

RACE: Reinforced Cooperative Autonomous Vehicle Collision Avoidance

Yali Yuan[†], Robert Tasik[†], Sripriya Srikant Adhatarao, Yachao Yuan, Zheli Liu *Member, IEEE*,
and Xiaoming Fu, *Senior Member, IEEE*

Abstract—With the rapid development of autonomous driving, collision avoidance has attracted attention from both academia and industry. Many collision avoidance strategies have emerged in recent years, but the dynamic and complex nature of driving environment poses a challenge to develop robust collision avoidance algorithms. Therefore, in this paper, we propose a decentralized framework named *RACE: Reinforced Cooperative Autonomous Vehicle Collision Avoidance*. Leveraging a hierarchical architecture we develop an algorithm named *Co-DDPG* to efficiently train autonomous vehicles. Through a security abiding channel, the autonomous vehicles distribute their driving policies. We use the relative distances obtained by the opponent sensors to build the VANET instead of locations, which ensures the vehicle's location privacy. With a leader-follower architecture and parameter distribution, RACE accelerates the learning of optimal policies and efficiently utilizes the remaining resources. We implement the RACE framework in the widely used TORCS simulator and conduct various experiments to measure the performance of RACE. Evaluations show that RACE quickly learns optimal driving policies and effectively avoids collisions. Moreover, RACE also scales smoothly with varying number of participating vehicles. We further compared RACE with existing autonomous driving systems and show that RACE outperforms them by experiencing 65% less collisions in the training process and exhibits improved performance under varying vehicle density.

Index Terms—Autonomous Driving, Deep Reinforcement Learning, Collision Avoidance, Privacy, VANET

I. INTRODUCTION

The 2018 global road traffic accident report from the World Health Organization [1] revealed that an estimated 1.35 million people died worldwide due to road accidents. Even though the rate of death relative to the world's population remains a constant, the United Nations' goal to realize 50% reduction in road accidents by 2020 remains a distant dream. Moreover, there is no observed reduction in road accidents in low income countries. The cost to endure traffic accidents per year is approximated to 518 billion US dollars [2], [3]; approximately 1% of the gross national product (GNP) in low income countries. Most of the road or traffic accidents are caused

by collisions [4]. Substituting the error-prone human driving with autonomous vehicles using artificial intelligence has the potential to circumvent the collision problems. Nevertheless, developing efficient computing models for robust autonomous vehicles that replicate the desired driving behaviour to avoid the collisions is still an ongoing research challenge [5].

Widespread research in the field of collision avoidance in autonomous driving has led to numerous solutions ranging from control theoretic formalization [6] and optimal control methods [7] to potential field-based and rule-based techniques [8]. Recent advancements in the field of machine learning with imitation learning and deep reinforcement learning have facilitated improved data-driven approaches to minimize collisions. For example, in [9], Min et al. predicted collisions using neural networks with optimized genetic algorithms and back propagation and in [10], Chen et al. explored collision avoidance using deep reinforcement learning with multiple agents. However, these approaches cannot effectively scale to accommodate the increasing road safety requirements. Further, many collision avoidance algorithms use precise information (*e.g.*, location) of the agents¹ [11], [12] to avoid collisions and thereby compromise the privacy of participating agents during autonomous driving.

Moreover, we observed that, autonomous vehicles learn their surrounding environment and optimal driving strategies through a large amount of onboard sensors, such as lidar, radar and acceleration sensors. However, because of the limited communication range of the vehicles, the collected information from one vehicle is not sufficient to fulfill large-scale road safety requirements and improve the collision performance efficiently. Therefore cooperation among the vehicles is a promising solution to overcome the above-mentioned shortcoming. In cooperated vehicle communication, the vehicles can communicate with each other. Each vehicle can learn their neighbor's surrounding environment. As a result, the cooperated vehicles communication enables the vehicles to incorporate essential environmental information and thereby improve their autonomous driving behaviors. However, Kaushik et al. [13] show that, cooperated vehicle communication strategy faces the challenges of enforcing vehicles within lanes and prevent unnecessary over-takings. They proposed parameter sharing to resolve these issues, but failed to take the scalability of their solution into account since with increasing number of agents the collisions increased exponentially. Some recent works [14], [15], [16] also proposed centralized solutions,

Yali Yuan, Robert Tasik, Sripriya Srikant Adhatarao and Xiaoming Fu are in Computer Networking Group, Institute of Computer Science, University of Goettingen, 37077 Goettingen, Germany. Emails: yali.yuan, fu, sripriya-srikant.adhatarao@cs.uni-goettingen.de, robert.tasik@stud.uni-goettingen.de.

Yachao Yuan is in Smart Mobility Research Group, Institute of Information Systems, University of Goettingen, 37075 Goettingen, Germany. Email: yachao.yuan@uni-goettingen.de.

Zheli Liu is in college of computer and control engineering, Nankai University, Tianjin, China. Email: liuzheli@nankai.edu.cn.

[†] Yali Yuan (corresponding author) and Robert Tasik have equal contributions and are both first authors.

¹Please note we use the terms agent and vehicle interchangeably.

where a server deployed at the central location computes efficient collision avoidance strategies for all participating agents. The centralized server maintains the comprehensive knowledge about all the agents as they send their current state information to the server and wait for the most intelligent driving strategy from the server. However, such centralized algorithms cannot effectively scale under heavy workload, especially when many agents are simultaneously requesting for driving instructions. Moreover, such systems heavily depend on the availability of a secure and private communication network between the agents and the server, which is hard to guarantee in an agile driving environment [17]. Therefore, further research is required to provide optimal autonomous driving solutions that satisfy the desired driving requirements, especially efficient collision avoidance solutions which also fulfill the privacy needs of the users [18], [19].

Based on the above discussion, we propose **RACE: Reinforced Cooperative Autonomous Vehicle Collision Avoidance**, a robust and efficient autonomous driving framework to overcome the above mentioned shortcomings. In RACE, we provide a hierarchical multi-agent deep reinforcement learning model that exploits the well-known Deep Deterministic Policy Gradients (DDPG) [20] algorithm to effectively avoid collisions during autonomous driving. We leverage the cooperated vehicle communication model, and develop a private and dynamic ad-hoc vehicular network algorithm based on VANETs, to build hierarchical vehicular networks with leading and following agents. Using the distribution of parameters from the leading agent to the following agents, we foster a resource efficient real-time deep reinforcement collision avoidance algorithm.

RACE is designed to meet real-world autonomous driving design goals and hence it can rapidly adapt to the changes in the driving environments, such as sparsity or density of vehicles at intersections, and reduce the number of collisions. During learning, RACE enforces varying levels of stringent penalties when collisions or potential situations, which may lead to collisions, occur; in order to build a robust learning environment. Employing the TORCS [21] simulator, we implement RACE in a non-deterministic real-world car driving scenario in a racing environment. With evaluations using the TORCS simulator, we show that RACE effectively learns from the environment and its neighboring agents. As the agents learn, we observe an exponential decrease in the number of collisions and a stable increase in rewards during learning. We compare RACE with existing baselines from Pinxin et al. [22] and show that RACE outperforms existing works and dramatically reduces the collisions by 65% and improves resource utilization by exploiting the multi-level agents. Interestingly, we also observed that due to the varying levels of penalties designed in RACE, with sufficient learning, the vehicles chose to go off track in a potential collision scenario rather than collide with other vehicles or nearby objects to minimize the risk and liability. Thus, RACE provides a robust learning environment with the proposed Co-DDPG algorithm, which converges quickly and distributes the best parameters among the participating agents. The main contributions in this work include:

- A novel hierarchical collision avoidance framework named *RACE* to administer an efficient autonomous driving environment. Through leveraging the deep reinforcement learning and the cooperated vehicle communication, RACE provides an efficient autonomous driving system that achieves the desired driving goals and ensures the location privacy of vehicles.
- A dynamic hierarchical ad-hoc mobile network creation algorithm based on VANETs to create and maintain vehicular networks during autonomous driving. The network is used to communicate necessary information to the participating agents in order to ensure smooth driving experience.
- An improved deep reinforcement learning algorithm named *Co-DDPG* to efficiently learn the driving patterns. Employing parameter distribution, the learned driving patterns are communicated with neighbors to utilize the resources efficiently and share the learning experiences among agents.
- Extensive evaluations using TORCS simulator to demonstrate the benefits of RACE in comparison to existing solutions in an autonomous driving environment.

