CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

Final Report 115
April 2026

High-Speed Driverless Racing: Lessons for Autonomous Mobility

Justin Gan
Samuel Labi

PURDUE UNIVERSITY

Report No. 115
Project Start Date: June 2023
Project End Date: December 2025
April 2026

High-speed driverless racing: Lessons for autonomous mobility

by

Justin Gan
Undergraduate Researcher

Samuel Labi
Professor

Purdue University

ACKNOWLEDGEMENTS AND DISCLAIMER

Funding for this research was provided by the Center for Connected and Automated Transportation under Grant No. 69A3552348305 of the U.S. Department of Transportation, Office of the Assistant Secretary for Research and Technology (OST-R), University Transportation Centers Program. The authors acknowledge the support of organizations and individuals that provided valuable material, institutional, or external resources for this project, including 9-12 (Chad Pittman, Kent Anderson, Doug Rapp), Professor Dan DeLaurentis (Executive Vice President of Research, Purdue University), and Luna Lu (Vice President of Industry Partnerships, Purdue University). We also acknowledge the overall support, direct and indirect, of Purdue's erstwhile team (Black-and-Gold Racing) founders and members, including Eric J. Dietz, Joe Pekny, Mike Saxon, Aly El Gamal, Shyam Kannan, Rohan Manna, Haoguang Yang, Manuel Valencia Mar, Andres Moreno, Alec Pannunzio, Sashank Modali, Richard Ajagu, Jiqian Dong, and Zainab Saka; Deepak Benny of the Build America Center; Dean Arvind Raman, Eckhard Groll, Shaoshuai Mou, Shreyas Sundaram, and Purdue AI Racing professors and students; Paul Mitchell and his IAC team in Indianapolis; Jack Silberman (Triton Racing, UC San Diego), and Danny White (former Purdue Motorsports Director). The companies, institutions, professors, and individuals that provided valuable in-kind cost share or other material or external support for this project, are acknowledged including the Indiana DOT Research Division (Barry Partridge and Samy Noureldin), National University of Singapore (Ghim Ping Ong), ETH Zurich (Bryan Adey), Nanyang Technical University (Feng Zhu), IIT Tirupathi (Krishna Prapoorna), and TU Delft (Goncalo H. de Almeida Correia) and other institutions are also acknowledged.

Suggested APA Format Citation:
Gan, Y.J., and Labi, S. (2026). High-speed driverless racing: Lessons for autonomous mobility, CCAT Report #115, The Center for Connected & Automated Transportation, Purdue University, W. Lafayette, IN.

Note: LLMs served as one of tools used for cross-checking some of the information presented in this report.

Contacts

For more information

Samuel Labi, Ph. D.	**CCAT**
Purdue University	University of Michigan Transportation Research Institute
550 Stadium Mall Dr.	2901 Baxter Road
HAMP G175A	Ann Arbor, MI 48152
Phone: (765) 494-5926	uumtri-ccat@umich.edu
Email: labi@purdue.edu	(734) 763-2498

Technical Report Documentation Page

<table>
<tr><td>1. Report No. 115</td><td>2. Government Accession No.</td><td>3. Recipient's Catalog No.</td></tr>
<tr><td colspan="2" rowspan="2">4. Title and Subtitle
High-speed driverless racing: Lessons for autonomous mobility</td><td>5. Report Date: April 2026</td></tr>
<tr><td>6. Performing Organization Code
N/A</td></tr>
<tr><td colspan="2">7. Author(s)
Justin Y. Gan, Samuel Labi</td><td>8. Performing Organization Report No.
N/A</td></tr>
<tr><td colspan="2" rowspan="2">9. Performing Organization Name and Address
Center for Connected and Automated Transportation, Purdue University, 550 Stadium Mall Drive, W. Lafayette, IN 47907; and University of Michigan, Ann Arbor, 2901 Baxter Rd, Ann Arbor, MI 48109</td><td>10. Work Unit No.</td></tr>
<tr><td>11. Contract or Grant No.
Contract No. 69A3552348305</td></tr>
<tr><td colspan="2" rowspan="2">12. Sponsoring Agency Name and Address
U.S. Dept of Transportation, Office of the Assistant Secretary for Research & Tech., 1200 New Jersey Ave. SE, Washington, DC 20590</td><td>13. Type of Report and Period Covered
Final report, April 2026</td></tr>
<tr><td>14. Sponsoring Agency Code: OST-R</td></tr>
<tr><td colspan="3">15. Supplementary Notes
Conducted under the U.S. DOT Office of the Assistant Secretary for Research and Technology's (OST-R) University Transportation Centers (UTC) program.</td></tr>
<tr><td colspan="3">16. Abstract
Driverless racecar competitions are often perceived as recreational. However, their academic, industry, and research benefits are potentially tremendous, as several significant lessons can be learned from these competitions to advance autonomous mobility. This study was motivated (and its conduction facilitated) by the active involvement of the authors in several autonomous racing initiatives in the United States. The study first synthesized and analyzed information from existing literature related to the key modules of autonomous vehicle (AV) operations as part of AV racing competitions and identified some lessons that could be learned in the context of each module – perception and sensing, data fusion, path planning and decision making, vehicle dynamics and control, and hardware and software safety. After synthesizing such information, the study establishes that high-speed AV tests and competitions where multiple teams push the limits of sensing, decision-making, and control serve as testbeds for creating, developing, and refining AV technologies. This is because the AV racing events have used controlled environment platforms to demonstrate what autonomous systems can achieve under extreme conditions, and allowing testing of edge-case scenarios and performance extremes that would be too risky or impractical on public roads. Further, by demonstrating what autonomous systems can achieve under extreme conditions and edge cases, the AV racing events have shown that they foster innovation in a safe and controlled environment. The identified lessons from the racing competitions can serve as a knowledge base to facilitate safe and efficient AV operations on high-speed road transportation corridors such as freeways. The discussion includes caveats regarding the dichotomies between AV operations guideway (racetrack vs. roadway) and AV vehicle designs (race car designs vs. standard automobile and truck design), and how these differences temper the translation of lessons learned from AV racing to real-world (roadway) AV operations. In sum, the study outcomes support the notion that high-speed AV events can continue to serve as proving grounds for testing edge cases of autonomous operations to facilitate safety and movement efficiency in autonomous driving in the real world. In the prospective era of AV operations on public roads, the lessons learned can help enhance AV design and sensing/communication capabilities for roadway operations specifically. The lessons can also promote transportation mobility, safety, reliability, and economic productivity associated with autonomous mobility at high-speed road corridors, particularly freeways.</td></tr>
<tr><td>17. Key Words
Autonomous vehicles, Self-driving competition, High speed, Racing, Edge cases</td><td colspan="2">18. Distribution Statement
No restrictions.</td></tr>
<tr><td>19. Security Classif. (of this report)
Unclassified</td><td>20. Security Classif. (of this page)
Unclassified</td><td>21. No. of Pages
134   22. Price</td></tr>
</table>

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ACRONYMS

A2RL	Abu Dhabi Racing League
AAA	American Automobile Association
AASHTO	American Association of State Highway and Transportation Officials
ACC	Adaptive Cruise Control
ADAS	Advanced Driver Assistance Systems
AEB	Automatic Emergency Braking
ATIS	Advanced Traveler Information Systems
AV	Autonomous Vehicle
BEV	Bird's Eye View
DDT	Dynamic Driving Task
EKF	Extended Kalman Filter
EOAM	Emergency Obstacle Avoidance Maneuvers
ESC	Electronic Stability Control
FoV	Field of View
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HDV	Human-Driven Vehicle
HSS	Hardware and Software Safety
IAC	Indy Autonomous Challenge
IMU	Inertial Measurement Unit
IRL	Infrastructure Readiness Level
IOO	Infrastructure Owner and Operator
LiDAR	Light Detection and Ranging
LKA	Lane Keeping Assist
MAP	Mean Average Precision
MPC	Model Predictive Control
ODD	Operational Design Domain
OEM	Original Equipment Manufacturer
NHTSA	National Highway Traffic Safety Administration
ROS	Robotic Operating System
RTK	Real-time Kinematic (positioning)
SAE	Society of Automotive Engineers
SDV	Software Defined Vehicle
SLAM	Simultaneous Localization and Mapping
SPM	Stakeholder Participation Model
SSD	Stopping Sight Distance
TRFC	Tire-Road Friction Coefficient
TUM	Technical University of Munich
URWB	Ultra-Reliable Wireless Backhaul
V2V	Vehicle-to-Vehicle
V2X	Vehicle-to-Everything

LIST OF KEY TERMS

Adaptation	The real-time adjustment of algorithms, sensors, and control systems of an AV to enable navigation in unpredictable environments, improve safety, and enhance passenger trust.
Attitude	An AV’s spatial orientation – roll, pitch, and yaw angles – which are critical for stability, control, and navigation at high speeds. Unlike human driving where “feel” is used, autonomous systems depend on a specific class of sensors (including GNSS, gyroscopes, and accelerometers) that measure these angles and detect when the AV is slipping into errant orientation such as oversteering or understeering.
Automated Driving System (ADS)	A collection of hardware and software that perform the entire dynamic driving task.
Drifting	Navigating a vehicle such that it maintains control while moving sideways with high slip angles.
Dynamic Driving Task (DDT)	A set of real-time functions (operational and tactical) that operate a vehicle continuously.
Fully Autonomous Phase (FAP)	A prospective future period where 100% of vehicles in the traffic stream on a specific corridor, or in a wider network or region have full or near full autonomy (Levels 4, 5, or both).
Ghosting	A state where the AV loses communication with outside world due to natural conditions or human-made actions (negligent or malicious).
Infrastructure Owner and Operator (IOO)	In the context of highway transportation, IOOs are organizations that own and/or operate the physical (and in some cases, cyber) infrastructure that support travel. IOOs include public and private organizations such as State and local DOTs, transit agencies, and operators of toll roads and bridges.
Infrastructure Readiness Level (IRL)	A quantitative statement of highway infrastructure system maturity in terms of its capability to support a certain specified level of AV operations. This may include the physical infrastructure, safety, governmental and institutional processes, data management (acquisition, sharing and privacy), systems interoperability and communications.
Operational Design Domain (ODD)	A set of conditions under which an autonomous system or AV-related component is designed to function with minimal hiccups. ODD may differ across the different levels of automation.
Original Equipment Manufacturer (OEM)	The companies that manufacture, sell, lease, or maintain AVs or AV-supporting components (including sensors). OEMs define vehicle ODDs and are responsible for safe operations of the vehicle within the ODD, irrespective of what the IOO has provided.
Pose	The position and orientation of the AV during motion.

Production vehicle	A mass-produced identical automobile manufactured for public sale and legal use on public roads, in contrast to racecars that operate on racetracks and road courses. May be human driven or autonomous.
Production AV	A production vehicle that is autonomous and operates on roads, streets and freeways, in contrast to an autonomous racecar that operates on racetracks and road courses.
Safety	In the context of autonomous driving, a "safe" autonomous system is one whose hardware and software are designed and operated in ways that reduce uncertainty, address highly dynamic environments, or have minimum vulnerability to human error in software design and hardware operations.
Software Defined Vehicle (SDV)	A vehicle whose software (and not just hardware) controls and enhances its features, allowing for continuous updates, new functions, and personalization remotely via over-the-air (OTA) updates, similar to a smartphone.
Sensor Type	Type or class of sensor based on the underlying technology, for example, camera, LiDAR, GNSS, radar, and IMU.
Stakeholder (AV)	A group of individuals, organizations, interest groups, or entities that are expected to (a) gain some utility (mobility, safety, etc.) or suffer some adversity (inequity, etc.) due to AV operations, (b) be involved in the manufacture, sale, maintenance, or R&D of AVs and/or their supporting components, (c) be responsible for the planning, design, and operations & maintenance of infrastructure intended to support AVs, or (d) be involved in AV-related policy formulation or regulation.
Transition Period	The period between negligible AV market penetration (MP) and 100% MP; characterized by mixed traffic (AV and HDV) operations on the road space and a growing AV proportion.

(Sources: SAE, 2016; SAE, 2018; SAE, 2021; Gopalkrishnan et al. 2015, Saka and Labi, 2023; Labi et al., 2023)

CHAPTER 1 INTRODUCTION

"Autonomous racing … accelerates the evolution of autonomous technology in the most demanding conditions possible." - H.E. Faisal Al Bannai, Advisor to UAE President for Strategic Research & Advanced Technology Affairs.

1.1 Background

The first prearranged race between self-powered road vehicles over a prescribed route is said to have occurred in 1867 at Ashton-under-Lyne in Greater Manchester, England, with the carriages achieving an average speed of 8 mph. This was followed by several early motor competitions including the 1894 Paris–Rouen Horseless Carriages Contest in France, the 1978 Wisconsin Great Race from Green Bay to Madison, and the 1985 Chicago-Evanston race (Backus, 2004; Gifford, 2006; Shahbakhshi, 2022). The 1894 Paris event was the first to involve vehicles that would later become known as automobiles. Later versions evolved into racing institutions with periodic races, including the Aspendale motor race (Australia), the Grand Prix at Le Mans circuit (France), the International Federation of the Automobile's Formula One (F1), Indianapolis 500 (Indy 500, USA), and NASCAR racing in the United of States (Atkins, 2022). Over time, these events reflected progressive improvements in vehicle technology and often served as catalysts for technological innovation in the automotive sector.

These events, dubbed "motorsports," are more than a sport. They have consistently served as proving grounds for automobile innovation since the invention of the automobile. The urge to win races has helped spawn race vehicle innovations which ultimately translated into improvements in production vehicles used on public roads (Jiang, 2021; Perrone, 2025). Such transfer of technology is also evident in the current era where society is on the verge of autonomous mobility, and as vehicles transform from primarily mechanical systems to highly integrated cyber-physical systems or robots on wheels. This technology, from 2005 to the current day (2026), has progressed from isolated academic experiments to collegiate competitions and then to international sporting leagues.

(a) An official autonomous racecar of IAC, 2021

(b) Waymo-operated AV in San Francisco, 2023

Image sources: (a) ultimate-robot-archive.fandom.com/wiki/Dallara_AV-21
(b) wikipedia.org/wiki/Waymo#/media/File:Waymo_Jaguar_I-Pace_in_San_Francisco_2023_dllu.jpg

Figure 1.1 Autonomous racecar platform and that of a production AV

Autonomous vehicle (AV) racing events feature fully driverless vehicles of various designs and sizes: the mini-sized F-1/10th vehicle, the moderate-size go-kart, and large vehicles such as the Dallara IL-15 based IAC race car and Dallara EAV24 A2RL racecar. These vehicles compete at high speeds on racetracks, road courses, and in some cases, guideways defined by improvised guardrails or tubing. These events are not merely spectator sports; they serve as testbeds for validating various components of self-driving technology in ways that are relatively safe and cost-effective. Given that the hardware and software systems of autonomous racecars are similar to those used in commercially available AVs (Figure 1.1), the experience of the former could provide valuable lessons for the latter. Key AV racing series are described below.

The current study represents the latest in a series of CCAT-sponsored research studies on production AVs, with recent studies addressing AI, perception, control algorithms, and infrastructure preparations. For example, regarding AI applications in production AVs, Dong et al. (2022) developed AI based systems to support safe and efficient operations on normal roadways; Du et al. (2022) developed an AI-based fog-cloud control framework; and Dong et al. (2023a) used explainable AI to enhance the perception capabilities of production AVs. Also on perception and sensing, Feng and Tarko (2023) developed a cooperative perception system or AVs, and Saka and Labi (2023) investigated alternative locations for sensor placement on an AV. Regarding control of production AVs, Zhou et al. (2022)'s cooperative control mechanism and Du et al. (2022)'s framework facilitated AV platooning and addressed large-scale network multi-level control, respectively, of production AVs. Regarding infrastructure preparations, Ramakrishnan et al. (2021) investigated the impacts of autonomous truck platooning on road infrastructure; Seilabi et al. (2022) investigated road lane allocation to normal vehicles and production AVs; and Labi et al. (2023) identified areas of highway infrastructure design and management to facilitate the deployment of production AVs. Manion and Durango-Cohen developed a framework to optimize investments in the development of transportation infrastructure considering AVs. Regarding testing and evaluation, Liu and Feng (2020) and Liu and Feng (2022) developed evaluation methods and a scenario generation system, respectively, for AV testing.

1.2 The Evolution of Autonomous Racing

The key autonomous racing competitions are discussed in this section. Figure 1.2 presents examples of the venues and vehicles for these events.

The DARPA Grand Challenges (2004-2007)
The DARPA Grand Challenges represent prize competitions for AVs funded by the US Department of Defense's Advanced Research Projects Agency (DARPA, 2014). Authorized by the U.S. Congress, the series promotes transformative, high-reward innovations that translate fundamental discoveries into practical military and civilian applications. The DARPA challenges established baselines for autonomous driving in diverse environments, including off-road and urban settings. The 2005 race involved navigating tunnels, sharp turns, and rugged desert terrain. In winning the competition, the Stanford University-Volkswagen team, led by Sebastian Thrun, laid the groundwork for sensor fusion stacks used in modern ADS, demonstrating the viability of modern sensor-fusion architectures (AU Motorsport, 2005). The 2007 Urban Challenge (at George Air Force Base, California) featured a 60-mile urban road course with rules governing traffic, roadway navigations with other vehicles (merging, etc.), and obstacle avoidance.

The Indy Autonomous Challenge (IAC)

IAC was primarily founded and organized by Energy Systems Network (ESN), an Indianapolis-based non-profit focused on accelerating energy and transportation technologies. The Indiana Economic Development Corporation and Lilly Endowment Inc. were also key partners in the founding and early support. The initiative, a branded project of the Central Indiana Corporate Partnership, was established to foster public-private partnerships and challenge university teams to create software for fully autonomous racecars (IAC, 2025). The IAC initially used the Dallara AV-21, a modified Indy Lights chassis, enabling high-speed autonomous racing using an internal combustion powertrain. The IAC platform subjects sensors and compute units to the vibration and thermal profiles of a high-performance internal combustion engine, offering valid durability data for standard automotive environments. By 2024, the vehicles were modified to use Dallara AV-24 specifications, which featured a re-engineered software stack, sensor suite, and computing hardware. IAC events often include simulator challenges and real-life races at test tracks and road courses. To date, the IAC events remain among the most prolific and impactful demonstrations of high-speed autonomous driving.

The Formula Student Driverless

This is a European competition involving student-built prototypes that race in a limited environment, funded by the Center for Connected & Autonomous Vehicles (CCAV) and Innovate UK. The initiative trains the next generation of workforce in autonomous systems to acquire skills and experience sought by the autonomous mobility industry and OEMs. The competition's budgetary constraints encourage student engineering teams to develop cost-effective, innovative solutions in sensor fusion and lightweight computing, both of which are relevant to the context of cost-conscious mass production of vehicles in the real world. The overall winner is selected based on vehicle design, performance, sales presentation, costs, and business plan (BMW, 2025). This series has become a recruitment pipeline for major European OEMs (Scania, 2025).

The Abu Dhabi Autonomous Racing League (A2RL)

This is an autonomous racing championship organized by the United Arab Emirates' Advanced Technology Research Council, occurring at Abu Dhabi's Yas Marina circuit. University teams develop AI software to race standardized, driverless racecars at high speeds. A2RL's event uses a chassis designated as EAV-24 (derived from the Japanese Super Formula's SF-23 chassis (Autosport, 2025). The A2RL website describes the event as one that pushes the boundaries of robotics, AI, and autonomous systems, ultimately intended for applications beyond motorsport, such as passenger transportation and logistics.

The Autonomous Karting Series/Purdue University (AKS/Purdue)

This is an international collegiate competition that started at Purdue University in 2018, born out of the electric vehicle autonomous grand prix. The participating teams design, build, and race their autonomous go-karts. The series, hosted by a co-author of this report and held at Purdue University's Grand Prix Racetrack in West Lafayette, Indiana, minimizes the use of standardized components, thereby encouraging teams to develop their unique solutions and pursue diverse development paths, and ensuring that the competition remains a cutting-edge showcase for new ideas and technologies in building and running the go karts. Most of the teams are multidisciplinary, with students from robotics, electrical and computer engineering, civil engineering, mechanical engineering, and technology programs of the participating institutions.

(a) The 2005 DARPA race concluded through Beer Bottle Pass, a winding mountain pass

(b) IAC teams at Indianapolis Motor Speedway, October 2021

(c) Roborace vehicle DevBot 2.0 (Technical University of Munich) competing in Season Alpha

(d) Official racecar of the Abu Dhabi Racing League (A2RL)

(e) Two competing vehicles interact at an intersection during the 2007 DARPA Urban Challenge Finals

(f) Roborace's autonomous racecar, August 2016

Image sources: (a) wikipedia.org/wiki/DARPA_Grand_Challenge_(2005)#/media/File:BeerBottlePass.JPG
(b) wikipedia.org/wiki/Indy_Autonomous_Challenge#/media/File:IAC+Group+Team+PD5_9137.jpg
(c) wikipedia.org/wiki/Roborace#/media/File:2019-04-Roborace_Monteblanco_Day02_Monday-PM-29.jpg
(d) ultimate-robot-archive.fandom.com/wiki/Abu_Dhabi_Autonomous_Racing_League,
(e) en.wikipedia.org/wiki/DARPA_Grand_Challenge_(2007)
(f) www.topgear.com/car-news/electric/roboraces-new-autonomous-racecar

Figure 1.2 Examples of autonomous racing venues and vehicles

RoboRacer
In 2015, Denis Sverdlov founded Roborace, a competition featuring autonomous electric-powered race vehicles, seeking to advance autonomous driving technology, improve road safety, and create immersive entertainment through AI-enabled competition. Originally launched in 2016 at the University of Pennsylvania, Roborace has expanded to several other universities worldwide, connecting a global community of engineers, researchers, and autonomous systems enthusiasts. The events intended to "foster interest, excitement, and critical thinking" in autonomous system development and operations (RoboRacer, 2025). The teams use a standard frame for the RoboRacer Autonomous Vehicle System. The vehicle is a small-scale autonomous vehicle approximately one-tenth the size of a Formula One vehicle, commonly referred to as the F1Tenth platform, which serves as a useful open-source platform that is useful for education and research in autonomous systems, including autonomous racing, reinforcement learning, robotics, and communication systems. The term "F1Tenth" is used solely in reference to an open-source academic project and does not indicate any connection to Formula 1 Motorsport.

Discussion
Besides the events discussed above, there have been other similar events on a smaller scale through regional, national, and international challenges, or within universities and high schools all over the world (Karaman et al., 2017; Mohanasundaram et al., 2023). These autonomous racing events feature racecars that were built specifically for the events. The racecars are equipped with state-of-the-art sensors (cameras, LiDAR, radars, IMUs, and GPS) and efficient onboard computers that run various software codes to determine the vehicle location, perceive the surrounding environment, plan the movement path, and control the vehicle to navigate the chosen path (Mar et al., 2025). The similarity (and in some cases, identicality) of hardware in the race cars to that found in existing public road AVs (such as Waymo), suggests that autonomous motorsports hardware performance can provide valuable lessons for the AV industry. This sets the stage for the study motivation as discussed in the next section.

1.3 Study Motivation

Over the decades, traditional human-driven high-speed racing has provided key stakeholders (the automotive industry, traffic regulators, driving policy makers, and road agencies), with critical lessons in safety innovation (helmets, seat belts, crumple zones), technology integration (turbocharging, fuel injection), and vehicle operation principles (alertness and situational awareness, and capability to quickly adapt to dynamic conditions on the roadway). This has spawned several improvements in vehicle design (interior and exterior), driving licensing rules and good practices, and roadway design. For example, crumple zones, reinforced cockpits, seatbelts, and airbags were developed initially for racing, but now they help protect everyday drivers. Also, racetracks have evolved to include runoff zones and advanced crash barriers to improve impact absorption. Regarding braking technology and engine technology, the literature suggests that features including disc brakes, direct fuel injection, turbocharging, and variable valve timing were perfected in racing environments before becoming standard in production vehicles.

Given such a long tradition in human-driven motorsports where several innovations from racing were eventually transferred to production vehicles, there exists great potential to reap the benefits of such symbiosis in the emerging era of autonomous mobility. The organizers of autonomous racing competitions recognize that their events are held not just for entertainment but to facilitate technology transfer. As competing teams push their standard-hardware AVs to their limits, they realize that they can be competitive only when they have developed software codes that maximize the efficacy of their capabilities for localization, sensing, control, and path planning of a vehicle in a controlled racing environment that poses extreme conditions and limitations typically encountered as edge cases in normal traffic driving.

The dichotomy of the AV "laboratory" (that is, autonomous car racetracks) and real-world AV deployment (normal driving in traffic) is featured primarily in the distinctiveness of their operational design domains (ODDs). On one hand, autonomous motorsport occurs at high speeds on highly controlled and well-defined test tracks and the vehicle carries no passengers; as such, teams are encouraged to push the boundaries of vehicle dynamics, sensor latency, and decision-making at the limits of the track environment. On the other hand, normal real-world AV operations occur with passenger vehicles in complex, low-speed urban environments and therefore are characterized by cautious, regulation-heavy development, with priority given to passenger comfort (including minimal jerk and sharp braking), fuel efficiency, and absolute safety. As technological transfer accelerates from the racetrack to the highway, it is expected that some of the boundaries that straddle this dichotomy will progressively thin out.

As experimental AV "laboratories," autonomous racing events provide significant technological and socio-economic benefits to all stakeholders (the auto industry, technology companies and hardware developers, road use regulators, and road infrastructure designers and management agencies), as illustrated in Figure 1.3. The life cycle of normal product development typically takes several years, and significant resources are expended during the validation stage of the development cycle. On the other hand, AV racing events offer rapid prototyping and obviate the long, multiple validation cycles and stringent safety requirements. Autonomous racing provides a controlled, concentrated, extreme, and adversarial testing environment where software and hardware stacks are validated. This helps vehicle manufacturers, traffic policymakers, and road infrastructure agencies earn insights into rare, high-risk scenarios that cannot be easily replicated at public roads without public safety endangerment, thereby addressing the elusive statistical "long tail" associated with autonomous driving validation. These innovations potentially inform the design of commercial autonomous cars. Lam (2018) argued that the automotive industry stands to earn significant benefits from autonomous racing events through advances in self-driving technology and safety. Also, in the May 11, 2023, edition of Autoweek, another auto technology expert, Chris Langrill, stated: "lessons learned in the [IAC] series can help make both race-car drivers and commuters on the interstate safer" (Langrill, 2023).

There exist several potential benefits, disbenefits, and issues associated with AV user trust (which will influence user demand and patronage), AV commercialization (for example, shared AVs), and AV occupant safety (which is a government responsibility and part of the stewardship of road infrastructure agencies and road designers). As such, it is expected that successful technology transfer from the racetrack to production vehicles operating on public roadways will be driven by the perspectives of the key stakeholders that may be placed into 4 broad categories, as shown in Figure 1.3 (adapted from Saeed, 2019):

- **Industry** (vehicle and AV component manufacturers and developers, and service providers),
- **Road users** (prospective AV users, and HDV users who will share roadways with AVs),
- **Government legislative and executive bodies** (at federal, state, and local levels of

government) responsible for AV-related regulations and policy formulation,

- **Road agencies and private-sector IOOs** are responsible for infrastructure development (design, construction, monitoring, and maintenance).

It is expected that the efforts of these stakeholders can be geared towards complementing each other, constituting a tapestry upon which AV adoption will thrive.

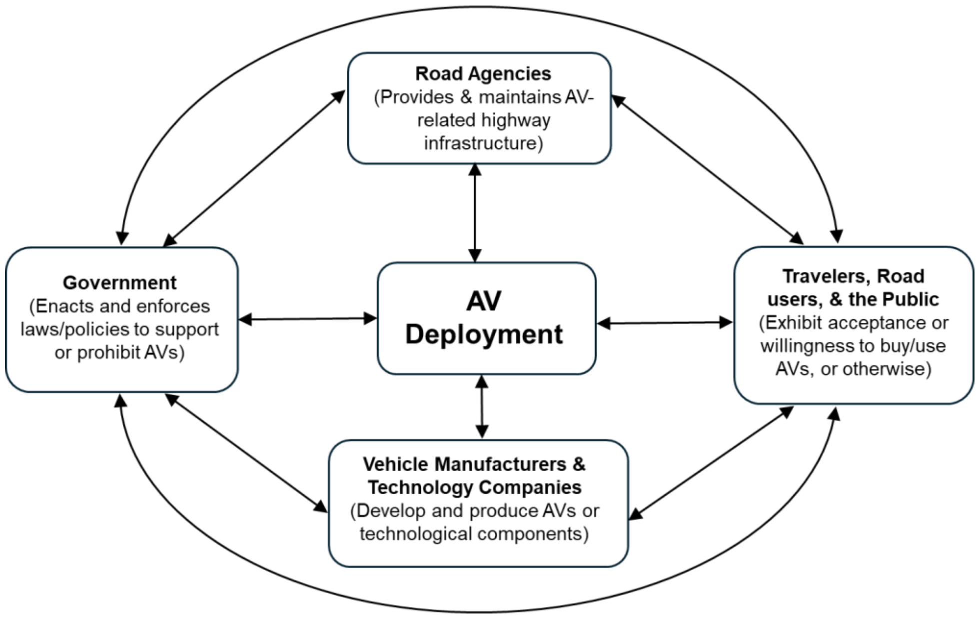

Figure 1.3 The AV stakeholder quad.

1.4 Problem Statement and Objectives

As discussed in Section 1.3, there exists a need to identify and analyze lessons from autonomous racing, particularly the implications for public-road operations. There are significant potential benefits from investigating extreme road traffic situations, including loss of tire-pavement traction, high speeds, cooperative or adversarial interactions among high-speed vehicles, and other edge cases that are characteristic of autonomous racing events. Such insights could help AV manufacturers design the vehicles to facilitate safety and efficiency. Further, from the perspective of infrastructure preparedness, it is important to investigate how high-speed autonomous racing could provide insights to enhance highway geometric design and the management of physical and cyber infrastructure to support AV operations on public roads. For example, the synthesis of insights related to road-tire friction estimation, game-theoretic trajectory planning, and hardware ruggedization as experienced in high-speed AV racing events could help inform both vehicle designers and road designers to enhance normal-road AV operations. That way, the resilience and precision associated with AV race operations can be incorporated into the automated driving systems of future AVs that use public roads and highways. Overall, it is needed to demonstrate how high-speed tracks autonomous racing events could, in many ways, serve as proving grounds for autonomous transportation. The present study was commissioned by the USDOT, through its CCAT University Transportation Center, to address these research needs.

1.5 *Latus Nota*

It must be mentioned that this study was also motivated (and its conduction facilitated) by the active involvement of the authors in several autonomous racing initiatives in the United States. These include the Indy Autonomous Challenge (where the present report's co-author was team director and lead PI of the Black and Gold Team which subsequently became the Purdue AI Racing Initiative), see Appendix 2A; the Autonomous Karting Series, AKS (where the present report's co-author is team advisor and partial sponsor of two teams: Autonomous Motorsports Purdue (see Appendix 2B) and IEEE Autonomous); and Roboracer competitions (where the present report's co-author is team co-PI of the Boiler Autonomy team). In addition, as part of the education mission associated with this project, the co-author has launched several new undergraduate courses (Purdue University College of Engineering's Vertically Integrated Projects (VIP) program) and a graduate course (see Chapter 12 for details) at Purdue University. These courses serve as formal educational counterparts to these hands-on racing initiatives and to serve as a platform where some of the lessons from AV racing for normal roadway operations could be identified and discussed in a formal educational setting. In addition, the annual student-run Next-generation Transport Systems (NGTS) conference and the accompanying Autonomy Week at Purdue provide avenues for information sharing and networking among the students that participate (or are interested) in these autonomous racing competitions.

1.6 Organization of this Report

Chapter 2 presents a systematic review of the literature on high-speed autonomous racing. Chapters 3-7 present a more detailed review of published research on autonomous racing, identifying and synthesizing such past research on vehicle localization (Chapter 3), perception and sensing (Chapter 4), data fusion (Chapter 5), path planning and decision-making (Chapter 6), vehicle dynamics and control (Chapter 7), and hardware and software safety and reliability (Chapter 8). Chapter 9 addresses lessons on road infrastructure design and management. Chapter 10 presents broader lessons, including the AV deployment and the transition period, autonomy components standardization, workforce development including talent recruitment, and the role of autonomous racing on AV safety assurance. Chapter 11 presents the conclusions, limitations, and future work, while Chapter 12 presents a synopsis of performance indicators established by USDOT. Chapter 13 presents the potential outcomes and outputs of the study, and their impacts.

CHAPTER 2 AUTONOMOUS RACING RESEARCH – FREQUENCY AND TRENDS

2.1 Introduction

Autonomous racing lies at the intersection of robotics, control theory, and artificial intelligence, where high-speed vehicles are used to push the boundaries of decision-making, safety, and adaptability for autonomous system operations. The domain is increasingly gaining prominence through competitions including the F1Tenth Autonomous Racing Championship, the Indy Autonomous Challenge, and Roborace. These initiatives not only promote innovation in algorithms used in the various modules of autonomy but also provide scalable platforms for benchmarking technologies relevant to real-world autonomous mobility. This chapter presents a systematic review of autonomous racing literature, with particular emphasis on publications between 2017 and 2025 that addressed the various modules of autonomous racing.

Research articles were identified from key academic databases, including SpringerLink, IEEE Xplore, ScienceDirect, Scopus, Web of Science, and arXiv. Titles, abstracts, and conclusions were screened to ensure relevance to autonomous racing. The extracted data included: year of publication, research topic, hardware platform, venue, and racing series. Each study was categorized by topic using keywords identified in the titles. A sample list of publications classified under the major modules of autonomous racing is presented in Table 2.1. Articles classified under more than one topic were counted in each topic category with which they were associated. Regarding the inclusion and exclusion criteria, peer-reviewed papers and conference proceedings focused on the key modules of autonomous racing received high priority. The search did not exclude preprint articles. Regarding data extraction and quality assessment, each study was independently reviewed and rated for completeness, clarity of evaluation, and rigor of validation.

This chapter presents global publication statistics, country-level topic distributions, and cumulative trends for each module of autonomous racing. In this chapter, the term topic and module are used synonymously.

2.2 Global Publication Trends

Figure 2.1 presents the cumulative number of autonomous racing papers published over the 1989-2025 period. The figure reveals two key inflection points in autonomous racing research. The first spike (2019-2020) aligns with increased adoption of simulator-based research and systematized platforms such as F1Tenth. The dramatic surge beginning in 2021 corresponds with the launch of the Indy Autonomous Challenge and simultaneous breakthroughs in AI, including more stable reinforcement learning (RL) algorithms, and transformer-based vision models. The sustained rapid growth from 2022 onward reflects a shift toward data-driven autonomy and improved computational accessibility.

Figure 2.2 presents the total number of autonomous racing papers by country. The country distribution highlights the dominance of researchers based in the United States and Germany. The U.S. spike corresponds to major academic involvement in F1Tenth and the Indy Autonomous Challenge while Germany's output reflects its contributions through Formula Student Driverless.

The emergence of China, Switzerland, and India is suggestive of a growing global participation driven by AI investments and racing initiatives, and the broader expansion of the robotics research ecosystem worldwide.

Figure 2.3 suggests that research in the autonomous racing domain is still dominated by work in control systems which comprise nearly half of all published papers (42.8%) in this domain. This dominance highlights how early autonomous racing research focused on stabilizing high-speed vehicles, designing robust controllers, and achieving consistent tracking performance laying the foundation upon which more advanced autonomy could be developed. The substantial size of the "Other" category (27.9%) suggests a strong shift toward hybrid or emerging approaches. The contents of the articles here are such that there is an increasing trend to combine classical control with reinforcement learning, sim-to-real transfer, or novel optimization techniques, reflecting a growing transition from purely model-based methods to data-driven or mixed-intelligence systems. The module with the next highest count is path planning (14.1%), driven by increasing interest in real-time aware-driven racing strategies, overtaking maneuvers, and predictive optimization. It is observed that these publications are often based on advances in GPU-based optimization and faster planning algorithms that became practical mostly after 2020. Therefore, while the field of planning and controls remains grounded in traditional control theory, it is rapidly expanding toward learning-based autonomy and strategic decision-making, mirroring broader trends in robotics and AI.

Figure 2.4 presents the countries that seem to exhibit distinct research specializations across major modules of autonomous racing. In the U.S., the research is characterized by a balance among perception, planning, and control, reflecting computational resources and competition-driven development. It seems that within the study period, research articles from Germany tend to emphasize model-based approaches that align with automotive engineering strengths. Articles from Switzerland generally shows capabilities in advanced perception research. In India and China, the research, it seems, focuses on core control and safety topics, as the autonomy ecosystems in these countries continue to expand.

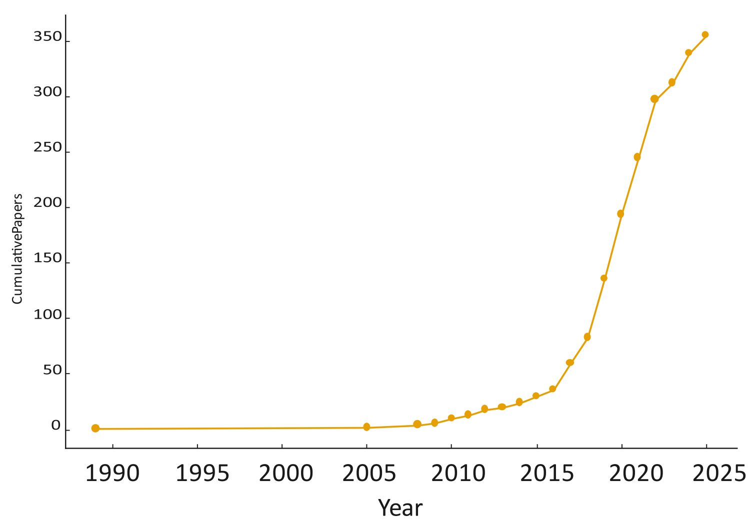

Figure 2.1 Growth in the number of published autonomous racing papers, 1998-2025.

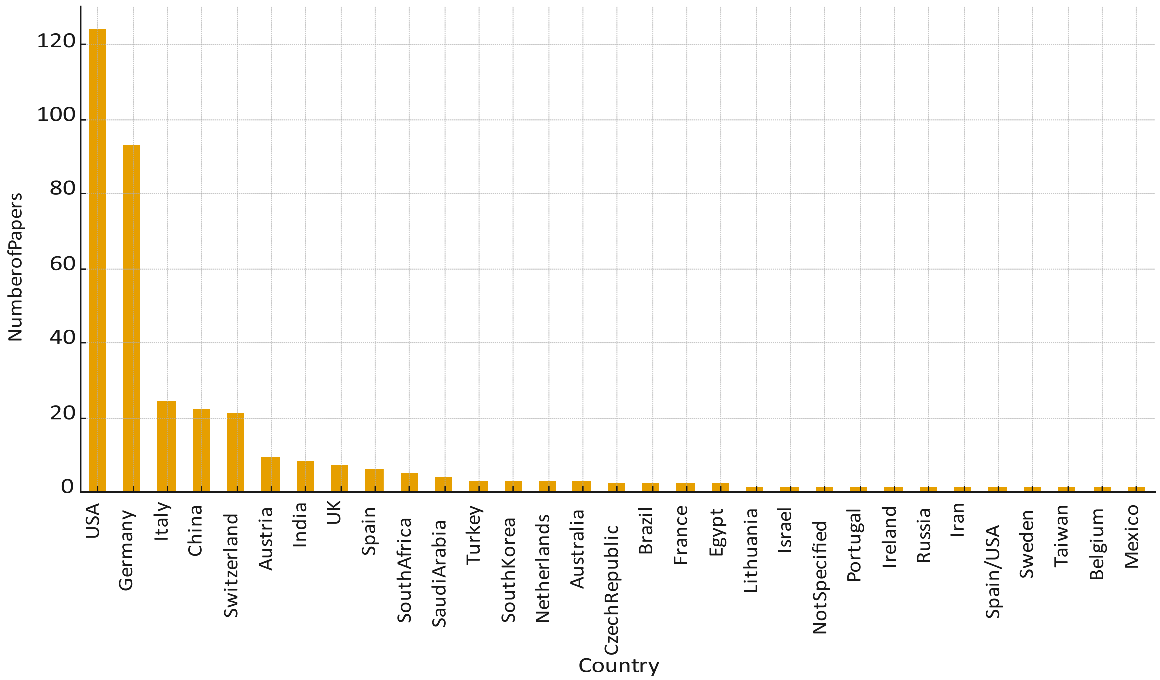

Figure 2.2 Distribution of the number of published articles by country

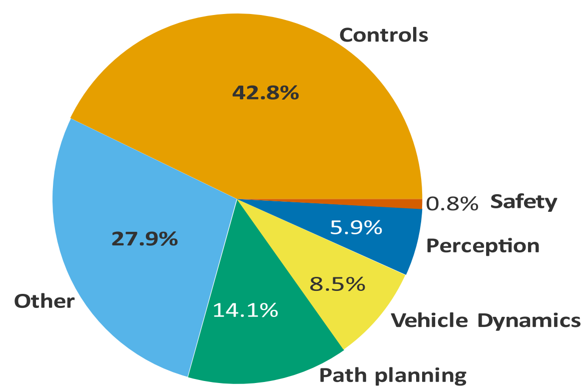

Figure 2.3 Overall distribution of autonomous racing research across the modules

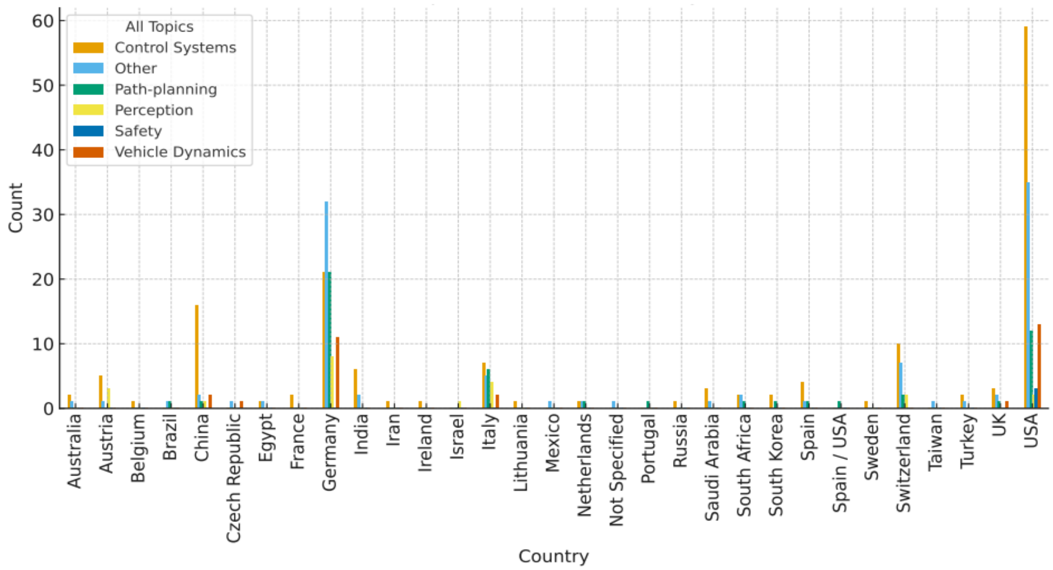

Figure 2.4 Distribution of research topics by country (1989-2025)

2.3 Topic Counts Per Year

Figure 2.5 presents the distribution of research topics addressed in published articles across each publication year. Noticeable spikes can be observed in years following key milestones related to AI and robotics advancements or adoption in autonomous racing competitions. The period following 2020 shows a growth in the number of planning and perception papers, correlating with advances within the period: the inception of transformer-based perception models, reinforcement learning stability, and increased reliance on sim-to-real pipelines during the COVID-era remote experimentation. The publication growth in all the autonomous racing modules in the post-2021 period aligns strongly with the emergence of the Indy Autonomous Challenge which significantly stimulated research contributions from multiple institutions.

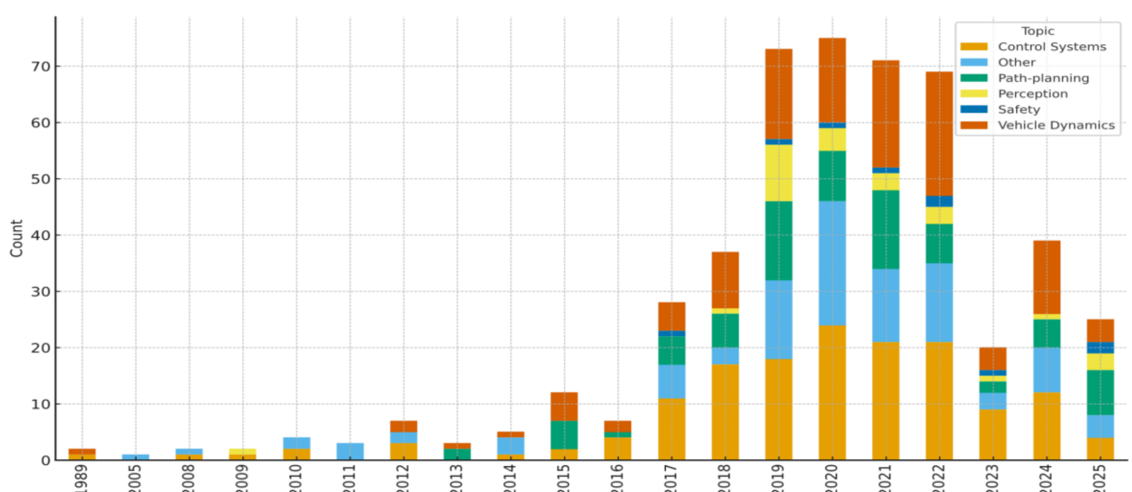

Figure 2.5 Distribution of subject area per year (1987–2025)

The chart suggests that autonomous racing publications began as early as 1989, grew gradually until 2016, followed thereafter by rapid growth, peaking around 2020 before declining modestly thereafter. The earliest publications addressed control systems and vehicle dynamics. In control systems, the number of articles increased modestly with the study period under investigation with some fluctuation in 2008–2012, dropped to nearly none in 2013, then rose to about four articles in 2016, peaked at 25 articles in 2020, and declined gradually to 5 articles in 2025. Path planning publications first appeared in 2013 with modest counts, grew unevenly from 2015 onward, peaked at about 12 articles in 2021, and then declined, with fluctuations, to about 7 articles in 2025. The trend for perception articles is somewhat similar but with longer gaps: 1 article in 2009, followed by a long hiatus until 2018 (1 article), and then a peak of 10 papers in 2019, and finally, fluctuating (but generally declining) to 3 articles in 2025. Like perception, vehicle dynamics publications exhibited a long gap after an early paper in 1989, then re-emerged around 2012 (2 articles) and grew unevenly to approximately15-20 articles annually in the 2019-2022 period, accompanied by a fluctuating decline thereafter. Safety-related publications appeared relatively late in the study period: beginning in 2017 and remaining limited, typically at 1 or 2 articles annually, with some years showing zero publication count.

