# Autonomous navigation in urban scenarios: From street driving to last-mile delivery

Iván del Pino* and
Angel Santamaria-Navarro†

**Abstract** Autonomous driving aims to enhance road safety, reduce traffic congestion, and improve transportation efficiency through intelligent vehicle automation. In this chapter, we describe the evolution of the autonomous navigation of vehicles in urban environments. Starting from the first projects that presented their potential applications, we describe the history and current state of these autonomous vehicles, from those driving in public roads to autonomous last-mile delivery robots. Here, we present the main challenges and solutions of these projects and analyze their current technological trends. Further, we present the robot Ona, a prototype designed for last-mile delivery operations, providing details of its configuration, a comparison with respect of other existing solutions, and describing its navigation stack intricacies.

## 1 Introduction

In this chapter, we will focus on autonomous robots that are capable of navigating in urban scenarios autonomously and, specifically, those known as "Public-area mobile robots" or PMRs. This concept, defined in ISO DTS 4448: Intelligent Transport Systems - Public-area mobile robots [1], refers to terrestrial robots that, regardless of their means of locomotion (wheels, legs, etc.), perform tasks requiring mobility in public spaces such as sidewalks, bike lanes, hospitals, parks, or airports. Therefore,

---

Iván del Pino
Institut de Robòtica i Informàtica Industrial (IRI), C/ Llorens i Artigas 4-6, 08028 Barcelona, Spain, e-mail: `idelpino@iri.upc.edu`

Angel Santamaria-Navarro
Institut de Robòtica i Informàtica Industrial (IRI), C/ Llorens i Artigas 4-6, 08028 Barcelona, Spain, e-mail: `asantamaria@iri.upc.edu`

* ORCID: 0000-0003-0516-7443

† ORCID: 0000-0002-6328-1423

they share space with pedestrians of all ages and physical and cognitive abilities, and even their pets [2]. This definition explicitly excludes drones and industrial robots that operate in warehouses, factories, farms, or mines. However, the boundary with autonomous passenger transport vehicles is much more blurred, as road vehicles sometimes need to interact with pedestrians (for example, when picking up or dropping off passengers), and PMRs may also need to navigate areas designated for vehicular traffic, depending on the specific application. For this reason, although our focus will be on PMRs, we will discuss some topics related to driver-less cars when required. Given the breadth of the topic, we do not aim at covering all areas related to PMRs (e.g., ethical or legal issues, privacy considerations, infrastructure, communication challenges, etc.). Instead, we present an easy-to-follow description of the main related projects, together with their challenges and solutions, directed to the robotics research community.

Specifically, this chapter is structured as follows. First, in Section 2, we provide a historical context by reviewing the key milestones that have led to the current state of outdoor mobile robotics. Next, in Section 3, we examine the main challenges of autonomous passenger transportation in urban environments and explore the most common and advanced solutions to address them. We review the cases of companies in the robotic-taxi ("robotaxi") field, like Waymo, Cruise or Baidu. Subsequently, in Section 4 we focus on the autonomous last-mile delivery sector, exploring the different challenges and approaches used by companies in the sector. We also spot similarities and differences between self-driving and autonomous last-mile navigation. In Section 5, we present our autonomous navigation system developed for the robot Ona, designed for last-mile logistics applications. The results obtained with our system are analyzed in Section 6, where we evaluate its performance against the challenges posed. Finally, Section 7 presents the conclusions, highlighting the lessons learned and potential future research directions.

## 2 Evolution of Autonomous Navigation in Urban Environments

Mobile robotics has come a long way from its humble beginnings in small laboratory experiments to the current state of fully autonomous delivery robots and robotaxi companies operating commercially in various cities. This section briefly comments on the principal milestones along this path.

### 2.1 The origins

The research origins in autonomous robot navigation can be traced back to the mid-1960s with the Shakey project [3]. This project, active between 1966 and 1972, was conducted at the Artificial Intelligence Center of the Stanford Research Institute. The objective was to create a system capable of autonomously navigating and interacting with objects in the environment to accomplish missions specified by a person using natural language [4]. Due to the limitations in computing power and miniaturization of computers at that time, a radio-connected system was designed. Data capture from the environment (primarily using a television camera and two rangefinders) and low-level motor control (a differential drive plus a pan and tilt platform) were performed on the mobile platform: the Shakey robot. In contrast, data interpretation and reasoning took place on an external computer, which was essentially a supercomputer occupying several cabinets in the lab. The environment comprised a series of rooms connected by door-sized openings and populated with a few simple geometric objects (cubes and triangular prisms). Although not dynamic, these objects could be manually moved to block passage between rooms, requiring the system to re-plan the sequence of actions to accomplish the navigation mission. The research conducted within the Shakey project made significant contributions to mobile robotics, such as the $A^*$ heuristic planning algorithm [5] and the Hough transform [6] for detecting lines and curves in images. This project laid the foundation for the subsequent decades of research, influencing numerous areas related to Artificial Intelligence.

## 2.2 First outdoor navigation systems

While early indoor navigation research emerged in the 1960s, it wasn't until the 1980s that outdoor navigation systems began to take shape. This advancement was driven by increased computational power and the availability of more sophisticated sensors, such as LiDAR [7], [8]. A notable example is the autonomous navigation system developed at Carnegie Mellon University (CMU) [9]. For these investigations, they developed the Terregator [10], a six-wheeled skid-steer platform, weighting seven hundred kilograms, being almost two meters long, and equipping a gasoline-powered electric generator. This robot provided the robustness, autonomy, and power necessary to develop an autonomous navigation system in unstructured outdoor environments. Similar to Shakey, this platform had an on-board low-level processing system that generated the control signals needed to execute the maneuvers requested via a radio link. The navigation system was designed to operate within the university's pedestrian network, employing a topological map of intersections to divide navigation missions into route segments (e.g., "follow the path", "turn at the intersection"). These segments were further subdivided into driving units, short segments of a few meters where local planning and obstacle avoidance occurred. For localization, the system relied on dead reckoning algorithms combined with relative positioning techniques based on visual features like paths and intersections. A pioneering aspect of the CMU system was its distributed architecture, which allowed for parallel and asynchronous processing of various tasks (planning, perception, control) coordinated by a central process. This architecture, similar to modern systems like ROS [11], enabled automatic transformation of data between different reference frames.

## 2.3 The beginnings of road navigation

At the hardware level, one of the limitations for autonomous navigation research during the early decades was the significant volume, weight, and power consumption required by the computers of that time to execute algorithms of certain complexity within a reasonable time frame. This made it unfeasible for them to be embedded in mobile robots. Thus, the systems consisted of a rover whose sole mission was to collect data using sensors and execute low-level motor control, while the processing was carried out in an external laboratory with which a radio link was maintained. This architecture, which allowed the state of the art to advance during the early years, had limitations as experimentation was restricted to the vicinity of the laboratories, which was insufficient for developing many applications such as autonomous driving systems on roads. The solution adopted by several groups in the late 1980s was to create mobile laboratories, where large vehicles were used as robotic platforms on which a complete laboratory with sufficient computing power could be installed to begin real-time autonomous driving experiments. Thus, at CMU, the Navlab platform [12] was developed from a modified Chevrolet van, allowing it to be driven both manually and autonomously. This mobile laboratory provided enough space for four fully equipped workstations for onboard researchers and all the necessary equipment for data acquisition, processing, and platform control. In parallel with the work at CMU, Ernst Dickmanns' team at the Bundeswehr University of Munich also employed a mobile laboratory approach in developing their VaMoRs prototype, which was built from a large Mercedes-Benz van modified to allow platform control using computer vision techniques [13]. By recognizing landmarks like road markings and edges, these systems achieved high-speed autonomous control. Throughout the 1990s, autonomous highway driving continued to advance, with systems like those developed by Dickmanns' team accumulating thousands of kilometers of autonomous driving experience [14], [15]. These systems achieved speeds of up to 180 km/h and demonstrated complex maneuvers like lane changes and merging, thanks to real-time vehicle detection and tracking modules based on vision [16]. The integration of GPS technology, which reached full operational capability in 1995 [17], enabled autonomous vehicles to execute missions based on digital maps [16].

## 2.4 The DARPA Challenges

In the first decade of the 21st century, outdoor autonomous navigation competitions organized by the Defense Advanced Research Projects Agency (DARPA) significantly boosted research in this field [18]. The first DARPA Grand Challenge, held in 2004, was a test where teams had to guide their prototypes fully autonomously through a more than 200 km route in the Mojave Desert (California, USA). Although none of the 17 finalist teams managed to complete the route in this first edition, the event was considered a success due to the scientific and technical advances achieved, as well as the great media attention it generated. The following year, the second edition of the DARPA Grand Challenge was held, in which 5 of the 23 finalist teams managed to complete the route. Stanford University's prototype, Stanley, won the competition by completing the route in less than 7 hours, reaching an average speed of approximately 30 km/h [19]. These first editions of the Challenge were mainly developed on wide dirt tracks with gentle curves, although they also included paved sections and more challenging sections with uneven terrain. The route to follow was provided to the teams two hours before the start of the race in a specific data format (RDDF). This format contained a series of geographic points (waypoints) with information about the width of the road and speed limits in each section. By providing the route, the organization simplified the global planning process, allowing teams to focus on road and obstacle detection, as well as high-speed vehicle control, as detailed in the article on Stanley: "The race was primarily a test of high-speed road finding, obstacle detection, and avoidance in desert terrain" [19].

Each team designed its own software architecture. Although some teams opted for minimalist solutions (such as Golem [20], which ran all algorithms on a single laptop), most implemented distributed systems in which multiple computers ran different software modules in parallel. These parallel architectures required careful data management, leading to the development of sophisticated communication techniques, such as converting raw data into encoded messages and precisely synchronizing data using timestamps. In addition, teams paid special attention to the modularity and extensibility of their architectures. For example, the CIMAR team used the JAUS architecture [21], while the MITRE Meteor team adopted an agent-based approach [22]. These designs facilitated software development and modification. At the subsystem level, most of the participating autonomous vehicles shared a similar structure. They typically incorporated modules for (i) localization, combining differential GPS, IMU, and estimation (sensor fusion) algorithms to determine the vehicle's position. (ii) Perception modules were used to detect roads, obstacles, and other environmental elements. (iii) Planning modules generated safe and efficient trajectories, while (iv) control modules executed these planned maneuvers. Teams extensively utilized data logs and simulators to develop and refine their algorithms.

The natural continuation of the two DARPA Grand Challenges came in 2007 with the DARPA Urban Challenge, which extended the original problem of navigation on desert roads and highways to an urban context where vehicles had to comply with California's traffic laws [23]. The challenge consisted of a series of autonomous driving missions (totaling approximately one hundred kilometers and to be completed within a maximum of six hours) driving alongside other vehicles, sharing the environment with a maximum of sixty cars: ten autonomously and fifty manually driven. According to the competition rules, the only traffic signs present would be stop signs (whose position was known in advance) and the presence of other dynamic obstacles other than cars was ruled out, simplifying perception tasks by not needing to recognize traffic lights or signs, or having to detect pedestrians or cyclists. To complete the missions, vehicles had to be capable of performing a series of basic maneuvers, such as changing lanes, merging, negotiating intersections (where they could turn left, right, or go straight) or navigating a parking lot to park the vehicle in a specific space [24]. The organization provided teams with a description of the environment in a file called the Road Network Definition File (RNDF), which contained a series of highly accurate GPS waypoints, and information about lane markings, stop signs, widths of each section, etc [25].