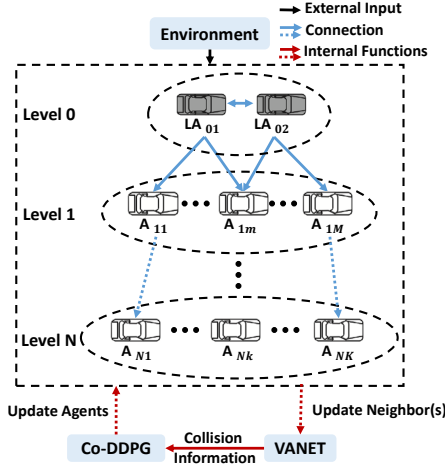
II. SYSTEM DESIGN

In this section, we elaborate the design goals of autonomous driving and the rationale behind them. This is followed by a detailed description of the RACE framework and a use case example. Figure 1 shows an overview of the proposed RACE framework along with its internal system components to realize the design goals of an autonomous driving system.

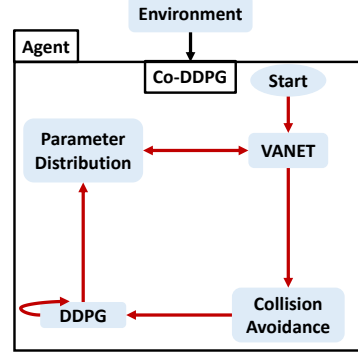
A. Design goals

An autonomous driving system should fulfill the following design goals:

- Risk mitigation: It is crucial for an autonomous driving system to avoid any potential collision and damage in order to offer safe transportation with minimal risk.
- Robustness: An autonomous driving system should be able to overcome technical difficulties arising during driving and opt for safer alternatives to avoid impending danger.
- Uncertainty: An autonomous driving system should be able to account for unpredictable and unexpected complications arising from the surrounding environment and neighboring vehicles.
- Resource utilization: An autonomous driving system should optimally consume the available resources, such as computational resources and energy. This is crucial to ensure an efficient and safe driving experience. Moreover, many resources are often depleted at a faster rate during driving and are often difficult to replenish in some scenarios.
- Efficiency: In autonomous driving, with increasing number of agents the complexity of machine learning models also increases. The learning algorithm should be able to scale easily and handle the learning complexity efficiently in order to generate accurate driving actions in real-time.



(a) RACE Framework.



(b) System Overview of Agent.

Figure 1: RACE Architecture.

- **Location Privacy:** As the information is distributed in an open access environment in VANET, drivers' location privacy is critical and has to be satisfied to provide a reliable system.

B. Architectural Components

The proposed RACE framework, which satisfies the above mentioned design goals is illustrated with an example scenario in Figure 1 along with an overview of the system in every autonomous vehicle. The framework in Figure 1(a) is composed of N layers organized hierarchically where, LA_{ij} represents the leading agent LA_j in the i th layer. Correspondingly, $A_{(i+1)m}$ is the following agent A_m in the $(i+1)$ th layer of the leading agent LA_{ij} . The agents at level 0 are always considered as leading agents, whereas the agents at subsequent layers can take up the role of both leading and following agents based on their level in the hierarchy. The main components of this framework are described as follows:

- **Environment:** It represents the collective information regarding various road conditions gathered by the on board sensors in the autonomous vehicles *e.g.*, distance towards the track edge, speed, angle, etc. We build dynamic VANETs to facilitate the agents to communicate with each other and update their environmental conditions.
- **VANET:** This component creates a dynamic vehicular ad-hoc network, *i.e.* VANET to rapidly find any nearby vehicle(s) and create a network for communication and efficient driving. Essentially, every agent identifies its neighboring autonomous vehicles and adds them to its local list of neighbors to which it is already connected. This module is mainly responsible for opening, leaving and closing the connections of an agent with its neighbors in a VANET.
- **Hierarchical Agent Network:** This is a submodule of the VANET. It is created to dynamically allocate agents with their respective roles in the VANET. The roles are designated based on varying levels of their available resources to ensure efficient operation. This hierarchical network is composed of many layers with various agents found on different layers. However, we can broadly




Figure 2: Parameter distribution.

classify the agents into two categories, namely leading agents and following agents. A leading agent represents an autonomous vehicle at a layer with maximum available resources in comparison to its following agents located at the lower layers. Every agent stores its corresponding following agents in a list similar to the neighbors list used during the VANET creation.

- **Collision Avoidance:** This component tries to minimize the collisions during driving by allotting penalties for undesired driving behaviours such as crashing or driving too closely with other participating vehicles and rewards for good driving behaviours like keeping on the road, avoiding collisions and minimizing the arrival time.
- **Co-DDPG:** Cooperative Deep Deterministic Policy Gradients (Co-DDPG) algorithm is an enhanced version of the standard DDPG algorithm designed in RACE. Along with autonomous learning, it also facilitates collision avoidance with neighboring agents during deep reinforcement learning within a VANET.
- **Parameter Distribution:** The parameter distribution enables leading agents in RACE to accumulate deep learning parameters through Co-DDPG, such as weights, biases and rewards and distribute them to their following agents.

C. Parameter Distribution

One major component in the system design of Co-DPPG is the *Parameter Distribution* module. In RACE, the parameter

distribution is facilitated by the creation of VANETs, which ensure that all participating vehicles can communicate with each other. The crux of parameter distribution is based on comparing the *average reward* accumulated by each agent for every action executed during autonomous driving. As this reward in reinforcement learning represents the incentives for good driving behaviour, it enforces and enriches the desired learning process. Therefore, we consider the average reward as a suitable metric to design the parameter distribution in RACE. Please note that in RACE, the following agents can choose to either learn or wait for distributed parameters from their leading agent(s) based on their available resources. If a following agent chooses to wait, it updates its parameters with the distributed parameters from a leading agent with highest average reward. Otherwise, a following agent compares its average reward based on its learned parameters with that of the leading agent. The following agent will update its parameters with the distributed parameters, if the leading agent's average reward is higher than its current average reward. Otherwise, the following agent will forward its parameters to the leading agent to ensure that all agents learn the best policy. For instance in Figure 2, Agent 1 and Agent 2 distribute two different set of parameters to their following agent *i.e.*, Agent 3. Based on the average rewards associated with these two parameter sets, Agent 3 will select the parameters with the highest possible reward and incorporates it into its learning model.

D. Design

We acknowledge that it is challenging to implement a multi-agent autonomous driving learning environment. This is mainly due to the increasing complexity and uncertainty of the participating agents and their responses to other agents and the varying environmental conditions. Each agent explores the learning environment with an aim to find the best policy under current circumstances and thereby increase the complexity of its learning. Subsequently, this also influences the learning complexity of other participating agents. With increasing number of agents, the resulting resource consumption also increases and thus, an inefficient learning environment will negatively impact the desired goals of autonomous driving. Therefore, we introduce an efficient decentralized learning algorithm named Co-DDPG (see Section IV) based on deep reinforcement learning and parameter distribution in RACE to ensure adequate learning in a multi-agent environment that satisfies the desired learning goals.

The RACE framework shown in Figure 1(a) employs a hierarchical agent network, where each agent runs the Co-DDPG algorithm with its supporting system components shown in Figure 1(b). The Co-DDPG algorithm enables reinforcement learning parameters such as weights and biases to be shared among all the connected agents in the VANET. The initial parameters are distributed by a primary leading agent that instantiates the vehicular network. For instance in Figure 2, an initial network was created by the leading Agent 1 and leading Agent 2 and thus, they distribute the initial learning parameters to their subsequent following agents *i.e.*, Agent 3. Each agent's model parameters are updated whenever there

exists a better trained model in the network. With such a learning behaviour, all agents in the VANET benefit from distributed and learned experiences with each other. Some agents may receive parameters from two or more leading agents *e.g.*, Agent 3 in Figure 2 received two parameter distributions, one each from Agent 1 and Agent 2. In such a scenario, the agents will select the best parameters from the received choices. This ensures that, in RACE always the best collision avoidance parameters are shared among the agents to maintain a safe distance between the participating vehicles and to stay on the road by avoiding the edges of the road. Specifically, every agent in RACE is organized in to a hierarchical leader-follower architecture. The leading agents learn from the current environment and distribute their learned parameters, specifically the weights and bias, to their following agents. The following agents can choose to either learn from the environment similar to their leading agent(s) or wait and download the parameters from their leading agent(s) based on their available resources. If the following agents choose to learn from the environment, they compare their learned parameters with the received parameters from their leading agents and store the best policy. If the following agent learns the best policy, then the following agent communicates the best policy to its leading agent. This ensures that all agents in the current environment learn the best possible driving behaviours.

