Watch Out for the Safety-Threatening Actors: Proactively Mitigating Safety Hazards

Saurabh Jha, Shengkun Cui, Zbigniew Kalbarczyk, Ravi Iyer

Is Autonomous Driving Safe Enough? [2018 – 2023]

03/2018

Self-Driving Uber Car Kills Pedestrian in Arizona, Where Robots Roam

02/2020

Apple Engineer Killed in Tesla Crash Had Previously Complained About Autopilot
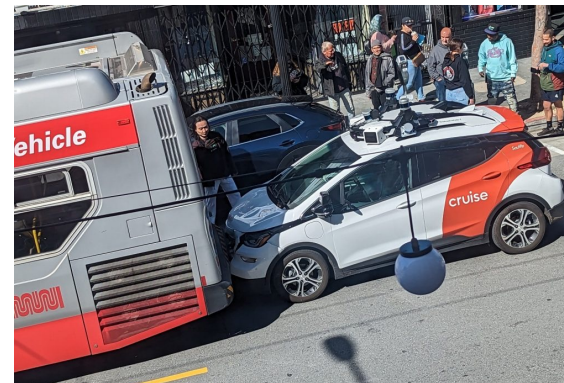
By Tom Krisher and Olga Rodriguez
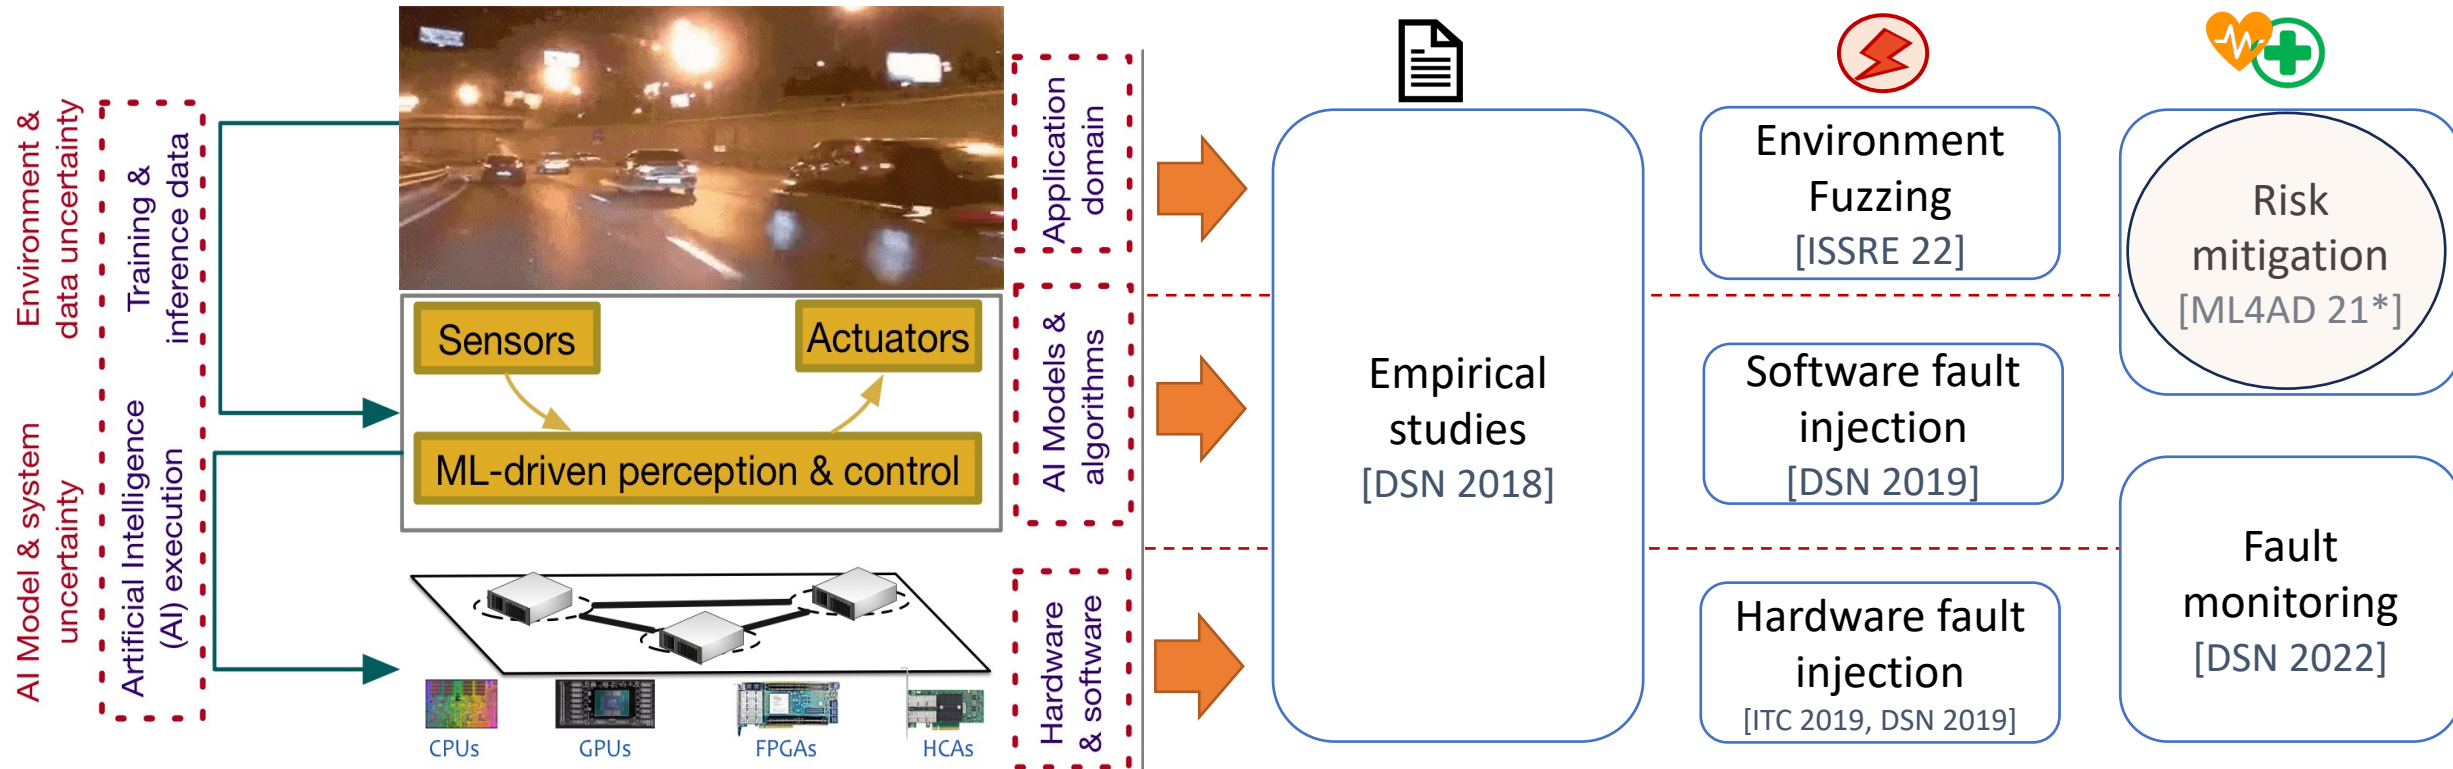The Associated Press — Feb 11, 2020 — Save Article

