

# CoTV: Cooperative Control for Traffic Light Signals and Connected Autonomous Vehicles Using Deep Reinforcement Learning

Jiaying Guo<sup>1b</sup>, Graduate Student Member, IEEE, Long Cheng<sup>2b</sup>, Senior Member, IEEE,  
and Shen Wang<sup>1b</sup>, Senior Member, IEEE

**Abstract**—The target of reducing travel time only is insufficient to support the development of future smart transportation systems. To align with the United Nations Sustainable Development Goals (UN-SDG), a further reduction of fuel and emissions, improvements of traffic safety, and the ease of infrastructure deployment and maintenance should also be considered. Different from existing work focusing on optimizing the control in either traffic light signal (to improve the intersection throughput), or vehicle speed (to stabilize the traffic), this paper presents a multi-agent Deep Reinforcement Learning (DRL) system called CoTV, which Cooperatively controls both Traffic light signals and Connected Autonomous Vehicles (CAV). Therefore, our CoTV can well balance the reduction of travel time, fuel, and emissions. CoTV is also scalable to complex urban scenarios by cooperating with only one CAV that is nearest to the traffic light controller on each incoming road. This avoids costly coordination between traffic light controllers and all possible CAVs, thus leading to the stable convergence of training CoTV under the large-scale multi-agent scenario. We describe the system design of CoTV and demonstrate its effectiveness in a simulation study using SUMO under various grid maps and realistic urban scenarios with mixed-autonomy traffic.

**Index Terms**—Deep reinforcement learning, multi-agent system, connected autonomous vehicles, mixed-autonomy traffic.

## I. INTRODUCTION

DEVELOPING the next generation Intelligent Transportation Systems (ITS) is one of the key ways to achieve the United Nations Sustainable Development Goals (UN-SDG) [1]. In particular, firstly, sustainable traffic requires higher efficiency to reduce enormous monetary losses caused by excessive traffic delays. Secondly, more eco-friendly driving should be encouraged to decrease fuel consumption and gas emissions (mainly CO<sub>2</sub>). Thirdly, traffic safety is one of the key indicators for sustainable traffic, inherently, which should

be enhanced by avoiding potential collisions to save lives. Last but not least, to achieve those sustainable traffic goals, easy-to-deploy ITS infrastructure is critical.

Most existing research in sustainable urban traffic control adjusts either traffic light signals or vehicle speed. Traffic light signal controllers dynamically select the best timing plan according to the real-time traffic. As shown in Fig. 1, this can directly increase the intersection throughput, thus reducing travel time as well as energy consumption and emissions. CAV can proactively control vehicles' acceleration, as shown in Fig. 1, to achieve more stable traffic nearby with relatively higher driving velocity (i.e., lower fuel consumption and gas emissions) and keep a safe distance [2] from the surrounding traffic (i.e., longer time-to-collision). Recent research from the transportation domain attempts to explore the potential of joint control for both traffic light signals and vehicle speed. Methodologies used in such research include mixed-integer linear programming [3], the enumeration method and the pseudo-spectral method [4]. However, these methods may not perform well in realistic traffic scenarios because their deterministic traffic control decisions are insufficient to deal with a fast-changing urban environment [5].

Unlike the aforementioned traditional methods, many researchers nowadays have demonstrated the great potential of DRL in solving traffic control challenges under complex urban scenarios. For instance, inspired by the traditional traffic signal control method MaxPressure [6], PressLight [7] can achieve even better traffic efficiency improvements under various urban scenarios using DRL. Moreover, the DRL-based traffic signal control can also reduce the waiting time of specific vehicles in emergency situations in which traffic condition varies quickly [8]. On the other hand, efficient and effective CAV speed control can stabilize traffic in many complex and changing road scenarios using DRL [9], which is traditionally difficult using optimization-based controllers. However, there is a lack of research using DRL for the joint control of both urban intersection signals and vehicle speed. This DRL-based joint control is challenging due to the difficulty of designing a proper cooperation scheme for two different agent types (i.e., traffic light controllers and CAV). Moreover, the high unpredictability of urban mixed-autonomy traffic makes it even harder to converge within a reasonable number of training iterations.

Manuscript received 31 January 2022; revised 20 July 2022, 8 December 2022, and 14 February 2023; accepted 3 May 2023. Date of publication 5 June 2023; date of current version 4 October 2023. This work was supported by the School of Computer Science and Beijing-Dublin International College, University College Dublin. The Associate Editor for this article was H. Farah. (Corresponding author: Shen Wang.)

Jiaying Guo and Shen Wang are with the School of Computer Science, University College Dublin, Dublin 4, D04 V1W8 Ireland (e-mail: jiaying.guo@ucdconnect.ie; shen.wang@ucd.ie).

Long Cheng is with the School of Control and Computer Engineering, North China Electric Power University, Beijing 102206, China (e-mail: lcheng@ncepu.edu.cn).

Digital Object Identifier 10.1109/TITS.2023.3276416

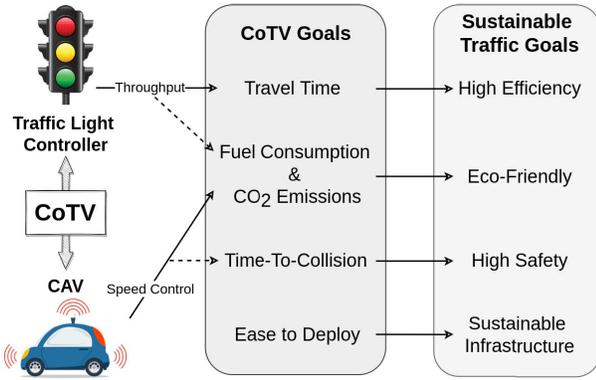


Fig. 1. The illustration of the motivation and goals of our proposed system CoTV. Traditionally, traffic light controllers can increase the intersection throughput, thus reducing the travel time and fuel. While CAV adjusts its speed to reduce the fuel, thus maintaining a safe time gap to its surrounding traffic. Our CoTV coordinates these two different types of agents to achieve a more comprehensive set of the goals of sustainable traffic.

To overcome the limitations mentioned above, we propose CoTV: a multi-agent DRL-based system that can cooperatively control traffic light signals and CAV, which is also open-sourced here.<sup>1</sup> CoTV well balances the advantages of both traffic light controllers and CAV to achieve more sustainable traffic, as shown in Fig.1. Concretely, the contributions of our work are as follows:

- **Effective cooperation schemes between CAVs and traffic light controllers.** Different from the methodology in the literature on Multi-Agent Reinforcement Learning (MARL) for traffic control, instead of using action-dependent design [10] (i.e., the action of one agent depends on the action of other agents in the shared environment), our cooperation schemes rely on the exchange of states between agents within the range of one intersection, including the traffic light controller and approaching CAVs. This so-called “action-independent MARL” [11] can work in our CoTV as the objective of traffic light controller and CAV for the traffic improvement are inherently complementary (i.e., rather than overlapping: all improving travel time or reducing fuel). Thus, CoTV takes advantage of the simplicity of “action-independent MARL” design on DRL training and keeps the effectiveness of CoTV in improving traffic under various scenarios. The cooperation schemes of CoTV are shown to facilitate training convergence, which is challenging for independent MARL that does not include any cooperation (either state or action). Specifically, our CoTV using Proximal Policy Optimization (PPO) [12] obtains up to 30% reduction in both travel time and fuel consumption & CO<sub>2</sub> emissions under varying CAV penetration rates.
- **Scalable to complex urban scenarios by avoiding cooperation with excessive CAV agents.** Compared with controlling all possible CAVs using MARL, the traffic light controller in our CoTV selects the closest CAV to the intersection on each incoming road as the CAV agent.

This idea is inspired by platooning can increase intersection throughput [13], as the leading vehicle in a certain road has the great potential to form a platoon with the rest vehicles on the same road. We also demonstrate that compared with coordinating all CAVs (CoTV\*), CoTV does not compromise efficiency improvement while significantly reducing the training time and resources used.

- **Efficient communication exchange schemes between CAV and traffic light controllers.** The amount of state information exchanged between CAV and traffic light controllers of CoTV is small. As shown in Fig.2, the communication schemes between two agents are designed to exchange the speed, acceleration, and location of CAVs and the current signal phase of traffic light controllers. The information exchanged requires less than 100 Kbps transmission rate, which can be well handled using Vehicle-To-Vehicle (V2V) and Vehicle-To-Infrastructure (V2I) communication infrastructure, as their IEEE 802.11p standard supports a bandwidth of 3 Mbps to 20 Mbps [14].