Essentially, with N number of layers organized hierarchically, the set of all layers is represented as $\mathcal{L} = \{L_0, L_1, \dots, L_i, \dots, L_N\}$. If the agents in layer L_i are designated as leading agents, then the agents in the following layer L_{i+1} are considered to be the followers of the agents in the layer L_i . The number of agents in each layer is denoted as M , where $M \in [0, N]$ and M, N are integers. The set of all agents in layer i is given as $\mathcal{V} = \{A_{i1}, \dots, A_{im}, \dots, A_{iM}\}$. Initially, two agents in the top most layer are selected as the leading agents. The agents run the Co-DDPG (see Section IV) learning algorithm to find the optimal driving strategy and how to interact with other neighboring vehicles. The following agents in the communication range of the leading agent estimate their remaining available resources and choose to either learn by running their Co-DDPG algorithm or wait for their leading agents to distribute the parameters. After the leading agents distribute their trained parameters, the following agents select the parameters from the leading agents and updates their current Co-DDPG parameters with the received parameters. Since we optimally organize the agents based on their available resources into leader or follower and further provide the followers with an option to wait in order to conserve their remaining resources, RACE guarantees optimal utilization of the available resources.

E. Use case Example

Let us consider a car driving on a road as an example use case to describe the operation of a system implemented using the proposed RACE framework. Initially, the road is empty, and at time T_i , where $i \in [0, N]$, a single agent, Agent 1, enters the road and starts driving. Since there are no neighboring agents to connect and synchronize the driving,

Algorithm 1 checkVANETradius()

```

1: for Neighbor in vehicle.Neighbor do
2:   if opponent < VANET_RADIUS then
3:     if not vehicle.openedVANET and not vehicle.insideVANET then
4:       vehicle.openVANET()
5:     else
6:       vehicle.identifyNeighbors()

```

Agent 1 learns its best driving behaviour primarily from Co-DDPG. At time T_{i+1} , a second agent, Agent 2, enters the road and soon comes in the connection range of Agent 1. Agent 1 and Agent 2 initialize the VANET module and thus create a VANET and assign the roles of leading and following agents based on their available resources. This process is repeated each time, a new vehicle comes in the VANET radius of the vehicles already in the VANET. The resulting neighbors are added to the neighbors list in every agent and used for communication as the learning continues. The participating agents continuously compute the distance between each other using the values from their neighboring agent's sensors and maintain a safe distance to avoid collision. The DDPG module in Co-DDPG collects the necessary state information from the environment, VANET and collision avoidance modules and generates the optimal driving actions along with the associated reward under current conditions. Using the parameter distribution module, the computed parameters from the leading agents are distributed to their following agents. Whenever a collision occurs, or if the vehicles overstep their safe distance threshold, a penalty with negative reward is awarded for that action. With a hierarchical networking architecture, RACE unites all the necessary autonomous driving components in Co-DDPG to provide an efficient learning platform that enhances the existing autonomous driving learning algorithms.

III. VANET

In this section, we describe the dynamic vehicular ad-hoc network creation algorithm, *i.e.*, the VANET component shown in Figure 1(b).

A. Opening, Leaving and Closing VANETs

The basic objective of a VANET during autonomous driving in RACE is achieved with three main functions: `openVANET()`, `leaveVANET()` and `closeVANET()`. If a vehicle scanning for neighbors identifies another vehicle in its VANET range, then a VANET connection is opened using the `openVANET()` function to initialize the connection. Contrary to `openVANET()`, as soon as a vehicle moves out of the VANET's range, this vehicle leaves the network using `leaveVANET()` function and disconnects with its neighbors located in that network. Subsequently, the neighbors of this vehicle remove its information from their list of connected vehicles. When all except one vehicle leaves the VANET, then the scanning vehicle closes the VANET using the `closeVANET()` function.

Algorithm 2 identifyNeighbors()

```

1: identifiedNeighbor = False
2: for VehicleX in LocalVehicleList do
3:   for VehicleA in LocalVehicleList do
4:     if VehicleX != VehicleA then
5:       if VehicleX not in VehicleA.Neighbors &&
         Distance between A and X < VANET_RADIUS then
6:         identifiedNeighbor = True
7:         Add VehicleX to VehicleA.Neighbors
8:         Add VehicleA to VehicleX.Neighbors

```

B. VANET Operation

In order to create a VANET, it is fundamental to identify the vehicles driving in the area of interest. Hence in RACE, we declare the desired VANETs range in meters and utilize the on-board sensors in the vehicles to compute the distance between vehicles located in this range.

In Co-DDPG, we define an algorithm called `checkVANETradius()` to find any neighboring vehicle in the specified VANET radius of a vehicle. A pseudocode of the algorithm is given in Algorithm 1. Essentially, every vehicle scans with `checkVANETradius()` the VANET radius *e.g.*, 200 meters to find any nearby vehicles to connect with. Since we follow a decentralized approach, every vehicle in RACE maintains a local list of connected vehicles called `LocalVehicleList`, which is a subset of the global list of all participating vehicles. If a vehicle is found within this range and there is no available VANET to connect with, then the scanning vehicle initially opens a VANET using the `openVANET()` function. This is followed by opening a connection with the identified neighbor for communication. This is accomplished by calling the method `identifyNeighbors()` shown in Algorithm 2. All vehicles, continue to scan the VANET radius to find potential neighbors to connect with. Further, the `identifyNeighbors()` method is also responsible for maintaining the presently connected vehicles within a list.

In addition to VANET creation, when a vehicle enters the VANET, it takes on the role of a leading or a following agent with each of its neighbors based on its available resources. Specifically, a new agent joining the VANET notifies its neighbours that it would like to be a leading or a following agent while its neighbours can choose to be either its leader or be a follower of the new agent. If a neighbouring agent is already a follower, then it notifies the new agent that it is already a following agent and provides the list of its leading agent(s). Since each vehicle available resources can vary, a following vehicle at level i may become a leading agent for a vehicle at level $i + 1$ if the resources of the agent at level $i + 1$ are less than that of agents located at the level i .

Utilizing the values gathered by the on-board sensors, the `identifyNeighbors()` method verifies the distance between the host vehicle and its neighbors and sends it to the collision avoidance component, which in turn sends this information to the Co-DDPG approach to learn collision avoidance during autonomous driving. We implemented an autonomous driving system using the RACE framework in the TORCS simulator by primarily using the *opponent range sensor*. According to the TORCS user manual [23], the op-




Figure 3: Distance computation in VANETs.

ponent range measurement is achieved using 36 minor range on-board sensors located on each agent to cover the complete 360 degree space around the vehicle. A pictorial representation of the opponents range sensor in TORCS is shown in Figure 3. Each of these sensors has a sensing range of up to 200 meters and span a coverage area of 10 degrees. Each sensor returns the distance between the hosting vehicle and any nearby vehicle(s) within its sensing range.

IV. Co-DDPG

In this section we describe the core of the proposed autonomous driving strategy designed in RACE, *i.e.*, the system model for deep reinforcement learning in Co-DDPG.

