Machine Learning for Autonomous Vehicle Fleet Coordination

By Dr. Ayşe Gülcü

Professor of Electrical and Electronics Engineering, Istanbul University, Turkey

1. Introduction

In recent years, the relevance of autonomous vehicles has increased greatly. In combination with the increase in shared mobility services, autonomous vehicles could provide a shift towards more sustainable transport systems. However, typical transport-related fundamental challenges such as scalability, large and sudden demand peaks, regional seasonal variations, and high goals for traffic system efficiency have persisted. To efficiently realize autonomous demand-responsive mobility concepts, research is needed on innovative methods for the coordination of vehicle fleets, introducing new ideas.

The recent advances in computer technology have created applications for machine learning research in many digital platforms, including intelligent and autonomous transportation systems. The classroom implementation of a relatively small-scale demand-responsive autonomous vehicle fleet constitutes a complex optimization problem that can be tackled with a learned agent. Applications at this level have the ability to impact society broadly. In traffic psychology, it has been repeatedly hypothesized that traffic automation would be beneficial. Advancements in this area have great potential in improving road traffic efficiency, safety, and comfort. In addition, the above coordination solutions are generic, with no application-specific requirement for on-road evaluation. In this work, we specifically focus on ML-based techniques to solve the vehicle coordination problem. In a broader sense, we try to answer the generic challenges in autonomous vehicle fleet coordination at microscopic levels, involving the assumptions of autonomy and closed vehicle fleets in urban areas.

The high accuracy of fleet routing, combined with the flexibility in scheduling, must be compensated by the real-time coordination of the fleet itself. The necessity of effective fleet coordination is vital to the optimal usage of the available infrastructure. To have a more holistic point of view, the on-demand public transit ecosystem includes requests from humans or virtualized sensors and should be able to serve multiple non-coordinated vehicle categories such as bicycles, scooters, or e-cabs, in addition to AVs. Further permitting single-passenger and group-based rides will enrich real-world scenarios from feasibility and effectiveness perspectives. In the absence of road transport infrastructure availability for the continuous flow of all the fleet, management is required. Overall, one of the contributions of this work is to provide realistic operations on-road without any method to precisely analyze the data since there is no human-driven vehicle coordination.

1.1. Background and Motivation

The promise of autonomous vehicles has enticed technologists and innovators for over half a century. With recent advances in fields such as robotics, machine learning, artificial intelligence, computer graphics, and automotive engineering, autonomous vehicles are beginning to transcend speculation and leave the experimental phase of development to enter the general public's national and international transportation systems. It is also at the forefront for many corporate interests. The public will have access to these vehicles via a ride-sharing app, and people will use it to take them to work, home, or recreational activities after a concert, game, or night at a bar. Action movies and popular press have investigated scenarios with varying climates from utopian hopes to dystopian fears regarding these vehicles and their use.

The public is justifiably interested in the development of autonomous vehicles because of the promises these vehicles hold. First and foremost, autonomous vehicles promise increased safety when they reach full autonomy. Transportation experts predict reductions in accident rates ranging from 80% to 90% once all vehicles on the road are fully autonomous. In addition to increased safety, autonomous vehicles also offer the potential for improved fuel efficiency through the optimization of acceleration, steady-speed driving, and electronic following, leading to a potential improvement in following distance and traffic flow. While a number of perception, prediction, and planning problems associated with driving intelligent agents, particularly in the car-following and highway merge scenarios, have been developed recently, concurrent research concerning the coordination of multiple agents has primarily focused on independent decisions or common communications for merged and reinforced rewards. Very little research addressing the importance of communication for differences of opinion exists. This is the specific gap in the autonomous vehicle coordination research we investigate, and we integrate current machine learning techniques to address this currently open issue.

Understanding the importance of contentious communication will provide a foundational understanding in the future for the potential impact of differences of opinion in a variety of social, acrimonious, and political systems.

