

SocialLight: Distributed Cooperation Learning towards Network-Wide Traffic Signal Control

Harsh Goel
University of Pennsylvania
Philadelphia, United States
harshg99@seas.upenn.edu

Yifeng Zhang
National University of Singapore
Singapore, Singapore
yifeng@u.nus.edu

Mehul Damani
Massachusetts Institute of Technology
Cambridge, United States
damanimehul24@gmail.com

Guillaume Sartoretti
National University of Singapore
Singapore, Singapore
guillaume.sartoretti@nus.edu.sg

ABSTRACT

Many recent works have turned to multi-agent reinforcement learning (MARL) for adaptive traffic signal control to optimize the travel time of vehicles over large urban networks. However, achieving effective and scalable cooperation among junctions (agents) remains an open challenge, as existing methods often rely on extensive, non-generalizable reward shaping or on non-scalable centralized learning. To address these problems, we propose a new MARL method for traffic signal control, SocialLight, which learns cooperative traffic control policies by distributedly estimating the individual marginal contribution of agents on their local neighborhood. SocialLight relies on the Asynchronous Actor Critic (A3C) framework, and makes learning scalable by learning a locally-centralized critic conditioned over the states and actions of neighboring agents, used by agents to estimate individual contributions by counterfactual reasoning. We further introduce important modifications to the advantage calculation that help stabilize policy updates. These modifications decouple the impact of the neighbors' actions on the computed advantages, thereby reducing the variance in the gradient updates. We benchmark our trained network against state-of-the-art traffic signal control methods on standard benchmarks in two traffic simulators, SUMO and CityFlow. Our results show that SocialLight exhibits improved scalability to larger road networks and better performance across usual traffic metrics.

KEYWORDS

Adaptive Traffic Signal Control, Autonomous Signal Control; Multi Agent Reinforcement Learning

ACM Reference Format:

Harsh Goel, Yifeng Zhang, Mehul Damani, and Guillaume Sartoretti. 2023. SocialLight: Distributed Cooperation Learning towards Network-Wide Traffic Signal Control. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 9 pages.

1 INTRODUCTION

In recent years, most cities around the world have seen growing traffic levels, and associated traffic congestion have started to show a

number of negative effects, both at the micro and macro level. At the micro level, passengers experience frustration due to delays and also face increased risks of collisions. At the macro level, unproductive time spent in traffic damages economic health, while wasted fuel and traffic jams increase air and noise pollution. As a result, there is a growing need for effective traffic signal control methods, which can play a significant role in alleviating traffic congestion.

Current traffic signal control methods can be broadly classified into two categories: fixed-time control and adaptive control. In fixed-time control, the duration of different traffic light phases are pre-determined, often optimized offline from historical data. However, urban traffic, on top of having considerable stochasticity, also shows significant temporal and spatial variations. For example, higher congestion is often seen due to temporal peaks at the end of a work day. Similarly, the spatial structure of traffic networks often gives rise to tidal patterns, with high congestion in particular lane directions. To account for this variability, adaptive traffic signal control (ATSC) methods aims at dynamically adjusting traffic signal phases online, based on current traffic conditions.

Multi-Agent Reinforcement learning (MARL) is one such adaptive and versatile data-driven method, which has recently shown great promise in ATSC and general control tasks [6, 10, 17]. ATSC is cast as a MARL problem in which each agent controls a single traffic intersection, based on locally-sensed real-time traffic conditions and communication with neighboring intersections. Thus, each agent learns a policy which maps the current traffic conditions at the intersection into control outputs (e.g., phase selection, phase duration). This lends it an advantage over conventional ATSC methods, which rely on complex dynamics models and heuristic assumptions. An alternative to MARL is to train a single centralized RL agent, which is responsible for controlling all traffic intersections. However, while centralization allows for direct maximization of a global reward/objective such as average trip time, training such a centralized method is infeasible in practice due to the exponentially growing joint action space, and the high latency associated with information centralization.

Although the MARL formulation of ATSC alleviates most issues associated with centralized methods, it introduces new challenges as the performance of control policies that optimize local objectives for each agent (intersection) will not be equivalent to that of a centralized global RL agent if the local objectives aren't well-aligned with the global (team-/network-level) one. Since, traffic

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

networks have complex spatio-temporal patterns and significant interdependence between agents, greedily optimizing each agent’s local reward usually does not optimize global (network-level) objectives. A possible solution to this is to directly sum each agent’s local reward with that of neighboring agents into a large *neighborhood reward*, which becomes a new, more global objective optimized by each agent. The idea here is to couple neighboring agents via their rewards, whereby improving their neighbors’ local rewards via their own actions directly affect their own long-term return. However, such a neighborhood reward has high variance since it is now conditioned on the actions of multiple agents, making it difficult for an agent to determine its true *marginal* contribution. A recent work introduced COMA [10], which learns a complex team-level network allowing agents to estimate their own marginal contribution to the team reward via counterfactual reasoning. Specifically, COMA’s centralized value function network uses both the global team reward, as well as the states and actions of all agents as input, which becomes exponentially harder to train in larger teams.

