Driving the Unknown

ODD-aware Decision Making for Autonomous Vehicles

Josep Maria Barbera, Jorge Villagra
{josep.barbera, jorge.villagra}@csis.es

Centre for Automation and Robotics (CAR, CSIC-UPM), Spain

autopia
Connected and automated driving

*In pursuit of more general, safe, and interpretable autonomous driving systems, this thesis aims to develop **a maneuver-planner architecture** that uses Reinforcement Learning (RL) **to adapt to the Operational Design Domain (ODD)***

Motivation

- **ODD** defines the precise conditions under which an autonomous vehicle can safely operate, encompassing:
  - **driving scenarios** (e.g., highway, urban, rural)
  - **environmental conditions** (e.g., weather, lighting)
  - **dynamic elements** (e.g., vehicles, pedestrians, obstacles)
- **ODD monitoring** allows the vehicle to **adapt its behavior** accordingly while also **quantifying risk**.
- **Partially Observable Markov Decision Processes (POMDPs)** make hidden states tractable (e.g., occlusions, sensor noise, environmental disruptions), while **RL** optimizes belief-state policies, enabling more robust and reliable decision-making under partial observability.

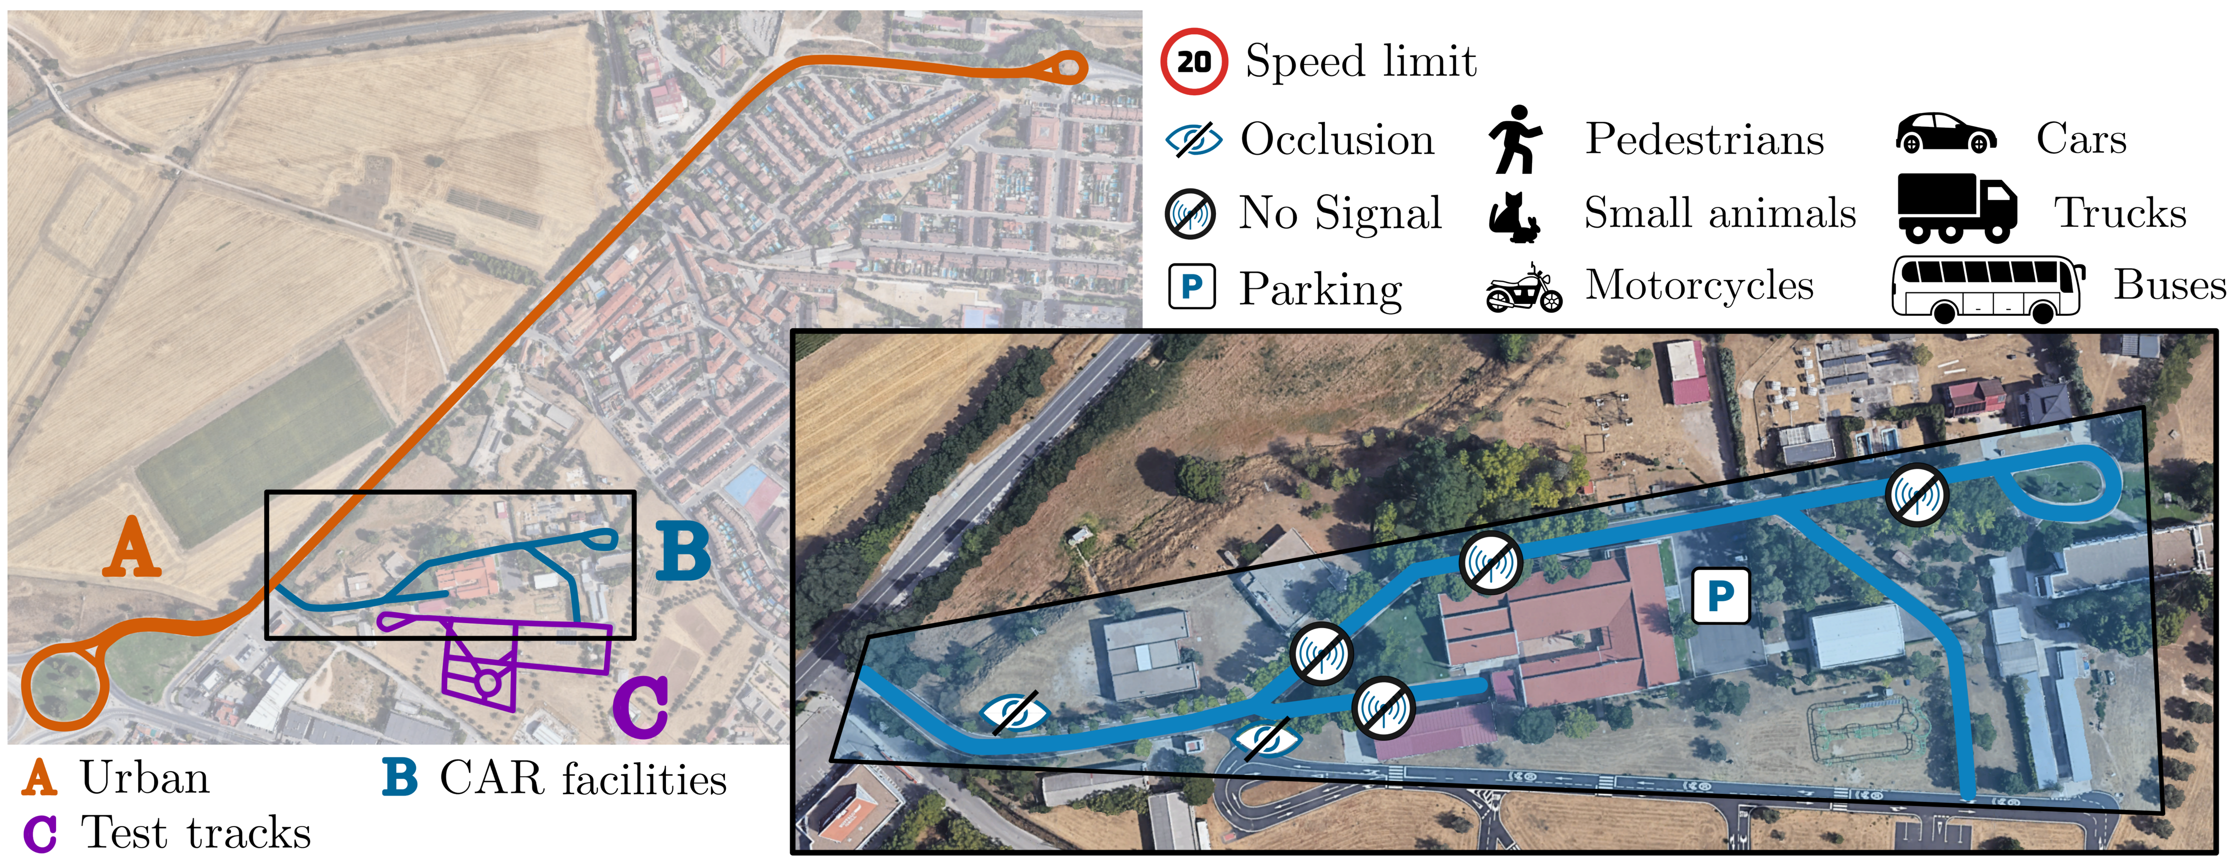

Figure 1: **ODD changes with context**. On the left, three areas are highlighted: urban (A), CAR facilities (B), and test tracks (C). On the right, **the CAR facilities ODD** is detailed, showing occlusions, signal interferences, a parking area, a 20 km/h speed limit, and dynamic elements such as pedestrians, animals, motorcycles, cars, trucks, and buses.