This paper extends our previous work [15] with the following improvements: 1) CoTV is more scalable as the number of its controlled CAVs is significantly reduced. 2) The state and reward of CoTV DRL agents are simplified, thus leading to more efficient agent communication for easy deployment. 3) CoTV is validated on more realistic urban scenarios under different CAV penetration rates. 4) Traffic safety of CoTV is also evaluated using the metric time-to-collision [16]. 5) Action-dependent and independent MARL approaches are compared to demonstrate the advance of CoTV in policy training and traffic improvements.

## II. RELATED WORKS

### A. Control for Either Traffic Light Signals or Vehicle Speed

Most existing research in sustainable urban traffic control adjusts either traffic light signals or vehicle speed. Sydney Coordinated Adaptive Traffic System (SCATS) [17] is one of the earliest and most widely applied traffic light signal control systems. It can dynamically select the best signal plan from pre-defined candidates that can achieve better intersection throughput. Varaiya [6] proposed a traffic light signal control scheme named MaxPressure, which was proven to maximize the throughput of the entire road network, with each traffic light controller receiving local traffic information. On the other hand, the field experiments in [18] prove that a well-designed speed control of CAV can stabilize traffic and is beneficial to reduce braking times and fuel consumption. Green Light Optimal Speed Advisory (GLOSA) system guides CAV to adjust its speed according to the current traffic signal phase and the distance to its approaching intersection [19]. However, these traffic control optimization approaches rely on deterministic formulations, which are not flexible enough to improve ever-changing realistic traffic.

DRL has been used to cope with complex traffic environments, promising better urban traffic. PressLight [7] is a DRL-based model using Deep Q-learning (DQN). It collects local real-time traffic information inspired by the traditional

<sup>1</sup><https://github.com/Guojjy/CoTV>

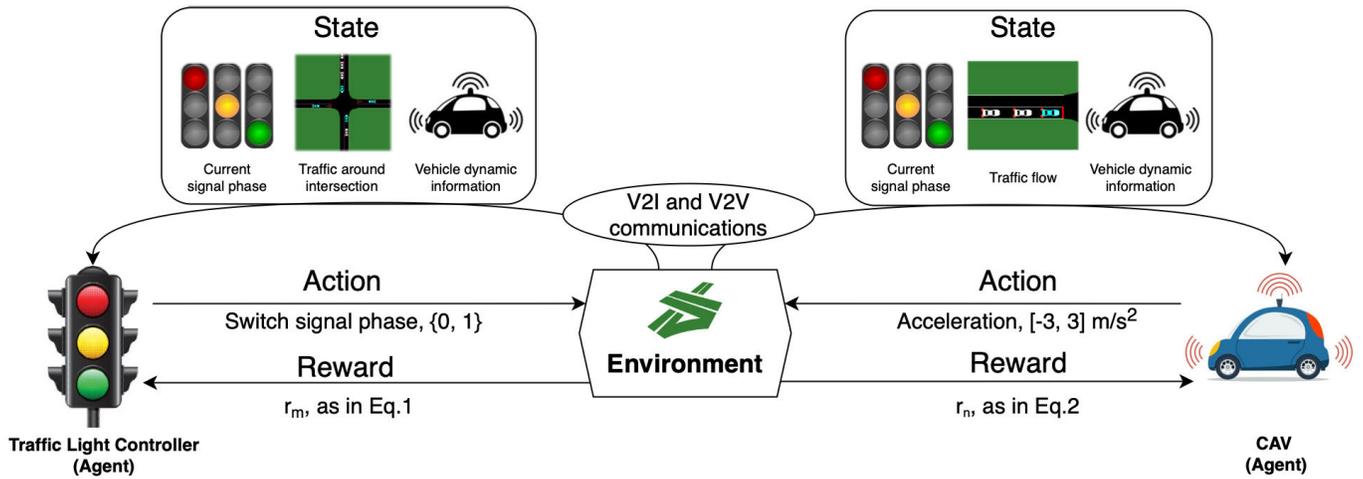


Fig. 2. The DRL design of CoTV. Two types of agents, traffic light controllers and CAV interact with the environment according to the state information exchanged via V2X communications.

method MaxPressure [6] while achieving more improvement on traffic efficiency than MaxPressure. Wu et al. [9] extended the field experiments in [18] using the Trust Region Policy Optimization (TRPO) method for training CAVs in a simulated experiment. The used DRL-based vehicle speed controller surpasses traditional optimization controllers on traffic improvement. Various scenarios using CAV had been tested in [20], including road merging and unsignalized intersections. The DRL-based speed control of CAV can optimize the vehicle trajectory of the whole trip and reduce the risk of collision all the time. Compared to traditional optimization methods with deterministic solutions, DRL methods, which are used in our proposed CoTV, learn from trial-and-error in the interaction with the dynamic environment, which is more capable of performing adaptive traffic control under fast-changing urban road scenarios.

### B. Joint Control for Traffic Light Signals and Vehicle Speed

Traditional optimization-based methods have been attempted to jointly control traffic light signals and vehicle speed. Yu et al. [3] developed mixed-integer linear programming for optimizing vehicle trajectories and traffic signals simultaneously at isolated intersections. The phase sequence and duration of traffic light signals are coordinated with vehicle arriving time to the intersections. A two-level model for traffic light controllers and CAV was proposed using the enumeration method and the pseudo-spectral method [4]. The first level is applied to coordinate CAV and traffic light controllers to minimize travel time, and the second level is used to regulate CAV trajectory to reduce fuel consumption. The same system targets were adopted in the cooperative optimization model [21]. The model uses a mixed-integer non-linear program, which has a high computational complexity.

To the best of our knowledge, DRL methods for the joint control of traffic light signals and CAV have not been well-studied. The joint control using DRL suffers many challenges, commonly in multi-agent systems [22]: (1) Traffic light controller or CAV agents proactively interact with the same

environment simultaneously, which makes it challenging for a stable training convergence. (2) A large number of agents (particularly CAV agents) cause scalability issues due to an excessive computational cost of joint action. (3) It is uncertain to which level should the reward of agents be assessed: individual junctions, regional areas, or global networks. For example, traffic light controllers explicitly coordinate traffic around intersections, and each CAV mainly affects its surrounding traffic. The proposed model in this paper attempts to overcome these difficulties and utilize the advantage of DRL methods to control traffic light signals and CAV cooperatively.

### C. Efficiency and Safety for Mixed-Autonomy Traffic

The development of CAV is thriving in both academia and industry, which is expected to improve traffic. However, the deployment must experience a gradual mixed transition from introductory, established, to prevalent [23] with the growth of CAV penetration rate. Existing work presents that CAV mixing in traffic still brings uncertainty. Mixed-autonomy experiments on motorways were conducted in [24], simplified from intersections with conflicting traffic movements. Similar work was tested in single-lane facilities, where CAV can enhance traffic safety by keeping a larger gap from the surrounding vehicles [25]. However, a low penetration rate (less than 10%) causes more conflicts in urban scenarios [26]. Reference [27] conducted experiments in a ring scenario, showing that the CAV penetration rate greater than 20% allows all vehicles to reach higher speeds and stabilize the flow. The penetration rate in 20% to 40% is possible to result in the near-maximum improvements [26]. Overall, a high penetration rate of CAV can bring traffic efficiency and safety improvement on mixed-autonomy traffic in various scenarios.

This work advances the state-of-the-art in assessing DRL-based mix-autonomy control under dynamic urban road scenarios with multiple intersections. Moreover, our system CoTV chooses only a small fraction of CAVs that cooperate with traffic light controllers, which have great potential to guide the rest of the vehicles. This makes the deployment of CoTV practical and easy-to-scale.

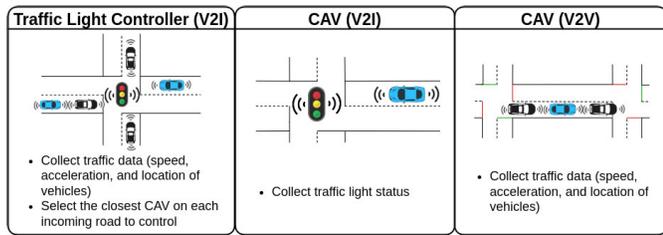


Fig. 3. V2X communication schemes in CoTV showing how traffic light controllers and CAVs use V2I and V2V. This implements state exchange and cooperative control. CAV agents of CoTV are highlighted in blue.

### III. SYSTEM OVERVIEW

#### A. System Design Goals

The proposed model CoTV aims to achieve the following goals, which are also shown in Fig.1:

- **Reduced travel time:** Travel time is the metric that end road users care about the most. Our system should reduce the travel time of a vehicle without changing its given route. This goal is traditionally achieved by traffic light signal control that can increase intersection throughput.
- **Lower fuel consumption and CO<sub>2</sub> emissions:** Sustainable traffic goals encourage eco-friendly driving behaviors. This goal is traditionally achieved by the speed control of CAV that can stabilize traffic flow. Smart traffic light control can also partly contribute to achieving this goal by reducing the number of stop-and-gos.
- **Longer time-to-collision:** Safety is a crucial consideration in sustainable traffic system design. Reducing the risk of collision can be achieved by maintaining a longer time-to-collision [16], with sufficient time to moderately decelerate. CAV can proactively keep a safe distance from the surrounding traffic. Thus, ITS using CAV has the potential to achieve higher traffic safety.
- **Easy to deploy:** Our system CoTV requires a V2X communication infrastructure to support information exchange over the cooperative control. Meanwhile, scalability issues should be addressed with the increasing number of agents (especially CAV agents). Efficient communication schemes among traffic light controllers and reduced CAV agents are the key to achieving this goal.