In this work, we propose to spatially distribute the global credit assignment problem into a collection of local marginal contributions calculations, as a natural means to balance the tradeoff between scalability and cooperative performance. To this end, we learn a shared value network, similar to COMA’s but only conditioned on each agent’s states/actions and that of its direct neighbors, which can be used for agents to estimate their own marginal contribution to the local neighborhood reward. By relying on a fixed number of neighbors, our locally-centralized value network allows for significantly improved scalability, while minimally affecting the quality of the learned solutions by leveraging the natural fixed structure of ATSC, where the natural flow of traffic means that neighboring agents must be more tightly coupled. We further introduce important modifications to the advantage calculation that helps improve and stabilize policy updates.

We present results of an extensive set of simulations conducted on a range of benchmark traffic networks using two standard traffic simulators, SUMO [14] and CityFlow [26]. We show that our framework - SocialLight - results in improved cooperation and natural scalability to larger networks compared to existing state-of-the-art ATSC baselines. To the best of our knowledge, we are also the first work to show effective performance on both these standard traffic simulators¹. Finally, through a series of ablation studies, we also show that the modified advantages in combination with the counterfactual baseline derived from COMA help improve the speed and stability of training in comparison to vanilla A3C/COMA.

2 RELATED WORK

2.1 Conventional Traffic Signal Control

Traffic signal control is a versatile problem with many possible objectives to optimize and different scopes of optimization. Conventional methods can be broadly categorized into adaptive or fixed-time control based on the ability of the method to adapt to current traffic conditions. They are also categorized based on the scope of their optimization - some methods only consider optimization over a single isolated traffic intersection while others consider

a network of traffic intersections (multiple intersections). Here, we briefly list out seminal works for each of these categories -

- **Single intersection, Fixed time:** The Webster method [13] obtains a closed-form solution for the optimal cycle length and phase split based on a set of modeling assumptions.
- **Single intersection, Adaptive:** SCATS [18] is a popular adaptive-control method, which has even been deployed in numerous urban cities around the world. It takes in pre-defined signal plans and iteratively selects from these traffic signals according to a defined performance measure.
- **Multiple intersections, Fixed time:** GreenWave [24] optimizes the timing offsets between different intersections to minimize the number of stops for vehicles traveling along a specific direction.
- **Multiple intersections, Adaptive:** Max-pressure control [28] addresses the risk of oversaturation at an intersection by balancing queue lengths between neighboring intersections.

This list is not exhaustive and for more details, we refer the reader to the recent survey by Wei et al. [31]. While conventional traffic control methods are currently the standard for real-world deployments, they rely on accurate traffic models.

2.2 RL-based Traffic Signal Control

Model-free RL is particularly suitable for ATSC due to its ability to learn from and find structure in large amounts of raw data. Early works in RL explored different ATSC problem formulations on simplified traffic environments. Out of these variants, the most common variant has been learning to select the next traffic light phase using a set of features describing the local traffic conditions. [11, 23, 33]. In contrast, Li et al. [15], Aslani et al. [1] and Casas et al. [3] focused on learning policies for selecting the traffic signal timing (also known as the phase duration). While most methods focused on learning policies from a low-dimensional feature space, Mousavi et al. [20] used a CNN to directly map from image snapshots (obtained using a simulator) to the policy for selecting the next traffic phase. Similarly, Wei et al. [32] used both extracted features and real-world images to learn the policy. Taking note of the fact that traffic intersections in the real world vary greatly, Oroojlooy et al. [22] proposed AttendLight, which uses two attention networks to learn a universal model applicable to intersections with any number of roads, lanes, phases (possible signals), and traffic flow.

Multiple works described above showed that deep RL agents can effectively control individual intersections. Hence, the focus of more recent works has shifted to developing methods for a network of traffic intersections that more closely resemble real-world traffic systems where different traffic intersections are highly interconnected. From lessons learned through conventional methods such as Greenwave [24], it is evident that coordination between these intersections is necessary to achieve effective performance for the network. Current coordination methods for multi-agent traffic signal control can be broadly grouped into two - joint action learners and independent learners. Joint action learners use a single global agent to control the traffic for all intersections [23, 34]. While joint action learners allow for direct optimization of a global objective, they find it difficult to scale beyond a few intersections. On the other hand, independent learners train an individual policy for each

¹To help the community standardize bench-marking on both simulators, our open-source code can be found at <https://github.com/marmotlab/SocialLight>

