







## Regular Article

# er.autopilot 1.0: The Full Autonomous Stack for Oval Racing at High Speeds

Ayoub Raji<sup>1,2</sup>, Danilo Caporale<sup>3</sup>, Francesco Gatti<sup>4</sup>, Andrea Giove<sup>5</sup>, Micaela Verucchi<sup>1,4</sup>, Davide Malatesta<sup>3</sup>, Nicola Musiu<sup>1</sup>, Alessandro Toschi<sup>1</sup>, Silviu Roberto Popitanu<sup>5</sup>, Fabio Bagni<sup>1,4</sup>, Massimiliano Bosi<sup>1,4</sup>, Alexander Liniger<sup>6</sup>, Marko Bertogna<sup>1,4</sup>, Daniele Morra<sup>5</sup>, Francesco Amerotti<sup>4</sup>, Luca Bartoli<sup>4</sup>, Federico Martello<sup>1</sup> and Riccardo Porta<sup>1</sup>

<sup>1</sup>University of Modena and Reggio Emilia, Italy

<sup>2</sup>University of Parma, Italy

<sup>3</sup>Technology Innovation Institute–Autonomous Robotics Research Center, UAE

<sup>4</sup>HIPERT srl

<sup>5</sup>University of Pisa, Italy

<sup>6</sup>Computer Vision Lab, ETH Zurich, Switzerland

**Abstract:** The Indy Autonomous Challenge (IAC) brought together for the first time in history nine autonomous racing teams competing at an unprecedented speed and in a head-to-head scenario, using independently developed software on open-wheel race cars. This paper presents the complete software architecture used by the team TII EuroRacing (TII-ER), covering all the modules needed to avoid static obstacles, perform active overtakes, and reach speeds above 75 m/s (270 km/h). In addition to the most common modules related to perception, planning, and control, we discuss the approaches used for vehicle dynamics modeling, simulation, telemetry, and safety. Overall results and the performance of each module are described, as well as the lessons learned during the first two events of the competition on oval tracks, where the team placed second and third, respectively.

**Keywords:** autonomous racing, control, motion planning, perception, extreme environments

## 1. Introduction

The introduction of Advanced Driver Assistance Systems (ADASs) and partially automated systems on commercial cars has reduced the number of motor vehicle crashes and deaths in the majority of high-income countries (Yellman, 2022). This trend could become even more effective in the coming decades thanks to speed limit regulations and the obligation for car manufacturers to include advanced safety systems, such as the Driver Alcohol Detection System for Safety and the Driver Drowsiness Detection System, on all their vehicles (Benson et al., 2018; Ecola et al., 2018).

---

Received: 1 October 2022; revised: 22 February 2023; accepted: 21 March 2023; published: 9 January 2024.

**Correspondence:** Ayoub Raji, University of Modena and Reggio Emilia, Italy; University of Parma, Italy, Email: [ayoub.raji@unimore.it](mailto:ayoub.raji@unimore.it)

This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Copyright © 2024 Raji, Caporale, Gatti, Giove, Verucchi, Malatesta, Musiu, Toschi, Popitanu, Bagni, Bosi, Liniger, Bertogna, Morra, Amerotti, Bartoli, Martello and Porta

DOI: <https://doi.org/10.55417/fr.2024004>

Nevertheless, currently a significant number of crashes and deaths are caused by harsh weather conditions, poor visibility, and loss of control, which are not likely to be preventable by current ADASs (Benson et al., 2018). This should serve to advance research on fully autonomous driving on highways, in high-speed scenarios, and in harsh road conditions.

Motorsport has always produced innovative technologies that, in many cases, were subsequently adapted to automobiles to improve safety and enhance performance. Examples are rear-view mirrors, seat belts, active suspensions, and engine recovery systems. Similarly to the integration of Motorsport technological innovations in human-driven urban cars, Autonomous Racing could help in the development and testing of self-driving capabilities in extreme cases on race tracks, with future applications in urban autonomous vehicles (Betz et al., 2019b).

Autonomous driving competitions have historically been very effective in fostering research and industrial interest to push self-driving technology beyond its limits. A first milestone was set in 2005 with the DARPA Grand Challenge, where multiple teams competed to drive autonomously off-road vehicles along a 132-mile path in the desert near the California/Nevada state line. The Stanford Racing Team won the \$2 million prize, completing the path in slightly less than 7 h. Researchers from the five teams that completed the challenge became involved as founders and chief researchers of companies that only a decade later would render urban autonomous vehicles a commercial reality.

In the autonomous racing domain, two notable initiatives were proposed in 2016. The f1tenth<sup>1</sup> initiative is an open source platform for the development and testing of autonomous driving software, consisting of 1:10 scale RC cars equipped with a LiDAR scanner, a stereo camera, and Nvidia computational boards. Annual international race events of the f1tenth are organized during the most important conferences in Robotics. Roborace<sup>2</sup> provides full-scale electric race cars able to achieve speeds of around 69.4 m/s (250 km/h). The competition is based on a championship formed by several real and virtual races. In 2017, the Formula SAE<sup>3</sup> created the new Formula Student Driverless (FSD) class where teams formed by students have to design and develop both the mechanics and software of a prototype capable of autonomously running in a closed loop track created by cones. In 2020, a new competition, the Indy Autonomous Challenge (IAC), was launched, with the aim to showcase multivehicle head-to-head races at the limits of handling in high-speed racetracks.

In this paper, we present `er.autopilot 1.0`, the complete software stack used during the IAC by the TII EuroRacing team, which accomplished the second and third position in the first two events. The system was demonstrated to be able to avoid static obstacles, perform active overtakes on other vehicles, and achieve speeds above 75 m/s. The aim of this work is not only to be a reference for the Autonomous Racing domain but also for other autonomous systems in edge-case scenarios for road vehicles and sport cars. Among the major contributions, we present a vehicle model identification approach in which a model-based controller is tuned using simulation tools without prior dynamic data of the vehicle. We exhaustively characterize the performance of each software and control module, and of the overall system, deriving the main lessons learned by analyzing the pros and cons of each solution.

In the remainder of Section 1, we give a brief introduction to the competition and the race car. Related research in the Autonomous Racing domain is discussed in Section 2. The full stack of `er.autopilot 1.0`, the underlying design principles, and the technological solutions adopted are presented in Section 3. In particular, the modules related to localization and perception are presented, including a LiDAR-based solution, and the description of different clustering and detection approaches. The software modules related to motion forecasting, planning, and control have already been presented in (Raji et al., 2022). We thus give here a brief description of their implementation and focus on their effects on the overall system and the final results obtained. Section 4 presents the simulation platforms used for testing. Telemetry and Visualization tools are illustrated in Section 5. The results of each module and of the overall system during the competition

<sup>1</sup> <https://f1tenth.org/>

<sup>2</sup> <https://roborace.com/>

<sup>3</sup> <https://www.fsaonline.com/>

are summarized in Section 6. Section 7 gathers an overview of the lessons learned by the team. The paper is concluded in Section 8, where potential improvements and future research directions are also presented.

### 1.1. Indy Autonomous Challenge

The IAC<sup>4</sup> is an international competition that brings together public-private partnerships and academic institutions to challenge university students around the world to imagine, invent, and validate a new generation of automated vehicle software to run fully autonomous race cars.

The challenge was carried out in two steps: a simulation race and a real race. Of the 30 teams from universities all around the world that participated in this competition, only 9 passed the simulation step. The first race, the Indy Autonomous Challenge powered by Cisco, was held on October 21, 2021 at the Indianapolis Motor Speedway (IMS), and the second one, the Autonomous Challenge @ CES, was held on January 7, 2022 at the Las Vegas Motor Speedway (LVMS). The car shakedown and a considerable part of the development before the IMS race was conducted at Lucas Oil Raceway (LOR).

The race at IMS was a solo time trial competition that consisted of a semifinal and final event. To get access to the race, the teams had to demonstrate a set of requirements during testing. This event, besides the time trial, included an obstacle avoidance challenge: two static obstacles were placed in the front stretch to prove that the car was capable of actively avoiding static obstacles. The final leg was limited to the three teams that achieved the best score in the semifinals. The entire run was formed of four warm-up laps and two performance laps. The winner was determined based on the highest average speed achieved during the two consecutive performance laps.

The race at LVMS consisted in a Passing Competition, where multiple rounds of head-to-head matches were conducted by two cars that had to take turns playing the role of defender and attacker, attempting to overtake at increasing speeds, until one or both cars were unable to complete a pass. In each round, the attacker had to follow the following four steps:

1. Reduce the gap with the defender.
2. Keep a longitudinal safety distance.
3. Overtake the defender once a passing zone is reached.
4. Switch the role to defender and reduce the velocity to a predetermined constant value.

A time trial event was created to determine the teams' seeding in the brackets of the Passing Competition.

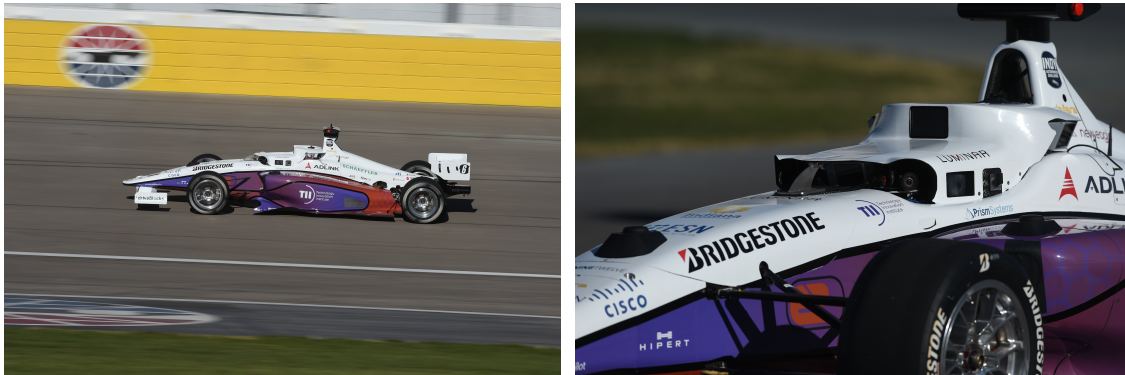
It is worth noting that the original plan for the IAC was to have a multivehicle race with 10 cars on the track already at IMS. Due to several challenges (from weather to logistics to teams facing new difficulties on the track over a short period of time), the race rules were modified to deliver a successful show where most teams could participate with their current level of readiness.

### 1.2. Dallara AV-21

Each team of the IAC participates with a Dallara AV-21, shown in Figure 1, a fully autonomous open-wheel race car based on the official Indy Lights IL-15. Unlike the original race car, the engine mounted is a turbo-charged Honda K20 with 390 horsepower. The mechanics, suspensions, and aerodynamics are adjusted for oval racing and high banked tracks with an asymmetrical setup.

On the perception side, the car is equipped with two GNSS modules, three LiDARs, six cameras (two cameras in a stereo setup and four to cover the 360 range), and three RADARs. The Novatel GNSS Pwrpak 7d receivers provide a centimeter precision localization of the car thanks to four antennas and RTK correction. The three solid-state LUMINAR H3 LiDARs have a range of 200 m and they operate at 20 Hz. The Aptiv ESR 2.5 frontal radar and the MRR side radars have a range

<sup>4</sup> <https://www.indyautonomouschallenge.com/>



(a) AV-21 with the TII EuroRacing livery.

(b) Close view on the onboard sensors.

**Figure 1.** Dallara AV-21.

of around 160 m and they provide the detected obstacles at 10 Hz. The six high-resolution RGB Mako G319C cameras from Allied Vision are mounted to have a view of almost 360 degrees around the car. The computing platform used on the vehicle is an ADLINK AVA-3501 consisting of an 8 core Intel Xeon E 2278 GE CPU, an NVIDIA RTX Quadro 8000 GPU, and 64 GB DDR4 RAM. For external communication, Cisco FM-4500 radio transceivers are used on the car and around the tracks in order to make telemetry data available to the crew team in the pit lane. Race Control signals can be exchanged by means of a MYLAPS RaceLink system mounted on the racetracks and a transponder mounted on the vehicle.

On the actuation side, the car has a Drive-by-Wire (DBW) system realized by Schaeffler to actuate the steering, the throttle pedal, the brake pedal, and the gearbox. The New Eagle GCM 196 Raptor control module is used as an interface between the DBW system, the computing platform in which the algorithms are executed, and the other units related to the engine and minor subsystems.

## 2. Related Work

Thanks to the availability of low-cost research race car prototypes and the media impact of autonomous driving challenges, the number of published works in this domain is progressively increasing. In (Betz et al., 2022b), the authors presented a survey on autonomous racing cars reviewing the most relevant publications detailing autonomous driving modules, vehicle modeling, simulation, and complete software architectures.

One of the first problems each team faces when working on an autonomous (race) car is obtaining an accurate localization and state estimation of the ego vehicle within the track. To solve the localization problem, Extended Kalman Filter (EKF) solutions based on a vehicle model are usually adopted to fuse the measurements from different onboard sensors, like GNSS, LiDAR, Inertial Measurement Unit (IMU), and wheel speed odometry (Wischniewski et al., 2019). LiDAR-only localization solutions are proposed in (Massa et al., 2020) and (Schratter et al., 2021), demonstrating an accuracy suitable for driving the Roborace DevBot 2.0 race car within 100 km/h.

On the perception side, the object detection problem in the Autonomous Racing field focused on camera-based solutions. In particular, cone detection using Convolutional Neural Network (CNN) methods on a camera frame has been used by several teams in the FSD competition (De Rita et al., 2019; Puchtler and Peinl, 2020; Vödisch et al., 2022). A more robust and redundant solution is presented in (Kabzan et al., 2019), including cone detection with feature-based and CNN-based approaches using both mono and stereo cameras, cone detection and color classification on LiDAR, and a sensor fusion solution based on the projection of the LiDAR in the camera reference system (Andresen et al., 2020). Finally, (Strobel et al., 2020) employ a CNN both for cone detection and key points estimation for localization purposes. The active detection of other dynamic agents in

a racing domain is also detailed in (Betz et al., 2022a), using an approach that is similar to that adopted in urban settings, where LiDAR, RADAR, and camera data are typically fused.

Global planning in the racing domain is approached by solving an optimization problem to retrieve the lowest lap time trajectory. Solutions based on minimum curvature and minimum time-based optimization considering the vehicle dynamics can be found in (Heilmeyer et al., 2020; Massaro and Limebeer, 2021).

For local planning, sampling-based methods are a popular option due to their effectiveness in various kinds of robotics problems related to obstacle avoidance. Different versions of the Rapidly-exploring Random Tree (RRT) method have been presented, especially for FSD and small-scale platforms, combined with local controllers, predictions using a vehicle model, and curve refinement (Arslan et al., 2017; Bulsara et al., 2020; Feraco et al., 2020). For full-scale vehicles and high-speed conditions, different local planners have been proposed, generating a graph of possible trajectories and choosing the one that minimizes a cost function defined on different criteria (Raji et al., 2022; Stahl et al., 2019). Other methods proposed optimization-based controllers considering obstacles and the free driveable area (Buyval et al., 2017; Liniger et al., 2015).

Regarding the controller module that tracks a certain reference path and speed profile, Model Predictive Control (MPC)-based methods had a great impact due to their advantages in considering complex systems with their inputs, outputs, and constraints, despite the burden on the algorithm design, modeling, and optimization for real-time usage. Some works focused on representing nonlinear models of the vehicles (Novi et al., 2020; Vázquez et al., 2020), while others use simpler models considering uncertainties and constraining the control on some physical parameters (Wischnewski et al., 2021). More classical controllers based on slip angle or feed-forward steering have reached similar performances in sports cars (Kapania and Gerdes, 2015; Laurence et al., 2017).

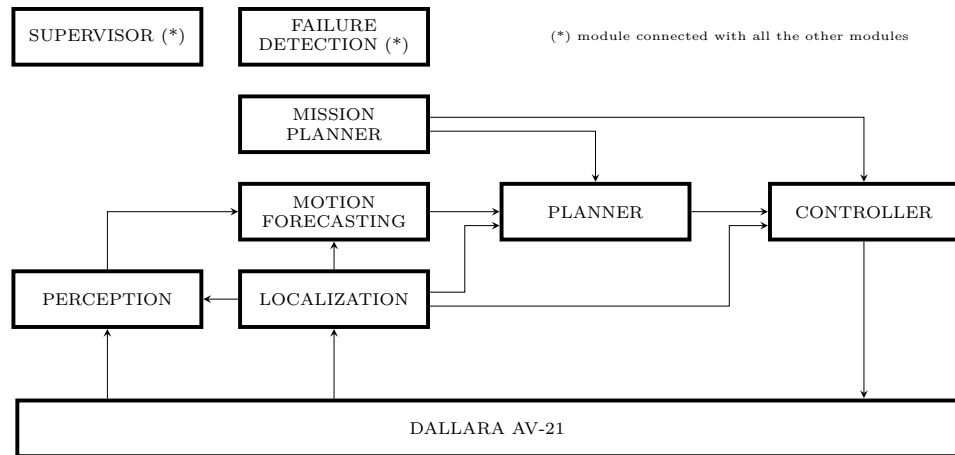
For what concerns the whole autonomous software stack, the literature presents several works related to the FSD competitions in which a brief description of each module is given (Chen et al., 2019; Culley et al., 2020; Nekkah et al., 2020; Tian et al., 2020). (Kabzan et al., 2019) can be considered the most complete system paper presenting implementation details for the common modules needed to succeed on the FSD events, describing the adopted testing framework and the weak points for future improvements. For full-scale race cars, (Betz et al., 2019a; Caporale et al., 2019) present their architecture for the Roborace competition giving particular attention to the motion planning and control modules, with very limited details on the object detection problem. Similar works not associated with any competition have been published. In (Funke et al., 2012), a system architecture is presented focusing on localization, path planning, and control of a commercial sports car at its limits, including the design of a safety module. (Funk et al., 2017) described the design of the hardware and software of an electric race car autonomously driven on a challenging Swiss mountain road.

