

Multi-Agent Deep Reinforcement Learning for Large-scale Traffic Signal Control

Tianshu Chu, Jie Wang, Lara Codecà, and Zhaojian Li *Member, IEEE*

Abstract—Reinforcement learning (RL) is a promising data-driven approach for adaptive traffic signal control (ATSC) in complex urban traffic networks, and deep neural networks further enhance its learning power. However, centralized RL is infeasible for large-scale ATSC due to the extremely high dimension of the joint action space. Multi-agent RL (MARL) overcomes the scalability issue by distributing the global control to each local RL agent, but it introduces new challenges: now the environment becomes partially observable from the viewpoint of each local agent due to limited communication among agents. Most existing studies in MARL focus on designing efficient communication and coordination among traditional Q-learning agents. This paper presents, for the first time, a fully scalable and decentralized MARL algorithm for the state-of-the-art deep RL agent: advantage actor critic (A2C), within the context of ATSC. In particular, two methods are proposed to stabilize the learning procedure, by improving the observability and reducing the learning difficulty of each local agent. The proposed multi-agent A2C is compared against independent A2C and independent Q-learning algorithms, in both a large synthetic traffic grid and a large real-world traffic network of Monaco city, under simulated peak-hour traffic dynamics. Results demonstrate its optimality, robustness, and sample efficiency over other state-of-the-art decentralized MARL algorithms.

Index Terms—Adaptive traffic signal control, Reinforcement learning, Multi-agent reinforcement learning, Deep reinforcement learning, Actor-critic.

I. INTRODUCTION

As a consequence of population growth and urbanization, the transportation demand is steadily rising in the metropolises worldwide. The extensive routine traffic volumes bring pressures to existing urban traffic infrastructure, resulting in everyday traffic congestions. Adaptive traffic signal control (ATSC) aims for reducing potential congestions in saturated road networks, by adjusting the signal timing according to real-time traffic dynamics. Early-stage ATSC methods solve optimization problems to find efficient coordination and control policies. Successful products, such as SCOOT [1] and SCATS [2], have been installed in hundreds of cities across the world. OPAC [3] and PROLYN [4] are similar methods, but their relatively complex computation makes the implementation less popular. Since the 90s, various interdisciplinary techniques

have been applied to ATSC, such as fuzzy logic [5], genetic algorithm [6], and immune network algorithm [7].

Reinforcement learning (RL), formulated under the framework of Markov decision process (MDP), is a promising alternative to learn ATSC based on real-world traffic measurements [8]. Unlike traditional model-driven approaches, RL does not rely on heuristic assumptions and equations. Rather it directly fits a parametric model to learn the optimal control, based on its experience interacting with the complex traffic systems. Traditional RL fits simple models such as piecewise constant table and linear regression (LR) [9], leading to limited scalability or optimality in practice. Recently, deep neural networks (DNNs) have been successfully applied to enhance the learning capacity of RL on complex tasks [10].

To utilize the power of deep RL, appropriate RL methods need to be adapted. There are three major methods: value-based, policy-based, and actor-critic methods. In value-based methods, such as *Q-learning*, the long-term state-action value function is parameterized and updated using step-wise experience [11]. *Q-learning* is off-policy¹, so it enjoys efficient updating with bootstrapped sampling of *experience replay*. However its update is based on one-step *temporal difference*, so the good convergence relies on a stationary MDP transition, which is less likely in ATSC. As the contrast, in policy-based methods, such as REINFORCE, the policy is directly parameterized and updated with sampled episode return [12]. REINFORCE is on-policy so the transition can be nonstationary within each episode. The actor-critic methods further reduce the bias and variance of policy-based methods by using another model to parameterize the value function [13]. A recent work has demonstrated that actor-critic outperforms *Q-learning* in ATSC with centralized LR agents [14]. This paper focuses on the state-of-the-art *advantage actor-critic* (A2C) where DNNs are used for both policy and value approximations [15].

Even though DNNs have improved the scalability of RL, training a centralized RL agent is still infeasible for large-scale ATSC. First, we need to collect all traffic measurements in the network and feed them to the agent as the global state. This centralized state processing itself will cause high latency and failure rate in practice, and the topological information of the traffic network will be lost. Further, the joint action space of the agent grows exponentially in the number of signalized intersections. Therefore, it is efficient and natural to formulate ATSC as a cooperative multi-agent RL (MARL) problem, where each intersection is controlled by a local RL

Tianshu Chu and Jie Wang are with the Department of Civil and Environmental Engineering, Stanford University, CA 94305, USA. Lara Codecà is with the Communication Systems Department, EURECOM, Sophia-Antipolis 06904, France. Her work was partially funded by the ANR project ref. No. ANR-11-LABX-0031-01 and EURECOM partners: BMW Group, IABG, Monaco Telecom, Orange, SAP, ST Microelectronics and Symantec. Zhaojian Li is with the Department of Mechanical Engineering, Michigan State University, East Lansing, MI 48824, USA. Email: cts198859@hotmail.com, jiewang@stanford.edu, Codecà@eurecom.fr, lizhaojl@egr.msu.edu.

¹In off-policy learning, the behavior policy for sampling the experience is different from the target policy to be optimized.

agent, upon local observation and limited communication. MARL has a long history and has mostly focused on Q-learning, by distributing the global Q-function to local agents. One approach is to design a coordination rule based on the tradeoff between optimality and scalability [16], [17]. A simpler and more common alternative is *independent* Q-learning (IQL) [18], in which each local agent learns its own policy independently, by modeling other agents as parts of the environment dynamics. IQL is completely scalable, but it has issue on convergence, since now the environment becomes more partially observable and nonstationary, as other agents update their policies. This issue was addressed recently for enabling experience replay in deep MARL [19].

To the best of our knowledge, this is the first paper that formulates independent A2C (IA2C) for ATSC, by extending the idea of IQL on A2C. In order to develop a stable and robust IA2C system, two methods are further proposed to address the partially observable and nonstationary nature of IA2C, under limited communication. First, we include the observations and *fingerprints* of neighboring agents in the state, so that each local agent has more information regarding the regional traffic distribution and cooperative strategy. Second, we introduce a *spatial* discount factor to scale down the observation and reward signals of neighboring agents, so that each local agent focuses more on improving traffic nearby. From the convergence aspect, the first approach increases the fitting power while the second approach reduces the fitting difficulty. We call the stabilized IA2C the multi-agent A2C (MA2C). MA2C is evaluated in both a synthetic large traffic grid and a real-world large traffic network, with delicately designed traffic dynamics for ensuring a certain difficulty level of MDP. Numerical experiments confirm that, MA2C outperforms IA2C and state-of-the-art IQL algorithms in robustness and optimality. The code of this study is open sourced².

