Intelligent Vehicle Dependability and Security: Challenges for Today (2026-30) and Tomorrow (2030-40)

Wilfried Steiner, TTTech Computertechnik AG

Kaunas, May/5, 2026

ifip WG.10.4

Today: Cars have not sufficiently been shown to be truly self-driving even within limited operational design domains (ODDs), revealing a remaining autonomy gap.

Tomorrow: While this autonomy gap is being closed, research should expand toward large-scale Systems of Autonomous Systems (SoAS), including self-driving cars and extending beyond them to other domains.

Overview

- Discussion of Self-Driving Cars
  - Status, Challenges, Role of IFIP WG 10.4

- Intelligent Vehicles in Other Domains
  - Similarities with and differences to Self-Driving Cars

***) The graphics are intended solely for dramatic effect and should not be interpreted as technically exact.**

ifip WG.10.4

SAE Levels of Autonomy

Alternative classifications have been proposed, e.g.,
- hands-on/off, eyes-on/off (3 cases)
- semi autonomous, fully autonomous

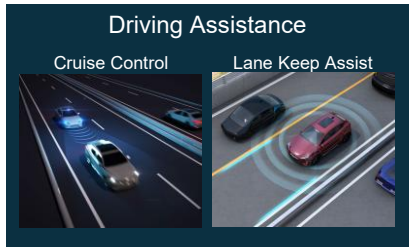

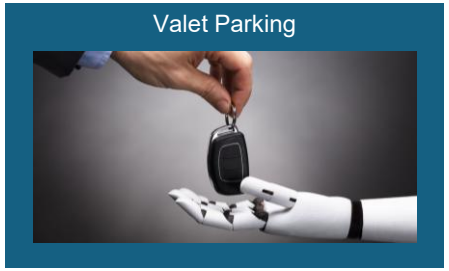

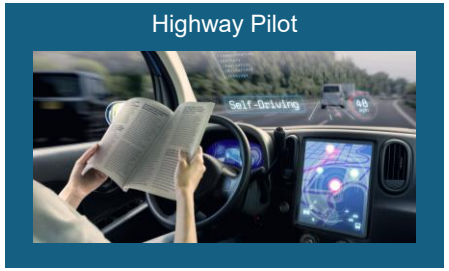

ODD ... Operational Design Domain

ifip WG.10.4

Current state of Self-Driving Cars

Data from Gartner's Hype Cycles 2010 - 2024

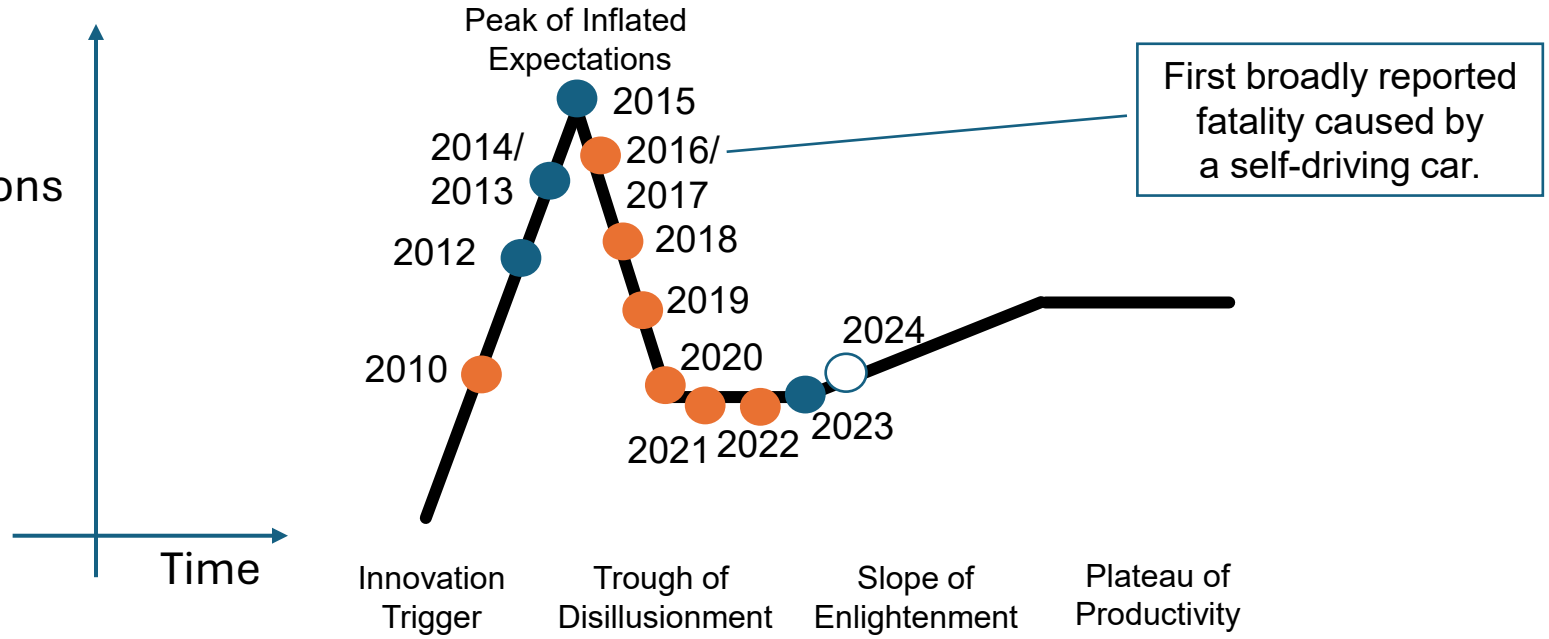

Estimated Time to reach the Plateau of Productivity:

- more than 10 years
- 5 -10 years
- 2 - 5 years

ifip WG.10.4

Current state: Selected accidents and incidents

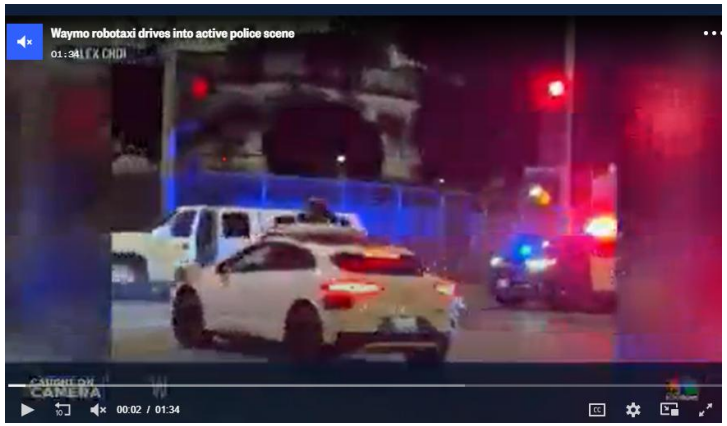

Dec/2025: A Waymo robotaxi carrying passengers entered an active LAPD standoff. No injuries were reported.

Photo by Oscar Palma.

