Individual Thesis

"Artificial Intelligence in Autonomous Vehicle Navigation"

Antreas Nathanail

UNIVERSITY OF CYPRUS

DEPARTMENT OF COMPUTER SCIENCE

May 2025

UNIVERSITY OF CYPRUS

DEPARTMENT OF COMPUTER SCIENCE

“Artificial Intelligence in Autonomous Vehicle Navigation”

Antreas Nathanail

Supervising Professor

Elpida Keravnou Papailiou

Thesis submitted in partial fulfilment of the requirements for the award of degree of Bachelor in Computer Science at University of Cyprus

May 2025

Acknowledgements

I would like to take this opportunity to express my deepest gratitude to all the people and organizations who have guided and supported me in preparing this diploma thesis. I am grateful for their motivation and input, which made its completion possible.

Above all, I am deeply grateful to my supervisor, Professor Elpida Keravnou-Papailiou, for entrusting me with this challenging topic, her useful feedback, and technical advice throughout the process of developing this thesis. Her profound knowledge in the arena of Artificial Intelligence, as well as their dedication towards academic merit, helped determine both the direction as well as quality of the work, and have been constant source of inspiration.

I would also like to recognize the department's faculty and staff, whose teachings as well as services over the years have formed the foundation of this work.

I am deeply grateful to the institutes, companies, their dedicated employees, and all individuals, both in Cyprus and around the world, who contributed to my research. Their knowledge and experience have been instrumental in shaping this study. The input from companies was essential for gaining industry insights, while the perspectives of the general public provided valuable understanding of societal attitudes toward autonomous vehicles.

I would like to extend special thanks to my fellow students and peers for their useful discussions, technical input, and moral support throughout. Working with them and exchanging ideas have both enriched and inspired me.

Lastly, I extend my gratitude to my family and friends for their unwavering support and encouragement throughout my academic journey. Their belief in my capabilities and constant motivation have been a source of strength and perseverance.

This diploma thesis represents the work of a collective effort, and I am gratefully thankful to all those who assisted in its completion.

Summary

This diploma thesis investigates the role of Artificial Intelligence (AI) and Machine Learning (ML) in the development of autonomous vehicles, focusing on both the technological foundations and societal implications of self-driving systems. Through a structured analysis of literature, technical systems, and research, including surveys conducted with industry professionals and the general public, this study examines how AI is shaping the future of mobility. The research explores key AI technologies such as deep learning, sensor fusion, and perception algorithms, as well as the planning and control models that enable autonomous decision-making. It also traces the historical evolution of autonomous vehicle technology, evaluates current environmental sensing techniques, and outlines the ethical, legal, and societal challenges these systems must overcome.

Findings from the industry survey highlight the main technical bottlenecks, model preferences, and ethical strategies employed by companies developing autonomous systems. Meanwhile, the public opinion survey reveals widespread concerns regarding safety, trust, and accountability, indicating a cautious but growing interest in the adoption of autonomous vehicles. This thesis underscores the need for ongoing interdisciplinary collaboration between AI developers, automotive manufacturers, regulators, and society at large to ensure that autonomous vehicles are not only technologically robust but also ethically and socially accepted. The results contribute valuable insights for researchers, policymakers, and technology developers working toward a safe and responsible future in autonomous transportation.

Περίληψη

Η παρούσα διπλωματική εργασία διερευνά τον ρόλο της Τεχνητής Νοημοσύνης και της Μηχανικής Μάθησης στην ανάπτυξη αυτόνομων οχημάτων, εστιάζοντας τόσο στα τεχνολογικά θεμέλια όσο και στις κοινωνικές επιπτώσεις των συστημάτων αυτόνομης οδήγησης. Μέσω μιας δομημένης ανάλυσης της βιβλιογραφίας, των τεχνικών συστημάτων και της έρευνας, συμπεριλαμβανομένων και των ερωτηματολόγιων που διεξήχθησαν με επαγγελματίες του κλάδου αλλά και στο ευρύ κοινό, η μελέτη αυτή εξετάζει πώς η Τεχνητή Νοημοσύνη διαμορφώνει το μέλλον της (αυτόνομης)

κινητικότητας. Γίνεται έρευνα σε βασικές τεχνολογίες Τεχνητής Νοημοσύνης, όπως η βαθιά μάθηση(deep learning), η σύντηξη αισθητήρων(sensor fusion) και οι αλγόριθμοι αντίληψης(perception algorithms), καθώς και σε μοντέλα σχεδιασμού και ελέγχου που επιτρέπουν την αυτόνομη λήψη αποφάσεων. Επίσης, γίνεται ανασκόπηση για την ιστορική εξέλιξη της τεχνολογίας των αυτόνομων οχημάτων, αξιολόγηση των τρέχουσων τεχνικών ανίχνευσης περιβάλλοντος και περιγραφή στις ηθικές, νομικές και κοινωνικές προκλήσεις που πρέπει να ξεπεράσουν αυτά τα συστήματα.

Τα ευρήματα των ερωτηματολόγιων που έχουν διεξαχθεί υπογραμμίζουν τις προτιμήσεις των μοντέλων, τις ηθικές στρατηγικές που χρησιμοποιούν οι εταιρείες, αλλά και τα κύρια εμπόδια που έχουν καθώς αναπτύσσουν αυτόνομα συστήματα. Ταυτόχρονα, η έρευνα κοινής γνώμης αποκαλύπτει εκτεταμένες ανησυχίες σχετικά με την ασφάλεια, την εμπιστοσύνη και την υπευθυνότητα, υποδεικνύοντας ένα επιφυλακτικό αλλά αυξανόμενο ενδιαφέρον για την υιοθέτηση αυτόνομων οχημάτων. Επίσης αυτή η έρευνα υπογραμμίζει την ανάγκη για συνεχή διεπιστημονική συνεργασία μεταξύ των προγραμματιστών Τεχνητής Νοημοσύνης, των κατασκευαστών αυτοκινήτων, των ρυθμιστικών αρχών και της κοινωνίας γενικότερα, ώστε να διασφαλιστεί ότι τα αυτόνομα οχήματα δεν είναι μόνο τεχνολογικά ισχυρά αλλά και ηθικά και κοινωνικά αποδεκτά. Τα αποτελέσματα προσφέρουν πολύτιμες γνώσεις για ερευνητές, υπεύθυνους χάραξης πολιτικής και προγραμματιστές τεχνολογίας που εργάζονται για ένα ασφαλές και υπεύθυνο μέλλον στις αυτόνομες μεταφορές.

Contents

Chapter 1 Introduction

1.1 Background

Autonomous vehicles (AVs) are set to transform the future of transportation giving the ability to vehicles to operate with minimal or no human interaction. The core technology behind this is Artificial Intelligence (AI), which enables the vehicles to perceive their environment, make decisions, and navigate autonomously. This is done through a combination of Machine Learning (ML), Deep Learning (DL), sensor fusion, computer vision, and decision-making algorithms, so Avs can react to our dynamic, usual, everyday life roads.

1.1.1 Personal Motivation

As a student of computer science, I have been fascinated by the rise of artificial intelligence in recent years. AI is evolving numerous industries, including transportation, and its potential in autonomous vehicles particularly sparks my interest. My personal passion for cars has also driven my interest in this field, so with a strong background in automotive expertise, I am excited to explore new ways AI can enhance vehicle autonomy. The combination of AI and cars brings together two of my greatest interests, and this research is an exciting opportunity for me to dive deeper into a subject that merges technology with one of the most essential aspects of modern life: transportation.

1.1.2 The Role of AI in Autonomous Vehicles

AI in autonomous vehicles involves various aspects that a human's brain does on its own, such as perception, decision-making, and control. Perception is about sensory input from the environment around the vehicle using cameras, radar, LiDAR, and other sensors. Machine learning and deep learning algorithms are used to process these inputs, allowing the vehicle to identify various things like objects, road signs, obstacles or other cars. Additionally decision-making algorithms are needed so the vehicle can determine the actions to do next, such as braking, accelerating, or steering, based on real-time data analysis. This sophisticated integration of AI technologies has enabled AVs to navigate through complex environments, handle traffic, and respond to potential hazards.

1.1.3 Key AI Technologies: Machine Learning, Deep Learning, and Sensor Fusion

ML algorithms are important for the AV systems to recognize patterns in road data, while DL are used for more advanced functions like object detection and semantic segmentation. Sensor fusion is also very crucial since it combines data from multiple sensors, creating an extensive understanding of the environment and reducing the possibility of errors caused by relying on a single data source. Combined, these technologies form the backbone of autonomous driving, allowing AVs to react intelligently in everyday life situations.

1.2 Problem Statement

Despite the advancement of AI, several significant challenges make it difficult for AI systems to always act correctly in real-world driving conditions. These issues could be related to successful object recognition, efficient real-time data processing, and decision-making under uncertainty. Autonomous vehicles should be able to operate in environments that are unpredictable and could vary from heavy traffic to extreme weather conditions. AI must have almost zero error dealing with these challenges, so AVs can be safe, reliable, and efficient to use. This research seeks to address these issues by investigating the current AI technologies used in AVs and identifying which systems and methods are most reliable for autonomous driving.

1.2.1 Challenges in AI Implementation

One of the key problems in AVs is achieving accurate and fast object recognition while ensuring that sensor data is processed in real-time. Path planning is another complex issue, requiring precise calculations to navigate through obstacles and follow the optimal route. Additionally, the prediction of the behaviour of other drivers, pedestrians, and animals in a chaotic environment remains an ongoing area of research.

1.2.2 AI for Autonomous Driving

Specifically, the study will focus on the following:

1. Algorithms and Real-Time Data Processing: AI algorithms are crucial since they empower perception, decision-making and control. Such algorithms are used to process data from various sensors, classify objects and determine a vehicle's next actions in real-time. Very impressive is the way AI algorithms handle the massive volumes of data generated by sensors, ensuring quick and accurate decision-making while driving, being able to ensure that AVs respond to dynamic road conditions instantaneously.

2. Object Recognition, Data Analysis, and Sensor Fusion: Machine learning (ML) and deep learning (DL) algorithms enable AVs to recognize and classify objects such as pedestrians, other vehicles, and road signs. This study will examine how these objects are detected, classified, and analysed using AI technologies. Furthermore, sensor fusion plays a critical role in AVs by combining data from cameras, LiDAR, radar, and other sensors to create an understanding of the vehicle's surroundings. The research will explore how all these are managed.

3. Path Planning and Predictive Capabilities: AVs should be able to determine the optimal path of the journey. The study will examine how AI systems can handle path planning, considering factors such as traffic rules, road conditions, and potential obstacles. Additionally, it will examine how AVs predict future events, such as the movements of nearby vehicles or pedestrians from the use of real-time and historical data. This predictive capability is critical in avoiding collisions and ensuring smooth navigation.

4. Decision-Making for Vehicle Control and Training of AI Systems: AI systems control the vehicle's actions, including acceleration, braking, and steering. Rapid and precise decisions are required to ensure the safety of passengers and other road users. This study will examine the way AI systems process sensor data and make decisions regarding the vehicle's control. Furthermore, the research will

explore the use of various training methodologies, such as reinforcement learning, supervised learning and simulations.

5. Achieving Efficiency in Autonomous Driving: AI systems are not only designed to ensure safety but also travel efficiency. Efficiency in AVs is not only fuel or energy consumption, but it also includes the optimization of the vehicle's driving strategies, like minimizing processing times, ensuring minimal risk and achieving a comfortable and pleasant experience for the users. The study will investigate how AI systems balance safety and efficiency, while operating smoothly and effectively.

1.2.3 Ensuring Safety, Reliability, and Ethics

The second major objective of this study is to explore how AI systems in AVs are designed to ensure safety and reliability, as well as how they handle ethical considerations in complex driving scenarios. The following areas will be covered:

1. Accident Prevention and System Failures: AI systems in AVs are developed to prevent accidents by predicting potential hazards and responding to them proactively. This study will examine how AI systems detect and react to dangerous situations, including other vehicles' sudden stops or pedestrians crossing the road. Moreover, the study will investigate how these systems are designed to manage hardware or software failures, ensuring that AVs are safe even when critical components malfunction.

2. Safety Checks and Handling Sensitive Data: Safety is paramount in the deployment of AVs. AI systems undergo rigorous testing to ensure they perform reliably under a wide range of driving conditions. This study will assess the methodologies used to validate the safety and reliability of AI systems. Additionally, AVs handle vast amounts of data, including sensitive personal information such as passenger behaviour or vehicle location. The study will explore how AI systems manage and protect such sensitive data, ensuring compliance with data privacy regulations and safeguarding against misuse.

3. Ethical Decision-Making in Complex Situations: AVs may face difficult ethical dilemmas, such as deciding between undesirable outcomes in emergency situations. We will explore how AI systems are programmed to handle these ethical challenges, ensuring that the vehicle makes decisions that prioritize human lives and minimize harm. The research will also investigate the ethical frameworks that guide the development of these systems, ensuring that AVs behave responsibly in life-threatening situations.

4. Legal and Accountability Issues: As AVs become more common, questions about legal liability and accountability in the event of an accident will become increasingly important. We will analyse the legal frameworks being developed to address issues of fault and responsibility when an AV is involved in an accident. The research will also explore how AI systems are designed to minimize the likelihood of accidents and how they document incidents to ensure transparency in the decision-making process.

1.3 Objective of the Thesis

In the context of the fast advancement of autonomous car technology and the increasing utilisation of Artificial Intelligence in this field, this thesis is framed with objectives for analysing both technological as well as societal components of the integration of AI. The primary goals of this research are to investigate the current position of AI in autonomous driving, determine technical as well as ethical issues, and collect inputs from both experts in the field as well as the public.

1.3.1 Research Goals

1. The role of AI in Autonomous Vehicles:
   The first objective is to explore how AI technologies like machine learning, deep learning, and sensor fusion are being utilized to enable perception, planning, and control of autonomous vehicles. This entails an in-depth examination of the technical infrastructure and the algorithms that support current self-driving systems.

2. Identifying Technical and Ethical Challenges:
   Another purpose of the thesis is to determine the existing challenges in adopting AI in autonomous vehicles, such as sensor limitations, quality of data, real-time processing limitations, as well as ethical decision-making in situations of importance. Safety, reliability, and accountability are crucial in autonomous decision-making systems.
3. Capturing Industry Views:
   With a focused industry survey, the current research aims to capture industry leaders' priorities, challenges, and strategic plans for developing autonomous car technologies. The results should provide insights about how industry experts view the maturity of AI techniques, model choices, data handling, and lastly, their opinion on the future of full autonomy.
4. Evaluating Public Opinion and Worries:
   One of the main goals is to indicate public perception of reliability, safety, and ethical considerations of AI in driverless cars. The results captured of this particular survey can be an indication of readiness for adoption as well as public expectations for developers and policymakers.

1.4 Significance of the Study

This study is significant since it offers a complete exploration of the role of Artificial Intelligence in autonomous vehicles, and it covers both technical foundations and societal implications. The analysis of key AI technologies and their application in perception, planning, and control helps the audience understand more clearly how autonomous systems operate. In addition, including the industry insights and public opinion helps to capture a balanced perspective on the challenges, expectations, and ethical concerns about the adoption of self-driving technologies.

1.5 Research Questions

The core of the study, which is Artificial Intelligence in autonomous vehicles are several research questions that steer the direction and scope of this study. These inquiries aim to examine the technical integration of AI in autonomous driving along with the societal readiness for their implementation, concentrating on practical challenges, industry standards, and public viewpoints.

1.5.1 Primary Question

The primary research question that this diploma thesis seeks to answer is:
"How is Artificial Intelligence currently applied in autonomous vehicles, and what are the main technical, ethical, and societal challenges hindering its full adoption?"

This central question is further examined through a series of sub-questions:

1. Technological Integration: What AI technologies are used in perception, motion planning, and control systems of autonomous vehicles, and how do they work?
2. Implementation Challenges: What are the primary technical and computational limitations in deploying AI in real-time driving environments?
3. Industry Perspective: How do companies working in the autonomous vehicle sector perceive the development status, model choices, and data requirements of AI systems?
4. Public Perception: What are the key safety, trust, and ethical concerns of the public in regards of AI-driven vehicles?
5. Ethical and Regulatory Issues: What legal and ethical dilemmas emerge from AI decision-making in critical or ambiguous traffic situations?

1.6 Scope and Limitations

It's important to define the scope and acknowledge the inherent limitations of this research to provide a focused and meaningful analysis. This section will outline the specific scope of the study and the constraints under which this research was conducted.

1.6.1 Focus of the study

The scope of the thesis addresses the following areas:

1. Technological focus:

   The focus here is particularly on the application of machine learning, deep learning, and sensor fusion in perception, planning, and control systems.
2. Industry Engagement:

   It includes insights from professionals and organizations actively working in the autonomous vehicle or AI development sector, offering real-world perspectives on current capabilities, challenges, and ethical concerns.
3. Public Perception:

   A significant portion of the study is devoted to understanding public attitudes toward AI-powered autonomous vehicles, especially regarding safety, ethical dilemmas, and trust in automation.
4. Geographical Context:

   While the technological analysis is globally relevant, and even though the surveys have answers from people all over the world, most of them are from Cyprus. So, this case study serves as societal readiness in a smaller European context.

1.6.2 Limitations

Despite efforts to ensure comprehensive and accurate research, this study encounters certain limitations:

1. Response Rate and Bias in Questionnaires:

   The primary data collection method relies on responses to questionnaires distributed to companies, institutes, and employees. The effectiveness of this method is subject to the response rate and potential response bias.
2. Dynamism of the Industry:

   The IT industry is marked by quick technological shifts and changing skill demands. Although the report strives for timeliness in its analysis, the dynamic nature of the industry makes some of its observations potentially outdated in short order.
3. Generalizability of Results:

Since the emphasis in this study is in Cyprus, the results and inferences of the current work could not be fully generalized to other areas or countries with varying economies, cultures, or educational environments.

4. Subjectivity in Qualitative Analysis:

   There are some components of the research, primarily those dealing with analysing the qualitative data from questionnaires, that have elements of interpretation involved, which can lead to subjectivity.

By acknowledging these constraints, the current study adopts an open and critical approach to its conclusions and recommendations. Readers need to consider these constraints when interpreting the results and conclusions drawn from this research.

1.7 Structure of the Thesis

- **Chapter 2: Evolution and Future of Autonomous Vehicles: Technological Advancements, Impacts, and Implications**
  - Provides a historical overview of autonomous vehicle development from the early 20th century to modern AI-integrated systems.
  - Discusses levels of vehicle autonomy and evaluates the environmental, economic, social, and ethical implications of widespread AV adoption.
- **Chapter 3: Environmental Perception and Sensing Techniques**
  - Explores the core sensing technologies (LIDAR, radar, cameras) that enable autonomous vehicles to perceive their environment.
  - Analyses the perception subsystem and key AI algorithms used for detection, tracking, and object classification.
- **Chapter 4: Decision-Making in Autonomous Vehicles**
  - Examines the planning and control architectures used in self-driving systems, including kinematic and inertial models.
  - Details motion planning techniques such as trajectory generation and graph search methods, along with control mechanisms for path stabilization.
- **Chapter 5: Methodology**
  - Describes the research framework used in this study, including the design and distribution of industry and public surveys.

    - Discusses data collection techniques, ethical considerations, and limitations of the research.
- **Chapter 6: Industry Insights on Autonomous Vehicle Development**
    - Presents survey results from AI and automotive industry professionals.
    - Identifies common technical challenges, ethical concerns, model usage, and future expectations for AI in AVs.
- **Chapter 7: Public Opinion Survey**
    - Analyses public perspectives on the safety, reliability, and ethical concerns related to autonomous vehicles.
    - Evaluates trust in AI, willingness to adopt the technology, and perceived legal responsibilities.
- **Chapter 8: Conclusion**
    - Summarizes key findings from the technical review and survey analyses.
    - Highlights the importance of interdisciplinary collaboration and provides recommendations for future research directions.
- **Bibliography**
    - Lists all references cited throughout the diploma thesis.

Chapter 2: Evolution and Future of Autonomous Vehicles: Technological Advancements, Impacts, and Implications

The journey of autonomous vehicles (AVs) is one marked by a gradual yet transformative integration of artificial intelligence (AI) and related technologies over nearly a century. This chapter traces the major historical milestones in AV development, emphasizing how AI has evolved from supporting simple automated systems to enabling advanced self-driving cars capable of making complex decisions in real time. The chapter will explore AV development's social, economic, environmental, and regulatory impacts, providing a well-rounded perspective on the current state of autonomous vehicles and their future trajectory.