II. BACKGROUND

A. Reinforcement Learning

RL learns to maximize the long-term return of a MDP. In a fully observable MDP, the agent observes the true state of the environment $s_t \in \mathcal{S}$ at each time t , performs an action $u_t \in \mathcal{U}$ according to a policy $\pi(u|s)$. Then transition dynamics happens $s_{t+1} \sim p(\cdot|s_t, u_t)$, and an immediate step reward $r_t = r(s_t, u_t, s_{t+1})$ is received. In an infinite-horizon MDP, sampled total return under policy π is $R_t^\pi = \sum_{\tau=t}^{\infty} \gamma^{\tau-t} r_\tau$, where $\gamma \in [0, 1)$ is a discount factor. The expected total return is represented as its *Q-function* $Q^\pi(s, u) = \mathbb{E}[R_t^\pi | s_t = s, u_t = u]$. The optimal Q-function $Q^* = \max_{\pi} Q^\pi$ yields an optimal greedy policy $\pi^*(u|s) : u \in \arg\max_{u'} Q^*(s, u')$, and Q^* is obtained by solving the Bellman equation $\mathcal{T}Q^* = Q^*$ [20], with a *dynamic programming* (DP) operator \mathcal{T} :

$$\mathcal{T}Q(s, u) = r(s, u) + \gamma \sum_{s' \in \mathcal{S}} p(s'|s, u) \max_{u' \in \mathcal{U}} Q(s', u'), \quad (1)$$

where $r(s, u) = \mathbb{E}'_s r(s, u, s')$ is the expected step reward. In practice, r and p are unknown to the agent so the above *plan-*

ning problem is not well defined. Instead, RL performs data-driven DP based on the sampled experience (s_t, u_t, s'_t, r_t) .

1) *Q-learning*: Q-learning is the fundamental RL method that fits the Q-function with a parametric model Q_θ , such as Q-value table [11], LR [9], or DNN [10]. Given experience (s_t, u_t, s'_t, r_t) , Eq. (1) is used to estimate $\hat{T}Q_{\theta^-}(s_t, u_t) = r_t + \gamma \max_{u' \in \mathcal{U}} Q_{\theta^-}(s'_t, u')$ using a frozen recent model θ^- , then temporal difference $(\hat{T}Q_{\theta^-} - Q_\theta)(s_t, u_t)$ is used to update θ . A behavior policy, such as ϵ -greedy, is used in Q-learning to explore rich experience to reduce the regression variance. Experience replay is applied in deep Q-learning to reduce the variance furthermore by sampling less correlated experience in each minibatch.

2) *Policy Gradient*: Policy gradient (REINFORCE) directly fits the policy with a parameterized model π_θ [21]. Each update of θ increases the likelihood for selecting the action that has the high “optimality”, measured as the sampled total return. Thus the loss is

$$\mathcal{L}(\theta) = -\frac{1}{|B|} \sum_{t \in B} \log \pi_\theta(u_t | s_t) \hat{R}_t, \quad (2)$$

where each minibatch $B = \{(s_t, u_t, s'_t, r_t)\}$ contains the experience trajectory, *i.e.*, $s'_t = s_{t+1}$, and each return is estimated as $\hat{R}_t = \sum_{\tau=t}^{t_B-1} \gamma^{\tau-t} r_\tau$, here t_B is the last step in minibatch. Policy gradient does not need a behavior policy since $\pi_\theta(u|s)$ naturally performs exploration and exploitation. Further, it is robust to nonstationary transitions within each trajectory since it directly uses return instead of estimating it recursively by \hat{T} . However it suffers from high variance as \hat{R}_t is much more noisy than fitted return $Q^{\pi_{\theta^-}}$.

3) *Advantage Actor-Critic*: A2C improves the policy gradient by introducing a *value* regressor V_w to estimate $\mathbb{E}[R_t^\pi | s_t = s]$ [13]. First, it reduces the bias of sampled return by adding the value of the last state $R_t = \hat{R}_t + \gamma^{t_B-t} V_w(s_{t_B})$; Second, it reduces the variance of sampled return by using $A_t := R_t - V_w(s_t)$, which is interpreted as the sampled advantage $A^\pi(s, u) := Q^\pi(s, u) - V^\pi(s)$. Then Eq. (2) becomes

$$\mathcal{L}(\theta) = -\frac{1}{|B|} \sum_{t \in B} \log \pi_\theta(u_t | s_t) A_t. \quad (3)$$

The loss function for value updating is:

$$\mathcal{L}(w) = \frac{1}{2|B|} \sum_{t \in B} (R_t - V_w(s_t))^2. \quad (4)$$

B. Multi-agent Reinforcement Learning

In ATSC, multiple signalized intersection agents participate a cooperative game to optimize the global network traffic objectives. Consider a multi-agent network $G(\mathcal{V}, \mathcal{E})$, where each agent $i \in \mathcal{V}$ performs a discrete action $u_i \in \mathcal{U}_i$, communicates to a neighbor via edge $ij \in \mathcal{E}$, and shares the global reward $r(s, u)$. Then the joint action space is $\mathcal{U} = \times_{i \in \mathcal{V}} \mathcal{U}_i$, which makes centralized RL infeasible. MARL, mostly formulated in the context of Q-learning, distributes the global action to each local agent by assuming the global Q-function is decomposable $Q(s, u) = \sum_{i \in \mathcal{V}} Q_i(s, u)$. To simplify the notation, we assume each local agent can observe the global state, and we will relax this assumption for ATSC in Section IV and V.

²See https://github.com/cts198859/deeprl_signal_control.

Coordinated Q-learning is one MARL approach that performs iterative message passing or control syncing among neighboring agents to achieve desired tradeoff between optimality and scalability. In other words, $Q_i(s, u) \approx Q_i(s, u_i) + \sum_{j \in \mathcal{V}_i} M_j(s, u_j, u_{\mathcal{V}_j})$, where \mathcal{V}_i is the neighborhood of agent i , and M_j is the message from neighbor j , regarding the impact of $u_i \in u_{\mathcal{V}_j}$ on its local traffic condition. [22] applied variable elimination after passing Q-function as the message, while [17] proposed a max-plus message passing. This approach (1) requires additional computation to obtain the coordinated control during execution, and (2) requires heuristics and assumptions to decompose Q-function and formulate message-passing, which can be potentially learned by IQL described shortly.

IQL is the most straightforward and popular approach, in which each local Q-function only depends on the local action, i.e., $Q_i(s, u) \approx Q_i(s, u_i)$ [18]. IQL is completely scalable, but without message passing, it suffers from partial observability and non-stationary MDP, because it implicitly formulates all other agents' behavior as part of the environment dynamics while their policies are continuously updated during training. To address this issue, each local agent needs the information of other agents' policies. [23] included the policy network parameter of each other agent for fitting local Q-function, i.e., $Q_i(s, u) \approx Q_{\theta_i}(s, u_i, \theta_{-i})$, while [19] included low-dimensional fingerprints, i.e., $Q_i(s, u) \approx Q_{\theta_i}(s, u_i, x_{-i})$. To handle the nonstationary transition in experience replay, *importance sampling* is applied to estimate the temporal difference of other agents' policies between the sampled time t and the updating time τ , as $\pi_{\tau, -i}(u_{-i}|s)/\pi_{t, -i}(u_{-i}|s)$, and to adjust the batch loss. MARL with A2C has not been formally addressed in literature, and will be covered in Section IV.