In terms of software architecture design, probably because the new Urban Challenge was announced just eighteen months after the last edition of the Grand Challenge, there were no major philosophical changes in their solution frameworks, but they created new modules and behaviors designed to address the new complexities, such as systems for detecting and tracking dynamic obstacles (necessary to interact with other agents in the environment) and new planning and control algorithms, since in this new edition (as mentioned above) in addition to driving following a lane, vehicles had to be able to execute a full range of basic maneuvers that would allow navigation both on roads and signposted urban streets and in other less structured areas such as parking lots that are not clearly delimited by markings and lanes [26]. This edition was won by the Carnegie Mellon University (CMU) team, with its Boss prototype [26].

Observing the software systems developed in the context of the different DARPA Challenges, it can be seen that many of the features implemented by the participating teams are very similar. In particular, the need to a) have a sensor interface to integrate data in a standardized way into system messages, b) have the ability to run numerous modules in parallel and to efficiently manage communications between them (since it is already necessary to share a large amount of information in real time in a secure and orderly manner in a system that can be distributed across a network of processors), c) have adequate user interfaces that allow interaction with the autonomous system and facilitate its configuration and debugging by allowing visualization of the data collected by the sensors and being able to consult and modify the relevant parameters, and d) have the ability to perform experimental tests of different system modules using pre-recorded datasets and simulators. These functionalities appeared very soon after in a very convenient way in the open-source operating system Robotic Operating System (ROS) [11], whose development began in 2007 at Willow Garage –a robotics incubator located in Silicon Valley– [27]. This company was created in 2006 and it is not surprising to know that its first project was precisely the participation in the DARPA Urban Challenge [28], although they did not pass the classification phase. Today, and thanks to its cooperative and open-source approach, the ROS operating system enjoys tremendous success, being used by tens of thousands of users worldwide [27] and having become a fundamental (and practically indispensable) piece in robotics research.

## 2.5 Google Self-Driving Car Project

A few years later, Sebastian Thrun and Anthony Levandowski from the Stanford Racing Team (winners of the 2005 DARPA Grand Challenge) initiated the Google Self-Driving Car Project. This initiative began in 2009 within Google X, an innovation lab at Google, with the aim of developing autonomous driving technology [29].

Early efforts focused on modifying existing vehicles, primarily Toyota Priuses, with an array of sensors including LiDAR, RADAR, and cameras. These vehicles were meticulously tested on public roads, accumulating millions of miles of autonomous driving data. In 2012, the first autonomous vehicle license in the United States was issued to a Google car in Nevada, marking a significant achievement in the field.

## 3 Autonomous Vehicles on Public Roads: Progress from 2012 to Present

In view of Google's successes with its autonomous driving program, other companies began to show interest in this area to compete in the robotaxi sector. In the following years, many actors entered the field, and numerous events occurred concurrently. For clarity, we will cite only the most relevant facts, explained below in chronological order but separately for each company.

### 3.1 Cruise LLC

Cruise was founded in 2013 as a startup dedicated to developing kits to modernize existing vehicles, equipping them with highway autopilot capabilities. In 2015, the company's focus shifted to developing software for fully autonomous vehicles. That same year, the company obtained a license to conduct autonomous driving tests in California, and in 2016, it was acquired by General Motors. After this acquisition, Cruise began an exponential growth phase, transforming from a small startup with about 40 employees into a mega company valued at over 30 billion dollars with thousands of employees by 2021 [30]. In September 2021, Cruise received a permit from the California DMV to provide driverless taxi rides without safety drivers. By February 2022, Cruise opened to the public and petitioned U.S. regulators to build self-driving vehicles without human controls. In June 2022, Cruise received California's first Driverless Deployment Permit, allowing it to charge fees for its service. The company operated 100 robotaxis in San Francisco by September 2022 and planned to expand the fleet to 5,000. Cruise also announced plans to expand to Phoenix, Arizona, and Austin, Texas, aiming for 1 billion in revenue by 2025 [31].

However, when Cruise's autonomous cars began commercial operations, safety issues started to emerge, including interference with police and fire services, collisions with other vehicles, and traffic blockages. These failures alerted authorities, but a strong response did not come until the first pedestrian accidents occurred. The most serious incident happened on October 2, 2023, when a driver struck a woman who was left lying in the lane where a Cruise driverless vehicle was traveling. The autonomous vehicle did not detect the woman, also running over her and subsequently dragging her about 6 meters before stopping, crushing the victim's leg. She had to be freed with hydraulic rescue tools [32]. This incident led to the suspension of the California Department of Motor Vehicles permit that allowed them to operate autonomously without a safety driver on board [33]. Finally, after a rapid process of decline, General Motors announced on December 10, 2024, that it would cease funding Cruise, focusing its efforts on advanced driver assistance systems (ADAS) instead of fully autonomous driving and robotaxi services [34], [35].

### 3.2 From Google to Waymo

The Google Self-Driving Car Project progressed with the development of "Firefly" a custom-built vehicle designed specifically for autonomous operation. This prototype lacked traditional driving controls like steering wheels and pedals. In 2015 this prototype achieved a significant milestone completing the world's first truly self-driving trip in Austin [36]. However, this approach was later abandoned in favor of retrofitting existing vehicle models, recognizing the importance of integrating autonomous technology into the existing automotive ecosystem. In 2016, Google restructured its self-driving car project by forming a separate company named Waymo. This move demonstrated a serious commitment to commercializing the technology beyond research and development. In 2017, Waymo introduced the fourth generation of the Waymo Driver with autonomously driven Chrysler Pacifica Hybrid minivans. By 2020, Waymo launched its first commercial application open to the general public in Chandler, Arizona, operating fully autonomously 24/7.

Currently, the Waymo Driver is in its fifth (Jaguar I-Pace, in the production phase) and sixth generations (Geely Zeekr, in the development and testing phase). It operates fully autonomously in San Francisco and Phoenix, while being developed in Los Angeles and Austin. The company is now scaling exponentially, serving one hundred thousand paid trips per week and accelerating its growth [37], [38].

### 3.3 Amazon Zoox

Zoox, founded in 2014 by Tim Kentley-Klay, an Australian artist and designer, and Jesse Levinson, who had been working on self-driving technology at Stanford University, set out with a bold vision to create a purpose-built, fully autonomous vehicle for ride-hailing services. The name "Zoox" is inspired by Zooxanthellae, a type of algae that shares a symbiotic relationship with coral, reflecting the company's philosophy of harmonious coexistence with its environment.

Zoox distinguished itself from other autonomous vehicle companies by focusing on designing a vehicle from the ground up specifically for autonomous operation, rather than retrofitting existing cars. Their vision encompassed creating a bi-directional, symmetrical electric vehicle capable of operating autonomously in complex urban environments. This unique approach allowed Zoox to optimize the vehicle's design for passenger comfort, safety, and efficient operation in cityscapes.

In 2018, Zoox achieved a significant milestone by becoming the first company to receive approval for providing public self-driving transport services in California. This permit allowed Zoox to test and demonstrate their autonomous technology in real-world conditions, a crucial step towards commercialization.

In June 2020, Zoox's innovative potential caught the attention of Amazon, which acquired the company for over 1.2 billion dollars. This acquisition provided Zoox with the resources and support needed to accelerate its development and bring its autonomous ride-hailing service to market. Under Amazon's ownership, Zoox continued to innovate, unveiling its fully autonomous vehicle in December 2020. The vehicle featured a unique design with no front or rear end, four-wheel steering, and the ability to travel in either direction, enhancing maneuverability and adaptability in urban settings.

Zoox's autonomous vehicle is equipped with advanced sensors, including LIDAR, RADAR, and cameras, providing a 360-degree field of view and enabling the vehicle to navigate complex traffic scenarios safely. The vehicle also incorporates robust safety features, such as redundant braking and steering systems, and a custom-designed airbag system that protects passengers in a variety of seating configurations.

Zoox's approach to autonomous mobility extends beyond just the vehicle; the company is also focused on creating an ecosystem that integrates seamlessly with urban infrastructure. This includes developing innovative solutions for ride-hailing, fleet management, and vehicle maintenance, ensuring a sustainable and scalable service.

Currently, Zoox is actively testing its vehicles in various cities, refining its technology, and preparing for a commercial launch, expected by the end of 2025 in Las Vegas, and planning to expand their business to other cities in the US like San Francisco, Miami, Austin, although the company has not set a time frame for those cities yet [39].

## 3.4 Baidu

Baidu, a prominent Chinese technology company, embarked on its journey into the autonomous vehicle industry around 2017. Known primarily as a search engine giant, Baidu leveraged its expertise in artificial intelligence and big data to make significant strides in the development of self-driving technology. The company's commitment to autonomous vehicles is reflected in the establishment of its Apollo platform, which serves as an open-source initiative to advance autonomous driving.

Baidu's Apollo platform, launched in April 2017, aims to provide an open, reliable, and scalable software and hardware platform for autonomous driving. By collaborating with over 200 global partners, including automakers, suppliers, and other tech companies, Apollo has become a comprehensive ecosystem for developing self-driving cars. The platform includes a range of tools and services, from data collection and simulation to high-definition mapping and vehicle control algorithms.

In 2018, Baidu received the first batch of licenses from the Chinese government to conduct open-road tests of autonomous vehicles. This milestone marked a significant step forward for the company, allowing it to gather valuable real-world data and improve its technology. The same year, Baidu's Apollo platform launched its first public robotaxi pilot program in the Chinese city of Changsha, providing autonomous ride-hailing services to residents.

Baidu continued to expand its autonomous vehicle operations, launching additional pilot programs in cities such as Beijing, Cangzhou, and Guangzhou. These programs have allowed Baidu to refine its autonomous driving technology and demonstrate its capabilities in diverse urban environments. In October 2020, Baidu made headlines by launching its fully autonomous ride-hailing service, Apollo Go, in Beijing. This service marked a significant milestone as it operated without a safety driver, showcasing Baidu's advancements in autonomous driving technology.

In addition to its domestic efforts, Baidu has pursued international collaborations to enhance its autonomous driving technology. The company partnered with Ford Motor Company in 2018 to test self-driving vehicles on Chinese roads, combining Baidu's AI and autonomous driving expertise with Ford's automotive knowledge. Baidu has also collaborated with Intel to develop computing solutions for autonomous vehicles, leveraging Intel's hardware capabilities to enhance the performance of Baidu's AI algorithms.

Baidu's autonomous vehicle research and development are supported by a dedicated team of engineers and researchers working on various aspects of self-driving technology. The company has invested heavily in areas such as computer vision, sensor fusion, and machine learning to improve the safety and reliability of its autonomous systems. Baidu's commitment to innovation is further evidenced by its involvement in the development of V2X (Vehicle-to-Everything) communication technology, which enables vehicles to interact with their surroundings, including other vehicles, infrastructure, and pedestrians.

