nuro

October 11, 2022

Autonomy Challenges

Aleksandr Petiushko

Nuro
Autonomy Interaction Research

Talk at BDD Seminar
Berkeley DeepDrive

Content

Part I: Policies

01 Self-driving policies as a reasonable gate?

02 Lack of simulators to trust

03 Lack of proxy metrics to trust

04 Autonomy stack: is it really a stack?

05 Academia and Industry

06 Conclusion

US Department of Transportation

USDOT: Automated Vehicles activities

Sep 2016 — Federal Automated Vehicles Policy: Accelerating the Next Revolution In Roadway Safety

Sep 2017 — Automated Driving Systems 2.0: A Vision for Safety

Oct 2018 — Automated Vehicles 3.0: Preparing for the Future of Transportation

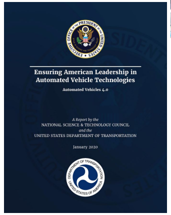

Jan 2020 — Automated Vehicles 4.0: Ensuring American Leadership in Automated Vehicle Technologies

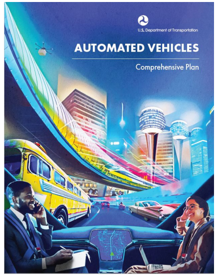

Jan 2021 — Automated Vehicles Comprehensive Plan

202X — YYY

Five Eras of Safety

According to National Highway Traffic Safety Administration (NHTSA)

1950-2000
Safety/Convenience Features

Cruise Control
Seat Belts
Antilock Brakes

2000-2010
Advanced Safety Features

Electronic Stability Control
Blind Spot Detection
Forward Collision Warning
Lane Departure Warning

2010-2016
Advanced Driver Assistance Features

Rearview Video Systems
Automatic Emergency Braking
Pedestrian Automatic Emergency Braking
Rear Automatic Emergency Braking
Rear Cross Traffic Alert
Lane Centering Assist

2016-2025
Partially Automated Safety Features

Lane Keeping Assist
Adaptive Cruise Control
Traffic Jam Assist

2025+
Fully Automated Safety Features

Everything?
** probably not only above things but even more and/or wider adoption*

Levels of Automation

0	1	2	3	4	5
No Automation	**Driver Assistance**	**Partial Automation**	**Conditional Automation**	**High Automation**	**Full Automation**
Zero autonomy; the driver performs all driving tasks.	Vehicle is controlled by the driver, but some driving assist features may be included in the vehicle design.	Vehicle has combined automated functions, like acceleration and steering, but the driver must remain engaged with the driving task and monitor the environment at all times.	Driver is a necessity, but is not required to monitor the environment. The driver must be ready to take control of the vehicle at all times with notice.	The vehicle is capable of performing all driving functions under certain conditions. The driver may have the option to control the vehicle.	The vehicle is capable of performing all driving functions under all conditions. The driver may have the option to control the vehicle.

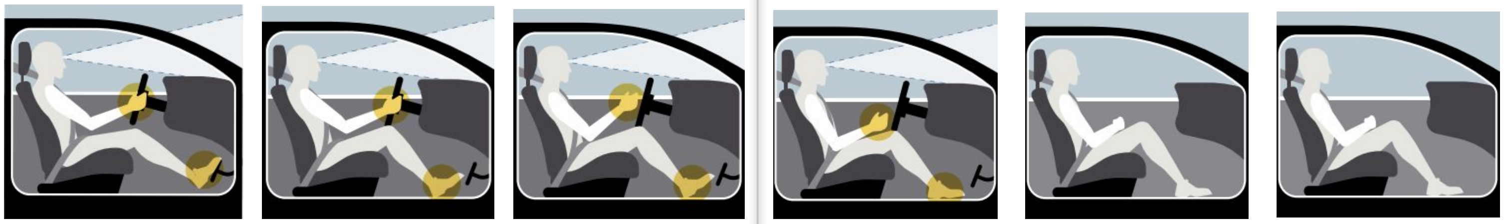

NHTSA: 1, 2 + SAE (Society of Automotive Engineers) J3016

AV Holistic Plan

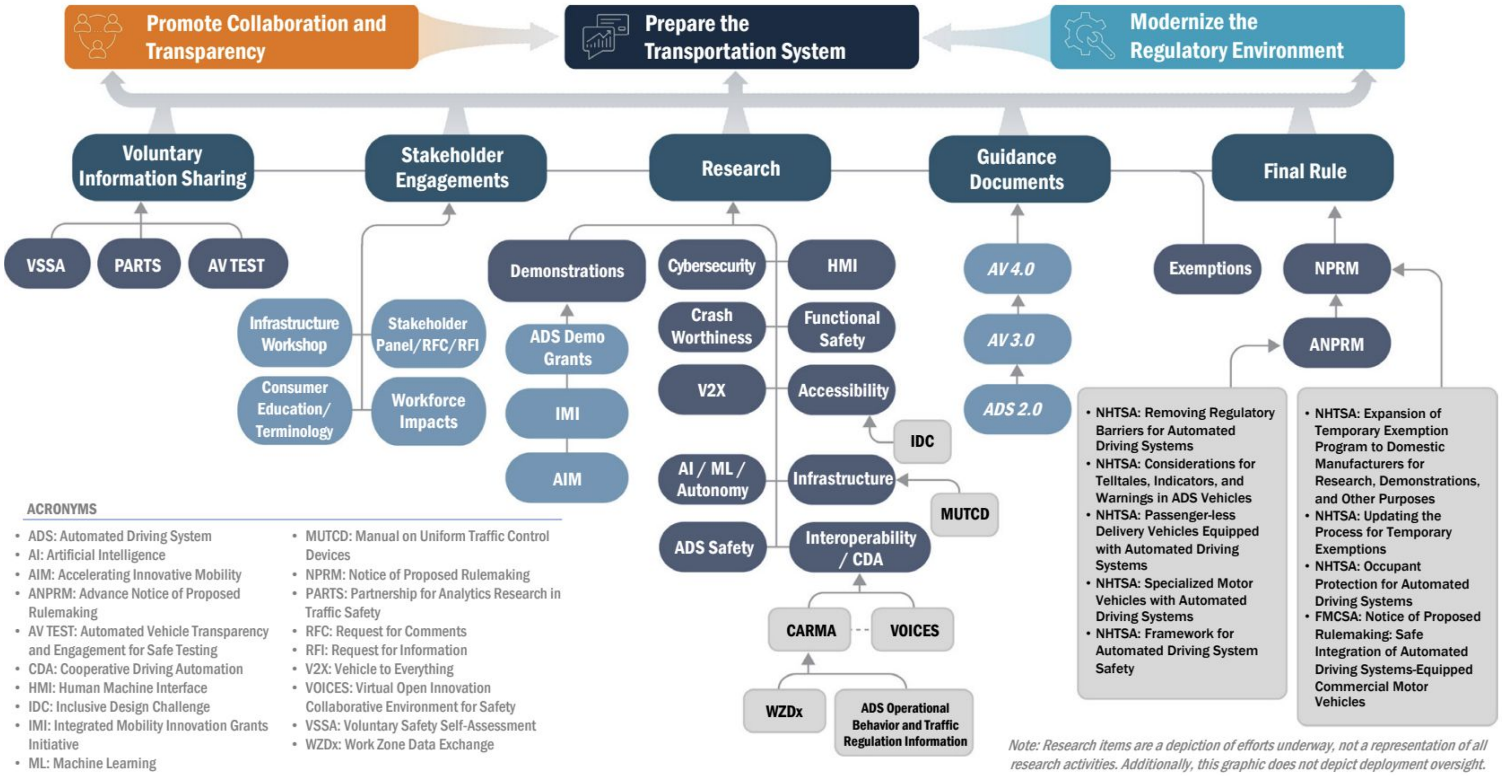

USDOT: [Automated Vehicles Comprehensive Plan]

State Regulations

CA DMV Autonomous Vehicle Testing Permit holders

- **Permit Type**
  - Testing with a Driver
    - 50 companies
    - ✅ **Nuro**
  - Driverless Testing
    - 7 companies
    - ✅ **Nuro**
  - Deployment
    - 3 companies
    - ✅ **Nuro**

CA and NV are the only states that allow deployment and require a permit.
* And NV's process is much simpler

California Department of Motor Vehicles (CA DMV)

State Regulations: metrics

Main metrics to report:

- Collisions
- Disengagements
- Mileage (in addition to Disengagement)

Article 3.7. Testing of Autonomous Vehicles

(Effective 4/13/2022)

§ 227.00. Purpose.

§ 227.02. Definitions.

§ 227.04. Requirements for a Manufacturer's Testing Permit.

§ 227.06. Evidence of Financial Responsibility.

§ 227.08. Instrument of Insurance.

§ 227.10. Surety Bond.

§ 227.12. Certificate of Self-Insurance.

§ 227.14. Autonomous Test Vehicles Proof of Financial Responsibility.

§ 227.16. Identification of Autonomous Test Vehicles.

§ 227.18. Manufacturer's Testing Permit and Manufacturer's Testing Permit - Driverless Vehicles.

§ 227.20. Review of Application.

§ 227.22. Term of Permit.

§ 227.24. Enrollment in Employer Pull Notice Program.

§ 227.26. Prohibitions on Operation on Public Roads.

§ 227.28. Vehicles Excluded from Testing and Deployment.

§ 227.30. Manufacturer's Testing Permit Application.

§ 227.32. Requirements for Autonomous Vehicle Test Drivers.

§ 227.34. Autonomous Vehicle Test Driver Qualifications.

§ 227.36. Autonomous Vehicle Test Driver Training Program.