#### B. System Components

Our proposed system assumes that all vehicles are connected, including CAV and Human-Driven Vehicles (HDV) (details can be found in Table II). The V2X communication is also assumed perfect without no packet loss and no latency. The main components of CoTV: traffic light controllers and CAV, as shown in Fig.2. The design of action, state, and reward for them are described as follows, while the V2X communication schemes involved are shown in Fig.3:

##### 1) Traffic Light Controller Agents:

- **Action:** We limit the action of traffic light controller to a binary set, where “1” represents switching to the next phase for the next time interval while “0” means to keep the current phase unchanged. As opposed to other

common action definitions in the literature, such as phase selection [11], the phase switch [28] we choose is more manageable for the DRL model training.

- **State:** The state of traffic light controller involves three parts: the current signal phase, traffic on the roads that this traffic light controller coordinates, and the status of the closest vehicle to the intersection on each incoming road. As shown in Fig.2, of all three parts of the traffic light controller’s state, the information of the last two parts is acquired by using the V2I communications infrastructure illustrated in Fig.3. The road traffic is presented by the number of vehicles on each road coordinated by the traffic light controller. These roads are divided into incoming roads and outgoing roads. The last part of the state includes speed, acceleration, distance to the intersection, and the road name where it is located for the closest vehicle to the intersection on each incoming road.
- **Reward:** The reward is the penalty of intersection pressure inspired by [7] and [6]. Intersection pressure is defined as the difference between the sum of the number of vehicles on the incoming roads  $N_{in}$  and the sum of the number of vehicles on the outgoing roads  $N_{out}$ . Then the intersection pressure is normalized by the maximum road capacity  $c$  to facilitate DRL training. The maximum road capacity  $c$  indicates the maximum number of vehicles on a single road in the given road network. It is calculated by dividing the length of the longest road by the minimum space required for a vehicle (i.e., the length of a single vehicle plus the minimum distance between two adjacent vehicles). The reward of a certain traffic light controller  $r_m$ , Eq.(1) becomes

$$r_m = -\frac{N_{in} - N_{out}}{c} \quad (1)$$

We also illustrate the above reward in Fig.4. The reward function is formulated to reduce travel time, one of the system goals, by increasing intersection throughput. Minimizing intersection pressure encourages vehicles to pass through the intersection quickly while considering the remaining capacity in the outgoing roads, thus improving green light efficiency and throughput [6]. We also simplify the calculation of intersection pressure without considering traffic movements (correspondence between incoming and outgoing roads) compared with [6] and [7]. Therefore, CoTV can be easily applied in various urban scenarios with multi-directional roads. Besides, we avoid using other common reward definitions in the current literature, such as queue length and waiting time [10], [28], which is precarious in different traffic flow conditions even without the influence of traffic lights.

##### 2) CAV Agents:

- **Action:** The action is set to be consistent with the literature [29], which is a continuous action space to represent the CAV acceleration in the range of  $[-3m/s^2, 3m/s^2]$ .
- **State:** The state explicitly includes speed and acceleration for itself and the vehicle preceding the CAV immediately, the distances to the preceding vehicle and the approaching

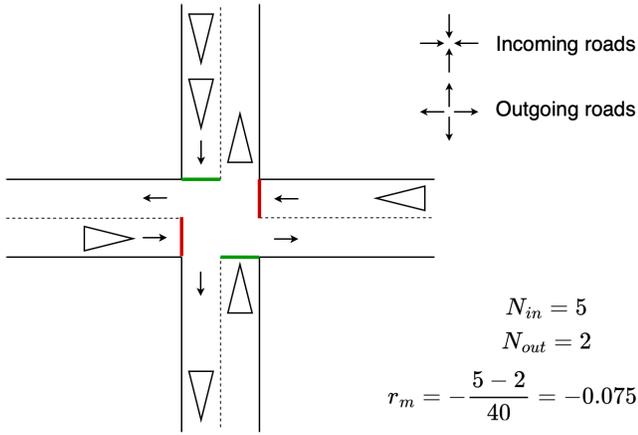


Fig. 4. The illustration of the reward of a traffic light controller  $r_m$ , assuming the maximum road capacity  $c = 40$ .

intersection, and the current signal status of the approaching traffic light controller. The CAV agent can receive state information from the vehicles on the same road and the approaching traffic light controller using V2V and V2I communication, as shown in Fig.3.

- **Reward:** The reward is penalized by the deviation of average speed  $v$  from the maximum speed limit  $v^*$ , plus the Euclidean norm of acceleration  $a$  after the normalization using the vehicle's maximum acceleration  $a^*$ , as shown in Fig.5. Speeds and accelerations in the reward are that of all vehicles  $K$  located on the same road as the CAV agent. The reward of certain CAV agent  $r_n$ , Eq.(2) becomes

$$\begin{aligned}
 r_n &= r_1 + r_2, \\
 r_1 &= -\frac{\sum_{j \in K} (v^* - v_j)}{v^* \times |K|}, \quad v_j \leq v^*, \\
 r_2 &= -\sqrt{\frac{\sum_{j \in K} (a_j/a^*)^2}{|K|^2}}, \quad a_j = \begin{cases} 0, & a_j < 0 \\ a_j, & a_j \geq 0 \end{cases} \quad (2)
 \end{aligned}$$

The first term of the reward function  $r_1$  encourages a higher average vehicle velocity but keeps it within the maximum speed limit. In this speed range, higher speed increases fuel economy, and potential collisions due to excessive speed can be avoided. The second term of the reward function  $r_2$  stabilizes acceleration to reduce fuel consumption, while also inducing a large time gap between adjacent vehicles [25] for enabling high-speed collision-free driving. Our reward function of CAV agents encourages better speed control, thus facilitating cooperative control of CoTV to achieve the reduction of fuel consumption and CO<sub>2</sub> emissions and the improvement of traffic safety.

### C. DRL Training Process

Algorithm 1 presents the training process of CoTV, and the outcome is the policy functions for traffic light agents  $M$  and CAV agents  $N$ ,  $\pi_{TL}$  and  $\pi_{CAV}$ . The trained policy  $\pi$  is expected to guide agents to select one appropriate action  $a$  given a certain state  $s$  for maximizing the accumulated

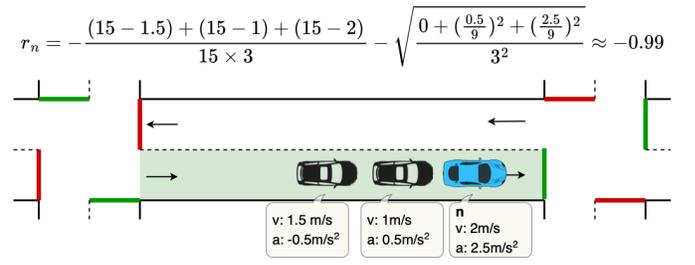


Fig. 5. Illustration of the CAV reward  $r_n$ , assuming the maximum speed limit  $v^* = 15m/s$ , and the vehicle's maximum acceleration  $a^* = 9m/s^2$ . CAV agents of CoTV are highlighted in blue.

value of reward  $r$ . We predefine the termination condition as the number of training iterations  $I$ . In each iteration, there are  $E$  episodes running in parallel, and each episode lasts  $H$  timesteps. DRL trajectory data  $\tau$  is collected per simulation timestep in each episode to extend the training batch  $B$ , and then sampled  $K$  times to update the traffic light and CAV policy parameter  $\theta_{TL}$  and  $\theta_{CAV}$  through gradient descent. Specifically, the traffic light controllers of CoTV select the closest CAV to the intersection on each incoming road as the CAV agent, as in Line 10 of Algorithm 1. These CAV agents have the potential to increase intersection throughput by forming a platoon with the rest vehicles on the same road [13]. The communication scheme of CoTV occur when the agent starts receiving state, in Line 13 of Algorithm 1.