10/2023

Cruise Stops All Driverless Taxi Operations in the United States

The move comes just two days after California regulators told the company to take its autonomously driven cars off the road.

AV Research Overview

Is Autonomous Driving Safe Enough? [DSN 2018]

SAMPLE OF DISENGAGEMENT REPORTS FROM THE CA DMV DATASET.

Manufacturer	Raw Disengagement Report (Log)	Category	Tags
Nissan	1/4/16 — 1:25 PM — **Software module froze**. As a result driver safely disengaged and resumed manual control. — City and highway — Sunny/Dry	System	Software
Nissan	5/25/16 — 11:20 AM — Leaf #1 (Alfa) — The AV **didn't see** the lead vehicle, driver safely disengaged and resumed manual control.	ML/Design	Recognition System
Waymo	May-16 — Highway — Safe Operation — Disengage for a **recklessly behaving** road user	ML/Design	Environment
Volkswagen	11/12/14 — 18:24:03 — Takeover-Request — **watchdog error**	System	Computer System

We use the "—" to denote field separators.
Note that log formats vary across manufacturers and time.
Bold-face text represents phrases analyzed by the NLP engine to categorize log lines.

- AVs 15-4000x worse than humans
- Failures equally attributed to hardware/software, **environment** and ML for Waymo
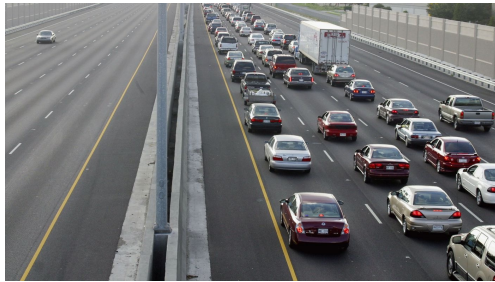
DSN 2018

How Do We Make Autonomous Driving Safer?