Due to the compelling design of the standard DDPG [20], we utilized it as the basis for learning in RACE. DDPG can efficiently learn competitive policies for all tasks considered in an autonomous driving environment using low-dimensional observations and handle continuous action spaces effectively. However, we model the autonomous driving learning strategy through enhancing the standard DDPG algorithm with hierarchical multi-agent parameter distribution to further improve the performance of learning. We refer to this enhanced version as Co-DDPG in this work and a pseudocode of the algorithm is provided in Algorithm 3. The actor network's update (line 23 in Algorithm 3) is based on the application of chain rules computed by the loss function (line 22 in Algorithm 3). The target network's weights are then updated according to $\theta' \leftarrow \tau\theta + (1 - \tau)\theta'$, where $\tau \ll 1$. The overall benefit in a multi-agent autonomous driving environment is achieved with the efficient collaboration of the surrounding components in the system design of Co-DDPG as shown in Figure 1(b). Co-DDPG receives input from the VANET and Collision Avoidance modules to enforce efficient driving behavior. We further improve the learning by allotting positive rewards for good driving behaviour and penalties for bad driving behaviour with negative rewards. Thereby, the system model of Co-DDPG is created in RACE using three essential concepts, namely the *state*, *action* and *reward*.

A. State Space

The *state* space represents the on-board sensors in a vehicle and is thus a representation of the autonomous driving environment. The *state* is defined as $\mathcal{S} = \{s_1, s_2, \dots, s_I\}$, where

Algorithm 3 Co-DDPG algorithm

- 1: Randomly initialize critic network $Q(s, a|\theta^Q)$ and actor $\mu(s|\theta^\mu)$ with weights θ^Q and θ^μ
- 2: Initialize target network Q' and μ' with weights $\theta^{Q'} \leftarrow \theta^Q$, $\theta^{\mu'} \leftarrow \theta^\mu$
- 3: Initialize action-value function Q with random weights
- 4: Initialize replay buffer R
- 5: **for** episode = 1 to M **do**
- 6: initializeVANET()
- 7: Initialize a random process \mathcal{N} for action exploration
- 8: Receive initial observation state s_1
- 9: **for** $t = 1$ to T **do**
- 10: setSensors(s_t)
- 11: Select action $a_t = \mu(s_t|\theta^\mu) + \mathcal{N}_t$ according to the current policy and exploration noise
- 12: Execute action a_t
- 13: checkCollisionRadius()
- 14: Observe reward r_t and observe new state s_{t+1}
- 15: checkVANETradius()
- 16: **if** $t == 100$ and VANET is open **then**
- 17: distributeParameters()
- 18: Store transition (s_t, a_t, r_t, s_{t+1}) in R
- 19: Sample a random minibatch of N transitions (s_i, a_i, r_i, s_{i+1}) from R
- 20: Set $y_i = r_i + \gamma Q'(s_{i+1}, \mu'(s_{i+1}|\theta^{\mu'}))|\theta^{Q'}$
- 21: Update critic by minimizing the loss: $L = \frac{1}{N} \sum_i (y_i - Q(s_i, a_i|\theta^Q))^2$
- 22: Update the actor policy using the sampled policy gradient:

$$\nabla_{\theta^\mu} J \approx \frac{1}{N} \sum_i \nabla_a Q(s, a|\theta^Q)|_{s=s_i, a=\mu(s_i)} \nabla_{\theta^\mu} \mu(s|\theta^\mu)|_{s_i}$$
- 23: Update the target networks:

$$\begin{aligned} \theta^{Q'} &\leftarrow \tau\theta^Q + (1 - \tau)\theta^{Q'} \\ \theta^{\mu'} &\leftarrow \tau\theta^\mu + (1 - \tau)\theta^{\mu'} \end{aligned}$$

I is the total number of states. The states represent a vehicle's sensors and forms the basis for the actions taken by the agents. As mentioned before, we collected the state information from the sensors available in the TORCS simulator. Please note that sensors provided in TORCS *e.g.*, range finding sensors operate in a similar fashion as the real world sensors implemented with LiDAR, radar or ultrasonic sensor technologies. Hence, the data produced by the sensors in TORCS enabled us to build a simulated learning environment akin to driving in the real world. Due to their prominent role, the sensors formed the basis for an autonomous vehicle to conduct any action in its environment.

The following list of sensors, which are also the state space variables, are used to form the state space in Co-DDPG:

- **angle**: It is the angle between the vehicle axis and the track axis during driving with values ranging between $-\pi$ to π and is measured in rad.
- **opponents**: It is a vector containing the 36 range finding sensors from TORCS. Each opponent sensor spans a coverage area of 10 degrees with a sensing range

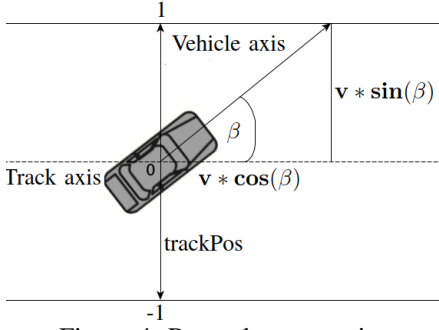


Figure 4: Reward computation.

of up to 200 meters. The vector returns the closest vehicle in the scanned area.

- **rpm**: It stands for the rotation per minute executed by the vehicle's engine.
- **v**: It denotes the vehicle's speed within the vehicle's longitudinal axis in km/h.
- **speedY**: It denotes the vehicle's speed within the vehicle's transverse axis in km/h.
- **speedZ**: It denotes the vehicle's speed within the vehicle's z-axis in km/h.
- **track**: It is a vector containing 19 range finding sensors that gather readings from the front of the vehicle with 90 degrees left and right from the vehicle's axis. Each sensor returns the distance between the vehicle and the track edge within a range of 200 meters.
- **trackPos**: It denotes the distance between the vehicle axis and the track axis. It returns a value of 0 when the vehicle is found on the track axis, -1 when the vehicle is on the right track edge and 1, when the vehicle is on the left track edge.
- **wheelSpinVel**: It represents a vector of four sensors capturing the rotation speed of the wheels.

B. Action Space

The *action* space is a three dimensional vector consisting of acceleration (i.e. accel), brake and steering. It is denoted as $\mathcal{A} = \{a_1, a_2, \dots, a_H\}$, where H is the total number of actions and is extracted from the following effectors:

- **accel**: It represents the simulated gas pedal and returns the values in the range of $[0, 1]$, where 0 represents no gas while 1 represents full gas.
- **brake**: It represents the simulated brake pedal and returns values in the range of $[0, 1]$, where 0 represents no brake while 1 represents full brake.
- **steering**: It represents the simulated steering values in the range of $[-1, 1]$, where -1 represents full right brake while 1 represents full left brake with an angle of 0.366519 rad.

C. Reward

The reward function in Eq. (1), which is our objective, is defined by considering three factors including collision avoidance, minimizing the arriving time of all vehicles and keeping vehicles on the road.

$$r_i^t = w_c c_i^t + w_h h_i^t + w_o o_i^t, \quad (1)$$



Figure 5: Experimental racing track.

where r_i^t denotes the reward r of the i th vehicle at timestep t . Essentially, higher the reward value, better is the autonomous driving behavior. c_i^t , h_i^t and o_i^t are the rewards generated by avoiding collisions, minimizing the arriving time and keeping the vehicles on the road of the i th vehicle at timestep t , respectively. $w_c \in [0, 1]$, $w_h \in [0, 1]$ and $w_o \in [0, 1]$ are the weights for them to define the contribution of each factor to the final reward function, where $w_c + w_h + w_o = 1$.

