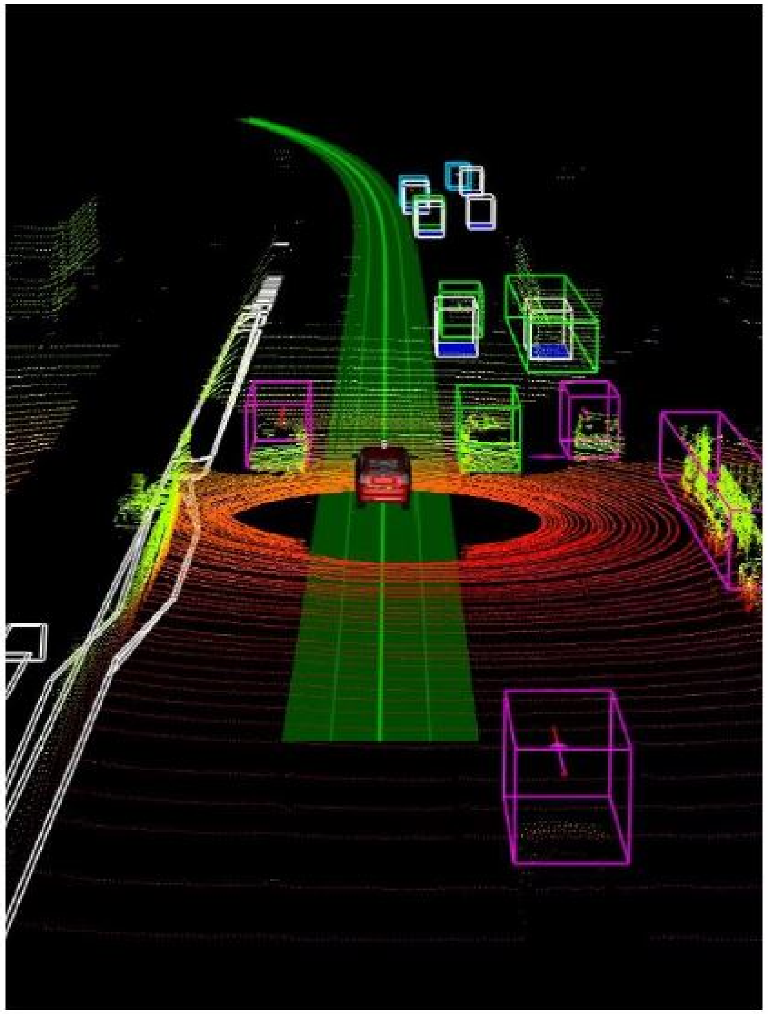

A2:

Full Autonomy

- **Localization and Mapping:**
  Where am I?
- **Scene Understanding:**
  Where/who/what/why of everyone else?
- **Movement Planning:**
  How do I get from A to B?
- **Human-Robot Interaction:**
  What is the physical and mental state of the driver?
- **Communicate:**
  How to I convey intent to the driver and to the world?

Is partially automated driving a bad idea? Observations from an on-road study

Article · April 2018 *with* 447 Reads

DOI: 10.1016/j.apergo.2017.11.010

Victoria Banks
14.44 · University of Southampton

Alexander Eriksson
11.13 · Swedish National Road and Transport Research Inst...

Jim O'donoghue

Neville A Stanton
43.23 · University of Southampton

WAYMO

Chris Urmson

Yes, with nothing to do, drivers quickly stop paying attention, get distracted, fall asleep

Public Perception of What Drivers Do in Semi-Automated Vehicles

What Does Data Say?

A look at several autonomous driving accidents

Based on work of Prof. Mark Lawford, McMaster University

1st Fatal Tesla Autopilot Crash

2016 - January 20th - Fatal - Tesla Model S(China)

1st Fatal Tesla Autopilot Crash
Analysis

- Model S was equipped with
  - a single forward facing radar,
  - a single forward facing camera,
  - a set of 12 ultrasonic sensors.
- Camera was used by MobileEye's EyeQ3 computing platform implementing a Deep Neural Network (DNN) for its object identification and detection
- Vehicle was also equipped with Tesla's Automatic Emergency Braking (AEB) system
  - AEB system required agreement between ***both*** the camera and the radar before any action was taken.
- Driver monitoring system consisted of a torque sensor in the steering wheel

2nd Fatal Tesla Autopilot Crash

2017 - May 7th - Fatal - Tesla Model S (Florida)

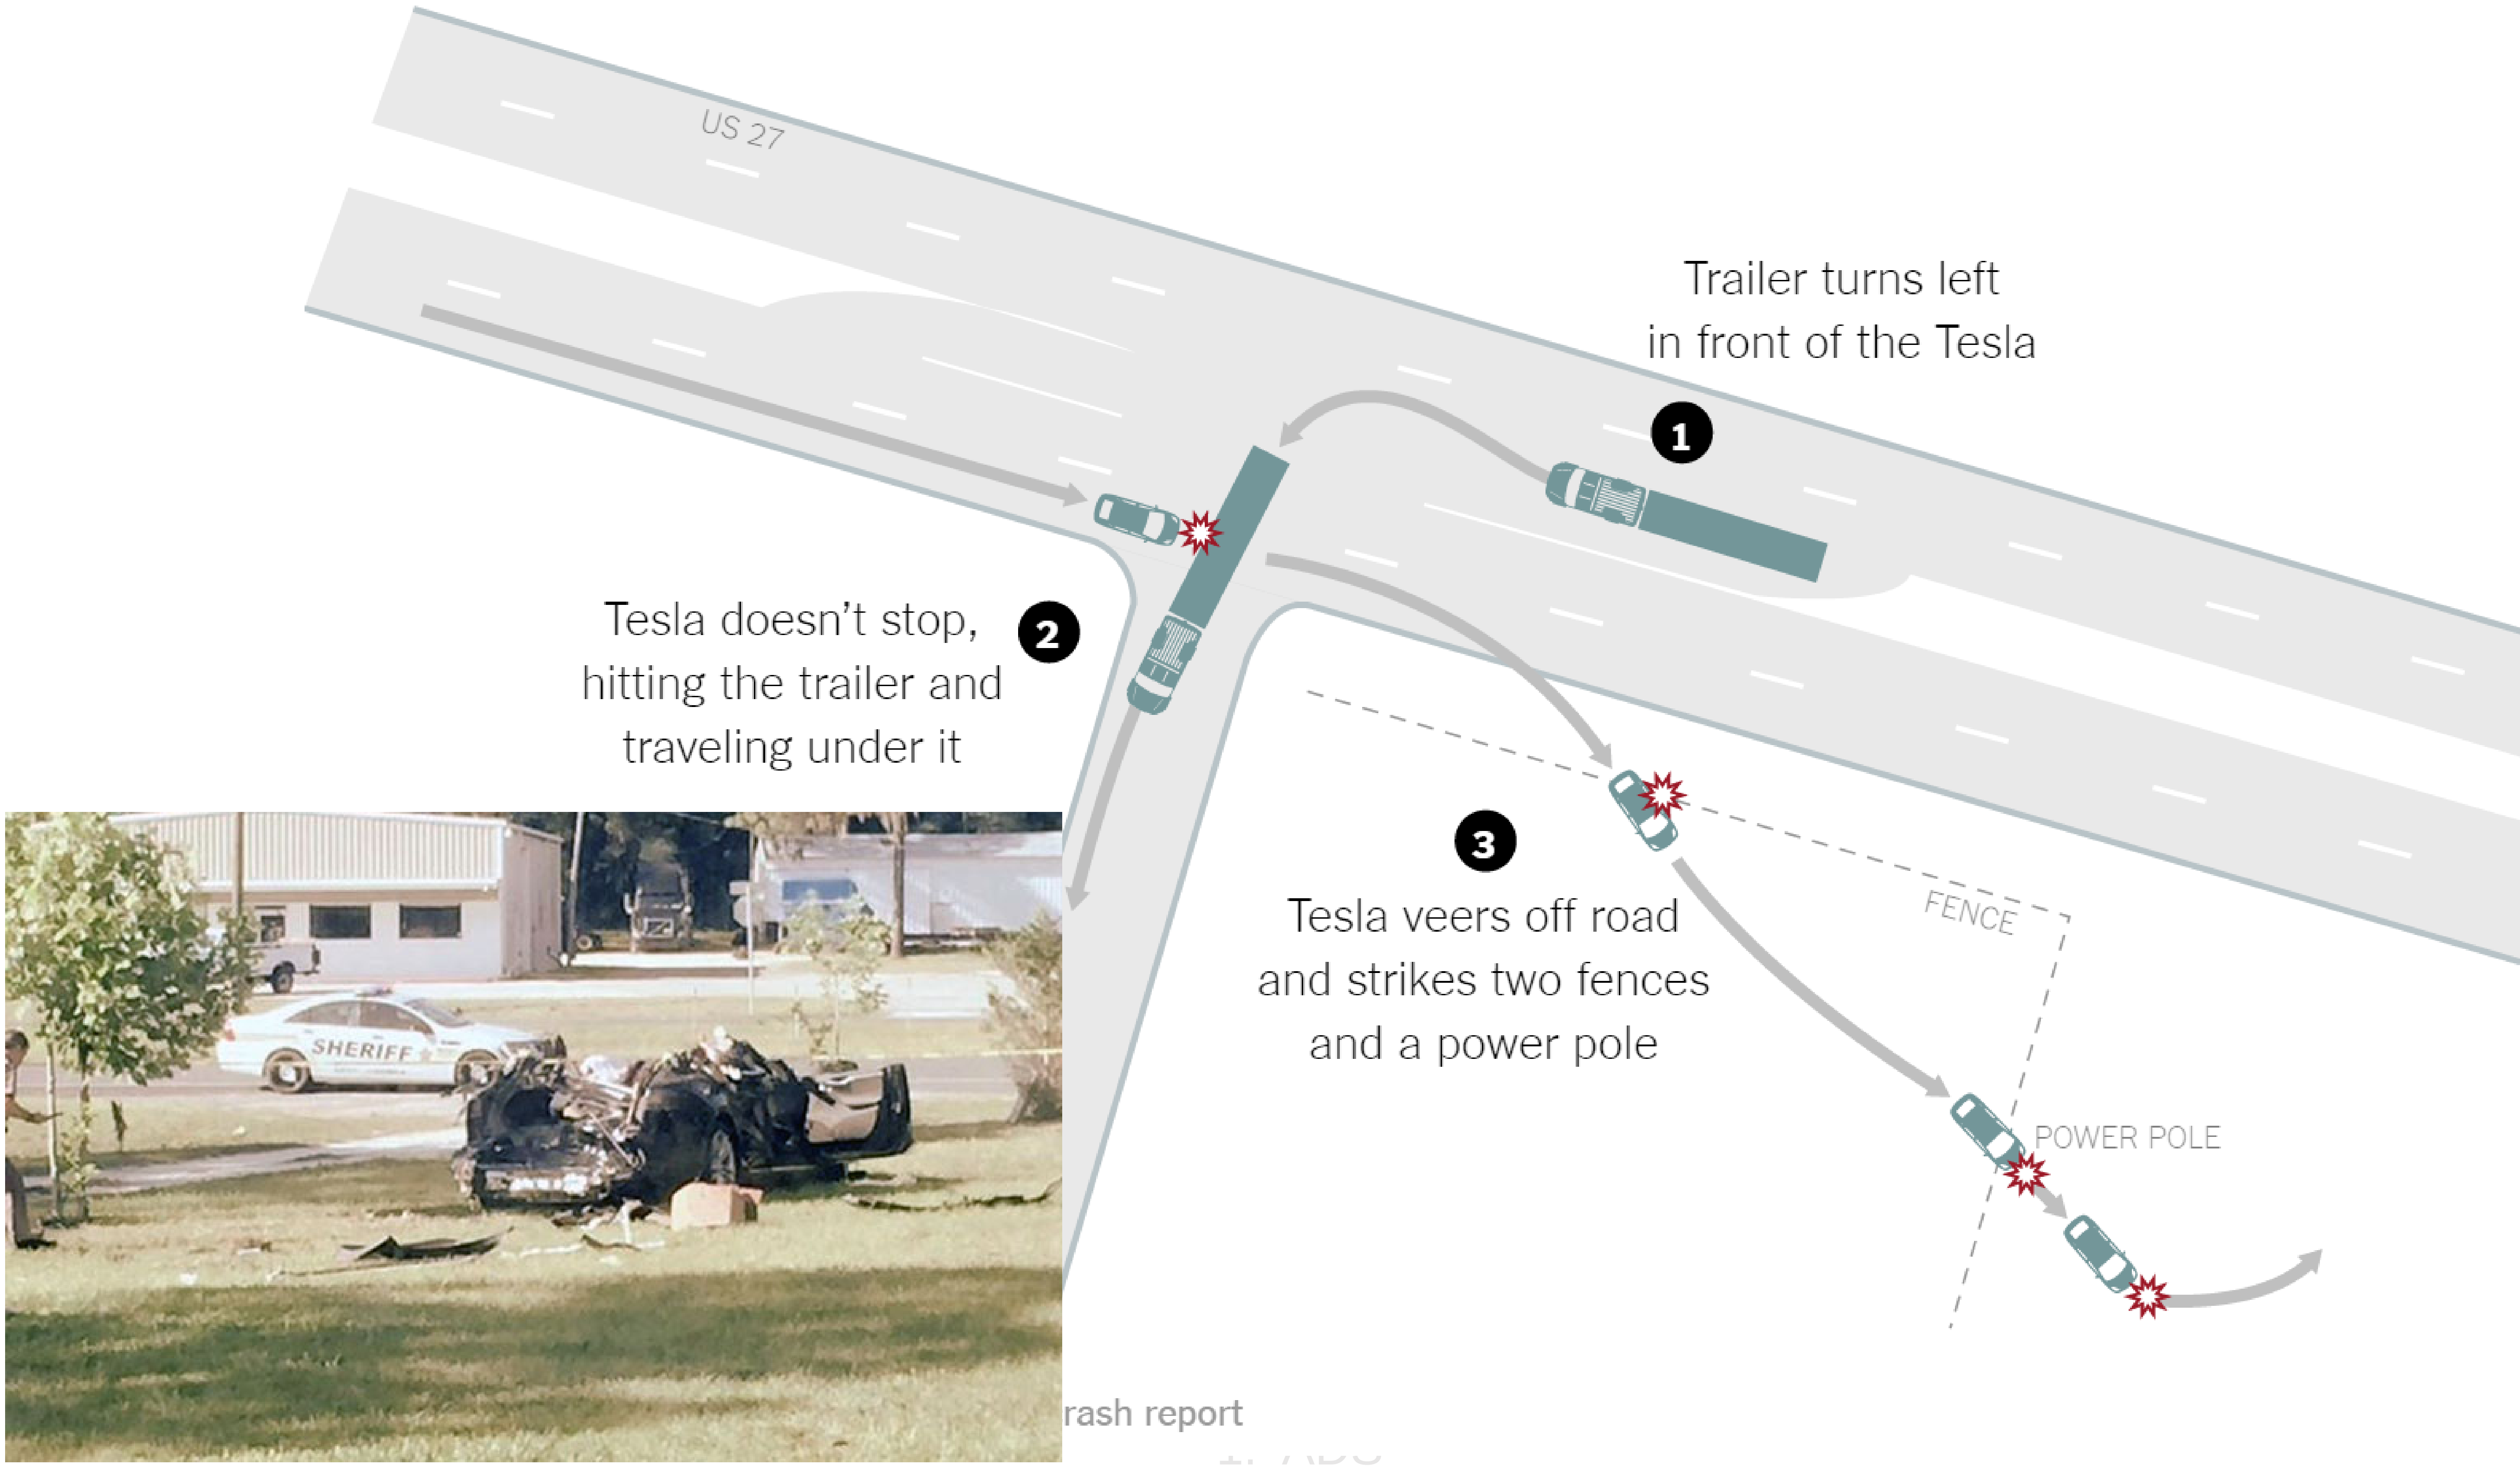

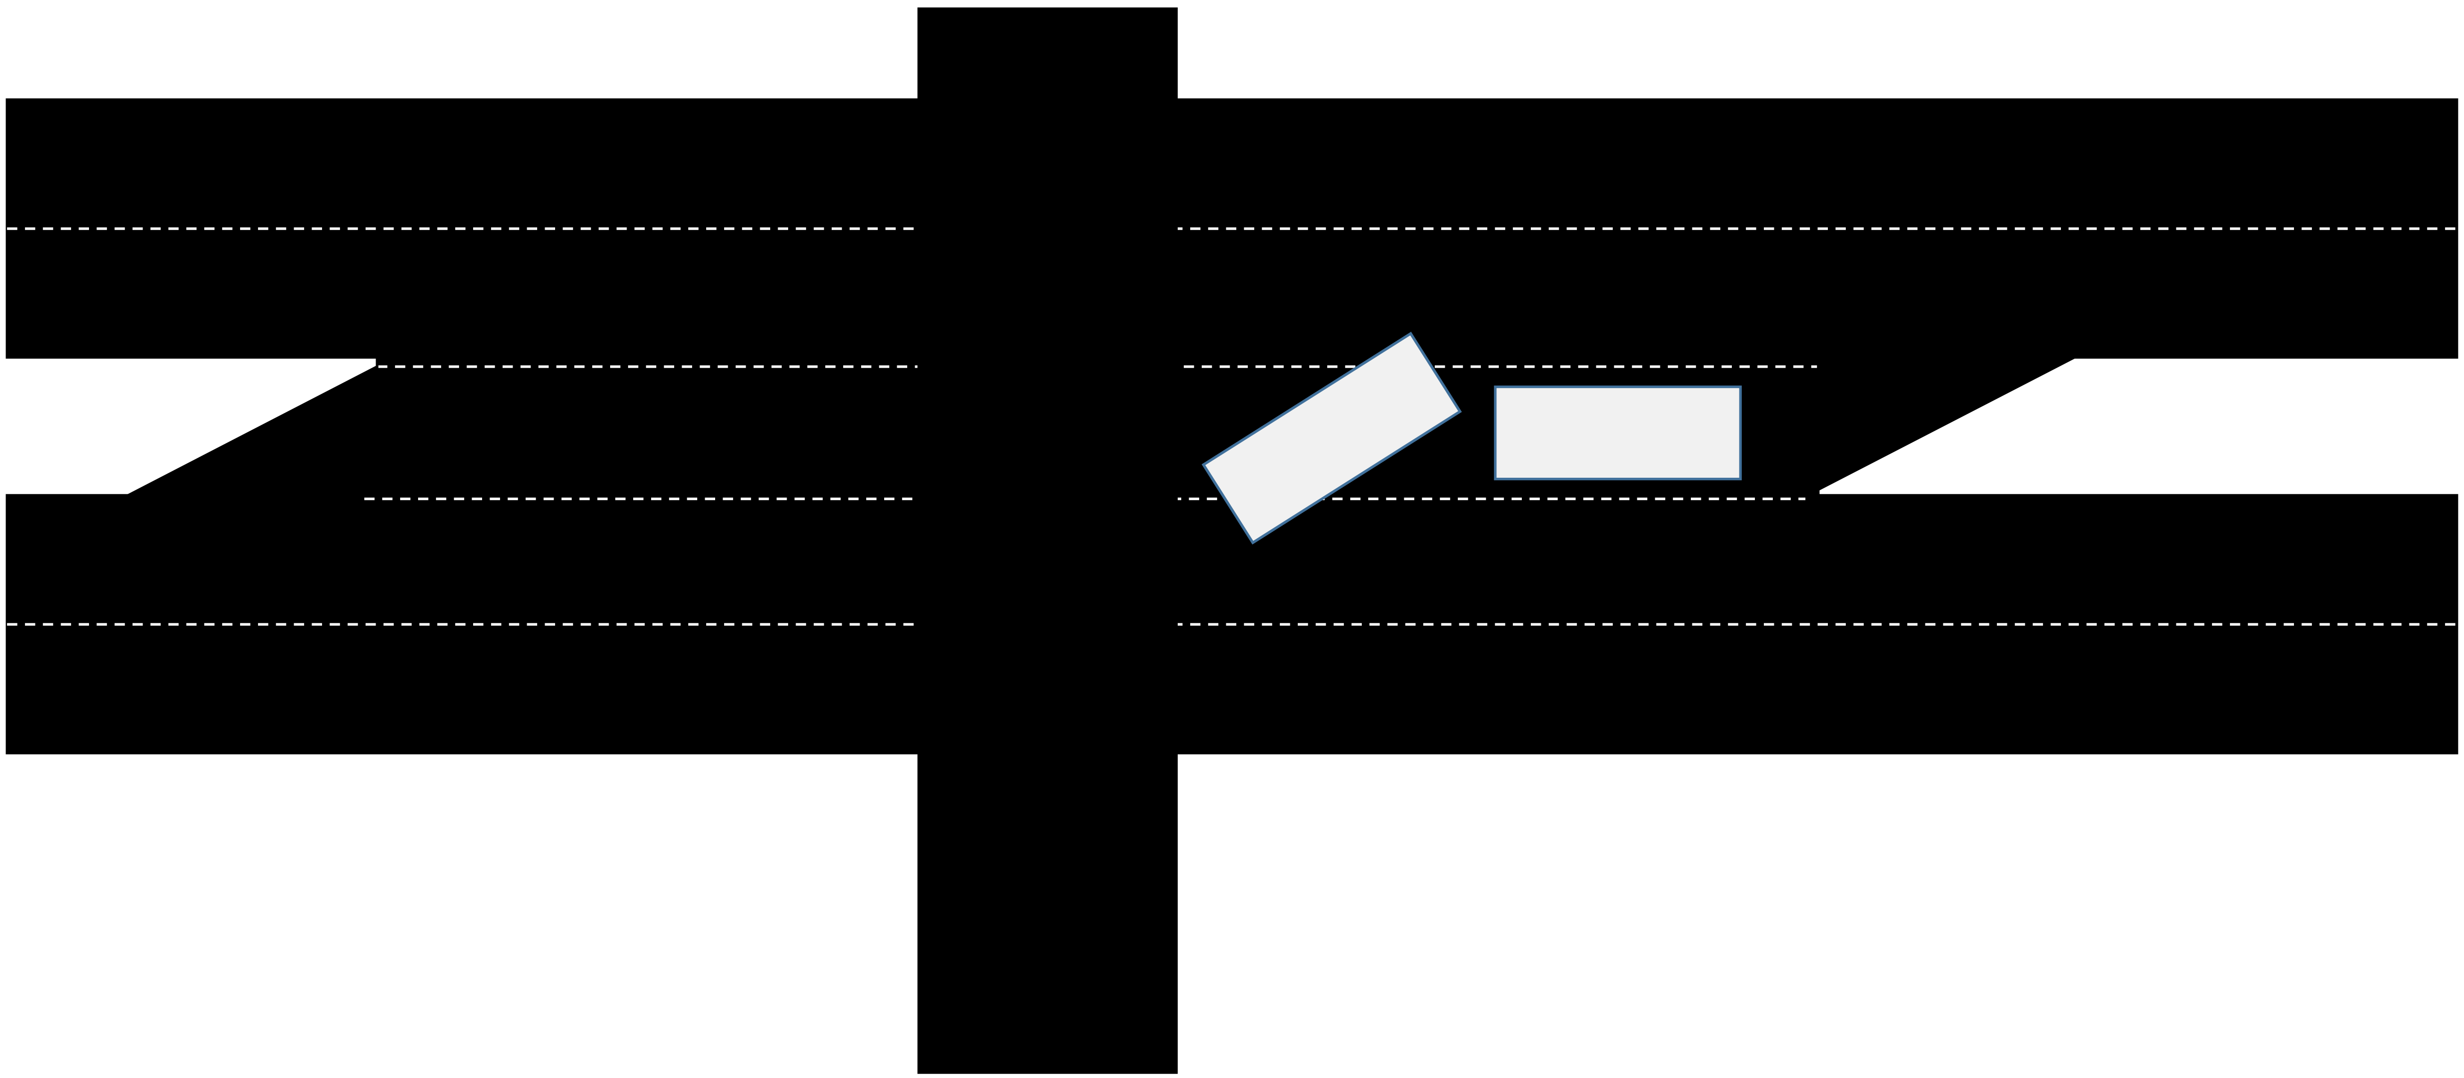

2nd Fatal Tesla Autopilot Crash
Analysis

- Similar Model S sensors and features to 1st Tesla Autopilot crash
- No braking or avoidance action prior to collision
- Tesla commented that the camera failed to detect the truck due to white color of the trailer against a brightly lit sky and a high ride height.
- They further commented that the radar filtered out the truck as an overhead road sign to prevent false braking.
- In both cases MobilEye commented that:
  - MobileEye's system was not designed to cover all accident scenarios and that Tesla was using it outside of its intended purpose.

3rd Tesla Autopilot Crash

2018 - January 22nd - Non-Fatal – Tesla Model S (California)

Tesla Collision with Fire Truck

- Tesla Model S in Autopilot mode was following a pickup truck in left lane
- Pickup changed lanes to avoid a stationary firetruck
- Tesla accelerated into the back of the firetruck at 65 m.p.h

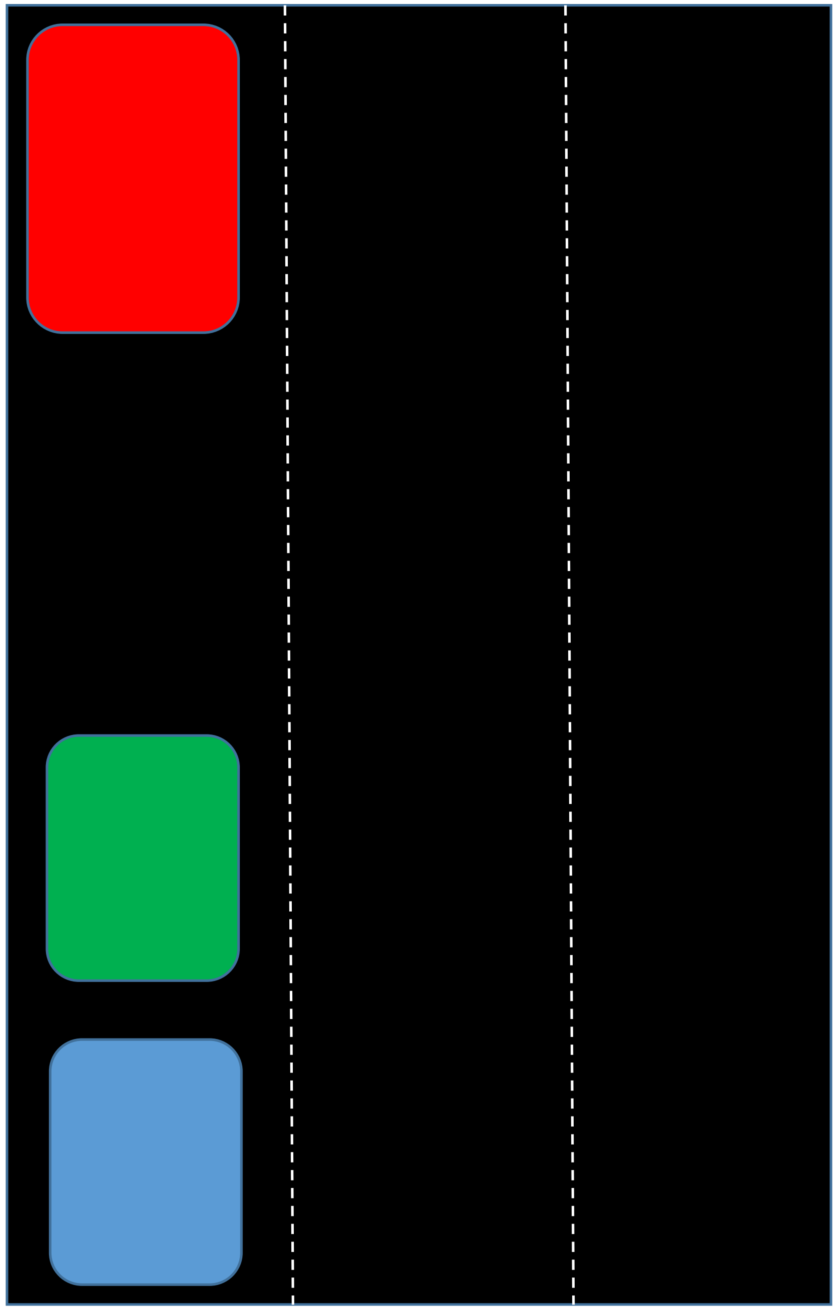

Similar Autopilot, lane changing lead vehicle & stationary vehicle failure

Autopilot, lane changing lead vehicle & stationary vehicle

- Tesla Model S Handbook states:

“Traffic-Aware Cruise Control cannot detect all objects and may not brake/decelerate for stationary vehicles, especially in situations when you are driving over 50 mph (80 km/h) and a vehicle you are following moves out of your driving path and a stationary vehicle or object is in front of you instead.”

