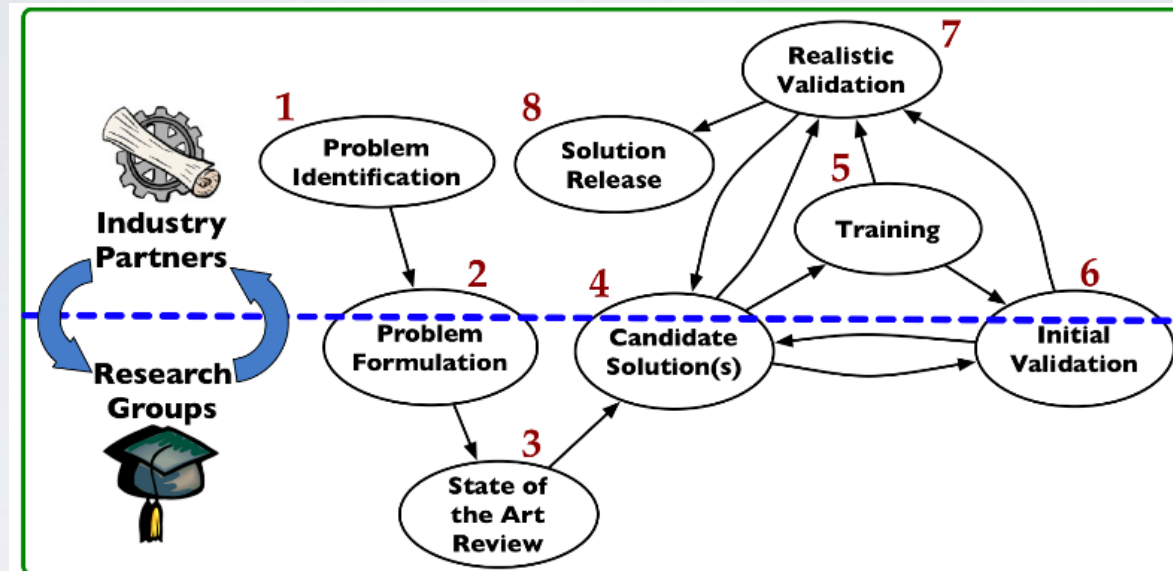


Automated Testing of Autonomous Driving Assistance Systems

Lionel Briand

SEMLA, Montreal, 2018

Collaborative Research @ SnT



SnT
securityandtrust.lu

- Research in context
- Addresses actual needs
- Well-defined problem
- Long-term collaborations
- Our lab is the industry



Software Verification and Validation @ SnT Centre

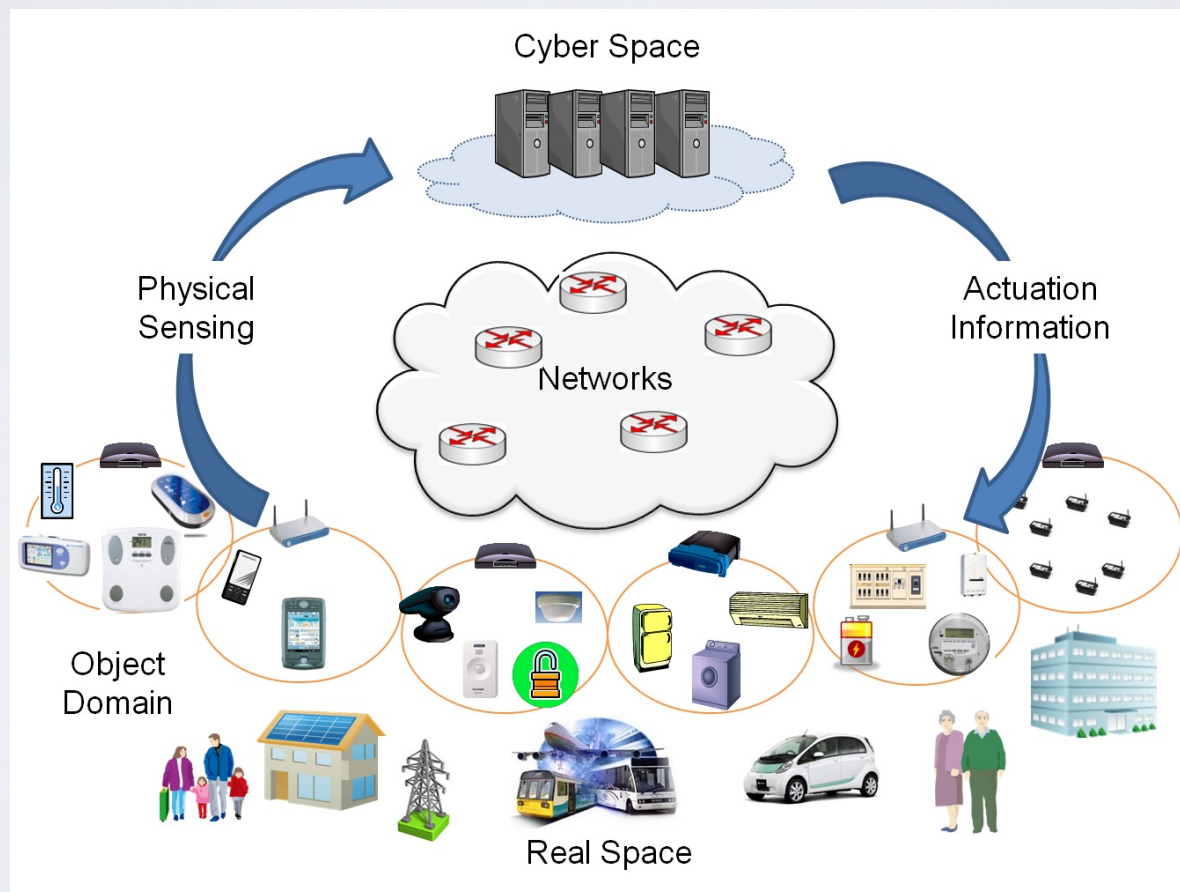
- Group established in **2012**
- **Focus:** Automated, novel, cost-effective V&V solutions
- ERC Advanced Grant
- ~ 25 staff members
- Industry and public partnerships



Introduction

Cyber-Physical Systems

- A system of collaborating computational elements controlling physical entities



Autonomous CPS

- Read sensors, i.e., collect data about their environment
- **Make predictions about their environment**
- **Make (optimal) decisions about how to behave to achieve some objective(s) based on predictions**
- Send commands to actuators according to decisions
- Often mission or safety critical

Advanced Driver Assistance Systems (ADAS)



Automated Emergency Braking (AEB)



Lane Departure Warning (LDW)



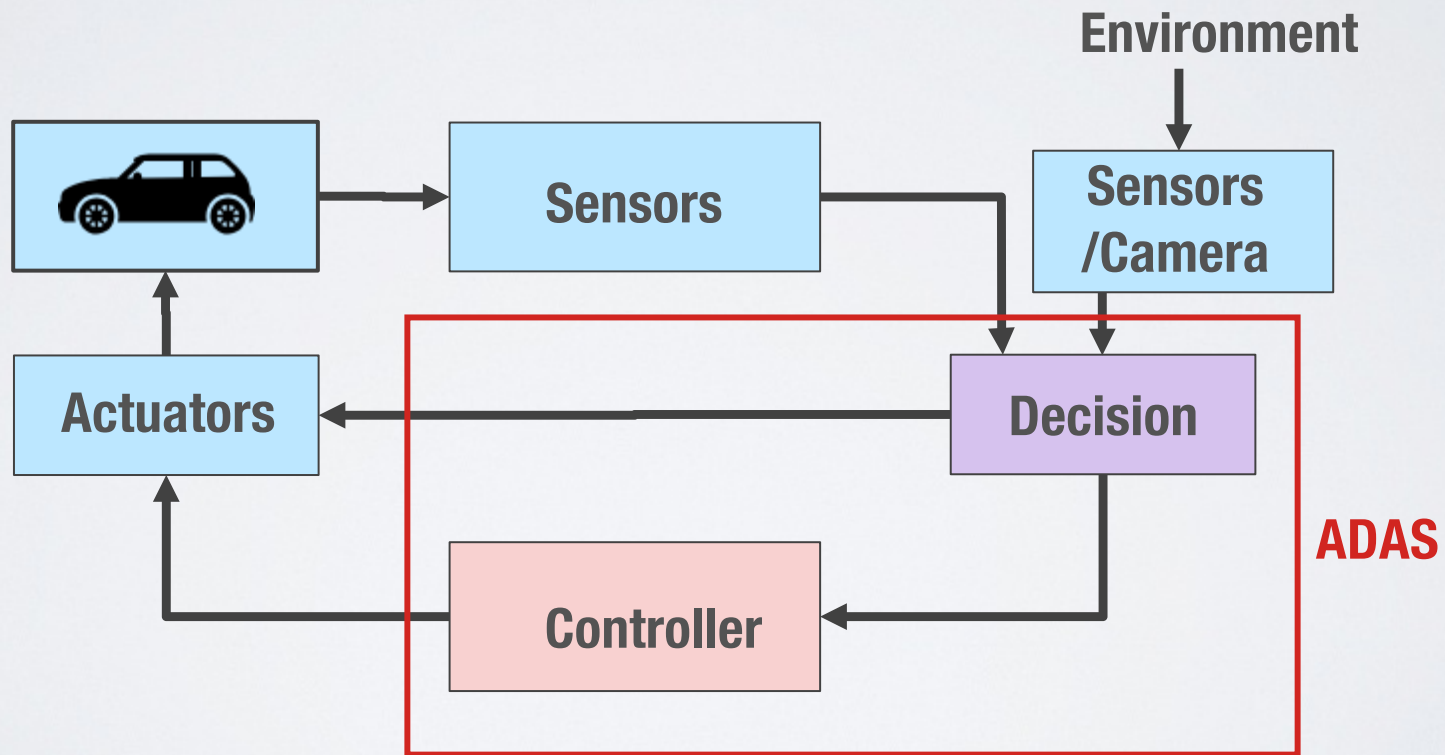
Pedestrian Protection (PP)



Traffic Sign Recognition (TSR)

Advanced Driver Assistance Systems (ADAS)