As of now, Baidu's autonomous driving efforts continue to gain momentum, with the company aiming to deploy its Apollo Go service in more cities and scale its operations. Baidu's vision is to create a safer, more efficient transportation system through the widespread adoption of autonomous vehicles, ultimately transforming urban mobility in China and beyond.

## 3.5 The Challenges of Autonomous Car Projects: Uber, Cruise, and Apple

The journey towards autonomous vehicles has been challenging for many companies, including Uber, Cruise, and Apple. Despite initial optimism and significant investments, these companies have faced numerous issues that led to the cancellation of their autonomous car projects.

**Accidents and Safety Issues:** Safety has been one of the biggest hurdles. Autonomous vehicles rely on complex systems of sensors, cameras, radar, and artificial intelligence to navigate roads. However, these systems are not perfect. Accidents involving autonomous cars have raised serious questions about their reliability and safety. For example, Uber's self-driving car was involved in a fatal accident in 2018, leading to a temporary halt in road tests and increased scrutiny from regulators. Similarly, Cruise faced controversy when one of its robotaxis struck and dragged a pedestrian in San Francisco in 2023, resulting in a temporary shutdown and suspension of its license.

**Economic Costs:** The high costs of developing autonomous vehicles have also been a major factor in the cancellation of these projects. Companies like Uber, Cruise, and Apple have spent billions of dollars on research and development, only to find that the path to profitability is longer and more expensive than anticipated. Uber's Advanced Technologies Group (ATG) spent hundreds of millions of dollars on its self-driving car project, which ultimately led to its sale to Aurora. Cruise, backed by General Motors, invested over $10 billion into its autonomous vehicle venture, but the company decided to stop funding the project due to the high costs and competitive market.

**Strategic Recalculations:** In addition to safety and economic challenges, strategic recalculations have also played a role in the cancellation of these projects. Companies have had to reassess their priorities and focus on areas where they can achieve quicker returns on investment. For example, Apple shifted its focus to developing generative AI after canceling its "Project Titan" electric car project, which had been in development for over a decade. Similarly, Cruise decided to focus on developing advanced driver-assistance systems rather than continuing with its robotaxi business.

## 3.6 Analysis of Current Technological Trends in Autonomous Cars

As private companies, the principal actors in this market are cautious about sharing the details of their technology. However, they must also provide some insight into their developments to attract potential customers, partners, and investors. Consequently, useful information can often be found on their websites and YouTube channels [38], [40], [41]. Notably, Waymo has a dedicated "Research" section on its webpage, compiling its publications and making them freely available for download. These papers are typically published in top robotics and computer vision conferences such as IROS, ICRA, CVPR, and ECCV. Additionally, some companies make open-source materials and datasets available to the public. For instance, the Apollo Auto repository [42] offers an open-source full navigation stack designed to advance research in autonomous driving. Although it is quite a complete framework for autonomous driving with an interesting set of tools, it is likely that the company is not making public its most advanced developments. Furthermore, both Waymo [43] and Baidu [44] have published annotated datasets intended to provide rigorous benchmarks for new algorithms.

Attending to the available information, several trends can be identified.

**Platforms:** At the platform level, the predominant approach has traditionally been to modify commercially available vehicles by equipping them with the necessary sensors, actuators, and processors for autonomous driving. Examples of this "classic" paradigm include Stanley (the robot that won the 2005 DARPA Grand Challenge), which was a Volkswagen Tuareg; the first generations of the Waymo Driver, which utilized Chrysler Pacifica and Lexus RC450 models, among others [45]; and the General Motors models employed by Cruise, such as the Chevrolet Bolt. Departing from this method, Waymo developed the Firefly prototype from scratch to conceptualize what future self-driving vehicles might look like. This prototype, which does not have a steering wheel or pedals, was intended as an experimental platform rather than for mass production, leading subsequent generations of the Waymo Driver to continue being implemented in commercially available vehicles. However, the trend towards producing vehicles specifically designed for autonomous operation has gained traction in recent years. Notable examples of this approach include Zoox's robotaxi, which is designed without manual controls and prioritizes the passenger experience with two rows of seats facing each other. Similarly, the sixth generation of the Apollo Go system, featuring the Apollo RT6 vehicle built by Baidu, has a removable steering wheel to provide more space for passengers.

**Hardware:** Currently, companies are striving to reduce costs without compromising safety. Consequently, there's a trend among autonomous vehicle manufacturers to produce their hardware equipment, such as sensors and processors, avoiding dependency on external suppliers [46], [47]. The consensus in the industry is to employ a combination of LiDAR, cameras, and RADAR as exteroceptive sensors: LiDAR offers precision in distance estimation; cameras provide the necessary data for semantic understanding of the environment; and RADAR sensors supply speed information, which is crucial (and robust enough) in adverse weather conditions. While the number of these sensors can vary (with attempts to reduce them to lower costs), safety remains the priority. It is essential to have 360-degree coverage around the vehicle and overlapping fields of vision among the sensors. For instance, the sixth generation of the Waymo Driver is equipped with 13 cameras, 4 LiDARs, and 6 radars, whereas the Baidu RT6 features 12 cameras, 8 LiDARs, and 6 radars, and Zoox's Robotaxi is equipped with 14 cameras, 8 LiDARs, and 10 radars [48]. Additionally, microphones and ultrasonic sensors are also employed at the exteroceptive level.

When it comes to computing hardware, companies tend to be reticent about sharing specific details. Zoox describes their solution as powerful computer systems with redundant backups [48]. Waymo mentions their use of cutting-edge CPU and GPU servers [49]. Baidu, on the other hand, highlights its computing power, boasting approximately 1200 Tera Operations Per Second (TOPS) [47].

**Software:** Similar to hardware, companies are reluctant to disclose the details of their navigation stacks. However, through CEO presentations and available materials from publications (especially the Waymo ones) and websites, several key trends can be identified:

In terms of localization, companies avoid using satellite positioning systems (GNSS) and instead create and continuously update highly precise and semantically rich maps of their intended operational areas [50]. It's important to note that the current systems of these companies are classified at SAE level L4 autonomy [51]. As geofenced systems, they can only operate within pre-mapped areas.

For tasks such as environmental perception, agent trajectory prediction, and the generation of trajectories for ego vehicles, there is a growing trend toward using data-driven approaches with Generative Artificial Intelligence. This marks a departure from traditional modular paradigms that separate the pipeline into blocks (perception, prediction, planning, and control). While the predominant trend has not yet shifted to end-to-end driving based on foundational AI models [52], we are currently at an intermediate stage. Encoder/decoder (AI) architectures are employed, where the encoder handles perception tasks [53], and the decoder generates path planning, trajectory predictions, and temporal evolution of the environment [38], [54].

## 4 Last-mile delivery robots

The logistics chain can broadly be categorized into two primary transportation phases: (a) long-range transportation, which involves the movement of packages between warehouses; and (b) short-range operations, commonly referred to as last-mile logistics, where goods are transported between a warehouse and the end customer. While long-range transportation primarily focuses on optimizing system efficiency and enhancing global coordination techniques, last-mile operations often necessitate precise maneuvering, interaction with individuals, and navigation of semi-structured terrains.

In the context of last-mile logistics, urban growth and increasing population density are exerting significant pressure on city logistics systems. Rapid urbanization, coupled with the proliferation and swift adoption of e-commerce, has substantially increased the volume of goods requiring delivery within metropolitan areas. This trend is expected to intensify in the coming years, driven by the increasing frequency and fragmentation of deliveries [55], [56]. Current logistics practices in urban centers contribute to a range of adverse effects, including heightened traffic congestion, increased air and noise pollution, and safety concerns for pedestrians, cyclists, and delivery personnel.

To address these challenges, researchers and institutions advocate for a more integrated logistics framework. Such a system would emphasize coordination among shippers, carriers, and freight movements, consolidating goods from different customers and carriers into shared *green* vehicles. Additionally, it aims to discourage business models reliant on exploitative, low-cost labor (e.g., e-riders) [57]. This approach seeks to enhance sustainability, improve efficiency, and mitigate the negative impacts of last-mile delivery on urban environments.

From a technical perspective, autonomous last-mile delivery robots share several similarities with autonomous passenger transport vehicles, but they also have some significant differences. Both delivery robots and autonomous vehicles operate in urban environments, where they encounter dynamic conditions, changing lighting, and inclement weather. They must be able to locate themselves, detect and predict the movements of other agents, plan trajectories, adapt dynamically to avoid collisions, and control their platforms smoothly and precisely in real-time.

However, there are notable differences. Delivery robots are usually much smaller and lighter than self-driving cars and are designed to navigate at low speeds primarily in pedestrian environments. This implies that they are not generally safety-critical applications, reducing the stringent demands placed on autonomous passenger transport. In pedestrian areas, the perception requirements for obstacle detection distance are much lower (on the order of meters or tens of meters) compared to self-driving vehicles, which require precise detections at distances greater than one hundred meters. Lastly, to be cost-effective, delivery robots typically have more basic sensors (and fewer of them) and less processing capacity than self-driving cars.

It is important to note that some delivery robots fall between small delivery robots that operate in pedestrian environments and autonomous passenger transport vehicles. These medium-sized robots have sufficient speed to navigate city streets alongside cars, trucks, and other road traffic, while also being capable of operating on sidewalks and in pedestrian areas.

In the following sections, we will examine some of the most notable delivery robots currently in commercial use. However, much like the self-driving car industry, companies developing these robots are discreet about their technology. Consequently, the available information and technical details are limited, often confined to what is shared in videos, websites, social media, news articles, and similar sources.

## 4.1 Kiwi Campus

Founded in 2016 by Felipe Chávez, Jason Oviedo, and Sergio Pachón at the University of California, Berkeley, Kiwi Campus emerged from a desire to revolutionize food delivery on college campuses. Kiwi Campus has experienced remarkable growth over the past decade. What began as a small startup at UC Berkeley has expanded to Australia, Colombia, Taiwan, and the United States. The company has completed over 150,000 orders and built more than 400 robots [58]. Recently, Kiwi Campus raised \$10 million in funding to support its scaling efforts [59].

Their flagship robot, the Kiwibot, is a compact vehicle approximately 60 cm long and weighing around 17 Kg, traveling at a top speed of 3 km/h, making it ideal for pedestrian environments. Technologically, the Kiwibot has advanced through levels of autonomy. Initially fully teleoperated, it progressed to autonomous navigation on sidewalks with teleoperator assistance for crossing streets, and it now has reached level 4 autonomy, according to the company [60]. However, teleoperators are still required under certain conditions such as maintaining connectivity, GPS signal, and having an adequate map for safe navigation. This incremental functionality, supported by teleoperators, has sparked controversy. Critics argue that advertising the robots as autonomous while requiring frequent human intervention is misleading.

Additionally, teleoperators in Colombia are paid significantly less than those in the United States, raising ethical concerns [61]–[63].