Why the acceleration?

- ACC is part of Autopilot
- Set max speed (normal cruising speed) & time gap (headway) when following a lead vehicle @speed < max speed

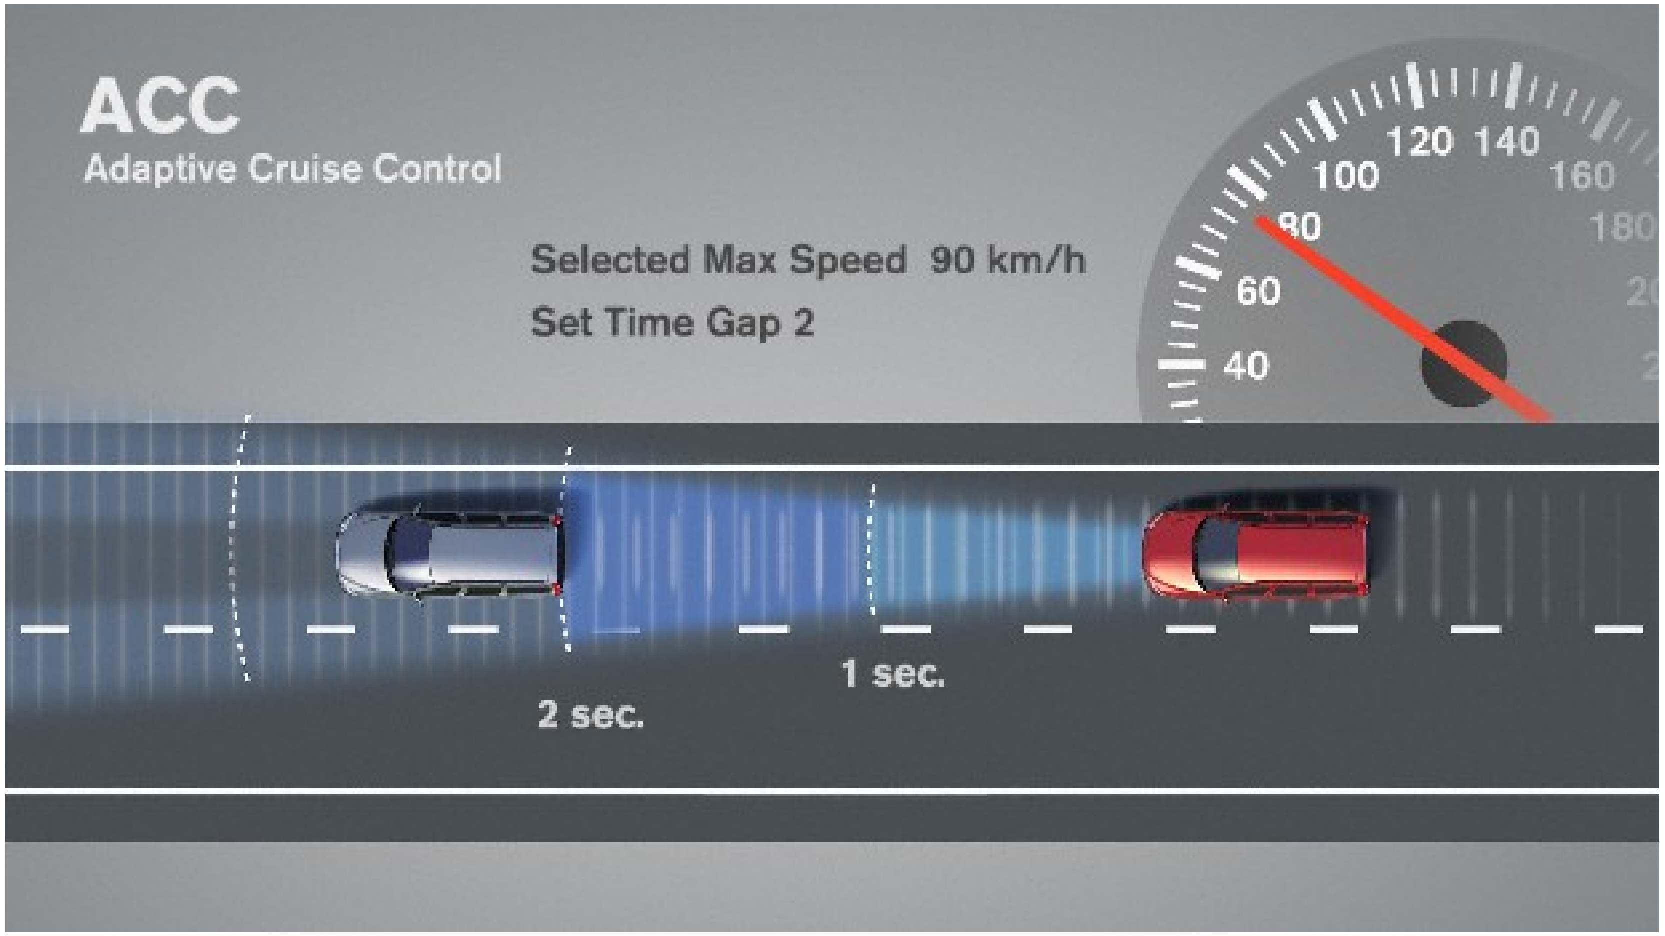

Hypothesis:

- When pickup changed lane distance to new lead vehicle (firetruck) increased
- ACC commanded acceleration to close the gap

Another Tesla Autopilot Crash show what this might be like at full speed

Uber Autonomous Vehicle Crash

2018 - March 18th - Fatal – Uber Volvo XC90 (Arizona)

Uber Accident Details

- Uber
  - Switched off Volvo's standard Aptiva/Intel Mobile Eye collision avoidance/mitigation system
    - Initially detected unknown object 6 seconds before impact
    - It decided it was a bicycle 1.3 second before impact and would have started braking
  - Why?
    - To reduce interference with their software? Avoid false positives?
    - Think of trying to making a right turn @Yonge & Dundas in Toronto
  - Also switched off Volvo's Driver Distraction Detection System

  - What's a poor autonomous vehicle to do?
    - Maybe requiring having these features turned on by an industry standard assurance case would help!

4th Tesla Autopilot Crash

2018 - March 23rd - Fatal - Tesla Model X (California)

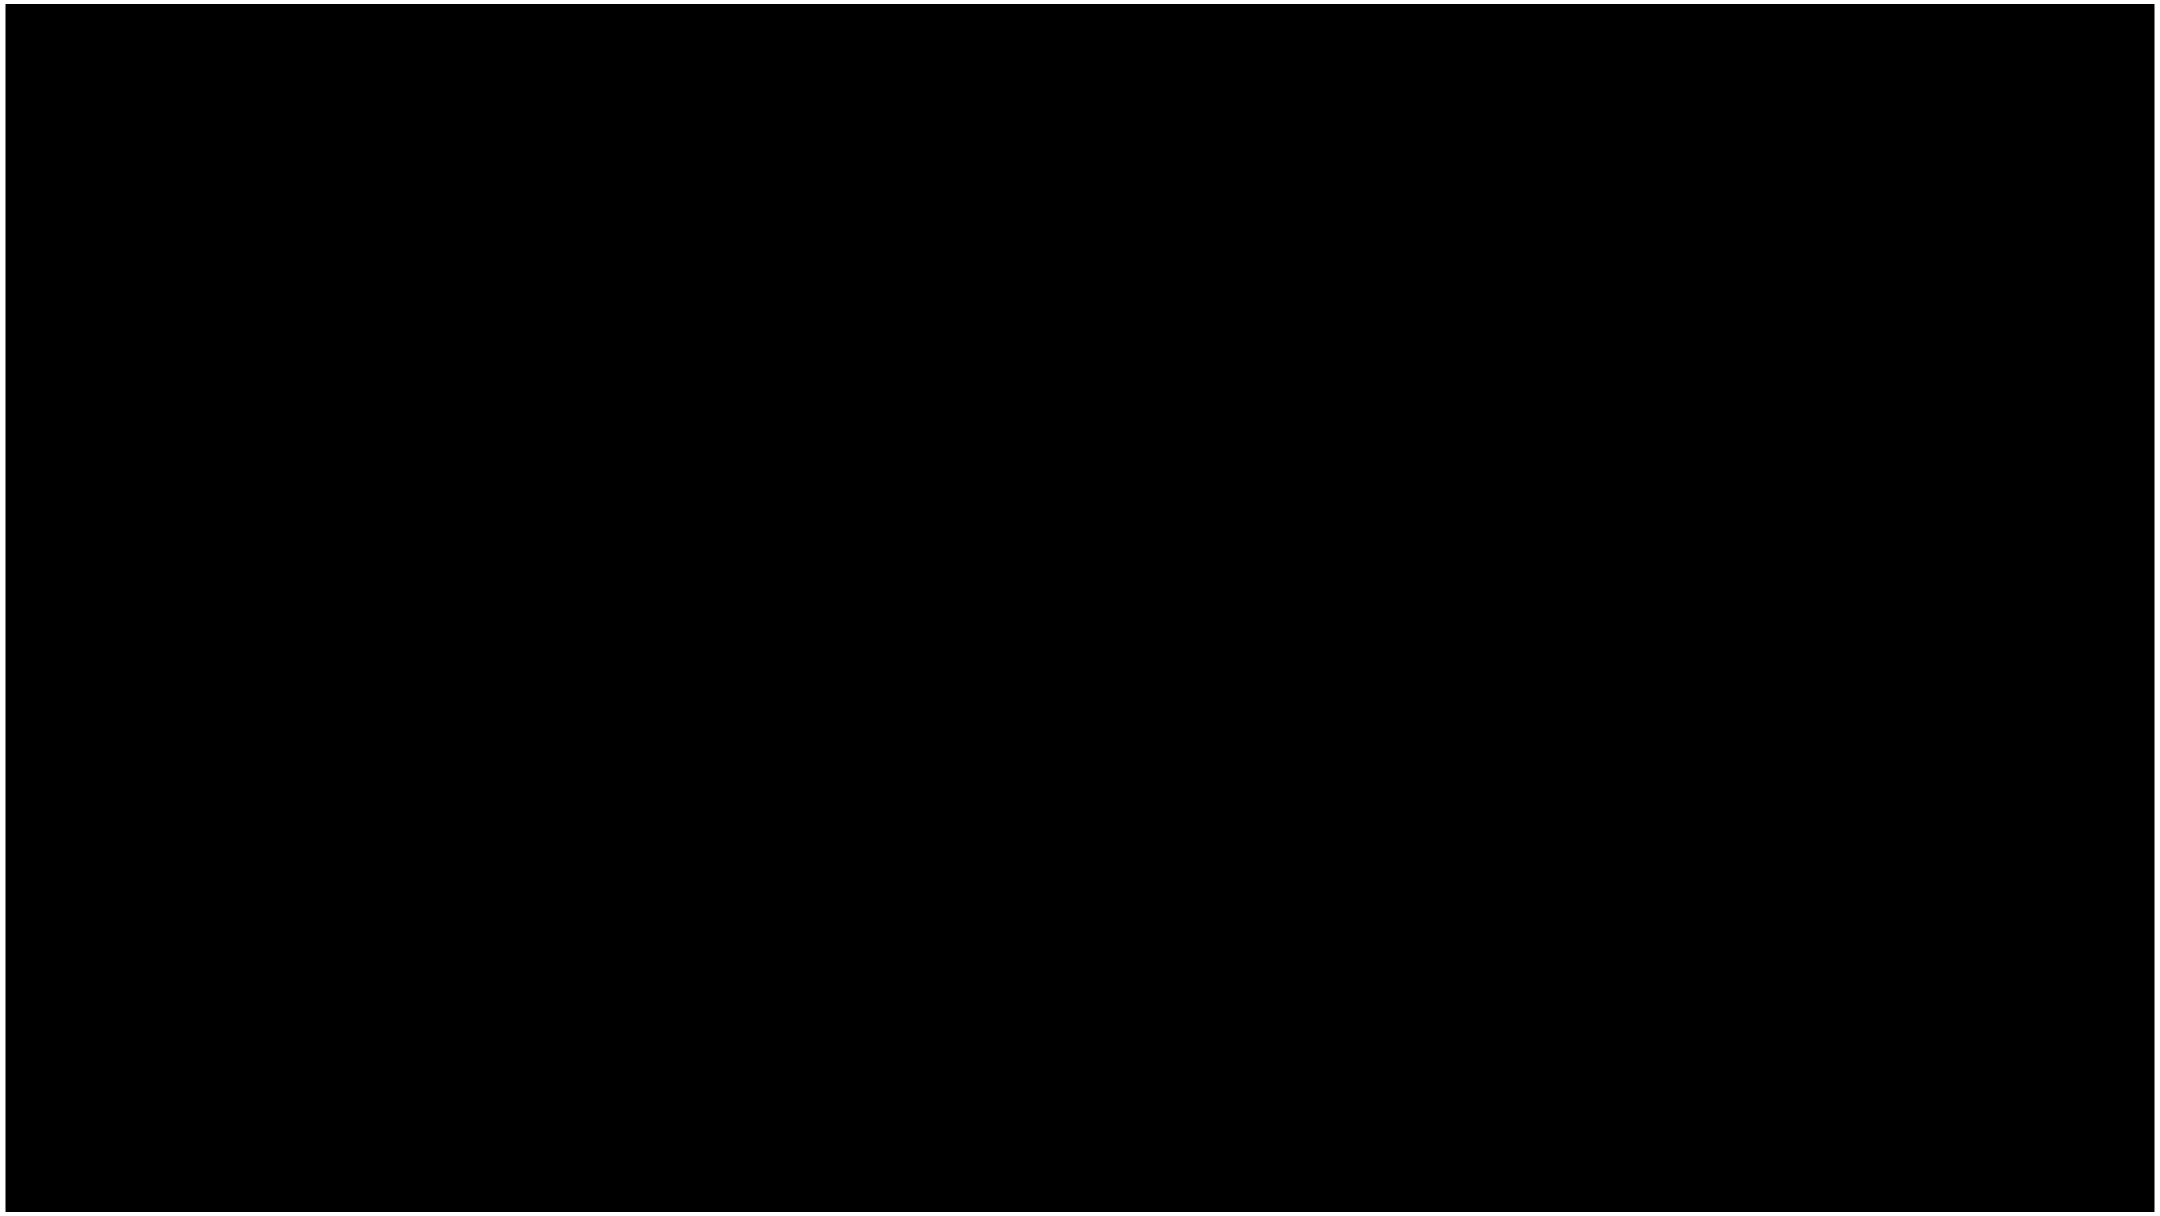

4^{th} Tesla Autopilot Crash

Analysis

NTSB preliminary report summary states:

- During the 60 seconds prior to the crash, the drivers hands were detected on the steering wheel on three separate occasions, for a total of 34 seconds;
- for the last 6 seconds prior to the crash, the vehicle did not detect the driver's hands on the steering wheel.
- At 8 seconds prior to the crash, the Tesla was following a lead vehicle and was traveling about 65 mph.
- At 7 seconds prior to the crash, the Tesla began a left steering movement while following a lead vehicle.
- At 4 seconds prior to the crash, the Tesla was no longer following a lead vehicle.
- At 3 seconds prior to the crash and up to the time of impact with the crash attenuator, the Tesla's speed increased from 62 to 70.8 mph, with no pre-crash braking or evasive steering movement detected.

4^{th} Tesla Autopilot Crash

Analysis

- Tesla stated after the accident:
  - “The driver had about five seconds and 150 meters of unobstructed view of the concrete divider with the crushed crash attenuator, but the vehicle logs show that no action was taken.“

- Oddly enough, Tesla failed to mention that the Tesla sensors and AEB had the exact same opportunity to see the concrete divider and react in a timely fashion to mitigate the outcome

A similar Tesla crash


Laguna Ca
1
133
Laguna Canyon Rd
2
133
Map
Google
3D
Imagery ©2018 Google, Map data ©2018 Google
Canada
Terms
Send feedback
5 m

Following lane marks – to an accident

1. Location of Police vehicle
2. Right hand lane marker as road starts to widen for turn lane

Probably during “rush hour” no vehicles park there

Main Fallacy in existing (implicit) Assurance Cases for ADAS

- The driver is going to catch the Machine Learning (ML) failures . . . without driver attentiveness monitoring!

Getting too (artificially) intelligent with safety

- Object identification is very useful
- Can help predict and plan in addition to help partially meet some safety goals
- Pedestrian detection is an example of how ML fails badly with the key safety requirement: “Don’t hit things!”

AI/ML Version:

“I don’t know what it is so it’s not there.”

vs

Safety Version:

“I don’t know what it BUT IT’S THERE!”

If ML Doesn’t Recognize It, It’s Not There

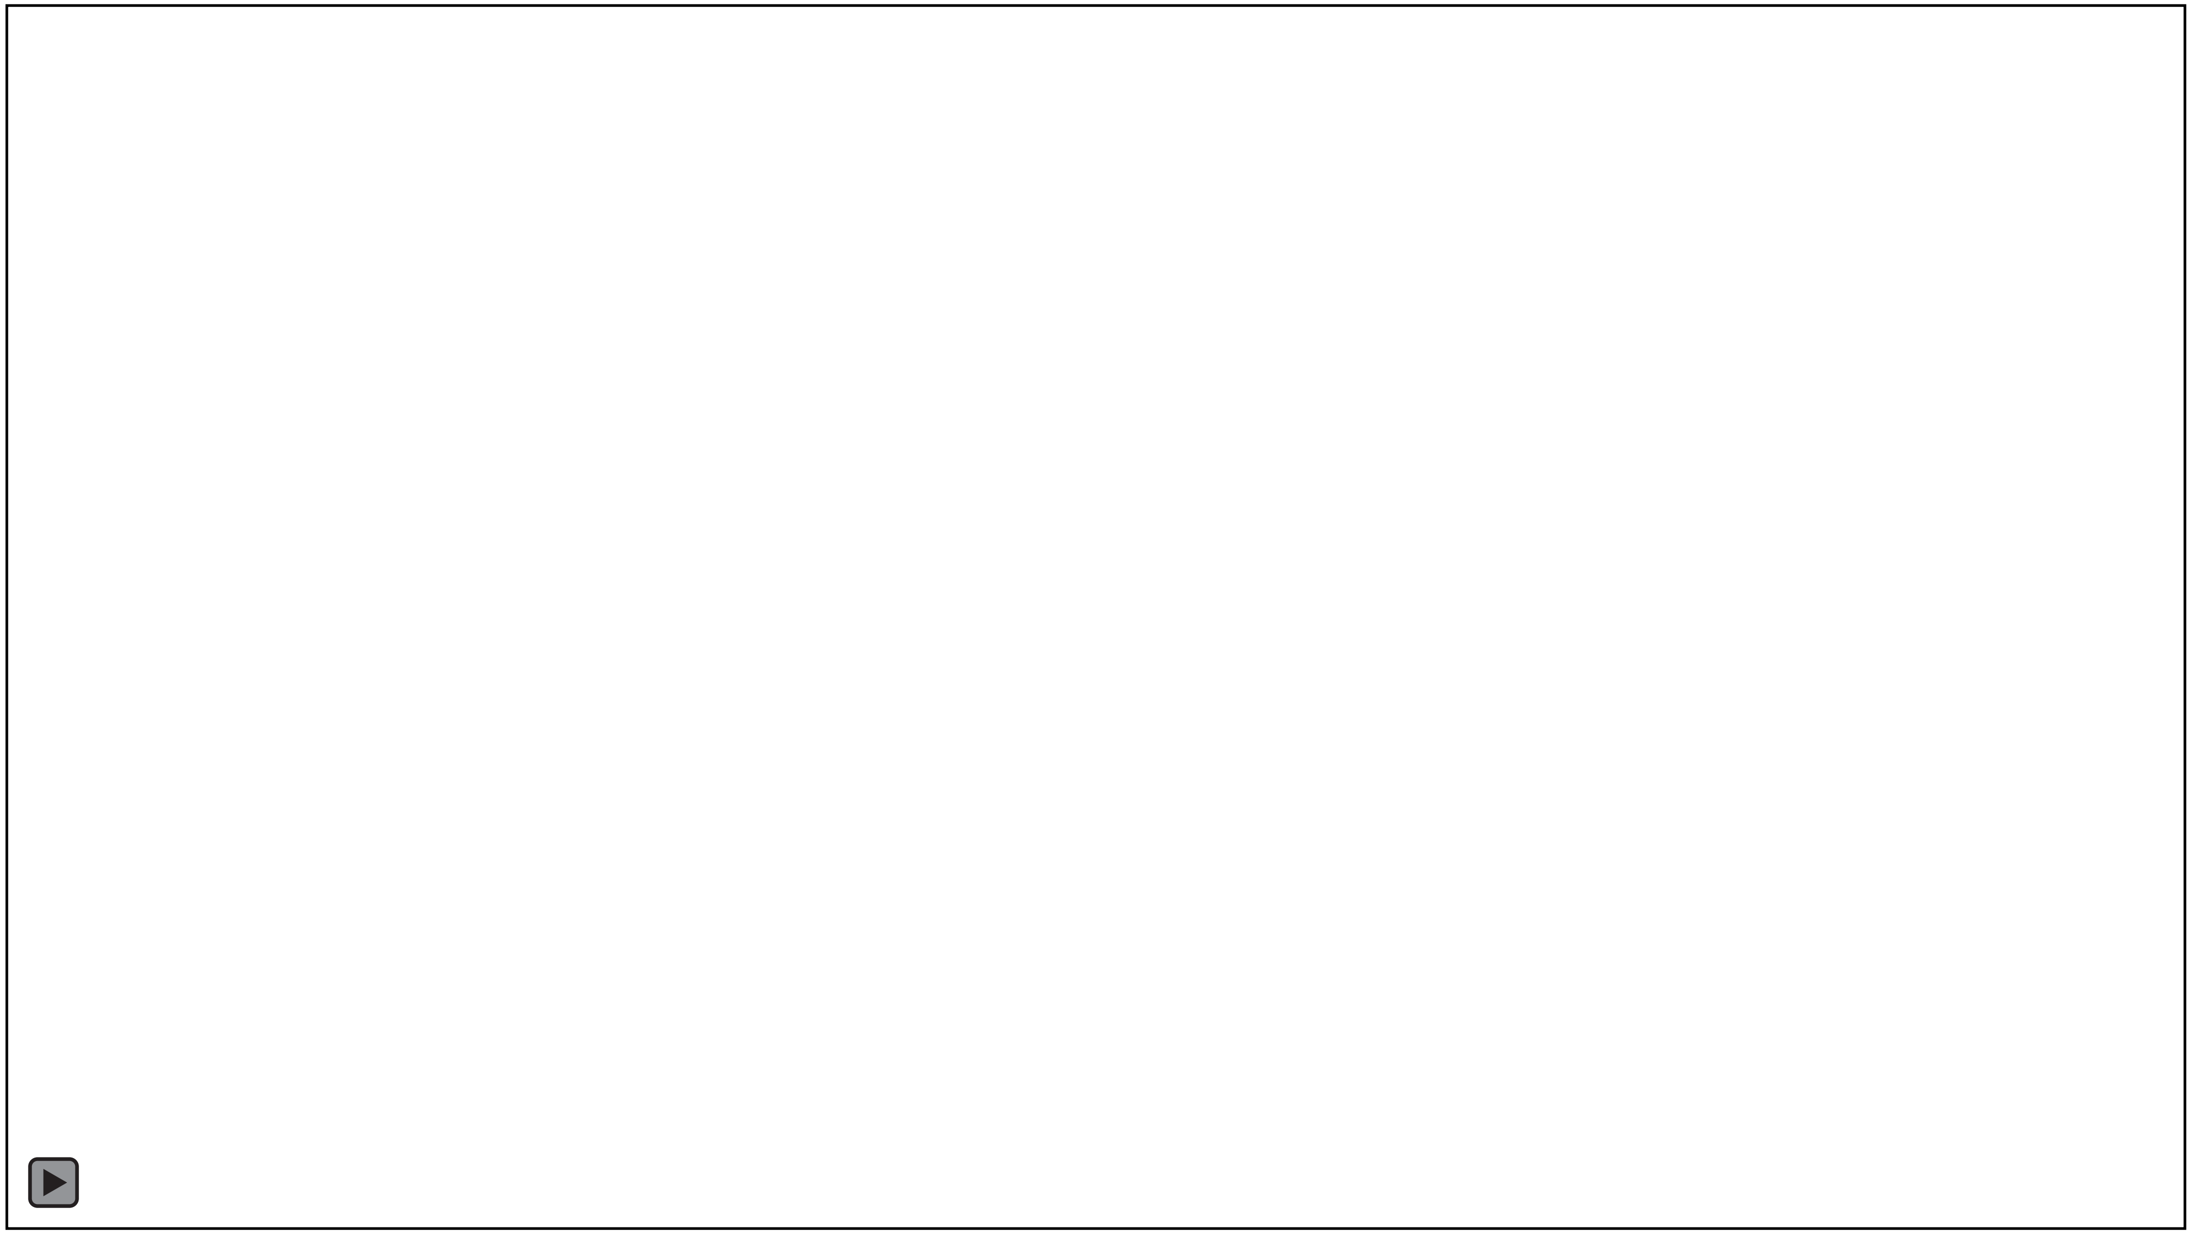

The trouble with AI in safety critical situations

- Using ML to deal with cross walks:
  - AI does a good job with this but not …

Lessons learned

- Production is currently taking precedence over safety and that is resulting in accidents
- The driver is not a sufficient mitigation without *real* driver attentiveness monitoring
- Interactions with other systems requirements is compromising safety (ACC acceleration in stopped vehicle accidents, interactions between control loops at different time scales)
- Current systems are not providing confidence information from ML components resulting in unsafe behaviour
  - When in doubt, slow down!
- New failure modes not discussed here – maintenance
  - replacing your windshield can now cause accidents due to sensor calibration errors!

Proper Monitoring of Driver Attentiveness

Super Cruise
Tested on Cadillac CT6

Automation System Rating

Super Cruise uses a camera to watch where the driver's eyes are looking.

- Capability & Performance
- Ease of Use
- Clear When Safe to Use
- Keeping Driver Engaged
- Unresponsive Driver

Autopilot
Tested on Tesla X/S/3

Automation System Rating

Autopilot performed well and is easiest to use in stop-and-go traffic.