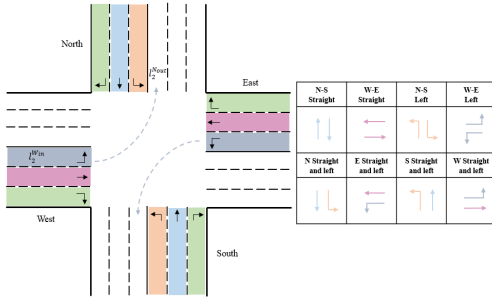


Figure 1: Single intersection with 8 traffic light phases

intersection, while considering other agents/intersections to be a part of the environment [6, 21, 22, 29, 30]. Wei et al. [29] proposed PressLight, which extends *max pressure* [28] to multi-agent RL by rewarding each agent for minimizing the pressure at an intersection. Scaling this, Chen et al [5] proposed MPLight, a deep MARL framework which uses parameter sharing to train policies using pressure-based objectives for large-scale networks. A fundamental challenge with achieving cooperation with independent learners is partial observability, as individual traffic intersections are unable to observe nearby intersections. To address this, Chu et al. [6] proposed MA2C, a MARL algorithm that adds fingerprints of its neighbors to each agent’s observation for improved observability, and a spatial discount factor to reduce learning difficulty. With a similar motivation, Nishi et al. [21] used a graph convolutional neural network (GCNN) to automatically extract traffic features between distant intersections. As an alternative to direct state augmentations or feature extractions, Wei et al. [30] proposed to learn a communication mechanism between different intersections. Their framework, referred to as Colight, uses graph attentional networks (GAT’s) to facilitate communication between intersections. More recently, Zhang et al. [37] further showed performance improvements by learning phase correlation with an attention mechanism over a queue length based state representation.

Finally, making use of the Centralized Training Decentralized Execution (CTDE) paradigm, some works have focused on learning a centralized critic to guide independent policy learning [4, 27].

3 BACKGROUND

3.1 Multi-Agent Reinforcement Learning

Multi-agent Reinforcement Learning generally optimize a global objective over a cooperative game involving numerous agents. Formally, a MARL problem can be formulated as a set of Decentralized Partially Observable Markov Decision Processes (Dec-POMDPs) [12] characterized by a tuple $G = (\psi, S, A, P, R, \rho, O, Z, \gamma)$, where ψ is a finite set of all agents ($|\psi| = n$, S is the state space, A the joint action space, defined as $A = A_1 \times A_2 \times \dots \times A_n$ and $P : S \times A \times S \rightarrow [0, 1]$ denotes the global transition dynamics. The reward function $R : S \times A \times S \rightarrow \mathbb{R}^n$ computes a set of private rewards $[r_i^t]$ for each agent i at each time-step t . These rewards can be global (i.e each agent receives the same global reward) or local via reward shaping. In partially observable settings each agent cannot access the true global state of the environment s_t . Instead, it draws an observation via observation models $O = [O_1, O_2, \dots, O_N]$ where $O_i : S \rightarrow Z_i$. Here $Z = [Z_1, Z_2, \dots, Z_n]$ are the agents’ observation spaces.

Let $\Pi^i : Z_i \times A_i \rightarrow [0, 1]$ denote the stochastic policy for agent i , then the joint policy of the multi-agent system is given by $\pi(a^t | s^t) = \prod_{i \in \psi} \pi_{\theta}^i(a_i^t | z_i^t)$ assuming the policy of each agent is parameterized by θ_i . The Multi-Agent RL objective thereby is to find an optimal joint policy π that formally maximizes the discounted returns over all agents $J(\pi) = E_{\tau \sim \pi} [\sum_{t=0}^{\infty} \sum_{i=0}^N r_i^t]$. Here, τ denotes the global trajectory $(s^0, a^0, s^1, a^1, \dots, s^t, a^t)$. Finally, ρ and γ represents the initial state distribution and the discount factor respectively.

This objective can be optimized over in a centralized manner by parameterizing the global policy $\pi(a^t | s^t)$. However, such a centralized approach usually scales poorly, given the exponentially growing state-action space of the agents. Independent learning algorithms [19, 25, 35] have been effective in many multi-agent settings to optimize $\pi_{\theta}^i(a_i^t | z_i^t)$ over the local returns $R^i(\tau) = \sum_{t=0}^{\infty} r_i^t$ with either a local observation value critic $V_i^{\pi_i}(z_i) = E_{\tau} [R^i(\tau) | z_i^0 = z_i]$ or $Q_i^{\pi_i}(z_i, a_i) = E_{\tau} [R^i(\tau) | z_i^0 = z_i, a_i^0 = a_i]$ to estimate local advantages $A_i^{\pi_i}(z_i, a_i) = Q_i^{\pi_i}(z_i, a_i) - V_i^{\pi_i}(z_i)$ for policy improvement. For cooperation, multi-agent actor-critic methods [9, 36] propose to learn centralized critics $Q^{\pi}(s, a)$ that optimize global returns $R(\tau) = E_{\tau} [\sum_{t=0}^{\infty} \sum_{i=0}^N r_i^t]$ with individual policies for decentralized execution in their environments.

3.2 Traffic Terminology

Definition 1 (Traffic movement): One way by which vehicles can traverse the intersection, i.e., from one incoming lane to one connected outgoing lane. The traffic movement m_{ij} between lane i and outgoing lane j is denoted as (l_i^{in}, l_j^{out}) , and the activation of the movement is defined as $m_{ij} = 1$.