2.4 Growth in the number of papers by research module

The charts in this section provide evidence of how research in each module of autonomous racing has grown over the study period, starting from the initial year of publication of an article in that module until 2025. As Figure 2.6 suggests, autonomous racing publications grew slowly between 1989 and 2016. After 2016, publications increased sharply across most modules. After 2016, the rate of growth increased sharply, with the controls module showing the sharpest increase, reaching approximately 150 articles in 2025. A notable trend is the apparent slowdown in publication growth across modules after 2023-2024. This interpretation is broadly consistent with Figure 2.5, which suggests a decline in annual publication counts after 2022. Hopefully, this is only a local dip, and the growth will resume after 2025 as IAC resumes organizing several competitions annually, as A2RL revs up its operations, and as other racing leagues continue to show growth in the number of participating teams and the number of events per year.

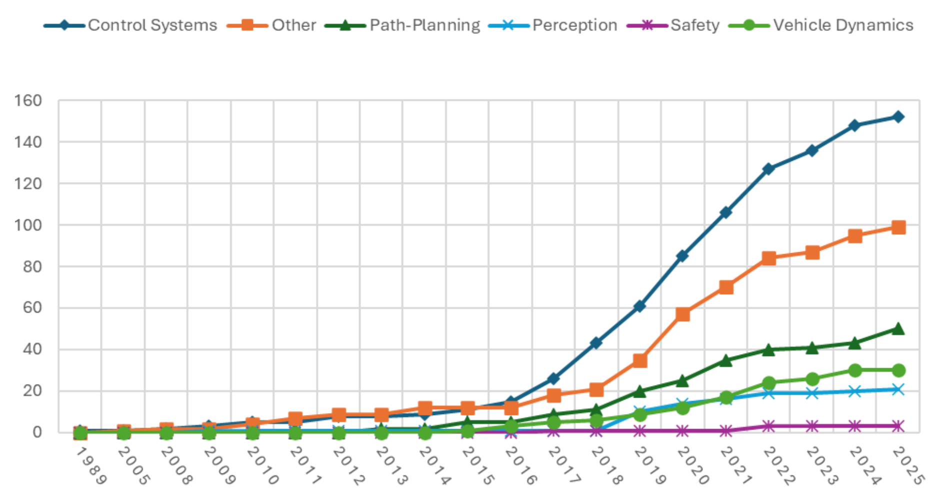

Figure 2.6 Growth in the number of papers (1989–2025) by module

2.5 Module Distribution by Country

Figures 2.7(a) to (i) present distribution of modules addressed in articles published by researchers in each of the countries that have made significant contributions in autonomous racing research.

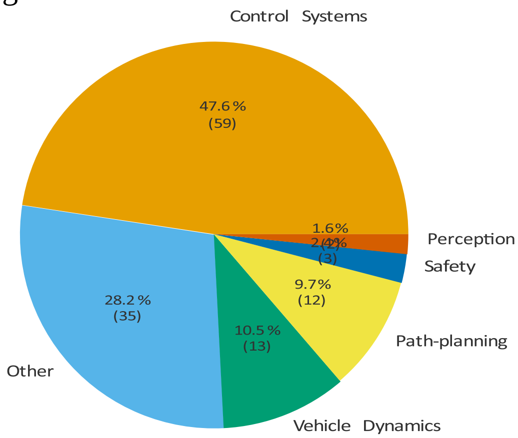

(a) United States (124 papers)

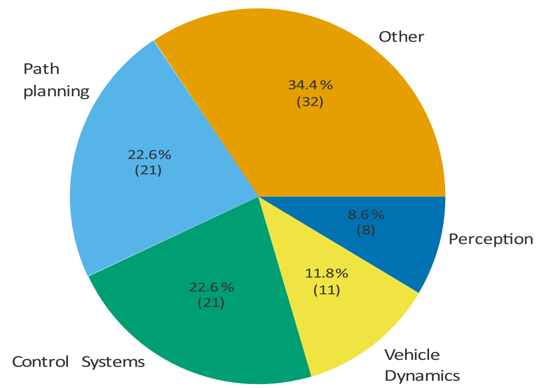

(b) Germany (93 papers)

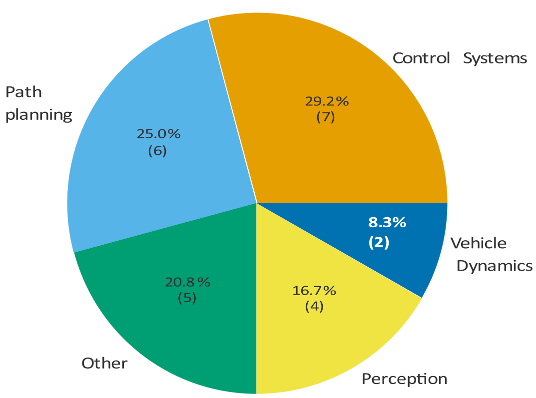

(c) Italy (24 papers)

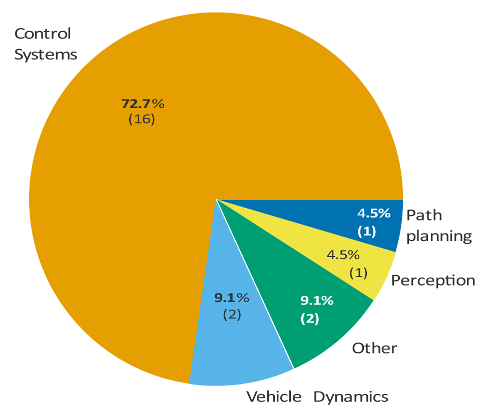

(d) China (22 papers)

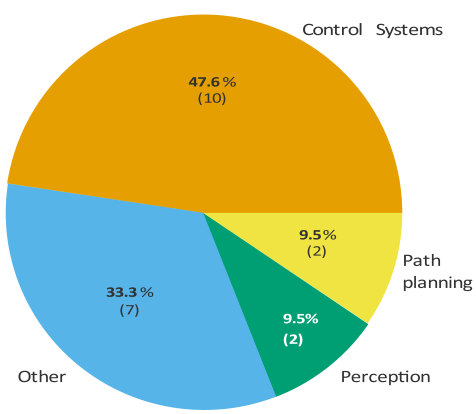

(e) Switzerland (21 papers)

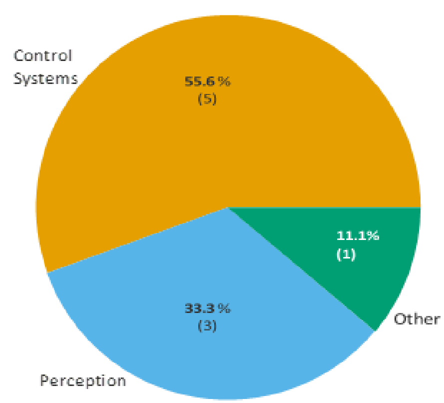

(f) Austria (9 papers)

Figure 2.7 Research topic distribution by country

CENTER FOR CONNECTED AND AUTOMATED TRANSPORTATION

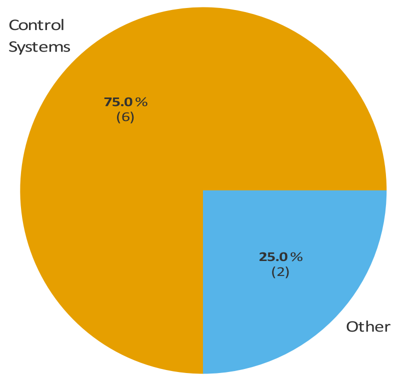

(g) India (8 papers)

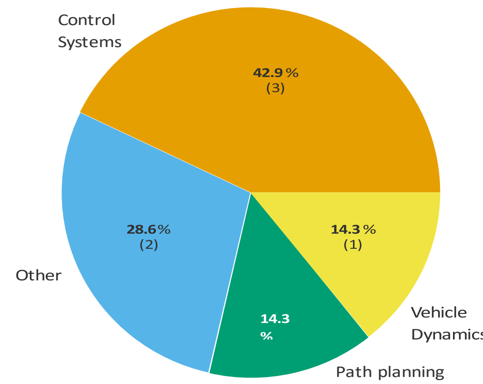

(h) United Kingdom (7 papers)

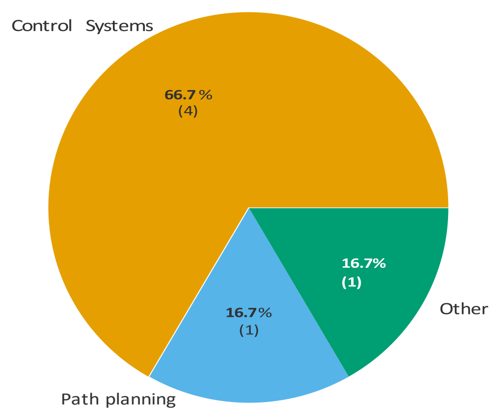

(i) Spain (6 papers)

Figure 2.7 (continued) Research topic distribution by country

The U.S. distribution (Figure 2.7(a)) suggests that research efforts in the United States are dominated by control systems, accounting for 47.6% of all autonomous racing research publications in that country. This strong dominance seems to reflect the focus, possibly, of F1Tenth and the Indy Autonomous Challenge racing teams in high-speed vehicle stabilization and robust controller design. The second-largest module, “Other” (28.2%), reflects a growing body of work in reinforcement learning, hybrid control and learning pipelines, and sim-to-real transfer studies. This trend highlights the availability of computational resources in U.S. institutions and a research culture that supports data-driven and experimental-autonomy methods probably to a greater extent compared to most other countries. Vehicle dynamics research account for 10.5% of the U.S. output, indicating sustained interest in modeling tire behavior, slip dynamics, and high-speed maneuvering. Although this proportion is smaller compared to controls-related research, it demonstrates the need for reliable physical models to support trajectory optimization and model predictive control. The remaining modules, including perception, planning, and safety, constitute relatively small shares of the overall research output but have shown notable growth in recent years as transformer-based perception models and faster real-time planning frameworks have become widely adopted. Overall, the U.S. distribution reflects a body of research that is still primarily control-focused while maintaining balanced contributions to learning-based autonomy, vehicle modeling, and advanced planning.

Regarding Germany (Figure 2.7(b)) the distribution suggests a dominance of the “Other” category, representing 34.4% of all autonomous racing research publications from that country. This category includes data-driven approaches, hybrid learning and control methods, sim-to-real transfer, and survey-style analyses that combine elements from multiple subfields. The prominence of this topic seems to indicate Germany’s growing interest in integrating classical engineering principles with machine learning in the context of high-speed autonomous racing. In that country, path planning and controls modules each account for 22.6% of the research output, reflecting balanced emphasis on robust control design and controls-planning integration. Germany possesses a long history of model-based engineering, and autonomous racing researchers in that country have made significant contributions to planning methods such as model predictive control, trajectory optimization, and multi-agent race strategy development. The controls module continues to serve as a pillar of autonomous racing research in Germany, evidenced particularly in the work of TUM and other institutions as part of the Indy Autonomous Challenge and Formula Student Driverless events. Vehicle Dynamics accounts for 11.8% of publications from Germany, reflecting continued interest in accurate modeling of slip behavior, tire forces, and dynamic limits in autonomous racing. These contributions have played key roles in the development of high-performance MPC and simulation tools. In Germany, the perception module accounts for 8.6% of the total research output which is relatively low compared with countries such as the United States and Switzerland. This suggests that during the study period, Germany’s autonomous racing research was more focused on planning and control pipelines than on large-scale perception or deep vision development. Germany’s teams have shown great success in competition events. For example, TUM autonomous racing team has shown excellent performance in IAC and A2RL events. Overall, the distribution reflects a research environment that blends strong classical engineering expertise with emerging learning-based methods. This balance between planning and control, coupled with a large proportion of hybrid or interdisciplinary work, appears to mirror the reputation and culture of German engineering institutions and their emphasis on both theoretical rigor and high-fidelity simulation.

The distribution for Italy (Figure 2.7(c)) suggests a clear emphasis on control systems and path planning, reflecting the country’s strong foundation in model-based engineering and control theory. Milan Polytechnic’s autonomous racing team (Polimove) has excelled in IAC events. Italian research institutions have a long tradition in automation, optimal control, and system modeling, which naturally translates into control-centric approaches for autonomous racing. These methods are particularly well-suited for high-speed driving, where stability, constraint handling, and performance near dynamic limits are critical. Specific problem contexts associated with the planning module, including race-line optimization, trajectory generation, and constraint-aware motion planning, have motivated researchers to leverage analytical vehicle models and well-established optimization frameworks. In autonomous racing, these approaches directly impact lap time and safety, therefore making the advancement of these approaches a research priority. Together, the dominance of control and planning in Italy’s published autonomous racing research suggests prioritization toward predictability, physical interpretability, and performance guarantees, favoring mathematically grounded solutions that integrate vehicle dynamics and control objectives. This orientation is consistent with Italy’s broader engineering culture and its decades-long history of high-performance automotive and motorsport development.

Regarding China, the distribution of autonomous racing research (Figure 2.7(d)) shows an overwhelming emphasis on control systems, with relatively little published work in path planning, perception, or hybrid categories in autonomous racing. This strong concentration reflects a research environment that prioritizes stability, robustness, and guaranteed performance, particularly for safety-critical autonomous systems. Control-oriented approaches allow researchers to leverage well-established analytical tools and physical vehicle models which are highly valued in large academic and government-funded research programs where reliability and formal validation are essential. The rather limited presence of path planning and perception literature in Chiana so far, suggests that many autonomous racing studies in that country currently focus on lower-level vehicle execution rather than high-level decision-making or data-intensive perception pipelines. In racing contexts, control systems can be developed, verified, and benchmarked in tightly constrained environments, making them attractive for rapid experimentation and repeatable evaluation. Overall, the dominance of control-related research suggests that China's autonomous racing research is driven by an emphasis on deterministic behavior, strong theoretical foundations, and performance guarantees, with vehicle dynamics and planning often treated as supporting components rather than primary research targets.

Autonomous racing research papers from Switzerland are topically distributed as shown in Figure 2.7(e), suggesting a strong emphasis on control systems, with a significant portion of publications falling into the "Other" category. This pattern may be a reflection of that country's institutional strengths in robotics, control theory, and systems engineering, as well as a research culture that prioritizes precision, reliability, and rigorous validation. Swiss research institutions are internationally recognized for foundational work in control and robotics, which naturally translates into control-focused contributions within autonomous racing. The relatively large "Other" category suggests a concentration of interdisciplinary and systems-level research, including hybrid learning–control methods, simulation frameworks, benchmarking studies, and experimental platform development. Swiss research environments are associated with emphasis on methodological rigor, reproducibility, and high-quality experimental validation, which encourages work that spans multiple subfields rather than narrowly focusing on a single component. Autonomous racing, in this context, serves as a structured testbed for validating broader robotics and autonomy innovations in the country. The low presence of path planning and perception research compared to other countries suggests that autonomous racing research in Switzerland may be prioritizing execution-level performance and system reliability over high-level decision making or data-intensive perception pipelines. Planning and perception are often treated as enabling modules that support control-centric experimentation rather than primary research drivers. This seems to be consistent with a research philosophy that favors tightly controlled experimental setups and well-defined problem formulations. Overall, Switzerland's autonomous racing research distribution reflects a research ecosystem oriented toward high-fidelity experimentation, integrated system design, and theoretically grounded control, where autonomous racing is used as a controlled, reliable, and repeatable platform for advancing robotics and autonomy research.

Regarding Austria, the topical distribution (Figure 2.7(f)) suggests a research profile centered on control systems and execution-level autonomy, consistent with the country's strong background in mechatronics, embedded systems, and real-time control. Austrian universities and research centers have a reputation for precise modeling, controller implementation, and hardware-aware algorithm design, which naturally translates into control-focused contributions in autonomous racing. In high-speed racing environments, these approaches are particularly valuable, as vehicle stability, actuator constraints, and timing determinism are critical to safe and competitive operation. The relatively limited focus on perception and planning suggests that

Austrian autonomous racing research may be treating these components as enabling technologies rather than primary research drivers. Instead, many studies appear to focus on how controllers behave under uncertainty, model mismatch, and tight dynamic constraints. Autonomous racing, in this context, serves as a compact and demanding validation platform for embedded control strategies, allowing researchers to evaluate robustness and real-time feasibility without relying on large-scale sensing pipelines or data-heavy learning frameworks.

India's topical distribution (Figure 2.7(g)) reflects a research ecosystem that is rapidly expanding its engagement with autonomous racing, with emphasis on control systems, safety, and foundational capabilities. This pattern is consistent with the widespread adoption of low-cost, open-source platforms such as F1Tenth, which have low entry barriers and allow institutions to prioritize core control and stability problems. As a result, many autonomous racing research contributions from India focus on reliable vehicle execution, controller tuning, and safety-oriented design rather than on computationally intensive perception or large-scale learning approaches. The comparatively limited presence of perception and advanced planning likely reflects practical constraints related to sensor cost, data availability, and computational resources. This does not imply a lack of sophistication; instead, it highlights a strategic focus on building robust baselines and reproducible experimental pipelines. Autonomous racing serves as a steppingstone for broader autonomous systems research, enabling Indian institutions to establish strong control and validation expertise before scaling toward more complex perception and multi-agent decision-making problems.

In the United Kingdom, the topical distribution of autonomous racing research (Figure 2.7(h)) suggests a balanced engagement across control, planning, and interdisciplinary work, apparently reflecting a research culture that seems to emphasize simulation, algorithm rigor, and cross-domain transfer. U.K. research groups often approach autonomous racing as a testbed for general autonomy problems, generally drawing inspiration from aerospace, robotics, and unmanned systems research. This background encourages exploration of planning and decision-making methods alongside classical control, particularly in simulated or hybrid testing environments. The presence of both planning and perception-oriented contributions suggests that U.K. research frequently targets system-level autonomy rather than isolated controller performance. Autonomous racing platforms are often used to validate planning strategies, uncertainty-aware decision-making, and scalable simulation pipelines. This positions researchers in the U.K. as significant contributors to methodological advances that generalize beyond racing, using high-speed autonomy as a structured but challenging experimental domain.

Spain's topic distribution (Figure 2.7(i)) is strongly concentrated in control systems, accounting for four of the six publications from that country. This suggests that the country's autonomous racing research has primarily focused on vehicle stabilization, trajectory tracking, and execution-level performance, areas that form the technical foundation for any high-speed autonomous system. Control-centric studies provide a practical entry point into autonomous racing, as they allow researchers to validate algorithms under extreme dynamics without requiring large-scale sensing infrastructure or extensive training data. The limited presence of "path planning" and "other" topics indicates that higher-level decision making and interdisciplinary approaches remain secondary research directions. Instead, planning and system-level integration are often treated as supporting components within control-focused frameworks rather than as standalone contributions. Overall, the distribution of modules reflects a research ecosystem that is still consolidating core control expertise in autonomous racing, positioning itself to expand toward more advanced planning and hybrid methodologies as experimental infrastructure and participation in racing competitions continue to mature.

2.6 Summary and Discussion

This section synthesizes visual and quantitative analyses and connects publication dynamics to broader technological, institutional, and methodological developments in autonomous racing. The combined trends reveal not only when and where the field expanded, but also the ecosystem pressures, available tools, and research incentives that shaped the trajectories of different modules of autonomous racing.

2.6.1 Global Publication Dynamics

The cumulative publication trend (Figure 2.1) illustrates a maturing research domain influenced by major technological and competitive milestones. The first inflection point (2019/2020) coincides with the widespread adoption of accessible, modular simulation platforms such as Gazebo-based F1Tenth environments, CARLA extensions for racing, and ROS2-based tooling. These platforms lowered the barrier to entry for institutions without access to high-speed racecar hardware. Additionally, COVID-19 restrictions indirectly accelerated simulator-driven research by forcing experimentation into virtual environments, leading to a surge in reproducible algorithm development. The much steeper rise beginning in 2021 strongly aligns with the launch of the Indy Autonomous Challenge (IAC). Unlike prior competitions, the IAC introduced a much larger vehicle (Dallara AV21 and AV24) and promoted the development of high-speed racing as a research platform, motivating tertiary institutions to invest in simulation, perception, controls, and planning in extreme dynamic conditions on the racetrack. The availability of standardized hardware platforms and shared development pipelines encouraged multi-institutional collaborations, spawning a significant spike in the number of research publications related to hardware and simulation. In this regard, the inception of A2RL competitions helped to complement the wider mission of the IAC. From 2022 onward, the field benefited from leveraging breakthroughs in artificial intelligence, particularly stable model-free reinforcement learning algorithms, transformer-based detection architectures, and improved scene representations, which collectively enabled researchers to explore more complex tasks such as long-horizon planning, multi-agent racing, and robust sim-to-real pipelines.

2.6.2 Country-Level Trends and Alignment with Capabilities

The distribution of publications by country (Figure 2.2) reflects how institutional strengths and national research ecosystems have shaped autonomous racing research output. The dominance of the United States could be attributed to the country's extensive participation in F1Tenth and the IAC, well-funded robotics laboratories with dedicated autonomy programs, and access to high-performance computing resources that encourage perception and learning-based research. Germany's high output directly correlates with its strong automotive engineering heritage, producing many works on model-based control and vehicle dynamics. Formula Student Driverless also plays a major role, as German teams tend to publish detailed system architecture and control methodologies.

The country-specific topical distributions (Figures 2.7(a) to (i)) reveal noticeable specialization patterns in autonomous racing research across the countries, as exemplified with a few examples presented below:

- United States: A relatively balanced distribution across perception, planning, and control modules could be manifesting the interdisciplinary nature of U.S. autonomous racing research laboratories and their integration of AI-driven approaches with hardware implementations.
- Germany: Strong emphasis on vehicle dynamics and control reflects the country's automotive engineering culture and long-standing capabilities in optimal control and real-time model predictive control (MPC), and the technical excellence of institutions like the Technical University of Munich (TUM).
- Switzerland: High engagement in perception aligns with global leadership in robotics, SLAM, and 3D vision of Swiss institutions including ETH Zurich.
- China and India: Increasing participation reflects expanding investments in AI and robotics. Early publications focus on classical control and safety, suggesting that hardware acquisition and simulation infrastructure are still scaling upward.

These patterns suggest that autonomous racing is not merely a robotics challenge but a reflection of national or institutional research priorities (explicit or implicit), the technological research infrastructure available in that country, and the trajectory of autonomous systems research in that country's tertiary engineering institutions.

2.6.3 Topic Trajectories and Technological Correlates

The topic-specific charts (Figure 2.6) demonstrate how the number of publications has grown over the study period (1989-2025). These growths could be attributed to external technological developments, competition milestones, and algorithm advances.

<u>Control Systems.</u> Research publications in control systems appeared earliest compared to other modules, and grew steadily, driven by advanced autonomous racing's fundamental requirement of reliable vehicle stabilization and tracking from the baseline. Growth spikes from 2018-2020 align with easier integration of nonlinear MPC frameworks and fast quadratic programming solvers. The renewed upsurge in controls-focused research articles in 2021 seems to reflect the IAC event's requirements for high-speed stability, emergency avoidance, and robust performance under extreme dynamics. These requirements catalyzed controls-related innovations for autonomous racing.

<u>Vehicle Dynamics.</u> Autonomous racing vehicle dynamics research has grown significantly from its early years, corresponding with efforts to accurately model tire slip, aerodynamics, and racecar-specific effects. These models became essential for MPC, trajectory optimization, and simulation of fidelity improvements. The growth parallels the improvements in higher-fidelity simulators, including custom Formula Student Driverless's simulators and IAC's AV21-specific physics models.

<u>Path Planning.</u> Planning research surged after 2020, catalyzed by advances in sampling-based and optimization-based planners, including cross-entropy method (CEM) planners and model predictive path integral (MPPI) control. As racing circuits require high-speed, high-curvature decision making, researchers began to combine trajectory optimization with learned heuristics and predictive models, enabling more competitive racing behavior and improved robustness to disturbances.

Perception. Autonomous racing perception research exhibits a relatively late but steep growth curve. This mirrors global progress in deep vision, including transformer-based networks, LiDAR-camera fusion, and bird's eye-view (BEV) representations. As robustness improved, researchers extended perception capabilities beyond simple track-marker detection into semantic track understanding and multi-object prediction for purposes of adversarial racing.

Software and Hardware Safety. Safety research in autonomous racing has grown gradually but steadily, reflecting a shift from pure performance to reliability and formal guarantees. As competitions matured, their real-world race competitions began encountering track failures such as sensor dropout, actuator saturation, and model mismatch. These challenges motivated the racing teams to pursue research in control barrier functions, reachability-based safety verification, fail-safe architectures, and redundancy in sensing and communication.

Other Topics. The "Other" category grew significantly between 2021 and 2023, coinciding with the emergence of stable reinforcement learning algorithms (e.g., SAC, TD3, PPO), domain randomization, and sim-to-real training pipelines. Many papers explored hybrid systems where classic MPC or rule-based controllers are combined with learned policies. This reflects a trend toward integrated autonomy, in a departure from purely model-based or purely data-driven approaches.

2.6.4 Integrated Interpretation of Trends

From the perspective of the methodological contributions made in autonomous racing research literature, the publication trends suggest that the field has undergone, and is still undergoing, a clear evolution across 5 phases/ Each phase exhibits its unique growth trend, emphasis areas, and general characteristics:

1. 1989-2017 (The Inception Era): Scattered research publications in various areas of autonomous racing, mostly in controls. The number of research articles grew slowly but steadily over this period.
2. 2017-2019 (The Foundations Era): Autonomous racing publications emphasized control systems and vehicle dynamics, with limited contributions in sensing, perception, and planning, and virtually none in learning-based methods.
3. 2019-2021 (The Simulator Expansion Era): The literature showed greater emphasis on simulation, with significant growth in planning and decision-making emerging contributions in deep learning, particularly for perception.
4. 2021-2023 (The Competition-Driven Acceleration Era): Driven primarily by the Indy Autonomous Challenge and related events, this era saw sharp growth in autonomous racing research, catalyzed by multi-agent racing (which enhanced perception research), high-speed stability (which spawned research innovations in controls systems), end-to-end pipeline frameworks, and software and hardware safety and validation.
5. 2023-2025 (The AI Integration Era): Transformers, reinforcement learning, domain adaptation, and sim-to-real pipelines became more central, often combined with classical control. Several autonomous racing research publications have begun to incorporate machine learning and other AI tools in the modules of perception, control, path planning, and decision making, and in end-to-end architectures.

Arguably, the charts suggest that the growth of autonomous racing research is governed by several factors including computing power, hardware standardization, AI and learning, and most importantly yet implicitly, the allure of the wider field of autonomous transportation systems and their concomitant virtues of safety, mobility, and economic productivity. The volume, topical distribution, and growth of autonomous racing research publications suggest that autonomous racing events continue to serve as auspicious laboratories. At these events, researchers from multiple disciplines (computer science, computer engineering, electrical engineering, etc.), institutions, and countries worldwide not only compete but also collaborate to advance racecar autonomy. This ultimately provides advances towards autonomous mobility for the benefit of society.

2.6.5 Closing Words

Autonomous racing research has evolved rapidly since 2017, with advances in control, perception, and reinforcement learning. Many of the racing teams have published their findings in conference proceedings and technical journals (see Table 2.1). So, autonomous racing research has come a long way. Nevertheless, challenges still persist in areas such as replicability, benchmarking, AI-based learning, cybersecurity, and real-world validation. Future research should investigate the use of AI tools to enhance the key modules of automated driving: perception and sensing, data fusion, path planning and decision-making, and vehicle dynamics and control. Racing teams should continue to emphasize cross-compatible simulation frameworks with standardized physics engines. Other research directions should include the continued development of open user-friendly race-data repositories to help bridge simulation-to-reality gaps and encourage a shift from module-specific research toward end-to-end pipelines powered by AI. Autonomous racing leagues are expected to continue expanding to multi-agent racing for competitive (and ultimately, collaborative) decision-making. By presenting a synthesis of existing research efforts across various modules of autonomous racing, along with their trends and distributions, this chapter aims to promote the realization of autonomous racing's progress and potential.

Table 2.1 A sample publications related to the various modules of autonomous racing

Module	References
Controls	Alcala et al. (2020b), Al-Sunni et al. (2025), Adamy, (2022), Anderson et al. (2016), Arab and Yi (2021), Atreya et al. (2022), Brown & Gerdes (2020), Brunner et al. (2017), Buyval et al. (2017), Caporale et al. (2018), Carrau et al. (2016), Cataffo et al. (2022), Corriou (2017), de Bruin et al. (2018), Du et al. (2022), Eguchi & Date (2023), Escoriza et al. (2021), Falb (2019), Fröhlich et al. (2022), Funke et al. (2017), Gandhi et al. (2021), Gao et al. (2018), Goblirsch et al. (2024), Gundu et al. (2019), He et al. (2022a), He et al. (2022b), Herrmann et al. (2019), Hermansdorfer et al. (2019), Hindiyeh and Gerdes (2014), Hobbs (2025), Hou et al. (2022), Ji et al. (2018), Jianwang et al. (2021), Joa et al.(2020), Kalaria et al. (2022), Kapania & Gerdes (2015a); Kapania & Gerdes (2015b); Ke et al. (2019), Kloeser et al. (2020), Korkmaz et al. (2018), Kritayakirana & Gerdes (2010), Kritayakirana & Gerdes (2012a), Kritayakirana & Gerdes (2012b), Li et al. (2022), . Liniger et al. (2015), Liniger et al. (2017), Liniger et al. (2019), Liu et al. (2020), Liu et al. (2021), Lu et al. (2023), Lu et al. (2021), Manna et al. (2022), Micheli et al. (2022), Ni et al. (2019), Niu et al. (2022), Notomista et al. (2020), Novi et al. (2019), Pagot et al. (2020), Park & Gerdes (2017), Raji et al. (2022), Rosolia & Borrelli (2018), Shiotsuka et

	al. (2024), Srinivasan et al. (2020), Stahl & Betz (2020), Stahl et al. (2020), Sun et al. (2023), Tătulea-Codrean et al. (2020), Tsai et al. (2019), Vallon & Borrelli (2020), Verschueren et al. (2016), Voser et al. (2010), Wachter et al. (2019, Wachter et al. (2020), Williams et al. (2016), Williams et al. (2018a), Williams et al. (2018b), Williams et al. (2018c), Wischnewski et al. (2020), Wischnewski et al. (2021), Wischnewski et al. (2022), Wohner et al. (2021), Yang et al. (2022), You & Tsiotras (2017), You & Tsiotras (2021), Yuan et al. (2019), Zarrouki et al. (2024a), Zarrouki et al. (2024b), Zarrouki et al. (2024c), Zarrouki et al. (2024d), Zhang et al. (2022c), Zhang et al. (2022d), Pour et al. (2021), Rosolia et al. (2017), Kabzan et al. (2019), Williams et al. (2017), Thakkar et al. (2024), Wurman et al. (2022), Wadekar et al. (2021), Coulter (1992), Hoffmann et al. (2007), O'Kelly et al. (2020).
Hardware and software safety	Evans (2023), Hafemann et al. (2023), Enisz et al. (2024), Massa et al. (2020), Rathore et al. (2022), Aptiv (2025), Ha et al. (2020), Hermansdorfer et al. (2020), Hilary (2021), Ivanov (2020), Acharya & Mekker (2022), Nakanishi and Auza (2023), Zenoh (2021), Iclodean et al. (2020), Nekkah et al. (2022). Patton et al. (2021), Yang et al. (2024).
Sensing and perception	Buckman et al. (2022), Bulsara et al. (2020), Cai et al. (2021), Cai et al. (2020), Cellina et al. (2025), Chamain et al. (2021), Dat et al. (2025), De Rita et al. (2019), Dhall et al. (2019), Dikici et al. (2025), Drage et al. (2014), Drews et al. (2019), Du et al. (2024), Evans (2021a), Evans (2021b), Falanga and al. (2019), Franke (2017), Fursa et al. (2022), Goh & Gerdes (2016), Herrman et al. (2021), Sauerbeck et al. (2023), Betz et al. (2022), Lim & Park (2025), Li et al. (2025), Szabo & Pup (2025), Kabzan et al. (2020), Buckman et al. (2022), Reveille Racing & Dietz (2023), Xia & Chen (2024), Kang et al. (2019), Pendleton et al. (2017), Kapocsi et al. (2025), Saka and Labi (2023), Hafemann et al. (2023), Zhao et al. (2025), New Eagle (2022), Mahmoud et al. (2020), Mar & Dietz (2024), Puchtler & Peinl (2020), Ramachandran et al. (2022), Renzler et al. (2020), Strobel et al. (2020), Zhuo et al., 2023, Wang et al. (2021a).
Decision-making, including path planning	Alcala et al. (2020a), Alrifaee et al. (2021), Alrifaee & Maczijewski, (2018), Andrei (2022), Arslan et al. (2017), Banjanovic-Mehmedovic et al. (2016). Bertogna (2022), Bevilacqua et al. (2017), Bhargav et al. (2022), Biswas et al. (2021), Bonab & Emadi (2019), Cai et al. (2024), Cao et al. (2024), Cardamone et al. (2010), Cartró et al. (2024), Chen et al. (2025), Christ et al. (2019), Coulter (1992), Dong et al. (2023a), Dong et al. (2023b), Dong et al. (2022), Evans et al. (2021a); Fehér et al. (2020), Feraco et al. (2020), Fisac et al. (2019), Fu et al. (2018), Garlick et al. (2022), Gosala et al. (2019), Hang et al. (2020), Herrmann et al. (2020), Hernandez et al. (2018), Keanly and Engelbrecht (2023), Kegelman et al. (2016), Kessler et al. (2020), Langmann et al. (2025), Le Floch et al. (2017), Le Large et al. (2021), Lovato & Massaro (2021), Ma et al. (2022), Micheli et al. (2022), Mirebeau et al. (2022), Mostaghim (2022), Ni et al. (2017), Ni et al. (2019), Park & Gerdes (2015), Piccinini et al. (2022), Piccinini et al. (2025), Raji et al. (2022), Reiter et al. (2021), Rowold et al. (2024), Rucco et al. (2015), Sridevi et al. (2021), Srinivasan et al. (2021), Subosits & Gerdes (2019), Subosits & Gerdes (2021), Trauth et al. (2024a), Trauth et al. (2024b), Werner et al. (2025b), Wolfgang Kuhn (2017), Zhang et al. (2022a), Zubaca et al. (2020b), Garlick & Bradley (2022), Javed et al. (2022), Jeon et al. (2013), Joglekar (2021), Thakkar et al. (2024), Rowold et al. (2022), Li et al. (2024), Remonda et al. (2021), Toaz & Bopardikar (2025), Theodosis & Gerdes (2011), Theodosis & Gerdes (2012).
Fusion	Abdulmaksoud & Ahmed (2025), Acharya and Mekker (2022), Patz et al. (2008), Kabzan et al. (2020), Massa et al. (2020), Peng et al. (2021), Betz et al. (2023), Karle et al. (2023), Lee et al (2024), Suprapto (2024), Cellina et al. (2025), Heilmeier et al. (2020b), Heilmeier et al. (2020c), Montani et al. (2021), Palafox et al. (2019), Palatti et al. (2021), Valls et al. (2018).

Other topic or multi-module	Agnihotri et al. (2020), APTIV (2020), Alves et al. (2019), Babu & Behl (2020), Bak et al. (2022), Baumann et al. (2024), Behl (2025), Bell (2020), Betz et al. (2019), Betz et al. (2022), Betz et al. (2023), Bolla et al. (2022), Bosello et al. (2022), Braghin et al. (2008), Burton & Gajjar (2024), Butz & Lonneker (2009), Cacciaguerra et al. (2005); Camara & Fox (2022), Caporale et al. (2019), Charles et al. (2025), Chisari et al. (2021), Chatzopoulos et al. (2019), De Caro & Corti (2018), Eken et al. (2020), Evans et al. (2021a), Evans et al. (2023a), Evans et al. (2023b), Evans et al. (2023c), Foris et al. (2020), Fox et al. (2018), Fuchs et al. (2021), Funke et al. (2012), Giesen et al. (2022), Goldfain et al. (2019), Gosala et al. (2020), Guckiran and Bolat (2019), Guo & Wu (2019), Hartman et al. (2023), Hafemann et al. (2023), Herrmann et al. (2022), Jaritz et al. (2018), Kaufeld et al. (2024), Khan et al. (2022), Koirala et al. (2024), Kong et al. (2021), Lee & Lee (2023), Liu (2024), Loiacono et al. (2010), Marley et al. (2021), Mehta & Mistry (2017), Naik (2019), Niu et al. (2020), Noorvand et al. (2017), Oliveira et al. (2018), Pan et al. (2018), Perot et al. (2017), Peyron et al. (2021), Quadflieg et al. (2011), Raji et al. (2020), Remonda et al. (2021), Reyes & Alexander (2020), Rito Lima et al. (2020), Sa et al. (2020), Salem et al. (2017), Salem et al. (2018), Salem et al. (2019), Scholze et al. (2024), Schwarting et al. (2021), Shah et al. (2023), Smith & Mistry (2020), Song et al. (2021a), Song et al. (2021b), Spielberg et al. (2019), Suresh Babu & Behl (2022), Tătulea-Codrean et al. (2020), Talvala (2011), Trauth et al. (2023a), Trauth et al. (2023b), Wang et al. (2021b), Werner et al. (2024), Werner et al. (2025a), Yamaguchi et al. (2022), Yoo and Langari (2012), You & Tsiotras (2018), Zhang et al. (2022b), Zubaca et al. (2020a), Tian et al. (2018).
Physical and Cyber Infrastructure	Carreras et al. (2018), Chen et al. (2023), Colonna et al. (2018), Farah et al. (2018), Georgouli et al. (2021), Gopalakrishna et al. (2015), Gungor & Al-Qadi (2020).
Localization	Enisz et al. (2024), Gotlib et al. (2019), Kumar & Muhammad (2023), Muhammad, (2012), Shin et al. (2020), Stahl et al. (2019), Goldfain et al. (2019), Wischnewski et al. (2019), Nobis & Betz (2019), Kabzan et al. (2020), Schratter et al. (2021), Sauerbeck et al. (2022), Williams et al. (2017), Betz et al. (2023), Baumann et al. (2024), Massa (2020), Kabzan (2020), Li et al. (2020), Karle (2023), Betz (2023), Li et al. (2020).

CHAPTER 3 LOCALIZATION

"In order to get anywhere, you must first know where you currently are." — Michael O'Neill.

3.1 Introduction

Prior to deciding where to go (path planning), how to get there (controls), or even interpreting the surrounding environment (sensing and perception), an autonomous vehicle must first determine where it is located currently. Knowledge of the position of the AV at any given time is essential for accurate navigation and motion planning because it ensures that vehicle decisions are grounded in the actual operating environment rather than in uncertain assumptions about its location. In autonomous vehicle operations, localization is commonly defined as the process of determining the AV's precise location (often to the nearest inch) within a high-definition (HD) map, allowing for higher-level decision tasks such as navigation, perception, and planning can be performed reliably. Kumar and Muhammad (2023) described localization as the process of estimating the AV's pose (orientation and position) and the uncertainty associated with that pose in a reference frame. Localization, which is essential for safe autonomous vehicle navigation in diverse and complex environments, relies on a fusion of data from one or more sensor types, such as cameras, LiDAR, GNSS, radar, and IMU, with map-matching techniques.

The AV performs localization not only at the beginning of a trip but at continuously throughout the trip. This is because, the AV must continuously account for changes in its environment, such as entering a new location, encountering updated map conditions, or detecting the positions of stationary or moving objects on the roadway. There are several approaches for AV localization, and the sensing types and computational requirements vary across them. Common approaches used or considered in autonomous racing include map-based localization, dead reckoning using IMU, simultaneous localization and mapping (SLAM), beacon-based localization, vision-based localization, wireless communication or RFID-based localization, and machine-learning-based methods (Muhammad, 2012; Kumar and Muhammad 2023). Figure 3.1 presents an example of a localization algorithm architecture.

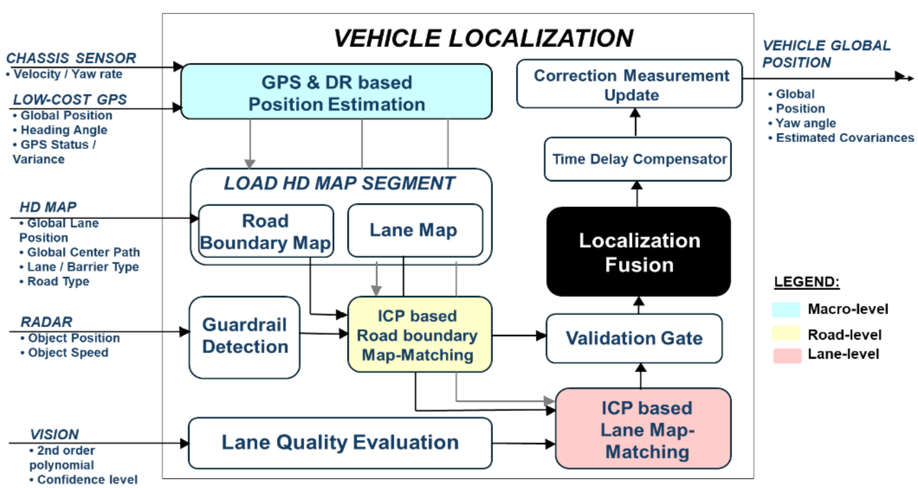

Figure 3.1 An example of a localization algorithm architecture (Shin et al., 2020)

3.2 A Review of Selected Literature

Table 3.1 presents selected localization studies in autonomous racing, summarizing their primary sensors, localization algorithms, and key lessons for real-world high-speed AV operations. One of the most cited works in autonomous racing localization is that of Stahl et al. (2019) who used ROS to develop a LiDAR-based localization approach for high-speed autonomous racecars. Their innovation was founded on the notion that LiDAR-measured pose (the position and orientation of the AV during motion) could be used as parallel redundant information and fused into the overall state estimate. This improved the robustness and real-time performance of the AV operations.

For localization in AutoRally, a popular platform for aggressive-driving and racing research (Figure 3.2), Goldfain et al. (2019) fused GNSS and IMU data to support operation in off-road and high-slip conditions, with particular attention to time synchronization and estimator stability under fast transients.

In a study that coupled localization and vehicle-dynamics state estimation for autonomous racing, Wischnewski et al. (2019) investigated estimator behavior under high acceleration and unstable communication. The novelty of their algorithm lay in the use of model-based filtering coupled with racing-grade dynamic models. The authors identified a set of parameters that must be estimated beyond the AV pose itself, including velocity components and slip-related states, in order to maintain the stability of the localization process and its output.

The insightful study by Nobis and Betz (2019) evaluated SLAM alternatives for outdoor racing-scale environments, with emphasis on performance and failure modes relevant to racing, including high speed, limited compute, and degraded features. The authors presented a valuable systematic comparison of alternative SLAM approaches in a high-speed racing environment.

To maintain AV pose within handling constraints, Kabzan et al. (2020) developed a localization method based on SLAM/state estimation for cone-delimited racetracks, fusing outputs from vehicle sensors with map information. Their study demonstrated how real-time racing localization can be engineered end-to-end in a competitive race environment.

Massa et al. (2020) developed a localization pipeline for autonomous racing in a GNSS-denied environment. They used multiple planar LiDAR sensors and onboard IMUs. Their combined map-based LiDAR localization with iterative closest point (ICP)-type scan matching against a prebuilt track map, together with temporal smoothing and fusion through a multi-rate Kalman-filter framework. Attitude monitoring was supported by an IMU orientation filter to stabilize scan alignment under racing dynamics. The major contribution of the study was the demonstration of robust, low-latency localization under aggressive racecar maneuvers.

Recognizing the severely limited preparation time available on racetracks during competitions, Schratter et al. (2023) developed a racing-ready mapping and localization pipeline that spans high-precision LiDAR mapping and online LiDAR localization. Specifically, the localization method relied on LiDAR scan matching against a high-quality map, with adaptations made to account for high operating speeds and sparse environmental features.

Sauerbeck et al. (2022) developed a multi-sensor localization architecture that fuses independent longitudinal and lateral estimators to promote AV stability at high speeds. The authors implemented a combined LiDAR-camera localization approach that decoupled lateral and longitudinal states using a track-bound coordinate system. Their approach integrated; (a) a modified vision SLAM to estimate longitudinal progress (distance traveled), and (b) 3D LiDAR detection of side barriers to estimate the lateral offset.

Williams et al. (2017) addressed the ecosystem for autonomous racing research using Georgia Tech's AutoRally platform (Figure 3.2) and simulation interfaces. The authors treated localization and a state estimation as critical prerequisites for the repeatable validation of racing strategies. Their localization approach establishes how GNSS/IMU-based state estimation and simulation ground truth can be used to support real-to-sim of evaluation racing algorithms.

Betz et al. (2023) described a full autonomous racing stack that provides redundancy and graceful degradation at high speed. The stack leverages outputs from multiple sensors that support localization, for example, dual GPS, LiDAR localization, and multiple IMUs. Their localization framework involves sensor-fusion state estimation (multi-sensor pose/velocity estimation) with independent pipelines that feed into a supervisory fusion/selection layer. One contribution of the study is its clearly articulated description of how racing teams enhance the efficiency of the localization tasks.

Enisz et al. (2024) enhanced localization robustness by using a multi-model EKF suite and a logic-based arbitration process for autonomous racecars. The authors used four EKF localization variants and a supervisory selection algorithm that switches among them based on the prevailing speed regime and the availability and quality of GNSS. In their approach, the authors used a kinematic model for low-speed EKFs and a dynamic vehicle model for high-speed EKFs, thereby adapting the localization framework to changing driving conditions.