1.2. Scope and Objectives

This work focuses on deploying machine learning techniques and coordination patterns that are centralized but scalable to a large autonomous vehicle (AV) fleet, either operated by a single entity or as a mixed fleet. The aim is to advance from theoretical concepts of fleet coordination towards scenarios closer to real-world applications. Thus, the mentioned demographic focus allows for elaboration on the actual state of the art of fleet composition, geographic constraints, and regulations in operation and testing, additional equipment availability, and the expected behavior of vehicle agents. The main objective of this work is to develop and verify improvements in the performance of traffic light management, such as reduced traffic time or increased satisfaction of traffic participants. Optimization of operations, i.e., reduced fuel consumption, less harmful emissions, and efficient transportation services management, is also set. The project results should indeed contribute to enhancing safety, i.e., reductions in collision numbers and general operating conditions. In terms of contributions to the state of the art, we would like to bridge the gap between a sophisticated and non-scalable realistic transition of alternatives with primary or sole reliance on vehicle-to-infrastructure communication towards a dense centralized network of vehicle agents directed and constrained to provide a desired general behavior of traffic networks instead of micro-optimization of their own interests. This can speed up other research in the direction of adding a mechanism for adaptation to the presence of human drivers, infrastructure, and policies or grouping of AVs to perform specific missions. Unfortunately, the research is subject to constraints not only theoretical, such as technology development, security, and implementation of fleet coordination concepts, but also related to the policy and legal environment. Indeed, we assumed the operation of the whole fleet of vehicles under control in the tested environments. The methods described benefited from this assumption as they were primarily designed to optimize central control, but on the positive side, it also proved the potential for scalability when the whole fleet or its part switched to autonomous mode. We chose to focus on methods and techniques applicable to relatively large real-life operations of multiple AVs, taking into account regulation in city traffic and the assumption of a daily maximum of motor vehicles throughout the project. The technologies used and tested are similar to those used in navigation systems or applications in which the vehicle indicates a desired destination and is routed to it. In addition, we have modeled and verified a path planning mechanism from the perspective of optimization and coordination between vehicles, not only at road intersections but through entire network lanes. From a theoretical perspective, to the best of our knowledge, the work done is a development of several well-known solutions and theories in the context of vehicle path coordination towards a fleet of contrasting composition and constraints. The use of various abstraction levels, such as the basic concept of a desired control point for a network of vehicle constraints, provided the basis for composing effective vehicle management strategies allowing for both collision avoidance and respect for traffic rules and network operations. In our work, vehicle take-over and give-back are explored.

2. Fundamentals of Autonomous Vehicles

Autonomous vehicles are self-driven vehicles that are capable of sensing and interpreting their environment to make real-time decisions about actions, such as steering, braking, and acceleration. These vehicles combine different types of sensors, such as cameras, radars, ultrasonic sensors, and lidars, to perceive the surrounding environment. Subsequently, the vehicles use various sophisticated algorithms to produce control signals that are used to maneuver in the environment. As a result, these vehicles have great potential in enhancing the safety, mobility, and eco-friendliness of transportation systems. The issues of vehicle-fleet coordination, for instance, are a novel and challenging research area in autonomous multi-vehicle systems. This text is focused on the application of machine-learning techniques in addressing these and allied concerns.

The notion of a fully autonomous vehicle, often synonymous with unmanned ground vehicles, comes with predefined characteristics. Unlike traditional car-like robots, these vehicles are not tethered. They do not need to be manually controlled or teleoperated by humans. Fully autonomous vehicles must be endowed with perception capabilities to sense and interpret their environment in order to plan and make actions in real-time. These vehicles could be designed in fun or functional fashions. Empirically, fully autonomous vehicles have been grouped under different taxonomies largely depending on the level of autonomy or application in question. Firstly, in terms of the level of autonomy, autonomous driving systems or vehicles are categorized into six different levels of driving automation. Level 0 corresponds to no automation. Level 5 designates full levels of autonomy. Other levels represent intermediary or so-called active safety features where an autonomous system takes over only for certain functions. Secondly, fully autonomous vehicles could also be categorized based on their usages in such areas under study as transportation functions, industrial logistics, or so-called robotic mission-specific functions. Many of these autonomous vehicles are developed to sense and move in structured outdoor or indoor environments that were initially enhanced to be more perceptually clear. These autonomous vehicles could be seen meandering in such environments as deserts, forests, highways, pavements, pedestrian zones, shopping malls, nuclear power plants, research facilities, airports, warehouses, waterways, or museum premises, or sometimes intelligent homes. These vehicles are aided by the use of different types of sensors, the interpretation of the sensory data, and the execution of algorithms realized in software and hardware to perform various sets of actions. The sensors could be over the ground, a mechanical representation of the six degrees of freedom compliant sensors with a suspended spring base; in the ground, the tractive and non-tractive wheels, the castor, ball or tank tracks, or the magnetic wheels; or the aerial vehicles. The central processing units used in these autonomous systems are usually centralized within the software and hence affect the processor-hardware design. This integration between processor (hardware) and the right algorithms (software) is what gives any autonomous vehicle or system its robustness.

2.1. Definition and Characteristics

In this paper, we focus on the coordination of autonomous vehicle (AV) fleets. Thus, we provide the definition and principal characteristics to lay down the groundwork for the cooperative management of AVs. An autonomous vehicle (AV) is a vehicle that can guide itself without human conduction. Despite the fact that the driver or passenger must provide some information to an autonomous vehicle, as this will affect the cost and arrival time of the trip, the decision about the route is at the discretion of each AV based on arrival time and working order.