2.1 Historical Overview

The development of autonomous vehicles has followed a long trajectory, beginning with early experiments in the 1920s and evolving through various phases. Initial efforts focused on remote-controlled cars and automated roadways, laying the groundwork for future innovations. Over time, advancements in electronic systems, sensors, and artificial intelligence (AI) transformed these vehicles, allowing them to navigate complex environments with increasing autonomy. By the late 20th and early 21st centuries, AI-driven technologies became integral, enabling vehicles to perform tasks like navigation, obstacle avoidance, and decision-making with minimal human input. Today, AI continues to play a central role in pushing the boundaries of autonomous driving.

The historical review presented in this subsection is taken from [1].

2.1.1 Early Experiments in Autonomous Driving (1920s-1950s)

The quest for autonomy in vehicles began as early as the 1920s, long before the term "artificial intelligence" had even been coined. One of the earliest appearances was the development of the Linriccan Wonder in 1926, a radio-controlled car that was the first step towards autonomous driving. The Linriccan Wonder was operated remotely using radio signals, showing the potential of vehicles that could be controlled without a human driver. A modified version of this car, called the Phantom Auto, was showcased later that year, adding to the fascination with the possibility of driverless cars.

In the 1930s, the concept of automated roadways was introduced, at the 1939 World's Fair, Norman Bel Geddes showed his vision of autonomous electric cars, which would be powered and controlled by circuits embedded in the road. The idea of roads "communicating" with vehicles became a base for later developments in AV technology.

By the 1950s, RCA Laboratories expanded on this vision with the creation of a model car guided by wires embedded on a laboratory floor. This innovation set the stage for real-world experimentation, especially in Nebraska, where engineers tested a prototype in a closed highway, using signals from detector circuits embedded in the pavement to guide vehicles.

2.1.2 The Rise of Electronic and Vision-Based Systems (1960s-1980s)

The 1960s marked a shift towards more sophisticated electronic guidance systems. Projects like the Ohio State University's driverless car experiment and the UK's Citroën DS tests showed the feasibility of using magnetic cables and sensors to control vehicle movements. The Citroën DS, in particular, was able to travel at speeds of up to 130 km/h, maintaining a steady course in all weather conditions, showcasing the potential for reliable automation in real-world settings.

However, a major downward trend in autonomous driving technology occurred in the 1980s with the introduction of vision-guided systems. This period saw significant integration of AI, as systems began to incorporate machine vision, LIDAR, radar, and GPS technologies. One of the most important breakthroughs came in 1987 when Mercedes-Benz, under the guidance of Ernst Dickmanns and his team at the Bundeswehr University in Munich, developed a vision-guided robotic van. The vehicle was able to autonomously navigate streets at speeds up to 63 km/h, marking a major milestone in AV development.

In the same decade was also the launch of EUREKA's Prometheus Project, which ran from 1987 to 1995, with over $1 billion USD invested in research on autonomous driving. This project emphasized on the importance of integrating various technologies such as computer vision, sensor data fusion, and AI for real-time decision-making, setting the stage for future developments.

2.1.3 The Role of AI in Modern Autonomous Vehicles (1990s-2000s)

In 1990s, projects such as DARPA's Autonomous Land Vehicle (ALV) and Carnegie Mellon University's NavLab pushed the boundaries in the use of AI in autonomous vehicles. These projects made a significant use of AI-driven systems for navigation, obstacle avoidance, and route planning.

In 1991, Mercedes-Benz, under Ernst Dickmanns, showcased the twin robot vehicles VaMP and Vita-2, which could autonomously drive on highways in heavy traffic. These vehicles demonstrated not only autonomous lane following but also the ability to pass other cars and change lanes, all with minimal human intervention. This marked a new era in AV technology, with AI playing a critical role in real-time decision-making [1].

By 1995, the NavLab project at Carnegie Mellon achieved a 98.2% autonomous cross-country drive during the "No Hands Across America" experiment. This remarkable achievement highlighted the growth of neural networks in controlling vehicle functions like steering. While throttle and braking were still controlled manually, this project showed that AI could successfully handle complex tasks like lane keeping and obstacle avoidance.

One of the most noteworthy advancements in this period was the ARGO Project by the University of Parma. Launched in 1996, this project demonstrated a fully autonomous 1,900 km journey through northern Italy, with AI-based stereoscopic vision used for detecting lane markings and ensuring safe travel. The vehicle maintained a 94% autonomous mode, showcasing the increasing reliability of AI in handling real-world driving conditions.

2.1.4 AI and Autonomous Public Transport Systems (2000s-Present)

The early 2000s marked the beginning of widespread use of autonomous technologies in public transportation systems. The Netherlands, for instance, introduced the ParkShuttle, an autonomous public road transport system that operated on predefined routes without human drivers. Around the same time, the U.S. government also began using autonomous vehicles for military purposes, with projects such as Demo III, funded by DARPA, focusing on ground vehicles capable of navigating complex terrains using AI-based control systems.

As AI and machine learning algorithms became more advanced, autonomous vehicles became increasingly capable of handling real-time data processing and making complex decisions. This was evident in the 2004 and 2005 DARPA Grand Challenges, where AI-enabled vehicles navigated difficult off-road courses, avoiding obstacles and completing predefined routes. These competitions were a turning point in AI integration into AV systems, demonstrating the effectiveness of machine learning in enhancing vehicle autonomy.

By 2010, AI started to get integrated in the development of semi-autonomous and fully autonomous features in commercial vehicles. Tesla’s Autopilot, introduced in 2015, became one of the most prominent examples of AI-driven autonomy in consumer vehicles. Using a combination of computer vision, LIDAR, radar, and AI-based decision-making, Tesla's Autopilot could autonomously steer, brake, and accelerate under certain conditions, demonstrating the feasibility of integrating AI into mass-market vehicles.

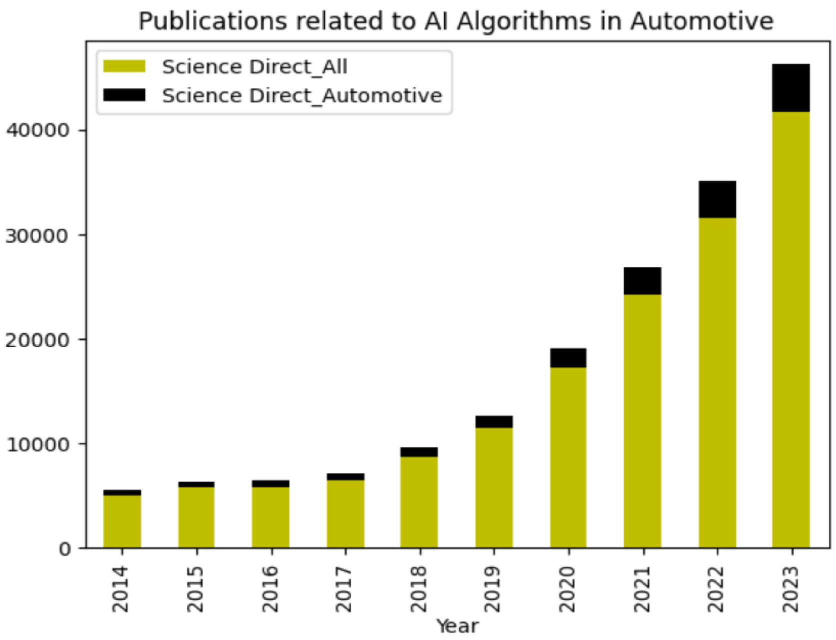

Figure 1.1: Distribution of Publications Related to AI for AVs

The above graph in Figure 1.1 shows that there were relatively few publications between 2014 and 2018. However, in 2018, there was a sharp increase in publications, reflecting the growing interest and development of autonomous vehicles with advanced self-driving features. Based on this trend, it is expected that the number of related publications will continue to rise exponentially in the future [6].

2.2 Levels of Autonomy in Vehicles

The term "autonomous" carries various definitions, and this ambiguity has prompted experts in the field of autonomous vehicles to classify autonomy into distinct levels. These levels represent the extent of automated system control and the degree of driver involvement required in operating a vehicle. These levels range from Level 0 (no automation) to Level 5 (full automation), like we see on figure 2.1. Below is an overview of these levels and their implications for AV development [8].

2.2.1 Level 0-2: Driver Assistance Systems

At level 0, it’s Manual Driving, since the vehicle provides no autonomous capabilities. The driver is in full control, and the system offers only optional warnings or alerts for potential hazards. Moving to level 1 is the Advanced Driver Assistance System (ADAS). Here, the driver retains full control of the vehicle, but automated systems assist with specific tasks like steering or acceleration. The car is working alongside the driver to enhance safety and convenience. At Level 2 is Partial Automation. The vehicle's automated systems can take over complete control of the car, but under certain conditions, but the driver must be ready to take control at any moment. [8]

2.2.2 Level 3-4: Conditional and High Automation

At level 3 is Automated Driving System (ADS). In this level the vehicle can operate completely on its own, handling all driving tasks without requiring constant driver input. However, drivers must remain ready to take control if the system alerts them to intervene in specific situations and at any time. Moving to level 4 it's High Automation. The vehicle can perform all driving tasks and monitoring the environment without any need for driver engagement or observation in most scenarios. However, in rare or unforeseen circumstances, the driver may still need to regain control[8].

2.2.3 Level 5: Full Automation

Level 5 autonomy represents the ultimate goal of AV development, where vehicles can operate in any environment without human intervention. In other words, the vehicle's Automated Driving System (ADS) is functioning as a virtual driver. Achieving this level requires overcoming significant challenges as we will see later in this research [8].

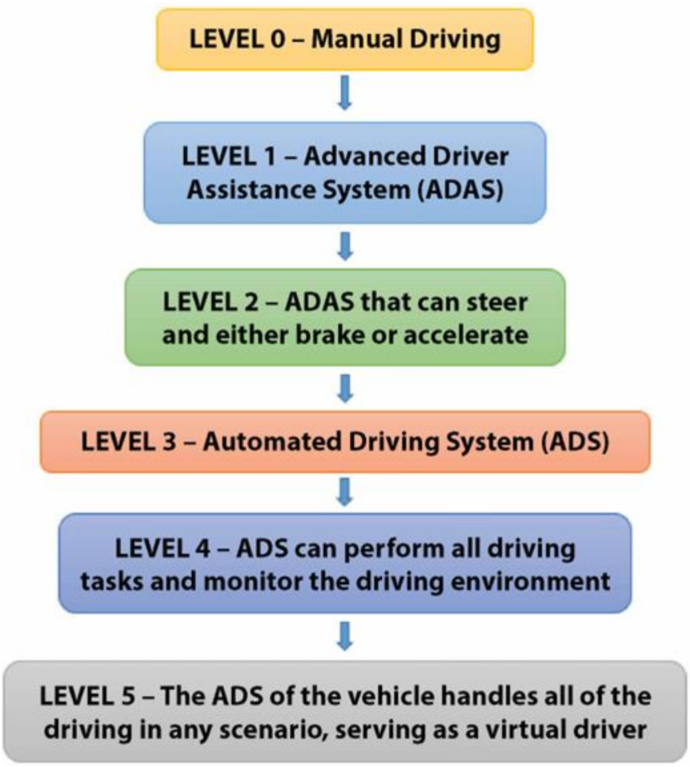

Figure 2.1: Autonomous levels diagram

2.3 Future Prospects: Fully Autonomous Vehicles and AI's Role

The future of self-driving cars looks very promising, with new advancements expected to change transportation forever. Using smart AI systems, these vehicles will soon handle navigation, decisions, and driving with little or no help from humans. This will make roads safer, transportation faster, and driving more reliable. These fully self-driving cars will be capable of handling busy city streets and mixed traffic and are expected to be available in the next ten years.

Big car companies are already making progress in this area. For example, Audi and Nissan started offering features like automatic steering and lane-keeping back in 2015. Volvo plans to make cars that aim to eliminate serious injuries or deaths in crashes. Tesla has developed autopilot features and is working toward cars that can drive themselves 90% of the time. Google has been developing self-driving cars that could take over all driving tasks.

Of course, AI plays a key role in making this happen, since it offers real-time data from sensors and advanced technology for smart decision-making. Experts predict that by

2035, most cars will be fully self-driving, with no need for human control. While challenges like creating rules and addressing safety concerns remain, the progress so far shows a bright future for autonomous vehicles, where accidents are fewer, and travel becomes easier for everyone [1].

2.4 Environmental Impact and Sustainability

Autonomous vehicles (AVs) can significantly enhance environmental sustainability by reducing transportation emissions, which account for around 24% of global greenhouse gas (GHG) emissions. In the U.S., transportation is the largest GHG contributor, with passenger cars being a primary source. AVs optimize traffic flow, reduce idling, and enable fuel-efficient driving patterns, cutting pollutants like CO, NOx, and CO2. Shared autonomous vehicles (SAVs) may further reduce private car ownership and urban congestion. Systems like adaptive cruise control (ACC) and cooperative adaptive cruise control (CACC) can reduce fuel consumption by up to 47%. Shared electric AVs can cut GHG emissions by as much as 94% compared to traditional vehicles. However, the increased demand for travel could offset these gains through higher vehicle miles travelled (VMT) and congestion. Despite all the above, if AVs integrate with electrification and smart traffic management, they can offer significant potential to reduce transportation's environmental footprint. Figure 2.2 illustrates factors that contribute to reducing greenhouse gas (GHG) emissions, these are gases that trap heat in the Earth's atmosphere, contributing to global warming and climate change [4].

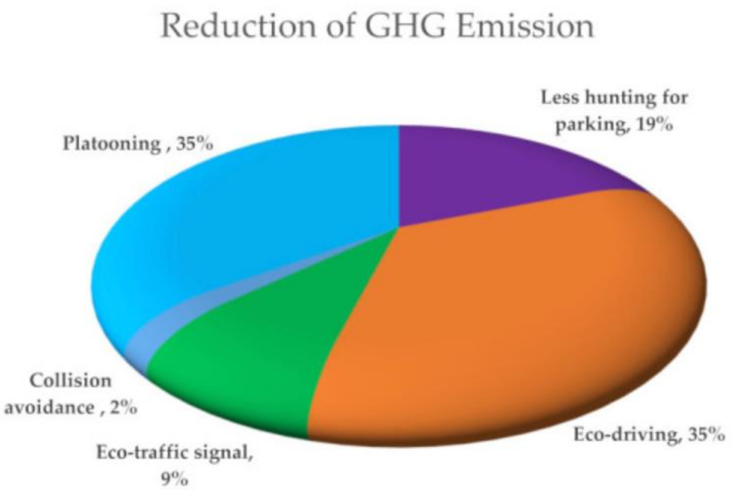

Figure 2.2: Percentage Effect for Factors Reducing GHC Emissions

2.5 Economic Transformation and Job Market Implications

With the integration of autonomous vehicles (AVs) the economy and job market are likely to reshape. Economically, AV technologies, particularly advanced driving systems like Level 3 (L3) and Level 4 (L4) autonomy, are projected to generate \$300 billion to \$400 billion in revenue by 2035. Automakers and technology providers most probably are going to adopt new business models, invest in new technologies, and address market demands to capitalize on these opportunities. However, this transformation would bring significant challenges to the job market. Roles in driving-dependent industries, such as trucking and delivery, will face displacement as automation reduces the need for human drivers. While AV adoption may create new opportunities in AI development, system maintenance, and technology-based jobs, people working in automotive would need to transition from traditional roles to these modern ones. Additionally, industries like insurance and vehicle repair may see major shifts due to fewer accidents. Figure 2.3 below illustrates an approximation of the revenue that each level can create in the future. It's clear that as the years pass, higher levels bring substantially more income while lower levels have a slight downward trend [3].

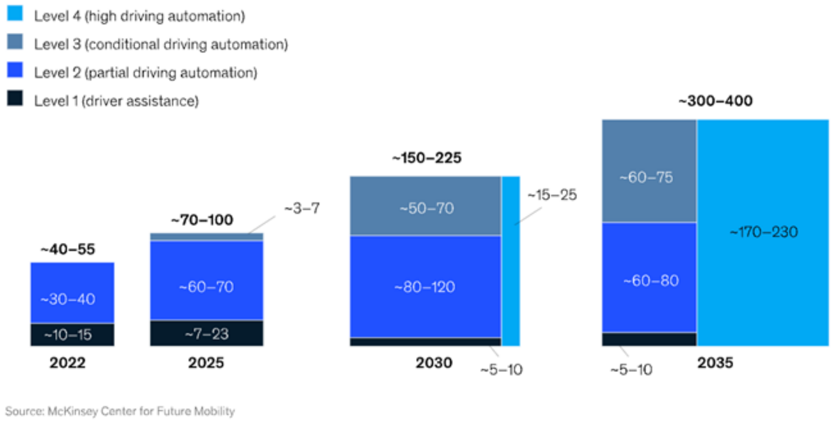

Figure 2.3: Advance Driver-Assistance Systems (ADAs) & Autonomous Driving Revenue in Billion\$

2.6 Safety and Health Benefits

Autonomous vehicles promise to improve road safety by reducing human errors that account for most traffic accidents. Widespread AV adoption could lead to a significant reduction in road fatalities and injuries. This improvement in safety could also lead to lower healthcare costs associated with accident-related injuries. Moreover, by reducing congestion and preventing collisions, AVs could help minimize harmful emissions, positively affecting public health [2].

2.7 Accessibility and Urban Development

Autonomous vehicles (AVs) have the potential to improve accessibility for people with disabilities, the elderly, and others unable to drive, offering door-to-door services, like autonomous taxi, that enhance independence. Beyond mobility, AVs could transform urban development by reducing reliance on private car ownership and parking, freeing up space for green areas and improving cityscapes. Shared mobility concepts are central to this vision, promoting efficient vehicle use and fewer cars on the road.
As Professor Malene Freudendal-Pedersen (Professor in Urban Planning) said, "We, as a society, have to get used to sharing things." [5].

2.8 Ethical and Regulatory Challenges

Autonomous vehicles (AVs) bring important ethical and legal questions that must be addressed before they can be widely adopted. One major issue is how AVs should respond to unavoidable accidents. Programming them to prioritize saving lives or choosing between passengers and pedestrians raises ethical dilemmas. Avoiding biases in these decisions, such as preferring certain individuals based on age or other traits, is crucial.

On the regulatory side, AVs must operate within laws designed for human drivers, creating challenges when vehicles need to break rules to prevent harm. Also, there are questions like, whether the responsibility of accidents falls on manufacturers, programmers, or users. Strict safety standards and clear regulations are essential to build trust and ensure these systems are safer or just safe as human drivers [7].

2.9 Conclusion

The history and progression of autonomous vehicle development highlight the transformative role AI has played in evolving AV capabilities. From radio-controlled prototypes in the 1920s to sophisticated AI-driven systems today, each milestone has brought AV technology closer to a future where fully autonomous vehicles become commonplace. The societal, economic, environmental, and ethical impacts of AVs are vast, reshaping not only transportation but also the way we live and work.

Chapter 3: Environmental Perception and Sensing Techniques

3.1 Introduction

For AVs, the demand for robust and accurate sensing and perception technologies is based on the complexity of autonomous navigation, particularly in unstructured environments such as urban traffic. Because of that, DARPA came into the picture. DARPA is the Defense Advanced Research Projects Agency, a U.S government agency that is responsible for national security technology. It is known for creating innovative projects, including early research in autonomous vehicles through its well-known DARPA Challenges. DARPA played a huge role in the development of AVs, which started from driving in the desert and later moved to more complex city environments. This evolution required better sensors and smarter perception systems. While other types of autonomous robots exist, like those for air (drones) or water (submarines), AVs have their own problems, like the need to drive fast in places that can be chaotic and unpredictable, like city streets with many different people, vehicles, and obstacles [10].

Figure 3.1 below represents a high-level diagram of how autonomous vehicles work, their architecture, and their main subsystems. Sensors capture information about the environment surrounding the vehicle and send it to the perception subsystem. This subsystem extracts useful information and converts it into data that the planning and control subsystem can understand, as this data is then transmitted there. Using navigation algorithms and the data mentioned above, the references used by the actuator control subsystem are generated, so the machine can make actions and interact with the environment. The levels of precision and complexity of this data depend on the different algorithms being employed [10].