For F1Tenth head-to-head racing, Baumann et al. (2024) documented a reproducible stack in which localization is performed on a pre-mapped track. They used an odometry filter and a map-referenced localization module designed for repeatable racing conditions and rapid deployment. The approach is useful for localization in adversarial autonomous racing environments where latency budgets, modularity, and robustness are critical.

Figure 3.2 The AutoRally platform for testing aggressive high-speed self-driving, developed at Georgia Tech, led by Brian Goldfain, Paul Drews, and James Rehg (Williams et al., 2017)

Table 3.1 Selected localization studies: Primary sensors, localization algorithms, and lessons for high-speed AV operations on real-world roadways

Paper	Primary sensors	Algorithm used for localization	Key takeaways/lessons
Stahl (2019)	LiDAR (+ others)	LiDAR poses as a redundant stream fused into the estimator	Practical high-speed LiDAR localization integration
Wischnewski (2019)	Vehicle sensors + GNSS/IMU	Model-based filtering for racing dynamics state + localization	Highlights the need for dynamics-aware estimation near limits
Goldfain (2019)	GNSS + IMU	Kalman-filter family GNSS–IMU fusion for aggressive dynamics	Reference platform/architecture for high-dynamics estimation
Nobis & Betz (2019)	Varies	Comparative SLAM evaluation under racing constraints	Direct comparison of SLAM choices for racing
Massa (2020)	LiDAR + IMU	Map-based LiDAR localization (ICP-style scan matching) + multi-rate KF; IMU orientation filter	GNSS-denied racing localization with LiDAR-map alignment
Kabzan et al. (2020)	Multi-sensor	Racing SLAM/state estimation integrated with full stack	End-to-end racing localization embedded in a full autonomy stack
Li et al. (2020)	Vision	Visual model-predictive localization	Predictive localization concept transferable to car racing stacks
Schratter (2023)	LiDAR (+ GPS baseline)	High-precision mapping + online map-based LiDAR localization	End-to-end racing mapping→localization workflow
Betz (2022)	—	Taxonomy of localization approaches used in racing	Helps structure method selection and gaps
Williams (2022)	GNSS + IMU (+ sim truth)	State estimation is framed as a core primitive for evaluation	Methodological framing for consistent racing evaluation
Sauerbeck (2022)	Camera + 3D LiDAR	Track-bound coordinates; modified OpenVSLAM (longitudinal) + LiDAR barrier relative-pose (lateral)	Decouple lateral/longitudinal to stabilize at IMS speeds
Karle (2023)	Multi-modal (LiDAR/radar/cam)	EKF-based fusion (tracking-focused) with racing-grade consistency handling	Transferable fusion techniques used alongside ego-localization
Betz (2023)	GPS(2) + LiDAR + IMUs(2)	Multi-source localization with redundancy + fusion/selection logic	System-level blueprint for redundant racing localization
Enisz (2024)	GNSS + onboard sensors	Bank of EKFs (kinematic vs dynamic) + supervisory switching	Robustness via multi-model EKF arbitration
Baumann (2024)	Typical F1TENTH (LiDAR/IMU/odometry)	Pre-mapped localization + odometry filter (system architecture)	Practical, reproducible racing localization module design

CHAPTER 4 PERCEPTION AND SENSING

"… advancing AI to build the world's most trusted driver." Dragomir Anguelov, Waymo.

4.1 Introduction

The advent of Automated Driving Systems (ADS) opened a new frontier for the automotive industry, offering a new era of autonomous mobility with potential benefits for occupant safety, mobility efficiency, economic productivity, and freight security. However, perception and sensing capabilities for autonomous driving have not yet fully met expectations, and this limitation has hindered user trust and slowed progression to higher levels of automation. As such, the detection and tracking of objects remain a formidable challenge for AI applications in perception (Mumuni & Mumuni, 2022). In high-speed complex and dynamic driving environments, object appearances and surrounding scenes change rapidly, thereby reducing the reliability of any single perception model for roadway object characterization.

At present, AVs use a portfolio of sensors, such as LiDAR, cameras, radar, and related devices (Figure 4.1) together with algorithms that process sensor data to develop, interpret, and update, in real time, a three-dimensional characterization of the surrounding environment. Radar sensors use radio waves to detect objects, and estimate features such as relative velocity, while cameras capture visual information from the driving environment. Using laser beams, LiDAR measures distances to objects, thereby creating a three-dimensional representation of the driving environment. Working together, these sensors provide a comprehensive representation of the surrounding road and traffic environment in real time. The autonomous system uses algorithms to analyze sensor data, identify objects, track their movements, and support navigation decisions by accounting for the positions, speeds, and directions of these objects.

Sensing is essential for reliable and efficient AV operation because is supports perception, situational awareness, and providing a basis for advanced decision-making capabilities. Together with vehicle-to-infrastructure (V2I) communication, the AV is empowered to fully characterize its environment. On autonomous racing tracks, most perception tasks are performed through onboard sensors, with V2I playing a relatively minor role. For example, Roborace's *DevBot 2.0* is equipped with multi-sensor suite of devices, including AI cameras, LiDARs, radars, ultrasonic sensors, and high-precision GNSS, all feeding data to an onboard NVIDIA AI computer for high-speed racing perception and decisions (Figure 4.2) (Scarborough, 2021; CSC, 2025).

This chapter presents a systematic review of literature on sensing and perception, followed by general lessons learned regarding autonomous driving sensing and perception, and finally, lessons specific to sensor ruggedization.

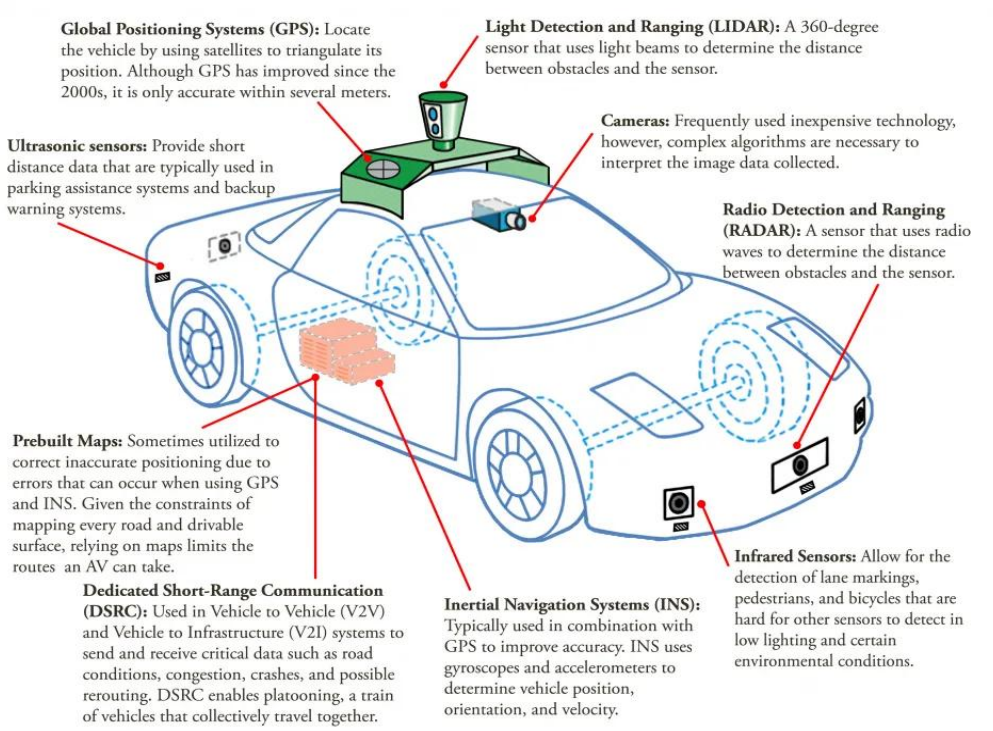

Figure 4.1 Basic sensing hardware of a typical AV (Image source: Center for Sustainable Systems, 2025)

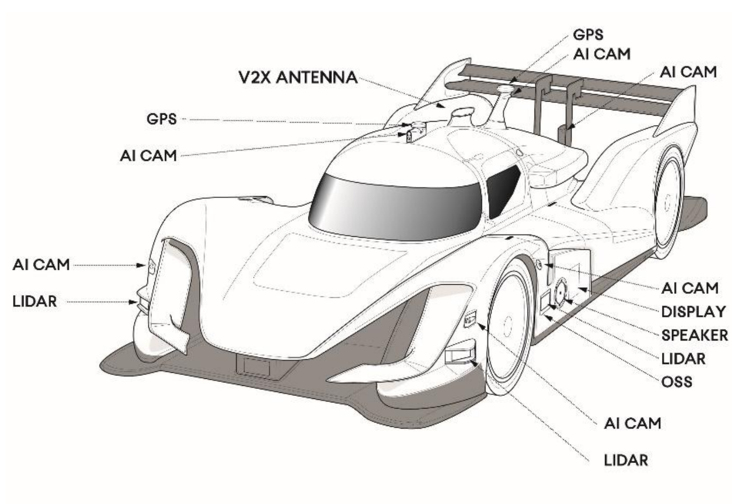

Figure 4.2 Sensing hardware of *DevBot* 2.0 autonomous vehicle (Scarborough, 2021)

4.2 Review of Relevant Perception and Sensing Literature

Pendleton et al. (2017) provided a foundational review of the entire AV software stack. In the perception section, they described the transition from raw sensor data to semantic understanding, discussing LiDAR and camera signal processing techniques, including occupancy grid mapping and feature extraction. Their paper explicitly established the fundamental "Sense-Plan-Act" architecture that most racing teams have since adapted to their platforms. The architecture signaled the transition from standard road-following modules to track-boundary estimation and opponent tracking algorithms specialized for high-speed interactions.

Autonomous racing data collection is expensive and risky. In acknowledging this limitation, Kang et al. (2019) surveyed datasets and virtual environments and highlighted the importance of high-fidelity simulators, including CARLA and AirSim, and datasets that include extreme weather or lighting environments. The paper argues that for racing perception, "closed-loop" testing in virtual environments is essential to validate how perception latency affects vehicle stability before deployment on a physical track.

In their survey of vision-based high-speed autonomous driving research, Kabzan et al. (2020) described the complete software stack of the AMZ Driverless Formula Student car, providing a valuable reference for the architecture of autonomous racing perception systems. The perception pipeline uses learning-based methods (camera) and complementary LiDAR processing to detect and classify track delimiters (cones). The outputs are fused into a consistent map/track representation that feeds localization and planning. The perception outputs are combined with vehicle-state estimation to maintain robust operation under the aggressive dynamics that characterize autonomous racing environments.

In what could be considered the first holistic survey of the autonomous racing field, Betz et al. (2022) covered platforms including Roborace, the Indy Autonomous Challenge, and Formula Student Driverless. The perception section of their survey categorized algorithms into mapping, localization, and object detection, and the authors noted a distinct shift from classical feature-based methods to deep learning approaches. The authors also discussed how racing environments differ from urban driving by requiring "long-tail" perception capabilities to handle extreme vehicle dynamics. They identified that while SLAM is standard, racing-specific SLAM must be tightly coupled with vehicle dynamics to prevent drift at high lateral accelerations.

Buckman et al. (2022) introduced the MiniCity platform to evaluate perception algorithms on 1/10th scale autonomous racing vehicles. While the setting is urban, the platform uses racing-grade hardware (LiDAR, camera) and high-speed reduced-scale dynamics for stress testing. Their study evaluated how perception error could cascade to downstream planning tasks. They found that in high-agility scenarios, the consistency of the perception update rate is often more important for stability compared to the absolute precision of single-frame detection.

Sauerbeck et al. (2022) and Sauerbeck et al. (2023) provided valuable lessons from their research on racetrack data. These articles presented comprehensive analyses of the perception pipeline used by the TUM Autonomous Motorsport team during a Indy Autonomous Challenge event. The authors described a multimodal sensor-fusion methodology that combines three LiDARs, six cameras, and four radars to achieve 360-degree coverage at speeds up to 174 mph. The primary algorithms include a custom LiDAR clustering method for object detection and a camera-based deep learning model for classification, fused via EKF. These studies, which were motivated by the need to maintain accurate perception and state awareness in racing environments characterized by severe motion dynamics and sensor failure modes, highlighted several critical lessons. For example, the authors suggested that at high speeds, the latency of perception algorithms is as critical as their accuracy. This suggests a trade-off between model complexity and accuracy: lighter but faster models, like varying YOLO iterations, and compact YOLO variants, typically outperform heavier, more accurate models in closed-loop racing scenarios.

Vödisch et al. (2022) used the FSOCO dataset (Formula Student Objects-in-Context) to enable perception during high-speed racing. The study developed a benchmark framework centered on Formula Student Driverless-style objects (notably cone-based track delimitation). The contribution of their paper lies in its standardization of training/evaluation data for perception models such as detection and segmentation algorithms. This facilitated reproducible development of racing perception pipelines.

Reveille Racing & Dietz (2023) documented the control and perception strategy used in the simulation phase of the IAC. The team used a PID-based controller fed by perfect ground-truth perception in simulation to establish a baseline for racing at high speeds (up to 195 mph). The study results serve as a useful benchmark for perception developers by quantifying the maximum allowable perception error and latency before the vehicle becomes unstable at very high speeds.

Betz et al. (2023) described the synthesis of a full-scale autonomous racing system, including the perception/estimation elements required for high-speed operation on structured tracks. Emphasizing that high-accuracy environmental representation is critical for racing, the authors discussed integration constraints, such as latency, robustness, and compute limit, that shape perception design decisions in motorsport conditions.

Zhuo et al. (2023) adapted and improved a YOLOv5s-style detector for the Formula Student Autonomous China setting, with modifications regarding small-object detection, precision/recall balance, and real-time performance. The validation results were strong, and the perception model produced a stable stream of cone detections for downstream track reconstruction and autonomous racing control.

Karle et al. (2023), in addressing multi-vehicle racing, developed a modular multi-sensor fusion and tracking pipeline that fuses heterogeneous detections and produces a consistent list of tracked objects for motion prediction and planning. The EKF-based core estimator included a delay-compensation mechanism to reduce the impact of perception latency. Designed to run efficiently on embedded compute, the model was validated at high speed through the Indy Autonomous Challenge.

Xia and Chen (2024) investigated AV trajectory planning, perceived collision risks, and the expectations of AV users. Their survey shed light on the intersection of planning and perception, particularly with respect to perceived risk of collision. In the context of high-speed driving and racing, they discussed how perception algorithms need to quantify uncertainty to assist in trajectory planning. The paper compared graph-search and sampling-based planning methods, noting that in dynamic racing scenarios, perception algorithms must provide not only object positions but also probabilistic velocity vectors, enabling the trajectory planner to generate safe, collision-free paths within the handling limits of the AV.

Szabo and Pup (2025) recognized that the physical configuration of a perception system influences algorithms performance and investigated this issue using Formula Student Driverless cars. Their study analyzed the placement of dual LiDAR sensors and stereo cameras to maximize the Field-of-View (FoV) and minimize blind spots caused by the vehicle's chassis. The authors demonstrated that optimal sensor placement is a prerequisite for robust perception algorithms, as even the most advanced cone-detection software performs poorly when sensor occlusion is not addressed. Their findings provided a geometric framework for initializing the perception stack in small-scale racing vehicles. Related work by Saka and Labi (2023) also used CARLA simulation tools and a multi-criteria approach to optimize sensor placement including the minimization of blind spots among other objectives. Lim and Park (2025), motivated by multi-vehicle racing (overtaking), investigated the challenge of detecting opponent vehicles at long ranges (over 328 ft), which is crucial for high-speed overtaking maneuvers. They proposed a novel two-stage algorithm that first uses a lightweight 2D Bird's Eye View (BEV) clustering method to reduce computational load compared to traditional 3D Euclidean clustering, followed by a machine learning classifier specifically trained on sparse, long-range LiDAR point clouds. The method demonstrated over 80% detection accuracy for vehicles beyond 164 ft., with a processing time of only 25ms, enabling real-time reaction at racing speeds.

Table 4.1 Synthesis of Sensing/Perception Algorithms in Autonomous Racing: Selected Studies

Paper	**Primary Sensor(s)**	**Algorithm/Approach**	**Key Application / Finding**
Sauerbeck et al., 2023	LiDAR, Camera, Radar	EKF fusion, Deep learning	Latency is as critical as accuracy; lighter models are generally superior.
Betz et al., 2022	Multiple sensor types	Survey of users of multiple sensor types	Racing requires specific "long-tail" handling and tight SLAM-control coupling.
Lim & Park, 2025	LiDAR	2D BEV Clustering + ML Classifier	80% acc @ >50m, 25ms latency; uses 2D clustering to save compute resources
Li et al., 2025	Camera	YOLOv13-Cone-Lite	Modified YOLO structure + removed NMS for 68Hz detection.
Szabo & Pup., 2025	LiDAR, Stereo Camera	Geometric analysis	Optimized placement is a prerequisite for algorithm success.
Kabzan et al., 2020	LiDAR, Camera	FastSLAM + YOLO	The "Standard Model" for fusing cone detection into a track map.
Pfeiffer et al., 2022	Camera	Attention prediction (imitation)	Mimicking human gaze reduces the compute load by focusing on gates.
Buckman et al., 2022	LiDAR, Stereo Camera	Comparison on a 1/10th scale	Perception update consistency matters more than single-frame precision.
Reveille Racing & Dietz, 2023	Simulation	PID (perception baseline)	Established max allowable perception latency/error for 300kph stability.
Xia & Chen, 2024	Multi-modal	Risk assessment / planning	Perception must output velocity vectors & risk probabilities, not just position.
Kang et al., 2019	Datasets	Validation frameworks	Emphasizes the importance of high-fidelity sim for validating perception algorithms.
Pendleton et al., 2017	LiDAR, Camera	Occupancy grids / signal processing	Establishes the core pipeline (Sense-Plan-Act) used in racing.

Li et al. (2025) developed a specialized object detection algorithm optimized for Formula Student Driverless competitions in which tracks are delineated to small traffic cones. The authors modified the YOLOv13 architecture by integrating a novel component to enhance feature extraction for small objects and removing the Non-Maximum Suppression (NMS) post-processing step to reduce latency. Their model achieved a frame rate of 68 Hz on embedded hardware (specifically, Jetson Orin NX) and improved Mean Average Precision (MAP) by 4.5% compared to the standard model. Their proposed algorithm helps achieve a balance between the high-speed requirements of racing and the detection of small, distant track markers.

Wang et al. (2025) proposed a lightweight deep detector tailored to racing cones (track boundaries), using a YOLO-family architecture optimized for accuracy/latency tradeoffs. The key contribution of their study was a robust cone detection suitable for race control loops, emphasizing "deployability" on resource-constrained platforms typical in student and prototype racing stacks.

Several journal articles have also shown how perception sensors can be used to support localization. These include the work of Massa et al. (2020), who developed a LiDAR-based localization pipeline for autonomous racing cars in GNSS-denied environments, and Schratter et al. (2023) who developed a practical LiDAR-based pipeline for map construction, map referencing, and accurate localization under high dynamic loads. Table 4.1 presents selected autonomous racing research studies on sensing and perception algorithms.

4.3 Some General Lessons Learned

One important lesson from autonomous racing is that advanced sensing and perception capabilities can overcome extremely demanding operating conditions, thereby driving innovation in AV perception for public-road environments. In these competitions, AVs must perceive their surroundings and react at speeds of roughly 90–180 mph (approximately 150–300 km/h), which exceed typical highway speeds (in contrast, freeway speeds generally range from 70–80 mph in the US, and slightly lower or higher in other countries, depending on the road class). At such high speeds on the racetrack, teams strive to avoid collisions between the AV and the track boundaries. In cases of head-to-head competitions (such as A2RL's multiple cars at a time and IAC's two cars at a time, as of the time of reporting), additional pressure falls on teams to refine their perception capabilities to the highest practical level possible, because collisions between competing cars could damage both vehicles. Because racecars are expensive to acquire and repair, teams are strongly incentivized to develop highly capable sensor suites and perception algorithms.

Racing AVs are typically equipped with multi-modal sensor arrays, such as HD cameras, LiDAR, radar, and precise GPS/GNSS, to detect obstacles and track vehicle position with very low latency. For example, Roborace's *DevBot* has 6 AI cameras, 5 LiDAR arrays, 2 radar arrays, 2 optical speed sensors, 18 ultrasonic sensors, and a high-precision GNSS, all feeding data to an NVIDIA Drive PX2 computer that plans the fastest line around the track (Stewart, 2017; Scarborough, 2021. Autonomous racing has shown that optimal sensor configuration, including number, types, and placement of sensors on the vehicle, is of great importance. Racing teams and organizers seek to maximize perception and minimize blind areas, and to also control costs, all under the demanding constraints of high-speed operation. Kapocsi et al. (2025) investigated optimal placement of LiDAR and stereo cameras on autonomous vehicles used in Formula Student racing. Saka and Labi (2023) and Hafemann et al. (2023)'s multiple-criteria frameworks help evaluate and optimize AV sensor coverage. Their approaches decomposed the road environment into discretized grid cells, estimated coverage, measured redundancy and blind spots, and developed quantitative metrics for sensor evaluation. Zhao et al. (2025) investigated an AV's aerodynamic penalties associated with sensors mounted externally on the AV. More generally, designers of production AVs can learn important lessons about sensor types and placement from the experience of autonomous racing engineers. However, because autonomous racecars differ markedly from production AVs in shape, size, and height, not all sensor-configuration lessons from the former will transfer directly to the latter.

The success of high-speed perception hardware and algorithms in IAC events has helped validate sensor-fusion algorithms and object-detection systems for production AVs that must maintain accuracy at high speeds. For example, 4D imaging radar units with ranges of about 300 m were tested in IAC racecars and were later deployed on production vehicles (Chudzinski, 2023).

4.4 Lessons Related to Sensor Ruggedization

Sensor suites on autonomous racecars are subjected to severe environmental stressors, including vibration, heat, and g-forces. For example, in the 2022 Indy Autonomous Challenge, sensors were subjected to sustained speeds exceeding 170 mph and intense vibration profiles arising from the internal combustion engine and rough track surfaces (New Eagle, 2022). Therefore, autonomous racing offers an opportunity to assess any degradation in sensor performance due to these stressors. This provides insight into the design of sensor hardware for production AVs, particularly with respect to longevity and reliability. Indeed, IAC stated on its website that the transition of IAC's sensor hardware from the 2021 to the 2023 version reflected the organization's commitment to "cultivating the technology pipeline between autonomous race cars and autonomous consumer vehicles" (IAC, 2022). With respect to vibration and sensor durability, early generations of autonomous hardware reportedly failed because connectors vibrated loose or internal measurement units became desynchronized. Vendors have since hardened their sensor units against these extreme conditions, and many of these designs are likely more robust than would normally be required for typically suburban-road vibration profiles. With respect to motion blur and rolling-shutter effects at high speed, autonomous racing engineers must address the high angular velocities that can distort camera feeds. Even a brief millisecond of distortion could lead to vehicle or vulnerable-road-user (VRU) misidentification. To mitigate this issue, racing teams have developed advanced deblurring and denoising algorithms and have deployed global-shutter cameras which capture the entire image simultaneously rather than line-by-line. According to the literature, these technologies are increasingly appearing in high-end automotive ADAS packages to improve clear detection of cross-traffic features at highway speeds.

CHAPTER 5 DATA FUSION

"All knowledge begins with the senses, proceeds then to the understanding, and ends with reason."
— Immanuel Kant.

5.1 Introduction

Data fusion is a critical aspect of autonomous vehicles. On a typical AV, there are not only several types of sensors (cameras, LiDAR, radar, and ultrasonic devices) but also multiple units of each type. The use of multiple sensor types helps compensate for the limitations of any single type, while multiple units of the same type help reduce blind-spot areas. The goal is to create a unified and reliable 360-degree representation of the AV's surroundings. Accordingly, data fusion improves perception reliability, reduces uncertainty, and enhances safety by combining complementary and, in some cases, even redundant data to overcome individual sensor limitations.

The benefits of fusion include enhanced safety, improved object detection and classification, improved localization, and more reliable decision-making for navigation-related operations. The challenges of data fusion include high computational demands, algorithm complexity, synchronization difficulties, and the high costs associated with fusion hardware.

Cameras provide high-resolution color and semantic information (in the real-world driving environment, for example, traffic lights and road signs); LiDAR offers precise 3D depth and distance measurements; and radar excels at estimating speed and operating under poor weather conditions. Sensor-data fusion methods may be categorized into three levels : (a) low-level fusion, which combines raw sensor data or signals and tends to offer high accuracy at the cost of greater computational demand; (b) mid-level fusion, which combines features extracted from different sensor before object recognition; and (c) high-level fusion, which combines independent object detections or decisions from different sensors to confirm, track, and support decisions from different sensors to confirm, track, and support decisions while reducing errors, particularly, false positives.

5.2 A Review of Past Work

Multiple sensors generate streams of different data types and formats containing a rich variety of information. At high speeds, often exceeding 100mph, the challenge is to combine these heterogeneous data streams quickly enough to support safe navigation decisions. In this regard, several autonomous racing teams have published studies based on their track experiences and technical contributions. This section reviews selected peer-reviewed studies on the critical issue of data fusion, particularly, the integration sensor data for perception and state estimation, in high-speed autonomous racing.

Patz et al. (2008) described a Unified Planning and Control architecture in which perception data were fused to create a drivability map rather than merely a geometric map, thereby facilitating aggressive maneuvering in complex environments. They used Particle Filter for localization and Motion Primitive Generation for planning. This work drew on the experience of the MIT team (Talos) in the DARPA Urban Challenge and helped establish foundations later relevant to autonomous racing.

Kabzan et al. (2020) described a complete system architecture spanning perception, planning, and control. The paper also made important contributions from the perspective of data fusion. Using the Formula Student Driverless platform, they integrated multiple onboard sensors support robust state estimation and environmental comprehension. The study showed how fusion methods translate into reliable performance in a competitive race environment.

Massa et al. (2020) provided a LiDAR-centric approach for racing localization in track environments where GNSS is unavailable or unreliable. Their model, designed for high-speed racing conditions in the absence of GNSS, used a hybrid probabilistic fusion method in which an EKF for continuous state estimation was combined with particle-filter/Monte Carlo localization against a LiDAR map informed by track priors.

Peng et al. (2021) developed methods for fusing information from three sensor types, Camera, LiDAR, and IMU, to address challenges in vehicle odometry on a Formula Student Driverless platform. Their algorithm involved multi-sensor odometry and state estimation, and the authors evaluated the extent to which multiple sensors improved system robustness compared to single sensor odometry. They argued that this capability is critical in autonomous racing conditions where lighting, motion blur, and aggressive dynamics could impair the efficacy of monocular or LiDAR-only pipelines.

Betz et al. (2023) described an innovative end-to-end autonomous racing stack. The authors described how multiple sensing types (GNSS, IMU, LiDAR, and radar) were combined to support localization, perception, and state estimation, and ultimately, to enable racing decision-making under tight time constraints in real time. In their architecture, data fusion was centered primarily on filter-based state estimation within a modular autonomy stack designed for high-speed operation. Karle et al. (2023) proposed a system for robust, low-latency fusion suitable for high speeds and aggressive maneuvers. Their approach used a multi-modal object-tracking stack that fused detections from multiple sensors to support motion prediction and trajectory planning in autonomous racing. Their system uses EKF-based multi-target tracking with data association, with explicit attention to real-time constraints and racing-relevant failure modes, including latency, missed detections, and mis-associations.

Lee et al. (2024) presented a state-estimation framework designed for autonomous racing. Also, by combining multiple sensor types, their architecture was designed to manage tradeoffs between compute capability and latency. Based on Kalman filtering with multi-rate and multi-sensor fusion, the estimator was designed to promote robustness in a highly dynamic environment. A key contribution of their paper was the systemic co-design of the estimator fidelity and resource scheduling; that way, the stack remains stable and reliable even at high speeds.

Suprapto (2024) integrated object detection with path planning for electric autonomous vehicles. This study demonstrated a fusion framework in which camera data processed by YOLO and distance sensor data inputs were combined to dynamically replan routes around obstacles in real-time, a critical capability for avoiding collisions on a track.

Cellina et al. (2025) used latency-aware EKF with reprocessing (for delay compensation) in an innovative sensor data fusion system and a-priori knowledge of the racetrack to track opponent racecars at speeds up to 171 mph. Their model addressed the specific challenge of out-of-sequence measurements caused by high-speed motion and sensor latency, thereby supporting precise and reliable overtaking.

Amaksoud and Ahmed (2025) reviewed state-of-the-art Transformer-based deep learning models for fusing LiDAR, camera, and radar data. The authors highlighted the use of transformer-based deep fusion techniques, showing how attention mechanisms enable the system to weigh the relative importance of inputs from different sensor types in a dynamic manner. This capability is especially important in competitive racing, where motion occurs at high speed, and in unstructured environments where static rules often perform poorly.

In addition, there were other autonomous racing research studies were not primarily focused on sensor fusion but nevertheless made valuable contributions to fusion-related thinking or developed approaches that highly depend on reliable fusion outputs. Heilmeier et al. (2020c) study focused on minimum curvature trajectory planning and control for AV race cars. Although the paper's key contributions were in the modules of planning and control, it made valuable contributions to data fusion because its racing pipeline operationalizes map and state inputs that had to be produced reliably by upstream estimation at racing speeds. Montani et al. (2021) also proposed a hierarchical architecture for racing autonomy in which real-time planning and tracking were layered to manage computational burden and stability. These hierarchical decision layers require consistent and timely state estimates and track representations, typically supported by filter-based estimation and synchronized sensing. Table 5.1 presents selected research articles on sensor data fusion.

Table 5.1 Selected research articles related to sensor data fusion.

Source	**Fusion-related Approach**
Karle et al. (2023)	Multi-modal fusion + multi-target tracking (KF/EKF, data association)
Lee et al. (2024)	Multi-sensor state estimator (Kalman filtering, multi-rate fusion) + compute scheduling
Betz et al. (2023)	Integrated localization/state estimation (filter-based fusion)
Kabzan et al. (2020)	Competition-validated estimation + perception integration (filter-centric)
Massa et al. (2020)	Multi-rate EKF + Monte-Carlo localization (MCL/particle filter)
Heilmeier et al. (2020)	Planning via QP; relies on stable state/map inputs
Montani et al. (2021)	Hierarchical real-time architecture; upstream filter-based estimation assumed
Peng et al. (2021)	Camera–LiDAR–IMU odometry/information fusion

5.3 Discussion: Sensor Fusion in Denied Environments

Autonomous racing teams must contend with intermittent sensor-data loss and degraded signal availability at sections of a track and under specific operating conditions. The nature of such denial or degradation of operating conditions is influenced by the sensor type in question. For GPS/GNSS data, such degradation may result from occlusion due to the presence of physical obstructions such as buildings or trees, track geometry such as banking, concrete walls, or fences, atmospheric interference (ionospheric conditions), electronic interference including jamming/spoofing devices, nearby transmitters, cyber-attacks, solar activity (flares on the sun's surface), hardware issues such as antenna/cable failure, and receiver limitations. All these tend to weaken or corrupt the signals needed for accurate positioning. Consequently, racing teams have found that reliance on a single sensor type can be highly risky in a racing environment. With respect to fusion for localization, racing teams have developed robust localization stacks that fuse LiDAR matching (SLAM), visual odometry, and Inertial Measurement Units (IMUs). For example, TUM's Autonomous Motorsport team validated a fusion approach that prioritizes different sensors based on real-time confidence scores. This approach is directly applicable to commercial CAVs navigating ghost zones that include "urban canyons" (where tall buildings block GPS signals). The ability to localize using only landmarks such as track walls or city buildings helps ensure continuity of operation. Sauerbeck et al. (2023) further argued that two specific sensor types may be sufficient for limited ODDs, but that broader sensor fusion is likely inevitable for future Level 5 autonomous mobility.

With respect to adversarial perception- that is, estimation of an opponent vehicle's state, including velocity, heading, and acceleration- it is critical to predict the paths of competing vehicles accurately and continuously on the track. In real world of autonomous mobility, a similar capability is essential because erratic drivers or motorcycles weaving through traffic must be tracked reliably. Racing teams have found that Extended Kalman Filter (EKF) and Particle Filter techniques can be effective for tracking opposing racecars moving at very high speed, often above 150+ mph (Sauerbeck et al., 2023).

CHAPTER 6 PATH PLANNING AND DECISION-MAKING

"Autonomous racing is rapidly becoming the proving ground for pitting self-driving car AI systems against each other, while steadily advancing the state of the art in perception, planning, and control." Madhur Behl, Principal, Cavalier Racing Team, University of Virginia.

6.1 Introduction

Path planning and decision making in autonomous systems are two intertwined processes that enable an autonomous vehicle to navigate from a start point to an end point efficiently and safely. Decision-making serves as the high-level control layer to determine **what action to take** (e.g., merging, stopping, changing lanes) based on real-time data, environmental conditions, and overall mission objectives. Planning (or motion planning) establishes **how to execute that action,** generating a precise, safe trajectory for the actuators (and the control system executes the action through the actuators). Some researchers consider path or motion planning to be a subset of decision-making. AV decision-making is characterized by, (i) data-driven decision processes based on real-time information, traffic conditions, safety considerations, and the predicted behavior of other agents, such as AVs and pedestrians on the roadway, and (ii) adaptability to dynamic and uncertain environments. Chen et al. (2020) offered the following classification of the levels of decision making (LODD) in AVs:

- Strategic: relatively long-term high-level driving decisions or tasks, including route choice in a road network (LODD I).
- Tactical: relatively mid-term mid-level driving decisions or tasks, including whether to overtake another vehicle (LODD II).
- Operational: relatively short-term low-level real-time driving decisions/tasks (LODD III).

Strategic driving decisions (and tasks) are high-level, general trip-planning choices made before or during a journey, such as choosing a route, deciding on a mode of transport, or adapting plans due to a crash or bad weather. These are more applicable in real-world driving and generally not applicable in autonomous racing. *Tactical driving decisions* (and tasks) involve executing maneuvers in response to immediate road conditions to achieve a goal, such as changing lanes, maintaining a safe following distance, obeying traffic signals, and avoiding obstacles. Many of these are applicable in head-to-head autonomous racing, while a few are applicable in single-car racing. *Operational driving decisions* and tasks are those that, in the context of human driving, are habitual and require minimal conscious thought during normal driving conditions, including immediate risk assessment, turning the steer to maintain lane position or change direction smoothly; pressing the accelerator pedal to increase speed; pressing the brake to reduce speed or to stop the vehicle; activating ancillary controls such as defrosters, wipers, turn signals, or headlights. LODD I generally involve a sequence of interdependent decisions, whereas LODD II typically involves single-step decision. It is worth noting, however, that there may be variations in the literature regarding what is described as strategic, tactical, and operational; for example, what some articles may consider as operational, other papers may consider as strategic.

Path planning, which may be considered part of strategic level of decision-making (LODD I), can generally be described as the process of identifying feasible routes and the optimal route from a given origin to a given destination, while avoiding obstacles and not violating specific constraints. As such, the key aspects of path planning are (a) path optimality (the best path is chosen based on shortest distance or time, or lowest energy consumption, for example), and (b) collision avoidance (ensuring that obstacles such as other AVs and fixed obstacles on or by the guideway are avoided). Path planning may be global (often executed using a comprehensive pre-existing map of the environment to compute an optimal route beforehand) or local (using real-time sensors to adjust in a dynamic environment and navigate sudden unforeseen obstacles (e.g., swerving to avoid an obstacle on the roadway).

The computational approaches for AV planning and decision making are (Yuan et al., 2023): *Rule-based models* (using pre-defined rules and logic, but limited in complex scenarios); *AI-based models* (using reinforcement learning and deep learning to learn complex behaviors and trajectories directly from data) (past experiences), to improve adaptability and efficiency in challenging environments); and *Hybrid models* (combining the strengths of traditional algorithms with AI for robust, adaptive solutions).

6.2 Review of Literature

Liniger and Lygeros (2020) addressed head-to-head racing, framing tactical decisions as a noncooperative, non-zero-sum game with collision-avoidance constraints and progress-based payoffs. By constructing multiple game formulations (e.g., follower-only collision handling vs. both players constrained; payoffs that encourage blocking/defending), their approach established equilibria (including Stackelberg/Nash variants) in a moving horizon manner. In the context of path planning and navigation, the contribution of this article is the explicit treatment of interactive multi-agent behavior (overtake/defend) as a structured planning problem, instead of heuristic rules layered on top of a single-vehicle racing line. Table 6.1 presents a synthesis of the selected planning and decision-making algorithms discussed in this section.

Kabzan et al. (2019) integrated a nominal vehicle model with a Gaussian Process (GP) regression model that learns the residual model mismatch (error) in real time to address trajectory prediction fidelity under high nonlinearity (tire/handling constraints) using this refined model within MPC optimization to plan trajectories, pushing the vehicle close to its physical limits within handling constraints. The authors validated their approach using a full-size Formula Student driverless vehicle, achieving a 10% reduction in lap times compared to a standard MPC by effectively adapting, in real time, to track conditions and tire dynamics.

In a subsequent study that similarly used GP, Hewing et al. (2018) designed a "cautious" MPC that used GP regression not only to estimate model deviation, but also to quantify the uncertainty associated with that deviation. The controller propagates such uncertainty throughout the prediction horizon and generates bifurcated chance-constrained trajectories: conservative when confidence is low, and aggressive when confidence is high. Such a flexible approach helps balance high-performance racing behavior with robust safety guarantees, ensuring the vehicle keeps within friction limits in uncertain environments.

Alcala et al. (2020a) combined an offline time-optimal trajectory planner with an online LPV-MPC tracking/controller that accommodates constraints and remains computationally feasible. Their contribution towards the planning field is the innovative LPV formulation that is considered suitable for aggressive driving, where full nonlinear MPC may be too expensive. In a subsequent paper, the Alcala et al. (2020b) extended the LPV concept into LPV-based motion/trajectory planning (LPV-MPP) that explicitly accounts for static obstacles by adapting lateral bounds (track corridor) as a function of obstacle position. The optimization was formulated as a quadratic program that maximizes velocity over a horizon, subject to time-varying corridor constraints derived from obstacle geometry. Their study is an example of trajectory planning (time-parameterized) instead of geometric-only path planning.

Rosolia and Borrelli (2020) considered each autonomous racing lap as a minimum-time iterative task and introduced a novel control algorithm that improved performance monotonically by storing prior safe trajectories and using them to update terminal constraints/costs for the next lap. In their model, the path and trajectory planning component is embedded inside the control algorithm. The controller iteratively refines a more efficient racing policy by learning from previously feasible racing trajectories rather than relying solely on a static model.

Heilmeier et al. (2020c)'s planning-and-control stack established a minimum-curvature trajectory using quadratic optimization (QP) and then tracked it near handling limits with a controller designed for high lateral accelerations. The planning component is a raceline method that minimizes curvature and enables high-speed cornering. In the literature on autonomous racing navigation, this paper has been widely cited for its contribution to QP-based racing line generation (coupled with limit-handling tracking) that emphasizes reliable performance when the vehicle operates close to physical limits.

Wang et al. (2021) used competitive simulation and real-world (1:10 scale) racing to develop a game-theoretic model that interactively plans vehicle movements, with opposing racecars actively responding to the ego vehicle's movements. The authors proposed the racing problem as a non-cooperative game, seeking a Nash equilibrium trajectory. The model incorporates a sensitivity term in the optimization that the ego vehicle uses to predict an opposing vehicle's response to avoid a collision. The model uses an iterative best-response scheme coupled with Model Predictive Control (MPC) to generate aggressive yet safe maneuvers, including overtaking and blocking. Their model was demonstrated to significantly outperform passive MPC baselines.

Wurman et al. (2022) developed a new algorithm (the Quantile-Regression Soft Actor-Critic (QR-SAC)) that highlighted the potential of model-free RL for complex, high-speed, multi-agent racing scenarios using a high-fidelity simulator. The algorithm, integrated with a specialized reward function and a mixed-scenario training curriculum, demonstrated performance surpassing top human racing drivers in head-to-head competition. In contrast to traditional path planners that use explicit geometric rules, QR-SAC learned high-level tactical behaviors, including blocking, drafting, and opportunistic overtaking.

Garlick & Bradley (2022) proposed a data-driven alternative to computationally expensive optimal control methods for establishing the best racing line. Their approach focuses on real-time racing line prediction using machine learning, trained on datasets of racing lines generated by traditional optimal control methods. Instead of solving a heavy optimal control problem on board, a feed-forward neural network predicts a near-optimal racing line quickly enough for real-time use. Their approach illustrates the concept of replacing expensive optimization-based planning with a learned surrogate that supports fast replanning across varied tracks. The real-time generation of near-optimal racing lines is critical for on-board path planning where computational resources are limited and rapid replanning is essential.

Rowold et al. (2022) proposed local planning for oval-track racing using efficient spatiotemporal graph search, where candidate maneuvers are evaluated in a time-augmented state space to handle dynamic obstacles (other cars) and produce collision-free yet competitive trajectories. Unlike standard methods for trajectory search, their approach is optimized for the specific constraints of oval racing, such as maintaining high speeds and banking management. The algorithm first generates a coarse trajectory that guides the vehicle around dynamic obstacles, then smooths the trajectory for execution, thereby providing a robust alternative to traditional purely optimization-based planners in multi-car racing environments. Thus, the approach mitigates complexities associated with spatiotemporal search methods using an interval-based graph representation.

Evans et al. (2023b) compared alternative deep reinforcement learning architectures in the F1Tenth context: (a) “full planning” (replacing global+local planner), (b) trajectory tracking (replacing the local planner), and (c) end-to-end planning (replacing the entire pipeline). The paper’s key contribution is its elucidation of the best opportunities for learning within the stack. It also sheds light on shifts in performance/robustness depending on whether the learned policy is responsible for planning decisions versus merely tracking a pre-planned trajectory.

Thakkar et al. (2024), in recognition of the complexity posed by multi-agent racing organizational rules and team strategies, introduced a bi-layer hierarchical control framework: first, a high-level game-theoretic Tactical Planner, and second, a low-level optimization-based path planner. The first layer solves a simplified discrete game to make strategic decisions (for example, overtake left, defend position) while ensuring compliance with complex racing rules (for example, blocking restrictions). The second layer executes these decisions using trajectory optimization. The layer separation technique allows the system to handle the long horizons required for strategy within constraints. The key contribution is the explicit operationalization of rules, coordination, and adversarial interaction in planning, and making it possible to enable the ego vehicle’s decisions, including coordinated overtaking and defense.

Li et al. (2024) focus on the fundamental problem of minimizing lap time purely using optimization. The authors generated a global time-optimal trajectory that accounts for complex vehicle dynamics and tire physics, and coupled it with a hierarchical tracking controller that adjusts for real-time errors. In highlighting the trade-off between the computational cost (of optimal lap generation) and the requirements associated with real-time tracking, the authors demonstrated effective performance in handling the vehicle within handling constraints.

Zhang et al. (2025) proposed a global minimum-time trajectory planning approach for autonomous racing. Their approach uses high-precision offline mapping, explicitly accounting for curvature and distance. Their offline/global planner represented a departure from local reactive planning and targeted a time-optimal trajectory consistent with the track geometry.

Toaz and Bopardikar (2025), proposed a Vector Cost alternative to standard game-theoretic weighting. Instead of scalarizing and weighing the multiple criteria (for example, speed, safety, and jerk) into a single number for each alternative, the authors treat them as a vector in a bi-matrix game. This facilitates the identification of the Pareto-optimal Nash equilibrium that prevents “worst-case” outcomes (such as crashes), which often occur in scalarized games when weights are poorly tuned. Applied to racing, this method allows an autonomous agent to aggressively overtake opponent racecars while guaranteeing a certain level of safety, providing a good balance between the tuning extremes of timidity and recklessness faced by standard game-theoretic planners.

Table 6.1 Synthesis of Selected Literature on Planning

Paper's Author & Year	**Algorithm Category**	**Specific Algorithm Used**	**Key Contribution / Feature**
Wang et al. (2021)	Game Theory	Game-Theoretic Planner + MPC	Modeled opponent interaction (Nash Equilibrium) to enable "assertive" overtaking and blocking.
Garlick & Bradley (2022)	Machine Learning	Neural Network (Supervised)	Replaces slow offline optimization with a fast Neural Network to predict optimal racing lines.
Thakkar et al. (2024)	Hybrid / Hierarchical	Hierarchical Game Theory	Separates strategic logic (discrete game) from path execution (continuous optimization) for rule compliance.
Rowold et al. (2022)	Graph Search	Spatiotemporal Graph Search	Specialized graph search for high-speed oval overtaking; handles banking and dynamic obstacles.
Li et al. (2024)	Optimization	Direct Collocation + Hierarchical Tracking	Strictly time-optimal global planning coupled with a fast local tracker for limit-handling.
Remonda et al. (2021)	Analysis	Deep RL vs. Human	Comparative study showing RL exploits physics better but lacks human-like smoothness/predictability.
Toaz & Bopardikar (2025)	Game Theory	Vector Cost Bimatrix Game	Solves multi-objective racing games without scalar weights, preventing "worst-case" crash scenarios.

6.3 Opportunities Made Available by Autonomous Racing

Rule-based models use a library of rules developed considering experience, knowledge, and traffic rules (Park & Kee, 2021; Yuan et al., 2018; Zeng et al., 2021). The major limitation of rule-based models is their weak generalization capabilities, stemming from their inability to account for all driving scenarios and their poor performance in unfamiliar driving conditions where such rules do not apply (Yuan et al., 2023). AI-based approaches use learning techniques, including reinforcement learning or deep neural networks, trained to optimize decision-making at any level using past data (Fehér et al., 2020; Biswas et al., 2021; Wang et al., 2020). However, like most AI models, these approaches have some major limitations. First, they need a large amount of training data, which is typically expensive to assemble. Specifically, very sparse data are available on real-world autonomous driving under different conditions, particularly edge cases (obviously due to passenger safety concerns). As such, autonomous vehicle racing events, which are not constrained by regulatory boundaries, provide a unique opportunity to produce adequate training data that is needed to train the AI models. Secondly, existing AI-based approaches are generally uninterpretable, which impairs their efficacy in unaccustomed or hypothetical scenarios (Yuan et al., 2023). Again, autonomous racing produces data that could potentially address this issue.