Oct/2025: The KitKat cat was under a Waymo as it began to depart from the curb, despite a woman bending over the car. KitKat was later pronounced dead at a veterinary clinic.

2025-ongoing: Waymo robotaxi bypassing a stopped school bus.
Jan/2026: Waymo struck a child.

Dec/2025: power-outage in San Francisco causes stuck Waymo cars.

How “self-driving” are cars today?

- Tele-operator
  - A human outside the vehicle remote controls the vehicle.
- Safety driver
  - A human fallback inside the vehicle intervenes in critical situations.
- Trail car
  - A human safety driver operates from a separate vehicle.
- Tele-assistance
  - A human assists the vehicle, upon vehicle’s request.

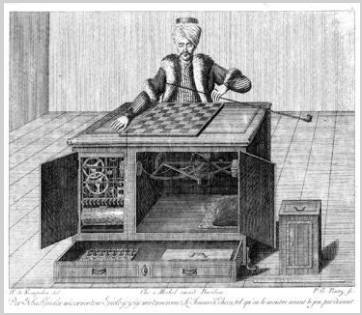

By Wolfgang von Kempelen - Copper engraving from book Original Uploader was Schaelss (talk) at 11:12, 7. Apr 2004., Public Domain, https://commons.wikimedia.org/w/index.php?curid=424092

Average miles driven before human safety driver intervention:

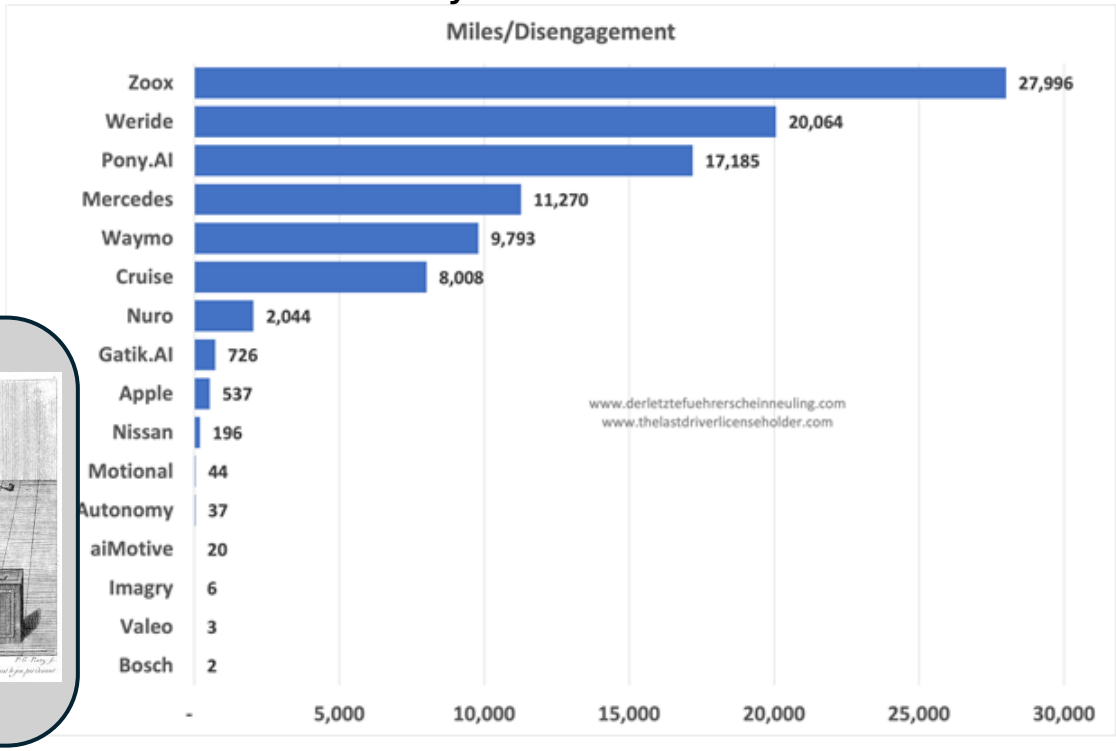

2024 Disengagement Reports from California

*) The computer-based driving system “disengages”.

ifip WG.10.4

Self-driving cars and the public opinion (simplification)