§ 227.38. Manufacturer's Permit to Test Autonomous Vehicles that DO Not Require a Driver.

§ 227.40. Refusal of Autonomous Vehicle Testing Permit or Testing Permit Renewal.

§ 227.42. Suspension or Revocation of Autonomous Vehicle Testing Permit.

§ 227.44. Demand for Hearing.

§ 227.46. Reinstatement of Testing Permit.

§ 227.48. Reporting Collisions.

§ 227.50. Reporting Disengagement of Autonomous Mode.

§ 227.52. Test Vehicle Registration and Certificates of Title.

§ 227.54. Transfers of Interest or Title for an Autonomous Test Vehicle.

California Department of Motor Vehicles (CA DMV)

International Standards

- International Electrotechnical Commission
- Functional **Safety** of Electrical/Electronic/Programmable Electronic Safety-related Systems (IEC 61508)

Likelihood of occurrence

Category	Definition	Range (failures per year)
Frequent	Many times in lifetime	$> 10^{-3}$
Probable	Several times in lifetime	10^{-3} to 10^{-4}
Occasional	Once in lifetime	10^{-4} to 10^{-5}
Remote	Unlikely in lifetime	10^{-5} to 10^{-6}
Improbable	Very unlikely to occur	10^{-6} to 10^{-7}
Incredible	Cannot believe that it could occur	$< 10^{-7}$

Consequences

Category	Definition
Catastrophic	Multiple loss of life
Critical	Loss of a single life
Marginal	Major injuries to one or more persons
Negligible	Minor injuries at worst

Risk Analysis

Risk class matrix

	Consequence			
Likelihood	Catastrophic	Critical	Marginal	Negligible
Frequent	I	I	I	II
Probable	I	I	II	III
Occasional	I	II	III	III
Remote	II	III	III	IV
Improbable	III	III	IV	IV
Incredible	IV	IV	IV	IV

- **Class I**: Unacceptable in any circumstance;
- **Class II**: Undesirable: tolerable only if risk reduction is impracticable or if the costs are grossly disproportionate to the improvement gained;
- **Class III**: Tolerable if the cost of risk reduction would exceed the improvement;
- **Class IV**: Acceptable as it stands, though it may need to be monitored.

Wiki on IEC 61508

International Standards

- International Organization for Standardization
- Road vehicles – Functional **safety** (ISO 26262)

ASIL = S x E x C

		C1	C2	C3
S1	**E1**	QM	QM	QM
S1	**E2**	QM	QM	QM
S1	**E3**	QM	QM	ASIL A
S1	**E4**	QM	ASIL A	ASIL B
S2	**E1**	QM	QM	QM
S2	**E2**	QM	QM	ASIL A
S2	**E3**	QM	ASIL A	ASIL B
S2	**E4**	ASIL A	ASIL B	ASIL C
S3	**E1**	QM	QM	ASIL A
S3	**E2**	QM	ASIL A	ASIL B
S3	**E3**	ASIL A	ASIL B	ASIL C
S3	**E4**	ASIL B	ASIL C	**ASIL D**

Autonomous Driving: **ASIL D** => acceptable probability of system / component failure of one in a hundred million

Severity Classifications (S):
- S0 No Injuries
- S1 Light to moderate injuries
- S2 Severe to life-threatening (survival probable) injuries
- S3 Life-threatening (survival uncertain) to fatal injuries

Exposure Classifications (E):
- E0 Incredibly unlikely
- E1 Very low probability (injury could happen only in rare operating conditions)
- E2 Low probability
- E3 Medium probability
- E4 High probability (injury could happen under most operating conditions)

Controllability Classifications (C):
- C0 Controllable in general
- C1 Simply controllable
- C2 Normally controllable (most drivers could act to prevent injury)
- C3 Difficult to control or uncontrollable

Safety integrity level (**SIL**)

SIL	Low demand mode: average probability of failure on demand	High demand or continuous mode: probability of dangerous failure per hour
1	$\geq 10^{-2}$ to $< 10^{-1}$	$\geq 10^{-6}$ to $< 10^{-5}$
2	$\geq 10^{-3}$ to $< 10^{-2}$	$\geq 10^{-7}$ to $< 10^{-6}$
3	$\geq 10^{-4}$ to $< 10^{-3}$	$\geq 10^{-8}$ to $< 10^{-7}$ (1 dangerous failure in 1140 years)
4	$\geq 10^{-5}$ to $< 10^{-4}$	$\geq 10^{-9}$ to $< 10^{-8}$

Automotive Safety integrity level (**ASIL**) vs SIL

Domain	Domain-Specific Safety Levels					
Automotive (ISO 26262)	QM	ASIL A	ASIL B	ASIL C	ASIL D	-
General (IEC 61508)	-	SIL-1	SIL-2		SIL-3	SIL-4

Wiki on IEC 61508 and ASIL (I)

All these regulations are about physical (onroad) metrics.

How to ensure the **safe & fast development** cycle?

Part II: Simulators

01 Self-driving policies as a reasonable gate?

02 Lack of simulators to trust

03 Lack of proxy metrics to trust

04 Autonomy stack: is it really a stack?

05 Academia and Industry

06 Conclusion

Simulators

Q: How to **safely test** the autonomous capabilities?

A: Using the **simulator**!

Main challenges:

- Sensors simulation
- **Behavior simulation**

CARLA simulator

+ NVIDIA DRIVE Sim, Deepdrive, LGSVL, SUMMIT, Flow, …
+ Internal and specific to any AV company simulators

Kaur, Prabhjot, et al. "A survey on simulators for testing self-driving cars." 2021

Simulators reliability

Reliability questions:

- How to guarantee the **generalization** of simulation results?
- Can we really rely on any **metrics inside** the simulation?