None of the mentioned works consider multiagent scenarios, or they assume the information on other vehicles will be given. (Betz et al., 2022a) presented the software architecture and methodology of the TUM Autonomous Motorsport team for the IAC. The authors described each module of the stack adopted during the first two events, including the head-to-head race. A framework report can be found in (Urmson et al., 2008), detailing the architecture of the vehicle that won the 2007 DARPA Urban Challenge. In this work, we present a complete autonomous stack for a multiagent scenario, including additional details that we consider fundamental for high-speed racing.

### 3. Software Stack

The `er.autopilot 1.0` software stack consists of multiple modules following the Perceive-Plan-Act paradigm. In Figure 2, a block diagram with a high-level overview of the modules is shown.

A localization module produces the state estimate of the ego vehicle used by all the other modules. By getting the data from the sensors of the AV-21, the Perception stack is able to detect the other vehicles and objects in the environment. This information is used mainly by the Motion Forecasting and Planner modules to predict the opponent’s movement and generate a local path to avoid collision



**Figure 2.** Diagram block with the software modules of `er.autopilot`. The modules marked with an asterisk are connected with all the other modules.

or perform an overtake. A Mission Planner is integrated into the software stack as a behavioral planner able to get the signals from Race Control and from an internal state machine in order to give high-level decisions to the Planner and Controller, such as entering or exiting the pit lane and starting to overtake the opponent. Lastly, the Controller module is the one responsible for generating the correct actuation commands to be sent to the vehicle.

The modules, represented as nodes, communicate with each other using the ROS2 framework and Eclipse Cyclone DDS as middleware. Nodes are compiled into a shared library loaded at runtime, which makes it possible to run multiple nodes in separate processes or as a single process. A base class has been defined for some nodes. For the control methods, the base class contains the callbacks needed to receive the vehicle state from the localization node and the actuation commands feedback, as well as some common methods. Each implementation extends the base class with the additional callbacks, methods, and interfaces to the libraries needed. Indeed, we decided to use the ROS2 nodes as wrappers for the communication, leaving the pure algorithmic parts as stand-alone software to be interfaced with. The only logics implemented on the control base node are the safety checks on the lateral and heading errors of the path tracking performance. Three thresholds have been defined:

- **max\_error:** Above this value, the target speed of the vehicle is linearly decreased.
- **max\_error\_soft:** After passing this threshold, the target speed is set to zero and the controller performs a soft stop.
- **max\_error\_hard:** Reaching this value, a hard brake is actuated over the command of the controller.

All the code executed on the car is written directly in C++ or generated from Matlab Simulink using the C-code generation tool. Python and Julia languages have been used for offline scripts related to data analysis, trajectory optimization and refinement, and visualization. A Docker container has been created for easy deployment on the vehicle and on the developers' machines, as well as on our online GitLab pipeline for basic testing.

### 3.1. Localization of the ego vehicle

The purpose of the localization module is to provide an estimate of the ego vehicle state using the available information from the sensor data or other software components. Several architectures were investigated before converging to the final one, which is presented hereafter.

The vehicle is equipped with two GNSS modules. As explained in [Section 1.2](#), each of these provides the RTK-corrected position of the primary antenna, together with other information like

the estimated speed and heading, and additional information on the quality of the position solution also called the fix. Each receiver module is connected to an IMU that provides linear acceleration and angular rates, and is used to provide some prefiltered signals useful for the user (such as estimated roll, pitch, and yaw angles). We decided to ignore these prefiltered signals and only use sensor raw data to retain a finer control on the localization pipeline. For a detailed description of the antenna’s capabilities, please refer to the producer website.<sup>5</sup>

In the design, we had to consider several vehicle and operational domain requirements:

- The two receivers, called *top* and *bottom* receiver, had the antennas mounted on the main longitudinal axis and lateral axis of the car, respectively. We found the relative positioning of the antennas to have a slight effect on the quality of the fix, hence different weighting was considered for the two.
- The ego vehicle estimation has to be consumed by other modules, the fastest ones being the planning module (running at 50 Hz) and the control module (running at 100 Hz).
- The RTK correction was not always reliable, and the same holds for the GNSS signal, which is a common problem for these kinds of systems [for a more in-depth discussion, please refer to (Massa et al., 2020)].
- The vehicle can reach a maximum velocity of around 300 km/h, hence latency should be reduced to a minimum to guarantee a tight correspondence between the car’s position on the track and its latest available estimate.

For all the above reasons, we chose to equip our car with a robust localization filtering scheme based on an EKF and to develop a LiDAR based localization system to further enhance the system’s robustness in case the GNSS signal is lost.

### 3.1.1. GNSS Localization

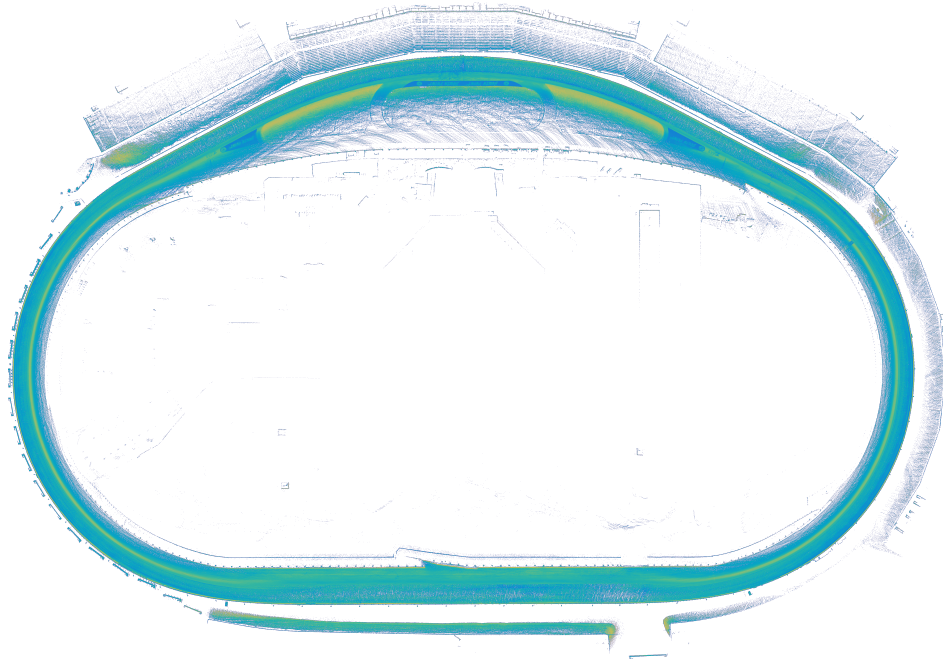
The model used for the estimation is a simple kinematic unicycle in global coordinates, Equation 1. Given the absence of a side-slip angle sensor on the car, it is difficult to obtain a reliable estimate for this important quantity. Despite this simplification, which has been taken into account in the motion controller, the state estimate was accurate enough for the localization even at high speed,

$$\begin{aligned}
 \dot{x}(t) &= v_x(t) \cos(\theta(t)) + \nu_x(t), \\
 \dot{y}(t) &= v_x(t) \sin(\theta(t)) + \nu_y(t), \\
 \dot{\theta}(t) &= \omega_z(t) + \nu_\theta(t), \\
 \dot{v}_x(t) &= a_x(t) + \nu_{v_x}(t), \\
 \dot{a}_x(t) &= \nu_{a_x}(t).
 \end{aligned} \tag{1}$$

The filter is implemented in such a way that it can manage the asynchronous data sources at the fastest possible rate, which is 250 Hz as per the IMU inputs. Model predictions are computed at the same frequency, while corrections are applied whenever new inputs are available. Measurements can come from different sources, such as GNSS- or LiDAR-based localization for the position and/or heading, wheel speed or GNSS speed for the vehicle velocity, and IMU for the yaw rate. The quality of the incoming signals is evaluated partially at the sensor level (e.g., before sending the readings to the EKF) and partially at the filter level when receiving the data, before using it to compute a correction (e.g., by checking the consistency of the GNSS fix with the last estimated vehicle state, or by discarding measurements that are too old due to a lack of real-time processing capability on the onboard computer).

Despite its simplicity, this model allowed us to obtain a precise estimate of the car’s position, heading, and longitudinal acceleration.

<sup>5</sup> <https://novatel.com/products/receivers/enclosures/pwrpak7d>



**Figure 3.** Top view of the LiDAR map obtained for the LVMS circuit. The color used for the points in the cloud is determined by the intensity value of each point.

### 3.1.2. LiDAR Localization

For the LiDAR vehicle localization, point clouds are first synchronized and merged together. Additionally, each cloud is individually motion-compensated using IMU data.

Mapping is done offline, on a log covering the whole track at a slow speed, to enhance the map quality. LiDAR clouds are aligned using a LiDAR Odometry and Mapping (LOAM) method aided by vehicle odometry and GPS. The obtained map is later globally optimized using GTSAM (Dellaert, 2012)<sup>6</sup> to maintain the shape and minimize the distance to the GPS trajectory. The mapping process produces a georeferenced point cloud, which is then used by the LiDAR localization method. A top view of the resulting map is shown in Figure 3.

From the LiDAR depth map, vertical objects are extracted filtering the image by the pixel normal value. The filtered point cloud is used to localize the car on a two-dimensional (2D) top-down map of the circuit. The 2D map is a likelihood field of the filtered point clouds. A particle filter approach is used to localize the car on the 2D map. The particle filter is parallelized on GPU evaluating the probability of each point of each particle. The extrapolated LiDAR localization can be used by the EKF as an alternative to (or together with) the GPS position.

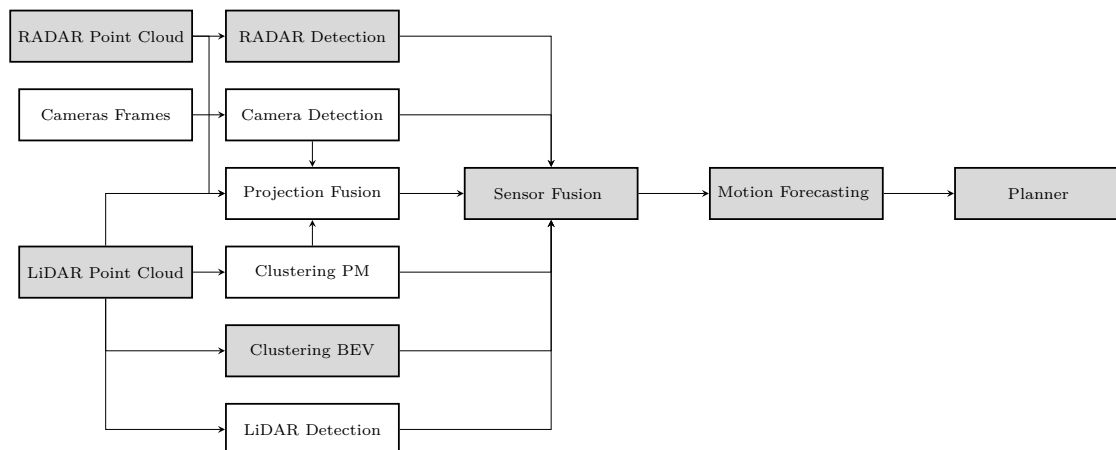
## 3.2. Perception

The complete perception scheme of our stack is depicted in Figure 4. All the blocks are hereafter discussed, but only the gray ones have been used for the races at IMS and LVMS.

The white blocks are implemented and are working, but we excluded these blocks for the following two main reasons. The first is the unsuitable accuracy; in fact, we did not trust some pipelines (e.g. LiDAR Detection) due to the high false positive rate. The second, related to the camera pipelines, is a bandwidth issue we experienced while reading the six cameras with all the other sensors. It occurred several times when reading the cameras' streams led to a higher drop of LiDAR packets

<sup>6</sup> <https://github.com/borglab/gtsam>





**Figure 4.** Perception scheme of `er.autopilot`.

or even to LiDAR failure that was irreversible and needed a system reboot. Considering that the LiDAR is the sensor with the most reliable detections, we sacrificed the cameras so as not to have such a problem during racing days. We will investigate how to solve this issue to exploit those high potential sensors as well.

### 3.2.1. Drivers, Settings, and Calibration

Before introducing the algorithms that run on the sensors' data, the acquisition, configuration, and calibration of the sensors need to be discussed.

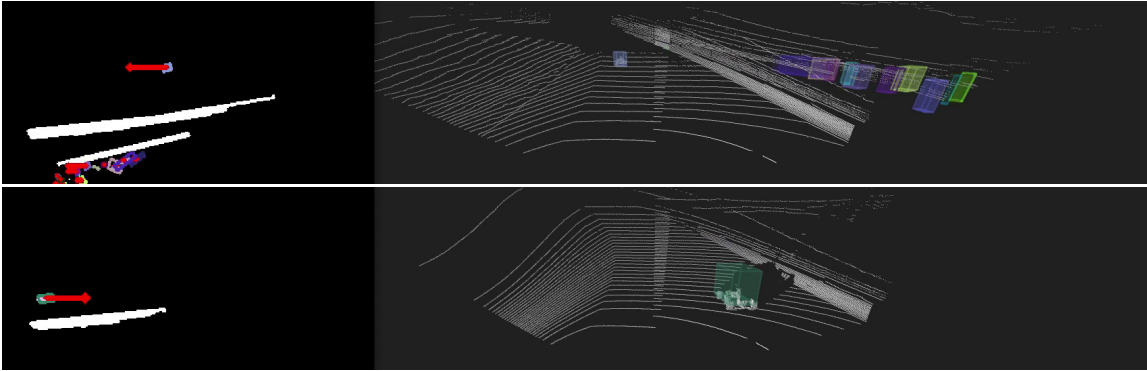
The raw data coming from the LiDARs, cameras, and RADAR (and also GNSS) are collected using our own-made drivers instead of the official ROS2-based ones provided by the sensors' manufacturers. This has been done to manage all the low-level data we consider useful that are not contemplated in the ROS2-based drivers and to reduce the delay between sensor data reading and usage, especially for big data such as frames or point clouds, on our perception algorithms.

We decided to use all six Mako cameras, with a resolution of  $1032 \times 772$  pixels and a frequency of 10 FPS. We utilized a limited resolution and frequency in our setup to prevent band saturation. We used all three Luminar LiDARs, setting a field of view (FOV) of 15 deg, the Gaussian pattern, the center at 0 deg at IMS (1 deg at LVMS due to banking), and a frequency of 20 Hz. We employed a Gaussian pattern to increase the point density on the horizon, while the FOV and layer number were optimized to maintain a 20 Hz frequency while preserving the point density. For the RADAR, we decided to employ only the frontal one, using the default settings and a frequency of 10 Hz, the only available option for the given sensor.

Thereafter, we took care of the sensors' calibration. First, we performed intrinsic camera calibration exploiting an  $8 \times 6$  checkerboard pattern printed on a rigid panel and the Kalibr tool (Oth et al., 2013). Then we implemented and performed camera-LiDAR extrinsic calibration with the same pattern, matching the checkerboard detected from the camera frames with the ones recognized in the LiDAR depth images. At the end of these procedures, we knew both the intrinsic parameters of the cameras and the transformation with respect to all the involved sensors.

### 3.2.2. LiDAR Clustering—Bird's-eye-view approach

The LiDAR Clustering Bird's-eye-view (BEV) pipeline takes the LiDAR point cloud as input, removes the ground, and gives as output clusters of tracked objects. It only executes on the CPU, which makes it robust to GPU failures. To remove the ground, the normals of the point cloud have been exploited. For each point, its  $x$ ,  $y$ , and  $z$  normals have been computed and the points have been filtered on the norm value on the vertical axis ( $z$ ). Additionally, all the points higher than a certain threshold (i.e., 3 m) and the points belonging to the ego vehicle have been removed.



**Figure 5.** Results of the clustering BEV pipeline. The algorithm employs a low-resolution representation of the BEV, depicted on the left, where the white lines indicate the track walls, and their respective clusters are ignored. Instead, the rectangles represent the clusters, color-coded based on their local tracker, and the red arrows indicate their speed vectors. On the right side, the clusters are projected back onto the cloud.

Once the cloud is processed, a BEV image of the remaining points is built. On that image we run a Connected Component algorithm to group the points into objects. That computes the clusters that we can reproject on the point cloud.

For a more stable detection, we also inserted a tracker in the pipeline. The tracker tracks the position, and in particular the center of the cluster, using an EKF, and it matches the objects in different iterations using a nearest-neighbor technique.

Some visual results are depicted in Figure 5 while overtaking a vehicle (top), where the vector representing the velocity (red arrow in the BEV) is negative, and while overtaken (bottom), where the velocity of the opponent is positive.

### 3.2.3. LiDAR Clustering—Point Map approach

The LiDAR clustering Point Map (PM) pipeline takes the LiDAR point cloud as input, removes the ground, and gives as output clusters of objects. It executes both in GPU and CPU, and it is an alternative to the other clustering algorithm.

At first, the point cloud is converted into a Point Map, an image that contains at least three channels that, instead of representing RGB values, are the position of the single point in the 3D environment  $x, y, z$ . Besides position, information such as intensity, time of flight, and ring index can also be included in the PM channel. The whole pipeline then uses this converted PM.