We choose PPO algorithm [12] for the following reasons. The PPO algorithm has the advantage of being easy to implement and achieving monotonic reward improvement. DQN is a widely used algorithm to train traffic light controllers [7], [28], as it is efficient in the discrete actions (e.g., a binary set of signal phase adjustment). However, DQN does not perform well on continuous actions (e.g., vehicle acceleration of any real number within a certain range) [30]. In contrast, PPO can perform well for scenarios with discrete actions or continuous actions. On the other hand, compared with traffic light signals that have a pre-defined phase sequence, the initial driving behavior of DRL-controlled CAV has lots of unreasonable stop-and-go and standstill. The constrained policy update of PPO aims to improve reward monotonically, which is more stable to train CAV and better than Asynchronous Advantage Actor-Critic used in [11]. Although TRPO can also constraint the policy update, PPO is easy to implement and simpler to sample data, which helps the cooperation of traffic light controllers and CAV.

When interacting with the environment, CoTV applies parameter sharing [9] to all agents of the same type in the multi-agent DRL system, which can converge the training process faster and benefit from shared experience, especially in large-scale applications [31].

### D. Considerations for "Easy-to-Deploy"

Firstly, CoTV is designed to be deployed in the existing adaptive traffic light systems (e.g., SCATS, SCOOTs, etc.), which are based on major junctions of urban scenarios. This deployment strategy covers broader arterial roads that carry the majority of traffic by not installing controllers for all

**Algorithm 1** Training Process of CoTV using PPO**Require:**

- 1: Obtain the set of traffic light agents to control,  $M$
- 2: Set the number of episodes in parallel to  $E$ , and the time horizon for each episode to  $H$
- 3: Initialize the policy parameter for one type of agent,  $\theta_{TL}$  for traffic light controllers and  $\theta_{CAV}$  for CAV, through parameter sharing
- 4: Initialize sample batch  $B = \emptyset$
- 5: Set the number of epochs for mini-batch updates in one iteration as  $K$

**Ensure:**

- 6: **for** iteration = 1,2,..., $I$  **do**
- 7:   **for** episode = 1,2,..., $E$  **do in parallel**
- 8:     **for** timestep  $h = 0,1,\dots,H$  **do**
- 9:       **for** each traffic light agent  $m$  in  $M$  **do**
- 10:         Add the closest CAV  $n$  to the intersection on each incoming road to the CAV agent set  $N$
- 11:       **end for**
- 12:       **for** each agent  $i$  in  $M + N$  **do**
- 13:         Run policy  $\pi_{TL}$  or  $\pi_{CAV}$  in the environment
- 14:         Collect trajectories  $\tau = (s_{h-1_i}, a_{h_i}, s_{h_i}, r_{h_i})$
- 15:         Extend  $B$  with  $\tau$
- 16:       **end for**
- 17:     **end for**
- 18:     Compute advantage estimates  $\hat{A}_1, \dots, \hat{A}_H$
- 19:   **end for**
- 20:   Update  $\theta_{TL}$  and  $\theta_{CAV}$  in the policy  $\pi_{TL}$ ,  $\pi_{CAV}$  using advantage estimates  $\hat{A}$ , with  $K$  epochs to sample mini-batches from  $B$ , and then reset  $B = \emptyset$
- 21: **end for**

intersections. Lane-changing operations are not considered in the action space of CoTV agent but permitted in the evaluation of CoTV shown in Section V. Secondly, compared with controlling all possible CAVs with DRL, the traffic light controller of CoTV selects only the closest CAV to the intersection on each incoming road to cooperate, which can significantly reduce training time and resources used in the process thus alleviating scalability issues. Meanwhile, the cooperation schemes among agents (i.e., the traffic light controller and the approaching CAV agents) only rely on the information exchange of states, not actions. This means the action for a certain agent is selected independently from other agents' actions. Therefore, CoTV avoids the exponentially increased complexity of joint actions for MARL using action-dependent design [22]. Besides, the amount of information exchanged in CoTV is small enough compared with high-dimensional transmission data (i.e., image representations to describe traffic features) [28], [32]. Specifically, as shown in Fig.3, the information of CAV involves speed, acceleration, and location. Their size is estimated to be approximately 40 Bytes if encoded using floating point numbers. While traffic light controllers send their current signal phase, which is about 8 Bytes if using integer numbers for encoding. This information plus headers will still be less than 100 Kbps. This transmission demand is met by the V2I and V2V

TABLE I  
TRAFFIC SETTINGS IN THE THREE TEST SCENARIOS

Scenario	Traffic generation duration (seconds)	Total number of vehicles
1x1 grid <sup>1</sup>	300	70
1x6 grid	300	240
Dublin	400	275

<sup>1</sup>  $a \times b$  grid,  $a$  is the number of horizontal roads,  $b$  is the number of vertical roads.

communications infrastructure [14] using IEEE 802.11p which is between 3 and 20 Mbps. Additionally, all the information exchanged using the vehicular network occurs within the range of a single intersection (i.e., the single-hop range that is about 300 meters), which can improve the robustness of CoTV instead of heavily relying on a large scale (i.e., using multi-hop transmission) of network conditions [11].

## IV. EVALUATION METHODOLOGY

### A. Simulation Scenarios

The simulation platform used in this work is Simulation of Urban MObility (SUMO),<sup>2</sup> which is one of the most widely used open-source microscopic traffic simulators. Our model design and implementation are based on FLOW,<sup>3</sup> which provides DRL-related API to work with SUMO dynamically.

We clarify some concepts relating to the time horizons. We set 1 *simulation timestep* equal to 1 simulation second. One *episode* refers to a full run of a single simulation scenario, which is set to 720 simulation timesteps. At the end of each *iteration*, CoTV starts to update the parameters of the PPO algorithm used, after 18 episodes run in parallel. All numerical experiment results shown in the next section are averaged from these 18 episodes. In total, we terminate the *training process* of CoTV after 150 iterations.

For testing scenarios, firstly, we demonstrate the effectiveness of CoTV under a simple  $1 \times 1$  grid map with a single intersection. Then, we show CoTV can be scalable to more consecutive intersections under a  $1 \times 6$  grid map. Lastly, we validate the effectiveness of CoTV using a subset of the realistic urban scenario of Dublin city, Ireland. Table I summarizes the settings of traffic in each scenario.

1) *1 × 1 Grid Map*: In our  $1 \times 1$  grid map, each edge has two roads in opposite directions. To make this map closer to the real urban scenario, we set the road length as 300 meters and the maximum speed limit as 15 m/s (=54km/h). As shown in Fig.6, we generate different go-straight traffic flows in four directions: N→S (from north to south), S→N, W→E (from west to east), and E→W. This traffic generation method is inspired from [11]. The origin and destination of each vehicle are at the end of the road at the perimeter of the network. The vehicle generation duration for each flow is approximately 300 seconds. The traffic flows N→S and W→E are relatively heavier than the other two. Specifically, the traffic flow rates in the number of vehicles per hour per road are: 288 (N→S),

<sup>2</sup><https://www.eclipse.org/sumo/>

<sup>3</sup><https://flow-project.github.io>

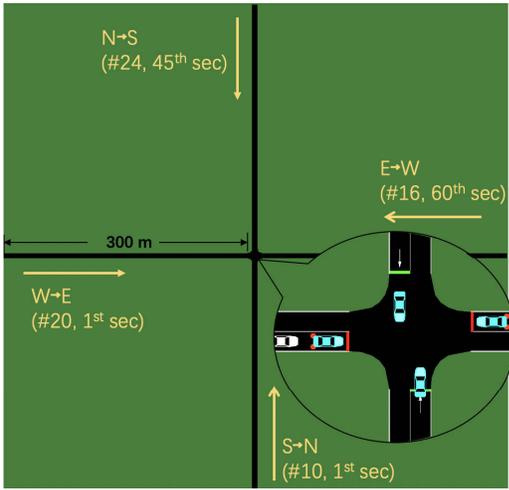


Fig. 6. The settings of traffic generation for  $1 \times 1$  grid scenario. For example,  $W \rightarrow E$  (#20, 1st sec) means there are 20 vehicles sequentially generated from the first simulation second.

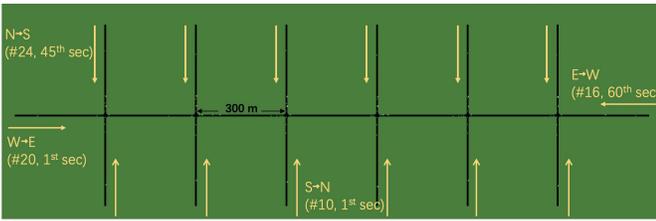


Fig. 7. The settings of traffic generation for  $1 \times 6$  grid scenario. The same settings (the number of vehicles generated, the simulation time to start traffic generation) apply for the traffic flow in the same direction.

240 ( $W \rightarrow E$ ), 192 ( $E \rightarrow W$ ), and 120 ( $S \rightarrow N$ ), respectively. The two traffic flows,  $S \rightarrow N$  and  $W \rightarrow E$ , are generated at the beginning of each episode. Then, the  $N \rightarrow S$  flow vehicles start to enter the network sequentially on the 45th second. After one minute, the traffic flow of  $E \rightarrow W$  appears. The speed of each vehicle when entering the network is random. Thus, the total number of vehicles is 70 in the scenario.

