AN OPEN, TRANSPARENT, INDUSTRY-DRIVEN APPROACH TO AV SAFETY

Jack Weast

Sr. Principal Engineer, Intel

VP, Automated Vehicle Standards, Mobileye

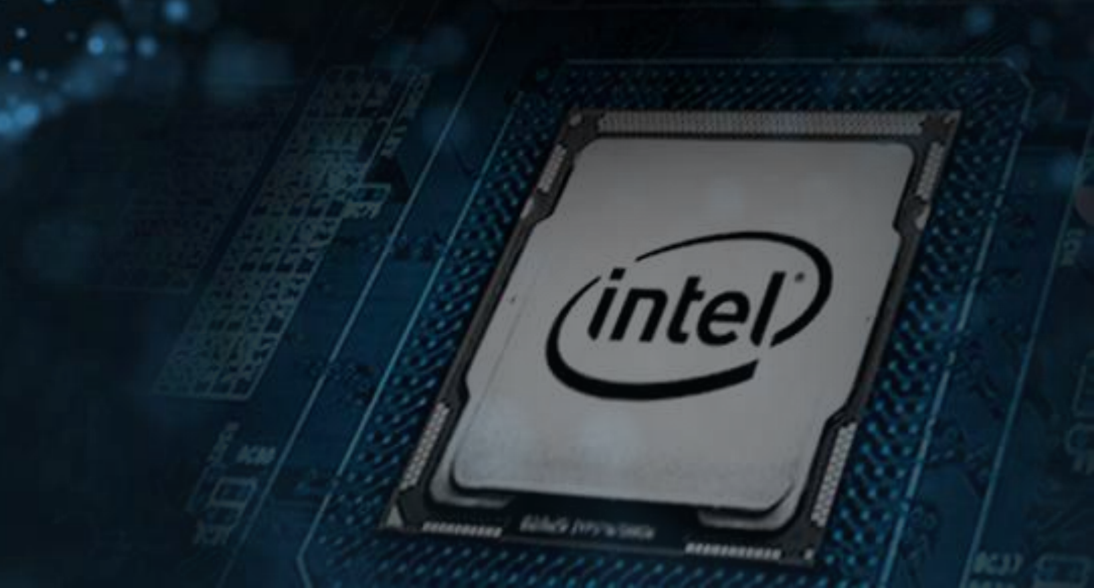

HUMAN DRIVING TODAY

The balance between safety & efficiency

HOW WOULD YOU DEFINE "DRIVING SAFELY" FOR AN AV?

A statistical argument

Self-driving cars should be statistically better than a human driver

MILES DRIVEN

The more miles I drive without a crash, the safer I am

Miles driven here

Not the same as here

HOW WOULD YOU DEFINE "DRIVING SAFELY" FOR AN AV?

A catch-all

Avoid collisions at all costs

THE AV MUST AVOID COLLISIONS AT ALL COSTS

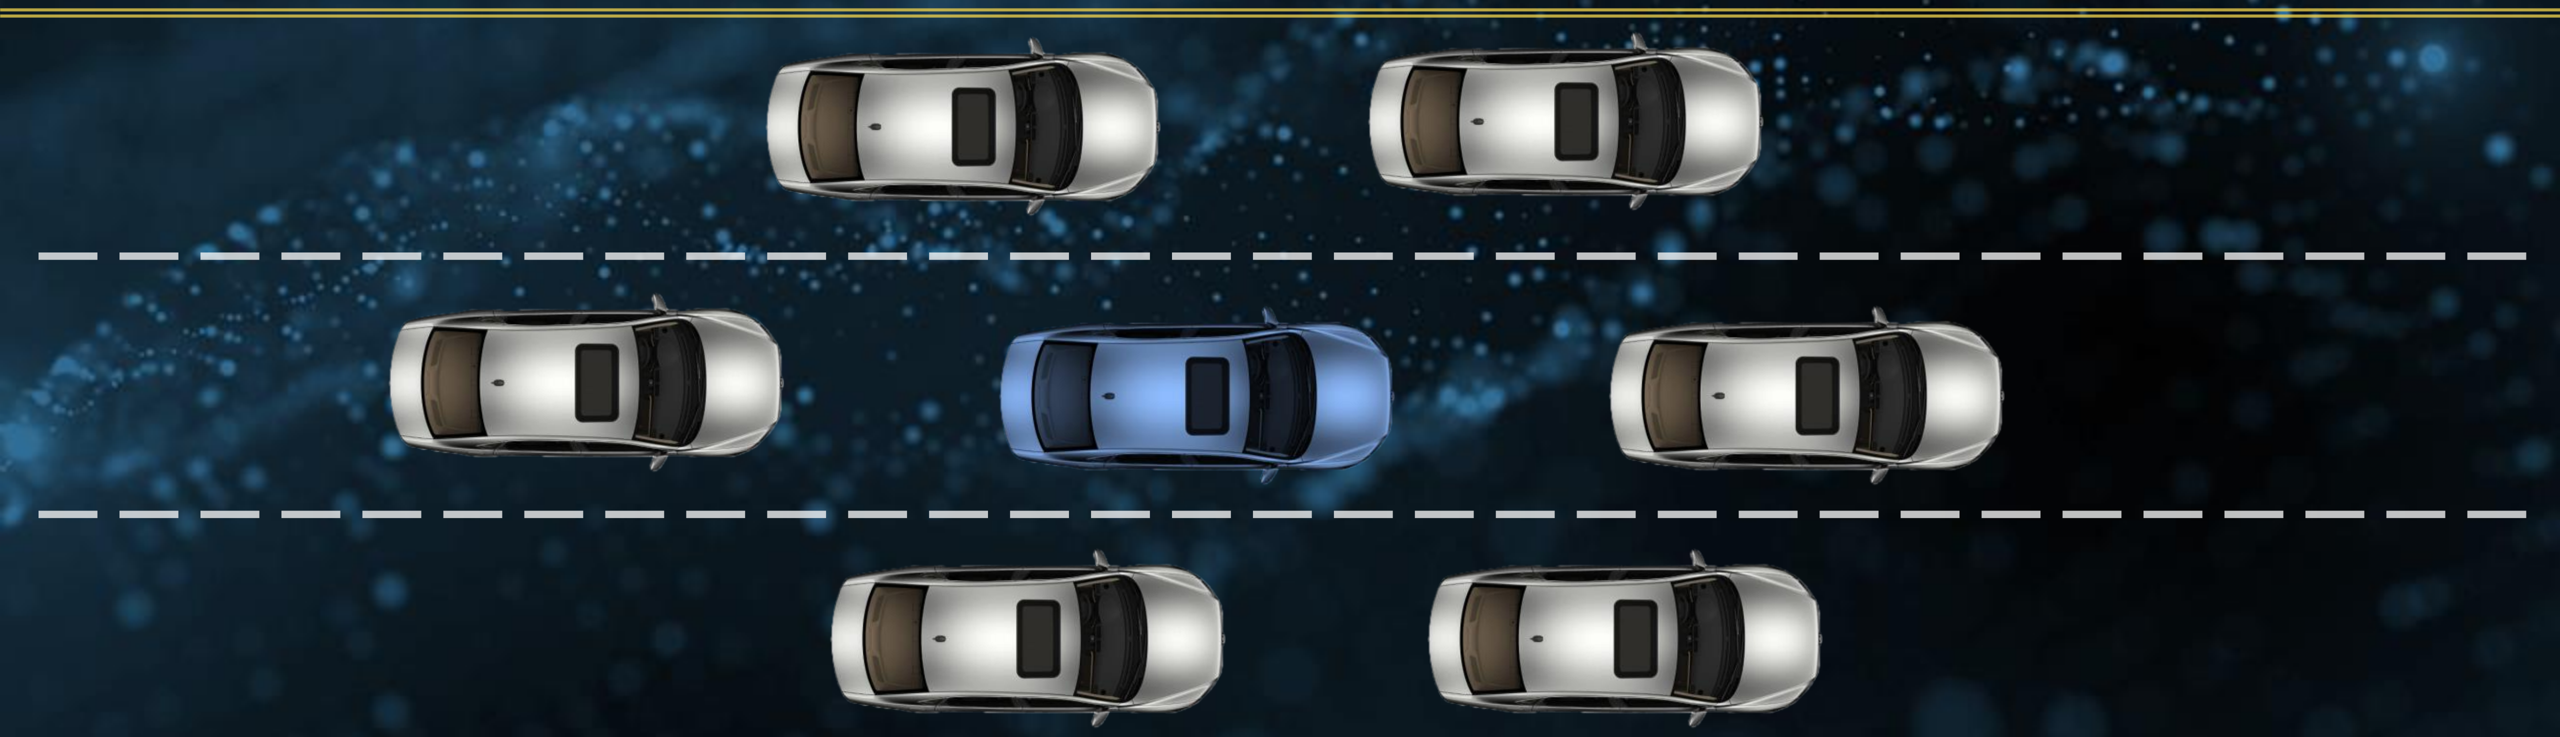

So what do we do?

What do humans do?

EXPLICIT TRAFFIC RULES

Establish **priority of road agent interests** to avoid collisions

- Come to complete stop at red lights
- Don't cross a double-yellow line
- Obey posted speed limits
- Yield to other road users when posted

Set limits on vehicle operation

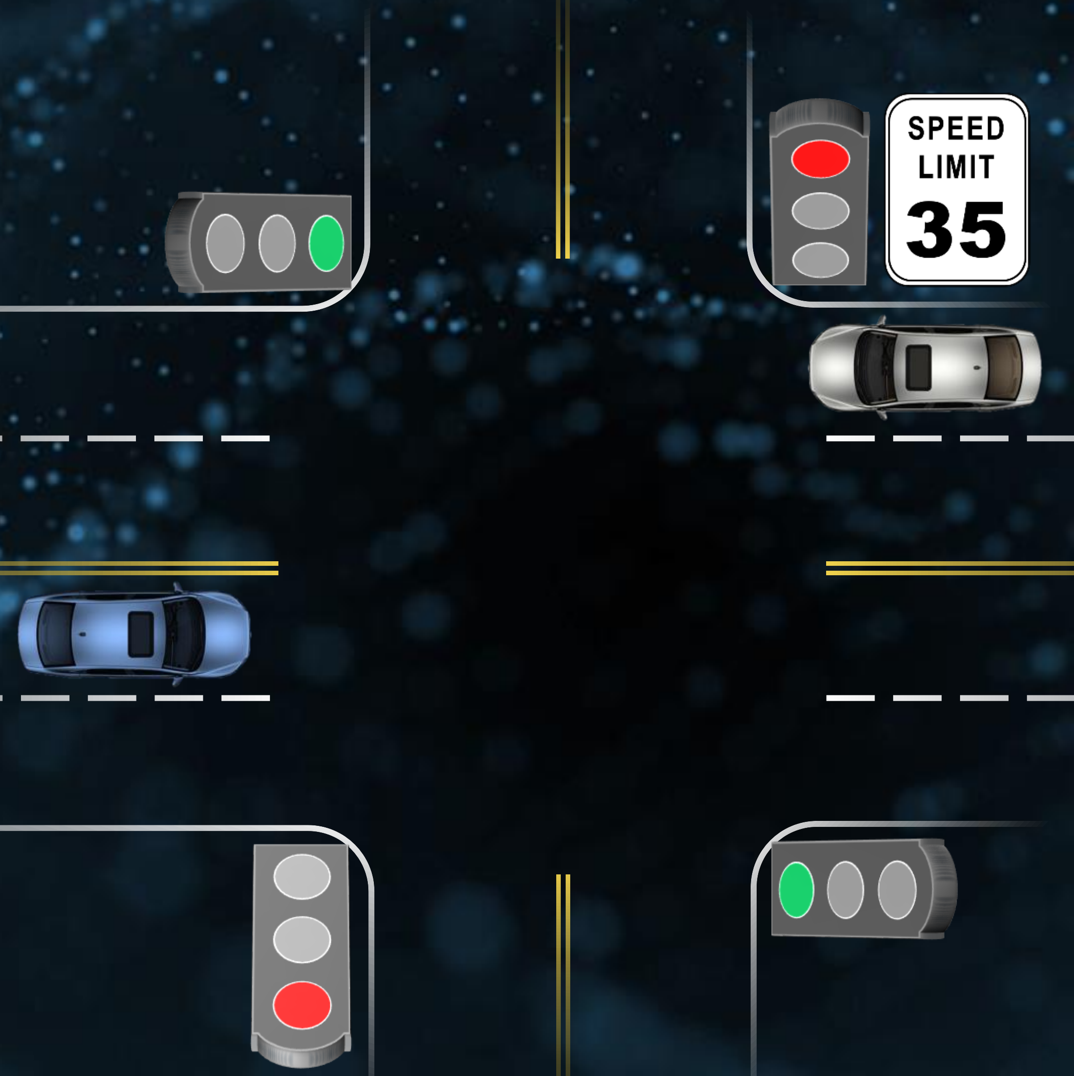

IMPLICIT RULES OF THE ROAD

A general set of principles applied by the driver

- Keep a safe distance from the car in front of you
- Drive cautiously under limited visibility
- Don’t drive slow in the fast lane
- Don’t cut off other drivers

Flexible, culturally dependent

IMPLICIT RULES OF THE ROAD

Essential for Navigating Complex Scenarios

RESPONSIBILITY SENSITIVE SAFETY

An open, transparent, technology neutral **safety model** for autonomous driving

RSS digitizes the implicit rules of the road, **providing a check on AV decision-making, and a technology-neutral performance benchmark for regulators**

RULES OF RSS

Rules to verify AV safety & performance

1. Do not hit someone from behind
2. Do not cut-in recklessly
3. Right-of-Way is given, not taken
4. Be careful in areas with limited visibility
5. If you can avoid a crash without causing another, you must

RESPONSIBILITY SENSITIVE SAFETY (RSS)

FORMALIZE

Human notions of
safe driving

IDENTIFY

A Dangerous Situation

EXECUTE

The Appropriate Response

Keep a safe distance longitudinally & laterally

Safe distance compromised in both directions

Brake to restore safe longitudinal distance

PART 1: DEFINING A SAFE STATE

A formal version of the 3-second-rule from Driver's Ed

WHAT MAKES A SAFE STATE?