This definition highlights the three main characteristics of what is known as an autonomous vehicle or intelligent vehicle: sensory perception, decision-making, and navigational capabilities.

Based on the level of automation (LoA) defined by the International Society of Automotive Engineers (SAE), the types of vehicles are classified according to the amount of human intervention needed in the driving task. There are six levels of automation, ranging from level zero to level five and varying from a fully manual to a fully autonomous system. These six levels are based on the driving domain, performance and safety of task execution, and fallback performance of the human driver. Overall, LoA can be seen as a policy of increasing convenience for AV passengers' trips that are defined according to five operational profiles within the main LoA, i.e., levels three to five. These operational profiles list the degree and range of human intervention, and the characteristics of fallback dealing. It is important to clarify these AV principles, because if we do not address the working method of the AVs, there might be an issue with the assumptions laid out in Section 1.

2.2. Types of Autonomous Vehicles

There are a number of ways to classify different types of autonomous vehicles, which reflect their diverse functionalities and potential applications. In the research literature, the most commonly used classification is based on the level of automation: for example, autonomous vehicles can be fully autonomous with no human intervention, or require some level of control from a human driver, designating them as semi-autonomous. Prevailing classification methods in industry define autonomous vehicles as a part of a larger network of connected machines, humans, and other entities. Established ways of categorizing include distinguishing between vehicle-to-vehicle and vehicle-to-infrastructure communication. Examples of autonomous vehicles are abundant, with applications spanning private, public, military, and industrial use. Various projects have produced small-scale autonomous personal vehicles, such as self-driving cars. A large proportion of research focuses on developing fully autonomous systems for persistently explicit purposes, such as autonomous public transport shuttles or delivery drones for last-mile transport of goods.

In a relevant industrial context, a potential client of an autonomous vehicle is a fleet management company that may have different cars being different kinds of vehicles as a part of their fleet. The choice of vehicle (and hence the underlying vehicle type) influences the coordination strategy of the fleet management algorithms. There are several factors that influence the choice of a particular vehicle type: its cost in depreciation and liability, regulatory compliance, and technology readiness, among others. Given this insight, we propose a classification scheme that aims to incorporate both the commercial priorities and the physical architecture of different vehicle designs. We hope that this classification scheme sheds light not only on the autonomous vehicle landscape but also on the landscape of proposals available for a fleet management company. We believe that this enhanced understanding of the problem will improve coordination strategy performance.

3. Fleet Coordination Challenges

A fleet of autonomous vehicles (AVs) poses significant complexity for those looking to manage them. The vast majority of existing deployments result in fleets existing in a world with non-fleet traffic; this provides benefits to the fleet that are not always possible for management. By and large, their interactions occur at low density in shared spaces, at slow to moderate speeds, and may be impacted by pedestrians or complex elements of the urban fabric. If the vehicles are to cooperate, their management systems need to be able to request and provide space, time, and other inputs. The vehicles will need to interleave, mix, and measure for inherently different flows, processes, and vehicle types. The fundamental challenge of managing these systems relates to the fact that if a single element behaves poorly, the entire system might encounter issues with traffic management or service delivery. Vehicles moving at lower speeds or with different densities can have a longer impact on their surroundings, potentially giving them greater heft as bottlenecks or obstructions than the same numbers or increased capacities of cars.

How do you: Increase capacity? Ensure safe and efficient interoperability with non-fleet users? Guarantee optimal flow with specific given intermediary infrastructure? Positively affect overall fleet dynamics with vertical or horizontal elements of coordination? The ability to manage fleets of this nature relies on a global view of traffic involving AVs and non-AV systems such as people, personal transport, and other vehicles. Many of the early works on AVs suggest that the integration benefits will come from the ability of fleet systems to interact with a broader traffic management system, typically an urban or highway management system designed to ensure the safe and efficient movement of vehicles and goods. The vast majority of existing systems are built on technology in the infrastructure to directly impact the vehicle equipment and overall vehicle movement, although some new technologies in infrastructure or in specific field use cases rely on wireless, long-range, or local range communication between vehicles and system managers. Such systems form the basis of much current research, and increasingly, academic interest is being expanded with technological development towards deeper study of the new types of artificial intelligence tools and measures that will be required to integrate very large fleets into logistically dense traffic zones or corridors. In more than 32 countries worldwide, research and investment are being made to introduce wireless communication and vehicle adaptation that will allow — among other things — the styling position of vehicles and direction of flow. There's a gap of about 10 years until wider commercial introduction suitable for genuinely federal introduction of research developments and similar that will impact the dynamics of space-filling.