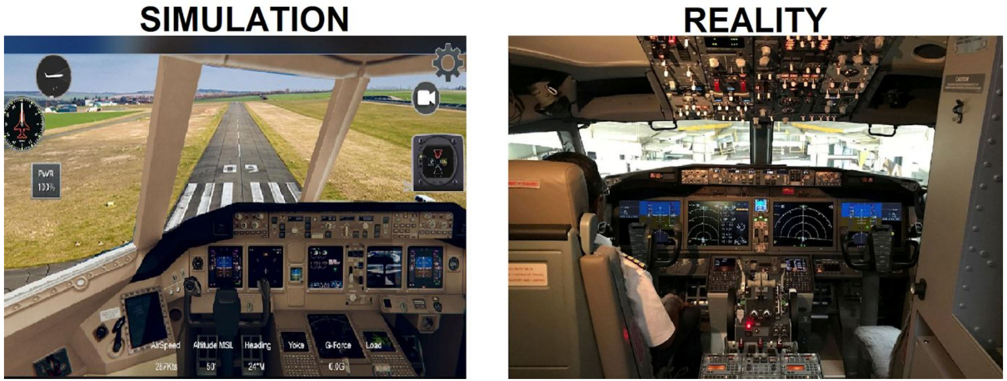

Medium.com: Simulation vs Reality in Marketing

Paperswithcode.com: Domain (distribution) shift

Towards Reliability

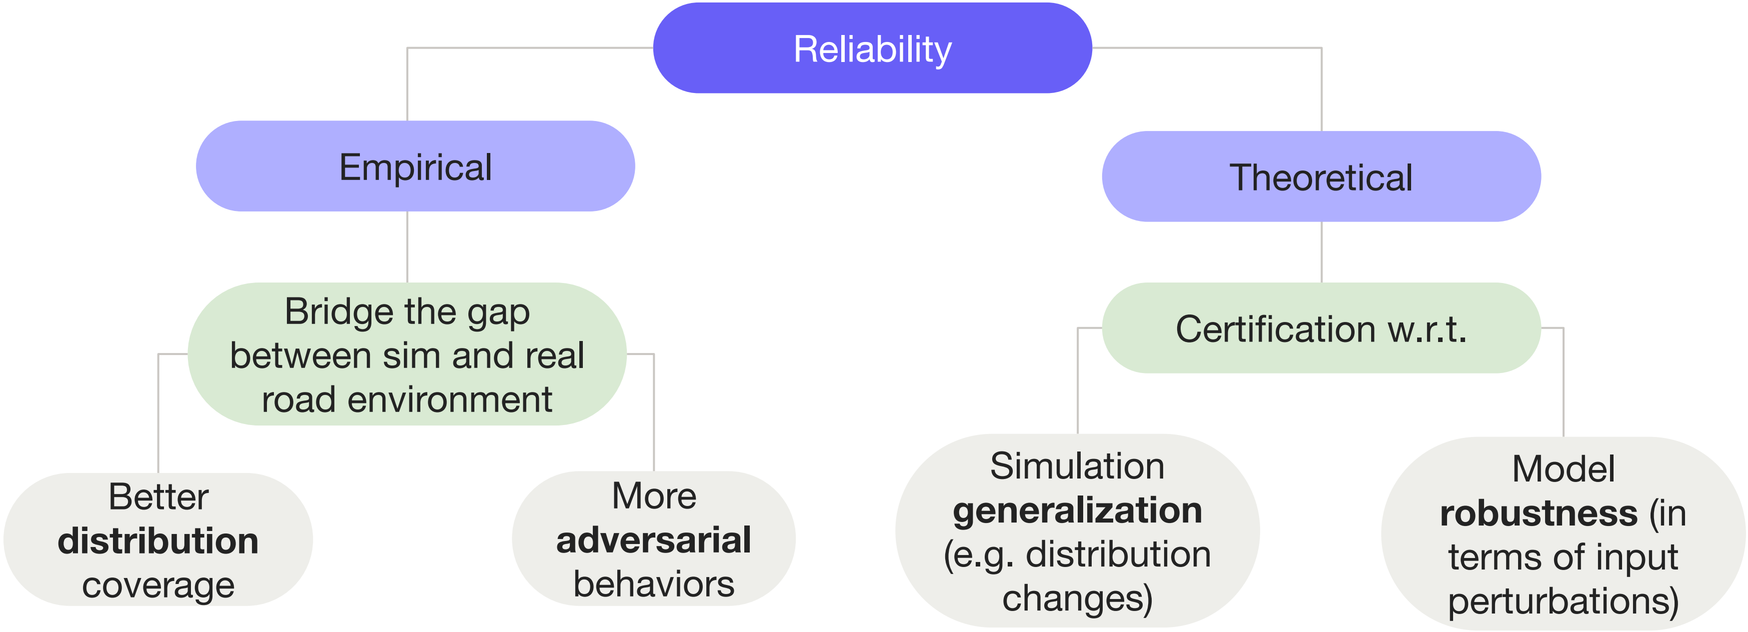

Towards Reliability

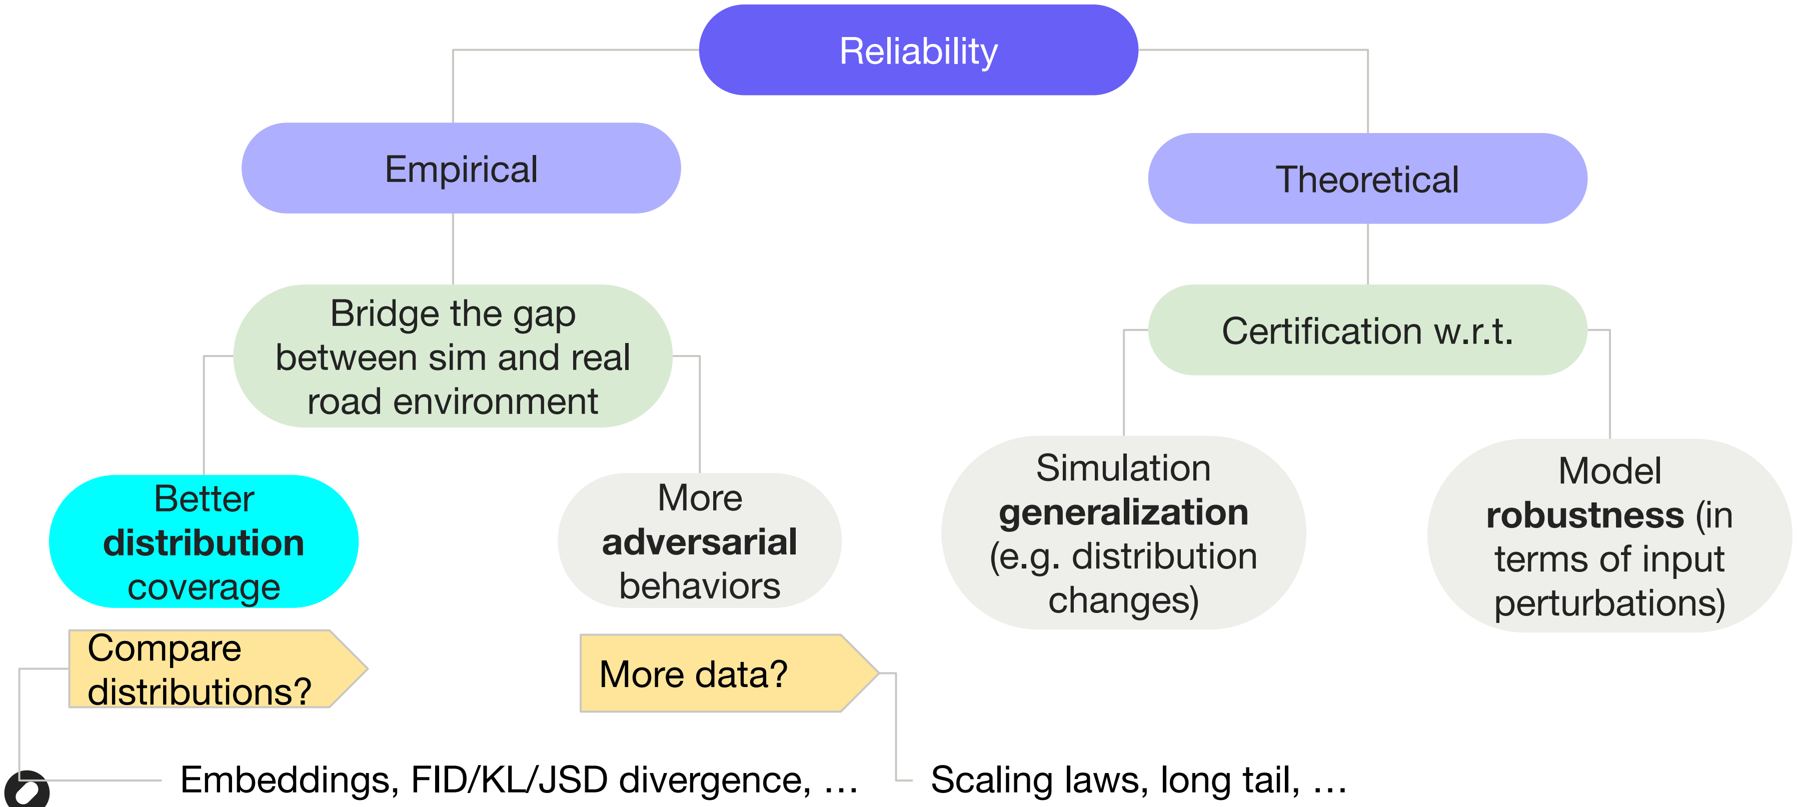