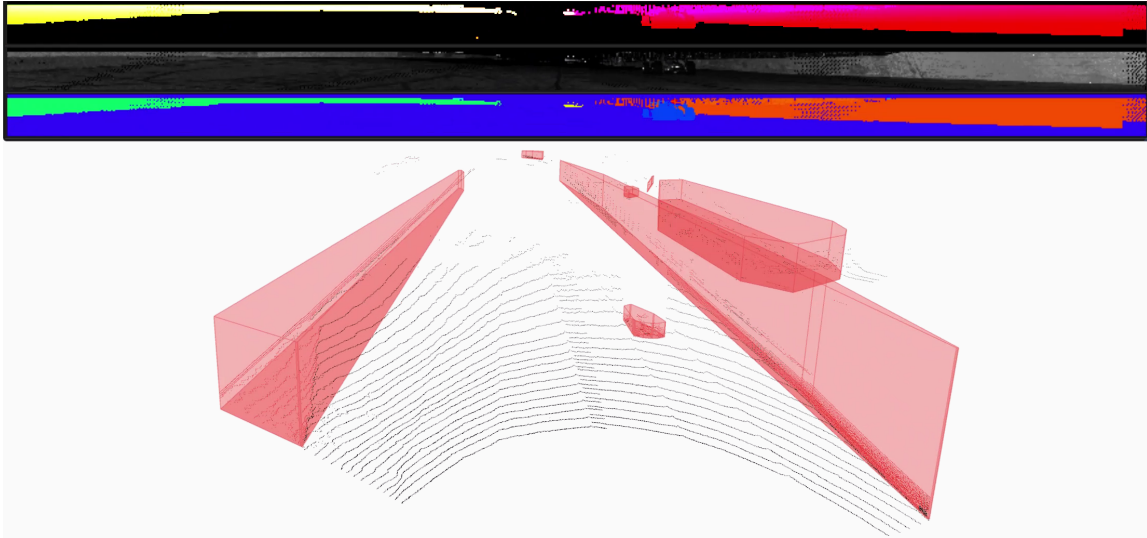
In this approach, introduced in (Costi, 2022), the ground is removed with an upgraded version of the Line Fit Ground Segmentation (Himmelsbach et al., 2010). Then the clustering is computed using again a Connected Components approach, but in this case on the PM rather than on a BEV. The clustering algorithm is subdivided into two main steps. The first one works directly on the Point Map exploiting a Connected Component algorithm to compute neighbors and label them with the same ID, executing entirely on the GPU. In the second step, running on CPU instead, the neighbors' data extracted from the PM are elaborated to aggregate neighbors in different clusters. The results are reported in Figure 6.

### 3.2.4. Camera detection

The camera detection pipeline takes the camera frames as input and gives as output tracked detected vehicles. It executes both in CPU and GPU, but most of the computation is performed on the latter.

The detection of the other AV-21 vehicles is performed using an Object Detection Convolutional Neural Network. We adopted YOLOv4 (Bochkovskiy et al., 2020), implemented via tkDNN (Verucchi et al., 2020), a custom framework that optimizes its performance on Nvidia GPUs.

To correctly detect open-wheel race cars, we trained the model on an open source data set, Deep Drive BDD100k (Yu et al., 2020), to learn road objects (cars, bikes, pedestrians, and so on). Then



**Figure 6.** Results of the clustering PM pipeline. On the top, there are three LiDAR PM, in order without the ground, the plain PM, and the one with clusterized components. On the bottom, the clusters are projected on the cloud.

we collected data on different tracks (LOR, IMS, LVMS), down-sampled the different logs on the different cameras, and manually labeled almost 400 images using the LabelImg tool.<sup>7</sup> Using those labeled images, we fine-tuned the network only for the car class. Eventually, the network only detects open-wheel race cars.

Once trained, the network has been deployed on tkDNN<sup>8</sup>, which uses TensorRT<sup>9</sup> and CUDA kernels to optimize each network layer. The visual result of the detection trained network for all the cameras is reported in Figure 7.

To estimate the distance of the object, a simple geometric approach reported in Equation 2 has been adopted, based on the intrinsic calibration of the cameras (in particular, their `focal_length`), the actual height of the vehicle in mm `object_h_mm`, and the height of the detected vehicle in pixels `object_h_pixel`.

$$\text{dist} = \frac{\text{object\_h\_mm} \cdot \text{focal\_length}}{\text{object\_h\_pixel}} \quad (2)$$

We also implemented an NN-based method to estimate the distance, with a simple Encoded-Decoder approach, but the results were unsatisfactory.

Finally, the same tracker used for the LiDAR clustering BEV pipeline has been used to track the vehicles in the frames.

### 3.2.5. LiDAR detection

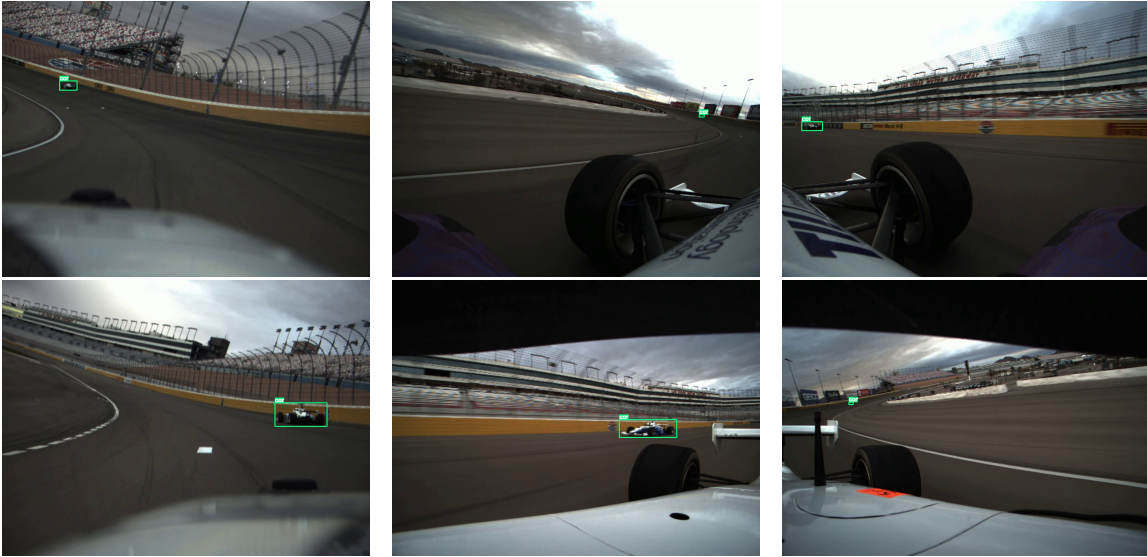
A very similar approach to the camera detection has also been applied to the LiDAR-based detection. From the point clouds, we have constructed LiDAR images based on the intensity of the points. Using this format, we have collected a data set of almost 600 images, and we have manually labeled them.

We modified Deep Drive BDD100k images to monochrome images, adapting the format to the LiDAR images one (16:1). We then trained YOLOv4 on the modified BDD100k to then fine-tune it on the 600 labeled images. The results of vehicle detection are reported in Figure 8.

<sup>7</sup> <https://github.com/heartexlabs/labelImg>

<sup>8</sup> <https://github.com/ceccocats/tkDNN>

<sup>9</sup> <https://developer.nvidia.com/tensorrt>



**Figure 7.** Results of vehicle detection on the six different camera views.



**Figure 8.** Results of vehicle detection on frontal (top) and lateral (bottom) LiDAR intensity images.

In this case, the objects' distance is given by the LiDAR, so there is no need for estimation. For the tracking, yet again we adopted the EKF tracker.

### 3.2.6. RADAR Detection

The RADAR detection pipeline takes the RADAR point cloud as input and gives as output tracked moving objects, executing on CPU.

The point cloud given by the RADAR is already processed, and it is not possible to retrieve the raw data. Therefore, we only applied filtering to the input data, considering only the stable moving objects lying inside the track boundaries.

### 3.2.7. Projection Fusion

The Projection Fusion pipeline takes as input the LiDAR point cloud, the RADAR point cloud, the detected vehicles from the camera, and the clusters of the LiDAR Clustering PM pipeline. It gives as output detected vehicles that have been recognized at least by two different sensors. It only executes on the CPU.

The proposed approach exploits camera projection to properly fuse detected objects from camera images with 3D estimations. The algorithm tries to estimate the 3D location in world coordinates for each detected vehicle. The projection converts vertices from the world coordinate system to the camera pixel coordinates system with Equation 3.

$$\begin{bmatrix} u \\ v \\ z \end{bmatrix} = K[R|T] \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}, \quad (3)$$

where  $K$  is the camera intrinsic matrix,  $[R|T]$  is the camera extrinsic calibration matrix,  $[u, v, z]$  is the undistorted point in camera pixel coordinates, and finally  $[X, Y, Z]$  are the real-world coordinates. It is worth mentioning that accurate intrinsic and extrinsic calibrations are required to reach satisfactory results.

The LiDAR point cloud is fused with camera detections with the following steps: (i) Equation 3 is applied to all the points of a cloud, for each camera. (ii) All the LiDAR points projected outside the camera bounding box are filtered out, therefore a frustum of 3D points is considered for each object. (iii) A single point in the frustum is chosen as the Point-Of-Interest (POI), sorting all the points by distance and picking the nearest (the first element), the median (the element in the middle of the array), or another custom array position. (iv) Finally, the location of the object is estimated as the average between all the points in the neighborhood of the POI.

The RADAR point cloud is fused differently with the camera detections. From the RADAR point clouds, we have a single point for each object, therefore a single point is fused with each camera bounding box with the following steps: (i) RADAR points are projected on-camera images. (ii) The matching cost is estimated as  $\frac{|cx-x|}{(w/2)}$ , where  $cx$  is the horizontal coordinate of the camera box center,  $x$  is the horizontal coordinate of the RADAR projected point, and  $w$  is the width of the camera box. (iii) The Hungarian algorithm (Kuhn, 1955) is applied to matchboxes and RADAR points using the cost previously computed. (iv) Finally, incorrect matches are filtered out via a user-defined threshold on the cost.

LiDAR clusters are considered as already detected 3D objects from another source. Similarly to RADAR-camera fusion, a single cluster is fused with each camera-detected object, with the following steps: (i) Each point of the cluster is projected on the camera, and the bounding box of these points is calculated. (ii) The matching cost is computed as the inverse of Intersection-over-Union (IoU) between the camera box and the cluster projected box. (iii) Hungarian matching is then in charge of the matching, and finally, (iv) bad matches are filtered out with a threshold.

At this point, there are multiple pairs of objects from the camera and another 3D source, and two steps are yet to be performed: fusion among all the cameras and fusion of all the object pairs. There are two straightforward cases: (i) if the objects belong to the same camera, the objects are fused with the same bounding box, and then it is the same object; otherwise, it is not; (ii) if detected vehicles from multiple cameras are fused with the same cluster, then it is the same object. For the other cases, aggregation is not as simple, and the proposed approach exploits 3D coordinates from fused 3D sources. A matrix cost is computed considering 3D box reprojection and Euclidean distance, combined with a weighted sum. Matches with costs exceeding a threshold parameter are discarded.

At this stage, a list of aggregated objects containing all the camera boxes and all the LiDAR, RADAR, and cluster positions has been obtained. Every 3D pose of the aggregated object has an associated score, calculated proportionally to the focal length of the camera fused with it. Higher confidence is given to objects detected from cameras with higher focal lengths because the field of view is narrower and the boxes are bigger. This score is used as a weight for a weighted average that gives the final aggregated object 3D position in world coordinates.

### 3.2.8. Sensor Fusion Module

All the presented pipelines flow into a Sensor Fusion module. This acts as an aggregator for all the different detection pipelines active on the machine. In particular, it deals with the transformation of the raw detections from local to global coordinates, the association between the new detections and the ones already tracked, and the prediction of their movement.

The node aggregates several pipelines of detections from various sensors located at different positions in the car, and each of them produces detections in its own reference system. To aggregate them together, we decided to transform every detection to global coordinates via Equation 4, using the ego vehicle position computed by the localization node,

$$T_{\text{global\_obj}} = T_{\text{loc}} T_{\text{sensor}} T_{\text{local\_obj}} \quad (4)$$

Each  $T_i$  is a  $4 \times 4$  transformation matrix.  $T_{\text{global\_obj}}$  is the object pose in global coordinates,  $T_{\text{loc}}$  is the car pose given by the localization node,  $T_{\text{sensor}}$  is the sensor pose relative to the car CoG, and  $T_{\text{local\_obj}}$  is the pose of the detected object locally to its sensor.

The estimated position of each detection is endowed with uncertainty, due to the method or the accuracy of the sensor itself, as well as the position estimated by the localization module. Therefore, when applying Equation 4, the localization error is propagated into the position estimation of the detections. For this reason, it is also necessary to propagate the localization covariance on the detection covariance.

The object tracking can work in two coordinate systems: (i) Cartesian  $(x, y)$  and (ii) Frenet  $(s, d)$ , in which the central trajectory of the track is used as a reference. In each case, a Kalman filter with a material point model is used. For the Cartesian version, the state is  $[x \ y \ V_x \ V_y]^T$  and the correction  $[x \ y]^T$ , while for the Frenet version, the state is  $[s \ d \ V_s \ V_d]^T$  and the correction  $[s \ d]^T$ . The Kalman filter calculations are the same in both cases; a simple Cartesian to Frenet transformation is applied on the input and vice versa on the output.

To summarize, the algorithm is composed of the following steps:

1. Kalman prediction.
2. Filtering of detections if (i) outside the track, (ii) inside the area of the ego vehicle, (iii) they have high covariance.
3. Detection association, with Hungarian matching and Mahalanobis distance.
4. Kalman correction.
5. Creation of new tracklets, i.e., tracked objects, if far from existing tracklets.
6. Removal of tracklets with too large covariance (not corrected for too long).

All the tracklets that are active are unique detections that are then passed to the motion forecasting module and finally to the planner.

### 3.3. Planning

To be able to safely avoid static obstacles and perform overtakes at high speeds, as requested by the IAC competitions at IMS and LVMS, in addition to the global planner, which produces offline the optimal racing line on each track, we implemented the modules needed to predict the other agent's movement and generate a local trajectory considering all the static and moving obstacles in the surroundings. We give here an overview of the proposed solution, while a more detailed report can be found in (Raji et al., 2022).

Considering the race rules limiting the defender to keep the inner line, one strategy could have been to switch directly to a racing line positioned in the outer lane of the track as soon as the ego car becomes the attacker. On the one hand, this approach can be considered safer since the two vehicles keep separate lines for most of the time except during the line switching performed in safe moments. On the other hand, staying in the outer lane where usually there is more dirt could result in less grip at high speeds and longer distances with respect to the defender's inner line. Considering our research interest in creating solutions suitable for unconstrained racing scenarios with more than two vehicles on a track, we decided to perform the overtakes once in proximity to the opponent keeping the same racing line followed while defending.

#### 3.3.1. Global Planner

A minimum-time optimization problem is solved for the global planning, formulating the nonlinear problem in JuMP and solving it using IPOPT. The dynamics of the vehicle, presented in Section 3.4, are transformed in the spatial domain discretizing the continuous space model with a discretization distance. The cost function is defined as

$$J_{\text{opt}}(x_k, u_k) = -\dot{s}_k + u^T R u + B(x_k), \quad (5)$$

where  $x$  is the state vector,  $u$  is the inputs vector,  $\dot{s}$  is the progress rate,  $B(x_k) = q_B \alpha_r^2$  is a regularization term that penalizes the rear slip angle  $\alpha_r$ , and  $u^T R u$  regularizes the inputs rates. The overall problem is defined as

$$\begin{aligned} \min_{X,U} \quad & \sum_{k=0}^N J_{opt}(x_k, u_k) \\ \text{s.t.} \quad & x_{k+1} = f_s^d(x_k, u_k), \\ & f_s^d(x_N, u_N) = x_0, \\ & x_k \in X_{track} \quad x_k \in X_{ellipse}, \\ & a_k \in \mathbf{A}, u_k \in \mathbf{U}, k = 0, \dots, N, \end{aligned}$$

where  $X = [x_0, \dots, x_N]$ , and  $U = [u_0, \dots, u_N]$  are the state and input sequences respectively.  $X_{ellipse}$  is a constraint on the friction ellipse, and  $X_{track}$  represents a track constraint.  $\mathbf{A}$  and  $\mathbf{U}$  are, respectively, box constraints on the physical inputs  $a$  and their rate of change  $u$ .

### 3.3.2. Motion Forecasting

The goal of motion forecasting is to estimate the future trajectory of the vehicles detected by the perception module. The estimated trajectories are then used by the motion planning algorithm to avoid collisions.

For each obstacle, the perception module provides a unique identifier and its position in a Cartesian frame. Given the sequence of the position of an obstacle, the goal is to predict its future trajectory. We employed a Kalman filter with a model defined in a Frenet frame.

Given the position of the  $i$ th obstacle in a Cartesian frame  $x_i(k), y_i(k)$ , the position of the obstacle in the Frenet frame  $s_i(k), n_i(k)$  is computed. The model of the obstacle is defined as

$$\dot{s}_i(k+1) = \dot{s}_i(k) \tag{6}$$

$$n_i(k+1) = n_i(k) \tag{7}$$

Equation 6 states that the longitudinal speed of the obstacle is constant, whereas Equation 7 indicates that the lateral displacement from the reference path is constant.