Co-DDPG obtains penalties by defining a c_i^t to avoid collisions, and is given as follows,

$$c_i^t = \begin{cases} -(D_{min,i}^t - d_i^t) \cdot w_p & d_i^t \leq D_{min,i}^t, \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

where $D_{min,i}^t$ calculated in Eq. (3) is the safe distance between two vehicles at timestep t , w_p is the punished value per distance, d_i^t is the distance between two vehicles at timestep t . For instance, for two vehicles i and j , d_i^t defined as $d_i^t = \sqrt{(x_i^t - x_j^t)^2 + (y_i^t - y_j^t)^2}$. Please note that we use the opponent sensors in the vehicles to get d_i^t (See Section III-B) and $D_{min,i}^t$ is the safe distance that should be maintained between vehicles. Similar to [24], we design the vehicle deceleration rate as $a_{cur}^t \in [a_{min}, a_{max}]$, whereas the safe distance between vehicles v_i^t and v_j^v is calculated as follows,

$$D_{min,i}^t = \frac{T_s v}{A} + \frac{v_i^t}{B|a_0^t|} - \frac{v_t^2}{B|a_{max}^t|}, \quad (3)$$

where T_s is the delay for detecting an emergency by on board sensors in the vehicles. This also incorporates the autonomous vehicle's brakes response time. We set $A = 3.6$ to convert the vehicle speed v from km/h to m/s while $B = 25.92$. B is defined as $2|a| \times (A)^2$ and it converts v^2 from $(km/h)^2$ to $(m/s)^2$.

h_i^t is the reward to minimize the traveling time, defined by

$$h_i^t = \begin{cases} 10w_p & \text{arrival to the destination,} \\ w_p |v_i^t \times \sin(\beta_i^t)| & \text{otherwise} \end{cases} \quad (4)$$

where w_p is the reward value as in Eq. (2).

o_i^t is an incentive awarded to the vehicle to encourage good driving behaviour whereby the vehicle drives within the road, the factors of which are described in Figure 4. It is calculated as follows,

$$o_i^t = v_i^t \times |trackPos_i^t| + (|v_i^t \times \sin(\beta_i^t)| - v_i^t \times \cos(\beta_i^t)), \quad (5)$$

where $v_i^t \times \cos(\beta_i^t)$ depicts the value of the longitudinal velocity of vehicle i at timestep t , which we want to minimize. The term $|v_i^t \times \sin(\beta_i^t)|$ computes the transverse velocity and it is aimed to be as high as possible, since the vehicle should stay on the road. Lower values indicate that the vehicle is going off the road edge. The third term, $v_i^t \times |trackPos_i^t|$, also represents stable driving on the center of the road and hence it should be as high as possible. In Co-DDPG, whenever $|v_i^t \times \sin(\beta_i^t)|$ and $v_i^t \times |trackPos_i^t|$ are low, the reward is reduced by awarding penalties.

V. EVALUATION

In this section, we evaluate the benefits of RACE through demonstrating the performance of our system implemented using the RACE framework with the TORCS simulator. We further compare the performance of this system with existing baselines and show that RACE effectively outperforms the existing systems with respect to collisions and scalability; the pivotal requirements of autonomous driving.

A. Assumptions

The following assumptions were made during the experiments to fulfill the basic requirements of autonomous driving. Security and privacy are intrinsic components of the autonomous driving system, hence we ensure security by following the proposal from Samara et al. [25]. As for the privacy, RACE can guarantee the location privacy by using the relative distance to build VANET instead of the absolute location of the vehicles.

We observed that underlying communication technology was necessary to ensure efficient operation of the overall system. Even though many communication technologies exist and newer technologies like 5G are gaining momentum [26], [27], the RACE framework is agnostic to these technologies. Therefore, we made one assumption that any underlying communication technology that fulfills the requirements for performance, latency and parameter distribution in an autonomous driving environment could be used in RACE.

Additionally, we made another assumption that VANETs created in the autonomous driving environment were robust in nature, wherein they effectively minimized connectivity failures and supported efficient data transfer among the participating agents. This was necessary to maintain stable networks especially during long training/driving sessions.

B. TORCS Simulator

As mentioned earlier, we implemented a system using the RACE framework in the widely used *TORCS* simulator developed by Wymann et al. [21]. Over the years, *TORCS* has become a profound testing environment for autonomous driving research [28]. Intrinsically, by including the basic vehicular dynamics and physics, *TORCS* provides realistic physical surroundings that are suitable for simulating desired driving environments. The participating vehicles also known as players are characterized as *robots* in *TORCS* and they are primarily external entities loaded by the simulation. In

particular, *TORCS* allows developers to create their own artificial intelligence agents to simulate desired experimental environments. This is accomplished by a low-level API located at the driver level that enables partial access to the state of simulation including the status of the ongoing race or robot information such as agent's speed, agent's distance from the track edge or the distance to other participating agents. Nevertheless, *TORCS* limits access to most of the components in the simulation and thus poses a challenge to even partially observe basic driving problems [21]. Despite such limitations, many *TORCS* robots were created over the years by researchers and developers to further examine the paradigms of Artificial Intelligence and Machine Learning. Many of these research efforts have enhanced the *TORCS* system to a point that it is now possible to execute driving scenarios that surpasses human-level driving.

C. Experimental Setup

We perform extensive experiments to measure three most important metrics of autonomous driving namely collisions, scalability and latency of the system during driving. We use the racing environment with *CG track 2* in *TORCS* shown in Figure 5 as an example to illustrate the performance of RACE. Since *CG track 2* is one typical track example in *TORCS* [29]. The width and length of it are about 3186 m and 15 m, respectively.

We gradually introduced vehicles on to the racing track and observed their performance with respect to the effectiveness of their learning behaviour. For example, in these simulations, 10 Normal Driving Vehicles (NDVs) are running on the road. Then, we vary the number of Autonomous Driving Vehicles (ADV) equipped with Co-DDPG to test the proposed RACE framework. We set weights of the reward function in Eq. (1) as $w_c = 0.6$ and $w_h = w_o = 0.2$. These values are obtained by comparing the performance from various simulations. The reward value of w_p in Eq. (2) is set to 1000 and the communication range R of each vehicle is $R \in [0, 200]$ meter. The speed range of vehicles is from 40 km/h to 60 km/h. Other parameters' values, like steering, acceleration, brake, angle, etc., are listed in detail in the *TORCS* documentation [23]. We implemented the RACE framework in TensorFlow where the Co-DDPG used two hidden layer neural networks as a non-linear function to achieve the optimal policy. There are 300 to 400 neurons in each layer, respectively. The learning rate is $1e-4$ with a batch size of 32. We use the decentralized sensor-level collision avoidance policy for multi-robot systems named POMDP proposed by Pinxin et al. [22] as the baseline to compare the performance of RACE. POMDP uses a multi-stage deep reinforcement learning framework to help multiple robots to learn an optimal collision avoidance strategy using the policy gradient approach.

One episode in the experiment denotes one instance on the real world racing track and thus one complete race from the start line until the finish line. However, there are instances when the race ends abruptly, such as when an agent turns back due to collision or leaves the track edge. As RACE is the foundation for the experiment, every leading agent

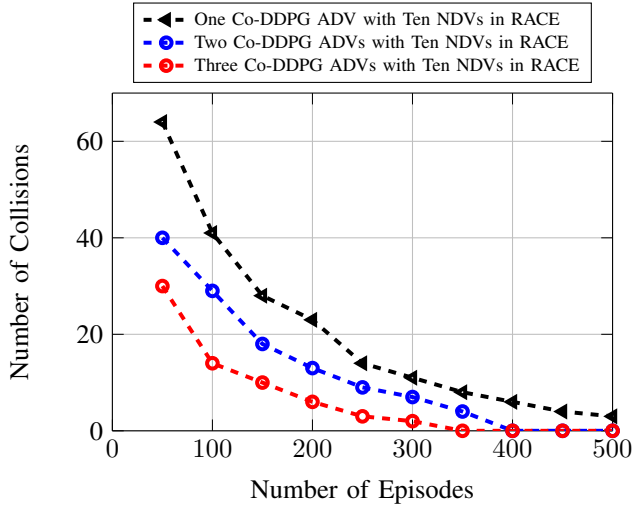


Figure 6: The total number of collisions of ADVs over different number of episodes during the training process.

trains with the Co-DDPG algorithm and shares the parameters with its following agents in the VANET. One main vehicle is depicted as the leading agent that distributes its learned model parameters within one local VANET with its following agents and the following agents remain idle and await the learned parameters from their leading agent. Since we designed a non-deterministic autonomous driving environment, we executed each experiment ten times and computed the average results from all runs.

D. RACE Performance Analysis

Using the RACE system in TORCS, the agents learn the best driving behavior during the simulations to master collision avoidance policy. The agents' goal is to dynamically enhance their driving behaviour by learning to avoid collision with other agents and objects in their vicinity. Meanwhile, they also minimize their arrival time. It is essential for an autonomous driving system to scale efficiently as the number of vehicles fluctuate. Therefore, we measure the scalability of RACE under varying density of participating ADVs. We used ten NDVs along with one, two and three number of ADVs installed with Co-DDPG to measure the performance of RACE.

The performance of ADVs is estimated by primarily examining the number of collisions experienced by the agents and the reward achieved by agents during the experiments, which is shown in Figure 6 and 7. The One CO-DDPG, Two Co-DDPG and Three Co-DDPG in the results denote ten NDVs with one, two and three ADVs operating with Co-DDPG in RACE, respectively.