Why did the rear car brake?

How Do We Make Autonomous Driving Safer?

Attention required increases with the increase in uncertainty of another actor's behavior
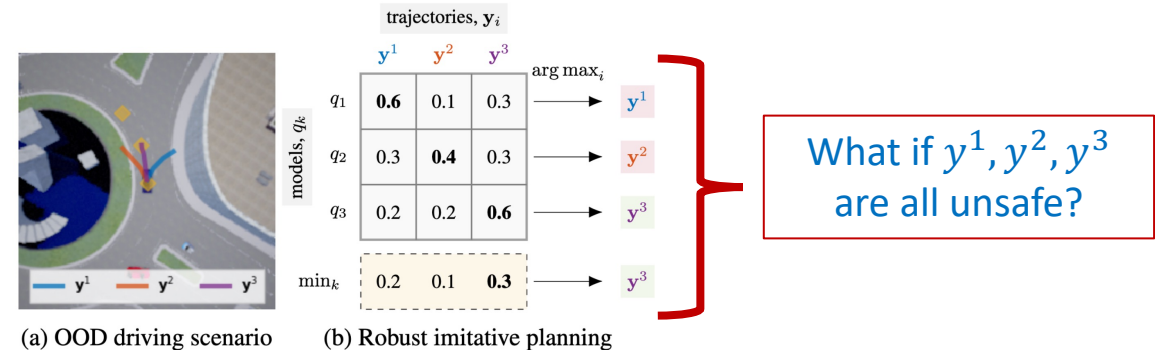
Ensuring Safety – Traditional Methods

- By avoiding collision trajectories
  - Time to collision
  - Intel Responsibility Sensitive Safety (RSS)
  - Nvidia Safety Force Field (SFF)

- **Does not proactively reduce risk**
  - **Often too late to avoid accident**

Suddenly accelerates
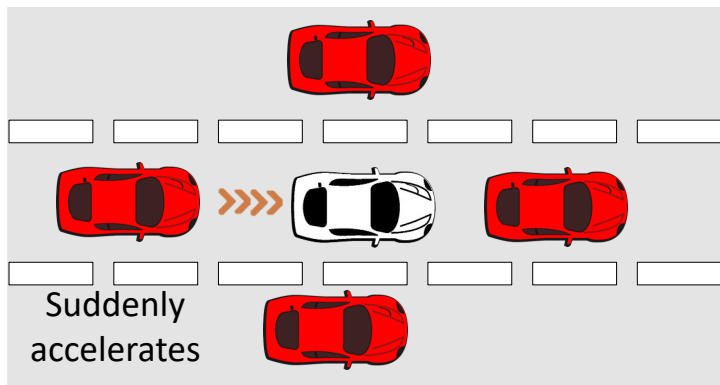
RSS/SFF cannot avoid accident!

- By learning from data
  - Reinforcement learning
  - Imitation learning
  - Adaptation to out-of-training-distribution

- **Depends on training data quality**
 - **Cannot handle rare driving scenarios**

trajectories, $\mathbf{y}_i$

	$\mathbf{y}^1$	$\mathbf{y}^2$	$\mathbf{y}^3$	$\arg\max_i$
q_1	**0.6**	0.1	0.3	$\mathbf{y}^1$
q_2	0.3	**0.4**	0.3	$\mathbf{y}^2$
q_3	0.2	0.2	**0.6**	$\mathbf{y}^3$
$\min_k$	0.2	0.1	**0.3**	$\mathbf{y}^3$

models, q_k

(a) OOD driving scenario   (b) Robust imitative planning

What if y^1, y^2, y^3 are all unsafe?

RIP agent (Filos et. al.) crashes under an OOD scenario in CARLA simulation

Ensuring Safety with Inter-actor Interactions

- By prioritizing detection and prediction accuracy for more influential actors (to planning)
  - Planner Objective Sensitivity (Ivanovic et al. 2022)
  - Planner KL-Divergence (Philion et al. 2020)