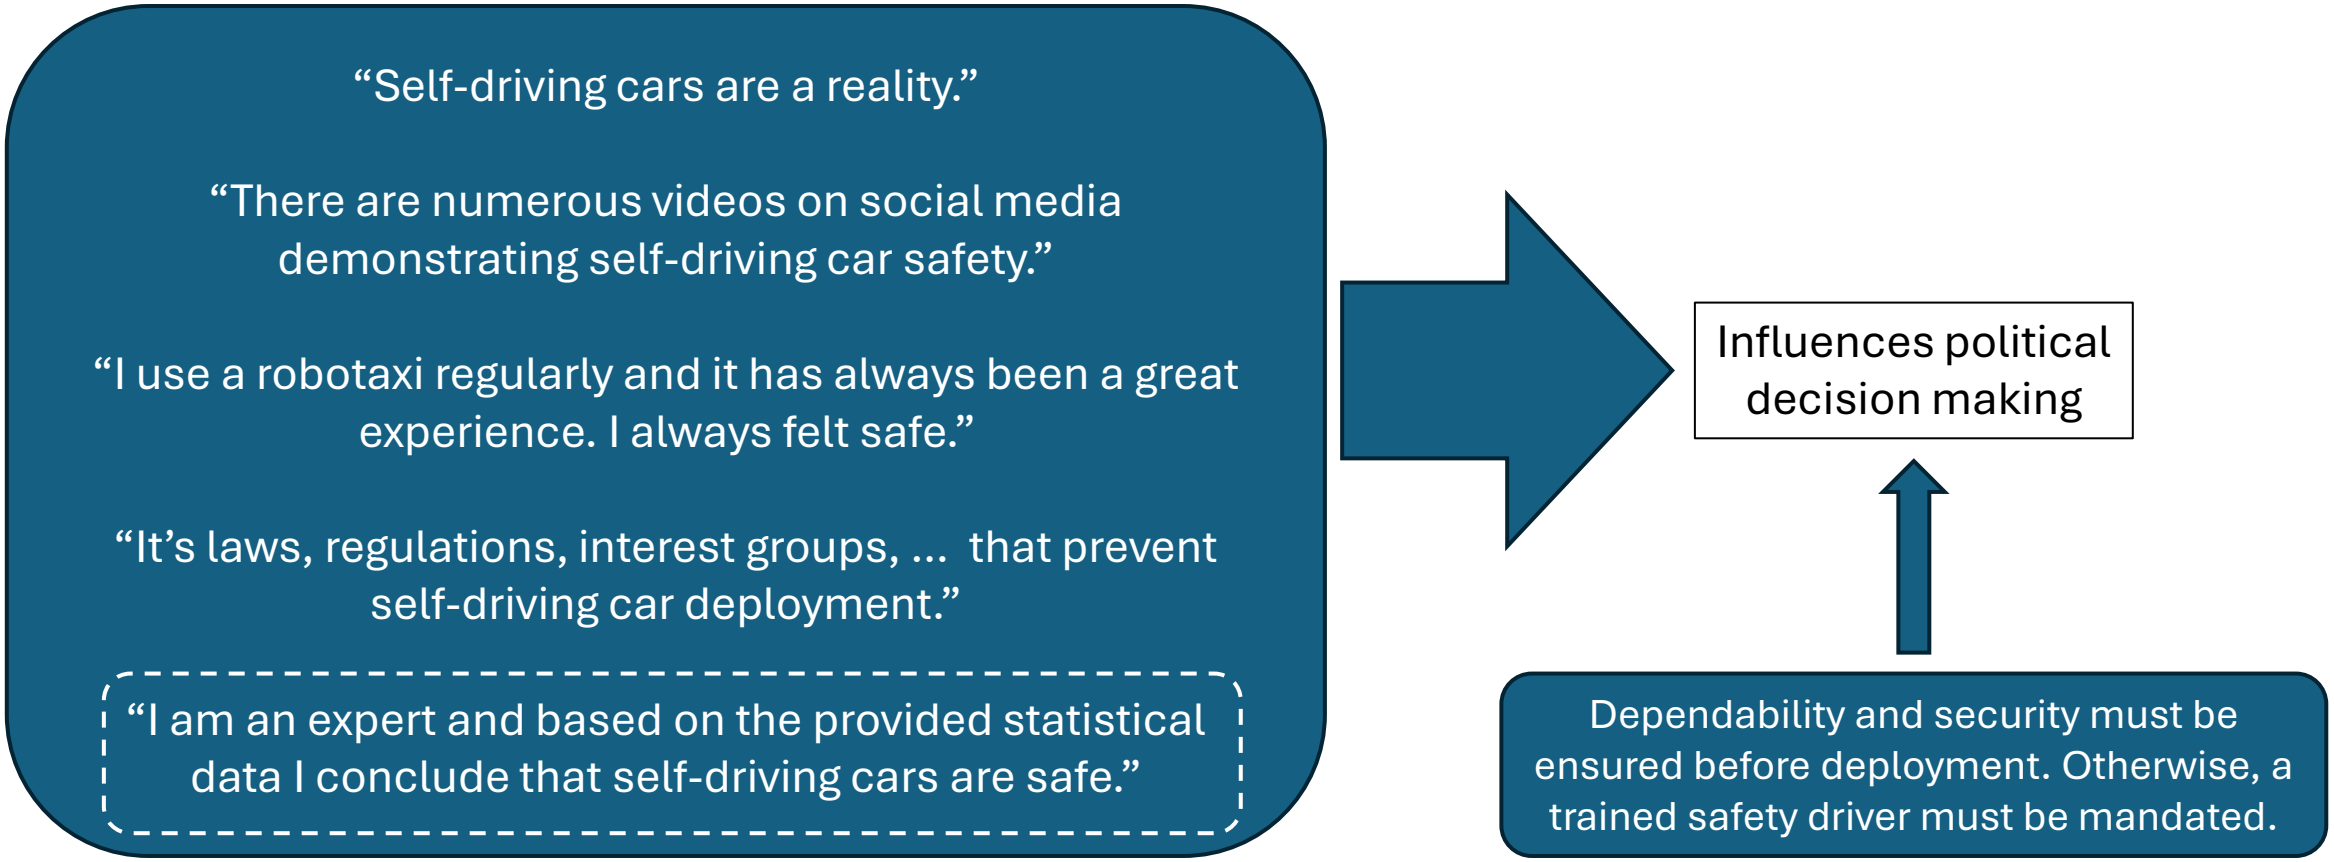

ifip WG.10.4

Current state: Non-dependability expert assessments

- Dec/2, 2025: NY Times, The Data on Self-Driving Cars Is Clear. "We Have to Change Course", Jonathan Slotkin, Neurosurgeon.
  https://www.nytimes.com/2025/12/02/opinion/self-driving-cars.html?smid=url-share

"If Waymo's results are indicative of the broader future of autonomous vehicles, we may be on the path to eliminating traffic deaths as a leading cause of mortality in the United States. While many see this as a tech story, I view it as a public health breakthrough."

Note: These media articles have high general public impact, e.g., referenced as expert opinion in an Austrian newspaper link to article at Der Standard

Dependability Experts' critique:

- Waymo cars are tele-assist (there is a publicly unknown number of people helping).
- Waymo's numbers are not indicative compared to human drivers (very beneficial ODD, low speed).
- 100M miles likely not enough: number of miles driven by itself is an insufficient metric – how many edge cases were encountered, and the vehicle response are more informative.
- There are general limits of testing.

Near-term challenges (2026-30)

- The Intelligent Vehicle itself
- The infrastructure that the Intelligent Vehicle is relying on
- New technologies in development and design
- Legislation and policy