This simple model exploits the fact that the only objects of interest on the track are other cars that will follow a racing line similar to the one that the ego car is following. For this reason, we decided to define the model in a Frenet frame that uses the race line as the reference path. Moreover, we can assume that the cars will run all the time at an almost constant speed because the oval shape of the tracks involved does not require as many decelerations and accelerations as in a course road track. From Equations 6, 7 the state space model used in the Kalman filter can be derived.

To summarize, at each step, for every obstacle, the following steps are applied to predict its future trajectory:

1. The new measurement  $\hat{s}_i(k), \hat{n}_i(k)$  is computed from  $\hat{x}_i(k), \hat{y}_i(k)$ .
2. Using the new measurement, the Kalman filter is updated with a prediction step, followed by a correction step.
3. The future trajectory of the obstacle is predicted by applying  $m$  consecutive prediction steps to the Kalman filter.
4. The trajectory is converted back into the Cartesian frame.

### 3.3.3. Local Planner

The local planner is an extension of (Werling et al., 2010), as it computes the trajectory generation in a Frenet coordinate frame, where the following adjustments have been implemented to satisfy the

needs of our racing scenario:

- The main reference used for the Frenet frame is the optimal racing line generated by the global planner.
- The time interval  $T$  between each node of the trajectories is kept constant since the controller requires a trajectory with a fixed length in time.
- The collision check of the trajectories set is performed in the Frenet frame to avoid converting the trajectories into a Cartesian frame. Rather than doing the checks on the polynomials, we sampled each trajectory in a finite number of points by a time interval  $\Delta_T$ .
- An improved collision check method is used in which a soft constraint is added to avoid the edge cases when only hard constraints are considered and to be safer in case of noises in the localization and the control loop. For each trajectory  $\tau_i$ , a collision coefficient  $\gamma_i \in [0, 1]$  is computed, where  $\gamma_i = 0$  indicates that the trajectory is not colliding with any obstacle, whereas  $\gamma_i = 1$  indicates that the trajectory is violating the safety margins (hard constraint). Then the total cost computed is

$$C_{tot,i} = k_{lat}C_{lat,i} + k_{lon}C_{log,i} + k_{soft}\gamma_i^2 \quad (8)$$

with  $k_{lat}, k_{lon}, k_{soft} > 0$ .

To compute  $\gamma_i$  we decided to exploit the Euclidean distance from the safety margin. For every trajectory  $\tau_i$  the minimum distance  $d_i$  from the safety margin is computed. Then  $\gamma_i$  is defined as

$$\gamma_i = \max \left\{ 1 - \frac{d_i}{\Delta_{soft}}, 0 \right\} \quad (9)$$

where  $\Delta_{soft} > 0$  is a parameter to enlarge or reduce the effect of the soft constraint.

- The initial conditions required to generate the set of trajectories are calculated by projecting the car's position on the best trajectory at the previous step. At the very first step of the planner instead, the initial conditions are calculated purely on the car's position.
- A different distance keeping mode. The desired speed used in the generation of the longitudinal movements is calculated by a simple proportional controller, which considers the opponent's speed, the desired distance to keep, and the current distance. This mode is used when the rules or the race control are not permitting us to perform an overtake.

Along with the main planner, a simple emergency planner is run, whose solution is used when the Supervisor module detects a failure in the main planner or in the modules it depends on. The emergency planner continuously extends the last feasible planned path with a smooth polynomial in the Frenet coordinate frame, being able to make the car move to the inner side of the track or enter the pit lane, if requested by the race control.

### 3.3.4. Mission Planner

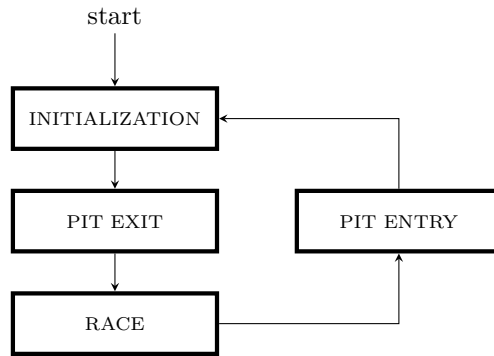
The mission planner is responsible for generating the reference signals and instructions used by the Local Planner based on the position of the car on the track, the phase of the race, and the flags received from Race Control. In particular, it controls when the car can enter or exit the pit lane, if the car has to perform the warm-up lap, which is the maximum allowed speed, whether the car is allowed to overtake or not, and which is the minimum distance that the car has to maintain from the opponent if it is not allowed to overtake.

A fundamental requirement for the mission planner is to be easy to change because it needs to rapidly adapt to changes in race rules. For this reason, we used a Finite-State Machine (FSM) to define the logic. A high overview of the FSM is shown in Figure 9.

For the implementation, we relied on `scxmlcc`,<sup>10</sup> an open-source tool that autogenerates the C++ code of the state machine from an XML file that defines the states, the events, and the transitions.

<sup>10</sup> scxmlcc: <https://github.com/jp-embedded/scxmlcc>





**Figure 9.** Overview of the mission planner state machine. Each of the shown states groups more specific states and transitions, which depend on the car’s position, actuators, and engine status, and external signals sent by Race Control.

### 3.4. Modeling and Control

Given the current state of the car, the controller computes the actuation commands to track the reference trajectory produced by the local planner illustrated in Section 3.3. At high speed, a Model Predictive Controller (MPC) is used to control simultaneously the steering, the throttle, and the brake. The vehicle model used in the MPC is a dynamic single-track model identified from a high fidelity multibody simulation. A kinematic model was also developed to work accurately at low speed. Due to the limited testing time, the integration between the kinematic and the dynamic models was not implemented, and the control at low speed (below 100 kph) has been delegated to a Pure Pursuit algorithm and a PID controller, respectively, for the steering and the pedals. A hysteresis and a consistency check on the steering wheel commands of the two controllers are applied to switch safely between the two solutions based on the operational conditions. The gearbox is controlled via a state machine that, based on the RPM of the engine, selects the appropriate gear.

#### 3.4.1. Modeling

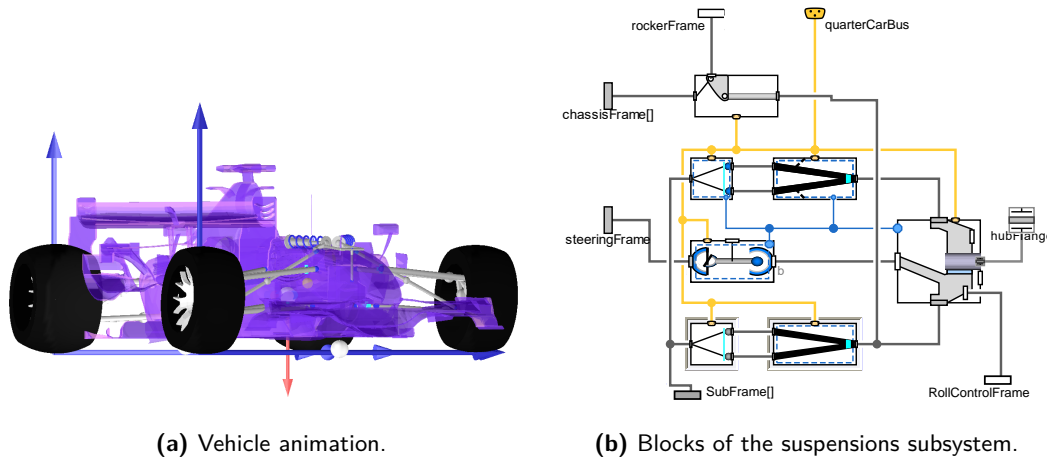
Prior to the testing time and physical access to the car, we developed a multibody model of the AV-21 on Dymola (Dempsey, 2006) using the VeSyMA - Motorsports libraries provided by Claytex.<sup>11</sup> The first set of parameters are derived from data provided by the IAC organizers and the vehicle’s component manufacturers. The remaining unknown details have been estimated from available information on similar vehicles and commercial racing-game simulators like RFactor 2.<sup>12</sup> The model has been refined after gathering experimental data on the track. In Figure 10, an overview of the model on Dymola is shown, in which the front suspensions subsystem blocks scheme is highlighted.

Particular attention has been given to the following components:

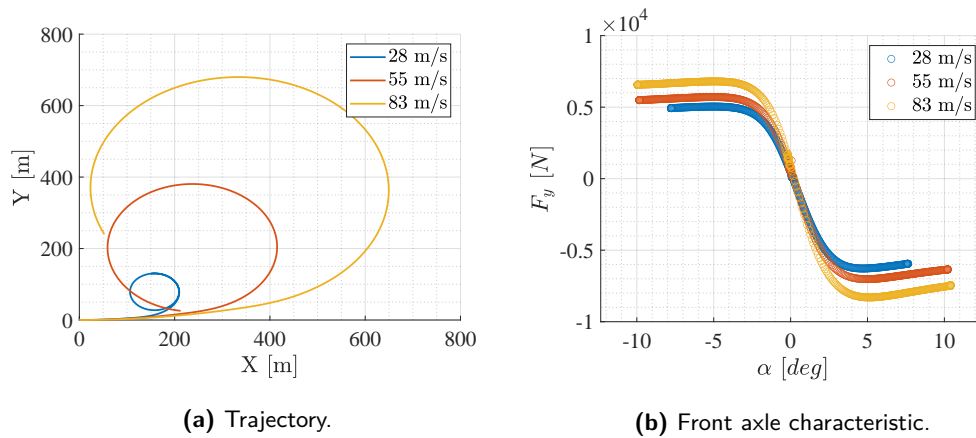
- Tires: the force-slip model is based on a Pacejka Magic Formula 6.2 (Pacejka and Bakker, 1991) with Kelvin-Voigt spring-damper vertical load, combined slip parameters and neglecting the relaxation length. Inflation pressure and camber angle are considered asymmetrical for the right and left sides due to the setup for an oval track.
- Powertrain: a 3D diagram with engine torque, RPM, and the throttle position has been implemented starting from the engine test bench data. The gear ratio, final drive, and shift time have been defined from telemetry and manufacturers’ data.

<sup>11</sup> <https://www.claytex.com/>

<sup>12</sup> <https://www.studio-397.com/>



**Figure 10.** Multibody model built on Dymola using the Claytex VeSyMA Motorsports libraries.

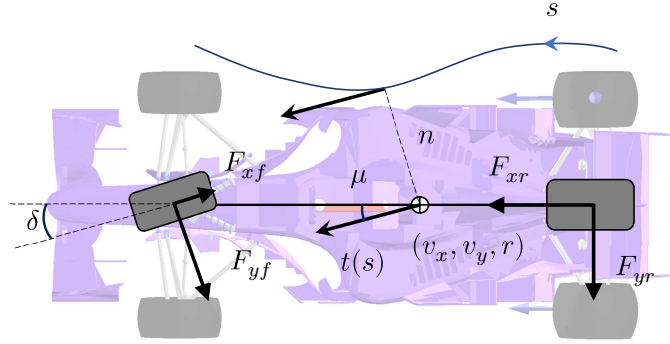


**Figure 11.** Virtual test results of the ramp steer manoeuvre at three different speeds.

- **Aerodynamics:** a simple model with drag and downforce coefficients is used. The center of pressure is positioned between the front and rear axle to define the correct aerobalance.
- **Suspensions:** the modeled components include the double wishbone geometry with a vertical antiroll bar, the rocker, and the shock absorber. Stiffness and damping are defined for all components.
- **Body:** the sprung and unsprung masses are modeled considering the center of gravity and cross-weight to validate the experimental data of static load balance. The inertia matrix is defined as well.

The multibody model has been used to produce manoeuvres that cannot be easily replicated on the real vehicle due to the lack of suitable space and limited testing time. In Figure 11 we report the ramp steer manoeuvres produced at different speeds. The wheels' steering angle is set to vary from zero to the maximum value at the rate of 1 deg/s, while the vehicle speed remains constant.

A single-track model on curvilinear coordinates, shown in Figure 12, has been identified from the high fidelity multi-body model, including the forces due to the road bank angle, the aerodynamic effects, the longitudinal forces on the rear axle generated by the turbo-charged engine, and the tire forces represented with a simplified Pacejka Magic Formula considering the vertical load and



**Figure 12.** Dynamic single-track model on curvilinear coordinates, where  $s$ ,  $n$ , and  $\mu$  are the progress along the path, the orthogonal deviation from the path, and the local heading.  $F_{x_{f,r}}$  and  $F_{y_{f,r}}$  are the longitudinal forces and lateral forces arising from the interaction between the tires and the ground.  $(v_x, v_y, r)$  identify the motion field of the center of gravity, and  $\delta$  is the steering angle.

camber angle as well as the combined slip effects. The equations of motion and the explanation of the modeled forces are presented in (Raji et al., 2022).

### 3.4.2. Pure Pursuit Controller

An extension of (Coulter, 1992) has been developed. The target point is chosen at a curvilinear distance `lookahead` from the projection of the car's position on the local path, hence the reference curvature is obtained as

$$k_{pp} = 2\psi_{\text{target}}/\text{lookahead},$$

where  $\psi_{\text{target}}$  is the angle of the target point position with respect to the  $x$ -axis of the local reference frame. The curvature is then converted to a steering angle at the wheels using the classical kinematic steering model:

$$\delta_{\text{wheel}} = \arctan(k_{pp} \cdot \text{wheelbase}).$$

The `lookahead` is updated at each step depending on the current speed and lateral error, in both cases with a contribution proportional to a reference value.

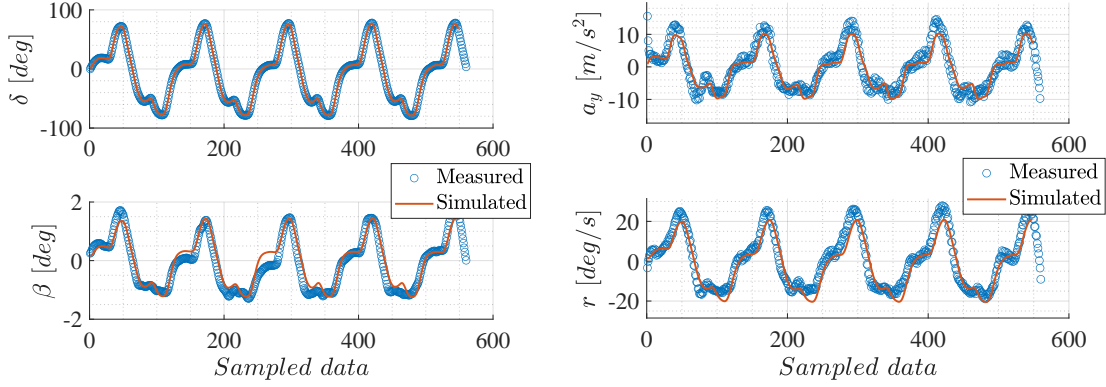
### 3.4.3. Warm-up manoeuvre

Without tire warmers and considering the low ambient temperatures during the race events, which were a maximum of 12.2°C (54°F) on October 21, 2021, at IMS and 17.2°C (63°F) on January 7, 2022, at LVMS, it was extremely important to heat the tires as much as possible. One approach, followed by the majority of the teams, consisted in incrementally increasing the speed of the vehicle during the first couple of warm-up laps following a normal raceline and running at least one lap at a speed higher than 50 m/s (180 km/h), where it has been demonstrated that the tires' temperatures increase rapidly. However, this approach produces higher energy and therefore higher temperature on the rear axle with respect to the front axle.

For this reason, during the first couple of laps, we performed an open-loop warm-up manoeuvre at 25 m/s (90 km/h) which consists of a series of  $\pm 80$  deg steering wheel angle commands on top of the Pure Pursuit algorithm.

The manoeuvre has been produced considering the following parameters:

- `steer_val`: the steering wheel angle that should be commanded during the manoeuvre;
- `step_duration`: the amount of time during which `steer_val` is kept;
- `step_gap`: the amount of time between two `step_duration`, during which the steering controller is not overridden;



(a) The steering angle  $\delta$  (a model input) and the sideslip angle  $\beta$  (an output signal).

(b) Lateral acceleration and yaw rate.

**Figure 13.** Virtual test results on the warm-up manoeuvre compared with the experimental data.

- `curvature_threshold`: a value for checking whether to reduce `steer_val` based on the curvature of the path in front of the car.

This solution aims to increase the temperature of the front tires, which is important to reduce the probability of occurring in an understeering condition, before setting a higher speed and following the global trajectory to increase more homogeneously the temperature on all the tires; some results of this manoeuvre will be discussed in Section 6.3.

The stability of the vehicle during the warm-up manoeuvre has been evaluated on the multibody model developed on Dymola. In Figure 13 simulation data are compared with measurements acquired during tests with an additional optical speed sensor.

#### 3.4.4. Model Predictive Controller

The MPC used at high speed is an extension of (Vázquez et al., 2020), where the optimization problem is formulated as in Section 3.3.1 using the model discretized in time  $f_t^d(x_t, u_t)$ . The main differences from the original work are:

- A more complex model
- The path and velocity produced by the Frenet-based planner are used as a reference to be tracked, considering the cost function

$$J_{MPC}(x_t, u_t) = -\dot{s}_t + q_n n_t^2 + q_\mu \mu_t^2 + q_v |s_{v,t}| + u^T R u + B(x_t),$$

where in addition to the terms used in Equation 5, it includes the path following weights  $q_n$  and  $q_\mu$ , and a velocity tracking weight  $q_v$  on the slack variable  $s_{v,t}$ .

- The optimization problem is solved using HPIPM (Frison and Diehl, 2020).
- An automatic differentiation library, CppADCodeGen,<sup>13</sup> is exploited to obtain the derivatives of the nonlinear differential equations of the model producing the source code, which is statically compiled offline and linked dynamically at runtime. This led to a speedup of the MPC keeping the computational execution below the 10 ms on high-end Intel processors such as E-2278GE, i7-10750H, and similar.