Decisions are made over time based on sensor data



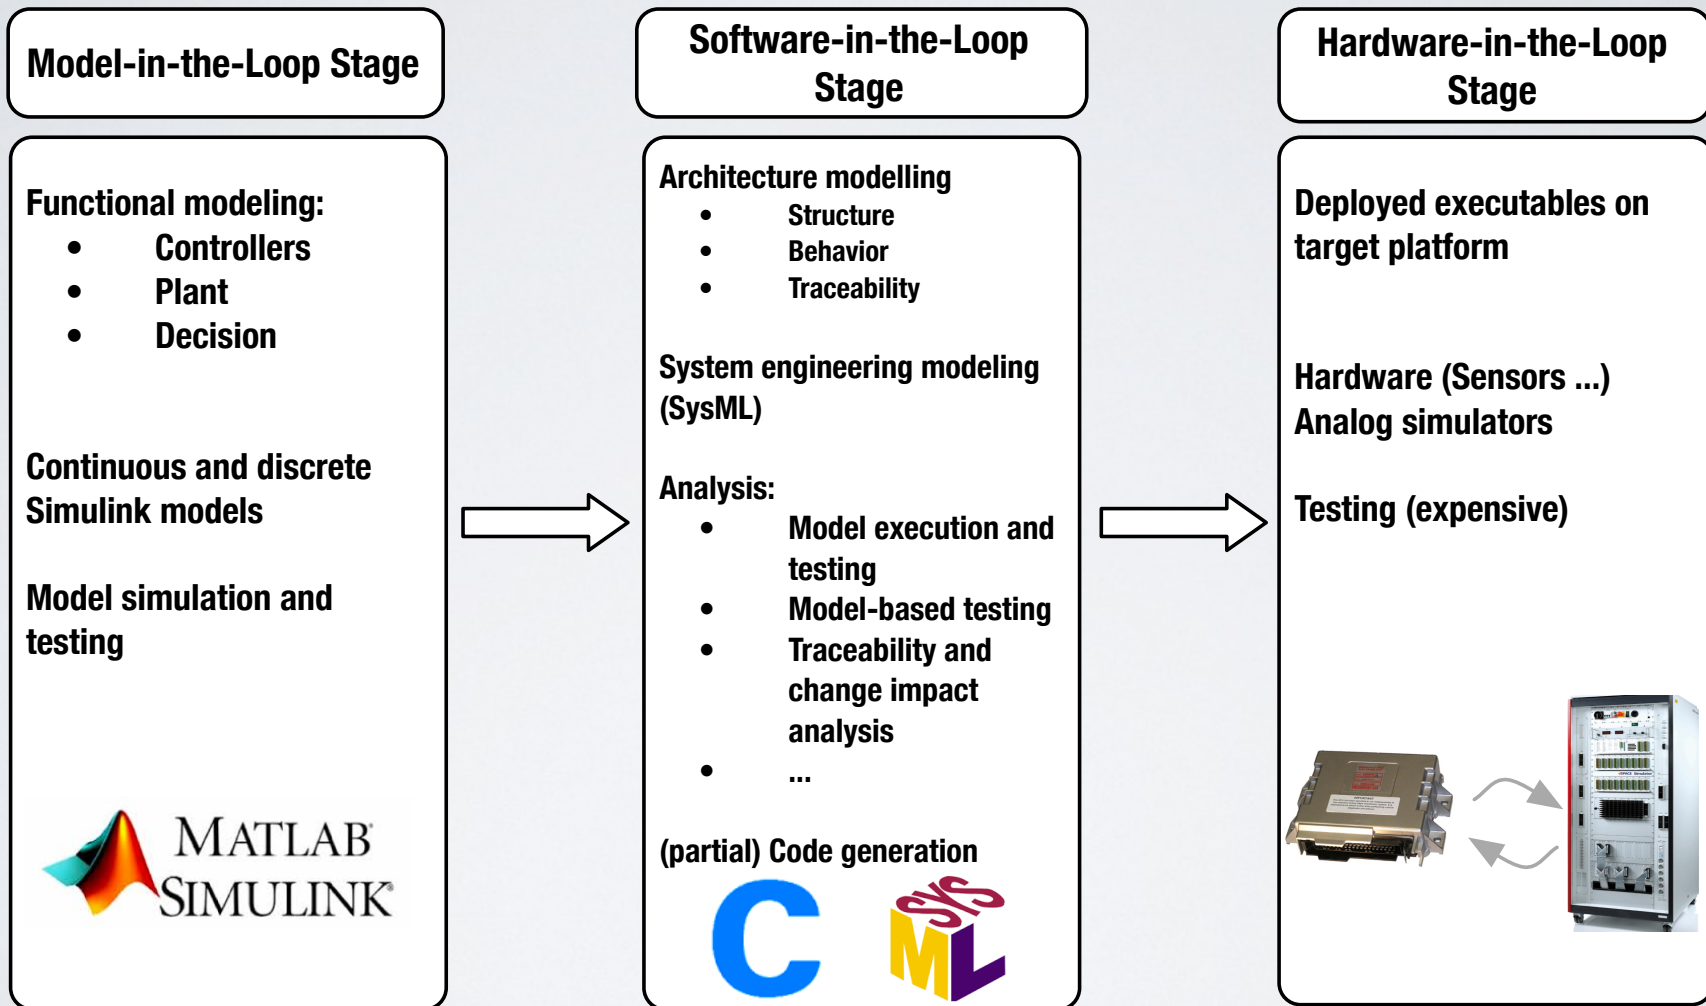
A General and Fundamental Shift

- Increasingly so, it is easier to **learn behavior from data** using machine learning, rather than specify and code
- Deep learning, reinforcement learning ...
- **Example: Neural networks (deep learning)**
- Millions of weights learned
- No explicit code, no specifications
- **Verification, testing?**

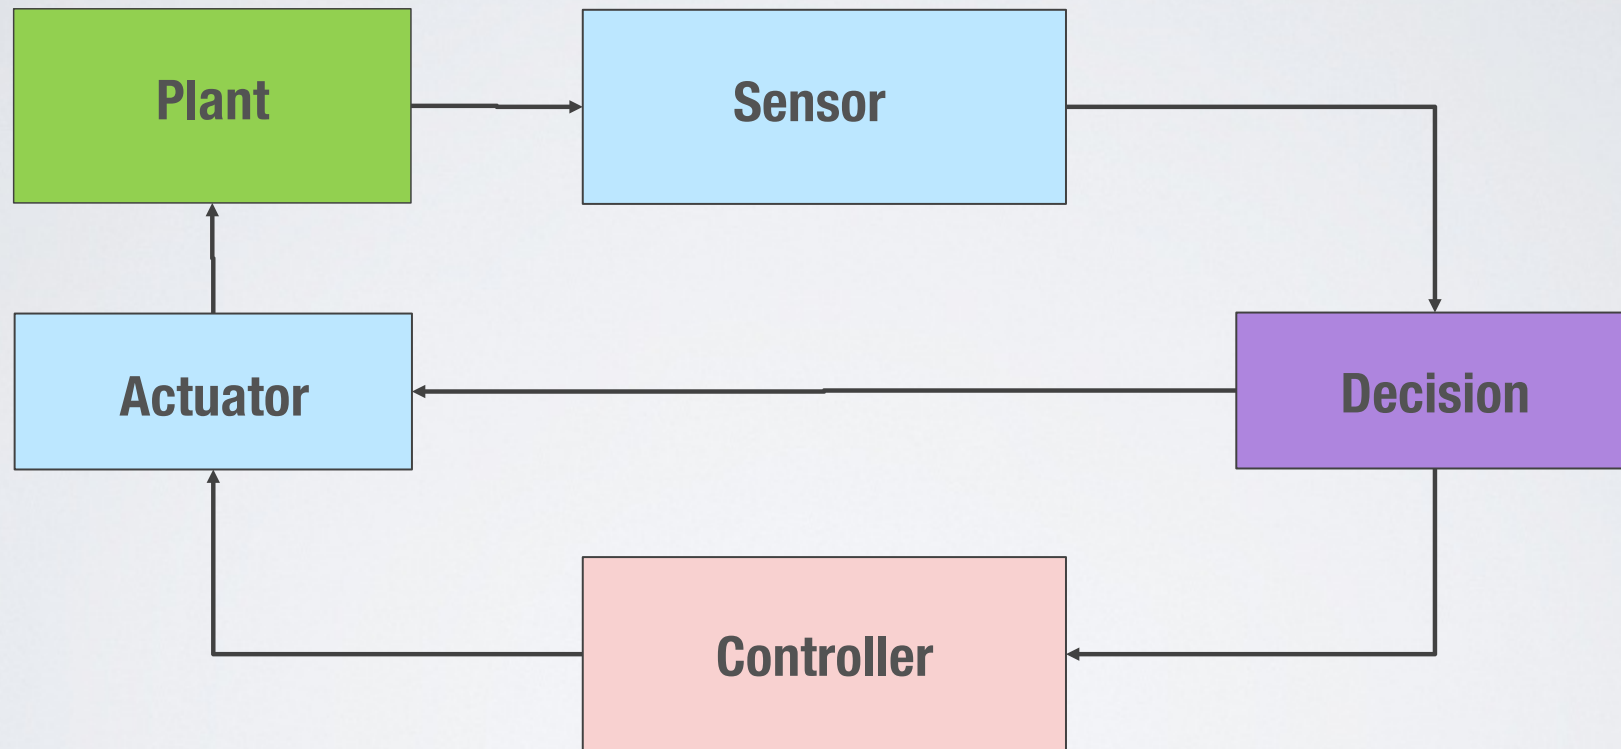
Testing Implications

- **Test oracles?** No explicit, expected test behavior
- **Test completeness?** No source code, no specification

CPS Development Process



MiL Components



Opportunities and Challenges

- Early functional models (MiL) offer **opportunities for early functional verification and testing**
- But a challenge for constraint solvers and model checkers:
 - **Continuous mathematical models**, e.g., differential equations
 - Discrete software models for code generation, but with **complex operations**
 - Library functions in **binary code**

Automotive Environment


- **Highly varied environments, e.g., road topology, weather, building and pedestrians ...**
- **Huge number of possible scenarios, e.g., determined by trajectories of pedestrians and cars**
- **ADAS play an increasingly critical role**
- **A challenge for testing**

Testing Advanced Driver Assistance Systems



Objective

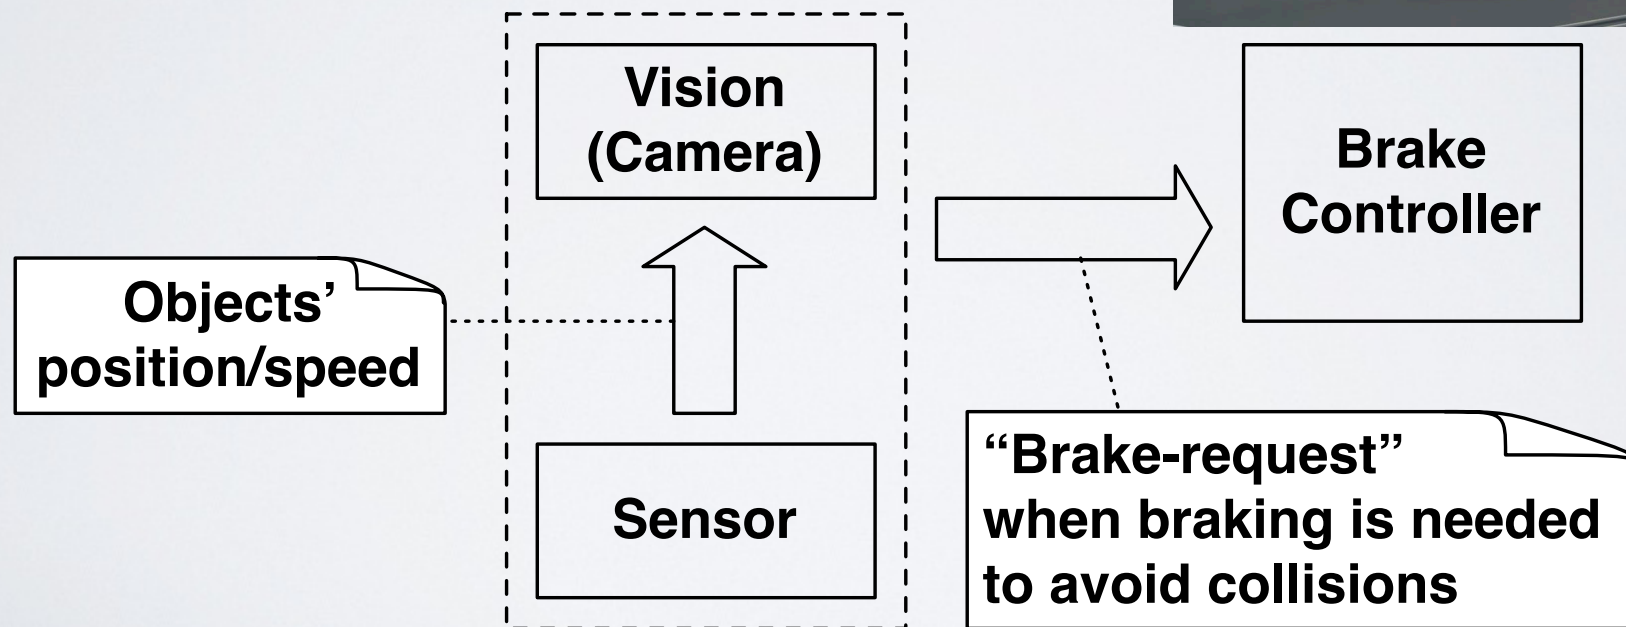
- **Testing ADAS**
 - **Identify and characterize most critical/risky scenarios**
 - **Test oracle: Safety properties**
 - **Need scalable test strategy due to large input space**