2) *1 × 6 Grid Map*: The  $1 \times 6$  grid scenario is shown in Fig.7, which contains six intersections extending the  $1 \times 1$  grid map with 5 more consecutive intersections. The road setting and traffic flow configurations are similar to the settings of the  $1 \times 1$  grid scenario. The increased vertical ( $N \rightarrow S$  and  $S \rightarrow N$ ) roads are allocated the corresponding traffic flow. A total of 240 vehicles are generated in this scenario.

3) *Dublin Map*: Fig.8 illustrates the selected six signalized intersections area in the city of Dublin. These intersections are the main ones connected by arterial roads, maximizing traffic improvement while considering “easy-to-deploy” with minimized infrastructure upgrades as mentioned in Section III-D. A variety of roads are introduced, including exclusive go-straight, exclusive turn, and multi-directional roads. Meanwhile, intersections come in different shapes and sizes, including one three-leg with four signal phases (the rightmost one in Fig.8); the four-leg intersection is the majority, three have four phases, and the other has six phases; the most complex intersection is 5-leg with 6 phases (the third one from the left of Fig.8). The scenario is extracted from the



Fig. 8. The selected six signalized intersections area in the city of Dublin (a regional road, R111, in South Dublin). The highlighted roads are our selected testing scenario (six intersections are highlighted using red circles).

open data in [26] to simulate the real-world traffic in Dublin city. We extracted dynamic traffic generated from 10 AM for 400 seconds, consisting of 275 vehicles allowed to drive straight, turn left or right at intersections. Each vehicle has a dedicated trip.

### B. Evaluation Metrics

We evaluate the sustainable traffic improvements of each scenario using the following metrics:

- **Travel time (seconds)**: Travel time of each vehicle is the time cost in the road network until finishing the designated trip. The average travel time is calculated on vehicles completing their trips in a scenario, which is the common measure to evaluate traffic efficiency [10].
- **Delay (seconds)**: Delay is the difference between the actual travel time and the ideal travel time (i.e., time spent when driving at the maximum permitted speed) for each trip. This value indicates the space in which the traffic efficiency can be further optimized to its upper-bound. This metric could be more noticeable than travel time to reflect the improvement of traffic efficiency [24].
- **Fuel consumption (l/100km)**: Fuel consumption is the average amount of fuel consumed in liters every 100 kilometers traveled. Low fuel consumption can be generally achieved by high vehicle speed and gentle change of acceleration [33]. In our experiments, fuel consumption, as well as the  $CO_2$  emission described later, is calculated using HBEFA3/PC\_G\_EU4 model (i.e., a gasoline-powered Euro norm 4-passenger car modeled using the HBEFA3 [34]), which is the default vehicle emission model in SUMO.<sup>4</sup> This model mainly considers the instantaneous speed and acceleration of a vehicle.
- **$CO_2$  emissions (g/km)**:  $CO_2$  emissions are measured by the average amount of carbon dioxide emitted in grams per kilometer traveled by all vehicles. As the primary component of greenhouse gas emissions,  $CO_2$  emissions are required to be reduced to achieve sustainable traffic.
- **Time-To-Collision (TTC)**: TTC is a widely-used safety indicator [16], estimating the time required for a car to hit its preceding one. We use the default threshold of TTC in SUMO, 3 seconds,<sup>5</sup> which means a possible collision is

<sup>4</sup><https://sumo.dlr.de/docs/Models/Emissions.html>

<sup>5</sup>[https://sumo.dlr.de/docs/Simulation/Output/SSM\\_Device.html](https://sumo.dlr.de/docs/Simulation/Output/SSM_Device.html)

TABLE II  
SIMULATION SETTINGS OF DIFFERENT VEHICLE TYPES

All vehicles are CONNECTED	
HDV (Non-CAV)	CAV
<ul style="list-style-type: none"> <li>• Can NOT be controlled by CoTV</li> <li>• IDM car-following model [35]</li> </ul>	<ul style="list-style-type: none"> <li>• Can be controlled by CoTV<sup>1</sup></li> <li>• IDM, if not controlled by CoTV</li> </ul>
$Penetration\ Rate = \frac{ CAV }{ HDV  +  CAV } \times 100\%$	

<sup>1</sup> Denoted as CoTV\* when all CAVs are controlled by our system. This scenario of 100% penetration rate is a different case as CoTV only controls CAV that is the closest to the incoming intersection on each road.

recognised when the time gap between the two adjacent cars is less than 3 seconds. The value of TTC is literally the total number of such possible rear-end collisions for a given time horizon.

### C. Compared Methods

To evaluate the effectiveness of our system CoTV, the compared methods are described as follows:

- **Baseline:** The Baseline traffic light signals have a static timing plan that does not change with the varying traffic, thus does not require V2X communications to collect vehicle information. All vehicles are HDV that are simulated by IDM car-following model [35] as shown in Table II, which is also used for simulating HDV in [29]. The Baseline scenario simulates most existing urban scenarios, which do not have any traffic light controllers and CAVs controlled by DRL. A cycle of the static traffic light signal plan contains four phases in order: Green-NS (green light for the flow N→S and S→N), Yellow-NS, Green-WE, and Yellow-WE. The duration of the green light is 40 seconds (default value in SUMO). The yellow light duration typically lasts from 3 to 6 seconds [10], so we set 3 seconds as the yellow light duration, which is also the default setting in SUMO. Thus, the length of a cycle is 86 seconds (40+3+40+3). The traffic light signal plans of the Dublin scenario dataset vary by different intersections with green light phase duration ranging from 37 to 42 seconds, and yellow light phase lasts 3 seconds. Some intersections have a short green light for turn-right with 6 seconds.
- **FlowCAV:** FlowCAV [29] is a state-of-the-art DRL-based model to control the speed of a CAV to improve fuel efficiency and reduce emissions. Each CAV observes its preceding vehicle and then regulates its speed. The reward of a single CAV is evaluated globally by the average speed and acceleration of all vehicles in the given network. In this scenario, all traffic light signals are static. There is only one CAV agent per road, which leads the following vehicles on the same road.
- **PressLight:** PressLight [7] is a state-of-the-art DRL-based model to control traffic light signals to improve intersection throughput. Its state includes the number of vehicles on the incoming roads and outgoing roads. The reward design utilizes the “pressure” to improve intersection throughput, which is inspired by [6]. All

vehicles are HDV and connected, as shown in Table II, which are periodically broadcast their up-to-date status (e.g., location, speed, acceleration), any agents within the communication range can aggregate them as the real-time traffic information.

- **GLOSA:** This is a non-DRL method for jointly controlling traffic light signals and CAVs. The GLOSA system<sup>6</sup> can adjust CAV speed considering the current traffic light phase and the current status of CAV. In our experiment, we integrate it with adaptive traffic light controllers<sup>7</sup> to achieve joint control. Thus, phase switching is actuated after detecting a sufficient time gap between successive vehicles, resulting in dynamic phase durations. All vehicles in this scenario are CAVs.
- **I-CoTV:** I-CoTV combines independent policy training on the two types of agents as a common and straightforward way to develop MARL. There is no cooperation design between agents in either state or action, distinct from CoTV (action-independent MARL with cooperation schemes in the state exchange). Hence, the state of traffic light controllers involves two parts: its current signal phase and traffic on the roads it coordinates. It does not include any instantaneous vehicle information compared to CoTV. Correspondingly, the state of CAV agent only consists of the speed, acceleration, and location of itself and its preceding vehicle, without the current signal of the approaching traffic light from agent communication. Introducing I-CoTV aims to demonstrate that the efficient cooperation schemes of CoTV facilitate training convergence.
- **M-CoTV:** M-CoTV is the action-dependent MARL version of CoTV that trains the policies of traffic light controllers and CAVs considering both the action and state of another agent type within the range of one intersection. Introducing M-CoTV aims to demonstrate that CoTV takes advantage of the simplicity of action-independent MARL on policy training while efficiently achieving traffic improvements.
- **CoTV\*:** CoTV\* has all features of CoTV, except that the traffic light controller of CoTV\* interacts with *all* CAVs instead of only the *closest* one to the intersection on each incoming road. Introducing CoTV\* aims to demonstrate the improvement of CoTV in alleviating scalability issues.

## V. EVALUATION RESULTS

### A. Traffic Efficiency & Safety

1) *Comparison With State-of-the-Art Methods:* Table III shows the traffic improvements of CoTV under 100% CAV penetration rate, the same for FlowCAV and GLOSA (while 0% CAV penetration rate for PressLight and Baseline scenario as no need for vehicle speed control).