Challenges regarding the Intelligent Vehicle itself

- Cyber World, Physical-World, and the Cyber-Physical Interface

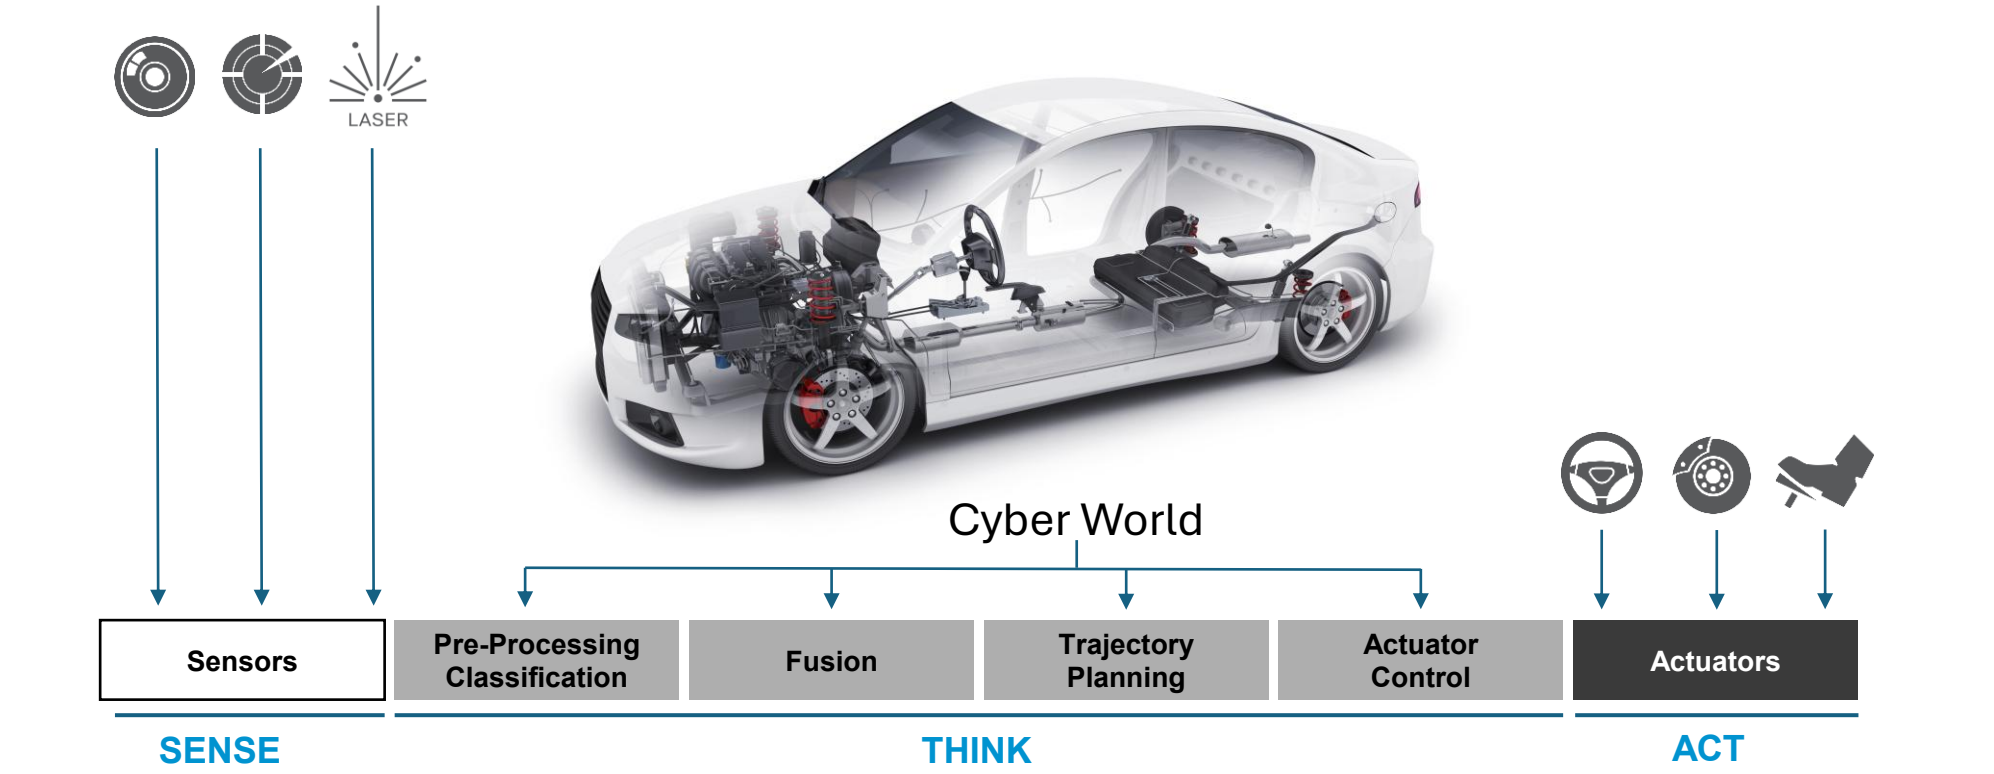

ifip WG.10.4

The Importance of the World Model

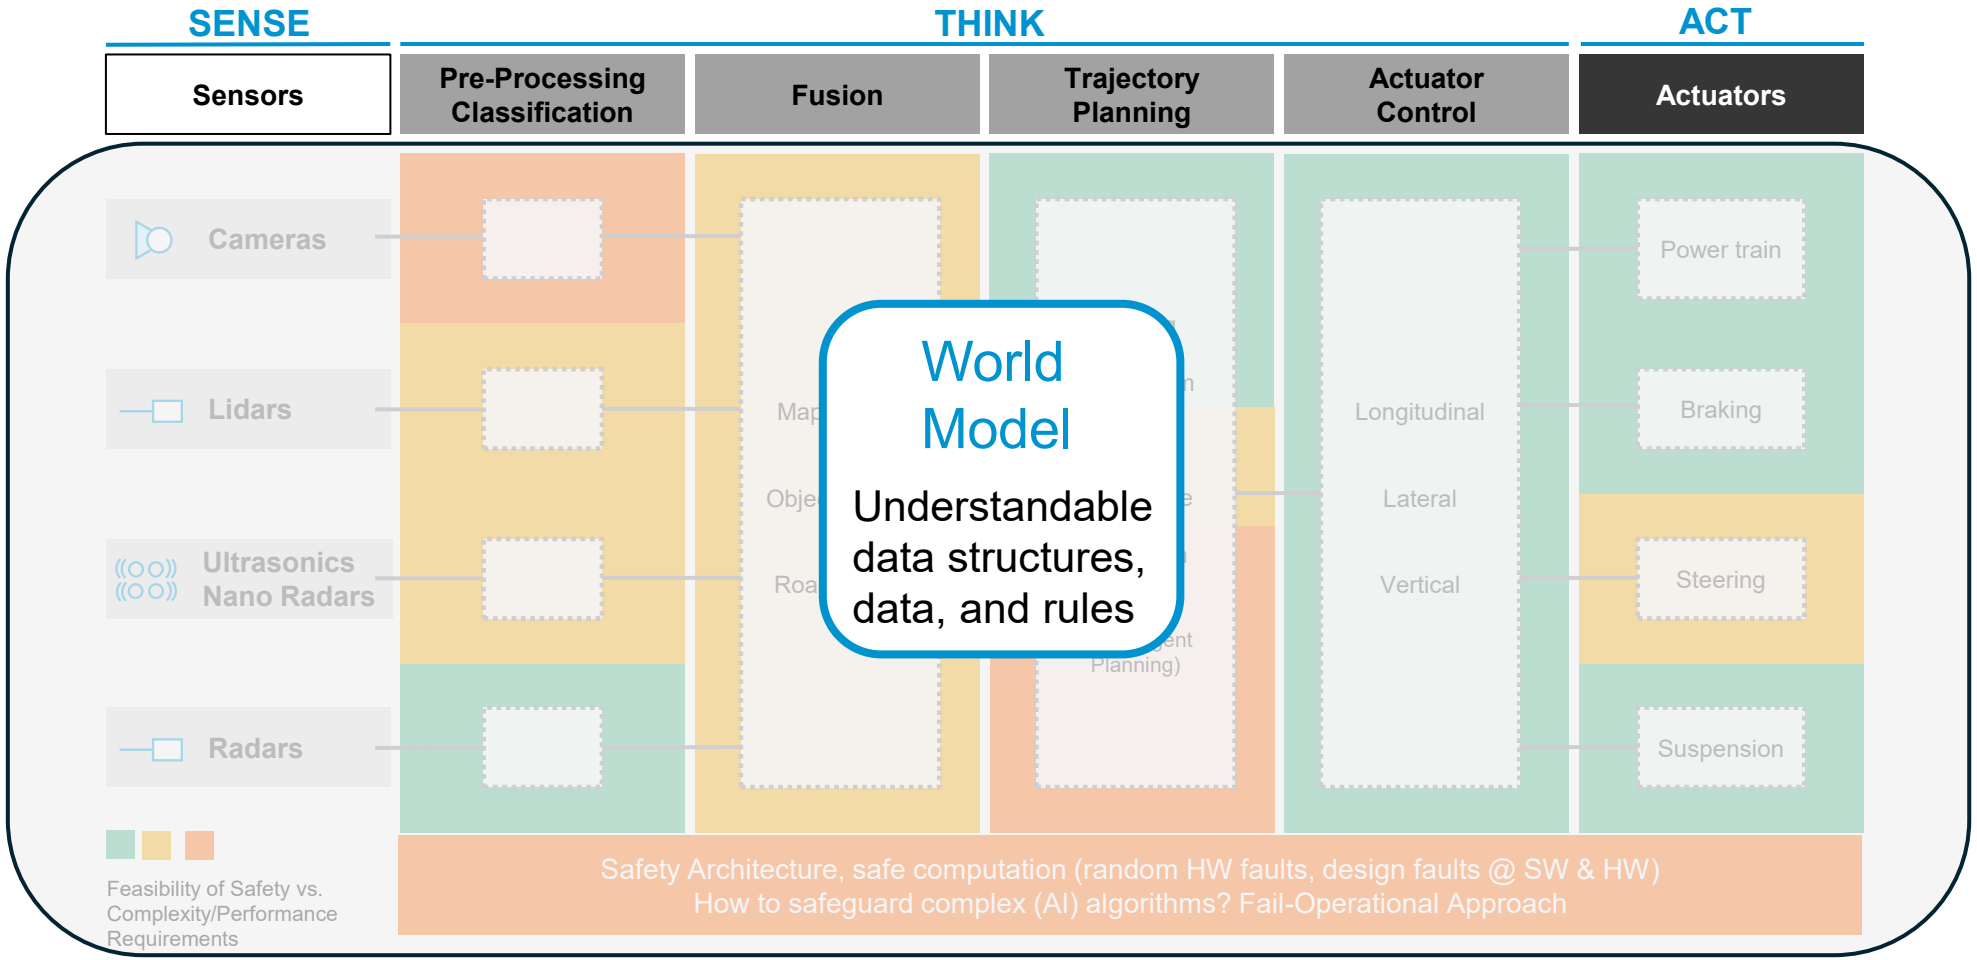

ifip WG.10.4

Two Approaches:
Distributed Decomposable Architecture vs. Monolithic End-to-End

This must run on a fault-tolerant computer.

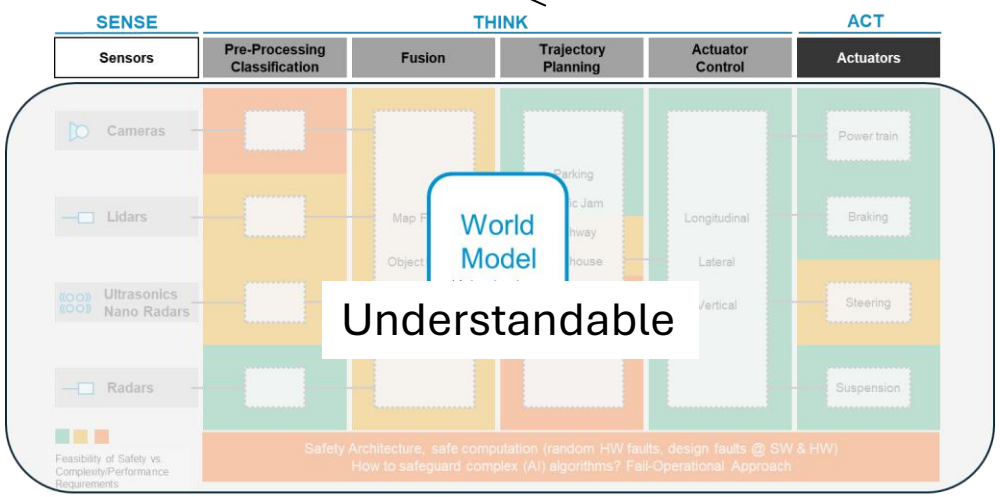

Distributed Decomposable Architecture

Likely the better choice for safety and security assurance.

Enables Continuous Validation, e.g., using predictive processing*

- at t1: compute model state at t2