3.1. Shared Spaces vs. Complex Environments

Shared spaces like residential areas, city centers, and road junctions are populated by people conducting their daily activities. Such places, public by nature, host many peripheral events including street vendors and kids playing. These generate attention at different scales of time and space, reflecting on driving speed or dynamic road width, as well as disrupted traffic regularity. Such a characteristic makes traffic in this class of shared space semi-predictable at an event level. On the other hand, a large portion of the considered urban street network hosts regular traffic patterns that are hard but feasible to predict if one knows the departure time of generator locations. These networks are, independently from the traffic conditions, qualitatively different from shared spaces as human activity is regular and prohibited from leaving sidewalks and crosswalks. There also exist extensive plans for these networks at a continental and global scale.

The most convenient paths in the aforementioned cases are detour-deprived: they are locally traffic-free. Uniformly 2, 3, or n detour-deprived paths may not exist in planar or vertical grids due to the impossibility of detouring all carriers involved. However, when carriers inhabit a multilayer network or a zoning road network, a plurality of detour-deprived paths exists, with detours being locally detour-deprived. The latter is not proof but only a sketch to provide insight into the flexibility of the AI module. The kind of AI module exhibited prohibits the extension of FF-based AI modules to complex environments, such as shared spaces, while shared traffic AI modules cannot handle multilayered environments. A new reasoning system must be researched that is universal in the scope of environments and is capable of activating shared and shared-modified AI modules, depending on the distribution of carriers in z-planar and layers that are quantitatively comparable.

3.2. Traffic Management and Safety

Traffic management is important for safety, as congestion can pose a significant safety risk. By controlling the movement of vehicles, traffic management systems are used to maximize throughput while minimizing queue length, congestion, and delay. Traffic management systems make use of data, analytics, and real-time monitoring infrastructures to enhance situational awareness for the vehicles and infrastructure alike. Moreover, the data captured sometimes allows for predictive analytics – for instance, the prediction of traffic patterns and volumes for the purpose of designing traffic management policy. Mechanisms for adjusting and managing traffic flows in real time typically involve making short-term decisions about routing, lane allocation, traffic signal timing, or speed limits. However, because of the limitations of reactive approaches, demand-responsive traffic management techniques have been developed, aiming to manage the traffic by influencing the demand. In terms of routing techniques, machine learning has a good ability to predict traffic queues and adaptively assist routing strategy. In addition to traffic management for traffic safety, there are other safety topics linked to autonomous fleets. When a vehicle, either autonomous or driven by a human driver, enters a viewpoint of objects in a shared environment, safety is all about how to react to the other objects in this environment. The other vehicle has to obey the traffic policy to proceed and give right-of-way to others. In some autonomous vehicle fleets, various autonomous vehicles are collectively making decisions in order to minimize the cost, such as the time to reach the target, in the combined conflict resolution. In the scenario we are concerned with now, since all the vehicles are autonomous vehicles under centralized decision-making, all the vehicles follow the traffic policy, and the traffic policy obeys the rules of traffic regulations. We should be concerned about how to ensure that the system that governs all the vehicles' decision-making processes is completely safe, considering all of the situations defined in the scenarios. In addition to decision-making, the safety of fleet control also contains coordination mechanisms of data fusion and sharing, requests broadcasting, and registration, and so on. Pursuits in the aspect of safety coordination are of particular interest, and more solid research constructed on the basis of standard regulation and mechanisms is in urgent need. Considering all the traffic safety curriculum that we have stated, for completely autonomous vehicle fleet control, it must combine with the traffic regulation authorities to make standardized regulations. In terms of traffic management, a traffic management authority is introduced due to its critical role in the security of traffic. Within the other safety aspects of the whole system discussed, effective regulation between the authorities and the local centralized components is required, or other risks that can arise unexpectedly, such as strikes among dispatchers of both authorities and centralized fleet components, for traffic management in emergency cases. In addition to these cases, the advanced integrated display system of traffic control with higher situational awareness is the tendency of development for secure control of traffic. When multiple fleets inhabit the same environment, implementation should be taken for coordination across fleets. The decision of fair sharing of the road should be given to cross-fleet centralized dispatchers; and centralized enterprise authorities from different fleets should have the right to perform roadmap scheduling on the road for multiple fleets; and when multiple vehicles from different fleets show centralized and coordinated behaviors on the road, the decision should be made by the centralized regulatory and dispatch authorities from all the fleets.