First and foremost, keep a **safe distance** from others

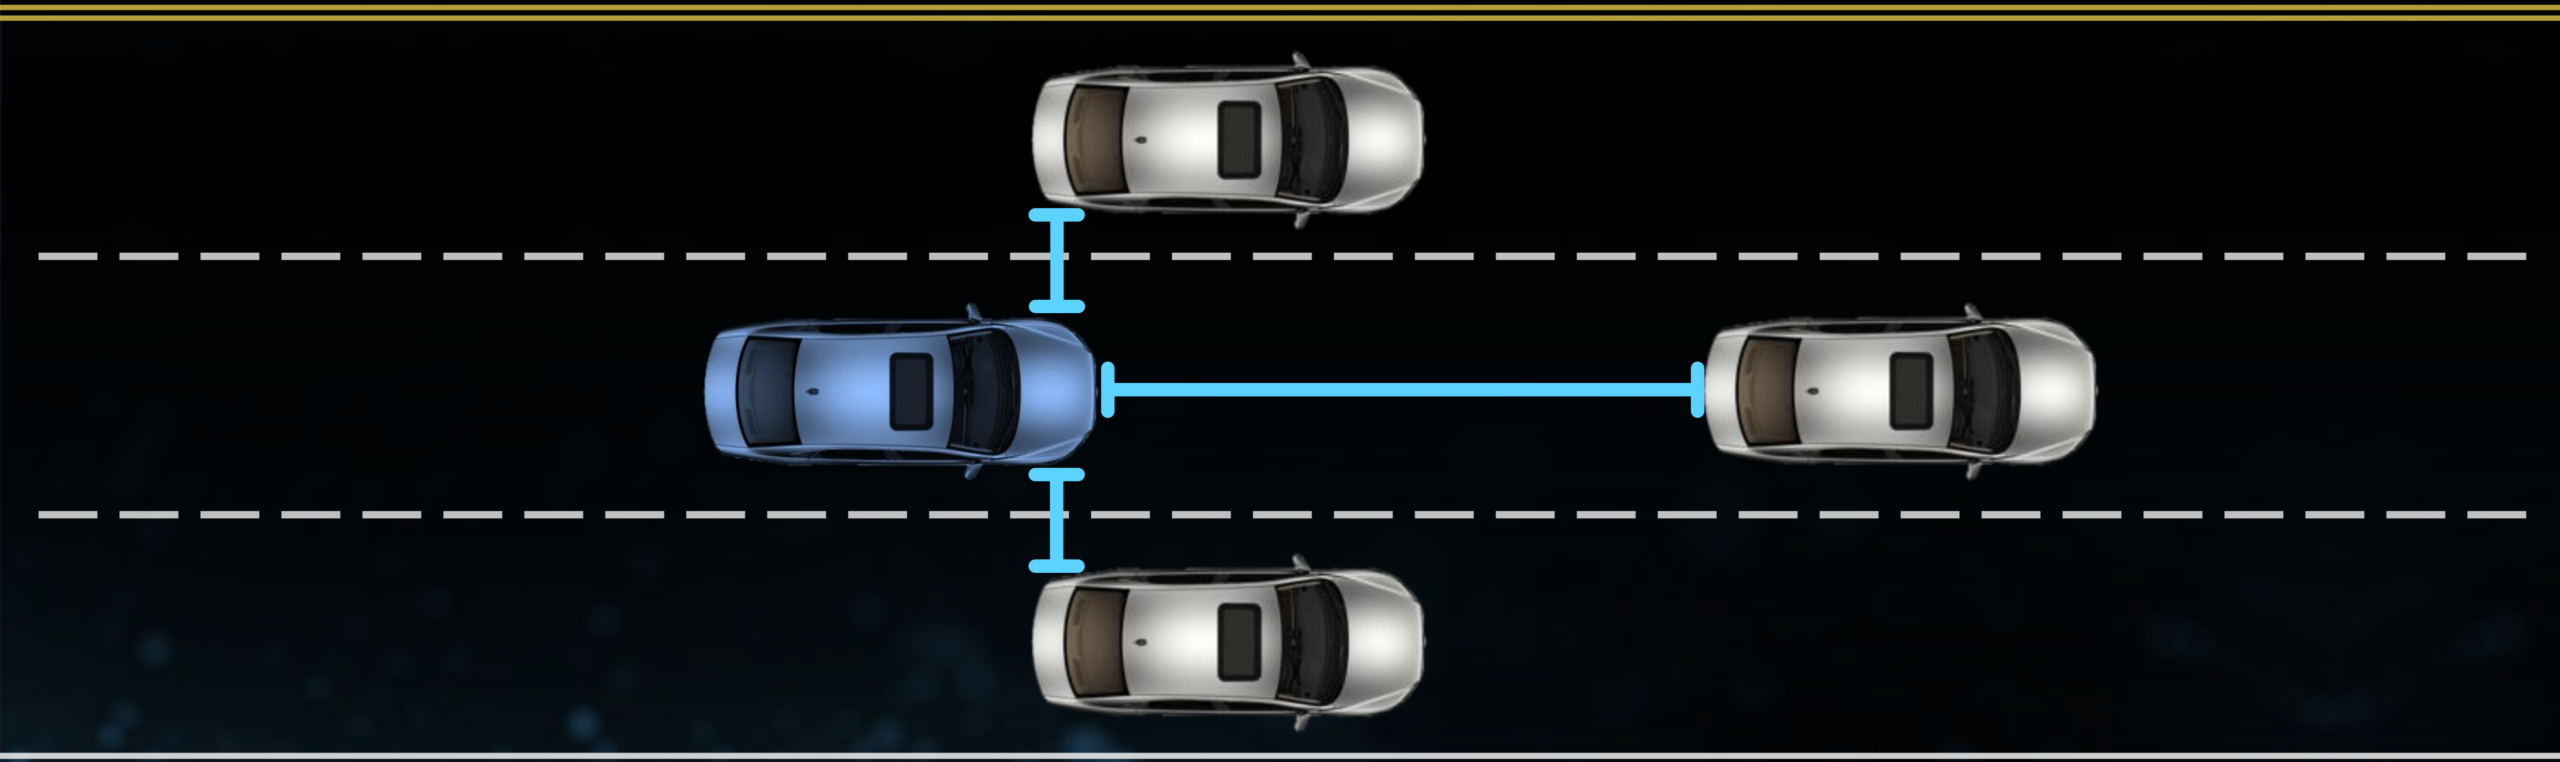

WHAT DETERMINES SAFE DISTANCE?

If the silver car slams on the brakes,
how much space does the following car need to avoid hitting it?

We determine safe distance
based on the following parameters

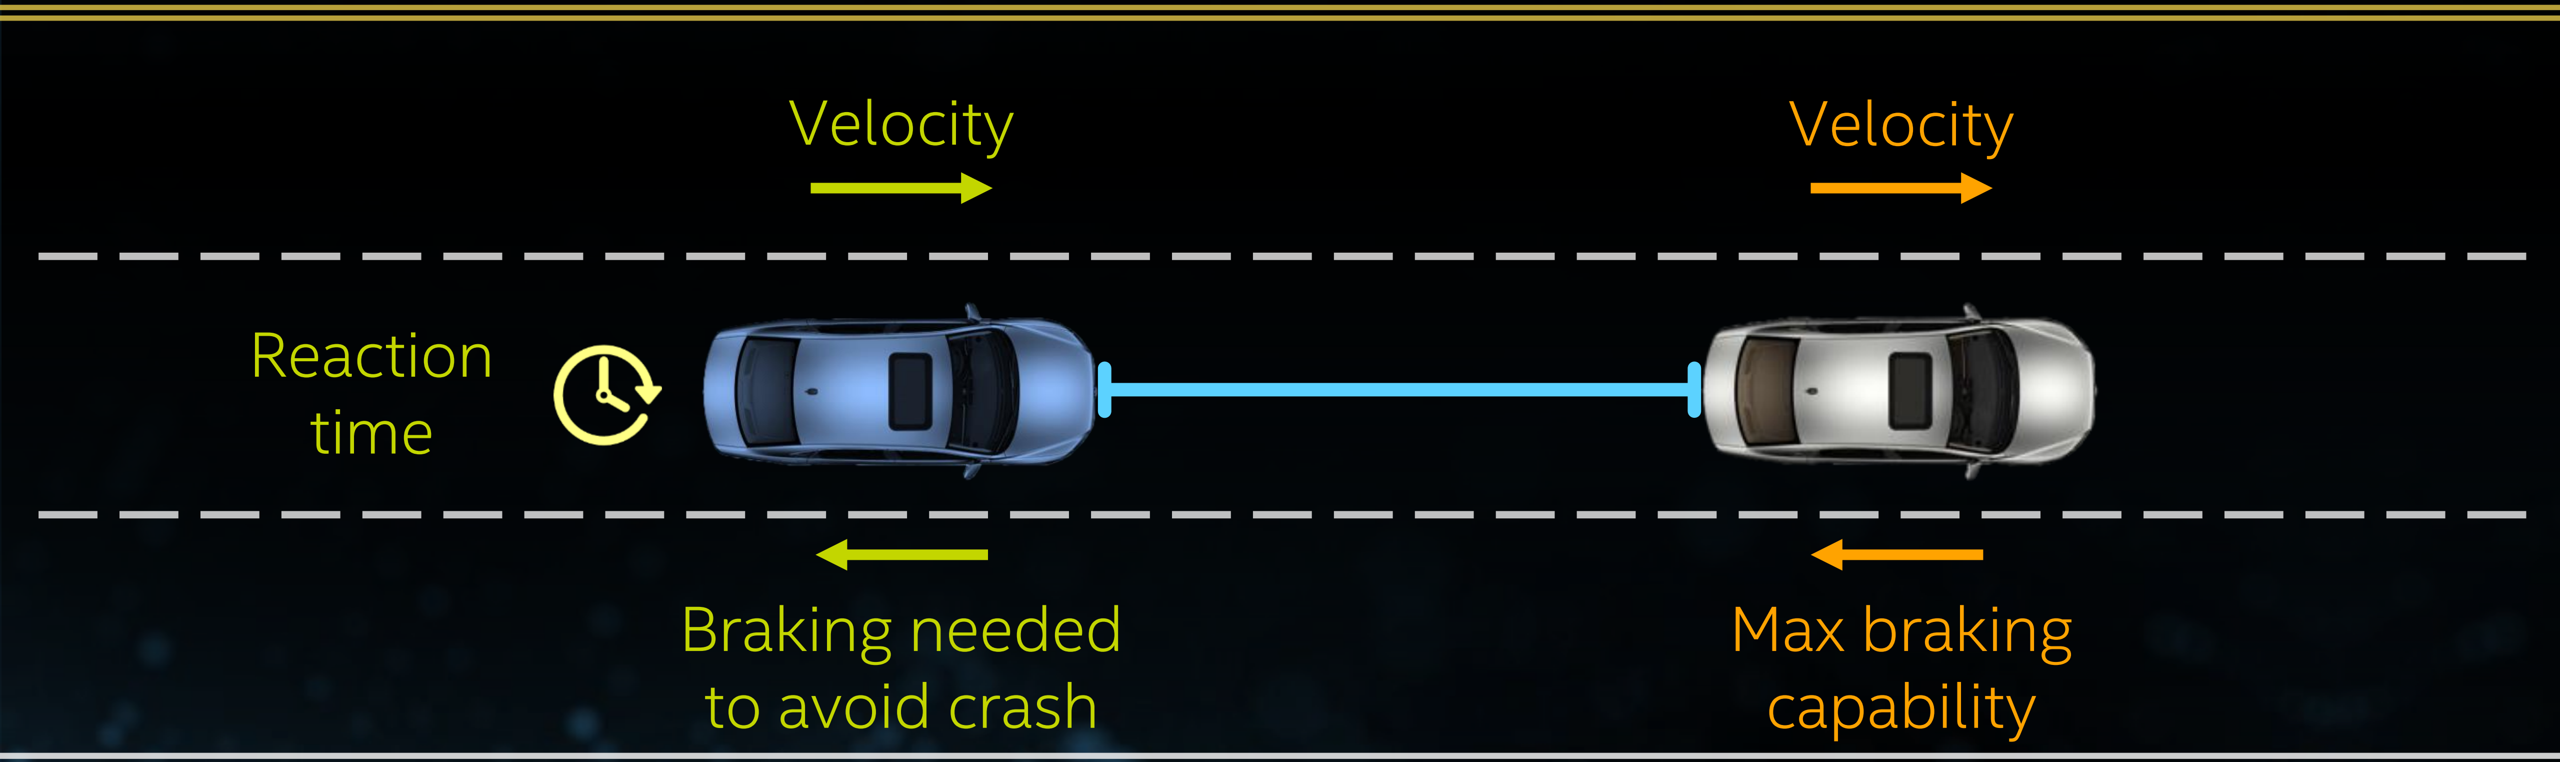

WHAT DETERMINES SAFE DISTANCE?

$$d_{min} = \left[v_r \rho + \frac{1}{2} \alpha_{max} \rho^2 + \frac{(v_r + \rho \alpha_{max})^2}{2\beta_{min}} - \frac{v_f^2}{2\beta_{max}} \right]_+$$