- Capability & Performance
- Ease of Use
- Clear When Safe to Use
- Keeping Driver Engaged
- Unresponsive Driver

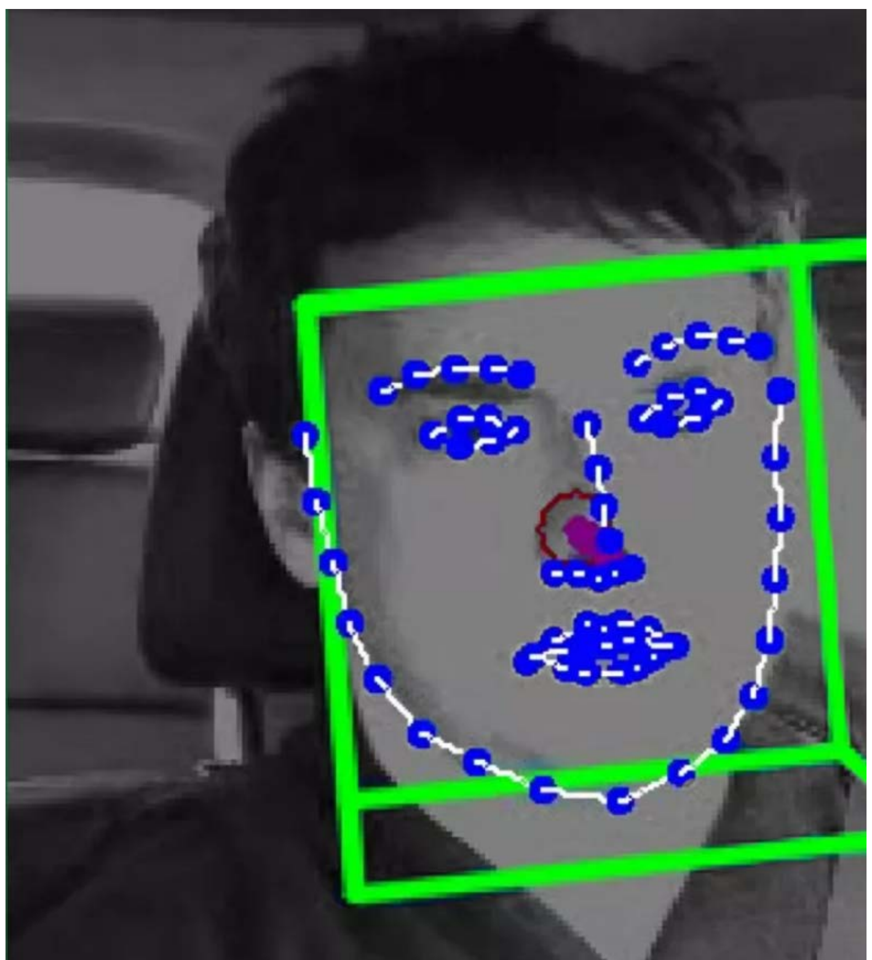

Also work from MIT (see Lecture 2 of MIT course on Deep Learning and Self-Driving)

Self-Driving Car Tasks

- Localization and Mapping – Where am I
- Scene Understanding – Where is Everyone Else?
- Movement Planning – How to get from Point A to Point B
- Driver State – What is the Driver Up to?
  - Essential if driver is part of the loop!
- Safety Monitoring

Safety Assurance of ADS

Source: Krzysztof Czarnecki, Waterloo

Operational Design Domain (ODD)

SAE J3016 Levels of Driving Automation

A set of **conditions** under which
the driving automation can operate a vehicle

Time of day
day
night

Types of roads
residential
urban
highway

Geographic area

Traffic conditions
stop-and-go
free flowing

Weather conditions
clear
raining
snowing
icy






Dynamic Driving Task (DDT) Fallback

Who performs the DDT
in the case of **system malfunction** or
when **leaving the ODD**?

Driver

Computer

Automated Driving Systems (ADS)

	3	4	5
SAE J3016 Levels of Automation	Conditional Automation	High Automation	Full Automation
Example	ADS for stop-and-go	Shuttle in geo fenced area	Robo Taxi anywhere
ODD	**limited**		**unlimited**
Fallback	Driver	ADS	

ADS Hazard Sources

Mature best practices

Mechanical faults

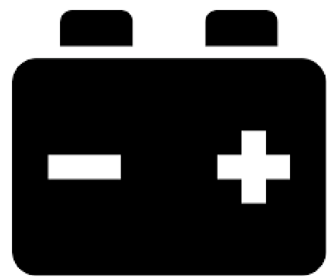

Electrical faults

ISO 26262

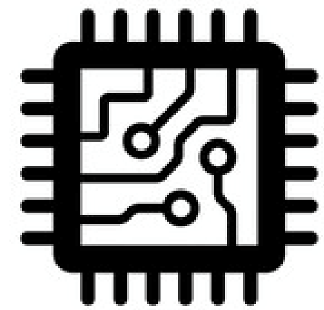

Computer HW faults

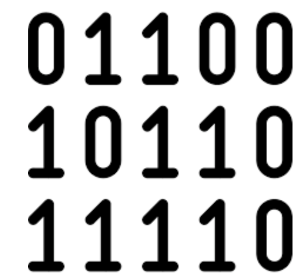

Computer SW faults

(ISO / PAS 21448)