4. AI Techniques for Fleet Coordination

AI Techniques for Fleet Coordination. Fleet operations can be significantly enhanced and optimized with the use of artificial intelligence (AI). These include machine learning and optimization techniques. Hence, autonomous vehicles can learn from their and the entire fleet's past behavior. There is a rich literature regarding the use of AI in the scheduling and control of transportation systems in general and in fleet management in particular. Here, after a brief discussion of the general survey on AI in the control and operation of transportation systems, we focus on two important learning techniques that hold many potentials for the coordination of autonomous vehicles in a platoon and their deployment in the more general problems involving a fleet of autonomous vehicles, namely, reinforcement learning and deep learning. Reinforcement learning essentially allows an autonomous vehicle to choose the actions that optimize the system performance based on the feedback it receives from the environment. Deep learning can be effectively used for pattern recognition such as human behavior modeling, dynamic environment prediction, and state-of-the-art path planners.

Reinforcement learning (RL) is an artificial intelligence technique in which the agent learns to make decisions by interacting with an environment and receiving a reward signal for its behavior. The physical environment is encapsulated in a state, action, reward, state Markov decision process. Assume that the agent operates in a finite time horizon. At each time step, the agent observes the state of the system, operates an action, receives a reward, and moves to the next state according to a transition model. The agent has to decide on an action that balances the trade-offs between immediate and future rewards. The transition model and reward function are usually assumed to be unknown to the agent. The temporal-difference method naturally falls out of the framework and is also known as on-policy learning. In this learning technique, the agent attempts to update the value function to reflect the actual state value. The method utilizes this difference to adjust the value of Q(s, a). This method initializes the Q-value estimates to zero explicitly and then iteratively updates the value following the update rule. In each iteration, the algorithm samples an episode while performing the policy to visit the states, actions, and receive the rewards. In the mathematical expression, the Q-value Q(s, a) is modified to: where $0 \leq \alpha \leq 1$ is the learning rate, which determines the impact of the new knowledge on the existing estimate. In general, the method reflects the fact that the estimated function has a direct relationship to the true value function. Therefore, the state-action value function is modified toward the best current estimate of the future return. The discount factor represents the proportion of the future rewards collected by the agent in the long or short term. The average reward level is denoted by $R(s, a) = E[rt | St = s, At = a]; s, a \in S \times A$.

4.1. Reinforcement Learning

4. Discussion 4.1. Reinforcement Learning Reinforcement learning (RL) is a subfield of artificial intelligence that focuses on enhancing the autonomous decision-making capabilities of agents. By implementing an interactive learning process based on rewards, RL algorithms allow decision processes to freely explore the best actions within a certain context. Throughout this learning curve, the agent can observe the rewards coming from the environment as a direct outcome of the actions carried out. Consequently, they can find the optimal strategy to

solve a certain problem by acting based on this feedback. The receiving state also depends on some internal weight factor, which in turn depends on the previous state and the arriving input. By including this feedback loop, reinforcement learning can guarantee adaptive behaviors in complex and unknown contexts. This approach makes decision-making processes invariant over time and free from predefined rule sets. Although their correct training should be performed over a wide variety of traffic scenarios, RL-based systems show clear potential to be easily adaptable and effective in practical traffic management. Reinforcement learning becomes particularly effective in cooperative and/or adversarial multi-agent systems. Some approaches try to learn or maximize global return, which is a combination of rewards that comes from the decisions of many agents, while others implement communication protocols to exchange information in order to optimize local rewards. A basic approach for managing the routing of vehicles is represented by deep Q-networks (DQN) as two-compartment models. One part is implemented by a convolutional neural network (CNN), which is used to approximate the Q-function, while the other part is represented by a special kind of memory, called replay memory, used for storing the tuples to optimize CNN during the training phase. A simulator was built, which utilized the low-level controller of a vehicle simulator, while the high-level controller was replaced with RL DQN for coordinated vehicle driving. An important issue related to the training processes of concerned DQN models refers to the need to be able to face a large variety of operational conditions, from very good weather conditions to odometric sensors blocked by dirt or snow. This seems to be the reason for the use of simulator communities in recent years, which are also able to simulate a variety of communication systems. The functionalities of explored RL movements seem to be sufficiently significant and interesting as lessons learned for the efficient design of vehicle routing techniques. Instead of collision avoidance, the use of reinforcement learning regards the selection of the vehicle's route toward the target. As routing agents, the selected vehicles running in a highway environment, using an adaptive cruise control (ACC) mechanism as default. Each agent has to choose between passing in front of the other ones or changing lanes and overtaking. Results show an efficacy that strongly depends on the studied scenario. Reinforcement learning could represent a feasible solution for the dynamic adaptation of vehicle routes toward the user-defined targets in a traffic environment. This assumes an up-to-date routing engine, capable of keeping track of the changing conditions and contributing dynamically to generate route-data responses based on the current traffic and weather conditions. It goes without saying that recent technological improvements in the field of sensors or communication systems will likely increase the level of more complex and coordinated traffic movements. Consequently, smart learning algorithms will be more important in order to fully exploit these potentials.