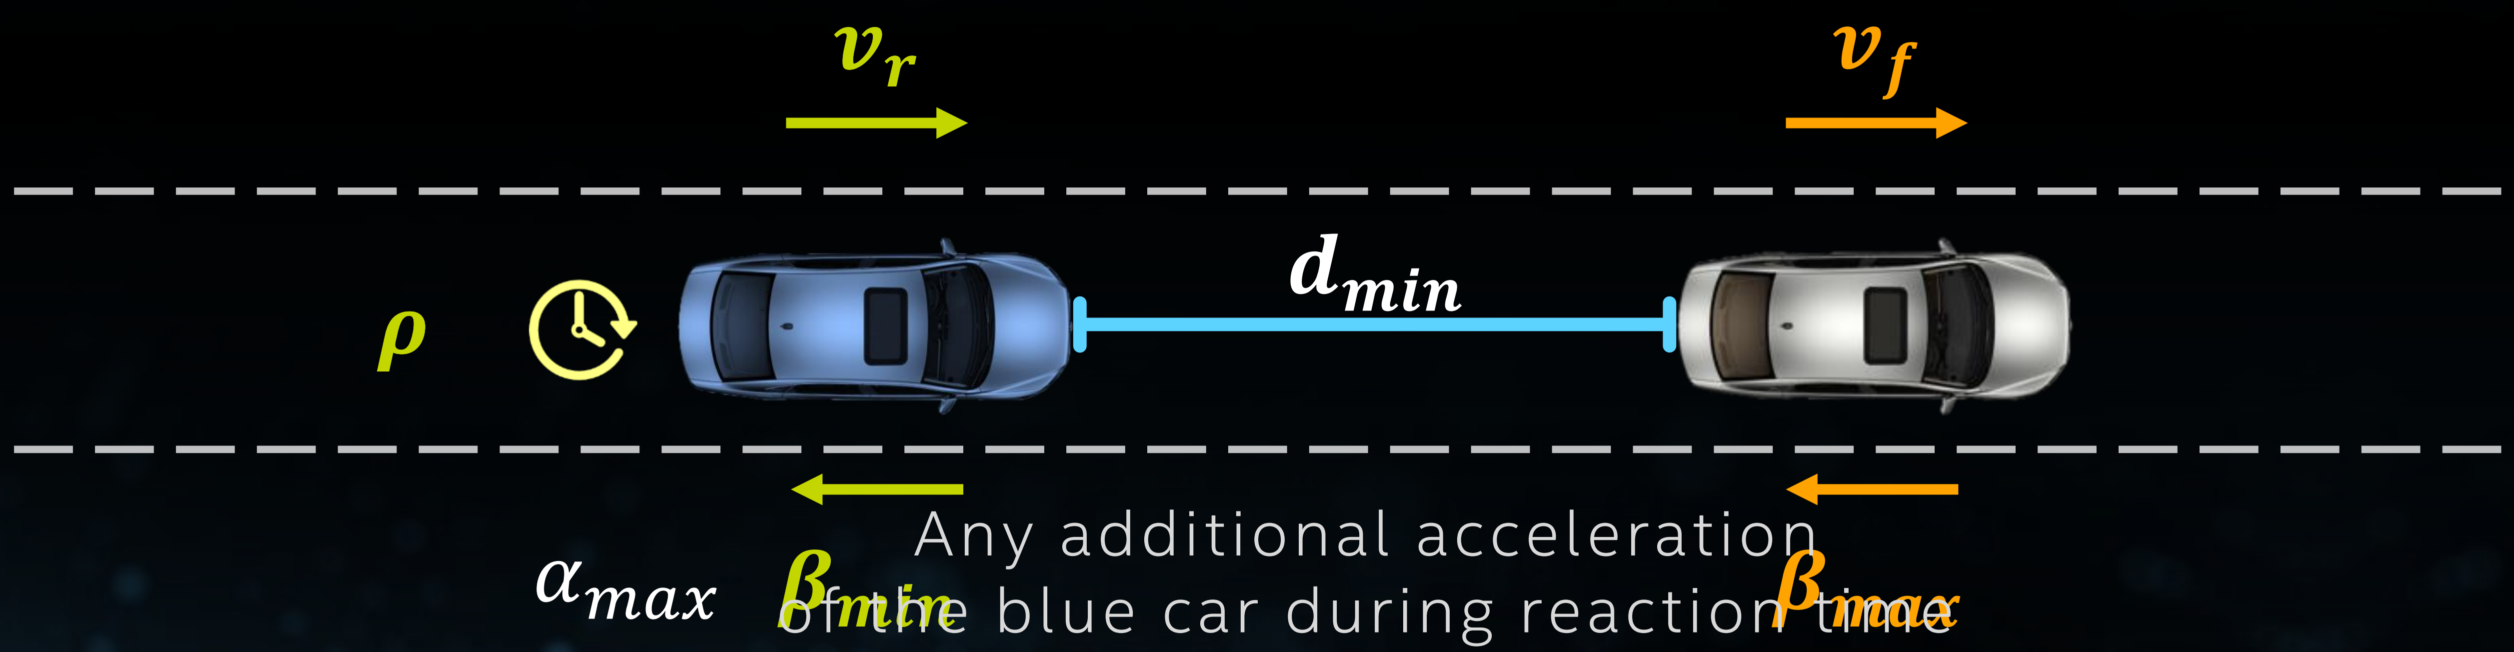

Any additional acceleration of the blue car during reaction time

DEFINE SAFE LATERAL DISTANCE

More complicated than longitudinal

We rarely stay perfectly centered in our lane

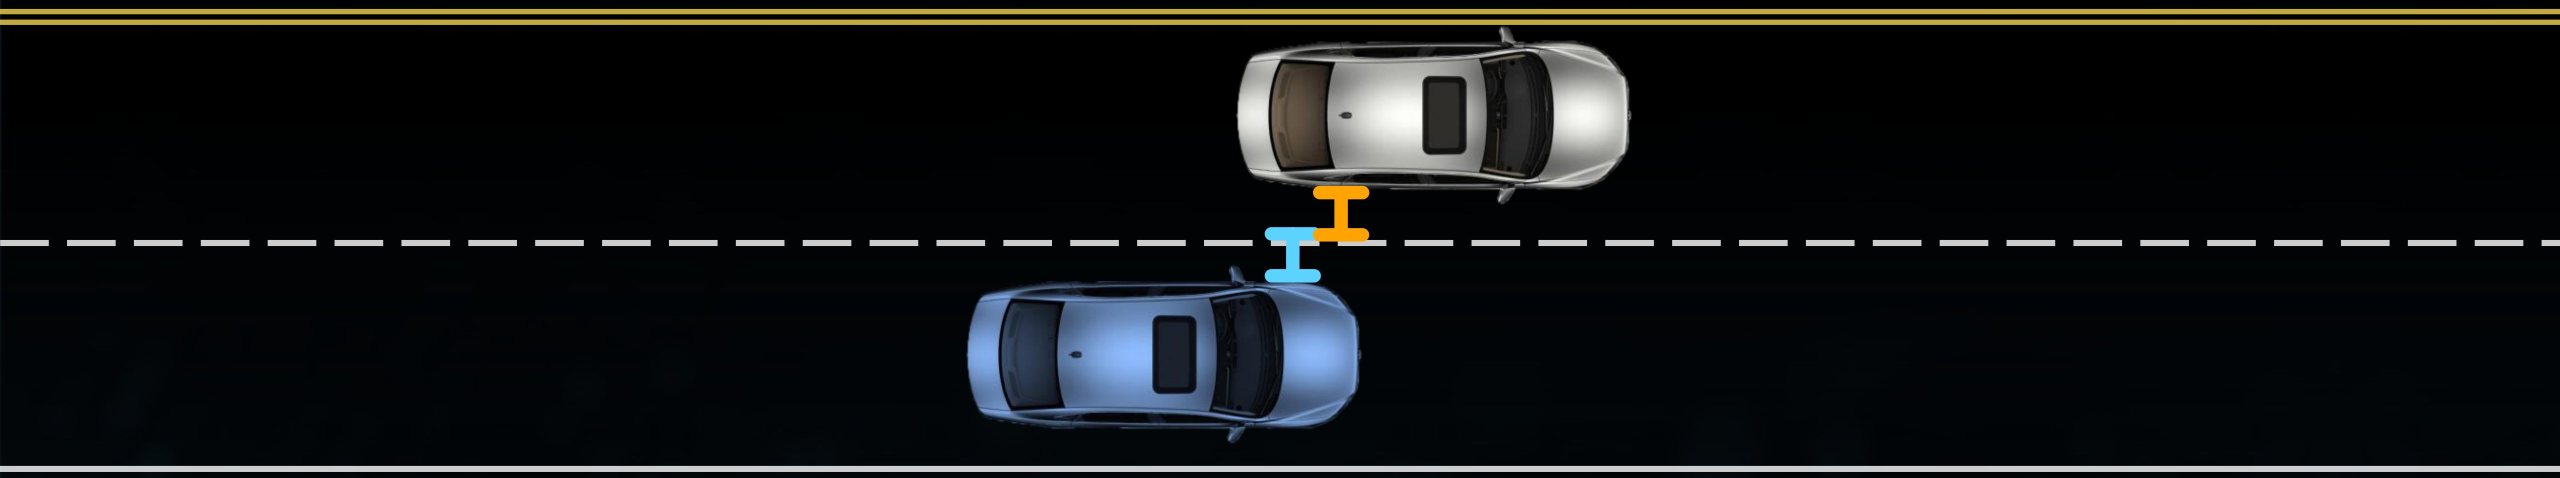

DEFINE SAFE LATERAL DISTANCE

$$d_{min} = \boldsymbol{\mu} + \left[\left(\frac{v_1 + v_{1,\rho}}{2} \right) \rho + \frac{v_{1,\rho}^2}{2\beta_{1,lat,min}} - \left(\left(\frac{v_2 + v_{2,\rho}}{2} \right) \rho + \frac{v_{2,\rho}^2}{2\beta_{2,lat,min}} \right) \right]$$

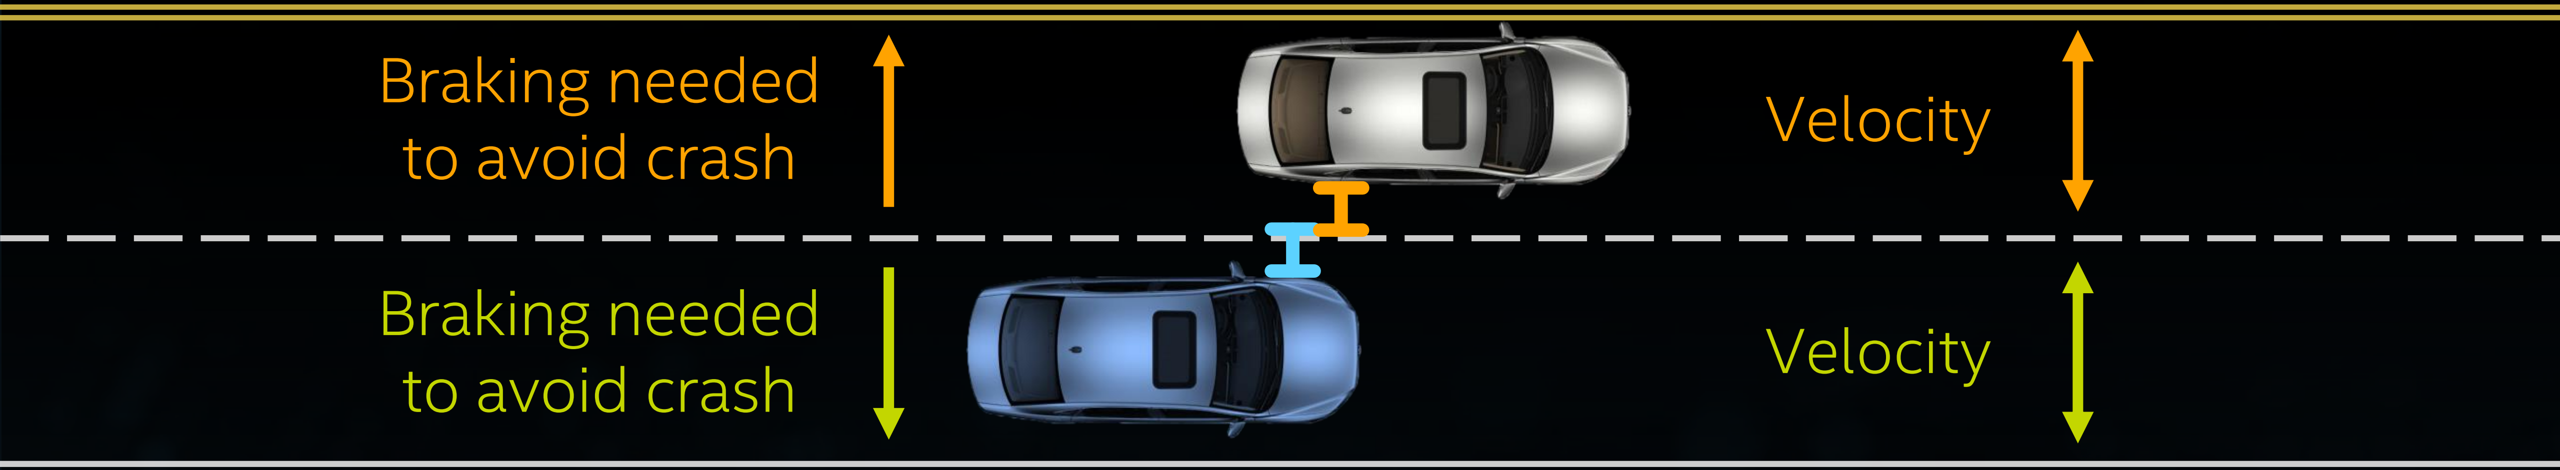

We also have a "lane within the lane" ($\boldsymbol{\mu}$)

The max movement allowed within the lane without compromising safety

From our safe distance formula, we can infer a tipping point between safety & danger

PROOF PART 2: DANGER THRESHOLD

THE DANGER THRESHOLD

The moment just before we reach an unsafe distance longitudinally and laterally

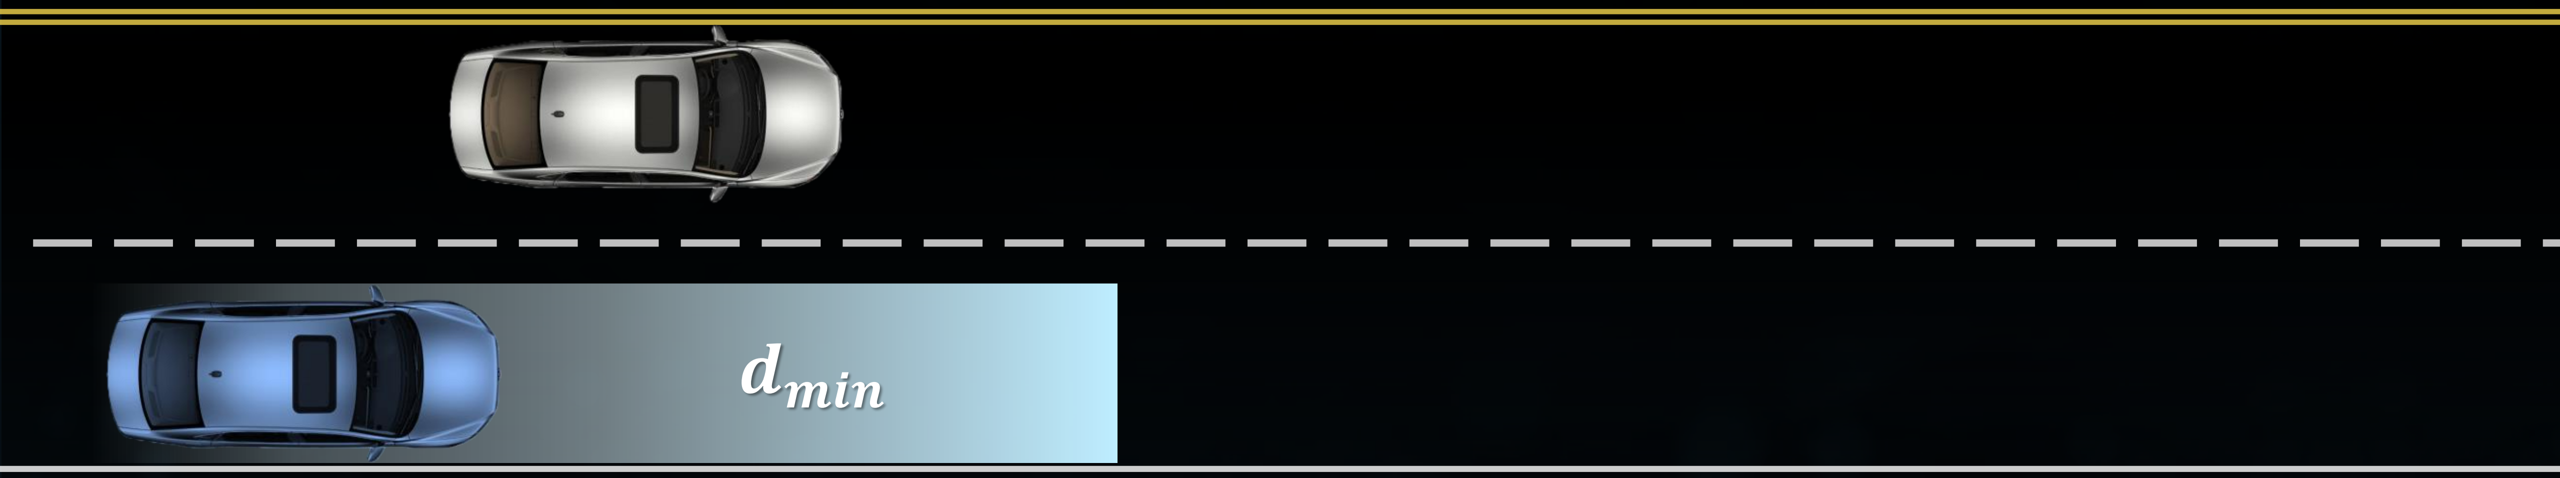

PROOF PART 3: PROPER RESPONSE

Once we cross the Danger Threshold, we must take action to restore safe distance, otherwise we remain exposed to a potentially unavoidable crash

PROPER RESPONSE – LONGITUDINAL DANGER

Though the silver car initiated the dangerous situation,
the blue car still ought to brake to return to a safe distance

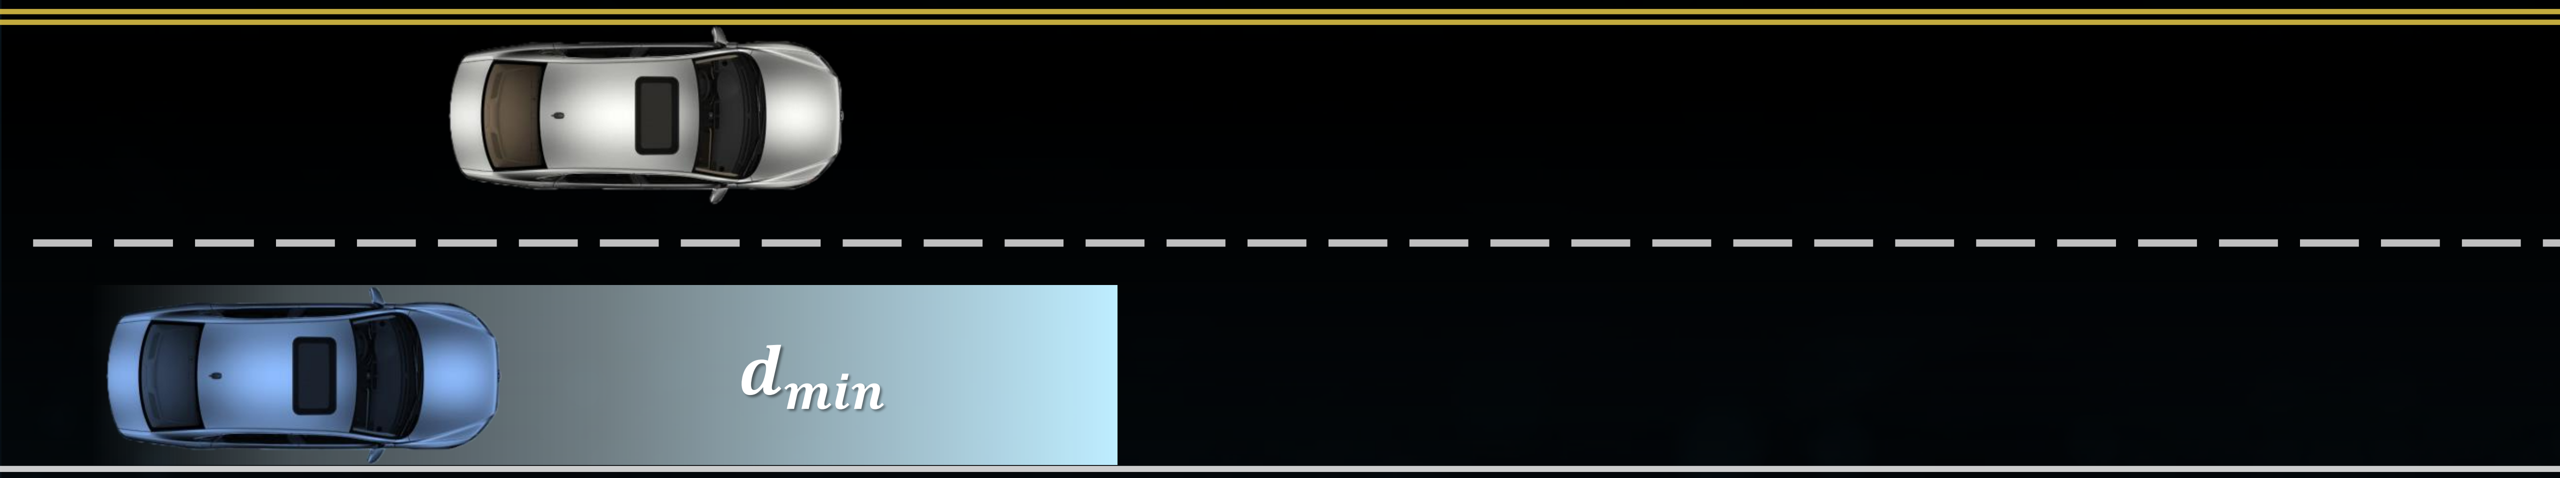

PROPER RESPONSE: LATERAL

In a dangerous lateral situation,
both vehicles may need to react to avoid a crash

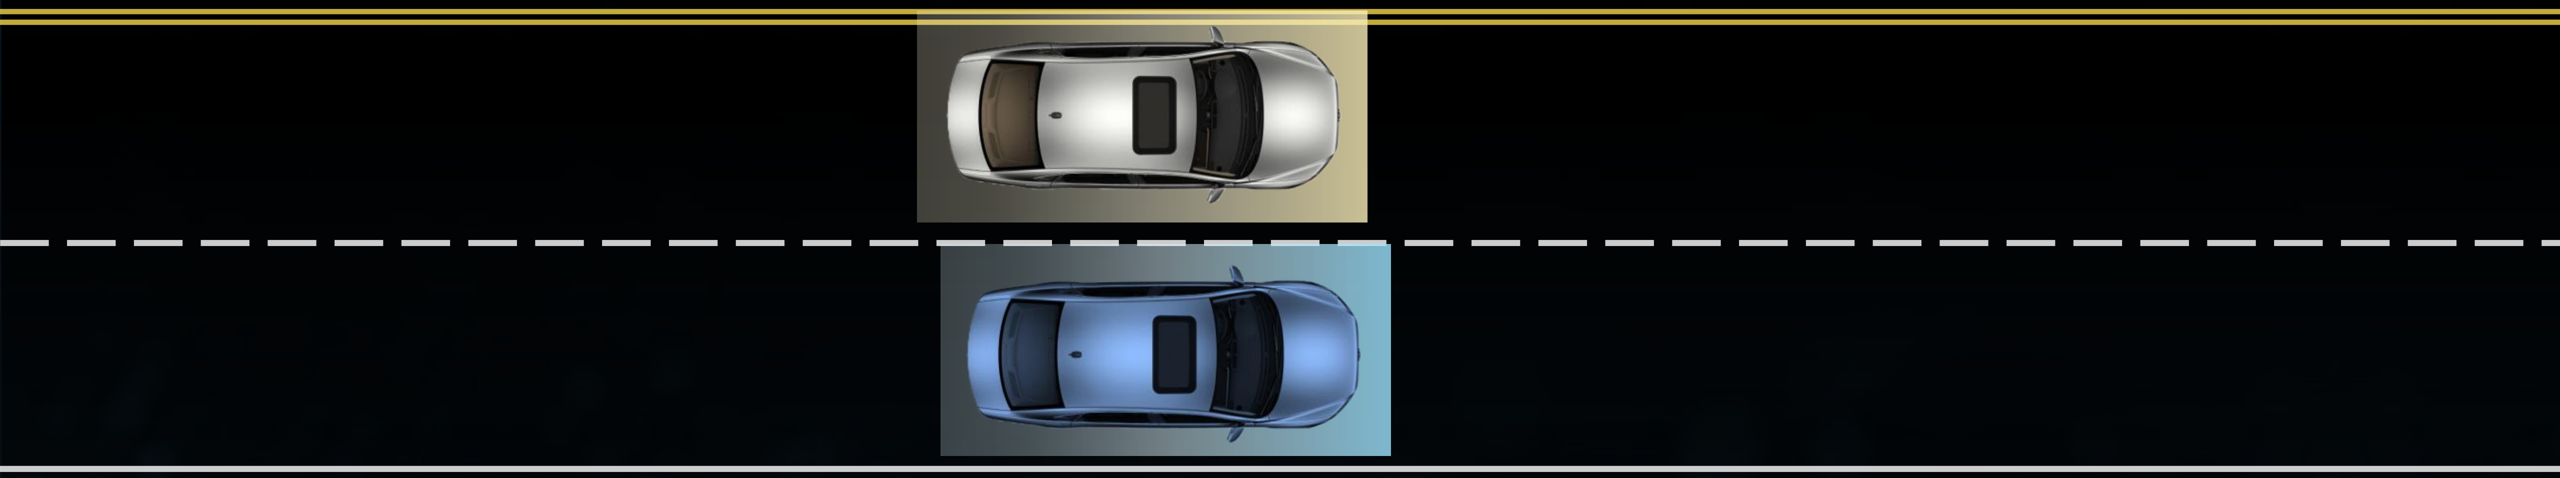

Does it work?