Sensor noise & limitations

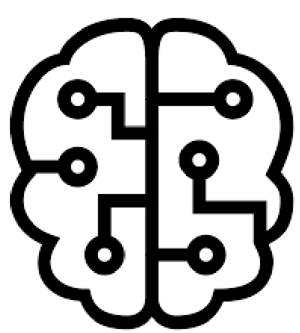

Machine learning errors

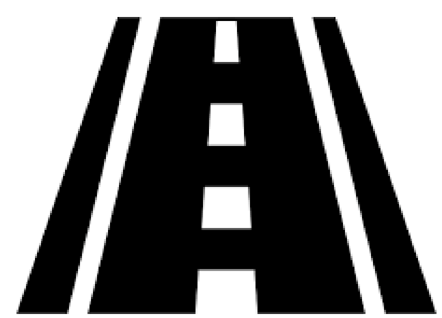

Inadequate driving behavior

DDT fallback failures

SAE J3061

Cyber attacks

Assurance: ISO 26262

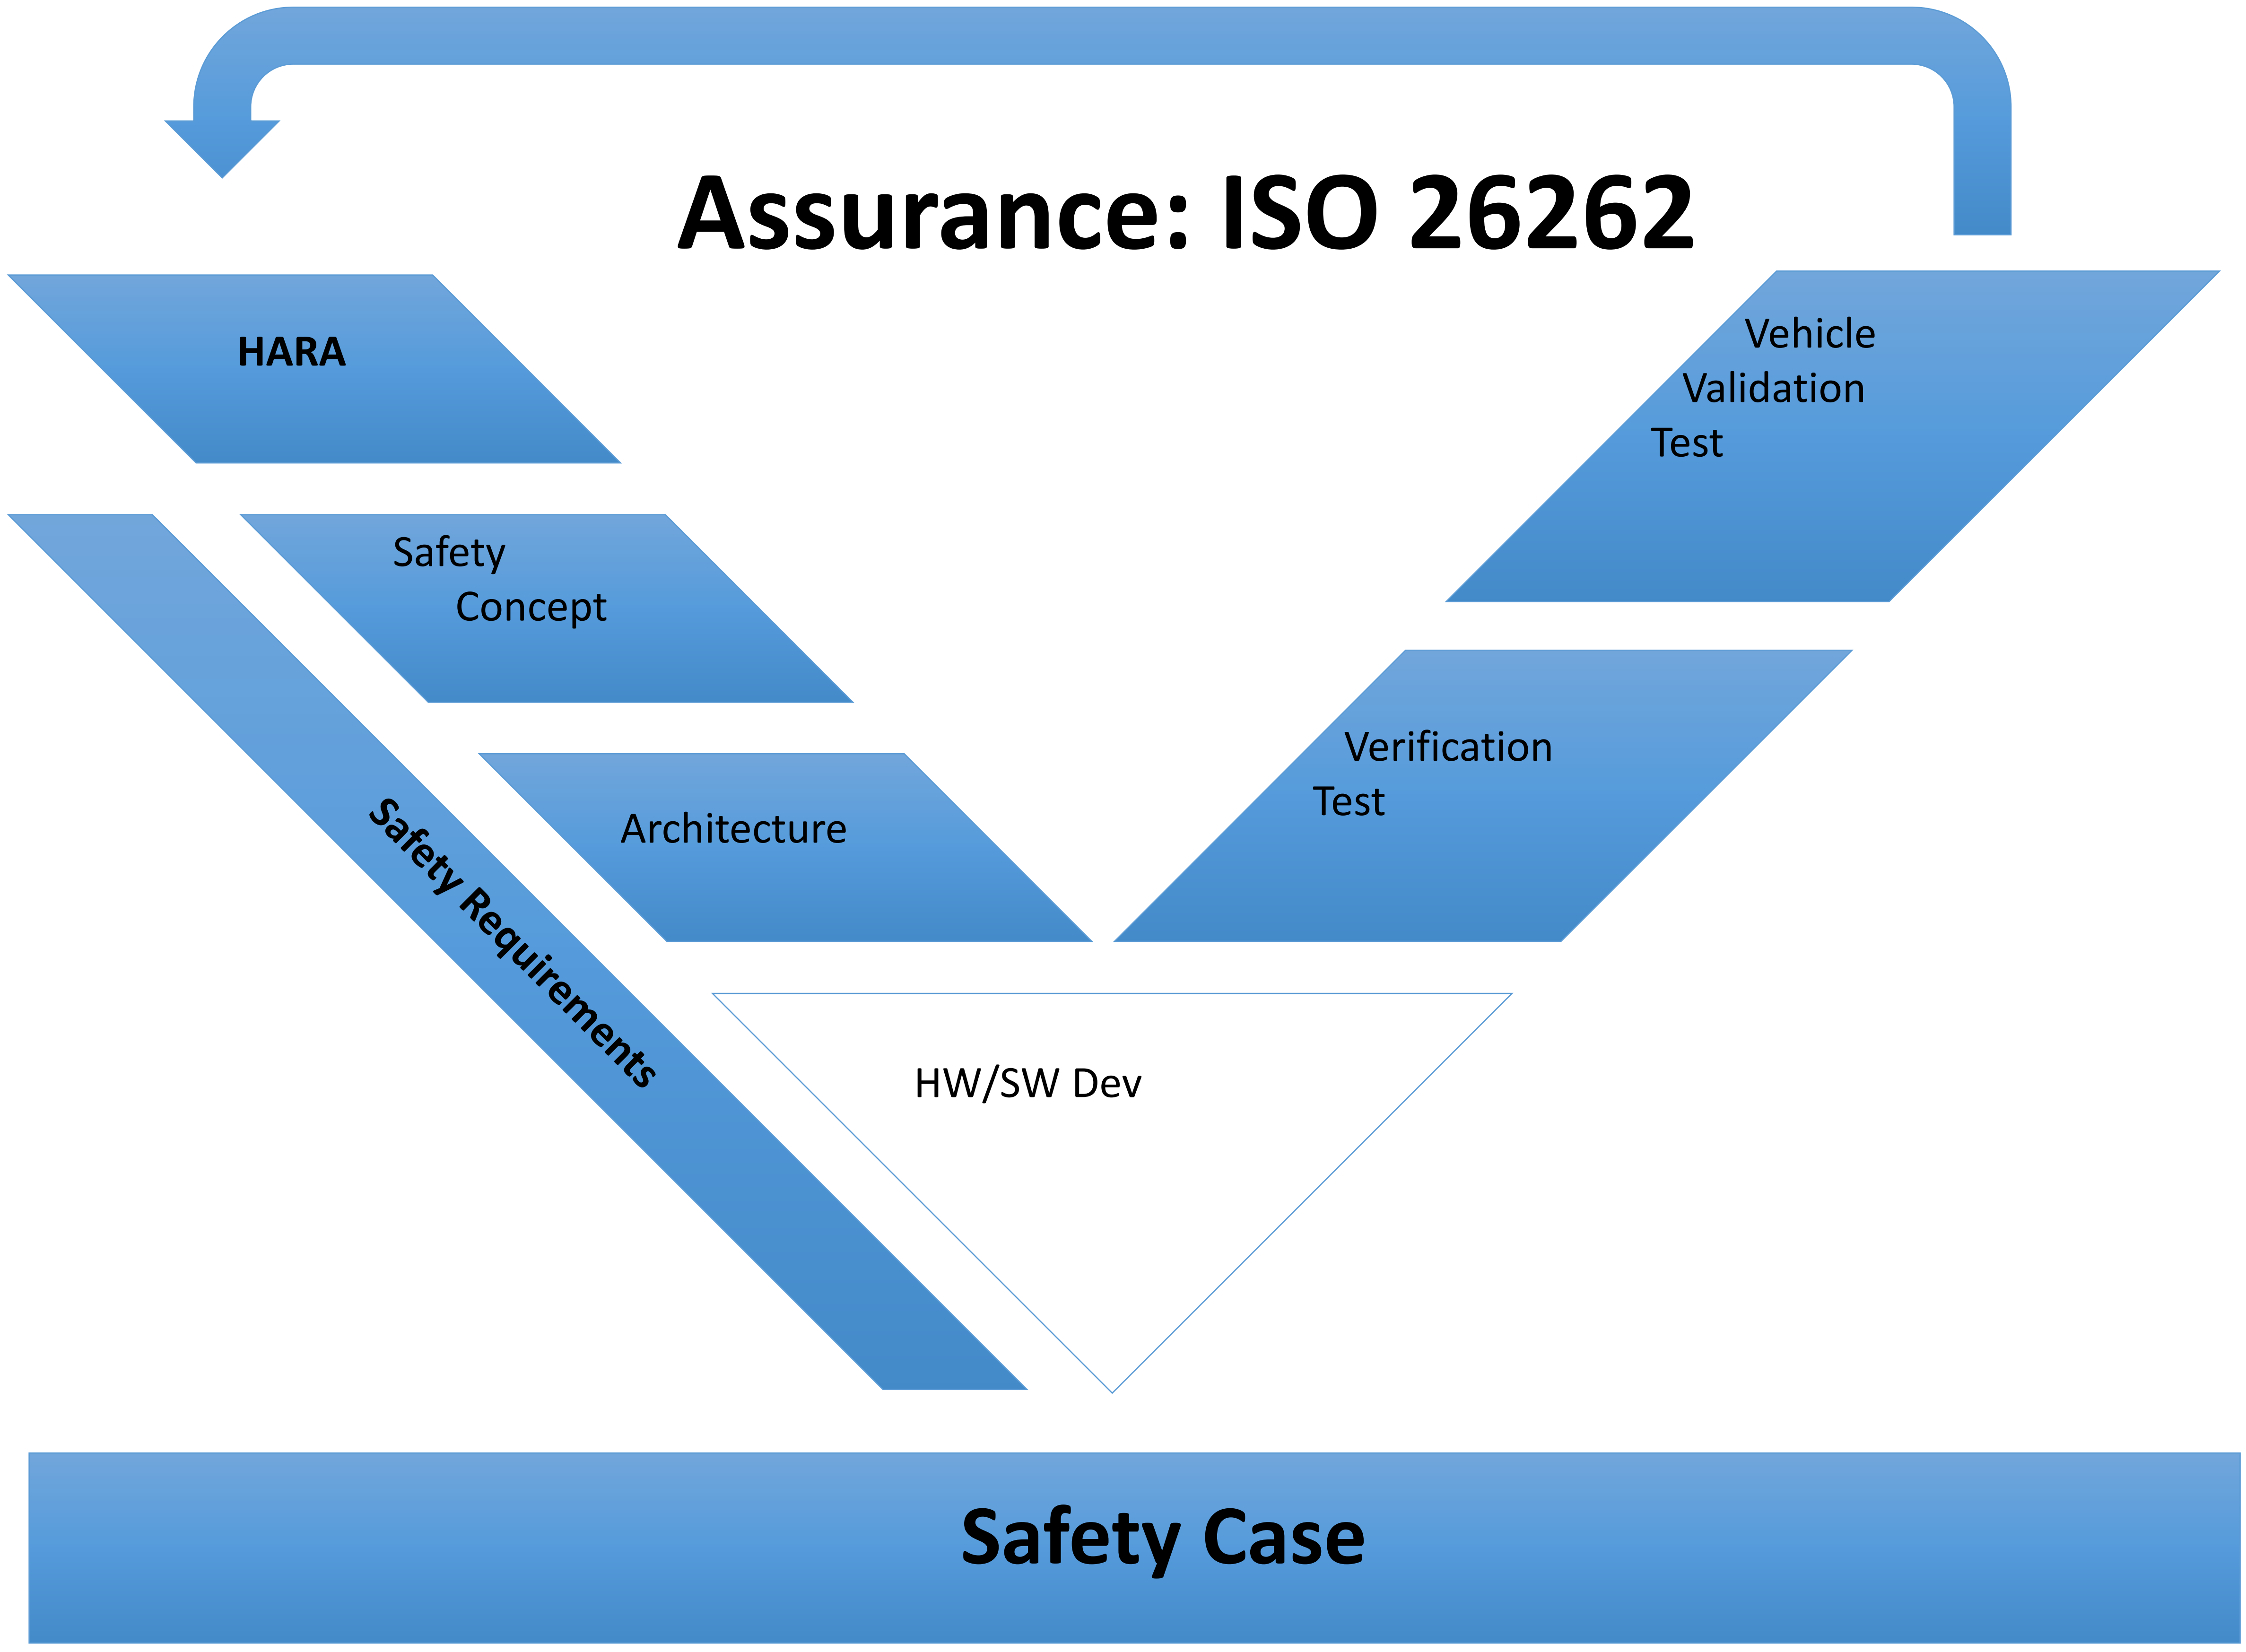

ADS Hazard Sources

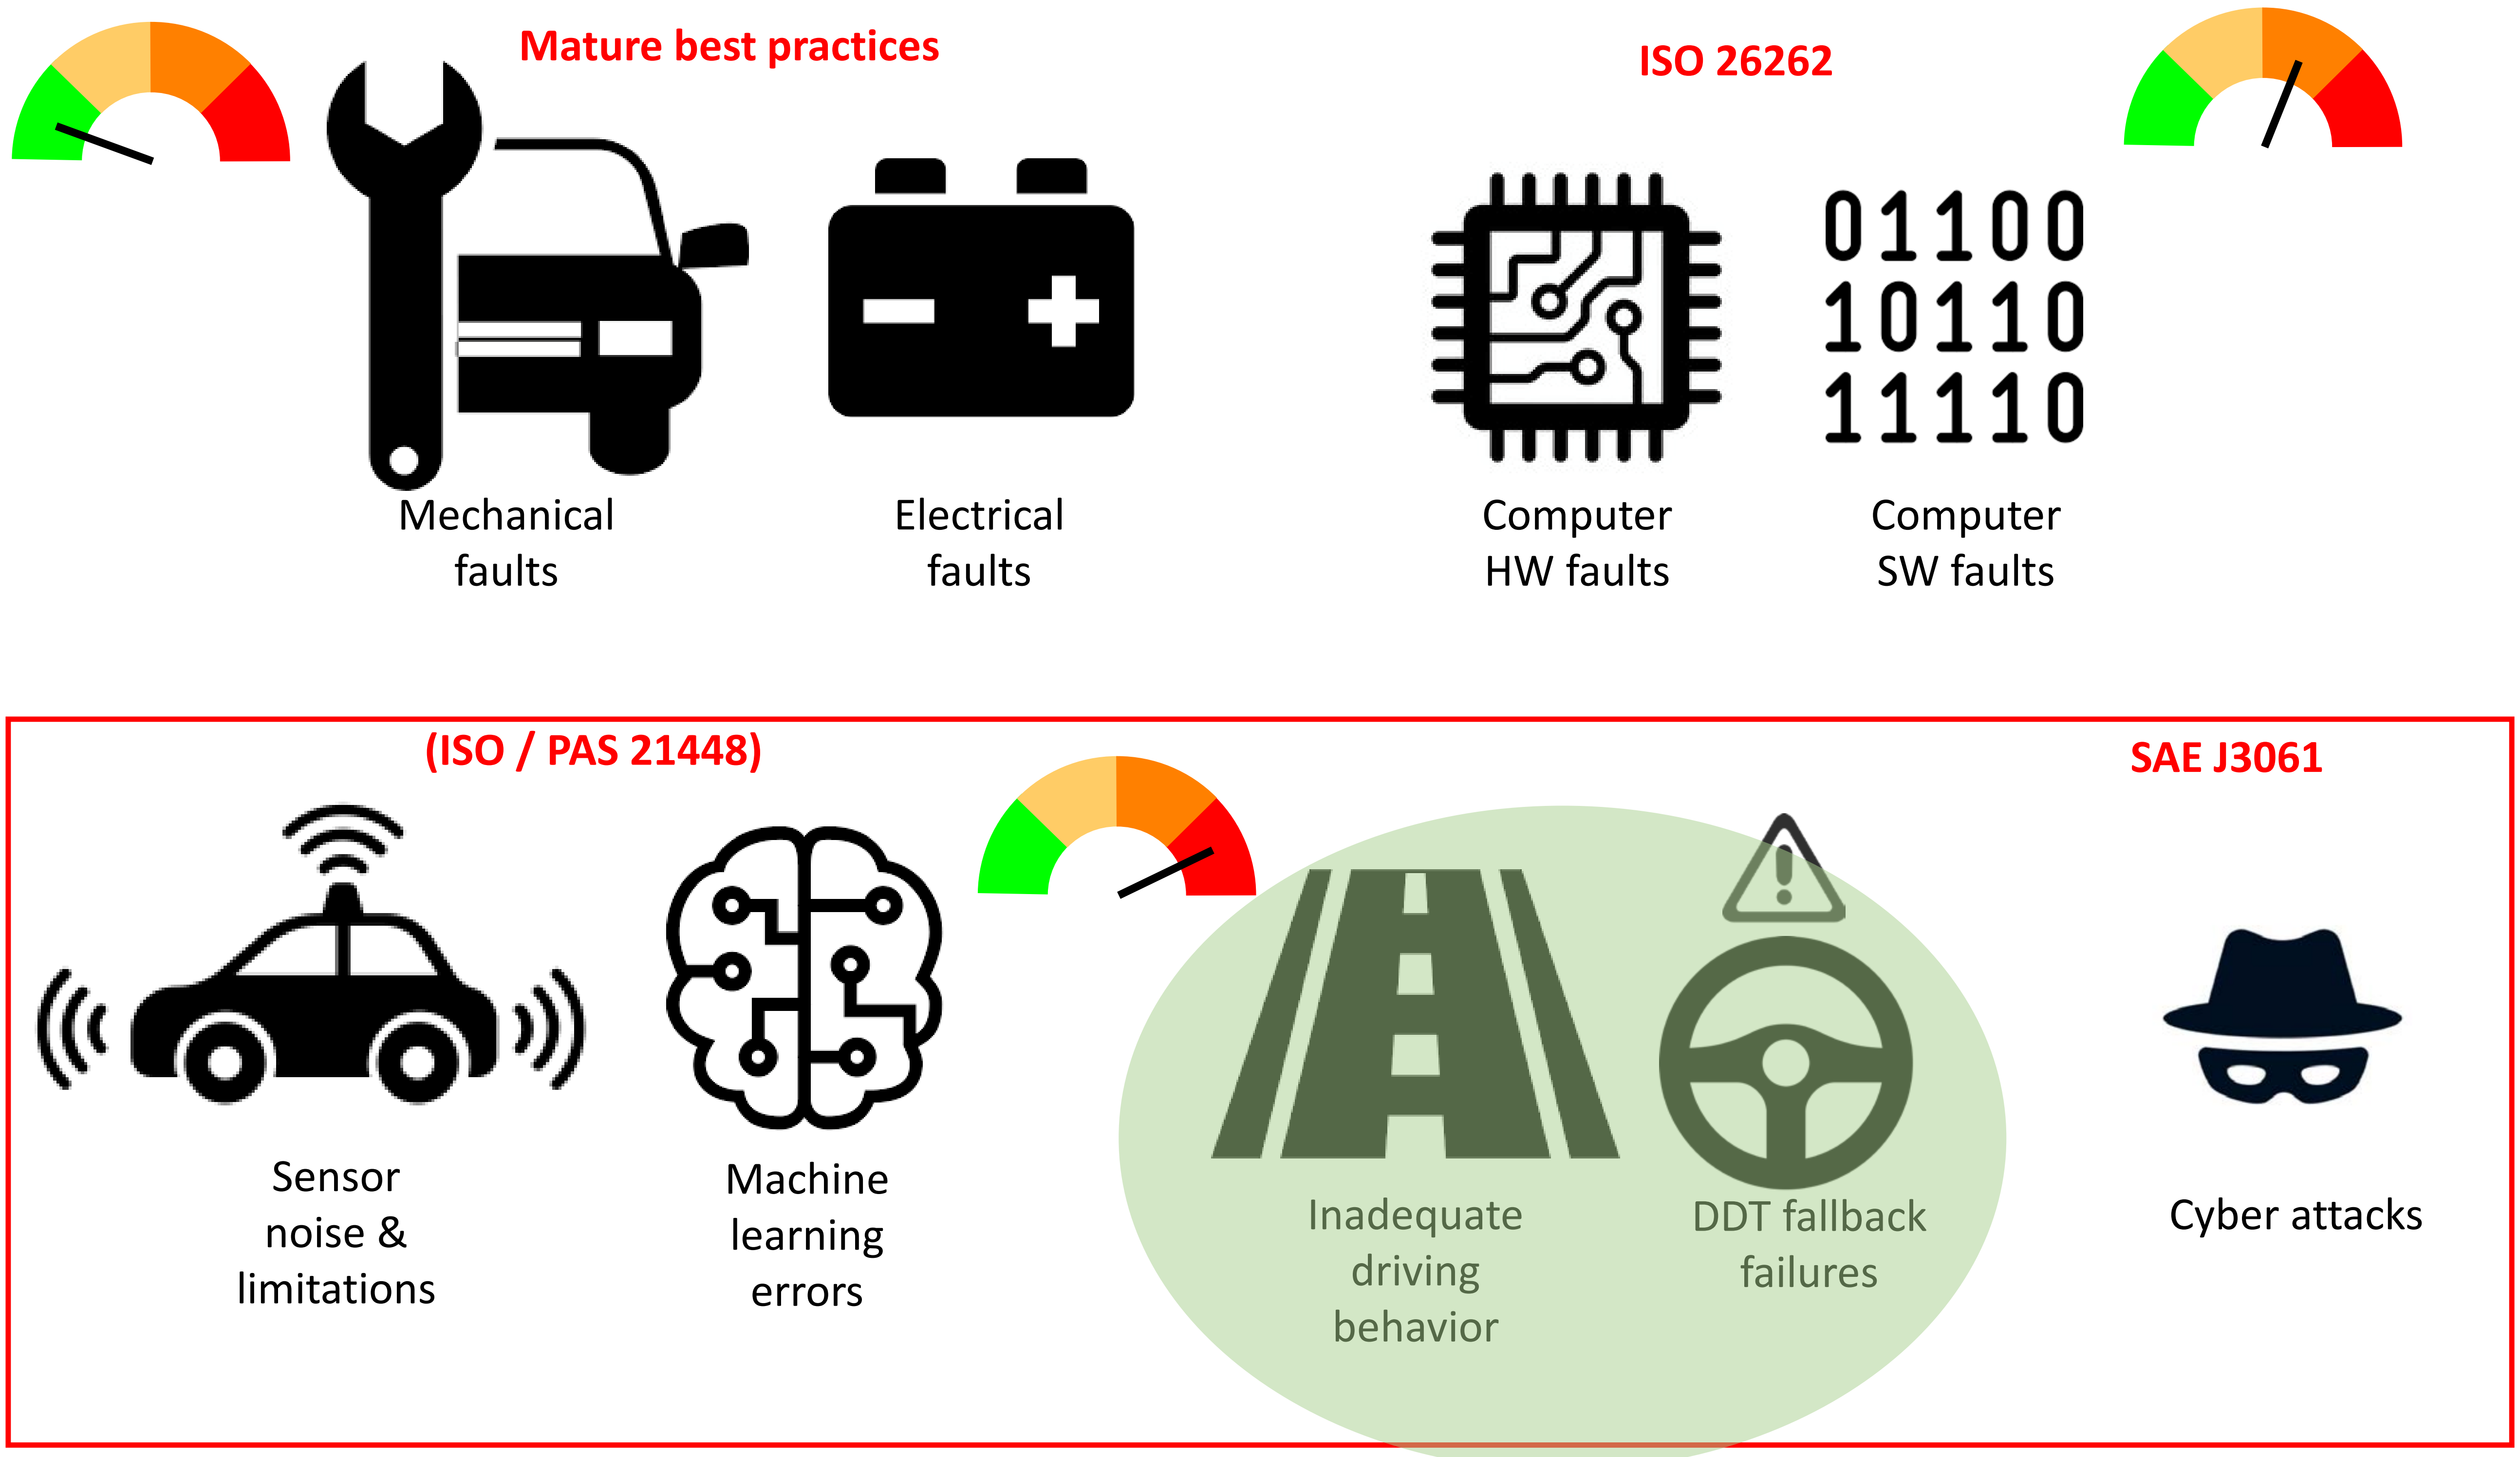

Safety

Absence of <u>unreasonable</u> risk of mishap

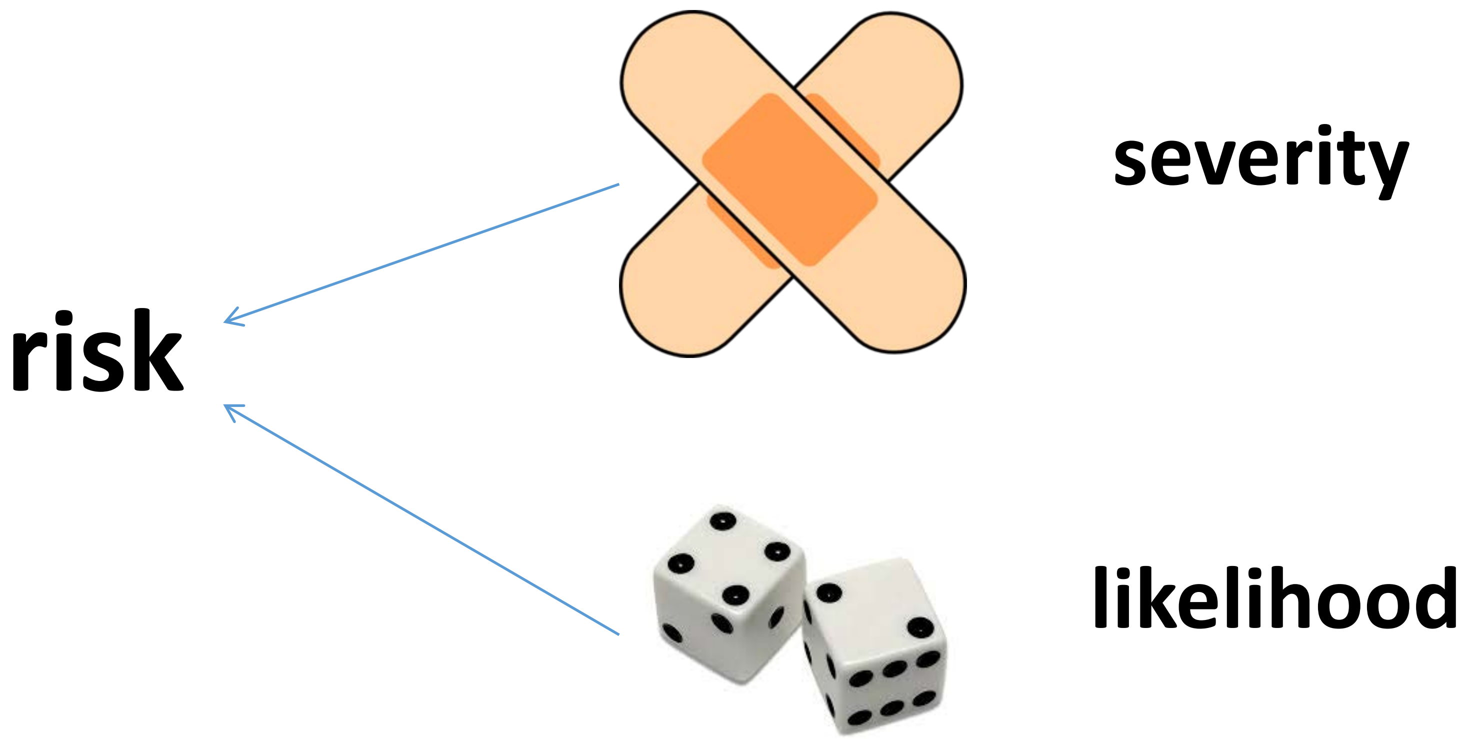

Driving Behavior Safety

Absence of unreasonable crash risk due to ADS driving behavior

Noncollisions

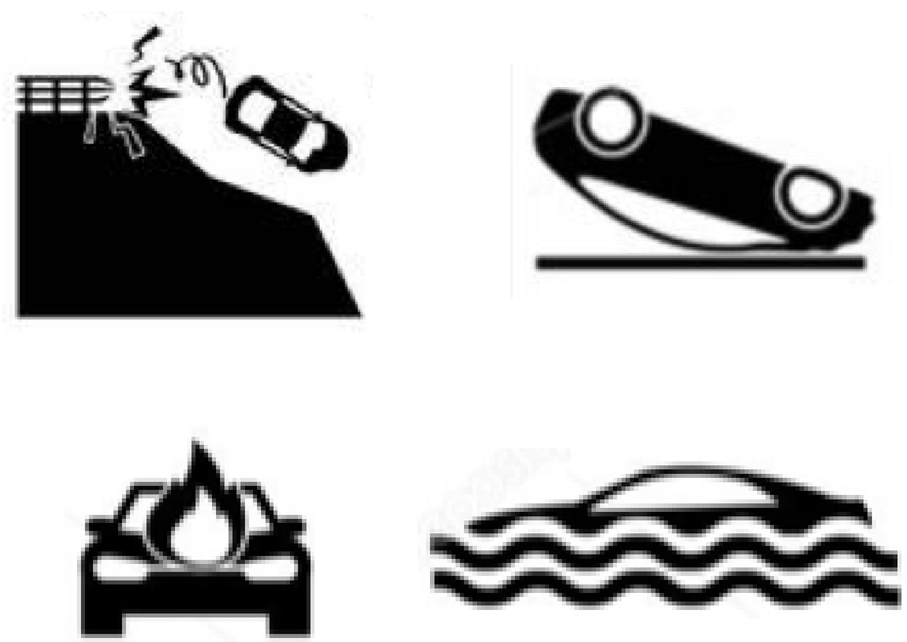

Collisions

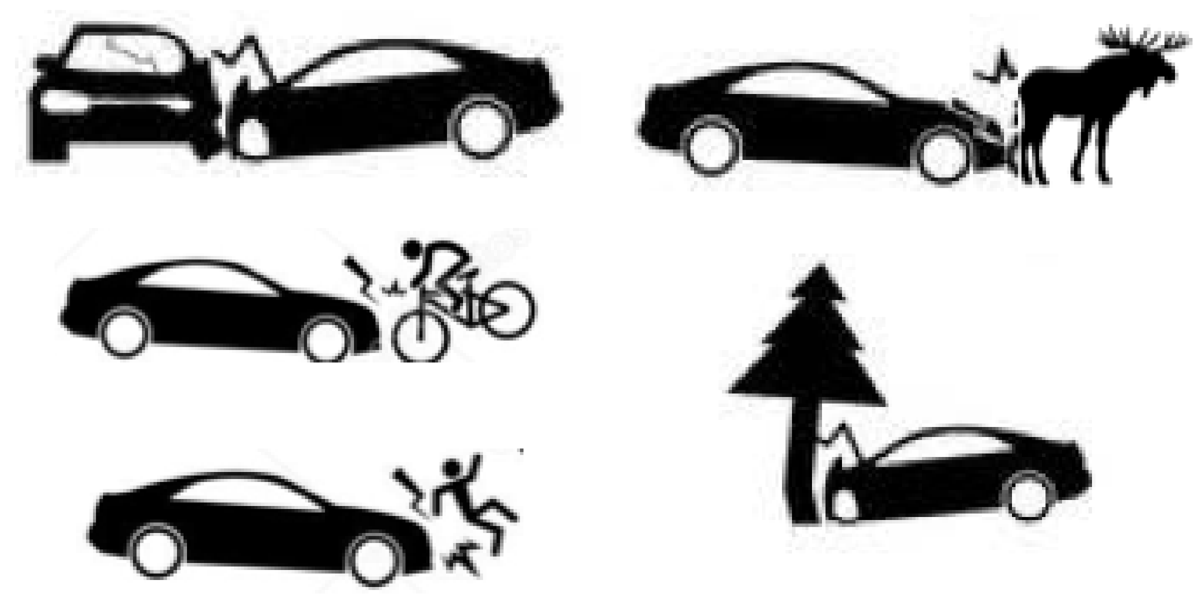

Factors Influencing Risk Acceptability

- Risk level
- Risk reduction cost
- Benefit of the risky functionality (risk taking)
- Best practice (state of technology)
- Replacement risk
- Who controls risk
- Perception/public opinion

Assurance Target

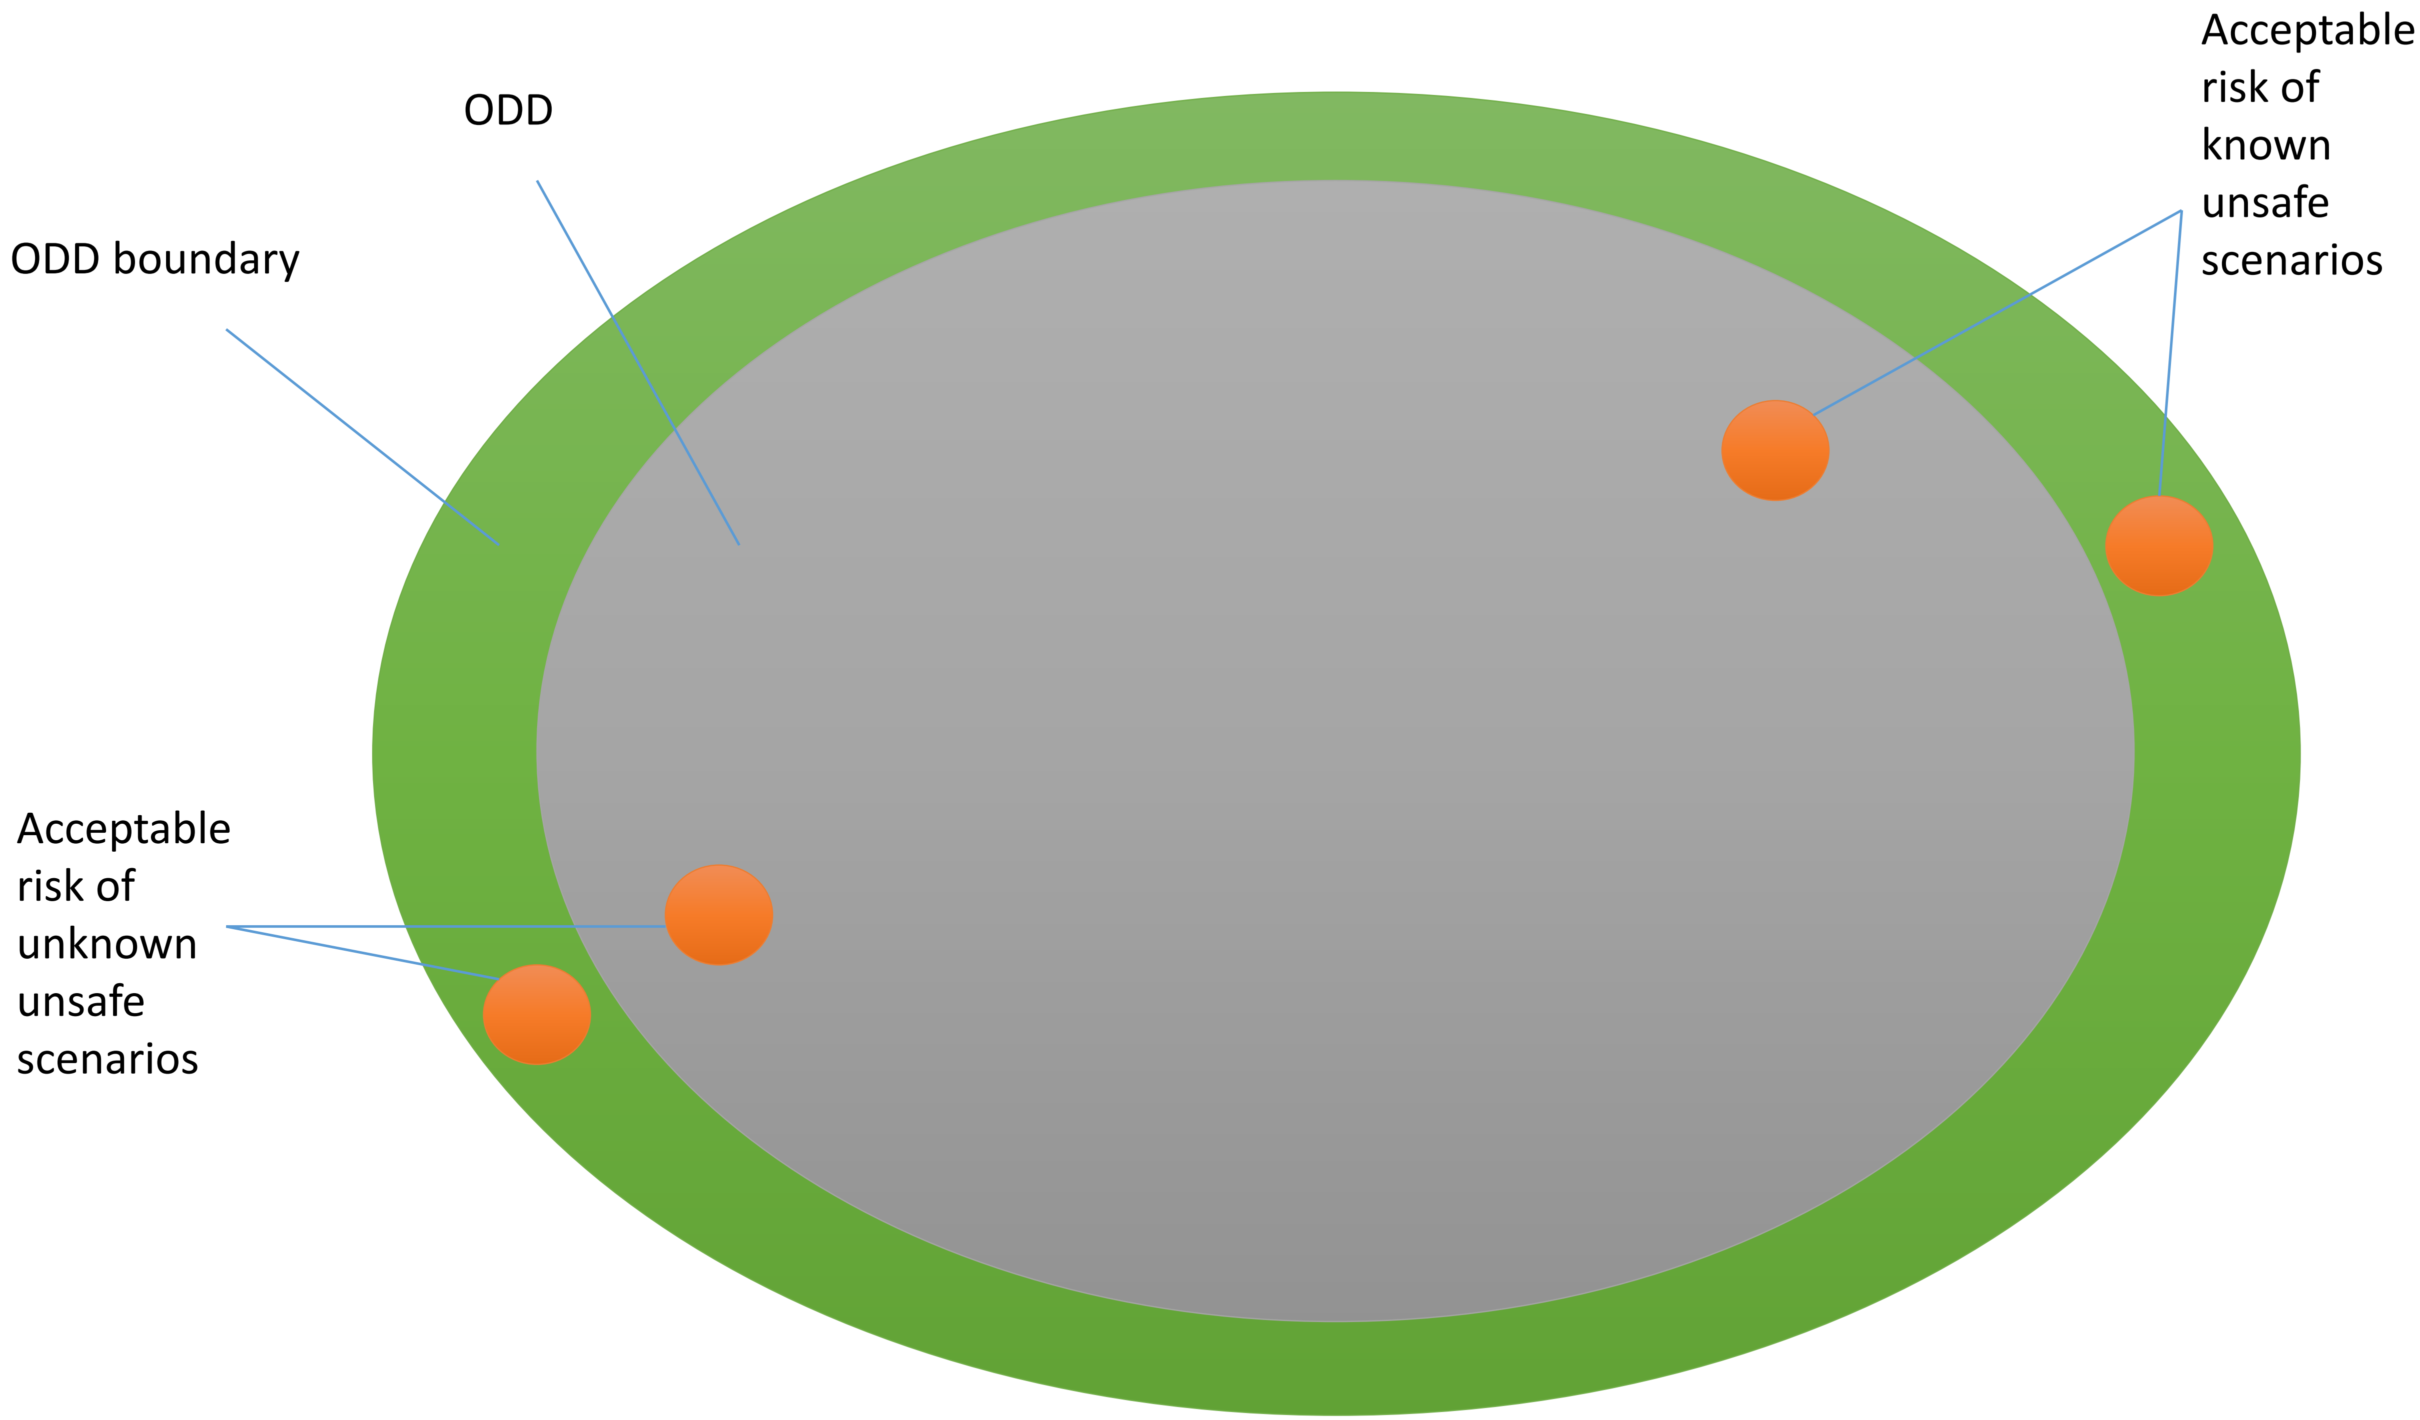

Responsibility-Driven Safety

- Normal driving scenarios
  - Must not cause unacceptable risk increase
  - Low/high demand (incl. other road user errors)

- Emergency scenarios
  - Near-crash
    - Must avoid crash if it can
  - Crash
    - Must mitigate if it can
    - Dilemmas often addressed by blame assignment
  - Fallback
    - Must minimize overall risk

(related: Responsibility-Sensitive Safety, https://arxiv.org/pdf/1708.06374)

Blame vs. Injury Risk

GM Cruise Chevy vs. motorcycle crash
https://www.dmv.ca.gov/portal/wcm/connect/1877d019-d5f0-4c46-b472-78cfe289787d/GMCruise_120717.pdf?MOD=AJPERES

Blame vs Injury Risk (from the Accident Report)

A Cruise autonomous vehicle ("Cruise AV"), operating in autonomous mode in heavy traffic, was involved in a collision while traveling east on Oak Street just past the intersection with Fillmore Street. The Cruise AV was traveling in the center of three one-way lanes. Identifying a space between two vehicles (a minivan in front and a sedan behind) in the left lane, the Cruise AV began to merge into that lane. At the same time, the minivan decelerated. Sensing that its gap was closing, the Cruise AV stopped making its lane change and returned fully to the center lane. As the Cruise AV was re-centering itself in the lane, a motorcycle that that had just lane-split between two vehicles in the center and right lanes moved into the center lane, glanced the side of the Cruise AV, wobbled, and fell over. At the time of the collision, the Cruise AV was traveling with the flow of traffic at 12mph, while the motorcycle was traveling at approximately 17mph. The motorcyclist got up and walked his vehicle to the side of the road, where the parties exchanged information. 911 was called pursuant to Cruise policy. The motorcyclist reported shoulder pain and was taken to receive medical care, and a police report was taken. As reported in Traffic Collision Report#l70989746, the motorcyclist was determined to be at fault for attempting to overtake and pass another vehicle on the right under conditions that did not permit that movement in safety in violation of CVC 21755(a).

High-Level Behavior Safety Requirements (Normal Driving)

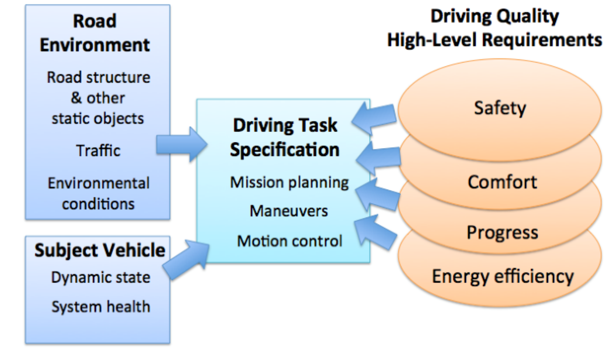

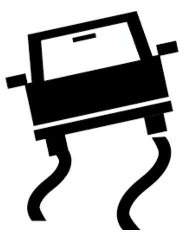

1. Vehicle stability

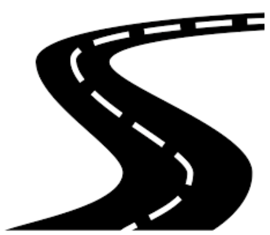

2. Assured clear distance ahead

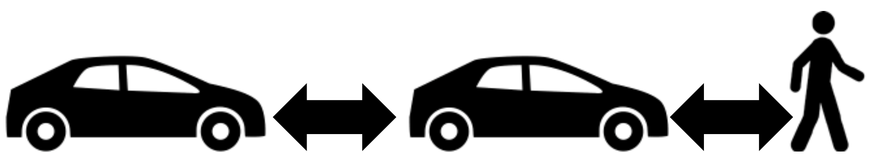

3. Minimum separation

4. Traffic regulations

5. Informal traffic rules (best practices)

Behavioral Safety: 1. Vehicle Stability

Skid stability

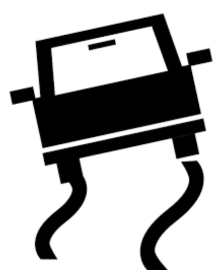

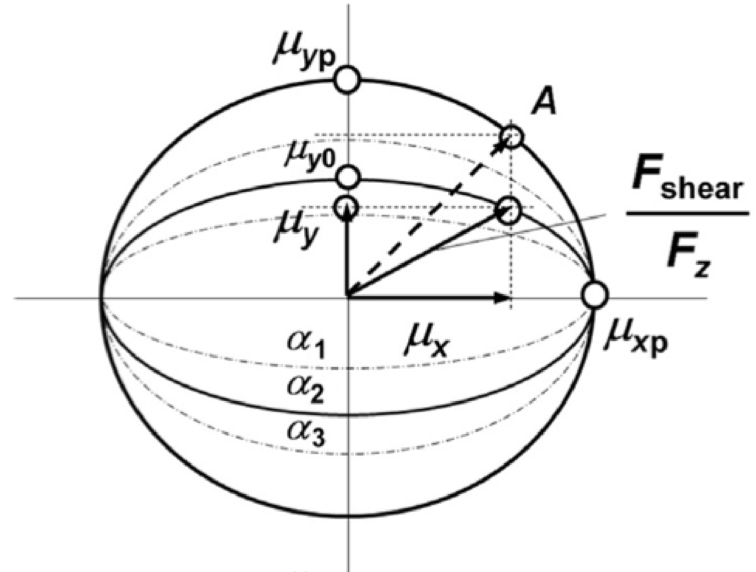

Friction ellipses

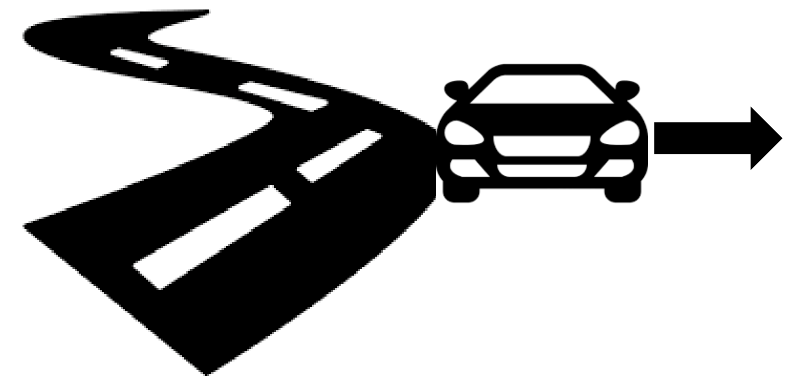

$e + \mu_y = v^2 / 127R$

Roll stability

Untripped

Tripped

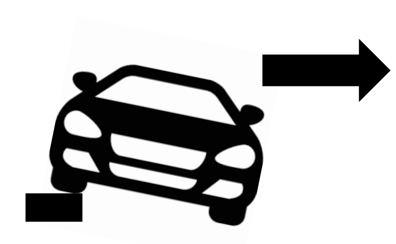

Behavioral Safety: 2. Assured Clear Distance Ahead (ACDA)

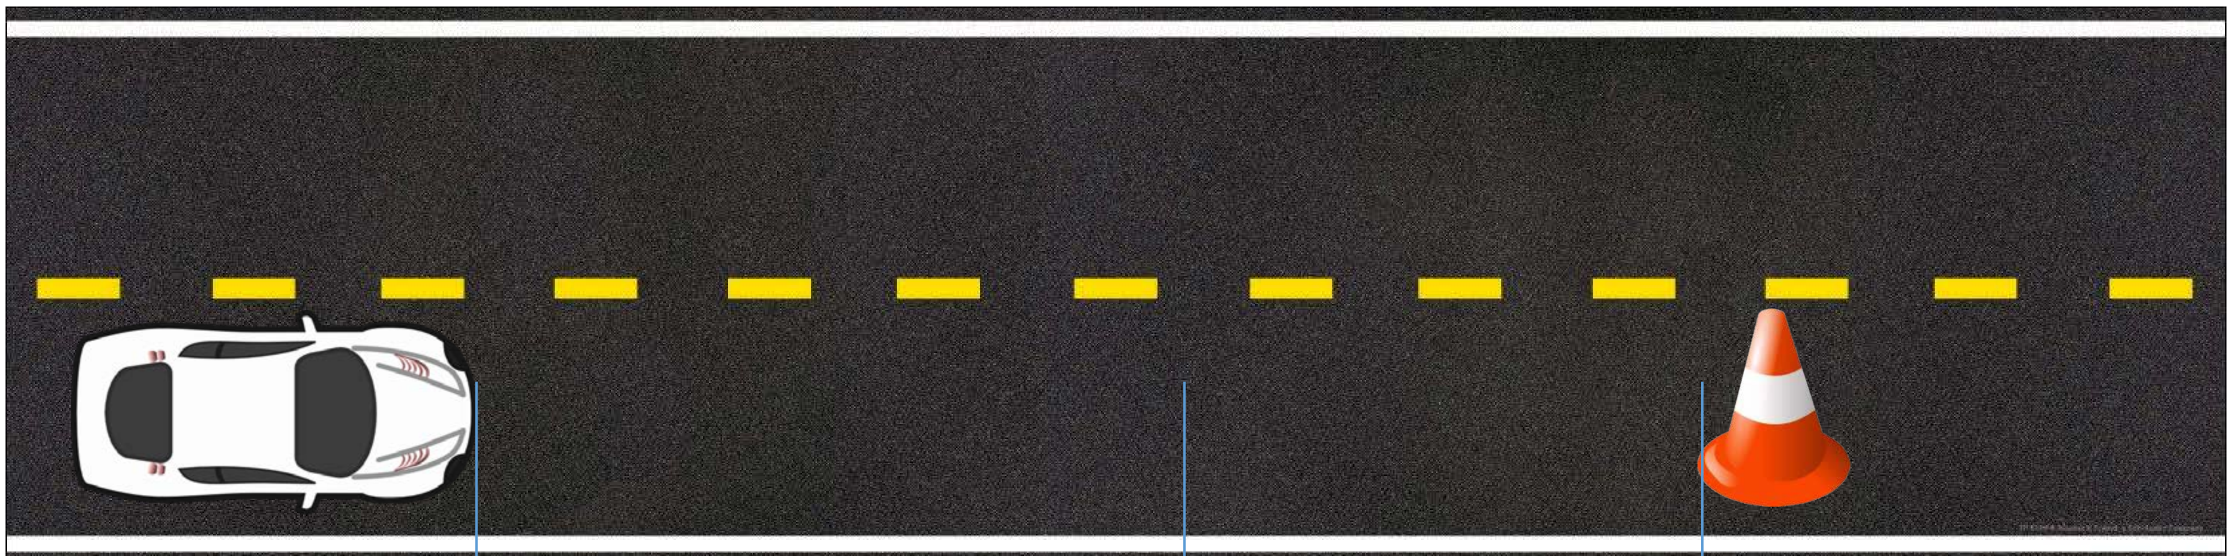

Stopping sight distance
(Perception-reaction time and braking distance)