Automated Emergency Braking System (AEB)



Decision making



Example Critical Situation

“AEB detects a pedestrian in front of the car with a high degree of certainty, but an accident happens where the car hits the pedestrian with a relatively high speed”



Testing ADAS

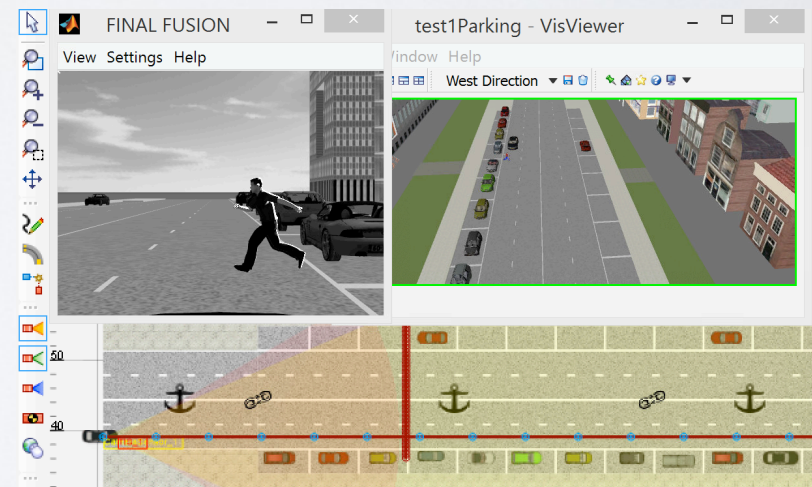
On-road testing



Time-consuming

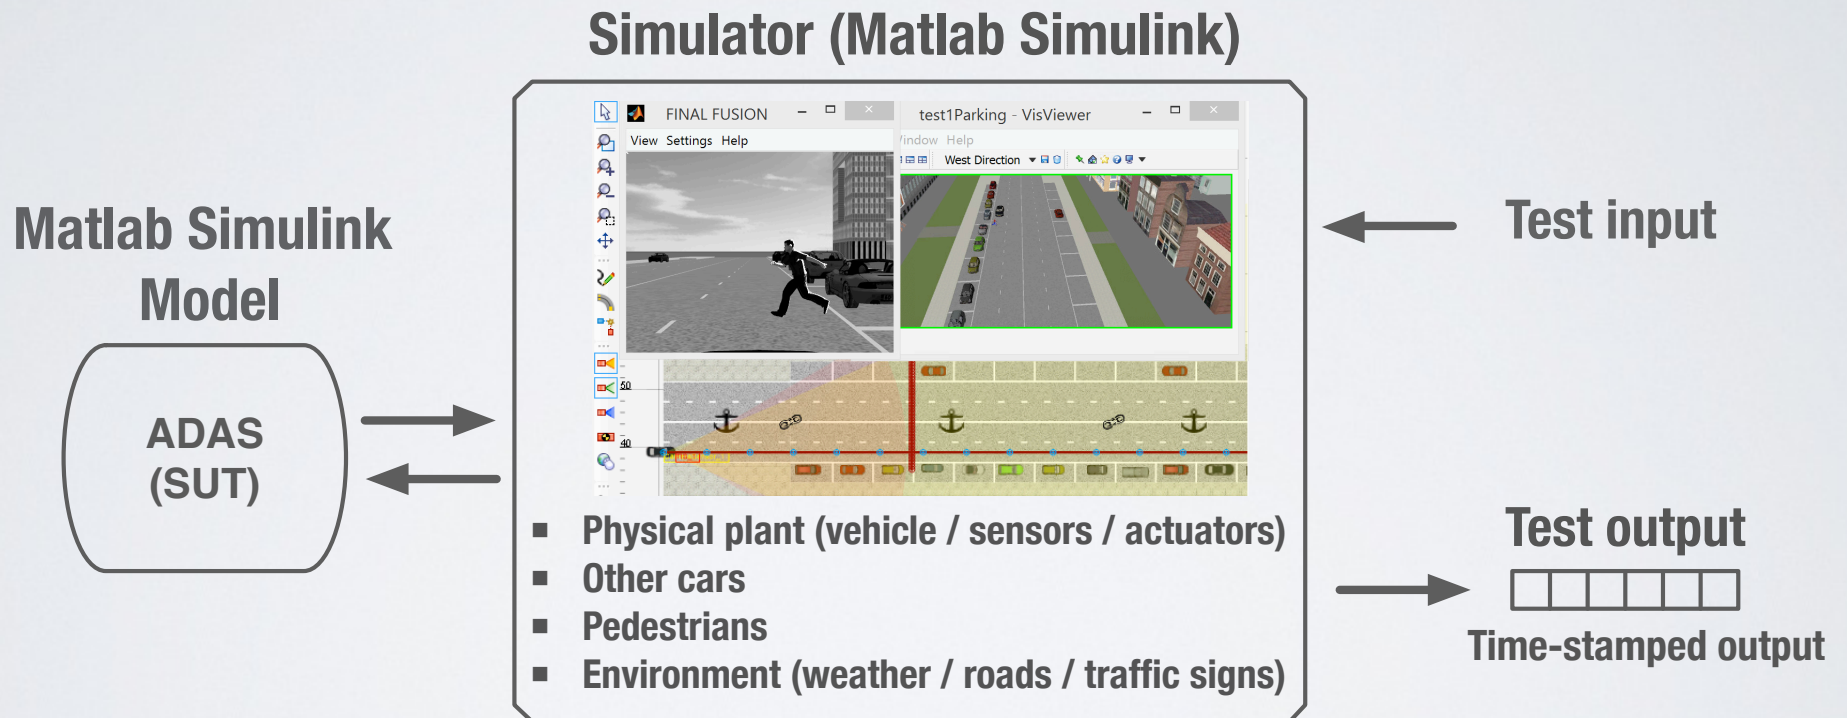
Expensive

Simulation-based (model) testing



**A simulator based on
physical/mathematical models**

Model Testing ADAS



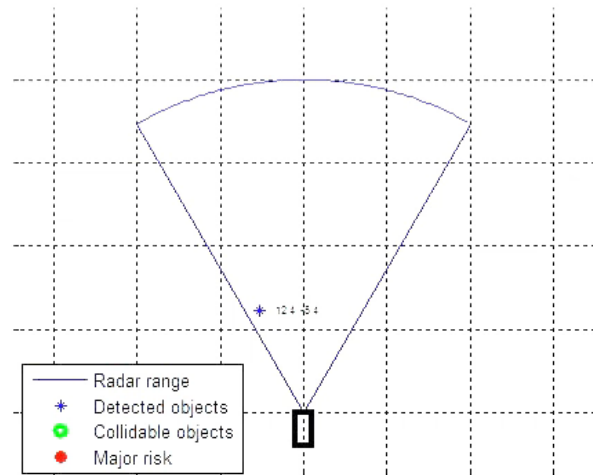
Physics-Based Simulations

FINAL FUSION

View Settings Help



PPS Radar



Experiment7 - VisViewer

File Window Help

South Direction



Our Goal

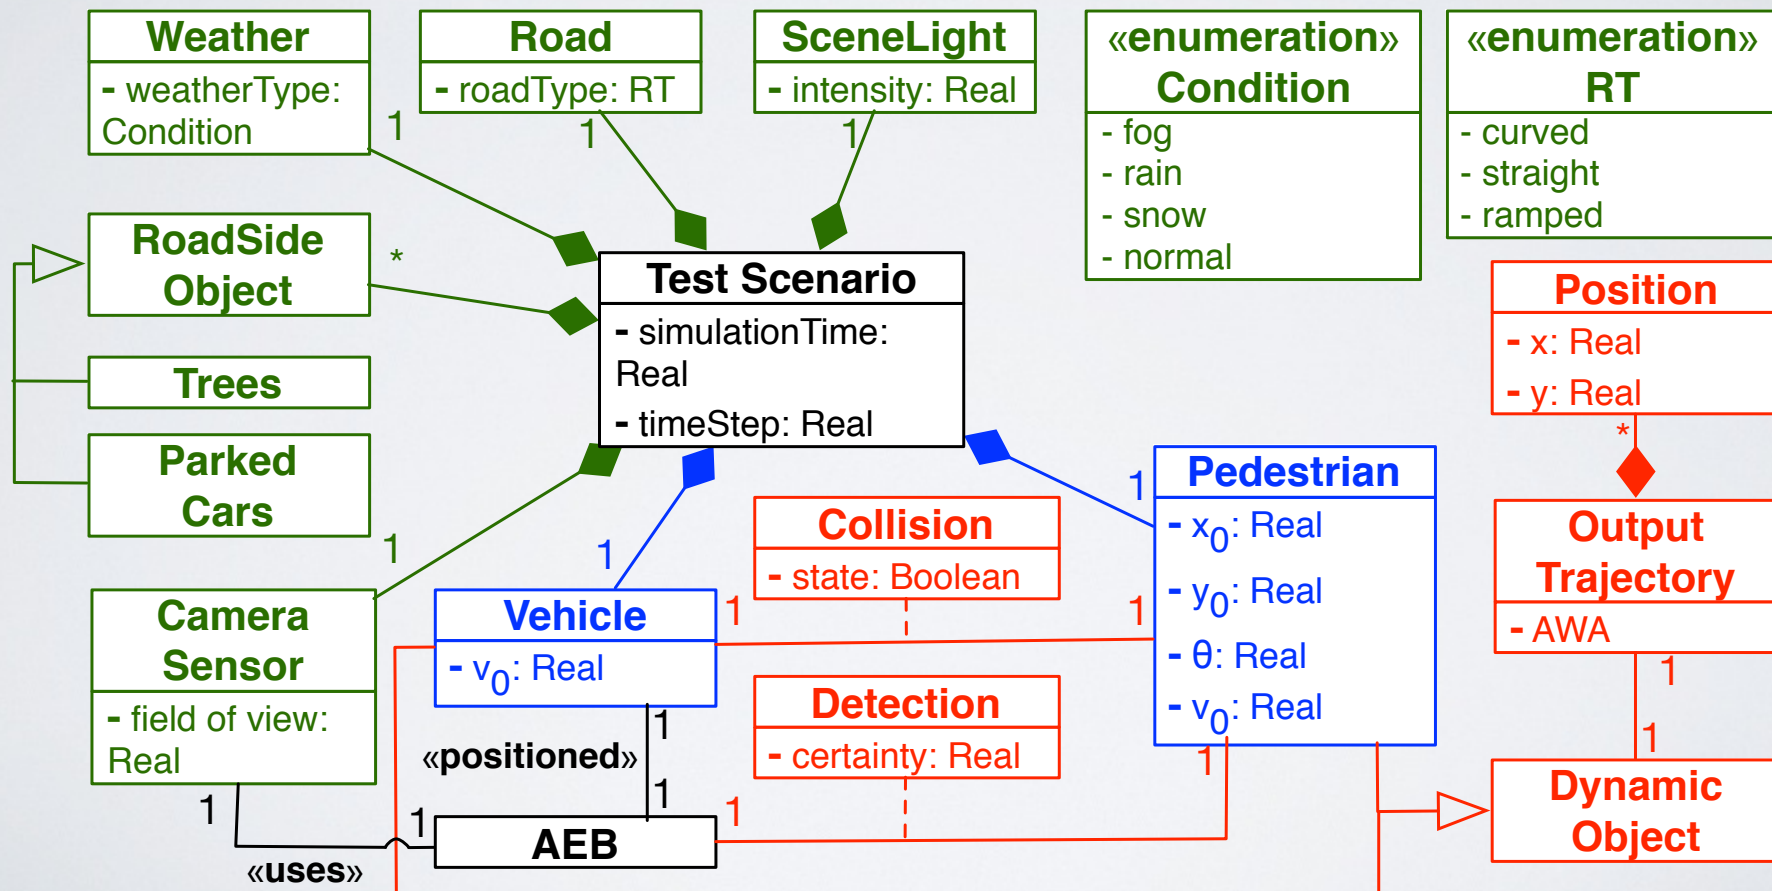
- **Developing an automated testing technique for ADAS**
 - To help engineers efficiently and effectively **explore** the complex test input space of ADAS
 - To **identify** critical (failure-revealing) test scenarios
 - **Characterization of input conditions** that lead to most critical situations