In Figure 6, the average number of collisions over a span of 500 episodes during the training process is presented. The value of each collision point in this figure is calculated by summing the collision number within every 50 episodes. We observe from the results that as the number of episodes increased, the average number of collisions in all scenes are decreased. The number of collisions experienced by the Two Co-DDPG is between that of One Co-DDPG and Three

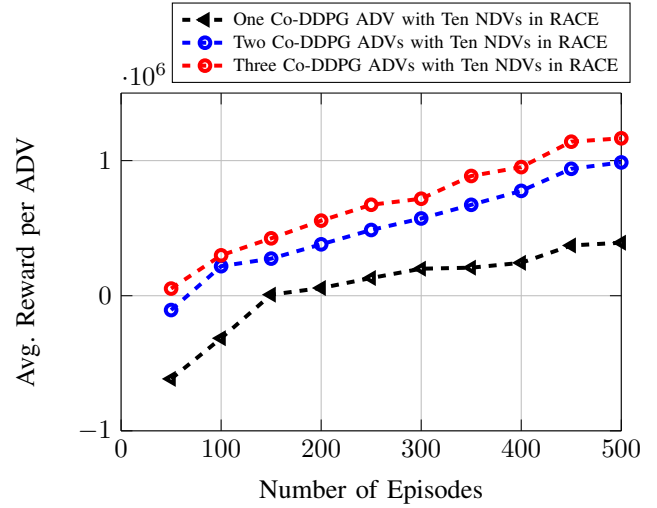


Figure 7: The average reward per ADV over different number of episodes during the training process.

Co-DDPG, as expected. Compared to Three Co-DDPG, the average number of collisions experienced by One Co-DDPG is around 211% higher. This is because in RACE, the parameter distribution strategy (See Section II-C) in Co-DDPG is employed to improve the performance of ADVs. We select the best driving policy for each vehicle by comparing the Co-DDPG rewards of neighboring autonomous driving vehicles in a VANET. The highest reward represents the best driving behaviours of the autonomous vehicles. Additionally, we also notice that longer training time is needed by One Co-DDPG to achieve zero collision in Figure 6. Whereas, in comparison with Two DDPG, the Three DDPG ADVs is 13% faster in reaching the zero collision goal during training phases.

The Co-DDPG rewards gain by the autonomous driving vehicles is measured and the results are shown in Figure 7. The results describe the summation of the reward recorded during every 50 episodes. We notice that as the number of episodes increases, the average reward increases in RACE. In the beginning, the reward is low in all scenes, since Co-DDPG of each autonomous driving vehicle is initialized with random learning parameters in the initial stage. Therefore, the vehicles may not select the correct action for their next movement, which leads to many collisions. Co-DDPG punishes the failure actions by decreasing the rewards to help Co-DDPG learn from the bad behavior. Then, during next episode, it has the ability to choose the correct action based on its previous experiences. As the learning experiences of the vehicles increase with further episodes, the rewards start to increase. This is because the vehicles learned to make a suitable action to avoid collisions. In such scenarios, incentives are awarded for encouraging the desired driving behaviours and thereby minimize any collision. We observe that One Co-DDPG has the lowest reward over a span of 500 episodes compared to that in Two Co-DDPG and Three Co-DDPG. The reward per autonomous driving vehicle in Three Co-DDPG is about 33% higher than that in Two Co-DDPG on average. Moreover, our algorithm also awards penalties when there is a possibility for collision *e.g.*, when the agents drive too close to each other. Interestingly, we observe that as they learned better policies, the agents choose to go

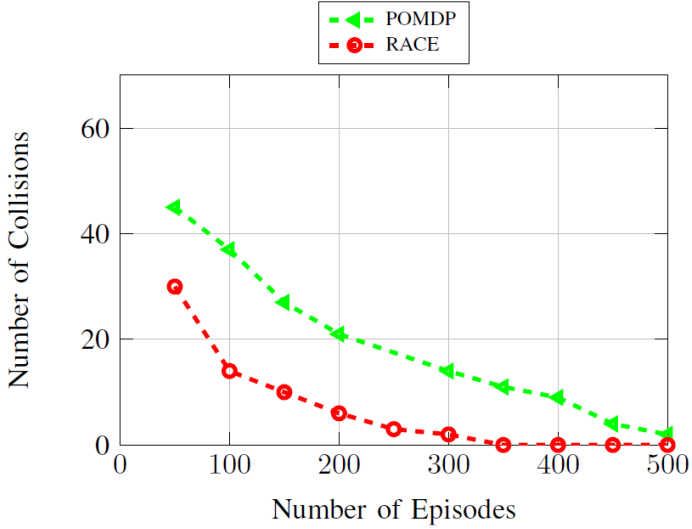


Figure 8: The comparison of total number of collisions of ADVs over the number of episodes in the training process.

off track rather than collide with other vehicles.

E. Comparison Analysis

In this section, we compare the performance of POMDP [22] with the proposed RACE framework. We utilize three metrics namely, collisions, reward and latency for this comparison.

The comparison results showing the number of collisions in RACE and POMDP over different number of episodes are shown in Figure 8, where three ADVs installed Co-DDPG and ten NDVs are used. Each value of point denotes a summation of the number of collisions experienced during every 50 episodes by all ADVs. We can see that number of collisions in RACE is around 65% lower than that in POMDP on average during the training process. In addition, compared to POMDP, RACE learned the best collision avoidance policy by 42% faster on average.

Figure 9 illustrates the average reward gained by every ADV in both POMDP and RACE. The simulation vehicles are same as in Figure 8. We add every 50 episodes' reward of all ADVs first. Then the summation reward value divides the number of ADVs to get the average reward point value in this figure. With increasing number of episodes, the average reward per ADV for both POMDP and RACE increases, which means that both of them can improve their driving policy by learning from the environment. However, compared to POMDP, RACE achieved 215% higher rewards. This demonstrates that RACE has the ability to learn better driving policies. Because RACE not only uses the parameter distribution strategy to speed up the learning ability but also employs the adaptive dynamic reward function to process the collision scenario. Besides, the learning process in RACE is more stable than in POMDP.

Figure 10 shows the latency per ADV over different models, where 1Co-D, 2Co-D and 3Co-D represents One Co-DDPG ADV with ten NDVs in RACE, Two Co-DDPG ADVs with ten NDVs in RACE and Three Co-DDPG ADVs with ten NDVs in RACE. Latency is the time duration from the trained Co-DDPG in one ADV collecting states information from

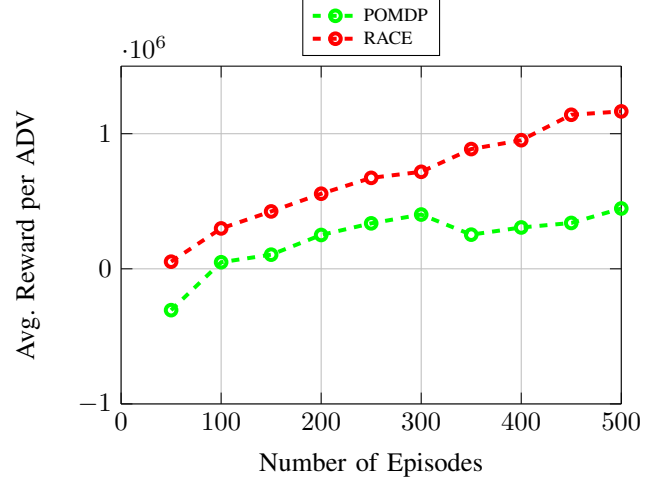


Figure 9: The comparison of average reward per ADV over the number of episodes in the training process.