Towards Reliability

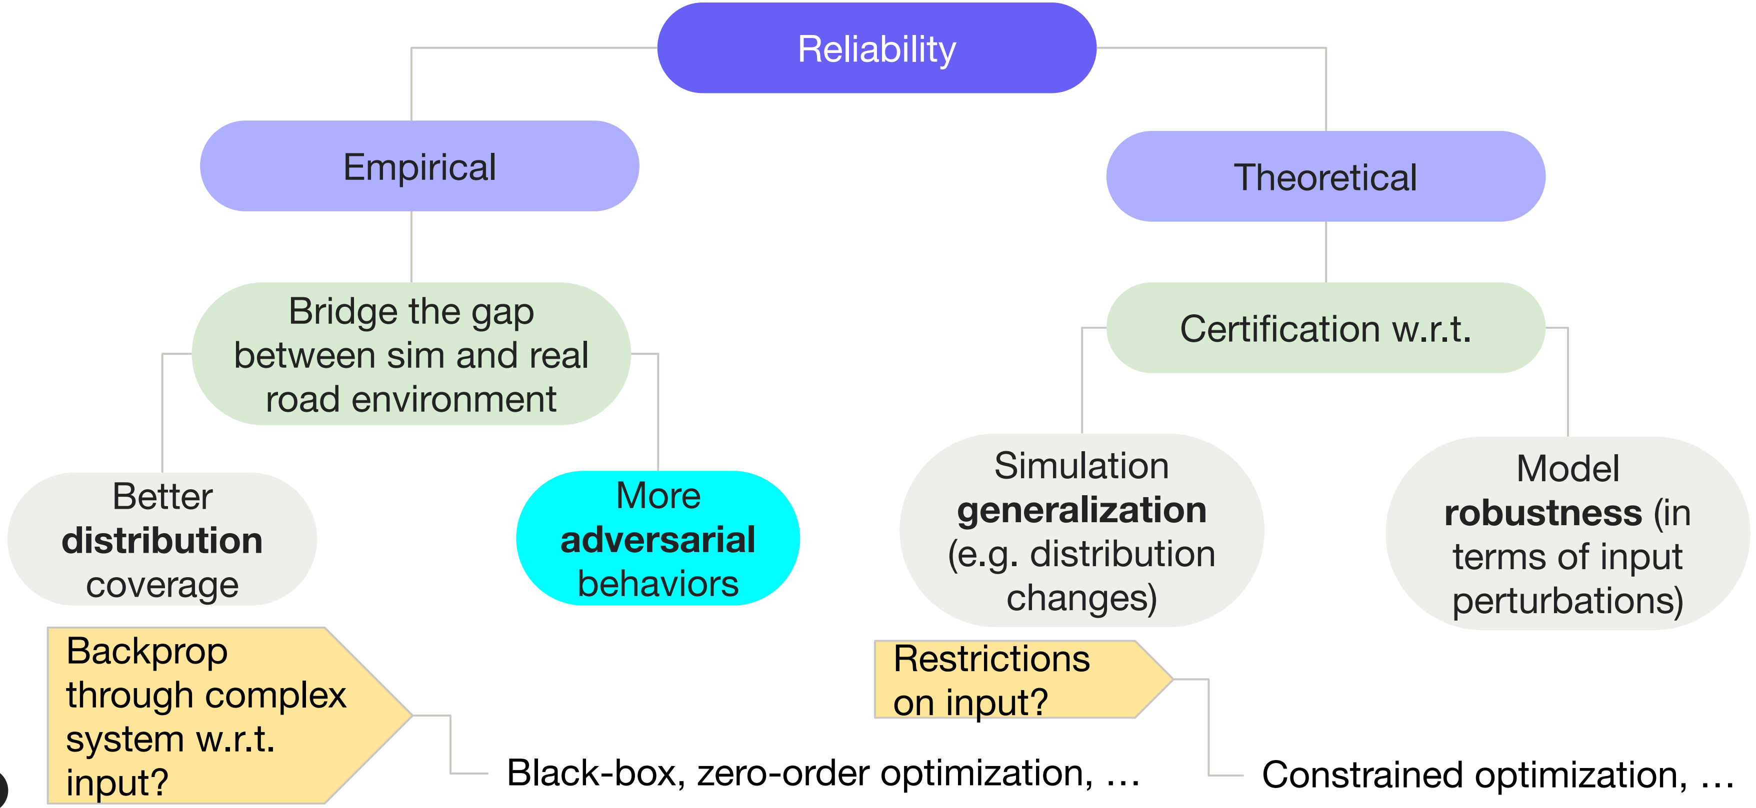

Towards Reliability

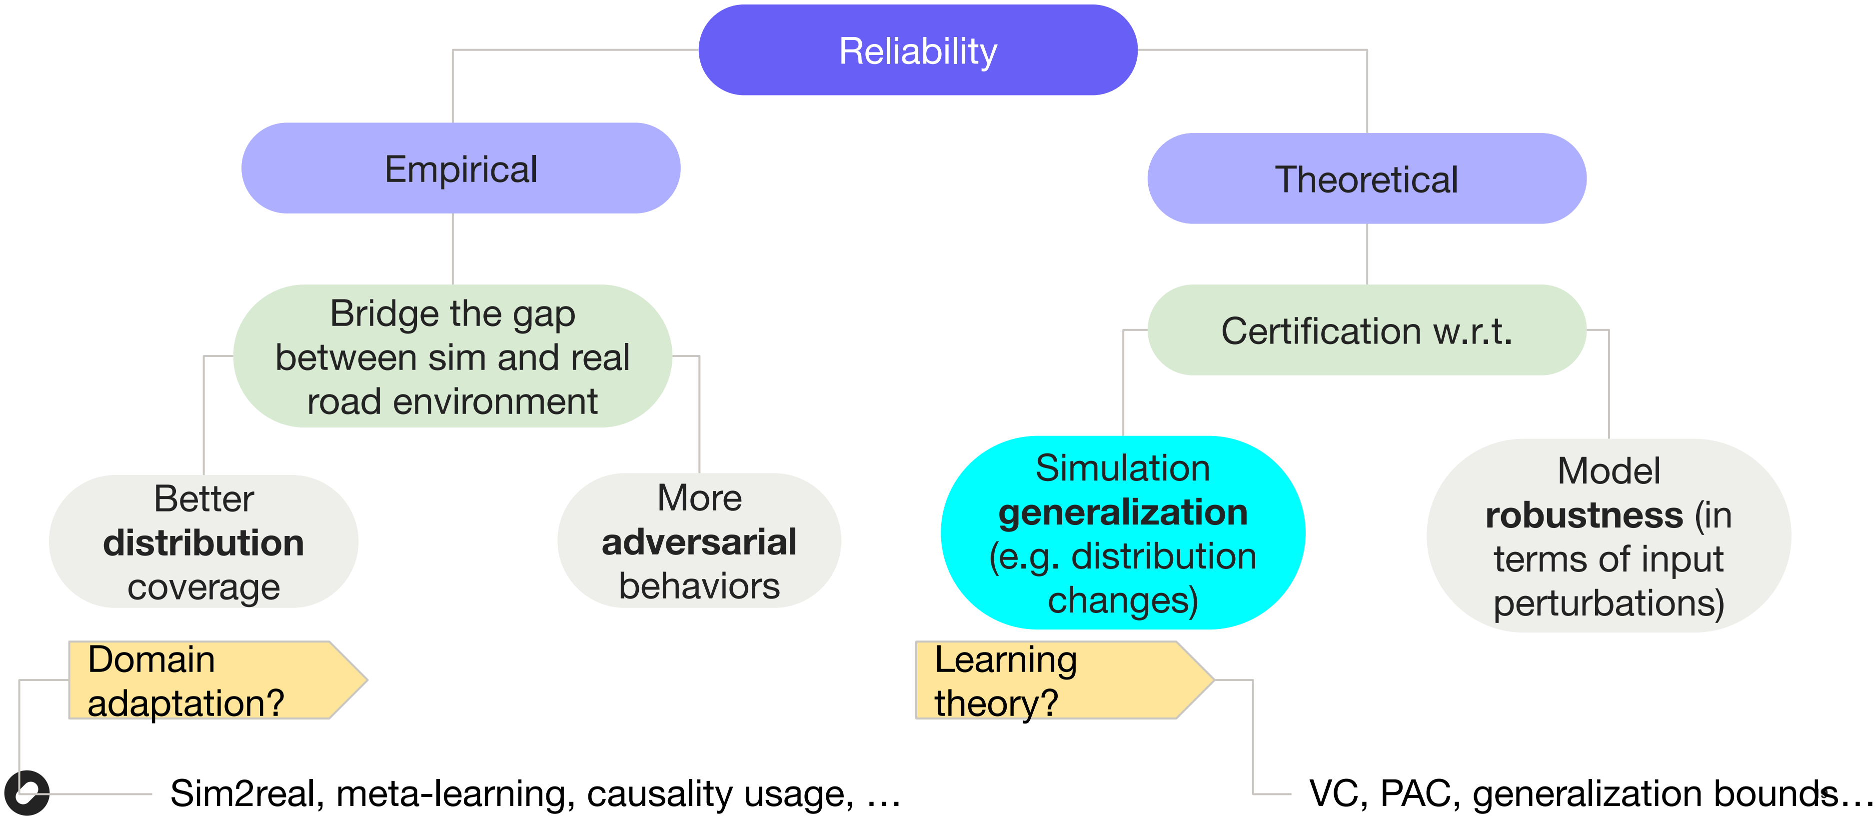

Towards Reliability

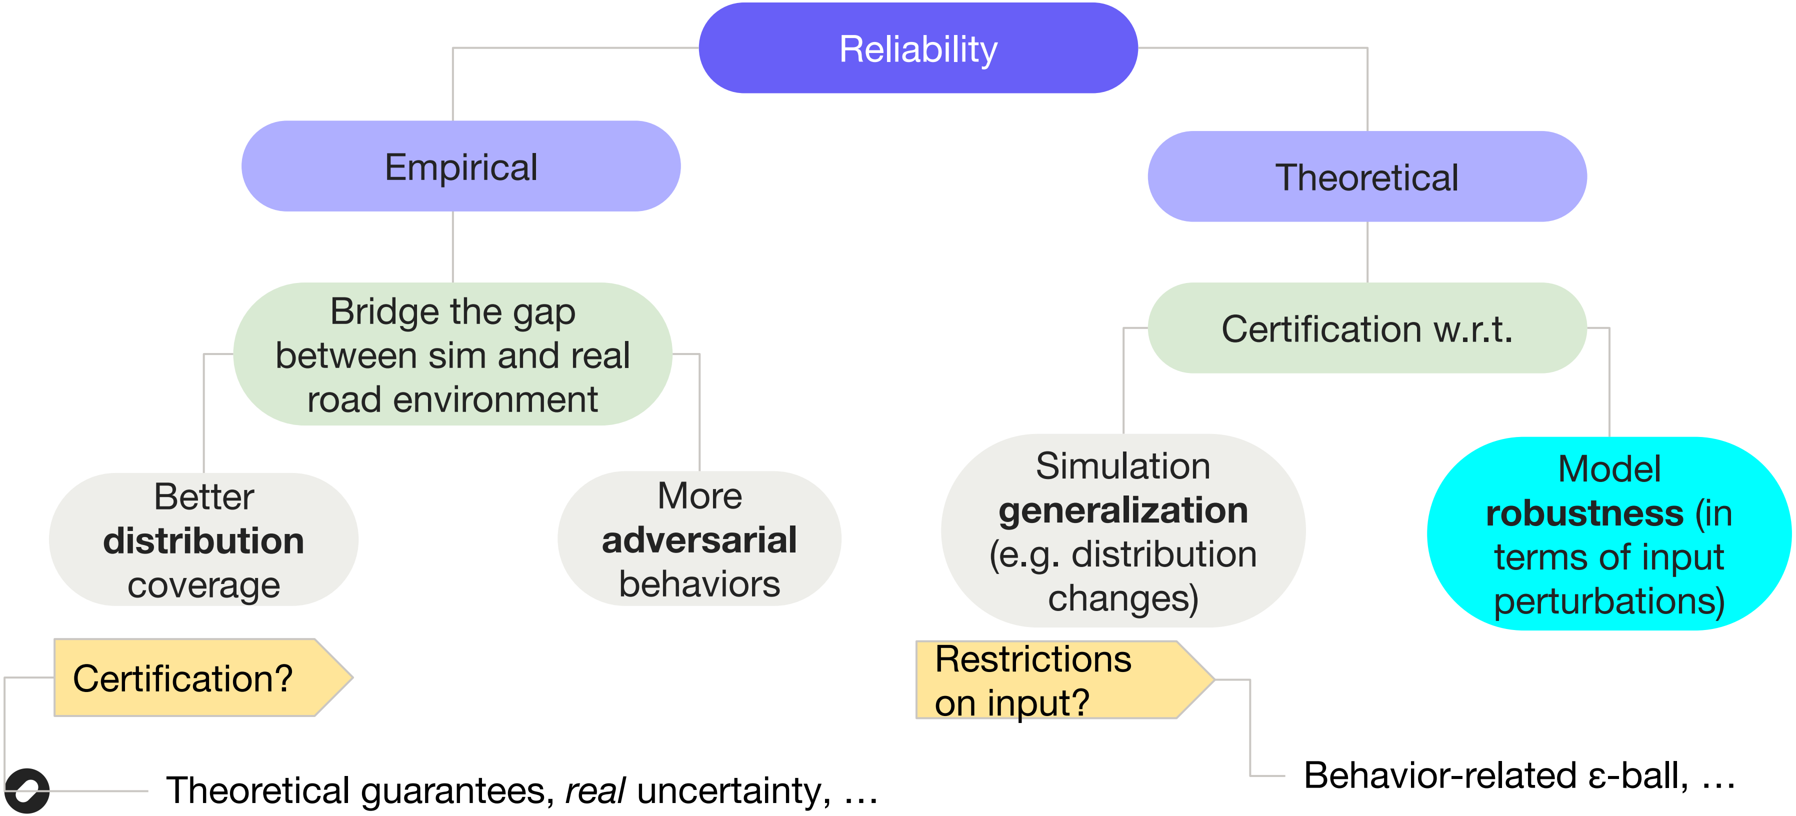

How to ensure the safe & fast development cycle?