III. RELATED WORK

The implementation of RL has been extensively studied in ATSC. Tabular Q-learning was the first RL algorithm applied, at an isolated intersection [24]. Later, LR Q-learning was adapted for scalable fitting over continuous states. [25] and [26] designed heuristic state features, while [27] integrated macroscopic fundamental diagram to obtain more informative features. However, LR was too simple to capture the Q-function under complex traffic dynamics. Thus kernel method was applied to extract nonlinear features from low-dimensional states [28]. Kernel method was also applied in LR actor-critic recently, under realistically simulated traffic environments [14]. Alternatively, natural actor-critic was applied to improve the fitting accuracy of LR in ATSC [29]. Deep RL was implemented recently, but most of them had impractical assumptions or oversimplified traffic environments. [30] and [31] verified the superior fitting power of deep Q-learning and deep deterministic policy gradient, respectively, under simplified traffic environment. [32] applied deep Q-learning in a more realistic traffic network, but with infeasible microscopic discrete traffic state. [33] explored A2C with different types of state information. Due to the scalability issue, most centralized RL studies conducted experiments in either isolated intersections or small traffic networks.

Despite rich history of RL, only a few studies have addressed MARL in ATSC, and most of them have focused on Q-learning. [34] applied model-based tabular IQL to each intersection while [35] extended LR IQL to dynamically clustered regions to improve observability. [36] studied LR IQL and its on-policy version independent SARSA learning, with observability improved by neighborhood information sharing. [37] applied deep IQL, with the partial observability addressed with *transfer planning*, but the states were infeasible and the simulated traffic environments were oversimplified in that study. Alternatively, coordinated Q-learning was implemented with various message-passing methods. [38] designed a heuristic neighborhood communication for tabular Q-learning agents, where each message contained the estimated neighbor policies.

[39] proposed a junction-tree based hierarchical message-passing rule to coordinate tabular Q-learning agents. [40] applied the max-sum communication for LR Q-learning agents, where each message indicated the impact of a neighbor agent on each local Q-value.

To summarize, traditional RL has been widely applied in ATSC, and some works have proposed realistic state measurements and decent traffic dynamics based on insightful domain-specific knowledge. However, benchmark traffic environments are still missing for fair comparisons across proposed algorithms. On the other hand, there is still no comprehensive and realistic studies proposed for implementing deep RL in practical ATSC. MARL has been addressed in a few works, mostly under the context of Q-learning.

IV. MULTI-AGENT ADVANTAGE ACTOR-CRITIC

MARL is mostly addressed in the context of Q-learning. In this section, we first formulate IA2C by extending the observations of IQL to the actor-critic method. Furthermore, we propose two approaches to stabilizing IA2C as MA2C. The first approach is inspired from the works for stabilizing IQL [19], [23], where the recent policies of other agents are informed to each local agent. The second approach proposes a novel spatial discount factor to scale down the signals from agents far away. In other words, each local value function is smoother and more correlated to local states, and each agent focuses more on improving local traffic condition. As a result, the convergence becomes more stable even under limited communication and partial observation. This section proposes a general MA2C framework under limited neighborhood communication, while its implementation details for ATSC will be described in Section V.

A. Independent A2C

In a multi-agent network $G(\mathcal{V}, \mathcal{E})$, i and j are neighbors if there is an edge between them. The neighborhood of i is denoted as \mathcal{N}_i and the local region is $\mathcal{V}_i = \mathcal{N}_i \cup i$. Also, the distance between any two agents $d(i, j)$ is measured as the minimum number of edges connecting them. For example, $d(i, i) = 0$ and $d(i, j) = 1$ for any $j \in \mathcal{N}_i$. In IA2C, each agent learns its own policy π_{θ_i} and the corresponding value function V_{w_i} .

We start with a strong assumption where the global reward and state are shared among agents. Then centralized A2C updating can be easily extended to IA2C, by estimating local return as

$$R_{t,i} = \hat{R}_t + \gamma^{t_B-t} V_{w_i^-}(s_{t_B} | \pi_{\theta_{-i}^-}). \quad (5)$$

The value gradient $\nabla \mathcal{L}(w_i)$ is consistent since \hat{R}_t is sampled from the same stationary policy π_{θ^-} . To obtain policy gradient $\nabla \mathcal{L}(\theta_i)$, $V_{w_i} : \mathcal{S} \times \mathcal{U}_i \rightarrow \mathbb{R}$ is served as the estimation of marginal impact of π_{θ_i} on the future return. However, if each agent follows Eq. (5), each value gradient will be identical towards the global value function V^π instead of the local one $V^{\pi_i} = \mathbb{E}_{\pi_{-i}} V^\pi$. As far as θ_{-i} is fixed, π_{θ_i} will still converge to the best according policy under this updating, and optimal policy $\pi_{\theta_i^*}$ can be achieved if $\theta_{-i} = \theta_{-i}^*$. However, when θ_{-i} is actively updated, the policy gradient may be inconsistent across minibatches, since the advantage is conditioned on changing $\pi_{\theta_{-i}}$, even it is stationary per trajectory.

Global information sharing is infeasible in real-time ATSC due to the communication latency, so we assume the communication is limited to each local region. In other words, local policy and value regressors take $s_{t,\mathcal{V}_i} := \{s_{t,j}\}_{j \in \mathcal{V}_i}$ instead of s_t as the input state. Global reward is still allowed since it is only used in offline training. Eq. (5) is valid here, by replacing the value estimation of the last state with $V_{w_i^-}(s_{t_B,\mathcal{V}_i} | \pi_{\theta_{-i}^-})$. Then the value loss Eq. (4) becomes

$$\mathcal{L}(w_i) = \frac{1}{2|B|} \sum_{t \in B} (R_{t,i} - V_{w_i}(s_{t,\mathcal{V}_i}))^2. \quad (6)$$

Clearly, V_{w_i} suffers from partial observability, as s_{t,\mathcal{V}_i} is a subset of s_t while $\mathbb{E} R_{t,i}$ depends on s_t . Similarly, the policy loss Eq. (3) becomes

$$\mathcal{L}(\theta_i) = -\frac{1}{|B|} \sum_{t \in B} \log \pi_{\theta_i}(u_{t,i} | s_{t,\mathcal{V}_i}) A_{t,i}, \quad (7)$$

where $A_{t,i} = R_{t,i} - V_{w_i^-}(s_{t,\mathcal{V}_i})$. The nonstationary updating issue remains, since $R_{t,i}$ is conditioned on the current policy $\pi_{\theta_{-i}^-}$, while both θ_{-i}^- and w_i^- are updated under the previous policy $\pi_{\theta_{-i}^-}$. This inconsistency effect can be mitigated if each local policy updating is smooth, *i.e.*, the KL divergence $D_{KL}(\pi_{\theta_{-i}^-} || \pi_{\theta_{-i}^-})$ is small enough, but it will slow down the convergence. Partial observability also exists as $\pi_{\theta_i}(\cdot | s_{t,\mathcal{V}_i})$ cannot fully capture the impact of \hat{R}_t .