Considering the uncertainties of the dynamics of the actuator at speeds never tested before, and to cope with a potential model mismatch related to the force offset of the asymmetrical setup of the

<sup>13</sup><https://github.com/joaoleal/CppADCodeGen>

AV-21, we decided to set a high value to the costs on the physical inputs  $a$  and their rate of change  $u$ . This has been done expecting to avoid critical oscillations and to keep a smooth movement in exchange for a slower system and a potentially higher path tracking error.

### 3.4.5. Controller Mux

The Pure Pursuit controller (Section 3.4.2) and the Model Predictive controller (Section 3.4.4) run in parallel, each producing a control command. Both of these commands are sent to a Controller Mux node, which selects one of the two sources based on their priority and availability, and routes its message to the hardware.

In `er.autopilot 1.0`, the Model Predictive controller has the highest priority, followed by the Pure Pursuit controller. In the rare event in which the Model Predictive controller fails to find a solution or it does not provide a control message at the required rate, the Controller Mux switches to the Pure Pursuit controller. As mentioned in Section 3.4.4, the MPC is not used at low speed, therefore the Controller Mux uses the Pure Pursuit for this case as well. The decision is made by checking a flag sent by the controllers indicating whether their commands should be applied.

During a switch, the commands are interpolated to smoothly match the ones of the new command source. This is done to avoid sudden changes in the control commands that could lead to undesired behavior.

## 3.5. Supervisor and Safety Layer

### 3.5.1. Supervisor

The supervisor module coordinates all the software modules. In particular, it takes part in the start-up sequence of the car and commands an emergency stop if an anomaly is detected by the failure detection module. It listens to the Mission Planner presented in Section 3.3.4, the Race Control, and to the joystick used by the pit crew to trigger a manual emergency stop. Besides this main Supervisor module, `er.autopilot 1.0` uses a concept we called MicroSupervision. Each ROS2 node of the software stack has some checks on the availability of the most important topics, such as the vehicle state (position and velocity), the commands feedback, and the status of other modules. In particular, the controller base node has the possibility to directly stop the car. This can be seen as a redundancy in the general safety system of the architecture.

### 3.5.2. Failure detection

The failure detection module is responsible for detecting anomalies in the system. One of its tasks is to monitor the signal of all the car's sensors to check if the values are in the nominal ranges or if the sensors give the correct outputs, excluding, for example, *NaN* or values with a wrong scale/range. Some sensible parameters related to the engine, transmission, fuel, and battery have additional checkups related to the optimal operating range in order to guarantee peak performance. When some of these sensors are out of their optimal values, but in acceptable ranges for a certain amount of time, the failure detection module sends a warning to the supervisor. On the other hand, if the sensors reach critical values, the emergency signal is triggered.

Another task of this module is to monitor the status of all the software stack in order to notify the supervisor if some of them trigger an error state or stop working. In the latter, the failure detection module monitors the time stamps of the messages sent for communication purposes in order to check for timeout conditions, and as a redundancy it takes advantage of the QoS (Quality of Service) API exposed by the ROS2 middleware interface (namely DDS<sup>14</sup>) in order to have confirmation for a potential crash of the modules.

Furthermore, it checks that the connection between the car and the base station is alive. If an anomaly is detected, an error is sent to the supervisor module, which reacts accordingly.

---

<sup>14</sup>Data Distribution Service: <https://www.omg.org/spec/DDS/>

## 4. Simulation

Two simulation environments have been used to test the software stack, considering the following criteria:

- Vehicle Dynamics fidelity: the simulated vehicle handling should behave similarly to the real one, and it should be easy to test different road friction coefficients, tire temperatures, and tracks.
- Simulation to Reality gap: should be limited to the differences between the reality for the steps followed on the real car for pit entry and exit, and the signals sent between the race control and the pit crew. The sensors' interfaces and communication protocols used should be replicated as well.
- Ease of use: each team's developer should be able to run the entire software stack and the simulator on the same machine, and easily restart and change the simulation scenario.

A single simulator that satisfies all these conditions is still in development, where the aim is to include the multibody model developed on Dymola into the simulator described in Section 4.2.

### 4.1. AssettoCorsa

AssettoCorsa<sup>15</sup> is a racing-game simulator developed by Kunos Simulazioni. It is popular for its realistic dynamics and because it can be easily extended with custom vehicles, tracks, and plugins. Furthermore, the simulator exposes an interface in Python to retrieve in real time detailed data related to the running vehicles, such as position, velocities, accelerations, tires, and aerodynamics.

We developed additional interfaces to send the actuation commands, and we created the ROS2 wrappers to use the same messages of the real system. We started with custom mods of the Dallara IL-15 and the oval tracks available online. The car model has been adjusted to replicate the engine map, setup, and tire model of the Dallara AV-21. Our Motion Planning and Control algorithms have been heavily tested in this simulation environment, in which it is possible to easily produce challenging scenarios changing several parameters, such as the road friction coefficient, car setup and stability, wind, and slipstream effects, as well as running against multiple AIs or human-driven agents. A Windows machine is dedicated to running the racing game and publishing the ROS2 messages, whereas a separate machine with Linux runs a version of the `er.autopilot 1.0` software disabling some nodes related to the Perception, and adapting the parameters related to Race Control and other communications that are not replicated on the simulator.

Further details on the interfaces, customization, and potential contribution of this simulation platform will be presented in a separate work.

### 4.2. Unity-based semi-HiL Simulator

Besides AssettoCorsa, we decided to implement a lightweight Unity-based simulator to test all the software stack onto, which has the same interface as the real car, and is easy to install and use.

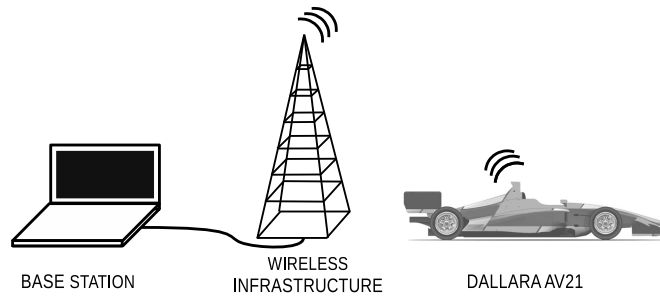
It is the semi Hardware In the Loop (HiL) approach, given that the communication is the same as the car at the lowest level possible, in particular:

- The Raptor and MyLaps communicate via a virtual CAN interface as the real car;
- The GPS is simulated and sent via TCP, using messages formatted as for the real Novatel GNSS modules;
- The LiDAR is simulated and sent via UDP, using messages formatted as the real Luminar.

Moreover, the race track in the simulator is georeferenced as well, therefore the GPS positions coincide with reality.

---

<sup>15</sup> <https://www.assettocorsa.it/>



**Figure 14.** The Dallara AV-21 communicates with the base station via a wireless infrastructure made up of multiple antennas placed around the track. The base station is connected to the infrastructure using an Ethernet cable.

The car dynamics are provided by the NWH Vehicle Physics 2,<sup>16</sup> and then they have been tuned to match our vehicle settings as closely as possible. Despite the ease of adjusting the vehicle model, it has not been possible to reach fidelity on the lateral dynamics as accurately as on AssettoCorsa or Dymola.

The simulator was designed to be used on the real hardware in an HIL fashion, through CAN and Ethernet connections. Nonetheless, it can run in a Software In The Loop (SIL) fashion on any high-end laptop, such as those used by the team for development. In particular, this simulator has been used by the developers to implement and validate the correctness of modules related to the system integration, such as Mission Planner (Section 3.3.4), Controller Mux (Section 3.4.5), Supervisor, and Safety (Section 3.5). Finally, there is also the possibility to run the simulator automatically, with a predefined mission, and headless, without visualization. We exploited this feature to include a simulation test in our GitLab pipelines.

## 5. Telemetry and Visualization

The Dallara AV21 car constantly communicates with an on-ground computer called a base station. The base station and the car communicate via a wireless infrastructure that is made up of multiple antennas placed around the track (Figure 14).

From the base station, it is possible to send commands to the car and monitor all the signals that are relevant to evaluate the performance of the car. From a joystick connected to the base station, it is possible to reduce the speed of the car and command an emergency stop. If communication with the base station is lost, the car performs a graceful stop.

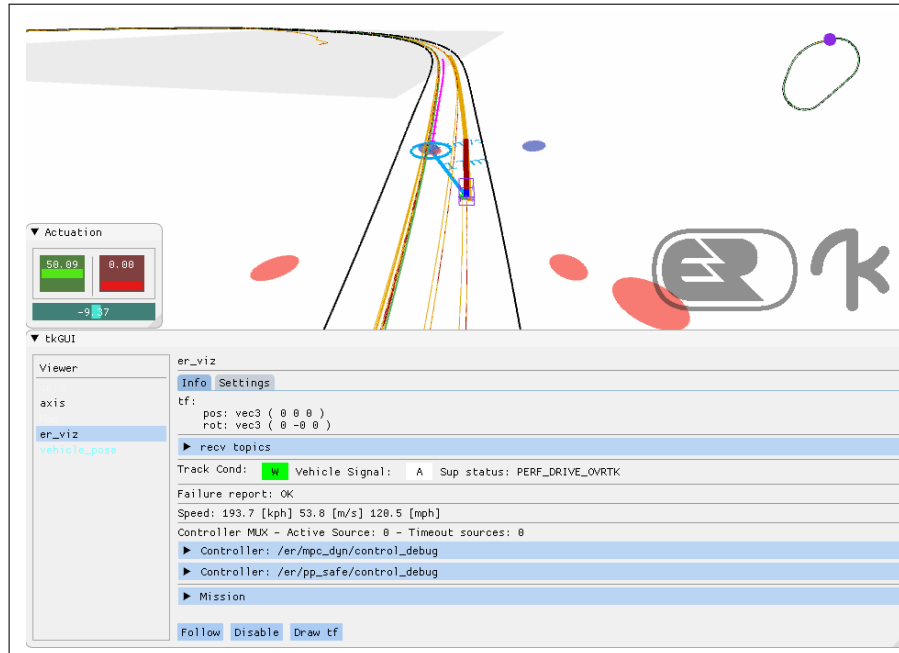
### 5.1. Telemetry Data

To overcome the issue of the limited bandwidth of the wireless infrastructure, the TII Euroracing team has implemented a proprietary protocol based on UDP that drastically reduced the amount of data going through the infrastructure. All the signals coming from the car are downsampled to 5 Hz and compressed before sending them to the base station.

### 5.2. Visualization

While the car is running, on the base station the proprietary software *er.viz* (Figure 15) allows the team to visualize the car position, the planned trajectory, and the detected obstacles, along with the car speed.

<sup>16</sup> <http://nwhvehiclephysics.com>



**Figure 15.** `er.viz` is the visualization tool used on the base station to monitor the performance of the car. The car position is shown in the center of the window; the detected car is marked with a blue circle; the desired trajectory is drawn with a thick yellow line, while the motion forecasting of the detected car is drawn in purple.

Along with `er.viz`, on the base station the open-source software *PlotJuggler*<sup>17</sup> is used to plot the signals in real time. The most relevant signals monitored by the team during a run are the lateral error from the desired trajectory, the steering and throttle commands, the tires' temperatures, and the covariances of the localization.

## 6. Results

In this section, we report the main results for each of the presented software modules.

### 6.1. Localization

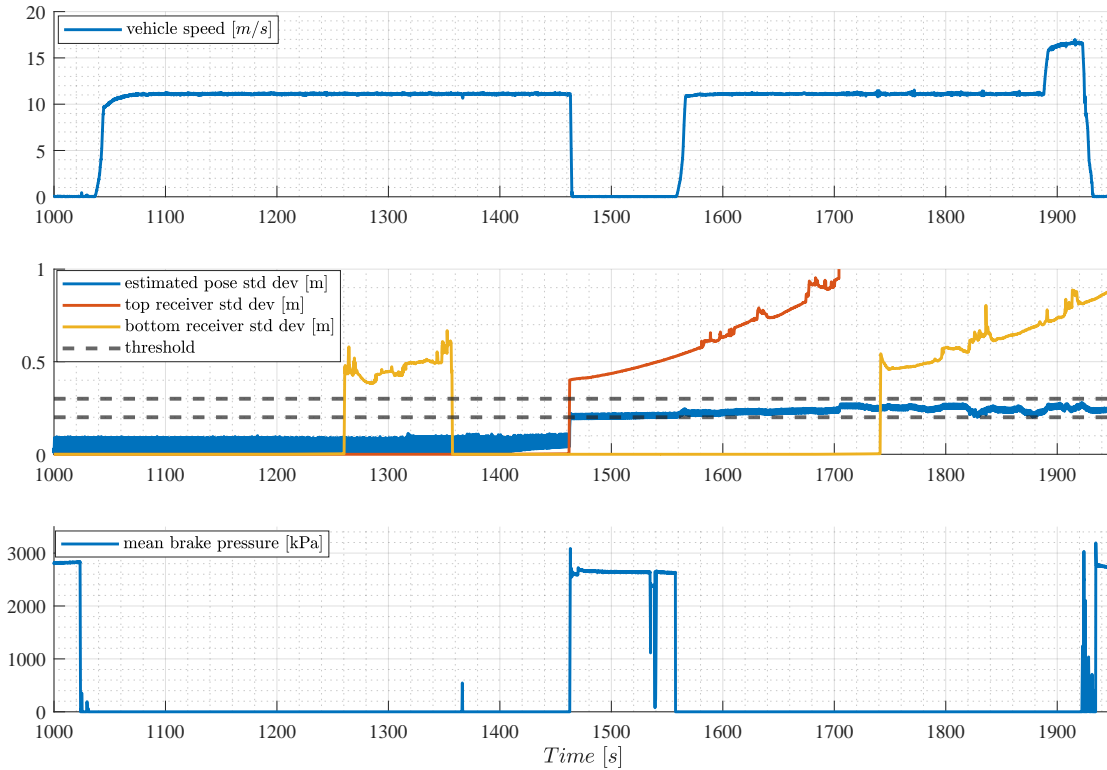
In evaluating the quality of the localization system, we could not rely on a ground truth system for comparison. We proceeded, keeping in mind the following objectives:

1. Empirically compare the ego vehicle estimate with the GNSS raw inputs, also considering their covariance.
2. Evaluate the performance of the estimator in case the RTK signal was lost on one or both receivers.
3. Understand the practical performance of the localization system and derive safety thresholds based on its accuracy and tolerance to sensor malfunctions.

In Figure 16 we show the result of some of these tests conducted at LOR, where each sensor RTK correction is disabled/enabled in different combinations.

<sup>17</sup>PlotJuggler: <https://www.plotjuggler.io/>



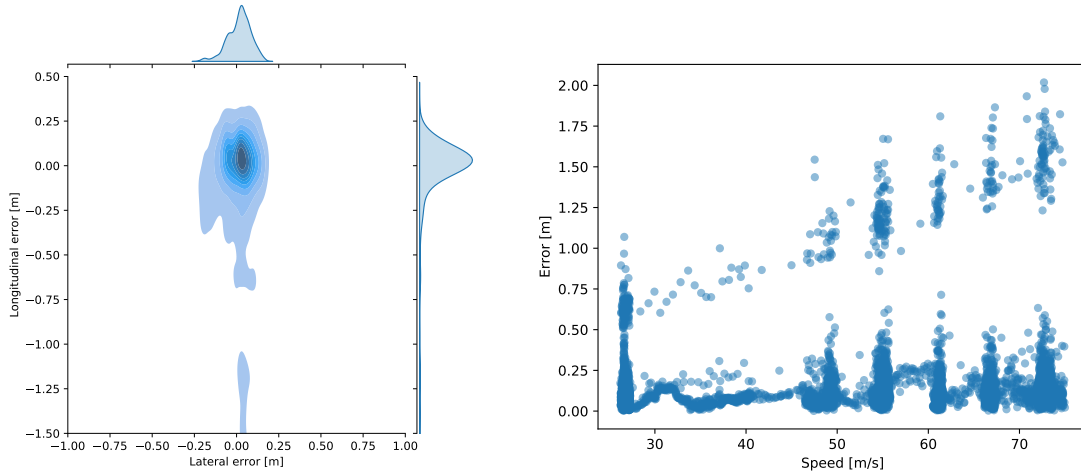


**Figure 16.** Ego-vehicle localization in the absence of RTK correction. i) In the first test, the RTK correction is disabled for the bottom receiver around 1250 s. Given the tuning used in the test, this does not affect the accuracy of the estimation. The correction is then reapplied. At around time 1450 s, the top receiver RTK correction is disabled, which immediately brings the pose estimate standard deviation (average over latitude and longitude) over the first threshold of 0.2 m triggering a safe braking of the car. ii) The safety threshold is increased to 0.3 m and the test continues. It is visible how the top receiver accuracy quickly degrades. Once also the bottom receiver RTK correction is disabled again at around 1750 s, the accuracy is still within range and the car keeps moving. The car is then stopped manually and the test is considered successful.

The importance of these tests lies in the fact that they gave us confidence in the accuracy of the car pose estimates, which could then be used to calibrate the safety thresholds for the safe-stop signals in the failure detection module and to confidently increase the operating speed during the test sessions. The LiDAR localization system described in [Section 3.1.2](#) is treated as an additional, virtual, sensor in the EKF, providing position, heading, and velocity estimates with their confidence at around 25 Hz. Given the limited amount of track time we had to test this critical software component, we decided not to use it in the final races in order to advance its development and present more results in future works.