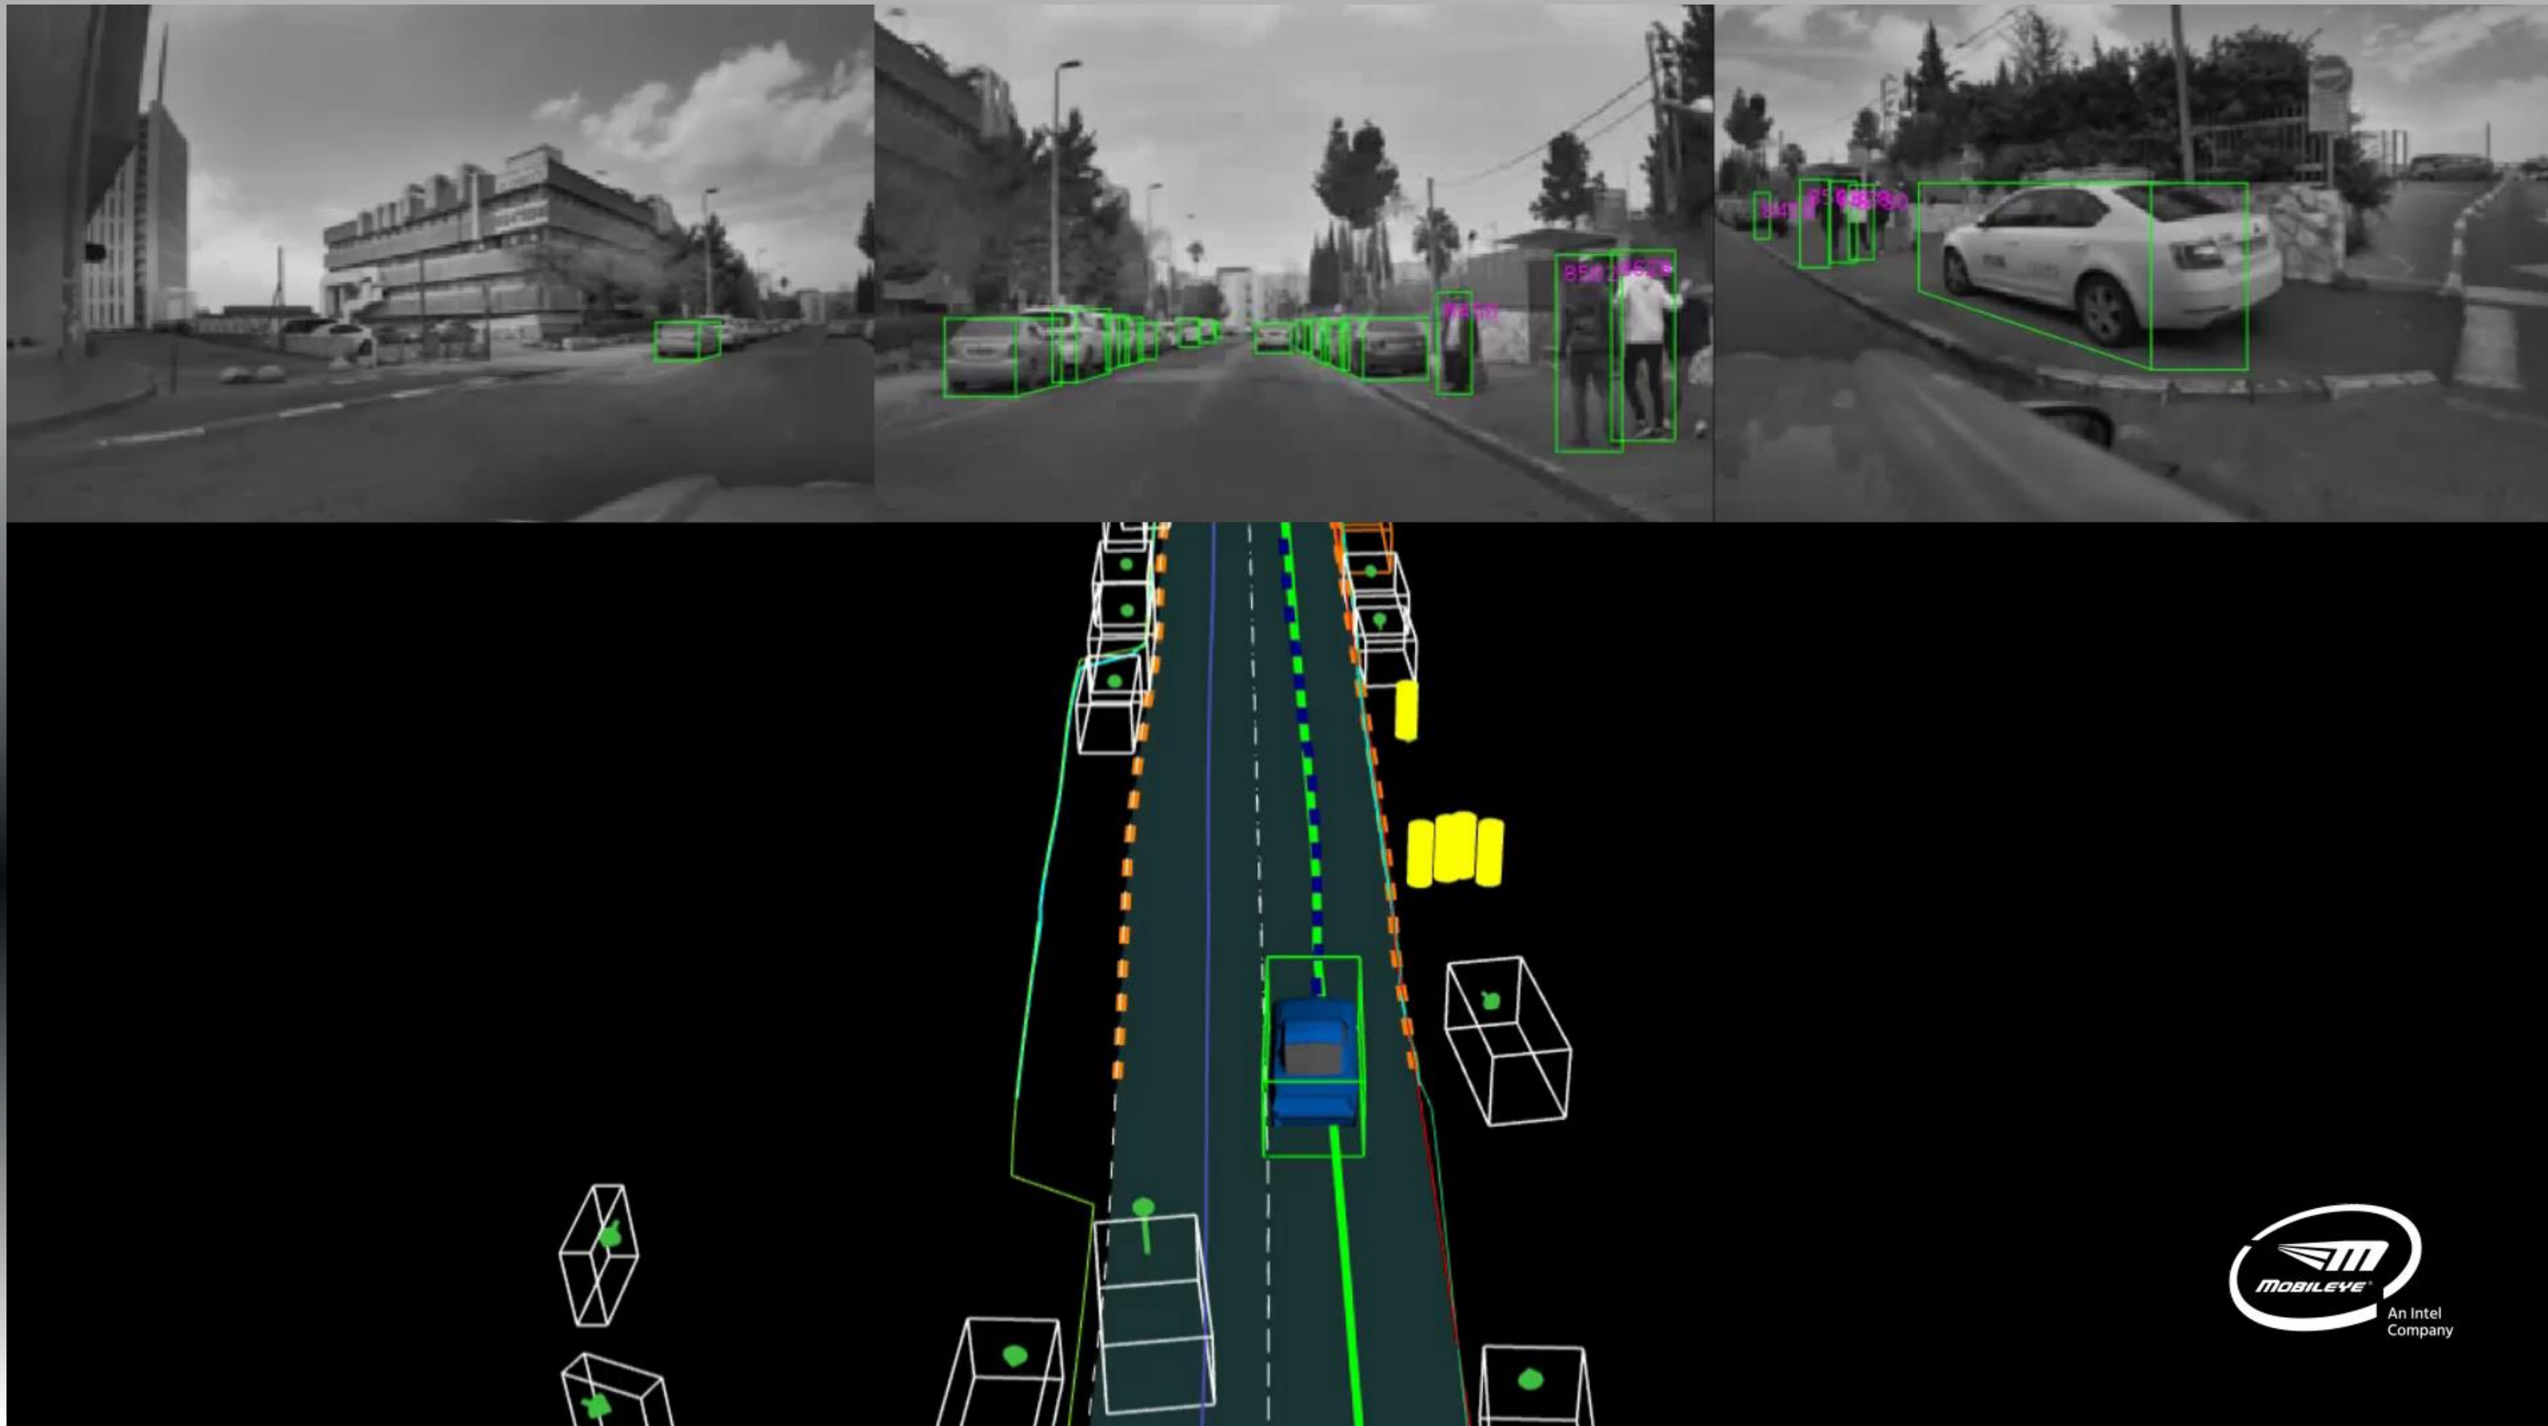
MOBILEYE
An Intel
Company

What’s the catch?

$$d_{min} = \left[v_r\rho + \frac{1}{2}\alpha_{max}\rho^2 + \frac{(v_r + \rho\alpha_{max})^2}{2\beta_{min}} - \frac{v_f^2}{2\beta_{max}} \right]_+$$

WHAT IS B_{MAX}?

Values for braking, acceleration, reaction time are not static, but dynamic based on the situation.