The Kiwibot navigation system relies on precise positioning through GNSS using RTK (Real Time Kinematics) techniques [64]. This differential positioning requires a base station with a known position, allowing the rovers to calculate their location with centimeter-level precision when phase ambiguity is resolved [65]. The RTK system also requires a communication channel between the base and the rover. Given the typical operating environment of university campuses, which usually cover a few kilometers, RTK is a suitable choice, offering accurate localization without needing exhaustive maps. However, satellite positioning systems can face challenges such as signal blockage by tall buildings, structures like bridges or tunnels, or tree canopies, making localization difficult or even impossible.

This RTK-based approach reduces the need for local perception, saving on sensors and processing hardware. Specifically, Kiwibots use RGB and RGBD cameras and 2D LiDAR [66] as exteroceptive sensors to detect objects and agents in their surroundings. These sensors help them navigate without leaving the sidewalk, avoid obstacles, and interpret traffic signs. According to their blog [66], the Mask R-CNN algorithm [67] is used for segmentation, although the exact role of each sensor in sensory fusion is not detailed.

In general operation, users place food orders through an app from restaurants affiliated with the Kiwibot program. A Kiwi Campus employee collects the orders and loads them into the robots. The robots then navigate the campus following predefined paths, which can be modified to avoid obstacles reaching any point on campus. This approach is similar to the system presented in [68], albeit without LiDAR-based SLAM. If an issue arises during navigation, the robots can be remotely teleoperated. This operational model requires company employees on campus to manage order collection and robot loading, maintenance personnel for infrastructure checks, and a team of qualified remote personnel to take control when required.

## 4.2 Starship

Starship Technologies is a pioneer in the development of autonomous delivery robots, founded in 2014 by Janus Friis and Ahti Heinla, who are also known for their role as co-founders of Skype. The company is headquartered in San Francisco, California, with engineering hubs in Tallinn, Estonia, and Helsinki, Finland. Their main purpose is the creation of small, self-driving robots designed to navigate urban environments and deliver packages, groceries, and food efficiently. These robots are intended to improve convenience while reducing environmental impact.

The company officially launched pilot programs in 2016 in several countries, including the United States and the United Kingdom, putting their robots to the test under real-world conditions. By 2017, Starship entered commercial operations, taking a significant step in making autonomous delivery services available to the public. In 2018, it introduced delivery services in Milton Keynes, UK, collaborating with major retailers such as Co-op and Tesco. This laid the foundation for broader adoption. One of the company's most notable expansions occurred in 2019 when it partnered with Sodexo to start with food delivery at George Mason University in Virginia, US. This initiative marked the largest deployment of autonomous robot food delivery on a university campus at the time.

Currently, Starship Technologies is operating in multiple cities in the US, UK, Germany, Finland, and Estonia, raising a total of 230 million dollars since its creation in 2014. Now, after more than 7 million deliveries, this company has been consolidated as one of the principal actors in the autonomous last-mile logistics sector.

Concerning the robot [69], it is a fully electric, small six-wheeled skid-steer platform able to navigate sidewalks in all weather conditions, make driveway crossings, and climb curbs when required. Moreover, it provides a relatively large capacity (up to three shopping bags) while being light enough (25 Kg at a maximum speed of 6 Km/h) not to pose a danger to pedestrians, and other agents present in urban scenarios. At its current development stage, Starship offers Level 4 autonomy (geofenced navigation and a backup teleoperation system for manually solving unexpected events). All of the robot's hardware is built in-house by the company and its design is focused on achieving the required functionality while keeping the costs low. It has a sensor suite consisting of radars, ultrasonic sensors, and twelve cameras (including time-of-flight cameras). They claim they process this sensory data with neural networks to identify dynamic agents and navigate while avoiding collisions, although there are not many technical details available. Regarding localization, they use multilayered maps: some of them are automatically built by the robots using proprietary visual SLAM techniques, and others are manually built and labeled by humans in a mapping phase that is done previously to any autonomous operation. It is stated that GPS is also used, but its role is not clearly explained in none of the consulted sources. Finally, the robots are integrated into a whole system that includes a mobile phone application for the customers to use, a fleet management system that optimizes routes and job assignments (which robot does each delivery), and a teleoperation team that takes control of the robots when required.

## 4.3 Scout

Amazon Scout was a last-mile delivery project by Amazon. It was active from 2017, when Amazon acquired the startup Dispatch to use their technology and expertise to create the Scout robot, to 2023 when the project was canceled (among other high-risk high-gain projects) due to strategic changes in Amazon's investments [70]. Regarding the robot itself, there are not many technical details published, but we can say that the robot was a 6-wheeled fully-electric platform quite similar in concept to the Starship robot. It was intended for sidewalk navigation, at pedestrian speeds and used computer vision and neural networks for localization and obstacle avoidance, discarding the use of GNSS sensors due to their lack of reliability [71].

## 4.4 Serve

Serve Robotics is a company dedicated to autonomous delivery solutions, aiming to improve urban logistics with sustainable technology. Established in 2017 as the robotics division of Postmates, it develops zero-emission robots designed to navigate public spaces safely. In 2020, following Uber's acquisition of Postmates, Serve Robotics became an independent company. In January 2022, the company launched Level 4 self-driving robots, capable of operating autonomously within designated areas. Regarding funding, Serve Robotics has raised a total of over $56 million in funding as of August 2023. This includes a $30 million financing round announced in 2023, which was supported by strategic investors such as Uber and NVIDIA [72].

The robots are four-wheeled, fully electric, and skid steered. From the perception side, these robots equip an Ouster LiDAR sensor placed on the top rear of the vehicle, cameras at the front, and several ultrasonic sensors. They seem to avoid the use of GNSS favoring instead the use of LiDAR-based localization and mapping. From the computing side, the robots use the NVIDIA Jetson platform, and the HRI is supported by a touch screen placed on top of the platform, just below the LiDAR [73]. At the system level, they have a mobile phone application integrated with Uber Eats [74] where customers can make orders and follow their evolution, also there are supervision centers where robots can be supervised/teleoperated based on the visual feedback the robot sends to the station in real time. To complete the system, there are fleet management algorithms and a number of robots that can be sent to do the actual delivery, navigating through sidewalks.

## 4.5 Camello

The Camello Autonomous Delivery Robot, developed by the Singapore-based company OTSAW, is an effort to address last-mile delivery challenges. This robot is tailored to deliver a variety of goods, including groceries, food, and packages, efficiently and safely in urban environments. Its design incorporates advanced technologies such as artificial intelligence (AI) and sensors, enabling it to navigate through outdoor spaces and interact seamlessly with its surroundings.

OTSAW also offers the Camello robot through a “Robot as a Service” (RaaS) model, which lowers the barrier to entry for businesses by providing access to the technology on a subscription basis. This model includes setup, training, maintenance, and software updates, ensuring that the robot remains at peak performance.

## 4.6 Roxo

The Roxo project, developed by FedEx, was an ambitious initiative aimed at revolutionizing last-mile delivery through the use of autonomous robots. Known as the

FedEx SameDay Bot, Roxo was designed to operate on sidewalks, bike lanes, and roadsides within a three-to-five-mile radius of a retailer's location. Its primary purpose was to provide efficient, reliable, and environmentally friendly delivery services directly to customers' doors.

Roxo was equipped with advanced technology, including an array of sensors for 360-degree awareness, cameras for navigation, and a LiDAR system for mapping its surroundings. It also featured a stabilization system based on the iBot wheelchair base, developed by DEKA Research & Development Corp., which allowed it to traverse curbs, steps, and uneven terrain. The robot was battery-powered, zero-emission, and capable of carrying up to 100 pounds of cargo.

Despite its innovative design and potential, the Roxo project faced challenges in meeting FedEx's near-term value requirements. In 2022, FedEx decided to discontinue the project, citing a shift in corporate strategy. While the project was ultimately shelved, it provided valuable insights into robotic technology and autonomous delivery systems, contributing to the broader development of last-mile delivery solutions.

## 4.7 Clevon 1

Clevon, an autonomous delivery innovator, has achieved significant milestones since its establishment. Originally a spin-off from Cleveron, a leader in robotic click-and-collect solutions, Clevon became an independent entity in 2022. That same year, the company was listed on the Nasdaq First North Baltic exchange, marking a key financial milestone and solidifying its position in the autonomous delivery market.

In May 2022, Clevon introduced the third-generation Clevon 1 autonomous robot carrier. This all-electric vehicle features advanced capabilities, including real-time mapping, telesupervision, and autonomous driving. These innovations allow the Clevon 1 to quickly adapt to new service areas while maintaining efficient delivery operations. The vehicle's unique combination of teleoperation oversight, cameras, and radar ensures safe navigation in urban environments. Clevon expanded its global footprint in 2022 by entering the North American market. The company established its U.S. headquarters at the AllianceTexas Mobility Innovation Zone (MIZ) in the Dallas-Fort Worth Metroplex. As part of this expansion, Clevon conducted its first North American autonomous delivery, successfully delivering over 100 meals in under two hours in partnership with the city of Arlington, Texas.

The company has also formed strategic partnerships to enhance its operations. In 2022, Clevon collaborated with DHL Express Estonia, Colruyt Group in Belgium, and IKI stores in Lithuania, deploying the Clevon 1 in diverse environments. These partnerships demonstrate the vehicle's versatility and reliability in real-world applications.

## 4.8 Nuro

Nuro, Inc., founded in 2016 by Jiajun Zhu and Dave Ferguson, is a leading robotics company specializing in autonomous delivery vehicles. Both founders were former engineers at Google's self-driving car project, Waymo, bringing extensive expertise in autonomous technology to their venture. Headquartered in Mountain View, California, Nuro focuses on creating small, electric, self-driving vehicles designed specifically for last-mile delivery.

The company's journey began with its first product, the R1, which was unveiled in 2018. This compact, cargo-only vehicle was designed to transport goods such as groceries and prescriptions efficiently and safely. Nuro's innovative approach quickly attracted attention, and in 2018, the company secured $92 million in funding from Greylock Partners and Gaorong Capital. This marked the beginning of its rapid growth. In 2019, Nuro achieved a major milestone by raising $940 million in funding from the SoftBank Vision Fund, which valued the company at $2.7 billion. This significant investment enabled Nuro to expand its fleet, grow its team, and establish partnerships with major companies. Notable collaborations included partnerships with Kroger for grocery delivery and Domino's Pizza for autonomous pizza delivery. In 2020, the company raised an additional $500 million in a Series C funding round led by T. Rowe Price, bringing its valuation to $5 billion. That same year, Nuro acquired Ike Robotics, a self-driving trucking startup, further enhancing its technological capabilities. In 2021, Nuro announced plans to invest $40 million in a manufacturing facility and test track in Nevada, solidifying its commitment to scaling production.

The company has also achieved regulatory milestones. Nuro became the first autonomous vehicle company to receive an exemption from the National Highway Traffic Safety Administration (NHTSA), allowing it to operate vehicles without traditional controls like steering wheels and pedals. This exemption was a critical step in advancing the deployment of autonomous delivery vehicles. Despite its successes, Nuro has faced challenges, including workforce reductions in 2022 and 2023 as part of strategic adjustments. Nevertheless, the company remains a key player in the autonomous delivery sector, with a focus on leveraging technology to transform last-mile logistics.