However, we report our current results in the charts of [Figure 17](#). On the left, we depict a joint plot with the distribution of the lateral and longitudinal error in meters. The error has been computed considering the RTK-corrected GPS position as ground truth, on the qualification log for the LVMS race where the car exits the pit, increments the speed up to 75 m/s, and goes back to the pit. The origin in the plot represents the car’s Center of Mass. The results are very promising: the maximum lateral error is 25 cm, while the longitudinal one is on average around 30 cm, with few peaks close to 2 m.

The chart on the right shows the error in m ( $y$ -axis) at the different speeds in m/s ( $x$ -axis). From this second plot, we can notice that the error peaks increase while increasing the speed, as one could expect. We can then notice that the 2m errors are experienced only over 70 m/s, while at lower speeds the maximum error is always lower. Moreover, the chart shows two distinct patterns



**Figure 17.** Results of the LiDAR localization at LVMS. On the left is the joint plot of the lateral error (x-axis) and the longitudinal one (y-axis). On the right is the plot of the error module (y-axis) at different speeds.

**Table 1.** Accuracy results of the Object Detection of the camera detection pipeline (Section 3.2.4).

Train-set Size	Val-set Size	Classes	AP 0.5	AP 0.75	AP 0.5:0.95
350	50	1	0.92	0.67	0.60

of errors: frequent low errors at the bottom of the chart, and higher errors in the upper part. These patterns appear to be correlated with straight sections and curved exits, respectively.

## 6.2. Object Detection and Tracking

To correctly evaluate the perception pipeline, we have to consider both execution time and accuracy.

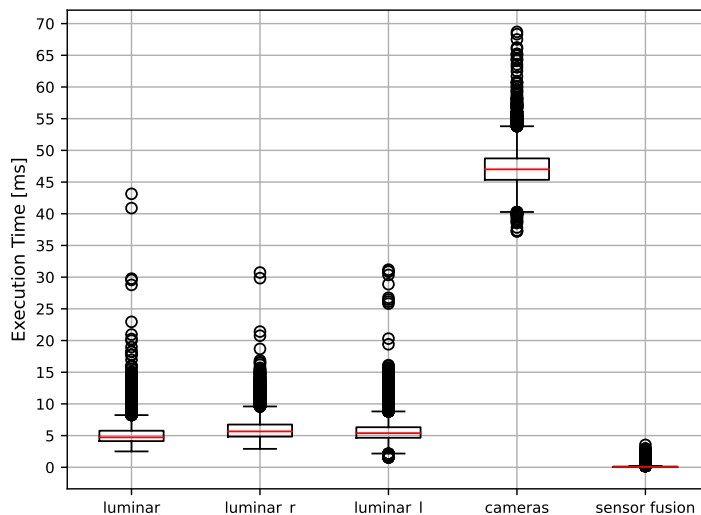
Given that there is not an official data set on which to evaluate our methods, we will report the accuracy results on the data we have manually labeled. For the same reason, we sometimes report both boxplot graphs and tables to serve as reference in future works.

For the Object Detection NN of the camera detection pipeline (Section 3.2.4), the accuracy results have been reported in Table 1. The data set statistics reported refer to the manually labeled images used to fine-tune the YOLOv4 pretrained on BDD100K. The Average Precision (AP) obtained for open-wheel race cars of our solution is 92% when using a confidence threshold of 0.5.

Regarding the execution time, Figure 18 shows the boxplots of the LiDAR clustering BEV (Section 3.2.2), camera detection (Section 3.2.4), and sensor fusion (Section 3.2.8) for 16 600 iteration of one complete log run where these pipelines were active. The execution times refer to the vehicle’s workstation. The three Luminar LiDAR are handled by three separated nodes, whose boxplots are `luminar`, `luminar_l` and `luminar_r`, while all the cameras are elaborated by a single node to have a single batched inference (therefore there is only one boxplot). Please keep in mind that the execution times provided were calculated while the computer was operating at full capacity, running the entire stack, rather than in isolation. As a result, the execution times for the same algorithm (such as `luminar`, `luminar_l`, and `luminar_r`) may fluctuate slightly.

Moreover, these pipelines’ statistics are also reported in Table 2. On average, the BEV clustering with its internal tracking is performed in 5.5 ms, the vehicle detection and tracking on six camera frames in 47 ms, and all the data are fused in 0.14 ms. Unfortunately, we have not recorded accurate profile data for the other pipelines.

Finally, we evaluated the result of the Sensor Fusion pipeline (Section 3.2.4), which is the module in charge to give the input to the Planner, when merging detections from the LiDAR clustering



**Figure 18.** Boxplots of the execution times of the LiDAR clustering BEV (Section 3.2.2), camera detection (Section 3.2.4), and sensor fusion (Section 3.2.8) pipelines, over 16 600 iterations.

**Table 2.** Execution time statistics of the LiDAR clustering BEV (Section 3.2.2), camera detection (Section 3.2.4), and sensor fusion (Section 3.2.8) pipelines, over 16 600 iterations.

	avg [ms]	max [ms]	min [ms]
luminar	5.08	43.14	2.51
luminar_r	5.93	30.73	2.91
luminar_l	5.66	31.18	1.53
cameras	47.05	68.69	37.21
sensor fusion	0.14	3.54	0.04

BEV (Section 3.2.2) and the Radar (Section 3.2.6) pipelines on the head-to-head run against TUM, during the LVMS event.

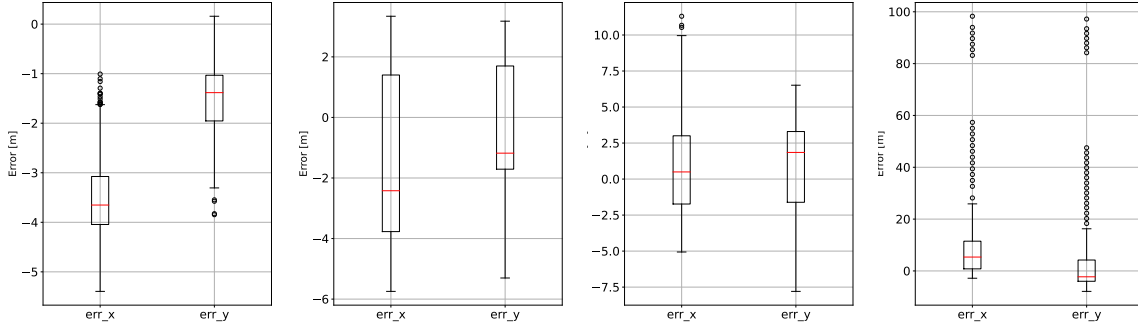
To compute that, we compared the GPS position of TUM’s car, using the output of the Novatel top (in front of the car), and the GPS position of our vehicle, using the output of the Novatel top (in front of the car), and the detection computed by the Sensor Fusion module, which is the center of the detected objects, aligning our data with the opponents’ ones with the Novatel’s time stamp. We then converted the GPS positions into local coordinates to compare them with the detections, and we only considered the case in which the opponent’s car is not behind ours.

It is important to notice that the considered ground truth (TUM’s position) is a point in the front of the car, while the detection is usually a point in the back of the car, and that the car is almost 5 m long and 2 m wide. Keeping that in mind, to compute the goodness of the pipeline, we considered various scenarios in which the two cars’ Euclidean distance was in a certain range.

This having been said, we evaluated the ranges [0–10] m, [10–25] m, [25–50] m, [50–100] m, [100–150] m, and [0–150] m. Table 3 reports the complete evaluation for those ranges, including True Positives (TP), False Negatives (FN), False Positives (FP), precision (p), recall (r), and longitudinal and lateral errors statistics, such as minimum (min), maximum (max), average (avg), and median (med). It is worth noting that the recall is 1.0 within 50 m, and the worse value is 0.65 in the [100–150] m range; on the other hand, the precision is around 1 when the distance is greater than 50 m. For smaller distances, the precision is worse due to the error of the detection, e.g., considering the range [0–10] m if the ground truth position is around 11.0 m (out of the considered range), while its detection is around 9 m and that is considered FP.

**Table 3.** Results of the Sensor Fusion (Section 3.2.8) output (center of detection) compared with the TUM position (Novatel top). The distance between TUM's car and ours is in the range [Min dist-Max dist]m, and all the error statistics (minimum, maximum, average, and median) are in meters.

Min dist	Max dist	TP	FN	FP	p	r	$min_x$	$max_x$	$avg_x$	$med_x$	$min_y$	$max_y$	$avg_y$	$med_y$
0.00	10.00	141.00	0.00	57.00	0.71	1.00	0.48	7.13	-2.71	-2.86	0.05	2.10	-1.06	-1.10
10.00	25.00	298.00	0.00	69.00	0.81	1.00	1.01	5.39	-3.40	-3.65	0.04	3.85	-1.48	-1.38
25.00	50.00	609.00	0.00	80.00	0.88	1.00	0.71	5.74	-1.85	-2.42	0.09	5.30	-0.88	-1.18
50.00	100.00	2,814.00	101.00	34.00	0.99	0.97	0.00	11.30	0.54	0.50	0.00	7.80	0.72	1.85
100.00	150.00	598.00	327.00	0.00	1.00	0.65	0.01	98.28	10.40	5.36	0.12	97.20	4.40	-2.26
0.00	150.00	4,460.00	428.00	0.00	1.00	0.91	0.00	98.28	0.41	-0.66	0.00	97.20	0.49	-0.10



**Figure 19.** Statistics of the longitudinal and lateral error of the Sensor Fusion's output (Section 3.2.8) compared with the TUM position (Novatel top). The distance between TUM's car and ours is, from left to right, in the ranges [0–25] m, [25–50] m, [50–100] m, and [100–150] m.

In addition to the table, the lateral and longitudinal errors of the detections are depicted in Figure 19, in particular for the ranges [10–25] m, [25–50] m, [50–100] m, and [100–150] m. From the range [10–25] m, the average longitudinal error is  $-3.40$  m, which is almost the distance between the top position (ground truth) and the rear one (center of detection seen from behind). More generally, the maximum longitudinal error is around  $10.4$  m and the lateral error is around  $4.4$  m (without considering the actual size of the cars) within  $100$  m of distance. It has greater outliers, up to  $98.28$  m of longitudinal error and  $97.2$  m of lateral error when the distance between the two vehicles is over  $100$  m; the cause can be found in the only presence of radar detection, less reliable than LiDAR, at that range and the divergence of the Kalman Filter on the banking.

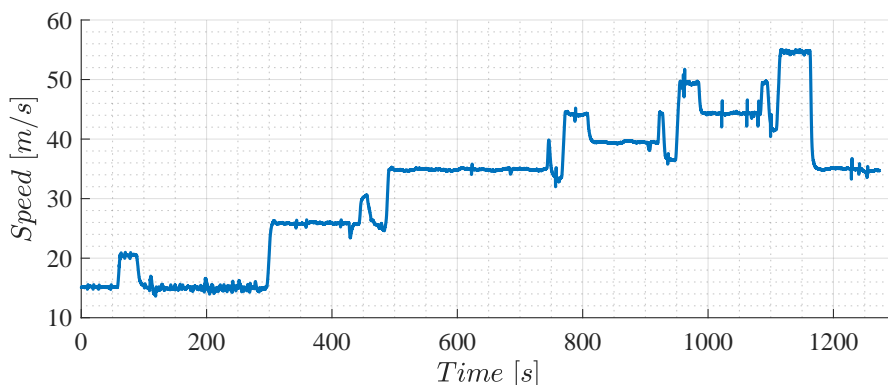
### 6.2.1. Motion Forecasting

A complete analysis of the performance of the motion forecasting module would not have been possible without the help of the TUM team, which provided us with the GPS log of their car. The GPS position is extremely accurate thanks to the RTK correction, so it has been used as the ground truth to evaluate the output of the motion forecasting module. Without using the log of another car, the only way to analyze the performance would have been the usage of the perception module output as ground truth. However, the estimation error of the motion forecasting module would have been affected by the error of the perception module.

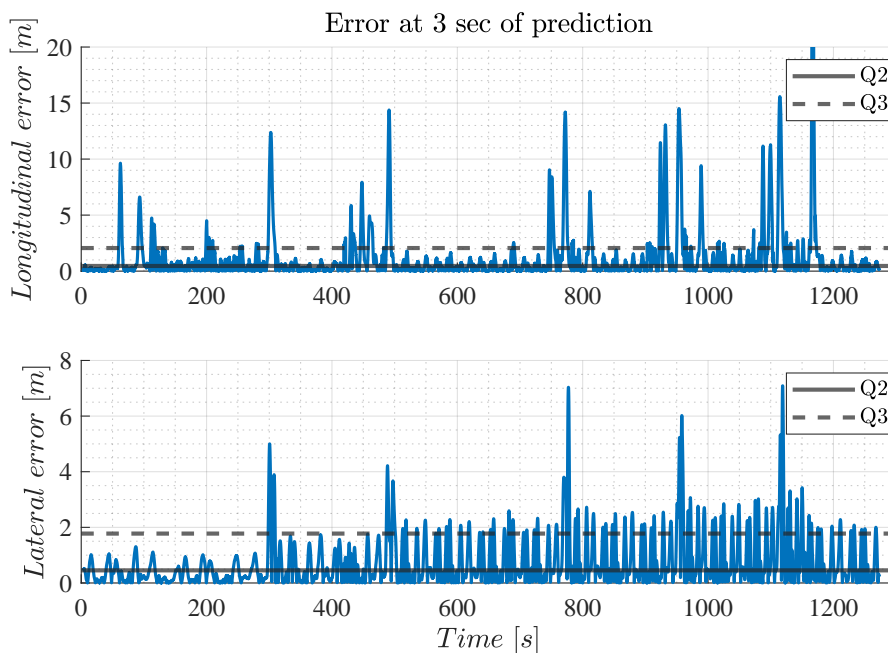
In the data set used, the TUM car did several laps at different speeds, ranging from  $10$  m/s up to  $55$  m/s (Figure 20). The position of the car at each step has been used directly as input to the motion forecasting module.

The predicted trajectory of the motion forecasting module has a length of  $3$  s,<sup>18</sup> so to evaluate the accuracy of the prediction, the prediction error is computed at different prediction lengths. Moreover, the error is split along the longitudinal and lateral axes. In Figure 21, the error at  $3$  s of

<sup>18</sup> Length in time: the last point of the predicted trajectory is the predicted position of the car in  $3$  s.



**Figure 20.** Speed profile in the log provided by the TUM team. The data set contains multiple laps at different speeds.



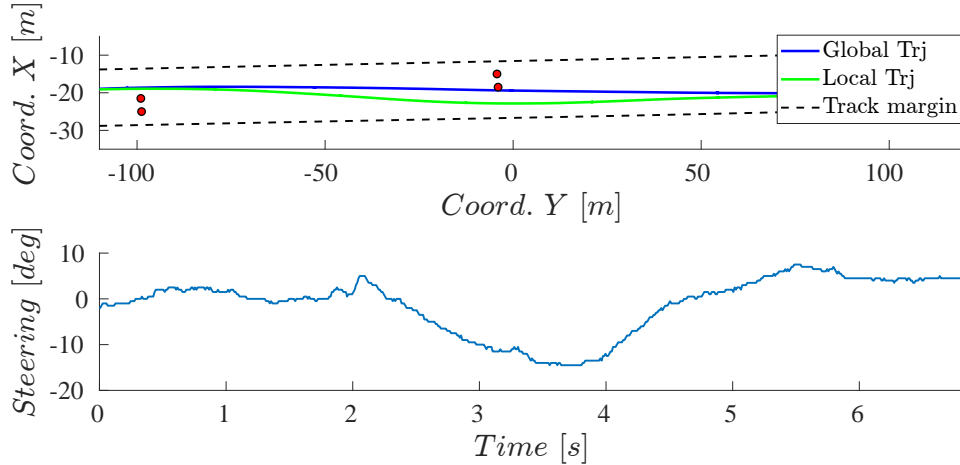
**Figure 21.** On the top graph, the longitudinal error of the motion forecasting is illustrated at 3 s of prediction, whereas the bottom one contains the lateral error. The output of the motion forecasting module is 3 s long. For convenience, the statistical quantities are also reported in Table 4.

prediction is shown, along with the median (Q2) and the 75th percentile (Q3). Table 4 summarizes the statistical characteristic of the error at 1, 2, and 3 s of prediction.

By combining the speed profile (Figure 20) and the prediction error (Figure 21), we can see that the error increases with the speed. Furthermore, the outliers are mostly due to the acceleration of the car. This phenomenon is expected because the model Equation 6 used to make the prediction assumes that the car is moving at a constant speed. However, the overall prediction error is promising: at 3 s of prediction the Q3 of the error is less than half of the car length on the longitudinal component and slightly more than half of the width of the car on the lateral component. This result lies in the same order of magnitude of the estimates performed relying only on the logs of our car, and it helped us in tuning the safety thresholds in the planning stack.