In head-to-head autonomous racing, the strategic level of decision-making can be considered the same across all participating vehicles. However, regarding tactical and operational decisions (such as how to optimize a racing line, when to overtake an opponent, or how to react if a car ahead suddenly swerves into the ego AV's path), each team has its decision-making model. As summarized in Table 6.1 and discussed in Chapter 2, there is evidence in the literature that significant advancements have been achieved in autonomous racing path planning and decision making. At the current time, autonomous racing AI models show great potential to fuse and process sensor data, making tactical decisions such as overtaking, reacting to the movement of competing vehicles, and handling adversarial situations in real time (Roles, 2025). The planning and decision-making models apply game theory and reinforcement learning. Production AVs can apply these models to resolve challenges they face on public roads, such as merging into traffic or avoiding an errant driver – albeit at lower speeds. Therefore, by mastering overtaking on a racetrack and managing interactions with other vehicles, autonomous vehicles can learn strategies for negotiating that make conventional autonomous driving safer and more efficient.

Regarding motion prediction of other vehicles (key prerequisite for tactical decision-making), autonomous racing competitions have yielded algorithms that predict the movements of neighboring vehicles and adjust the ego vehicle's path to avoid collisions or exploit movement opportunities. Such models are directly applicable to advanced driver-assistance systems. In his online article, Roles (2025) opined that the AI systems used by autonomous racecars can help predict the movements of competing vehicles on the racetrack. This has clear real-world applications for predicting the behavior of not only often unpredictable HDVs but also other AVs in the immediate vicinity of the ego AV.

6.4 Game Theory and Trajectory Planning

Game theory is the study of interactions between (or among) rational decision-makers, analyzing conflict and cooperation in situations where outcomes depend on multiple choices (Myerson, 1991). According to researchers (Adler et al., 2021; Li, 2022), the application of game theory in transportation goes back to the mid-20th century, with significant contributions from several key individuals: Frank Knight (who provided an early conceptualization of traffic equilibrium), John von Neumann and Oskar Morgenstern (whose work laid the mathematical foundation for game theory, making it applicable to various fields including economics and, later, transportation), John F. Nash Jr. (whose Nash equilibrium concept was applied for non-cooperative games that are widely used in transportation analysis, particularly in route planning models where no driver can unilaterally improve their travel time by changing routes), J.G. Wardrop (who introduced his principles of traffic equilibrium in 1952, which are considered equivalent to the Nash equilibrium under certain conditions and form the basis of many traffic assignment models), and Martin Beckmann, C.B. McGuire, and Christopher Winsten (who applied equilibrium principles to a comprehensive transportation network model, a foundational work that effectively marks the first systematic application of game theory concepts to transportation problems).

In the field of autonomous transportation, game theory began to be systematically applied in the early 2010s, evolving from basic traffic interactions to complex scenarios involving human-like decision-making and legal frameworks, with many researchers contributing to various aspects of this domain. Yoo and Langari (2012) used Stackelberg games to model lane-changing/merging; Van Arem and his team introduced differential game models for car-following and lane-changing (Wang et al., 2015); Charles Fox and his team at the University of Leeds developed models for priority negotiation at intersections between AVs and pedestrians/vehicles (Camara et al., 2018) and other applications (Camara & Fox, 2022). Christian Gerdes and his team applied game-theoretic planning for AVs in competitive (racing) scenarios (Wang et al., 2019). In 2023, Dejan Milutinovic (University of California, Santa Cruz) solved a decades-old game theory problem (the "wall pursuit" game) that advanced decision-making for autonomous systems, including driverless vehicles (Cerf, 2023).

The key components of game theory, in the context of this report, are: the players, or decision-makers, (the autonomous vehicles that need to make strategic, tactical, or operational decisions), actions (the possible actions each vehicle could undertake), payoffs (the beneficial or adverse outcomes a vehicle receives in response to a combination of actions). Additional key features of game theory include rationality (the assumption that each vehicle acts in its own self-interest to maximize its payoff) and equilibrium (a stable state where no vehicle has an incentive to change its action, given how other vehicles are acting, for example, the Nash Equilibrium. Other applications of game theory in autonomous systems were discussed by Dennis et al. (2022), Hang et al. (2020), Fisac et al. (2019), Fox et al. (2018), Smirnov et al. (2021), and Wang et al. (2020). As learning-based approaches and multi-car racing become more prevalent in autonomous racing, it is expected that AI and game theoretic approaches will be used increasingly.

In traditional autonomous driving systems, obstacles are often treated as either static (e.g., a large rock in the middle of the road) or moving along fixed paths (e.g., a neighboring vehicle). In autonomous racing, an obstacle may be another autonomous racecar actively trying to block or overtake the ego car. These scenarios call for game-theoretic planning. Racing algorithms often model such scenarios as a Stackelberg game, where the ego vehicle (leader) makes a move, anticipating the follower's optimal response. This hierarchical optimization allows the car to plan a trajectory that influences the other driver's behavior. Researchers have shown that Stackelberg-based merging algorithms significantly outperform rule-based systems in dense traffic by reducing the likelihood of "freezing robot" situations (where the AV is too "polite" because it is risk-averse and thus unable to merge into traffic) (Du et al., 2024). Also, safe overtaking on a two-lane road (one lane in each direction) requires the ego vehicle to estimate whether the oncoming car is too close and if the car being overtaken will accelerate. Even though racing occurs on one-direction tracks, racing algorithms continuously calculate time-to-collision and time-to-overtake (Kalaria et al., 2024), thereby providing logic for safely executing these maneuvers. This could be applied in AV operations on public roads.

6.5 Discussion

In this chapter, we discussed how path planning and decision-making continue to play essential roles in autonomous racing and will be critical components of ADS for real-world driving on public roads. In autonomous racing literature, several planning and decision-making algorithms have been developed (many of them developed through researchers' experiences on the racetrack during competitive events). At these races, teams identify optimal ways to achieve track objectives, make real-time choices to navigate complex, dynamic driving environments, adapt to obstacles, race rules, and new data, and often, blend traditional methods such as the Dijkstra algorithm with advanced AI-like reinforcement learning for safe, efficient movement of the autonomous system. AI-based approaches use learning techniques, including reinforcement learning and deep neural networks, to optimize decision-making across multiple levels of historical data. However, AI models require enormous volumes of training data whose acquisition is typically expensive. In real-world autonomous driving, there is limited data across diverse driving conditions, particularly, edge cases. This is not surprising, obviously, given passenger safety and freight security concerns, as well as legal/ethical frameworks. As such, autonomous racing events, which are largely unconstrained by such boundaries, provide a unique opportunity to generate the training data needed to train AI models for safe and efficient path planning and navigation of the autonomous system.

CHAPTER 7 CONTROLS

"You can't always control the wind, but you can control your sails." – Bob Chope.

7.1 Introduction

Control generally refers to the power or authority to guide, regulate, or manage a system. In the world of HDVs, control is carried out by the human driver. Autonomous control systems represent an evolution of conventional control systems, with the addition of intelligent components. Thus, in the context of AVs, control can be defined as the mechanism by which an autonomous vehicle uses its intelligent hardware and software to carry out independent management, direction, and regulation of its behavior, actions, and decision-making processes without continuous human intervention.

One of the most direct forms of such technology transfer lies in the domain of control. In the current era, most newly produced human-driven vehicles (and a significant proportion of vehicles in the extant traffic stream) possess some level of ADS capabilities. In vehicles equipped with ADS, failure may occur when the vehicle exceeds the linear region of tire dynamics, for example when the tires begin to lock, slide, or lose grip. Common modes of ADS failure include: **Lane keeping failure** (the ADS may fail to keep the car centered, jerk the wheel, or disengage entirely during aggressive maneuvers); **Adaptive cruise control failure** (the vehicle may brake or accelerate unevenly, or fail to maintain safe spacing/distance when the leading car changes lanes abruptly); **Automatic emergency braking failure** (the ADS might not be able to stop the vehicle in time).

There are at least 4 reasons for such behavior under boundary or limit-handling conditions:
(a) Control System Lag: At the physical limits of the tires, the ADS's ability to react, compute, and command steering or braking may be unable to keep up with the rapid, nonlinear changes associated with those conditions.
(b) Tire Slip and Grip Limits: During the negotiating of sharp curves or sudden braking, the tires enter a nonlinear "slip angle" region where tire grip diminishes. This causes unpredictable vehicle behavior that standard ADS struggles to predict or control effectively.
(c) Limitations of Traditional Models: ADS often uses basic, linear models of vehicle response (for example, how much steering input equals how much turn). These models break down when the vehicle tires slip, lock up, or experience extreme G-forces.
(d) Sensor Limitations: Cameras, radar, and LiDAR perceive the surrounding environment. However, if a wheel is misaligned (due to, for example, bumps or wear), the resulting vehicle-state inconsistency may cause the system to misinterpret the road conditions or vehicle position.

Evidence from the autonomous racing literature suggests that racing control algorithms can perform well in this non-linear region of tire dynamics. Racing pushes vehicles to physical extremes; hard braking, rapid acceleration, and cornering at or beyond the traction limit, and in doing so, enables autonomous racing teams to develop valuable techniques for handling emergencies and avoiding crashes. As such, the teams develop their algorithms to handle the forces efficiently at the limits of adhesion. These algorithms are tuned to maintain control of the vehicle even near physical limits (an effort that is possible only in a racetrack environment). In fact, during such races, the competing autonomous racecars routinely drive at the edge of tire friction limits.

This has spurred research into advanced control techniques, including model predictive control and advanced stability-control methods, that can keep a vehicle stable even when skidding or drifting. For example, inspired by rally racing and drifting at racetracks, Toyota Research Institute's Nonlinear Model Predictive Control (NMPC) model, which combines vehicle dynamics and control design, intentionally induces and controls drift to transition smoothly between normal driving and controlled slides when faced with road surface hazards like black ice (Rosen, 2022). It is reported that this technique is also used by professional race drivers to regain control of their vehicle. Therefore, racing-level control algorithms can significantly enhance road safety by handling scenarios where a human might lose control.

The vehicle-control lessons from autonomous racing at the limits discussed above can foster the development of novel safety initiatives in the context of consumer vehicle design and road infrastructure design and operations. In the rest of this chapter, we present more details related to vehicles dynamics and control, including the estimation of surface friction potential, emergency obstacle avoidance maneuvers, longitudinal and lateral control strategies, and hydroplaning and wet weather management. We identify lessons from autonomous racing in these areas of vehicle dynamics and control that have been, or could be, transferred to production AV operations on public highways.

7.2 A Review of Literature

This review discusses a sample of academic papers covering the spectrum of control algorithms for autonomous racing, ranging from classical geometric methods to advanced optimization-based methods, deep reinforcement learning and AI based techniques. Table 7.1 summarizes selected control algorithms, their core mechanisms, computational demands, and best-use cases in autonomous racing.

7.2.1 Classical & Geometric Baselines

Coulter (1992)'s Pure Pursuit algorithm is used as a baseline in many modern racing studies (including F1Tenth). This algorithm calculates the curvature required to navigate the racecar from its current position to a "lookahead point" (a fixed distance away on the reference path). In the context of autonomous racing, the "lookahead distance" is dynamically adjusted based on speed. This control method has negligible computational expense; however, it tends to "cut corners" excessively (when the lookahead distance is too long) or oscillate (when the lookahead distance is too short). Hoffmann et al. (2007) described the "Stanley" controller used by the Stanford University team in DARPA's Grand Challenge. In contrast to Pure Pursuit, the controller described by Hoffmann et al. (2007) uses a nonlinear feedback law based on "cross-track deviation" (that is, the distance from the prescribed path) and "heading deviation" (that is, the angle difference). The approach has been shown to exhibit local stability and ultimately converges to the prescribed path with zero deviation. In autonomous racing, the Stanley controller is often preferred over Pure Pursuit for its higher path-tracking accuracy on straightaway sections, though it requires careful tuning to avoid instability at high speeds.

7.2.2 Optimization-Based Control (MPC & Variants)

Liniger et al. (2015), in their seminal paper that used 1:43 scale RC cars, introduced MPCC, shifting the paradigm from merely following a fixed path to explicitly maximizing progress along the track. In other words, rather than minimizing deviation from a pre-specified reference line, their algorithm introduced a "progress" state variable, optimizing the movement by trading off between the risks and rewards of strict adoption of the reference line versus deviation from the line to earn some speed benefits. Thus, the algorithm establishes the optimal racing line in real-time within the constraints of the tires' friction circle limits. Rosolia et al. (2017a) developed a learning model predictive control (LMPC) approach that does not rely on a perfect physics model but rather leverages data from preceding laps to build a "safe set" represented as a convex hull of successful states together with an associated "value function". As the car completes more laps, the controller learns to reduce lap times with due regard for the physical limits of the vehicle and tires. Williams et al. (2017), using the AutoRally platform in a dirt-track racing environment, developed the Model Predictive Path Integral (MPPI) method, a particle-filter-like approach that simulates thousands of noisy trajectories forward in time. A weighted average of these trajectories (with weights derived from an information-theoretic cost function) enabled the racecar to handle aggressive slip angles on loose gravel surfaces. The authors demonstrated that their MPPI control algorithm handled drifting on dirt and other highly nonlinear dynamics, and, in general, exhibited superior performance relative to the traditional gradient-based MPC. In recognition of the limitations of physics-based models, Kabzan et al. (2020) used a full-size Formula Student platform to develop a hybrid control approach. This involved integrating a nominal vehicle model with a data-driven Gaussian Process (which learns the mismatch or error between the physical model's predictions and the actual car's behavior in real-time). By allowing the controller to operate closer to traction limits, the "corrected" model deployed in the MPC solver reduced lap times by approximately 10% compared to standard MPC. Samada et al. (2023) presented an MPC-based Reference Governor (MPC-RG) for autonomous racing that accounts for energy consumption. Their approach combines the MPC-RG with an LQR and a Kalman filter for disturbance compensation and estimation of unmeasured states, thereby guaranteeing recursive feasibility and constraint adherence while incorporating energy awareness and obstacle avoidance. Instead of only focusing on lap time minimization, their approach seeks to achieve a speed-vs-energy balance. Muraleedharan et al. (2022) addressed the computational bottleneck associated with non-linear MPC by using massive parallelization on GPUs. In contrast with gradient-based solvers, the RMPC algorithm establishes a single optimal control sequence by simultaneously sampling several thousand potential control trajectories. RMPC then evaluated the alternative trajectories using a cost function (weighted combination of lap time and safety) and selected the best trajectory. The authors stated that their approach helps the controller to avoid getting stuck in local minima.

7.2.3 Control using Reinforcement Learning & AI

Wurman et al. (2022) presented "GT Sophy," an AI agent that beat world-class human drivers in the Gran Turismo simulator. The control algorithm uses a specific form of Deep Reinforcement Learning (DRL) called Quantile Regression DQN. Unlike standard RL which estimates the average future reward, this method estimates the distribution of possible rewards, allowing the agent to better analyze the risk/reward tradeoffs between alternative maneuvers. GT Sophy learned complex driving tactics like slipstreaming (following closely behind a lead vehicle to reduce drag and fuel consumption) and trail-braking purely through trial-and-error, without hard-coded rules. Yi et al. (2024) proposed a framework that adds a safety layer to a standard DRL agent using a state-mapping mechanism to normalize track observations and an action-mapping technique that strictly enforces physical traction constraints. This helps ensure that even when the neural network suggests an unsafe action (for example, steering too hard at high speed), the final control output is constrained to a feasible range, thereby integrating the aggressiveness of RL with the stability offered by control theory. In their end-to-end deep learning control method for autonomous racing, Weiss and Behl (2020) explored the so-called “pixels-to-torque” approach, which involves a Convolutional Neural Network (CNN) that maps raw image data to steering and throttle commands, bypassing the traditional pipeline of Perception-Planning-Control. The authors trained their algorithm using imitation learning from a human driver's behavior, followed by DRL refinement, demonstrating that end-to-end systems are feasible for autonomous racing. Qiao et al. (2025), using the F1/10th platform, presented a “lightweight” learning algorithm for head-to-head racing that leverages a Gated Recurrent Unit neural network to capture temporal dependencies (memory) of the car’s state. The algorithm transforms raw LiDAR scans into “spatial pressure tokens,” allowing the agent to continuously remember the relative velocities of opponent racecars and enabling strategic overtaking maneuvers.

7.2.4 Hybrid Controllers

Wang et al. (2021b) proposed a bi-level control architecture, with an upper level that used Game Theory to make strategic movement decisions (yield, block, overtake) based on the opponent’s likely actions and a lower level that used a standard MPC to execute the chosen strategy. This separation allowed the car to “outsmart” opponents rather than just outpace them, simulating the tactical decision-making of a human race driver. This study offered a valuable parallel for real world autonomous driving: “it is not just about driving fast; it is also about interacting with other vehicles on the guideway.” O’Kelly et al. (2020b) documented standard baseline algorithms for the popular F1Tenth racing platform. Their study described the "Follow-the-Gap" method, a purely reactive algorithm that requires no map; the LiDAR scan is processed to find the largest "gap" (open space), and the steering angle is established to aim at the center of that gap. This method has been demonstrated to be simple yet effective, particularly on unfamiliar tracks, because it is successful in avoiding obstacles (walls, opponents, etc.) without requiring complex path planning. Xia et al. (2024)’s survey compared discrete planning methods against continuous control, and their findings suggest that while graph methods guarantee finding a path if one exists, they tend to produce “jerky” outcomes that are problematic in situations with handling limits. One of the paper’s key contributions is the conclusion that optimization-based smoothing (like MPC) is generally required as a post-processing step for any graph-based planner in racing. Liu et al. (2025) reviewed algorithms for stabilizing vehicles in unstable equilibrium states (high slip angles). Their article highlights the concept of Equilibrium Point Control where the controller linearizes the dynamics around a specific “drift state” (e.g., a 30-degree slip angle) rather than a straight-line state. This allows the racecar to maintain a continuous drift, a technique that helps the racecar to rotate quickly around tight hairpins (see Figure 7.1).

7.2.5 Summary

Table 7.1 presents the control algorithms, their primary mechanisms, computational requirements, and best-use cases.

Table 7.1 Control algorithms, their primary mechanism, computational requirements, and best-use case

Algorithm	**Paradigm & Key Mechanism (KM)**	**Dynamic Handling (DH) & Computational Cost (CC)**	**Best Use Case**
MPCC Liniger et al. (2015)	Optimization. KM: Optimization of progress along the centerline	DH: High: Explicitly models forces/friction. CC: High: Requires convex solvers	Time-trials; finding the optimal racing line dynamically.
LMPC Rosolia & Borrelli (2017)	Data-Driven / Optimization KM: Iterative learning from safe sets of previous laps	DH: Very High: Learns limits from history. CC: High: Solving MPC + convex hull constraints.	Repetitive track racing where the car improves lap-by-lap.
GP-MPC (Kabzan et al., 2019)	Hybrid (AI + Optimization) KM: Residual learning of model error via Gaussian Process	DH: High: Adapts to model errors online. CC: Very High: GP inference is expensive.	Changing conditions (wet/dry) or undefined tire models.
MPPI (Williams et al., 2017)	Sampling. KM: Information-theoretic sampling of noisy trajectories	DH: High: Handles high non-linearity (drifting). CC: Medium/High: Highly parallelizable (GPU).	Dirt/Rally racing; highly non-linear traction regimes.
Game-Theoretic MPC (Thakkar et al., 2024)	Game theory. KM: Hierarchical: Strategic discrete decisions + MPC execution	DH: High: GPU-accelerated sampling. CC: Medium (on GPU): Fast parallel evaluation.	Head-to-head racing; overtaking logic.
GT Sophy (DQN) Wurman et al. (2022)	Deep RL. KM: Distributional RL estimating reward variance.	DH: Implicit: Learns physics via reward. CC: Low (Inference): Fast forward pass.	Complex strategy; overtaking humans; high-fidelity sims.
End-to-End CNN (Wadekar et al., 2021)	Deep Learning. KM: End-to-End CNN with noise injection.	DH: Implicit: Pixels to Torque. CC: Low (Inference): Fast forward pass.	Scenarios rely heavily on visual cues rather than maps.
Pure Pursuit (Coulter, 1992)	Geometric. KM: Geometric tracking of a lookahead point	DH: Low: Ignores dynamics/mass. CC: Negligible: Simple geometric math.	Low-speed movement; robust fallback; baseline testing.
Stanley (Hoffmann et al., 2007)	Geometric. KM: Non-linear feedback of cross-track/heading error.	DH: Medium: Better tracking than PP. CC: Negligible: Simple feedback law.	High-accuracy path tracking on straight/moderate curves.
Follow-the-Gap (O'Kelly et al., 2020)	Reactive. KM: Reactive steering towards the largest LiDAR gap.	DH: None: Reacts to free space. CC: Negligible: LiDAR array processing.	Unknown tracks; obstacle avoidance; "survival" racing.

7.3 Friction Potential Estimation

The Tire-Road Friction Coefficient (TRFC), often described as the "foundation of vehicle safety," is a key safety parameter in vehicle operations. In public road vehicles, this is often estimated reactively, for example, when an Electronic Stability Control (ESC) unit detects wheel slip and applies braking to correct yaw. In autonomous racing, knowledge of the TRFC must be predictive and precise to optimize lap times and prevent loss of control at high speed.

7.3.1 The Physics of the Friction Circle

The algorithms used by autonomous racing teams maximize the use of the "friction circle", the combined limit of longitudinal (braking/acceleration) and lateral (cornering) forces that are generated by a tire in such maneuvers. The maximum possible longitudinal acceleration or deceleration, as well as the lateral acceleration of a vehicle, is influenced largely by the friction potential. As such, in racing, the coefficient of friction dictates the vehicle's driving dynamics: underdriving causes a loss of race time, and overdriving causes a crash. As such, researchers suggest that to minimize lap times, it is important to determine the friction potential with high local resolution (TUM CAT, 2015; Laurense et al., 2017). In a typical racing application, the racecar enters a corner, and the planner calculates the exact maximum speed based on an estimated grip; if the estimation deviates by even 2%, the car may slide off the track or lose critical racing time (Laurense et al., 2017). In a road application, a production AV may approach an icy bridge or a wet curve; instead of waiting to slide, the system utilizes the same predictive TRFC algorithms to pre-adjust speed and steering gain before the friction change occurs. Therefore, knowledge of the friction potential enables effective control of the AV in safety-critical situations, ensuring the vehicle remains within the domain of safe handling (TUM CAT, 2015).

7.3.2 Advanced Techniques for Estimating the Friction

In contemporary racing research, at least four approaches exist, two of which are particularly prominent in current racing research to estimate friction, and the application of these methods is transitioning to production vehicles. These are:

<u>(a) Dynamics-based approach</u> (Wang et al., 2025): This uses the so-called "magic formula" or "brush" tire models to estimate friction from micro-slips detected during driving. Racing researchers suggest that integrating these estimates into Model Predictive Control (MPC) allows vehicles to operate safely on surfaces with variable friction. The available grip is estimated by comparing the expected vehicle reaction, derived from a bicycle model, with the actual reaction measured by IMU data.

<u>(b) Optical and sensor fusion approach</u> Developed by researchers including Hu et al. (2025), this method uses LiDAR and cameras to detect road pavement surface texture and contaminants (water, oil, ice, snow) to predict the expected level of tire grip. Thermal data is considered optional but potentially enhances the model's predictive power. By "measuring" the pavement friction before the tire touches the pavement, the vehicle can adjust its controls preemptively.

The other approaches are acoustic and slip-slope approaches. Table 7.2 summarizes the friction estimation methods and their utility across the two domains: racing and public road operations by production vehicles.

Table 7.2 Racing vs. Road Applications: Estimation Methods and Sensor Requirements

Estimation approach	Application during racing	Extended application to public road operations	Sensor required
Optical prediction	Proactively detects wet/dry patches to adjust corner entry speed, to surface texture/grip.	Detecting black ice or hydroplaning risk (e.g., Continental system) to warn the driver/ADAS.	Cameras, LiDAR, Thermal
Acoustic	Analyzes tire noise frequencies to determine surface texture/grip.	Winter road maintenance (snowplows) and wetness level estimation.	Microphones in wheel wells
Dynamics-based (Brush model)	Measures the relationship between lateral force and slip angle during cornering to determine surface texture/grip.	ESC and ABS systems adjusting thresholds based on real-time grip estimates.	IMU, Wheel speed sensors, Steering angle
Slip-slope	Analyzes the slip-force curve slope in the linear region (small steering inputs) to determine surface texture/grip.	Early warning systems for ADAS to increase following distance in low-grip conditions.	High-fidelity IMU

7.4 Emergency Obstacle Avoidance Maneuvers (EOAM)

A large fraction of traffic crashes has been attributed to human error, inattention, excessive perception time or reaction time, or, in some cases, inappropriate reaction in response to perceived threats (Lowe & Guvenc, 2023; Campbell et al., 2024). Split-second decisions made at freeway speeds, if made in error or too late, could cause single-car collisions, multi-car pileups, and other crash modes that are often fatal. High-speed autonomous racing researchers have shown that when needed, autonomous cars can execute aggressive crash-avoidance maneuvers more reliably and safely compared to human drivers, without destabilizing the vehicle (Behl et al., 2025). For example, in public road operations, the production AV might encounter a situation where a pedestrian suddenly crosses the road; to avoid a collision, the AV could execute a split-second swerving maneuver, such as a hard braking, emergency swerving, or near-drift recovery. The algorithms developed in racing provide confidence that such maneuvers can be executed without losing control. This concept, often referred to as "agile safety", is a direct result of autonomous racing experience. As researchers have stated, "*by enabling an autonomous agent to learn how to race, we are... learning how to react to complex dynamic situations where obstacles can appear in close proximity with only seconds to avoid a collision*." (Behl, 2025).

In normal driving, safety systems are often conservative; on the other hand, autonomous racing has shown that agility and safety are not contradictory if managed appropriately. Prospectively, Emergency Obstacle Avoidance Maneuvers (EOAM) (Lowe and Guvenc, 2023), useful for vehicles that use existing high-speed highways, could learn from high-speed racing where autonomous agents are invoked to execute evasive maneuvers that are not only appropriate but also have minimal delay in perception and reaction. In certain cases, reinforcement learning has been used to select the safest available maneuver among multiple options in near-crash situations. The example discussed below explains how an autonomous agent could avoid single- or multi-vehicle collisions through the adoption of stability control via drifting.

An experimental study at Stanford University using a modified (autonomous) DeLorean sports car demonstrated that an autonomous vehicle can be adept at drifting, maintaining control while the vehicle is sideways with high slip angles, as shown in Figure 7.1 (Goh et al, 2018). Drifting is not a standard highway maneuver; however, the underlying control theory can be vital for maintaining vehicle stability when slipping on ice. On public roads, if an AV hits a patch of black ice and starts yaw-spinning, the "drifting" algorithms developed for racing could be deployed to modulate throttle and steering, to stabilize the vehicle, thereby turning a potential rollover into a controlled recovery. Such active handling that uses positive torque to stabilize the chassis is considered an improvement over standard ESC, which primarily uses braking as a stabilization maneuver.

Autonomous vehicle drifting experiment at 50km/h,
showing transitions through +/- 40 degrees sideslip in a second. *(Image credit: Jonathan Goh. Source: https://news.stanford.edu/stories/2019/12/autonomous-delorean-drives-sideways-move-forward)*

Figure 7.1 Stability Control via Drifting

In head-to-head autonomous races such as A2RL, avoiding contact with other vehicles is important because collisions could end a vehicle's participation in the competition when its hardware is damaged beyond timely repair. Therefore, autonomous racing places great emphasis on its collision avoidance systems which include predictive models that anticipate adversarial vehicles' movements and reactive controllers that execute emergency evasion maneuvers. In public-road operations, AVs will often encounter such situations. Unfortunately, such extreme scenarios are hard to study on public roads. Through autonomous racing, engineers have learned how to detect the onset of a spin (loss of control) in an AV and how to recover from it. Transferring this knowledge to public road vehicles could prevent several crashes caused by spinning out on wet road pavements. Further, through autonomous racing experiences, the importance of redundant safety systems has been highlighted. For example, each racecar has a remote "kill" button and dedicated safety circuits to bring the racecar to a controlled stop in an emergency. In public road operations, the AV could be equipped with redundant braking and steering systems, and even the ability for remote supervision or intervention in emergency situations. Researchers agree that autonomous racing makes it possible to study challenging pre-crash situations in unforeseen ways and that autonomous racing offers a valuable test environment for the development of cutting-edge solutions in autonomous driving (Demeter & Hajgató, 2024).

7.5 Longitudinal and Lateral Control Strategies

Longitudinal control manages speed (acceleration and braking) for distance keeping and comfort, while lateral control manages steering for path/lane following, ensuring accuracy and stability. Together, these constitute the core of autonomous driving. To assist in these control functions, sensors perceive the road environment, and algorithms (PID, MPC, Reinforcement Learning, and so on) generate precise steering angles and throttle/brake commands for smooth and safe navigation, particularly through curves and road sections with obstacles.

In standard highway lane-keeping, lateral (steering) and longitudinal (throttle/brake) control are often decoupled, for example, while steering, speed is kept constant. This is considered by many researchers to be an unduly restrictive simplification. Autonomous racing researchers have shown that coupling these two control dimensions is not only possible but also beneficial at high speeds. Autonomous racing stacks often use Model Predictive Control (MPC) which optimizes steering and throttle simultaneously; for example, applying throttle induces rearward weight transfer, affecting front steering grip. Research has shown that integrating feedforward steering controllers that account for delay and yaw compensation yields stability comparable to that of complex predictive controllers (Kapania & Gerdes, 2015a; Demeter & Hajgató, 2024). It has been argued that applying such coupled control to production AV operations on public roads can improve occupant comfort and safety during high-speed merging or curve negotiation.

7.6 Feedback Control Methods: PID, MPC, and AI

Model Predictive Control (MPC) and Proportional-Integral-Derivative (PID) control represent two popular feedback control methods. MPC uses a system model to predict and optimize future behavior, handling complex, multivariable systems with constraints, whereas PID is simpler, reacts to current error for single-variable tasks, and is computationally cheaper. PID excels at basic regulation, whereas MPC offers superior performance in industrial processes by "seeing ahead" to optimize multiple interacting variables and respect operational limits, often acting as a supervisory layer above PID loops. While optimal MPC-based control is dominant in the autonomous racing research literature, Reinforcement Learning is gaining ground rapidly. Sony AI's *Gran Turismo Sophy* demonstrated that RL agents could beat human champions by discovering racing lines and braking points that humans missed (Vasco et al., 2024).

7.7 Loss of Control due to Hydroplaning

Hydroplaning represents a potentially catastrophic loss of control in which tire-pavement contact is severely compromised effectively. Racing teams that compete in wet conditions have no option but to manage these conditions. Technology companies and vehicle manufacturers continue to develop aquaplaning warning systems based on surround-view cameras and tire sensor data. The ability to detect impending hydroplaning is well documented in the literature (Fichtinger et al., 2021; Vilsan et al., 2025). Evidence from racing competitions and the industry indicates that the risk of aquaplaning could be measured through subtle changes in wheel speed variance or torque response which trigger preventive measures by the vehicle, like micro-braking to dry rotors or adjusting torque vectoring (Pinkow & Fillenburg, 2018). Reliable prediction of aquaplaning has direct safety implications, and the high-fidelity sensing required in racing events that take place even under such inclement weather conditions helps highlight the safety benefits of this technology for production AVs.

7.8 Discussion

The extreme stress associated with racing environments has spawned improvements in racing vehicle hardware that can scale up production AVs on public roads. Standard Advanced Driver Assistance Systems (ADAS), for example, LKA, ACC, and AEB, typically disengage or perform poorly when a vehicle exceeds the linear region of tire dynamics. This is because they rely on simplified models or misaligned sensors, and are unable to handle extreme dynamics, leading to system disengagement, jerky actions, or failure to intervene. This highlights the need for robust validation, calibration, and alignment to ensure accurate sensor and control performance. Fortunately, autonomous racing algorithms thrive in this non-linear region because they can manage the forces at the adhesion limits to maximize performance. The analysis of racing data has helped teams to refine tire models, leading to enhanced understanding of how an autonomous control system should adjust to road conditions. Insights including the important role of tire management and AI-facilitated estimation of tire grip, have been earned from autonomous racing experience and can help production AVs assess pavement friction in black-ice or rain-slick conditions and adjust their speed in a proactive manner. For racing competitiveness on track and for production AV passenger safety on real roads, time optimal control (Mesterton-Gibbons, 2009) is a vital issue in autonomous driving and is a primary motivation for most control algorithms developed in this domain (Metz and Williams, 1989).

The chapter also showed how autonomous racing events provide lessons that are valuable for transfer to production AVs on public roads in the contexts of emergency obstacle avoidance, longitudinal and lateral control, and hydroplaning and wet weather management. Regarding emergency obstacle avoidance, autonomous racing researchers have developed algorithms that provide confidence that such maneuvers can be executed without losing control, and that agility and safety are not contradictory if managed appropriately. Regarding crash avoidance and mitigation, through autonomous racing, engineers have learned how to detect the onset of a spin (loss of control) in an AV and how to recover from it. Transferring this knowledge to production AVs on public roads could prevent many crashes associated with spin-out events on wet road pavements. Further, through autonomous racing experiences, the importance of redundant safety systems has been highlighted, and they can be beneficial to society if applied to production AVs operating on roads. Regarding longitudinal and lateral control, autonomous racing researchers have shown that coupling these two control dimensions is not only possible but also beneficial at high speeds for safety and travel efficiency.

In summary, autonomous racing greatly accelerates the development of safety-critical control techniques for production AVs. By training a vehicle in an environment where they must routinely operate near the vehicle's physical limits while avoiding crashes, racing programs produce systems that are far better prepared for the “corner cases” of real-world driving. These competitions help prepare production AV systems for uncertain, highly dynamic, and otherwise extreme situations that may arise in the real world. As a result, when these technologies transition to consumer cars, everyday driving will be expected to become not only more efficient but also safer, as vehicles perform more effective crash-avoidance and crash-mitigating maneuvers and maintain stability where a human driver might fail.

CHAPTER 8 HARDWARE AND SOFTWARE SAFETY

"The goal of autonomous racing is to push AI to its limits to safely handle 'edge cases', ultimately making self-driving cars for public roads more robust and reliable."- Simon Sagmeister, former team principal, TUM Autonomous Motorsport

8.1 Introduction

"A racing car needs to be fast, but it also needs to be efficient and durable to complete the race smoothly. Engineering is concerned with maximizing the efficiency of the systems and making sure that the car's components can withstand the extreme conditions and stress of a race" (Meriño Córdoba et al., 2025). Racing introduces engineering challenges that, in turn, stimulate automotive innovation and accelerate the development of robust safety systems for the broader autonomous-vehicle industry. For example, the high stress situations experienced by racing vehicles motivate the improvements in AV hardware that can ultimately translate into safer AVs on public roads. In this chapter, we identify several areas related to hardware and software safety (HSS), and algorithm integrity where OEMs could learn from the autonomous racecar experiences. We begin with a review of autonomous racing research articles related to HSS, then discuss the overarching concept of a software-defined vehicle and finally show how autonomous racing provides valuable lessons regarding HSS, including software robustness and reliability, latency, and thermal management of vehicle components.

8.2 The Concept of Hardware and Software Safety

In the context of autonomous driving, a safe autonomous system is one whose hardware and software are designed and operated to reduce uncertainty, function reliably in highly dynamic environments, and minimize vulnerability to human error in both software design and hardware operations. Figure 8.1 presents major categories of autonomy-related threats to autonomous racecars operating on a racetrack.

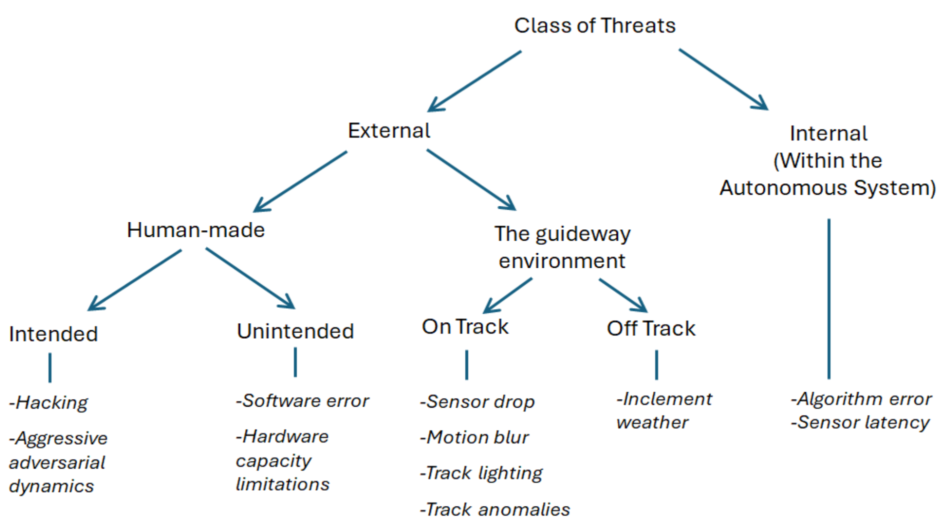

Figure 8.1 Categories of autonomy-related threats to autonomous racecars on a racetrack

From the perspective of autonomous systems, a threat may be external or internal. Internal threats are those that occur within the autonomous system and include sensor latency, algorithm errors (implementation and execution errors including as runtime, syntax, and semantic errors, and arithmetic overflows; and resource and system errors (memory leak, time limit exceedance, and stack overflow), and other shortcomings of the autonomy stack. External threats may arise either from human actions or from the guideway environment. External human-made threats may be intended (malicious), including hacking by malicious parties or aggressive and unsafe maneuvers by competing racecars, or unintended (software errors and hardware-capacity limitations occurring in systems external to the AV). For example, inadequate memory or computing resources of base station computers may prevent the transmission of massive streams of sensor data to or from the AV. Regarding autonomy-related external threats encountered by the autonomous racecar in the guideway environment, the source may be on the track (sensor drops, motion blur, track lighting or lack thereof, and track anomalies) or off the track (inclement weather, for example). Recognizing that no driving environment is perfectly deterministic and unexpected situations can arise at any time, racing teams seek to ensure that their AV software can address such anomalous situations as much as possible. Accordingly, a production AV should be capable of real-time adaptation in its sensing, control, and planning functions so that it can navigate unpredictable environments, improve safety, and enhance user trust. This capability is referred to as "adaptation".

8.3 A Review of the Literature

This section reviews selected peer-reviewed studies that address hardware and software safety (HSS), including data security and protection against attacks. The articles included in this review were chosen based on relevance to the topic of HSS in autonomous racing.

To reduce sensitivity to outliers, sensor dropouts, and mode changes in vehicle dynamics, Enisz et al. (2024) used multiple EKFs including parallel filters and/or filter switching strategies. They demonstrated how estimation redundancy improves stability in autonomous racecars operating near their limits. Massa et al. (2020) addressed the integrity of autonomous racing localization when GNSS is unavailable or unreliable, proposing a LiDAR-centric approach that combines two multi-rate EKFs with a Monte Carlo localization variant leveraging track priors. Peng et al. (2021), working in the Formula Student Driverless context, sought to improve robustness of multi-sensor odometry needed for autonomous racecars operating under threats associated with such variables as track lighting, motion blur, and aggressive dynamics from adversarial racecars on the track.

Autonomous race vehicles operate at high-speed while relying on dense, safety-critical in-vehicle networks under adversarial and fault-intolerant conditions. In this context, many studies not specifically focused on racing nevertheless provide contributions that are applicable to autonomous racing stacks. Rathore et al. (2022) discussed classes of algorithms used in practice, including cryptographic protection and detection of ML-based intrusions or anomalies, and proposed a multi-layer framework.

Hossain et al. (2020) proposed an intrusion detection system for CAN traffic using Long Short-Term Memory (LSTM) sequence models trained to detect attacks such as denial-of-service (DoS), fuzzing, and spoofing based on message timing and content patterns. The core algorithm is a deep sequence classifier for CAN message streams, evaluated on attack scenarios. For autonomous racing, this is directly relevant because a compromised CAN network could impair steering and braking commands as well as sensor availability, while an LSTM-based IDS offers a plausible approach to real-time monitoring of the autonomous racecar.

8.4 The Concept of a Software-Defined Vehicle

Over the past several decades, vehicles have been differentiated largely by their mechanical features, including torque and horsepower. In the current era, as driving becomes increasingly automated, consumers increasingly seek what could be termed software-defined vehicles; vehicles that possess a preponderance of software features including driver assistance, vehicle communications capabilities, and infotainment possibilities (Aptiv, 2020). In addition, V2X connectivity features are expected to accompany most future production of AVs (Ha et al., 2020; Acharya & Mekker, 2022; Nakanishi & Auza, 2023). As such, transportation vehicles are increasingly becoming part of the internet of things (IoT) ecosystem revolution, exchanging massive volumes of data with the cloud. Software will play a major role in managing such data to enhance the AV operations. These developments have given rise to the concept of a software-defined vehicle. A software-defined vehicle (SDV) is a vehicle where software, rather than just hardware alone, increasingly controls and enhances its features, allowing for continuous updates, new functions, and personalization remotely via over-the-air (OTA) updates, in a manner analogous to a smartphone (Aptiv, 2020; Slama et al., 2023; Siemens, 2025). SDVs shift from many individual electronic control units (ECUs) to a centralized, powerful computing architecture, enabling features (including advanced driver-assistance systems, powertrain management, and infotainment) to be upgraded and optimized throughout the vehicle's life, improving safety, efficiency, and user experience. It is expected that the transition from hardware with embedded software to SDVs and services will shape future mobility (Bosch, 2025). Autonomous racing provides a valuable test environment for validating certain SDV features before they become mainstream.

8.5 Software Robustness and Reliability

The extreme demands of autonomous racing motivate racing teams to ensure high-performance computing and reliable software for their vehicles. At racing events, teams typically run complex autonomy stacks, often based on ROS 2 and Autoware, on high-performance CPUs and GPUs operating near capacity to handle massive data throughput in real time. For this reason, racing teams strive to optimize their codebases to improve efficiency, real-time communication, and fail-safe capabilities. The takeaway for production AVs is clear: software needs to be updated and optimized regularly and made robust to handle real-world complexities. Further, the racing competitions provided avenues for the teams to develop methods to achieve low-latency processing and to prevent software crashes or vehicle collisions under pressure. For example, teams learned to limit non-essential data traffic so only high-priority data streams could come through. These lessons inform how software for production AVs could be made more robust to enhance safety on public highways (Zenoh, 2021). In general, autonomous racing events held worldwide have served as extreme stress tests for the entire autonomous driving stack, revealing and fixing software weaknesses that could surface subsequently on public roads in the prospective era of widespread autonomous mobility.

Anecdotal evidence from autonomous racing teams indicates that they place a high premium on computational efficiency to process raw LiDAR point clouds locally and in real time on the racetrack. To do this, they often used compact, high-performance edge computing units such as the ADLINK AVA computer series. For perception and control tasks, many teams eschewed cloud processing (which they consider too slow for real-time safety-critical maneuvers), instead opting for optimized embedded local hardware for efficient onboard execution. According to ADLINK (2021), this local processing capability was essential for production AVs to function safely in areas with poor cellular connectivity. As V2X capabilities become more robust and reliable, it is possible that some racing architectures will make greater use of cloud-supported or distributed computing. Such developments could provide useful decision-support evidence for OEMs as they evaluate trade-offs among onboard, edge, and cloud-based architectures in the transition towards production AVs, specifically, SDVs.

8.6 Latency

Strongly related to software robustness and reliability is the concept of latency. At 190 mph, a racecar travels approximately 278 feet per second. A latency of 100 milliseconds, or one-tenth of a second, means that there is a 0.1-second delay for data to travel from the racecar to a server and back. At 190 mph, the racecar moves 28 feet before the software processes a new sensor frame. Clearly, racing demands ultra-low latency pipelines if the racecar must perceive and respond quickly to roadway opportunities, and particularly, threats.

At present, automobile design practices are undergoing a major shift in automotive electronics, and autonomous racing is serving as a valuable test bed for OEMs in this regard. High-performance steering, and other critical systems, in future Software-Defined Vehicles (SDVs) demands actuator update rates in the kilohertz (kHz) range for stability, far exceeding typical Controller Area Network (CAN) bus speeds. This requires high-bandwidth solutions like Automotive Ethernet to handle the immense data flow for precise, real-time control, which is particularly crucial to safe and efficient advanced driver-assistance systems (ADAS) and autonomous driving (Manna, 2022). The trend, which has significant implications for production AVs, is driven by at least four factors:

(a) *Limitations of Controller Area Network (CAN) capabilities*: A standard CAN architecture is not designed for the massive data volumes and low latency required for complex, real-time feedback loops in autonomous vehicles.

(b) *Stability at High Speeds*: At high speeds or in emergency situations, the AV needs to make rapid steering adjustments to maintain vehicle stability. Slower updates could lead to lag, making the car feel sluggish or unstable.

(c) *The Role of Automotive Ethernet*: Ethernet provides the gigabit speeds needed to transmit the massive data expected from various sensors (LiDAR, camera, radar) and to send rapid control commands to actuators (steering, braking). This must occur with minimal delay if the AV must be promptly responsive to threats.

(d) *SDV & Electrical Architecture*: The evolution towards advanced (high-bandwidth) zonal architectures using Ethernet is foundational for SDVs. This is expected to facilitate software-based management of vehicle dynamics, leading to more agile and intelligent vehicles.

In effect, autonomous racing, with its extreme demands, serves as a safe and cost-effective proving ground where advanced communication architectures can be tested for latency, performance, verification, and validation before they become mainstream in production AVs.

8.7 Thermal Management of Vehicle Components

The thermal management of racecar processors and brakes is important to ensure their reliability under extreme and prolonged stress and to prevent them from exceeding their thermal limits. Just as racecars rely on sophisticated electronics to function reliably throughout a race, a production AV computing system must perform continuous and complex computations without throttling or failing. This requires robust, often liquid-based, cooling systems similar in principle to those used in high-performance autonomous racing applications (Iclodean et al., 2020). In a future era of widespread autonomous mobility, such thermal management approaches may prove important for reliable, continuous operation of production AVs under demanding traffic conditions.

CHAPTER 9 ROAD INFRASTRUCTURE PREPARATIONS – DESIGN AND MANAGEMENT

"The future belongs to those who prepare for it today." – Malcolm X.