Definition 2 (Traffic signal phase): A set of simultaneously allowed traffic movements, allowing only vehicles under these activated traffic movements to traverse the intersection. We denote the signal phase as $p = \{m_{ij} | m_{ij} = 1\}$, where $i \in L_{in}$ and $j \in L_{out}$.

Definition 3 (Traffic Agent and traffic network): A traffic agent is in charge of one intersection and relies on the real-time traffic conditions within its own area to control the signal phases. A traffic network is a multi-agent network $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where the vertices \mathcal{V} are traffic agents and the edges \mathcal{E} define the road network connecting them. An agent i has an immediate neighbor j if $E_{ij} = 1$. In practice, this means that agents i and j are directly connected.

Fig. 1 depicts an example single intersection which is composed of twelve incoming lanes and twelve outgoing lanes. Considering a single connection between incoming and outgoing lanes, i.e., each incoming lane is only connected to one outgoing lane, there is a total of twelve movements (left-turn, go-straight, and right-turn for each direction). Therefore, we can define eight phases, shown in the right side of the Fig. 1. Currently, the W-E left-turn phase is activated at the intersection, allowing vehicles at the left-turn lanes of the west and east directions to move.

3.3 Traffic as a MARL Problem

3.3.1 Problem Definition. Given the current traffic conditions, the goal of traffic agents in the network is to select their own optimal signal phase a_t for a fixed phase duration, until the next decision time-step $t + 1$, to maximize a global cumulative objective.

3.3.2 Observations. The true global state for a traffic systems comprises the position and current travel time of each vehicle in the

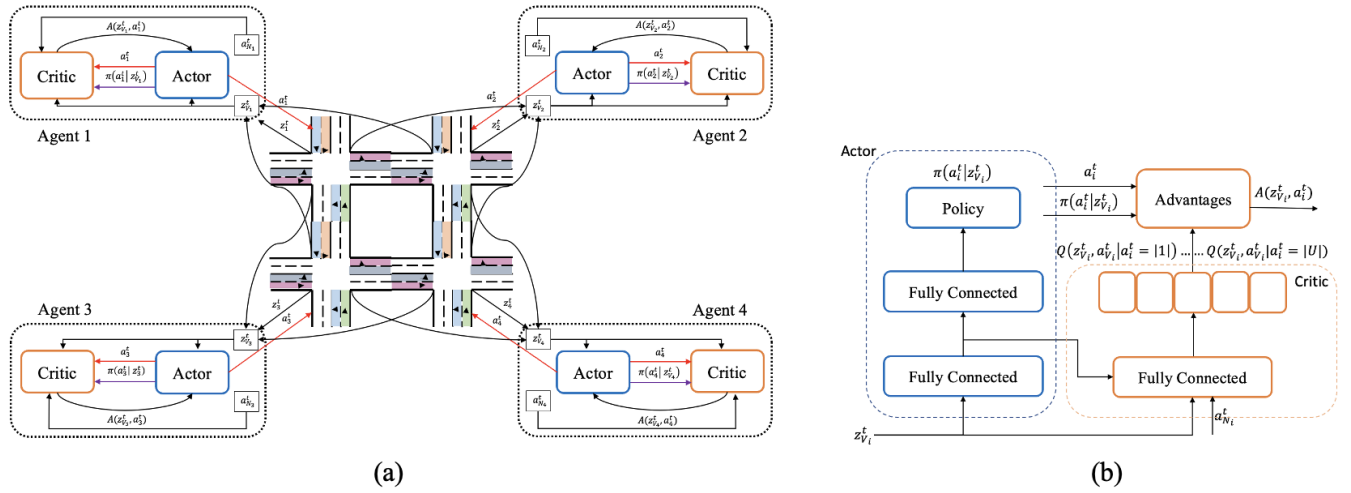


Figure 2: Illustration of SocialLight: Fig a) highlights the distributed training framework where each agent maintains an actor network and a local critic network with parameter sharing. The local critic is conditioned on the augmented observations and the actions of neighboring agents to marginalize individual credit via a counterfactual baseline. This baseline is then used to compute the individual advantages used for individual policy improvement. Fig b) Details the actor and critic network architectures. (Image helped prepared by Stuti Mittal)

network, as well as the current traffic phase of each agent. However, obtaining this information is infeasible in practice. Instead, each agent is only allowed to observe a portion of this full state, which contains the incoming queue lengths, current local traffic phase, average waiting time of queued vehicles, and pressure, which can all be locally measured via Induction Loop Detectors commonly found in modern traffic networks. In this work, aligned with recent work in the community, we use a combination of these features to define each agent’s local state.

3.3.3 Actions. We let agents directly select one of the 8 traffic phases, and execute that phase for a pre-specified duration (i.e., there is not fixed cycle among phases), for maximal adaptability. Note other works have also considered switching phases in a fixed cycle without specified phase duration [8], or setting phase durations within a fixed cycle length[2].