**Table 4.** Statistical quantities of the error of prediction at various prediction steps.

	Longitudinal Err. (m)			Lateral Err. (m)		
	1 s	2 s	3 s	1 s	2 s	3 s
min	0.0	0.0	0.0	min	0.0	0.0001
Q1	0.04	0.09	0.16	Q1	0.0659	0.1939
Q2	0.107	0.26	0.49	Q2	0.1646	0.454
Q3	0.27	0.71	1.34	Q3	0.4308	1.2053
max	7.91	19.75	34.19	max	3.7161	7.0894

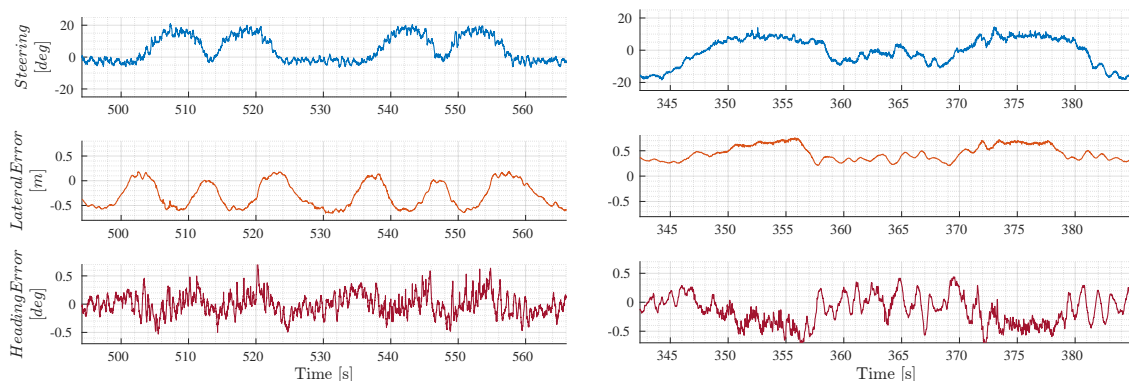
**Figure 22.** Avoiding pylons during the Semifinal at IMS. The location of four water-filled pylons is marked with red dots. The local planner is informed by the detection stack and reacts with a smooth trajectory to avoid them.

### 6.3. Planning and Control

Thanks to the optimization-based global planner, it has been possible to generate a feasible raceline that does not take into consideration only the time minimization but also the rate of change of the steering wheel angle. Considering the uncertainties of the actuators, the model mismatch on the controller, and the need for a feasible path at all the possible velocities, we preferred to follow a smoother and safer trajectory than the potential minimum lap trajectory.

The Frenet-based planner and the Control module have been tested thoroughly at various speeds, as presented in (Raji et al., 2022).

In particular, the planner has been capable of generating smooth overtaking trajectories with a speed that went from 22 m/s ( $\sim 80$  km/h) up to 64 m/s ( $\sim 230$  km/h), as well as performing safely static obstacle avoidance, as can be seen in Figure 22. During the high-speed laps at LVMS, when the speed varied from 72 m/s to 75.5 m/s, the MPC reached a maximum lateral error of 1 m and an RMS value of 0.5 m, while the heading error was between 0.7 deg and  $-1.0$  deg. On the other side, when the speed varied from 61.45 m/s to 63.16 m/s during the final at IMS, the maximum lateral error was  $-0.67$  m and the RMS value  $-0.29$  m, while the heading error was between 0.6 deg and  $-0.5$  deg with a mean of  $-0.016$ . In Figure 23, we show the tracking performance during the two events at similar speeds over one lap. The error on the lateral tracking has been influenced by the choices on the regularization terms and the mismatch in the tire model. The differences, mainly on the lateral error, are caused by changes in the weights' values in the MPC cost function we made, keeping unvaried the parameters of the Pacejka Magic Formula, from the first identification explained in Section 3.4, due to insufficient time to safely validate on track a set of potentially more accurate coefficients.



(a) MPC Tracking Errors during a 215 kph lap at IMS. (b) MPC Tracking Errors during a 220 kph lap at LVMS.

**Figure 23.** Comparing controller tracking performance at a similar speed in different circuits. Despite different tuning values being used, consistent behavior can be observed.

**Table 5.** Tires temperatures during the final at IMS. In the absence of tire warmers, and despite a strategy comprising warm-up manoeuvres and fast laps, the achieved temperatures after six laps are far from the nominal ones.

	Pit-exit Temp. [°C]	Temperature Gap [°C]			Maximum Temp. [°C]
		Warm-up (1 lap)	Target 60 m/s (2.5 laps)	Target 66 m/s (2 laps)	
Front Left	25.5	+ 9.9	+ 9.6	+ 2.8	51
Front Right	20.3	+ 10.5	+ 18.4	+ 4	53.9
Rear Left	21.6	+ 7.3	+ 11.3	+ 1.8	45.2
Rear Right	18.8	+ 8.4	+ 21	+ 5.8	57.8

The results of the Warm-up manoeuvre described in [Section 3.4](#) are presented in [Table 5](#) and [Figure 24](#), where data from the final event at IMS are used. It has been possible to increase the temperature on the front tires before increasing the speed, as we were aiming for. As we saw in the simulations, this manoeuvre leads to a controlled front slip which can generate front temperature. It is noticeable that the rear tires' temperature increased when the car ran at 60 m/s while the temperature of the front tire began to flatten. The greater temperature on the right side is due to the banking effect and the high load transfer, which increases the vertical load on those wheels. This leads to deformation of the tire carcass, which is converted into energy, and therefore heat. This effect is greater on the rear right due to the car balance (both weight and aerodynamic) and the combined force.

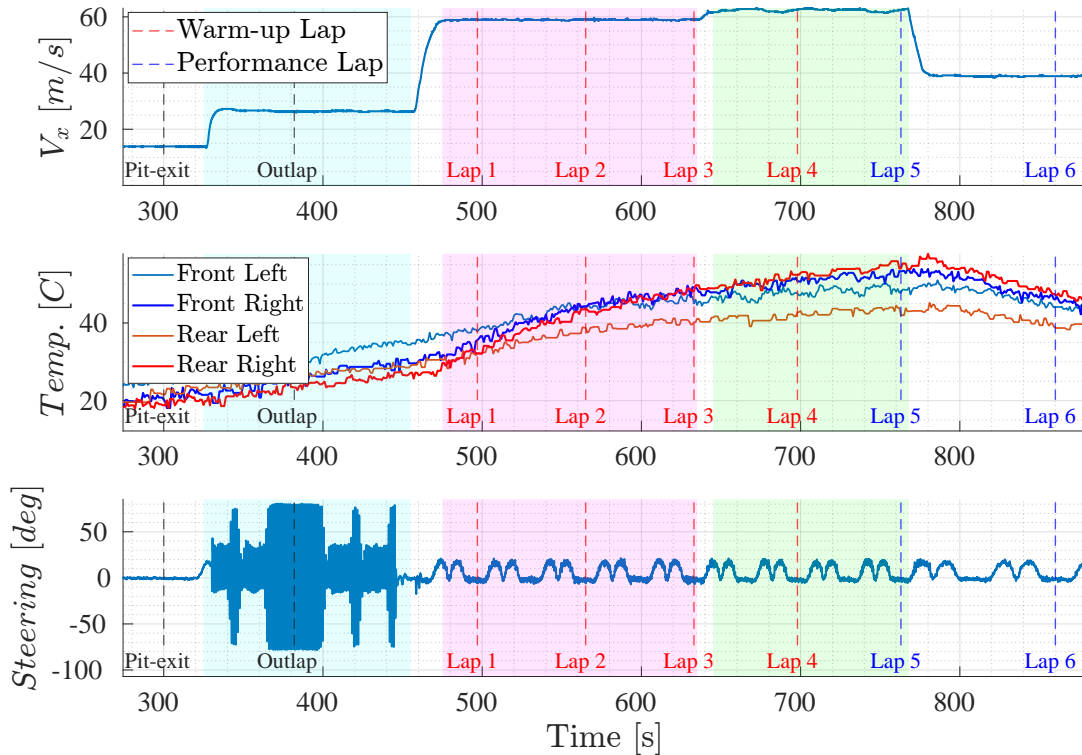
## 6.4. Overall Results in the Competition

Using the presented software stack, the team achieved the second and third position, respectively, at the two main events of the competition. This section gives an overview of the overall results, explaining the strategy we used and the failures that limited our final placements.

### 6.4.1. Indy Autonomous Challenge powered by Cisco

In [Table 6](#) we summarize the results of the teams qualified for the Seminal and Final of the race at IMS and that they have been able to complete at least one of the rounds.

The order of the runs for the Semifinal was set by draw. The PoliMove Autonomous Racing team was the first, followed by TUM, KAIST, Cavalier Autonomous Racing, and finally us. Running after these teams gave us the advantage to know their results and setting a target speed high enough to conclude first without taking too many risks. Given the first position gained in the Semifinal, we



**Figure 24.** Tires temperatures during the final at IMS. From the pit-exit and during the outlap, the warmup manoeuvre is performed by alternating left and right steering. At higher speed, we relied purely on friction to increase the temperature. Given the low track temperature, in the few laps available for the race it was not possible to heat the tires to their nominal temperature (around 80°C).

**Table 6.** Semifinal and Final results of teams qualified for the race at IMS. DNF stands for Did Not Finish, DNQ for Did Not Qualify. In bold, the best average speed of the top three teams is used to determine the final positions.

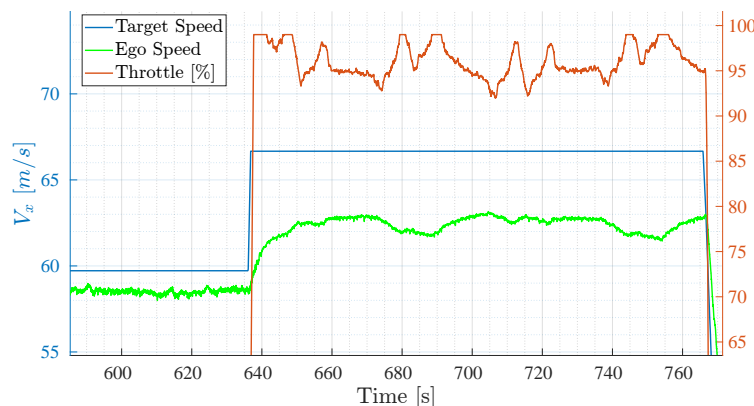
	Semifinal		Final	
	Average Speed [m/s]	Position	Average Speed [m/s]	Position
<b>TUM Autonomous Motorsport</b>	57.774	2	<b>60.772</b>	1
<b>TII EuroRacing</b>	<b>58.628</b>	1	51.83	2
<b>PoliMOVE</b>	<b>55.634</b>	3	DNF	3
<b>KAIST</b>	37.71	4		DNQ
<b>Cavalier Autonomous Racing</b>	53.592	DNF		DNQ

started our run for the Final after TUM and PoliMove. The format consisted of four warm-up laps followed by two performance laps, in which the average speed was the criteria used to declare the winner, and one cool-down lap before coming back to the pit.

Originally, our strategy, tested in simulation the night before the race, consisted in the following:

- Four Warm-up laps: two laps using the Warm-up manoeuvre presented in Section 3.4, one lap at 200 km/h, and the last one at 240 km/h.
- Two Performance laps at a speed equal to or higher than 240 km/h considering the other teams' result.
- One Cool-down lap at 150 km/h.





**Figure 25.** Fast laps in IMS Final. While the controller was requesting full throttle or a high value, the oscillations are due to a nonideal tuning of the turbocharger and a temporary malfunction of its mechanics.

During the time available before our turn for the Final round, we considered performing the weaving manoeuvre to heat first the front tires at 95 km/h just on the first warm-up lap, followed by two laps at 215 km/h and a final warm-up lap at 240 km/h. This choice had been applied and tested in our SiL simulator a few minutes before the run without the attention properly needed to guarantee its correctness. This resulted in what we internally called the “Million dollar bug,” considering the first place prize. Taking into account the TUM’s result of an average speed of 218.8 km/h and a hardware failure that led PoliMove to keep their average speed of 200.03 km/h gained during the Semifinal, we decided to keep the target speed for the performance laps at a speed of 240 km/h around the entire track. As can be seen in Figure 25, despite a throttle pedal command higher than 90%, the speed did not increase more than 227.4 km/h due to a momentary malfunction on the turbo-charger. After completing the first performance lap, the car reduced its speed to 150 km/h, as depicted in Figure 24, which was the target for the cool-down lap. During the last-minute change in the code, we erroneously set the cool-down lap speed at the end of the fifth lap instead of the sixth. Concluding the last lap at the same speed as the previous one would have guaranteed an average speed high enough to win, but the error led us to an average speed lower than that achieved in the Semifinal, positioning us second.

#### 6.4.2. Autonomous Challenge @ CES

At the second event, the order was set by draw as well and we ended up starting first at the time trial part in which the lap time was the criteria for determining the seeding in the bracket of the Passing Competition explained in Section 1. The run consisted of a maximum of 10 laps where the teams were not constrained to follow any kind of format. We decided to perform the outlap and the first lap doing the weaving manoeuvre at 100 km/h, followed by a series of laps at incrementally higher speed. During the seventh lap, we set a target speed of 77.7 m/s (280 km/h) but the vehicle reached a maximum of 75.5 m/s (272 km/h), despite a full throttle command as presented in (Raji et al., 2022), achieving a lap time of 33.99 s. Once the race car came back to the pit lane, the mechanics found that a cable related to the power train had been detached and caused the speed limitation. The results of the time trial are summarized in Table 7.

In the Passing Competition, we faced TUM for the Semifinal of the event. We were able to pass the rounds up to the defending speed level of 58 m/s overtaking at 63 m/s (226.8 km/h). Table 8 compares the overtaking speed with the other teams.

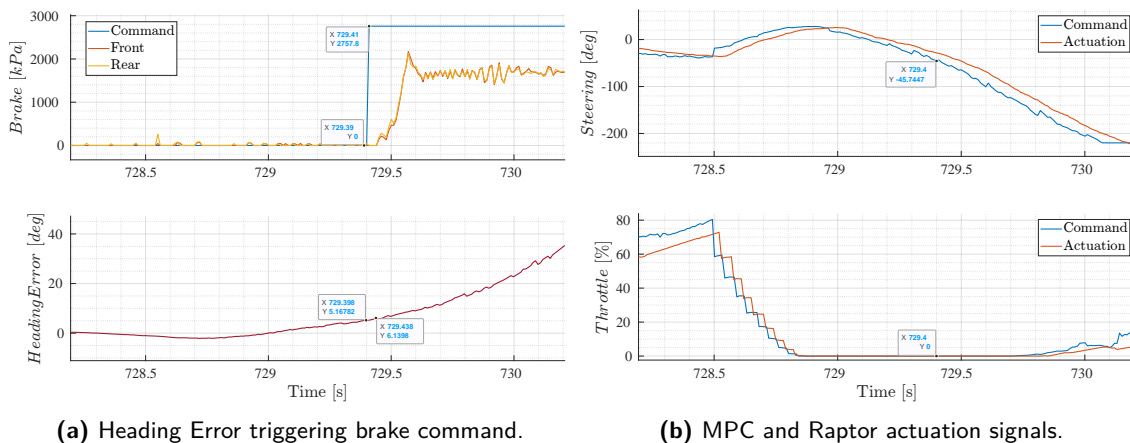
An edge case for the motion planning and control modules happened during the round at the defense speed of 60 m/s. A false detection of a standing obstacle by the Radar lets the planner generate a series of aggressive manoeuvres with a higher difference in terms of curvature and smoothness from each other. The MPC reacted with a much higher heading error than the one usually had during other tests on the track at higher speeds, triggering a hard brake by the safety

**Table 7.** Qualifying results of the race at LVMS; the four fastest teams in testing were admitted.

	Qualifying	
	Fastest Lap Time [s]	Position
<b>TII EuroRacing</b>	33.99	2
<b>TUM Autonomous Motorsport</b>	35.3	3
<b>PoliMOVE</b>	32.54	1
<b>Cavalier Autonomous Racing</b>		DNF

**Table 8.** Comparison of successful overtakes between all the qualified teams, during the Las Vegas Motor Speedway event. The defense speed is the speed of the defending car during the overtake.

Defence Speed (m/s)	TII EuroRacing	TUM	KAIST	PoliMOVE
36	✓	✓	✓	✓
45	✓	✓	✓	✓
51	✓	✓	✓	✓
56	✓	✓	×	✓
58	✓	✓		✓
60	×	✓		✓
62		✓		✓
65		✓		✓
67		×		✓

**(a)** Heading Error triggering brake command.**(b)** MPC and Raptor actuation signals.**Figure 26.** LVMS semifinal crash analysis. A false detection of a ghost car in front of the car triggered a steering correction and reduction in throttle request (b). Shortly after, the heading error with respect to the planned path went over the safety threshold, and the supervisor triggered an emergency braking (a). At over 200 kph, this resulted in the loss of control of the vehicle.

checks, as explained in Section 3. In Figure 26 it can be seen that as soon as the heading error passed the value of 6 deg, the brake pressure started to increase, bringing the car to a spin and causing it to hit the track boundaries despite the controller trying to steer to the right.

## 7. Lessons Learned

Robotics challenges are great instruments for pushing the integration of the latest research results in real-world applications, and they steer the efforts of research teams towards new results that can find

application outside the lab. In addition, the opportunity to compete against other teams is an important way to build and maintain an active research community and to exchange ideas at a fast rate.

We could easily argue that everyone who competes in a race wants to win. When this result is not achieved, it is very important to go back and make a rational analysis of the result. In our case, we summarize our analysis in the following points, hoping they might provide useful insights to the community.