- **Travel time & delay:** As shown in Table III, CoTV achieves the shortest travel time with up to 30% reduction compared to Baseline. PressLight and GLOSA achieve over 24% and 23% reduction, respectively. However,

<sup>6</sup><https://sumo.dlr.de/docs/Simulation/GLOSA.html>

<sup>7</sup>[https://sumo.dlr.de/docs/Simulation/Traffic\\_Lights.html#type\\_actuated](https://sumo.dlr.de/docs/Simulation/Traffic_Lights.html#type_actuated)

TABLE III

COMPARISON OF CoTV AGAINST BASELINE AND STATE-OF-THE-ART METHODS. PERCENTAGE CHANGES SHOWN ARE COMPARED TO BASELINE. THE BEST ACHIEVED MEASUREMENTS ARE IN BOLD

Method	Travel time (s)	Delay (s)	Fuel (l/100km)	CO <sub>2</sub> (g/km)	TTC
<b>1x1 grid</b>					
Baseline	59.67	18.76	9.29	216.05	354.00
FlowCAV	87.23	46.32	12.98	302.02	694.22
	+46.19%	+146.91%	+39.72%	+39.79%	+96.11%
PressLight	49.81	8.90	8.64	201.00	51.61
	-16.52%	-52.56%	-7.00%	-6.97%	-85.42%
GLOSA	50.95	10.03	8.65	201.34	65.83
	-14.61%	-46.54%	-6.89%	-6.81%	-81.40%
CoTV	48.42	7.50	8.44	196.42	30.11
	<b>-18.85%</b>	<b>-60.02%</b>	<b>-9.15%</b>	<b>-9.09%</b>	<b>-91.49%</b>
<b>1x6 grid</b>					
Baseline	89.99	33.13	9.54	221.97	1724.94
FlowCAV	172.12	118.30	17.60	409.44	5200.00
	+91.27%	+257.08%	+84.49%	+84.46%	+201.46%
PressLight	77.59	21.22	8.82	205.13	676.22
	-13.78%	-35.95%	-7.55%	-7.59%	-60.80%
GLOSA	68.91	12.05	7.59	176.49	252.94
	-23.42%	-63.63%	-20.44%	-20.49%	-85.34%
CoTV	65.56	8.70	7.27	169.19	68.28
	<b>-27.15%</b>	<b>-73.74%</b>	<b>-23.79%</b>	<b>-23.78%</b>	<b>-96.04%</b>
<b>Dublin</b>					
Baseline	59.33	29.17	10.98	255.53	1212.67
FlowCAV	59.39	29.43	11.16	259.70	1223.28
	+0.10%	+0.89%	+1.64%	+1.63%	+0.87%
PressLight	44.92	14.69	8.49	197.43	463.11
	-24.29%	-49.64%	-22.68%	-22.74%	-61.81%
GLOSA	45.40	15.06	8.46	196.92	545.50
	-23.48%	-48.37%	-22.95%	-22.94%	-55.02%
CoTV	41.76	11.48	7.97	185.42	195.94
	<b>-29.61%</b>	<b>-60.64%</b>	<b>-27.41%</b>	<b>-27.44%</b>	<b>-83.84%</b>

FlowCAV does not reduce travel time due to its static traffic light plan and the fact that CAV agents do not consider current traffic light signals. The results of FlowCAV in grid road maps are much worse than Baseline. The further improvement of CoTV demonstrates the advantages of cooperative traffic control compared with controlling traffic light signals only, meanwhile indicating that DRL-based approaches provide better adaptive traffic control than traditional approaches. Moreover, Fig.9 illustrates travel time of all vehicles can be reduced significantly and more densely distributed around a lower value under the three scenarios using CoTV, compared with other methods. The results in delay from Table III shows that CoTV reduces the travel time very close to its minimum possible value. Compared with other methods, CoTV can achieve up to about 74% reduction in delay.

- **Environmental indicators:** CoTV brings the best results of fuel consumption and CO<sub>2</sub> emissions, both achieving over 27% reduction shown in Table III. The reduced

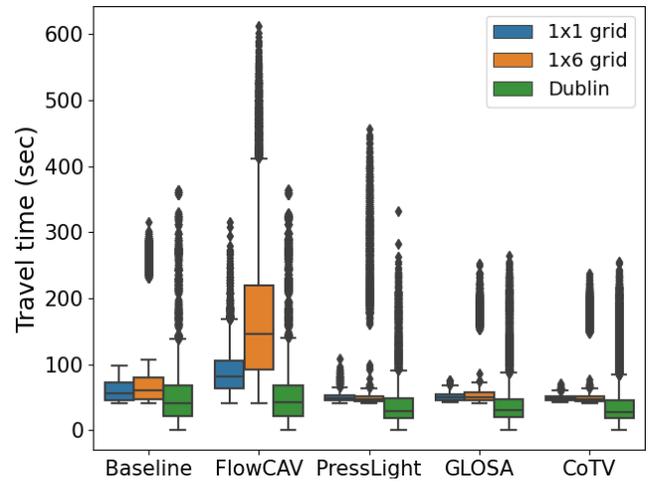


Fig. 9. Travel time distributions for four compared methods under three test scenarios. CoTV can reduce the travel time of all vehicles as all travel time values are densely distributed around the lowest values compared with other methods.

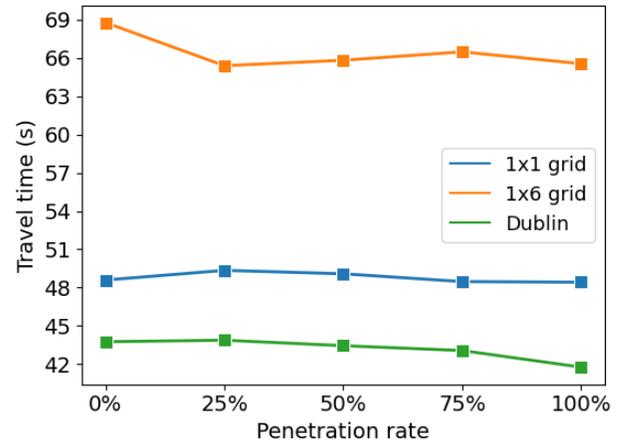


Fig. 10. The average travel time of CoTV under different penetration rates under both grid and Dublin scenarios. Travel time tends to decrease as the CAV penetration rate increases. Note that the travel time of CoTV shown in Table III is when penetration rate is 100%.

travel time of PressLight results in less fuel consumption. GLOSA obtains the second-best results due to the jointly optimised traffic light timings and vehicle speed. However, FlowCAV does not show any improvement on the two environmental indicators due to the complexity of urban scenarios containing intersections.

- **Traffic safety:** CoTV reduces TTC by over 96%, as shown in Table III. PressLight and GLOSA improve traffic safety as well. However, there is a great difference in TTC between PressLight and CoTV under the 1×6 grid scenario, and the result of CoTV under Dublin scenario is much better than the other two methods. Conversely, FlowCAV hurts traffic safety under the grid maps but not in Dublin scenario. The more realistic urban scenario brings explicit complexity to enhance safety. This also highlights the advantages of CoTV using DRL-based methods for the cooperative control.

2) *Robustness to Varying CAV Penetration Rates:* Fig.10 shows that the travel time of CoTV tends to decrease as the

TABLE IV

COMPARISON BETWEEN I-CoTV (INDEPENDENT, NO AGENT COOPERATION), M-CoTV (ACTION-DEPENDENT, AGENT COOPERATION IN ACTION AND STATE), AND CoTV, UNDER FULL-AUTONOMY TRAFFIC UNDER DUBLIN SCENARIO

Method	Travel time (s)	Fuel (l/100km)	TTC	Training time (h)
I-CoTV	49.21	9.19	489.78	1.36
M-CoTV	47.53	10.44	660.08	2.00
CoTV	41.76	7.97	195.94	1.33

CAV penetration rate increases under  $1 \times 1$  and  $1 \times 6$  grid maps and Dublin scenarios. Specifically, CoTV with 0% CAV penetration rate implicates no vehicle speed control, which can still be way better than the Baseline and PressLight shown in Table III. Therefore, CoTV with CAV speed control can get better results, which demonstrates the effectiveness of cooperative traffic control between two types of agents. Similar results are shown in other traffic metrics (e.g., delay, fuel consumption, etc). This demonstrates the practicability of CoTV when deployed in a realistic mixed-autonomy scenario.

3) *Comparison With Other MARL Methods:* To further demonstrate the effectiveness of our CoTV system design on cooperative control, we compare CoTV with two other common MARL methods, I-CoTV (independent, without any cooperation schemes) and M-CoTV (action-dependent, with cooperation schemes in action and state). Results under Dublin scenario with full-autonomy traffic are shown in Table IV. CoTV achieves the best results, while I-CoTV suffers from convergence issues, resulting in the worst traffic performance. M-CoTV fails to overcome high complexity from consideration of other agents' actions, which affects traffic improvements. In particular, the performance changes in fuel consumption and travel time are inconsistent in M-CoTV compared to I-CoTV. The training time of M-CoTV is also increased by about 50%. In addition, referring to Table III, M-CoTV and I-CoTV perform better than Baseline and FlowCAV but do not surpass PressLight and GLOSA.