## 4.9 Analysis of Current Technological Trends in Autonomous Delivery Robots

The autonomous delivery vehicle sector shares many similarities with the self-driving car industry. Many of the leading companies in this field originated as small university-affiliated startups, which were later acquired or supported by major corporations. For instance, Dispatch was acquired by Amazon to develop Scout, and Serve began as the robotics division of Postmates, a small delivery startup later purchased by Uber. Other companies, such as Starship, have pursued different growth strategies;

starting as a small startup, Starship expanded by partnering with supermarket chains like Co-op and Tesco in the UK, as well as food delivery services such as Foodora, which operates in several European countries including Germany, Austria, Sweden, and Finland. Overall, companies in this sector have attracted substantial investments over the past decade (ranging from tens to hundreds of millions of dollars) though these figures pale in comparison to the multi-billion-dollar investments seen in the self-driving car sector.

Operationally, both sectors exhibit notable similarities. Smartphone applications are predominantly chosen as the primary interface for users, while fleets of robots are managed by centralized systems designed to assign tasks efficiently, often minimizing client wait times. Teleoperation is also widely employed to address issues that robots are unable to resolve autonomously, though its usage appears more prevalent and frequent in delivery robots.

However, there are significant differences between the two sectors. The navigation requirements for delivery robots are far less stringent in terms of safety, given their lightweight and low-speed design, which makes them considerably less hazardous compared to passenger-carrying autonomous vehicles. Consequently, delivery robots are generally optimized for sidewalk navigation and crossing streets, a critical aspect of their operation, as they are at risk of being struck by road vehicles. It is not uncommon for remote operators to supervise or even manually perform such crossings via teleoperation.

These characteristics (e.g., lighter weight, lower speeds, less valuable cargo, or shorter travel distances) enable delivery robots to reduce both sensor costs and computational demands. This reduction can be advantageous during business scaling, as it lowers the cost per unit, particularly when compared to the robotaxi sector. Nonetheless, while more economical sensor and processing alternatives are generally pursued, sensor suites typically ensure 360-degree coverage and are multimodal, avoiding reliance on a single sensor type. Configurations often include a combination of cameras (RGB and RGBD), ultrasonic sensors, RADAR, LiDAR, and GNSS sensors, with specific technologies and the number of sensors varying by robot. Unlike self-driving cars, some delivery robots forego 3D LiDAR or RADAR altogether. For example, Kiwibots utilize a basic 2D LiDAR sensor, whereas Starship models employ radar and ultrasonic sensors while avoiding LiDAR. Those equipped with 3D LiDAR generally use low-resolution, cost-effective models such as the Velodyne VLP-16 (used in OTSAW's Camello robot). Recent advances in LiDAR technology and efforts to reduce costs—for instance, the merger of Ouster and Velodyne—have facilitated the development of competitive products like the Ouster Rev7. This sensor, recently incorporated into the third-generation Serve Robotics robot, offers 128 channels at a price point that was unachievable just a few years ago for sensors with more than 16 layers.

Localization strategies also vary among delivery robots. While Kiwibots and Clevon robots rely on RTK-based GNSS technology, other models favor mapping and localization techniques based on local sensors, such as LiDAR and cameras.

Broadly, both sectors (self-driving cars and autonomous delivery robots) face similar challenges in achieving economic viability despite differences in scale and safety requirements. Companies are employing various strategies to reduce costs and increase profitability, including:

- Reducing the number of sensors and processing hardware without compromising application success.
- Developing hardware in-house to minimize dependency on external suppliers.
- Partnering with technology companies for essential hardware (e.g., Serve Robotics with Ouster and NVIDIA).
- Limiting system autonomy, for instance, by increasing reliance on teleoperation (as seen with Kiwibots).
- Restricting operational areas, such as confining services to university campuses.
- Prioritizing GNSS sensors to reduce exhaustive mapping requirements for operational areas.
- Developing more efficient mapping techniques to expand businesses into new regions more rapidly and cost-effectively.

It is important to highlight that the most advanced systems in both self-driving cars and autonomous delivery robots currently operate at Level 4 automation, meaning they can only function in pre-mapped areas.

In conclusion, while technology has advanced significantly over the past decade, many challenges remain before these systems can deliver the societal impact and economic benefits that companies anticipate. Achieving greater scalability will be essential, whether through the development of Level 5 automation systems (a still distant goal) or by enhancing Level 4 systems to facilitate geographic expansion. Beyond technological hurdles, coordinated efforts from other stakeholders—such as public institutions, regulatory bodies, insurers, ethicists, and urban planners—will play a critical role in integrating these technologies into daily life.

## 5 Our autonomous last-mile delivery system: the Ona prototype

In this section, we present our last-mile autonomous delivery system, with a particular focus on our prototype robot, Ona [75]. To contextualize its development, we compare Ona with other prominent autonomous delivery robots discussed in the preceding section. These include: the Kiwibot, a four-wheeled device optimized for delivering small items; the Starship platform, designed for operation on sidewalks; Amazon's Scout, a compact six-wheeled platform; Postmates' Serve, which features increased ground clearance for carrying small items; OTSAW's Camello, developed for sidewalk navigation; FedEx's Roxo, Clevon 1, a lightweight and adaptable multi-purpose platform; and Nuro, a high-capacity autonomous robot designed to operate on streets or wide pedestrian areas.

Fig. 1: Examples of last-mile delivery robots under development

### 5.1 Comparison with other robots on the state of the art

The robots discussed share similarities because they are designed to solve comparable challenges. This means their software includes modules for tasks like perception, localization, mapping, planning (for both tasks and movements), motor control, and human-robot interaction (HRI). However, our robot Ona is more similar to Clevon 1 and Nuro in terms of its features and design. This is shown in Figure 1, where the platforms are arranged by size, from smallest (left) to largest (right).

Clevon 1 has been developed as a flexible solution for diverse business applications, particularly in goods delivery. Similar to Ona (and differing from Nuro) Clevon 1 positions its autonomy payload (e.g., sensors) on the lower section of the platform. However, Ona is engineered explicitly for navigation in pedestrian areas, incorporating features such as highly visible signaling lights and a large frontal screen to enhance human-robot interaction. Compared to Nuro, Ona is slightly smaller in size. While both Clevon 1 and Nuro are primarily tailored for street navigation, Ona is designed to operate efficiently in both street and pedestrian environments. This adaptability is supported by a novel navigation stack and a unique storage management system.

In addition to its technical features, Table 1 provides a comparative analysis of the capabilities of these robots in addressing last-mile delivery challenges [75]. This includes factors such as storage capacity, types of navigable areas, architectural barriers (e.g., steps or stairs), HRI mechanisms, and motion models. Although Ona shares similarities with Nuro, it introduces two distinctive features: the capacity for navigation in pedestrian zones and a hybrid motion train incorporating an Ackerman steering model (*Ack.*) and a differential drive (*D.D.*), enhancing its maneuverability. These advancements position Ona as a versatile and effective solution for last-mile delivery.

### 5.2 Ona’s main features

Ona is a mobile robotic platform designed for last-mile delivery, featuring a six-wheel drive system with steering and an Ackerman drive. The platform weighs approximately 200 kg and has dimensions of 1.8 × 1.1 × 1 meters, including the outer shell. As an all-electric vehicle, Ona offers more than five hours of continuous

| | Kiwibot | Starship | Scout | Serve | Camello | Roxo | Clevon 1 | *Ona* | Nuro |
|---|---|---|---|---|---|---|---|---|---|
| **Pkg. cap.** | ↘ | ↘ | ↘ | ↘ | → | → | ↗ | ↗ | ↗ |
| **Road nav.** | ✗ | ✗ | ✗ | ✗ | ✗ | ✗ | ✓ | ✓ | ✓ |
| **Ped. zone** | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✗ |
| **Step size** | $< 10cm$ | $< 10cm$ | $< 10cm$ | $< 10cm$ | $< 10cm$ | $< 10cm$ | $< 20cm$ | $< 20cm$ | $< 20cm$ |
| **Stairs** | ✗ | ✗ | ✗ | ✗ | ✗ | ✓ | ✗ | ✗ | ✗ |
| **HRI** | ↗ | → | → | → | ↘ | ↘ | ↘ | ↗ | ↗ |
| **Kinematics** | *Ack.* | *D.D.* | *D.D.* | *D.D.* | *D.D.* | *D.D.* | *Ack.* | *Ack.*<br>*D.D.* | *Ack.* |
| **Bi-dir. nav.** | ✗ | ✗ | ✗ | ✗ | ✓ | ✗ | ✗ | ✓ | ✓ |
| **LiDAR** | ✓ | ✗ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |

Table 1: Summary of the primary features of leading last-mile delivery robots from [75]. ‘Pkg. cap.’ denotes storage capacity. ‘Road nav.’ indicates if the robot supports street navigation, while ‘Ped. zone’ specifies suitability for pedestrian-only areas. ‘Step size’ shows the maximum obstacle height the robot can traverse, and ‘Stairs’ highlights its ability to climb consecutive steps. ‘HRI’ assesses Human-Robot Interaction interfaces, such as screens. ‘Kinematics’ details the robot’s motion model. ‘Bi-dir. nav.’ refers to whether the robot can move both forward and backward, reflecting hardware symmetry. Lastly, ‘LiDAR’ confirms if LiDAR is used in the robot’s perception systems.

operation on a single charge, which is sufficient for last-mile delivery tasks given the storage capacity detailed later in this text.

The platform employs a combination of four front traction-only wheels and two rear wheels that provide both traction and steering. This configuration results in some wheel skidding during turns, a characteristic inherent to the platform’s architecture. To enhance interaction with users and pedestrians, Ona is equipped with two 22-inch screens (at the front and rear) and a 7-inch lateral screen.

Ona integrates multiple sensors to ensure robust and reliable operation by leveraging various physical principles, including:

- **Wheel Encoders**.
- **Inertial Measurement Unit (IMU):** Provides angular velocity and linear acceleration estimates.
- **Global Navigation Satellite System (GNSS):** Enables accurate global positioning. Ona’s GNSS receiver can operate in Real-Time Kinematic (RTK) mode with a base station for improved accuracy but currently runs in standalone mode.

Fig. 2: Ona's storage management system with four internal compartments.

- **3D LiDARs:** Two 16-beam LiDAR sensors are installed diagonally on the robot (front-right and back-left corners). These sensors support odometry estimation and local map creation using point clouds. Further details on their integration are presented in Section 5.3.
- **Depth Cameras:** Three depth cameras, located at the front and sides, are oriented downward to capture high-resolution local terrain data. This facilitates traversability analysis, including obstacle detection and negative obstacle identification. Additional data processing details can be found in [76]. A forward-facing depth-color camera aids teleoperation, though its depth data has been excluded due to robustness issues.
- **Safety 2D LiDAR:** Used for hardware safety functions.
- **Bumpers:** Act as mechanical safety features.
- **Sonars:** Detect potential holes or obstacles near the wheels.
- **Emergency stops:** Ranging from Category 0 to Category 2.