$$D_{KL}(P \| Q)$$

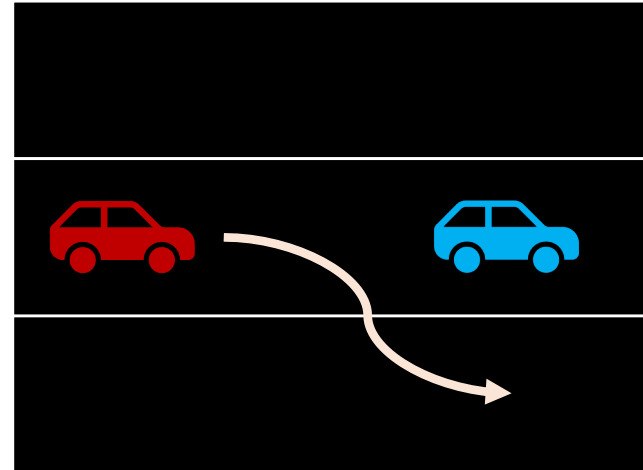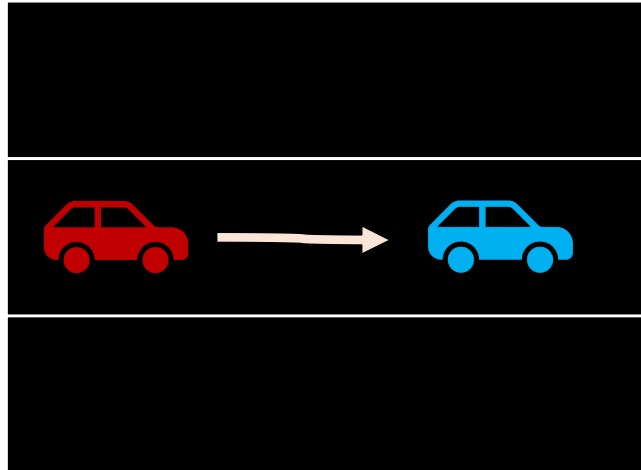Sensitive to the accuracy of p_θ -- the neural planner
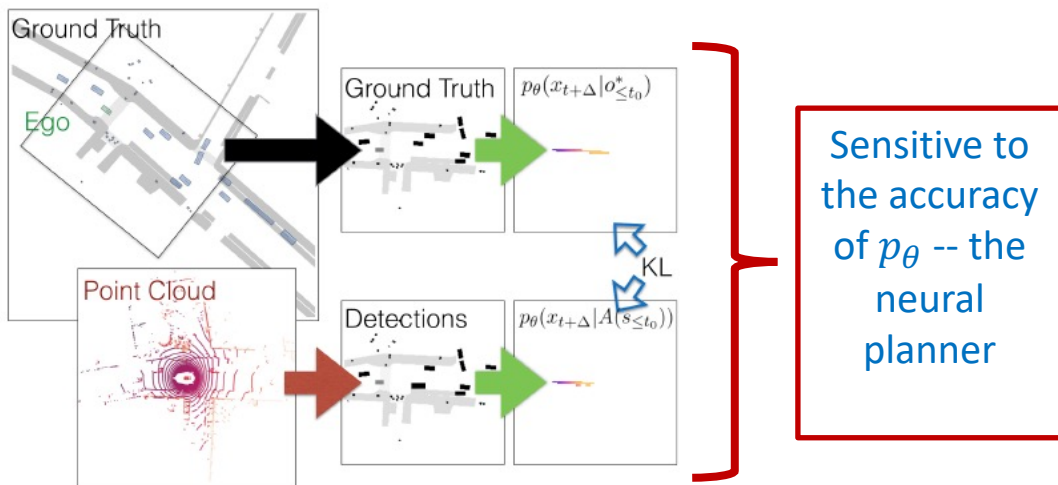
- By explicitly calculate "actor-actor interactivity" scores
  - Actor Interactivity Score (Tolstaya et al. 2021)

$$I(S^A, S^B) = \int_{\mathbf{s}^A} p(\mathbf{s}^A) D_{KL} \left[p(S^B|\mathbf{s}^A) \| p(S^B) \right]$$

What if the conditional predictions are inaccurate?
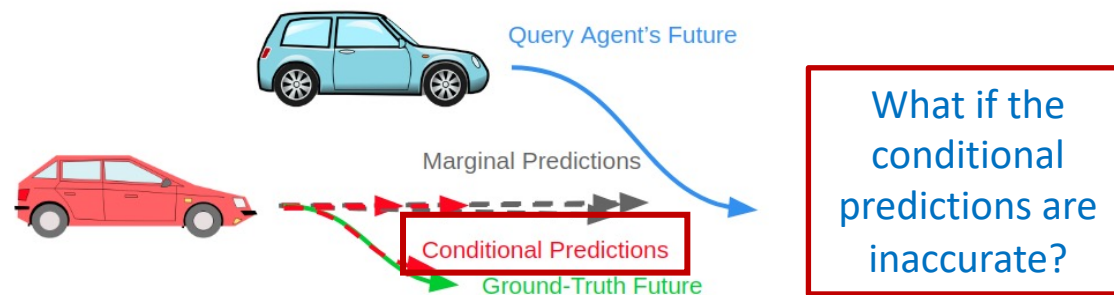
Counter Example Illustration