- at t2: check if sensor readings confirm computed state

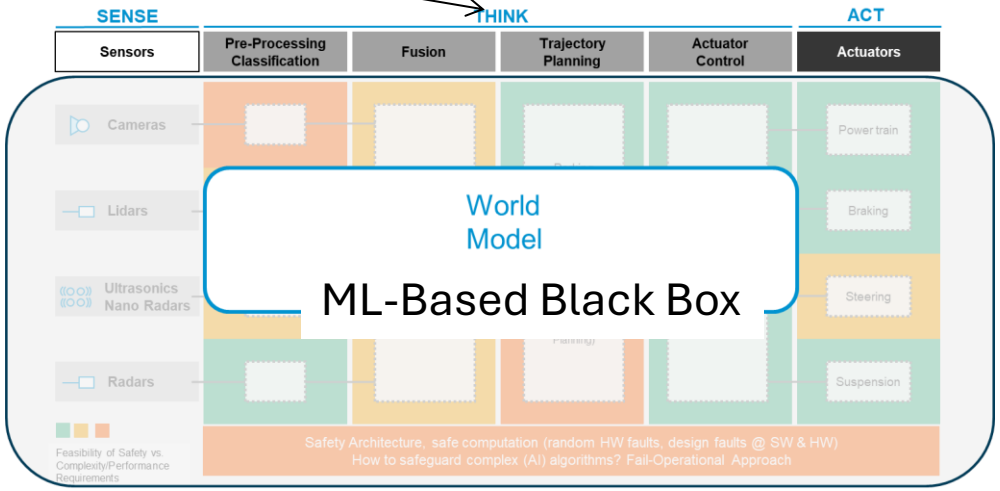

Monolithic End-to-End

Restricts V&V possibilities to the full system only.

* Rushby, J. (2022). *Models and their Validation and their Role in Perception And in Safe Autonomous Vehicles*. IFIP WG10.4 Virtual Meeting.
URL: http://www.csl.sri.com/users/rushby/abstracts/ifip-11may22

ifip WG.10.4

Challenges regarding the Intelligent Vehicles itself

Top-Down Design of a Decomposable Architecture

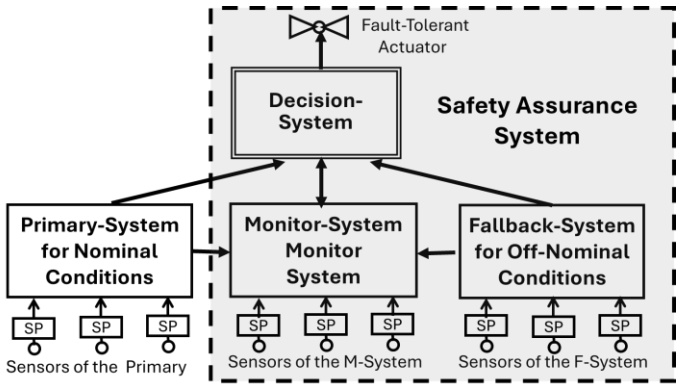

Figure depicts "logical" units. Ideally, they would be implemented as separate & diverse ECUs.