- Control solutions based on complex models of the vehicle bring difficulties related to the different conditions of the real world with respect to the estimated model. This mismatch could lead the researchers to move to different approaches based on Robust Control, where, rather than modeling the nonlinear dynamics of the system, the focus is on considering the uncertainty around a simpler model, or on bounding the control commands on a set of potential limits of longitudinal and lateral accelerations. In our work, we demonstrated that a model identified from simulation has been used in an MPC in two different tracks with different weather conditions and different scenarios, with similar performances adapting only the weighting terms of the optimization problem and the cost function. Although we did not update the model after the validation on the track due mainly to the limited testing time, we believe that the experience gathered on modeling, validating, and tuning the controller will help to speed up the process and enable us to refine the parameters for each new track and road condition in a timely manner. Further effort should be put into automatically learning and refining the model parameters at runtime, taking into account the stability and the tracking performance of the controller.
- Given the technical challenges that GNSS systems pose in practice (Massa et al., 2020), we believe that investing effort in GNSS-denied localization systems will be key in future racing championships. We believe that to produce a step-change in autonomous technology, a race car should be equipped with all its capabilities in the edge-vehicle, relying on external infrastructure only for interactions with race control and live telemetry streaming. A first step in this direction will be reducing the importance of GNSS modules in ego vehicle localization, up to the point at which these could become unnecessary.
- The crash that happened in the head-to-head race was caused technically by setting the threshold on the safety check related to the heading error without considering possible extreme scenarios for the perception and planning modules, which are difficult to address in a simulation. A threshold of 6 deg on the heading error can be considered a strict value since in simulation the controller was able to correctly react in similar conditions when the check was not enabled. Furthermore, it would have been trivial to understand the potential effects of a hard brake command at high speeds with tires below their nominal temperature. For this reason, the safety check would be reevaluated for future runs. Potentially, the hard brake would be limited to low speeds and to the occasions when the control commands are not properly applied to the vehicle.
- Differently from usual research, in our case the race and the competitive component of the challenge led us to make important decisions in a stressful environment and in a short time. In addition to technical errors, wrong or high-risk organizational and operational decisions could ruin the final result as well. In our experience, this was proven in the case of the Million dollar bug, where we neglected the high probability of human error in applying and testing a last-minute change in the software.

## 8. Conclusions

In this work, we presented the complete software stack implemented by the TII EuroRacing team. Each module has been described including technical results as well as the overall achievements. Insights on the aspects we considered crucial for reaching speeds above 75 m/s (270 km/h), avoiding static obstacles, and racing in a head-to-head scenario have been given. We explained the failures that caused us not to achieve the first position in the final events, with important considerations

that could be beneficial to the other teams and researchers competing in challenges of different robots and fields.

With new Autonomous Racing challenges planned for the coming years, the team is working on `er.autopilot 2.0`. The updated software should improve the robustness of the sensors fusion on the detection module to cope with potential false detection and disturbances on the racetrack. Despite the implementation of a light warm-up manoeuvre to increase the temperature of the front tires, we will consider more aggressive manoeuvres to be performed in closed-loop control to speed up the warm-up procedure and reach the ideal temperature. A key point to achieve this goal is to improve the path-tracking performance of the MPC. Thanks to the data gathered at high speeds, it has been possible to confirm the correctness of an updated version of the multibody model developed in Dymola, which has been used to identify the single-track model parameters for the controller.

Another important feature will be the capability of running on GNSS-denied solutions. In addition to the LiDAR-based localization, we will consider the integration of a pure local control method. In (Lee et al., 2022), the authors presented a resilient navigation method based on following a distance from the wall of the track, using the LiDAR sensor and a variant of the RANSAC algorithm (Fischler and Bolles, 1981). Similarly, we will implement LiDAR-based and camera-based navigation for emergency situations.

Future applications will be on running the algorithms in free racing scenarios where more than two vehicles are allowed to drive without strict limitations on the possible trajectories to follow. For this purpose, it will be important to develop a Local Planner capable of generating aggressive, but feasible and safe, paths and velocity profiles. This would be possible by considering a dynamic model on the planning module and improving the integration between Planning and Control.

## Acknowledgments

We would like to thank all the students and researchers who contributed in small part to the development. In particular, we thank André Jesus, Abdurrahman İşbitirici, Mankaran Singh, Andrea Serafini, Francesco Moretti, Maciej Dziubiński, and Vallabh Ansingkar. Thanks to Claytex for the VeSyMA Motorsports libraries, and MegaRide<sup>19</sup> for the support in the tire model identification. Thanks to SpinItalia<sup>20</sup>, particularly Francesco La Gala, for the insight on the GNSS failure at LVMS.

We would also like to thank the IAC organization and all their partners for making this work possible.

Lastly, we would like to acknowledge the collaborative work done by all the teams during the first months at LOR and IMS. In particular, the work carried out by Alexander Wischnewski from TUM Autonomous Motorsport and Will Bryan from Autonomous Tiger Racing.

## ORCID


Ayoub Raji  <https://orcid.org/0000-0003-4188-8854>

Daniilo Caporale  <https://orcid.org/0000-0003-2665-3950>

Micaela Verucchi  <https://orcid.org/0000-0003-3898-8571>

Alessandro Toschi  <https://orcid.org/0009-0002-4497-0589>

Alexander Liniger  <https://orcid.org/0000-0002-7858-7900>

Marko Bertogna  <https://orcid.org/0000-0003-2115-4853>

## References

- Andresen, L., Brandemuehl, A., Honger, A., Kuan, B., Vödisch, N., Blum, H., Reijgwart, V., Bernreiter, L., Schaupp, L., Chung, J. J., et al. (2020). Accurate mapping and planning for autonomous racing. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4743–4749. IEEE.

<sup>19</sup> <https://www.megaride.eu/>

<sup>20</sup> <https://www.spinitalia.com/>

- Arslan, O., Berntorp, K., & Tsiotras, P. (2017). Sampling-based algorithms for optimal motion planning using closed-loop prediction. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4991–4996.
- Benson, A., Tefft, B., Svancara, A., & Horrey, W. (2018). Potential reductions in crashes, injuries, and deaths from large-scale deployment of advanced driver assistance systems. AAA Foundation for Traffic Safety, Washington, DC.
- Betz, J., Betz, T., Fent, F., Geisslinger, M., Heilmeier, A., Hermansdorfer, L., Herrmann, T., Huch, S., Karle, P., Lienkamp, M., Lohmann, B., Nobis, F., Ögretmen, L., Rowold, M., Sauerbeck, F., Stahl, T., Trauth, R., Werner, F., & Wischnewski, A. (2022a). Tum autonomous motorsport: An autonomous racing software for the indy autonomous challenge.
- Betz, J., Wischnewski, A., Heilmeier, A., Nobis, F., Stahl, T., Hermansdorfer, L., & Lienkamp, M. (2019a). A software architecture for an autonomous racecar. In *2019 IEEE 89th Vehicular Technology Conference (VTC2019-Spring)*, pages 1–6.
- Betz, J., Wischnewski, A., Heilmeier, A., Nobis, F., Stahl, T., Hermansdorfer, L., Lohmann, B., & Lienkamp, M. (2019b). *What Can We Learn from Autonomous Level-5 Motorsport?: chassis.tech plus*, pages 123–146.
- Betz, J., Zheng, H., Liniger, A., Rosolia, U., Karle, P., Behl, M., Krovi, V., & Mangharam, R. (2022b). Autonomous vehicles on the edge: A survey on autonomous vehicle racing. *IEEE Open Journal of Intelligent Transportation Systems*, 3, 458–488.
- Bochkovskiy, A., Wang, C.-Y., & Liao, H.-Y. M. (2020). Yolov4: Optimal speed and accuracy of object detection. *arXiv:2004.10934*.
- Bulsara, A., Raman, A., Kamarajugadda, S., Schmid, M., & Krovi, V. (2020). Obstacle avoidance using model predictive control: An implementation and validation study using scaled vehicles. *SAE Technical Paper Series*, 1.
- Buyval, A., Gabdulin, A., Mustafin, R., & Shimchik, I. (2017). Deriving overtaking strategy from nonlinear model predictive control for a race car. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2623–2628.
- Caporale, D., Settini, A., Massa, F., Amerotti, F., Corti, A., Fagiolini, A., Guiggiani, M., Bicchi, A., & Pallottino, L. (2019). Towards the design of robotic drivers for full-scale self-driving racing cars. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 5643–5649.
- Chen, T., Li, Z., He, Y., Xu, Z., Yan, Z., & Li, H. (2019). From perception to control: an autonomous driving system for a formula student driverless car. *arXiv:1909.00119*.
- Costi, G. (2022). Realtime lidar point cloud clustering and segmentation for automotive. MoreThesis Unimore.
- Coulter, R. C. (1992). Implementation of the pure pursuit path tracking algorithm. Technical Report CMU-RI-TR-92-01, Carnegie Mellon University, Pittsburgh, PA.
- Culley, J., Garlick, S., Esteller G., Georgiev, P., Fursa, I., Vander Sluis, I., Ball, P., & Bradley, A. (2020). System design for a driverless autonomous racing vehicle. In *2020 12th International Symposium on Communication Systems, Networks and Digital Signal Processing (CSNDSP)*, pages 1–6.
- De Rita, N., Aimar, A., & Delbruck, T. (2019). Cnn-based object detection on low precision hardware: Racing car case study. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 647–652.
- Dellaert, F. (2012). Factor graphs and gtsam: A hands-on introduction. Technical report, Georgia Institute of Technology.
- Dempsey, M. (2006). Dymola for multi-engineering modelling and simulation. In *2006 IEEE Vehicle Power and Propulsion Conference*, pages 1–6.
- Ecola, L., Popper, S., Silbergliitt, R., & Fraade-Blanar, L. (2018). The road to zero: A vision for achieving zero roadway deaths by 2050. *Rand Health Quarterly*, 8, 11.
- Feraco, S., Luciani, S., Bonfitto, A., Amati, N., & Tonoli, A. (2020). A local trajectory planning and control method for autonomous vehicles based on the rrt algorithm. In *2020 AEIT International Conference of Electrical and Electronic Technologies for Automotive (AEIT AUTOMOTIVE)*, pages 1–6.
- Fischler, M. A. & Bolles, R. C. (1981). Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6), 381–395.
- Frison, G. & Diehl, M. (2020). HPIPM: a high-performance quadratic programming framework for model predictive control. *IFAC-PapersOnLine*, 53(2), 6563–6569.
- Funk, N., Alatur, N., Deuber, R., Gonon, F., Messikommer, N., Nubert, J., Patriarca, M., Schaefer, S., Scotoni, D., Bünger, N., Dube, R., Khanna, R., Pfeiffer, M., Wilhelm, E., & Siegwart, R. (2017). Autonomous electric race car design. *arXiv:1711.00548*.

- Funke, J., Theodosis, P., Hindiyeh, R., Stanek, G., Kritatakirana, K., Gerdes, C., Langer, D., Hernandez, M., Müller-Bessler, B., & Huhnke, B. (2012). Up to the limits: Autonomous audi tts. In *2012 IEEE Intelligent Vehicles Symposium*, pages 541–547.
- Heilmeier, A., Wischnewski, A., Hermansdorfer, L., Betz, J., Lienkamp, M., & Lohmann, B. (2020). Minimum curvature trajectory planning and control for an autonomous race car. *Vehicle System Dynamics*, 58(10), 1497–1527.
- Himmelsbach, M., Hundelshausen, F. V., & Wuensche, H.-J. (2010). Fast segmentation of 3d point clouds for ground vehicles. In *2010 IEEE Intelligent Vehicles Symposium*, pages 560–565. IEEE.
- Kabzan, J., Valls, M. d. l. I., Reijgwart, V., Hendriks, H. F. C., Ehmke, C., Prajapat, M., Bühler, A., Gosala, N., Gupta, M., Sivanesan, R., Dhall, A., Chisari, E., Karnchanachari, N., Brits, S., Dangel, M., Sa, I., Dubé, R., Gawel, A., Pfeiffer, M., Liniger, A., Lygeros, J., & Siegart, R. (2019). Amz driverless: The full autonomous racing system. *Journal of Field Robotics*, 37, 1267–1294.
- Kapania, N. R. & Gerdes, J. C. (2015). Design of a feedback-feedforward steering controller for accurate path tracking and stability at the limits of handling. *Vehicle System Dynamics*, 53(12), 1687–1704.
- Kuhn, H. W. (1955). The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2), 83–97.
- Laurense, V. A., Goh, J. Y., & Gerdes, J. C. (2017). Path-tracking for autonomous vehicles at the limit of friction. In *2017 American Control Conference (ACC)*, pages 5586–5591.
- Lee, D., Jung, C., Finazzi, A., Seong, H., & Shim, D. H. (2022). Resilient navigation and path planning system for high-speed autonomous race car. *arXiv:2207.12232*.
- Liniger, A., Domahidi, A., & Morari, M. (2015). Optimization-based autonomous racing of 1:43 scale rc cars. *Optimal Control Applications and Methods*, 36(5), 628–647.
- Massa, F., Bonamini, L., Settimi, A., Pallottino, L., & Caporale, D. (2020). Lidar-based gnss denied localization for autonomous racing cars. *Sensors*, 20(14).
- Massaro, M. & Limebeer, D. (2021). Minimum-lap-time optimisation and simulation. *Vehicle System Dynamics*, 59, 1–45.
- Nekkah, S., Janus, J., Boxheimer, M., Ohnemus, L., Hirsch, S., Schmidt, B., Liu, Y., Borbély, D., Keck, F., Bachmann, K., & Bleszynski, L. (2020). The autonomous racing software stack of the kit19d. *arXiv:2010.02828*.
- Novi, T., Liniger, A., Capitani, R., & Annicchiarico, C. (2020). Real-time control for at-limit handling driving on a predefined path. *Vehicle System Dynamics*, 58(7), 1007–1036.
- Oth, L., Furgale, P., Kneip, L., & Siegart, R. (2013). Rolling shutter camera calibration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1360–1367.
- Pacejka, H. B. & Bakker, E. (1991). The magic formula tyre model. *Vehicle System Dynamics*, 21, 1–18.
- Puchtlar, P. & Peinl, R. (2020). Evaluation of deep learning accelerators for object detection at the edge. In *KI 2020: Advances in Artificial Intelligence: 43rd German Conference on AI, Bamberg, Germany, September 21–25, 2020, Proceedings*, pages 320–326, Berlin, Heidelberg. Springer-Verlag.
- Raji, A., Liniger, A., Giove, A., Toschi, A., Musiu, N., Morra, D., Verucchi, M., Caporale, D., & Bertogna, M. (2022). Motion planning and control for multi vehicle autonomous racing at high speeds. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)*, pages 2775–2782.
- Schratter, M., Zubača, J., Lassnig, K., Renzler, T., Kirchengast, M., Loigge, S., Stolz, M., & Watzenig, D. (2021). Lidar-based mapping and localization for autonomous racing. Opportunities and Challenges with Autonomous Racing : 2021 ICRA Workshop; Conference date: 31-05-2021.
- Stahl, T., Wischnewski, A., Betz, J., & Lienkamp, M. (2019). Multilayer graph-based trajectory planning for race vehicles in dynamic scenarios. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 3149–3154.
- Strobel, K., Zhu, S., Chang, R., & Koppula, S. (2020). Accurate, low-latency visual perception for autonomous racing: Challenges, mechanisms, and practical solutions. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1969–1975.
- Tian, H., Ni, J., Li, Z., & Hu, J. (2020). Autonomous formula racecar: Overall system design and experimental validation. *arXiv:2009.00385*.
- Urmson, C., Anhalt, J., Bagnell, D., Baker, C., Bittner, R., Clark, M., Dolan, J., Duggins, D., Galatali, T., Geyer, C., Gittleman, M., Harbaugh, S., Hebert, M., Howard, T., Kolski, S., Kelly, A., Likhachev, M., Mcnaughton, M., Miller, N., & Ferguson, D. (2008). Autonomous driving in urban environments: Boss and the urban challenge. *Journal of Field Robotics*, 25, 425–466.

- Verucchi, M., Brilli, G., Sapienza, D., Verasani, M., Arena, M., Gatti, F., Capotondi, A., Cavicchioli, R., Bertogna, M., & Solieri, M. (2020). A systematic assessment of embedded neural networks for object detection. In *2020 25th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*, volume 1, pages 937–944. IEEE.
- Vázquez, J. L., Brühlmeier, M., Liniger, A., Rupenyan, A., & Lygeros, J. (2020). Optimization-based hierarchical motion planning for autonomous racing. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2397–2403.
- Vödisch, N., Dodel, D., & Schötz, M. (2022). FSOCO: The formula student objects in context dataset. *SAE International Journal of Connected and Automated Vehicles*, 5(1).
- Werling, M., Ziegler, J., Kammel, S., & Thrun, S. (2010). Optimal trajectory generation for dynamic street scenarios in a frenét frame. In *2010 IEEE International Conference on Robotics and Automation*, pages 987–993.
- Wischniewski, A., Euler, M., Gümüs, S., & Lohmann, B. (2021). Tube model predictive control for an autonomous race car. *Vehicle System Dynamics*, pages 1–23.
- Wischniewski, A., Stahl, T., Betz, J., & Lohmann, B. (2019). Vehicle dynamics state estimation and localization for high performance race cars\*\*research was supported by the basic research fund of the institute of automotive technology of the technical university of munich. *IFAC-PapersOnLine*, 52(8):154–161. 10th IFAC Symposium on Intelligent Autonomous Vehicles IAV 2019.
- Yellman, M. A. (2022). Motor vehicle crash deaths—United States and 28 other high-income countries, 2015 and 2019. *MMWR. Morbidity and Mortality Weekly Report*, 71.
- Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., & Darrell, T. (2020). Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2636–2645.

**How to cite this article:** Raji, A., Caporale, D., Gatti, F., Giove, A., Verucchi, M., Malatesta, D., Musiu, N., Toschi, A., Popitanu, S. R., Bagni, F., Bosi, M., Liniger, A., Bertogna, M., Morra, D., Amerotti, F., Bartoli, L., Martello, F., & Porta, R. (2024). er.autopilot 1.0: The full autonomous stack for oval racing at high speeds. *Field Robotics*, 4, 99–137.

**Publisher's Note:** Field Robotics does not accept any legal responsibility for errors, omissions or claims and does not provide any warranty, express or implied, with respect to information published in this article.