B. Multi-agent A2C

In order to stabilize IA2C convergence and enhance its fitting power, we propose two approaches. First, we include information of neighborhood policies to improve the observability of each local agent. Second, we introduce a spatial discount factor to weaken the state and reward signals from other agents. In IQL, additional information is included to represent the other agents' behavior policies. [23] directly includes the Q-value network parameters θ_{-i}^- , while [19] includes low-dimensional fingerprints, such as ϵ of the ϵ -greedy exploration and the number of updates so far. Fortunately,

the behavior policy is explicitly parameterized in A2C, so a natural approach is including probability simplex of policy $\pi_{\theta_{-i}^-}$. Under limited communication, we include sampled latest policies of neighbors $\pi_{t-1,\mathcal{N}_i} = [\pi_{t-1,j}]_{j \in \mathcal{N}_i}$ in the DNN inputs, besides the current state s_{t,\mathcal{V}_i} . The sampled local policy is calculated as

$$\pi_{t,i} = \pi_{\theta_i^-}(\cdot | s_{t,\mathcal{V}_i}, \pi_{t-1,\mathcal{N}_i}). \quad (8)$$

Rather than long-term neighborhood behavior, the real-time recent neighborhood policy is informed to each local agent. This is based on two ATSC facts: 1) the traffic state is changing slowly in short windows, so the current step policy is very similar to last step policy. 2) the traffic state dynamics is Markovian, given the current state and policy.

Even if the local agent knows the local region state and the neighborhood policy, it is still difficult to fit the global return by local value regressor. To relax the global cooperation, we assume the global reward is decomposable as $r_t = \sum_{i \in \mathcal{V}} r_{t,i}$, which is mostly valid in ATSC. Then we introduce a spatial discount factor α to adjust the global reward for agent i as

$$\tilde{r}_{t,i} = \sum_{d=0}^{D_i} \left(\sum_{j \in \mathcal{V} | d(i,j)=d} \alpha^d r_{t,j} \right), \quad (9)$$

where D_i is the maximum distance from agent i . Note α is similar to the temporal discount factor γ in RL, rather it scales down the signals in spatial order instead of temporal order. Compared to sharing the same global reward across agents, the discounted global reward is more flexible for the trade-off between greedy control ($\alpha = 0$) and cooperative control ($\alpha = 1$), and is more relevant for estimating the ‘‘advantage’’ of local policy π_{θ_i} . Similarly, we use α to discount the neighborhood states as

$$\tilde{s}_{t,\mathcal{V}_i} = [s_{t,i}] \cup \alpha [s_{t,j}]_{j \in \mathcal{N}_i}. \quad (10)$$

Given the discounted global reward, we have $\hat{R}_{t,i} = \sum_{\tau=t}^{t_B-1} \gamma^{\tau-t} \tilde{r}_{\tau,i}$, and the local return Eq. (5) becomes

$$\tilde{R}_{t,i} = \hat{R}_{t,i} + \gamma^{t_B-t} V_{w_i^-}(\tilde{s}_{t_B,\mathcal{V}_i}, \pi_{t_B-1,\mathcal{N}_i} | \pi_{\theta_{-i}^-}). \quad (11)$$

The value loss Eq. (6) becomes

$$\mathcal{L}(w_i) = \frac{1}{2|B|} \sum_{t \in B} \left(\tilde{R}_{t,i} - V_{w_i}(\tilde{s}_{t,\mathcal{V}_i}, \pi_{t-1,\mathcal{N}_i}) \right)^2. \quad (12)$$

This updating is more stable since (1) additional fingerprints π_{t-1,\mathcal{N}_i} are input to V_{w_i} for fitting $\pi_{\theta_{-i}^-}$ impact, and (2) spatially discounted return $\tilde{R}_{t,i}$ is more correlated to local region observations $(\tilde{s}_{t,\mathcal{V}_i}, \pi_{t-1,\mathcal{N}_i})$. The policy loss Eq. (13) becomes

$$\begin{aligned} \mathcal{L}(\theta_i) = & -\frac{1}{|B|} \sum_{t \in B} \left(\log \pi_{\theta_i}(u_{t,i} | \tilde{s}_{t,\mathcal{V}_i}, \pi_{t-1,\mathcal{N}_i}) \tilde{A}_{t,i} \right. \\ & \left. - \beta \sum_{u_i \in \mathcal{U}_i} \pi_{\theta_i} \log \pi_{\theta_i}(u_i | \tilde{s}_{t,\mathcal{V}_i}, \pi_{t-1,\mathcal{N}_i}) \right) \end{aligned} \quad (13)$$

where $\tilde{A}_{t,i} = \tilde{R}_{t,i} - V_{w_i}(\tilde{s}_{t,\mathcal{V}_i}, \pi_{t-1, \mathcal{N}_i})$, and the additional regularization term is the entropy loss of policy π_{θ_i} for encouraging the early-stage exploration. This new advantage emphasizes more on the marginal impact of π_{θ_i} on traffic state within local region.

Algorithm 1 illustrates the MA2C algorithm under the synchronous updating. T is the planning horizon per episode, $|B|$ is the minibatch size, η_w and η_θ are the learning rates for critic and actor networks. First, each local agent collects the experience by following the current policy, until enough samples are collected for minibatch updating (lines 3 to 8). If the episode is terminated in the middle of minibatch collection, we simply restart a new episode (lines 9 to 11). However the termination affects the return estimation. If T is reached in the middle, the trajectory return (line 14) should be adjusted as

$$\hat{R}_{t,i} = \begin{cases} \sum_{\tau=t}^{T-1} \gamma^{\tau-t} \tilde{r}_{\tau,i} & \text{before reset,} \\ \sum_{\tau=t}^{t_B-1} \gamma^{\tau-t} \tilde{r}_{\tau,i} & \text{after reset.} \end{cases} \quad (14)$$

If T is reached at the end, $\tilde{R}_{t,i} = \hat{R}_{t,i}$ (line 15), without the value estimation on future return in Eq. (11). Next, the minibatch gradients are applied to update each actor and critic networks (lines 16, 17), with constant or adaptive learning rates. Usually the first order gradient optimizers are used, such as stochastic gradient descent, RMSprop, and Adam. Finally, the training process is terminated if the maximum step is reached or a certain stop condition is triggered.

Algorithm 1: Synchronous multi-agent A2C

Parameters: $\alpha, \beta, \gamma, T, |B|, \eta_w, \eta_\theta$.
Result: $\{w_i\}_{i \in \mathcal{V}}, \{\theta_i\}_{i \in \mathcal{V}}$.

- 1 **initialize** $s_0, \pi_{-1}, t \leftarrow 0, k \leftarrow 0, B = \emptyset$;
- 2 **repeat**
 - /* explore experience */
 - 3 **for** $i \in \mathcal{V}$ **do**
 - 4 **sample** $u_{t,i}$ from $\pi_{t,i}$ (Eq. (8));
 - 5 **receive** $\tilde{r}_{t,i}$ (Eq. (9)) and $\tilde{s}_{t,i}$;
 - 6 **end**
 - 7 $B \leftarrow B \cup \{(t, \tilde{s}_{t,i}, \pi_{t,i}, u_{t,i}, \tilde{r}_{t,i}, \tilde{s}_{t+1,i})\}_{i \in \mathcal{V}}$;
 - 8 $t \leftarrow t + 1, k \leftarrow k + 1$;
 - /* restart episode */
 - 9 **if** $t = T$ **then**
 - 10 **initialize** $s_0, \pi_{-1}, t \leftarrow 0$;
 - 11 **end**
 - /* update A2C */
 - 12 **if** $k = |B|$ **then**
 - 13 **for** $i \in \mathcal{V}$ **do**
 - 14 **estimate** $\hat{R}_{\tau,i}, \forall \tau \in B$;
 - 15 **estimate** $\tilde{R}_{\tau,i}, \forall \tau \in B$;
 - 16 **update** w_i with $\eta_w \nabla \mathcal{L}(w_i)$ (Eq. (12));
 - 17 **update** θ_i with $\eta_\theta \nabla \mathcal{L}(\theta_i)$ (Eq. (13));
 - 18 **end**
 - 19 $B \leftarrow \emptyset, k \leftarrow 0$;
 - 20 **end**
- 21 **until** *Stop condition reached*;

V. MA2C FOR TRAFFIC SIGNAL CONTROL

This section describes the implementation details of MA2C for ATSC, under the microscopic traffic simulator SUMO [41]. Specifically, we address the action, state, reward definitions, the A2C network structures and normalizations, the A2C training tips, and the evaluation metrics.