9.1 Introduction

Throughout the history of road transportation, highway design has evolved significantly, catalyzed by changes in vehicle technology, pavement materials science, and construction innovations. The motivation has been not only greater highway-infrastructure longevity but also enhanced travel efficiency and improved safety of vehicle occupants. As the world moves toward autonomous mobility, road agencies and transportation organizations increasingly acknowledge that current road-design standards, traffic rules, and operational protocols require adjustment before they can safely and effectively accommodate AV operations (AASHTO 2017; TRB 2014; FHWA, 2018). Accordingly, governments at the local, state, and federal levels seek guidance on what changes are needed, how extensive those changes should be, and when they should be implemented. As explained in this chapter, several researchers and practitioners have offered guidance on specific areas in road design where vehicle automation will influence the design of road infrastructure. Autonomous racing continues to provide needed validation of these prospective infrastructure needs in the era of partial or full automation. Chapters 3 to 8 of this report discuss autonomous racing lessons that apply primarily to vehicles and autonomous-driving systems; however, they provide little regarding any lessons to be learned regarding road infrastructure. Autonomous racing may offer some lessons for road-infrastructure design and management, including the introduction of new physical and cyber assets needed to support high-speed autonomous operations. Any such lessons could motivate future-proof investing in physical and cyber infrastructure in the current era to support emerging autonomous mobility in a proactive manner.

9.2 Physical Infrastructure

The operation of connected and autonomous vehicles on public roads is likely to prompt revisions to the design standards published in AASHTO's *Green Book* (A Policy on Geometric Design of Highways and Streets) for implementation on roadways in the era of full or partial autonomy, particularly at high-speed areas. Several researchers have already discussed these potential changes. Gopalakrishna et al. (2015) and Tengilimoglu et al. (2023) discussed the changes needed in road infrastructure to render it ready for AV operations. Colonna et al. (2018) investigated aspects of rural road design that need to be modified in an era of automated vehicles from a safety perspective, and Carreras et al. (2018) discussed road infrastructure support levels for automated driving. Pham et al. (2021) assessed the impacts of autonomous vehicles on road and pavement design.

Most racetracks used in autonomous vehicle competitions are oval-shaped (see Appendix Figure A.1(a)) and therefore provide a road layout that is different compared to public roads and highways. Autonomous racecourses (Appendix Figure A.1(b)) are closer to public roads in terms of design and operating conditions but still differ in important respects. Nevertheless, a few insights, mostly from a geometric perspective, can be extrapolated. Racetracks are designed with smooth curves, banking, and clearly delineated edges. Production AVs on public roads use these features (clear lane markings and consistent surfaces, etc.) to stay stable and oriented. In the era of autonomous mobility, it will become critical to maintain high-contrast highway lane markings on public roads. **This is a valuable lesson, underscored every time an autonomous racecar loses precious race time due to faded track markings or poor lighting.** To enhance roadways for autonomous mobility, road agencies will need to update standards to ensure sensor-friendly pavement markings and signage. Some agencies have already tested different paint materials and machine-readable signs (Saeed, 2019; Labi et al., 2023).

Regarding racing lines and road geometry, it is worth noting that race cars do not stay in the center of the lane. They minimize curvature by using the full width of the track, following an optimized path known as the "racing line". Autonomous racing teams devote substantial effort to optimizing their race lines, because following optimized race lines helps the autonomous racecar minimize lateral acceleration (g-force), thereby allowing for higher speeds and, in other contexts, greater stability at curved sections. In highway design, clothoids (Euler spirals) are used to smooth the transition from a straight road section to a curve in a manner that minimizes jerk (which reflects the acceleration change rate). Production AVs on public roads can use path-planning algorithms learned from racing to dynamically calculate an "optimal line" within a lane at curved sections to maximize vehicle stability.

In addition, to maintain their competitiveness, autonomous racing teams have conducted research on trajectory planning to maximize tire grip and minimize jerk. The lesson learned is that such trajectories can be followed to minimize jerk, thereby improving AV passengers' comfort and reducing motion sickness. A path planning algorithm that optimizes for the "limit of friction" at racetracks can be re-tuned to optimize for the "limit of comfort" (for example, in a manner that maintains the lateral acceleration below $0.2g$), ensuring a smooth ride even on winding roads. de Winkel et al. (2023) evaluated alternative passenger comfort standards for AVs based on jerk and acceleration. Similarly, Lee et al. (2024), in a paper covering lessons from an F1Tenth race, used deep reinforcement learning to investigate autonomous driving strategies aimed at minimizing jerk.

Secondly, AVs have shorter reaction times and more precise reactions compared to humans. Therefore, geometric design standards, which traditionally account for human limitations, perception-reaction time, driving behavior and human psychology, may require adjustment in an era of autonomous mobility, particularly under full market penetration of autonomy in a traffic stream. These standards include stopping sight distance, lane widths, horizontal curve radii, and so on. This assumption holds only in near-100% AV market penetration; otherwise, such designs may be unsafe for the remaining human-driven vehicles in the traffic stream.

Generally, a well-defined track promotes the performance of autonomous racecars. An infrastructure lesson from autonomous racing, therefore, is that consistency and predictability in road geometry will promote safe operations of production AVs. In addition, more compact road designs, including shorter merge sections and narrower lanes and parking spaces, could be adopted when AVs achieve near-total market penetration.

Current US highway lanes are typically 12 ft wide, much wider than a standard vehicle, to accommodate human error and wander. Unlike humans, autonomous driving systems achieve centimeter-level precision, as evidenced by empirical evidence of autonomous race cars traversing the same trajectories with minimal deviation lap after lap. This ADAS precision has a dichotomous effect on infrastructure design and performance. On the one hand, it may support reduced road lane widths. Due to such precision, a 12 ft lane can be considered wasteful in the era of full autonomy, and research suggests that AV lanes could be as narrow as 9-10 ft (Othman, 2021). For new roads, this means reduced cost of construction and maintenance (Labi et al., 2023; Cao et al., 2024); for existing roads: this means freeing road space for reallocation to other uses (additional lanes, buffer areas, or lanes for other modes). On the other hand, precision in ADAS vehicle trajectories will mean that the AVs will always travel along the exact same paths within the lane, leading to grooves (ruts) on the pavement surface (which, in turn, exacerbate hydroplaning and skidding crashes). Research has shown that programmed wander of the AV within a defined boundary around the trajectory can help mitigate this problem (Gungor & Al-Qadi, 2020; Okte & Al-Qadi, 2022; Georgouli et al., 2021). This implication extends beyond geometric design to pavement performance, as shown in Table 9.1. In autonomous racing, racing lines can be effectively programmed to introduce a "lateral offset" into their path. A central traffic controller can command different vehicles to track slightly different lines within the lane, distributing wear evenly across the pavement surface. This is a potential direct application of the precise path control concept developed in autonomous racing.

Stopping Sight Distance (SSD) is "the distance required to perceive an object in the roadway and bring the vehicle to a stop" (AASHTO, 2018). For manual driving, SSD is based on human reaction time (2.5s); for automated driving systems, reaction time is less than 100 milliseconds, far lower compared to manual driving. Changes in the stopping sight distance in the era of full or partial autonomy will influence the following elements of roadway design: vertical curves (crest and sags); road ancillary assets (median barriers, bridge abutments, crash barriers); storage and turn bay lengths at intersections (McDonald, 2017); and the frequency of warning signs and safety systems.

Overtaking Sight Distance (OSD) is the minimum length of clear roadway for a driver to safely pass a slower vehicle, accounting for their own reaction time, the space needed for the maneuver (including the slower vehicle's movement), and the distance an oncoming vehicle travel (AASHTO, 2018). This could be on a straight section, a vertical curve, or a horizontal curve. On curves, AVs may have an advantage over HDVs because their V2X capability can supplement line-of-sight limitations by providing information beyond the driver's direct field of view. Regarding horizontal curves specifically, overtaking or passing sight distance refers to the clear, unobstructed view needed for a driver to safely overtake another vehicle on a curved section of the road, limited by obstructions like cut slopes or buildings on the inside. As such, it is needed to calculate the necessary curve radius or obstruction offset to provide adequate stopping sight distance (SSD) or passing sight distance (PSD) for safe operation.

The Horizontal Sight Offset (HSO) is the lateral distance between the inside lane centerline and a sight obstruction (like a wall or slope) on a horizontal curve, ensuring that drivers have adequate clear view (sight distance) to stop safely. For this reason, Overtaking Sight Distance in a horizontal curve (OSD_H) is influenced by the quickness of human perception and reaction. The situation is similar for OSD in a vertical curve (OSD_V). OSDs are based on reaction time (for manual driving, human driver reaction time is 2.5s; for ADAS, reaction time (<100 milliseconds) is far lower compared to human drivers).Therefore, on AV-only roadways, there may be flexibility to adopt smaller curve radii or permit higher operating speeds than those used in conventional HDV-based design compared to current (traditional or HDV-only) designs (Welde & Qiao, 2020; Zhao et al., 2022).

Head-to-head autonomous racing has also offered valuable guidance for enhanced efficiency of an ego vehicle's movements by sensing all other vehicles in its neighborhood and predicting their movements. For production AVs that possess this capability, they will be able to operate at grades that are steeper, ramp terminals that are shorter, and merge areas that are shorter compared to those of human-driven vehicles (McDonald, 2017). AVs may also be able to operate with smaller gap acceptance requirements for crossing or turning movements, which could impact the design of medians and refuge islands, as well as with shorter queues, which could reduce required turn-bay lengths relative to those for human-driven vehicles. Table 9.1 synthesizes selected geometric design criteria may change in the era of full autonomy. The geometric design implications of AV operations can be profound. AVs, through sensing and V2X, can communicate and perceive obstacles (and other vehicles) beyond the traditional lines of sight. As such, for AV-only roadways, standard Stopping Sight Distance requirements that govern the design of vertical and horizontal curves could be relaxed, allowing for tighter highway geometries (an important consideration in constrained urban areas), and reducing construction costs. **However, to preserve HDV safety, roads that will have mixed traffic (HDV + AV) will need to have traditional or conservative geometric designs.**

As Table 9.1 shows, many of the geometric design elements currently governed by human perception-reaction limitations may be reinterpreted under AV-only operating conditions, although mixed-traffic environments will still require conservative design values.

Table 9.1 Road geometric design criteria in the HDV era versus the AV era

Design Element	**Current Standard (AASHTO Green Book)**	**Direct Observation/Lesson from AV Racing**	**Ultimate Lesson for Development of AV-friendly Infrastructure**
Lane Width	12 ft (3.6m) standard.	High precision of ADAS allows safe operations in 9-10 ft lanes.	More lanes can be added in the existing right-of-way; reduced costs of new construction.
Stopping Sight Distance (SSD)	Based on human reaction time (2.5s) + braking distance.	ADAS reaction time (<100ms) is far less than human drivers.	Fewer warning signs and safety systems.
Overtaking Sight Distance (OSD)	Based on human driver reaction time (2.5s), distances between vehicles, and other factors	ADAS reaction time (<100ms) is far less than human drivers. OSD is shorter for CAVs because V2X supplements line-of-sight limitations.	Reduction in curve radii; increased allowable speed on existing curves.
Superelevation (Banking)	Limited to prevent sliding if stopped.	Dynamic friction estimation allows safe traversal at higher limits.	Optimized banking for higher design speeds.
Pavement Design	Assumes random wheel distribution (wander) of HDVs.	Zero-error tracking of ADAS causes localized rutting.	Implementation of programmed wander to distribute pavement wear.

Sources: Gopalakrishna et al. (2015), Gungor and Al-Qadi (2020), Othman (2020), Welde and Qiao (2020), Georgouli et al. (2021), Okte and Al-Qadi (2022), Zhao et al. (2022), Labi et al. (2023).

9.3 Cyber Infrastructure

During autonomous races, each racecar is typically in continuous communication with a trackside network including the racing team's base station. This link is used for transmitting telemetry data, allowing the team to monitor their vehicle's status in real time, for purposes of assistance in localization and navigation, and, critically, for implementing safety commands such as emergency stops in hazardous situations. For example, in 2021, the Indy Autonomous Challenge used Cisco's Ultra-Reliable Wireless Backhaul (URWB) wireless network system around the track to ensure that the cars remain connected with minimal latency even at speeds over 180 mph (Zenoh, 2021; Caragata, 2022). Autonomous racing teams must account for the possibility that a vehicle traveling at high speed on the track may experience "ghosting", that is, loss of connectivity and communication with external systems. In such cases, the teams are unable to remotely stop their cars in emergency cases, and their car is left with only its sensors for navigation; this may be insufficient during emergency situations and could lead to a collision.

During autonomous racing events, all participating teams share the wireless network system around the track. As such, the race organizers often impose strict limits on bandwidth and encourage teams to practice prudent data management to prevent dropouts (Zenoh, 2021). Cisco's involvement in the IAC demonstrated that "seamless handoffs at high speeds between network towers eliminate dangerous delays, keeping the vehicle safely on the track" (Caragata, 2022).

Several racing teams have encountered the ghosting phenomenon, resulting in collisions with the racetrack wall, competing vehicles, and other objects. The implications of ghosting for public-road operations of production AVs are significant. A ghosted AV traveling at 75 mph on an urban freeway presents an obvious safety concern particularly when none of the occupants is eligible to take over the vehicle (for Level 3-4 AVs) or where takeover is not possible due to the absence of a steering wheel and brake pedal (Level 5 AVs). In the era of autonomous mobility, it is imperative that all public roads on which production AVs operate must be equipped with the necessary V2I infrastructure to ensure uninterrupted connectivity.

This is vital, because in the future of autonomous mobility, high-speed freeway sections must be able to serve very large volumes of connected autonomous vehicles daily, and these sections will need robust and reliable dedicated wireless infrastructure (5G or specialized networks) with adequate bandwidth to provide reliable service for all the vehicles at high vehicle speeds between cell towers or roadside units. Such service would include maintaining communications for the transmission of traveler information ATIS messages such as downstream hazard warnings, and congestion avoidance, map updates, and sharing/receiving V2V messages and signals.

In addition, having an emergency-stop capability or override via a communications network serves as a safety buffer layer for the production AV on public roads – similar to autonomous racing events where race officials can remotely disable an errant racecar: traffic management centers overseeing roadway operations would benefit from the capability to slow or stop errant autonomous vehicles (due to mechanical failure, software failure, or cyberattack). The autonomous racing experience has motivated engineers to design and implement communication networks that can meet these ultra-reliability and low-latency requirements, with direct lessons for AV deployment on highways.

Another area where autonomous racing has offered infrastructure-related lessons is localization. On the racetrack, racing teams have found that GPS alone is often insufficient for pinpoint accuracy, particularly on sections of track with limited sky visibility, such as nearby buildings, and under tunnels or bridges. According to Chudzinski (2023), IAC cars have used real-time kinematic (RTK) GPS positioning coupled with high-definition maps of the track to stay close to the optimal line; at an autonomous racing event at the Monza road course (which has overpasses), racing teams had to confront losing GPS signals and make their systems more reliant on maps and vision-based sensing. This mirrors the critical connectivity needs on public roads: urban canyons, tunnels, or rural areas with poor GPS signals translate into a need for production AVs to operate only at roads that have adequate communication infrastructure including roadside beacons, differential GPS base stations, or enhanced digital maps for localization.

Autonomous racing has helped highlight the importance of high-fidelity maps and the integration of onboard sensors for localization when external signals drop. Learning from this, highway agencies at all levels of government can invest in digital infrastructure, such as high-definition roadway mapping, radio beacons in GPS-denied zones, and corridor-level positioning aids to help the AV maintain accurate positioning continually. Experiences from autonomous racing underscores the importance of reliable connectivity, particularly where vehicles travel at high speeds with little tolerance for error. The clear lesson for autonomous mobility involving production AVs on public roads is that infrastructure supporting robust V2I (vehicle-to-infrastructure) communication is vital for highway connectivity and must be made available.

9.4 Discussion

The timeline and intensity of the needed modifications in infrastructure design and management in preparation for AVs will be influenced by factors associated with four key stakeholder groups: AV market adoption (reflecting the AV users), AV technology (reflecting the OEMs), budgets and investment levels for AV dedicated lanes or existing roadway retrofits to accommodate AVs (reflecting the road agencies), and the AV-related regulation and policies (reflecting the government legislature) (Saeed et al., 2020).

At a given level of AV market penetration (MP), the prevailing maximum level of vehicle automation (LOA) will influence the timing, scope, and intensity of infrastructure preparations needed to support AV operations at that MP and LOA. Generally, higher levels of MP and LOA translate into higher levels of infrastructure preparation. The SAE has defined six levels of driving automation from Level 0 (no driving automation) to Level 5 (full driving automation) (SAE, 2021).

Autonomous racing provides a microcosm of the emerging autonomous mobility ecosystem. Autonomous racing competitions continue to work with companies, including Cisco, to develop cyber (networking) infrastructure solutions. The experience of autonomous racing provides evidence of the need for infrastructure that not only protects AV occupants but also facilitates reliable communication and precise positioning of the autonomous system. Road transportation agencies and other infrastructure-owning and operating organizations (IOOs) seeking to “prepare the ground” for the emerging future of autonomous mobility, will need to undertake investments and work with the legislature and OEMs to identify, plan, design, and construct the needed supporting infrastructure; physical infrastructure (dedicated lanes, road markings, and roadside sensor installations) and cyber infrastructure (high density maps, communication networks, roadside units, V2X facilities, and AI-based decision systems). Therefore, IOOs can learn a lot from the organizers of autonomous racing competitions.

At the time of reporting, anecdotal evidence suggests that fully autonomous and connected taxis are operating on selected public roads in several major US cities, including Phoenix, Los Angeles, San Francisco, Atlanta, and Austin. As connected and automated vehicles become more common on roadways, lessons from the autonomous racetrack can guide policies that help prepare the physical infrastructure and cyber facilities, including communication spectrum allocation for V2X systems. Such preparations include revising standards for road condition and design through maintenance and CAV-ready audits.

CHAPTER 10 WIDER LESSONS, PROGRESSIVE AV DEPLOYMENT, AND SAFETY ISSUES

"Massive technological innovations are taking place in connected and autonomous transportation that will dramatically transform the way we live, work, shop, and entertain." -
Kumares C. Sinha, National Academy of Engineering.

10.1 Introduction

The preceding chapters of this report have provided insights into the lessons that stakeholders in autonomous mobility on public roads can learn from the experience of autonomous racing teams, organizers, and events. It is important to note that beyond the lessons related to the vehicle, autonomous driving system modules, and infrastructure as discussed in preceding chapters, autonomous racing has also played a critical role in shaping public perception, building trust in automation, educating an enthusiastic workforce, and fostering supportive regulation and policy related to AVs. These broader impacts are particularly important for infrastructure owners and operators (IOOs) which generally have a public fiduciary duty to ensure safe and efficient transportation and a specific duty to respond to public scrutiny, safety concerns, and growing expectations regarding readiness of government policy and infrastructure in preparation for the emerging era of autonomous mobility.

10.2 Wider Lessons

Some impacts of autonomous racing may be characterized more broadly as systemic benefits rather than lessons in a narrow technical sense. One is *user trust and public acceptance*: by visibly demonstrating autonomous systems capabilities, autonomous racing helps to familiarize the public with the idea of driverless vehicles operating safely at high speed. Head-to-head races are particularly helpful in this regard. These events can help overcome public skepticism, enhance user trust, and pave the way of broader adoption of AVs. The excitement of the races also helps *recruit talent* and inspire future engineers. Motorsports events are characterized by "adrenaline-inducing hype and large fan bases, thereby serving as great avenues for showing the public that automated vehicle technology is making huge strides" (New Eagle, 2022).

Mar et al. (2025) argued that the competitive nature of autonomous racing promotes innovation and fosters the emergence of specialized startup OEMs that develop high-end racing software and hardware. They argue that the increased demand for AI enhancements to autonomous driving modules (sensing and perception, localization, controls, and path planning, etc.) will facilitate AV component development that will benefit both motorsports and commercial autonomous systems in production vehicles. In addition, as reported in Mirage News (2025), Markus Lienkamp, professor and chair of automotive engineering at the Technical University of Munich, observed that autonomous racing promotes "the vision of bringing autonomous systems safely and efficiently onto the road." Similarly, in recognition of lessons learned from roboracing competitions, Stewart (2017) envisaged that innovations from track testing would eventually translate to production vehicles.

Another wider impact of autonomous racing is workforce development for autonomous mobility in the future, as discussed by several researchers including Nakamoto & Kobayashi (2019). Autonomous racing teams train highly skilled engineers in various aspects of vehicle automation. In most racing leagues, most team members are university-based, and the competitions serve as unparalleled hands-on learning experiences for students, including undergraduates, graduate students (MS and PhD), post-doctoral researchers, and research scientists. Many of these participants ultimately build careers in OEMs in the autonomous vehicle industry, AV startups, and related technology vendors. These organizations include autonomous systems startups and related initiatives that emerged from autonomous racing competitions and adjacent activities. Examples include Driveblock, Blickfeld, and Artisense in Germany; Aidoptation in the United States and Belgium; and Code 19 in the United States. Through these channels, transportation agencies and OEMs are benefitting from a steady stream of talent that understands both advanced technology and practical roadway operations (New Eagle, 2022). Moreover, the data and research output from racing are being published and shared publicly, for the benefit of researchers and OEM R&D personnel engaged in autonomous mobility.

Regarding workforce development in the autonomous mobility industry, Formula Student Driverless, a European competition funded partly by the Center for Connected & Autonomous Vehicles and Innovate UK, is probably the most impactful initiative. The competition involves student teams developing autonomous vehicle technologies for future of autonomous mobility applications. According to the organizers, a key goal of the competition is to train an autonomous systems workforce to acquire skills sought by the autonomous mobility industry and OEMs. The overall winner of the competition is selected based on vehicle design, performance, sales presentation, costs, and business plan (BMW Group, 2025). This series has become a recruitment pipeline for major European OEMs, including BMW and Volkswagen.

In addition, autonomous racing events serve as a regulatory sandbox for autonomous mobility policy. Because they operate in a controlled environment, organizers and participants in these events are provided with a controlled "laboratory' in which they can experiment with rules that might foreshadow future traffic regulations for AVs. For example, specifying and enforcing safe lateral and longitudinal distances between autonomous cars at high speeds; procedures for certifying an autonomous race car as "race-ready" for the track, analogous to production AV roadworthiness on public roads. Accordingly, lessons from racing may inform departments and bureaus of motor vehicles, licensing offices, and transport authorities as they develop minimum vehicle performance standards (sensing and V2X capabilities) and protocols for technical inspections that will be required of production AVs seeking registration. This is needed before AVs can be allowed to use public highways either in mixed traffic or on dedicated AV lanes. These experiences might inform how licensing bureaus or transport authorities conduct inspections and certifications of self-driving cars.

Further, autonomous racing is playing an important role in the future standardization of components and protocols for autonomous mobility. The competitive yet collaborative nature of autonomous races is fostering the creation of an internal community of researchers and engineers that share best practices through various platforms (including journals, conferences, exhibitions, and the races themselves) where autonomous vehicle safety and standards are discussed. These forums promote the transfer of technologies honed through racing events towards deployment in the industry. The organizers of racing competitions continue to influence the standardization of autonomous vehicle protocols by engaging industry partners (sensor makers, automakers, research funding organizations, federal agencies, etc.). These efforts are accelerating interactions among stakeholders across the autonomous mobility ecosystem, including race teams, researchers, OEMs, IOOs, and public agencies. Such an ecosystem, because of the uniqueness of racing as a test bed, is causing coalescence of these stakeholders around autonomous racing and will be characterized by spillovers (at the very least) of knowledge and talent from these competition events into the broader AV industry.

10.3 Prospective Patterns of AV Deployment

Chapter 1 of this report introduced the issue of the incremental deployment of vehicle autonomy. Two dimensions are expected to characterize such deployment. First is market penetration, which is expected to grow in a manner consistent with the classic innovation adoption curve: gradually at first, then rapidly, followed by slow growth as it approaches market saturation. This will be driven by factors related to the four key stakeholders: consumer attitudes and preferences, and trust in technology (AV users); supportive regulation and policy (government); AV technology (OEMs); and availability of supporting infrastructure (IOOs). Second is the level of autonomy (LOA) (SAE International, 2021), specifically the highest LOA possessed by autonomous vehicles deployed on public roads, recognizing that the traffic stream will comprise vehicles of various levels of automation. Market penetration and levels of autonomy (MPLA) diagrams (Labi, 2019; Saeed et al., 2020) can be used to describe how AV demand and technological level may evolve across these two dimensions.

The extent to which autonomous racing will provide lessons for autonomous mobility and the specific modules of autonomous driving that will benefit most will likely be influenced by the prevailing market penetration and the highest level of autonomy: the most significant lessons are likely to emerge during the early phase of the adoption curve, particularly when intermediate-to-high SAE automation levels begin to dominate the AV fleet on public roads. In addition, lessons related to the perception module are expected to be particularly impactful at this MPLA coordinate.

In this chapter, we discuss the concept of the AV transition period, its phases, and the issues related to the most sought-after outcome of automated transportation: traffic safety. These are important considerations in the overall fabric of the narrative that describes ways in which autonomous racing lessons could benefit autonomous mobility.

10.4 The AV Transition Period and its Phases

Researchers have argued that full autonomous operations will emerge on public roads incrementally rather than spontaneously, as customer demand, supportive regulation, technological readiness, and infrastructure readiness gradually align. This period is often referred to as the AV Transition Period, as illustrated in Figure 10.1. The attributes of the transition period (start year and demand pattern of AV deployment, year of full autonomous operations) will influence the extent to which lessons learned from autonomous racing are put into practice. During this transition period, it is anticipated that roadways will host a mixed traffic stream (that is, a stream comprising traditional vehicles (human driven or Level 0 autonomy), automated vehicles (Levels 1–3), and fully autonomous or FAV vehicles (Levels 4–5)) until a time when all vehicles on the road are fully autonomous. The scale and intensity of infrastructure modifications to support AV operations will influence (and will be influenced by) the market penetration and the highest prevailing level of autonomy. Researchers have suggested the following phases of the HDV-FAV transition period:

- Phase I: low FAV-HDV ratio: up to 25% of the traffic stream vehicles are FAV.
- Phase II: low-to-medium FAV-HDV ratio: 25-50% of the traffic stream vehicles are FAV.
- Phase III: mid-to-high FAV-HDV ratio, 50%-75% of traffic stream vehicles are FAV.
- Phase IV: high HDV-FAV ratio, 75%-almost 100% of traffic stream vehicles are FAV.

  The fully autonomous phase (FAP) is where 100% of traffic stream vehicles are FAVs.

Further, within these phases, there could exist sub-phases depending on the prevailing LOA distribution. Also, it is expected that the market penetration of Level 4 and Level 5 AVs will be low during the incipient years of the transition period (as AVs continue to be tested for commercial use and learn lessons from field validations, including autonomous racing), and the traffic stream will be dominated by Levels 1-2 AVs. Then AV, and, ultimately, FAV market penetration will increase gradually until they dominate the traffic stream.

With fully autonomous AV service already operating in deployment in Phoenix and selected other cities, it may be argued that the AV sunrise year has already begun. But when will HDV operations on public roads end? This answer will likely vary by road corridor rather than by the public road network. Litman (2014) suggested that after 2060, human driving will face restrictions if AV benefits are widely realized. An IHS Automotive study stated that by 2050, nearly all vehicles in the traffic stream will be expected to be autonomous (Diana, 2014). Kyriakidis et al. (2015) cited other studies that predicted high levels of AV market share by 2030.

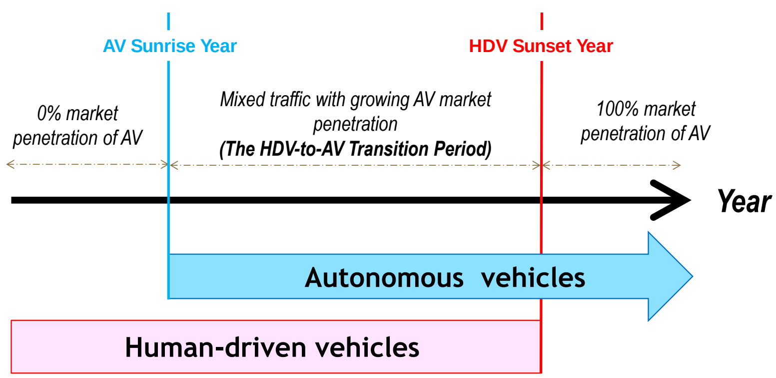

Figure 10.1 Projected (conceptual) timeline of HDV-only traffic, mixed traffic, and AV-only traffic distribution growth (adapted from Labi, 2019)

Other researchers have recognized the significant uncertainty associated with the expected length of the transition period, AV market penetration growth, and the level of automation prevailing at each year of the transition period. Such uncertainty, reflected in the significant variability in their prognostications, could be attributed to variations in user trust in automation, governmental support through policy, the willingness and capability of road agencies and other IOOs to provide the needed physical and cyber infrastructure, and OEMs' technological advancement of AV modules and systems (which is influenced by the efficacy of field validations including autonomous racing).

The transition to full autonomy is expected to be evolutionary and gradual, occasionally punctuated with technological, policy or consumer related hiccups (Labi et al., 2023). From the perspective of IOOs, trends in AV market penetration and levels of autonomy will be watched, as MPLA transitions will serve as the basis for investing in AV-supporting infrastructure (physical and cyber). Retrofitting existing infrastructure to serve AVs will need to be proactive and responsive. However, because of the nature of public-sector infrastructure project development processes and the funding limitations IOOs often face, roadway retrofits will likely be incremental, stepwise, and possibly slow. Currently, roadways are designed to serve HDVs. IOOs will need to develop long-term infrastructure plans which recognize that with increasing AV market penetration and levels of autonomy, there will be a need for physical and cyber infrastructure that can serve simultaneously a mixed vehicle stream comprising HDVs, automated vehicles (SAE Levels 1–3), autonomous vehicles (Levels 4 and 5) during the transition period, and eventually serve a fully autonomous fleet in the post-transition period.

10.5 The Role of Autonomous Racing on AV Safety Assurance

Autonomous racecars do not carry passengers; however, the lessons they provide are highly relevant to production AVs for which traffic safety remains the central motivation. Safety literature has long suggested that human error is a contributory factor in most traffic crashes and fatalities including Wierwille et al. (2002): 92% of accidents (Treat et al., 1979); 90–95% of accidents (Rumar, 1985); 45–75% (Hankey et al., 1999; Aberg & Rimmo, 1998); 75% of all crashes (Salmon et al., 2005). FHWA (1995) reported that only 7% of crashes are attributed to factors other than the human driver. Macadam (2003) and Gordon & Lidberg (2015) stated that most driver-attributed crashes are due to human driver attributes, including the driver's state of sobriety, extreme age, distractedness (including texting while driving), fatigue, inexperience, and poor vision.

Researchers expect that autonomous vehicles will significantly reduce crash frequency and severity and have found that the safety benefits of AVs are already being realized at even low levels of autonomy (Chen et al., 2020). For example, between 2008 and 2010, electronic stability control technology helped prevent over 2,200 fatalities (Sivinski, 2011). It is expected that the various modules of autonomous driving will continue to advance technologically to a level at which the AV will be capable of reliably handling anomalous traffic situations and irregular roadway conditions, including work zones, poor weather, and unexpected obstacles on the roadway.

It may be argued, albeit counter-intuitively, that during the early stages of the AV transition period as more AVs get deployed, the absolute number of AV-involved crashes may increase, even if crash rates decline. Extensive coverage of AV-involved crashes may exacerbate existing apprehensions about AVs, and consequently, policymakers may be reluctant to pass AV-supporting policies, thereby lengthening the AV transition period.

Autonomous racing plays a vital role in AV safety assurance. In the early 20th century, when automobiles had just been invented, automakers used races to help convince a skeptical public that automobiles were safe, and car racing served as a showcase not only for their cars but also for improvements in specific safety-related car components, including brakes and tires (Cavalier Autonomous Racing, 2025). More than a century later, society finds itself in a somewhat analogous position: as we find ourselves in the doorway to autonomous mobility, autonomous racing events offer a real-life laboratory for all stakeholders to witness the capability of AV technology to handle situations much more extreme compared to the most extreme roadway situations. It is impressive to see AVs traveling at 75 mph on a car manufacturer's test track. However, seeing multiple autonomous vehicles navigate a track at 150 mph without incident can have a far more profound effect on observers, sending the clear message that autonomy can be (and is) safe. Each successful race (with no accidents) is evidence that the systems can be dependable. This can ease the concerns of policymakers, regulators, and the public, potentially accelerating AV market penetration and promotion of AV-friendly regulation and policies. Cavalier Autonomous Racing (2025) defined "agile safety" as using all the full capacity of the vehicle to avoid a crash, as a direct contribution of autonomous motorsports to road-safety engineering.

AV-related crashes tend to attract significant media attention, and when they occur, they exacerbate existing public skepticism and erode public trust in AV safety. As such, AV regulators and customers continually seek safety assurance. Such assurance is achieved through AV testing on in-service roads or dedicated road networks (both of which are limited) and dedicated non-public test tracks. Autonomous racing events provide a valuable testing platform and environment in which enthusiastic race teams can validate novel sensors and algorithms across the various modules of autonomous driving. For example, several vendors of perception-related hardware have attested to the valuable role autonomous racing plays in this regard. High-speed competitions force AV systems to achieve higher levels of safety performance than ordinary driving, yielding a margin of safety and capability that can trickle down to real-world autonomous operations. The success of autonomous racecars in making split-second decisions at 150 mph on the racetrack suggests significant potential for improving the performance of production AVs operating on high speed expressways and freeways. If an AI system can safely pilot a car in the dynamic environment of a race, it may be better prepared to handle the comparatively more structured environment of public roads (Roles, 2025).

In a future in which AVs command a dominant share of the traffic stream, there will be a short period characterized by a jump in crash rate as both vehicle classes (HDVs and CAVs) "learn" how to adjust to each other. Labi et al. (2023) argued that after such a period of learning and adjustment, both classes of vehicles will be more cognizant of the limitations of the other and be accommodative of such limitations, leading to a likely reduction in crash rates. Also, with time, as the mobility and safety benefits of AVs become apparent, public skepticism toward AVs is expected to subside, which could help shorten the AV transition period.

10.6 The SAE level of autonomy (LOA) associated with autonomous racing

In autonomous racing, specially designed vehicles travel along a guideway also specially designed for these vehicles. There are two opposing schools of thought on whether the autonomous racecar is best described as a Level 4 or a Level 5 vehicle. Figure 10.2 presents the SAE J3016TM levels of driving automation (SAE, 2021), which provide the framework for interpreting the autonomy level associated with autonomous racing. A related question is whether the LOA of autonomous racing should be based on the vehicle features or on the nature of the guideway? It may be more appropriate to describe the racing environment, rather than the vehicle alone (that is, the vehicle-guideway system) as Level 4 operations or Level 5 operations? A review of standard autonomy documentation (SAE, 2021; NHTSA, 2022) suggests that two commonly invoked criteria in arguments about Level 5 operations: (a) the absence of any possibility of human intervention because there is no steering wheel or speed-control pedal, and (b) the capability to operate everywhere and under all conditions. There is a school of thought that the defining criterion is not the absence of a steering wheel or pedals but rather the vehicle's operational capability in each environment. Proponents of this notion posit that the autonomous racecar cannot operate everywhere (e.g., in inclement weather, local streets, and rough dirt terrain) and therefore cannot be labeled Level 5. The second school argues that autonomous racecars are consistent with Level 5 for two reasons: first, the autonomous racecar has no steering wheel or pedals; second, it can operate on any guideway designed for any racing operations and therefore can be described as Level 5. By a similar reasoning, an autonomous buggy could be described as a Level 5 buggy if it can travel on guideways where a human-driven buggy can travel (i.e., rough dirt terrain). Therefore, the inability of a Level 5 vehicle to operate on a guideway for which it was not designed does not invalidate its status as Level 5.

Ultimately, the classification of an autonomous racecar as Level 4 vs. Level 5 may depend on the perspective of the reader: intervention capability vs. host guideway/environment capability. During the race, the autonomous racecar generally operates at what many observers would regard as the highest level of autonomy (Level 5), because there is no capability for human intervention and control once the race begins. However, if the analyst considers true Level 5 as the capability to drive anywhere, on any guideway type, and under any environmental conditions, then from their perspective, the autonomous racecar does not pass their Level 5 litmus test, which would place it closer to Level 4 under that interpretation.

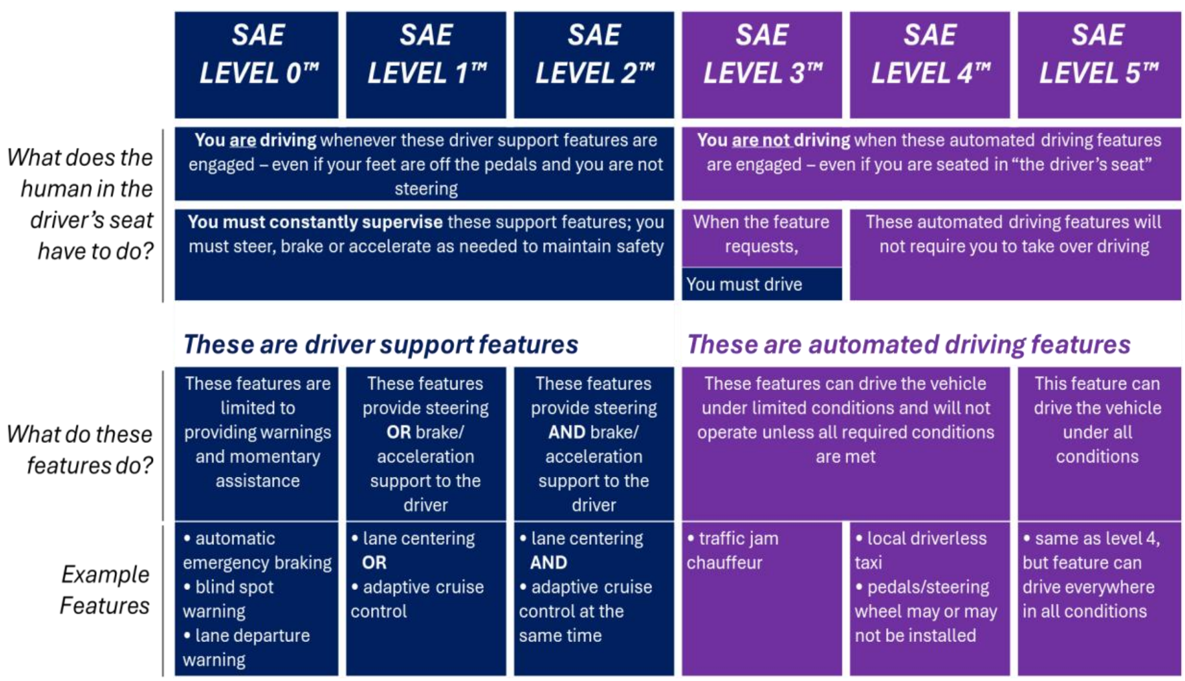

Figure 10.2 SAE J3016TM Levels of Driving Automation

10.7 Some reflections

This chapter examined the AV transition period, its phases, and the safety issues associated with automated transportation, while also highlighting the broader institutional and societal benefits of autonomous racing. During the AV transition period, the HDV sunset year, and the pattern of AV growth (in terms of market penetration and the maximum or dominant level of autonomy) will be influenced by variations associated with the four key stakeholders in AV deployment. As such, the transition period will be influenced by positive and negative stimuli that may emanate from the stakeholders, including OEMs, IOOs, customers, regulators, and policy makers. In addition, the initial stages of the transition period will be marked by a jump in crash rates as HDVs and CAVs learn to adjust to each other. After this initial adjustment period, crashes are expected to decline as AVs account for a larger share of the traffic stream and fewer vehicles are driven by humans. In seeking safety assurance for AVs at test sites, manufacturers have options, including in-service roads, AV-dedicated road networks, and AV test tracks. However, as an experimental test bed, autonomous racing may offer a more concentrated and potentially cost-effective validation environment than to the other test site types. Autonomous racing plays a vital role in AV safety assurance, as every accident-free race event is evidence of AV systems' reliability. This can help allay the concerns of policymakers, regulators, and the public, potentially accelerating AV market penetration and promotion of AV-friendly regulations and policies. The chapter also highlighted broader impacts of autonomous racing that extend beyond technical performance and contribute to the long-term development of autonomous mobility.

CHAPTER 11 CONCLUDING REMARKS

To make autonomous driving a reality, it is not sufficient to test algorithms in simulation or virtual reality. Testing them on a physical and realistic platform is of utmost importance." Fanta Camara, Institute for Safe Autonomy, University of York.

11.1 Summary

There currently are a limited number of international autonomous racecar competitions in existence. These may sometimes be perceived as being primarily for recreational purposes. However, their academic, industry, and research value are tremendous, as several significant lessons could be drawn from these competitions. This report presented and discussed lessons that high-speed driverless racing events and teams offer for the benefit of autonomous mobility (production AVs operating on public highways). The report started with a frequency and trend analysis of autonomous racing literature, showing the distribution of research papers across the various modules, over time, and across countries. The review showed that autonomous racing research started in 1989, growing gradually until 2017 after which it grew rapidly, with most publications in the controls topic area, followed by path planning, vehicle dynamics, and perception. The period 2019-2022 may be regarded as the golden years of autonomous racing research, with about 70 papers published annually; since then, the publication rate has declined. Most publications are from the United States, followed by Germany and Italy.

The report then presents lessons learned in vehicle dynamics and control. The chapter discusses the exceedance of the linear region of tire dynamics and the reasons for (and the common modes of) ADAS failure. The chapter discusses friction potential estimation (that is, the ability to estimate friction before a slide occurs) which can promote safety operations of the production AV in an era of autonomous mobility. By operating at physical limits, autonomous racing offers lessons on maintaining control in extreme situations (skids, sudden obstacles), and techniques such as autonomous drifting, emergency path diversion, spin recovery have been adopted into ADAS systems in a bid to reduce accidents on slippery or hazardous roads. Thus, the chapter showed how autonomous racing events provide lessons for production AVs on public roads in the contexts of emergency obstacle avoidance, longitudinal and lateral control, and hydroplaning and wet-weather management.

Regarding perception and sensing, the sensor suites installed on autonomous racecars are subjected to severe environmental stressors, including vibration, heat, and g-forces. The success of sensor installation techniques on high-speed race cars offers valuable lessons. Autonomous racing offers opportunities to assess degradation in sensor performance under these stressors, providing insights into sensor hardware design for production AVs to enhance their longevity and reliability. Racing teams have developed advanced de-blurring and de-noising algorithms and deployed global shutter cameras which capture the entire image simultaneously rather than line-by-line. According to the literature, these technologies are now becoming standard in high-end automotive ADAS packages to ensure lucid detection of cross-traffic features at highway speeds. Also, perception hardware and algorithms that maintain accuracy even at high speeds in autonomous racing has helped spawn, or at least validate, sensor-fusion algorithms and object detection systems for the benefit of production AVs.

The intertwined processes of decision making and planning enable an autonomous vehicle to navigate a trip efficiently and safely. By mastering overtaking on a racetrack and managing interactions with other vehicles, autonomous race cars have learned strategies for negotiating that can make production AV operations safer and more efficient. Regarding motion prediction of other vehicles (which is a key pre-requisite for tactical decision making), the autonomous racing competitions have yielded algorithms to predict the movements of neighboring vehicles and adjust the ego vehicle's path. Such models are directly applicable to advanced driver-assistance systems used in production vehicles. AI-based approaches use learning techniques and past data to optimize decisions, but they require massive volumes of expensive-to-acquire training data. There is very limited real-world autonomous-driving data available in different conditions, particularly for edge cases (obviously due to passenger safety concerns). As such, autonomous vehicle racing events which are less constrained by conventional regulatory boundaries provide a unique opportunity to produce adequate training data that are needed to train the AI models. The ability to plan trajectories using game theory in aggressive, high-speed traffic environments and on freeways has clear safety implications in the era of autonomous mobility.

In this report, we identify a few areas related to data fusion, hardware requirements, and algorithm integrity where OEMs could learn from autonomous racecar experiences. We discuss the overarching concept of a software-defined vehicle and then show how autonomous racing provides valuable lessons regarding software robustness and reliability, latency, sensor fusion in denied environments, and thermal management of vehicle components, all of which promote the realization of SDVs. The report also presents the lessons learned in terms of the need for specific physical roadway infrastructure designs and cyber infrastructure, and the wider lessons to be learned from autonomous racing. The precision demonstrated by autonomous racecars offers a path toward more efficient roadways, including narrower lanes, optimized pavement wear patterns, and more machine-readable digital signage.

This report examines separate modules of autonomous racing in dedicated chapters: control, perception, location, etc. It is worth noting, however, that the latest trend in autonomous racing is characterized by a move toward holistic system design; algorithms are no longer developed in isolation. The field is progressing from making cars that can drive themselves to creating integrated robotic systems that can perceive, reason, and act within physical limits.

11.2 Conclusions

This study found that the high-speed AV competitions, where multiple teams push the limits of sensing, decision-making, and control, serve effectively as testbeds for developing and refining AV technologies. These racing environments are particularly advantageous for accelerating innovation because they offer a controlled environment where extreme conditions permit testing of edge-case scenarios and performance extremes that would be too risky or impractical on public roads.

Autonomous motor racing is therefore proving to be profoundly influential for mainstream road transportation. The racetrack experiences provide not only a valuable knowledge base for facilitating safe and efficient AV operations on high-speed road transportation corridors, including freeways, but also lessons for designing AVs, developing AV software, and planning physical and cyber infrastructure. These insights are particularly valuable as governments and road agencies work with industry to safely integrate AVs into the existing transportation system.

In recognizing these lessons, however, appropriate caveats must be made regarding the differences between AV operations guideways (racetracks vs. roadways) and AV vehicle designs (race car designs vs. standard automobile and truck design), and how these differences limit the direct translation of lessons from AV racing to real-world (roadway) AV operations. The ODD for autonomous racing differs significantly from that of everyday driving. As such, researchers such as Mar et al. (2025) have cautioned that not all challenges related to autonomous racecars translate directly to production AVs.