4.2. Deep Learning

Deep learning, as a part of machine learning, encompasses several approaches to learning from vast data. Its main purpose is task-driven pattern recognition on a large scale, using a deep hierarchy of layers to learn a multilevel representation in hidden units. One of the most widely used deep learning methods is a Convolutional Neural Network (CNN) in the case of data that are grids or would have been shaped as grids if they had been regular. This approach is heavily utilized in processing visual and auditory stimuli. Currently, CNNs are a basic tool in automatic image recognition; efficient utilization of these methods in sensor sprite recognition and localization can be found in the development of modern autonomous vehicle technologies.

Deep learning makes it possible to process sensor data within an autonomous vehicle quickly and with no manual human intervention. The ability to understand the environment around a vehicle is crucial for several subtasks, such as interpreting paths or routes, obstacle avoidance, and trajectory planning when organizing a fleet of moving assets within a seaport facility. The learning stage, performed as an offline preparation, is a time-consuming process and requires expensive computational resources; thus, deep learning is not suitable for use in real-time operation. The parallel approach to path or route identification is an online adjustment of input and pre-processed computational algorithms utilized for autonomous vehicle movement, such as reinforcement learning. This approach facilitates the creation of autonomous vehicles capable not only of recognizing well-known paths but also of proposing new paths, providing an additional layer of flexibility that increases the resilience of the whole system. Autonomously operated fleets provide significant value to large commercial entities. However, there is no extensive literature on the technology cooperation of autonomous vessel fleet operations and deep learning. Up-to-date technologies and the capability to exchange real-time data streams cause the aforementioned technical challenges to evolve rapidly.

5. Case Studies and Applications

In this section, we cover four case studies. In each of the following applications, we discuss the chosen strategy, challenges, proposed solution, evaluation, and results. Briefly, Cricket Hopper simulated a fleet of interconnected autonomous vehicles coordinating through Model Predictive Control and developed a procedure to assess crowd manipulation techniques. As a complement to this work, a priority rule was implemented to help facilitate coordination in the Watertaxi Pad van Foreest. The Watertaxi application of QC-MDP is capable of finding sensible traveling speed and departure times. The use of GridWorld as an application was based on multi-grid reinforcement learning training to find aggregational inflow. Using a similar concept, Argo AI employs Sparse Proximal Multitask Learning to generate more realistic, traffic-aware simulations. Lastly, using a Pedestrian Activity Space Model, the UvA transport team developed a vision for modeling and simulating a mixed crowd of passengers.

Benefits, drawbacks, and limitations for each development are discussed in the conclusion. Case studies are featured in articles on logistics, the journal on efficient and safe flow of goods and people. Cricket Hopper - The application of MC-MDP for fleet coordination to simulate the movement of 26 users at Schiphol. Cricket Hopper used queuing theory to scale the demand from 26 people to 299, which is the number of occupants in the Watertaxi Pad van Foreest. We programmed a first-come, first-served model and transferred the optimization problem into a Q-cost problem. We compiled the code and computed a so-called 'optimal' queuing time minimizer for the current number of available taxis. To counteract the crowd, we then found the set Q of hu(t) such that the GC-criteria strategy yields at most 0.3 for the crowd manipulation.

5.1. Real-World Implementations

Real-life implementations of machine learning algorithms to coordinate autonomous vehicle fleets have been carried out in recent years in the automotive, public transport, and logistics sectors. Some representative examples are described and analyzed next.

The machine learning algorithms behind Sections 3 and 4 have been implemented in two functional prototypes capable of controlling real autonomous electric vehicles. Both prototypes operate under different SDCs supported by a specifically designed middleware with cloud integration.