How do we determine the **reasonable expectations of other agents?**

12 M/S² IS A BIG ASSUMPTION FOR BRAKING

Different braking capability means different stopping distances

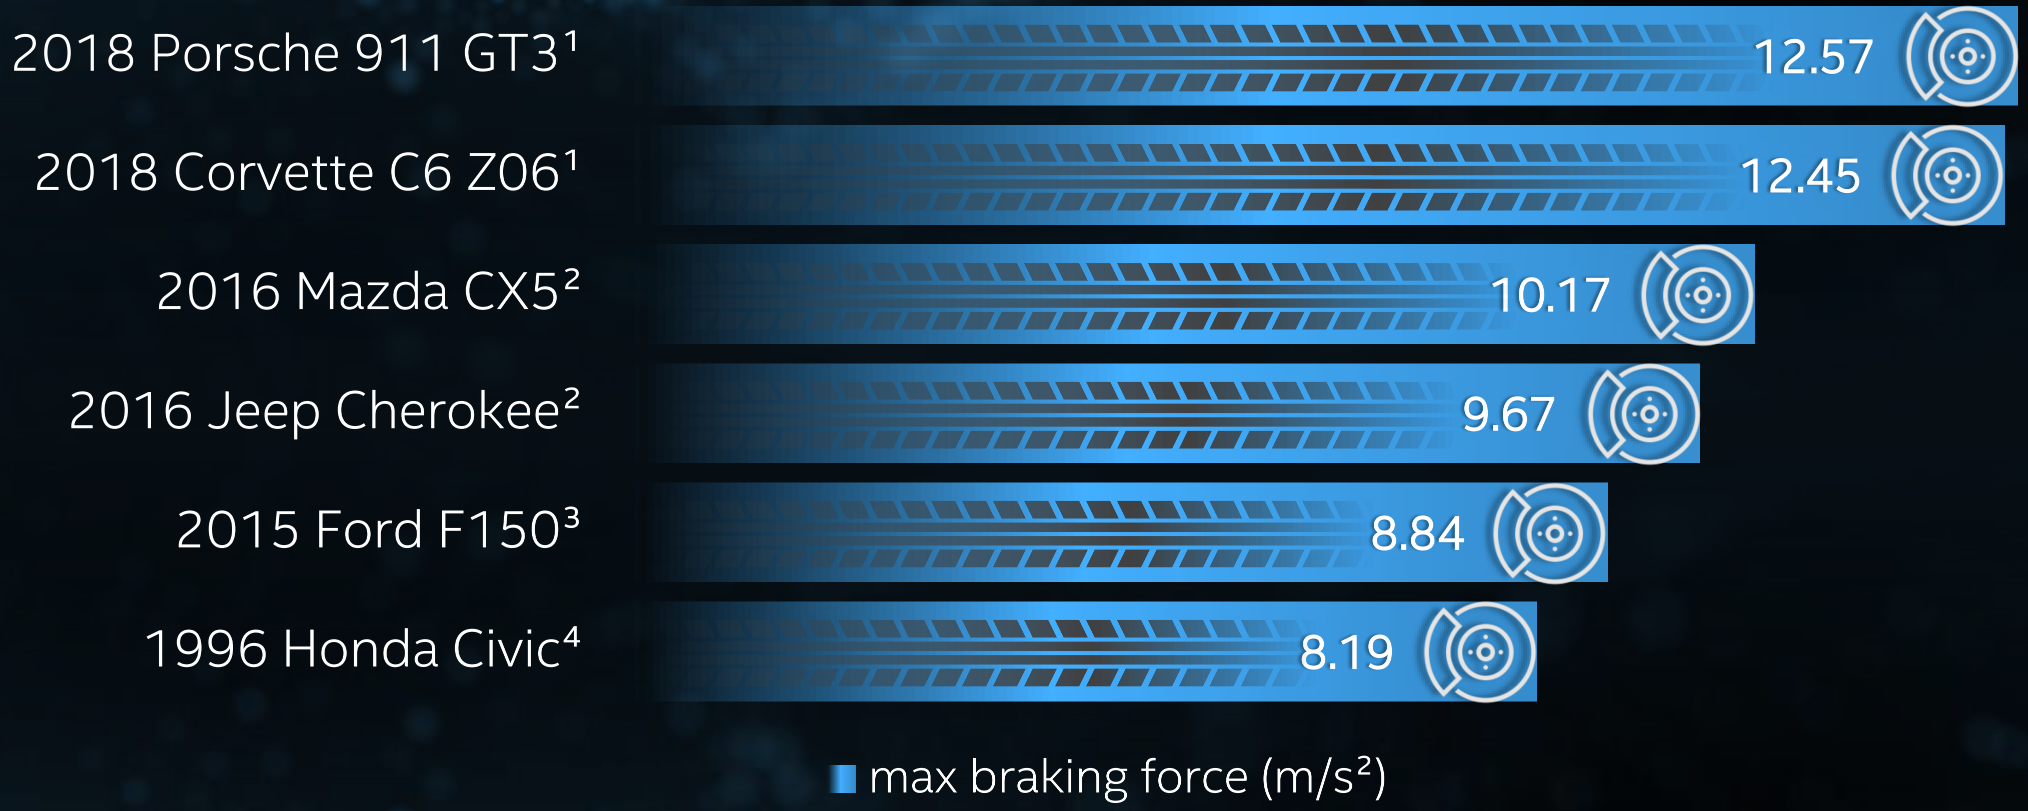

1 https://www.brembo.com/en/company/news/50-special 2 https://www.motortrend.com/cars/mazda/cx-5/2016/small-crossover-comparison-big-test/
3 https://special-reports.pickuptrucks.com/2015/01/2015-annual-physical-braking.html 4 https://www.motortrend.com/cars/honda/civic/1996/1996-honda-civic-ex-wrapup
Calculations were made using initial velocity, v_i (100kph or 60mph) and stopping distances, d, with the formula: force= v_i / (d*(2/v_i))

2 CARS TRAVELING 45 MPH: HOW MUCH SPACE IS SAFE?

36 ft (1-2 car lengths)

With **super**human reaction time and **super**car braking capability *for both vehicles*

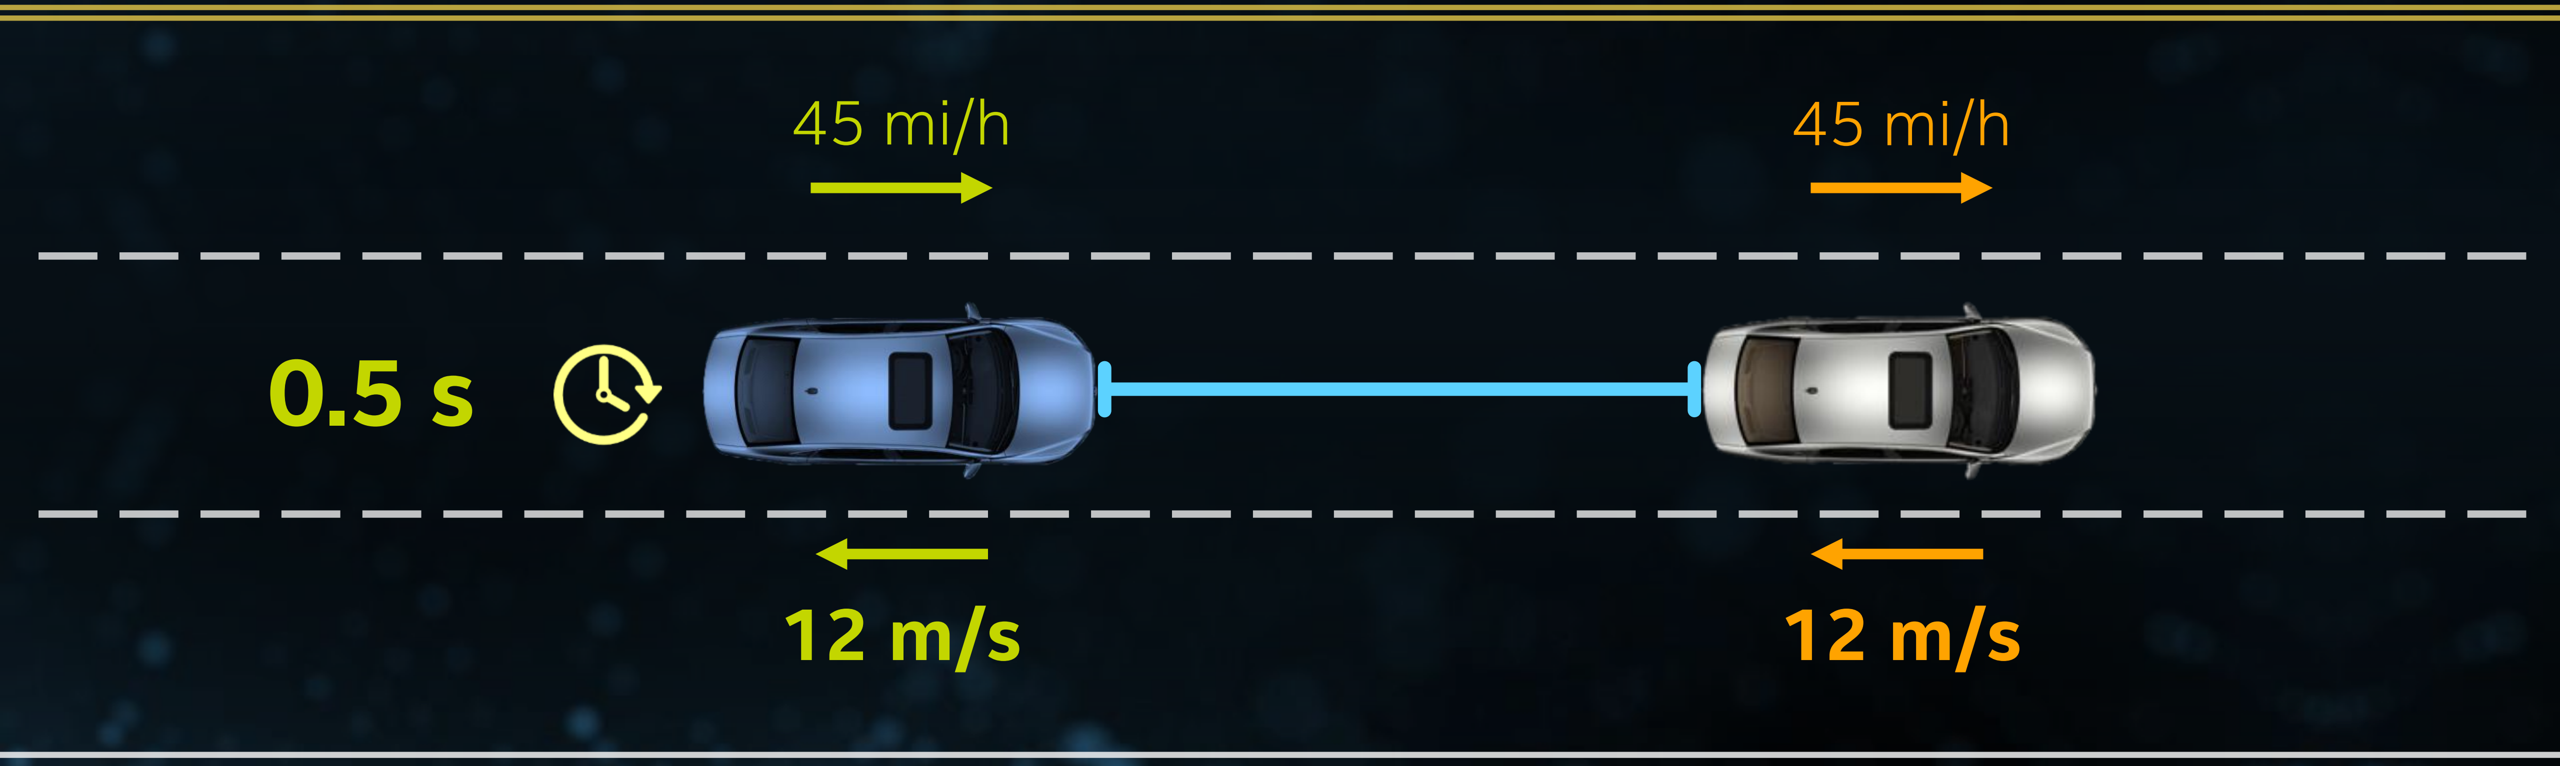

REACTION TIME PLAYS A HUGE ROLE

Human average is ~2.3 s[1]

AVs will be dramatically better (closer to 0.5 s)

Reaction Time	Braking (following car)	Braking (lead car)	**Safe Distance**
0.5 s	12 m/s^2	12 m/s^2	36 ft (2+ car lengths)
1.5 s	12 m/s^2	12 m/s^2	110 ft (7+ car lengths)
2.5 s	12 m/s^2	12 m/s^2	188 ft (9+ car lengths)

Examples assume a_{max} is 1.0 m/s^2, velocity of both cars is 20 m/s (~45mph), and that the average car length is 15 ft

1 https://copradar.com/redlight/factors/IEA2000_ABS51.pdf

WHAT HAPPENS AS WE CHANGE BRAKING CAPABILITY?

Better capability in the following car shrinks safe distance needed

Reaction Time	Braking (following car)	Braking (lead car)	**Safe Distance**
0.5 s	8 m/s^2	10 m/s^2	53 ft (3+ car lengths)
0.5 s	10 m/s^2	10 m/s^2	36 ft (2+ car lengths)
0.5 s	12 m/s^2	10 m/s^2	25 ft (1+ car lengths)

Examples assume a_{max} is 1.0 m/s^2, velocity of both cars is 20 m/s (~45mph), and that the average car length is 15 ft

WHAT HAPPENS AS WE CHANGE BRAKING CAPABILITY?

Better capability in the lead car grows safe distance needed

Reaction Time	Braking (following car)	Braking (lead car)	**Safe Distance**
2.3 s	10 m/s^2	8 m/s^2	159 ft (10+ car lengths)
2.3 s	10 m/s^2	10 m/s^2	175 ft (11+ car lengths)
2.3 s	10 m/s^2	12 m/s^2	186 ft (12+ car lengths)

Examples assume a_{max} is 1.0 m/s^2, velocity of both cars is 20 m/s (~45mph), and that the average car length is 15 ft

FASTER LATERAL ACTION → MORE DISTANCE NEEDED

Lane width in the United States ranges from 9-12 ft[1]

Reaction Time	Lateral acceleration assumption	Lateral braking to avoid crash	**Safe Distance**
0.5 s	0.8 m/s²	1.8 m/s²	4.4 ft
0.5 s	1.8 m/s²	1.8 m/s²	6.4 ft
0.5 s	3 m/s²	1.8 m/s²	10 ft

Examples assume $v_{1,p}$ and $v_{2,p}$ are $(v_1 + pa_{lat,max})$ and $(v_2 + pa_{lat,max})$, respectively, where $a_{lat,max}$ is 0.8 m/s^2, and both v_2 and v_1 are 1 m/s. μ is set to 0.5 m

1 https://safety.fhwa.dot.gov/geometric/pubs/mitigationstrategies/chapter3/3_lanewidth.cfm

LIMITED VISIBILITY & OCCLUDED AREAS

When sensing capabilities are physically limited,
We must exhibit caution

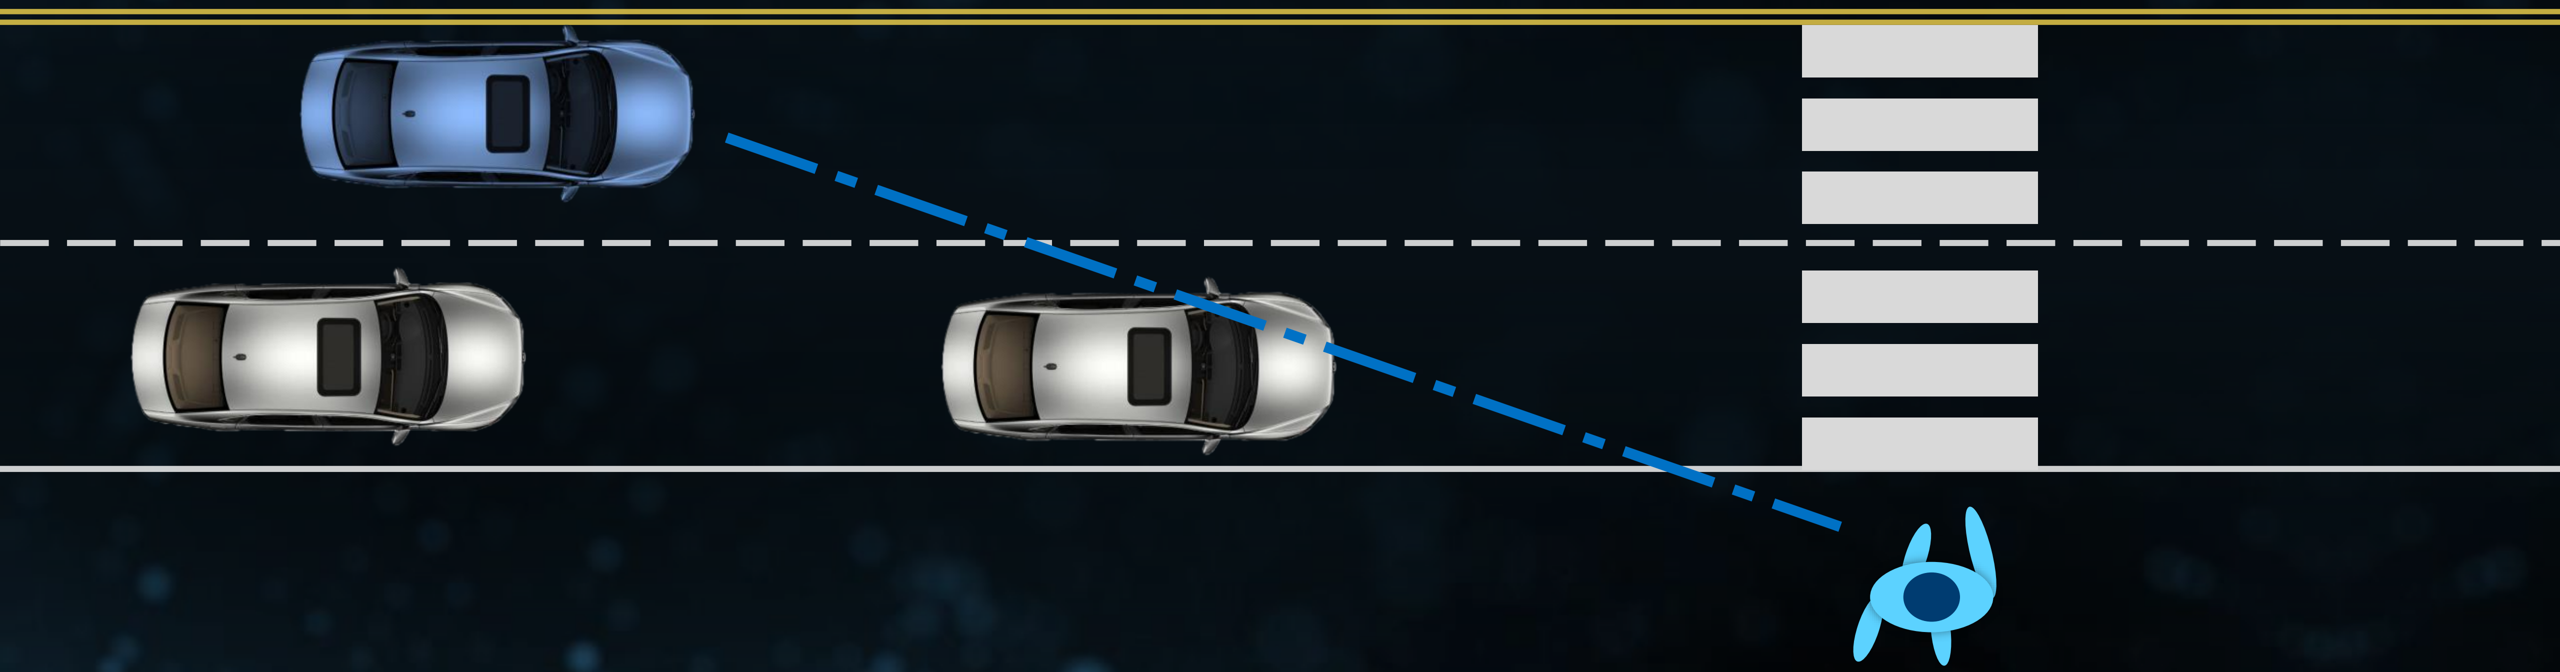

THE BLIND CORNER

If something obstructs our view, we may not see that we're about to cross the Danger Threshold