Perception distance
(Range + road geometry)

Limits safe speed

Behavioral Safety: 2. ACDA Perception Distance

Crests

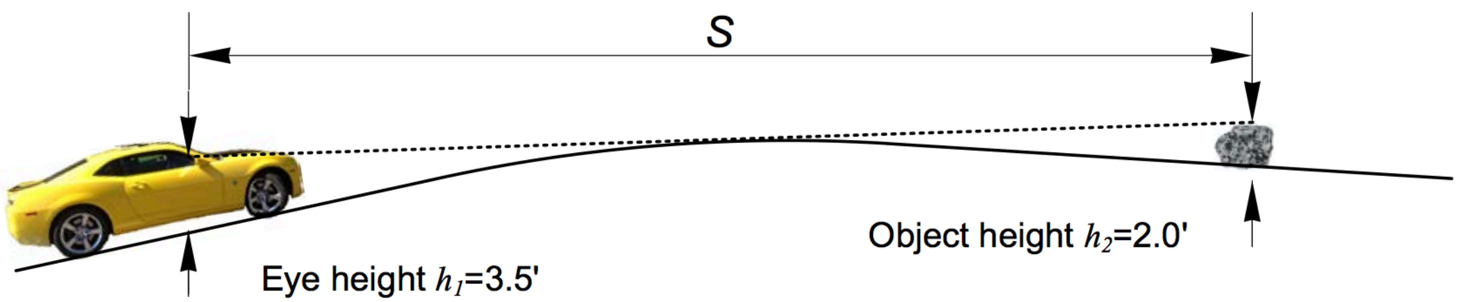

Curves

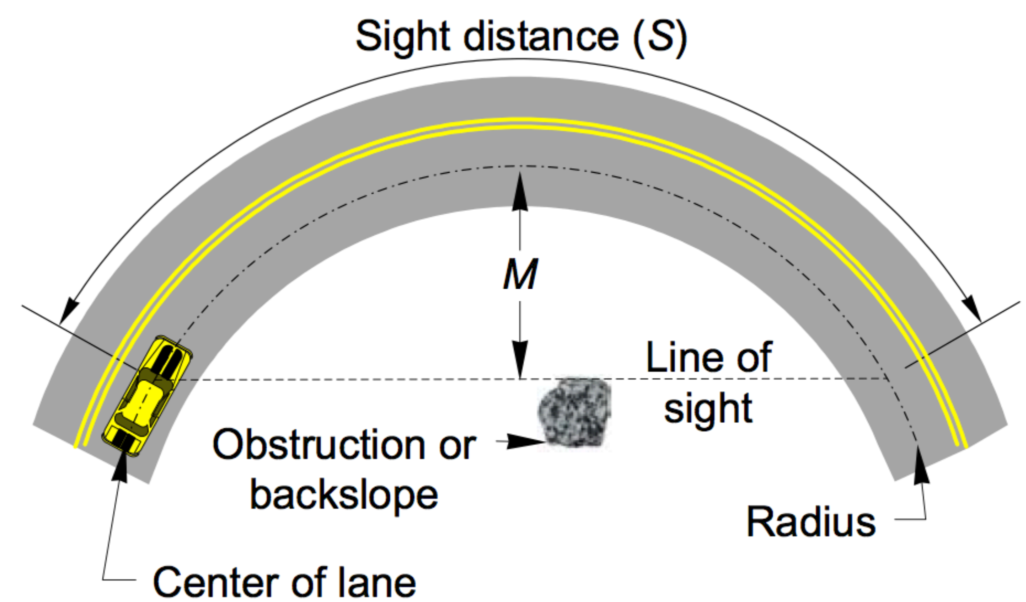

Intersections

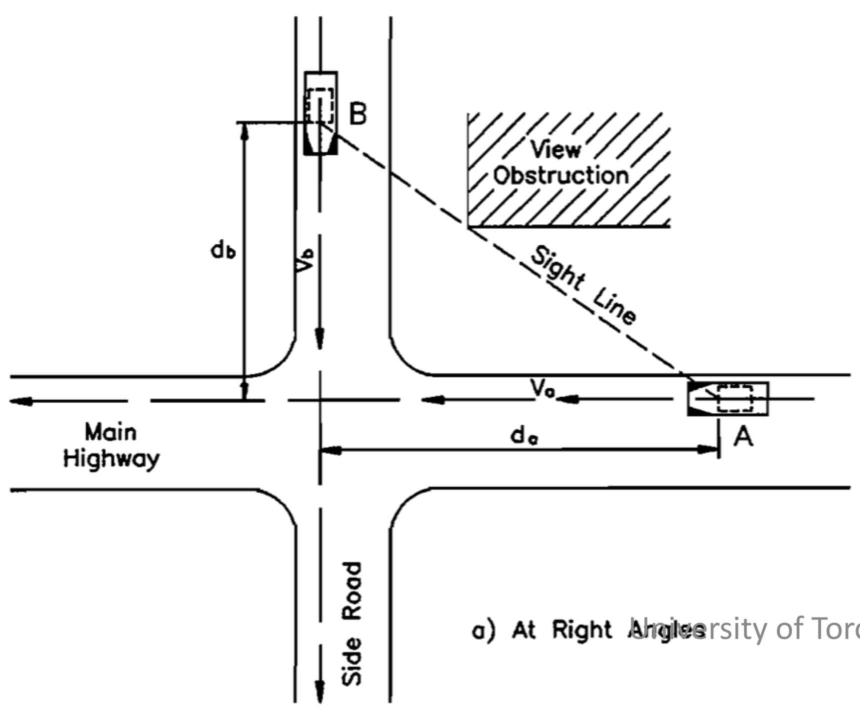

Overtaking

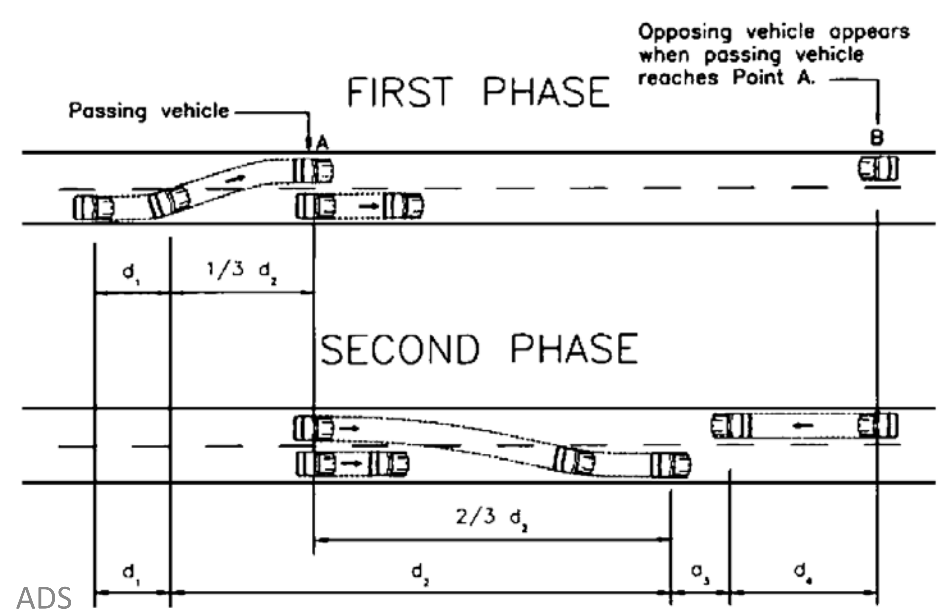

Behavioral Safety:
3. Minimum Separation

Separation in terms of **distance gap**, **time gap**, and **time-to-collision**

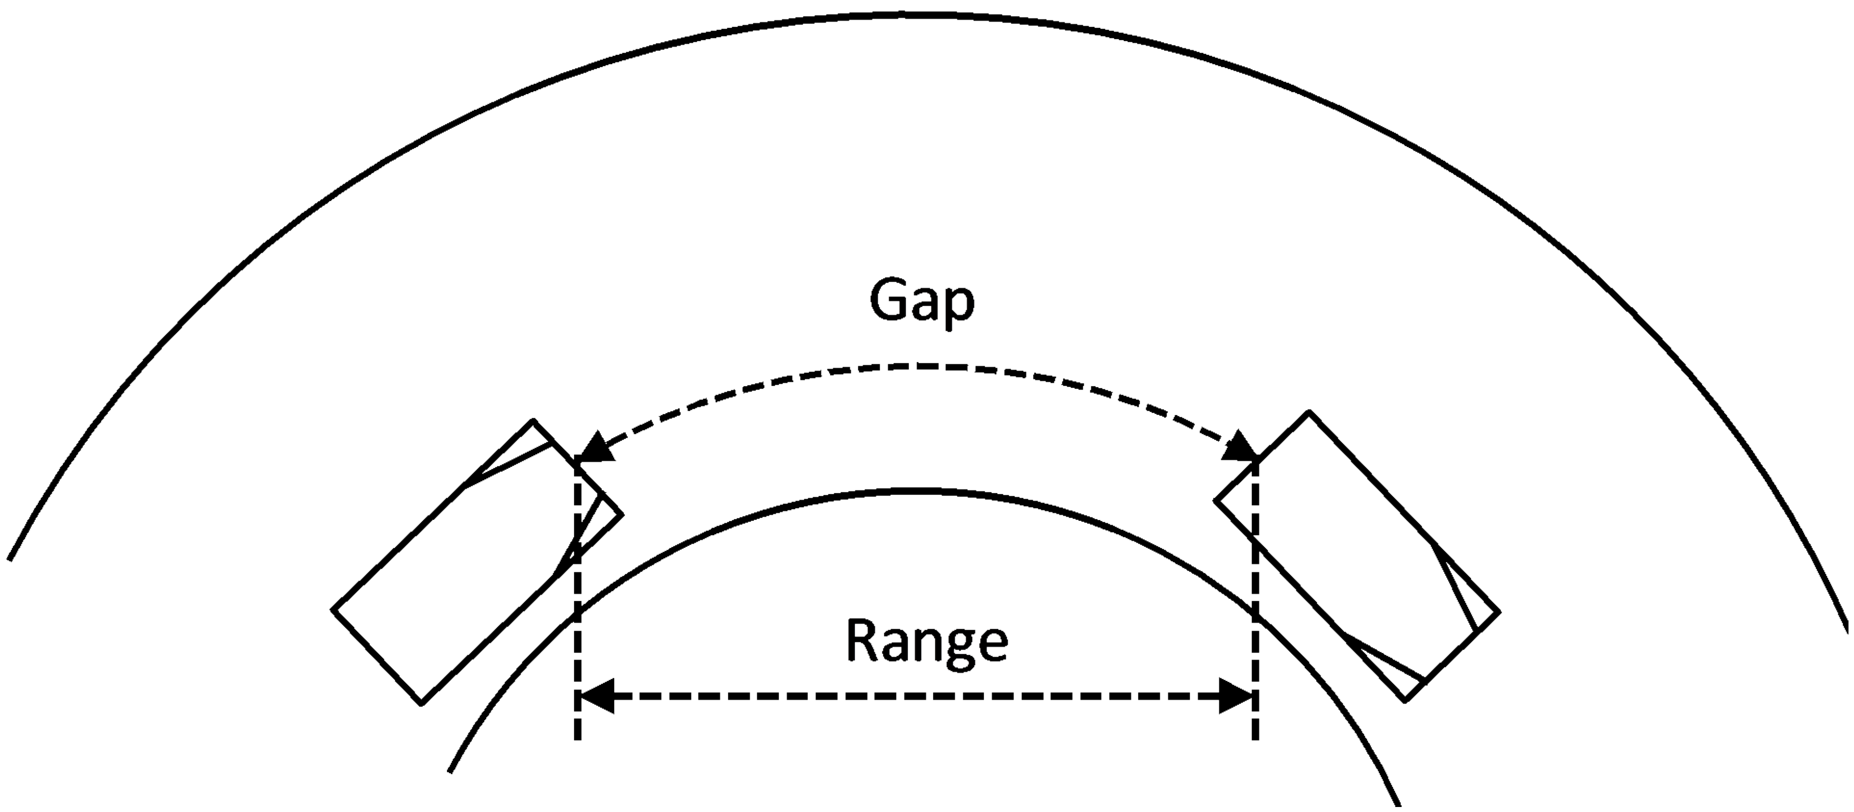

… and various maneuver-specific gaps, including following, overtaking, turning

Behavioral Safety:
4. Traffic Regulations

Safe speed (ACDA)

Yielding to other road users rules

Obeying regulatory traffic signs & signals

Where to drive

Reacting to emergency vehicles & school buses

U-turn prohibitions

Safe following gap

Passing rules

Signaling stops & turns

Parking restrictions

Use of passing beam

Required behavior at railway crossings

…

Behavioral Safety:
5. Informal Traffic Rules

2/3 – second rule

Responding to tailgaters

How early to signal turns

Delayed acceleration at signalized intersections

Lane selection

Anticipating aberrant behaviors of other road users

Responding to animals on the roadway

…

WISE Drive Documentation

WISE Drive comes with comprehensive documentation (over 350 pages) available from this page.

All eight documents in two zip archives: zip1, zip2

Driving Task Specification

Maneuver Catalog

K. Czarnecki. Automated Driving System (ADS) Task Analysis – Part 2: Structured Road Maneuvers. Waterloo Intelligent Systems Engineering Lab (WISE) Report, University of Waterloo, 2018, DOI: 10.13140/RG.2.2.23280.76800

Basic Motion Control Task Catalog

K. Czarnecki. Automated Driving System (ADS) Task Analysis – Part 1: Basic Motion Control Tasks. Waterloo Intelligent Systems Engineering Lab (WISE) Report, University of Waterloo, 2018, DOI: 10.13140/RG.2.2.29991.65447

Road Environment Specification

ODD Taxonomy

K. Czarnecki. Operational Design Domain for Automated Driving Systems – Taxonomy of Basic Terms. Waterloo Intelligent Systems Engineering Lab (WISE) Report, University of Waterloo, 2018, DOI: 10.13140/RG.2.2.18037.88803