In sum, the autonomous motorsport laboratory remains one of the most effective crucibles for validating the robustness and reliability of technological components needed for the next generation of road transportation: connected, autonomous, and electric. The integration of racing-validated technologies (including predictive friction estimation, game-theoretic planning, and ruggedized perception) will be a defining factor in transitioning autonomous high-speed vehicles from experimental novelty toward a new world of autonomous mobility.

The outcome of this research is a reinforced the understanding that autonomous racing offers valuable lessons for autonomous mobility. This outcome could have significant social (safety) and economic (mobility and productivity) implications in the prospective era of AV operations on public roads. The lessons learned can help not only in enhancing AV design, sensing, and communication capabilities for roadway operations, but also enhance mobility, safety, reliability, and economic productivity of AV operations at high-speed road corridors, including freeways.

11.3 Limitations of the Present Study

Despite all the valuable lessons learned from autonomous racing, the reality is that test tracks do not always reflect adequately the public road driving environment. This issue has been recognized by a few researchers including Doubek et al. (2021), Jiang (2021), and Mar et al. (2025). A racetrack is a specially designed course, often oval, featuring marked start/finish lines, safety features, and spectator areas, functioning as a venue for competition. Examples of oval courses that have hosted autonomous racing competitions, as shown in Figure 8.1, include the Indianapolis Motor Speedway (IMS), the Lucas Oil Speedway, and the Monza racetrack. While oval racing courses offer some infrastructure lessons, particularly at high-speed sections (for example, banked angles and transition clothoid design parameters), they still represent some departure from real roads. As such, the lessons of vehicle dynamics may not translate to autonomous mobility to the extent offered by the other modules of autonomous racing. Road racing courses, on the other hand, such as that of Putnam Park, Indiana, are much closer to the engineering design of real roads. Given the limitations of race ovals and road courses in mimicking real road conditions, driving simulators continue to serve as a useful platform for testing AVs (Chen et al., 2020; Camara, 2022). Secondly, in most autonomous racing events of today, each car operates independently (they are competitors, not collaborators). Further, in head-to-head racing, there exist adversarial interactions, far from collaboration. In contrast, in the real world, there is significant collaboration between HDVs as they give way to others to maintain road etiquette, at least in certain cultures. In the future of autonomous mobility, this is expected to continue, primarily as a deliberate effort to optimize systemwide traffic performance as part of a larger citywide V2X ecosystem managed by a traffic control center.

CHAPTER 12 SYNOPSIS OF PERFORMANCE INDICATORS

12.1 Part I of USDOT Performance Indicators

Over the study period for this project, one (1) transportation-related graduate course and four (4) transportation-related undergraduate courses were offered annually that were taught by the PI. During the study period, one (1) transportation-related advanced degree (doctoral) program and one (1) transportation-related M.S. program utilized the CCAT grant funds from this research project to support four (4) undergraduate students and three (3) graduate students.

12.2 Part II of USDOT Performance Indicators

Research Performance Indicators:
Four (4) journal publications and two (2) conference presentations related to the study topic were produced from this project. The findings of this research project were disseminated to 60 attendees (from industry, government, and academia) through two conferences.

Leadership Development Performance Indicators:
This research project generated 2 academic engagements and 9 industry engagements. The PI held positions in 2 national organizations that address issues related to this research project.

Education and Workforce Development Performance Indicators:
The lessons from this study were incorporated (or are being incorporated) in the syllabi for the Spring 2025 and Fall 2025 versions of the following courses at Purdue University:
(a) 3 VIP courses: Three (3) vertically integrated program courses under the auspices of the College of Engineering. (average 15 students at each course offering).

(b) CE 398: Introduction to Civil Engineering Systems, a mandatory undergraduate-level course at Purdue University's civil engineering program, (average 85 students at each course offering).

(c) CE 597: Emerging Transportation Technologies, a graduate-level course at Purdue University's civil engineering program, (average 8 students at each course offering).

These students will soon be entering the workforce. Thus, the research helped enlarge the pool of people trained to develop knowledge and utilize at least a part of the technologies developed in this research, and to put them to use when they enter the workforce. The methods, data, and/or results from this study will also be incorporated in future versions of the courses stated above.

The outputs, outcomes, and impacts are described in Chapter 13.

CHAPTER 13 STUDY OUTCOMES AND OUTPUTS

13.1 Outputs

13.1.1 Publications, conference papers, or presentations

(a) Publications

1. Seilabi, S.E., Pourgholamali, M., Miralinaghi, M., Correia, G., Labi, S. (2024). "Optimizing dedicated lanes and tolling schemes for connected and autonomous vehicles to address bottleneck congestion considering morning commuter departure choices," *Journal of Intelligent Transportation Systems*, doi.org/10.1080/15472450.2024.2408024
2. Chen, S., Zong, S., Chen, T., Huang, Z., Chen, Y., & Labi, S. (2023). "A taxonomy for autonomous vehicles considering ambient road infrastructure." *Sustainability*, 15(14), 11258. doi.org/10.3390/su151411258
3. Dong, J., Chen, S., Miralinaghi, M., Chen, T., Li, P., Labi, S. (2023). "Why did the AI make that decision? Towards an explainable artificial intelligence (XAI) for autonomous driving systems," *Transportation Research Part C: Emerging Technologies* 156, 104358, 10.1016/j.trc.2023.104358
4. Pourgholamali, M., Miralinaghi, M., Seilabi, S.E., Labi, S. (2023). "Sustainable deployment of autonomous vehicles dedicated lanes in urban traffic networks," *Sustainable Cities and Society* 99, 104969, doi.org/10.1016/j.scs.2023.104969

(b) Conference (c) Presentations

1. Gan, J., Moreno, A., Fasate, J., Ajagu, R., Labi, S. (2025). System-level dynamics of a remote-controlled vehicle: A black box approach to drivetrain modeling, Summer Undergraduate Research Symposium, Purdue University, West Lafayette, July 30, 2025.
2. Seilabi, S.E., Pourgholamali, M., Miralinaghi, M., Correia, G., He, Xiaozheng, Labi, S. (2024). "Managing dedicated lanes for connected and autonomous vehicles to address bottleneck congestion considering morning commuter departure choices," Paper Nr. 24-20938, *Annual Meeting of the Transportation Research Board,* Washington, DC.

13.1.2 Other outputs

(a) Editorial of a Special Issue of *Fundamental Research* Journal:

Huang, J., Labi, S., Kulcsar, B.A., Monreal, C.O., Wu, F., He, Z., Qu, X. (2024). "Editorial: Intelligent vehicles and smart transportation," *Fundamental Research*, 4(5), 979–980. doi: 10.1016/j.fmre.2024.08.004

(b) The research outcome (documentation of the lessons learned from AV racing) is being used in Purdue University's undergraduate courses (see Appendix A3) and graduate courses related to autonomous transportation:

CE 398 (Civil Engineering Systems),

VIP ARES (Autonomous Race-Engineering Systems)

VIP AMP (Autonomous Motorsports Purdue)
VIP ROBOAT (Autonomous Maritime Maneuvers (formerly NSWC AIMM))
CE 597 (Emerging Transportation Technologies: Automated, Connected, Shared, Electric, Airborne)

(c) Collaboration Partners:

Purdue AI Racing
Autonomous Karting Series (AKS): Andrew Goeden
Motorsports Consulting: Danny White
Us Navy and Artificial Intelligence Maritime Maneuver Indiana Collegiate Challenge (AIMM-ICC): Brittany Newell, Richard Volyles, Tim Murphy, Jason Thume, Amy Ruggles.

No patents have been filed for the research outcomes.

13.2 Outcomes

This project produced outcomes that could influence road agencies' transportation infrastructure preparedness for the era of full autonomous mobility. The study helped to increase understanding and awareness of the lessons that autonomous racing provides for an emerging future of autonomous mobility, in the areas of vehicle localization, sensing and perception, path planning and decision making, control, hardware and software safety, and physical and cyber infrastructure preparations. This also includes other wider lessons, such as public awareness and trust in automation, workforce development, and standardization of autonomous system components, traffic protocols, and certifications.

13.3 List of impacts

The impacts of this project are expected to include enhanced overall preparations towards the emerging future of autonomous mobility, particularly, at high-speed expressways and freeways. These impacts may be manifested through the effects on the transportation system or society in general, such as increased user trust in automation, reduced fatalities and increased travel efficiency. The research outcomes can increase the body of knowledge on autonomous mobility technologies, expand the pool of people trained to develop knowledge on autonomous mobility, and, overall, help improve the physical, institutional, and information resources associated with preparations toward a future of autonomous mobility. A list of specific impacts from this research project are:

- Stronger justification for highway agencies to collaborate with autonomous racing organizers to learn lessons provided by these competitions
- Enhanced motivation for highway agencies to make physical and cyber infrastructure investments as part of preparations towards the emerging era of autonomous mobility. This can benefit society through social, economic, and environmental dimensions.
- The undergraduate and graduate students who worked on this project will enter the workforce in 2026 and 2027 to help support the workforce associated with the development of autonomous mobility. One of the students who worked on this project is currently interning at TESLA.

- The project impacted on education, as parts of the research material were incorporated (to varying extents) into four undergraduate courses and one graduate course at Purdue University. The students in these courses, who have begun entering the workforce, benefited from the outcomes of this research through the classes they took. This helps to increase the size and quality of the potential workforce that could make use of the knowledge acquired through this project, for the benefit of society and the economy.

REFERENCES

AASHTO. (2017). Infrastructure needs for autonomous vehicles. Accessed May 16, 2025. https://www.tam-portal.com/project/infrastructure-needs-for-autonomous-vehicles/.

AASHTO (2018). A policy on geometric design of highways and streets, 7th Ed., American Association of State Highway and Transportation Officials, Washington, DC.

Abdulmaksoud, A., & Ahmed, R. (2025). Transformer-based sensor fusion for autonomous vehicles: A comprehensive review, in IEEE Access, 13, 41822-41838.

Aberg, L., & Rimmo, P. (1998). Dimensions of aberrant driver behavior. Ergonomics, 41(1), 39–56.

Alcala, E., Puig, V., & Quevedo, J. (2020a). LPV-MP planning for autonomous racing vehicles considering obstacles. Robotics and Autonomous Systems 124, 103392.

Alcala, E., Puig, V., Quevedo, J., & Rosolia, U. (2020b). Autonomous racing using linear parameter varying-model predictive control (LPV-MPC). Control Engineering Practice, 95, 104270.

Al-Sunni, M. F., Almubarak, H., Horng, K., & Dolan, J. M. (2025). LLA-MPC: Fast adaptive control for autonomous racing. 2025 IEEE/RSJ International Conf. on Intell. Robots & Systems, 1969-1976.

Acharya, S., Mekker, M. (2022). Importance of the reputation of data manager in the acceptance of connected vehicles, Communications in Transportation Research Vol. 2, 100053

Adler, N., Brudner, A., Proost, S. (2021). A review of transport market modeling using game-theoretic principles, European Journal of Operational Research 291(3), 808-829.

ADLINK (2021). Rise to the Indy Autonomous Challenge, Solution Brief, ADLINK Technology, material.adlinktech.com/ADLinkFile/Publication/1549/Solution_Brief-Rise_to_the_Indy_Autonomous_Challenge_05-07-2021.pdf, accessed on December 17, 2025,

Adamy, J. (2022). Nonlinear control of linear systems. Nonlinear Systems and Controls, 283-342.

Agnihotri, A., O'Kelly, M., Mangharam, R., & Abbas, H. (2020). Teaching autonomous systems at 1/10th-scale. Proceedings, 51st ACM Technical Symp. on Comp. Science Education, 657-663.

APTIV (2020). What Is a Software-Defined Vehicle? Inside Edge, March 19, 2020 Edition, https://www.aptiv.com/en/insights/article/what-is-a-software-defined-vehicle

Alrifaee, B., Scheffe, P., Kloock, M., & Henneken, T. M. (2021). Sequential convex programming methods for real-time trajectory optimization in autonomous vehicle racing. IEEE Trans. on Control Systems Technology.

Alrifaee, B., & Maczijewski, J. (2018). Real-time trajectory optimization for autonomous vehicle racing using sequential linearization. 2018 IEEE Intelligent Vehicles Symposium (IV), 476-483.

Alves, J. P., Fonseca Ferreira, N. M., Valente, A., Soares, S., & Filipe, V. (2019). Autonomous driving car competition. Advances in Intelligent Systems and Computing, 356-363.

Anderson, J. R., Ayalew, B., & Weiskircher, T. (2016). Modeling a professional driver in ultra-high performance maneuvers with a hybrid cost MPC. 2016 American Control Conf., 1981-1986.

Andrei, N. (2022). Sequential quadratic programming. springer optimization and applications, 521-567. Andresen, L., Brandemuehl, A., Honger, A., Kuan, B., Vodisch, N., Blum, H., Reijgwart, V., Bernreiter, L., Schaupp, L., Chung, J. J., Burki, M., Oswald, M. R., Siegwart, R., & Gawel, A. (2020). Accurate mapping and planning for autonomous racing. 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 4743-4749.

Arslan, O., Berntorp, K., & Tsiotras, P. (2017). Sampling-based algorithms for optimal motion planning using closed-loop prediction. 2017 IEEE International Conf. on Robotics & Automation (ICRA).

Arab, A., and Yi, J. (2021). Instructed reinforcement learning control of safe autonomous J-Turn vehicle maneuvers, Proceedings, 2021 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM), Delft, Netherlands, July 12-16, 2021, IEEE Press.

Atkins, H. (2022). 10 facts about the evolution of motorsport, History Hit, https://www.historyhit.com/facts-about-the-evolution-of-motorsport/

Atreya, P., Karnan, H., Sikand, K. S., Xiao, X., Rabiee, S., & Biswas, J. (2022). High-speed accurate robot control using learned forward kinodynamics and non-linear least squares optimization. 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 11789-11795.

Autosport (2025). A2RL to return for Season two at Yas Marina Circuit. www.autosport.com/general/news/a2rl-to-return-for-season-two-at-yas-marina-circuit/10724680/

Babu, V. S., & Behl, M. (2020). f1tenth.dev - An Open-source ROS based F1/10 autonomous racing simulator. 2020 IEEE 16th International Conf. on Automation Sci. & Eng. (CASE), 1614-1620.

Backus, R. (2004). The great race of 1878, Farm Collector, https://www.farmcollector.com/steam-traction/the-great-race-of-1878/?PageId=2

Bak, S., Betz, J., Chawla, A., Zheng, H., & Mangharam, R. (2022). Stress testing autonomous racing overtake maneuvers with RRT. 2022 IEEE Intelligent Vehicles Symposium (IV), 806-812.

Banjanovic-Mehmedovic, L., Halilovic, E., Bosankic, I., Kantardzic, M., Kasapovic, S., (2016). autonomous vehicle-to-vehicle (V2V) decision making in roundabout using game theory. Int J Adv Comput Sci Appl. 7:292–8. doi:10.14569/ijacsa.2016.070840

Baumann, N., Ghignone, E., Kühne, J., Bastuck, N., Becker, J., Imholz, N., Kränzlin, T., Lim, T. Y., Lötscher, M., Schwarzenbach, L., Tognoni, L., Vogt, C., Carron, A., & Magno, M. (2024). ForzaETH Race Stack—Scaled autonomous head-to-head racing on fully commercial off-the-shelf hardware. Journal of Field Robotics. https://doi.org/10.1002/rob.22429

Behl, M (2025). How can autonomous racing make autonomous driving safer? The Cavalier Autonomous Racing Team, https://autonomousracing.dev/research

Bell, J., Donikian, V., & Washington, G. (2020). Modeling and design of an autonomous dragster. SAE Technical Paper Series, 1.0.

Bertogna, M. (2022). Motion planning and control for multi vehicle autonomous racing at high speeds. 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC).

Betz, J., Heilmeier, A., Wischnewski, A., Stahl, T., & Lienkamp, M. (2019a). Autonomous driving—A crash explained in detail. Applied Sciences, 9(23), 5126.

Betz, J., Wischnewski, A., Heilmeier, A., Nobis, F., Stahl, T., Hermansdorfer, L., & Lienkamp, M. (2019b). A software architecture for an autonomous racecar. 2019 IEEE 89th Vehicular Technology Conference (VTC2019-Spring), 1-6.

Betz, J., Wischnewski, A., Heilmeier, A., Nobis, F., Hermansdorfer, L., Stahl, T., Herrmann, T., & Lienkamp, M. (2019c). A software architecture for the dynamic path planning of an autonomous racecar at the limits of handling. 2019 IEEE Intern. Conf. on Connected Vehicles & Expo, 1-8.

Betz, J., Zheng, H., Liniger, A., Rosolia, U., Karle, P., Behl, M., Krovi, V., Mangaram, R. (2022a). Autonomous Vehicles on the Edge: A survey on autonomous vehicle racing, IEEE Open Journal of Intelligent Transportation Systems, 3, 458-488.

Betz, J., Zheng, H., Liniger, A., Rosolia, U., Karle, P., Behl, M., Krovi, V., Mangram, R. (2022b). Autonomous vehicles on the edge: A survey on autonomous vehicle racing, in IEEE Open Journal of Intelligent Transportation Systems, 3, 458-488.

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. (2023a). TUM autonomous motorsport: An autonomous racing system. Journal of Field Robotics 40, 783–809.

Betz, T., Karle, P., Werner, F., & Betz, J. (2023b). An analysis of software latency for a high-speed autonomous race car—A Case Study in the Indy Autonomous Challenge. SAE International Journal of Connected and Automated Vehicles, 6(3), 283-296.

Betz, J., Betz, T., Fent, F., & Wischnewski, A. (2023c). TUM autonomous motorsport: An autonomous racing software for the Indy Autonomous Challenge. Journal of Field Robotics. https://doi.org/10.1002/rob.22153

Bevilacqua, M., Tsourdos, A., & Starr, A. (2017). Particle swarm for path planning in a racing circuit simulation. 2017 IEEE International Instrumentation and Measurement Tech. Conference 1-6.

Bhargav, J., Betz, J., Zehng, H., & Mangharam, R. (2022). Deriving spatial policies for overtaking maneuvers with autonomous vehicles. 2022 14th International Conference on Communication Systems and Networks (COMSNETS), 859-864.

Biswas, S., Anavatti, S. G., Garratt, M. A. (2021). Multi-objective mission route planning problem: a neural network-based forecasting model for mission planning, in IEEE Transactions on Intelligent Transportation Systems 22(1), 430-442.

BMW Group (2025). BMW and Formula Student Germany, accessed on December 17, 2025

Bolla, D. R., Singh, S., & Sarojadevi, H. (2022). Intelligent miniature autonomous vehicle. Lecture Notes in Networks and Systems, 701-713.

Bonab, S. A., & Emadi, A. (2019). Optimization-based path planning for an autonomous vehicle in a racing track. IECON 2019 - 45th Annual Conf., IEEE Industrial Electronics Society, 3823-3828.

Bosch (2025). The software defined vehicle. Bosch Mobility, https://www.bosch-mobility.com/en/mobility-topics/software-defined-vehicle/

Bosello, M., Tse, R., & Pau, G. (2022). Train in Austria, Race in Montecarlo: Generalized RL for Cross-Track F1^{tenth} LIDAR-Based Races. 2022 IEEE 19th Annual Consumer Communications &; Networking Conference (CCNC), 290-298.

Braghin, F., Cheli, F., Melzi, S., & Sabbioni, E. (2008). Race driver model. Computers and Structures, 86(13-14), 1503-1516.

Brown, M., & Gerdes, J. C. (2020). Coordinating tire forces to avoid obstacles using nonlinear model predictive control. IEEE Transactions on Intelligent Vehicles, 5(1), 21-31.

Brunner, M., Rosolia, U., Gonzales, J., & Borrelli, F. (2017). Repetitive learning model predictive control: An autonomous racing example. IEEE 56th Annual Conf. on Decision & Control.

Buckman, N., Hansen, A., Karaman, S., & Rus, D. (2022). Evaluating autonomous urban perception and planning in a 1/10th scale MiniCity. Sensors, 22(18), 6793.

Bulsara, A., Raman, A., Kamarajugadda, S., Schmid, M., & Krovi, V. N. (2020). Obstacle avoidance using model predictive control: an implementation and validation study using scaled vehicles. SAE Technical Paper Series, 1.0.

Burton, K., & Gajjar, D. (2024). Automated vehicles. Parliamentary Office of Science & Technology, UK Parliament POST. https://doi.org/10.58248/pb62.

Butz, M. V., & Lonneker, T. D. (2009). Optimized sensory-motor couplings plus strategy extensions for the TORCS car racing challenge. 2009 IEEE Symp. on Computational Intelligence, 317-324.

Buyval, A., Gabdulin, A., Mustafin, R., & Shimchik, I. (2017). Deriving overtaking strategy from nonlinear model predictive control for a race car. 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2623-2628.

Cacciaguerra, S., Ferretti, S., Roccetti, M., & Roffilli, M. (2005). Car racing through the streets of the web. Special interest tracks and posters of the 14th international conference, 884.

Cai, P., Mei, X., Tai, L., Sun, Y., & Liu, M. (2020). High-speed autonomous drifting with deep reinforcement learning. IEEE Robotics and Automation Letters, 5(2), 1247-1254.

Cai, P., Wang, H., Huang, H., Liu, Y., & Liu, M. (2021). Vision-based autonomous car racing using Deep imitative reinforcement learning. IEEE Robotics and Automation Letters, 6(4), 7262-7269.

Cai, L., Guan, H., Xu, Q. H., Jia, X., & Zhan, J. (2024). Game-theoretic decision-making method and motion planning for autonomous vehicles in overtaking. IEEE Transactions on Intelligent Transportation Systems, 25(8), 9693-9709.

Camara, F., & Fox, C. (2022). Unfreezing autonomous vehicles with game theory, proxemics, and trust, Front. Comput. Sci., Sec. Mobile and Ubiquitous Computing Vol. 4. https://doi.org/10.3389/fcomp.2022.969194

Campbell, J.L., Hoekstra-Atwood, L., Monk, C., Fraser, A.K., Torbic, D., & Potts, I. (2024). Diagnostic assessment and countermeasure selection: A toolbox for traffic safety practitioners, Washington, DC: The National Academies Press.

Cao, Y., Liao, Y., Lai, J., Shen, T., & Wang, X. (2024). Deviation characteristics of driving trajectories for autonomous vehicles and the design of dedicated lane widths. Sustainability, 16(21), 9155.

Caporale, D., Settimi, 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. 2019 International Conference on Robotics and Automation (ICRA), 5643-5649.

Caporale, D., Venturini, L., Fagiolini, A., Pallottino, L., Settimi, A., Biondo, A., Amerotti, F., Massa, F., Chatzikomis, C., Sorniotti, A., Gruber, P., Zanchetta, M., Willans, D., & Balcombe, B. (2018). Comparison of path tracking and torque-vectoring controllers for autonomous electric vehicles. IEEE Transactions on Intelligent Vehicles, 3(4), 559-570.

Camara, F. (2022) Inferring and operating pedestrian behavior models on autonomous vehicles. PhD Thesis, University of Leeds, Leeds, UK.

Caragata, S. (2022). Cisco returns as presenting sponsor of the Indy Autonomous Challenge 2022, Industries, blogs.cisco.com/industries/join-cisco-and-our-partners-at-the-indy-autonomous-

Cardamone, L., Loiacono, D., Lanzi, P. L., & Bardelli, A. P. (2010). Searching for the optimal racing line using genetic algorithms. Proceedings of the 2010 IEEE Conference on Computational Intelligence and Games, 388-394.

Carrau, J. V., Liniger, A., Zhang, X., & Lygeros, J. (2016). Efficient implementation of randomized MPC for miniature race cars. 2016 European Control Conference (ECC).

Cartró, J., Jackson, J., Sanchez, J., Fos, R., & Gutiérrez, J. (2024). Development of model predictive motion planning and control for autonomous vehicles. Proceedings, 230-246.

Carreras, A., Xavier, D., Erhart, J., Ruehrup, S., 2018. Road infrastructure support levels for automated driving, in: 25th ITS World Congress. 12–20.

Cataffo, V., Silano, G., Iannelli, L., Puig, V., & Glielmo, L. (2022). A nonlinear model predictive control strategy for autonomous racing of scale vehicles. 2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC), 100-105.

Cavalier Autonomous Racing (2025). How can autonomous racing make autonomous driving safer? University of Virginia, Charlottesville, VA, https://autonomousracing.dev/research

Cellina, M., Corno, M., Savaresi, S.M. (2025). Target tracking via LiDAR-RADAR sensor fusion for autonomous racing, https://arxiv.org/abs/2505.20043

Center for Sustainable Systems (2025). Autonomous Vehicles Factsheet. Pub. No. CSS16-18. University of Michigan. css.umich.edu/publications/factsheets/mobility/autonomous-vehicles-factsheet

Cerf, E. (2023). Researcher solves nearly 60-year-old game theory dilemma, UC Santa Cruz, https://news.ucsc.edu/2023/03/dejan-gametheory/

Charles, I., Maghsoumi, H., & Fallah, Y. (2025). Advancing autonomous racing: A comprehensive survey of the Roboracer (F1TENTH) Platform. 2025 6th International Conference on Artificial Intelligence, Robotics and Control (AIRC), 207-213.

Chen, J., Gao, H., Luo, Y., Yang, D., & Li, J. (2025). Efficient semantic trajectory similarity search. IEEE Internet of Things Journal, 12(2), 2219-2232.

Chisari, E., Liniger, A., Rupenyan, A., Van Gool, L., & Lygeros, J. (2021). Learning from simulation, racing in reality. 2021 IEEE International Conf. on Robotics & Automation (ICRA), 8046-8052.

Chamain, L. D., Qi, S., & Ding, Z. (2021). An end-to-end learning architecture for efficient image encoding and deep learning. 2021 29th European Signal Processing Conf. (EUSIPCO), 691-695.

Chatzopoulos, A., Papoutsidakis, M., Kalogiannakis, M., & Psycharis, S. (2019). Action research implementation in developing an open source and low cost robotic platform for STEM education. International Journal of Computer Applications, 178(24), 33-46.

Chen, S., Leng, Y., Labi, S. (2020). A deep learning algorithm for simulating autonomous driving considering prior knowledge and temporal information, Computer-Aided Civil and Infrastructure Engineering 35 (4), 305-321.

Chen, S., Zong, S., Chen, T., Huang, Z., Chen, Y., & Labi, S. (2023). A taxonomy for autonomous vehicles considering ambient road infrastructure. Sustainability, 15(14), 11258.

Chudzinski, M. (2023). How a 192-mph self-driving race car is accelerating autonomous technology, MotorTrend, www.motortrend.com/features/indy-autonomous-challenge-self-driving-racecars.

Corriou, J. P. (2017). Model Predictive Control. Process Control, 631-677.

Christ, F., Wischnewski, A., Heilmeier, A., & Lohmann, B. (2019). Time-optimal trajectory planning for a race car considering variable tire-road friction coefficients. Vehicle Sys. Dyn., 59(4), 588-612.

Colonna, P., Intini, P., Berloco, N., Ranieri, V. (2018). Connecting rural road design to automated vehicles: The concept of safe speed to overcome human errors, N.A. Stanton (Ed.), Advances in Human Aspects of Transportation, 571-582.

Coulter, R. C. (1992). Implementation of the pure pursuit path tracking algorithm (Technical Report CMU-RI-TR-92-01). Robotics Institute, Carnegie Mellon University. https://www.ri.cmu.edu/pub_files/pub3/coulter_r_craig_1992

DARPA (2014). The DARPA Grand Challenge: Ten Years Later, https://www.darpa.mil/news/2014/grand-challenge-ten-years-later

Dat, V. T., Bao, N. V. H., & Hung, P. D. (2025). HybridNets: End-to-end perception network. Pattern Recognition and Image Analysis, 35(2), 106-118.

de Bruin, T., Kober, J., Tuyls, K., & Babuska, R. (2018). Integrating state representation learning into deep reinforcement learning. IEEE Robotics and Automation Letters, 3(3), 1394-1401.

De Caro, S., & Corti, A. (2018). A planning and control system for self-driving racing vehicles. 2018 IEEE 4th International Forum on Research and Technology for Society and Industry (RTSI), 1-6.

De Rita, N., Aimar, A., & Delbruck, T. (2019). CNN-based object detection on low precision hardware: Racing car case study. 2019 IEEE Intelligent Vehicles Symposium (IV), 647-652.

de Winkel, K.N., Irmak, T., Happee, R., Shyroka, B. (2023). Standards for passenger comfort in automated vehicles: Acceleration and jerk, Applied Ergonomics 106 (2023) 103881

Demeter, Z., Hell, M., & Hajgató, G. (2024). Lessons learned from an autonomous race car competition. Engineering Proceedings, 79(1), 25.

Dennis, S., Petry, F., and Sofge, D. (2022). Game theory approaches autonomy. Front. Phys. 10:880706. doi: 10.3389/fphy.2022.880706

Dhall, A., Dai, D., & Van Gool, L. (2019). Real-time 3D traffic cone detection for autonomous driving. 2019 IEEE Intelligent Vehicles Symposium (IV), 494-501.

Diana, F. (2014). Autonomous vehicles: The automotive ecosystem, Frank Diana Blog, https://frankdiana.net/2014/03/03/autonomous-vehicles-

Dikici, O., Ghignone, E., Hu, C., Baumann, N., Xie, L., Carron, A., Magno, M., & Corno, M. (2025). Learning-based on-track system identification for scaled autonomous racing in under a minute. IEEE Robotics and Automation Letters, 10(2), 1984-1991. https://doi.org/10.1109/lra.2025.3527336.

Dong, J., Chen, S., Miralinaghi, M., Chen, T., Li, P., Labi, S. (2023a). Why did the AI make that decision? Towards an explainable artificial intelligence (XAI) for autonomous driving systems, Transportation Research Part C: Emerging Technologies 156, 104358

Dong, J., Chen, S., Labi, S. (2023b). Promoting CAV deployment by enhancing the perception phase of autonomous driving using explainable AI, CCAT Report #74, The Center for Connected and Automated Transportation, Purdue University, West Lafayette, IN.

Dong, J., Du, R., Ha, Y.J., Chen, S., Labi, S. (2022). Development of AI-based and control-based systems for safe and efficient operations of connected and autonomous vehicles, CCAT Report #41, The Center for Connected and Automated Transportation, Purdue University, West Lafayette, IN.

Doubek, F., Salzmann, F., & de Winter, J. (2021). What makes a good driver on public roads and race tracks? An interview study. Transportation Research F: Traffic Psych. & Beh. 80, 399-423.

Drage, T. H., Kalinowski, J., & Bräunl, T. (2014). Development of an autonomous formula SAE Car with Laser Scanner and GPS. IFAC Proceedings Volumes, 47(3), 2652-2657.

Drews, P., Williams, G., Goldfain, B., Theodorou, E. A., & Rehg, J. M. (2019). Vision-based high-speed driving with a deep dynamic observer. IEEE Robotics and Automation Letters, 4(2), 1564-1571.

Du, R., Ha, P.Y.J., Dong, J., Chen, S., Labi, S. (2022). Large network multi-level control for CAV and smart infrastructure: AI-based fog-cloud collaboration, CCAT Report #55, The Center for Connected and Automated Transportation, Purdue University, West Lafayette, IN.

Du, Z., Xie, H., Zhai, P., Yuan, S., Li, Y., Wang, J., Wang, J., & Liu, K. (2024). Game-based flexible merging decision method for mixed traffic of connected autonomous vehicles and manual driving vehicles on urban freeways. Applied Sciences, 14(16), 7375.

Eguchi, M., & Date, H. (2023). Dynamic obstacle avoidance for autonomous motorcycles on racing circuits by nonlinear model predictive control. 2023 23rd International Conference on Control, Automation and Systems (ICCAS), 1127-1132.

Eken, S., Şara, M., Satılmış, Y., Karslı, M., Tufan, M. F., Menhour, H., & Sayar, A. (2020). A reproducible educational plan to teach mini autonomous race car programming. International Journal of Electrical Engineering & Education, 57(4), 340-360.

Enisz, K., Szalay, I., Horváth, E. (2024). Localization robustness improvement for an autonomous race car using multiple extended Kalman filters. Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering, 239(9), 3771-3783.

Escoriza, A. L., Revach, G., Shlezinger, N., & van Sloun, R. J. G. (2021). Data-driven Kalman-based velocity estimation for autonomous racing. 2021 IEEE International Conference on Autonomous Systems (ICAS), 1-5.

Evans, B., Engelbrecht, H. A., & Jordaan, H. W. (2021a). Learning the subsystem of local planning for autonomous racing. 2021 20th International Conference on Advanced Robotics (ICAR), 601-606.

Evans, B., Engelbrecht, H. A., & Jordaan, H. W. (2021b). Reward signal design for autonomous racing. 2021 20th International Conference on Advanced Robotics (ICAR), 455-460.

Evans, B. D., Betz, J., Zheng, H., Engelbrecht, H. A., Mangharam, R., & Jordaan, H. W. (2023a). Bypassing the simulation-to-reality gap: Online reinforcement learning using a supervisor. 2023 21st International Conference on Advanced Robotics (ICAR), 325-331.

Evans, B.D., Jordaan, H.W., Engelbrecht, H.A. (2023b). Comparing deep reinforcement learning architectures for autonomous racing. Machine Learning with Applications, 14, 100496.

Evans, B. D., Jordaan, H. W., & Engelbrecht, H. A. (2023c). Safe reinforcement learning for high-speed autonomous racing. Cognitive Robotics, 3.0, 107-126.

Falanga, D., Kim, S., & Scaramuzza, D. (2019). How fast is too fast? The role of perception latency in high-speed sense and avoid. IEEE Robotics and Automation Letters, 4(2), 1884-1891.

Falb, P. (2019). Direct methods in control problems. IFAC-PapersOnLine. https://doi.org/10.1007/978-0-8176-4723-0.

Farah, H. Erkens, S.M.J.G. Alkim, T. van Arem, B. (2018). Infrastructure for automated and connected driving: State of the art and future research directions, G. Meyer, S. Beiker (Eds.), Road Vehicle Automation, 4, 187-197.

Fehér, A., Aradi, S., Bécsi, T. (2020). Hierarchical evasive path planning using reinforcement learning and model predictive control, IEEE Access, 8, 187470-187482.

Feng, Y., and Tarko, A. (2023). Development of a cooperative perception system, CCAT Report No. 70, Center for Connected & Automated Transportation, Purdue University, W. Lafayette, IN.

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. 2020 AEIT International Conference of Electrical and Electronic Technologies for Automotive (AEIT automotive). https://doi.org/10.23919/aeitautomotive50086.2020.9307439.

FHWA (Federal Highway Administration). (2018). National dialogue on highway automation. https://ops.fhwa.dot.gov/automationdialogue/. Accessed June 18, 2025.

FHWA. (1995). Interactive highway safety design model: Accident predictive module. Public Roads, 59(2).

Fiaformulae (2015). Formula E & Kinetik announce driverless support series. fiaformulae.com.

Fichtinger, A., Edelmann, J., Plochl M., and Holl, M. (2021). Aquaplaning detection using effect-based methods: an approach based on a minimal set of sensors, Electronic Stability Control, and Drive Torques, in IEEE Vehicular Technology Magazine, 16(3), 20-28, doi: 10.1109/MVT.2021.3085536

Fisac, J., Bronstein, E., Stefansson, E., Sadigh, D., Sastry, S., Dragan, A. (2019). Hierarchical game-theoretic planning for autonomous vehicles. In: 2019 international conference on robotics and automation, ICRA, 9590–9596.

Foris, A., Wagener, N., Boots, B., & Mazumdar, A. (2020). Exploiting singular configurations for controllable, low-power friction enhancement on unmanned ground vehicles. IEEE Robotics and Automation Letters, 5(2), 3556-3563.

Fox, C.W., Camara, F., Markkula, G., Romano, R.A., Madigan, R., Merat, N. (2018). When should the chicken cross the road? Game theory for autonomous vehicle-human interactions. In: Procs, 4th International Conference on Vehicle Technology & Intelligent Transport Systems, 429–31.

Franke, U. (2017). Autonomous driving. Computer Vision in Vehicle Technology, 24-54.

Fröhlich, L. P., Küttel, C., Arcari, E., Hewing, L., Zeilinger, M. N., & Carron, A. (2022). Contextual tuning of model predictive control for autonomous racing. 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 10555-10562.

Fu, M., Ni, J., Li, X., & Hu, J. (2018). Path tracking for autonomous race car based on G-G diagram. International Journal of Automotive Technology, 19(4), 659-668.

Fuchs, F., Song, Y., Kaufmann, E., Scaramuzza, D., & Durr, P. (2021). Super-human performance in Gran Turismo sport using deep reinforcement//doi.org/10.1109/lra.2021.3064284.

Funke, J., Brown, M., Erlien, S. M., & Gerdes, J. C. (2017). Collision avoidance and stabilization for autonomous vehicles in emergency scenarios. IEEE Transactions on Control Systems Technology, 25(4), 1204-1216.

Funke, J., Theodosis, P., Hindiyeh, R., Stanek, G., Kritatakirana, K., Gerdes, C., Langer, D., Hernandez, M., Muller-Bessler, B., & Huhnke, B. (2012). Up to the limits: Autonomous Audi TTS. 2012 IEEE Intelligent Vehicles Symposium.

Fursa, I., Fandi, E., Musat, V., Culley, J., Gil, E., Teeti, I., Bilous, L., Vander Sluis, I., Rast, A., & Bradley, A. (2022). Worsening perception: Real-time degradation of autonomous vehicle perception performance for simulation of adverse weather conditions. SAE International Journal of Connected and Automated Vehicles, 5(1), 87-100.

Gandhi, M. S., Vlahov, B., Gibson, J., Williams, G., & Theodorou, E. A. (2021). Robust model predictive path integral control: Analysis and performance guarantees. IEEE Robotics and Automation Letters, 6(2), 1423-1430.

Gao, Y., Shen, Y., Zhang, W., Wang, H., & Guvenc, L. (2018). Robust path tracking control for autonomous heavy vehicles. SAE Technical Paper Series, 1.0.
Garlick, S., & Bradley, A. (2022). Real-time optimal trajectory planning for autonomous vehicles and lap time simulation using machine learning. Vehicle System Dynamics, 60(12), 4269-4289.
Georgouli, K., Plati, C., Loizos, A. (2021). Autonomous vehicles wheel wander: Structural impact on flexible pavements, Journal of Traffic and Transportation Eng., 8(3), 388-398.
Giesen, J., Kuehne, L., & Laue, S. (2022). The GENO software Stack. Lecture Notes in Computer Science, 213-228.
Gifford, C. (2006). Racing: The Ultimate Motorsports Encyclopedia. Kingfisher.
Goblirsch, S., Weinmann, M., & Betz, J. (2024). Three-dimensional vehicle dynamics state estimation for high-speed race cars under varying signal quality. 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 3371-3378.
Goh, J., & Gerdes J. C. (2018). Drifting. The Night Chorus, 31-31.
Goh, J. Y., Goel, T., & Christian Gerdes, J. (2019). Toward automated vehicle control beyond the stability limits: drifting along a general path. Journal of Dynamic Systems, Measurement, and Control, 142(2). https://doi.org/10.1115/1.4045320.
Goh, J. Y., & Gerdes, J. C. (2016). Simultaneous stabilization and tracking of basic automobile drifting trajectories. 2016 IEEE Intelligent Vehicles Symposium (IV), 597-602.
Goldfain, B., Drews, P., You, C., Barulic, M., Velev, O., Tsiotras, P., & Rehg, J. M. (2019). AutoRally: An open platform for aggressive autonomous driving. IEEE Control Systems, 39(1), 26-55.
Gooding, R. (2020). Racing line. Electric and Hybrid Vehicle Technology International, 2020(2), 46-52.
Gordon, T. J., & Lidberg, M. (2015). Automated driving and autonomous functions on road vehicles. Vehicle System Dynamics, 53(7), 958–994.
Gopalakrishna, D., Carlson, P., Sweatman, P., Raghunathan, D., Brown, L., Serulle, N. (2015). Impacts of automated vehicles on highway infrastructure. Federal Highway Administration (FHWA)-HRT-21-015, Washington, DC.
Gosala, N., Buhler, A., Prajapat, M., Ehmke, C., Gupta, M., Sivanesan, R., Gawel, A., Pfeiffer, M., Burki, M., Sa, I., Dube, R., & Siegwart, R. (2019). Redundant perception and state estimation for reliable autonomous racing. 2019 International Conference on Robotics and Automation (ICRA). https://doi.org/10.1109/icra.2019.8794155.
Gosala, N., Gupta, M., Sivanesan, R., Dhall, A., Chisari, E., Karnchanachari, N., Brits, S., Dangel, M., Kapania, N. R., & Gerdes, J. C. (2020). Learning at the racetrack: Data-driven methods to improve racing performance over multiple laps. IEEE Transactions on Vehicular Technology, 69(8), 8232-8242.
Gotlib, A., Łukojć, K., & Szczygielski, M. (2019). Localization-based software architecture for 1:10 scale autonomous car. 2019 International Interdisciplinary PhD Workshop (IIPhDW), 7-11.
Guckiran, K., & Bolat, B. (2019). Autonomous car racing in simulation environment using deep reinforcement learning. 2019 Innovations in Intelligent Systems & Applications Conf., 1-6.
Guo, F., & Wu, Z. (2019). A deep reinforcement learning approach for autonomous car racing. Lecture Notes in Computer Science, 203-210.
Gundu, P. K., Vamvoudakis, K. G., & Gerdes, R. M. (2019). An intermittent learning algorithm for high-speed autonomous driving in unknown environments. 2019 American Control Conf., 4286-4292
Gungor, O. E. & Al-Qadi, I. L. (2020). Wander 2D: a flexible pavement design framework for autonomous and connected trucks. International Journal of Pavement Engineering, 1-16.
Hankey, J. M., Wielwille, W. W., Cannell, W. J., Kieliszewski, C. A., Medina, A., Dingus, T., & Cooper, L. M. (1999). Identification and evaluation of driver errors: Driver error taxonomy development (Report). Washington, DC: Virginia Tech Transportation Institute.
Ha, P., Chen, S., Du, R., Dong, J., Li, Y., Labi, S. (2020). Vehicle connectivity and automation: A sibling relationship, Frontiers in Built Environment 6, 590036, 1-12

Hafemann, P., Hahn, S. E., and Lienkamp, M. (2023). Optimizing autonomous vehicle sensor setups: A framework for coverage analysis, 2023 IEEE Symposium Sensor Data Fusion and International Conference on Multisensor Fusion and Integration (SDF-MFI), Bonn, Germany, 1-8, doi: 10.1109/SDF-MFI59545.2023.10361427.

Hang, P., Lv, C., Xing, Y., Huang, C., Hu, Z. (2020). Human-like decision making for autonomous driving: A noncooperative game theoretic approach. IEEE Trans Intell Transp Syst 22, 2076–2087. doi:10.1109/tits.2020.3036984

Hanover, D., Loquercio, A., Bauersfeld, L., Romero, A., Penicka, R., Song, Y., Cioffi, G., Kaufmann, E., & Scaramuzza, D. (2024). Autonomous drone racing: A survey. IEEE Transactions on Robotics, 40, 3044-3067.

Hartmann, G., Shiller, Z., & Azaria, A. (2022). Competitive driving of autonomous vehicles. IEEE Access, 10, 111772-111783.

He, L., Wang, Y., Gao, Z., & Chen, G. (2022a). A model predictive control for steady-state drifting and tracking of electronic-two-rear-wheel drive automobiles. IEEE Trans. on Intelligent Transportation Systems. https://doi.org/10.21203/rs.3.rs-1719854/v1.

He, S., Zeng, J., & Sreenath, K. (2022b). Autonomous racing with multiple vehicles using a parallelized optimization with safety guarantee using control barrier functions. 2022 International Conference on Robotics and Automation (ICRA), 3444-3451

Heilmeier, A., Graf, M., Betz, J., & Lienkamp, M. (2020a). Application of Monte Carlo methods to consider probabilistic effects in a race simulation for circuit motorsport. Applied Sciences, 10(12), 4229.

Heilmeier, A., Graf, M., & Lienkamp, M. (2018). A race simulation for strategy decisions in circuit motorsports. 2018 21st International Conference on Intelligent Transportation Systems (ITSC), 2986-2993. https://doi.org/10.1109/itsc.2018.8570012.

Heilmeier, A., Geisslinger, M., & Betz, J. (2019). A quasi-steady-state lap time simulation for electrified race cars. 2019 Fourteenth International Conference on Ecological Vehicles and Renewable Energies (EVER), 1-10. https://doi.org/10.1109/ever.2019.8813646.

Heilmeier, A., Thomaser, A., Graf, M., & Betz, J. (2020b). Virtual strategy engineer: Using artificial neural networks for making race strategy decisions in circuit motorsport. Applied Sciences, 10(21), 7805. https://doi.org/10.3390/app10217805.

Heilmeier, A., Wischnewski, A., Hermansdorfer, L., Betz, J., Lienkamp, M., & Lohmann, B. (2020c). Minimum curvature trajectory planning and control for an autonomous race car. Vehicle System Dynamics, 58(10), 1497-1527.

Herman, J., Francis, J., Ganju, S., Chen, B., Koul, A., Gupta, A., Skabelkin, A., Zhukov, I., Kumskoy, M., & Nyberg, E. (2021). Learn-to-race: A multimodal control environment for autonomous racing. 2021 IEEE/CVF International Conference on Computer Vision (ICCV), 9773-9782.

Herrmann, T., Christ, F., Betz, J., & Lienkamp, M. (2019). Energy management strategy for an autonomous electric racecar using optimal control. 2019 IEEE Intelligent Transportation Systems Conference (ITSC), 720-725.

Herrmann, T., Passigato, F., Betz, J., & Lienkamp, M. (2020). Minimum race-time planning-strategy for an autonomous electric racecar. 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC), 1-6.

Herrmann, T., Sauerbeck, F., Bayerlein, M., Betz, J., & Lienkamp, M. (2022). Optimization-based real-time-capable energy strategy for autonomous electric race cars. SAE International Journal of Connected and Automated Vehicles, 5(1), 45-59.

Hermansdorfer, L., Betz, J., & Lienkamp, M. (2019). A concept for estimation and prediction of the tire-road friction potential for an autonomous racecar. 2019 IEEE Intelligent Transportation Systems Conference (ITSC), 1490-1495. https://doi.org/10.1109/itsc.2019.8917024.