In summary, CoTV achieves the first three system goals, including reduced travel time, lower fuel consumption and CO<sub>2</sub> emissions, and longer time-to-collision. The cooperation schemes between CAV and traffic light controllers, which is the first contribution of this paper, can overcome the difficulties of DRL-based joint control in complex urban traffic scenarios.

### B. Improved Scalability by Reducing the CAVs Agents

The second contribution of CoTV is the improvement of the multi-agent system scalability by reducing the number of CAV agents controlled. Compared with CoTV\* that trains all possible CAVs, results from Table V indicate that CoTV can reduce the training time by up to 44%, while still having comparable (sometimes slightly better) improvement in both traffic efficiency and safety under Dublin scenario. Although CoTV\* obtains better results under the two grid maps than

TABLE V

COMPARISON BETWEEN CoTV AND CoTV\* (CONTROL ALL POSSIBLE CAVS) UNDER FULL-AUTONOMY TRAFFIC

	Method	Travel time (s)	Fuel (l/100km)	TTC	Training time (h)
<b>1x1 grid</b>	CoTV*	48.34	8.40	29.61	0.70
	CoTV	48.42	8.44	30.11	0.45
<b>1x6 grid</b>	CoTV*	65.34	7.11	61.78	2.57
	CoTV	65.56	7.27	68.28	1.65
<b>Dublin</b>	CoTV*	43.42	7.98	219.67	2.37
	CoTV	41.76	7.97	195.94	1.33

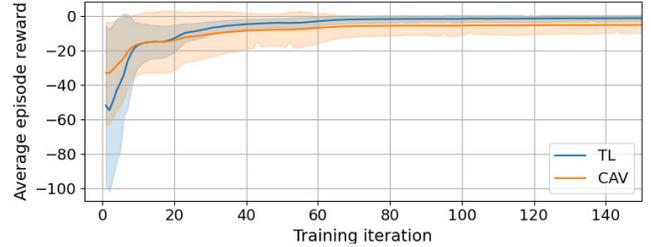


Fig. 11. Evolution of the average episode reward for traffic light controllers (TL) and CAV agent of CoTV under Dublin scenario. The shade represents the standard deviation value. After DRL training on CoTV, the rewards for both types of agents can converge to higher values and smaller standard deviations than in the initial stage.

CoTV, it is worth reminding that CoTV achieves this by only cooperating with the closest CAV on each incoming road for the traffic light controller. The closest CAV has the great potential to increase intersection throughput, which is similar to controlling the leading vehicle only for improving the traffic efficiency of a platoon [13]. The CAV as the leading vehicle is well controlled by CoTV, all its following vehicles are subsequently self-adjusted. Moreover, Fig. 11 indicates that two agent types of CoTV, traffic light controllers and CAVs, can be converged at a higher reward with a small standard deviation than the start after about 60 training iterations. Thus, CoTV can alleviate scalability issues, while also not compromise traffic improvement. The last goal of system design, easier to deploy, is achieved.

### C. Discussion: When CoTV Meets More Complicated Urban Scenarios

To further explore the deployment options of CoTV, we conduct experiments under a relatively large and dense urban scenario in Dublin city centre, which traditionally requires sophisticated coordination between adjacent traffic light controllers. The selected area covers nearly  $1 \text{ km}^2$  with 31 signalized intersections, as shown in Fig. 12. These intersections with different road shapes and traffic light signal cycles/phases are all controlled by CoTV. Table VI shows the traffic performance under this dense Dublin scenario at 100% CAV penetration rate. Although CoTV can get converged and obtain the best results in all evaluation metrics, which shows that CoTV can be deployed in both major and minor junctions, we still need further studies to find the optimal selection of key intersections to control to avoid costly deployment on all urban junctions.

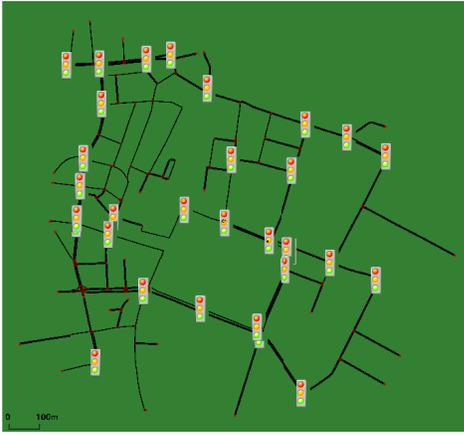


Fig. 12. The selected dense urban scenario in the city centre of Dublin. There are 119 intersections in total, including 31 signalized intersections. 321 vehicles are generated from 10 AM in 400 seconds.

TABLE VI

TRAFFIC PERFORMANCE UNDER A DENSE DUBLIN SCENARIO. PERCENTAGE CHANGES SHOWN ARE COMPARED TO BASELINE. THE BEST ACHIEVED MEASUREMENTS ARE IN BOLD

Method	Travel time (s)	Delay (s)	Fuel (l/100km)	CO <sub>2</sub> (g/km)	TTC
Baseline	125.40	78.84	11.22	261.09	2344.50
FlowCAV	127.71	81.14	11.06	257.03	1947.44
	+1.84%	+2.92%	-1.43%	-1.45%	-16.94%
PressLight	114.60	68.15	9.70	225.59	1800.83
	-8.61%	-13.56%	-13.55%	-13.60%	-23.19%
GLOSA	104.21	57.64	8.50	197.70	1193.61
	-16.90%	-26.89%	-24.24%	-24.28%	-49.09%
CoTV	103.18	56.60	8.40	195.39	787.29
	<b>-17.72%</b>	<b>-28.21%</b>	<b>-25.13%</b>	<b>-25.16%</b>	<b>-66.42%</b>

## VI. CONCLUSION AND FUTURE WORK

This paper proposes a multi-agent DRL system, CoTV, to control traffic light signals and CAV cooperatively for achieving sustainable urban traffic goals. CoTV can significantly improve traffic efficiency (i.e., travel time, fuel consumption, and CO<sub>2</sub> emissions) as well as traffic safety (i.e., time-to-collision), which outperforms other DRL-based systems that control either traffic light signal or vehicle speed, and non-DRL joint control method based on GLOSA. Moreover, the traffic light controllers in our CoTV utilize V2I communications infrastructure to only cooperate with the closest CAV (i.e., as the leader of a platoon) on each incoming road for alleviating the scalability issue of multi-agent DRL systems. This also eases the deployment and achieves the training process to converge within a moderate number of iterations. Experiments in various grid maps and realistic urban scenarios demonstrate the effectiveness of CoTV. Compared to the Baseline, CoTV can save up to 28% in fuel consumption and CO<sub>2</sub> while reducing travel time by up to 30%. The robustness of CoTV is also validated under different penetration rates of CAV. In the future, we plan to investigate further how to efficiently integrate the action-dependent MARL for CoTV.