ADS Hazard Sources

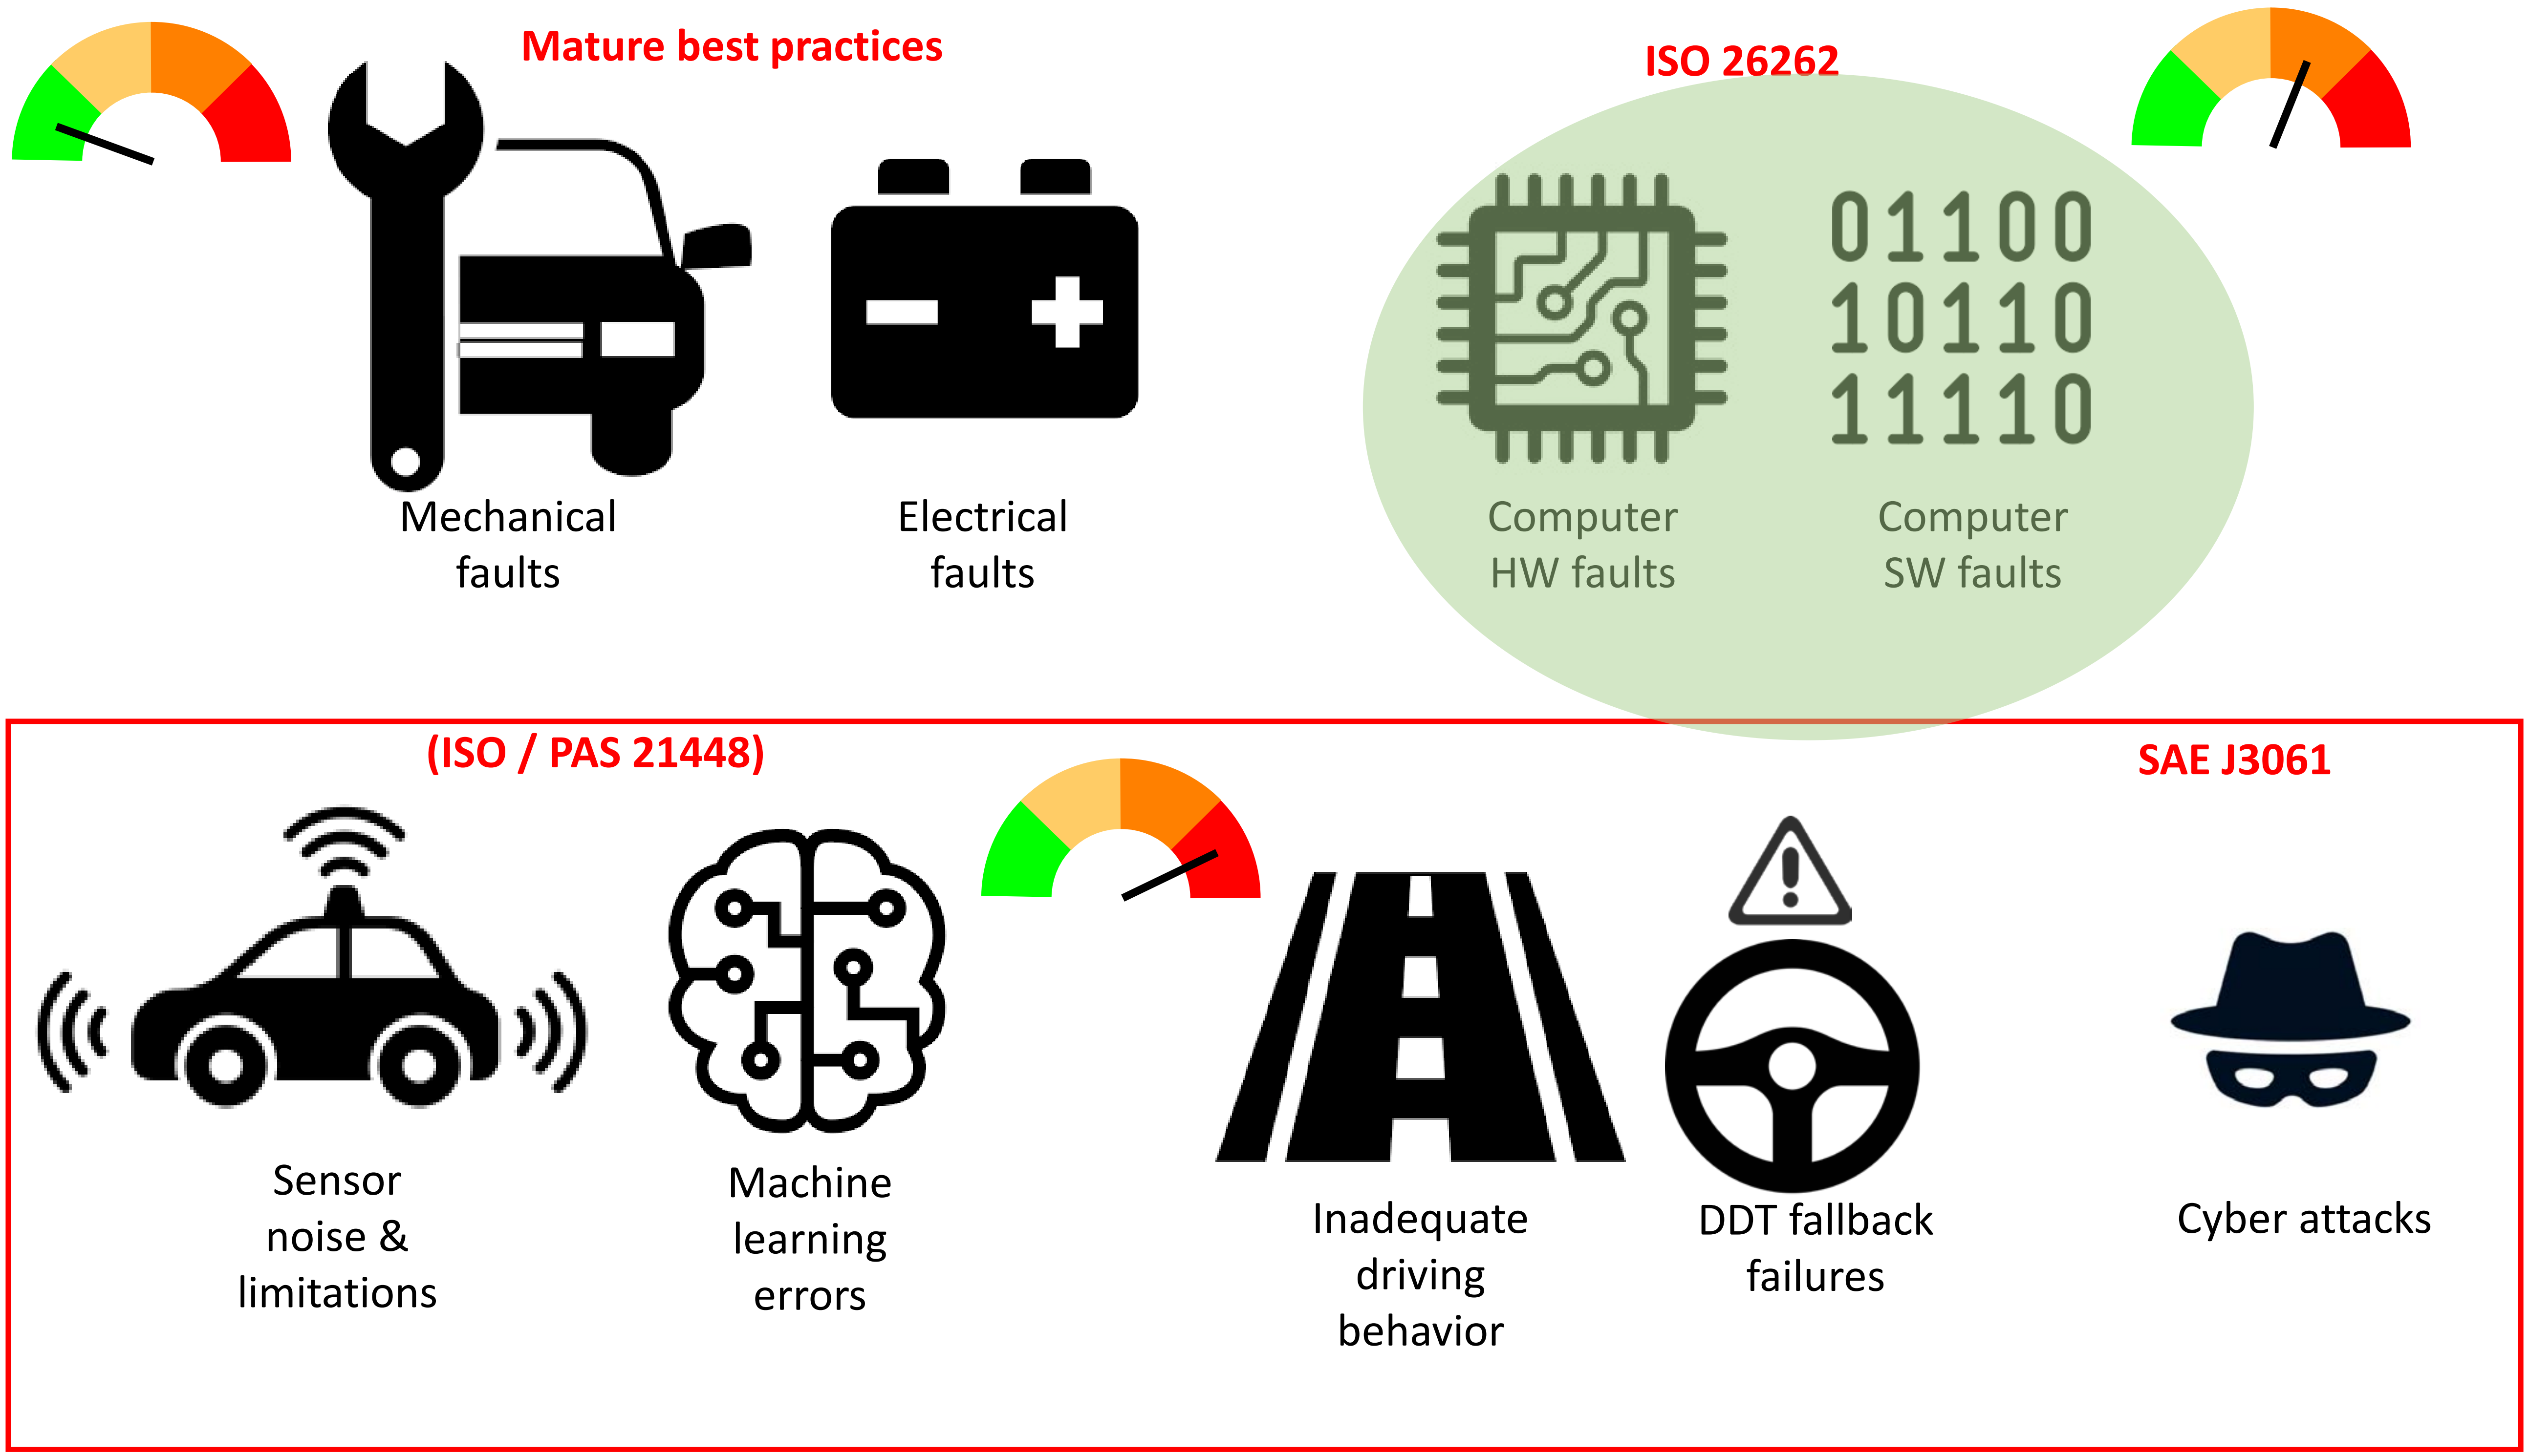

Fail-Operational ADS Architecture

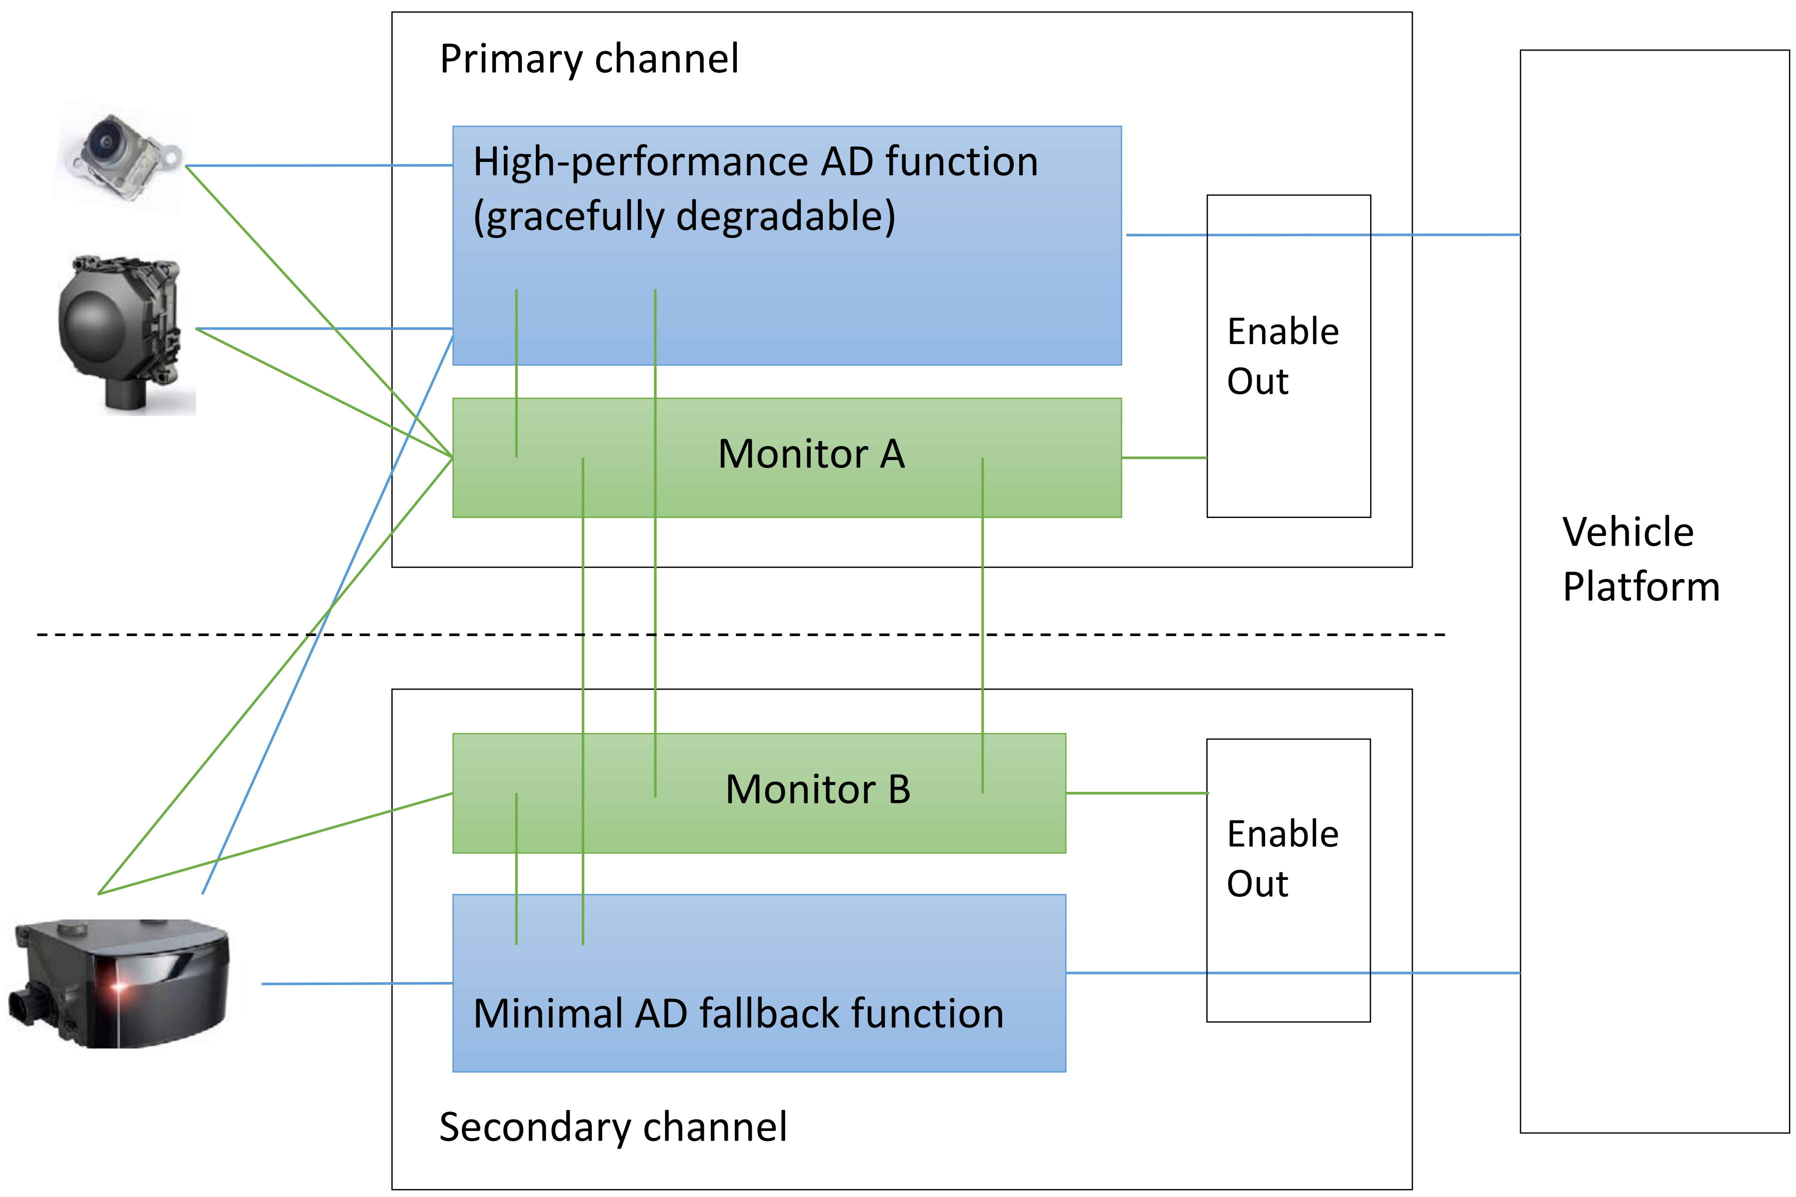

Dependability patterns:

- Redundancy
- Diversity
- Simplex
- Graceful degradation
- Monitoring of monitoring
- Minimized cost

No single-point failures

ODD vs. ROD

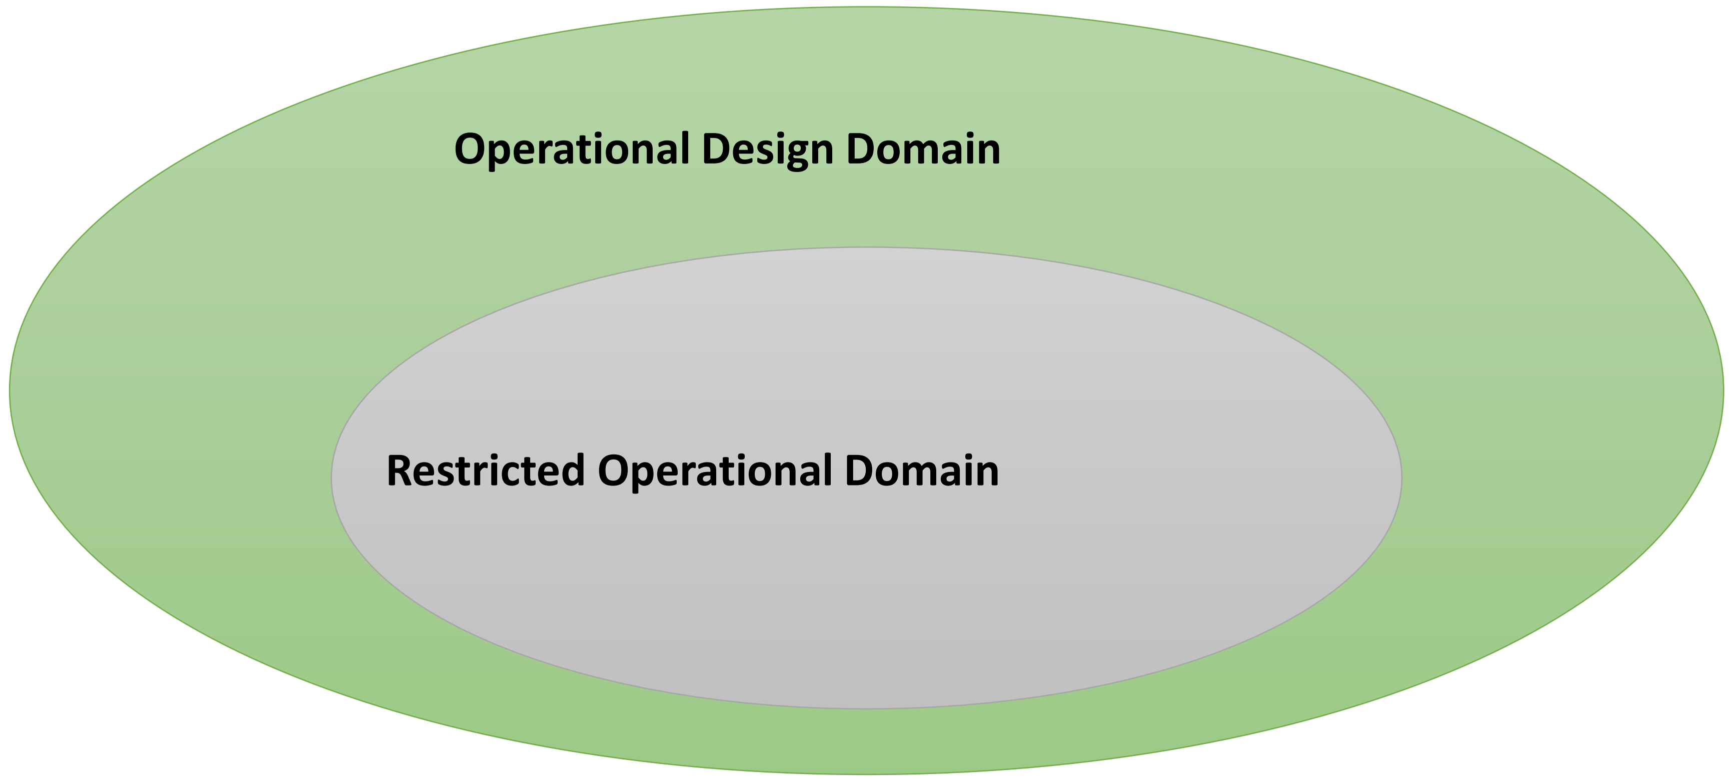

I Colwell, B Phan, S Saleem, R Salay, K Czarnecki. An Automated Vehicle Safety Concept Based on Runtime Restriction of the Operational Design Domain. IEEE Intelligent Vehicles Symposium (IV), 2018