ADAS Testing Challenges

- Test input space is **large, complex** and **multidimensional**
- **Explaining failures and fault localization** are difficult
- Execution of **physics-based simulation models** is computationally expensive

Test Inputs/Outputs

Environment inputs
 Mobile object inputs
 Outputs



Our Solution: Learnable Evolutionary Algorithms




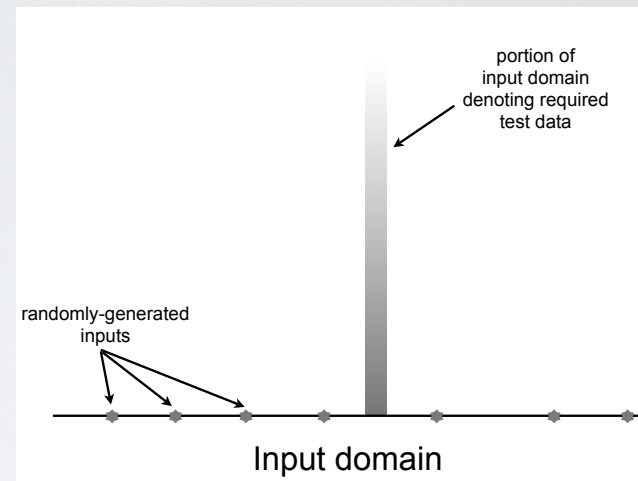
Learn regions likely to contain most critical (failure) test scenarios

Search for critical test scenarios in the critical regions, and help refine classification models

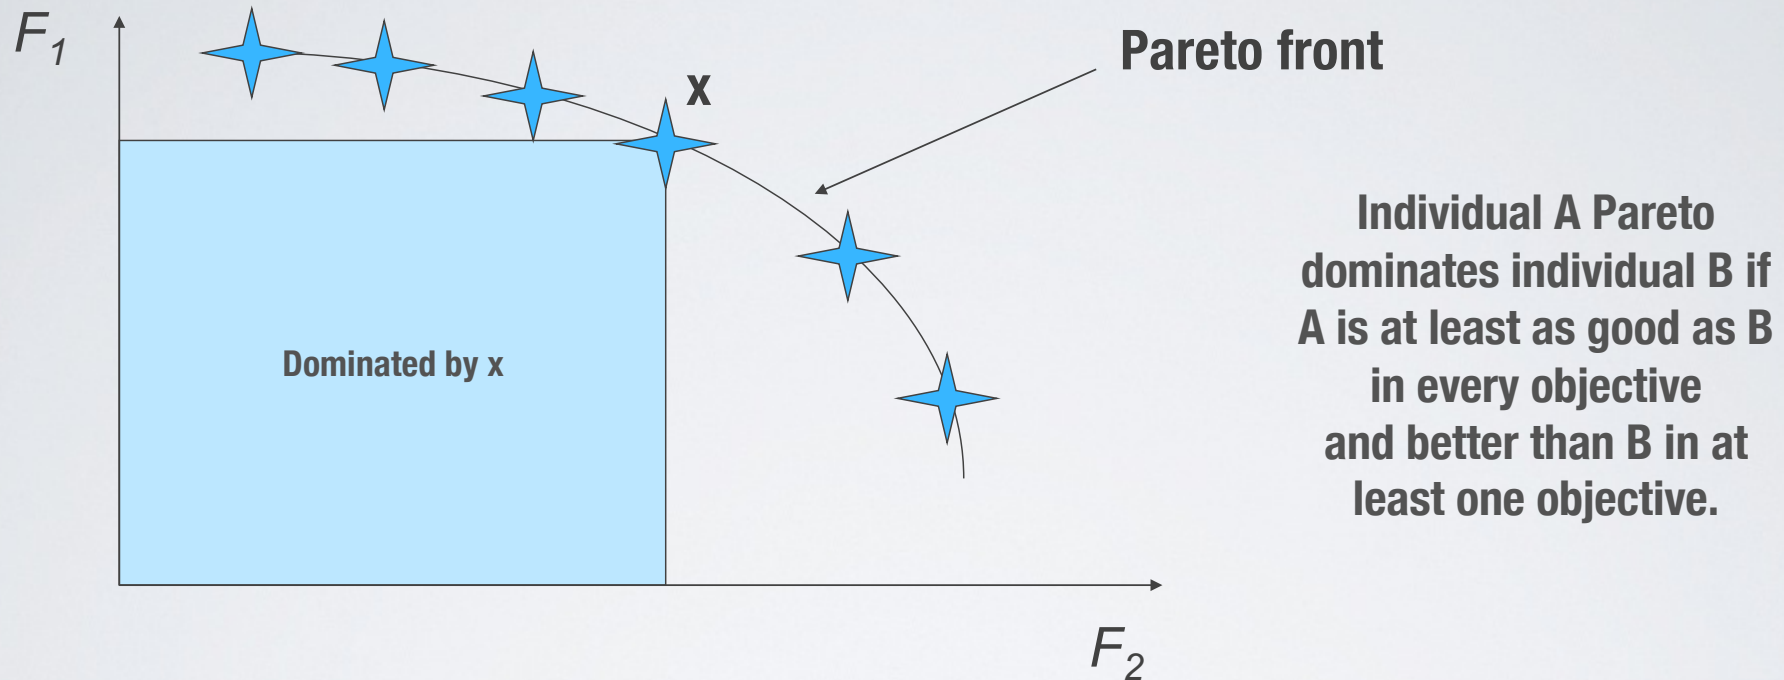
→ Machine-learning classification models are used to characterize failures and guide the search towards critical test scenarios faster

Search-Based Software Testing

- Express test generation problem as a **search problem**
- Search for **test input data** with certain properties, i.e., constraints
- **Non-linearity** of software (if, loops, ...): complex, discontinuous, non-linear search spaces (Baresel)
- Many search algorithms (**metaheuristics**), from local search to global search, e.g., Hill Climbing, Simulated Annealing and Genetic Algorithms



Multiple Objectives: Pareto Front



- A multi-objective optimization algorithm (e.g., NSGA II) must:
 - Guide the search towards the global Pareto-Optimal front.
 - Maintain solution diversity in the Pareto-Optimal front.

Our ADAS Testing

- We use **decision tree classification models**
- We use **multi-objective** search algorithm (NSGAI)
- **Objective Functions:**
 1. Minimum distance between the pedestrian and the field of view
 2. The car speed at the time of collision
 3. The probability that the object detected is a pedestrian
- Each search iteration **calls simulation** to compute objective functions
- Input values required to perform the simulation:

Precipitation

Fogginess

Road shape

Visibility range

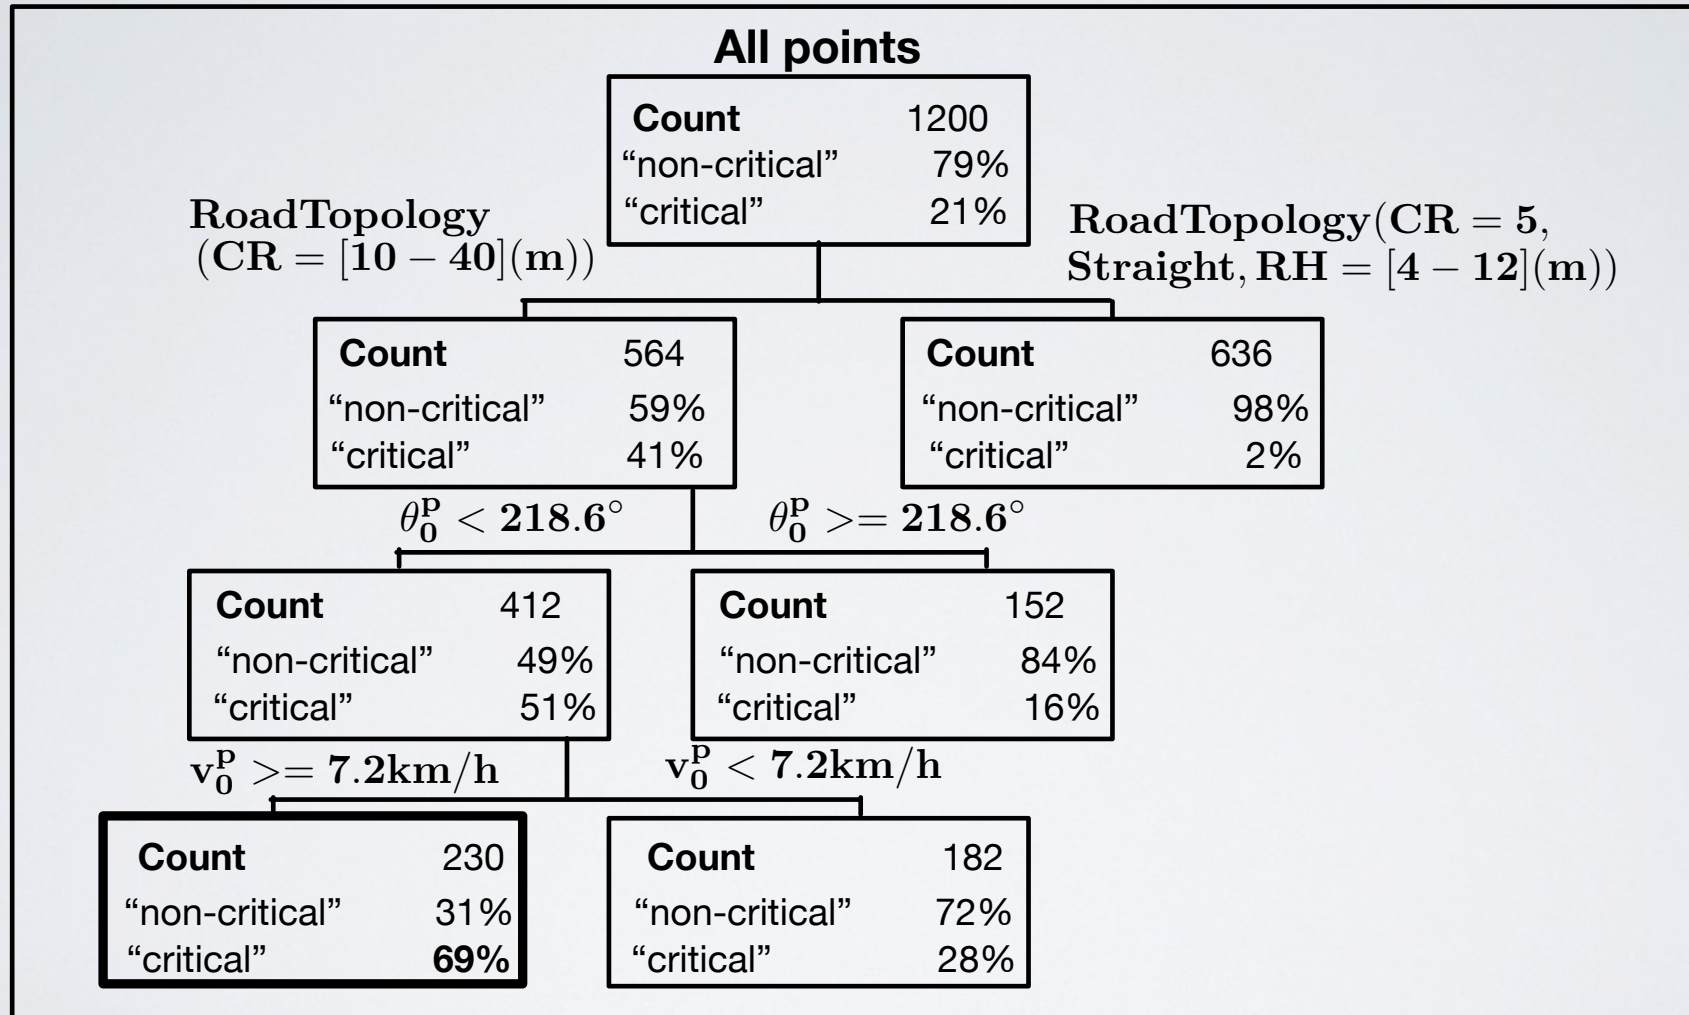
Car-speed

Person-speed

Person-position


Person-orientation

Decision Trees



Partition the input space into homogeneous regions

Genetic Evolution guided by Classification




- Initial input ✓
- Fitness computation ✓
- Classification ✓
- Selection ✓
- Breeding

NSGAI-DT


1. Generate an initial representative set of input scenarios and run the simulator to label each scenario as critical or non-critical

2. Build a decision tree model



3. Run the NSGAI search algorithm for the elements inside each critical leaf

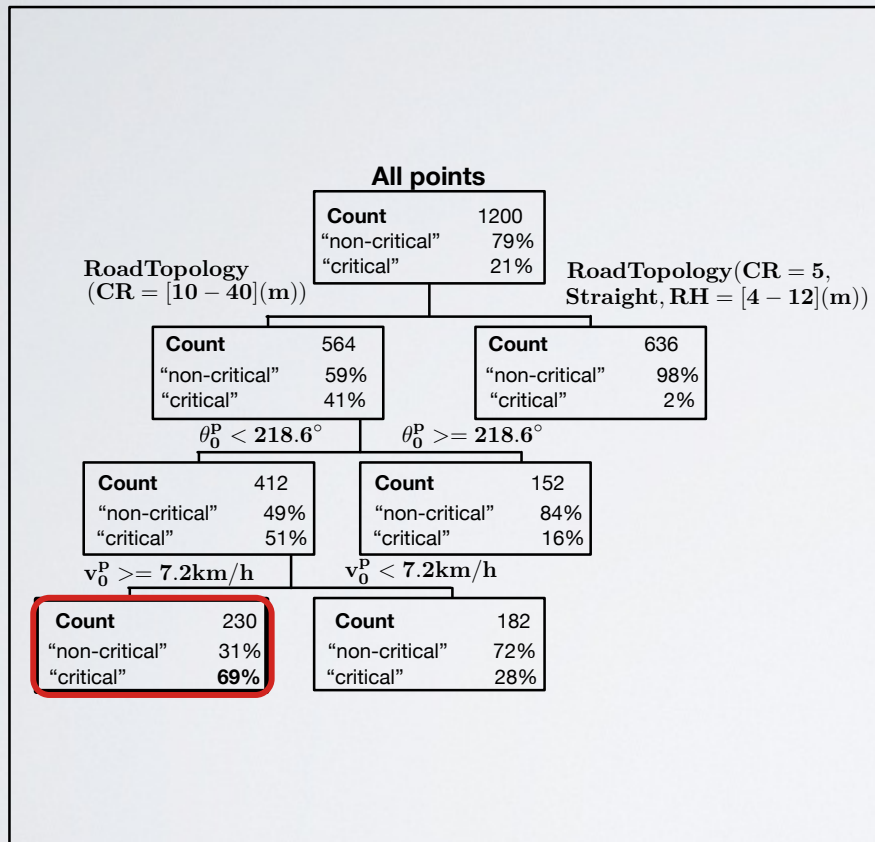
4. Rebuild the decision tree (step 2) or stop the process



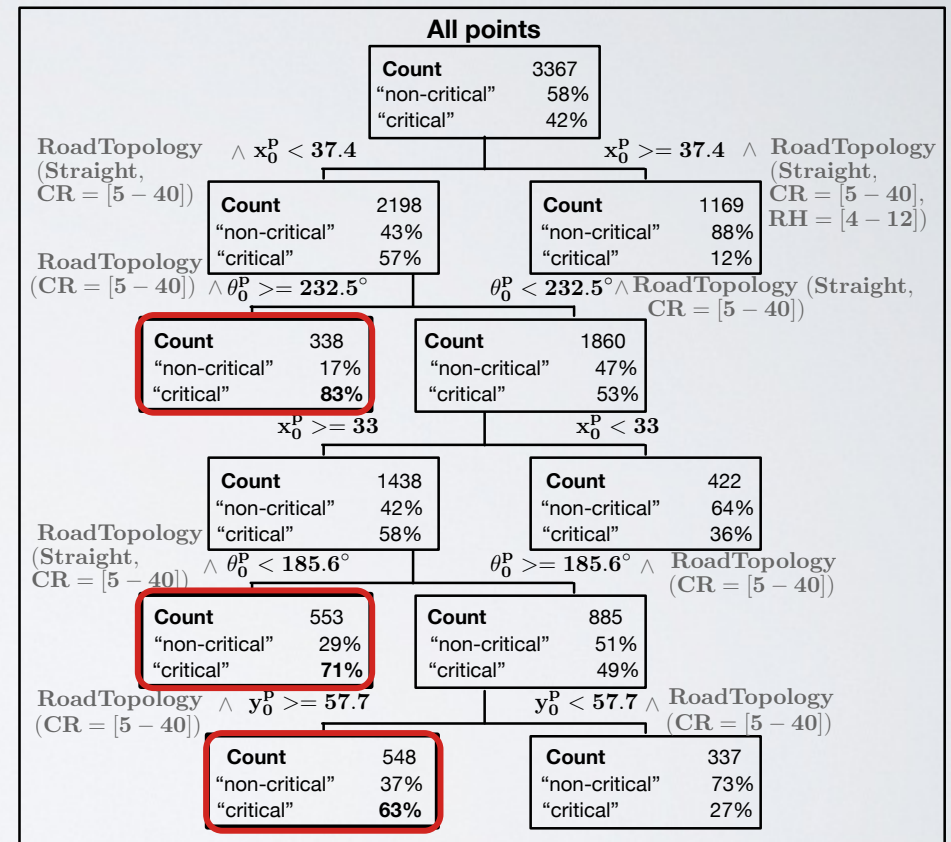
Region in the input space that is likely to contain more critical scenarios

Iterative Process

Initial Classification Model



Refined Classification Model




We focus on generating more scenarios in the critical region, respecting the conditions that lead to that region

We get a more refined decision tree with more critical regions and more homogeneous areas

Research Questions

- RQ1: Does the decision tree technique help **guide** the evolutionary search and make it more **effective**?
- RQ2: Does our approach help **characterize** and **converge** towards **homogeneous** critical regions?
- Failure explanation
- Usefulness (feedback from engineers)