Mapping to a system of ECUs, under constraints

E/E Architecture options

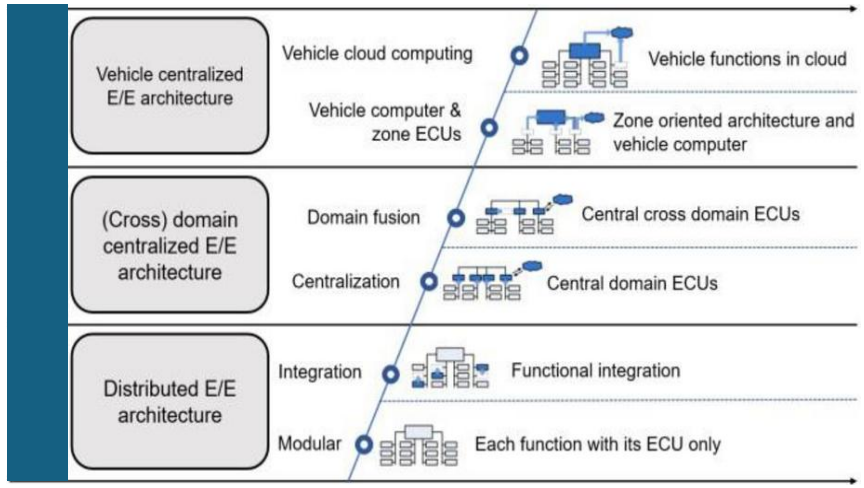

Figure depicts different ECU arrangement strategies.

Design Constraints (examples):

- Legacy re-use (full ECUs, technologies, etc.),
- SWaP-C,
- HW performance,
- Availability of certifiable HW & SW components,
- Scalable to support model variants
- …

Kopetz, H. (2021). An Architecture for Driving Automation. URL: https://www.the-autonomous.com/news/an-architecture-for-driving-automation

E/E … Electrical/Electronic
ECU … Electronic Control Unit

ifip WG.10.4

Challenges regarding the infrastructure that the Intelligent Vehicle is relying on

- Safety driver
  - A human fallback inside the vehicle intervenes in critical situations.
- Trail car
  - A human safety driver operates from a separate vehicle.
- Tele-assistance
  - A human assists the vehicle upon vehicle's request.
- Tele-operator
  - A human outside the vehicle remote controls the vehicle.

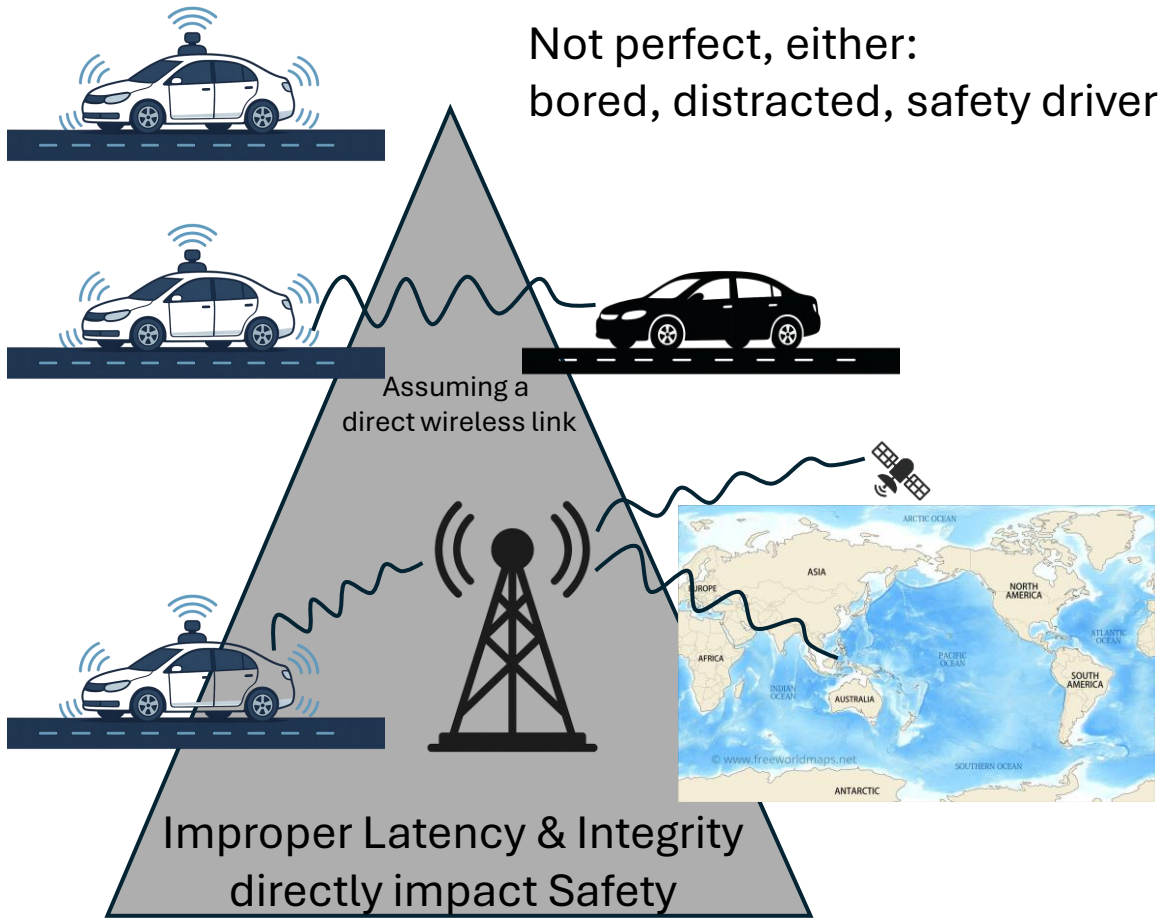

Challenges regarding new technologies in development and design

Positive example of GenAI use:

- KitKat experiment*): LLMs list and explain dependability issues of the KitKat incident

Negative example of GenAI use:

- Stadler Metro Car experiment*): LLMs wrongly assign new cars to different target customers, even when iteratively feeding them the respective other answers

*) see annex slides for experiment details

Continuous improvements and role of neuro-symbolic AI (symbiosis of ML and formal methods).

Still: for now, use Zero Trust methodology when dealing with LLMs. Never trust, always verify.