ROD Monitoring for Graceful Degradation

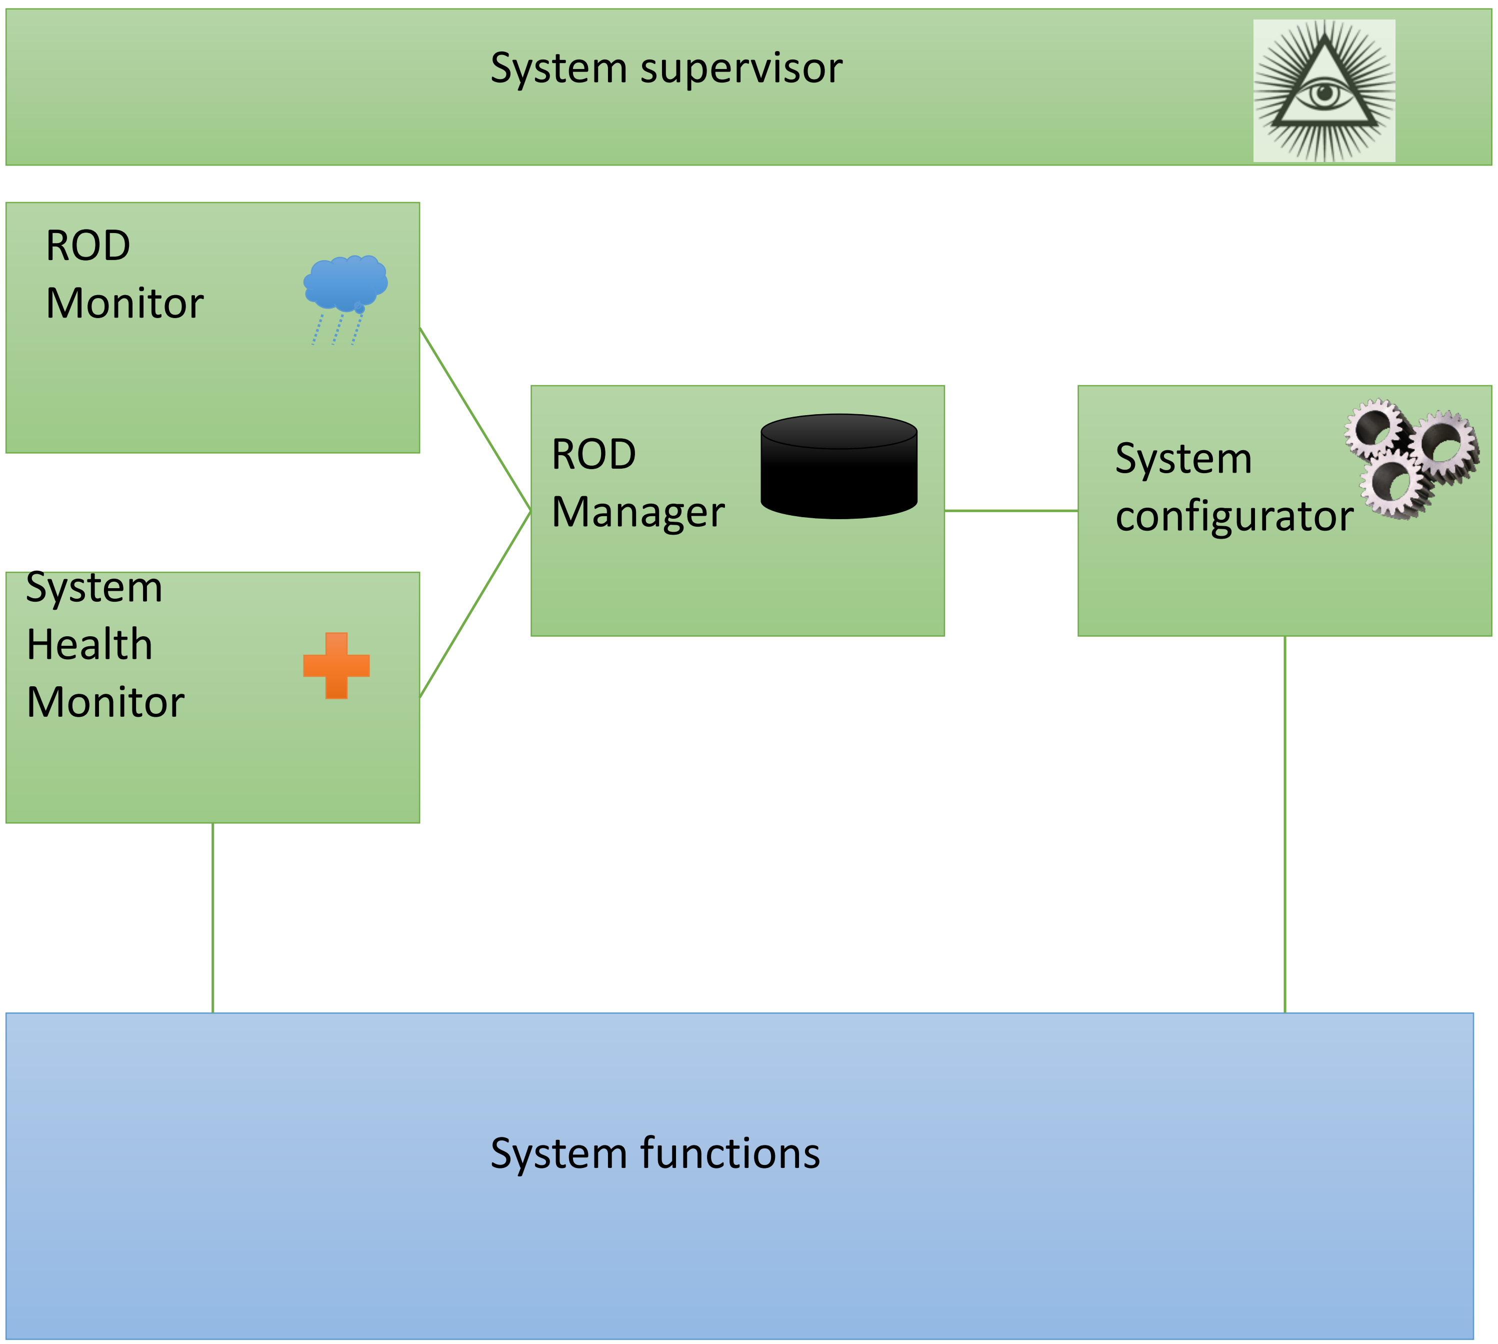

ADS Hazard Sources

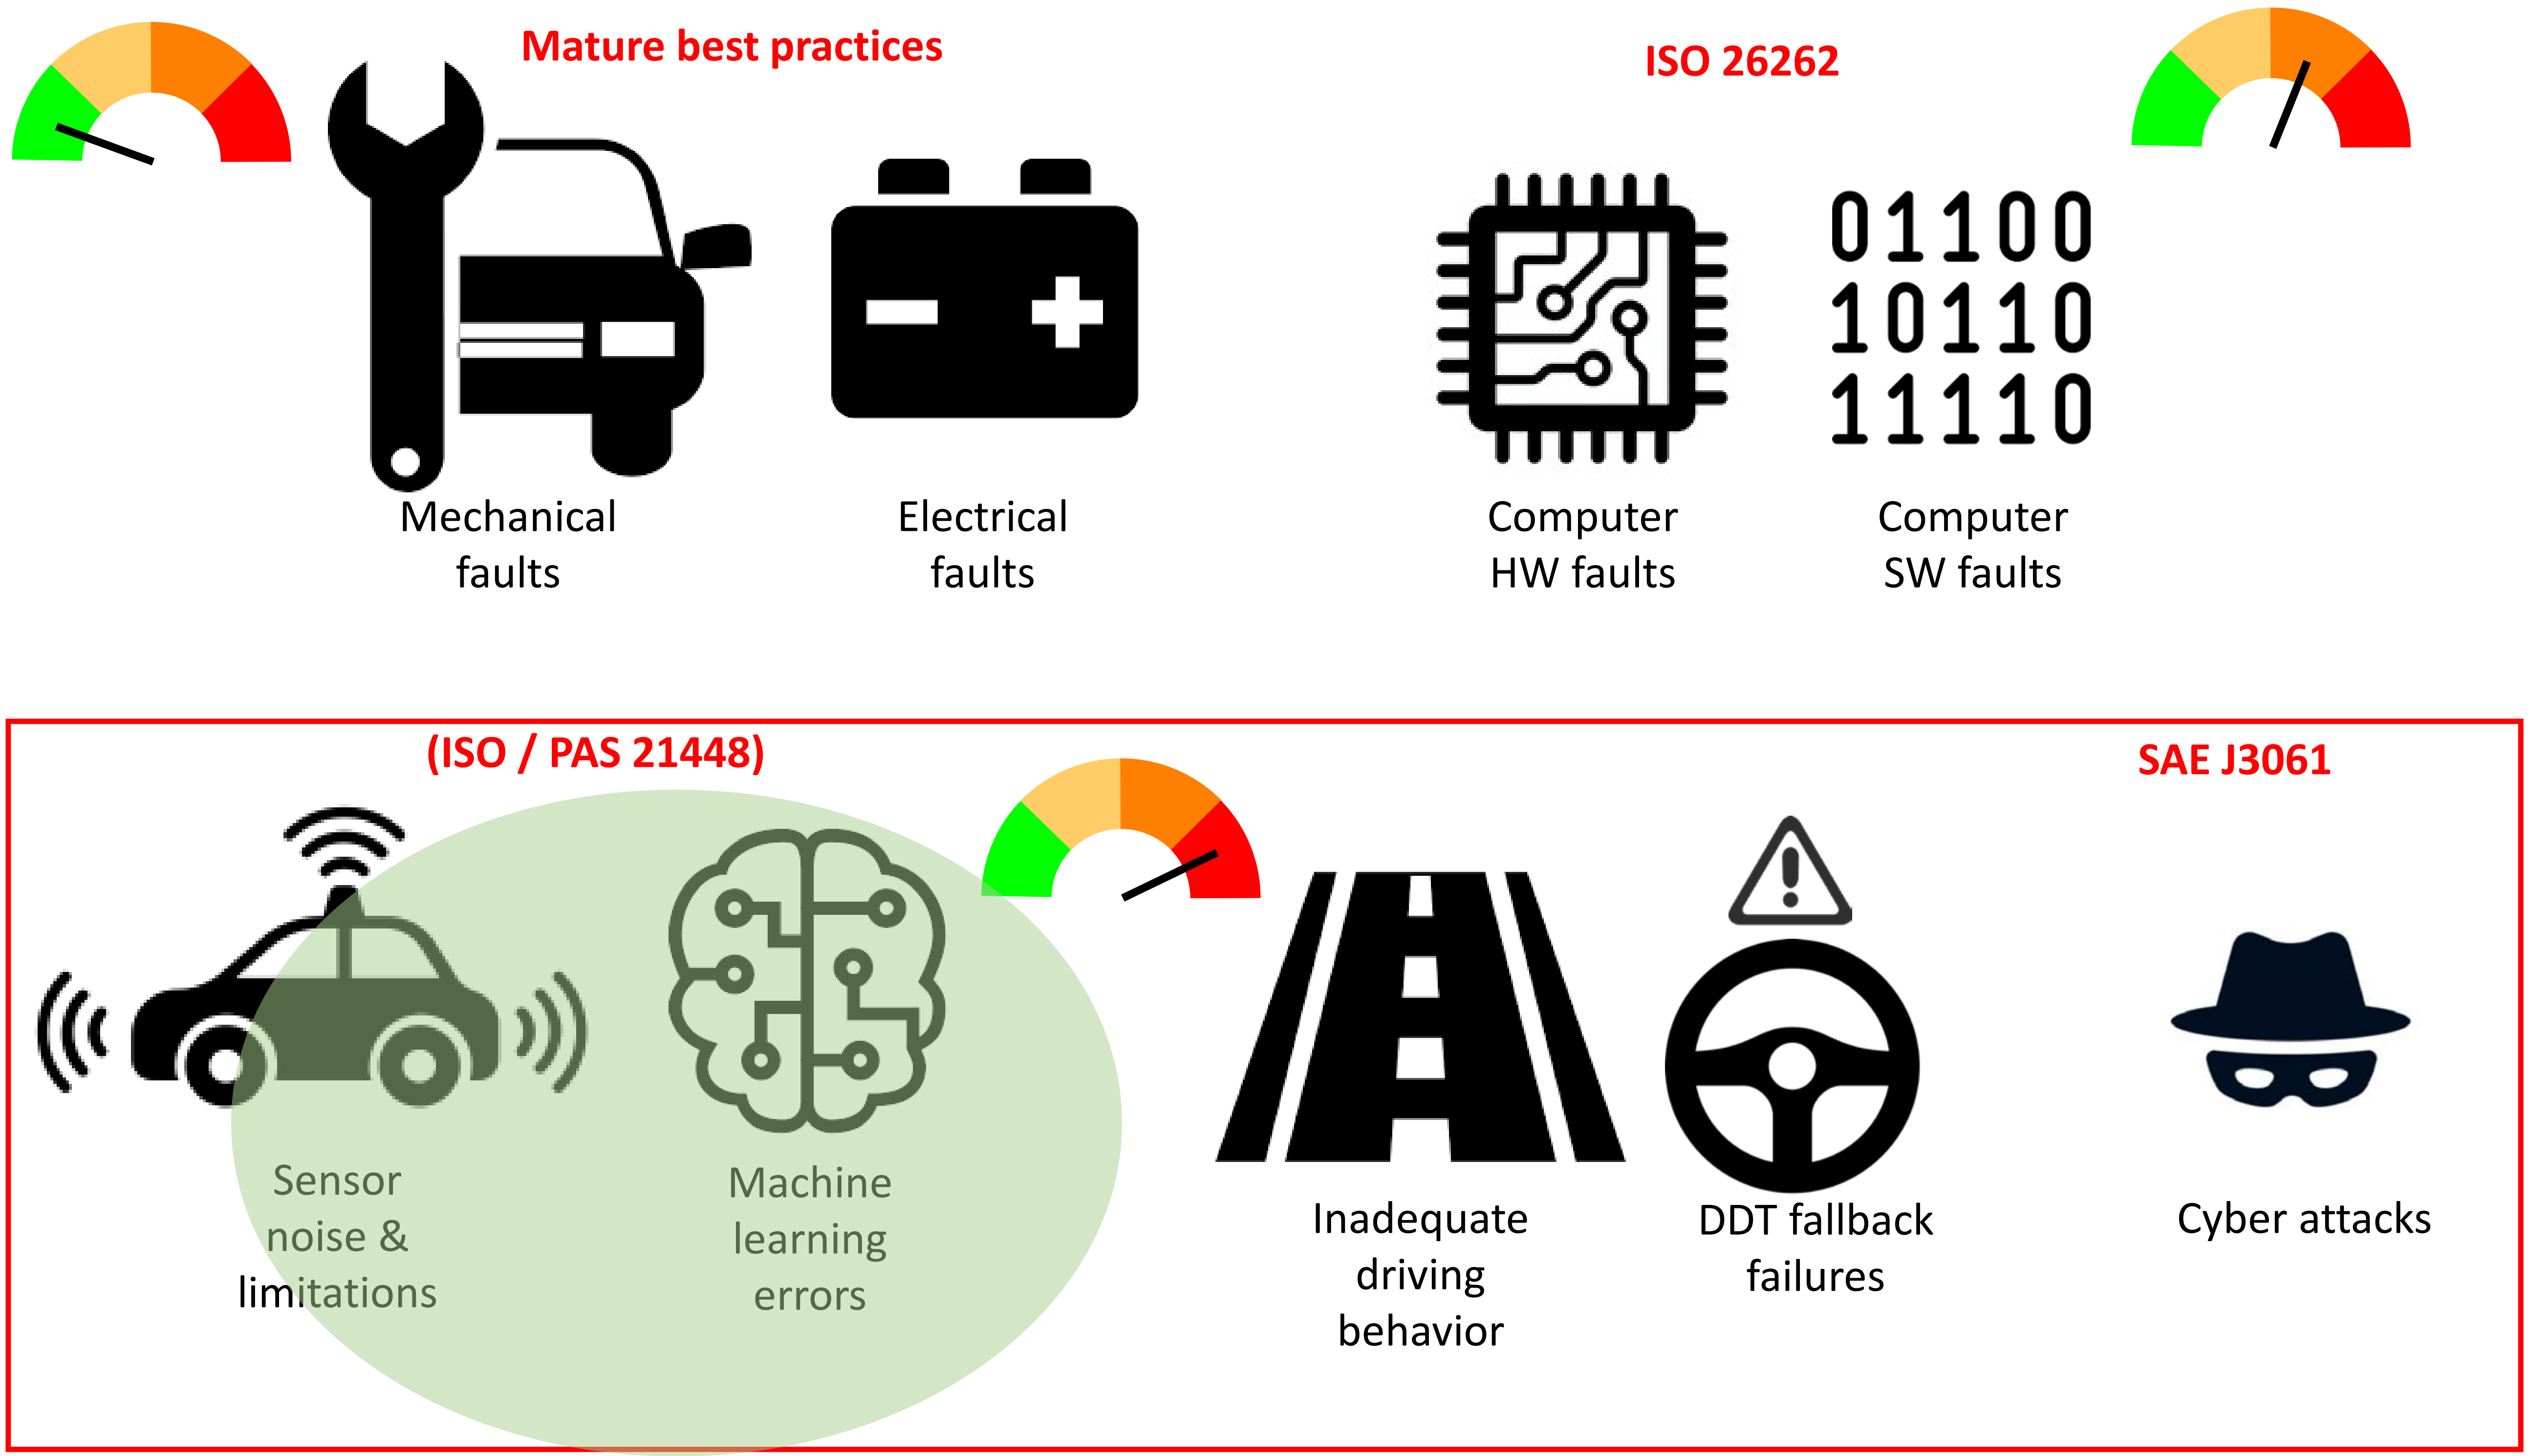

Challenges of Assuring Machine Learned Components

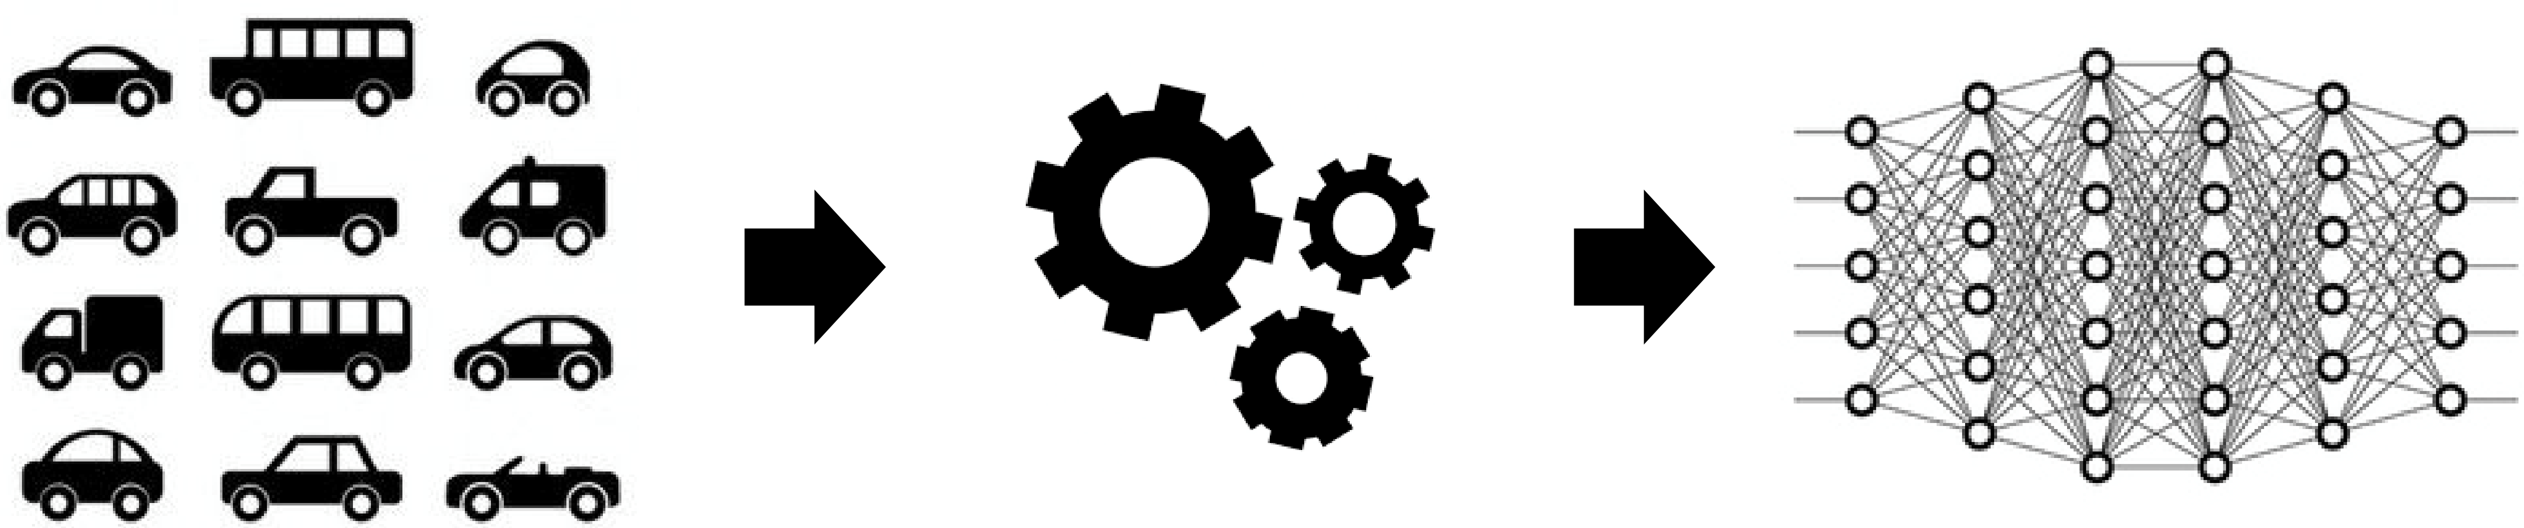

Lack of specification

Lack of inspectability

R. Salay, R. Queiroz, K. Czarnecki. An Analysis of ISO 26262: Machine Learning and Safety in Automotive Software. SAE, 2018-01-1075, 2018; preliminary version also available at https://arxiv.org/abs/1709.02435

Lack of Complete Spec Affects Verification and Testing (see Lecture 4 by R. Salay)

Best practices

- Spec notations
- Design guidelines
- Coding guidelines

Fault tolerance

- Error detection & handling

Verification methods

- Walkthroughs
- Inspections
- Formal verification
- Static code analysis

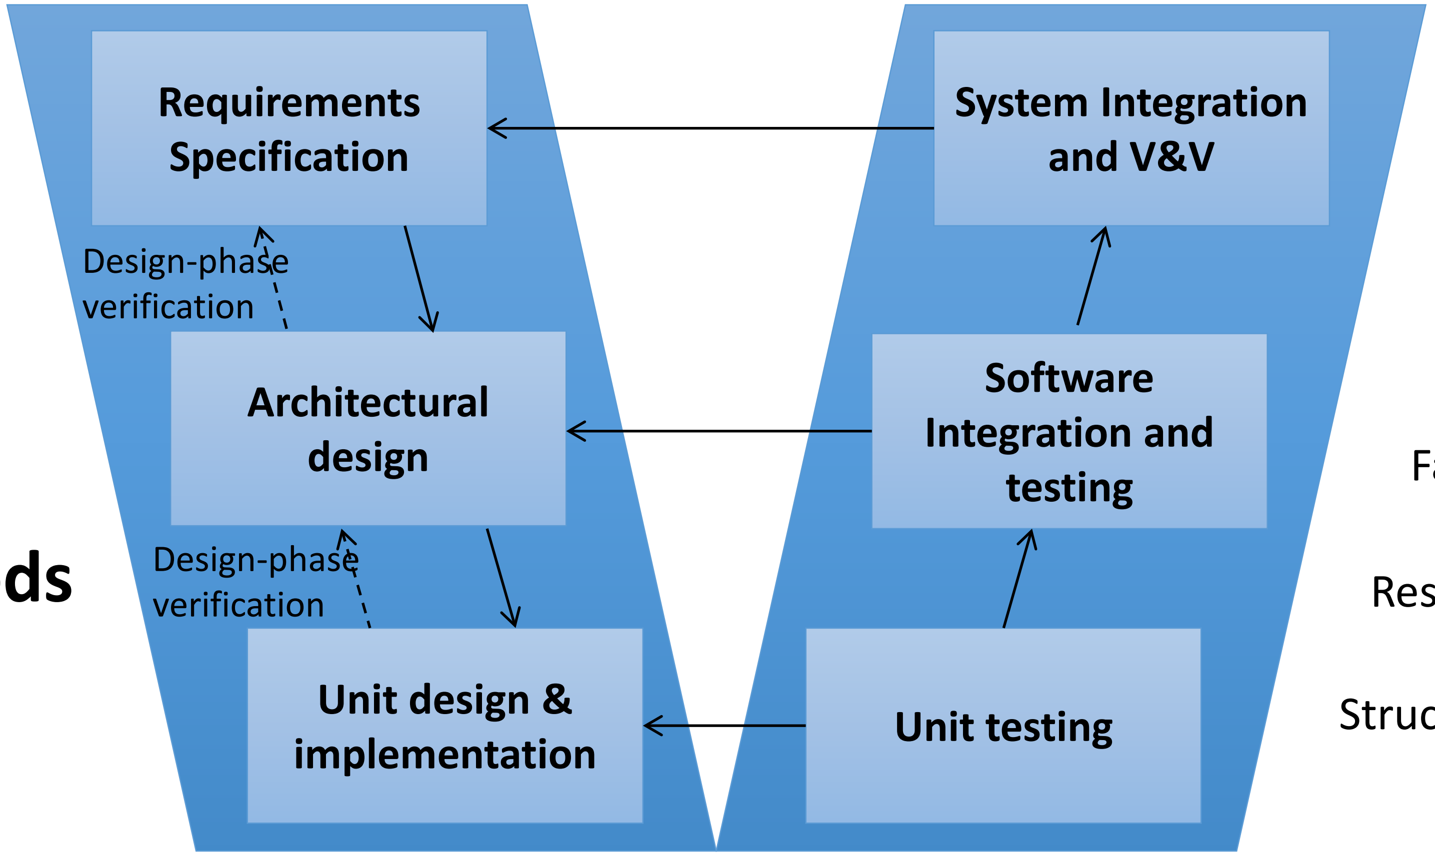

ISO 26262 Part 6

Testing methods

- Requirements-based testing
- Error guessing
- Interface test
- Fault injection test
- Resource usage test
- Structural coverage

Key Recommendations (see Lecture 4)

- Partial specifications
  - Assumptions, necessary/sufficient conditions, in- and eqivariants
  - Runtime monitoring, test generation, regularization
- Data requirements
  - Domain coverage (e.g., ontology)
  - Risk profiling

ADS Challenges

(an unsorted list)

Road User Intension

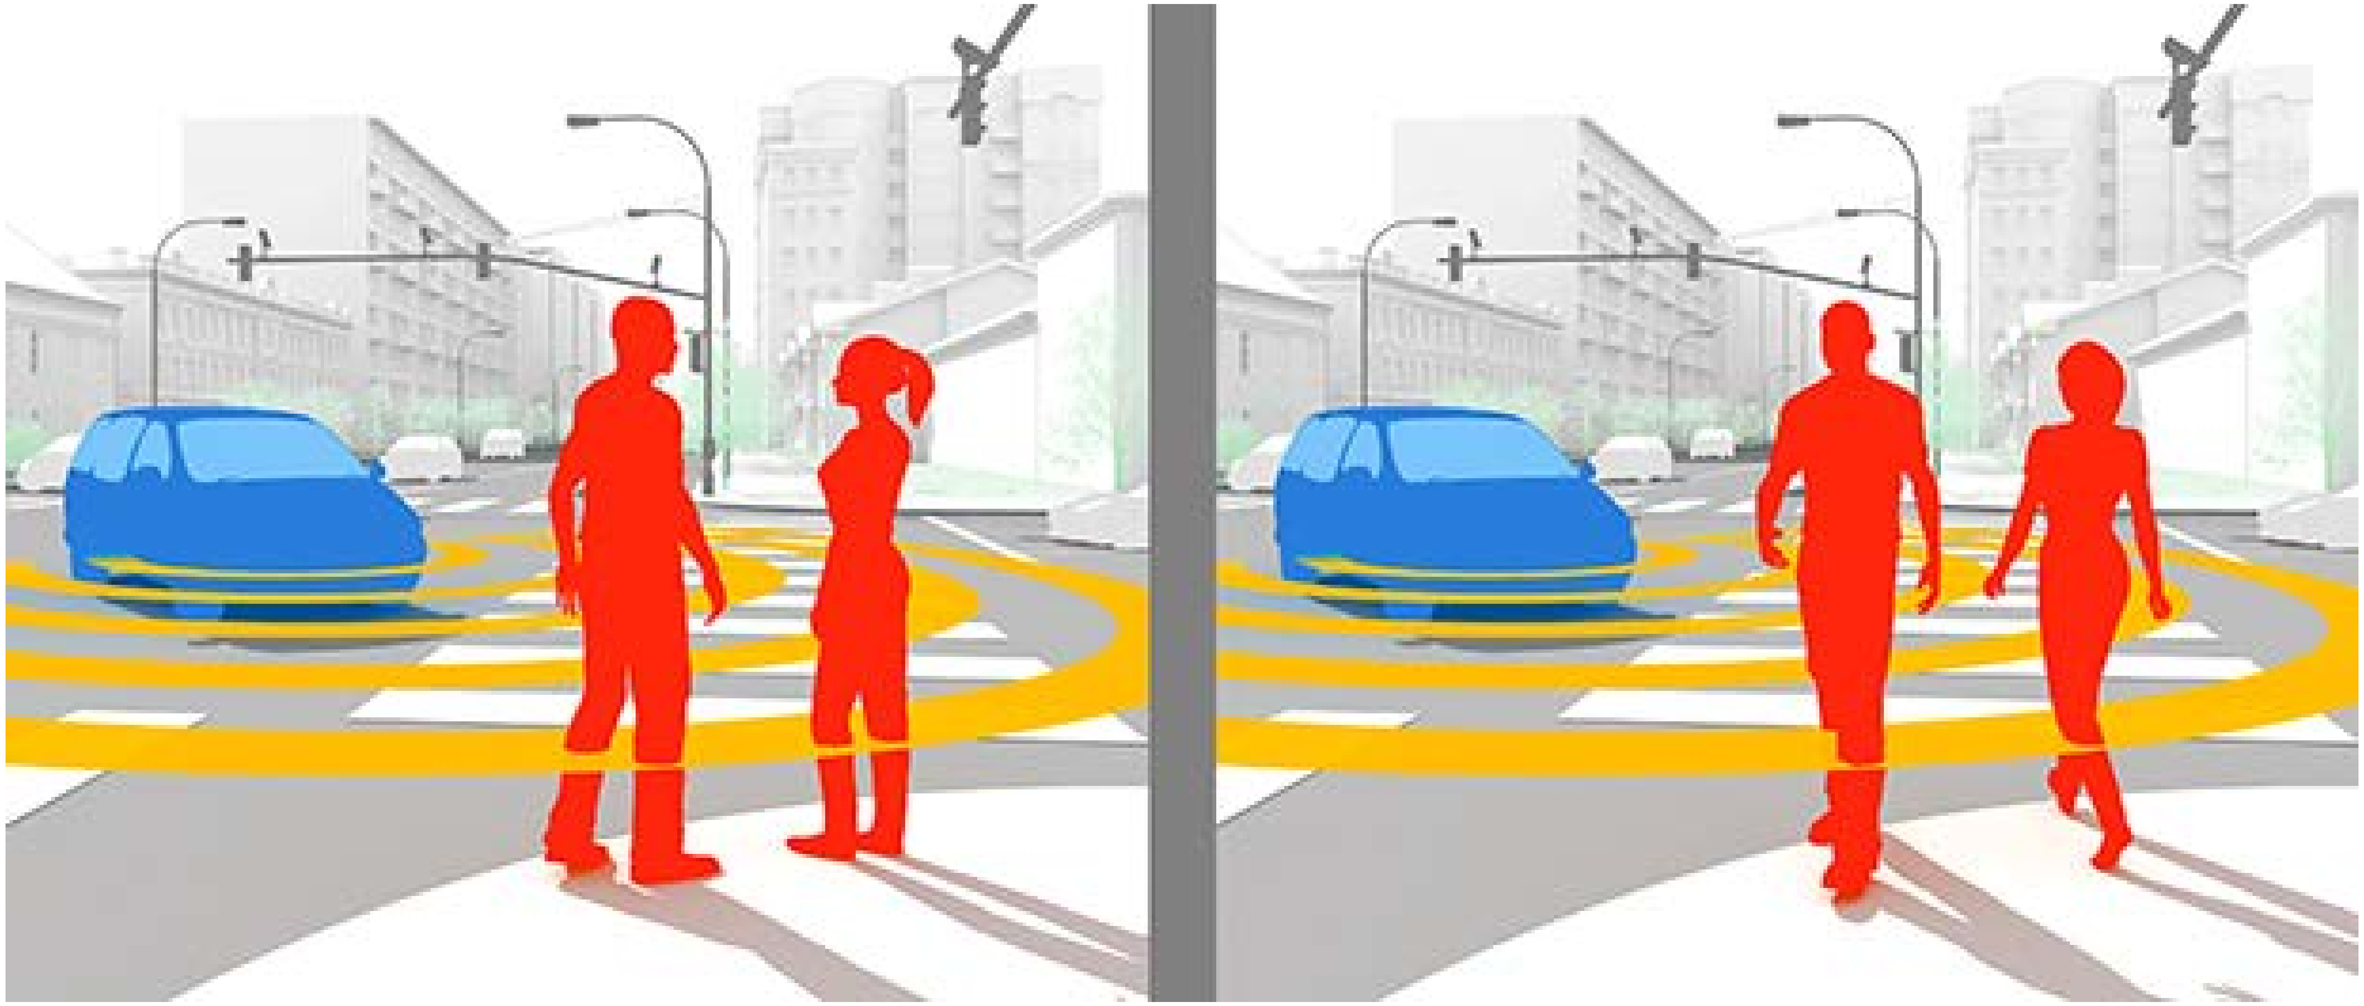

Will she cross the street?

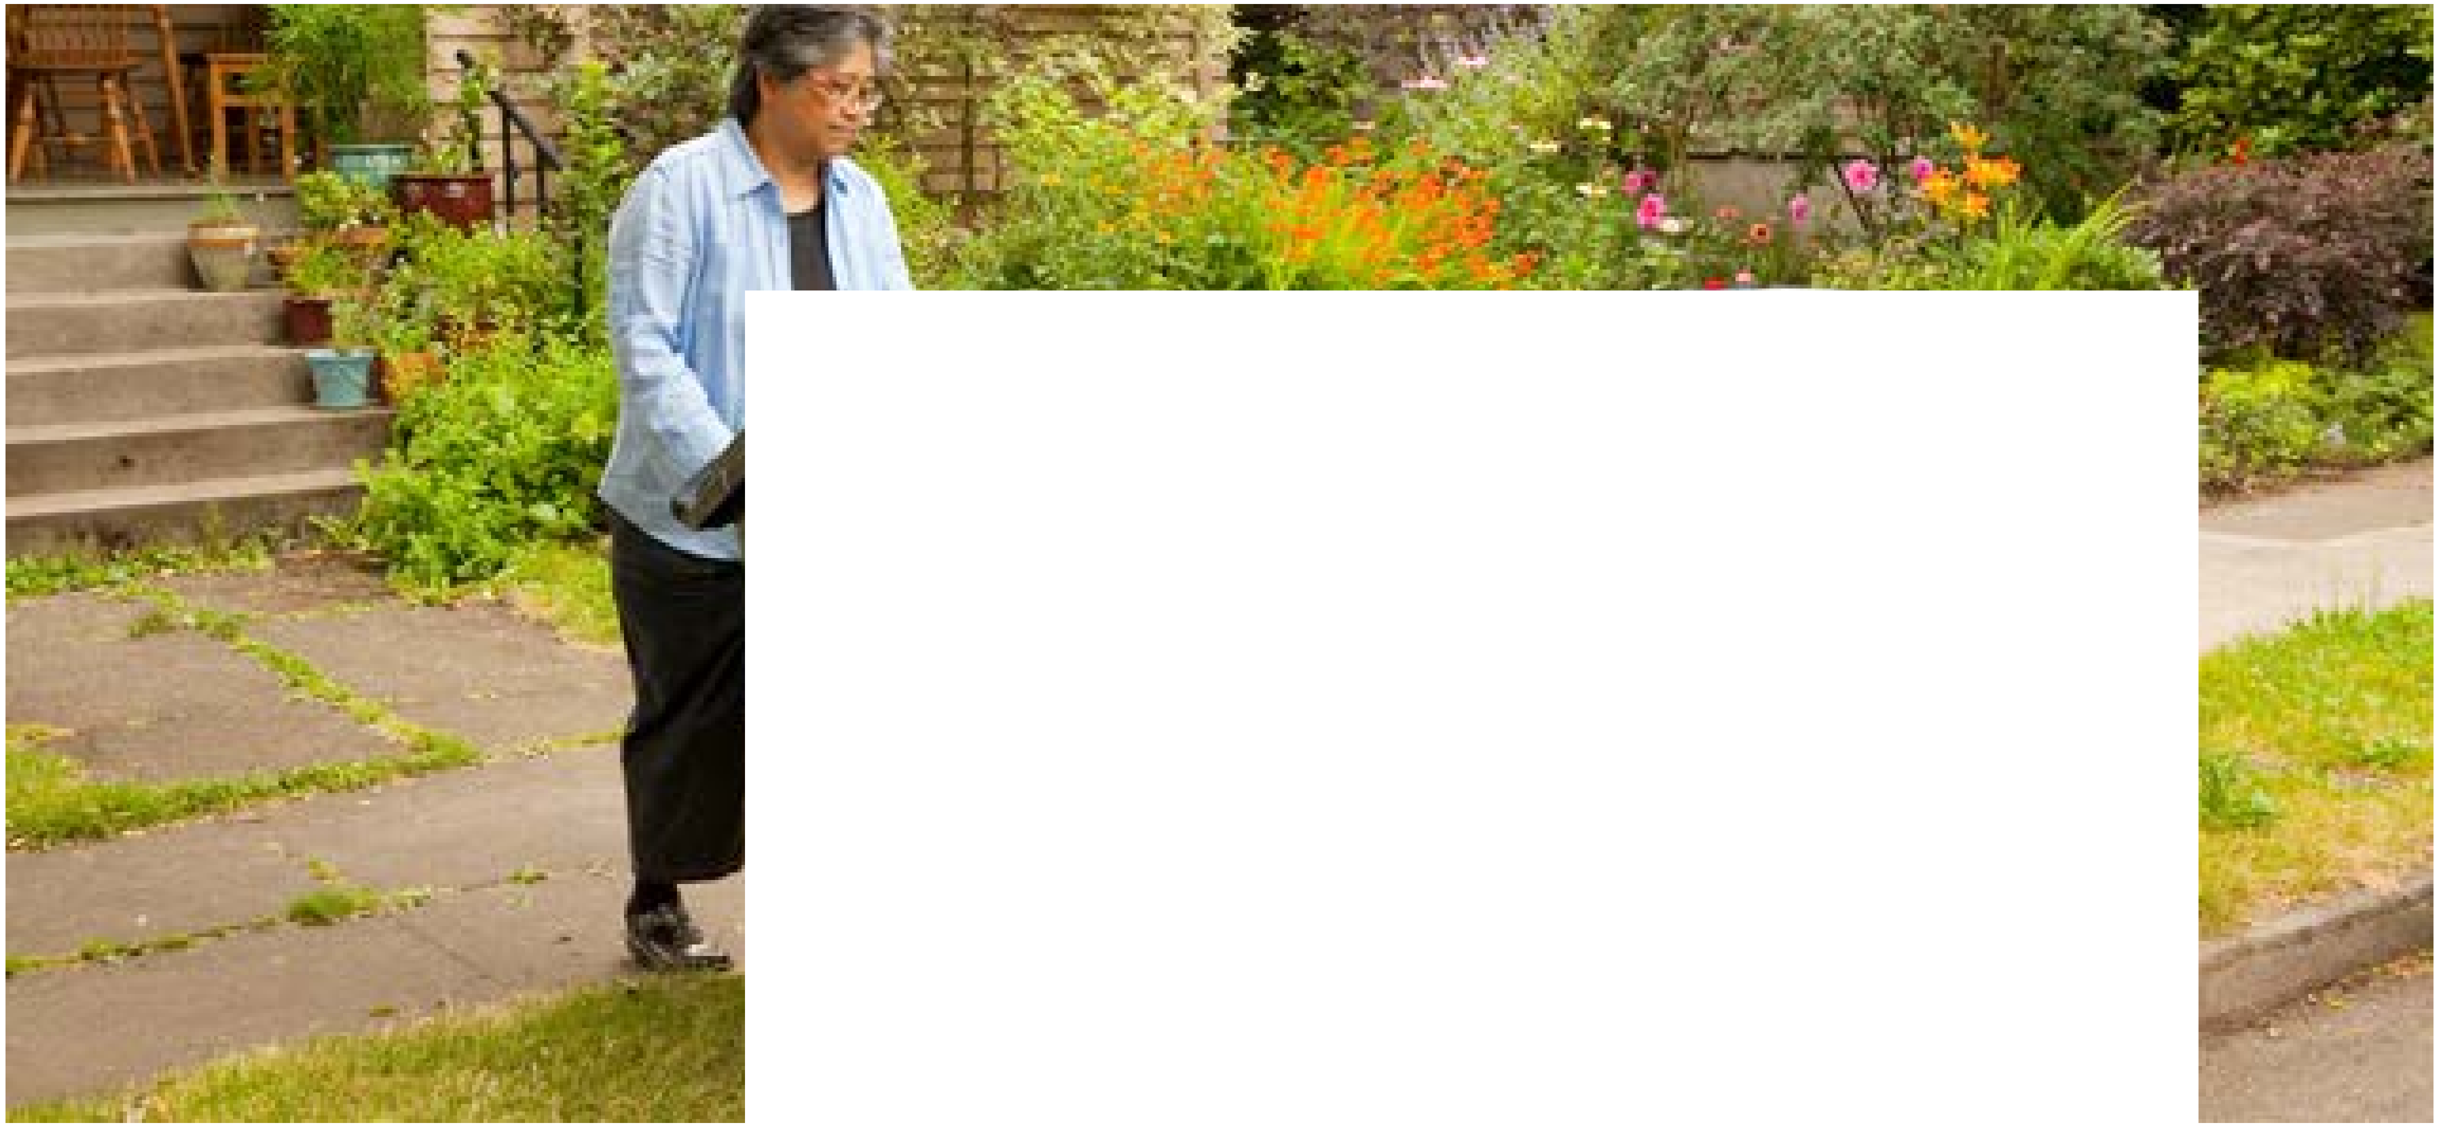

Will she cross the street?