The platform also features an innovative pickup and delivery system with:

- **Automatic Lateral Doors:** Open and close on command from the mission manager when the customer is ready.
- **Package Manipulation System:** Includes a belt mechanism for moving packages in and out, complemented by a two-axis manipulator for precise handling.
- **Storage Compartments:** Four dedicated parcels designed to securely hold packages during navigation. These elements are depicted in Figure 2.

To manage internal communications, Ona employs the Robot Operating System (ROS) middleware, along with Behavior Trees [77] to plan tasks and coordinate mission management, remote control, navigation systems, and hardware operations. Externally, Ona connects to a control center for logistics coordination and human operator interface via a "communications backpack" with parallel SIM cards supporting 4G, 5G, and global roaming across nearly 200 countries, including Europe.

Ona's integration into urban environments addresses various Human-Robot Interaction (HRI) roles, as detailed in [78]. These include supervisors handling logistics and routes, operators skilled in teleoperation and alarm management, technicians responsible for maintenance, and peers like end-users and teammates. Bystanders encompass pedestrians, vulnerable individuals, urban workers, and others coexisting with the robot. Ona interacts safely with these stakeholders using tools such as screen displays, emojis, and colored lights coded for navigation status (green for nominal, orange for warning, and red for critical). The pickup and delivery process integrates

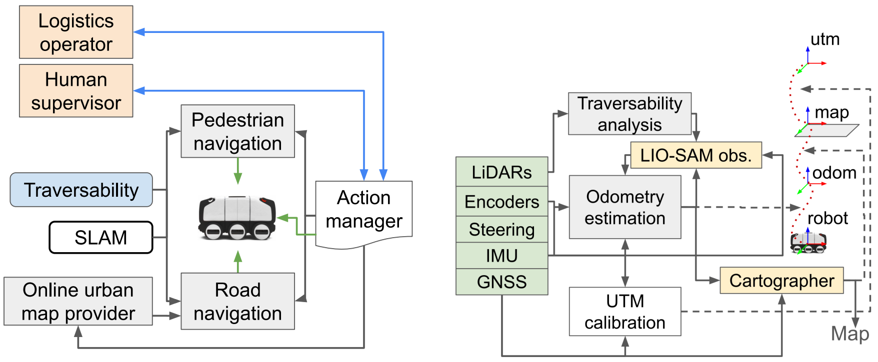

(a) Operational.

(b) Localization

Fig. 3: Main Software modules of Ona's autonomous navigation pipeline

HRI by sending authentication codes to end-users, ensuring secure and seamless interactions.

## 5.3 Ona's navigation stack

### 5.3.1 Simultaneous localization and mapping

To navigate autonomously, it is essential to estimate the robot's ego-motion and construct an accurate model of its surroundings by analyzing sensor observations [79]. These processes are managed by the SLAM module, depicted in Figure 3a and further detailed in Figure 3b.

Each sensor used in autonomous driving presents unique characteristics. A commonly employed visual sensor is the LiDAR (light detection and ranging), which provides a wide field of view but typically generates sparse data. This necessitates combining consecutive readings, using motion estimation, to represent the surrounding environment. Since LiDAR measurements are produced at a relatively low frame rate, their observations are often fused with data from an inertial measurement unit (IMU) to enhance temporal resolution.

To achieve this, Ona employs a custom sensor fusion architecture inspired by LIO-SAM [80], augmented with an additional observability module as developed in previous work [81]. The module, referred to as LIO-SAM Obs. in Figure 3a, utilizes LiDAR point clouds generated during traversability analysis (described in subsequent sections) and integrates these with IMU data. The output from this LiDAR-IMU fusion is then further refined by incorporating wheel odometry through an Extended Kalman Filter (EKF), as represented by the odometry estimation module in Figure 3a.

Although the odometry estimation provides reasonable ego-motion data, it is prone to localization drift and lacks global consistency. To address these limitations, a localization outer loop has been implemented to: **i)** achieve globally consistent localization for Ona by fusing odometry data with GNSS readings, and **ii)** create a local environmental map by combining motion estimates and laser scan data to produce a point cloud. This map is then utilized for navigation planning and control command generation. For global localization and local mapping, the system employs Cartographer [82], a real-time simultaneous localization and mapping (SLAM) framework.

To ensure the map adheres to Universal Transverse Mercator (UTM) consistency, an additional EKF-based UTM calibration module has been introduced. This module integrates GNSS readings and IMU-based heading estimations. The interconnections between the UTM calibration module, the Cartographer system, and the outer global localization loop are illustrated in Figure 3b.

### 5.3.2 Traversability analysis

Ground segmentation is essential for autonomous robots, as the quality of its output directly impacts higher-level processes such as obstacle detection, path planning, and localization. Here, we employ a probabilistic graph-based real-time ground segmentation algorithm [83] that processes instantaneous 3D point clouds to create a probabilistic model of the ground surface and its traversability in real-time. The algorithm constructs a graph, starting from a solid prior of the ground level at the origin and expanding to the limits of the point cloud. It fits small planes to the data using a probabilistic filter approach: extracting local data, rejecting outliers based on the prior, updating plane estimations with inliers, and propagating the posterior as the prior for subsequent nodes. This methodology explicitly accounts for model uncertainties during segmentation.

Additionally, during graph construction, statistical descriptors (including intensity information) are extracted to refine traversability analysis using a fully connected neural network. This network processes the descriptors to classify surfaces such as roads, sidewalks, and low vegetation areas, enhancing the segmentation accuracy and contextual understanding of the environment. For detailed insights into the algorithm and its evaluation, please refer to [83].

An example of the algorithm in action at The Barcelona Robot Lab (North Campus of UPC) is shown in Figure 4a, where ground and obstacle points are represented in blue and green, respectively, while roadmap vertices and edges are shown in brown.

The resulting traversability analysis is utilized by planning modules for both pedestrian and road navigation (see Figure 3a), and the segmented point cloud, including static objects above the ground, is fed into the LiDAR-IMU fusion algorithm (Figure 3b).

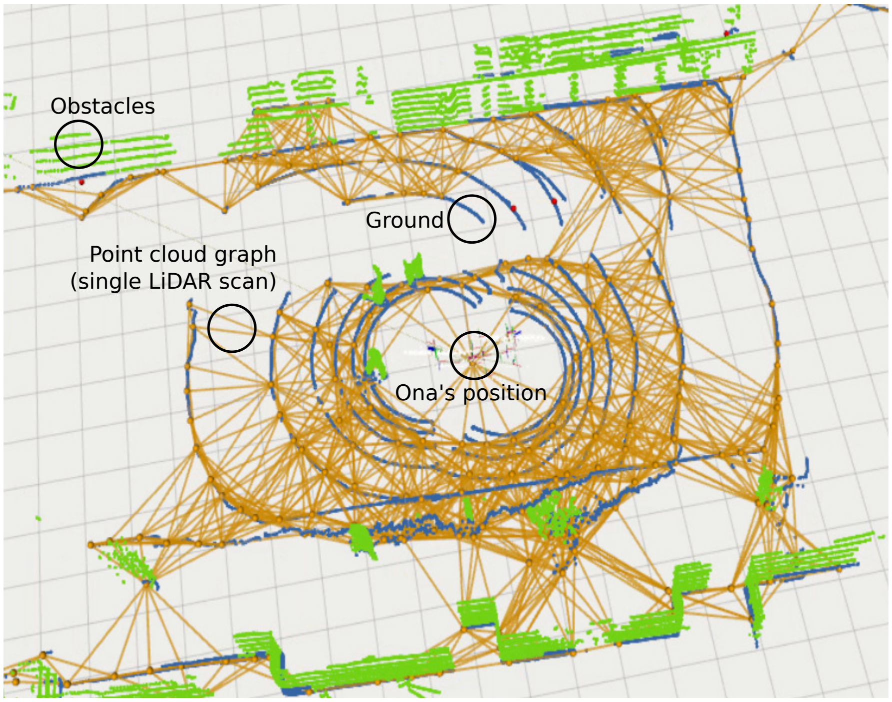

(a) Output sample of the ground segmentation algorithm used for traversability analysis, processing the merged instantaneous LiDAR readings of Ona, during an experiment in The Barcelona Robot Lab (North campus of the Universitat Politècnica de Catalunya).

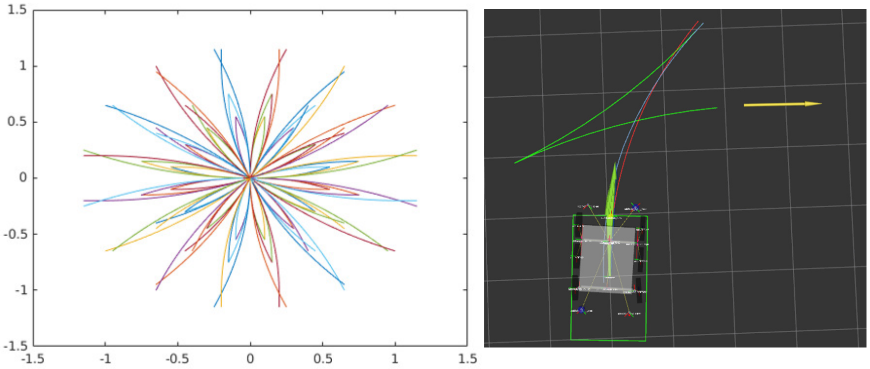

(b) Example of motion primitives (left) and a generated path (right) computed by the pedestrian area global planner (green lines: global path; blue line: current segment executed by the local planner; red line: actual motion command generated by the local planner).

Fig. 4: Perception (traversability) and planning details of Ona.

### 5.3.3 Global map handling and path planning

The mission manager provides Ona with the global UTM coordinates for pickup and delivery points. Ona computes a global path between these points using a map service and path planning algorithms [84]. With its state expressed in a global reference frame (via global localization), Ona accesses a global map provider to generate navigation trajectories (see Figure 3a).

For global path computation, Ona utilizes Open Street Map (OSM), an open-source resource. Ona maintains a local copy of the OSM map for its operational region, allowing for manual verification and corrections to ensure data accuracy—an effort required only once per area. A ROS node processes OSM data to generate routes, classifying segments into pedestrian or road paths (an existing distinction in OSM). This classification enables Ona to split the route into sequential trajectory goals and select specialized planners for road or pedestrian navigation [85], considering the structural variability of pedestrian areas and the presence of vulnerable road users (VRUs).

Once segmented into road and pedestrian paths (see Figure 3a), different planners compute the robot's global trajectory:

- **Road Navigation:** Generates the shortest path/lane following traffic rules defined in the map. The global path consists of lane-center lines and includes the transition from/to parking spots through drivable areas.
- **Pedestrian Area Navigation:** Generates paths from the robot's position to a goal using "motion primitives", short kinematically feasible motions (Figure 4b-right). This planner ensures smooth paths, considering robot orientation, and generates maneuvers when simple paths are unavailable (Figure 4b-left).