The ITS UTCAM is a broker device capable of processing cloud information that becomes an enhanced plan showing new opportunities or another option. The brokerage algorithm is an artificial intelligence (AI) expert system programmed in Prolog that has a database of facts and rules. As the number of facts and alternatives increases, a decision becomes more complex than human brain capacity. Instead of reasoning each fact and alternative, we need a technique to prune the search space following a greedy approach. They have designed a convolutional neural network (CNN) combining a few AI techniques, aiming to diverge human thinking and actions. The associations between some facts are a matter of expression; their weight is subjective and dynamic, i.e., related to the individual's personal context. Unlike implicit expertise, the database aims to make the decision considering (a) usability, (b) robustness, and (c) error performance. In order to decide through implicit expertise, we apply shallow CNN to filter certain facts or outcomes from the database. Our CNN contains three layers, with the first layer enabled to have hierarchical features. Every feature within CNN has the potential to connect to different parts of the database (facts and alternatives). The links are functional blocks as well as learning to reduce uncertainty. The prototype was designed to operate as FCDAI during operations in the city center where the electric buses interchange with riders. Overall, a learning of experiences in the field led to the conclusion that ambitious schedules and times for the field implementation had to be extended. Moreover, a gap in the system integration was identified to deduce that the service provider is not considering UTCM's vehicle occupancy efficient options.

The Traffic Advisor (TA) is an innovative service where an agent supports people's decisions by seating clever cars "X" together. The Multimodal Traffic Advisor (MTA) goes beyond as it is designed to look simultaneously for different transport modes. The TA/MTA will communicate with clever autonomous pods and with other CarXs. The task of our machine learning model is to determine in advance the "community of intelligence" and to which autonomous pods they should allocate when they address the MTA system. The Almaty visitors provided a two-year transversal study using students of the host university as a data source. Based on the feedback from the study and simulated functionalities carried out during the pandemic period, the MTA will be field-integrated during the international visitors' experiences.

5.2. Benefits and Limitations

5.2. Benefits and Limitations. Coordinating autonomous vehicle fleets would bring a number of benefits but also reveal several challenges. From a social perspective, vehicles would move more efficiently through better use of road infrastructure. This would yield lower operational costs, for example, less energy used and less waste generated. Coordinated vehicles can also be safer as they may make better-informed decisions to avoid aggressive maneuvers, given that machine learning techniques can reveal predictive analytics. From an environmental perspective, research has been looking at solutions for reducing expected carbon emissions in the road sector. Learning vehicle behaviors will lead to less aggressive or smoother vehicle operation, which results in a slide-to-stop reduction in maximum longitudinal deceleration, reduces the number of collisions that will occur, and mitigates some of the harmful effects of such construction.

Machine Learning for Autonomous Vehicle Fleet Coordination. From a technical point of view, to provide global situational awareness, agents will need to select part of the information received from their environment. As a result, privacy concerns and the possibilities of biases in the system must be evaluated. It is also important to note that psychological profiling—be it intended or incidental—is a step outside the typical role for an automated vehicle system, and such capabilities create unique challenges and considerations in developing machine learning for use in transportation networks. Data storage and computations are also still challenging given the large amount of data generated by the vehicles. This data also need to be complexly processed and analyzed. Other limitations include the need for advanced sensors onboard the vehicles, such as communication capabilities, and sufficient computational resources. The applicability of any machine learning algorithm to fleet coordination also relies on the quality of the data and its ability to generalize from past experiences to new situations. It should also be noted that machine learning models are "black boxes," and any controller learning based on their output may have difficulty relating it to system states. Finally, the vehicle operator must place trust in the coordinated vehicle fleet for the technology to function. Trust will need to be scalable, sophisticated, earned, manageable, authentic, protected, adaptable, and resilient. There are still no survey results quantifying trust levels in the public at the time of this review. Starting a new research field can meet societal-level resistance or blowback from deployed users if the socioeconomic costs are not explained or are not thought to be acceptable.

6. Future Direction

The future of research and development of machine learning for autonomous vehicle fleet coordination presents several opportunities. The technology of deep learning is evolving at a rapid pace, and in the near future, algorithms will be developed to facilitate the continual learning that current algorithms cannot handle. This will involve the continued acquisition of data as well as the evolution of artificial intelligence. Machine learning techniques are being developed for transferring learned skills and representations to new tasks that are not shown during training. As a result, the efficiency of an autonomous vehicle will evolve not just with experience, but also in conjunction with the continual advancement of machine learning algorithms. There is also potential for further integration of autonomous vehicles with transportation infrastructure. Indeed, machine learning for autonomous vehicle fleet coordination could be developed in conjunction with smart cities and their policies.

Machine learning techniques for autonomous vehicle fleet coordination are in their infancy, and there are numerous areas in need of further research. Chiefly, the issue of resource allocation between the autonomous vehicles in the fleet must be explored. Significant unknown factors could also be addressed in a more detailed model, such as the risk of taking on passengers or transporting goods, or the level of congestion within a city. The concern of ethics, in addition to the formidable regulatory framework required, is another current research area in need of study. Government, the motoring industry, and the automotive industry regulators must continue to work together to create engines, infrastructure, and internal safety policies in the most important fields. Industry stakeholders are also encouraged to collaborate on standardizing hardware, firmware, and more advanced fleet-coordinating software, in accordance with current best practices. Public acceptance, security, and the robustness of the algorithms specifically for autonomous vehicle fleet coordination will be difficult. Some familiar statistics provide an unclear roadmap for the future of autonomous vehicles, but a large number of the world's leading technology and automotive companies are already experimenting extensively with the technology. Further research is needed to explore how autonomous vehicle fleet coordination can become a worthy investment.