Traffic Lights in Toronto

Bad Weather Driving

“Plastic Bag” Problem

Edge Cases

Driving into a Tornado

Autonomous Trap 101

James Bridle

Crossing Double Yellow Lines

http://iamtraffic.org/engineering/crossing-double-yellow-line/

Place Charles de Gaulle, Paris

Busy City Traffic

Vehicle To Pedestrian Communication

Clamann et al. 2016

Daimler Prototype

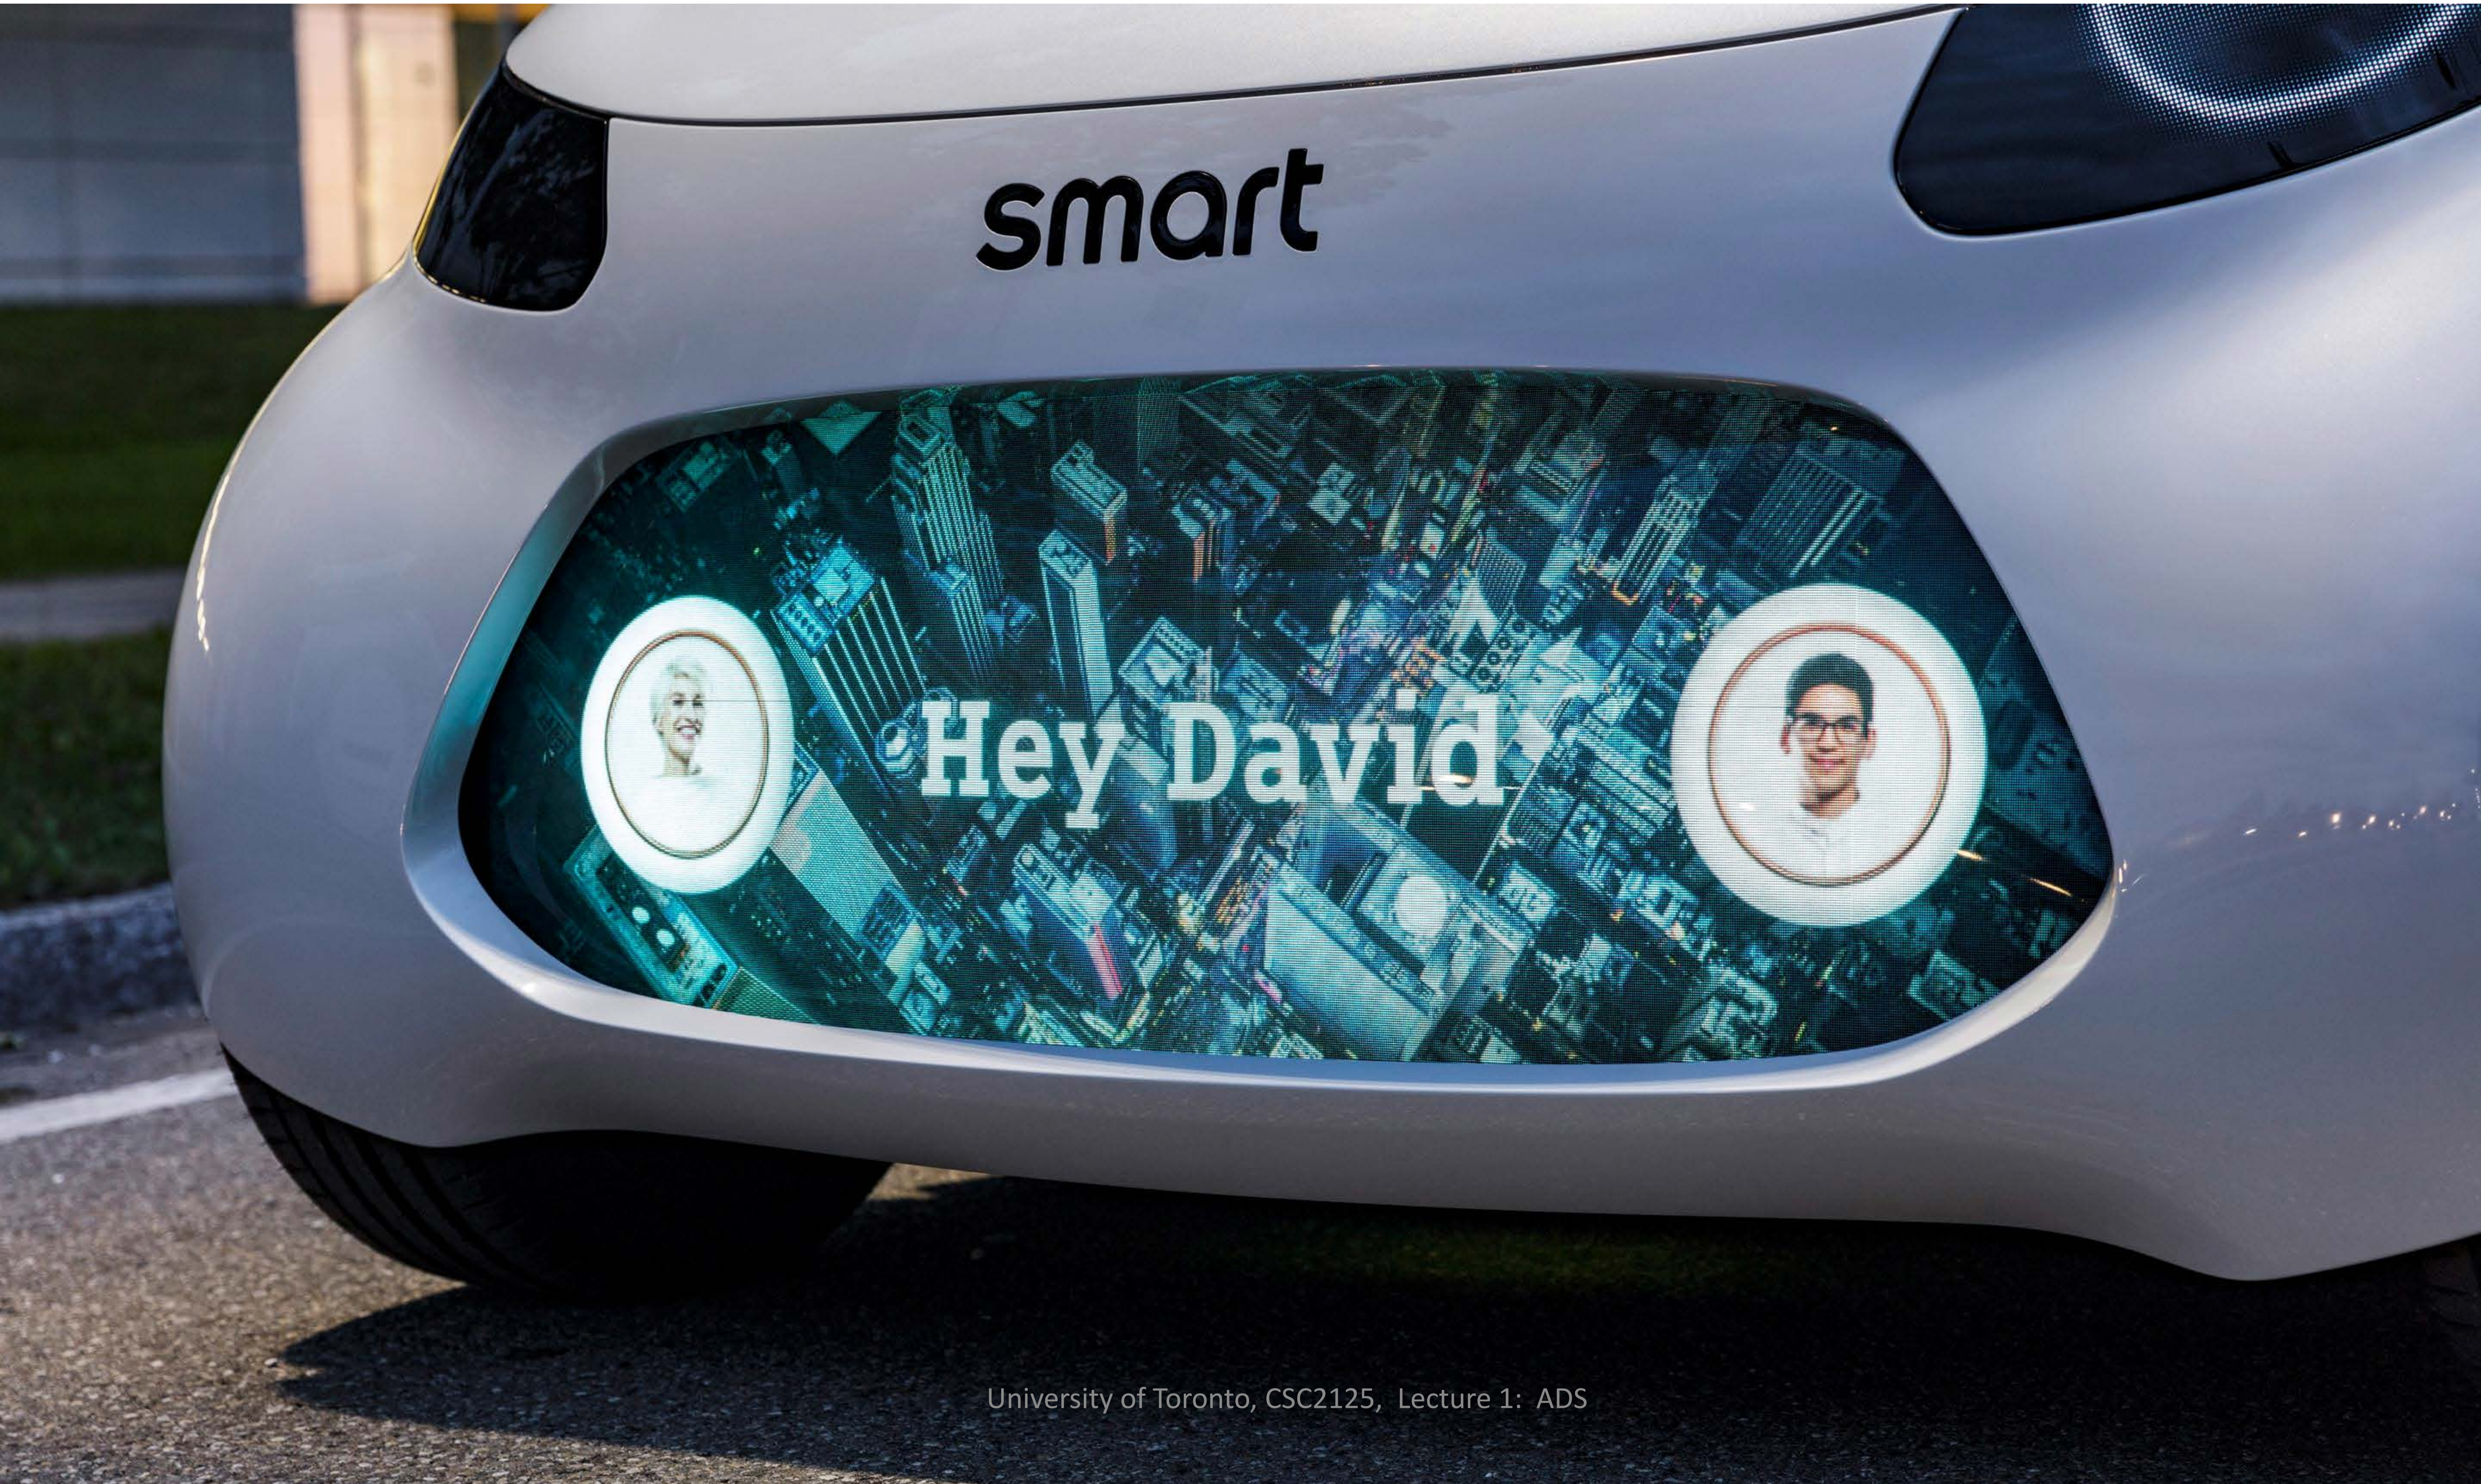

Unexpected Road Incursion by Pedestrians

Sudden Emergency Doctrine for human drivers

What is the expected standard for AVs?

Moral Machines

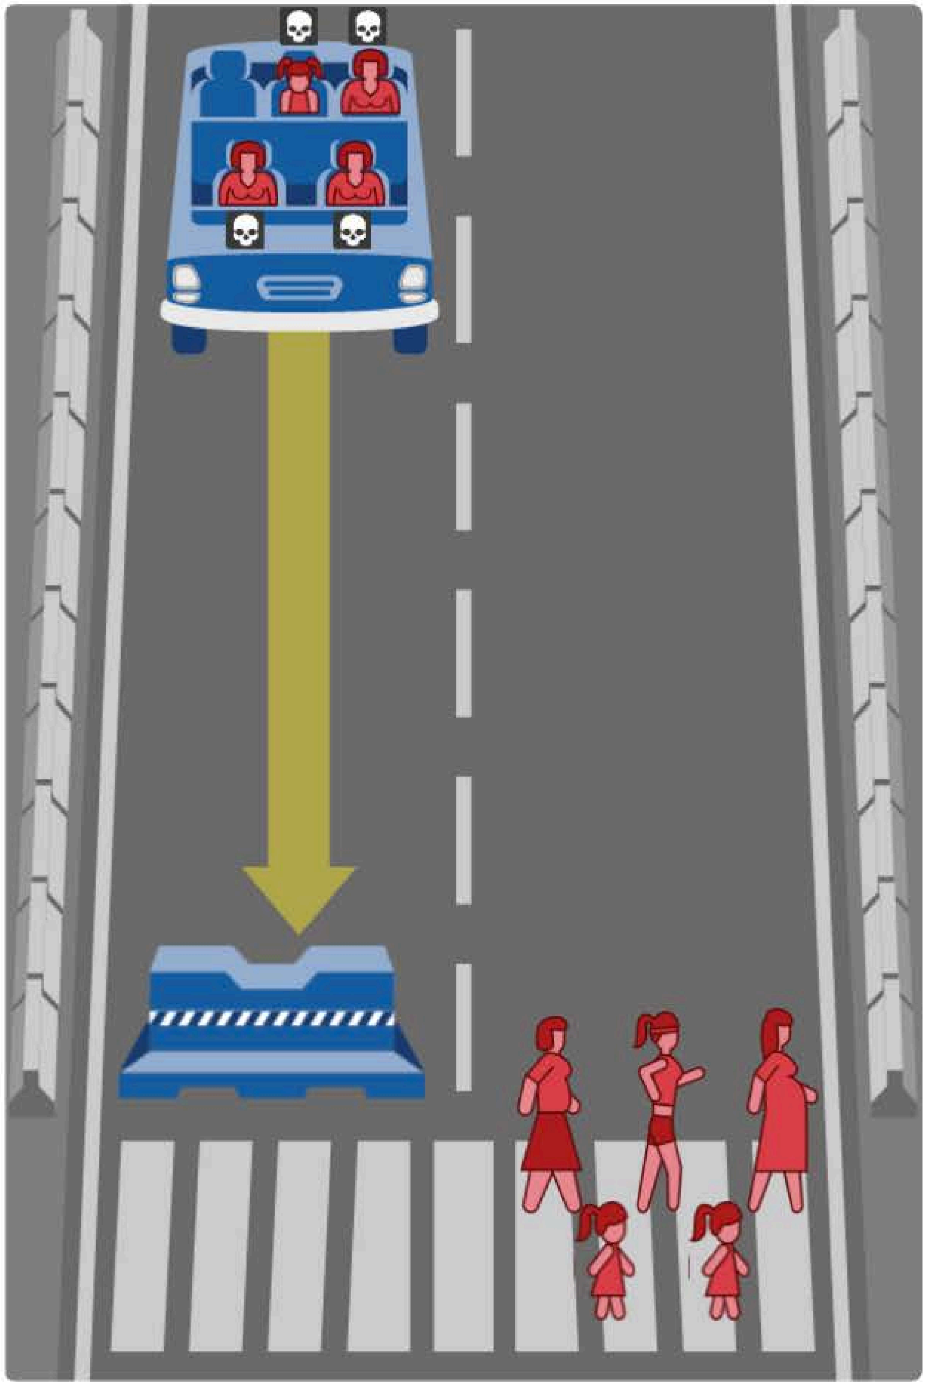

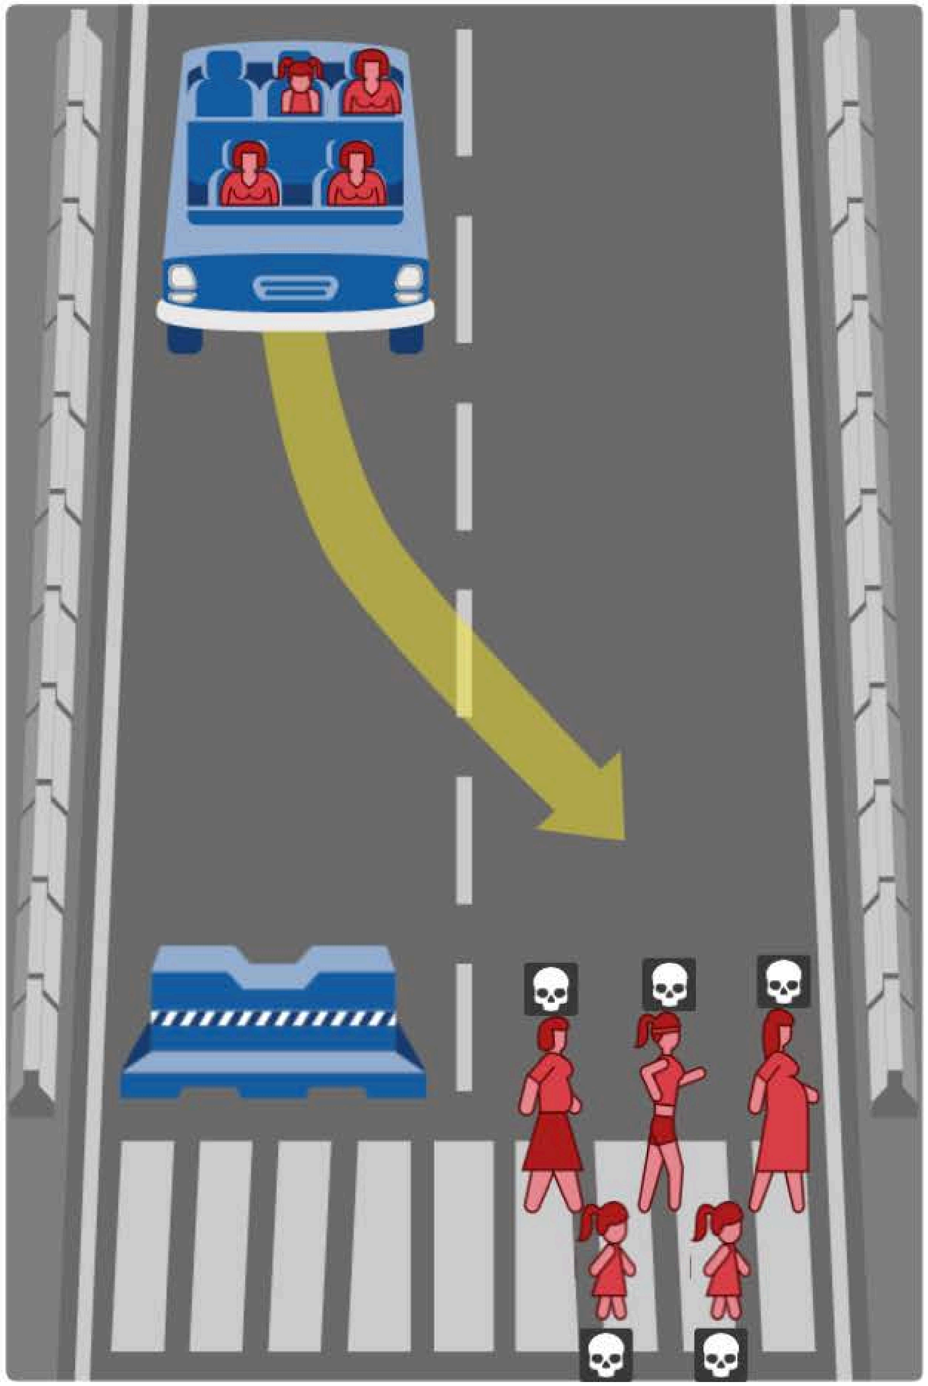

http://moralmachine.mit.edu/

Safety of Sensors and AI


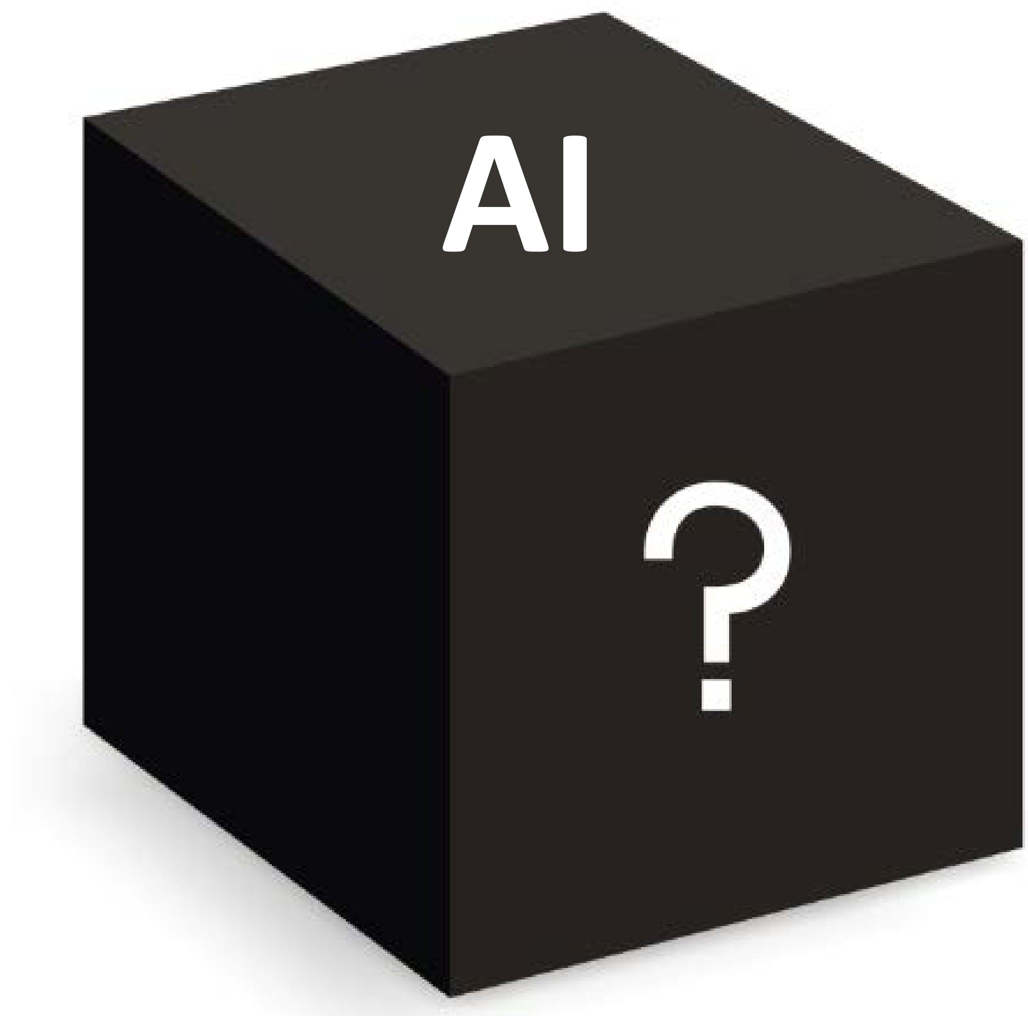

Testing Challenges

- 100 million miles driven between deadly crashes (US)
  - Crashes are rare events
  - Human drivers are extremely good, when they pay attention
- Showing equal performance by an AV with 95% confidence requires demonstrating 300 million miles driven without a deadly crash

California DMV Disengagement Reports

- Google (miles driven between disengagements):
  - 2015: 2000 miles
  - 2016: 5000 miles

Tesla Autopilot Data Collection and Testing

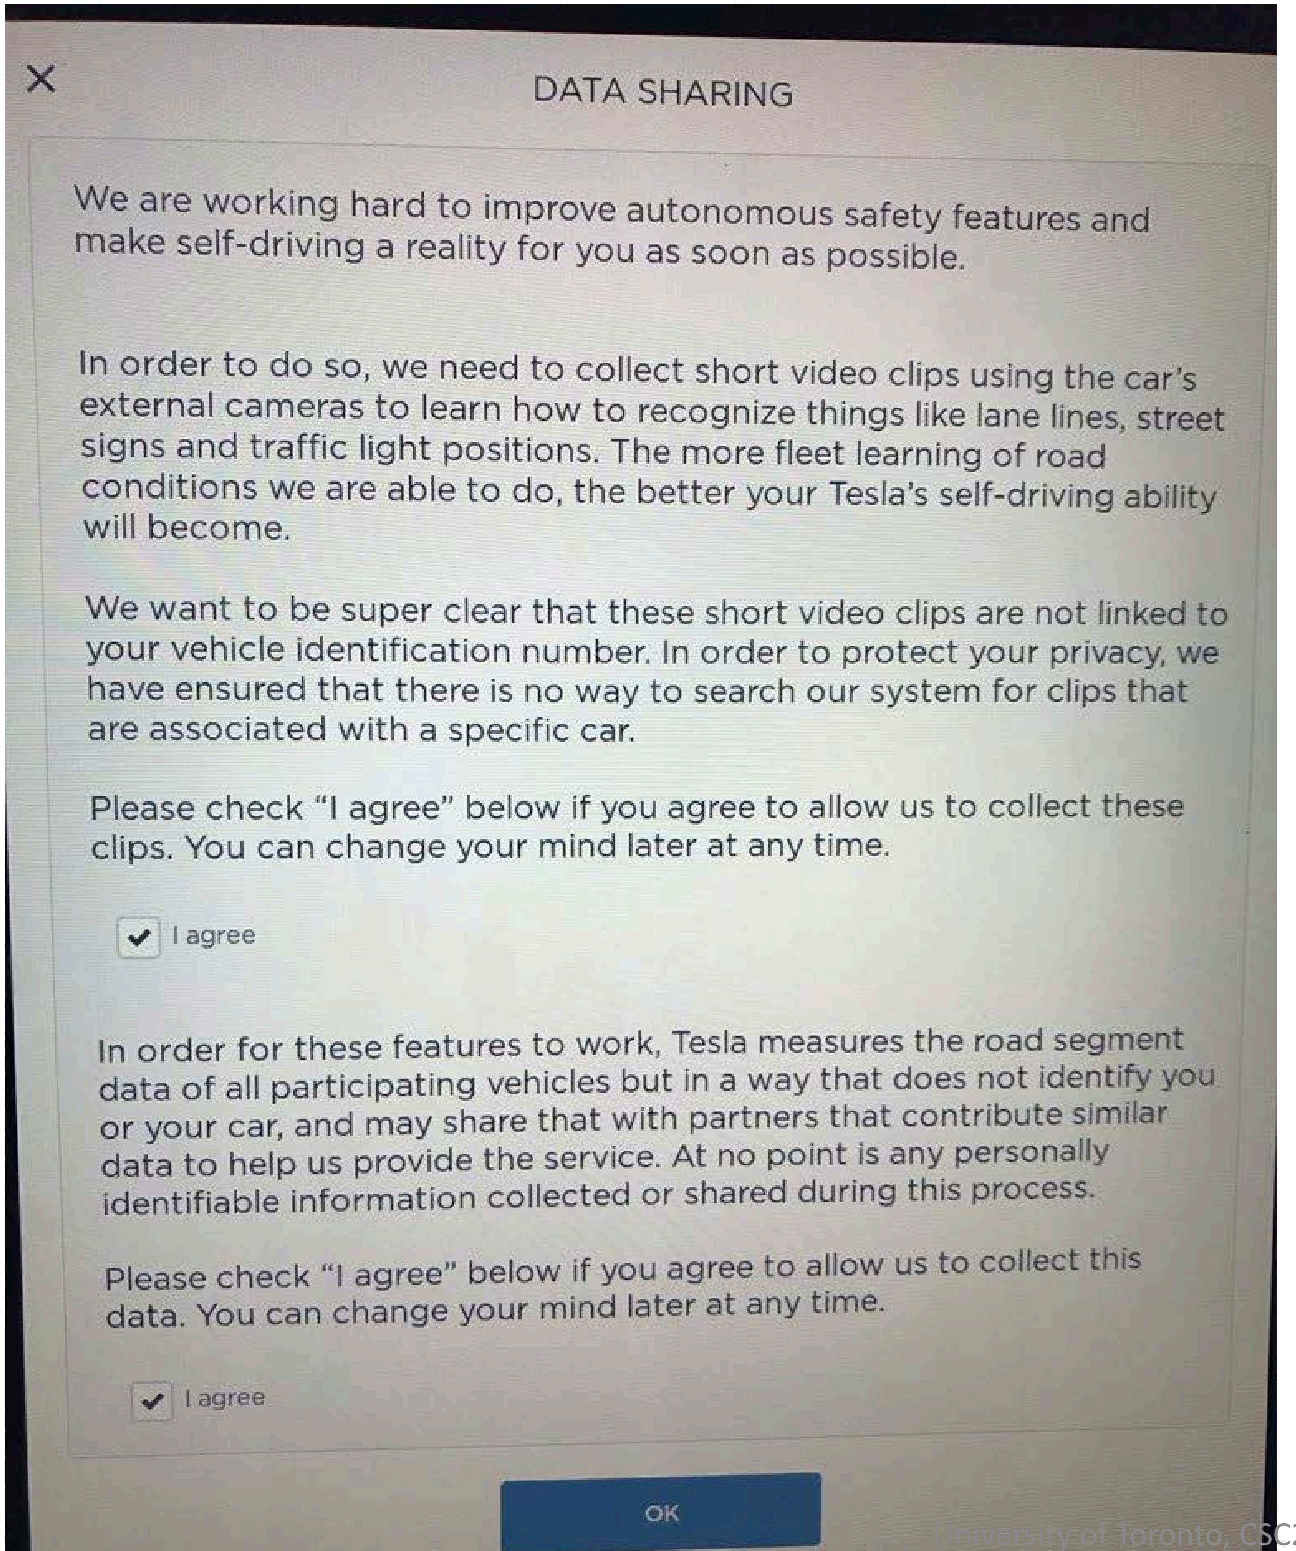

electrek.co/2017/05/06/tesla-data-sharing-policy-collecting-video-self-driving/

- In 2016, on average, 1 million miles per 10h data collected
  - Object lists
  - Driver inputs
  - Vehicle state
- Since May 5, 2017, Tesla asks for permission to gather video clips from their customers
- OtA Update staging
  - Dormant mode
  - Gradual release

Testing in Virtual World

Vehicle Physics and 3D Photo-Realistic Simulation
by Aleksandar Pocuc & Igor Ilic

V2X: Major Infrastructure Requirements