Part III: Metrics

01 Self-driving policies as a reasonable gate?

02 Lack of simulators to trust

03 Lack of proxy metrics to trust

04 Autonomy stack: is it really a stack?

05 Academia and Industry

06 Conclusion

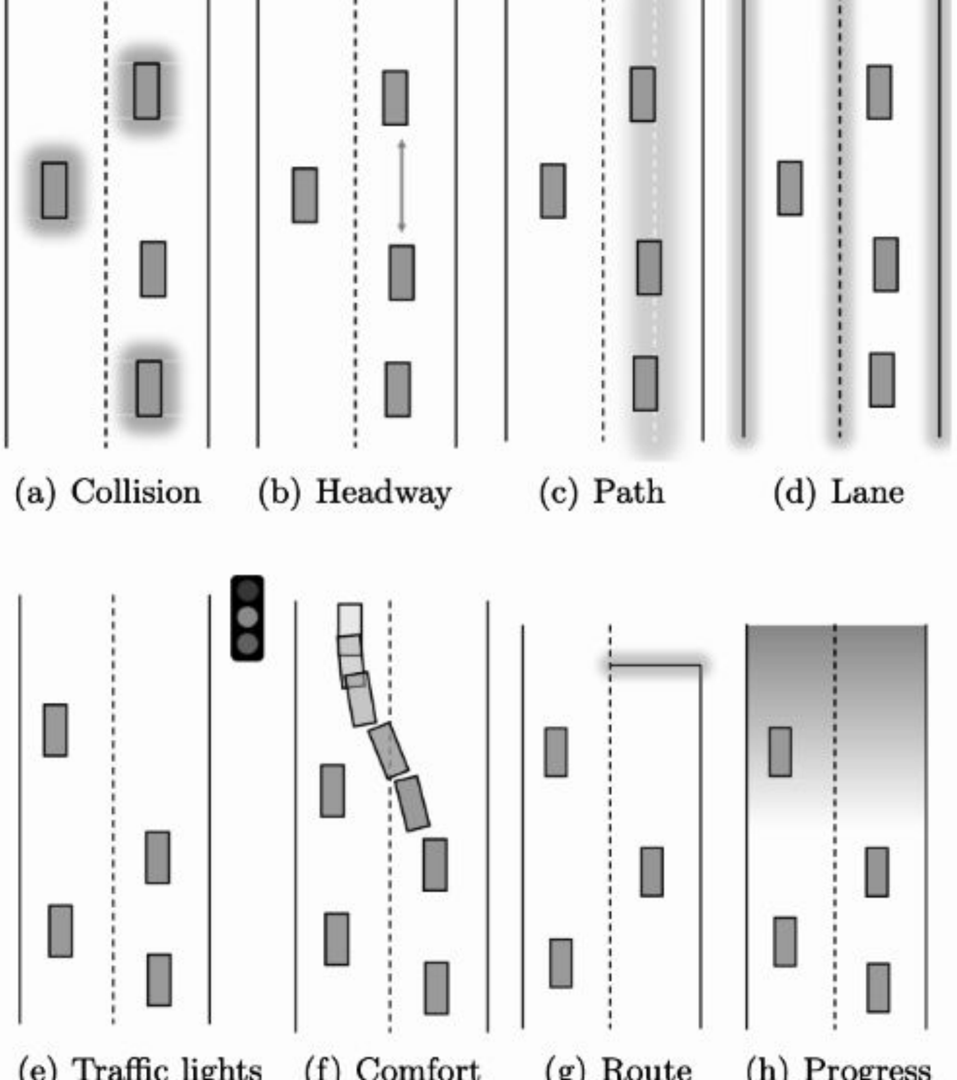

Metrics

Common metrics of AV:

- Miles per (critical) disengagement (**MPD**, **MPCD**)
- **Inverse**: number of disengagements per thousand of miles

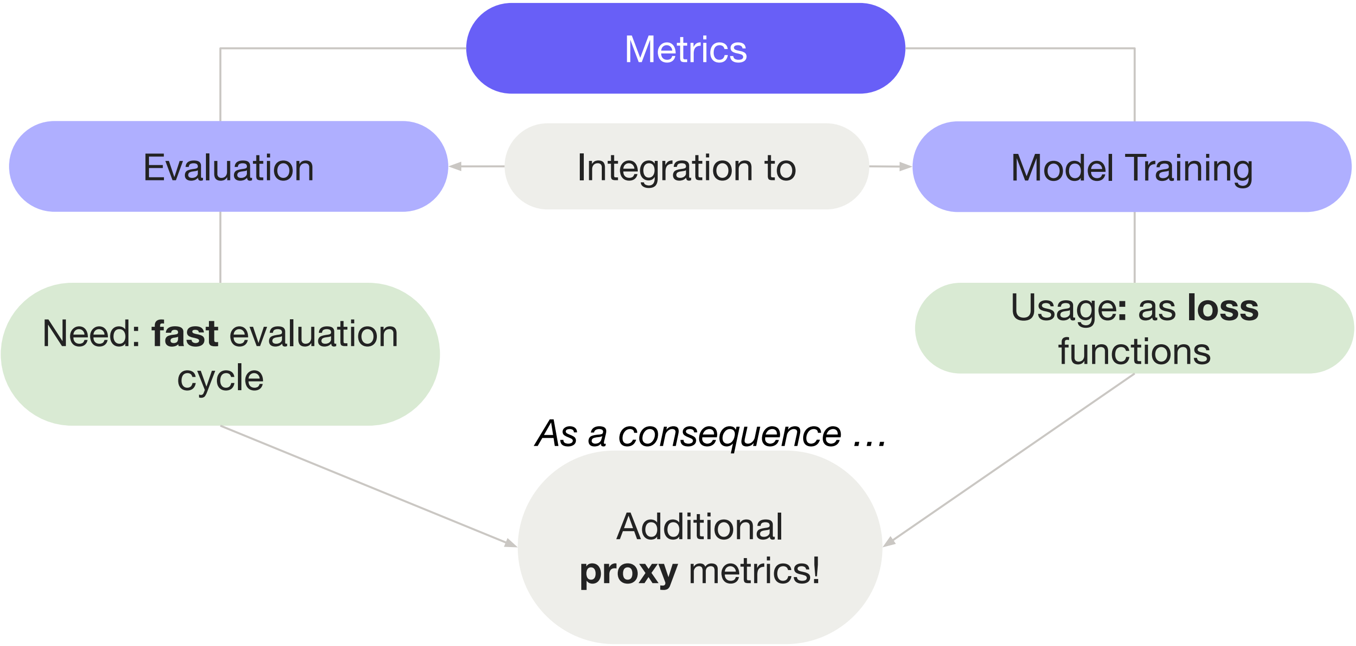

Metrics in the literature

<u>Proxy</u> metrics:

- Time to Collision
- Collision rate
- Off-road rate
- Off-route rate
- L2-based
- Comfort-based
    - Jerk
    - Lateral acceleration
- …

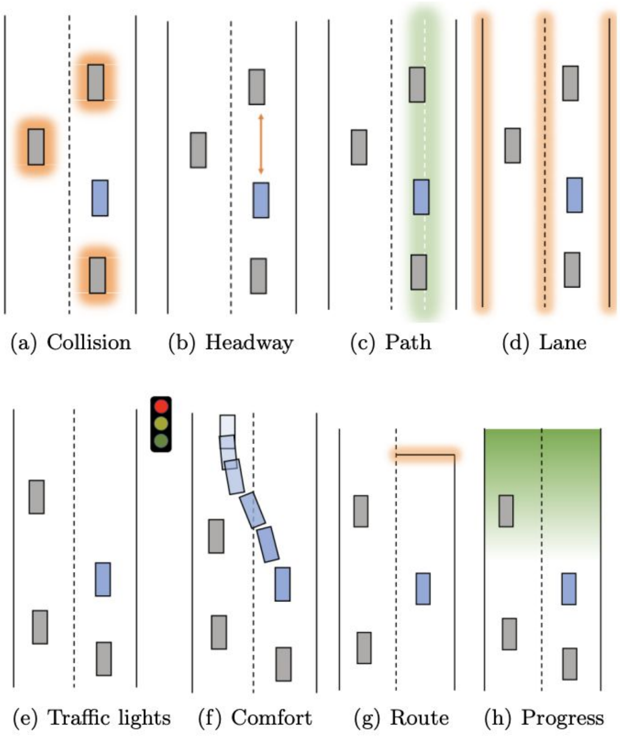

<u>Metrics</u>:

- **Open**-loop vs **Closed**-loop
    - L2-distance is not very important for closed-loop eval
- **Eval**-only vs **Train**+eval
    - The earlier to get the signal for the model, the better
- **Correlation** of MPCD/Disengagements with proxy metrics?
    - What are just regularization metrics for better train / faster eval?