RQ1: NSGAI-DT vs. NSGAI

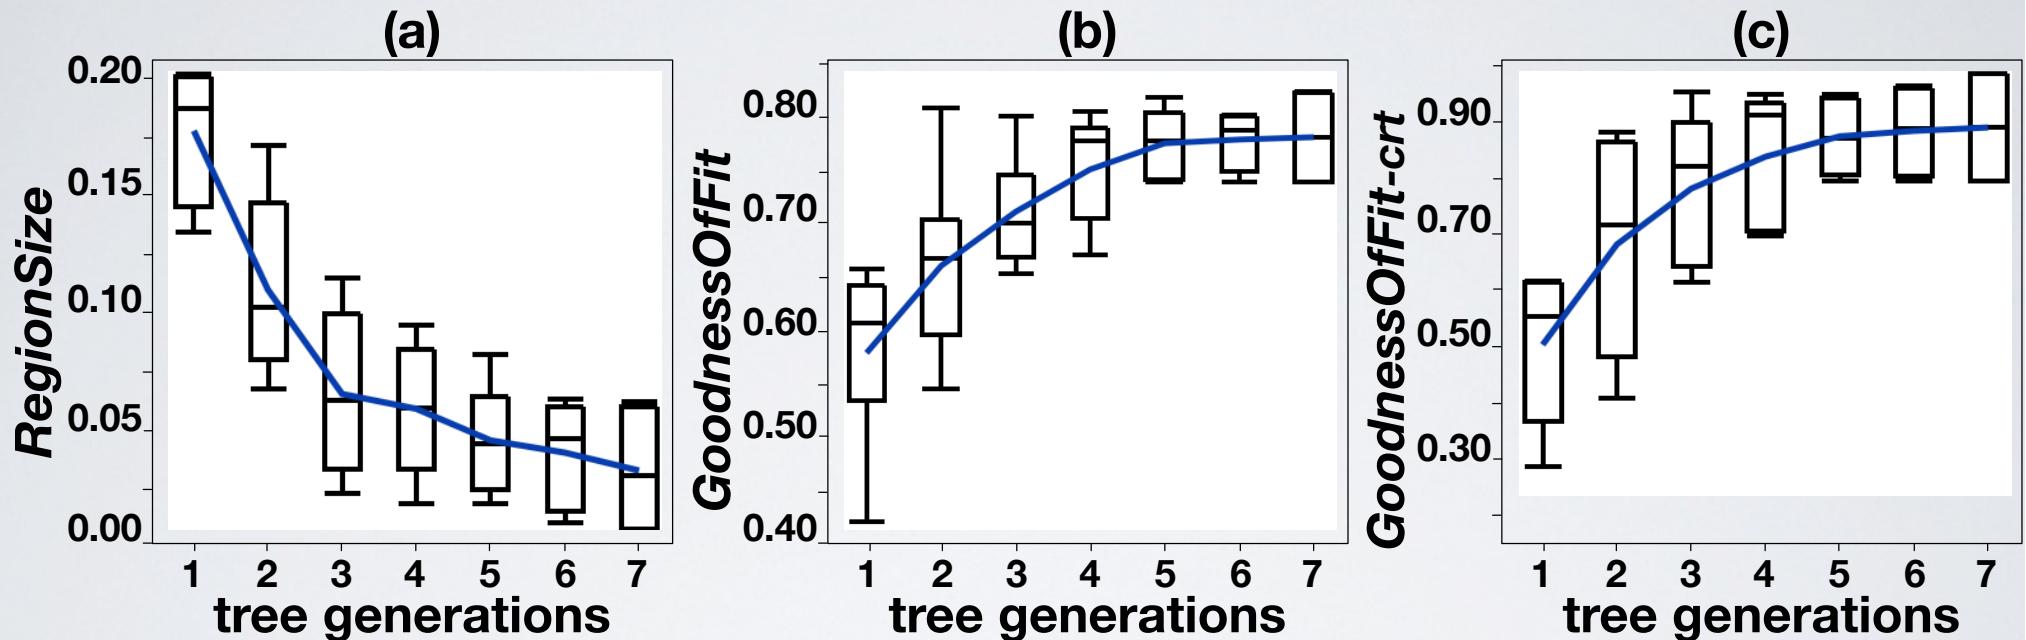


NSGAI-DT outperforms NSGAI

RQ1: NSGAI-DT vs. NSGAI

- NSGAI-DT generates 78% more **distinct, critical** test scenarios compared to NSGAI


RQ2: NSGAI-DT (evaluation of the generated decision trees)



The generated critical regions consistently become smaller, more homogeneous and more precise over successive tree generations of NSGAI-DT

Failure explanation

- A characterization of the input space showing **under what input conditions the system is likely to fail**
- Visualized by decision trees or dedicated diagrams
- **Path conditions in trees**



Usefulness

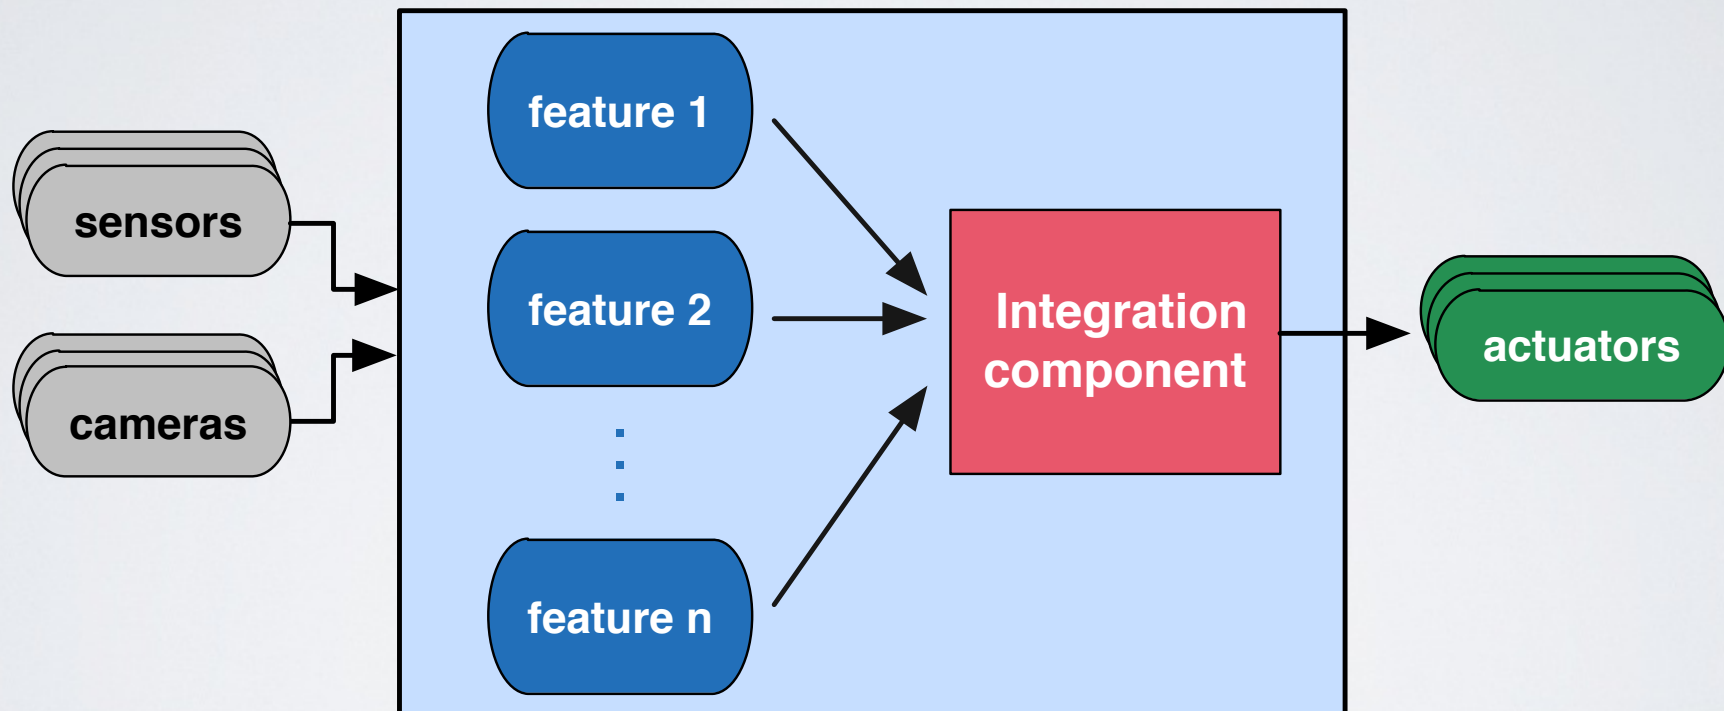
- The characterizations of the different critical regions can help with:
 - (1) **Debugging** the system model (or the simulator)
 - (2) **Identifying possible hardware changes** to increase ADAS safety
 - (3) **Providing proper warnings** to drivers

Automated Testing of Feature Interactions Using Many Objective Search



System Integration

System Under Test (SUT)

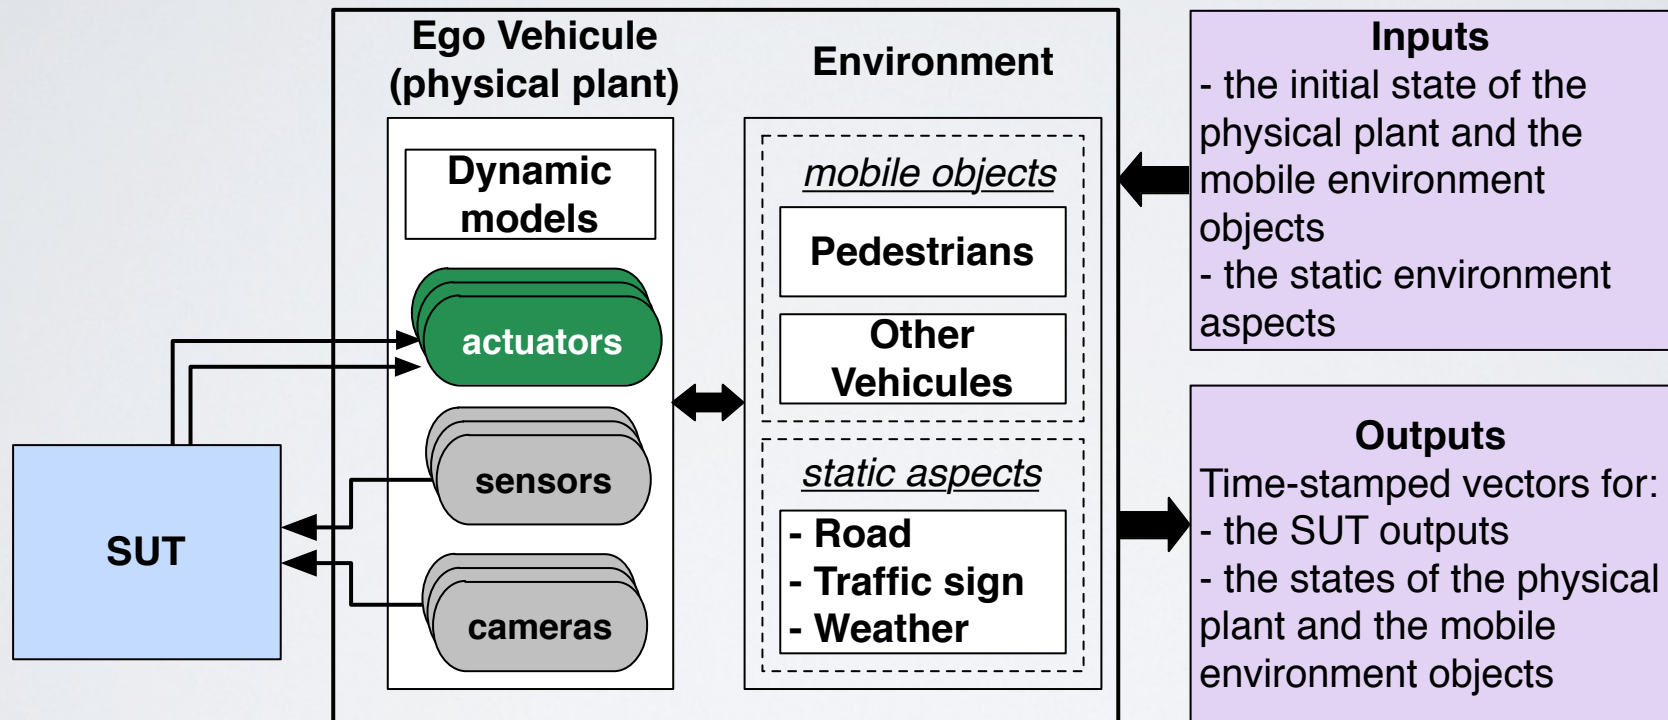


Case Study: SafeDrive

- Our case study describes an automotive system consisting of **four advanced driver assistance features**:
 - Cruise Control (ACC)
 - Traffic Sign Recognition (TSR)
 - Pedestrian Protection (PP)
 - Automated Emergency Breaking (AEB)