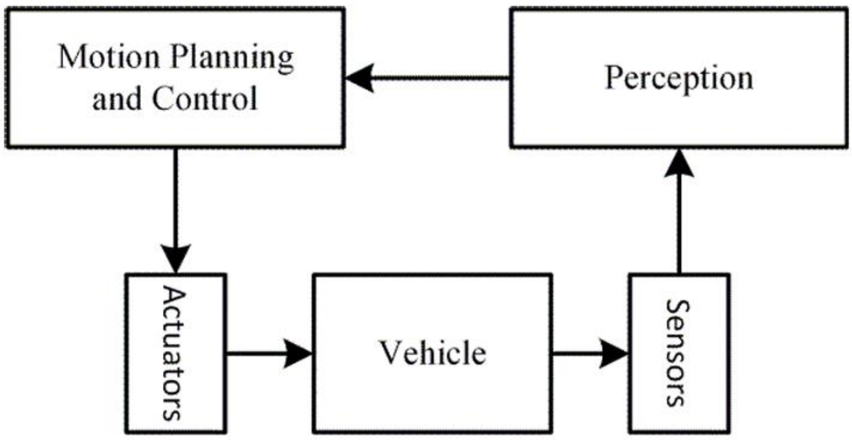

Figure 3.1: Control system architecture of an autonomous vehicle

3.2 Main Types of Electronic Sensing Devices

Electronic sensing devices capture real-time data from AVs' surroundings. The primary types of them utilized in AVs are:

3.2.1 Cameras [10]

Cameras give 2D visual data with a digital image of the covered region of space. They use passive light sensors, with each sensor capturing a pixel's value. Images can be monochrome or color, with resolutions from 640×480 to 1328×1048 pixels. Since image is in 2D advanced image processing is required to infer depth information. In figure 3.2 is displayed what the Cameras of a car capture.

In vehicles, cameras face two key challenges: low-light sensitivity and capturing both bright and dark areas (dynamic range). Sensitivity depends on sensor efficiency and signal amplification, which must enhance useful signals, not noise. For night vision, cameras should also detect infrared (IR) radiation which is what is reflected by the objects, often using external IR LEDs for better visibility. Far-IR (thermal radiation) cameras are those that measure radiation emitted by objects, but they are rarely used. The amount of light that every object reflects can range from 1 to 10 lux at night and up to

105 lux with sunlight, so high-quality automotive cameras need a dynamic range of about 120 dB (standard ones offer ~60 dB).

Figure 3.2: Camera captures on a car [13]

3.2.2 LIDAR Sensors (Light Detection and Ranging) [10]

LIDAR systems are critical for precise distance measurement and spatial mapping. These systems and others as well (like radar which will analyze below) are active sensors, and they are based on the principle of sending a pulse of wave (of certain nature and frequency), receiving it reflected by the obstacle and measuring the total time of flight. In figure 3.3 is displayed what LIDAR sensors of a car capture. What we need to consider is that their performance in range is heavily influenced by the reflectivity of target objects and environmental conditions, such as humidity. Typically, LIDARs are categorized based on range, Long-range LIDAR with a maximum range that exceeds 50 meters and Short-range LIDAR for smaller distances. Therefore, range is sometimes indicated by either its maximum value, or by a value at which poor reflecting objects (typically objects with 10% reflectivity) can be detected. For example, the maximum range can be 80 m, whereas its value for objects with 10% reflectivity is 30 m.

LIDAR systems scan environments by either oscillating a single laser beam or using multiple beams simultaneously, in this case both approaches are used simultaneously. These devices are further divided into, based on the shape of the environment part they scan:

- 2D LIDARs: Scan a single plane radially and often horizontally, covering a sector of a circle with typical Fields of View (FOV) ranging from 80° to 270°.
- 3D LIDARs: Cover volumetric space, described by both horizontal (azimuth) coverage (e.g., 270° or 360°) and vertical (elevation) coverage (e.g., 26.8°). For this to be achieved, several laser beams are positioned so that together they cover

the given part of the vertical plane, whereas the entire unit spins to cover the horizontal one. Because typically a laser beam covers less than one degree, a few dozens of laser beams are normally used (e.g. 32 or 64). For each of these beams, separate emitter and receiver units are needed.

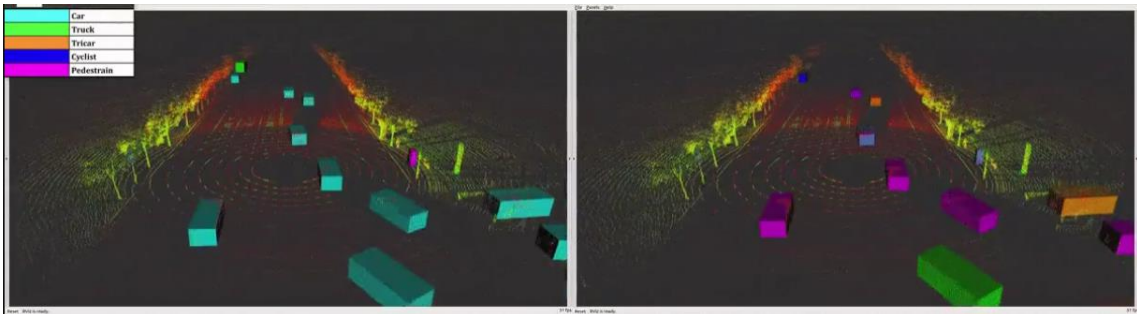

Figure 3.3: What an AV “sees” after the combination of camera and LIDAR inputs [13]

3.2.3 Radar Sensors [10]

Radar uses electromagnetic radiation in different frequency bands (e.g., 24 GHz, 77 GHz) and is less affected by weather. To achieve a good angular resolution, it emits the signal as narrow beams, and the resolution improves with other methods like direction of arrival or spatial spectrum measurement.

Radar types are divided by their range as well; those with a range up to 100-150m are considered short-range or mid-range, and those with a range up to 250 m are considered long-range. Worth mentioning is that only short-range radars can detect objects at short distances.

What makes radar crucial for AVs is the fact that they can determine the velocity of the detected objects, something that plays a significant role in decision-making.

Less prevalent is the usage of ultrasonic (US) and contact sensors. Most AVs use GPS for positioning purposes, the accuracy being supplemented by adding data from Inertial Measurement Units (IMUs). IMU readings may also be used to calibrate readings from other sensors.

LIDAR, radar, and ultrasonic sensors are considered active, since they transmit a wave, receive its reflection from an object, and measure the time taken for the wave to come back to calculate the distance.

The table below (Figure 3.4) summarizes the main properties of the sensors mentioned above. Based on that is clear that none of these technologies is ideal, so to achieve the best outcome is to use two or more simultaneously.

Criteria	LIDAR	Radar	Camera	US
Very short range (0-1m) detection	Poor	Only for short range radar	Ok	Very good
Short range (1-30 m) detection	Very good	Very good	Good	Poor
Long range (30-100+ m) detection	Medium	Very good	Poor	No
Angle < 10	Very good	Good	Good	Poor
Angular resolution	Very good	Good	Good	Poor
Velocity measurement	No	Yes	No	No
Operation in adverse weather conditions	Poor	Very good	Poor	Good
Operation at night	Very good	Very good	Limited	Very good

Figure 3.4: Main technologies for environment perception

3.3 Perception Subsystem

This subsystem's purpose is to determine what parameters to give in motion planning and control subsystems. These parameters are based on the information provided by various sensors and could be the coordinates of the road, the position and speed of obstacles, the position of pedestrian crossings, and much more. The more complexity the environment has, the more parameters are needed. In those highly unstructured environments, like usual urban traffic, the perception subsystem should detect and classify different types of objects around the vehicle and track them to predict their future location. Lastly, some AVs, either prototypes or final products with a targeted audience are designed to operate on specific occasions like off-road, or with another vehicle in front to follow. Evidently, in these cases the perception subsystem has less requirements. [10]

3.3.1 Perception Subsystem Architecture [10]

The subsystem is divided into three main sections, detection, classification and tracking like is presented in Figure 3.5 below. First is detection which handles the separation of multiple objects from the environmental data provided by the sensors. Sensors like those mentioned above (cameras or LIDAR), provide data as a matrix, with every point in them describing a specific point in the real world which is covered by a sensor. The number of those points depends on the resolution the sensor has. Next is classification, where the system decides in what class each object belongs to (e.g. pedestrian, car, traffic lights, etc.). Lastly, is tracking where the system calculates the changes of position for each object from one data sample to another.

To make the system faster and more efficient several tasks are done in parallel, this can be achieved with the use of other dedicated sensors. For instance, lane markers detection or traffic signs can have different cameras/LIDAR from other tasks.

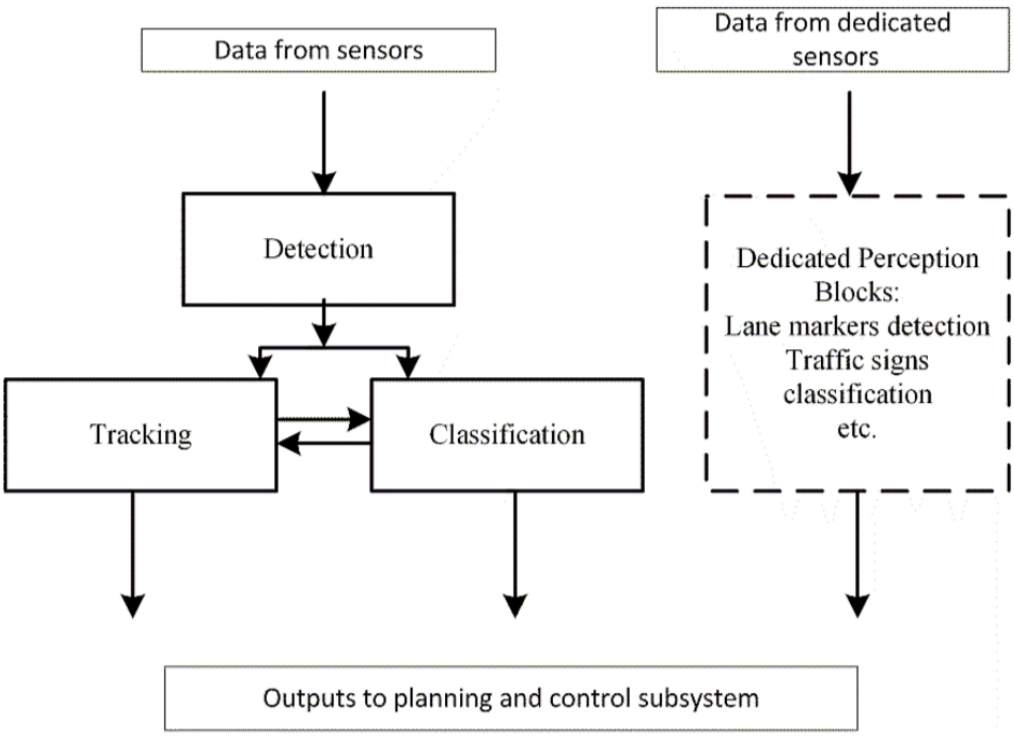

Figure 3.5: Architecture of perception system in AVs.

3.4 Perception Algorithms [10]

Below we will analyse the main algorithms that are used in AVs for object detection, classification and tracking. The data mainly comes from cameras and 3GLIDAR and as said before it has the form of a matrix.

Below, the term images is used to describe the matrix output of a camera or LIDAR.

3.4.1 Detection Algorithms [10]

The main objective of these algorithms is to distinguish the objects that exist in an image. This can be done in several ways including background subtraction and segmentation. Background subtraction as the name suggests is based on removing the image background, allowing only objects to remain. This method is very useful for the identification of moving objects, since on that occasion the background is relatively fixed through the frames, so the main differences between the frames is the position of the objects.

Segmentation on the other hand, separates the region of the image that has similar points, in other words it uses a certain quantification of similarity. One of those algorithms is mean-shift clustering. It starts with several possible starting centers spread all over the image. In each iteration it updates these centers by moving towards the average of the data points within some ellipsoidal shaped boundary around their current location. This move from the old to the new center is a mean-shift. This goes on until all the centers are stable and cease to move. During these iterations, some cluster centers might merge.

3.4.2 Tracking Algorithms [10]

These algorithms are responsible for detecting the movement of object points across different frames. There are several methods to achieve that, with Kalman filter being one of the most widely used. A Kalman filter begins with the basic structure of a linear observer but calculates the gain matrix (Kalman gain) as a solution of the Riccati equation, incorporating the noise and its covariance. Kalman filters are designed for state variables having a Gaussian distribution. If the distribution is non-Gaussian, a particle filter can be employed.

3.4.3 Classification Algorithms

Classification algorithms categorise input data into predefined classes or labels based on learned patterns. Below, we will discuss some of those algorithms, with some being foundational in the development of AVs and others being in-use now.

3.4.3.1 SVM (foundational)

A Support Vector Machine (SVM) binary classifier does not use any probabilistic methods, since it classifies which of two classes the input belongs to according to the training data. Because it's only able to divide two classes at a time, multiple SVMs must be employed for multi-class classification. In summary the algorithm starts with a set of training pairs (xi, yi) of dimension L and based on them, it computes ∞i coefficients (by solving a quadratic problem), and the support vectors S (of size Ns) are then given by the (x, y) elements corresponding to ∞I > 0. With these, another two values, w and b can be computed and finally give as y′ = sign(w*x′ + b), where x′ is a new input vector. What is described above is a linear problem, SVM can also handle non-linear problems by applying what is known as the “kernel trick” [10].

3.4.3.2 Boosting Algorithms (foundational) [10]

Boosting is a technique for enhancing accuracy by aggregating numerous weak classifiers into a strong, more accurate model. Boosting includes a variety of methods such as AdaBoost, GentleBoost, and JointBoost.

The theory here is that each of the weak classifiers performs a minimal amount better than guessing completely at random, however when many of these are put together, we end up with a very effective classifier. The weak classifiers are constructed iteratively. For each iteration after the training of one classifier, the algorithms find the errors and assigns more weight to them for the next classifier. This way, each new classifier pays attention to the previous one's errors, enhancing the overall accuracy. The number of weak classifiers may be fixed from the start, or it can be calculated dynamically during iterations to minimise errors, for example, until the classifier error cannot be further minimized.

3.4.3.3 Bayes Filter (foundational) [10]

This technique can be used when we have the conjoint of two probabilities. For instance, P(B|A) is the probability of B if A is true and P(A|B) is the probability of A if B is true. With Baye's theorem, we have a relationship between these two probabilities. For classification problems, we consider A as a prediction on old data, and B as new and with the use of the theorem the idea is to update the probability of A after seeing B. This update equals with the original probability of A, multiplied by the probability of B if A is true and divided by the probability of the new data B. In classification, the likelihood of an object belonging to class A and class B at the same time. This can be more complex if we have more classes or have just only two, A and B, where B means "not A". The algorithm begins by setting initial probabilities for A and B based on the statistical differences of each class. Then, each time we examine a feature, and the two probabilities get updated. The probability of each feature is determined during the training stage.

3.4.3.4 Decision Tree (foundational) [11]

Decision Tree (DT) is one of the earliest and most widely recognized ML algorithms. It functions by partitioning learning activities. The tree structure is constructed by iteratively partitioning the dataset into successively smaller subsets until all groups become clean or pure. In a typical DT, nodes are located at multiple levels, where the highest node, known as the root node, defines the beginning point. Each internal node applies to a test on an input feature or attribute. Depending on this test, the algorithm travels along a branch either to the left or right, moving on to its corresponding child node. The testing and branching continues until a leaf node has been reached and these leaf nodes, also known as terminal nodes, are the final decisions or classifications.

3.4.3.5 K-Nearest Neighbor (foundational) [11]

The K-Nearest Neighbors algorithm is a supervised learning approach. It works by computing the distance from new input data points to all other data points in the dataset. The 'K' in KNN stands for the number of closest neighbours that are used during the classification. The classification of a sample may change based on the selected value of 'K.' Finding the K closest neighbours means that we found the most similar records using common features, a process also known as distance calculation or similarity search. There are several methods to compute distance, such as Euclidean Distance (ED) and Manhattan Distance (MD). For classification problems, KNN gives a label to a new point by taking a majority vote amongst its closest neighbours (the label that appears more frequently in neighbours). Accuracy of the model is measured by comparing the predictions and estimations to the true labels of the available classes in the testing set.

3.4.3.6 Multi-Layer Perceptron (Popular as structural element for today) [11]

A Multi-Layer Perceptron, also known as an MLP, is a type of model trains functions through a dataset, is based on how a human brain how the brain evaluates and processes information. An MLP has three distinctive layers: an input layer, a single or multiple hidden layers, and an output layer. The input layer is the one that receives the feature values, which are then processed by the hidden layers, placed in the middle of input and output layer. In the hidden layers there are artificial neurons with weights, the algorithm calculates a weighted input sum, as $\sum(x_i w_i)$. These weights dictate how much each input impacts the output. The output from these weights then goes into an activation function, which scales the input into a bounded range, commonly [0, 1] or [−1, 1]. The result of the network is computed at the output layer by the multiplication of the total number of hidden layer neuron values with their assigned weights. MPL is also known as a feed-forward neural network (FNN), since data travels from input to output. It is a well-known machine learning algorithm that inspires a lot of neural networks.

3.4.3.7 CNN Convolutional Neural Network (Trending) [9]

With CNNs we can achieve machine learning classification when the input is an image which is very useful in autonomous cars. In real-life applications, cars have several kinds of sensors, but for now we consider that CNN only receives images. CNN is inspired by the biological visual perception, where specific neurons in our visual system are responsible for different areas and characteristics of the image we receive, and it has three basic operations/layers, convolution, pooling and fully connected.
A diagram of a CNN structure is shown in figure 3.7.

Convolution layer is where the filters are, these filters have height and width and are those which allow the algorithm to work as a biological visual system. The convolution operation converts regions of the image to tiny arrays of values known as kernels. The system takes information from various regions using these kernels and operationalizes it into a feature map.

CNNs can identify objects in images through this feature extraction process. Feature maps are created from each of these thumbnails, resulting from the convolution operation, as seen in figure 3.6, where the original image from the image data becomes smaller images with the same convolution. The parts of every thumbnail are processed by applying the same filter, enabling CNNs to identify unique object features in an image from the differences that these thumbnails have between them. These thumbnails are organized in a kernel.

The pooling layer performs the sub-sampling operation, that reduces the input spatially and hence reduces the parameters. The down-sampling operation decreases the space of feature maps produced by convolution, but without losing significant information. The output of these feature maps goes through a deep neural network in small arrays, then is transformed into a one-dimensional array. The output of the network finally comes in the form of a classification probability of each image, which is done from fully connected layers.

Backpropagation is a method that is used as a form of error correction, since it helps to minimize errors from earlier stages. It is a supervised learning algorithm with multi-layered perceptrons. The output error, computed by comparing the output obtained through forward propagation with its actual value, is employed to update weights in backpropagation. The process iterates until the rate of errors falls to a desired value.

Figure 3.6: The original image becomes smaller images with the same convolution

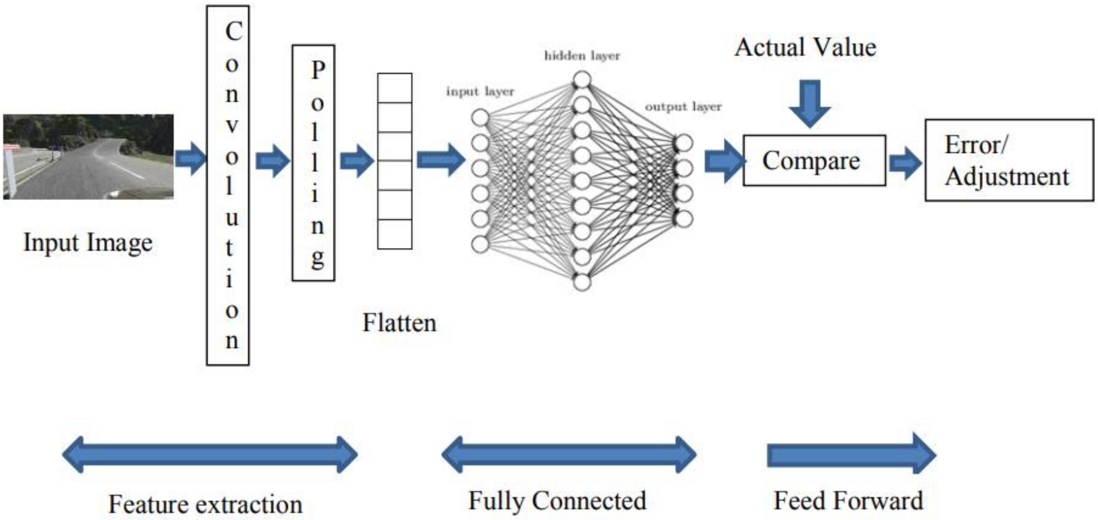

Figure 3.7: CNN architecture

3.5 Closing

The full potential of autonomous vehicle perception is realized when data from multiple sensors (e.g. cameras, LIDAR, radar) is combined with the appropriate techniques to create a unified, detailed understanding of the environment. As shown in Figure 3.8, this fusion enables object detection, 3D spatial localization, and semantic scene segmentation to work together, providing the vehicle with the contextual awareness needed for safe navigation in complex urban settings.