- Interaction between actors did not change
- $D_{KL} >> 0$

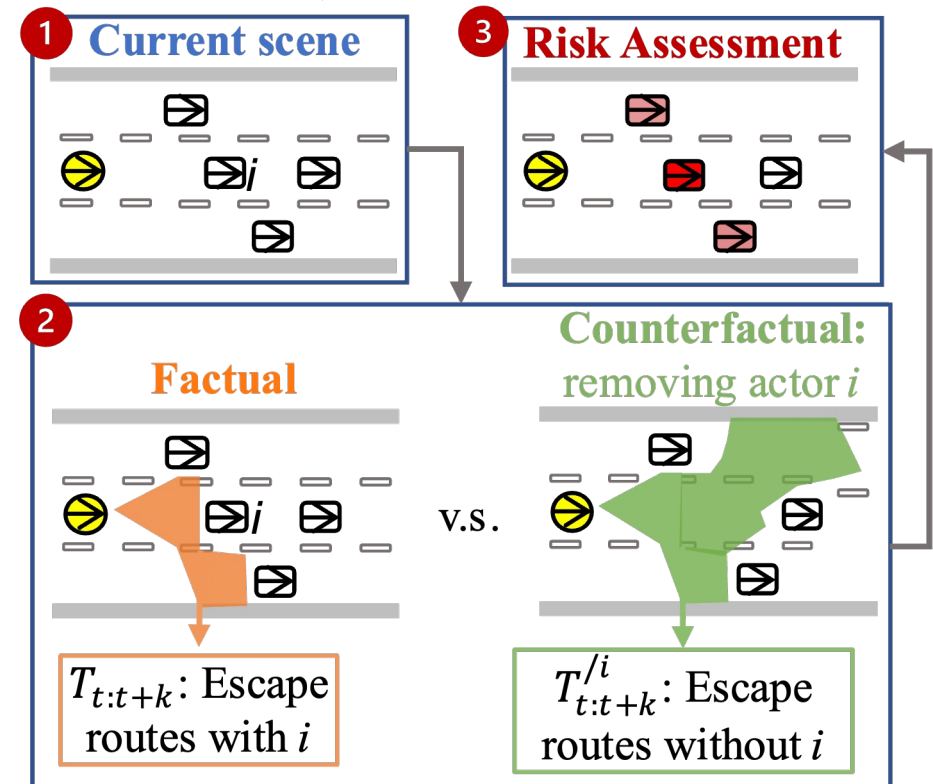# AD Safety & Risk Assessment

Human intuitions

1. Actively ensure "backup plans" (aka "escape routes")
2. Handle uncertainty

Research Question:
How do we design risk metric that embeds these intuitions!
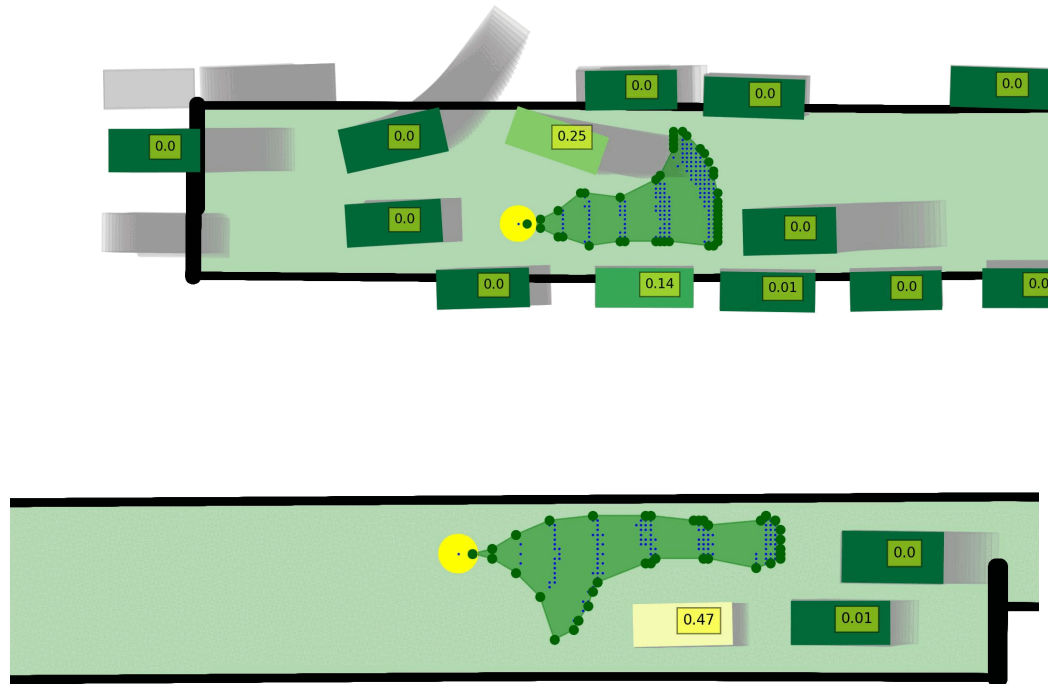
Analytical, no learning needed!