The global path is computed once per mission. To follow it, Ona employs a reactive local planner, especially adapted for VRUs, within the ROS move-base architecture[3].

Unlike the global planner, the local planner re-plans multiple times per second to adapt to dynamic situations, such as VRU presence. Using traversability analysis, Ona creates global and local cost maps as occupancy grids, enabling the local planner to identify suitable paths in each control cycle. The local planner incorporates Ona's Ackerman motion model, enhanced with custom Spline trajectories for smoother paths matching its physical capabilities. Recovery behaviors are implemented to handle unfeasible plans, including re-planning tasks. If no valid plan is found, the task manager flags assistance from an operator, leveraging Behavior Trees [77].

### 5.3.4 Remote monitoring and teleoperation

Ona supports bidirectional communication with an external control center based on ROS2 (Figure 3a). This communication system addresses key constraints:

---

[3] http://wiki.ros.org/move_base

- **Bandwidth (BW)**: Data from Ona is compressed to a maximum upload rate of $1Mbit/s$, consisting of robot state messages and the front image for visual feedback. The download rate reaches $8Mbit/s$, allowing for control commands. These rates are optimized via a ROS2 QoS network controller.
- **Latency**: To enable effective teleoperation, latency is maintained below $300ms$, adjustable through the a ROS2 throughput controller.

While 4G networks are widely available in major and medium-sized cities in western countries, the deployment of Ona requires initial verification to ensure stable coverage, as 5G signaling may not be uniformly distributed.

### 5.3.5 Logistic's mission manager

Ona can easily be integrated into any major logistics operator Software. Currently, Ona can send to the mission (business) manager any information related to its state, the state of the parcels or the delivery/pickup status. Unlike other last-mile robots, Ona features a unique automated compartment mechanism for fetching parcels and opening/closing doors. This simplifies communication between the mission manager and Ona, reducing it to essential details such as pickup/delivery coordinates (UTM) and pickup/delivery synchronization. Additionally, Ona provides live navigation data to the mission manager for tracking purposes.

## 5.4 Ona Experimental Results

The validation of the Ona platform was carried out in three distinct scenarios. The initial tests were conducted at the Barcelona Robot Lab (Universitat Politècnica de Catalunya) to prepare the platform and evaluate the preliminary sensor configurations and methodologies. Following these preparatory tests, the official presentation of Ona took place in Esplugues de Llobregat in June 2022. This event was attended by notable Spanish national authorities, including the Minister of Transportation. Given that Esplugues de Llobregat is among the most densely populated areas in Europe, the test results demonstrated highly satisfactory performance, supporting the feasibility of deploying autonomous last-mile delivery robots in such urban environments. A second pilot demonstration was conducted in Debrecen, Hungary, in November 2022. This event was attended by the city's mayor and featured a complete delivery mission. Notably, Ona navigated a route shared with tram vehicles and a significant number of pedestrians, further validating its capabilities in mixed-traffic scenarios.

Further details regarding the preparation and implementation of these validation experiments are provided in [75]. This includes information on the use of simulators and digital twins, specific examples of routes and environments, as well as findings from technological acceptance studies conducted through on-site surveys with bystanders.

## 6 Conclusions

This Chapter contains a comprehensive and up-to-date overview of autonomous navigation vehicles in urban environments, focusing on applications such as autonomous passenger transportation and last-mile delivery robotics. We begin with a historical perspective on the origins of outdoor autonomous navigation, discussing the key challenges these systems have faced and the algorithms and design decisions that have progressively overcome them. Following this historical overview (primarily centered on research from academic setting) we proceed to describe various industry-driven initiatives in detail. These initiatives, ranging from "robotaxi" services to autonomous delivery robots, have emerged since the 2010s, leveraging the scientific and technical knowledge developed in previous decades.

Building on these descriptions, we conduct an analysis to identify the differences and similarities among existing systems, which in turn allows us to outline the current challenges and solutions in the field of autonomous robotics for urban environments. In the final part of the chapter, we provide a detailed description of the "Ona" last-mile autonomous logistics system, where the authors of this chapter have played a significant role in the design, implementation and validation of its autonomous navigation stack [75]. This description serves as a practical example of the authors' perspective on the chapter's main topic, offering insights into contemporary approaches to autonomous navigation in urban settings.

## Acknowledgements

This work has been partially supported by Horizon Europe under the grant agreement No. 101168042 (TRIFFID); by the Spanish Ministry of Science and Innovation under the project LENA (PID2022-142039NA-I00, funded by MCIN/AEI/10.13039/501100011033 and by "ERDF A way of making Europe"); by the BotNet project (23S06128-00), funded by the Ajuntament de Barcelona and Fundació la Caixa; and by the Consolidated Research Group RAIG (2021 SGR 00510) of the Departament de Recerca i Universitats de la Generalitat de Catalunya.

## References

[1] B. Grush, "A regulatory roadmap for public-area mobile robots," in *2024 IEEE International Conference on Smart Mobility (SM)*, IEEE, 2024, pp. 129–134.

[2] B. Grush, *Discovery Guide to PMRs: Opportunities and Challenges with Public-area Mobile Robots*. Urban Robotics Foundation, 2018.

[3] N. J. Nilsson, "A mobile automaton: An application of artificial intelligence techniques," SRI INTERNATIONAL MENLO PARK CA ARTIFICIAL INTELLIGENCE CENTER, Tech. Rep., 1969.

[4] B. Kuipers, E. A. Feigenbaum, P. E. Hart, and N. J. Nilsson, "Shakey: From conception to history," *Ai Magazine*, vol. 38, no. 1, pp. 88–103, 2017.

[5] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

[6] R. O. Duda and P. E. Hart, "Use of the hough transformation to detect lines and curves in pictures," *Communications of the ACM*, vol. 15, no. 1, pp. 11–15, 1972.

[7] M. Hebert, "Outdoor scene analysis using range data," in *Proceedings. 1986 IEEE International Conference on Robotics and Automation*, IEEE, vol. 3, 1986, pp. 1426–1432.

[8] I. S. Kweon, M. Hebvert, and T. Kanade, "Sensor fusion of range and reflectance data for outdoor scene analysis," The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, Tech. Rep., 1988.

[9] Y. Goto and A. Stentz, "The cmu system for mobile robot navigation," in *Proceedings. 1987 IEEE International Conference on Robotics and Automation*, IEEE, vol. 4, 1987, pp. 99–105.

[10] L. Champeny-Bares, S. Coppersmith, and K. Dowling, "The terregator mobile robot," CARNEGIE-MELLON UNIV PITTSBURGH PA ROBOTICS INST, Tech. Rep., 1991.

[11] M. Quigley, K. Conley, B. Gerkey, *et al.*, "Ros: An open-source robot operating system," in *ICRA workshop on open source software*, Kobe, vol. 3, 2009, p. 5.

[12] C. Thorpe, M. H. Hebert, T. Kanade, and S. A. Shafer, "Vision and navigation for the carnegie-mellon navlab," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 10, no. 3, pp. 362–373, 1988.

[13] E. D. Dickmanns and A. Zapp, "Autonomous high speed road vehicle guidance by computer vision1," *IFAC Proceedings Volumes*, vol. 20, no. 5, pp. 221–226, 1987.

[14] E. D. Dickmanns, R. Behringer, D. Dickmanns, *et al.*, "The seeing passenger car'vamors-p'," in *Intelligent Vehicles' 94 Symposium, Proceedings of the*, IEEE, 1994, pp. 68–73.

[15] T. Jochem and D. Pomerleau, *No hands across america official*, Press Release, 1995.

[16] E. D. Dickmanns *et al.*, "Vehicles capable of dynamic vision," in *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI)*, vol. 97, 1997, pp. 1577–1592.

[17] E. Kaplan and C. Hegarty, *Understanding GPS: principles and applications*. Artech house, 2005.

[18] G. Seetharaman, A. Lakhotia, and E. P. Blasch, "Unmanned vehicles come of age: The darpa grand challenge," *Computer*, vol. 39, no. 12, pp. 26–29, 2006.

[19] S. Thrun, M. Montemerlo, H. Dahlkamp, *et al.*, "Stanley: The robot that won the darpa grand challenge," *Journal of field Robotics*, vol. 23, no. 9, pp. 661–692, 2006.

[20] R. Mason, J. Radford, D. Kumar, *et al.*, "The golem group/university of california at los angeles autonomous ground vehicle in the darpa grand challenge," *Journal of Field Robotics*, vol. 23, no. 8, pp. 527–553, 2006.

[21] S. Rowe and C. R. Wagner, "An introduction to the joint architecture for unmanned systems (jaus)," Cybernet Systems Corporation, Tech. Rep., 2008.

[22] M. Minsky, *Society of mind*. Simon and Schuster, 1988.

[23] C. Urmson *et al.*, "Self-driving cars and the urban challenge," *IEEE Intelligent Systems*, vol. 23, no. 2, pp. 66–68, 2008.

[24] B. J. Patz, Y. Papelis, R. Pillat, G. Stein, and D. Harper, "A practical approach to robotic design for the darpa urban challenge," *Journal of Field Robotics*, vol. 25, no. 8, pp. 528–566, 2008.

[25] M. Montemerlo, J. Becker, S. Bhat, *et al.*, "Junior: The stanford entry in the urban challenge," *Journal of field Robotics*, vol. 25, no. 9, pp. 569–597, 2008.

[26] C. Urmson, J. Anhalt, D. Bagnell, *et al.*, "Autonomous driving in urban environments: Boss and the urban challenge," *Journal of Field Robotics*, vol. 25, no. 8, pp. 425–466, 2008.

[27] M. Quigley, B. Gerkey, and W. D. Smart, *Programming Robots with ROS: a practical introduction to the Robot Operating System*. O' Reilly Media, Inc., 2015.

[28] S. Cousins, "Willow garage retrospective [ros topics]," *IEEE Robotics & Automation Magazine*, vol. 21, no. 1, pp. 16–20, 2014.

[29] Wikipedia. "Waymo." (2025), [Online]. Available: `https://en.wikipedia.org/wiki/Waymo` (visited on 01/24/2025).

[30] M. Wayland. "Microsoft is investing and partnering with gm's cruise on self-driving cars." (2021), [Online]. Available: `https://www.cnbc.com/2021/01/19/microsoft-invests-partners-with-gms-cruise-on-self-driving-cars.html` (visited on 01/19/2021).

[31] A. J. Hawkings. "Cruise continues to burn gm's cash as robotaxis expand to daylight hours." (2023), [Online]. Available: `https://www.theverge.com/2023/4/25/23697406/cruise-robotaxi-gm-daylight-sf-earnings` (visited on 04/25/2023).

[32] J. Parker and N. Mishanec. "Driver hits woman in s.f., then cruise driverless car runs her over." (2023), [Online]. Available: `https://www.sfchronicle.com/bayarea/article/woman-run-autonomous-vehicle-san-francisco-18403044.php` (visited on 10/03/2023).

[33] H. Field. "California dmv suspends cruise's self-driving car permits, effective immediately." (2023), [Online]. Available: `https://www.cnbc.com/2023/10/24/california-dmv-suspends-cruises-self-driving-car-permits.html` (visited on 10/24/2023).