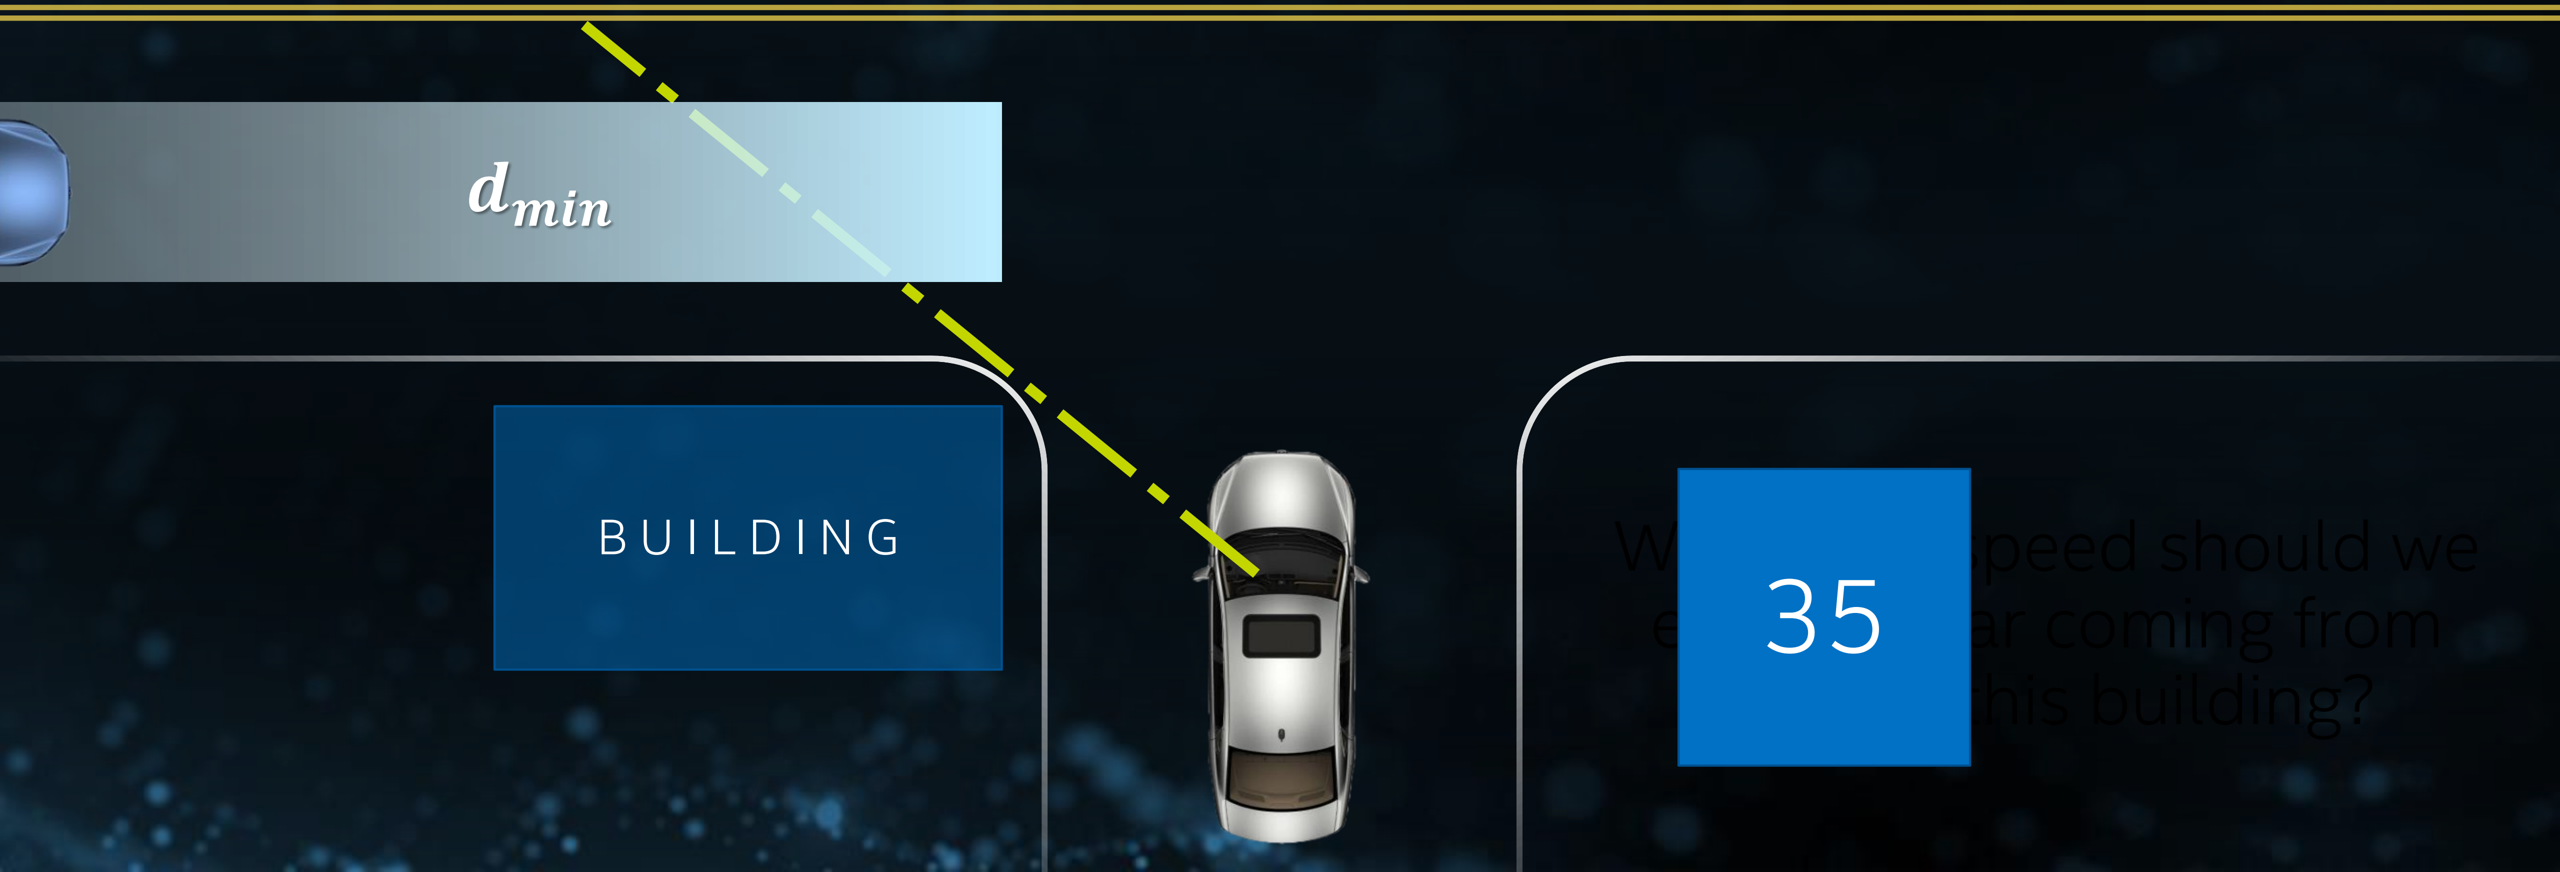

NEIGHBORHOODS WITHOUT SIDEWALKS

Are likely to have people walking along
& playing in the street

NEIGHBORHOODS WITH SIDEWALKS

Pull people away from the street, allowing cars to more safely operate at higher speeds