Hermansdorfer, L., Betz, J., & Lienkamp, M. (2020). Benchmarking of a software stack for autonomous racing against a professional human race driver. 2020 Fifteenth International Conference on Ecological Vehicles and Renewable Energies (EVER). https://doi.org/10.1109/ever48776.2020.9242926.

Hernandez, B., & Giraldo, E. (2018). A review of path planning and control for autonomous robots. 2018 IEEE 2nd Colombian Conference on Robotics and Automation (CCRA), 1-6.

Hewing, L., Liniger, A., & Zeilinger, M. N. (2018). Cautious NMPC with gaussian process dynamics for autonomous miniature race cars. 2018 European Control Conference (ECC), 1341-1348.

Hilary, C. B. (2021). Cloud testing perspectives with continuous integration. International Journal of Advanced Research in Science, Communication and Technology, 458-463.

Hindiyeh, R. Y., & Christian Gerdes, J. (2014). A controller framework for autonomous drifting: Design, stability, and experimental validation. Dynamic Systems, Measurement, and Control, 136(5).

Hobbs, C. (2025). Formal verification. Embedded Software Dev. for Safety-Critical Systems, 220-248.

Hoffmann, G. M., Tomlin, C. J., Montemerlo, M., & Thrun, S. (2007). Autonomous automobile trajectory tracking for off-road driving: Controller design, experimental validation and racing. American Control Conference, 2296–2301. https://doi.org/10.1109/ACC.2007.4282788

Hossain, M. D., Inoue, H., Ochiai, H., Fall, D., and Kadobayashi, Y. (2020). LSTM-based intrusion detection system for in-vehicle CAN bus communications, IEEE Access, 8, 185489-185502.

Hou, X., Zhang, J., He, C., Ji, Y., Zhang, J., & Han, J. (2022). Autonomous driving at the handling limit using residual reinforcement learning. Advanced Engineering Informatics, 54.0, 101754.

Huang, J., Labi, S., Kulcsar, B.A., Monreal, C.O., Wu, F., He, Z., Qu, X. (2024). Editorial: Intelligent vehicles and smart transportation, Fundamental Research 2024 Oct 1;4(5), 979–980.

Huang, Z., Mo, X., & Lv, C. (2022). Multi-modal motion prediction with transformer-based neural network for autonomous driving. 2022 International Conference on Robotics and Automation (ICRA), 2605-2611.

Huilong Yu, Castelli-Dezza, F., & Cheli, F. (2017). Optimal powertrain design and control of a 2-IWD electric race car. 2017 Intern. Conf. of Electrical & Electronic Tech. for Automotive, 1-7.

IAC (2022). Indy Autonomous Challenge showcases world's fastest autonomous racecar powered by Luminar; announces long-term partnership, www.indyautonomouschallenge.com/luminar-la

IAC (2025). Indy Autonomous Challenge - Official Website. Indy Autonomous Challenge - Official Website. 2023-02-13. Retrieved 2025-12-06.

Iclodean, C., Cordos, N., & Varga, B. O. (2020). Autonomous shuttle bus for public transportation: A review. Energies, 13(11), 2917.

Ivanov, R., Carpenter, T. J., Weimer, J., Alur, R., Pappas, G. J., & Lee, I. (2020). Verifying the safety of autonomous systems with neural network controllers. ACM Transactions on Embedded Computing Systems, 20(1), 1-26.

Jahncke, F., & Betz, J. (2025). MIND-Stack: Modular, interpretable, end-to-end differentiability for autonomous navigation. 2025 IEEE Intelligent Vehicles Symposium (IV), 1179-1186.

Jaritz, M., de Charette, R., Toromanoff, M., Perot, E., & Nashashibi, F. (2018). End-to-end race driving with deep reinforcement learning. 2018 IEEE International Conference on Robotics and Automation (ICRA), 2070-2075.

Javed, S., Hassan, A., Ahmad, R., Ahmed, W., Alam, M. M., & Rodrigues, J. J. P. C. (2022). UAV trajectory planning for disaster scenarios. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.4232781.

Jeong hwan Jeon, Cowlagi, R. V., Peters, S. C., Karaman, S., Frazzoli, E., Tsiotras, P., & Iagnemma, K. (2013). Optimal motion planning with the half-car dynamical model for autonomous high-speed driving. 2013 American Control Conference, 188-193.

Jianwang, H., Ramirez-Mendoza, R. A., & Xiaojun, T. (2021). Robust analysis for data-driven model predictive control. Systems Science & Control Engineering, 9(1), 393-404.

Ji, X., He, X., Lv, C., Liu, Y., & Wu, J. (2018). Adaptive-neural-network-based robust lateral motion control for autonomous vehicle at driving limits. Control Engineering Practice, 76.0, 41-53.

Jiang, A. (2021). Research on the development of autonomous race cars and impact on self-driving cars. Journal of Physics: Conference Series, 1824(1), 012009.

Joglekar, A., Deshpande, B., Basuthakur, M., & Krovi, V. N. (2021). Fusing offline and online trajectory optimization techniques for goal-to-goal navigation of a scaled autonomous vehicle. SAE Technical Paper Series, 1.

Joa, E., Cha, H., Hyun, Y., Koh, Y., Yi, K., & Park, J. (2020). A new control approach for automated drifting in consideration of the driving characteristics of an expert human driver. Control Engineering Practice, 96.0, 104293

Kalaria, D., Lin, Q., & Dolan, J. M. (2022). Delay-aware robust control for safe autonomous driving. 2022 IEEE Intelligent Vehicles Symposium (IV), 1565-1571.

Kalaria, D., Maheshwari, C., Sastry, S. (2024). Real-time algorithms for game-theoretic motion planning and control in autonomous racing using near-potential function, https://arxiv.org/html/2412.08855v1

Kang, Y., Yin, H., & Berger, C. (2019). Test your self-driving algorithm: An overview of publicly available driving datasets and virtual testing environments. IEEE Transactions on Intelligent Vehicles, 4(2), 171–185.

Kapania, N. R., & Gerdes, J. C. (2015a). Design of a feedback-feedforward steering controller for accurate path tracking and stability at the limits of handling. Vehicle Sys. Dyn. 53(12), 1687-1704.

Kapania, N. R., & Gerdes, J. C. (2015b). Path tracking of highly dynamic autonomous vehicle trajectories via iterative learning control. 2015 American Control Conference (ACC), 2753-2758.

Kapocsi, M., Orova, L. I., & Pusztai, Z. (2025). Optimal sensor placement for autonomous formula student vehicles: A field-of-view analysis of dual LIDAR and stereo camera configurations. Engineering Proceedings, 113(1), 27.

Karaman, S., Anders, A., Boulet, M., Connor, J., Gregson, K., Guerra, W., Guldner, O., Mohamoud, M., Plancher, B., Shin, R., & Vivilecchia, J. (2017). Project-based, collaborative, algorithmic robotics for high school students: Programming self-driving race cars at MIT. 2017 IEEE Integrated STEM Education Conference (ISEC), 195-203. https://doi.org/10.1109/isecon.2017.7910242.

Karimi Pour, F., Theilliol, D., Puig, V., & Cembrano, G. (2021). Health-aware control design based on remaining useful life estimation for autonomous racing vehicle. ISA Transactions, 113, 196-209.

Karle, P., Fent, F., Huch, S., Sauerbeck, F., & Lienkamp, M. (2023). Multi-modal sensor fusion and object tracking for autonomous racing. IEEE Trans. on Intelligent Vehicles, 8(7), 3871–3883.

Kaufeld, M., Trauth, R., & Betz, J. (2024). Investigating driving interactions: A robust multi-agent simulation framework for autonomous vehicles. 2024 IEEE Intelligent Veh. Symp., 803-810.

Kaufmann, E., Bauersfeld, L., Loquercio, A., Müller, M., Koltun, V., & Scaramuzza, D. (2023). Champion-level drone racing using deep reinforcement learning. Nature, 620(7976), 982-987.

Kabzan, J., Hewing, L., Liniger, A., & Zeilinger, M. N. (2019a). Learning-based model predictive control for autonomous racing. IEEE Robotics and Automation Letters, 4(4), 3363-3370.

Kabzan, J., de la Iglesia Valls, M., Reijgwart, V., Hendrikx, H. F. C., Ehmke, C., Prajapat, M., Bühler, A., Gosala, N. B., 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., & Siegwart, R. (2019b). AMZ Driverless: The full autonomous racing system. Journal of Field Robotics. https://doi.org/10.1002/rob.21977

Kabzan, J., Valls, M. I., Reijgwart, V. J. F., Hendrikx, H. F. C., Ehmke, C., Prajapat, M., Bühler, A., Kapania, N. R., Subosits, J., & Gerdes, J. C. (2015). A sequential two-step algorithm for fast generation of vehicle racing trajectories. Volume 3: Multiagent Network Systems; Natural Gas and Heat Exchangers; Path Planning and Motion Control; Powertrain Systems; Rehab Robotics; Robot Manipulators; Rollover Prevention (AVS); Sensors and Actuators; Time Delay Systems; Tracking Control Systems; Uncertain Systems and Robustness; Unmanned, Ground and Surface Robotics; Vehicle Dynamics Control; Vibration and Control of Smart Structures/Mech Systems; Vibration Issues in Mechanical Systems. https://doi.org/10.1115/dscc2015-9757.

Ke, F., Zhao, D., Sun, G., & Feng, W. (2019). Precise evaluation for continuous action control in reinforcement learning. Proceedings of the 2019 3rd High Performance Computing and Cluster Technologies Conference, 67-70.

Kegelman, J. C., Harbott, L. K., & Gerdes, J. C. (2016). Insights into vehicle trajectories at the handling limits: analyzing open data from race car drivers. Vehicle System Dynamics, 55(2), 191-207.

Keanly, O., & Engelbrecht, J. A. A. (2023). Optimization-based path planning and collision avoidance for autonomous racing. MATEC Web of Conferences, 388, 04018.

Khan, L. U., Ndikumana, A., Tran, N. H., & Hong, C. S. (2022). Edge intelligence for autonomous driving cars. towards a wireless connected world: Achievements and New Tech., 223-259.

Kessler, T., Esterle, K., & Knoll, A. (2020). Linear differential games for cooperative behavior planning of autonomous vehicles using mixed-integer programming. 2020 59th IEEE Conference on Decision and Control (CDC), 4060-4066.

Kloeser, D., Schoels, T., Sartor, T., Zanelli, A., Prison, G., & Diehl, M. (2020). NMPC for racing using a singularity-free path-parametric model with obstacle avoidance. IFAC-PapersOnLine, 53(2), 14324-14329.

Koirala, P., & Fleming, C. (2024). F1tenth autonomous racing with offline reinforcement learning methods. 2024 IEEE 27th International Conference on Intelligent Transp. Sys., 3922-3929.

Kong, Q., Zhang, L., & Xu, X. (2021). Lane keeping algorithm for autonomous driving via safe reinforcement learning. Lecture Notes in Computer Science, 439-450.

Korkmaz, B., Etlik, U. B., Beke, A., & Kumbasar, T. (2018). Fuzzy logic based self-driving racing car control system. 2018 6th International Conference on Control Engineering & Information Technology (CEIT), 1-6. https://doi.org/10.1109/ceit.2018.8751900.

Kritayakirana, K., & Gerdes, J. C. (2010). Autonomous cornering at the limits: Maximizing a “g-g” diagram by using feedforward trail-braking and throttle-on-exit. IFAC Procs. 43(7), 548-553.

Kritayakirana, K., & Gerdes, J. C. (2012). Autonomous vehicle control at the limits of handling. International Journal of Vehicle Autonomous Systems, 10(4), 271.

Kritayakirana, K., & Gerdes, J. C. (2012). Using the center of percussion to design a steering controller for an autonomous race car. Vehicle System Dynamics, 50(sup1), 33-51.

Kumar, D., and Muhammad, N. (2023). A survey on localization for autonomous vehicles, IEEE Access, 10.1109/ACCESS.2023.3326069

Lam, J. (2018). Roborace hopes to normalize autonomous vehicles through racing sports, Autonomous Racing Leagues, Trendhunter, https://www.trendhunter.com/trends/roborace

Langrill, C. (2023). What you're getting all wrong about the Indy Autonomous Challenge, Autoweek, https://www.autoweek.com/racing/more-racing/a43863454/what-youre-getting-all-wrong-about-the-indy-autonomous-challenge/

Labi, S., Saeed, T.U., Pourgholamali, M., Saka, Z.A., Sinha, K.C. (2023). Design and management of highway infrastructure to accommodate CAVs, CCAT Report #29, The Center for Connected and Automated Transportation, Purdue University, West Lafayette, IN.

Labi, S. (2019). Preparing our infrastructure for connected and autonomous vehicle operations” In Proc., Int. Conf. on Smart Cities. Iowa City, IA: International Society for Maintenance and Rehabilitation of Transport Infrastructures (iSmarti), Univ. of Iowa.

Langmann, A., Ögretmen, L., Werner, F., & Betz, J. (2025). Online velocity profile generation and tracking for sampling-based local planning algorithms in autonomous racing environments. 2025 IEEE Intelligent Vehicles Symposium (IV), 632-639.
Laurense, V. A., Goh, J. Y., & Gerdes, J. C. (2017). Path-tracking for autonomous vehicles at the limit of friction. 2017 American Control Conference (ACC), 5586-5591.
Le Floch, B., How, J., Breger, L., & Stoeckle, M. (2017). Trajectory planning for autonomous parafoils in complex terrain. 24th AIAA Aerodynamic Decelerator Systems Technology Conference. https://doi.org/10.2514/6.2017-3220.
Le Large, N., Bieder, F., & Lauer, M. (2021). Comparison of different SLAM approaches for a driverless race car. tm - Technisches Messen, 88(4), 227-236.
Lee, D., Nam, H., Ryu, C., Nah, S., Moon, S., & Shim, D. H. (2024a). Enhancing state estimator for autonomous racing: Leveraging multi-modal system and managing computing resources. IEEE Transactions on Intelligent Vehicles.
Lee, G. S., & Lee, J. C. (2023). Indoor autonomous driving system using reinforcement learning. Journal of the Korea Academia-Industrial cooperation Society, 24(10), 863-868.
Lee, J., Eom, C., Lee, D., Kwon, M. (2024b). Jerk-minimized autonomous driving strategy with deep reinforcement learning, 2nd Workshop on Multi-Agent Dynamic (MAD) Games, 2024 IEEE International Conference on Robotics and Automation (ICRA), May 13, 2024. Yokohama, Japan, https://icra2024-madgames.f1tenth.org/papers/lee.pdf
Li, J.-T., Chen, C.-K., & Ren, H. (2024). Time-optimal trajectory planning and tracking for autonomous vehicles. Sensors, 24(11), 3281.
Li, S., van der Horst, E., Duernay, P., de Wagter, C., & de Croon, G. C. H. E. (2020). Visual model-predictive localization for computationally efficient autonomous racing of a 72-g drone. Journal of Field Robotics, 37(4), 667–692. https://doi.org/10.1002/rob.21956
Li, X., Zhang, Y., & Liu, Z. (2025a). YOLOv13-Cone-Lite: An enhanced algorithm for traffic cone detection in autonomous formula racing cars. Applied Sciences, 15(17), 9501.
Li, N., Goubault, E., Pautet, L., & Putot, S. (2022). A real-time NMPC controller for autonomous vehicle racing. 2022 6th International Conf. on Automation, Control & Robots (ICACR), 148-155.
Li, Y. (2022). A review of how game theory is applied in transportation analysis, 2022 5th International Conference on Financial Management, Education and Social Science (FMESS 2022), Hohhot, China, Clausius Scientific Press.
Li, Z., Jin, G., Yu, R., Chen, Z., Li, N., Han, W., Xiong, L., Leng, B., Hu, J., Kolmanovsky, I., & Filev, D. (2025b). A survey of reinforcement learning-based motion planning for autonomous driving: Lessons learned from a driving task perspective. arXiv. https://arxiv.org/abs/2503.23650
Lim, N. Y., & Park, T. H. (2025). Long-range LiDAR vehicle detection through clustering and classification for autonomous racing. IEEE Access, 13.0, 30406-30419.
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.
Liniger, A., & Lygeros, J. (2019). Real-time control for autonomous racing based on viability theory. IEEE Transactions on Control Systems Technology, 27(2), 464-478.
Liniger, A., & Lygeros, J. (2020). A noncooperative game approach to autonomous racing. IEEE Transactions on Control Systems Technology, 28(3), 884-897.
Liniger, A., & Lygeros, J. (2015). A viability approach for fast recursive feasible finite horizon path planning of autonomous RC cars. Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control, 1-10.
Liniger, A., Zhang, X., Aeschbach, P., Georghiou, A., & Lygeros, J. (2017). Racing miniature cars: Enhancing performance using Stochastic MPC and disturbance feedback. 2017 American Control Conference (ACC), 5642-5647.

Litman, T. (2014). Autonomous Vehicle Implementation Predictions. Victoria Transport Policy Institute. <http://www.vtpi.org/avip.pdf> (March 10, 2015).

Liu, H. X. & Feng, Y. (2020). Accelerated training for connected and automated vehicles based on adaptive evaluation method. Technical Report. UMTRI-2023-5. DOI: 10.7302/7017

Liu, H. X., & Feng, Y. (2022). DeepScenario: City scale scenario generation for automated driving system testing & evaluation. Technical Report No. UMTRI-2023-3. DOI: 10.7302/7018

Liu, J. (2024). The development and application of autonomous driving technology. Applied and Computational Engineering, 43(1), 160-166.

Liu, Y., Ma, F., Mei, X., Xue, B., Wu, J., & Zhang, C. (2025). Autonomous drifting like professional racing drivers: A survey. Technologies, 5(2), 33. https://doi.org/10.3390/technologies5020033

Liu, Y., Shi, L., Xu, W., Xiong, X., Sun, W., & Qu, L. (2020). Design of driverless racing chassis based on MPC. 2020 Chinese Automation Congress (CAC), 6061-6066.

Liu, W., Liu, C., Chen, G., & Knoll, A. (2021). Gaussian process based model predictive control for overtaking in autonomous driving. Frontiers in Neurorobotics, 15.0. https://doi.org/10.3389/fnbot.2021.723049.

Loiacono, D., Prete, A., Lanzi, P. L., & Cardamone, L. (2010). Learning to overtake in TORCS using simple reinforcement learning. IEEE Congress on Evolutionary Computation, 1-8.

Lovato, S., & Massaro, M. (2021). Three-dimensional fixed-trajectory approaches to the minimum-lap time of road vehicles. Vehicle System Dynamics, 60(11), 3650-3667.

Lowe, E., & Guvenc, L. (2023). Autonomous vehicle emergency obstacle avoidance maneuver framework at highway speeds. Electronics, 12(23), 4765. https://doi.org/10.3390/electronics12234765

Lu, Y., Yang, B., Li, J., Zhou, Y., Chen, H., & Mo, Y. (2023). Consecutive inertia drift of autonomous rc car via primitive-based planning and data-driven control. 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 4835-4840.

Lu, Z., Zhang, C., Zhang, H., Wang, Z., Huang, C., & Ji, Y. (2021). Deep reinforcement learning based autonomous racing car control with priori knowledge. 2021 China Automation Congress (CAC), 2241–2246. https://doi.org/10.1109/CAC53003.2021.9728289

Ma, H., Meng, F., Ye, C., Wang, J., & Meng, M. Q. H. (2022). Bi-risk-RRT Based efficient motion planning for autonomous ground vehicles. IEEE Trans. on Intelligent Vehicles, 7(3), 722-733.

Macadam, C. C. (2003). Understanding and modeling the human driver. Vehicle System Dynamics, 40(1–3), 101–134.

Mahmoud, Y., Okuyama, Y., Fukuchi, T., Kosuke, T., & Ando, I. (2020). Optimizing deep-neural-network-driven autonomous racecar using image scaling. SHS Web of Conferences, 77.0, 04002. https://doi.org/10.1051/shsconf/20207704002.

Manion, P., & Durango-Cohen, P.L. (2025). Optimizing investments in design, maintenance, and operations of transportation infrastructure to capitalize on the promise of connected autonomous vehicles, Technical Report CCAT-NU-2025-2, Northwestern University, Evanston, IL.

Manjunathan, A., Suresh Kumar, A., Udhayanan, S., Thirumarai Selvi, C., & Stonier, A. A. (2021). Design of autonomous vehicle control using IoT. IOP Conference Series: Materials Science and Engineering, 1055(1), 012008.

Manjanoor, M. A. (2021). Towards autonomous drone racing. IEEE Int. Conf. on Robotics and Automation (ICRA). https://doi.org/10.31274/cc-20240624-1511.

Manna, R. K. (2022). Control challenges for high-speed autonomous racing: Analysis and simulated experiments, SAE International Journal of Connected & Automated Vehicles 5(1),101-114.

Manna, R. K., Gonzalez, D. J., Chellapandi, V., Mar, M., Kannan, S. S., Wadekar, S., Dietz, E. J., Korpela, C. M., & El Gamal, A. (2022). Control challenges for high-speed autonomous racing: analysis and simulated experiments. SAE International Journal of Connected and Automated Vehicles, 5(1), 101-114.

Mar, M., & Dietz, E. (2024). A qualitative review of full-sized autonomous racing vehicle sensors: A case study. Proceedings of the 10th International Conference on Vehicle Technology and Intelligent Transport Systems, 311-318.

Mar, M. Chellapandi, V. P. Yuan, L. Wang Z. and Dietz, E. (2025). Advanced sensor configurations for high-speed autonomous racing vehicles, in IEEE Journal of Selected Areas in Sensors 2, 136-149.

Marley, M., Skjetne, R., Basso, E., & Teel, A. R. (2021). Maneuvering with safety guarantees using control barrier functions. IFAC-PapersOnLine, 54(16), 370-377.

Massa, F., Bonamini, L., Settimi, A., Pallottino, L., & Caporale, D. (2020). LiDAR-based GNSS denied localization for autonomous racing cars. Sensors, 20(14), 3992.

McDonald, D. (2017). How might connected vehicles and autonomous vehicles influence geometric design? Hanson Engineering, Transportation Short Course, Texas A&M University, College Station, TX. static.tti.tamu.edu/conferences/tsc17/presentations/design/mcdonald.pdf, accessed December 17, 2025.

Mehta, H. P., & Mistry, V. K. (2017). Driverless car. IJARCCE, 6(5), 39-42.

Meng, D., Jiang, Y., Chu, H., Tian, M., & Gao, B. (2024). Data-driven tire forces estimation for autonomous vehicle applications. 2024 43rd Chinese Control Conference (CCC), 6421-6426.

Meriño Córdoba, V.H., García Lirios, C., Vallejo Montoya, E.A., Urzola Berrio, H.E. (2025) Artificial intelligence, sports, and dentistry: Synergistic advances and future applications. Biomedical Research and Clinical Reviews, 10(1).

Mesterton-Gibbons, M. (2009). Time-optimal control. The Student Mathematical Library, 149-158. https://doi.org/10.1090/stml/050/18.

Metz, D., & Williams, D. (1989). Near time-optimal control of racing vehicles. Automatica, 25(6), 841-857.

Micheli, F., Bersani, M., Arrigoni, S., Braghin, F., & Cheli, F. (2022). NMPC trajectory planner for urban autonomous driving. Vehicle System Dynamics, 61(5), 1387-1409.

Mirage News (2025). TUM triumphs again at Abu Dhabi autonomous race, Mirage Technology, www.miragenews.com/tum-triumphs-again-at-abu-dhabi-autonomous-race-1571240/

Mirebeau, J., Gayraud, L., Barrère, R., Chen, D., & Desquilbet, F. (2022). Massively parallel computation of globally optimal shortest paths with curvature penalization. Concurrency and Computation: Practice and Experience, 35(2).

Muhammad, N. (2012). Contributions to the use of 3D LiDARs for autonomous navigation: Calibration and qualitative localization,'' M.S. Thesis, INSA de Toulouse, France.

Mohanasundaram, N., Sai Eshwar Supreet, P. A. (2023). Semi autonomous miniature car. 2023 2nd International Conference on Advancements in Electrical, Electronics, Communication, Computing and Automation (ICAECA), 1-6. https://doi.org/10.1109/icaeca56562.2023.10199902.

Montani, M., Ronchi, L., Capitani, R., & Annicchiarico, C. (2021). A hierarchical autonomous driver for a racing car: Real-Time Planning and Tracking of the Trajectory. Energies, 14(19), 6008.

Mostaghim, S. (2022). Computational intelligence methodologies for multi-objective optimization and decision-making in autonomous systems. Women in Engineering and Science, 377-392.

Mumuni, A., Mumuni, F. (2022). Robust appearance modeling for object detection and tracking: A survey of deep learning approaches, Progress in Artificial Intelligence 11, 279–313.

Muradi, M., & Wanka, R. (2020). Sample-based motion planning for multi-robot systems. 2020 6th International Conference on Control, Automation and Robotics (ICCAR), 130-138.

Muraleedharan, A., Okuda, H., Suzuk, T. (2022). Real-time implementation of randomized model predictive control for autonomous driving IEEE Transactions on Intelligent Vehicles, 7(1), 11-20.

Myerson, R. B. (1991). Game theory: Analysis of conflict. Harvard University Press, Cambridge, MA. Fudenberg, D., Tirole, J. (1991). Game theory, 1st Edition, MIT Press, Cambridge, MA.

Naik, A. (2019). Autonomous car. International Journal for Research in Applied Science and Engineering Technology, 7(4), 2167-2172.

Nakanishi YJ, Auza PM. Connected vehicles and driving automation systems. In Springer Handbook of Automation 2023 Jun 17 (1079-1113). Cham: Springer International Publishing.

Nakamoto, N., & Kobayashi, H. (2019). Development of an open-source educational and research platform for autonomous cars. IECON 2019 - 45th Annual Conference of the IEEE Industrial Electronics Society, 6871-6876.

Nekkah, S., Janus, J., Boxheimer, M., Ohnemus, L., Hirsch, S., Schmidt, B., Liu, Y., Borbély, D., Keck, F., Bachmann, K., & Bleszynski, L. (2022). The autonomous racing software stack of the KIT19d. SAE International Journal of Connected and Automated Vehicles, 5(1), 73-86.

New Eagle (2022). 4 Lessons Learned from the Indy Autonomous Challenge, neweagle.net/blog/4-lessons-learned-from-the-Indy-autonomous-challenge/

NHTSA (2022). Levels of Automation. https://www.nhtsa.gov/sites/nhtsa.gov/files/2022-05/Level-of-Automation-052522-tag.pdf

Ni, J., Hu, J., & Xiang, C. (2019). Robust path following control at driving/handling limits of an autonomous electric racecar. IEEE Transactions on Vehicular Technology, 68(6), 5518-5526.

Niu, J., Hu, Y., Li, W., Huang, G., Han, Y., & Li, X. (2022). Closing the dynamics gap via adversarial and reinforcement learning for high-speed racing. 2022 International Joint Conference on Neural Networks (IJCNN), 01-10. https://doi.org/10.1109/ijcnn55064.2022.9892424.

Niu, J., Hu, Y., Jin, B., Han, Y., & Li, X. (2020). Two-stage safe reinforcement learning for high-speed autonomous racing. 2020 IEEE International Conf. on Sys. Man, & Cybernetics, 3934-3941.

Nobis, F., Betz, J., Hermansdorfer, L., & Lienkamp, M. (2019). Autonomous racing: A comparison of SLAM algorithms for large scale outdoor environments. Proceedings of the 2019 3rd International Conference on Virtual and Augmented Reality Simulations, 82-89.

Notomista, G., Wang, M., Schwager, M., & Egerstedt, M. (2020). Enhancing game-theoretic autonomous car racing using control barrier functions. 2020 IEEE International Conference on Robotics and Automation (ICRA), 5393-5399.

Novi, T., Liniger, A., Capitani, R., & Annicchiarico, C. (2019). Real-time control for at-limit handling driving on a predefined path. Vehicle System Dynamics, 58(7), 1007-1036.

Noorvand, H., Karnati, G., & Underwood, B. S. (2017). Autonomous vehicles. Transportation Research Record 2640(1), 21-28.

Ni, J., & Hu, J. (2017). Path following control for autonomous formula racecar: Autonomous formula student competition. 2017 IEEE Intelligent Vehicles Symposium (IV), 1835-1840.

O'Kelly, M., Zheng, H., Jain, A., Auckley, J., Luong, K., & Mangharam, R. (2020a). TUNERCAR: A Super-optimization toolchain for autonomous racing. 2020 IEEE International Conference on Robotics and Automation (ICRA), 5356-5362.

O'Kelly, M., Zheng, H., Karthik, D., & Mangharam, R. (2020b). F1TENTH: An open-source evaluation environment for continuous control and reinforcement learning. Proceedings of Machine Learning Research, 123, 77–89. https://proceedings.mlr.press/v123/o-kelly20a.html

Okte, E., & Al-Qadi, I. L. (2022). Impact of autonomous and human-driven trucks on flexible pavement design. Transportation Research Record 2676(7), 144-160.

Oliveira, R., M. Rocha, F. H., Ott, L., Guizilini, V., Ramos, F., & Grassi, V. (2018). Learning to race through coordinate descent bayesian optimization. 2018 IEEE International Conference on Robotics and Automation (ICRA), 6431-6438.

Othman, K. (2021). Impact of autonomous vehicles on the physical infrastructure: Changes and challenges. Designs, 5(3), 40. https://doi.org/10.3390/designs5030040

Pagot, E., Piccinini, M., & Biral, F. (2020). Real-time optimal control of an autonomous RC car with minimum-time maneuvers and a novel kineto-dynamical model. 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2390-2396. https://doi.org/10.1109/iros45743.2020.9340640.

Palafox, P. R., Betz, J., Nobis, F., Riedl, K., & Lienkamp, M. (2019). SemanticDepth: Fusing semantic segmentation and monocular depth estimation for enabling autonomous driving in roads without lane lines. Sensors, 19(14), 3224.

Palatti, J., Aksjonov, A., Alcan, G., & Kyrki, V. (2021). Planning for safe abortable overtaking maneuvers in autonomous driving. 2021 IEEE International Intelligent Transportation Systems Conference (ITSC), 508-514.

Pan, Y., Cheng, C. A., Saigol, K., Lee, K., Yan, X., Theodorou, E., & Boots, B. (2018). Agile autonomous driving using end-to-end deep imitation learning. Robotics: Science and Systems XIV. https://doi.org/10.15607/rss.2018.xiv.056.

Park, H., & Gerdes, J. C. (2015). Optimal tire force allocation for trajectory tracking with an over-actuated vehicle. 2015 IEEE Intelligent Vehicles Symposium (IV). https://doi.org/10.1109/ivs.2015.7225820.

Park, H., & Gerdes, J. C. (2017). Analysis of feasible tire force regions for optimal tire force allocation with limited actuation. IEEE Intelligent Transportation Systems Magazine, 9(3), 75-87.

Park, C., Kee, S. C. (2021). Online local path planning on the campus environment for autonomous driving considering road constraints and multiple obstacles, Appl. Sci. 11(9), 3909.

Patton, R., Schuman, C., Kulkarni, S., Parsa, M., Mitchell, J. P., Haas, N. Q., Stahl, C., Paulissen, S., Date, P., Potok, T., & Snyder, S. (2021). Neuromorphic computing for autonomous racing. International Conference on Neuromorphic Systems 2021, 1-5.

Patz, B. J., Papelis, Y., Pillat, R., Stein, G., Harper, D. (2008). A practical approach to robotic design for the DARPA Urban Challenge, Journal of Field Robotics 25(8), Special Issue on the 2007 DARPA Urban Challenge Part I, 528-566.

Pedone, S., & Fagiolini, A. (2020). Racecar longitudinal control in unknown and highly varying driving conditions. IEEE Transactions on Vehicular Technology, 69(11), 12521-12535.

Pendleton, S., Andersen, H., Du, X., Shen, X., Meghjani, M., Eng, Y., Rus, D., & Ang, M. (2017). Perception, planning, control, and coordination for autonomous vehicles. Machines, 5(1), 6.

Peng, W. Z., Ao, Y. H., He, J. H., & Wang, P. F. (2021). Vehicle odometry with Camera-Lidar-IMU information fusion and factor-graph optimization. J. of Intelligent & Robotic Systems, 101(4).

Perez, D., Saez, Y., Recio, G., & Isasi, P. (2008). Evolving a rule system controller for automatic driving in a car racing competition. 2008 IEEE Symposium On Computational Intelligence and Games, 336-342. https://doi.org/10.1109/cig.2008.5035659.

Perot, E., Jaritz, M., Toromanoff, M., & De Charette, R. (2017). End-to-end driving in realistic racing game with deep reinforcement learning. 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). https://doi.org/10.1109/cvprw.2017.64.

Perrone, A. (2025). Innovations from racing that found their way on our roads, Endurance, www.endurancewarranty.com/learning-center/motorsports/

Peyron, M., Fillion, A., Gürol, S., Marchais, V., Gratton, S., Boudier, P., & Goret, G. (2021). Latent space data assimilation by using deep learning. Quarterly Journal of the Royal Meteorological Society, 147(740), 3759-3777.

Pham, S., Truong, L., Nguye, N. (2021). Impacts of autonomous vehicles on road and pavement design, Australasian Transport Research Forum 2021 Procs., 8-10 Dec, Brisbane, Australia

Piccinini, M., Taddei, S., Betz, J., & Biral, F. (2025). Kineto-dynamical planning and accurate execution of minimum-time maneuvers on three-dimensional circuits. 2025 IEEE International Conference on Robotics and Automation (ICRA), 1-7. https://doi.org/10.1109/icra55743.2025.11127446.

Piccinini, M., Larcher, M., Pagot, E., Piscini, D., Pasquato, L., & Biral, F. (2022). A predictive neural hierarchical framework for on-line time-optimal motion planning and control of black-box vehicle models. Vehicle System Dynamics, 61(1), 83-110.

Pinkow, S., and Fillenburg, S. (2018). Continental develops aquaplaning warning system, https://www.continental.com/en/press/press-releases/aquaplaning-warning-system/

Pour, K. F., Theilliol, D., Puig, V., & Cembrano, G. (2021). Health-aware control design based on remaining useful life estimation for autonomous racing vehicle. ISA Transactions, 113, 196-209.

Pourgholamali, M., Miralinaghi, M., Seilabi, S.E., Labi, S. (2023). Sustainable deployment of autonomous vehicles dedicated lanes in urban traffic networks, Sustainable Cities and Society 99, 104969.

Puchtler, P., & Peinl, R. (2020). Evaluation of deep learning accelerators for object detection at the edge. Lecture Notes in Computer Sci., 320-326. https://doi.org/10.1007/978-3-030-58285-2_29.

Qiao, Q., Li,H., Cao, Z., Liu, H. (2025). End2Race: Efficient End-to-End Imitation Learning for Real-Time F1Tenth Racing, Preprints, https://arxiv.org/html/2509.16894v1

Quadflieg, J., Preuss, M., & Rudolph, G. (2011). Driving faster than a human player. Lecture Notes in Computer Science, 143-152. https://doi.org/10.1007/978-3-642-20525-5_15.

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. 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC). https://doi.org/10.1109/itsc55140.2022.9922239.

Raji, A., Liniger, A., Giove, A., Toschi, A., Musiu, N., Morra, D., Verucchi, M., Caporale, D., & Rong, G., Shin, B. H., Tabatabaee, H., Lu, Q., Lemke, S., Mozeiko, M., Boise, E., Uhm, G., Gerow, M., Mehta, S., Agafonov, E., Kim, T. H., Sterner, E., Ushiroda, K., Reyes, M., Zelenkovsky, D., & Kim, S. (2020). LGSVL simulator: A high fidelity simulator for autonomous driving. 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC), 1-6. https://doi.org/10.1109/itsc45102.2020.9294422.

Ramachandran, S., Horgan, J., Sistu, G., & McDonald, J. (2022). Fast and efficient scene categorization for autonomous driving using VAEs. 24th Irish Machine Vision and Image Processing Conference, 9-16. https://doi.org/10.56541/suhe3553.

Ramakrishnan, A., Alrajhi, A., Okte, E., Ozer, H., & Al-Qadi, I.L. (2021). Truck-platooning impacts on flexible pavements: Experimental and mechanistic approaches, Technical Report Nr. ICT-21-038). Illinois Center for Transportation, Champaign, IL.

Rathore, R. S., Hewage, C., Kaiwartya, O., & Lloret, J. (2022). In-vehicle communication cyber security: challenges and solutions. Sensors, 22(17), 6679.

Reiter, R., Kirchengast, M., Watzenig, D., & Diehl, M. (2021). Mixed-integer optimization-based planning for autonomous racing with obstacles and rewards. IFAC-PapersOnLine, 54(6), 99-106.

Remonda, A., Veas, E., & Luzhnica, G. (2021). Comparing driving behavior of humans and autonomous driving in a professional racing simulator. PLOS ONE, 16(2), e0245320. https://doi.org/10.1371/journal.pone.0245320.

Renzler, T., Stolz, M., Schratter, M., & Watzenig, D. (2020). Increased accuracy for fast moving LiDARS: Correction of distorted point clouds. 2020 IEEE International Instrumentation and Measurement Technology Conference (I2MTC), 1-6. https://doi.org/10.1109/i2mtc43012.2020.9128372.

Reveille Racing, & Dietz, E. (2023). A case study from the Indy Autonomous Challenge's simulation race. Journal of Technology, Management, and Applied Engineering, 39(1).

Reyes, K. G., & Alexander, F. J. (2020). Autonomous experimental design and execution. Handbook on Big Data and Machine Learning in the Physical Sciences, 241-270.

Rito Lima, I., Haar, S., Di Grassi, L., & Faisal, A. A. (2020). Neurobehavioural signatures in race car driving: a case study. Scientific Reports, 10(1). https://doi.org/10.1038/s41598-020-68423-2.

Roboracer (2025). The leading platform for autonomous racing education and research, https://roboracer.ai/
Roles, J. (2025). How self-driving technology is changing motorsport, Abu Dhabi Racing League Blog, Sept 22, 2025, https://a2rl.io/blog/18/How-Self-Driving-Technology-is-Changing-Motorsport
Rosen, W. (2022). Toyota Research Institute Pushes vehicle's capabilities to advance active safety, Toyota USA Newsroom, https://pressroom.toyota.com/toyota-research-institute-pushes-vehicles-capabilities-to-advance-active-safety/
Rosolia, U., & Borrelli, F. (2018). Learning model predictive control for iterative tasks: A data-driven control framework. IEEE Transactions on Automatic Control, 63(7), 1883–1896.
Rosolia, U., & Borrelli, F. (2020). Learning how to autonomously race a car: A predictive control approach. IEEE Transactions on Control Systems Technology, 28(6), 2713-2719.
Rosolia, U., Carvalho, A., & Borrelli, F. (2017a). Autonomous racing using learning Model Predictive Control. 2017 American Control Conference (ACC), 5115-5120. https://doi.org/10.23919/acc.2017.7963748.
Rosolia, U., Zhang, X., & Borrelli, F. (2017b). Robust learning model predictive control for iterative tasks: Learning from experience. 2017 IEEE 56th Annual Conference on Decision and Control (CDC). https://doi.org/10.1109/cdc.2017.8263812.
Rowold, M., Langmann, A., Lohmann, B., & Betz, J. (2024). Open-loop and feedback Nash trajectories for competitive racing with iLQGames. 2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC), 1827-1834. https://doi.org/10.1109/itsc58415.2024.10920189.
Rowold, M., Ögretmen, L., Kerbl, T., & Lohmann, B. (2022). Efficient spatiotemporal graph search for local trajectory planning on oval racetracks. Actuators, 11(11), 319.
Rumar, K. (1985). The role of perceptual and cognitive filters in observed behavior. In L. Evans & R. Schwing (Eds.), Human behavior and traffic safety (pp. 164–165). New York, NY: Plenum Press.
Rucco, A., Notarstefano, G., & Hauser, J. (2015). An efficient minimum-time trajectory generation strategy for two-track car vehicles. IEEE Transactions on Control Systems Technology, 23(4), 1505-1519. https://doi.org/10.1109/tcst.2014.2377777.
Sa, I., Dubé, R., Gawel, A., Pfeiffer, M., Liniger, A., Lygeros, J., & Siegwart, R. (2020). AMZ driverless: The full autonomous racing system. Journal of Field Robotics, 37(7), 1267-1294.
SAE (2016) Taxonomy and definitions for terms related to on-road motor vehicle automated driving systems. SAE International. Issued in 2014-01 and revised in 2016-09. On-Road Automated Vehicle Standards Committee. Retrieved from https://saemobilus.sae.org/content/j3016_201609
SAE International (2018). J3016_201806: Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles (Warrendale: SAE International, 15 June 2018), https://www.sae.org/standards/content/j3016_201806/.
SAE (2021). Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles, Standard J3016_202104, Quality, Testing and Safety, Automated and Autonomous Systems. https://www.sae.org/standards/j3016_202104-taxonomy-definitions-terms-related-driving-automation-systems-road-motor-vehicles
Saeed, T.U. (2019). Road infrastructure readiness for autonomous vehicles. PhD Dissertation, Purdue University, W. Lafayette, IN, USA.
Saeed, T.U., Alabi, B., Labi, S. (2020). Preparing road infrastructure to accommodate connected and automated vehicles: System-level perspective, Journal of Infrastructure Systems 27 (1), 06020003-1-6.
Saka, Z., Labi, S. (2023). Investigation of AV operational issues using simulation equipment, CCAT Report #34, The Center for Connected and Automated Transportation, Purdue University, West Lafayette, IN.

Salem, M., Mora, A. M., & Merelo, J. J. (2019). Beating uncertainty in racing bot evolution through enhanced exploration and pole position selection. 2019 IEEE Conference on Games (CoG), 1-8.

Salem, M., Mora, A. M., & Merelo, J. J. (2018). The evolutionary race: improving the process of evaluating car controllers in racing simulators. 2018 IEEE Conference on Computational Intelligence and Games (CIG), 1-8. https://doi.org/10.1109/cig.2018.8490364.

Salem, M., Mora, A. M., Merelo, J. J., & García-Sánchez, P. (2017). Driving in TORCS using modular fuzzy controllers. Lecture Notes in Computer Science, 361-376. https://doi.org/10.1007/978-3-319-55849-3_24.

Salmon, P. M., Regan, M. A., & Johnston, I. (2005). Human error and road transport: Phase one–Literature review (Report No. 256). Clayton, Australia: Monash University Accident Research Centre.

Samada, S. E., Puig, V., & Nejjari, F. (2023). Robust MPC-RG for an autonomous racing vehicle considering obstacles and the battery state of charge. Control Eng. Practice, 141, 105730.

Samper Mejia, J. P., Theodosis, P. A., & Gerdes, J. C. (2013). Using a path-fitting algorithm to analyze the racing techniques of a skilled driver. Volume 3: Nonlinear Estimation and Control; Optimization and Optimal Control; Piezoelectric Actuation and Nanoscale Control; Robotics and Manipulators; Sensing; System Identification (Estimation for Automotive Applications, Modeling, Therapeutic Control in Bio-Systems); Variable Structure/Sliding-Mode Control; Vehicles and Human Robotics; Vehicle Dynamics and Control; Vehicle Path Planning and Collision Avoidance; Vibrational and Mechanical Systems; Wind Energy Systems and Control. https://doi.org/10.1115/dscc2013-4106.

Sanket, S., Sinha, A., Varakantham, P., Andrew, P., & Tambe, M. (2020). Solving online threat screening games using constrained action space reinforcement learning. Proceedings of the AAAI Conference on Artificial Intelligence, 34(02), 2226-2235, doi.org/10.1609/aaai.v34i02.5599.

Sauerbeck, F., Baierlein, L., Betz, J., & Lienkamp, M. (2022). A combined LiDAR-camera localization for autonomous race cars. SAE International J. of Connected & Automated Vehicles, 5(1), 61–71.

Sauerbeck, F., Huch, S., Fent, F., Karle, P., Kulmer, D., & Betz, J. (2023). Learn to see fast: Lessons learned from autonomous racing on how to develop perception systems. IEEE Access, 11, 44034–44050.

Savoca, A., Barazzetta, J., Pesenti, G., & Manca, D. (2018). Model predictive control for automated anesthesia. Computer Aided Chemical Engineering, 1631-1636.

Scania Group (2025). The race is on with KTH Formula Student self-driving vehicle, accessed on December 17, 2025, https://www.scania.com/group/en/home/newsroom/news/2019/the-race-is-on-with-KTH-Formula-Student-self-driving-vehicle.html

Scarborough, C. (2021). Autonomous racing is here! How Roborace is helping develop technology for the future, https://motorsport.tech/roborace/how-roborace-is-helping-develop-technology

Scheffe, P., Henneken, T. M., Kloock, M., & Alrifaee, B. (2024). Sequential convex programming methods for real-time optimal trajectory planning in autonomous vehicle racing. 2024 IEEE Intelligent Vehicles Symposium (IV), 3144-3144. doi.org/10.1109/iv55156.2024.10588847.

Scholze, H., Anand, R., Kornitsky, Y., Loustau, L. (2024). Racing redefined: the autonomous revolution, www.kearney.com/industry/automotive/article/racing-redefined-the-autonomous-revolution

Schratter, M., et al. (2023). LiDAR-based mapping and localization for autonomous racing. https://linklab-uva.github.io/icra-autonomous-racing/contributed_papers/paper4.pdf.

Seward, D. (2014). Racing car basics. Race Car Design, 1-32. https://doi.org/10.1007/978-1-137-03015

Schwarting, W., Pierson, A., Karaman, S., & Rus, D. (2021). Stochastic dynamic games in belief space. IEEE Transactions on Robotics, 37(6), 2157-2172. https://doi.org/10.1109/tro.2021.3075376.

Seilabi, S.E., Pourgholamali, M., Miralinaghi, M., Correia, G., Labi, S. (2024). Optimizing dedicated lanes and tolling schemes for connected and autonomous vehicles to address bottleneck congestion considering morning commuter departure choices, Journal of Intelligent Transportation Systems, doi.org/10.1080/15472450.2024.2408024

Seilabi, S.E., Pourgholamali, M., Wang, J., Miralinaghi, M., Sundaram, S., Labi, S. (2022). Lane management in the era of CAV deployment, CCAT Report #72, The Center for Connected and Automated Transportation, Purdue University, West Lafayette, IN.

Shah, P., Maheshwari, M., Ramane, S., Chandra, P., Valvi, M. S., & Mehendale, N. (2023). Autonomous racing vehicle. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.4464632.