Sadat, Abbas, et al. "Perceive, predict, and plan: Safe motion planning through interpretable semantic representations." 2020.

Do we really need to stick to the classical Autonomy **Stack**?

Part IV: Stack

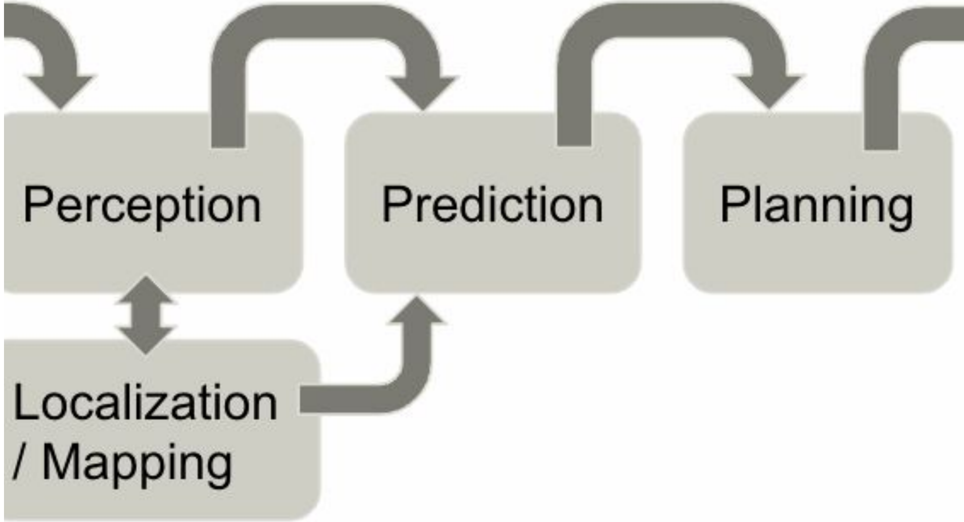

Stack

Classical **modular** structure

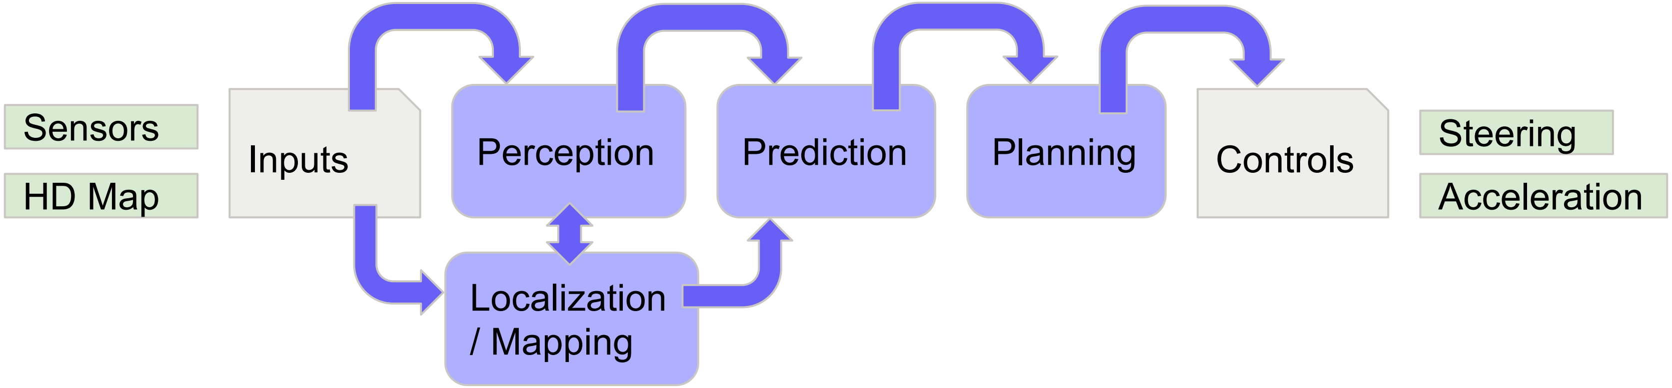

Each module:

- Has its **own** training / validation **data**
- Can be developed **independently**

Stack: unification?

Modular system being very useful still has **cons**:

- **Sub-optimal** optimization and performance
- **Hard to propagate** uncertainty estimations

Would be **helpful**:

- To **propagate** the learning **signal** through the **whole** stack
- (Probably) **not to do end2end** approach like *Behavior Cloning* (or even *Imitation Learning*)

Is it **real**?

- The “**Theorem of existence**” provides the way to incorporate the non-differentiable modules into the pipeline
  - Although done for some narrow class of tasks

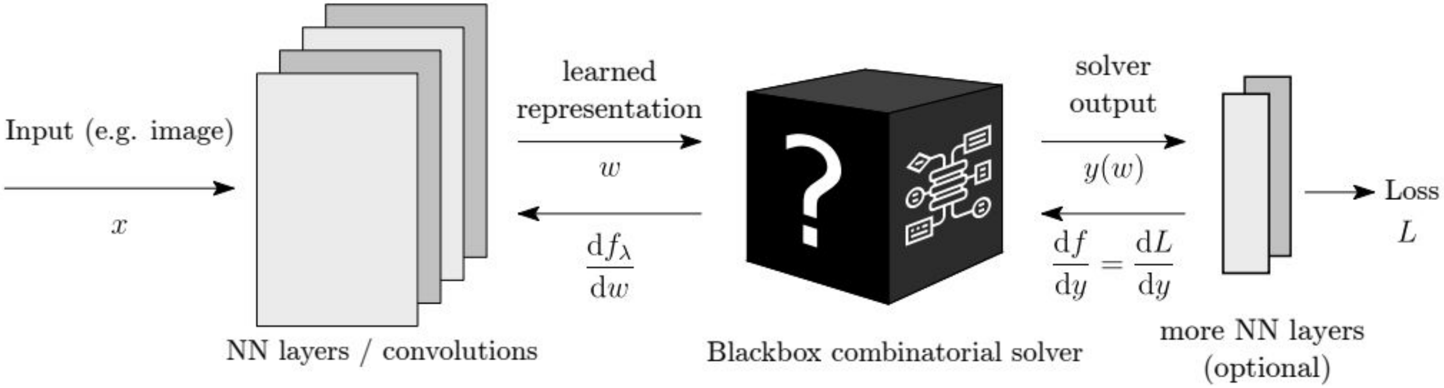

Vlastelica, Marin, et al. "[Differentiation of blackbox combinatorial solvers](#)." 2019

Stack: unification I

Combine: **Perception** + **Prediction**

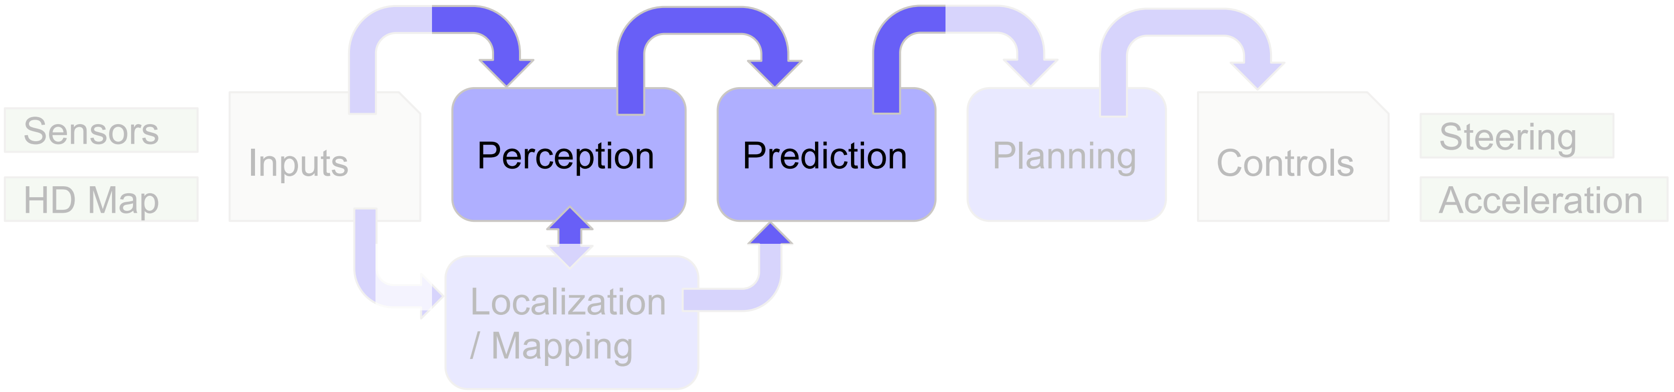

Luo, Wenjie, et al. "[Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net]." 2018

Stack: unification II

Combine: **Prediction + Planning**

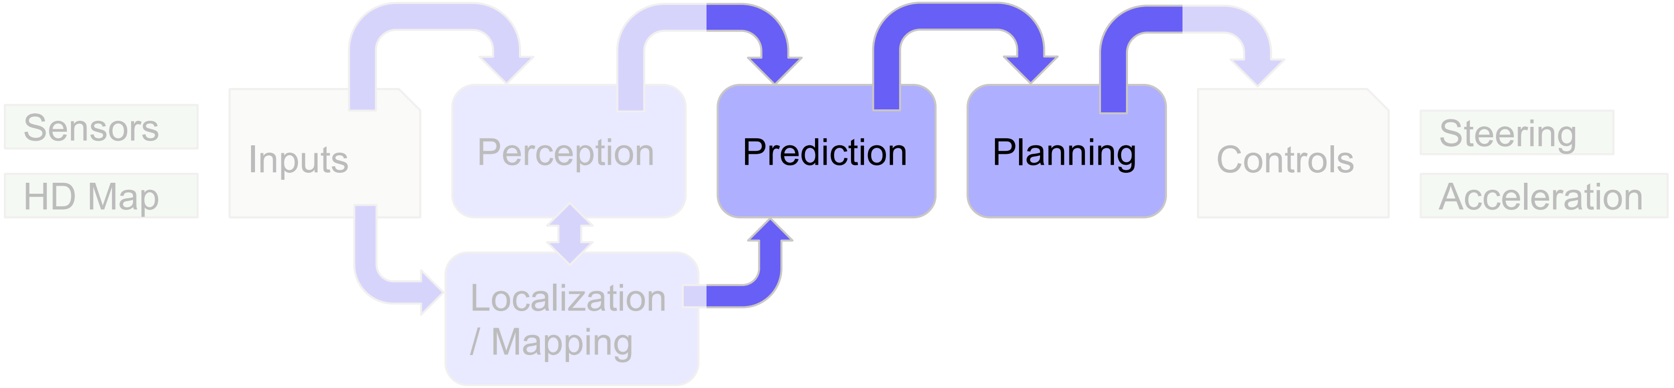

Liu, Jerry, et al. "Deep structured reactive planning." 2021.