APPLYING RSS

RSS can fit in the vehicle, in our testing,
and in our lexicon of vehicle safety

RSS INSIDE THE VEHICLE

Validation as a doer-checker

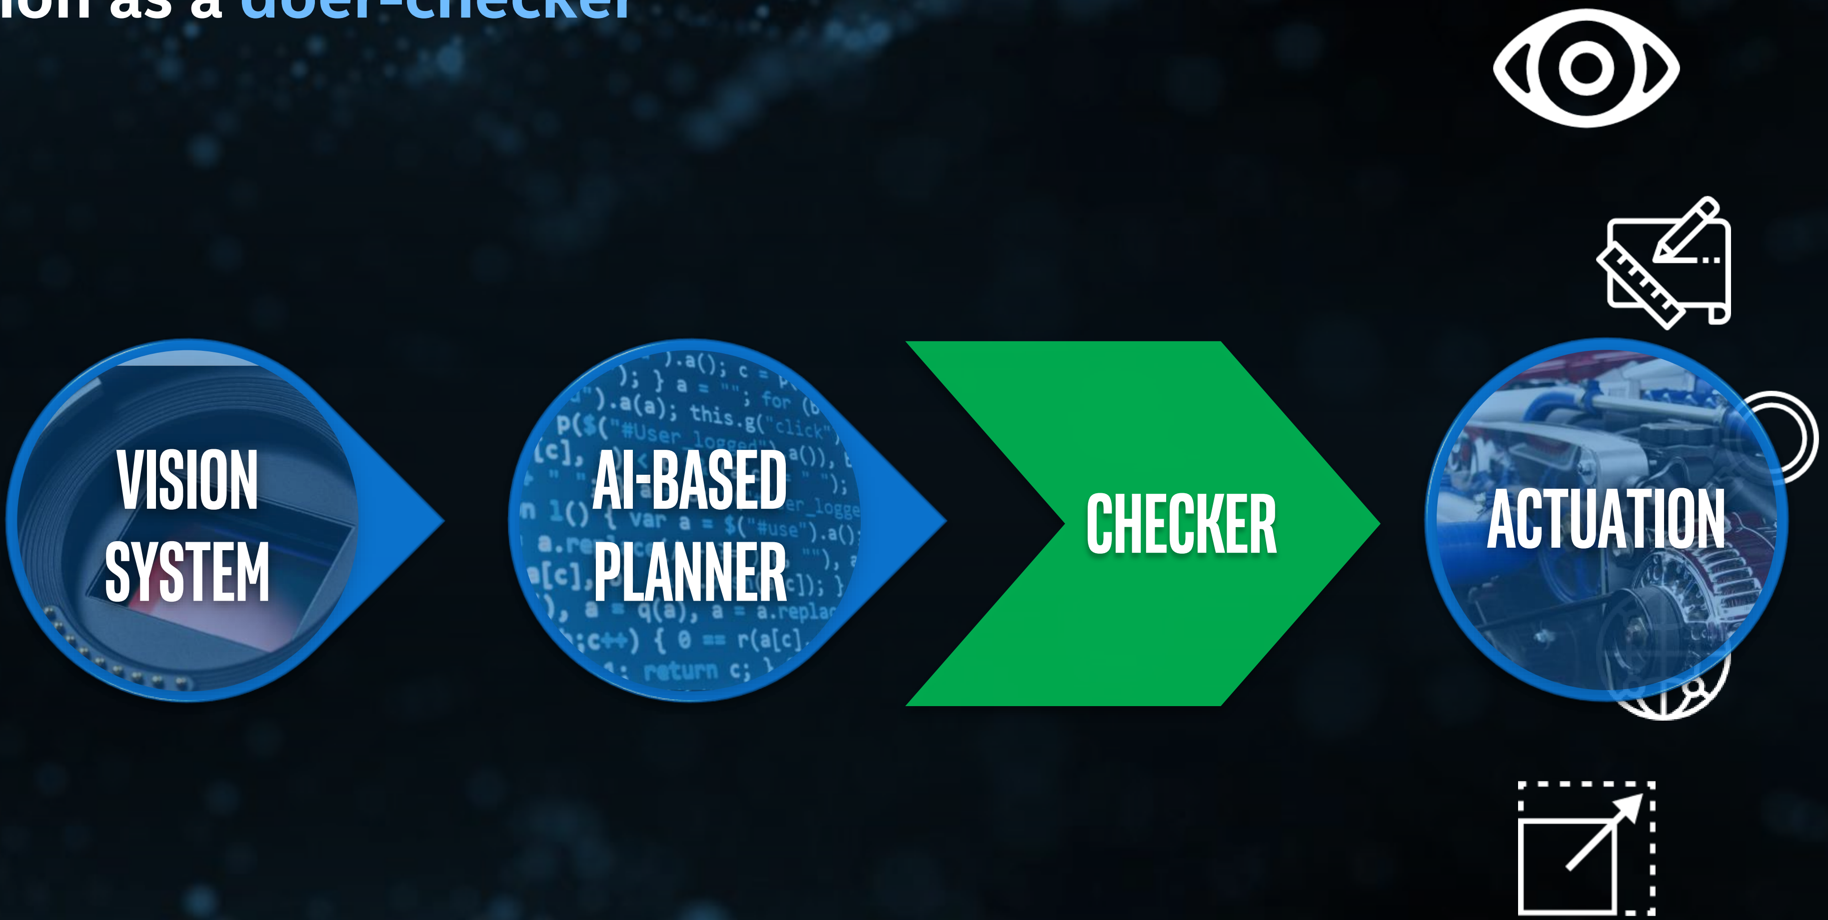

THE DOER-CHECKER

Assesses danger, validates planner's decisions, and triggers proper responses

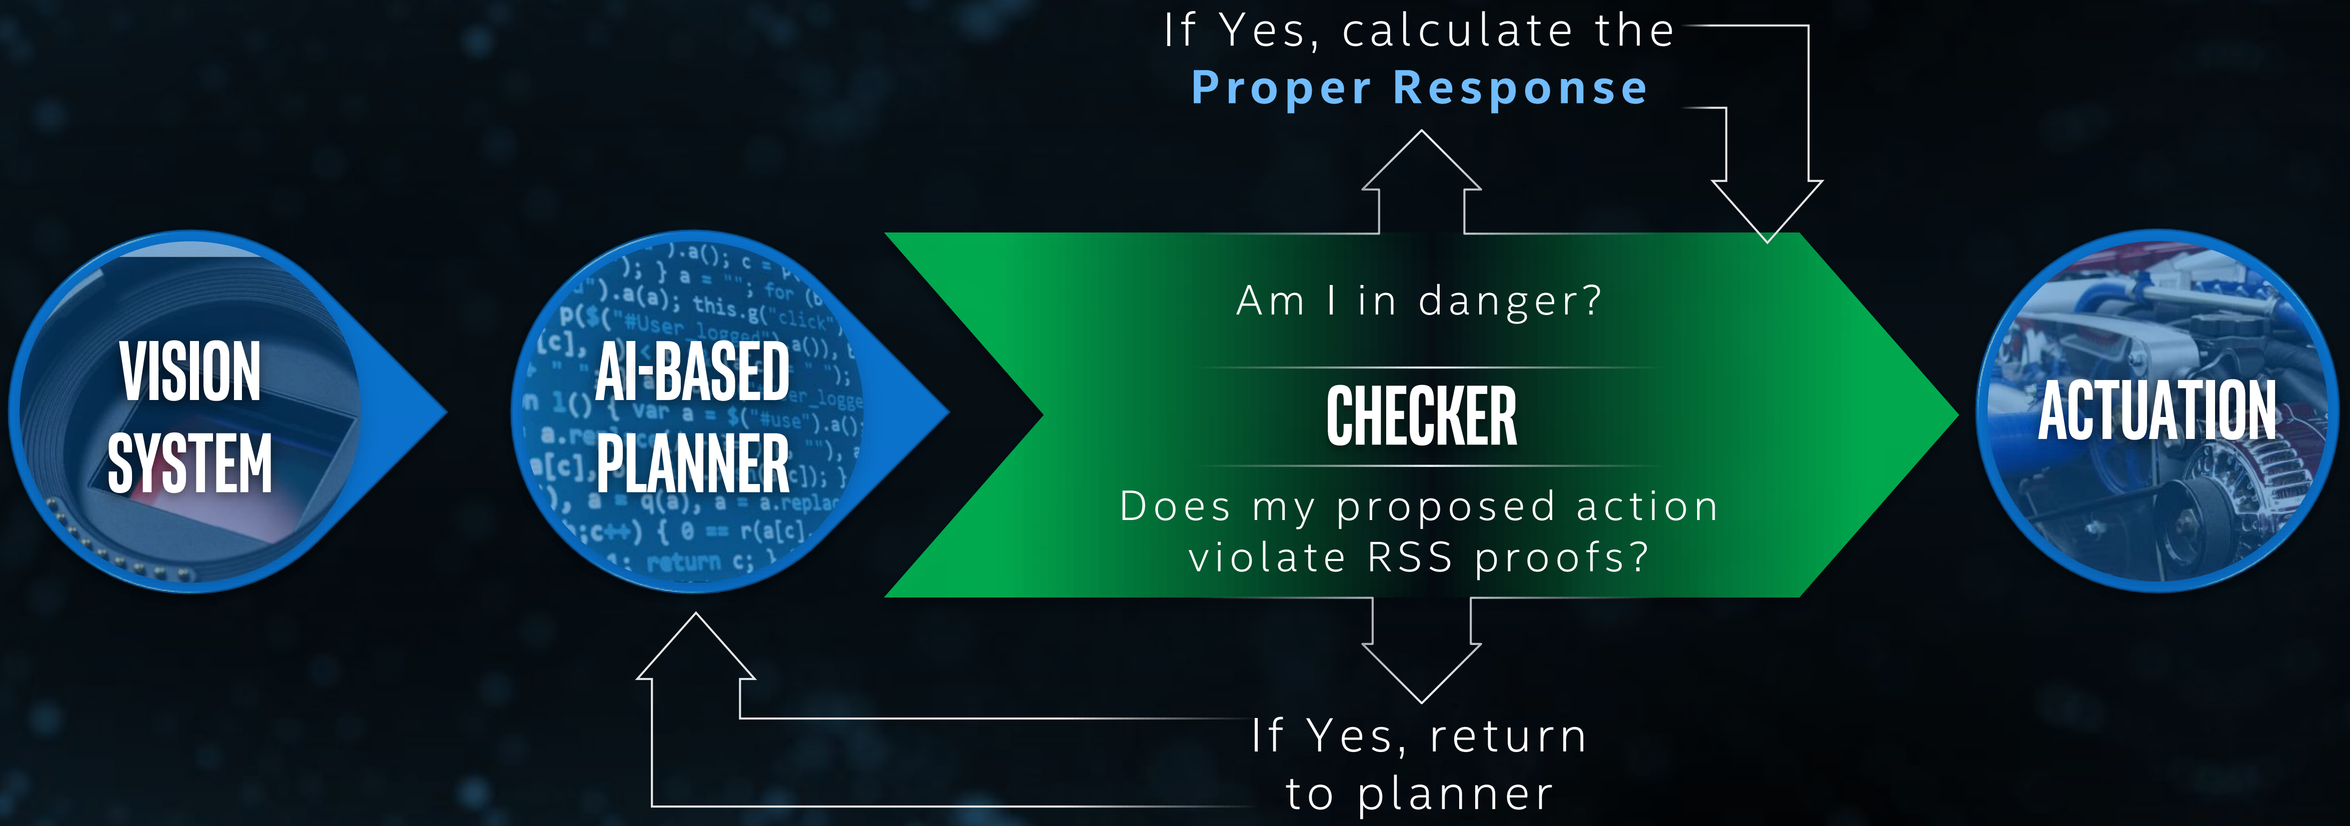

RSS OUTSIDE THE VEHICLE

Validating vehicle behavior on a test track

ARIZONA
COMMERCE AUTHORITY

GOVERNMENT

PRIVATE
INUDSTRY

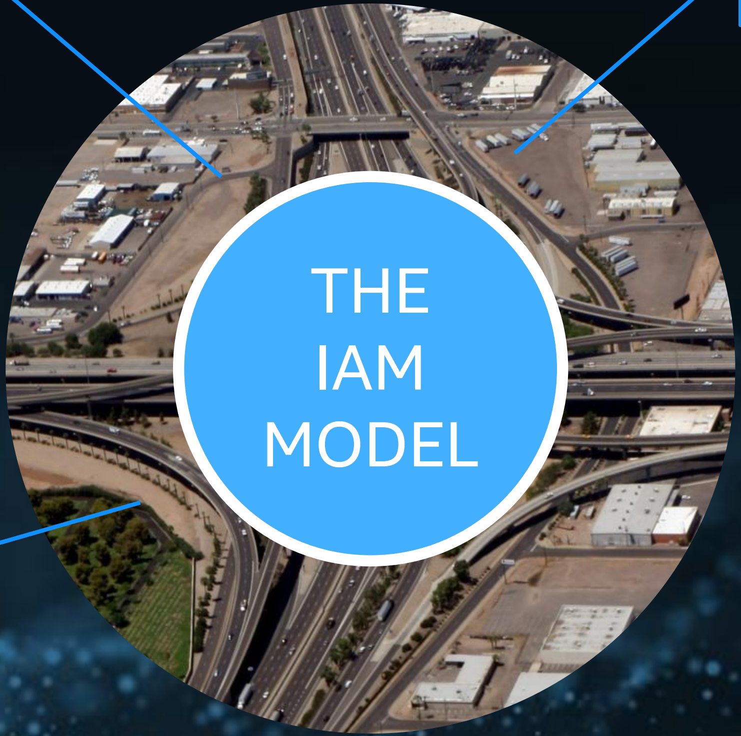

ACADEMIA

A consortium of **industry**, **academia**, and **government**

INSTITUTE FOR AUTOMATED MOBILITY

NORTHERN
ARIZONA
UNIVERSITY

RSS & THE INSTRUMENTED INTERSECTION

With the rise of intelligent sensors
& edge computing

we can enable intersections
to analyze driver behavior
using RSS

PROACTIVE REGULATION OF AV

An opportunity to get ahead of the curve

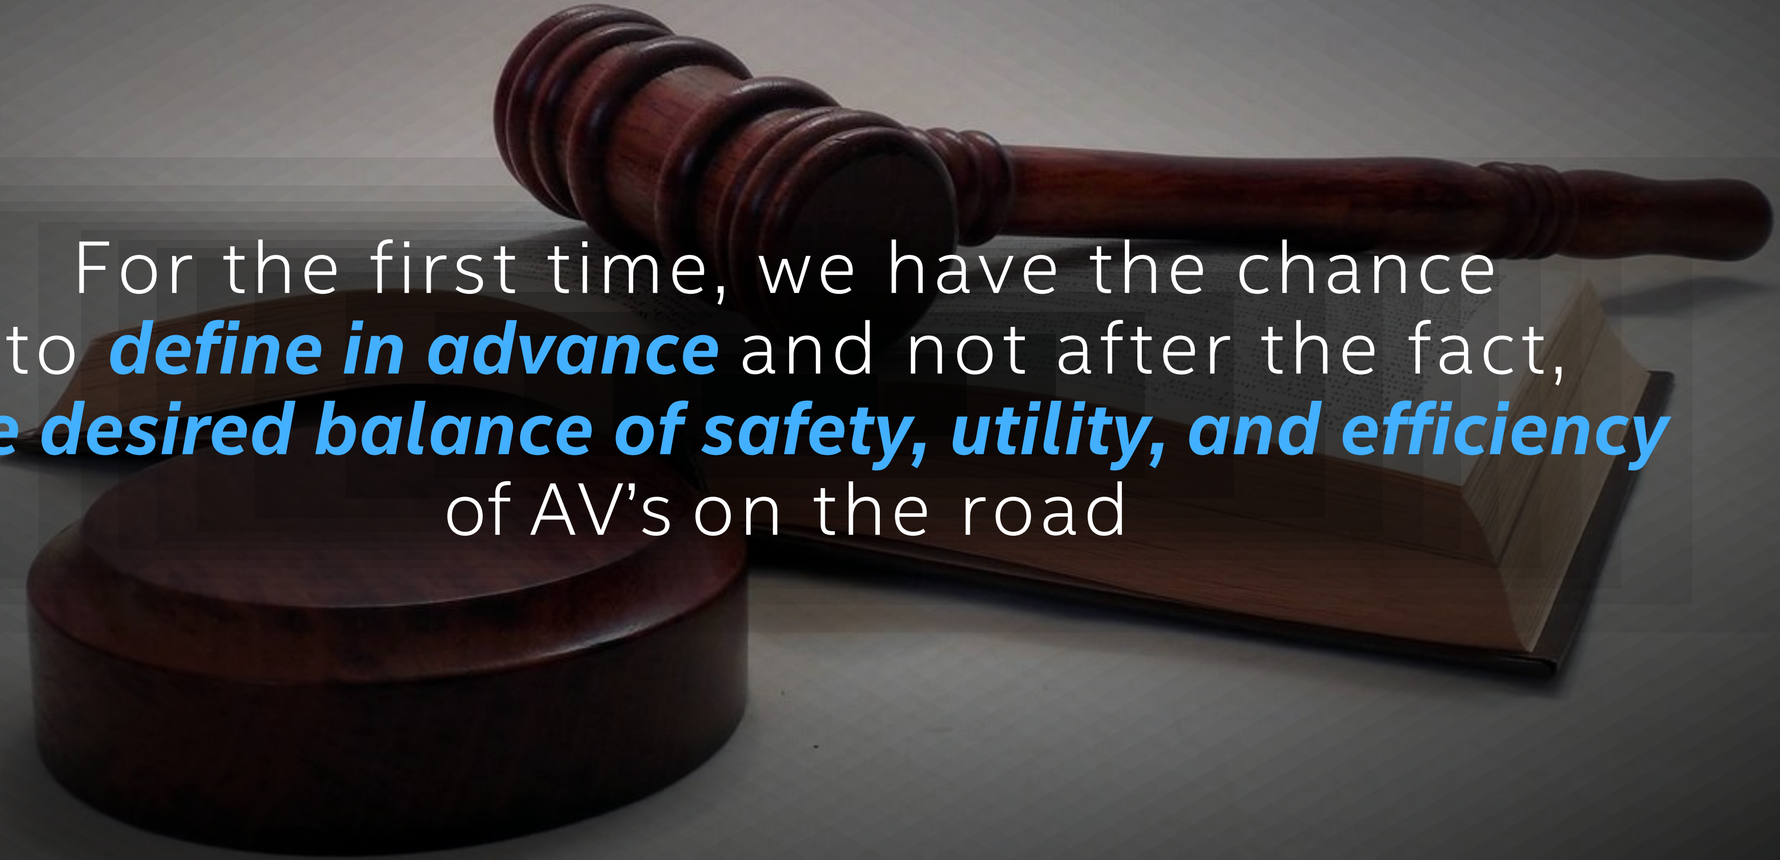

For the first time, we have the chance to ***define in advance*** and not after the fact, ***the desired balance of safety, utility, and efficiency*** of AV's on the road

RSS MOMENTUM

The model is gaining traction across the globe

"Our team recognizes the value and critical role that Mobileye's RSS model plays in safely deploying autonomous driving. ***Apollo platform will integrate RSS to successfully enable safe driving today, and drive further autonomous research on China's roadways.***"

– Weihao Gu
General Manager, Intelligent Driving Unit
Baidu

INTEL PARTNERS WITH BAIDU

To develop an RSS-based AV driving policy

Baidu 百度 | apollo

FCA
FIAT CHRYSLER AUTOMOBILES

intel

BMW

DAIMLER

infineon

APTIV

here

Continental

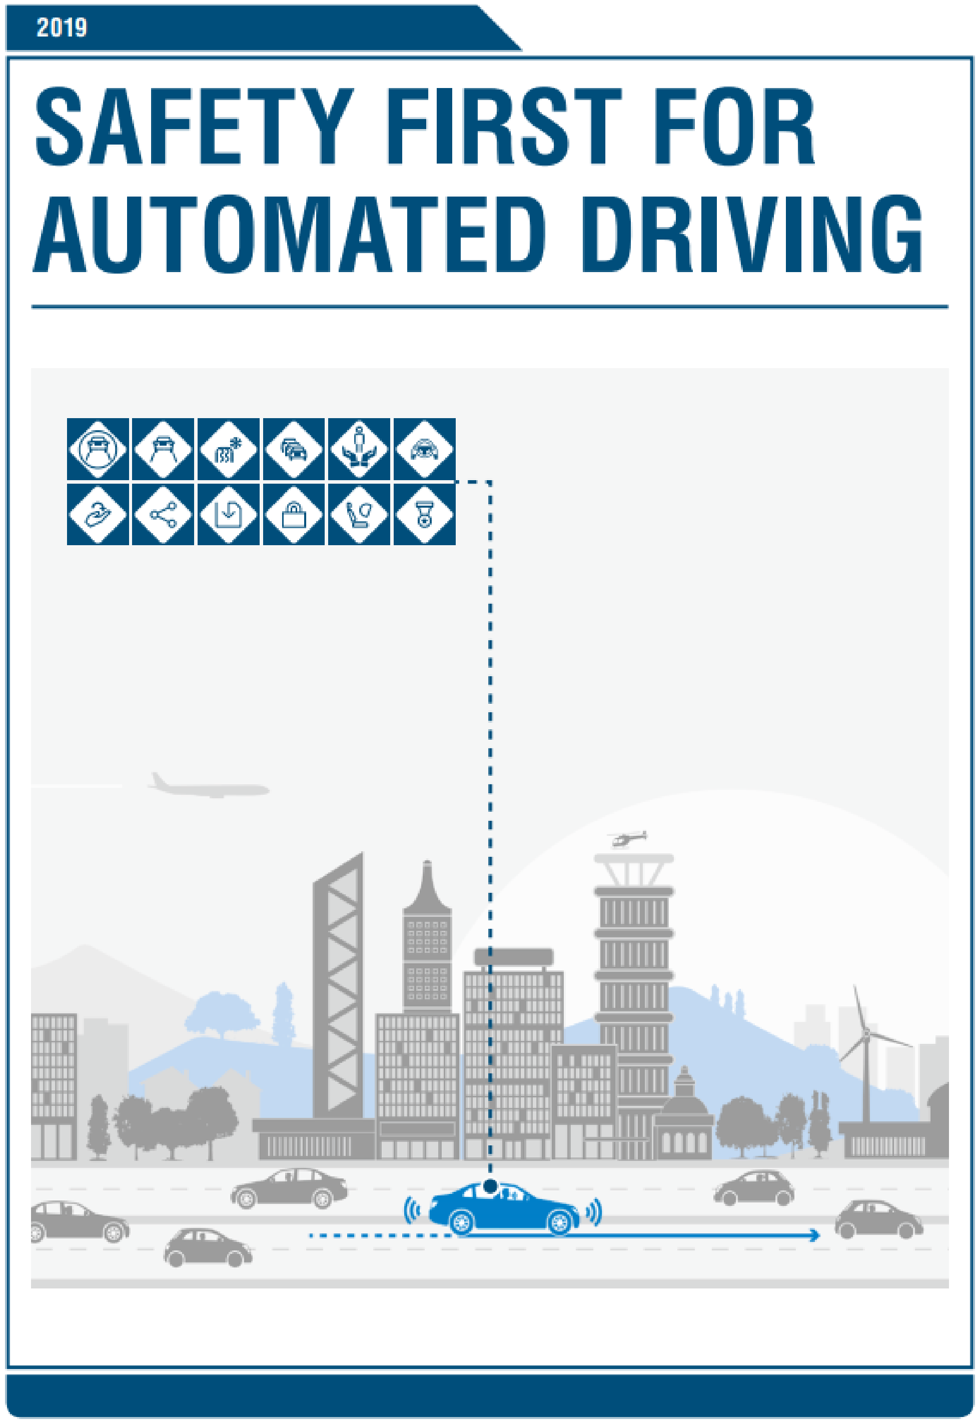

SAFECOMP PAPER

Autonomous Vehicles Meet the Physical World: RSS, Variability, Uncertainty, and Proving Safety

Philip Koopman, Beth Osyk, Jack Weast

Edge Case Research, Pittsburgh PA, USA
Intel, Chandler, AZ, USA
koopman@cmu.edu, bosyk@ecr.guru, jack.weast@intel.com

Abstract. The Responsibility-Sensitive Safety (RSS) model offers provable safety for vehicle behaviors such as minimum safe following distance. However, handling worst-case variability and uncertainty may significantly lower vehicle permissiveness, and in some situations safety cannot be guaranteed. Digging deeper into Newtonian mechanics, we illustrate complications that result from considering vehicle status, road geometry and environmental parameters. We propose a Micro-Operational Design Domain (μODD) approach to subdividing the operational space as a way of improving permissiveness. Confining probabilistic aspects of safety to μODD transitions permits proving safety (when possible) under the assumption that the system has transitioned to the correct μODD for the situation. Each μODD can additionally be used to encode system fault responses, take credit for advisory information (e.g., from vehicle-to-vehicle communication), and anticipate required responses for likely emergent situations. Finally, we augment the original RSS dmin equation to cover additional cases.

AND MORE...

Calibration and Evaluation of Responsibility-Sensitive Safety Model on Autonomous Car-Following Maneuvers Using Naturalistic Driving Study Data

Xiaoyan Xu
Graduate Research Assistant
College of Transportation Engineering, Tongji University
4800 Cao'an Road, Shanghai 201804, China

Safety Evaluation of Responsibility-Sensitive Safety (RSS) on Autonomous Car-Following Maneuvers Based on Surrogate Safety Measurements

Chen Chai, *Member, IEEE*, Xianming Zeng, Xiangbin Wu and Xuesong Wang*

A Situation-Aware Collision Avoidance Strategy for Car-Following

Li Li, *Fellow, IEEE*, Xinyu Peng, Fei-Yue Wang, *Fellow, IEEE*, Dongpu Cao, *Member, IEEE*, and Lingxi Li, *Senior Member, IEEE*

***Abstract*—In this paper, we discuss how to develop an appropriate collision avoidance strategy for car-following. This strategy aims to keep a good balance between traffic safety and efficiency while also taking into consideration the unavoidable uncertainty of position/speed perception/measurement of vehicles and other drivers. Both theoretical analysis and numerical testing results are provided to show the effectiveness of the proposed strategy.**

***Index Terms*—Collision avoidance, safety, traffic efficiency, uncertainty.**

Re
of Re
avoid
drivir
were
car-fc
avoid
releas
ever,

Towards Standardization of AV Safety Assurance: C++ Library for Responsibility Sensitive Safety

Bernd Gassmann[1], Fabian Oboril[1], Cornelius Buerkle[1], Shuang Liu[1], Shoumeng Yan[1], Maria Soledad Elli[2], Ignacio Alvarez[1], Naveen Aerrabotu[2], Suhel Jaber[2], Peter van Beek[2], Darshan Iyer and Jack Weast[2]

***Abstract*— The need for safety assurances in Automated Driving (AD) is becoming increasingly critical with the accelerating deployment of this technology. Beyond functional safety, indus-**

RSS is a technology-neutral model for safety that can be used to define and measure whether an AV is driving safely. *RSS* formalizes an interpretation of "common sense" and defines what it means for an AV to drive safely on its own and how it should exercise reasonable caution to protect against the unsafe driving behavior of others. The paper presented *RSS* as a mathematical model that formalizes this interpretation for automated driving vehicles and aims to satisfy the need for sound (i.e. law-abiding), useful (vs. overly conservative) and efficiently verifiable driving policies for automated driving.