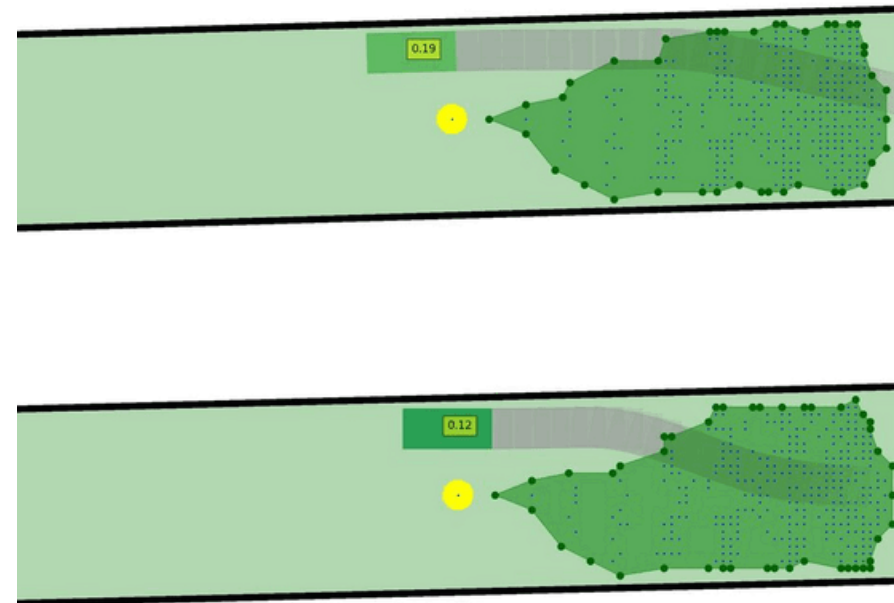$$\text{risk} \propto |T^{/i}_{t:t+k}| - |T_{t:t+k}|$$