3.3.4 Rewards. The global objective for traffic control is to minimize the cumulative trip time of all vehicles. However, optimizing over total trip time is hard, since vehicles usually accumulate local delays as they pass through multiple junctions in their journey. Thus, different local reward structures are often used instead, such as cumulative delay, queue length, pressure, or waiting time, which align well enough with the global trip time objective. In this work, we specifically opt to use queue lengths, to implicitly maximize throughput at each intersection. This has been recently shown to be superior compared to other local reward formulations [37].

4 SOCIALLIGHT

SocialLight introduces a new learning mechanism where individual traffic agents learn to cooperate by marginalizing out their true contributions to a neighborhood reward within an independent learning framework. We outline the components proposed in SocialLight to train traffic light control policies for a junction. We first provide an intuition that motivates SocialLight. Then we present the adaptations within the asynchronous actor-critic framework

that are inspired by COMA [9] to address these challenges. Finally we introduce modifications to the vanilla COMA formulation of the advantages to enhance training stability and improve convergence.

4.1 Notation

Given a traffic network $\mathcal{G}(\mathcal{V}, \mathcal{E})$, the local neighborhood for agent i is denoted as $\mathcal{V}_i = i \cup \mathcal{N}_i$, where \mathcal{N}_i are the agents in the immediate neighborhood of agent i . Each agent learns a policy network $\pi_\theta(a_i^t | z_{\mathcal{V}_i}^t)$ and a critic network $Q_\phi(z_{\mathcal{V}_i}^t, a_{\mathcal{V}_i}^t)$, where $z_{\mathcal{V}_i}^t$ is the augmented observation comprising the observation of each agent and its neighbors, i.e., $z_{\mathcal{V}_i}^t = [z_i^t] \cup_{j \in \mathcal{N}_i} [z_j^t]$. The $a_{\mathcal{V}_i}^t$ denotes the joint action of the agent and its neighbors i.e. $a_{\mathcal{V}_i}^t = [a_i^t] \cup_{j \in \mathcal{N}_i} [a_j^t]$.

An agent i receives an individual reward r_i^t which can be any reward function computed using local traffic conditions such as queue length over its incoming lanes or local max pressure. Through reward sharing, an individual agent would sum up its own reward with those received by its neighbors. The neighborhood reward for the agent i is defined as $r_{\mathcal{V}_i}^t = \sum_{j \in \mathcal{V}_i} r_j^t$.

4.2 Changes to the Policy Gradient

We leverage the locally centralized critic to compute advantages by marginalizing individual contributions via a counterfactual baseline that is inspired by COMA for policy improvement. COMA learns a central critic $Q(s^t, a^t)$ over the global state s^t and global action $a^t = \cup_{i \in \mathcal{N}} [a_i^t]$ for all N agents. The counterfactual baseline and COMA advantages for a global state is given by $Q(s^t, a^t) - E_{a_i^t} Q(s^t, a^t)$.

In contrast our method computes a local critic $Q_\phi(z_{\mathcal{V}_i}^t, a_{\mathcal{V}_i}^t)$ for each agent i.e. conditioned on the joint observation space and action space of the neighborhood \mathcal{V}_i to marginalize individual contribution over the neighborhood reward. Hence in our setting, a naive COMA update is given by

$$\hat{A}(z_{\mathcal{V}_i}^t, a_i^t) = Q_\phi(z_{\mathcal{V}_i}^t, a_{\mathcal{V}_i}^t) - \sum_{a_i^t} \pi_\theta(a_i^t | s_{\mathcal{V}_i}^t) Q_\phi(s_{\mathcal{V}_i}^t, a_{\mathcal{V}_i}^t), \quad (1)$$

where the second term is a counterfactual baseline that marginalizes an agent's expected contributions by fixing the actions of neighboring agents. However, we observed that using original COMA advantages made learning unstable especially with multiple agents updating shared parameters. While the counterfactual baseline's expected contribution to the gradient is zero, we suspect that the instability primarily arises from the high bias in the critic's gradients during the first few epochs of training. To reduce this bias, we introduce two modifications. First, we take inspiration from Temporal Differences (TD), which reduces bias by estimating advantages from rolling out the trajectory and bootstrapping the critic estimates at a future state. However, unbiasing advantages with the estimated return from the full trajectory roll-out comes at the cost of larger variance in the policy gradient. This bias-variance trade-off problem is then further addressed via standard Generalized Advantage Estimation (GAE) over the modified COMA advantages.