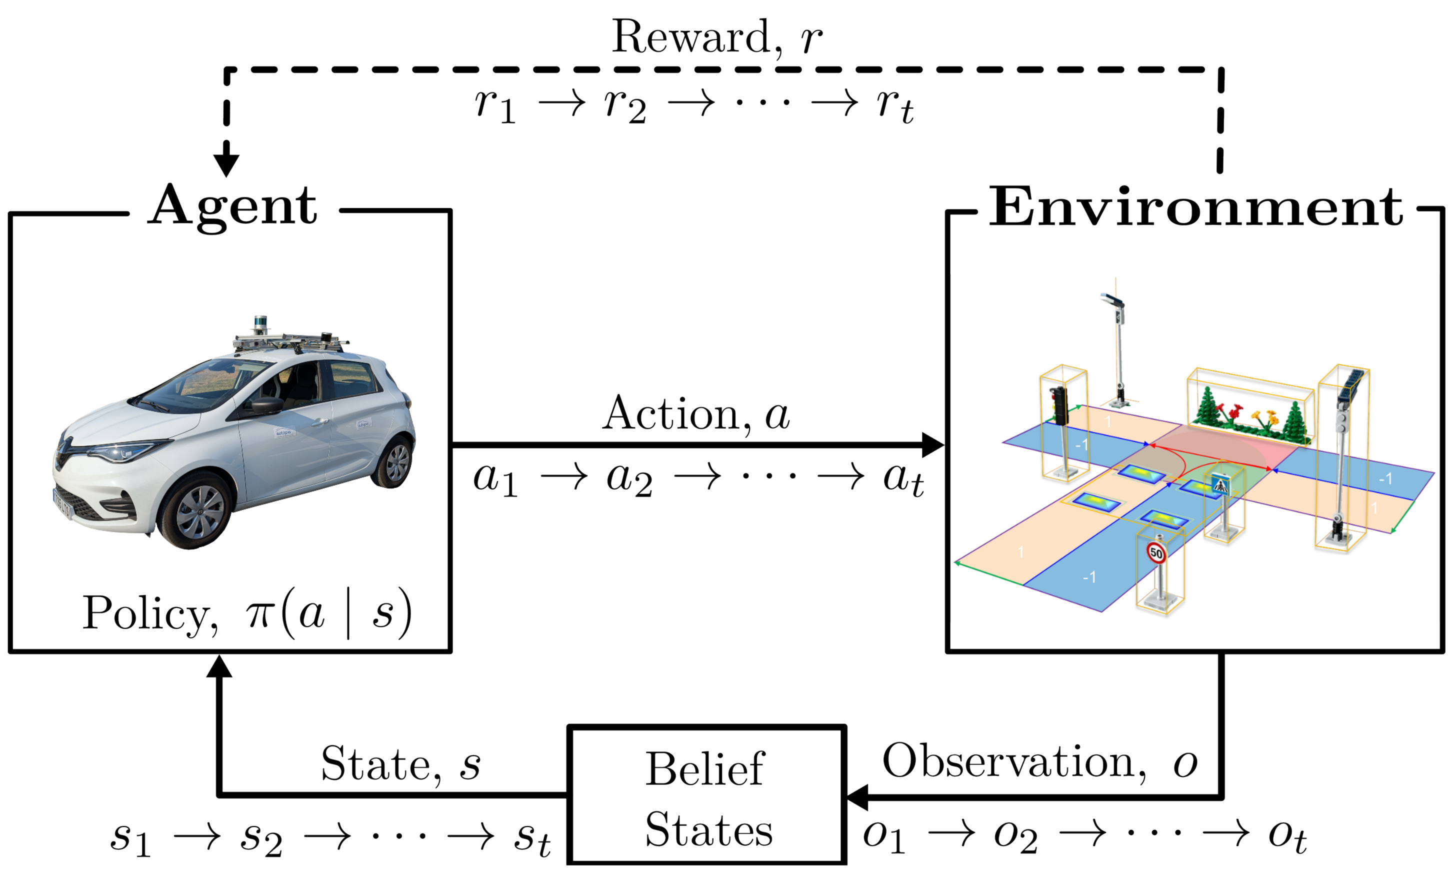

Figure 2. **RL General Framework with Belief States**. Main diagram inspired by [1]; environment icon from [2].