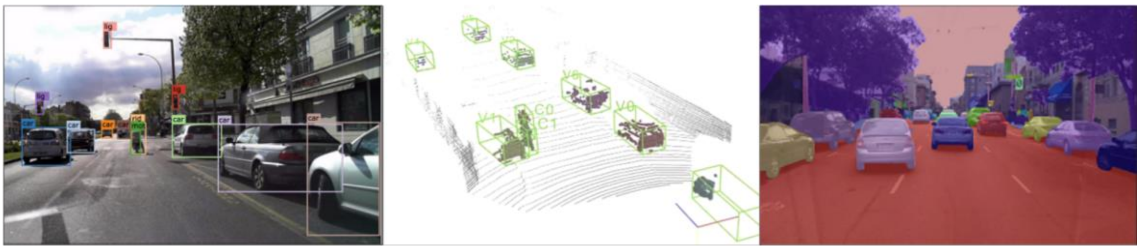

Figure 3.8: Combined perception outputs showing object detection, LIDAR-based 3D bounding, and semantic segmentation [14]

Section 4: Decision making in autonomous vehicles

4.1 Overview

For autonomous vehicles to truly drive independently, they must make decisions like real people. To do that, they rely on a layered architecture to navigate safely and effectively through urban environments. This hierarchy also lets these machines make decisions in a modular, interpretable, and computationally manageable manner. It includes four major components, route planning, behavioral decision making, motion planning, and vehicle control.

At the highest level is Route Planning, which with simple words, determines an optimal road sequence to travel to the destination. The real-life road network is illustrated as a directed graph where edge weights correspond to travel expenses (travel time, distance, or fuel). Traditional algorithms like Dijkstra’s or A* algorithms are inefficient for these applications since the size of these maps is too large for them to handle. Preprocessed hierarchical routing techniques are used for modern solutions like Avs, because they can calculate continent-scale routes in milliseconds, making them highly suitable for real-time autonomous applications.

Next is the Behavioral Decision Making layer, whose job is to understand real-time traffic context, such as intersections, pedestrians, and surrounding vehicles and then choose the appropriate high-level maneuvers. These would be lane following, stopping at intersections, overtaking and more. This is where finite state machines may be used for modelling transitions between behaviours based on environmental conditions that the sensors captured. Probabilistic models like Markov Decision Processes (MDPs) and partially observable MDPs (POMDPs) are also used to handle the uncertainty of the behaviour of other agents on the roads.

Third in the hierarchy is Motion Planning, which generates a feasible trajectory or path based on the dynamic and kinematic limits of the vehicle and it does that considering a predefined driving behaviour profile. This path should be collision-free, guarantee passenger comfort, and respect road regulations. Since this is computationally complex, approximate methods such as variational optimization, graph-based search, and sampling-based approaches are typically employed.

At the lowest of the hierarchy is Vehicle Control, which is the one that executes the planned trajectory using the feedback from controllers that compute steering, throttle and breke commands. This level corrects model uncertainties as well as external disturbances to maintain tracking accuracy. It is essential that the controllers used are robust, low-latency and capable of remaining stable even in extreme conditions.

All the information in the following section is from [12].

4.2 Planning and Control Models

How effective planning and control is of AVs depends on the model that is used to predict or simulate the motion. There should be a balance between fidelity and computational efficiency. The two categories most used to achieve the above, are kinematic models that perform well in low-speed maneuvers and dynamic models that include inertial effects and offer more accurate behaviour under higher accelerations.

4.2.1 Kinematic Single-Track Model

This model is also known as the bicycle model and is a simplified representation of a vehicle. If we consider it as two-wheeled system (one front wheel and one rear wheel aligned with the centerline of the vehicle), it operates under the no-sleep assumption, which means there is no lateral slipping when the tires roll, but they can rotate freely about their axes of rotation. To model steering, the front wheel has an added degree of freedom where it is allowed to rotate about an axis normal to the plane of motion.

With the visual representation of figure 4.3, we can see a quick overview of the algorithm. The location of the rear and front wheels can be denoted from the vectors p_r and p_f in a stationary or inertial coordinate system with basis vectors (e^x, e^y, e^z). The heading θ is an angle that describes the direction that the vehicle is facing. The motion of the rear wheel along the e^x-direction is x_r. Similarly, for e^y-direction is y_r. The forward speed is v_r. In terms of the scalar quantities xr, yr, and θ, the differential constraint is figure 4.1 and alternatively, considering speed V_f the differential constraint is figure 4.2.

$$\begin{aligned}\dot{x}_r &= v_r\cos(\theta),\\ \dot{y}_r &= v_r\sin(\theta),\\ \dot{\theta} &= \tfrac{v_r}{l}\tan(\delta).\end{aligned}$$

Figure 4.1

$$\begin{aligned}\dot{x}_f &= v_f\cos(\theta+\delta),\\ \dot{y}_f &= v_f\sin(\theta+\delta),\\ \dot{\theta} &= \tfrac{v_f}{l}\sin(\delta),\end{aligned}$$

Figure 4.2

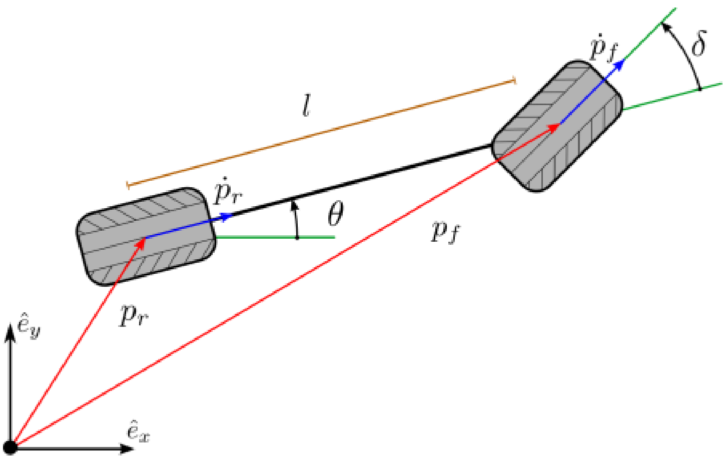

Figure 4.3: High-level representation of the model’s logic

4.2.2 Inertial Effects

When the vehicles reach higher speeds or do aggressive maneuvers, the no-slip assumption from before does not hold. For modelling these scenarios, dynamic inertial models are used. We use these to target basic momentum principles, the main of them being that the acceleration is proportional to the force generated by the ground on the tires.

Let P_c denote the position of the vehicle’s center of mass. The motion is governed by figure 4.4.

$$\begin{aligned}m\ddot{p}_c &= F_f + F_r,\\ I_{zz}\ddot{\theta} &= (p_c - p_f)\times F_f + (p_c - p_r)\times F_r\text{F}\end{aligned}$$

Figure 4.4

Where F_f and F_r are the forces at the front and rear wheels, m is the mass of the vehicle and I_{zz} is the moment of inertia(inactivity).

The slip dynamics are modeled by relating the tire forces to slip velocities and slip angles. A widely used model for tire-ground interaction is “Pacejka”, which expresses the traction force as a function of slip that has nonlinear characteristics, which produce a friction ellipse. This limits the maximum lateral and longitudinal force that can be generated emphasizing the tradeoff between steering and acceleration.

These dynamic models really shine in applications like high-speed driving, emergency maneuvering and stability control systems (e.g. ABS, ESC). While these models improve accuracy, they significantly increase computational complexity, especially for planning and real-time control.

4.3 Motion Planning

Autonomous vehicle motion planning is the computation of an optimal, feasible, and safe path or trajectory that directs the vehicle's state from where it is to where it needs to be. The process needs to consider both environmental geometry as well as vehicle dynamic constraints, traffic regulations, and interaction and coordination with dynamic agents. The motion plan is used as input by the vehicle's control system.

There exist two types of motion planning approaches: path planning, where the geometric path is calculated neglecting temporal constraints, and trajectory planning, where temporal evolution of the vehicle is considered. Autonomous driving systems, due to urban complexity and uncertainties, utilize optimization, graph-based, and sampling-based planning in combination.

4.3.1 Path Planning

Path planning seeks to discover a smooth curve in the vehicle's configuration space between the start and goal states such that obstacles are avoided and kinematic constraints are satisfied. The output path is a geometric curve, typically represented as a continuous function $\sigma(\alpha)$:[0,1]$\rightarrow X$, where X is the vehicle's configuration space (e.g., position and orientation in 2D or 3D space).

In contrast to trajectory planning, path planning does not consider timing execution and instead, it gives a geometric path that can be followed at a later stage when a velocity profile or delegating temporal consideration is assigned to the vehicle's controller.

4.3.1.1 Problem Formulation

The optimal path planning problem is defined over a tuple
$(X_{\text{free}}, x_{\text{init}}, X_{\text{goal}}, D, J)$, where:

- X_{free} is the free configuration space, representing all configurations that do not result in collisions.
- $X_{\text{init}} \in X$ is the initial configuration.
- $X_{\text{goal}} \subset X$ is the goal region.
- $D(x,x',x'',\ldots)$ defines differential constraints (e.g., bounds on curvature, smoothness).
- $J(\sigma)$ is the cost functional, such as path length or energy.

The task is to find a path $\sigma*(\alpha)$ such that:

- $\sigma(0) = x_{\text{init}}$,
- $\sigma(1) \in X_{\text{goal}}$,
- $\sigma(\alpha) \in X_{\text{free}}$, $\forall\alpha \in [0,1]$,
- $D(\sigma(\alpha),\sigma'(\alpha),\ldots)$ is satisfied $\forall\alpha$,
- and $J(\sigma)$ is minimized.

4.3.1.2 Complexity

The optimal planning problem mentioned above is PSPACE-hard, that means that it is computationally intractable in the worst case, especially when differential constraints are considered (nonholonomic systems). Although in some special cases the solution would be polynomial-time, even in simplified settings (e.g., 2D environments without curvature constraints) the problem might still be hard to solve precisely. For instance:

- Finding the shortest path for a holonomic vehicle (with no constraints) in a 2D polygonal environment can be done in $O(n^2)$ time using visibility graphs.
- Finding the shortest curvature-bounded path for a car-like robot in a polygonal environment is NP-hard.
- Exact solutions for obstacle-free curvature-bounded paths exist, e.g., Dubins curves and Reeds–Shepp paths.

4.3.1.3 Numerical Methods

Because of the complexity of exact solutions, realistic systems depend on approximate numerical schemes, typically categorized as:

- Variational approaches: Parameterize the path as a finite-dimensional vector and then apply nonlinear optimization. They are fast, however sometimes they converge to local optima only and require an appropriate initial guess.

- Graph search approaches: Discretize the configuration space as a graph, with vertices as discrete configurations and edges as valid transitions. Dijkstra and A* algorithms can then be employed. Such approaches can escape local minima but have limitations due to discretization granularity.

- Incremental search (sampling method): Techniques such as RRT (Rapidly-exploring Random Trees) or PRM (Probabilistic Roadmaps) sample the space to build a tree of graph of reachable configurations. They are scalable to high dimensions and probabilistically complete.

Each of the above comes with trade-offs in optimality, completeness, computational cost, and applicability to differential constraints. Variational approaches, for example, can fine-tune paths discovered by sampling-based planners, whereas strong guarantees can be provided by graph-based algorithms when resolution is high enough.

4.3.1.4 Constraints

In path planning for autonomous vehicles, typical constraints can be:

- Nonholonomic constraints (e.g., cars cannot move sideways).
- Curvature bounds (e.g., minimum turning radius).
- Obstacle avoidance, both static and dynamic.
- Road boundary constraints, such as staying within a lane.

4.3.1.5 Motion Planning Closing

Modern autonomous agents tend to combine path planning algorithms to try and benefit from the strength of all. An initial coarse plan can be calculated with graph search and then fine-tuned by a local variational planner based on real-time sensor measurements. This hybrid approach allows for efficient and safe navigation in dynamic and unpredictable urban environments.

4.3.2 Trajectory Planning

Trajectory planning acts as an extension of path planning by incorporating time as a dimension to generate a time-parameterized sequence of configuration for the vehicle to follow. A major difference to path planning is that it considers the changes of the vehicle over time, so it can reason about vehicle dynamics, dynamic obstacles, and timing-dependent constraints, whereas path planning yields a spatial path without temporal constraints.

4.3.2.1 Problem Formulation

An optimal trajectory planning problem is defined over a tuple like path planning:
$(X_{\text{free}}, x_{\text{init}}, X_{\text{goal}}, D, J, \text{T})$, where:

- X_{free} is the time-dependent free space (e.g., due to dynamic obstacles).
- $X_{\text{init}} \in X$ is the initial configuration.

- $X_{goal} \subset X$ is the goal region at final time T.
- $D(x,x',x'',\ldots)$ defines differential constraints (e.g., due to dynamic obstacles).
- $J(\pi)$ is the cost functional (e.g., time, energy, comfort).
- T is the fixed planning horizon

The task is to find an optimal trajectory $\pi^*(t)$ with:

- $\pi(0) = X_{init}$,
- $\pi(t) = X_{goal}$,
- $\pi(t) \in X_{free}(t)\ \forall t \in [0,T]$,
- $D(\pi(t),\pi'(t),\pi''(t),\ldots)D$ holds $\forall t$,
- $J(\pi)$ is minimized

4.3.2.2 Complexity

The complexity of trajectory planning is at the same level as path planning and generally even more so. With that being said, planning collision-free trajectories for a holonomic robot in a dynamic 2D environment is NP-hard. Robot planning, in the presence of moving and rotating 3D obstacles for robots with 2 or more degrees of freedom, is PSPACE-hard. These results clearly show that when dynamic obstacles and time constraints are involved, the complexity of the problem increases significantly.

4.3.2.3 Numerical Approaches

In real-world applications of automated driving systems, planning of trajectories is typically addressed through numerical optimization.

The optimization includes, that the path is segmented into a finite number of control points. An objective function encodes tracking accuracy, smoothness, comfort and safety. Lastly, there are constraints about collision avoidance, speed limits, bounds on curvature, and system dynamics.
These are then resolved by techniques such as nonlinear programming (NLP), quadratic programming (QP), or gradient-based optimization.

Alternatively, sampling-based methods like RRT* can be modified to accommodate trajectory constraints by changing the state representation and having time-aware collision checks.

4.3.3 Variational Methods

Variational methods are methods that treat the motion planning problem as an optimization problem over functions. It specifically tries to find a trajectory that minimizes a given functional cost and satisfies the systems dynamics and constraints at the same time. We can refer to that as trajectory optimization.

These approaches are especially effective in optimizing trajectories to satisfy requirements such as smoothness, control limits, and dynamic feasibility. Additionally, many companies prefer to use them in autonomous vehicle systems because they are very flexible and fast when working in known or structured environments.

4.3.3.1 Problem Formulation

Given a trajectory (like the one we mentioned above) variational methods attempt to solve [Figure 4.5]:

$$\underset{\pi \in \Pi(\mathcal{X}, T)}{\arg\min} \quad J(\pi)$$

Figure 4.5

subject to:

- Boundary conditions: $\pi(0) = x_{init}$, $\pi(T) \in X_{goal}$
- Equality constraints: $f(\pi(t), \pi'(t), \ldots) = 0, \forall t \in [0,T]$
- Inequality constraints: $g(\pi(t), \pi'(t), \ldots) \leq 0 \ \forall t \in [0,T]$
- $J(\pi)$ may represent travel time, path length, control effort, or a combination of them.

4.3.3.2 Constraint Handling

As said above these methods are used to handle constraints at the same time as optimization. Two of those are the following:

- Penalty methods that add a penalty term to the cost function for constraint violations [Figure 4.6]:

$$\tilde{J}(\pi) = J(\pi) + \frac{1}{\varepsilon}\int_0^T \Big[\|f(\pi,\pi',\ldots)\|^2 + \\ \|\max(0, g(\pi,\pi',\ldots))\|^2\Big]\,\mathrm{d}t\,.$$

Figure 4.6

- Barrier functions that impose infinite cost when constraints are violated, encouraging feasibility from the start [Figure 4.7]:

$$\tilde{J}(\pi) = J(\pi) + \varepsilon\int_0^T h(\pi(t))\mathrm{d}t\,,$$

where the barrier function satisfies $g(\pi) < 0 \Rightarrow h(\pi) < \infty$, $g(\pi) \geq 0 \Rightarrow h(\pi) = \infty$, and $\lim_{g(\pi)\to 0}\{h(\pi)\} = \infty$.

Figure 4.7

4.3.3.3 Discretization Approaches

To convert the infinite-dimensional problem into a tractable optimization problem, the trajectory π(t) is divided into two other methods, Direct and Indirect.

Direct methods approximate the trajectory using a finite-dimensional basis, such as figure 4.8:

$$\pi(t) \approx \tilde{\pi}(t) = \sum_{i=1}^{N} \pi_i \phi_i(t),$$

Figure 4.8

Some examples are Collocation methods that evaluate dynamics and constraints at selected time points. Euler and Runge-Kutta that integrators approximate the evolution of system dynamics. Finally, Pseudospectral methods that use orthogonal polynomials (e.g., Chebyshev, Legendre) to approximate trajectories with high accuracy.

These approaches result in a nonlinear program (NLP) over coefficients π_i and are solved using standard optimization techniques.

Indirect Methods are based on Pontryagin's Minimum Principle, and derive necessary conditions for optimality, that leads to a two-point boundary value problem involving system states and co-states. Lastly, even though these methods offer better insight into optimality, they are numerically more fragile and less commonly used in real-time systems.

4.3.3.4 Strengths and Weaknesses

Starting with the advantages of variational methods, worth mentioning is that they have fast stabilization to locally optimal solutions. In addition, they are well-suited when we want to fine-tune trajectory smoothness and control constraints. Also, those methods can be directly incorporated with vehicle dynamics and environmental constraints simultaneously.

On the other hand, a limitation of them is the fact that they converge to local minima quite often. Additionally, they are sensitive to initial guesses and most of the time require a precomputed efficient path. Lastly, there is no global optima unless they are combined with other methods (e.g. graph search).

4.3.4 Graph Search Methods

Graph search algorithms are a category of motion planning algorithms that discretize the vehicle configuration space into a finite graph and apply search algorithms to this graph to compute feasible or optimal paths. These techniques the local minima problems by performing a global search in a discretized representation of space. Every node in the graph represents a vehicle configuration, and each edge represents a feasible motion segment combining two configurations. The vehicle constraints, such as nonholonomic nature and obstacle avoidance, must be satisfied by the motion segments.

4.3.4.1 Formulation

A graph $G = (V,E)$ if defined where:

$V \subset X$ is a set of discrete vehicle configurations (vertices), E is a set of edges (o_i,d_i,σ_i) where:

o_i : origin vertex,

d_i : destination vertex,

σ_i : path segment connecting o_i to d_i satisfying $\sigma_i(0) = o_i * \sigma_i(1) = d_i$ and all motion and collision constraints.

The feasible path is then found by joining the path segments of the chosen edges in the specified graph. The optimality of the resulting path is limited by the granularity of the graph and the design of motion primitives.

4.3.4.2 Lane Graph

Hand-crafted lane graphs are a reasonable approach for structured urban environments. Here, the edges represent a maneuver or a lane-following graph like the change of a lane or a turn at an intersection. Also, the graph is created with a combination of the map data (e.g. road network topology) and manual design. These graphs are effective for standard navigation scenarios. One of their advantages is that they are fast and efficient in planning for nominal conditions and integrate well with routing and behavioral layers. A drawback is that they are inflexible in the presence of unexpected obstacles (e.g., a disabled vehicle blocking a lane), requiring fallback to general-purpose planners.