7. Conclusion

In the following text, we presented our results and insights obtained from investigating the coordination of autonomous vehicle fleets with machine learning for decision-making and control. The introduction reminded us the reasons why coordination is a necessity for AV fleets, and current challenges. We then presented some theoretical approaches and our applications-oriented approach to various problems including fleet rebalancing, platooning, and intersection crossing permission negotiation. Safe, efficient and trustworthy transportation systems of the future are impossible without addressing the coordination challenges of future autonomous vehicle (AV) fleets in advance. One of the primary contributions of this paper is to reaffirm the justification of our investigation, starkly supported by the fact that our world foresees a variety of upcoming transportation modes and systems which urgently need coordination tools and methods. Hundreds of billions of dollars' worth of investment is being made as the coming overtures of the fourth industrial revolution begin to reveal themselves. As more and more people bring home money from the spectrum of multiple employment opportunities, the goals, behaviors, and expectations of the general population will rapidly accelerate. The skills/hobbies combos that they will bring into the future workforce have never been tested by time. The life experience generation of the future workforce is anything but predictable. Given this reality, we must rethink how we train people to be effective workers in the coming age of uncertainty by those who hold the mantel of destiny with regards to training and employment. This statement contributes importance to our already-coordinated study. The need for further investigation cannot be overemphasized. Resounding and precedential transformations are expected to be forthcoming for the automotive industry as the economy of scale-based growth is materialized. In future systems, planning and management of large AV fleets according to the requirements of the entire transportation system or society is essential. Close-knit collaboration is needed between all stakeholders to understand these facts and requirements, aspects of AI and machine learning that provide us with tremendous support, responsibility, and ethical questions that cannot be ignored. We call for a new era of intentionally planned, ethical, and conscious AI and machine learning development.

Reference:

1. Tamanampudi, Venkata Mohit. "Automating CI/CD Pipelines with Machine Learning Algorithms: Optimizing Build and Deployment Processes in DevOps Ecosystems." *Distributed Learning and Broad Applications in Scientific Research* 5 (2019): 810-849.

2. Pal, Dheeraj Kumar Dukhiram, et al. "AI-Assisted Project Management: Enhancing Decision-Making and Forecasting." Journal of Artificial Intelligence Research 3.2 (2023): 146-171.

3. Kodete, Chandra Shikhi, et al. "Determining the efficacy of machine learning strategies in quelling cyber security threats: Evidence from selected literatures." Asian Journal of Research in Computer Science 17.8 (2024): 24-33.

4. Singh, Jaswinder. "The Rise of Synthetic Data: Enhancing AI and Machine Learning Model Training to Address Data Scarcity and Mitigate Privacy Risks." Journal of Artificial Intelligence Research and Applications 1.2 (2021): 292-332.

5. Alluri, Venkat Rama Raju, et al. "Serverless Computing for DevOps: Practical Use Cases and Performance Analysis." Distributed Learning and Broad Applications in Scientific Research 4 (2018): 158-180.

6. Machireddy, Jeshwanth Reddy. "Revolutionizing Claims Processing in the Healthcare Industry: The Expanding Role of Automation and AI." Hong Kong Journal of AI and Medicine 2.1 (2022): 10-36.

7. Tamanampudi, Venkata Mohit. "Autonomous AI Agents for Continuous Deployment Pipelines: Using Machine Learning for Automated Code Testing and Release Management in DevOps." Australian Journal of Machine Learning Research & Applications 3.1 (2023): 557-600.

8. J. Singh, "How RAG Models are Revolutionizing Question-Answering Systems: Advancing Healthcare, Legal, and Customer Support Domains", Distrib Learn Broad Appl Sci Res, vol. 5, pp. 850–866, Jul. 2019

9. S. Kumari, "AI-Enhanced Mobile Platform Optimization: Leveraging Machine Learning for Predictive Maintenance, Performance Tuning, and Security Hardening ", Cybersecurity & Net. Def. Research, vol. 4, no. 1, pp. 29–49, Aug. 2024

10. Tamanampudi, Venkata Mohit. "Leveraging Machine Learning for Dynamic Resource Allocation in DevOps: A Scalable Approach to Managing Microservices Architectures." Journal of Science & Technology 1.1 (2020): 709-748.