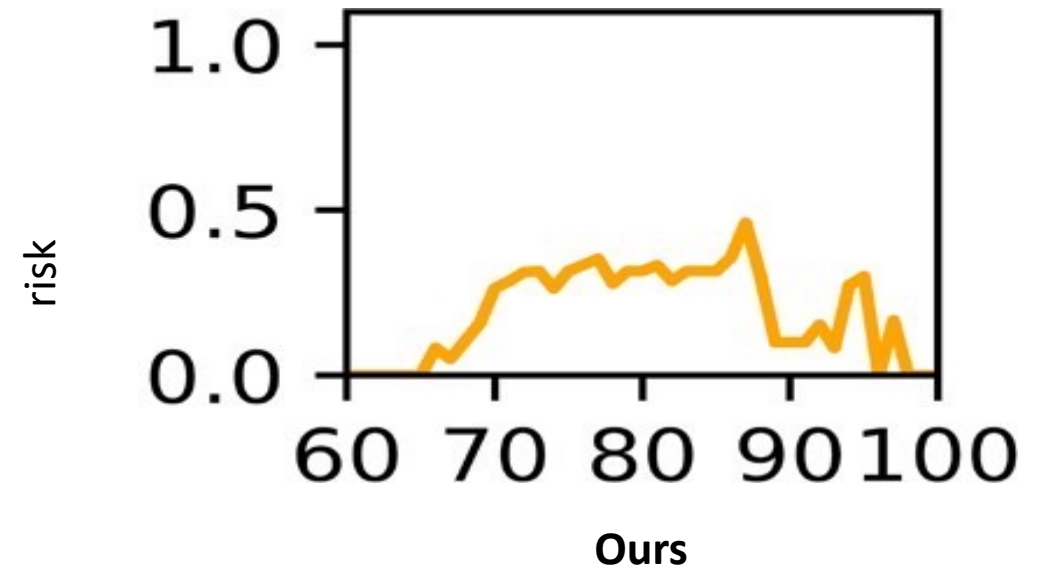Motivated from Barlow & Proschan work [1975]
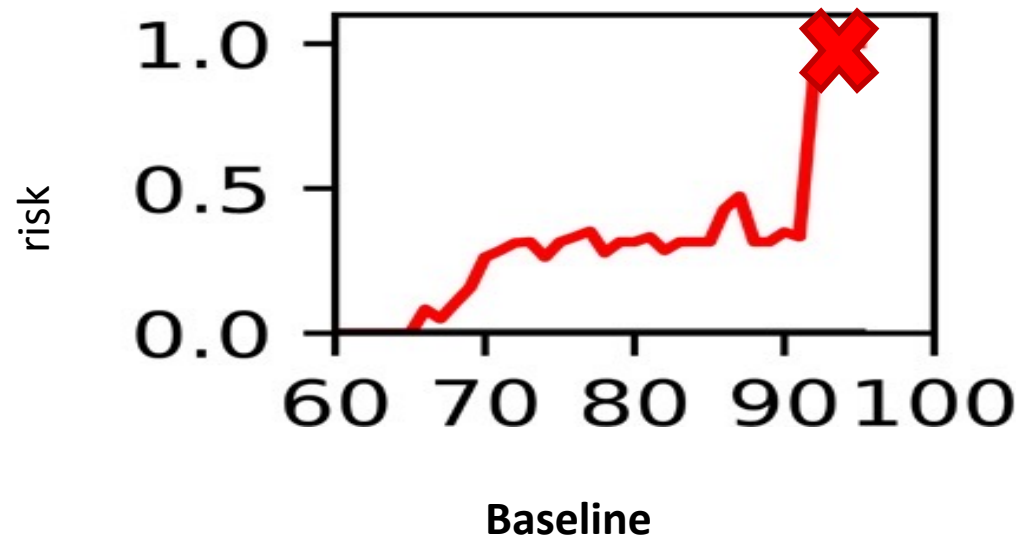
Risk Assessment in Action

Argoverse (Chang et al. 2019) Real-world Dataset

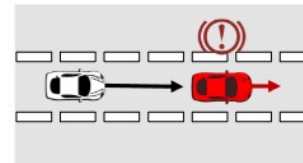CARLA Simulator with High-risk Scenarios

Our Solution: Risk-aware Safety Blanket
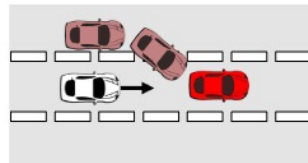
Proactive Reduce Risk for Mitigating Accidents

Baseline

Ours

Proactively avoids **trajectories of no return** by reducing risk!

Results
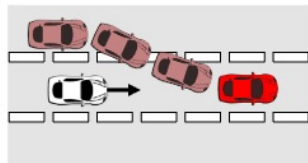
Agent	Ghost cut-in	Lead cut-in	Lead slowdown
LBC + Ours	**267**	**3**	**15**
LBC	519	170	118

Agent	Ghost cut-in	Lead cut-in	Lead slowdown
RIP + Ours	**65**	**265**	**129**
RIP	478	671	440

 # collisions in 1000 scenarios per typology (lower is better)

Significant reduction in accidents

Conclusion and Future Work

- Defining risk metric that captures escape path

- Future work
  - Scenario mining and assessment
  - Adversarial attack in high-risk scenarios
  - Integration with the planners

Driving Scenario(s) → **Risk analysis** → **Target scenes and actors** →

BFI: Smart Fault Injection (DSN 2019) — offline — What faults to inject into the AV system?