4.3.4.3 Geometric Methods

These methods work by modelling obstacles and configuration space using polyhedral or polygonal bodies. They compute roadmaps that capture the topology of the free space and guarantee completeness, meaning they will return a path if it exists. Some examples include Visibility graphs that join all pairs of visible vertices between polygonal obstacles. Also, Voronoi Diagrams, which maximize clearance from obstacles and Cell Decompositions that partition space into manageable areas.

These approaches are effective in low-dimensional spaces like planar worlds with no differential constraints.

One drawback is that they have difficulty in dealing with differential constraints naturally, such as minimum turning distance, and they are also unsuitable to scale to kinodynamic and/or high-dimensional systems.

4.3.4.4 Sampling-based Methods

Sampling-based planners like Probabilistic Roadmaps (PRM) and Rapidly-exploring Random Trees (RRT) avoid explicit geometric modelling by sampling and collision-checking the configuration space.

A critical building block of such planners is the steering function, represented by steer (x, y), which calculates a feasible (but possibly incomplete) motion from configuration x towards y, respecting vehicle constraints at the same time. Straight-line paths, Dubins paths, and motion primitives are a few examples. Another critical component is the collision checker col_free(σ), which tests whether a motion segment σ lies in the free space X_{free}

There are several types of steering strategies. One is random steering, which applies random control inputs. Then, heuristic steering, this one chooses actions from a library of known motion primitives. Exact steering computes precise paths through known solutions such as Dubins or Reeds–Shepp. Lastly, optimal steering takes the best available connection in terms of cost.

Sampling-based planners build a roadmap or a tree incrementally by sampling new states and attempting to add them to existing nodes and finally adding edges when a connection is considered possible.

Some popular variations are PRM to construct a global roadmap for multi-query scenarios, PRM* to improve PRM with asymptotic optimality, RRT to build a tree rapidly exploring the space (single-query), RRT* to refine RRT by rewiring to improve path cost over time, and SST* to avoid the need for exact steering and maintain a sparse tree.

Some of these variations have theoretical properties such as probabilistic completeness, meaning it will find a path if one exists even if it must work for infinite time. Also, asymptotic optimality means that it will converge to the optimal path as the number of samples grows.

4.3.4.5 Strategies for Graph Search

Once a graph is constructed, either geometrically, manually, or via sampling, then standard graph search algorithms are used:

- Dijkstra's algorithm: Finds the shortest path to all nodes; guarantees optimality but can be slow.
- A*: Enhances Dijkstra with a heuristic that estimates the cost-to-go; guarantees optimality if the heuristic is admissible.
- Weighted A*: Inflates the heuristic to allow bounded suboptimality in exchange for faster computation.
- Anytime A*: Finds an initial suboptimal path quickly and refines it as time allows.
- ARA* (Anytime Repairing A*): Performs a series of searches with decreasing heuristic inflation, reusing prior computations.
- D* and D* Lite: Designed for real-time replanning in dynamic environments; they efficiently update paths when the environment changes.
- Any-angle planners (e.g., Theta*): Allow edge shortcuts in grid-based graphs to produce smoother paths.
- Field D*: Interpolates path costs to produce smooth trajectories directly from grid representations.

These strategies are crucial for urban driving scenarios where updates to the map or environment are frequent (e.g., dynamic obstacles, traffic updates).

4.3.4.6 Practical Deployments

Below, we will discuss industrial efforts that have successfully integrated motion planning algorithms into autonomous vehicles:

1. CMU's "BOSS": This is a vehicle that belongs to Carnegie Mellon University and was the winner of the DARPA Urban Challenge. Its planning system is a hybrid layered architecture. Firstly, for Nominal driving which used variational trajectory generation for lane following and navigation on structured roads. For complex maneuvers like parking or turning in tight areas "BOSS" used a 4D lattice graph planner with Anytime D* search. This 4D space included x,y,θ and velocity. This combination allowed Boss to achieve fast, smooth driving in open roads and precise maneuvering in constrained environments.
2. Stanford's "Junior": This is an execution of a Hybrid A* planner that combined a discrete search over motion primitives and continuous state-space updates to refine orientation and curvature. From the above Hybrid A* ensured that trajectories are dynamically feasible and also produce smooth transitions between maneuvers. In addition, it used curve interpolation to create paths from the discrete plan. This deployment demonstrated strong performance in both structured and unstructured road segments.
3. Victortango's "Odin": This is Virginia Tech's vehicle, which used a maneuver-based graph search approach. It defined a discrete set of maneuvers that were possible such us drive forward, stop and turn. Then it constructed a maneuver graph that was based on the vehicle's state and map geometry and finally it applied A* search to find optimal sequences of maneuvers. This allowed Odin to reason at a higher level about navigation, particularly for behavioral decisions at intersections and merges.

These implementations demonstrate how motion planning algorithms can be adapted to meet the real-world demands of autonomous driving systems.

4.4 Vehicle Control

Within the motion planning hierarchy of AVs, the vehicle control module is in the lowest layer. It is its responsibility to calculate real-time actuator commands, which will help the AV to follow a planned trajectory or path. This can be considered a challenging task because of the nonlinearity of vehicle dynamics, physical actuator limitations, and the uncertainty introduced by modelling approximations and environmental disturbances.

AV control strategies are usually developed in reference to the kinematic single-track model and are classified based on the objective of the model, this could be path stabilization (spatial reference only) or trajectory tracking (spatiotemporal reference). There are also higher fidelity approaches that consider constraints explicitly via predictive control or model adaptation via linear parameter-varying (LPV) control.

4.4.1 Kinematic Model Path Stabilization

The path stabilization includes the computation of control inputs in a way that the vehicle asymptotically converges to a predefined geometric path, independent of time. If (x,y,θ) is the vehicle configuration and $\sigma(a)$ is the reference path in configuration space, then the control objective is to create a law of feedback $\delta = \mu(x,y,\theta,\sigma)$, where δ is the angle of steering that minimizes the deviation between the actual and desired paths.

4.4.1.1 Common Strategies

1. The pure pursuit algorithm is a geometric method which works by selecting a look-ahead point on the reference path at fixed distance L from the current vehicle position. The control law's aim is to minimize the required curvature to steer the vehicle to a curved path that passes through the look-ahead point. This method is intuitive and computationally effective but has no formal guarantees of convergence. The performance depends on how the look-ahead distance L is chosen, which is usually adjusted by trial and error based on the vehicle's speed. The method works well at low speeds, but it can have trouble following the path accurately on sharp turns.
2. The Rear-Wheel Position Feedback is a method that models the vehicle the axle as a reference point and defines the lateral deviation and orientation error considering the path. To stabilize the vehicle, a nonlinear feedback controller is designed based on the errors mentioned before. This approach has formal convergence properties, and it is shown to be asymptotically stabilizing under

certain curvature and speed conditions. It is used for differential geometric path tracking formulation and is used extensively in theoretical analysis.

3. Front-wheel Position Feedback is an alternative to stabilize the front axle considering the path. This theory might be less common than the rear axle formulations, whereas these approaches are valid under specific curvature limitations and they might be better aligned with the physical location of the steering actuator. On the other hand, practical adoption is limited.

4.4.1.2 Kinematic Trajectory Tracking Control

Trajectory tracking control involves stabilizing the vehicle along a time-parametrized trajectory $\pi(t)$, which includes both spatial and temporal evolution. Unlike path tracking, it requires synchronization between the vehicle state and a moving reference. There are two analytical approaches for this problem that are worth mentioning:

1. Control Lyapunov Function (CLF) Based Design:
   This approach is about the construction of a scalar function $V(x,t)$ that is the tracking of the error state, and this function is always positive except at zero error and decreases over time as the system moves along its path. Also, $\delta(t)$ is a control law that is synthesized like $V (x,t) < 0$, which ensures the convergence of the tracking error. This approach almost guarantees stability, is robust to bounded disturbances, and it works with nonlinear systems such as car-like vehicles. However, it usually requires symbolic derivation and precise state estimation.
2. Output Feedback Linearization:
   This approach uses feedback transformations to cancel system nonlinearities and convert the system into a linear form with respect to outputs. Then, linear control methods (e.g., PID, pole placement) are employed for controller synthesis. The greatest strength of this is its ability to leverage classical linear control methods, however it requires accurate models and full state observability. Lastly, incorrect model structure or parameterization may reduce performance.

4.4.1.3 Predictive Control Approaches

Model Predictive Control (MPC) solves an open-loop optimal control problem every time step across a finite time horizon based on a model of the vehicle dynamics and constraints. At the start, the first control action is implemented, and the algorithm repeats in the next step in a receding horizon manner.

1. Unconstrained MPC with Kinematic Models: This version employs a kinematic vehicle model with no explicit constraints (aside from those embedded in the cost function). It minimizes a quadratic cost that penalizes tracking error and control effort. This formulation is manageable and has smoother control than geometric feedback laws. Nevertheless, it is not able to ensure satisfaction of constraints such as steering saturation or obstacle avoidance.
2. Path Tracking Controllers: These are intended to minimize lateral and heading errors in relation to a path in space, usually in Frenet coordinates. The controller outputs a speed profile and steering actions that maintain the vehicle close to the reference. These types of MPC formulation are specifically appropriate for road-following and lane-keeping tasks and offer good performance in structured environments.
3. Trajectory Tracking Controllers: In this approach, MPC minimizes deviation from a fully time-parameterized trajectory, incorporating velocity and acceleration profiles. The optimizer can also consider constraints such as acceleration bounds, actuator limits, and obstacle avoidance. This method offers high fidelity control but requires significant computational resources and a reliable prediction model. Its success in field deployments depends on solver efficiency and model accuracy.

4.5 Conclusion

Over the past 30 years, driverless vehicle technology has advanced rapidly, driven by improvements in hardware and major theoretical progress in motion planning and control. These vehicles are complex systems broken down into a hierarchy of decision-making problems, each feeding and leading into the next one. This modular approach allows researchers to apply specialized methods from different fields but integrating them requires careful attention to both semantic compatibility and computational efficiency.

For example, a powerful planning algorithm may require an equally advanced, resource-heavy controller, while simpler control methods might demand more detailed planning.

Chapter 5: Methodology

This chapter outlines the methodology employed in this research to investigate industry and public perspectives on autonomous vehicles. The study utilized two structured surveys: one focusing on public opinions regarding AI-powered self-driving cars and another targeting industry professionals working in autonomous vehicle development. This methodology section details the design, distribution, and analysis of these surveys.

5.1 Research Design

The research was structured around two distinct questionnaires, each designed to collect specific insights. The Industry Survey targets professionals in the autonomous vehicle sector, gathering insights on key challenges such as data collection, computational limitations, model interpretability, regulatory concerns, and emerging industry trends. The Public Opinion Survey was built to examined general perceptions, trust levels, and ethical concerns of the public regarding AI-driven self-driving cars. This approach facilitated a multi-faceted analysis of the challenges, advancements, and societal impacts of autonomous vehicles.

5.1.1 Industry Survey

The industry survey was designed to capture insights from professionals involved in autonomous vehicle development. The key areas of focus include any significant challenges that a company may face in AI development, which machine learning model is most used by them, methods used to maintain and verify the quality of training data and also some questions target their opinion as professionals on various aspects, like ethical situations. These insights provide a valuable perspective on the technical and ethical considerations shaping the future of autonomous vehicle technology.

5.1.2 Public Opinion Survey

The public opinion survey was designed to assess general attitudes towards AI-powered autonomous vehicles, focusing on both individuals familiar with autonomous vehicle technology and those with little to no prior knowledge. The survey explored different perceptions of AI’s ability to handle complex driving scenarios, trust in AI-driven vehicles compared to human drivers and asked what would increase their trust. In addition, it examined their opinions and suggestions regarding safety and ethics. With this I ensured a well-rounded understanding of public perspective and the factors influencing trust in autonomous vehicle technology.

5.2 Sampling and Distribution

A total of 70 different automotive companies all over the world were selected for the Industry Insights survey. It was given priority to the companies with high AI achievements like Tesla, but the survey was sent to other automotive companies as well. The Public Opinion survey was sent to many people to ensure broader perspectives. The Industry Insights survey was distributed by emails found in their official websites and when the email was not presented the companies gave me the option to fill an online form for my request. I also went directly to the members of a company in Cyprus (BMW) and hand them a device to answer the questions. The survey for Public Opinions was distributed mainly by text messages and on some occasions with email.
Respondents were assured of the anonymity of their responses to encourage honest feedback.

5.3 Data Collection and Analysis

Responses were collected over a three-month period, with follow-up reminders sent to them once a month to increase participation rates. The data was analysed using both numerical and descriptive methods. Numerical analysis helped identify trends, patterns, and differences in opinions about autonomous vehicle development, while descriptive analysis focused on open-ended responses to better understand ethical concerns, industry expectations, and trust in AI systems.

5.4 Ethical Considerations

This study was carried out in accordance with the ethical guidelines set by the University of Cyprus, guaranteeing the confidentiality and anonymity of all participants. The responder’s email addresses associated with the questionnaire were not disclosed to me and the contact details of the respondents were obtained from publicly accessible websites.

5.5 Limitations

While this study provides valuable insights, certain limitations exist:
The sample may not fully represent the entire industry or global public opinion. Also, self-reported data may introduce biases in perceptions of autonomous vehicle technology. Additionally, rapid technological advancements may render some findings outdated over time. Additionally, the survey may not have accounted for regional variations in industry insights, since each company’s progress can differ significantly based on geographic location, and access to technological advancements.

Chapter 6: Industry Insights on Autonomous Vehicle Development

6.1 Introduction

In this chapter I present findings from a survey sent to professionals in the autonomous vehicle (AV) field. The survey was mainly conducted to identify key challenges, technological trends, and future directions in AI-driven AV development. Furthermore, I explore some bottlenecks in AI deployment, ethical frameworks, collaboration dynamics, and anticipated breakthroughs.

6.2 Presentation of Survey Results

In the following analyses whatever is encapsulated in “” means that are direct words from the survey responses.

6.2.1 Question 1: What is the primary focus of your company in the autonomous vehicle sector?

Analysis:
This question was answered by six responders. Three of them focus on car’s safety. One of these three responses focuses on sensors or software fails “Ensuring vehicle safety when sensors or software fail.”. However the other two while focusing on safety they value people’s comfort, efficiency and luxury as well “Make cars safer, free up people's time, Improve mobility for everyone, and Transform the modern luxury experience for customers.”, “Our vision is to develop autonomous solutions that emphasize reliability, connectivity, and sustainability, while ensuring that the driving pleasure BMW is known for remains a core element, even in an autonomous future”. One answer quoted that their main goal is “Customer satisfaction”. There are also two more technical objectives, where one team’s goal is “enhancing real-time path planning algorithms” and the other’s team goal is “developing advanced perception systems using multi-sensor fusion”.

6.2.2 Question 2: What is your role within the company?

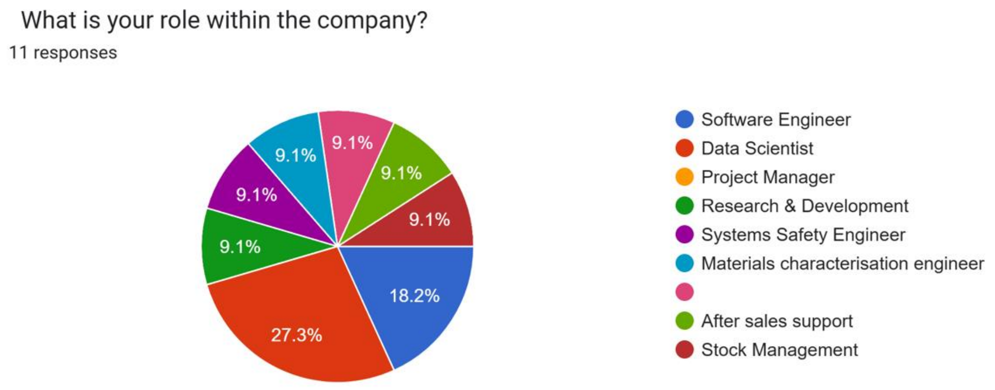

Analysis:

This question got 11 responses. The distribution of roles of the responders is an important factor to consider, since people's opinions could differ based on their skills. While software engineers (18.2%) and data scientists (27.3%) form the core of the answers following, opinions from other roles are important as well. Systems safety engineers (9.1%) focus on reducing risks and ensuring AI systems work correctly, while the presence of R&D specialists (9.1%) and materials characterization engineers (9.1%) reflects ongoing efforts to refine hardware and optimize autonomous system performance. Additionally, after-sales support (9.1%) and stock management (9.1%) indicate some views are based on sustainability.

6.2.3 Question 3: What are the primary bottlenecks your company faces in developing AI for autonomous vehicles?

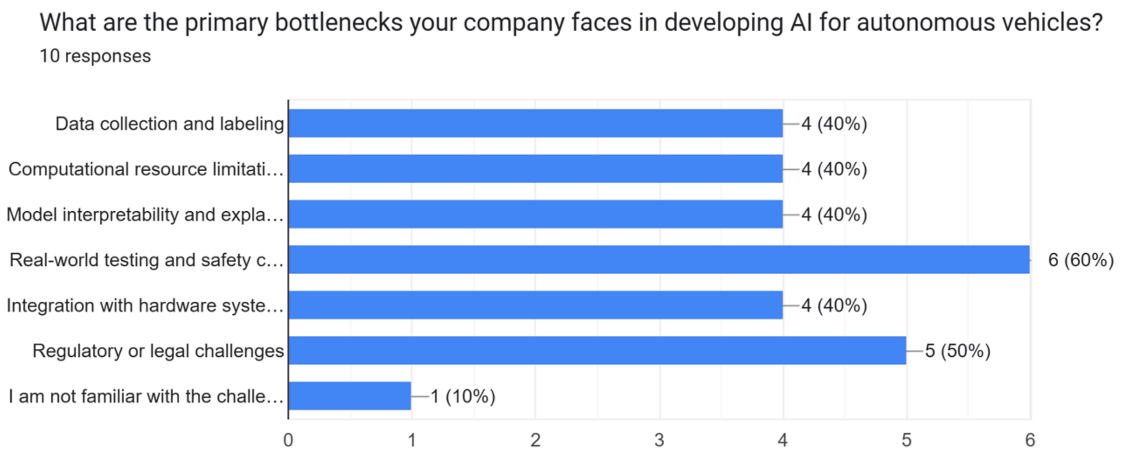

Analysis:

This question got 10 responses. Real-world testing (60%) is the most cited bottleneck, mainly because of the difficulty to replicate unpredictable scenarios in simulations. (e.g., extreme weather, erratic pedestrian behaviour). In addition, regulatory obstacles (50%) are caused due to inconsistent global standards and liability frameworks, which could delay deployment timelines. Data collection and labelling (40%) challenges arise from the fact that many data is very complex and difficult to manage, while Hardware integration (40%) issues are also present due to retrofitting AI systems into legacy vehicle architectures and ensuring compatibility with cost-effective sensors. Lastly, computational limitations (40%) relate to balancing real-time processing demands with energy efficiency, particularly for edge devices.

6.2.4 Question 4: Are there any challenges specific to your company's operations that are not widely discussed in the industry? Please elaborate.

Analysis:

This question got 10 responses. Two of them stated that are not sure or they don't face any. Other two of the responders specified that addressing unpredictable human behaviours and environmental changes are their challenges, this might be the case because simulations cannot produce correctly this kind of actions yet "Addressing unpredictable human behaviour in mixed autonomy environments (e.g., pedestrians jaywalking, aggressive drivers). Current simulations often oversimplify human unpredictability", "Our company faces difficulties in adapting AI models to rapidly changing environmental conditions". Additionally, one of the answers above and three other answers mentioned that balancing inexpensive computing resources (or computing power in general) with AI models and algorithms, as well as safety, is difficult as well "Another unique challenge is optimizing AI performance for cost-effective sensors while maintaining safety and accuracy.", "Balancing the cost of high-performance computing resources with the scalability required for global deployment.", "We often struggle with how to make sure AI models perform well on cheaper hardware without losing accuracy.", "Managing latency in real-time systems during high-speed manoeuvres, for example highway merges.". Adding to the previous issue, there is also a mention on integrating modern technologies into legacy vehicle architectures "Also, we face issues integrating AI systems into legacy vehicle architectures, which often requires custom solutions.". One expert stated that sensor degradation due to extreme weather conditions is a problem need to be solved “Handling sensor degradation in extreme weather conditions (e.g., heavy rain, snow).”. In addition, two responders are facing challenges on laws and driving rules, some of them related to the changes in different regions and others in general “Another problem is keeping up with changing laws and rules in different regions.”, “Another challenge is ensuring that our AI models can adapt to different countries' traffic rules without retraining them from scratch.”. In the previous answers data collection is mentioned where it stated that data for rare driving situations is not enough “One issue we face is collecting enough high-quality data in rare driving situations, like accidents or unusual weather.”. BMW expert’s answer was more general mentioning several problems stated above.