A. MDP Settings

Consider a simulated traffic environment over a period of T_s seconds, we define Δt as the interaction period between RL agents and the traffic environment, so that the environment is simulated for Δt seconds after each MDP step. If Δt is too long, RL agents will not be adaptive enough. If Δt is too short, the RL decision will not be delivered on time due to computational cost and communication latency. Further, there will be safety concerns if RL control is switched too frequently. To further guarantee the safety, a yellow time $t_y < \Delta t$ is enforced after each traffic light switch. A recent work suggested $\Delta t = 10s$, and $t_y = 5s$ [14]. In this paper, we set $\Delta t = 5s$, $t_y = 2s$ to ensure more adaptiveness, resulting in a planning horizon of $T = T_s / \Delta t$ steps.

1) *Action definition:* There are several standard action definitions, such as phase switch [38], phase duration [14], and phase itself [26]. We follow the last definition and simply define each local action as a possible phase, or red-green combinations of traffic lights at that intersection. This enables more flexible and direct ATSC by RL agents. Specifically, we pre-define \mathcal{U}_i as a set of all feasible phases for each intersection, and RL agent selects one of them to be implemented for a duration of Δt at each step.

2) *State definition:* After combining the ideas of [26] and [14], we define the local state as

$$s_{t,i} = \{\text{wait}_t[l], \text{wave}_t[l]\}_{j \in \mathcal{E}, l \in L_{j_i}}, \quad (15)$$

where l is each incoming lane of intersection i . wait [s] measures the cumulative delay of the first vehicle, while wave [veh] measures the total number of approaching vehicles along each incoming lane, within 50m to the intersection. Both wait and wave can be obtained by near-intersection induction-loop detectors (ILD), which ensures real-time ATSC. To simplify the implementation, we use `laneAreaDetector` in SUMO to collect the state information.

3) *Reward definition:* Various objectives are selected for ATSC, but an appropriate reward of MARL should be spatially decomposable and frequently measurable. Here we refine the definition in [26] as

$$r_{t,i} = - \sum_{j \in \mathcal{E}, l \in L_{j_i}} (\text{queue}_{t+\Delta t}[l] + a \cdot \text{wait}_{t+\Delta t}[l]), \quad (16)$$

where a [veh/s] is a tradeoff coefficient, and queue [veh] is the measured queue length along each incoming lane. Notably the reward is post-decision so both queue and wait are measured at time $t + \Delta t$. We prefer this definition over others like cumulative delay [38] and wave [14], since this reward is directly correlated to state and action, and emphasizes both traffic congestion and trip delay.

B. DNN Settings

1) *Network architecture*: In practice, the traffic flows are complex spatial-temporal data, so the MDP may become nonstationary if the agent only knows the current state. A straightforward approach is to input all historical states to A2C, but it increases the state dimension significantly and may reduce the attention of A2C on recent traffic condition. Fortunately, *long-short term memory* (LSTM) is a promising DNN layer that maintains hidden states to memorize short history [42]. Thus we apply LSTM as the last hidden layer to extract representations from different state types. Also, we train actor and critical DNNs separately, instead of sharing lower layers between them. Fig. 1 illustrates the DNN structure, where *wave*, *wait*, and *neighbor* policies are first processed by separate fully connected (FC) layers. Then all hidden units are combined and input to a LSTM layer. The output layer is softmax for actor and linear for critic. For DNN training, we use the state-of-the-art orthogonal initializer [43] and RMSprop as the gradient optimizer.

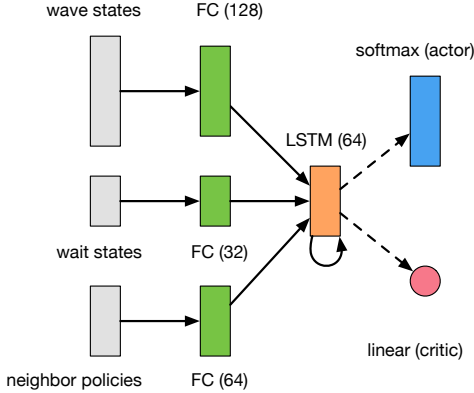


Fig. 1: Proposed DNN structures of MA2C in ATSC. Hidden layer size is indicated inside parenthesis.

2) *Normalization*: Normalization is important for training DNN. For each of *wave* and *wait* states, we run a greedy policy to collect the statistics for a certain traffic environment, and use it to obtain an appropriate normalization. To prevent gradient explosion, all normalized states are clipped to $[0, 2]$, and each gradient is capped at 40. Similarly, we normalize the reward and clip it to $[-2, 2]$ to stabilize the minibatch updating.

VI. NUMERICAL EXPERIMENTS

MARL based ATSC is evaluated in two SUMO-simulated traffic environments: a 5×5 synthetic traffic grid, and a real-world 30-intersection traffic network extracted from Monaco city [44], under time-variant traffic flows. This section aims to design challenging and realistic traffic environments for interesting and fair comparisons across controllers.

A. General Setups

To demonstrate the efficiency and robustness of MA2C, we compare it to several state-of-the-art benchmark controllers. IA2C is the same as MA2C except the proposed stabilizing methods, so it follows the updating rules Eq. (6) and Eq. (7).

IQL-LR is LR based IQL, which can be regarded as the decentralized version of [26]. IQL-DNN is the same as IQL-LR but uses DNN for fitting the Q-function. To ensure fair comparison, IQL-DNN has the same network structure, except the LSTM layer is replaced by a FC layer. This is because Q-learning is off-policy and it becomes meaningless to use LSTM on randomly sampled history. Finally, Greedy is a decentralized greedy policy that selects the phase associated with the maximum total green wave over all incoming lanes. All controllers have the same action space, state space, and interaction frequency.

We train all MARL algorithms over 1M steps, which is around 1400 episodes given episode horizon $T = 720$ steps. We then evaluate all controllers over 10 episodes. Different random seeds are used for generating different training and evaluation episodes, but the same seed is shared for the same episode. For MDP, we set $\gamma = 0.99$ and $\alpha = 0.75$; For IA2C and MA2C, we set $\eta_\theta = 5e-4$, $\eta_w = 2.5e-4$, $|B| = 120$, and $\beta = 0.01$ in Algorithm 1; For IQL, we set the learning rate $\eta_\theta = 1e-4$, the minibatch size $|B| = 20$, and the replay buffer size 1000. Note the replay buffer size has too be small due to the partial observability of IQL. Also, ϵ -greedy is used as the behavior policy, with ϵ linearly decaying from 1.0 to 0.01 during the first half of training.