## REFERENCES

- [1] U. Desa et al., "Transforming our world: The 2030 agenda for sustainable development," United Nations, New York, NY, USA, Tech. Rep. A/RES/70/1, 2016.
- [2] A. Shetty, M. Yu, A. Kurzhanskiy, O. Grembek, H. Tavafoghi, and P. Varaiya, "Safety challenges for autonomous vehicles in the absence of connectivity," *Transp. Res. C, Emerg. Technol.*, vol. 128, Jul. 2021, Art. no. 103133.
- [3] C. Yu, Y. Feng, H. X. Liu, W. Ma, and X. Yang, "Integrated optimization of traffic signals and vehicle trajectories at isolated urban intersections," *Transp. Res. B, Methodol.*, vol. 112, pp. 89–112, Jun. 2018.
- [4] B. Xu et al., "Cooperative method of traffic signal optimization and speed control of connected vehicles at isolated intersections," *IEEE Trans. Intell. Transp. Syst.*, vol. 20, no. 4, pp. 1390–1403, Apr. 2019.
- [5] X. Di and R. Shi, "A survey on autonomous vehicle control in the era of mixed-autonomy: From physics-based to AI-guided driving policy learning," *Transp. Res. C, Emerg. Technol.*, vol. 125, Apr. 2021, Art. no. 103008.
- [6] P. Varaiya, "Max pressure control of a network of signalized intersections," *Transp. Res. C, Emerg. Technol.*, vol. 36, pp. 177–195, Nov. 2013.
- [7] H. Wei et al., "PressLight: Learning max pressure control to coordinate traffic signals in arterial network," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2019, pp. 1290–1298.
- [8] G. Benedetti, M. P. Fanti, A. M. Mangini, and F. Parisi, "Application of deep reinforcement learning for traffic control of road intersection with emergency vehicles," in *Proc. IEEE Int. Conf. Syst., Man, Cybern. (SMC)*, Oct. 2021, pp. 182–187.
- [9] C. Wu, A. Kreidieh, K. Parvate, E. Vinitsky, and A. M. Bayen, "Flow: A modular learning framework for mixed autonomy traffic," 2017, *arXiv:1710.05465*.
- [10] H. Wei, G. Zheng, V. Gayah, and Z. Li, "A survey on traffic signal control methods," 2019, *arXiv:1904.08117*.
- [11] T. Chu, J. Wang, L. Codecà, and Z. Li, "Multi-agent deep reinforcement learning for large-scale traffic signal control," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 3, pp. 1086–1095, Mar. 2020.
- [12] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017, *arXiv:1707.06347*.
- [13] J. Lioris, R. Pedarsani, F. Y. Tascikaraoglu, and P. Varaiya, "Doubling throughput in urban roads by platooning," *IFAC-PapersOnLine*, vol. 49, no. 3, pp. 49–54, 2016.
- [14] P. K. Singh, S. K. Nandi, and S. Nandi, "A tutorial survey on vehicular communication state of the art, and future research directions," *Veh. Commun.*, vol. 18, Aug. 2019, Art. no. 100164.
- [15] J. Guo and S. Wang, "Poster: Can traffic lights and CAV work together using deep reinforcement learning?" in *Proc. IEEE Veh. Netw. Conf. (VNC)*, Nov. 2021, pp. 127–128.
- [16] D. N. Lee, "A theory of visual control of braking based on information about time-to-collision," *Perception*, vol. 5, no. 4, pp. 437–459, Dec. 1976.
- [17] A. G. Sims and K. W. Dobinson, "The Sydney coordinated adaptive traffic (SCAT) system philosophy and benefits," *IEEE Trans. Veh. Technol.*, vol. VT-29, no. 2, pp. 130–137, May 1980.
- [18] R. E. Stern et al., "Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments," *Transp. Res. C, Emerg. Technol.*, vol. 89, pp. 205–221, Apr. 2018.
- [19] H. Suzuki and Y. Marumo, "A new approach to green light optimal speed advisory (GLOSA) systems for high-density traffic flow," in *Proc. 21st Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2018, pp. 362–367.
- [20] E. Vinitsky et al., "Benchmarks for reinforcement learning in mixed-autonomy traffic," in *Proc. Conf. Robot Learn.*, 2018, pp. 399–409.
- [21] M. Tajalli and A. Hajbabaie, "Traffic signal timing and trajectory optimization in a mixed autonomy traffic stream," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 7, pp. 6525–6538, Jul. 2022.
- [22] P. Hernandez-Leal, B. Kartal, and M. E. Taylor, "A survey and critique of multiagent deep reinforcement learning," *Auto. Agents Multi-Agent Syst.*, vol. 33, no. 6, pp. 750–797, Nov. 2019.
- [23] J. Olstam, F. Johansson, A. Alessandrini, P. Sukennik, J. Lohmiller, and M. Friedrich, "An approach for handling uncertainties related to behaviour and vehicle mixes in traffic simulation experiments with automated vehicles," *J. Adv. Transp.*, vol. 2020, pp. 1–17, Sep. 2020.
- [24] I. Postigo, J. Olstam, and C. Rydergren, "Effects on traffic performance due to heterogeneity of automated vehicles on motorways: A microscopic simulation study," in *Proc. 7th Int. Conf. Vehicle Technol. Intell. Transp. Syst.*, 2021, pp. 142–151.

- [25] A. Sharma, Z. Zheng, J. Kim, A. Bhaskar, and M. M. Haque, "Assessing traffic disturbance, efficiency, and safety of the mixed traffic flow of connected vehicles and traditional vehicles by considering human factors," *Transp. Res. C, Emerg. Technol.*, vol. 124, Mar. 2021, Art. no. 102934.
- [26] M. Guérliau and I. Dusparic, "Quantifying the impact of connected and autonomous vehicles on traffic efficiency and safety in mixed traffic," in *Proc. IEEE 23rd Int. Conf. Intell. Transp. Syst. (ITSC)*, Sep. 2020, pp. 1–8.
- [27] H. Wei, X. Liu, L. Mashayekhy, and K. Decker, "Mixed-autonomy traffic control with proximal policy optimization," in *Proc. IEEE Veh. Netw. Conf. (VNC)*, Dec. 2019, pp. 1–8.
- [28] H. Wei, G. Zheng, H. Yao, and Z. Li, "IntelliLight: A reinforcement learning approach for intelligent traffic light control," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2018, pp. 2496–2505.
- [29] C. Wu, A. Kreidieh, K. Parvate, E. Vinitzky, and A. M. Bayen, "Flow: Architecture and benchmarking for reinforcement learning in traffic control," 2017, *arXiv:1710.05465*.
- [30] T. P. Lillicrap et al., "Continuous control with deep reinforcement learning," 2015, *arXiv:1509.02971*.
- [31] C. Chen et al., "Toward a thousand lights: Decentralized deep reinforcement learning for large-scale traffic signal control," in *Proc. AAAI Conf. Artif. Intell.*, 2020, vol. 34, no. 4, pp. 3414–3421.
- [32] X. Liang, X. Du, G. Wang, and Z. Han, "A deep reinforcement learning network for traffic light cycle control," *IEEE Trans. Veh. Technol.*, vol. 68, no. 2, pp. 1243–1253, Feb. 2019.
- [33] K. Ahn, H. Rakha, A. Trani, and M. Van Aerde, "Estimating vehicle fuel consumption and emissions based on instantaneous speed and acceleration levels," *J. Transp. Eng.*, vol. 128, no. 2, pp. 182–190, Mar. 2002.
- [34] M. Keller, S. Hausberger, C. Matzer, P. Wüthrich, and B. Notter, "Handbook of emission factors for road transport (HBEFA) 3.1," Quick Reference, INFRAS, Zurich, Switzerland, Tech. Rep. I-20/2009, 2010.
- [35] M. Treiber and A. Kesting, "The intelligent driver model with stochasticity—new insights into traffic flow oscillations," *Transp. Res. Proc.*, vol. 23, pp. 174–187, Jan. 2017.



**Jiaying Guo** (Graduate Student Member, IEEE) received the B.Sc. degree from University College Dublin, Ireland, and the B.Eng. degree from the Beijing University of Technology, China, in 2020. She is currently pursuing the Ph.D. degree with the School of Computer Science, University College Dublin. Her research interests include the multi-agent systems of intelligent transportation, reinforcement learning, mix-autonomy traffic, and connected autonomous vehicles.



**Long Cheng** (Senior Member, IEEE) received the B.E. degree from the Harbin Institute of Technology, China, in 2007, the M.Sc. degree from the University of Duisburg-Essen, Germany, in 2010, and the Ph.D. degree from the National University of Ireland, Maynooth, in 2014. He was an Assistant Professor with Dublin City University and a Marie Curie Fellow with University College Dublin. He has also worked at many organizations, like Huawei Technologies Germany, IBM Research Dublin, TU Dresden, and TU Eindhoven. He is currently a Full Professor with the School of Control and Computer Engineering, North China Electric Power University, Beijing. He has published more than 80 papers in journals and conferences, like IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, IEEE/ACM TRANSACTIONS ON NETWORKING, IEEE TRANSACTIONS ON COMPUTERS, IEEE TRANSACTIONS ON SERVICES COMPUTING, IEEE TRANSACTIONS ON AUTOMATION SCIENCE AND ENGINEERING, IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN OF INTEGRATED CIRCUITS AND SYSTEMS, IEEE TRANSACTIONS ON CLOUD COMPUTING, IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS, IEEE TRANSACTIONS ON NETWORK AND SERVICE MANAGEMENT, IEEE NETWORK, HPCA, CIKM, ICPP, CCGrid, and Euro-Par. His research interests include distributed systems, deep learning, cloud computing, and process mining. He is the Co-Chair of the *Journal of Cloud Computing*.



**Shen Wang** (Senior Member, IEEE) received the M.Eng. degree from Wuhan University, China, and the Ph.D. degree from Dublin City University, Ireland. Some key industry partners of his applied research are IBM Research Brazil, Boeing Research and Technology Europe, and Huawei Ireland Research Centre. He is currently an Assistant Professor with the School of Computer Science, University College Dublin, Ireland. He has been involved with several EU projects as a co-PI, the WP, and a task leader of big trajectory data streaming for air traffic control and trustworthy AI for intelligent cybersecurity systems. His research interests include connected autonomous vehicles, explainable artificial intelligence, and security and privacy for mobile networks. He was a recipient of the IEEE Intelligent Transportation Systems Society Young Professionals Travelling Fellowship in 2022.