6.2.5 Question 5: How does your company address ethical considerations, such as decision-making in life-critical situations?

Analysis:

This question got 9 responses. Four responses highlighted the emphasis on simulation, and they are using it to “perfect” decisions in life-critical scenarios and prioritize safety “Mainly with the simulation of life-critical situations to refine decision-making algorithms, prioritizing overall passenger and pedestrian safety”, “using simulation to test ethical dilemmas”, “We test how our AI handles dangerous situations in simulations and make sure it’s designed to minimize harm”, “We employ scenario-driven testing to evaluate ethical decision-making in simulated environments”. In one of the answers above an expert mentioned the adoption of the “ASIL-D” which is a safety standard for decision-making frameworks as he/she said and another one also quoted “Transparency is key for this kind of technology”. Transparency was mentioned as well, very likely to be clear who’s at fault after an incident. Four responders, including one from above, specified that they are working with third party companies for these kinds of decisions “Implementing a transparent "black box" system that logs AI decisions during accidents, subject to third-party companies for analysis”, “We also talk to outside experts and communities to ensure fairness and transparency in how decisions are made”, “We also review ethical issues with experts...”, “We collaborate with experts to ensure our decision-making aligns with global ethical standards and focuses on accident prevention and transparency”. An approach in another answer is giving the drivers the option to intervene in critical scenarios for now, probably to judge a person, not a machine, since the main objective is to quote “make the autonomous side react like a human driver”. Lastly, another one said these are addressed “With sense of responsibility”.

6.2.6 Question 6: What types of machine learning models are most used in your autonomous systems?

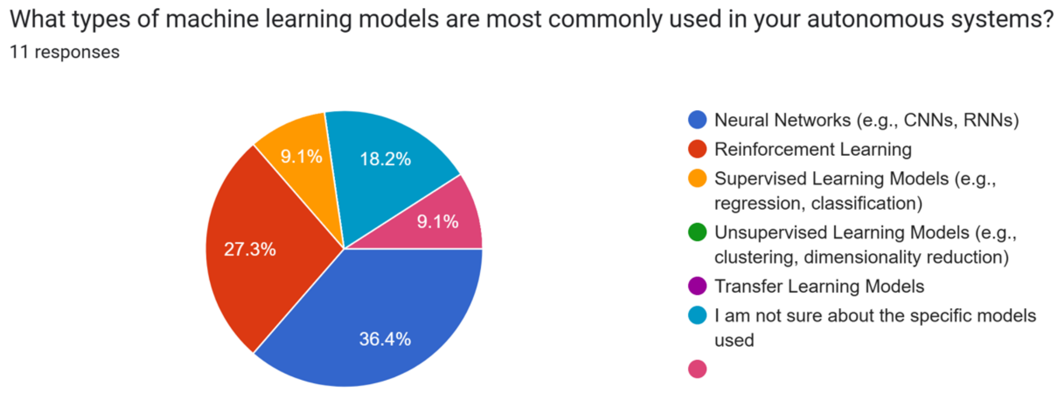

Analysis:

This question got 11 responses. For this question Neural Networks dominated (36.4%), apparently because they are used for perception tasks like object detection. Reinforcement learning (27.3%) is also highly used very likely for creating decision-making systems, enabling vehicles to learn optimal actions through trial-and-error in simulated environments. A minority use supervised learning (10%) and transfer learning models, I assume for simpler classification tasks, such as identifying traffic signs. All models play their role in the development of AVs, since each one is used for different purpose.

6.2.7 Question 7: How is the quality of your training data maintained and verified?

Analysis:

This question received 7 responses. Four specified that labelling is being done using AI models and tools, since is faster, however two of them also mentioned manual/human checks are done as well “use AI tools to label it faster”, “By manual and automated labelling processes”, “We use both AI and human checks to make sure the data is labelled correctly”. Also, two answers highlighted that these data is regularly updated and refined “We get and refine new data regularly”, “With regularly check”. Also, two experts (one from above) stated the partnership with third party companies, but for different reasons. One is to “ensure data integrity”, and the other is to label datasets. One of them mentioned also “Automated anomaly detection using pipelines”. Furthermore, a responder’s company “ensures the quality of training data by sourcing diverse” and that this data is “high-resolution datasets from real-world driving scenarios and simulations”. Continuing with the same company, their data is “rigorously annotated, validated, and continuously updated to reflect various conditions” and all of that is to ensure “accuracy and reliability”. Lastly, a specialist that is not sure how the data is maintained, thus he knows that “the data is collected from sensors of the test vehicles”.

6.2.8 Question 8: What strategies are employed to optimize computational efficiency in real-time AI operations?

Analysis:

This question received 9 responses, though three of them either couldn’t answer or were not sure. Many responses, exactly four, mention edge computing, that basically means simpler data is processed locally in the car (or really close to it) instead of server and supercomputers “We use simpler versions of our AI models (like pruning or compression) and process data directly on the car’s computer instead of relying on cloud servers”. Most probably these companies, including BMW use edge computing to achieve low latency, improved reliability and enhanced security. This is also mentioned as “Hybrid cloud-edge processing” and at the same answer it explains that “lightweight models run locally on vehicles, while complex computations offload to cloud servers”. Also in one of the responses above it says that a car used for this “is not a regular one, the interior is stripped and it's full of machines and computers”. Additionally, the above also specified that to achieve edge computing, they “use model compression” and more generally “parallel processing on dedicated AI accelerators to ensure real-time performance”. BMW and a company (again from above) also stated the optimization of algorithms “to reduce latency without sacrificing accuracy”. Lastly, two other replies mentioned the “co-design” between hardware and software and also the optimization of their “AI by reducing its size and only keeping the most important parts of the model”. They “use hardware that’s designed to process AI tasks faster, like GPUs and special chips.”

6.2.9 Question 9: Are there specific breakthroughs in AI that your team/company is excited about, or that you believe will significantly advance autonomous vehicle technology?

Analysis:

This question got 8 responses. Two responders are excited about new AI models that can process data from a lot of sources simultaneously, leading to more accurate decisions “AI models can process data from cameras, radars, and other sensors together, which makes the car’s decisions more accurate.”, “multimodal AI systems capable of integrating data from vision, LiDAR, and radar”. One of the experts above and one other expert are excited about the development of computational power, with one of them referring to quantum computing as well “general development of chips and computational that can handle more complex scenarios and complete training faster.”, “innovations in quantum computing, which could revolutionize route optimization and real-time decision-making.”. Two other specialist is thrilled with “multi-agent reinforcement learning” and “federated learning”, as they said, these will make vehicles cooperate without centralized control. One of them also quoted “This mimic human-like negotiation in lane changes and merges.”. Additionally, the same expert that is excited about federated learning also mentioned “neuromorphic computing” which helps to “enhance energy efficiency in AI hardware”. There is also a mention on “Vision transformers”, which as being said it improves “object detection accuracy in complex environments”. BMW specialist’s answer was more general as he/she excited in breakthroughs in “deep learning, sensor fusion, and advance reinforcement learning”. Lastly, a response did not state any specific breakthroughs, thus it highlighted that the progress of AI alone, will make “the decision making side and the predictive modelling side” easier.

6.2.10 Question 10: What are the most critical milestones in your opinion for achieving full autonomy?

Analysis:

This question got 9 responses. Three of the responses highlighted that “public trust” is mandatory for this achievement, probably supporting that society is the one to determine it. One of them stated that people must ensure that “are comfortable with the choices being made for us” and this will be done with “extensive learning”. The same one also quoted that “I'm not sure full autonomy will ever be feasible as it stands as people wont fully trust it.”. Two of the above mentioned “regulatory approval” as well and also “approval from governments”. The previous view can surely put together with the opinions of two other specialist that stated the build of “public infrastructure that supports autonomous vehicles” and also “Standardizing ethical decision-making frameworks globally (e.g., ISO/PAS 21448 for SOTIF)”. Continuing with the same way of thinking another response specified “Universal standards for edge-case scenario testing and global regulatory alignment.”. The remaining answers both mentioned “V2X communication protocols”, “perfecting sensor accuracy” and “developing robust AI decision-making systems”. Lastly, level 4 autonomy in urban areas is stated as well.

6.2.11 Question 11: In your opinion, what role will collaboration between car manufacturers, AI companies, and governments play in shaping the future of autonomous vehicles?

Analysis:

This question got 8 responses. More than half of the answers (5) highlighted that with governments participating in this collaboration will set clear rules, regulations and universal safety standards between our society and this technology, with other words, build an infrastructure for self-driving cars to enter our lives. Some of the above as stated as “establishing universal safety standards and ensuring interoperability across systems.”, “setting regulatory frameworks”, “solving legal and regional issues and creating a unified approach to safety and innovation.”, “…infrastructure needed for self-driving cars, like connected traffic lights and smart roads”. Also, one of the experts above mentioned that “car manufacturers and AI companies must work together to align technological innovation with these frameworks”. Another specialist supports that gove “Governments should fund public testing areas for self-driving cars, while companies work together on common safety standards.”, one example he/she mentioned is the creation of “virtual city models to test AI in realistic simulations.”. Additionally, a response stated the need of all the above working together, but that to happen they all need to “embrace autonomy”. Adding to that, this particular person wrote that “If half the cars on the roads are autonomous it makes the non-autonomous ones unpredictable so everyone will need to work together to gather and share real life data for this to work.”. Lastly, a more general response specified that this collaboration could bring “Shared open-source datasets” and “accelerate industry-wide innovation”.

6.2.12 Question 12: How do you see the role of AI in autonomous vehicle innovation evolving over the next 5-10 years?

Analysis:

Most responses highlight a significant growth in AI's prediction, learning, and efficiency over the next 5-10 years. Many believe that AI will have the ability to predict problems before they happen, improving safety and decision-making. Sensor fusion and real-time learning from vehicles are also expected with the development of technology. Some responses noted that with deep learning system would be less rule-based allowing AI to handle complex situations more naturally. Energy efficiency and regulatory compliance are also mentioned, with AI models becoming more optimized and explainable. Others mention integration with 5G, will improve vehicle communication and coordination. In conclusion, while some of the responses are sceptical about achieving full autonomy soon, most of them agree that AI will make autonomous systems smarter, safer, and more efficient, with full autonomy being the main goal.

6.2.13 Question 13: Are there any other insights you'd like to share regarding autonomous vehicle development?

Analysis:

Most responses do not provide additional insights, but two ideas are mentioned. Firstly, transparency is seen as a way to help people feel more comfortable with self-driving technology, such as openly reporting accidents and human interventions. Secondly, is mentioned that cybersecurity will play a key role to protect vehicles from hacking and data breaches. This is not only crucial for personal data protection, but also for protection from people that want to interfere with the cars to create accidents and harm other individuals.

6.3 Summary of Findings

This survey has brought to light several key findings regarding the professional and insight views on autonomous vehicles. Here is a summary of the findings:

1. **Primary Focus for Companies:** The responders were mostly focused on their costumer's safety and then their sustainability and comfort. Of course, there was an equal number of other experts that focused on more technical stuff.

2. **Technical Challenges:** The challenges that appeared the most in the survey were real-world conditions testing and simulation, the collection of data and labelling to ensure that is high-quality, and hardware integration and computational limitations. From all the above we can understand the complexity and difficulty of developing those machines.

3. **Machine Learning Models:** Neural Networks and Reinforcement Learning are the models selected the most, however from what I comprehend all ML models are important and have their own role for the development of autonomy.

4. **Ethical and Regulatory Needs:** Companies mostly target those through ethical dilemmas simulations, making transparent frameworks and working with third party companies.

5. **Collaboration Dynamics:** Cross-industry partnerships appear to be very important for the development of a complete infrastructure that supports AVs. Additionally, collaboration is mentioned for data maintenance/evaluation and also for ethical and regulatory problems.

6. **Computational Efficiency:** Edge computing is the technique mentioned the most here, that leads to the need for simpler models (algorithms) to run locally on the cars. More complex computations happen on supercomputers.

7. **Key Milestones:** Regulatory approval and public trust are essential achievements for full AV adoption. After that, a public infrastructure with global framework would be the next milestone to chase.

8. **Societal Acceptance:** To achieve public trust is essential for companies to be transparent with people, offer automation gradually, and try their best to prove AVs reliability.

These insights highlight the multifaceted challenges and opportunities in AV development, emphasizing the need for balanced progress across technology, regulation, and societal engagement to realize the promise of autonomous mobility.

Chapter 7: Public Opinion Survey

7.1 Introduction

In this chapter, I present the key findings from the survey conducted for the general public to understand their perspectives on autonomous vehicles. The analysis explores respondent's awareness of AI in self-driving cars, their trust in the technology, safety concerns, ethical considerations, and willingness to adopt autonomous vehicles in Cyprus. Additionally, the findings highlight the factors that influence public confidence, the challenges that need to be addressed, and the potential impact of AI-driven transportation on society.

7.2 Presentation of Survey Results

In the following analyses whatever is encapsulated in "" means that are direct words from the survey responses.

7.2.1 Question 1: Age

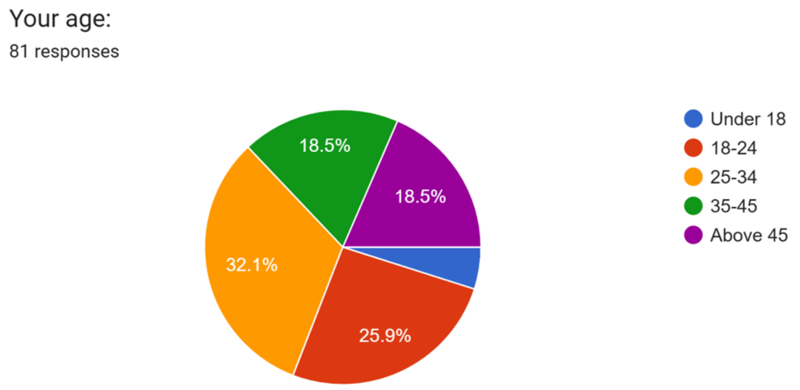

Analysis:

The largest group was 25-34 years old (32.1%), followed by 18-24 (25.9%). Both 35-45 and Above 45 had 18.5% each, while Under 18 had the fewest responses. This indicates that the majority of respondents are young adults, with a balanced representation among older age groups.

7.2.2 Question 2: Gender

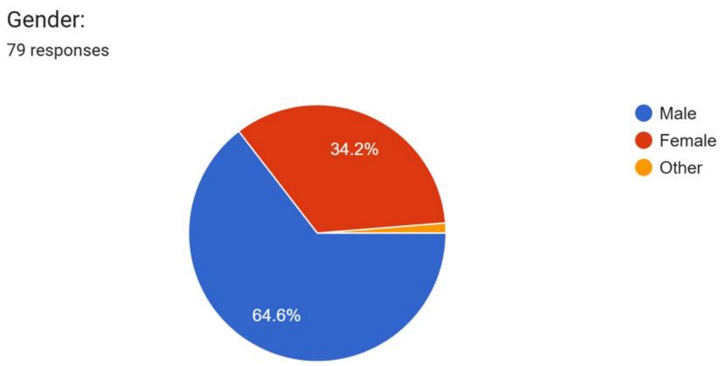

Analysis:

The majority identified as Male (64.6%), followed by Female (34.2%), with a small percentage selecting Other. This indicates a significantly higher male representation among respondents.

7.2.3 Question 3: Education Level

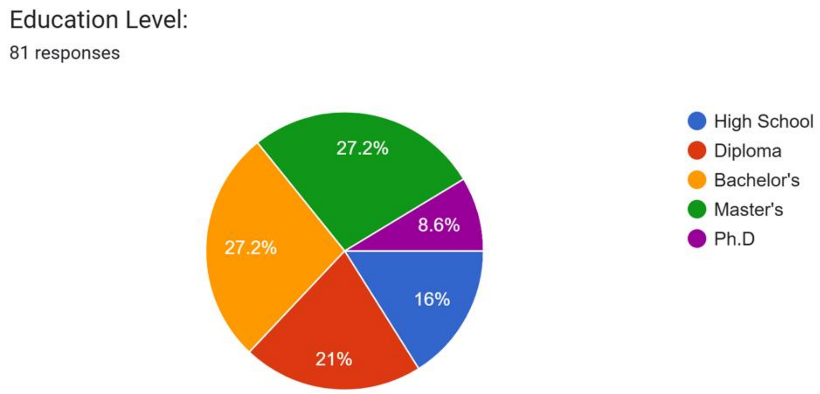

Analysis:

The highest percentage of respondents hold a Bachelor’s (27.2%) or master’s degree (27.2%), followed by those with a Diploma (21%). High school graduates make up 16%, while Ph.D. holders represent the smallest group at 8.6%. This suggests a well-educated respondent base, with a significant portion having higher education degrees.

7.2.4 Question 4: Country

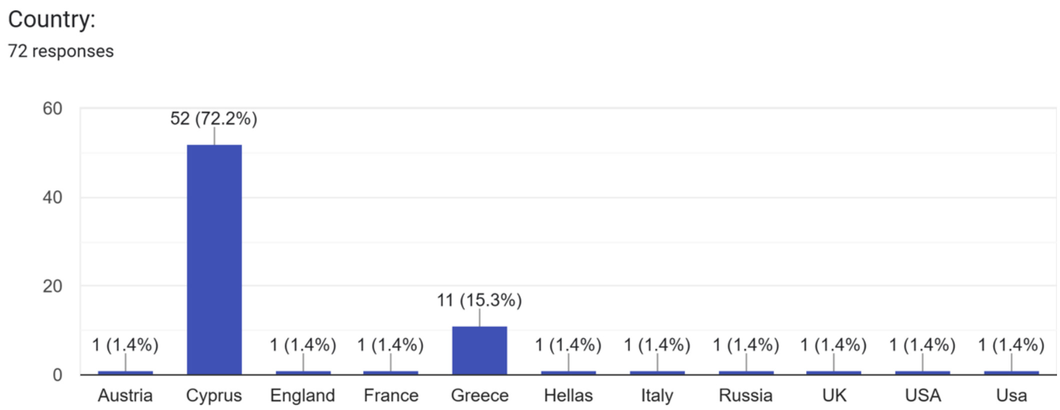

Analysis:

About the responder's country of origin, 72.2% of the respondents indicated they are from Cyprus. Another 15.3% of the responses were from Greece, while others come from various locations in the world. This highlights that most of the following opinions come mainly from Cypriot and Greek cultures, while a small portion of them are from other countries.

7.2.5 Question 5: How familiar are you with the use of AI algorithms in automated (self-driving) vehicles?

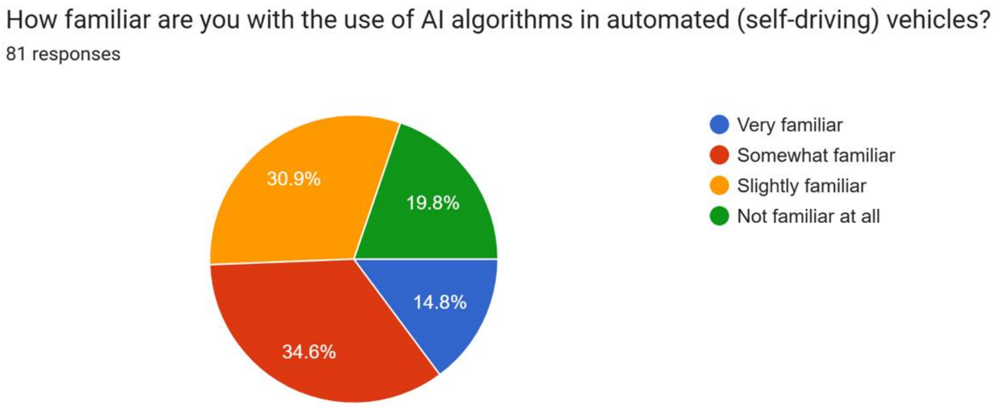

Analysis:

Regarding familiarity with AI in self-driving cars. The largest group (34.6%) is somewhat familiar, followed by slightly familiar (30.9%). Additionally, (19.8%) are not familiar at all, while only (14.8%) consider themselves very familiar. This suggests that while most of the responders have awareness about AVs, a significant portion lacks deep familiarity with AI in autonomous vehicles.