B. Synthetic Traffic Grid

1) *Experiment settings*: The 5×5 traffic grid, as illustrated in Fig. 2, is formed by two-lane arterial streets with speed limit 20m/s and one-lane avenues with speed limit 11m/s. The action space of each intersection contains five possible phases: E-W straight phase, E-W left-turn phase, and three straight and left-turn phases for E, W, and N-S. Clearly, centralized RL agent is infeasible since the joint action space has the size of 5^{25} . To make the MDP problem challenging, four time-variant traffic flow groups are simulated. At beginning, three major flows F_1 are generated with origin-destination (O-D) pairs $x_{10}-x_4$, $x_{11}-x_5$, and $x_{12}-x_6$, meanwhile three minor flows f_1 are generated with O-D pairs x_1-x_7 , x_2-x_8 , and x_3-x_9 . After 15 minutes, the volumes of F_1 and f_1 start to decrease, while their opposite flows (with swapped O-D pairs) F_2 and f_2 start to be generated, as shown in Fig. 3. Here the flows define the high-level traffic demand only, and the route of each vehicle is randomly generated during run-time. Regarding MDP settings, the reward coefficient a is 0.2veh/s, and the normalization factors of *wave*, *wait*, and reward are 5veh, 100s, and 2000veh, respectively.

2) *Training results*: Fig. 4 plots the training curve of each MARL algorithm, where the solid line shows the average reward per training episode

$$\bar{R} = \frac{1}{T} \sum_{t=0}^{T-1} \left(\sum_{i \in \mathcal{V}} r_{t,i} \right), \quad (17)$$

and the shade shows its standard deviation. Typically, training curve increases and then converges, as RL learns from cumulated experience and finally achieves a local optimum. In Fig. 4, IQL-DNN is failed to learn, while IQL-LR achieves a performance as good as MA2C does. This may because DNN




Fig. 2: A traffic grid of 25 intersections, with an example intersection shown inside circle. Time variant major and minor traffic flow groups are shown as solid and dotted arrows.




Fig. 3: Traffic flows vs simulation time within the traffic grid.




Fig. 4: MARL training curves for synthetic traffic grid.

overfits the Q-function using the partially observed states, misleading the exploitation when ϵ decreases. On the other hand, MA2C shows the best and the most robust learning ability, as its training curve steadily increases and then becomes stable with narrow shade.




Fig. 5: Average queue length in synthetic traffic grid.

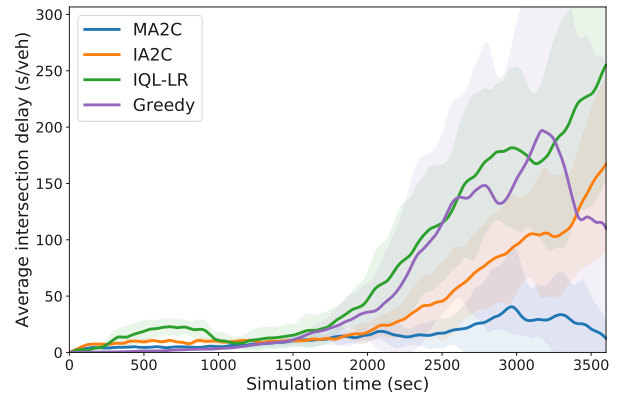


Fig. 6: Average intersection delay in synthetic traffic grid.

3) *Evaluation results:* IQL-DNN is not included in evaluation as its policy is meaningless. The average \bar{R} over evaluation are -414, -845, -972, and -1409, for MA2C, IA2C, Greedy, and IQL-LR. Clearly MA2C outperforms other controllers for the given objective. Also, IQL-LR is failed to beat IA2C in this evaluation over more episodes, which may due to the high variance in the learned policy. Fig. 5 plots the average queue length of network at each simulation step, where the line shows the average and the shade shows the standard deviation across evaluation episodes. Both IQL-LR and IA2C fail to learn a sustainable policy to recover the congested network near the end. As contrast, a simple greedy policy does well for queue length optimization by maximizing the flow at each step, but its performance variation is high (see wide shade). MA2C learns a more stable and sustainable policy that achieves lower congestion level and faster recovery, by paying a higher queue cost when the network is less loaded at early stage.

Fig. 6 plots the average intersection delay of network over simulation time. As expected, both IQL-LR and IA2C have monotonically increasing delays as they are failed to recover from the congestion. However, IA2C is able to maintain a more slowly increasing delay when the traffic grid is less saturated, so its overall performance is still better than the greedy policy, which does not explicitly optimize the waiting time. MA2C is able to maintain the delay at low level even at the peak.

TABLE I: ATSC performance in Monaco traffic network. Best values are in bold.

Metrics	Temporal Averages					Temporal Peaks				
	Greedy	MA2C	IA2C	IQL-LR	IQL-DNN	Greedy	MA2C	IA2C	IQL-LR	IQL-DNN
reward	-41.8	-31.4	-54.6	-109.8	-151.8	-86.4	-78.7	-117.9	-202.1	-256.2
avg. queue length [veh]	0.51	0.29	0.52	1.19	1.57	1.08	0.75	1.16	2.21	2.69
avg. intersection delay [s/veh]	65.5	23.5	60.7	100.3	127.0	272	104	316	202	238
avg. vehicle speed [m/s]	6.06	6.81	5.36	4.04	2.07	14.96	14.26	14.26	14.26	13.98
trip completion flow [veh/s]	0.54	0.67	0.62	0.44	0.28	2.10	2.40	2.10	1.60	1.20
trip delay [s]	144	114	165	207	360	2077	1701	2418	2755	3283

C. Monaco Traffic Network

1) *Experiment settings*: Fig. 7 illustrates the studied area of Monaco city for this experiment, with a variety of road and intersection types. There are 30 signalized intersections in total: 11 are two-phase, 4 are three-phase, 10 are four-phase, 1 is five-phase, and the rest 4 are six-phase. Further, in order to test the robustness and optimality of algorithms in challenging ATSC scenarios, intensive, stochastic, and time-variant traffic flows are designed to simulate the peak-hour traffic. Specifically, four traffic flow groups are generated as a multiple of “unit” flows of 325veh/hr, with randomly sampled O-D pairs inside given areas (see Fig. 7). Among them, F_1 and F_2 are simulated during the first 40min, as $[1, 2, 4, 4, 4, 4, 2, 1]$ unit flows with 5min intervals, while F_3 and F_4 are generated during a shifted time window from 15min to 55min.




Fig. 7: Monaco traffic network, with signalized intersections colored in blue. Four traffic flow groups are illustrated by arrows, with origin and destination inside rectangular areas.

As the MDP becomes challenging in this experiment, we remove `wait` terms in both reward and state, making the value function easier to be fitted. As the price, MARL algorithms will not learn to explicitly optimize the delay. The normalization factor of `wave` is 5veh, and that of reward is 20veh per intersection involved in reward calculation.