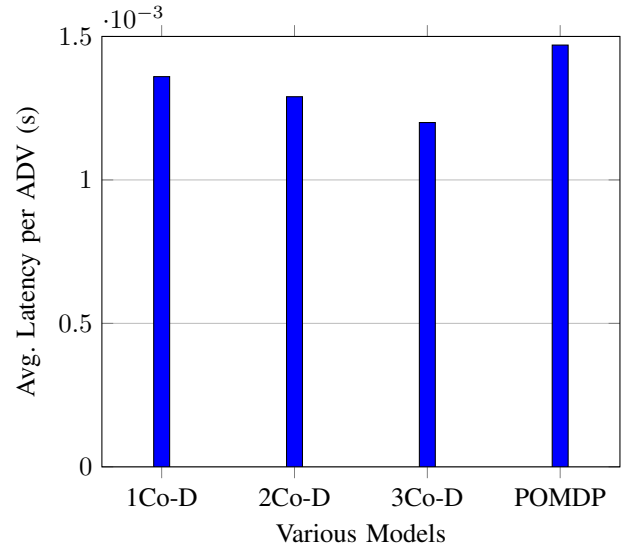


Figure 10: The comparison of latency over different models in the testing process.

the environment to the time of obtaining its action, such as push the brake. We observed that 3Co-D in RACE has the lowest latency among all models. 2Co-D in RACE performed better than 1Co-D in RACE. Since RACE utilizes the leading-following framework and parameter distribution strategy, a higher number of ADVs with Co-DDPG leads to learning better driving policies. Compared to POMDP, the latency per ADV was about 19% lower in RACE with 3Co-D.

VI. RELATED WORK

In this section we describe the related work through broadly classifying the relevant literature into two main categories, namely collision avoidance and multi-agent collision avoidance in an autonomous driving environment.

A. Collision Avoidance

Collision avoidance for autonomous vehicles and robots is an actively researched topic. Various methods including

rule-based [30], optimization-based [14], traditional machine learning [31], [32] and deep learning methods [22], [33] have been proposed in recent years to minimize collisions during autonomous driving. However, rule-based solutions do not easily adapt to different scenarios, while optimization-based methods may rapidly converge to local minimal solutions and hence, have relatively low computational efficiency. In addition, a study by Namazi et al. [34] shows that traditional machine learning-based solutions are not suitable for a complex and dynamic environment such as autonomous driving. Leveraging deep learning especially the Convolutional Neural Networks (CNNs), Lv et al. [35] handled collision avoidance by predicting the traffic flow while Chen et al. [10] utilized DRL with multi-agents settings to avoid collisions. In addition, Cheng et al. [36] formulated an automated enemy avoidance problem with Markov Decision Process and resolved it with temporal-difference reinforcement learning. However, these proposals cannot effectively minimize the collisions since the participating robots do not communicate with each other. Chen et al. [37] predicted rear-end collisions using GA-optimized neural networks. Collisions between cyclists and heavy goods vehicles were prevented in [38] by solving the side-to-side collisions problem. However, these approaches did not consider collisions from any other directions. On the other hand, Long et al. [22] used deep learning to avoid collisions among robots and Lu et al. [39] investigated interaction between three pedestrians and one robot while Sangiovanni et al. [40] considered an unpredictable object as an obstacle to avoid collision. All of the above mentioned proposals failed to consider dynamic and complex environments with possibly high speed vehicles.

B. Multi-Agent Collision Avoidance

In the following we discuss the recent proposals that avoid collisions in a multi-agent environment.

In [7], [41], authors utilize connected networks and simulated overtaking maneuvers to avoid collisions. However, these proposals mainly focus on a specific type of collision and use complex mathematical models instead of machine learning methods. Hence they require expert knowledge to design efficient mathematical models. In addition, Long et al. [33] proposed sensor-level collision avoidance policies. However, the supervised policies face limitations during learning w.r.t training dataset. In [42], Lou et al. proposed a formation control law using a combination of consensus-based formation control method and a collision avoidance algorithm, while Lu et al. [39] proposed a human-aware algorithm to find a smooth and collision-free path. However, these algorithms are computationally intensive and were not tested in complex environments and hence their effectiveness is difficult to assess.

In [43], [44], authors proposed centralized solutions. However, centralized learning produces obvious limitations for effective transmissions of the policies. The high computational burden of an optimization-based centralized scheme makes the deployment of the control system on real platforms challenging. On the other hand, Chen et al. [10] developed a decentralized multi-agent collision avoidance algorithm where

two agents were simulated to navigate toward their own goal positions and learn a value network that encodes the expected time to goal. However, cooperative information among robots is not accounted for in the solution and the design is not suitable for high speed scenarios. The Optimal Reciprocal Collision Avoidance (ORCA) framework from Van et al. [45] along with its extensions [46], [47] is widely used in crowd simulation and multi-agent systems. It provides the necessary conditions for multiple robots to avoid collisions with each other in a short time span, and easily scales to handle large systems with many robots. However, these methods are sensitive to the uncertainties in the real world driving scenarios as they assume that each robot has accurate information about the surrounding agents' positions, velocities and shapes.

Based on the above mentioned literature and our extensive studies, we believe that RACE provides a more private, dynamic and efficient solution for collision avoidance during autonomous driving. In comparison to existing approaches, RACE is more dynamic than rule-based methods and has less computational complexity than optimization-based and other above mentioned deep learning methods. Hence, RACE optimally enhances autonomous driving behaviors of vehicles and effectively minimizes collisions during driving.

VII. CONCLUSION

In this work, we studied the requirements of autonomous driving in detail and provided the necessary design goals. We proposed an efficient and decentralized framework called RACE to fulfill these goals. We developed an efficient collision avoidance algorithm named Co-DDPG which exploits a hierarchical multi-agent environment in VANET. The rewards in Co-DDPG are defined by considering the collision avoidance, minimizing arrival time and maintaining vehicles on the road. Its value measures the quality of driving behavior e.g., a higher reward represents a better driving behaviour. The parameter distribution strategy is employed to expedite the learning process. Using a location-private communication strategy with a leader-follower multi-agent architecture, we distribute the best policies in the VANET and conserve important resources during driving. With extensive experiments in TORCS, we showed that RACE effectively reduced the number of collisions and scaled effortlessly with increasing number of autonomous vehicles. We also demonstrated the meticulous learning in RACE, where vehicles learned to go off-track instead of colliding with other vehicles. This improved the driving behaviour and minimized the associated risk and liability. As part of future work, we plan to enhance RACE through incorporating images or videos of the surroundings by camera during driving into the learning process. Additionally, we consider to improve RACE to adapt to the scenario, where the road has no speed range requirement.

ACKNOWLEDGEMENT

This project has received funding from the European Union's Horizon 2020 COSAFE project under grant agreement No 824019 and the China Scholarship Council under Grant ID 201706050095.