Shahbakhshi, A. (2022). The History of Car Racing, Ferro and Company, https://www.ferrowatches.com/blogs/news/the-history-of-car-racing

Shin, D., Park, K.-m., & Park, M. (2020). High definition map-based localization using ADAS environment sensors for automated driving vehicles. Applied Sciences, 10(14), 4924.

Shiotsuka, T., Bando, M., & Hokamoto, S. (2024). Data-driven modeling for multirotor autonomous control. AIAA SCITECH 2024 Forum. https://doi.org/10.2514/6.2024-0568.

Sivinski, R. (2011). Crash prevention effectiveness of light-vehicle electronic stability control: An update of the 2007 NHTSA evaluation (No. HS-811 486). Scotts Valley, CA: CreateSpace Independent Publishing Platform. https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/811486

Smirnov, N., Liu, Y., Validi, A., Morales-Alvarez, W., Olaverri-Monreal, C. (2021). A game theory-based approach for modeling autonomous vehicle behavior in congested, urban lane-changing scenarios. Sensors 21, 1523–32. doi:10.3390/s21041523

Smith, J., & Mistry, M. (2020). Online simultaneous semi-parametric dynamics model learning. IEEE Robotics and Automation Letters, 5(2), 2039-2046. https://doi.org/10.1109/lra.2020.2970987.

Siemens (2025). Software-defined vehicles, Automotive Digital Thread, www.sw.siemens.com/en-US/digital-thread/mbse/software-defined-vehicle/

Slama, D., Nonnenmacher, A., Irawan, T. (2023). The Software-Defined Vehicle, O'Reilly Media, Santa Rosa, CA.

Song, Y., Lin, H., Kaufmann, E., Durr, P., & Scaramuzza, D. (2021). Autonomous overtaking in Gran Turismo Sport using curriculum reinforcement learning. 2021 IEEE International Conference on Robotics and Automation (ICRA). https://doi.org/10.1109/icra48506.2021.9561049.

Song, Y., Steinweg, M., Kaufmann, E., & Scaramuzza, D. (2021). Autonomous drone racing with deep reinforcement learning. 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 1205-1212. https://doi.org/10.1109/iros51168.2021.9636053.

Spielberg, N. A., Brown, M., Kapania, N. R., Kegelman, J. C., & Gerdes, J. C. (2019). Neural network vehicle models for high-performance automated driving. Science Robotics, 4(28). https://doi.org/10.1126/scirobotics.aaw1975.

Spisak J. et al. (2022). Racing. Encyclopedia of animal cognition and behavior, 5843-5843.

Sridevi, M., Sugirtha, T., Kiran, B. R., & Yogamani, S. (2021). Localization and mapping for autonomous driving. Autonomous Driving and Advanced Driver-Assistance Systems (ADAS), 411-436. https://doi.org/10.1201/9781003048381-23.

Srinivasan, S., Nicolas Giles, S. N. G., & Liniger, A. (2021). A holistic motion planning and control solution to challenge a professional racecar driver. IEEE Rob. & Autom. Lett, 6(4), 7854-7860.

Srinivasan, S., Sa, I., Zyner, A., Reijgwart, V., Valls, M. I., & Siegwart, R. (2020). End-to-end velocity estimation for autonomous racing. IEEE Robotics and Automation Letters, 5(4), 6869-6875.

Stahl, T., Wischnewski, A., Betz, J., & Lienkamp, M. (2019a). ROS-based localization of a race vehicle at high-speed using LIDAR. E3S Web of Conferences, 95, 04002. https://doi.org/10.1051/e3sconf/20199504002.

Stahl, T., & Diermeyer, F. (2021). Online verification enabling approval of driving functions—Implementation for a Planner of an Autonomous Race Vehicle. IEEE Open Journal of Intelligent Transportation Systems, 2, 97-110.

Stahl, T., Wischnewski, A., Betz, J., & Lienkamp, M. (2019b). Multilayer graph-based trajectory planning for race vehicles in dynamic scenarios. 2019 IEEE Intelligent Transportation Systems Conference (ITSC). https://doi.org/10.1109/itsc.2019.8917032.

Stahl, T., & Betz, J. (2020). An open-source scenario architect for autonomous vehicles. 2020 Fifteenth International Conference on Ecological Vehicles and Renewable Energies (EVER), 1-8. https://doi.org/10.1109/ever48776.2020.9243029.

Stahl, T., Eicher, M., Betz, J., & Diermeyer, F. (2020). Online verification concept for autonomous vehicles – Illustrative study for a trajectory planning module. 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC), 1-7. https://doi.org/10.1109/itsc45102.2020.9294703.

Stewart, J. (2017). Meet the self-driving car built for human-free racing, Transportation, Wired, wired.com/2017/02/meet-self-driving-car-built-human-free-racing/

Strobel, K., Zhu, S., Chang, R., & Koppula, S. (2020). Accurate, low-latency visual perception for autonomous racing: Challenges, mechanisms, and practical solutions. 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 1969-1975.

Suresh Babu, V., & Behl, M. (2022). Threading the needle—Overtaking framework for multi-agent autonomous racing. SAE International Journal of Connected & Automated Veh., 5(1), 33-43.

Subosits, J. K., & Gerdes, J. C. (2019). From the Racetrack to the Road: Real-time trajectory replanning for autonomous driving. IEEE Transactions on Intelligent Vehicles, 4(2), 309-320.

Subosits, J. K., & Gerdes, J. C. (2021). Impacts of Model fidelity on trajectory optimization for autonomous vehicles in extreme maneuvers. IEEE Trans. on Intelligent Veh., 6(3), 546-558.

Suprapto, B. Y., Dwijayanti, S., Hafiz, D. M. A., Ardandy, F. A., and Jonathan, J. (2024). Designing an autonomous vehicle using sensor fusion based on path planning and deep learning algorithms, in SAIEE Africa Research Journal, 115(3), 86-98.

Sun, X., Zhou, M., Zhuang, Z., Yang, S., Betz, J., & Mangharam, R. (2023). A benchmark comparison of imitation learning-based control policies for autonomous racing. 2023 IEEE Intelligent Vehicles Symposium (IV), 1-5. https://doi.org/10.1109/iv55152.2023.10186780.

Szabo, D., & Pup, D. (2025). Optimal sensor placement for autonomous formula student vehicles: A field-of-view analysis of dual LIDAR and stereo camera configurations. Machines, 13(1), 27.

Tătulea-Codrean, A., Mariani, T., & Engell, S. (2020). Design and simulation of a machine-learning and model predictive control approach to autonomous race driving for the F1/10 Platform. IFAC-PapersOnLine, 53(2), 6031-6036.

Talvala, K. L., Kritayakirana, K., & Gerdes, J. C. (2011). Pushing the limits: From lane keeping to autonomous racing. Annual Reviews in Control, 35(1), 137-148.

Tengilimoglu, O., Carsten, O., Wadud, Z. (2023). Infrastructure requirements for the safe operation of automated vehicles: Opinions from experts and stakeholders, Transport Policy 133, 209-222.

Thakkar, R. S., Samyal, A. S., Fridovich-Keil, D., Xu, Z., & Topcu, U. (2024). Hierarchical control for head-to-head autonomous racing. Field Robotics, 4.0, 46-69.

Theodosis, P. A., & Gerdes, J. C. (2012). Nonlinear optimization of a racing line for an autonomous racecar using professional driving techniques. Volume 1: Adaptive Control; Advanced Vehicle Propulsion Systems; Aerospace Systems; Autonomous Systems; Battery Modeling; Biochemical Systems; Control Over Networks; Control Systems Desig. https://doi.org/10.1115/dscc2012-movic2012-8620.

Theodosis, P. A., & Gerdes, J. C. (2011). Generating a racing line for an autonomous racecar using professional driving techniques. ASME 2011 Dynamic Systems and Control Conference and Bath/ASME Symposium on Fluid Power and Motion Control, Volume 2, 853-860.

Tian, H., Ni, J., & Hu, J. (2018). Autonomous driving system design for formula student driverless racecar. 2018 IEEE Intelligent Vehicles Symposium (IV), 1-6. https://doi.org/10.1109/ivs.2018.8500471.

Toaz, B. R., Bopardikar, S. D. (2025). Vector cost bimatrix games with applications to autonomous racing, Pre-prints, https://arxiv.org/abs/2507.05171

Tramacere, E., Luciani, S., Feraco, S., Bonfitto, A., & Amati, N. (2021). Processor-in-the-loop architecture design and experimental validation for an autonomous racing Vehicle. Applied Sciences, 11(16), 7225.

Trauth, R., Hobmeier, A., & Betz, J. (2024a). A reinforcement learning-boosted motion planning framework: Comprehensive generalization performance in autonomous driving. 2024 IEEE Intelligent Vehicles Symposium (IV), 2413-2420.

Trauth, R., Kaufeld, M., Geisslinger, M., & Betz, J. (2023a). Learning and adapting behavior of autonomous vehicles through inverse reinforcement learning. 2023 IEEE Intelligent Vehicles Symposium (IV), 1-8. https://doi.org/10.1109/iv55152.2023.10186668.

Trauth, R., Moller, K., Würsching, G., & Betz, J. (2024b). FRENETIX: A high-performance and modular motion planning framework for autonomous driving. IEEE Access, 12.0, 127426-127439. https://doi.org/10.1109/access.2024.3436835.

Trauth, R., Werner, F., & Wischnewski, A. (2023b). TUM autonomous motorsport: An autonomous racing software for the Indy Autonomous Challenge. Journal of Field Robotics, 40(4), 783-809.

Treat, J. R., Tumbas, N. S., McDonald, S. T., Shinar, D., Hume, R. D., Mayer, R. E., Stansifer, R.L., Castellan, N. J., McDonald, S.T.(1979). Tri-level study of the causes of traffic accidents. (Technical Report No. DOT HS 805 086). Prepared by the Institute for Research in Public Safety, for the U.S. Department of Transportation, Washington, D.C. https://rosap.ntl.bts.gov/view/dot/25794/dot_25794_DS1.pdf?

TRB (Transportation Research Board). 2014. Road vehicle automation. Burlingame, CA: TRB.

Tsai, C. C., Yu, C. C., & Tsai, C. T. (2019). Adaptive ORFWNN-Based Predictive PID Control. International Journal of Fuzzy Systems, 21(5), 1544-1559.

TUM CAT (2015). Estimation and prediction of the tire-road friction potential of an autonomous racecar, Future Transport Mobility, Chair of Automotive Technology (CAT), https://www.mos.ed.tum.de/en/ftm/main-research/intelligent-vehicle-systems/estimation-and-prediction-of-the-tire-road-friction-potential-of-an-autonomous-racecar/. Accessed Dec 17, 2025.

Vallon, C., & Borrelli, F. (2020). Task decomposition for iterative learning model predictive control. 2020 American Control Conference, 2024-2029. doi.org/10.23919/acc45564.2020.9147625.

Valls, M. I., Hendrikx, H. F., Reijgwart, V. J., Meier, F. V., Sa, I., Dube, R., Gawel, A., Burki, M., & Siegwart, R. (2018). Design of an autonomous racecar: Perception, state estimation and system integration. 2018 IEEE International Conference on Robotics and Automation (ICRA), 2048-2055. https://doi.org/10.1109/icra.2018.8462829.

Vasco, M., Seno, T., Kawamoto, K., Subramanian, K., Wurman, P., Stone, P. (2024). A super-human vision-based reinforcement learning agent for autonomous racing in Gran Turismo, Reinforcement Learning Conference (RLC), 2024

Verschueren, R., Zanon, M., Quirynen, R., & Diehl, M. (2016). Time-optimal race car driving using an online exact hessian based nonlinear MPC algorithm. 2016 European Control Conference (ECC), 141-147. https://doi.org/10.1109/ecc.2016.7810277.

Vilsan, A., Sandu, C., Anghelache, G., & Warfford, J. (2025). Methodology for real-time hydroplaning risk estimation using an intelligent tire system: An analytical approach. Sensors, 25(23), 7299.

Vödisch, N., Dodel, D., & Schötz, M. (2022). Formula student objects in context (FSOCO): A dataset and benchmark for autonomous racing perception. SAE International Journal of Connected and Automated Vehicles, 5(1), 23-32.

Voser, C., Hindiyeh, R. Y., & Gerdes, J. C. (2010). Analysis and control of high sideslip maneuvers. Vehicle System Dynamics, 48(sup1), 317-336.

Wachter, E., Schmeitz, A., Bruzelius, F., & Alirezaei, M. (2020). Path control in limits of vehicle handling: A sensitivity analysis. Lecture Notes in Mechanical Engineering, 1089-1095. https://doi.org/10.1007/978-3-030-38077-9_126.

Wachter, E., Alirezaei, M., Bruzelius, F., & Schmeitz, A. (2019). Path control in limit handling and drifting conditions using State Dependent Riccati Equation technique. Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Eng., 234(2-3), 783-791.

Wadekar, S. N., Schwartz, B. J., Kannan, S. S., Mar, M., Manna, R. K., Chellapandi, V., Gonzalez, D. J., & Gamal, A. E. (2021). Towards end-to-end deep learning for autonomous racing: On data collection and a unified architecture for steering and throttle prediction. arXiv. Https://arxiv.org/abs/2105.01799

Wagener, N., Cheng, C. A., Sacks, J., & Boots, B. (2019). An online learning approach to model predictive control. Robotics: Science and Systems XV. https://doi.org/10.15607/rss.2019.xv.033.

Wang, M., Hoogendoorn, S., Daamen, W., van Arem, B., Happee, R. (2015). Game theoretic approach for predictive lane-changing and car-following control. Trans. Res C: Emerging Tech., 58, 73–92.

Wang, M., Wang, Z., Talbot, J., Christian Gerdes, J., & Schwager, M. (2019). Game theoretic planning for self-driving cars in competitive scenarios. Robotics: Science and Systems XV.

Wang, J., Chi, W., Li, C., Wang, C., Meng, M. Q. (2020). Neural RRT*: Learning-based optimal path planning, in IEEE Transactions on Automation Science and Engineering 17(4), 1748-1758.

Wang, N., Shi, X., Zhong, K., Zhang, X., & Chen, W. (2021a). A path correction method based on global and local matching for robotic autonomous systems. J. of Intelligent & Robotic Systems, 104(1).

Wang, M., Wang, Z., Talbot, J., Gerdes, J. C., & Schwager, M. (2021b). Game-theoretic planning for self-driving cars in multivehicle competitive scenarios. IEEE Trans. on Robotics, 37(4), 1313-1325.

Wang, C., & Li, Y. (2022). Motion prediction for autonomous vehicles based on EfficientNet-B1. 2022 4th International Conference on Communications, Information System and Computer Engineering (CISCE), 648-651. https://doi.org/10.1109/cisce55963.2022.9851007.

Wang, D., Kaufeld, M., & Betz, J. (2025a). DualAD: Dual-layer planning for reasoning in autonomous driving. 2025 IEEE/RSJ Intern. Conf. on Intelligent Robots & Systems (IROS), 12057-12063.

Wang, R., Sun, H., Lin, Y., Zhang, X., & Yin, Y. (2025b). Improving generalization in meta-learning via meta-gradient augmentation. Procs., 34th International Joint Conference on Artificial Intelligence, 6388-6396. https://doi.org/10.24963/ijcai.2025/711.

Wang, Y., Hu, S., Wang, X., Gao, Y., Zhang, W., Chen, Y., Chen, Y., Lin, H., Gao, T., Chen, J., Gong, X., Wang, B., Liu, W. (2025c). YOLOv13-Cone-Lite: A lightweight cone detection network for autonomous racing. Applied Sciences 15(17), 10.3390/app15179501.

Wang, Y., Liu, J., Yuan, X., Yang, J. (2025d). Safe reinforcement learning for competitive autonomous racing: Integrated state–action mapping and exploration guidance. Actuators. 14(7):315.

Wang, Y., Wang, J., Yang, Y., Li, Z., & Zhao, X. (2023). An end-to-end deep reinforcement learning model based on proximal policy optimization algorithm for autonomous driving of off-road vehicle. Lecture Notes in Electrical Engineering, 2692-2704.

Wang Z, Spica R, Schwager M. (2020). Game theoretic motion planning for multi-robot racing. Distributed Autonomous Robotic Syst 9:225–38. doi:10.1007/978-3-030-05816-6_16.

Weiss, T., & Behl, M. (2020a). DeepRacing: A framework for autonomous racing. 2020 Design, Automation & Test in Europe Conference & Exhibition (DATE), 1163-1168.

Weiss, T., & Behl, M. (2020b). DeepRacing: Parameterized trajectories for autonomous racing. arXiv. https://doi.org/10.48550/arXiv.2005.05178

Welde, Y., Qiao, F. (2020). Effects of autonomous and automated vehicles on stopping sight distance and vertical curves in geometric design,13th Asia Pacific Transportation Development Conference, American Society of Civil Engineers, Resilience and Sustainable Transportation Systems, 715 – 724.

Werner, F., Sagmeister, S., Piccinini, M., & Betz, J. (2025a). A quasi-steady-state black box simulation approach for the generation of G-G-G-V diagrams. 2025 IEEE Intell. Veh. Sym., 2503-2509.

Werner, F., Schwehn, A. K., Lienkamp, M., & Betz, J. (2025b). GripMap: An efficient, spatially resolved constraint framework for offline and online trajectory planning in autonomous racing. 2025 IEEE Intelligent Vehicles Symposium (IV), 103-110. https://doi.org/10.1109/iv64158.2025.11097550.

Werner, F., Oberhuber, R., & Betz, J. (2024). Accelerating autonomy: Insights from pro racers in the era of autonomous racing - An expert interview study. 2024 IEEE Intelligent Vehicles Symposium (IV), 329-336. https://doi.org/10.1109/iv55156.2024.10588407.

Wierwille, W. W., Hanowski, R. J., Hankey, J. M., Kieliszewski, C. A., Lee, S. E., Medina, A., Dingus, T. A. (2002). Identification and evaluation of driver errors: Overview and recommendations (Report No. FHWA-RD-02-003). Washington, DC: Federal Highway Administration.

Williams, G., Drews, P., Goldfain, B., Rehg, J. M., & Theodorou, E. A. (2016). Aggressive driving with model predictive path integral control. 2016 IEEE International Conference on Robotics and Automation (ICRA), 1433-1440.

Williams, G., Goldfain, B., Drews, P., Rehg, J.M., Theodorou, E.A. (2017a). Autonomous racing with AutoRally vehicles and differential games. arXiv preprints, https://arxiv.org/abs/1707.04540.

Williams, G., Aldrich, A., & Theodorou, E. A. (2017b). Model predictive path integral control: From theory to parallel computation. Journal of Guidance, Control, and Dynamics, 40(2), 344–357. https://doi.org/10.2514/1.G001921

Williams, G., Wagener, N., Goldfain, B., Drews, P., Rehg, J. M., Boots, B., & Theodorou, E. A. (2017c). Information theoretic MPC for model-based reinforcement learning. 2017 IEEE International Conference on Robotics and Automation (ICRA), 1714-1721. https://doi.org/10.1109/icra.2017.7989202.

Williams, G., Drews, P., Goldfain, B., Rehg, J. M., & Theodorou, E. A. (2018a). Information-theoretic model predictive control: Theory and applications to autonomous driving. IEEE Transactions on Robotics, 34(6), 1603-1622.

Williams, G., Goldfain, B., Drews, P., Saigol, K., Rehg, J., & Theodorou, E. (2018b). Robust sampling based model predictive control with sparse objective information. Robotics: Science and Systems XIV. https://doi.org/10.15607/rss.2018.xiv.042.

Williams, G., Goldfain, B., Drews, P., Rehg, J. M., & Theodorou, E. A. (2018c). Best response model predictive control for agile interactions between autonomous ground vehicles. 2018 IEEE International Conference on Robotics and Automation (ICRA), 2403-2410. https://doi.org/10.1109/icra.2018.8462831.

Wischnewski, A., Euler, M., Gümüs, S., & Lohmann, B. (2021). Tube model predictive control for an autonomous race car. Vehicle System Dynamics, 60(9), 3151-3173.

Wischnewski, A., Betz, J., & Lohmann, B. (2019a). A model-free algorithm to safely approach the handling limit of an autonomous racecar. 2019 IEEE International Conference on Connected Vehicles and Expo (ICCVE), 1-6. https://doi.org/10.1109/iccve45908.2019.8965218.

Wischnewski, A., Betz, J., & Lohmann, B. (2020). Real-time learning of non-gaussian uncertainty models for autonomous racing. 2020 59th IEEE Conference on Decision and Control (CDC), 609-615. https://doi.org/10.1109/cdc42340.2020.9304230.

Wischnewski, A., Geisslinger, M., Betz, J., Betz, T., Fent, F., Heilmeier, A., Hermansdorfer, L., Herrmann, T., Huch, S., Karle, P., Nobis, F., Ögretmen, L., Rowold, M., Sauerbeck, F., Stahl, T., Trauth, R., Lienkamp, M., & Lohmann, B. (2022). Indy Autonomous Challenge - Autonomous race cars at the handling limits. Proceedings, 163-182. https://doi.org/10.1007/978-3-662-64550.

Wischnewski, A., Stahl, T., Betz, J., & Lohmann, B. (2019b). Vehicle dynamics state estimation and localization for high performance race cars. IFAC-PapersOnLine, 52(8), 154-161.

Wolfgang Kuhn, P. (2017). Methodology for the numerical calculation of racing lines and the virtual assessment of driving behavior for training circuits for the automobile industry. Transportation Research Procedia, 25.0, 1416-1429.

Wurman, P. R., Barrett, S., Kawamoto, K., MacGlashan, J., Subramanian, K., Walsh, T. J., Capobianco, R., Devlic, A., Eckert, F., Fuchs, F., Gilpin, L., Kapturowski, P., Kompella, V., Lin, H., MacAlpine, P., Mirowski, P., Muller, T. M., Nakagawa, S., Ortega, P. A., ... Kitano, H. (2022a). Outracing champion Gran Turismo drivers with deep reinforcement learning. Nature, 602(7896), 223–228.

Wurman, P. R., Stone, P., & Spranger, M. (2022b). Challenges and opportunities of applying reinforcement learning to autonomous racing. IEEE Intelligent Systems, 37(3), 20-23.

Wohner, B., Sevilla, F., Grueter, B., Diepolder, J., Afonso, R., & Holzapfel, F. (2021). Hierarchical nonlinear model predictive control for an autonomous racecar. 2021 20th International Conference on Advanced Robotics (ICAR), 113-120.

Xia, T., & Chen, H. (2024). A survey of autonomous vehicle behaviors: Trajectory planning algorithms, sensed collision risks, and user expectations. Sensors, 24(15), 4808.

Xiao, X., Biswas, J., & Stone, P. (2021). Learning inverse kinodynamics for accurate high-speed off-road navigation on unstructured terrain. IEEE Robotics and Automation Letters, 6(3), 6054-6060.

Yamaguchi, T., & Matsumoto, T. (2022). Autonomous control for cooperative operation between energy storage systems. 2022 International Power Electronics Conference (IPEC-Himeji 2022- ECCE Asia), 2351-2358. https://doi.org/10.23919/ipec-himeji2022-ecce53331.2022.9806848.

Yang, J., Wang, Y., Dong, L., & Yuan, X. (2024). Reinforcement learning with safe action generation for autonomous racing. 2024 IEEE Congress on Evolutionary Computation (CEC), 1-8. https://doi.org/10.1109/cec60901.2024.10611776.

Yang, B., Lu, Y., Yang, X., & Mo, Y. (2022). A hierarchical control framework for drift maneuvering of autonomous vehicles. 2022 International Conference on Robotics and Automation (ICRA), 1387-1393. https://doi.org/10.1109/icra46639.2022.9812110.

Yoo J.-H. and Langari, R. (2012). Stackelberg Game based model of highway driving, in ASME Dynamic Systems and Control Conference (DSCC), Ft. Lauderdale, FL, 499-508.

You, C., & Tsiotras, P. (2018). Real-time trail-braking maneuver generation for off-road vehicle racing. 2018 Annual American Control Conference (ACC), 4751-4756. https://doi.org/10.23919/acc.2018.8431620.

You, C., & Tsiotras, P. (2017). Vehicle modeling and parameter estimation using adaptive limited memory joint-state UKF. 2017 American Control Conference (ACC), 322-327.

You, C., & Tsiotras, P. (2021). High-speed cornering for autonomous off-road rally racing. IEEE Transactions on Control Systems Technology, 29(2), 485-501.

Yuan, M., Chen, Z., Yao, B., Hu, J. (2018). A general online trajectory planning framework in the case of desired function unknown in advance, in IEEE Trans. on Ind. Informatics, 15(5), 2753-2762.

Yuan, Q., Yan, F., Yin, Z., Chen, L., Hu, J., Wu, D., Li, Y. (2023). Decision-making and planning methods for autonomous vehicles based on multistate estimations and game theory, Advanced Intelligent Systems 5(11), 2300177.

Yuan, Y., Zhang, K., Cheng, Q., Liu, Z., & Wang, W. (2019). Traffic states recognition and prediction based on floating car data. CICTP 2019, 2236-2248. https://doi.org/10.1061/9780784482292.195.

Zadok, D., Hirshberg, T., Biran, A., Radinsky, K., & Kapoor, A. (2019). Explorations and lessons learned in building an autonomous formula SAE Car from simulations. Procs., 9th International Conference on Simulation and Modeling Methodologies, Tech. & Applications, 414-421.

Zakaria, W. N. W., Mahmood, I. A. T., Shamsudin, A. U., Rahman, M. A. A., & Tomari, M. R. M. (2022). ROS-based SLAM and path planning for autonomous unmanned surface vehicle navigation system. 2022 IEEE 5th International Symposium in Robotics and Manufacturing Automation (ROMA), 1-6. https://doi.org/10.1109/roma55875.2022.9915665.

Zang, Z., Tumu, R., Betz, J., Zheng, H., Mangharam, R. (2022). Winning the 3rd Japan automotive AI challenge–autonomous racing with the Autoware Auto Open Source Software Stack. arXiv preprint arXiv:2206.00770.

Zarrouki, B., Spanakakis, M., & Betz, J. (2024a). A safe reinforcement learning driven weights-varying model predictive control for autonomous vehicle motion control. 2024 IEEE Intelligent Vehicles Symposium (IV), 1401-1408. https://doi.org/10.1109/iv55156.2024.10588747.

Zarrouki, B., Wang, C., & Betz, J. (2024b). A stochastic nonlinear model predictive control with an uncertainty propagation horizon for autonomous vehicle motion control. 2024 American Control Conference (ACC), 5466-5473. https://doi.org/10.23919/acc60939.2024.10645032.

Zarrouki, B., Wang, C., & Betz, J. (2024c). Adaptive stochastic nonlinear model predictive control with look-ahead deep reinforcement learning for autonomous vehicle motion control. 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 12726-12733. https://doi.org/10.1109/iros58592.2024.10801876.

Zarrouki, B., Nunes, J., & Betz, J. (2024d). R²NMPC: A real-time reduced robustified nonlinear model predictive control with ellipsoidal uncertainty sets for autonomous vehicle motion control. IFAC-PapersOnLine, 58(18), 309-316. https://doi.org/10.1016/j.ifacol.2024.09.048.

Zeng, D., Yu, Z., Xiong, L., Zhao, J., Zhang, P., Li, Y., Xia, L., Wei, Y., Li, Z., Fu, Z. (2021). Driving-behavior-oriented trajectory planning for autonomous vehicle driving on urban structural road, Procs., Institution of Mech. Eng. Part D: J. of Automobile Engineering 235(4), 975-995.

Zenoh (2021). Indy Autonomous Challenge (IAC): Experiences from the Trenches, https://zenoh.io/blog/2021-09-28-iac-experiences-from-the-trenches/

Zhao, Y., Ying, X., & Li, J. (2022). Research on geometric design standards for freeways under a fully autonomous driving environment. Applied Sciences, 12(14), 7109.

Zhao, T., Chen, G., Wang, X., Yong, E., & Qian, W. (2021). Aerodynamic modeling using an end-to-end learning attitude dynamics network for flight control. Acta Mechanica Sinica, 37(12), 1799-1811.

Zhao, J., Su, S., Liu, X., Wang, J., Tang, D., Wang, Y. (2025). Numerical and experimental investigations of the aerodynamic drag characteristics and reduction of an autonomous vehicle, Physics of Fluids, 37(1), 015152.

Zhang, D., Yuan, Y. & Zhu, F. (2025a) Global minimum time trajectory planning considering curvature and distance for track racing. Scientific Reports 15, 37264.

Zhang, K., & Tang, X. (2025b). Risk-aware motion planning for collision avoidance under uncertainty. IEEE Int. Conf. on Robotics and Automation (ICRA). https://doi.org/10.2139/ssrn.5193272.

Zhang, M., Liu, Y., Wang, Y., Li, F., & Chen, L. (2022a). Real-time path planning algorithms for autonomous UAV. 2022 China Automation Congress (CAC), 3818-3823. https://doi.org/10.1109/cac57257.2022.10054770.

Zhang, R., Mak, S., & Dunson, D. (2022b). Gaussian process subspace prediction for model reduction. SIAM Journal on Scientific Computing, 44(3), A1428-A1449.

Zhang, R., Hou, J., Chen, G., Li, Z., Chen, J., & Knoll, A. (2022c). Residual policy learning facilitates efficient model-free autonomous racing. IEEE Robotics & Automation Lett., 7(4), 11625-11632.

Zhang, T., Sun, Y., Wang, Y., Li, B., Tian, Y., & Wang, F. Y. (2024d). A survey of vehicle dynamics modeling methods for autonomous racing: Theoretical models, Physical/Virtual Platforms, and Perspectives. IEEE Transactions on Intelligent Vehicles, 9(3), 4312-4334.

Zhang, Z., He, H., Guo, J., & Han, R. (2020). Velocity prediction and profile optimization based real-time energy management strategy for Plug-in hybrid electric buses. Applied Energy, 280.0, 116001.

Zheng, H., Betz, J., Ramamurthy, A., Jin, H., & Mangharam, R. (2022). Combinatorial and parametric gradient-free optimization for cyber-physical system design. 2022 IEEE Workshop on Design Automation for CPS and IoT (DESTION), 34-41. https://doi.org/10.1109/destion56136.2022.00012.

Zhou, A., Wang, C., Wang, J., Wang, Y., Peeta, S. (2022). Cooperative control mechanism for platoon formation of connected and autonomous vehicles, CCAT Report #25, Center for Connected and Automated Transportation, Purdue University, West Lafayette, IN.

Zhou, B., Zarrouki, B., Piccinini, M., Hu, C., Xie, L., & Betz, J. (2025). Safe reinforcement learning with a predictive safety filter for motion planning and control: A drifting vehicle example. 2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 4196-4203. https://doi.org/10.1109/iros60139.2025.11246272.

Zhuo, J., Li, G., & He, Y. (2023). Research on cone bucket detection algorithm based on improved YOLOv5s. World Electric Vehicle Journal, 14(10), 269.

Zubaca, J., Stolz, M., & Watzenig, D. (2020a). Extended H_∞ Filter adaptation based on innovation sequence for advanced ego-vehicle motion estimation. 2020 IEEE 3rd Connected and Automated Vehicles Symposium (CAVS), 1-5. https://doi.org/10.1109/cavs51000.2020.9334568. https://doi.org/10.15607/rss.2019.xv.048.

Zubaca, J., Stolz, M., & Watzenig, D. (2020b). Smooth reference line generation for a racetrack with gates based on defined borders. 2020 IEEE Intelligent Vehicles Symposium (IV), 604-609. https://doi.org/10.1109/iv47402.2020.9304722.

APPENDIX 1 - PHOTOS OF SELECTED RACE TRACKS AND RACE COURSES

(a) The Indianapolis Motor Speedway Complex

(b) The Las Vegas Motor Speedway

(c) Lucas Oil Indianapolis Raceway Park

(d) Texas Motor Speedway, Fort Worth, Texas

Sources: (a) www.visitindy.com/directory/indianapolis-motor-speedway/
(b) www.espn.com/espnw/news-commentary/story/_/id/2346802/las-vegas-motor-speedway
(c) www.iracing.com/tracks/indianapolis-motor-speedway/
(d) athlonsports.com/racing/nascar/texas-motor-speedway-2025-track-profile,

Figure A.1(a) Oval racetracks that have hosted autonomous racing events

(a) Yas Marina Circuit, Abu Dhabi, UAE

(b) Circuito Monteblanco, Spain. Home of most Roborace events

(c) Anglesey Circuit, Ty Croes, Wales

(d) Moulay El Hassan Formula E circuit, Morrocco

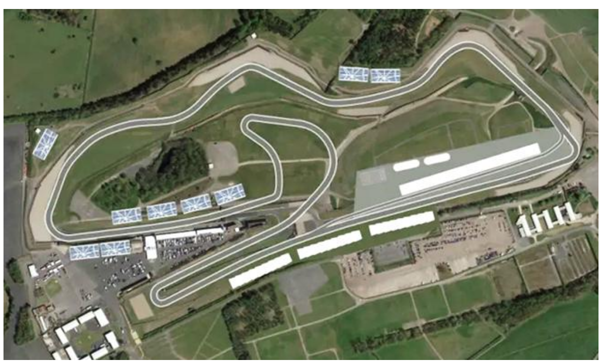

(e) Donington Park Circuit, Donington, England

(f) Monza National Racetrack, Monza, Italy

Sources: (a) https://www.musco.com/project/yas/
(b) bikeport.no/eu/monteblanco-motogp-portimao-november-3-10-2021/
(c) racecollective.com/tracks/anglesey-circuit/
(d) tilke.de/portfolio/circuit-moulay-el-hassan/
(e) allalongtheracetrack.co.uk/2021/07/31/donington-park-donington/
(f) rookief1.com/2011/09/09/monza-what-you-need-to-know/

Figure A.1(b) Road courses that have hosted autonomous racing events

APPENDIX 2 - PHOTOS AT RACE EVENTS WITH PARTICIPATION OR HOSTING BY THE REPORT AUTHORS

Figure A.3(a) Co-author of this report with Purdue's University's full-size autonomous racecar at the Las Vegas Motor Speedway in 2024. Purdue University participated in this IAC event.

Figure A.3(b) Purdue's University's autonomous go-kart racecar at the Purdue Grand Prix Track in 2026. Purdue's Lyles School of Civil and Construction Engineering hosted this AKS event.

APPENDIX 3 - NEW UNDERGRADUATE COURSES AT PURDUE MOTIVATED BY THIS PROJECT

1. ARES: AUTONOMOUS RACE-ENGINEERING SYSTEMS (VIP COURSE)

Introduction: The ARES (Autonomous Race-Engineering Systems) course explores robotics, perception, controls, AI-ML methods, and adversarial racing strategies to develop solutions to address a variety of multi-disciplinary problems in high-speed autonomous Racing (AR) with real world applications across multiple modes and platforms.

Course Advisors: Samuel Labi

Course Instructors and Mentors: Andres Moreno, Joseph Reed, Richard Ajagu

Couse Description: This VIP course serves as a testbed for advancing research and innovation in self driving technology through the study of Autonomous Racing Systems and its real world applications. The course explores latest advancements in robotics, AI-ML methods, Race Theory and Engineering skills to develop solutions to address a variety of multi-disciplinary problems across multiple modes and platforms in Autonomous Racing (AR).

Areas of interest: (a) controls, localization, path planning and navigation in high-speed AR environments, adversarial decision making, and safety-critical edge cases, and/or
(b) other aspects/applications of AR beyond the track, to include perspectives of road infrastructure design, human factors, economics, and other real-world contexts of AR systems.
Different Topics that can be explored on this course: Cyber-Physical AI, Deep Learning, Control Systems, High-speed autonomy, Multi-agent Adversarial Strategy and Decision-making, Overtaking, Sensor Fusion, ROS, Platform Optimization, Simulation (Digital Twin), Reinforcement Learning, Software Development, System Integration, Real-world Applications & Systems Analysis, Infrastructure Design, and the Economics and Business applications of AR Systems.

Eligibility: All levels (freshmen, sophomores, juniors and seniors) in any department.

Application Platforms: Full size racecars, Go Karts, Small-size road/track race cars, F1/10th, Boats, Submarine, Drones and Rally/Offroad.

2. AUTONOMOUS MOTORSPORTS PURDUE (AMP) (VIP COURSE)

Introduction: Bridging motorsports and autonomy, AMP VIP works together with the AMP Student Club to build and autonomize go-kart platforms for autonomous go-kart racing. One of the main competitions AMP participates in is the Autonomous Kart Series (AKS), an inter-college go-kart race that takes place at the Purdue Gran Prix track, where teams test their autonomous kart and compete, pushing the limits of their creation at the track. The building, development and implementation of such platforms involves challenging hardware, software and mechatronics problems that students can choose to work on with highly applicable skills sought after by industry.

Course Advisors: Samuel Labi

Course Instructors and Mentors: Andres Moreno, Joseph Reed, Richard Ajagu

Description: AMP VIP works together with the AMP Student Club to build and autonomize go-kart platforms for autonomous go-kart racing. One of the main competitions AMP participates in is the Autonomous Kart Series (AKS), an inter-college go-kart race that takes place at the Purdue Grand Prix track, where teams test their autonomous kart and compete, pushing the limits of their creation at the track. The building, development and implementation of such platforms involves challenging hardware, software and mechatronics problems that students can choose to work on with highly applicable skills for industry. Moreover, this VIP gives students the opportunity to pursue research of their interest related to autonomous kart racing, where high-speeds, precision, accuracy, robustness and safety are imperative.

Areas of Interest: Design, construction, optimization and "autonomization" of go-kart platforms; Digital Twin and Simulation of autonomous go-karts; Cyber-physical AI; Software development for path planning, perception, controls, localization, sensor fusion, ML, AI, RL; Racing strategies. (Raceline optimization, overtaking strategies, braking and throttling strategies); Business applications in AMP (sponsorship, business strategy, marketing, publicity)

Pre-requisites: The following skill is required to enter this VIP: Willingness to learn.
The following skills are NOT a requirement for entering, but will be aligned with this VIP: Mechanical, electrical, electromechanical, and/or mechatronics design; Software development (Robot Operating System (ROS), Matlab, Simulink, Python, C++, AI, ML); Ability to create, improve and properly use simulation tools; Building circuits, mechanical assemblies, CAD; Good teamwork and communication; Experience in robotics, autonomous vehicles.

3. ROBOAT: AUTONOMOUS MARITIME MANEUVERS (FORMERLY NSWC AIMM)

Introduction: RoBoat participates in the US Navy's AIMM ICC competition, where students fit sensors on an autonomous boat and develop the software to pass a series of challenges posed in the competition. This project involves hands on problems, development of the sensor suite for the boat, algorithms capable of autonomizing the boat, including perception, path planning, controls, localization, integration and simulation.

Course Advisors: Brittany Newell and Samuel Labi

Course Instructors and Mentors: Andres Moreno, Joseph Reed, Richard Ajagu

Description: RoBoat participates in the AIMM ICC competition, where students install sensors on a boat and develop the software to pass the series of challenges. This project involves hands on problems, development of the sensor suite for the boat, algorithms capable of autonomizing the boat, including perception, path planning, controls, localization, integration and simulation. This experience provides in-depth learning in important fields like machine learning, artificial intelligence, circuits, mechatronics, and naval engineering, under the fun and excitement of a competition. Offering opportunities for students to delve into hot research topics highly applicable in modern-day technical area related industry.

Website: For more information on the competition, visit the website:
https://www.trine.edu/innovation-one/aimm/index.aspx?_ga=2.32049085.2083888718.1750167314-1593673814.1750167314

Areas of Interest: Naval engineering; autonomous driving and robotics; design and building of an autonomous boat; sensor suite, at hardware level and software level; physical AI; software development: robot operating system, ML, deep learning, path planning, localization, controls, system dynamics, perception; artificial intelligence; machine learning; maritime maneuvering.

APPENDIX 4 PUBLICATION ABSTRACTS

Seilabi, S.E., Pourgholamali, M., Miralinaghi, M., Correia, G., Labi, S. (2024). Optimizing dedicated lanes and tolling schemes for connected and autonomous vehicles to address bottleneck congestion considering morning commuter departure choices, *Journal of Intelligent Transportation Systems*, doi.org/10.1080/15472450.2024.2408024

The introduction of connected and autonomous vehicles (CAVs) provides a significant opportunity to address the persistently increasing problem of urban traffic congestion. By virtue of their connectivity and automation features, CAVs can reduce vehicle headways, thereby increasing road capacity and enhancing throughput. It has been hypothesized that CAV-infrastructure design policies can influence traveler behavior in ways that could reduce congestion. This research focuses on the potential of using CAV-dedicated lanes (CAVL) to alleviate traffic congestion in a bottleneck corridor that serves both human-driven vehicles (HDVs) and CAVs. We delve into investigating the impacts of CAVLs on the departure time and lane choices of morning commuters. The study first expresses traffic equilibrium conditions as a linear program with complementarity constraints. Then, a system-optimal commute congestion management design is formulated to minimize the overall system cost, which consists of queuing delays and early and late arrival costs. The results of the computational experiments suggest that: (i) the CAV technological advancements can significantly reduce traffic congestion under CAVL deployment with an almost similar effect as a tolling policy; and (ii) the lower value of time for CAV commuters leads them to depart closer to their desired arrival time without a tolling policy, which could significantly increase the bottleneck traffic congestion that commuters experience, particularly HDVs.

Chen, S., Zong, S., Chen, T., Huang, Z., Chen, Y., & Labi, S. (2023). A taxonomy for autonomous vehicles considering ambient road infrastructure. *Sustainability*, 15(14), 11258. https://doi.org/10.3390/su151411258

To standardize definitions and guide the design, regulation, and policy related to automated transportation, the Society of Automotive Engineers (SAE) has established a taxonomy consisting of six levels of vehicle automation. The SAE taxonomy defines each level based on the capabilities of the automated system. It does not fully consider the infrastructure support required for each level. This can be considered a critical gap in the practice because the existing taxonomy does not account for the fact that the operational design domain (ODD) of any system must describe the specific conditions, including infrastructure, under which the system can function. In this paper, we argue that the ambient road infrastructure plays a critical role in characterizing the capabilities of autonomous vehicles (AVs) including mapping, perception, and motion planning, and therefore, the current taxonomy needs enhancement. To throw more light and stimulate discussion on this issue, this paper reviews, analyzes, and proposes a supplement to the existing SAE levels of automation from a road infrastructure perspective, considering the infrastructure support required for automated driving at each level of automation. Specifically, we focus on Level 4 because it is expected to be the most likely level of automation that will be deployed soon. Through an analysis of driving scenarios and state-of-the-art infrastructure technologies, we propose five sub-levels for Level 4 automated driving systems: Level 4-A (Dedicated Guideway Level), Level 4-B (Expressway Level), Level 4-C (Well-Structured Road Level), Level 4-D (Limited-Structured road Level), and Level 4-E (Disorganized Area Level). These sublevels reflect a progression from highly structured environments with robust infrastructure support to less structured environments with limited or no infrastructure support. The proposed supplement to the SAE taxonomy is expected to benefit both potential AV consumers and manufacturers through defining clear expectations of AV performance in different environments and infrastructure settings. In addition, transportation agencies may gain insights from this research towards their planning regarding future infrastructure improvements needed to support the emerging era of driving automation.

Dong, J., Chen, S., Miralinaghi, M., Chen, T., Li, P., Labi, S. (2023). Why did the AI make that decision? Towards an explainable artificial intelligence (XAI) for autonomous driving systems, *Transportation Research Part C: Emerging Technologies* 156, 104358

User trust has been identified as a critical issue that is pivotal to the success of autonomous vehicle (AV) operations where artificial intelligence (AI) is widely adopted. For such integrated AI-based driving systems, one promising way of building user trust is through the concept of explainable artificial intelligence (XAI) which requires the AI system to provide the user with the explanations behind each decision it makes. Motivated by both the need to enhance user trust and the promise of novel XAI technology in addressing such need, this paper seeks to enhance trustworthiness in autonomous driving systems through the development of explainable Deep Learning (DL) models. First, the paper casts the decision-making process of the AV system not as a classification task (which is the traditional process) but rather as an image-based language generation (image captioning) task. As such, the proposed approach makes driving decisions by first generating textual descriptions of the driving scenarios, which serve as explanations that humans can understand. To this end, a novel multi-modal DL architecture is proposed to jointly model the correlation between an image (driving scenario) and language (descriptions). It adopts a fully Transformer-based structure and therefore has the potential to perform global attention and imitate effectively, the learning processes of human drivers. The results suggest that the proposed model can and does generate legal and meaningful sentences to describe a given driving scenario, and subsequently to correctly generate appropriate driving decisions in autonomous vehicles (AVs). It is also observed that the proposed model significantly outperforms multiple baseline models in terms of generating both explanations and driving actions. From the end user's perspective, the proposed model can be beneficial in enhancing user trust because it provides the rationale behind an AV's actions. From the AV developer's perspective, the explanations from this explainable system could serve as a debugging tool to detect potential weaknesses in the existing system and identify specific directions for improvement.

Pourgholamali, M., Miralinaghi, M., Seilabi, S.E., Labi, S. (2023). Sustainable deployment of autonomous vehicles dedicated lanes in urban traffic networks, *Sustainable Cities and Society* 99, 104969

Autonomous vehicles (AVs) show promise for increasing roadway safety and capacity. During the AV transition era, which will be characterized by a mixed fleet of AVs and human-driven vehicles (HDVs), it is expected that the allure of these prospective benefits will motivate road agencies to allocate AV-dedicated lanes. This paper proposes a sustainability-driven AV dedicated lane and pricing policy (SALP) framework that addresses the three pillars of sustainable development—social, environmental, and economic. The framework is formulated as a bilevel problem where the upper-level model yields decisions on the timing, location, and quantity of AV-dedicated lanes and tolling levels to minimize total travel time, emissions, and electricity consumption costs (that is, the economic and environmental pillars). To alleviate potential inequity (the social pillar), two considerations are proposed: revenue neutrality to compensate for the increase in travel costs of travelers and an equity constraint to limit the exacerbation of HDV travel costs. At the lower level, travelers react to the decisions made at the upper level by choosing their vehicle types (AV vs. HDV) and routes. The SALP is solved using Genetic and Frank-Wolfe algorithms. The results of the numerical experiments suggest that the proposed SALP addresses all three pillars, as it yields significant reductions in total travel time, emissions, and electricity costs.