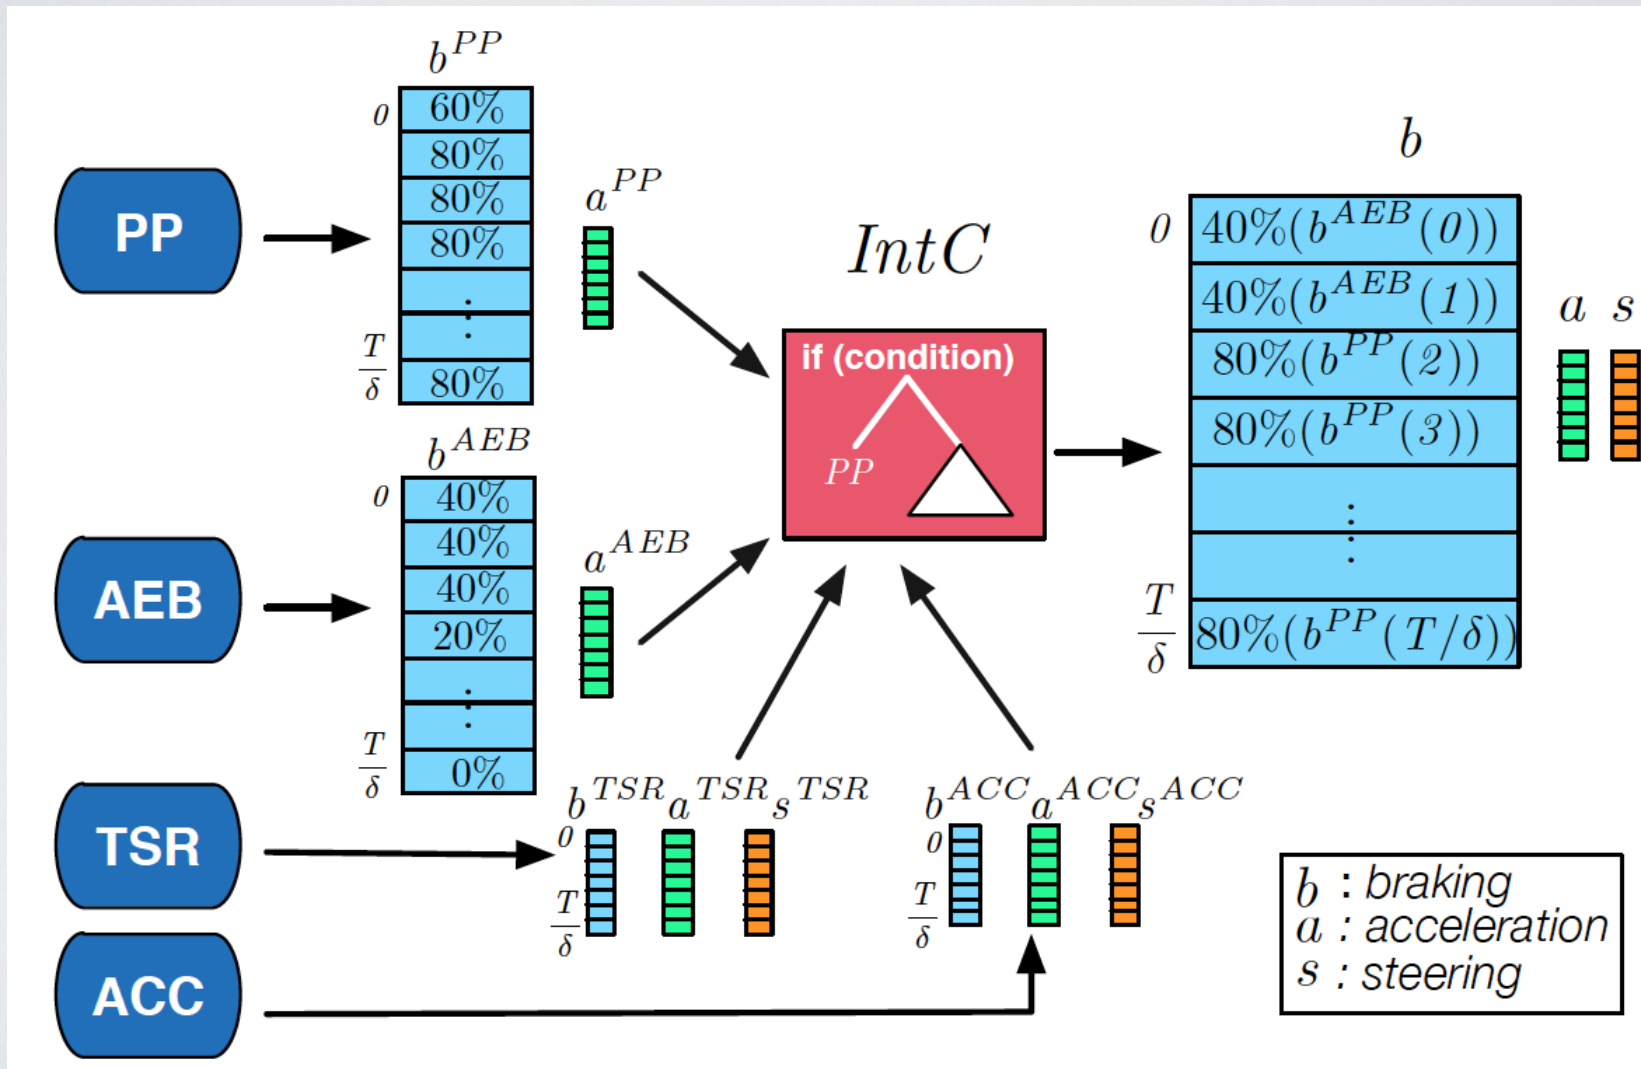
Simulation

Simulator



Feedback loop

Actuator Command Vectors



Safety Requirements

Feature	Requirement	Failure distance functions (FD_1, \dots, FD_5)
<i>PP</i>	No collision with pedestrians	$FD_1(i)$ is the distance between the ego car and the pedestrian at step i .
<i>AEB</i>	No collision with cars	$FD_2(i)$ is the distance between the ego car and the leading car at step i .
<i>TSR</i>	Stop at a stop sign	Let $u(i)$ be the speed of the ego car, at time step i , once it reaches a stop sign. If there is no stop sign, then $u(i) = 0$. We define $FD_3(i) = 0$ if $u(i) \geq 20km/h$. Otherwise, we define $FD_3(i) = \frac{1}{u(i)}$. If there is no stop sign, we have $FD_3(i) = 1$.
<i>TSR</i>	Respect the speed limit	Let $u'(i)$ be the difference between the speed of the ego car and the speed limit at step i if a speed limit sign is detected. If there is no speed limit sign $u'(i) = 0$. We define $FD_4(i) = 0$ if $u(i) \geq 20km/h$. Otherwise, we define $FD_4(i) = \frac{1}{u'(i)}$. If there is no speed limit sign, we have $FD_4(i) = 1$.
<i>ACC</i>	Respect the safety distance	$FD_5(i)$ is the absolute difference between the safety distance sd and $FD_2(i)$.

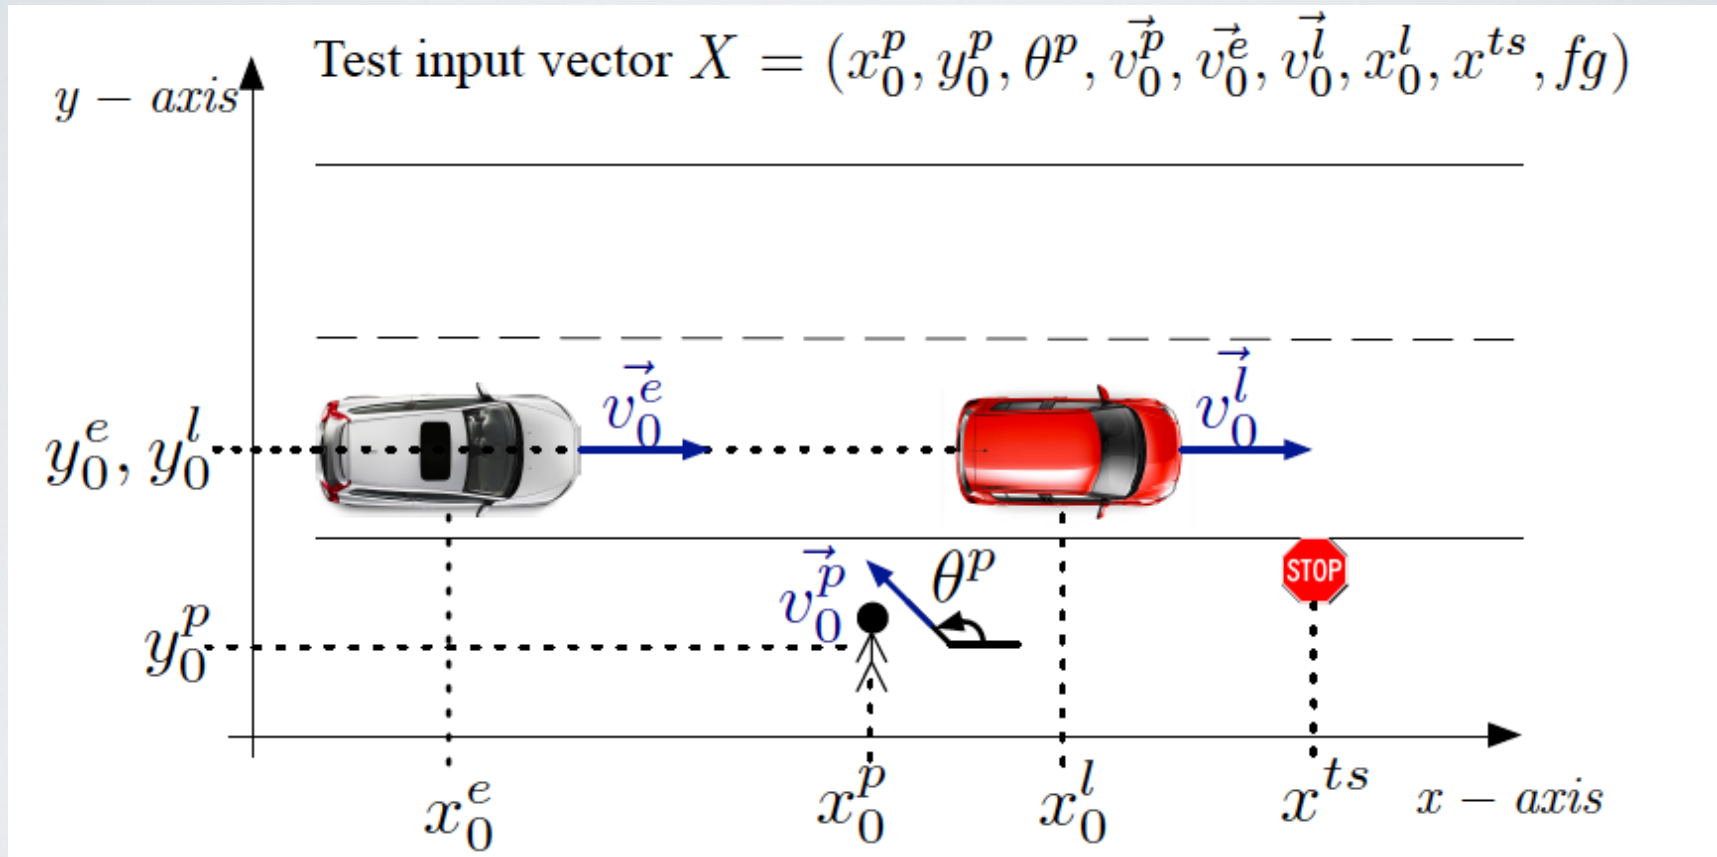
Features

- Behavior of features based on **machine learning algorithms processing sensor and camera data**
- **Interactions** between features may lead to **violating safety requirements**, even if features are correct
- E.g., ACC is controlling the car by ordering it to accelerate since the leading car is far away, while a pedestrian starts crossing the road. PP starts sending braking commands to avoid hitting the pedestrian.
- **Complex:** predict and analyze possible interactions at the requirements level in a complex environment
- **Resolution strategies** cannot always be determined statically and may depend on environment