7.2.6 Question 6: Do you believe that AI algorithms can handle complex driving scenarios (e.g., heavy traffic, adverse weather) effectively?

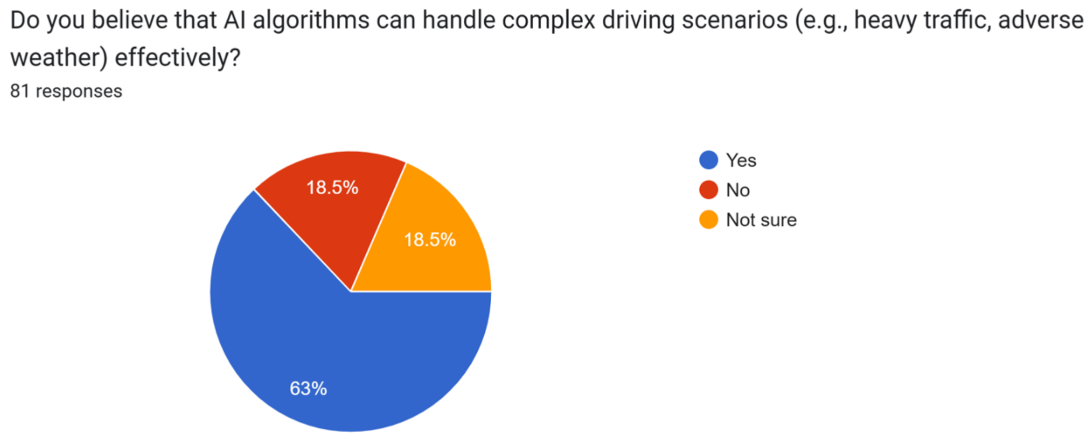

Analysis:

The public opinions about autonomous vehicles ability to handle complex driving scenarios is generally positive. The majority (63%) believe AI can manage challenges like heavy traffic and adverse weather effectively. However, 18.5% do not trust AI in these situations, while another 18.5% are uncertain. In conclusion there is a generally positive outlook, though some scepticism remains.

7.2.7 Question 7: What features would make you feel more comfortable with AI technology in automated vehicles?

Analysis:

This answer got 43 responses. Ten of the responses highlighted that they will feel more comfortable if the driver has the ability to observe in real time what the car is going to do and, in some cases, to understand the way that the machine thinks or works ("Something to make me able to predict and avoid an accident"). Half of those answers mentioned that this can be done using "voice explanation about the decision", while others stated, "augmented reality on the windshield presenting the next move". Augmented reality might be a better solution at the end of the day since voice explanations may interrupt conversations or radio music. The next most wanted feature, supported by six people, is a "button" or a way in general to switch quickly from autonomy to manual driving at any time. This obviously makes the drivers confident that they can "take control anytime" and their safety depends on their own actions. Additionally, three of the responders pointed out that they want "Specific sensors for heavy rain, fog or other low visibility situations", so they can "guarantee 100% safety in any condition", and also be sure about "Surroundings checking (e.g. pedestrians, kids etc.)". In additionally, there are two responses that want features in cat to specifically "make sure that other drivers that use it (the autonomous car) pay attention on the road and are not distracted". In my opinion these features are very crucial, since until full autonomy (level 5) drivers must be ready to intercept at any time. Other responses stated some quality of life features like "Suggest me a move and wait for me to approve it, e.g. change lanes in the highway" or "Being able to choose between modes of driving style and safety", "Automated parking" and also about the ride that "should feel smooth and natural, not weird or jerky.". Moreover, there were five people that did not mentioned features, thus specified that they want see "the way it was developed. What scenarios mainly used to be trained" or "See them navigate through complex and narrow roads". These people want "transparency" and "honest reports" on the machines that they will use every day. I believe this is a way to grow their trust for autonomous cars. Lastly, there was a more technical response that stated "Highspeed communication between other vehicles on the road to verify actions and gather information".

7.2.8 Question 8: Do you think AI-driven autonomous vehicles are safer than human-driven cars?

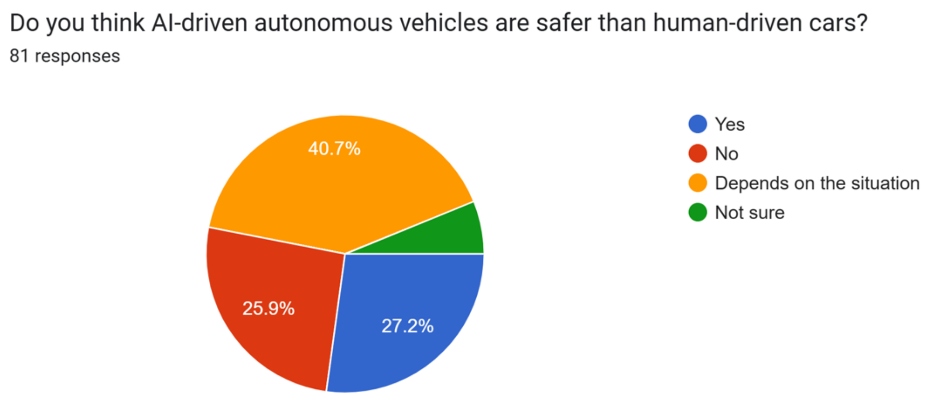

Analysis:

The majority of participants (40.7%) believe that safety depends on the situation that may occur, indicating that trust in AI varies based on driving conditions and external factors. Other respondents (27.2%) think autonomous vehicles are indeed "safer", while 25.9% disagree, believing that human-driven cars remain the safer option. A small portion (6.2%) are unsure about AI's safety capabilities. These results suggest that while some people recognize AI's potential to enhance road safety, many still see limitations and situational risks that require human judgment and intervention.

7.2.9 Question 9: How concerned are you about potential technology failures in autonomous vehicles?

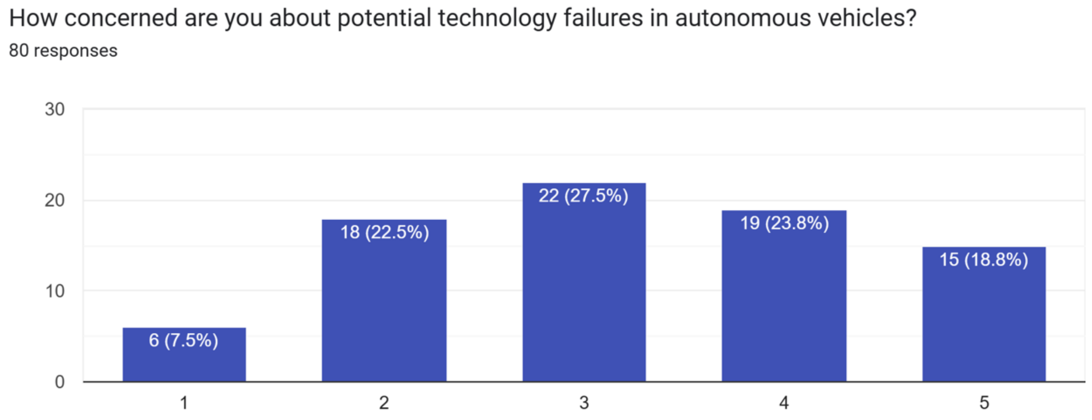

Analysis:

Most respondents (27.5%) express a moderate level of concern (level 3), indicating that they recognize potential risks, but they are not extremely worried. A significant portion (23.8%) reports a higher level of concern (level 4), while 18.8% was at the highest level of concern (level 5), suggesting substantial apprehension about technology failures. On the lower end, 22.5% are at level 2, showing mild concern, and 7.5% are at level 1, indicating minimal worry. These findings highlight that while some individuals are cautiously optimistic, a considerable number of participants have significant concerns about the reliability of autonomous vehicle technology, showing the need for further advancements and assurances in this area.

7.2.10 Question 10: Which of the following would increase your trust in the safety of autonomous vehicles?

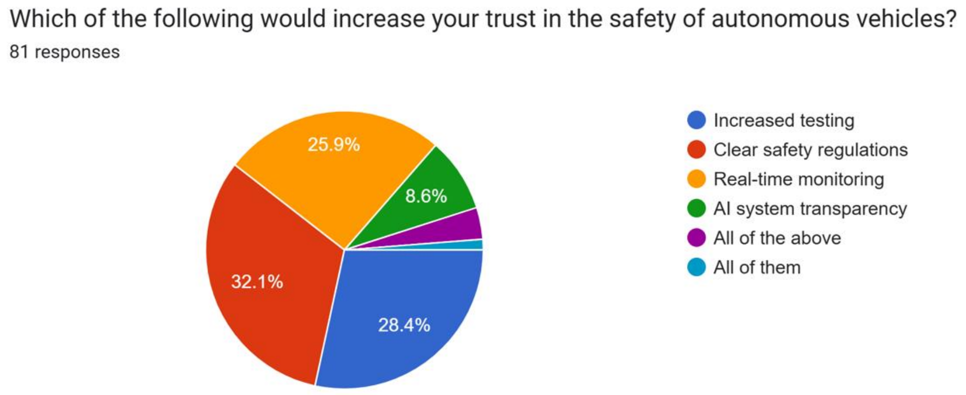

Analysis:

The majority of respondents (32.1%) believe that clear safety regulations would enhance their trust, indicating a desire for structured guidelines. A significant portion (28.4%) prioritizes increased testing highlighting the importance of precise analysis. Additionally, 25.9% of participants emphasized the need for real-time monitoring, suggesting that continuous supervision could boost confidence. A smaller percentage (8.6%) focus on AI system transparency, showing the value of understanding how AI decisions are made. These results reflect a multifaceted approach to building trust, with participants recognizing the importance of testing, regulation, monitoring, and transparency in ensuring the safety of autonomous vehicles.

7.2.11 Question 11: Would you feel safer if autonomous vehicles were equipped with AI systems that explain their decisions in real-time?

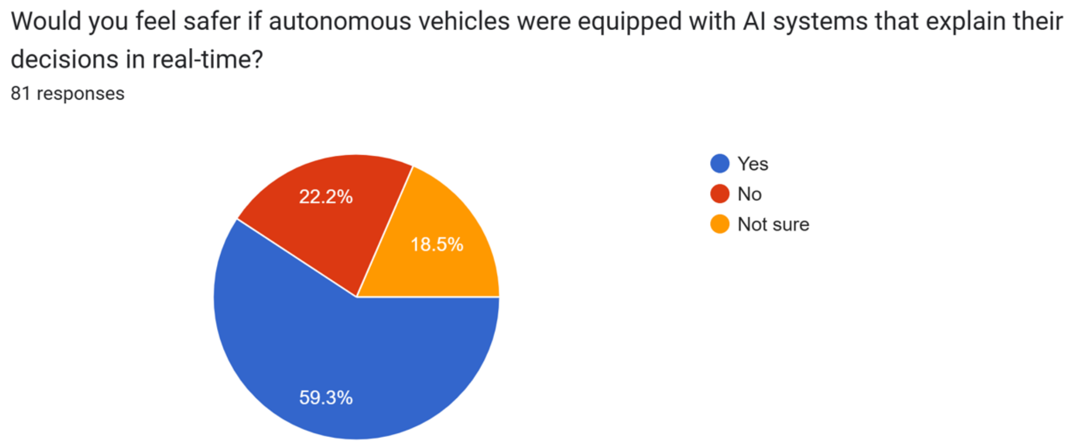

Analysis:

The majority of the participants (59.3%) answered yes, suggesting that they believe real-time explanations could enhance their trust and safety. Meanwhile, 18.5% of respondents are not sure, indicating uncertainty about the impact of such transparency on their sense of safety. Lastly, 22.2% responded no, implying that they do not think this feature would significantly affect their perception of safety. These results highlight that a large group of people see potential benefits in AI decision transparency, while a smaller amount is cautious or did not decide yet.

7.2.12 Question 12: Who do you think should be held legally responsible in the event of an accident involving an autonomous vehicle?

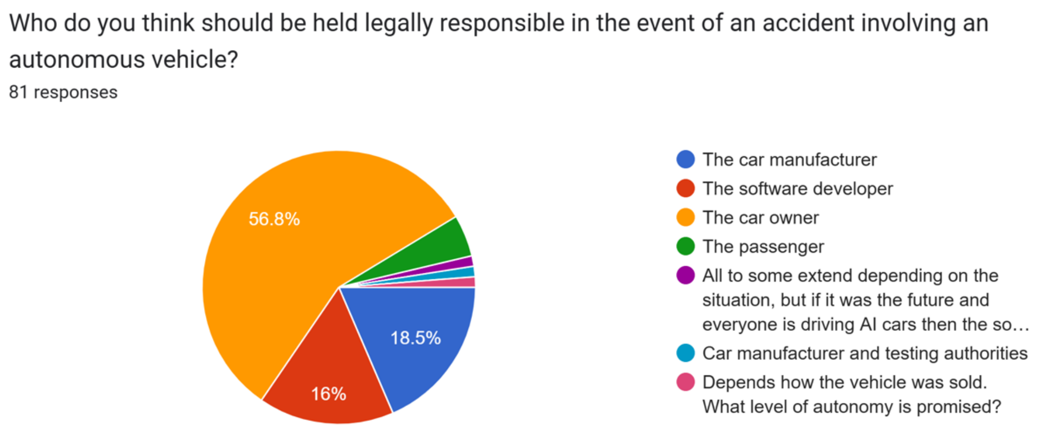

Analysis:

Most of the responses (56.8%) show that the car owner should be held legally responsible, the same that implies now with human driven cars. A smaller portion (18.5%) think the car manufacturer should be accountable, highlighting the importance of the AI systems development before is handed to people. Additionally, some of respondents (16%) suggest that the software developer should bear responsibility. Some participants also provided other responses, emphasizing that liability could depend on factors such as the level of autonomy promised before someone buys the car, or shared responsibility among multiple parties. These results reflect that this is a complicated decision, so many people decide to leave the rules as they are now.

7.2.13 Question 13: How concerned are you about the ethical decisions AI might make in critical situations (e.g., unavoidable accidents)?

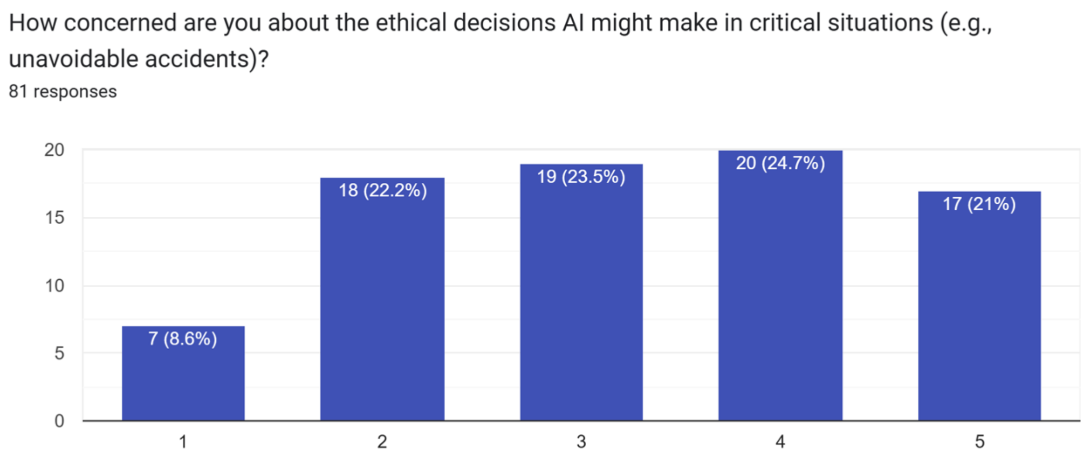

Analysis:

The majority of the respondents (24.7%) express a high level of concern (level 4), indicating significant concern about AI's ethical decision-making. Close behind at 23.5% are at level 3, showing moderate concern, while others (22.2%) are at level 2, with less worries. The 21% are at the highest level of concern (level 5), and 8.6% are at the exact opposite level 1, indicating minimal concern. These findings suggest that most of participants are deeply concerned about AI's ethical capabilities in critical situations, however there are others that are more optimistic or less worried, highlighting the complexity and varied perspectives on this issue.

7.2.14 Question 14: In your opinion, whose safety autonomous vehicle algorithms should prioritize in difficult scenarios:

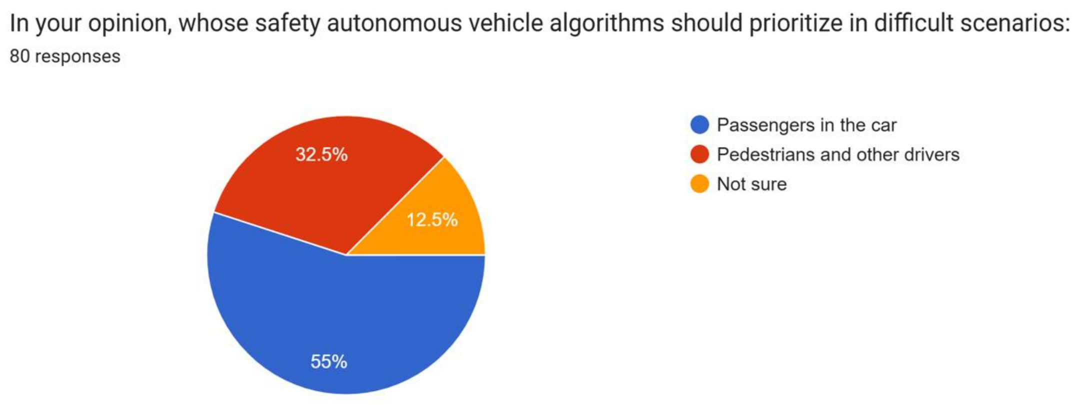

Analysis:

The majority (55%) believe that the algorithms should prioritize the passengers in the car, focusing on the safety of those inside the autonomous vehicle. Meanwhile, 32.5% of participants think the priority should be pedestrians and other drivers, emphasizing the importance of protecting those outside the vehicle. A smaller portion (12.5%) are not sure, indicating uncertainty about how such decisions should be made. In conclusion, the results indicate that people prioritize the safety of those inside the vehicle, perhaps reflecting the idea that if everyone focuses on their own protection, overall harm will be minimized.

7.2.15 Question 15: What is your biggest ethical concern regarding autonomous vehicles?

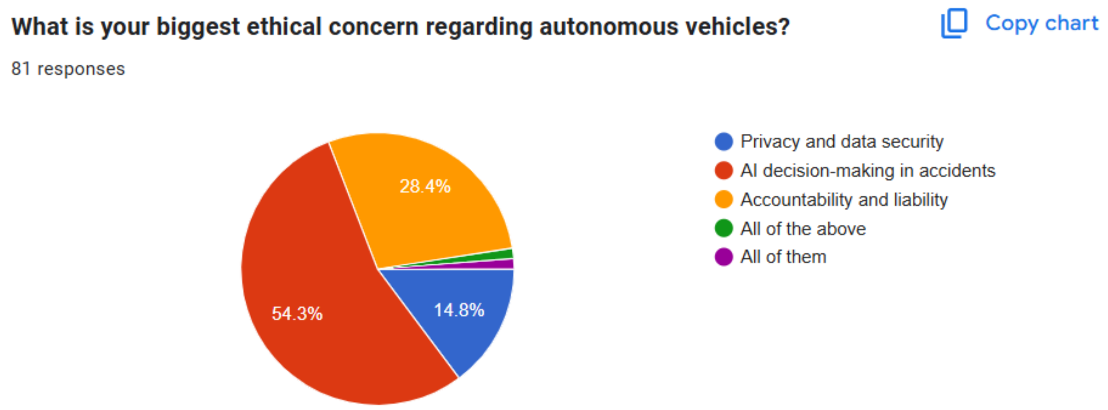

Analysis:

Most of the respondents (54.3%) are most concerned about AI decision-making in accidents, reflecting worries about how AI will handle critical and potentially life-threatening situations. Additionally, 28.4% of respondents are concerned about accountability and liability, emphasizing the need for clear guidelines on who is responsible in the event of an incident. Also, privacy and data security are the primary concern for 14.8% of participants, indicating significant apprehension about how personal data is handled and protected. These results underscore the multifaceted ethical challenges associated with autonomous vehicles, with the main concern being the decisions that our lives depend on.