Stack: unification III

Combine: **Perception** + **Prediction** + **Planning**

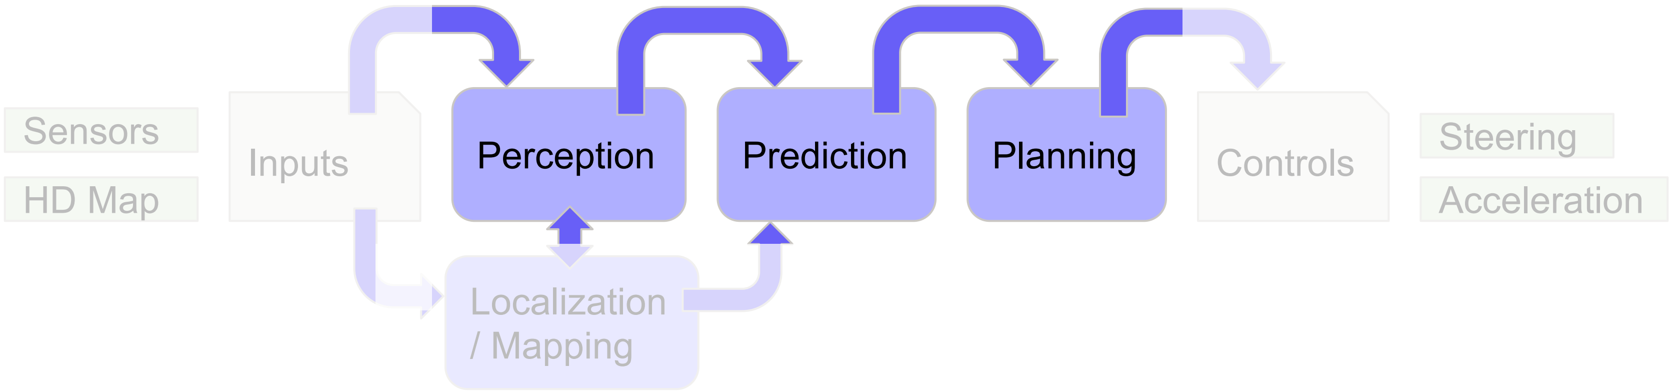

Sadat, Abbas, et al. "Perceive, predict, and plan: Safe motion planning through interpretable semantic representations." 2020.

Stack: unification IV

Combine: **Mapping** + **Perception** + **Prediction** + **Planning**

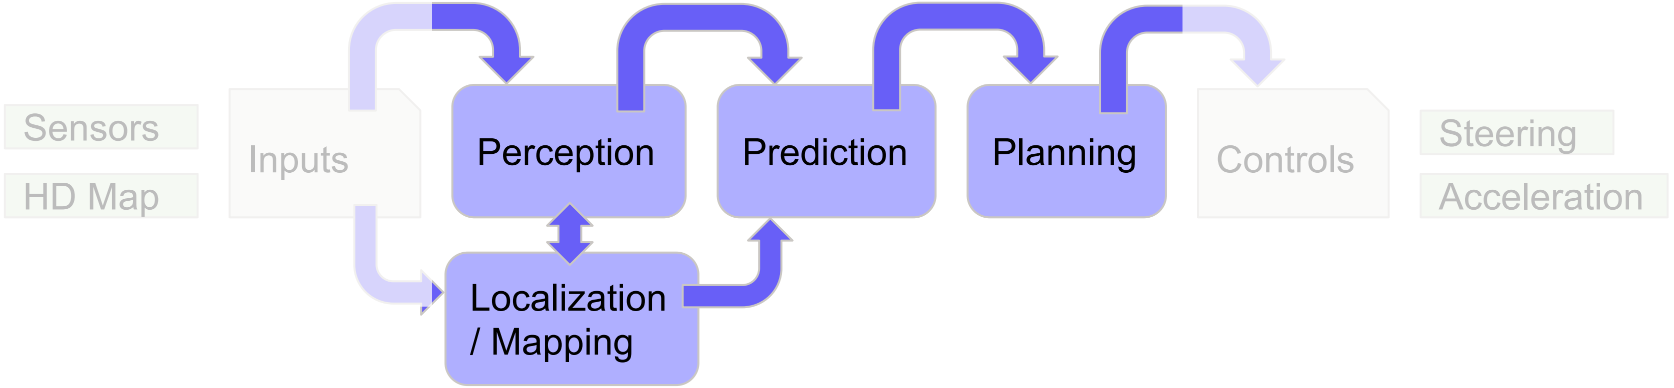

Casas, Sergio, et al. "Mp3: A unified model to map, perceive, predict and plan." 2021.

Stack and RL

Reinforcement Learning can be added for some of the modules combination

- Naturally integrates **planning**
- **State defines** the amount of input information (and the combination of modules as well)

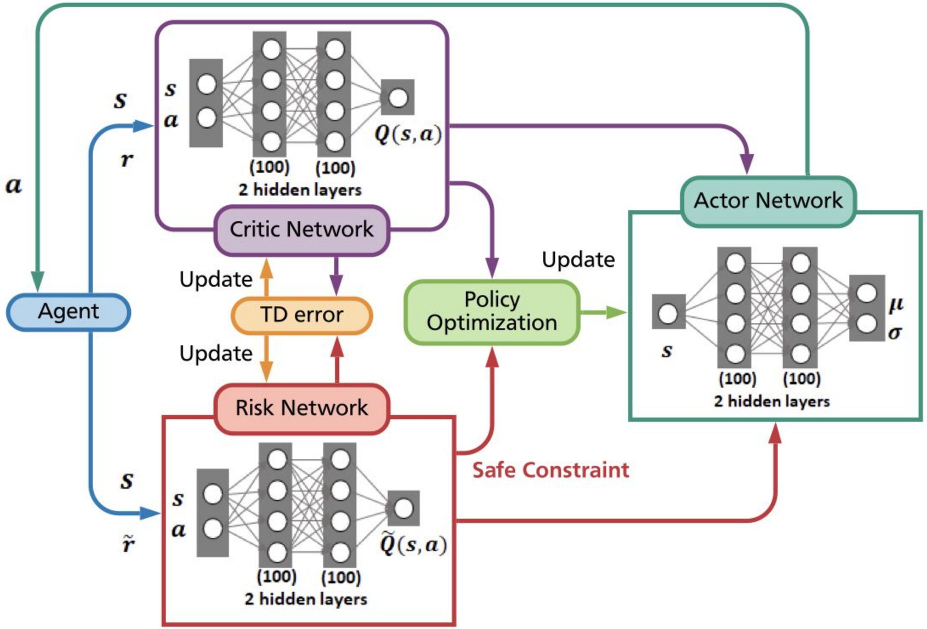

Wen, Lu, et al. "[Safe reinforcement learning for autonomous vehicles through parallel constrained policy optimization](#)." 2020.

Part V: Collaboration

AI: Academia + Industry

Academia I

To **provide**:

- **Closed-loop** simulation benchmark with clear API specifications to be used by any commercial company
  - Example: NIST FRVT and its API
    - First workshop on metrics by NIST

To **research**:

- **Different modes** of agents behavior
  - Average (max coverage) and critical (adversarial) types
  - **Integrate** it into closed-loop simulation (above) as separate tracks

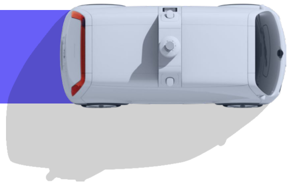

Academia II

To **investigate**:

- **How far** are the current simulators from **real life**?
  - Distribution shift estimation

To **prove**:

- What is the **influence** of this sim2real **distribution shift** analytically?
  - Sort of certification of numbers inside *specific* closed-loop simulation

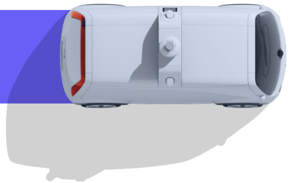

BDD as an example I

Multiple things have been **done…**

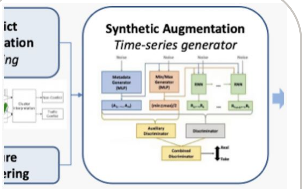

Analysis and Synthesis of Road Vehicle Conflicts using Vehicle Trajectory Data

BDD Project

Adversarial

Offline Reinforcement Learning for Data-Driven Autonomous Driving

BDD Project

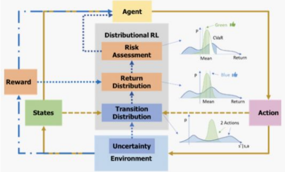

Figure 1. Distribution Reinforcement Learning (DRL) – Uncertainty Aware

Uncertainty-Aware Reinforcement Learning for Interaction-Intensive Driving Tasks

BDD Project

RL

Theoretical and Empirical Investigations of the Trade-offs between Open-loop and Closed-loop Planning

BDD Project

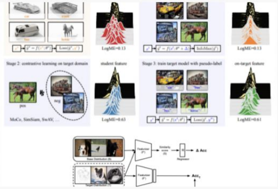

Simultaneous On-Target Domain Adaptation and Performance Prediction

BDD Project

Domain Adaptation

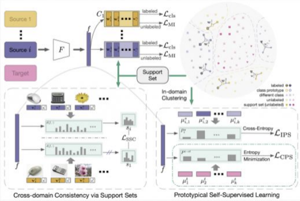

Cross-Domain Self-Supervised Learning for Adaptation and Generalization

BDD Project

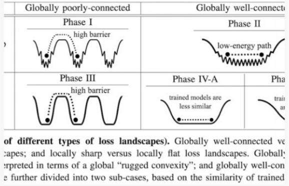

Measuring Prediction Trustworthiness and Safety Through Neural Network Loss Landscape Analysis

BDD Project

Generalization and Robustness

*Berkeley DeepDrive projects

BDD as an example II

Multiple things have been **done…**

- And much more ***still*** to do!