[34] GeneralMotors. "Gm to refocus autonomous driving development on personal vehicles." (2025), [Online]. Available: `https://news.gm.com/home.detail.html/Pages/news/us/en/2024/dec/1210-gm.html` (visited on 11/12/2024).

[35] TheGuardian. "General motors pulls plug on cruise, its self-driving robotaxi company." (2024), [Online]. Available: `https://www.theguardian.com/`

us-news/2024/dec/11/general-motors-self-driving-cruise-robotaxi (visited on 12/11/2024).

[36] Waymo. "From post-it note to prototype: The journey of our firefly." (2017), [Online]. Available: https://waymo.com/blog/2017/06/from-post-it-note-to-prototype-journey (visited on 01/27/2025).

[37] Waymo. "Our mission: To be the most reliable driver in the world." (2025), [Online]. Available: https://waymo.com/intl/es/about (visited on 01/27/2025).

[38] D. Dolgov. "The waymo way: Making autonomous driving a reality." (2024), [Online]. Available: https://youtu.be/s_wGhKBjH_U (visited on 09/17/2024).

[39] CNBC. "Why 2025 is set to be a crucial year for amazon's zoox robotaxi unit." (2025), [Online]. Available: https://www.cnbc.com/2025/01/17/amazon-zoox-plans-commercial-expansion.html (visited on 01/17/2025).

[40] K. Vogt. "Cruise under the hood." (2021), [Online]. Available: https://youtu.be/uJWN0K26NxQ (visited on 11/05/2021).

[41] Baidu. "Apollo day 2022 (autonomous driving tech deep dive)." (2022), [Online]. Available: https://www.youtube.com/live/JISnVvff7Jc (visited on 11/29/2022).

[42] Baidu. "Apollo github repository." (2025), [Online]. Available: https://github.com/ApolloAuto (visited on 02/12/2025).

[43] P. Sun, H. Kretzschmar, X. Dotiwalla, *et al.*, "Scalability in perception for autonomous driving: Waymo open dataset," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 2446–2454.

[44] X. Huang, X. Cheng, Q. Geng, *et al.*, "The apolloscape dataset for autonomous driving," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2018, pp. 954–960.

[45] T. Mawakana. "Waymo's roadmap for a multi-city av service." (2024), [Online]. Available: https://youtu.be/Qot1uX2g9jk (visited on 03/25/2024).

[46] Waymo. "Meet the 6th-generation waymo driver: Optimized for costs, designed to handle more weather, and coming to riders faster than before." (2024), [Online]. Available: https://waymo.com/blog/2024/08/meet-the-6th-generation-waymo-driver (visited on 08/19/2024).

[47] Electrek. "Baidu unveils apollo rt6 level 4 autonomous robotaxi with detachable steering wheel." (2022), [Online]. Available: https://electrek.co/2022/07/21/baidu-apollo-rt6-robotaxi (visited on 07/21/2022).

[48] Zoox. "The 'full-stack' behind autonomous driving." (2025), [Online]. Available: https://zoox.com/autonomy (visited on 02/17/2025).

[49] Waymo. "Waymo driver." (2025), [Online]. Available: https://waymo.com/intl/es/waymo-driver (visited on 02/19/2025).

[50] L. Chen, P. Wu, K. Chitta, B. Jaeger, A. Geiger, and H. Li, "End-to-end autonomous driving: Challenges and frontiers," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.

[51] B. Williams, "Automated driving levels," 2021.

[52] J.-J. Hwang, R. Xu, H. Lin, *et al.*, "Emma: End-to-end multimodal model for autonomous driving," *arXiv preprint arXiv:2410.23262*, 2024.

[53] J. Li, X. Shi, F. Chen, *et al.*, "Pedestrian crossing action recognition and trajectory prediction with 3d human keypoints," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2023, pp. 1463–1470.

[54] A. Seff, B. Cera, D. Chen, *et al.*, "Motionlm: Multi-agent motion forecasting as language modeling," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 8579–8590.

[55] A. Sharma, P. Zanotti, and L. P. Musunur, "Drive through robotics: Robotic automation for last mile distribution of food and essentials during pandemics," *IEEE Access*, vol. 8, pp. 127 190–127 219, 2020.

[56] M. Viu-Roig and E. J. Alvarez-Palau, "The impact of e-commerce-related last-mile logistics on cities: A systematic literature review," *Sustainability*, vol. 12, no. 16, 2020.

[57] C. Lemardelé, M. Estrada, L. Pagès, and M. Bachofner, "Potentialities of drones and ground autonomous delivery devices for last-mile logistics," *Transportation Research Part E: Logistics and Transportation Review*, vol. 149, p. 102 325, 2021, ISSN: 1366-5545.

[58] Sodexomyway. "Kiwibot." (2025), [Online]. Available: `https://southernct.sodexomyway.com/en-us/kiwibot` (visited on 02/18/2025).

[59] roboticsandautomationmagazine. "Kiwibot." (2023), [Online]. Available: `https://www.roboticsandautomationmagazine.co.uk/news/deliveries/kiwibot-secures-us10m-to-expand-robot-food-delivery-service.html` (visited on 03/03/2023).

[60] Kiwibot. "How does a kiwibot navigate autonomously?" (2023), [Online]. Available: `https://www.linkedin.com/pulse/how-does-kiwibot-navigate-autonomously-kiwibot` (visited on 01/31/2023).

[61] T. Barrett, "Kiwibots on kampus: The ethical concerns of autonomous vehicles in black spaces," *GRACE: Global Review of AI Community Ethics*, vol. 1, no. 1, 2023.

[62] roboticsandautomationmagazine. "Kiwi campus." (2025), [Online]. Available: `https://en.wikipedia.org/wiki/Kiwi_Campus` (visited on 01/21/2025).

[63] J. Steinhoff, J. Posada, and A. Delfanti, "Remote robotics, or the digital reembodiment of labout," 2024.

[64] Kiwibot. "Understanding rtk: A closer look at kiwibot's autonomous positioning system." (2024), [Online]. Available: `https://www.linkedin.com/pulse/understanding-rtk-closer-look-kiwibots-autonomous-positioning-twdze` (visited on 09/23/2024).

[65] J. Sanz, J. Juan, and M. Hernández-Pajares, "Gnss data processing, vol. i: Fundamentals and algorithms. noordwijk, the netherlands: Esa communications," ESTEC TM-23/1, 2013.[Online]. Available: http://www. navipedia. net/GNSS . . ., Tech. Rep., 2013.

[66] Kiwibot. “Decoding the streets: Kiwibot’s journey in segmentation.” (2025), [Online]. Available: `https://www.kiwibot.com/articles/kiwibots-journey-in-segmentation` (visited on 02/18/2025).

[67] K. He, G. Gkioxari, P. Dollár, and R. Girshick, “Mask r-cnn,” in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 2961–2969.

[68] M. Á. Muñoz–Bañón, I. del Pino, F. A. Candelas, and F. Torres, “Framework for fast experimental testing of autonomous navigation algorithms,” *Applied Sciences*, vol. 9, no. 10, p. 1997, 2019.

[69] S. Engineering. “Starship technologies autonomous delivery robots.” (2020), [Online]. Available: `https://youtu.be/ztTwXYL4Qaw` (visited on 11/30/2020).

[70] J. Vincent. “Amazon stops field tests of its delivery robot scout.” (2022), [Online]. Available: `https://www.theverge.com/2022/10/7/23392360/amazon-disbands-delivery-robot-scout-development` (visited on 10/07/2022).

[71] Wikipedia. “Amazon scout.” (2024), [Online]. Available: `https://en.wikipedia.org/wiki/Amazon_Scout` (visited on 03/17/2025).

[72] serverobotics. “Serve robotics inc. announces go-public transaction and $30 million financing to support scaling of robotic delivery.” (2023), [Online]. Available: `https://www.serverobotics.com/go-public-transaction` (visited on 08/10/2023).

[73] serverobotics. “Serve robotics becomes first autonomous vehicle company to commercially launch level 4 self-driving robots.” (2022), [Online]. Available: `https://www.serverobotics.com/level-4-autonomy` (visited on 01/13/2022).

[74] serverobotics. “Serve robotics and uber expand partnership to bring robotic delivery to uber eats.” (2021), [Online]. Available: `https://www.serverobotics.com/uber-eats-partnership` (visited on 11/08/2021).

[75] A. Santamaria-Navarro, S. Hernández, F. Herrero, *et al.*, “Toward the deployment of an autonomous last-mile delivery robot in urban areas: The ona prototype platform,” *IEEE Robotics & Automation Magazine*, 2024.

[76] A. Santamaria-Navarro, E. H. T. Avilés, M. Morta, and J. Andrade-Cetto, “Terrain classification in complex three-dimensional outdoor environments,” *J. Field Robotics*, vol. 32, no. 1, pp. 42–60, 2015. DOI: `10.1002/rob.21521`.

[77] M. Colledanchise and P. Ogren, *Behavior Trees in Robotics and Al: An Introduction*, 1st. USA: CRC Press, Inc., 2018, ISBN: 1138593737.

[78] A. Puig-Pey, J. L. Zamora, B. Amante, *et al.*, “Human acceptance in the human-robot interaction scenario for last-mile goods delivery,” in *IEEE International Conference on Advanced Robotics and Its Social Impacts*, 2023, pp. 32–29.

[79] C. Cadena, L. Carlone, H. Carrillo, *et al.*, “Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age,” *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.

[80] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, “Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping,” *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5135–5142, 2020.

[81] A. Tagliabue, J. Tordesillas, X. Cai, *et al.*, “Lion: Lidar-inertial observability-aware navigator for vision-denied environments,” in *Experimental Robotics*, B. Siciliano, C. Laschi, and O. Khatib, Eds., Cham: Springer International Publishing, 2021, pp. 380–390, ISBN: 978-3-030-71151-1.

[82] W. Hess, D. Kohler, H. Rapp, and D. Andor, “Real-time loop closure in 2d lidar slam,” *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1271–1278, 2016.

[83] I. d. Pino, A. Santamaria-Navarro, A. G. Zulueta, F. Torres, and J. Andrade-Cetto, “Probabilistic graph-based real-time ground segmentation for urban robotics,” *IEEE Transactions on Intelligent Vehicles*, pp. 1–14, 2024. DOI: `10.1109/TIV.2024.3383599`.

[84] R. L. Guimarães, A. S. de Oliveira, J. A. Fabro, T. Becker, and V. A. Brenner, “Ros navigation: Concepts and tutorial,” in *Robot Operating System (ROS): The Complete Reference (Volume 1)*, A. Koubaa, Ed. Cham: Springer International Publishing, 2016, pp. 121–160, ISBN: 978-3-319-26054-9.

[85] G. Ferrer and A. Sanfeliu, “Bayesian human motion intentionality prediction in urban environments,” *Pattern Recognition Letters*, vol. 44, pp. 134–140, 2014, Pattern Recognition and Crowd Analysis, ISSN: 0167-8655.