TD advantages do not require training an additional network. Hence, we propose a similar modification to the COMA advantages to resemble TD advantages as follows:

$$\begin{aligned} \hat{A}^1(z_{\mathcal{V}_i}^t, a_{t,i}) &= r_{\mathcal{V}_i}^t + \gamma \sum_{a_i^{t+1}} \pi_{\theta}(a_i^{t+1} | z_{\mathcal{V}_i}^{t+1}) Q_{\phi}(z_{\mathcal{V}_i}^{t+1}, a_{\mathcal{V}_i}^{t+1}) \\ &\quad - \sum_{a_i^t} \pi_{\theta}(a_i^t | z_{\mathcal{V}_i}^t) Q_{\phi}(z_{\mathcal{V}_i}^t, a_{\mathcal{V}_i}^t) \end{aligned} \quad (2)$$

Note that our key distinction lies in the estimate of the bootstrapped value which we do by evaluating the counterfactual baseline at the future state. In contrast, TD(1) advantages in a multi-agent setting bootstraps the discounted return over the expectation of the joint action over the policies of all agents $j \in \mathcal{V}_i$ that is given by $E_{\prod_{j \in \mathcal{V}_i} \pi(a_j^{t+1} | z_{\mathcal{V}_i}^{t+1})} Q(z_{\mathcal{V}_i}^{t+1}, a_{\mathcal{V}_i}^{t+1})$. However, this fails to capture the dependence of future returns on the future actions of neighboring agents. Hence, setting the bootstrapped return as the counterfactual baseline at the future state reduces the variability of future returns on the actions of future neighboring agents, thereby reducing the variance in the gradient updates to the policy network.

The above advantages via TD errors still uses a biased critic for the one-step lookahead during the first few epochs. Inspired by the success of GAE, we also use a GAE-type computation to trade off the bias and variance to the policy gradients, for improved learning stability. The generalized advantages are given as:

$$A^{GAE}(z_{\mathcal{V}_i}^t, a_i^t) = \sum_{l=0}^{\infty} (\gamma \delta)^l \hat{A}^1(z_{\mathcal{V}_i}^{t+l}, a_i^{t+l}) \quad (3)$$

Note that the Generalised Advantage estimate implicitly weights the n step TD advantages by a factor δ^n as follows:

$$\begin{aligned} \hat{A}^n(z_{\mathcal{V}_i}^t, a_i^t) &= \sum_{l=0}^{n-1} \gamma^l r_{\mathcal{V}_i}^{t+l} + \gamma^n \sum_{a_i^{t+n}} \pi_{\theta}(a_i^{t+n} | z_{\mathcal{V}_i}^{t+n}) Q_{\phi}(z_{\mathcal{V}_i}^{t+n}, a_{\mathcal{V}_i}^{t+n}) \\ &\quad - \sum_{a_i^t} \pi_{\theta}(a_i^t | z_{\mathcal{V}_i}^t) Q_{\phi}(z_{\mathcal{V}_i}^t, a_{\mathcal{V}_i}^t) \end{aligned} \quad (4)$$

The discounting factor δ regulates the bias variance tradeoff where the variance in the gradient estimator increases with the time horizon due to the influence of the returns on the actions of the team in the neighborhood. The policy gradient for a rollout of length T is thereby given as:

$$\nabla L_{\pi}(\theta) = \sum_{t=0}^T \nabla_{\theta} \log(\pi_{\theta}(a_i^t | z_{\mathcal{V}_i}^t)) A^{GAE}(z_{\mathcal{V}_i}^t, a_i^t) \quad (5)$$

4.3 Critic Training

The critic introduced in the above section estimates returns over a joint action space of the agents in the neighborhood \mathcal{N}_i . However, having the critic output $|\mathcal{A}|^n$ values, where $|\mathcal{A}|$ represents the size of the action space of one agent, is impractical. We address this problem by using a critic representation similar to COMA which allows for an efficient evaluation of the baseline. In this work, the critic is a neural network that takes local observations of the neighborhood $z_{\mathcal{V}_i}^t$ and the actions of other agents $a_{\mathcal{V}_i}^t \setminus \{a_i^t\}$ and outputs the a vector of length $|\mathcal{A}|$. Note that the actions of other agents are one hot encoded; the network is depicted in Fig. 2. In doing so, the advantage term can be computed in a single pass through a dot product between the outputs of the actor and critic networks.

To conform with the TD error used to compute advantages via the counterfactual baseline computed at the joint future agent state, the targets of the critic are modified similarly. Here the TD(1) error is given by

$$G_i^t = r_{\mathcal{V}_i}^t + \gamma \sum_{a_i^{t+1}} \pi_{\theta}(a_i^{t+1} | z_{\mathcal{V}_i}^{t+1}) Q_{\phi}(z_{\mathcal{V}_i}^{t+1}, a_{\mathcal{V}_i}^{t+1}) \quad (6)$$

More generally, the n step TD returns is formulated as

$$G_i^{t:t+n} = \sum_{l=0}^{n-1} \gamma^l r_{\mathcal{V}_i}^{t+l} + \gamma^n \sum_{a_i^{t+n}} \pi_{\theta}(a_i^{t+n} | z_{\mathcal{V}_i}^{t+n}) Q_{\phi}(z_{\mathcal{V}_i}^{t+n}, a_{\mathcal{V}_i}^{t+n}) \quad (7)$$