Opportunity: using LLM technology for development and design

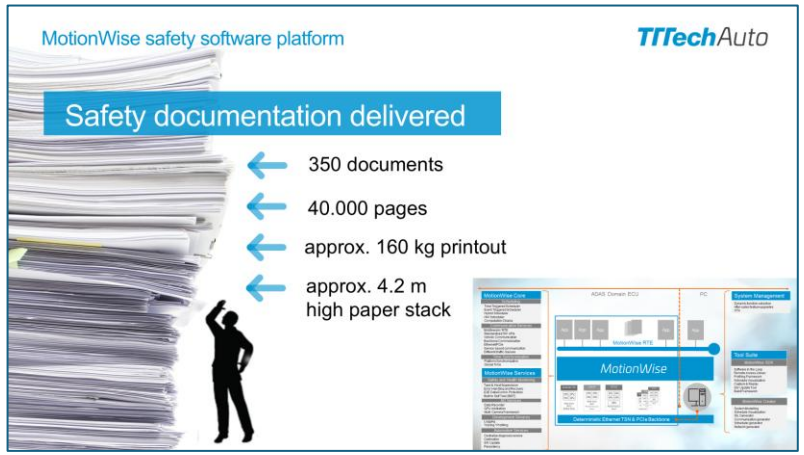

Typical traditional example. How much of such artefacts could be generated by GenAI?

Legislation and policy challenges

- Availability and transparency of field data regarding incidents and tele-assistant help of vehicles on the road today
  - While California collects some data, not all states where driverless-cars are allowed to experiment do so.
  - CA data is a good start but is not comprehensive enough to assess state of current operational safety and security.
  - It also does not shed light on how frequently tele-assistants have to help and what situations lead to self-driving systems being stuck.

- Fragmentation and localization of regulations and laws is problematic

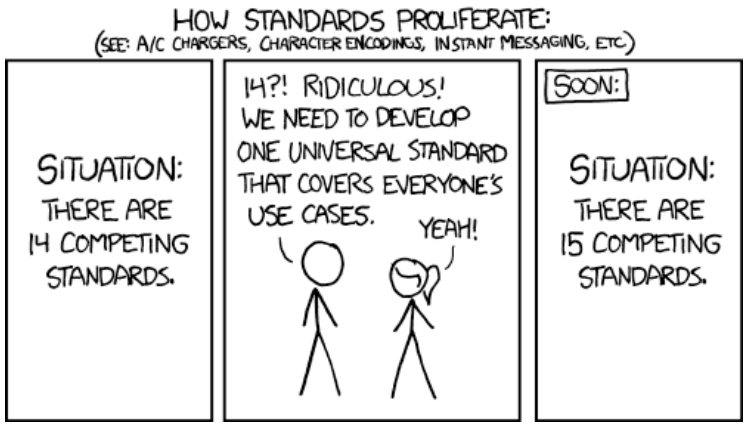

https://xkcd.com/927

Possible Role of IFIP WG 10.4 Today (2026-30)

- Develop quantitative acceptance criteria for solutions that use AI/ML.
  - (#miles driven is the insufficient metric – how many edge cases were encountered, and the vehicle response are more informative)
- Definition of the means of deciding whether the acceptance criteria are met.
  - Develop guidelines for collecting field data regarding incidents and tele-assistant help.
- Develop methods to construct the overall system assurance case.
- Give recommendations for preferred technologies to realize IVs and their usage that ease the assurance case argumentation.
  + Develop technologies that ease the assurance case argumentation.

ifip WG.10.4

Longer-term challenges (2030-40)

- Regulatory bodies establish safety and security standards.
- AI/ML algorithms can classify objects with extremely high level of fidelity.
- Operational Design Domain (ODD) Extension: Testing, validation, and verification methodologies can account for extreme cases of operational environment, including weather, obstacles, people, other vehicles, and emergency activities.
- **Impact of non-availability of fleets of autonomous systems.**

Possible Role of IFIP WG 10.4 Tomorrow (2030-40)

*) Assuming positive development of self-driving cars / robotaxis.

- Research dependability and security of a System-of-Autonomous-Systems that include:
  - fleets of self-driving cars,
  - autonomous duty and transportation vehicles,
  - infrastructure (including maintenance and construction machines),
  - unmanned aerial vehicles and vertical take-off and landing aircrafts,
  - etc.
- Address the implied scaling challenges in:
  - definition of quantitative acceptance criteria and their measurement
  - assurance case construction & system technology recommendations
  - etc.

ifip WG.10.4

Relation to Intelligent Vehicles In Markets Adjacent to Self-Driving Cars

Example Verticals

AVIATION	*SPACE & DEFENSE*	*MOBILE MACHINERY*
➜ Commercial aviation	→ Human and robotic exploration	→ Agricultural vehicles
➜ Defense aviation	→ Launch vehicles and space transportation	→ Construction machines
➜ Rotorcraft (vertical lift)	→ Satellites (Platform & Payload)	→ Municipal vehicles
➜ Advanced Air Mobility (AAM / eVTOL)	→ Land systems/ Land vehicles	→ Material handling