As highlighted by [5], *RSS* contributes to the overall safety of automated driving vehicles in the operational safety domain (complimentary to the functional safety domain). Advanced automated driving systems with capabilities beyond L3+, as described in [6], require significant investments in operational safety, in particular in the areas of scenario development and formal verification, testing and validation tools. Recent contributions in these areas such as [7] have expanded *RSS* with formalized components of dynamics and policies and highlight limitations of existing tools towards developing an automatic formal verification framework. A

2019 IEEE Intelligent Vehicles Symposium (IV)
Paris, France. June 9-12, 2019

Specifying Safety of Autonomous Vehicles in Signal Temporal Logic

Nikos Aréchiga[1]

***Abstract*— We develop a set of contracts for autonomous control software that ensures that if all traffic participants follow the contracts, the overall traffic system will be collision-free. We express our contracts in Signal Temporal Logic (STL), a lightweight specification language that enables V&V methodologies. We demonstrate how the specification can be used for evaluation of the performance of autonomy software, and We provide preliminary evidence that our contracts are not excessively conservative, i.e., they are not more restrictive than existing guidelines for safe driving by humans.**

The work of [12] develops a planner that generates "maneuver automata", and uses the theorem prover Isabelle to ensure that the maneuver automata satisfy specifications given in Linear Temporal Logic (LTL). Since LTL is primarily used to reason about discrete transition over automata, an additional reachable state computation is required to ensure that the continuous dynamics respect the high-level specification. In contrast, STL supports both discrete and continuous reasoning, which means that we do not need to

Announcing…

C++11 RSS LIBRARY

Standalone **Open Source Library** currently covering a subset of RSS rules (with development ongoing)

1. Longitudinal scenarios
   - Same and opposite direction
2. Lateral scenarios & Multilane roads
3. Intersection handling

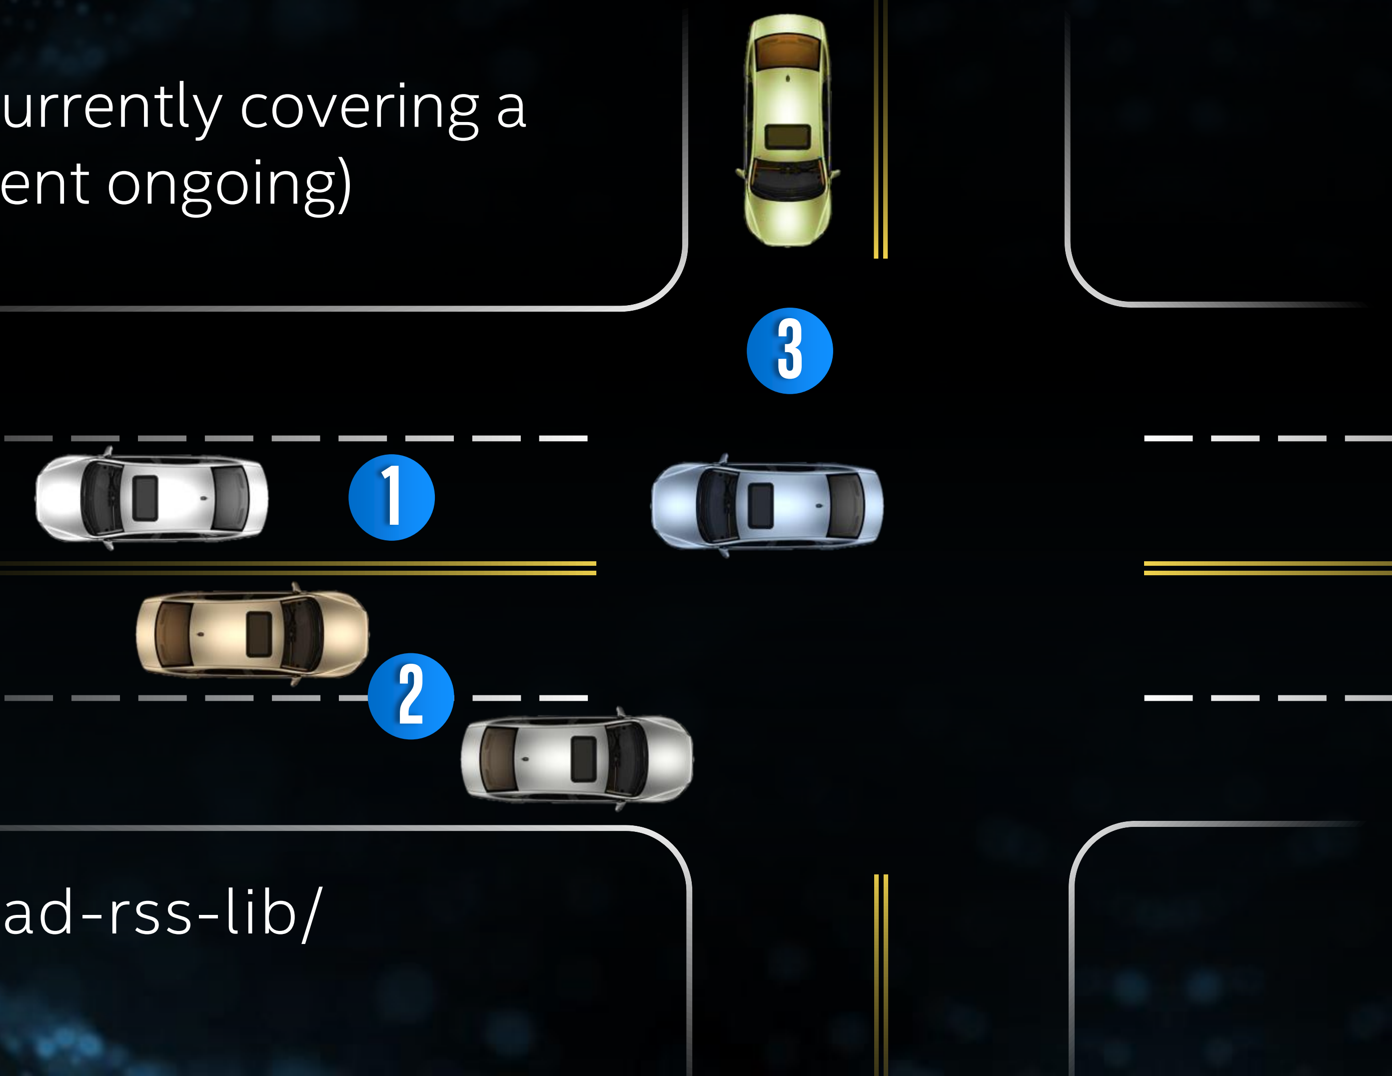

https://intel.github.io/ad-rss-lib/

C++ RSS LIBRARY OVERVIEW

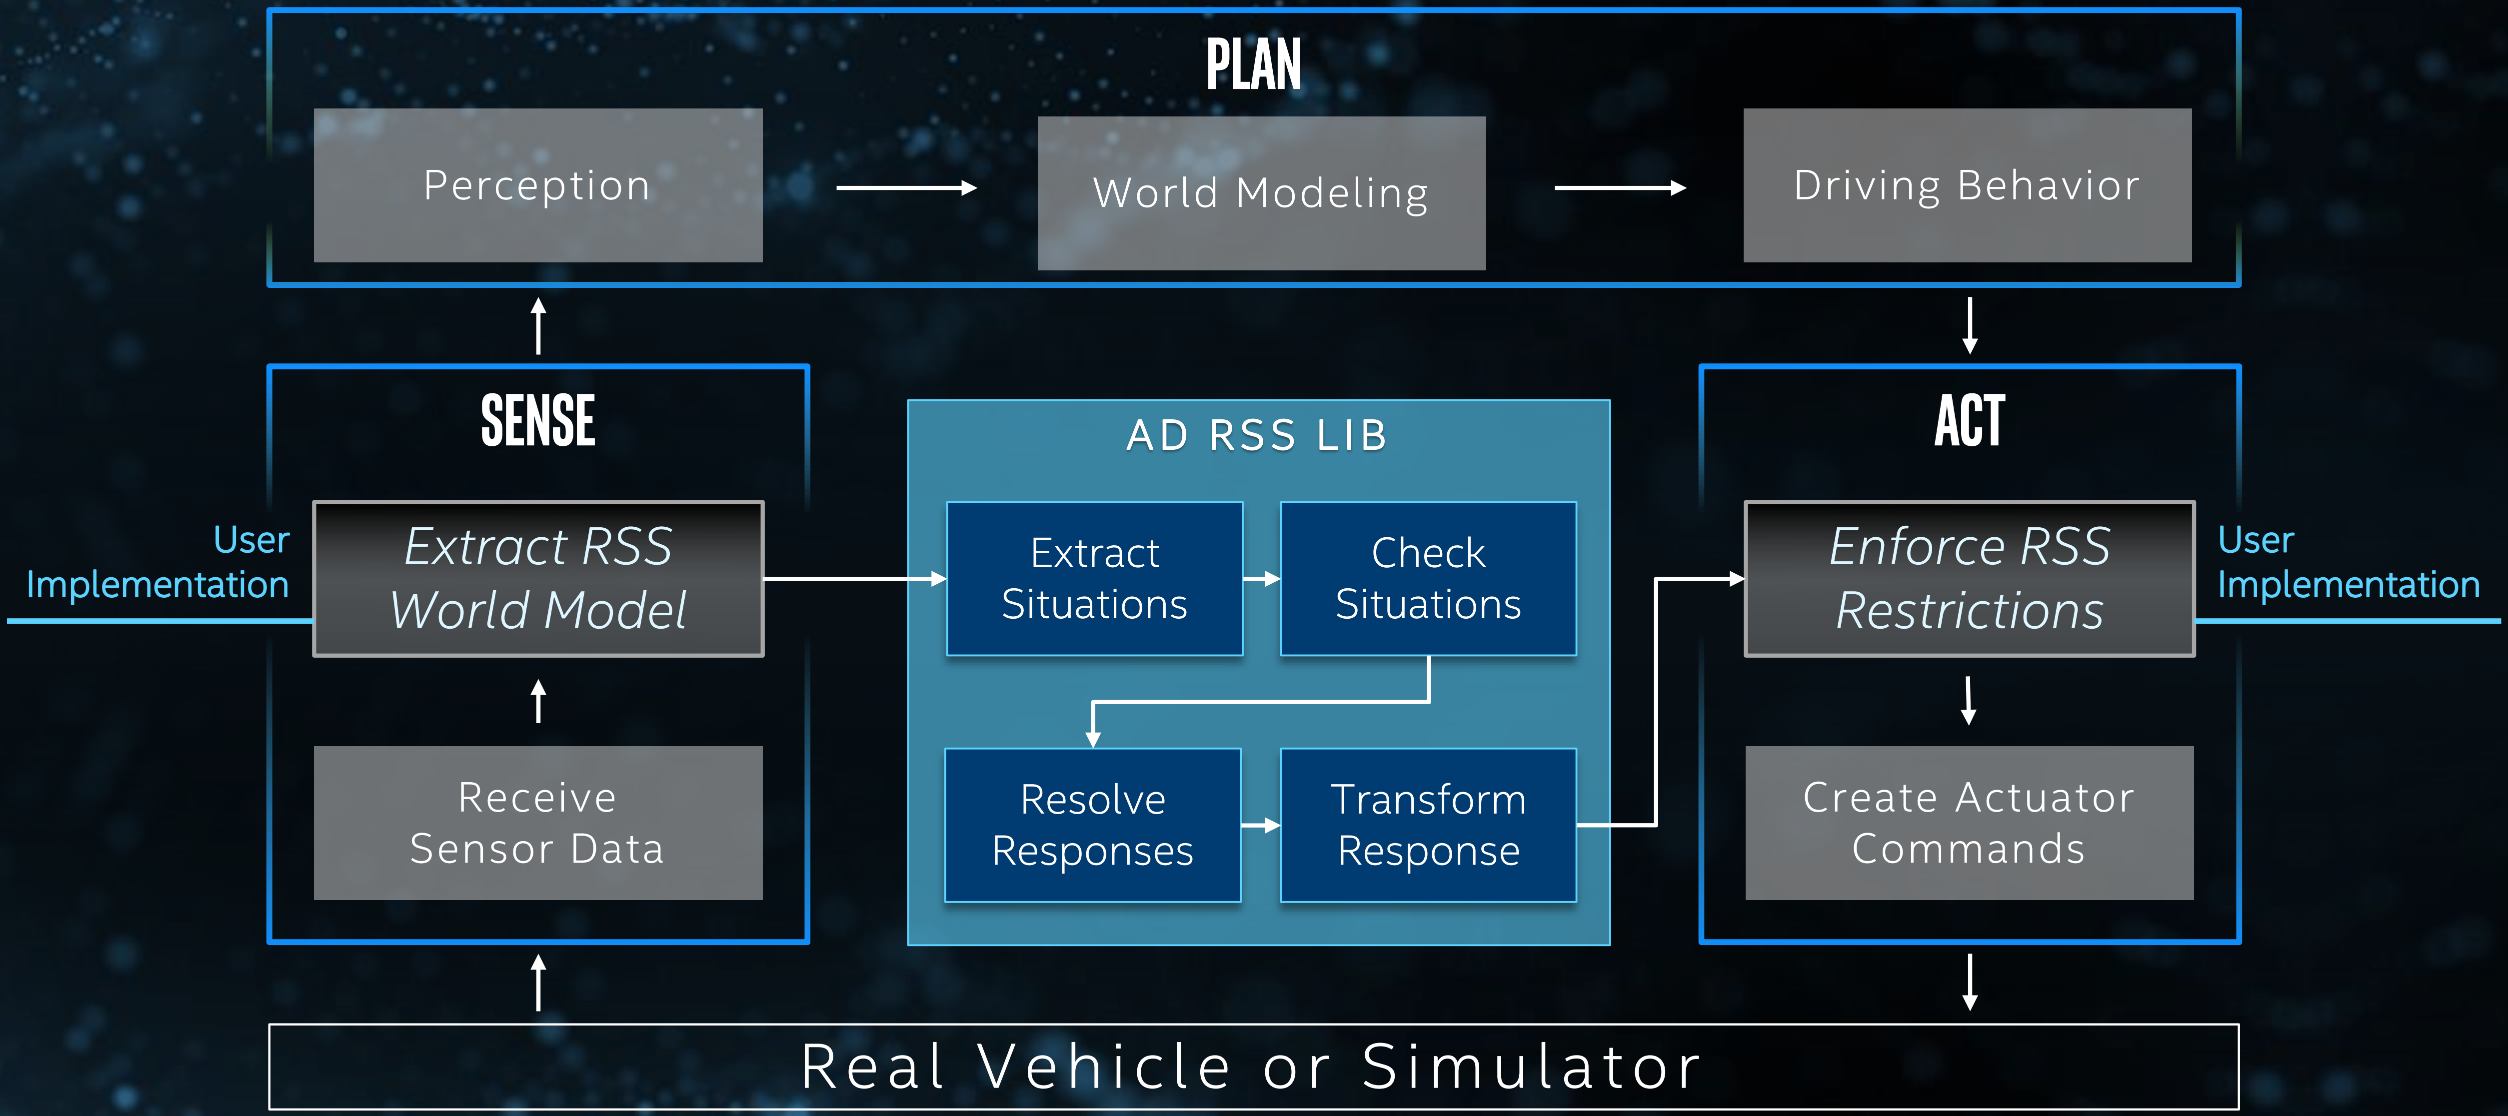

RSS LIBRARY & CARLA

NHTSA PRE-CRASH SCENARIOS IN CARLA

Scenario: red light / stop sign at T-Junction

V_2 has a stop sign…
And runs right through it

Using RSS, V_1 analyzes vehicle telemetry, identifies the danger, and avoids a crash

AV SAFETY: AN ISSUE LARGER THAN ONE COMPANY

What are we doing

INDUSTRY

Engaging with customers, competitors and consortia to have an open dialogue on AV safety

GOVERNMENT / NGO'S

Understanding government expectations on transparency and verification of AV safety

ACADEMIA

RSS Research Centers at Universities in USA, PRC and EU

REAL WORLD

Deploying RSS in our on AV Fleet in very challenging environments

An Intel
Company