The TD(λ) targets $G_{t,i}^{\lambda}$ can be computed via these modified n step returns and the critic network is regressed against these targets over a rollout of length T . The critic loss reads

$$L_Q(\phi) = \frac{1}{T} \sum_{t=0}^T (Q_{\phi}(z_{t+n, \mathcal{V}_i}, a_{t+n, \mathcal{V}_i}) - G_{t,i}^{\lambda})^2 \quad (8)$$

5 EXPERIMENTS

Our experiments aim to answer the following questions:

- (1) Does SocialLight improve over the current state of the art, when keeping standard definitions of the state space and reward?
- (2) What is the impact of the proposed individual contribution marginalization mechanism in comparison to standard reward sharing?
- (3) Do our modified advantages improve the stability of training, compared to simply applying COMA advantages locally?

To answer these questions, our experiments are conducted in both the SUMO and CityFlow traffic simulators [14, 26]. We first benchmark the performance of SocialLight on common baselines developed over both the synthetic SUMO and real world Cityflow traffic datasets and measure the same standard traffic metrics in both datasets.

We then perform an ablation to study the impact of individual contribution marginalization on the traffic performance over artificially generated traffic flows on a Manhattan road-map over

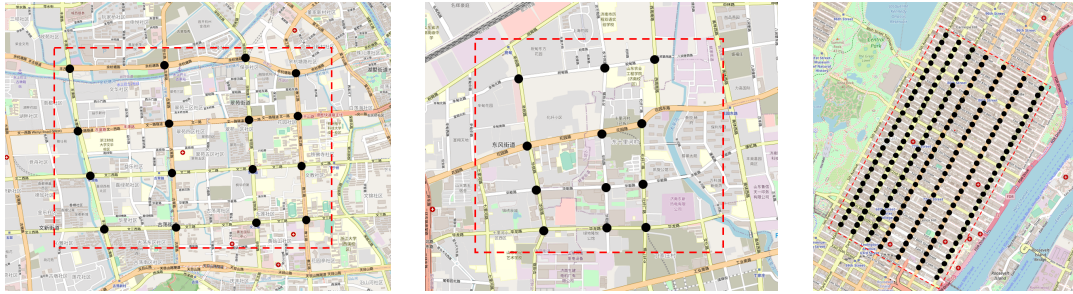


Figure 3: Real-world traffic road network for CityFlow dataset (From left to right, Hangzhou map, Jinan map and New York map)

Table 1: SocialLight vs Baselines on CityFlow

Methods	New York		Hangzhou		Jinan		
	1	2	1	2	1	2	3
Fixed Time	1397.37	1660.29	432.316	359.44	364.36	289.74	316.69
Max Pressure	1177.75	1535.77	262.35	348.68	275.79	223.06	234.93
Co-Light	1221.77	1476.18	271.07	297.26	276.33	237.14	278.16
MP - Light	1168.49	1597.24	343.47	282.14	300.93	259.10	261.45
Attention-Light	978.62	1571.68	259.62	284.75	268.68	212.76	216.34
SocialLight	760.94	1114.47	255.68	301.16	226.21	209.81	205.38
Performance Gain	22.23 %	24.53 %	1.54 %	-	15.68%	1.41 %	5.09 %

the SUMO simulator. By simulating a diverse range of traffic scenarios to train and test on, we can eliminate networks potentially over-fitting to a single scenario that may misrepresent our analysis. Moreover, we show the improvement over the learning stability of our network over both synthetic traffic and a real traffic datasets.

5.1 Description of Traffic Datasets

We conduct experiments on two different microscopic traffic simulators SUMO [16] and CityFlow [26] with synthetic and real-world datasets respectively. A traffic simulation on either simulator comprises a road network and a traffic flow dataset. Here the road network defines the positions of the intersections, the attributes of the roads (e.g., number and length of lanes, speed limits, and lane connections etc.) and the phase settings. The traffic flow datasets define the travel information of all vehicles, characterized by the origin destination (O-D) pair and time that the vehicle enters the network. Note that our simulations are conducted over homogeneous intersections that have the same settings for the roads and traffic light phases.

5.1.1 Synthetic Traffic Datasets. We use synthetic traffic dataset with a Manhattan road network on the SUMO simulator that is adapted from the benchmark method MA2C [7]. The road network is a 5×5 traffic grid network with 25 intersections. Each intersection is formed by two-lane streets (W-E) with speed limit 72 km/h and one-lane avenues with speed limit 40km/h.

The traffic flow datasets are artificially generated during run time with different fixed seeds, to fairly compare among algorithms. Aligned with previous work [7], all episodes consider a fixed traffic

flow composed of travel information (O-D pair) for all vehicles, while the seed sets the (random) initial position and speed of these vehicles.