RoboTack: Smart Malware (DSN 2020) — online — What, when, and how to activate an attack vector?
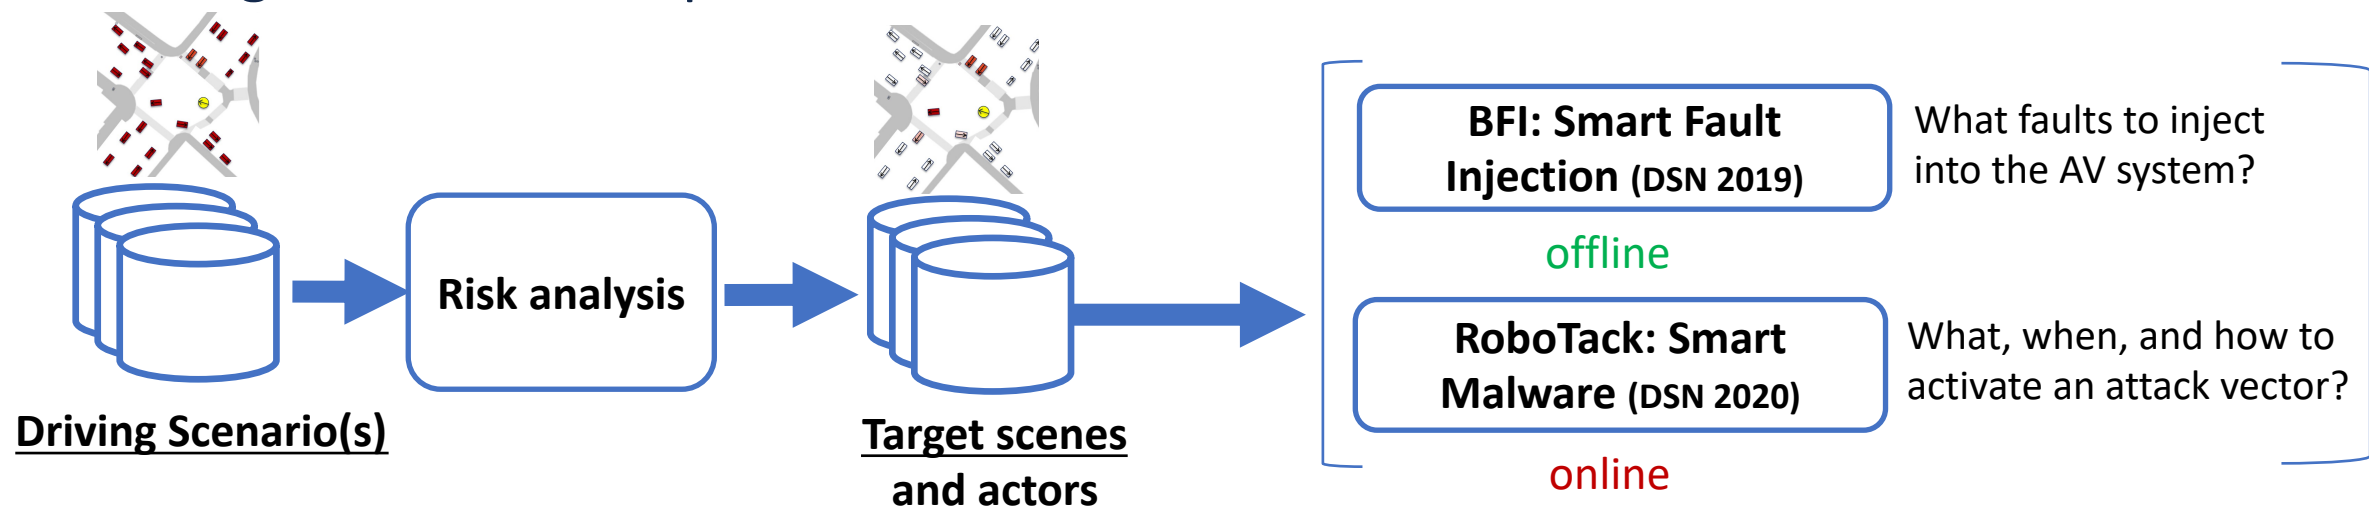
IBM in Autonomous Vehicle

RED HAT BLOG

The new standard: Red Hat In-Vehicle Operating System in modern and future vehicles

May 10, 2022 | Francis Chow

Submit a proposal to speak at Safety-Critical Software Summit at EOSS, Seattle, April 16-18 LEARN MORE >>

Advancing Open Source Safety-Critical Systems

The Enabling Linux In Safety Applications (ELISA) project aims to make it easier for companies to build and certify Linux-based safety-critical applications – systems whose failure could result in loss of human life, significant property damage or environmental damage. ELISA members are working together to define and maintain a common set of tools and processes that can help companies demonstrate that a specific Linux-based system meets the necessary safety requirements for certification.

Read the Whitepaper

See how Red Hat is taking the lead

Running containers inside software-defined vehicles

In the rapidly evolving landscape of software-defined vehicles, containers are emerging as an integral component. To meet the specific requirements of container use in cars, Red Hat has been in discussions with key players and is working to combine the best of both worlds: cars and the cloud.

Read the blog post →

Fueling automotive innovation through open source collaboration

Learn how an open source in-vehicle operating system (OS), running as a horizontal software platform, empowers rapid innovation.

Read the blog post →

Red Hat to deliver the first continuously certified Linux platform for vehicles

Red Hat is preparing for ISO 26262 certification, an international standard governing the functional safety of electronic systems on road vehicles.

Read the press release →