7.2.16 Question 16: How likely is to use an autonomous vehicle if it becomes widely available in Cyprus?

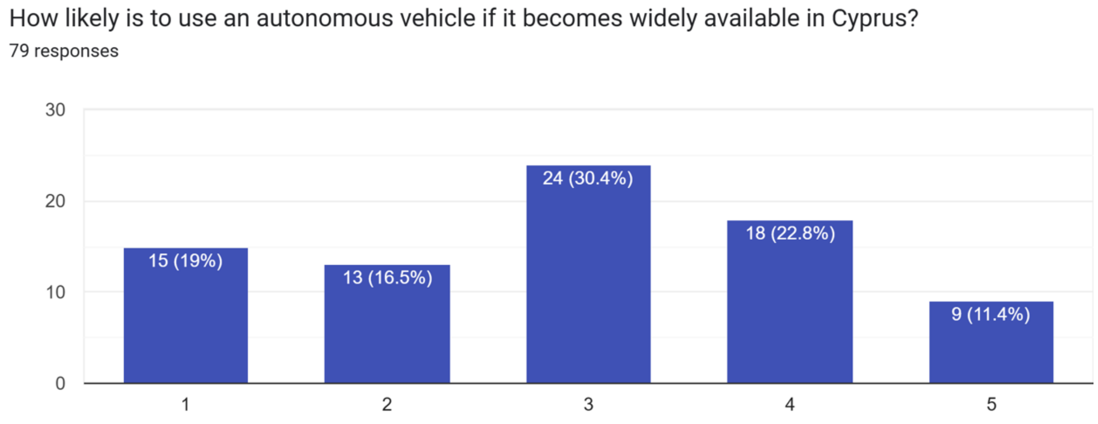

Analysis:

The least number of respondents (11.4%) are highly likely to use such vehicles (rating 5), indicating strong interest and openness to adopting this technology. Another 22.8% are somewhat likely (rating 4), while the largest amount (30.4%) are in the middle (rating 3). On the lower end, 16.5% are somewhat unlikely (rating 2), and 19% are very unlikely (rating 1) to use autonomous vehicles. These results suggest a mixed reception, with most respondents feeling neutral about autonomous vehicles, while a smaller but significant portion is either eager to adopt them or remains skeptical.

7.2.17 Question 17: What is your overall attitude towards autonomous vehicles?

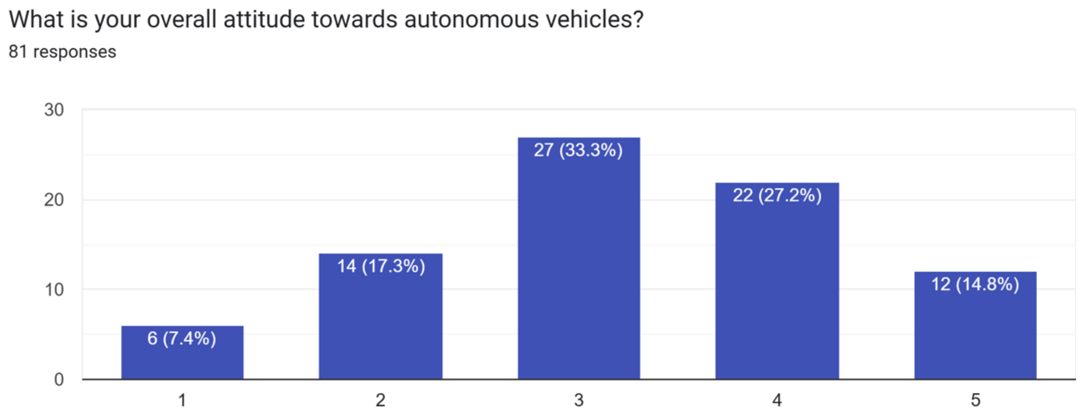

Analysis:

Most of respondents (33.3%) have a neutral attitude (rating 3), indicating a balanced or undecided perspective. A significant portion (27.2%) are somewhat positive (rating 4), while 14.8% are very positive (rating 5). On the more skeptical side, 17.3% are somewhat negative (rating 2), and 7.4% are very negative (rating 1). These results suggest a diverse range of attitudes, with a notable segment of participants holding neutral or positive views, while a smaller group remains cautious or negative about autonomous vehicles.

7.2.18 Question 18: Would you be willing to pay more for an autonomous vehicle equipped with advanced AI safety features?

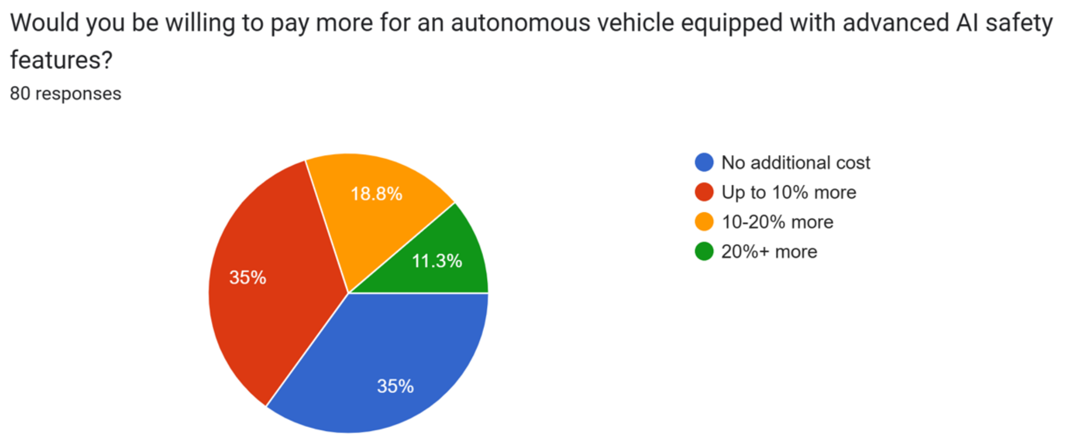

Analysis:

A significant portion of respondents (35%) are willing to pay up to 10% more for these features, indicating some interest in enhanced safety. Another 18.8% are willing to pay 10-20%, while 11.3% are open to paying more than 20%. These results suggest that while a notable segment of participants value advanced safety features and are willing to invest in them, there are others that would not pay extra for these features.

7.2.19 Question 19: Would you like to share any additional thoughts on automated vehicles, AI, or their future impact?

Analysis:

This question gathered 32 responses. Almost half of the responses (14) mentioned that they want to trust more these technologies and to do that they would like more testing and transparency in their development and test results. The above clearly stated "..,but we need time to trust it.”, “I think much more testing is needed..”, “I’m skeptical about the reliability of AI in unpredictable situations.”. Some of those 14 people were concerned in different ways like “people might rely too much on them” or they think that “privacy” could be an issue. Two other people are concerned about privacy as well especially hacking scenarios and gaining access to personal data, or malicious software that leads to crashes, they mention that “ethical programming is critical” and the need of “strong security”. On the other hand, seven respondents focused on useful directions that autonomous driving could leads us. Upon them there are thoughts about “AI taxes” taking over so ‘fewer people might own cars, and cities could change in a way to help the environment” and “AI may well replace drivers, such as trucks, in order to reduce companies' spending on truckers”. However the same ones mentioned that this also could lead to some people losing their jobs, or that insurance could be a problem for the users. Two others stated that this would be one of the biggest achievements in the world if fully automated cars become publicly available. Someone also specified that in some countries “the roads are not ready yet for these technologies”. This lead me to another answer that says that “road network needs to be prepared for this kind of vehicles” and if they are not “then more accidents are likely to happen.”. Lastly, a person pointed out that “These vehicles should be completely funded”.

7.3 Summary of Findings

This survey has brought to light several key findings regarding the current and future public views on autonomous vehicles, mostly in Cyprus. Here is a summary of the findings:

1. **Population Profile**
   The survey shows that most respondents are young and well-educated, since a high number of participants hold a bachelor’s or master’s degree. Most of the respondents are from Cyprus, with a smaller group from Greece and other parts of the world. Many participants have a moderate understanding of how AI works in self-driving cars. A good number say they are somewhat familiar with the technology, while only a few consider themselves very knowledgeable.

2. **Awareness of AI in Vehicles**

   Considering that these people are not experts, many of them are optimistic about AV's handling challenging and complex scenarios e.g. bad weather, traffic, although a portion of them still feels unsure about this technology's ability.

3. **Trust and Safety Concerns**

   When it comes to trust, from the survey we get that while many believe AI can perform well in complex driving situations, they also want more safety measures in place. They emphasize on the need for more testing, rules becoming clearer for them, and updates that happen real-time on how the AI makes decisions. They also really value the option for the drivers to take control manually if they think is necessary, showing that people want a balance between automated technology and human oversight to gain more trust.

4. **Ethical and Legal Issues**

   The findings highlight that most respondents think that if an accident happens, the car owner should be held responsible. Some also point to the software developer or even the manufacturer, depending on the situation. There are significant concerns on how AVs will make ethical decisions, especially in scenarios where accidents are realistically unavoidable. These concerns indicate that clear guidelines and regulations are crucial for the technology's acceptance.

5. **Adoption and Attitudes**

   The overall attitude towards autonomous vehicles is mixed. While a significant number of respondents are interested in using these vehicles and show a willingness to pay extra for enhanced safety features, many others remain cautious or unsure.

6. **Future Impact and Additional Thoughts**

   Despite the worries on safety and trust of AVs, many see the potential benefits of AI in reducing accidents and easing traffic congestion. However, there are also worries about issues like cybersecurity, ethical programming, and even economic impacts. In summary, while there is hope that autonomous vehicles can bring positive changes, the technology must address several safety and legal challenges to gain widespread public support.

Chapter 8: Conclusion

8.1 Summary of the Study

The thesis goal was to explore and categorize the key components of autonomous driving systems, focusing on artificial intelligence and machine learning techniques. From examining the structure of the autonomous vehicle decision-making specifically perception, localization, prediction, planning, and control, the aim was to provide a technical but accessible breakdown of the technologies that are shaping self-driving capabilities these days. Drawing from a broad set of academic sources, the study contextualized each module within the field of robotics and AI-based automation. This played a critical role for analysing both existing implementations and future trajectories of autonomous systems.

8.2 Emphasis on Research-Based Insights

The two research efforts in this thesis, which targeted both industry professionals and the public, offer a thorough perspective on the current state as well as societal reception of AI-driven autonomous vehicles.

From the industry side, the surveys capture major challenges, like real-world testing, data labelling, integration with legacy systems, and computational efficiency. In addition, experts also emphasized handling human unpredictability, adapting to various conditions, and meeting regulatory demands. Ethical decision-making is mainly approached through simulation and some external audits for accountability. Several AI models that are commonly used are neural networks for perception and reinforcement learning for decision-making. Edge computing emerged as a key strategy to compete with real-time limitations. Lastly, certain innovations like multimodal sensor fusion and federated learning were described as promising for advancing autonomy.

The public survey revealed overall mixed attitudes, since there is growing awareness of AVs; however, trust remains conditional. The participants stated that they want clear safety regulations, transparent AI behaviour, manual override options, and real-time feedback. There was given importance to Ethical concerns, liability questions, and system reliability as well. In conclusion, still, many people are cautiously optimistic and willing to adopt AVs if safety and transparency are ensured, so this shows that successful AV deployment requires both technical robustness and alignment with public values.

8.3 Future Research Directions

Regarding the combined findings, the following areas for future research are suggested:

- Human-Centered Interfaces: Further exploration is needed on explainable AI and real-time driver feedback mechanisms to increase user trust and situational awareness.
- Simulation and Edge-Case Testing: It is important that advanced simulation environments are created that mimic rare or unpredictable human actions, as both were highlighted from industry experts and public concerns.
- Collaborative models: Research must examine ways in which industry-government partnerships can clear safety, ethical standards, as well as legal accountability throughout regions.
- Efficient Onboard AI: The optimization of lightweight AI models and hardware-software co-design are still a problem for real-world large-scale scenarios.

Bibliography

[1] IEEExplore "Autonomous cars: Past, present and future"
https://ieeexplore.ieee.org/abstract/document/7350466 [accessed on 18/11/2024]

[2] ConsumerAffairs "How safe are autonomous vehicles?"
https://www.consumeraffairs.com/automotive/autonomous-vehicle-safety-statistics.html [accessed on 13/10/2024]

[3] McKinsey&Company "Autonomous driving's future: Convenient and connected"
https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/autonomous-drivings-future-convenient-and-connected [accessed on 18/11/2024]

[4] National Library of Medicine, PubMed Central "Impacts of Autonomous Vehicles on Greenhouse Gas Emissions—Positive or Negative?"
https://pmc.ncbi.nlm.nih.gov/articles/PMC8197118/ [accessed on 16/10/2024]

[5] Bosch (Malene Freudendal-Pedersen, Professor in Urban Planning at Aalborg University) "Impact of self-driving cars on society"
https://www.bosch.com/stories/impact-of-self-driving-cars-on-society/ [accessed on 22/10/2024]

[6] MDPI "Autonomous Vehicles: Evolution of Artificial Intelligence and the Current Industry Landscape"
https://www.mdpi.com/2504-2289/8/4/42 [accessed on 22/10/2024]

[7] Brookings Institution "Ethical challenges and autonomous vehicles"
https://www.brookings.edu/articles/the-folly-of-trolleys-ethical-challenges-and-autonomous-vehicles/ [accessed on 18/11/2024].

[8] Artificial Intelligence for Autonomous Vehicles. The Future of Driverless Technology.
First published 2024 by John Wiley & Sons, and Scrivener Publishing LLC, Edited by Sathiyaraj Rajendran, Munish Sabharwal, Yu-Chen Hu, Rajesh Kumar Dhanaraj and Balamurugan Balusamy.

[9] Journal of Physics:Conference Series "Autonomous car using CNN deep learning Algorithm"

https://iopscience.iop.org/article/10.1088/1742-6596/1869/1/012071/pdf [Accessed on 10/04/2025]

[10] Autonomous Vehicles Intelligent Transport Systems and Smart Technologies By Nicu Bizon Lucian Dascalescu and Naser Mahdavi Tabatabaei

Published by Nova Science Publishers, Inc. † New York

[11] An Object Classification Approach for Autonomous Vehicles Using Machine Learning Techniques by Majd Alqarqaz, Maram Bani Younes and Raneem Qaddoura https://www.mdpi.com/2032-6653/14/2/41 [Accessed on 12/04/2025]

[12] Motion Planning and Control Techniques for Self-driving Urban Vehicles by Brian Paden, Michal Cáp, Sze Zheng Yong, Dmitry Yershov and Emilio Frazzoli. https://arxiv.org/pdf/1604.07446 [Accessed on 24/04/2025]

[13] Self-Driving Cars With Convolutional Neural Networks (CNN) https://neptune.ai/blog/self-driving-cars-with-convolutional-neural-networks-cnn [Accessed on 27/04/2025]

[14] A Survey of Deep Learning Techniques for Autonomous Driving https://www.researchgate.net/publication/336638898_A_Survey_of_Deep_Learning_Techniques_for_Autonomous_Driving [Accessed on 27/04/2025]

Appendix A: Industry Insights on Autonomous Vehicles Survey

Industry Insights on Autonomous Vehicles

Thank you for participating in this survey. The aim is to understand the challenges, innovations, and trends in developing autonomous vehicles, particularly in relation to AI and machine learning. Your responses will contribute to valuable research in this field.

1. What is the primary focus of your company in the autonomous vehicle sector?

   ____

2. What is your role within the company?

   Mark only one oval.

   - ( ) Software Engineer
   - ( ) Data Scientist
   - ( ) Project Manager
   - ( ) Research & Development
   - ( ) Other: ______________________________

3. What are the primary bottlenecks your company faces in developing AI for autonomous vehicles?

   Check all that apply.

   - [ ] Data collection and labeling
   - [ ] Computational resource limitations
   - [ ] Model interpretability and explainability
   - [ ] Real-world testing and safety concerns
   - [ ] Integration with hardware systems
   - [ ] Regulatory or legal challenges
   - [ ] I am not familiar with the challenges
   - [ ] Other: ______________________________

4. Are there any challenges specific to your company's operations that are not widely discussed in the industry? Please elaborate.

5. How does your company address ethical considerations, such as decision-making in life-critical situations?

6. What types of machine learning models are most commonly used in your autonomous systems?

   Mark only one oval.

   - Neural Networks (e.g., CNNs, RNNs)
   - Reinforcement Learning
   - Supervised Learning Models (e.g., regression, classification)
   - Unsupervised Learning Models (e.g., clustering, dimensionality reduction)
   - Transfer Learning Models
   - I am not sure about the specific models used
   - Other: ____________________

7. How is the quality of your training data maintained and verified?

8. What strategies are employed to optimize computational efficiency in real-time AI operations?

9. Are there specific breakthroughs in AI that your team/company is excited about, or that you believe will significantly advance autonomous vehicle technology?

10. What are the most critical milestones in your opinion for achieving full autonomy?

11. In your opinion, what role will collaboration between car manufacturers, AI companies, and governments play in shaping the future of autonomous vehicles?

12. How do you see the role of AI in autonomous vehicle innovation evolving over the next 5-10 years?

13. Are there any other insights you'd like to share regarding autonomous vehicle development?

Appendix B: AI-Powered Self-Driving Cars: Public Views on Technology, Trust, and Ethical Boundaries Survey

AI-Powered Self-Driving Cars: Public Views on Technology, Trust, and Ethical Boundaries

This
survey is part of a diploma study examining public opinions on AI-driven automated vehicles. Your responses will provide valuable insights into areas like technology, safety, and ethics related to autonomous vehicles. The survey should take about 5-10 minutes to complete.

1. **Your age:**

   Mark only one oval.

   - Under 18
   - 18-24
   - 25-34
   - 35-45
   - Above 45

2. **Gender:**

   Mark only one oval.

   - Male
   - Female
   - Other

3. **Education Level:**

Mark only one oval.

- High School
- Diploma
- Bachelor's
- Master's
- Ph.D

4. **Country:**

____

5. **How familiar are you with the use of AI algorithms in automated (self-driving) vehicles?**

Mark only one oval.

- Very familiar
- Somewhat familiar
- Slightly familiar
- Not familiar at all

6. **Do you believe that AI algorithms can handle complex driving scenarios (e.g., heavy traffic, adverse weather) effectively?**

Mark only one oval.

- Yes
- No
- Not sure

7. **What features would make you feel more comfortable with AI technology in automated vehicles?**

8. **Do you think AI-driven autonomous vehicles are safer than human-driven cars?**

Mark only one oval.

- Yes
- No
- Depends on the situation
- Not sure

9. **How concerned are you about potential technology failures in autonomous vehicles?**

Mark only one oval.

	1	2	3	4	5	
Not						Very concerned

10. **Which of the following would increase your trust in the safety of autonomous vehicles?**

Mark only one oval.

- Increased testing
- Clear safety regulations
- Real-time monitoring
- AI system transparency
- Other: ____________________

11. **Would you feel safer if autonomous vehicles were equipped with AI systems that explain their decisions in real-time?**

Mark only one oval.

- Yes
- No
- Not sure

12. **Who do you think should be held legally responsible in the event of an accident involving an autonomous vehicle?**

Mark only one oval.

- The car manufacturer
- The software developer
- The car owner
- The passenger
- Other: ____________________

13. **How concerned are you about the ethical decisions AI might make in critical situations (e.g., unavoidable accidents)?**

Mark only one oval.

	1	2	3	4	5	
Not	◯	◯	◯	◯	◯	Very concerned

14. **In your opinion, whose safety autonomous vehicle algorithms should prioritize in difficult scenarios:**

Mark only one oval.

- ◯ Passengers in the car
- ◯ Pedestrians and other drivers
- ◯ Not sure

15. **What is your biggest ethical concern regarding autonomous vehicles?**

Mark only one oval.

- ◯ Privacy and data security
- ◯ AI decision-making in accidents
- ◯ Accountability and liability
- ◯ Other: ____________________

16. **How likely is to use an autonomous vehicle if it becomes widely available in Cyprus?**

Mark only one oval.

	1	2	3	4	5	
Very	◯	◯	◯	◯	◯	Very likely

17. **What is your overall attitude towards autonomous vehicles?**

Mark only one oval.

	1	2	3	4	5	
Very	◯	◯	◯	◯	◯	Very positive

18. **Would you be willing to pay more for an autonomous vehicle equipped with advanced AI safety features?**

Mark only one oval.

- ◯ No additional cost
- ◯ Up to 10% more
- ◯ 10-20% more
- ◯ 20%+ more

19. **Would you like to share any additional thoughts on automated vehicles, AI, or their future impact?**

____

____

____

____

____