2) *Training results*: Fig. 8 plots the training curves of MARL algorithms. IQL-DNN still does not learn anything, while IQL-LR does not converge, despite good performance in the middle of training. Again, both IA2C and MA2C converge




Fig. 8: MARL training curves for Monaco traffic network.

to reasonable policies, and MA2C shows a faster and more stable convergence.

3) *Evaluation results*: Table I summarizes the key performance metrics for comparing ATSC in evaluation. The measurements are first spatially aggregated at each time over evaluation episodes, then the temporal average and max are calculated. Except for trip delay, the values are calculated over all evaluated trips. MA2C outperforms other controllers in almost all metrics. Unfortunately, other MARL algorithms are failed to beat the greedy policy in this experiment. Therefore extensive effort and caution is needed before the field deployment of any data-driven ATSC algorithm, regarding its robustness, optimality, efficiency, and safety.

IQL algorithms are not included in following analysis as their training performance is as bad as that of random exploration. Fig. 9 and Fig. 10 plot the average queue length and average intersection delay over simulation time, under different ATSC policies. As expected, both IA2C and Greedy are able to reduce the queue lengths after peak values, and IA2C achieves a better recovery rate. However, both of them are failed to maintain sustainable intersection delays. This may be because of the “central area” congestion after upstream intersections greedily maximizing their local flows. On the other hand, MA2C is able to achieve lower and more sustainable intersection delays, by distributing the traffic more homogeneously among intersections via coordination with shared neighborhood fingerprints.

Fig. 11 scatters the output (trip completion) flow and network vehicle accumulation for different ATSC algorithms. Macroscopic fundamental diagram (MFD) is present for each controller: when the network density is low, output increases as accumulation grows; when the network becomes more sat-




Fig. 9: Average queue length in Monaco traffic network.




Fig. 10: Average intersection delay in Monaco traffic network.




Fig. 11: Output flow vs vehicle accumulation scatter in Monaco traffic network. Each point is aggregated over 5min.

urated, further accumulation will decrease the output, leading to a potential congestion. As we can see, compared to other controllers, MA2C is able to maintain most points around the “sweet-spot”, maximizing the utilization of network capacity.

VII. CONCLUSIONS

This paper proposed a novel A2C based MARL algorithm for scalable and robust ATSC. Specifically, 1) fingerprints of neighbors were adapted to improve observability; and 2) a


spatial discount factor was introduced to reduce the learning difficulty. Experiments in a synthetic traffic grid and a Monaco traffic network demonstrated the robustness, optimality, and scalability of the proposed MA2C algorithm, which outperformed other state-of-the-art MARL algorithms.

Non-trivial future works are still remaining for the real-world deployment of proposed MARL algorithm. These include 1) improving the reality of traffic simulator to provide reliable training data regarding real-world traffic demand and dynamics; 2) improving the algorithm robustness on noisy and delayed state measurements from road sensors; 3) building a pipeline that can train and deploy deep MARL models to each intersection controller for a given traffic scenario; 4) improving the end-to-end pipeline latency, with a focus on the inference time and memory consumption of model query at each intersection, as well as the communication delay among neighboring intersections for state and fingerprint sharing.

REFERENCES

- [1] P. B. Hunt, D. I. Robertson, R. D. Bretherton, and M. C. Royle, “The SCOOT on-line traffic signal optimisation technique,” *Traffic Engineering & Control*, vol. 23, no. 4, 1982.
- [2] J. Y. K. Luk, “Two traffic-responsive area traffic control methods: SCAT and SCOOT,” *Traffic engineering & control*, vol. 25, no. 1, pp. 14–22, 1984.
- [3] N. H. Gartner, *Demand-responsive Decentralized Urban Traffic Control*. US Department of Transportation, Research and Special Programs Administration, 1982.
- [4] J.-J. Henry, J.-L. Farges, and J. Tuffal, “The PRODYN real time traffic algorithm,” in *Proceedings of the IFAC/IFIP/FORS Conference On Control*, 1984.
- [5] B. P. Gokulan and D. Srinivasan, “Distributed geometric fuzzy multi-agent urban traffic signal control,” *Intelligent Transportation Systems, IEEE Transactions on*, vol. 11, no. 3, pp. 714–727, 2010.
- [6] H. Ceylan and M. G. Bell, “Traffic signal timing optimisation based on genetic algorithm approach, including drivers’ routing,” *Transportation Research Part B: Methodological*, vol. 38, no. 4, pp. 329–342, 2004.
- [7] S. Darmoul, S. Elkosantini, A. Louati, and L. B. Said, “Multi-agent immune networks to control interrupted flow at signalized intersections,” *Transportation Research Part C: Emerging Technologies*, vol. 82, pp. 290–313, 2017.
- [8] R. S. Sutton and A. G. Barto, “Reinforcement learning: an introduction,” *Neural Networks, IEEE Transactions on*, vol. 9, no. 5, pp. 1054–1054, 1998.
- [9] C. Szepesvári, “Algorithms for reinforcement learning,” *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 4, no. 1, pp. 1–103, 2010.
- [10] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al., “Human-level control through deep reinforcement learning,” *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [11] C. J. Watkins and P. Dayan, “Q-learning,” *Machine learning*, vol. 8, no. 3-4, pp. 279–292, 1992.
- [12] R. J. Williams, “Simple statistical gradient-following algorithms for connectionist reinforcement learning,” in *Reinforcement Learning*, pp. 5–32, Springer, 1992.
- [13] V. Konda and J. Tsitsiklis, “Actor-critic algorithms,” in *NIPS*, vol. 13, pp. 1008–1014, 1999.
- [14] M. Aslani, M. S. Mesgari, and M. Wiering, “Adaptive traffic signal control with actor-critic methods in a real-world traffic network with different traffic disruption events,” *Transportation Research Part C: Emerging Technologies*, vol. 85, pp. 732–752, 2017.
- [15] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, “Asynchronous methods for deep reinforcement learning,” in *International Conference on Machine Learning*, pp. 1928–1937, 2016.
- [16] C. Guestrin, M. Lagoudakis, and R. Parr, “Coordinated reinforcement learning,” in *ICML*, vol. 2, pp. 227–234, 2002.
- [17] J. R. Kok and N. Vlassis, “Collaborative multiagent reinforcement learning by payoff propagation,” *The Journal of Machine Learning Research*, vol. 7, pp. 1789–1828, 2006.