Objectives

- To **formulate a realistic and generalized decision-making problem** for autonomous vehicles that accounts for uncertainty, partial observability, and diverse urban driving conditions.
- To **train** adaptive **decision policies** that ensure safe and efficient behavior in complex, real-world urban and high speed environments.
- To **implement and assess a maneuver-planner architecture** that remains **effective across** diverse contexts:
  - **traffic density**
  - potential **hazards** (e.g., emergency vehicles, road works)
  - **degraded conditions** (e.g., sensor failures, poor connectivity, adverse weather)

Use Case Example

- Lane Change Maneuver: high level decision
  - {NOT CHANGE, CHANGE RIGHT, CHANGE LEFT}
- On highways (2 or more lanes)
- States inferred from observations (noisy measurements or occlusions)

Action	State Variables	ODD Examples
NOT CHANGE	p_{OV}	Traffic density
CHANGE LEFT	d_{OV}	Sensor failure (e.g., GPS loss)
CHANGE RIGHT	v_{OV}	Occlusion of other vehicles

Figure 3. **Action set**, corresponding **state variables**, and **ODD examples** for the lane change maneuver

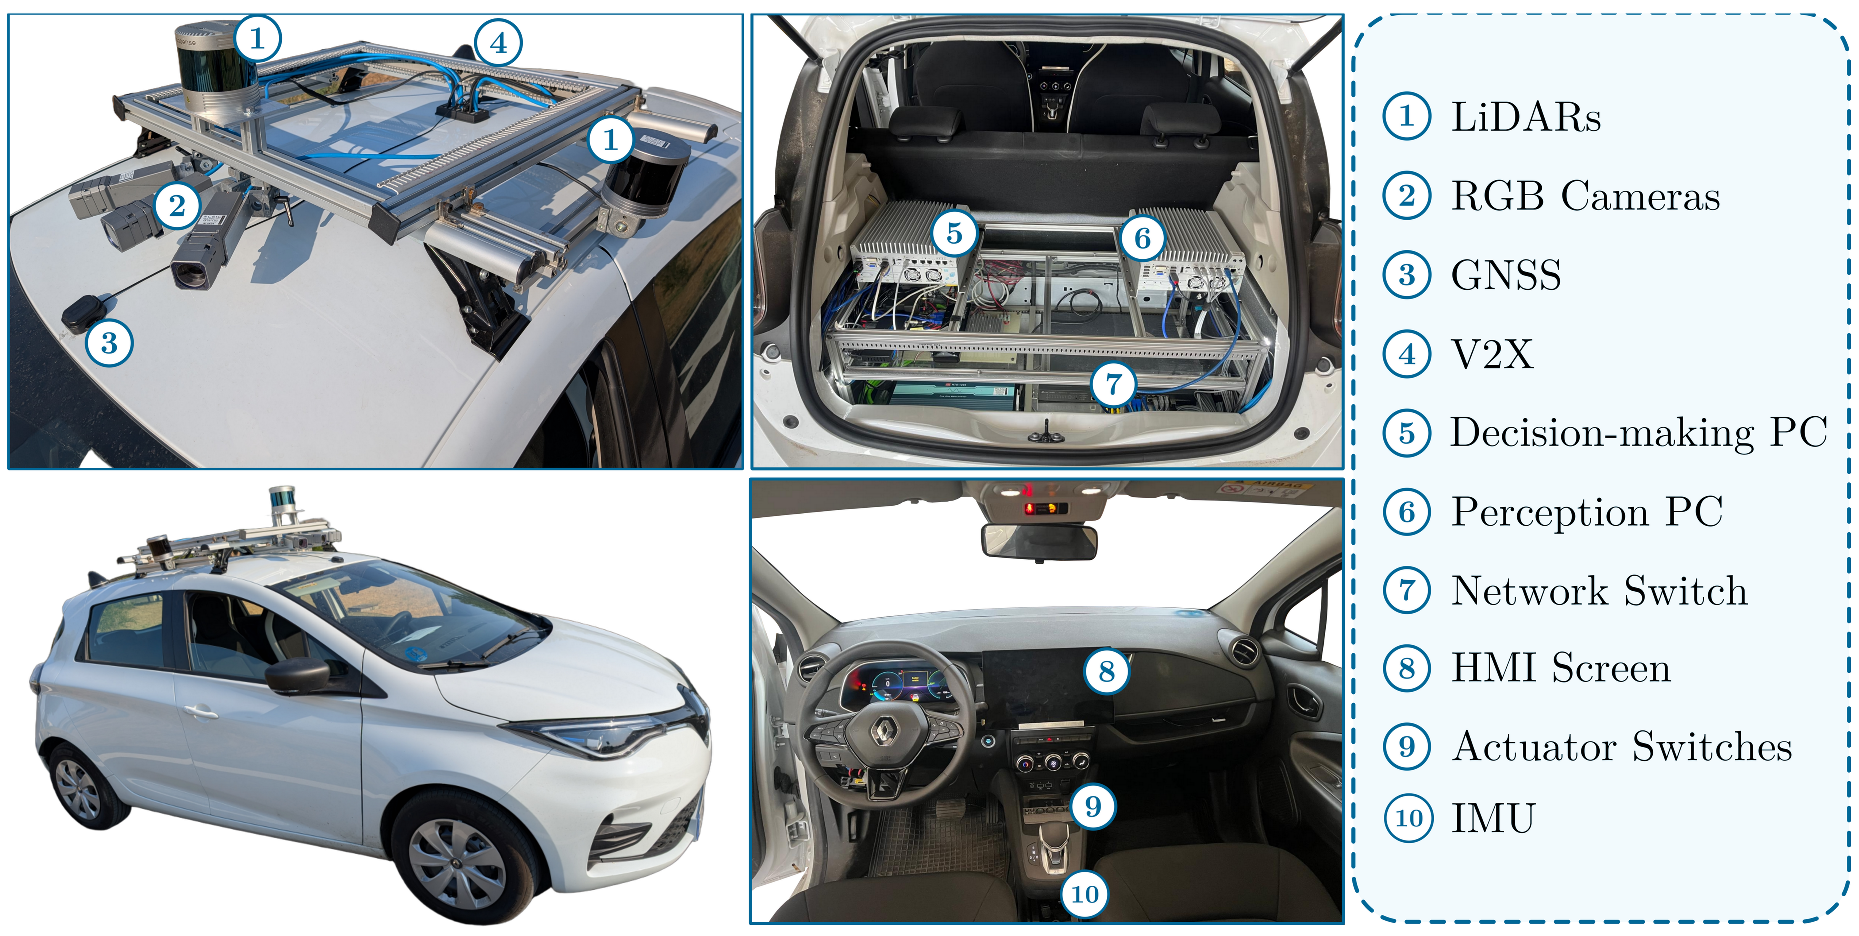