ifip WG.10.4

Mobile Machinery Example: Intelligent Snow-groomer

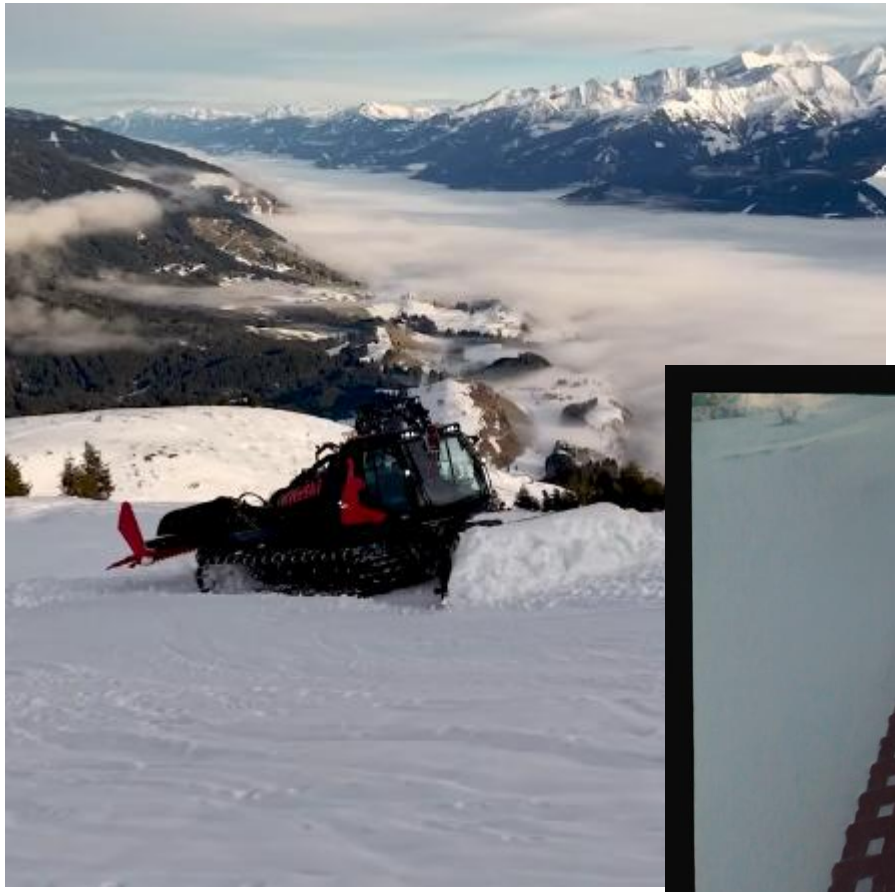

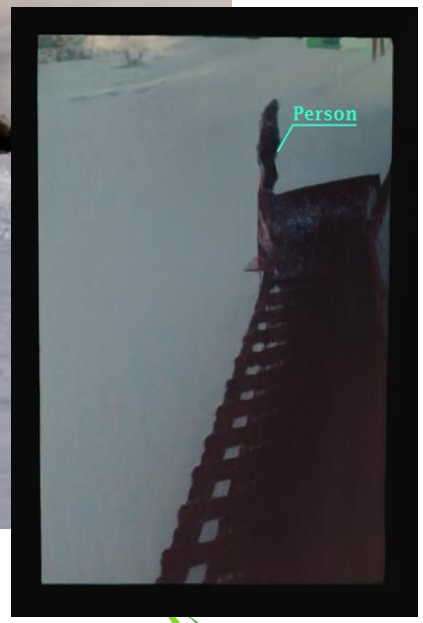

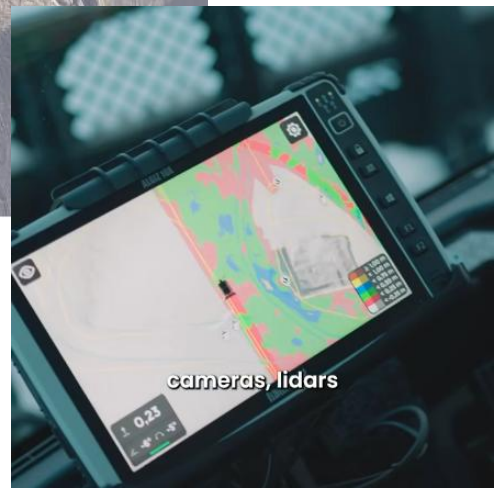

Similarities to Markets Adjacent to Self-Driving Cars

- Intelligent Vehicles in different domains seem similar but they can be very different.
  - E.g., radiation challenge in space / aerospace.
- Autonomous operations have been proven since decades in space, e.g., docking to the ISS, autopilots.
- In Aerospace: regulation and certification is in place and is well-coordinated.
  - It took significant time and is based on industry-wide cooperations w/ regulators.
- Are self-driving cars the most difficult Intelligent Vehicles?
  - Likely not in all aspects, but the specifics of their operational environment stand out (e.g., number of vehicles, pedestrians, bikes, etc.).

Automotive as an Enabler Example: LIDAR

Period	Typical LiDAR Type	Approx. Cost per Unit (USD)	Context / Milestones
2010–2012	Mechanical spinning LiDAR (e.g. Velodyne HDL-64E)	**$60k – $80k**	DARPA-era prototypes; roof-mounted sensors; research and robotaxis only [innoviz.tech], [fleetowner.com]
2013–2015	High-resolution mechanical LiDAR	**$30k – $75k**	Google/Waymo, Cruise, Uber AV programs; LiDAR = dominant perception sensor [innoviz.tech], [fleetowner.com]
2016–2018	Early solid-state / hybrid (MEMS) LiDAR	**$8k – $20k**	First automotive-grade designs; cost reduction via fewer moving parts [innoviz.tech], [spectrum.ieee.org]
2019–2020	Automotive-grade solid-state LiDAR	**$3k – $8k**	OEM pilot programs (Volvo, Audi); sensor miniaturization and scale effects [innoviz.tech], [fleetowner.com]
2021–2022	Production-intent solid-state LiDAR	**$1k – $3k**	Series-production contracts (Luminar, Innoviz, Hesai); roofline & grille integration [fleetowner.com], [innoviz.tech]
2023–2024	Mass-production solid-state LiDAR	**$500 – $1k**	L2+/L3 passenger cars (China & EU); multi-LiDAR vehicle setups become viable [fleetowner.com], [businessinsider.com]
2025 (today)	High-volume automotive solid-state LiDAR	**$150 – $500**	Near-radar pricing; standard component in advanced ADAS & autonomy stacks [fleetowner.com], [spectrum.ieee.org], [businessinsider.com]

Generated with MS Copilot

ifip WG.10.4

Today: Cars have not sufficiently been shown to be truly self-driving even within limited operational design domains (ODDs), revealing a remaining autonomy gap.

Tomorrow: While this autonomy gap is being closed, research should expand toward large-scale Systems of Autonomous Systems (SoAS), including self-driving cars and extending beyond them to other domains.

Acknowledgments

- Significant parts of this talk have been a direct result from discussions in regular online meetings with John Meyer, Jay Lala Chuck Weinstock, and Carl Landwehr.
- Further information on Intelligent Vehicles Dependability and Security can be found at https://ivds.dependability.org/

ifip WG.10.4

Annex

The KitKat Incident – Oct 27, 2025

- A Waymo robotaxi struck KitKat, the beloved neighborhood Bodega cat in San Francisco's Mission District, late on the night of Monday, October 27, 2025, at approximately 11:30 PM.
  - Videos show KitKat laying down in front of the Waymo's right front tire.
  - A bystander ran to try to get him to move. She leans in front of the car to get the cat to move but it would not. When she stood up and apparently tried to get the Waymo not to move, but it left immediately, running over and killing KitKat.

Dependability Issues related to KitKat

- We asked Gemini|ChatGPT|Claude|Copilot to discuss dependability issues highlighted by this incident
  - All provided reasonable answers, highlighting:
    - Perception System Failures
    - Decision-Making Under Edge Cases
    - Fail-Safe Mechanisms
    - Transparent Accountability
    - Community Trust

Changing it up – KitKat is a candy bar

- We then got a bit irreverent and told the engines that we actually meant the candy bar
  - The engines all highlighted the optimization for larger objects, missing smaller objects (road hazards).
    - The ability to detect and classify all objects regardless of size/perceived threat
    - Err on the side of caution
    - Adapt to edge cases (e.g., small animals, toys, candy bars)
  - Copilot: So while no candy bars have been harmed (that we know of), the question raises a sweet point: dependability isn't just about avoiding collisions — it's about understanding the world in all its messy, snack-sized detail.

Newcastle Metro Issue -- September 2025

- Gemini and Claude were provided with a photo of a brand new rail car parked outside of a Stadler plant in Switzerland and asked to identify the Metro System that it was for.
- They came up with radically different answers: Gemini – Sydney, Claude – Newcastle-upon-Tyne.
- We cross-fed the answers and both gave reasons why the other was wrong – over multiple iterations.
- The logo on the car was black and orange. We asked Gemini to shows a picture of the logo it was basing its answer on: blue and white! It finally admitted it was wrong.