Objective

- **Automated and scalable testing** to help ensure that resolution strategies are safe
- Detect **undesired feature interactions**
- **Assumptions:** IntC is white-box (integrator is testing), features were previously tested

Input Variables



Search

- **Input space is very large**
- **Dedicated search algorithm** (many objectives) directed/guided by test objectives (fitness functions)
- **Fitness (distance) functions:** reward test cases that are more likely to reveal integration failures leading to safety violations
- **Combine three types of functions:** (1) safety violations, (2) unsafe overriding by integration component (IntC), (3) coverage of the decision structure of IntC
- **Many test objectives** to be satisfied by the test suite

Failure Distance

- Goal: Reveal **safety requirements violations**
- **Fitness functions** based on the trajectory vectors for the ego car, the leading car and the pedestrian, generated by the simulator
- **PP fitness**: Minimum distance between the car and the pedestrian during the simulation time.
- **AEB fitness**: Minimum distance between the car and the leading car during the simulation time.

Distance Functions

Feature	Requirement	Failure distance functions (FD_1, \dots, FD_5)
<i>PP</i>	No collision with pedestrians	$FD_1(i)$ is the distance between the ego car and the pedestrian at step i .
<i>AEB</i>	No collision with cars	$FD_2(i)$ is the distance between the ego car and the leading car at step i .
<i>TSR</i>	Stop at a stop sign	Let $u(i)$ be the speed of the ego car, at time step i , once it reaches a stop sign. If there is no stop sign, then $u(i) = 0$. We define $FD_3(i) = 0$ if $u(i) \geq 20km/h$. Otherwise, we define $FD_3(i) = \frac{1}{u(i)}$. If there is no stop sign, we have $FD_3(i) = 1$.
<i>TSR</i>	Respect the speed limit	Let $u'(i)$ be the difference between the speed of the ego car and the speed limit at step i if a speed limit sign is detected. If there is no speed limit sign $u'(i) = 0$. We define $FD_4(i) = 0$ if $u'(i) > 0$. Otherwise, we define $FD_4(i) = 1$.
<i>ACC</i>	Adaptive Cruise Control	Let $u(i)$ be the speed of the ego car, at time step i , once it reaches a stop sign. If there is no stop sign, then $u(i) = 0$. We define $FD_5(i) = 0$ if $u(i) \geq 20km/h$. Otherwise, we define $FD_5(i) = \frac{1}{u(i)}$. If there is no stop sign, we have $FD_5(i) = 1$.

When any of the functions yields zero, a safety failure corresponding to that function is detected.

Unsafe Overriding Distance

- **Goal:** Find faults faults in integration component
- Reward test cases generating integration outputs **deviating from the individual feature outputs**, in such a way as to possibly lead to safety violations.
- **Example:** A feature f issues a braking command while the integration component issues no braking command or a braking command with a lower force than that of f .

Branch Distance

- Many decision branches in IntC
- Branch coverage of IntC
- **Fitness:** Approach level and branch distance d (standard for code coverage)
- $d(b,tc) = 0$ when tc covers b

Algorithm 1: Decision-making

```

Input: - Targetbrake
        /*Targetbrake = (brakepp, brakeAEB, brakeACC, brakeTSR)
        - Targetthrottle
        /*Targetthrottle = (throttlepp, throttleAEB, throttleACC, throttleTSR)
        /*LevelOfConfidence = (LCpp, LCAEB, LCACC, LCTSR)
        - Objectdistance, StopSign, TrafficSign, speedLeadCar, speedlimit

Output: - Brake, Throttle

1 begin
2   if (Targetbrake[1] > 0 and Pedestrian is detected and Objectdistance is Close)
3     then
4       Brake ← Targetbrake[1] /*b1
5     else
6       if (Targetbrake[2] > 0 and Car is detected and Objectdistance is Close)
7         then
8           Brake ← Targetbrake[2] /*b2
9         else
10          if TrafficSign is pedestrians crossing and LevelOfConfidence[1] is
11            Low then
12              Brake ← Targetbrake[1] /*b3
13            else
14              if StopSign then
15                Brake ← Targetbrake[4] /*b4
16              else
17                if (Targetbrake[1] > 0 or Targetthrottle[1] > 0) then
18                  Brake ← Targetbrake[1] /*b5
19                  Throttle ← Targetthrottle[1]
20                else
21                  if (Targetbrake[2] > 0 or Targetthrottle[2] > 0)
22                    then
23                      Brake ← Targetbrake[2] /*b6
24                      Throttle ← Targetthrottle[2]
25                    else
26                      if ((Targetbrake[4] > 0 or Targetthrottle[4] >
27                        0) and speedlimit < speedLeadCar) then
28                        Brake ← Targetbrake[4] /*b7
29                        Throttle ← Targetthrottle[4]
30                      else
31                        if (Targetbrake[3] > 0 or
32                          Targetthrottle[3] > 0) then
33                          Brake ← Targetbrake[3] /*b8
34                          Throttle ← Targetthrottle[3]
35                        else
36                          Brake ← 0 /*b9
37                          Throttle ← 0

```

Combining Distance Functions

- **Goal: Execute every branch of IntC such that while executing that branch, IntC unsafely overrides every feature f and its outputs violate every safety requirement related to f .**

$$\Omega_{j,l}(i) = \begin{cases} \overline{BD}_j(i) + \text{Max}(\overline{UOD}) + \text{Max}(\overline{FD}) & (1) \text{ If } j \text{ is not covered } (\overline{BD}_j(i) > 0) \\ \overline{UOD}_f(i) + \text{Max}(\overline{FD}) & (2) \text{ If } j \text{ is covered, but } f \text{ is not unsafely} \\ & \text{overridden } (\overline{BD}_j(i) = 0 \wedge \overline{UOD}_f(i) > 0) \\ \overline{FD}_l(i) & (3) \text{ Otherwise } (\overline{BD}_j(i) = 0 \wedge \overline{UOD}_f(i) = 0) \end{cases}$$

$$\Omega_{j,l} = \text{Min}_{i=0}^{\overline{\delta}} \Omega_{j,l}(i)$$

$$\Omega_{j,l}(tc) > 2$$

Indicates that tc has not covered branch j

$$2 \geq \Omega_{j,l}(tc) > 1$$

Branch covered but did not cause unsafe override of f

$$1 \geq \Omega_{j,l}(i) > 0$$

Branch covered, unsafe override, but did not violate requirement /

Search Algorithm

- Best **test suite** covers **all search objectives**, i.e., for all IntC branches and all safety requirements
- Not a Pareto front optimization problem
- **Objectives compete** with each others for each test case
- Example: cannot have the ego car violating the speed limit after hitting the leading car in one test case
- Tailored, **many-objective genetic algorithm**
- **Must be efficient (test case executions are very expensive)**

Search Algorithm

Algorithm 1: FITEST

Input: Ω : Set of objectives

N : Initial population size

Result: A : Archive

```
1 begin
2    $P \leftarrow \text{RANDOM-POPULATION}(N)$ 
3    $W \leftarrow \text{CALCULATE-OBJECTIVES}(P, \Omega)$ 
4    $[\Omega_c, T_c] \leftarrow \text{GET-COVERED-OBJECTIVE}(P, W)$ 
5    $A \leftarrow T_c$ 
6    $\Omega \leftarrow \Omega - \Omega_c$ 
7   while not (stop_condition) do
8      $Q \leftarrow \text{RECOMBINE}(P, N)$ 
9      $Q \leftarrow \text{CORRECT-OFFSPRINGS}(Q)$ 
10     $W \leftarrow \text{CALCULATE-OBJECTIVES}(Q, \Omega)$ 
11     $[\Omega_c, T_c] \leftarrow \text{GET-COVERED-OBJECTIVE}(P, W)$ 
12     $A \leftarrow A \cup T_c$ 
13     $\Omega \leftarrow \Omega - \Omega_c$ 
14     $F_0 \leftarrow \text{ENVIRONMENTAL-SELECTION}(P \cup Q, \Omega)$ 
15     $P \leftarrow F_0$ 
16     $N \leftarrow |F_0|$ 
17  return  $A$ 
```

Randomly generated TCs

Compute fitness

Archive covering tests


Tests are evolved

Crossover, mutation

Correct constraint violations

Fittest tests selected

Evaluation



Discussion

Observations

- **We are unlikely to have precise and complete requirements, we face great diversity in the physical environment, including many possible scenarios.**
- **It is possible, however, to define properties characterizing unacceptable situations (safety)**
- **Notion of test coverage is elusive: No specification or code/models for some key (decision) components based on ML**
- **We have executable/simulable functional models (e.g., Simulink) at early stages**

Conclusions

- We proposed **solutions** based on:
 - Efficient and realistic (hardware, physics) **simulation**
 - Metaheuristic **search**, e.g., evolutionary computing
 - Guided by fitness functions derived from properties of interest (e.g., safety requirements)
 - **Machine learning**, e.g., to speed up search, provide explanations to engineers
- No guarantees though

Generalizing

- Our work **easily generalizes** to many other cyber-physical systems
- Can a similar strategy be applied in **other domains** to test for bias or any other undesirable properties (e.g., legal), when system behavior is driven by machine learning?
- **Executable models** of environment and users?

Summary

- **Machine learning plays an increasingly prominent role in autonomous systems**
- **No (complete) requirements, specifications, or even code**
- **Some safety and mission-critical requirements**
- **Neural networks (deep learning) with millions of weights**
- **How do we gain confidence in such software in a scalable and cost-effective way?**

Related Testing Research

- Testing of hybrid controllers
- Testing timeliness requirements
- Testing for deadline misses (schedulability)
- HiL acceptance testing prioritization
- Testing for security vulnerabilities
- **Find publications on: svv.lu**

Acknowledgements

- **Raja Ben Abdessalem**
- **Shiva Nejati**
- **Annibale Panichella**
- **IEE, Luxembourg**

References

- **R. Ben Abdesslem et al., "Testing Advanced Driver Assistance Systems Using Multi-Objective Search and Neural Networks", IEEE ASE 2016**
- **R. Ben Abdesslem et al., "Testing Vision-Based Control Systems Using Learnable Evolutionary Algorithms", IEEE/ACM ICSE 2018**

Automated Testing of Autonomous Driving Assistance Systems

Lionel Briand

SEMLA, Montreal, 2018