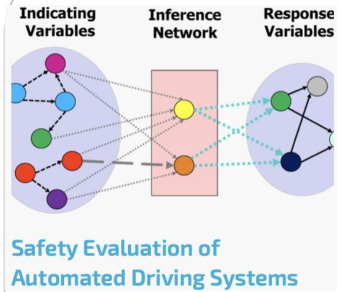

BDD Project

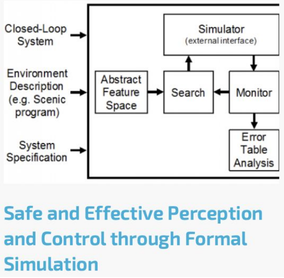

BDD Project

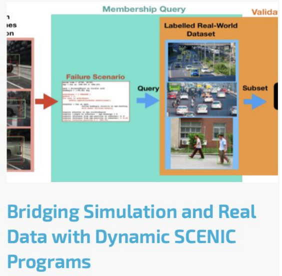

BDD Project

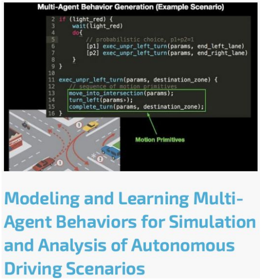

BDD Project

Metrics, Generalization and Verification

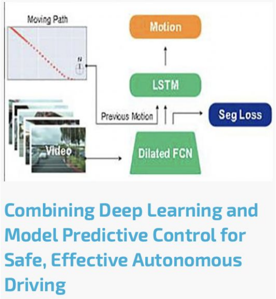

BDD Project

Stack

*Berkeley DeepDrive projects

Industry

To **propose**:

- Set of metrics to be used for evaluation

To **provide**:

- The training data
    - Not the evaluation!
    - Example: Waymo Open dataset

To **research**:

- The better technological stack / approaches for AV driving

AI: Academia + Industry

To **agree**:

- On **metrics** to **rely** on
  - Correlation with MPCD / Disengagements / others?

To **collaborate**:

- On **research** topics
- On **(sub)tasks** descriptions
- …

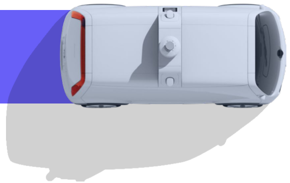

Part VI: Conclusion

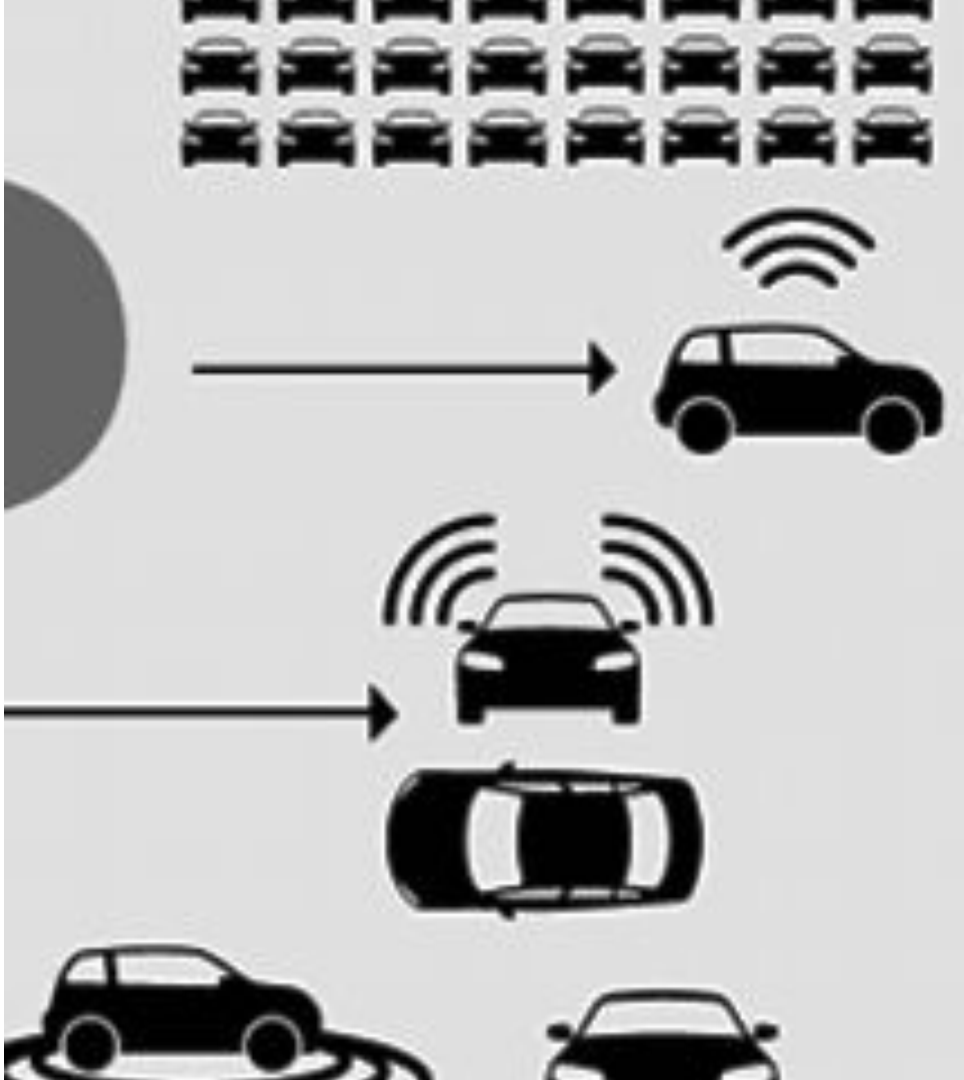

Conclusion

- Hard to use common AV metrics for research
- Current closed-loop evaluation is still imperfect
- Need to understand what are discrepancies w.r.t. the real environments (distribution shift) and how to certify the current results (analytical guarantee)
- Eventually the technological approach can be much (or even completely) different from the classical one
- Academia+Industry: synergy is inevitable

Bright Future

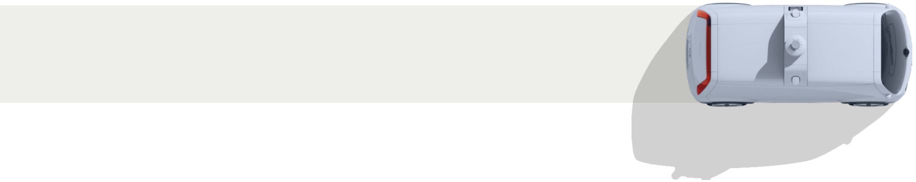

Great **change** of paradigm:

1. Be **as a human driver**:
   - **N** years?
2. Be **much better** as a human driver:
   - Is it really a jump of N→**NN** years?

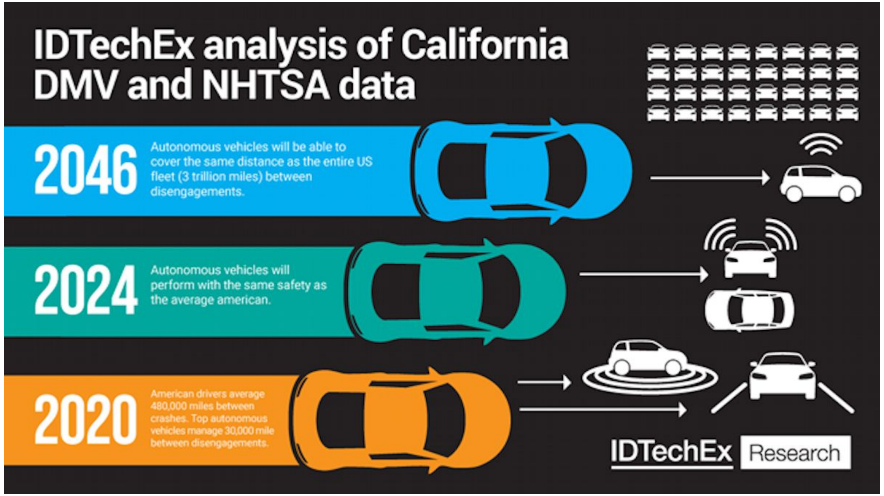

Source: IDTechEx


Thank
You.
nuro