REFERENCES

- [1] S. World Health Organization, Geneva, "Global status report on road safety 2018: Supporting a decade of action." Available: https://www.who.int/violence_injury_prevention/road_safety_status/2018/en/, 2018. [Online]. Accessed: August 9, 2019.
- [2] A. Mukhtar, L. Xia, and T. B. Tang, "Vehicle detection techniques for collision avoidance systems: A review," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 5, pp. 2318–2338, 2015.
- [3] World Health Organization, Geneva, Switzerland, "World report on road traffic injury prevention: Summary," 2004.
- [4] J. J. Rolison, S. Regev, S. Moutari, and A. Feeney, "What are the factors that contribute to road accidents? an assessment of law enforcement views, ordinary drivers opinions, and road accident records," *Accident Analysis & Prevention*, vol. 115, pp. 11–24, 2018.
- [5] J. Wang, J. Liu, and N. Kato, "Networking and communications in autonomous driving: A survey," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 2, pp. 1243–1274, 2018.
- [6] A. Liniger and J. Lygeros, "A noncooperative game approach to autonomous racing," *IEEE Transactions on Control Systems Technology*, 2019.
- [7] R. Deng, B. Di, and L. Song, "Cooperative collision avoidance for overtaking maneuvers in cellular v2x-based autonomous driving," *IEEE Transactions on Vehicular Technology*, 2019.
- [8] T. Statheros, G. Howells, and K. M. Maier, "Autonomous ship collision avoidance navigation concepts, technologies and techniques," *The Journal of Navigation*, vol. 61, no. 1, pp. 129–142, 2008.
- [9] M. Chen, Y. Tian, G. Fortino, J. Zhang, and I. Humar, "Cognitive internet of vehicles," *Computer Communications*, vol. 120, pp. 58–70, 2018.
- [10] Y. F. Chen, M. Liu, M. Everett, and J. P. How, "Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning," in *2017 IEEE international conference on robotics and automation (ICRA)*, pp. 285–292, IEEE, 2017.
- [11] L. Jiang, B. G. Cai, Y. P. Wang, and W. Jian, "Location-based cooperative vehicle collision avoidance for unsignalized intersections: A multi-sensor integration approach," in *International Conference on Connected Vehicles and Expo*, 2012.
- [12] H. Su, D. Ko, and S. An, "Power-aware location-based anti-collision protocol for rfid-sensor networks," in *International Conference on Ubiquitous and Future Networks*, 2009.
- [13] M. Kaushik, K. M. Krishna, *et al.*, "Parameter sharing reinforcement learning architecture for multi agent driving behaviors," *arXiv preprint arXiv:1811.07214*, 2018.
- [14] L. Riegger, M. Carlander, N. Lidander, N. Murgovski, and J. Sjöberg, "Centralized mpc for autonomous intersection crossing," in *2016 IEEE 19th international conference on intelligent transportation systems (ITSC)*, pp. 1372–1377, IEEE, 2016.
- [15] B. Alrifae, M. G. Mamaghani, and D. Abel, "Centralized non-convex model predictive control for cooperative collision avoidance of networked vehicles," in *2014 IEEE international symposium on intelligent control (ISIC)*, pp. 1583–1588, IEEE, 2014.
- [16] K. Loayza, P. Lucas, and E. Peláez, "A centralized control of movements using a collision avoidance algorithm for a swarm of autonomous agents," in *2017 IEEE Second Ecuador Technical Chapters Meeting (ETCM)*, pp. 1–6, IEEE, 2017.
- [17] J. Ning, J. Wang, J. Liu, and N. Kato, "Attacker identification and intrusion detection for in-vehicle networks," *IEEE Communications Letters*, 2019.
- [18] Q. Luo and J. Liu, "Wireless telematics systems in emerging intelligent and connected vehicles: Threats and solutions," *IEEE Wireless Communications*, vol. 25, no. 6, pp. 113–119, 2018.
- [19] Q. Luo, Y. Cao, J. Liu, and A. Benslimane, "Localization and navigation in autonomous driving: Threats and countermeasures," *IEEE Wireless Communications*, vol. 26, no. 4, pp. 38–45, 2019.
- [20] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.
- [21] B. Wymann, C. Dimitrakakis, A. Sumner, E. Espié, and C. Guionneau, "Tors: The open racing car simulator," 2015.
- [22] P. Long, T. Fanl, X. Liao, W. Liu, H. Zhang, and J. Pan, "Towards optimally decentralized multi-robot collision avoidance via deep reinforcement learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6252–6259, IEEE, 2018.
- [23] D. Loiacono, L. Cardamone, and P. L. Lanzi, "Simulated car racing championship: Competition software manual," *arXiv preprint arXiv:1304.1672*, 2013.
- [24] P. Tientrakool, Y.-C. Ho, and N. F. Maxemchuk, "Highway capacity benefits from using vehicle-to-vehicle communication and sensors for collision avoidance," in *2011 IEEE Vehicular Technology Conference (VTC Fall)*, pp. 1–5, IEEE, 2011.
- [25] G. Samara and Y. Al-Raba'nah, "Security issues in vehicular ad hoc networks (vanet): a survey," *arXiv preprint arXiv:1712.04263*, 2017.
- [26] X. Cheng, C. Chen, W. Zhang, and Y. Yang, "5g-enabled cooperative intelligent vehicular (5genciv) framework: When benz meets marconi," *IEEE Intelligent Systems*, vol. 32, no. 3, pp. 53–59, 2017.
- [27] K. Katsaros and M. Dianati, "Evolution of vehicular communications within the context of 5g systems," *Enabling 5G Communication Systems to Support Vertical Industries*, pp. 103–126, 2019.
- [28] A. E. Sallab, M. Abdou, E. Perot, and S. Yogamani, "Deep reinforcement learning framework for autonomous driving," *Electronic Imaging*, vol. 2017, no. 19, pp. 70–76, 2017.
- [29] B. Wymann, "Cg track 2." Available: http://www.berniw.org/trb/tracks/track_view.php?viewtrackid=11, 2017. [Online]. Accessed: May 12, 2019.
- [30] G. Lu, L. Li, Y. Wang, R. Zhang, Z. Bao, and H. Chen, "A rule based control algorithm of connected vehicles in uncontrolled intersection," in *17th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, pp. 115–120, IEEE, 2014.
- [31] J. Tian, M. Gao, and E. Lu, "Dynamic collision avoidance path planning for mobile robot based on multi-sensor data fusion by support vector machine," in *2007 international conference on mechatronics and automation*, pp. 2779–2783, IEEE, 2007.
- [32] Y.-H. Liu and C.-J. Shi, "A fuzzy-neural inference network for ship collision avoidance," in *2005 International Conference on Machine Learning and Cybernetics*, vol. 8, pp. 4754–4759, IEEE, 2005.
- [33] P. Long, W. Liu, and J. Pan, "Deep-learned collision avoidance policy for distributed multiagent navigation," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 656–663, 2017.
- [34] E. Namazi, J. Li, and C. Lu, "Intelligent intersection management systems considering autonomous vehicles: A systematic literature review," *IEEE Access*, vol. 7, pp. 91946–91965, 2019.
- [35] Y. Lv, Y. Duan, W. Kang, Z. Li, and F.-Y. Wang, "Traffic flow prediction with big data: a deep learning approach," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 2, pp. 865–873, 2015.
- [36] Q. Cheng, X. Wang, J. Yang, and L. Shen, "Automated enemy avoidance of unmanned aerial vehicles based on reinforcement learning," *Applied Sciences*, vol. 9, no. 4, p. 669, 2019.
- [37] C. Chen, H. Xiang, T. Qiu, C. Wang, Y. Zhou, and V. Chang, "A rear-end collision prediction scheme based on deep learning in the internet of vehicles," *Journal of Parallel and Distributed Computing*, vol. 117, pp. 192–204, 2018.
- [38] Y. Jia and D. Cebon, "Field testing of a cyclist collision avoidance system for heavy goods vehicles," *IEEE Transactions on Vehicular Technology*, vol. 65, no. 6, pp. 4359–4367, 2016.
- [39] X. Lu, Y. Cao, Z. Zhao, and Y. Yan, "Deep reinforcement learning based collision avoidance algorithm for differential drive robot," in *International Conference on Intelligent Robotics and Applications*, pp. 186–198, Springer, 2018.
- [40] B. Sangiovanni, A. Rendiniello, G. P. Incremona, A. Ferrara, and M. Piastra, "Deep reinforcement learning for collision avoidance of robotic manipulators," in *2018 European Control Conference (ECC)*, pp. 2063–2068, IEEE, 2018.
- [41] S. Y. Gelbal, S. Zhu, G. A. Anantharaman, B. A. Guvenc, and L. Guvenc, "Cooperative collision avoidance in a connected vehicle environment," tech. rep., SAE Technical Paper, 2019.
- [42] G. Lou and W. Yang, "Formation control and collision avoidance of multi-agent system with switching communication topology," in *Proceedings of the 2nd International Conference on Computer Science and Application Engineering*, p. 47, ACM, 2018.
- [43] J. Foerster, I. A. Assael, N. de Freitas, and S. Whiteson, "Learning to communicate with deep multi-agent reinforcement learning," in *Advances in Neural Information Processing Systems*, pp. 2137–2145, 2016.
- [44] J. Dentler, M. Rosalie, G. Danoy, P. Bouvry, S. Kannan, M. A. Olivares-Mendez, and H. Voos, "Collision avoidance effects on the mobility of a uav swarm using chaotic ant colony with model predictive control," *Journal of Intelligent & Robotic Systems*, vol. 93, no. 1-2, pp. 227–243, 2019.
- [45] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics research*, pp. 3–19, Springer, 2011.

- [46] J. Snape, J. Van Den Berg, S. J. Guy, and D. Manocha, "The hybrid reciprocal velocity obstacle," *IEEE Transactions on Robotics*, vol. 27, no. 4, pp. 696–706, 2011.
- [47] D. Bareiss and J. van den Berg, "Generalized reciprocal collision avoidance," *The International Journal of Robotics Research*, vol. 34, no. 12, pp. 1501–1514, 2015.