- [18] M. Tan, “Multi-agent reinforcement learning: Independent vs. cooperative agents,” in *Proceedings of the tenth international conference on machine learning*, pp. 330–337, 1993.
- [19] J. Foerster, N. Nardelli, G. Farquhar, P. Torr, P. Kohli, S. Whiteson, et al., “Stabilising experience replay for deep multi-agent reinforcement learning,” *arXiv preprint arXiv:1702.08887*, 2017.
- [20] R. Bellman, “A markovian decision process,” tech. rep., DTIC Document, 1957.
- [21] R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour, “Policy gradient methods for reinforcement learning with function approximation,” in *Advances in neural information processing systems*, pp. 1057–1063, 2000.
- [22] C. Guestrin, D. Koller, and R. Parr, “Multiagent planning with factored mdps,” in *NIPS*, vol. 1, pp. 1523–1530, 2001.
- [23] G. Tesauro, “Extending q-learning to general adaptive multi-agent systems,” in *Advances in neural information processing systems*, pp. 871–878, 2004.
- [24] M. Wiering, J. Van Veenen, J. Vreeken, and A. Koopman, “Intelligent traffic light control,” *Institute of Information and Computing Sciences. Utrecht University*, 2004.
- [25] C. Cai, C. K. Wong, and B. G. Heydecker, “Adaptive traffic signal control using approximate dynamic programming,” *Transportation Research Part C: Emerging Technologies*, vol. 17, no. 5, pp. 456–474, 2009.
- [26] L. Prashanth and S. Bhatnagar, “Reinforcement learning with function approximation for traffic signal control,” *Intelligent Transportation Systems, IEEE Transactions on*, vol. 12, no. 2, pp. 412–421, 2011.
- [27] T. Chu and J. Wang, “Traffic signal control with macroscopic fundamental diagrams,” in *American Control Conference (ACC), 2015*, pp. 4380–4385, IEEE, 2015.
- [28] T. Chu, J. Wang, and J. Cao, “Kernel-based reinforcement learning for traffic signal control with adaptive feature selection,” in *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, pp. 1277–1282, IEEE, 2014.
- [29] S. Richter, D. Aberdeen, and J. Yu, “Natural actor-critic for road traffic optimisation,” in *Advances in neural information processing systems*, pp. 1169–1176, 2007.
- [30] T. Chu, S. Qu, and J. Wang, “Large-scale multi-agent reinforcement learning using image-based state representation,” in *Decision and Control (CDC), 2016 IEEE 55th Conference on*, pp. 7592–7597, IEEE, 2016.
- [31] N. Casas, “Deep deterministic policy gradient for urban traffic light control,” *arXiv preprint arXiv:1703.09035*, 2017.
- [32] W. Genders and S. Razavi, “Using a deep reinforcement learning agent for traffic signal control,” *arXiv preprint arXiv:1611.01142*, 2016.
- [33] W. Genders and S. Razavi, “Evaluating reinforcement learning state representations for adaptive traffic signal control,” *Procedia computer science*, vol. 130, pp. 26–33, 2018.
- [34] M. Wiering, “Multi-agent reinforcement learning for traffic light control,” in *ICML*, pp. 1151–1158, 2000.
- [35] T. Chu, S. Qu, and J. Wang, “Large-scale traffic grid signal control with regional reinforcement learning,” in *American Control Conference (ACC), 2016*, July 2016.
- [36] H. A. Aziz, F. Zhu, and S. V. Ukkusuri, “Learning-based traffic signal control algorithms with neighborhood information sharing: An application for sustainable mobility,” *Journal of Intelligent Transportation Systems*, vol. 22, no. 1, pp. 40–52, 2018.
- [37] E. Van der Pol and F. A. Oliehoek, “Coordinated deep reinforcement learners for traffic light control,” in *In proceedings of NIPS*, vol. 16, 2016.
- [38] S. El-Tantawy, B. Abdulhai, and H. Abdelgawad, “Multiagent reinforcement learning for integrated network of adaptive traffic signal controllers (MARLIN-ATSC): methodology and large-scale application on downtown Toronto,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 3, pp. 1140–1150, 2013.
- [39] F. Zhu, H. A. Aziz, X. Qian, and S. V. Ukkusuri, “A junction-tree based learning algorithm to optimize network wide traffic control: A coordinated multi-agent framework,” *Transportation Research Part C: Emerging Technologies*, vol. 58, pp. 487–501, 2015.
- [40] T. Chu and J. Wang, “Traffic signal control by distributed reinforcement learning with min-sum communication,” in *American Control Conference (ACC), 2017*, pp. 5095–5100, IEEE, 2017.
- [41] D. Krajzewicz, J. Erdmann, M. Behrisch, and L. Bieker, “Recent development and applications of SUMO - Simulation of Urban MObility,” *International Journal On Advances in Systems and Measurements*, vol. 5, pp. 128–138, December 2012.
- [42] S. Hochreiter and J. Schmidhuber, “Long short-term memory,” *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [43] A. M. Saxe, J. L. McClelland, and S. Ganguli, “Exact solutions to the nonlinear dynamics of learning in deep linear neural networks,” *arXiv preprint arXiv:1312.6120*, 2013.
- [44] L. Codeca and J. Härrä, “Monaco SUMO Traffic (MoST) Scenario: A 3D Mobility Scenario for Cooperative ITS,” in *SUMO 2018, SUMO User Conference, Simulating Autonomous and Intermodal Transport Systems, May 14-16, 2018, Berlin, Germany*, (Berlin, GERMANY), 05 2018.
- [45] Z. Wang, V. Bapst, N. Heess, V. Mnih, R. Munos, K. Kavukcuoglu, and N. de Freitas, “Sample efficient actor-critic with experience replay,” *arXiv preprint arXiv:1611.01224*, 2016.




ing, and other engineering control systems.



and innovation, enterprise IT infrastructure management, smart manufacturing, smart infrastructures and smart city, and environmental informatics.



collaborates with the SUMO developers at DLR (German Aerospace Centre).



Systems, and Robotics and Automated Vehicles. Dr. Li was a recipient of the National Scholarship from China.

Tianshu Chu received the B.S. degree in physics from the Waseda University, Tokyo, Japan, in 2010, and the M.S. and Ph.D. degrees from the Department of Civil and Environmental Engineering, Stanford University, in 2012 and 2016. He is currently a data scientist at Uhana Inc. and an adjunct professor at the Stanford Center for Sustainable Development & Global Competitiveness. His research interests include reinforcement learning, deep learning, multi-agent learning, and their applications to traffic signal control, wireless network control, autonomous driving,

Jie Wang received his B.S. degree from Shanghai JiaoTong University, two M.S. degrees from Stanford University and University of Miami, and the Ph.D. degree in civil and environmental engineering from Stanford University, in 2003. He is currently an adjunct professor with the Department of Civil and Environmental Engineering, Stanford University, and the executive director of the Stanford Center for Sustainable Development & Global Competitiveness. His research interests include information and knowledge management for sustainable development and innovation, enterprise IT infrastructure management, smart manufacturing, smart infrastructures and smart city, and environmental informatics.

Lara Codeca received a Ph.D. Degree from the University of Luxembourg in 2016 and her Computer Sciences Master degree at the University of Bologna (Italy) in 2011. In 2011, she was a visiting fellow at Prof. Dr. Mario Gerla’s Vehicular Lab at the University of California, Los Angeles (UCLA). She is currently a post-doctoral fellow in the CATS group in EURECOM (France). Her research interests include (Cooperative) Intelligent Transportation Systems, Vehicular Traffic Modelling, and Big-data Analysis. She is active in the SUMO community and

Zhaojian Li is an Assistant Professor in the Department of Mechanical Engineering at Michigan State University. He obtained M.S. (2013) and Ph.D. (2015) in Aerospace Engineering (flight dynamics and control) at the University of Michigan, Ann Arbor. As an undergraduate, Dr. Li studied at Nanjing University of Aeronautics and Astronautics, Department of Civil Aviation, in China. Dr. Li worked as an algorithm engineer at General Motors from January 2016 to July 2017. His research interests include Learning-based Control, Nonlinear and Complex