Figure 4. Implementations and new architectures are validated in the **experimental platforms**

Methodology and Validation

- Development on **ROS2** using **Lanelet2** maps [3].
- Training and initial testing in **CARLA**, complemented by closed-loop simulation validation in **SCANeR** [4].
- Final test on **real vehicles**, progressively from area **C** to **B** and then **A**.

References

[1] S. L. Brunton, *Machine learning meets control theory*, Video, Reinforcement learning, Cassyni, 2021. [Online]. Available: https://doi.org/10.52843/cassyni.x2t0sp

[2] ASAM, *OpenODD: Concept Paper*, Oct. 2021. [Online]. Available: https://www.asam.net/standards/detail/openodd

[3] F. Poggenhans, J.-H. Pauls, J. Janosovits, S. Orf, M. Naumann, F. Kuhnt, et al., *Lanelet2: A high-definition map framework for the future of automated driving*, in Proc. IEEE Intell. Trans. Syst. Conf., Hawaii, USA, 2018.

[4] AVSimulation. (2025). SCANeR™ Studio [Computer software]. AVSimulation. https://www.avsimulation.com/

CAR CENTRO DE AUTOMÁTICA Y ROBÓTICA
CSIC
POLITÉCNICA
SHAPE FUTURE

Acknowledgements
Funding for this research was provided by the European Union's Horizon Europe programme under the SHAPEFUTURE project (Grant Agreement No. 101139996).