5.1.2 Real Traffic Datasets. The real traffic dataset considers three city networks - Jinan, Hangzhou and New York. These datasets serve as popular benchmarks for ATSC [30, 31, 38]. The road networks are extracted from a portion of real-world traffic map for simulations, and traffic flow datasets were compiled from cameras during different time periods. As shown in Figure 3, there are a total of 12 (3×4) intersections in the Jinan map, 16 (4×4) intersections in the Hangzhou map, and 192 (28×7) intersections in the New York map. The traffic flow datasets used in these different maps are described in [37]; there are three different flow datasets for Jinan, two for Hangzhou, and two for New York.

5.2 Traffic Performance on SUMO Synthetic Traffic Datasets

We compare our method SocialLight with current state-of-the-art baselines optimized on the various traffic datasets. Following the baselines [7] for a fair comparison, we keep the same experiment settings, as well as similar POMDP settings in terms of the definitions of actions and rewards. The state is modified to include current traffic phase.

5.2.1 Baseline methods.

- (1) **Greedy [13]:** Greedily chooses the phase associated with lanes with maximum incoming queue length.

Table 2: SocialLight vs Baselines on SUMO - Manhattan

Metrics(Average)	Greedy	IA2C	MA2C	IQL-LR	A3C	A3C (nr)	SocialLight
Queue Length	5.00 (2.86)	3.27 (2.00)	2.24 (1.28)	3.76(2.74)	1.71 (1.14)	1.11 (0.89)	0.74 (0.69)
Speed (m/s)	1.56 (1.33)	1.70 (1.32)	2.31 (1.22)	3.04(3.24)	4.55 (2.78)	4.94 (2.34)	5.36 (2.66)
Intersection Delay (s)	60.97 (47.15)	58.27 (46.08)	21.96 (19.83)	92.59(94.69)	38.23 (32.83)	25.43 (22.53)	10.07 (9.02)
Cumulative Delay (s)	595.78 (464.23)	446.85 (410.71)	325.02 (269.77)	214.521(342.95)	184.60 (297.99)	159.23 (226.64)	106.51 (165.18)
Trip Time (s)	885.47 (572.63)	704.62 (501.33)	597.88 (399.09)	462.862(453.69)	395.04 (387.06)	386.01 (305.74)	309.81 (243.95)

- (2) **IQL-LR** [7]: A linear regression based independent Q-learning (IQL) algorithm, where each local agent learns its own policy independently by considering other agents as part of the environment’s dynamics.
- (3) **IA2C** [7]: An extension of IQL-LR which relies on advantage actor-critic (A2C) algorithm instead of IQL.
- (4) **MA2C** [7]: A cooperative MARL algorithm, which includes the observations and fingerprints of neighboring agents in each agent’s state. to alleviate the instability caused by partial observability. MA2C further introduces a spatial discount factor to scale down the observation and rewards signals of the neighboring agents, to encourage agents towards neighborhood-level cooperation.
- (5) **A3C**: The distributed learning framework with parameter sharing relying on A3C algorithm, where each agent learns to maximize its individual objective.
- (6) **A3C(nr)**: The distributed learning framework where each agent is to maximize the neighborhood reward rather than individual reward.

5.2.2 Analysis. We first observe that SocialLight outperforms the heuristic based Greedy baseline, as well as the Deep RL baselines MA2C, IA2C, IQL-LR and even A3C with and without neighborhood reward over all traffic metrics (average queue length, speed, intersection delay, cumulative delay, and average total trip time). Prior methods IA2C, MA2C and IQL-LR are outperformed significantly by both A3C with and without neighborhood rewards, even though these are on policy actor-critic methods with the same POMDP settings. We believe that this may be due to the way in which neighboring states are aggregated into each agents’ individual state via *discounted* summation in these baselines, which results in poorer policies. SocialLight, on the other hand, significantly outperforms standard A3C methods with and without neighborhood reward in terms of average trip time and cumulative delays. This is most likely due to the proposed contribution marginalization scheme within the neighborhood reward which tightly couples the given agent with neighboring agents to maximize throughput across the neighborhood, thereby reducing travel time and delays.

5.3 Traffic Performance on CityFlow Real Traffic Datasets

We evaluate SocialLight on real traffic datasets. The experiment settings and POMDP settings are unchanged with respect to [37] for a fair comparison among all methods.

5.3.1 Baseline methods.

- (1) **FixedTime** [13]: Fixed time control considers a fixed cycle over phases with a pre-defined total phase length and pre-defined phase split over the total cycle length.
- (2) **MaxPressure** [28]: MP (max-pressure) control greedily selects the phase that can minimize the intersection pressure, where the pressure is calculated by the difference between the vehicles of incoming lanes and connected outgoing lanes.
- (3) **Co-Light** [30]: A state-of-the-art method that uses Graph Attention Neural Networks to accomplish junction level cooperation and has been trained via Deep Q learning.
- (4) **MP-Light** [5] MP-Light incorporates pressure in their states and reward to achieve state-of-the-art scalable ATSC over a city-level traffic network. MP-Light also applies the FRAP based training architecture that uses phase competition within various traffic movements to improve control performance.
- (5) **Attention-Light** Attention-Light [37] is a recent state-of-the-art model that incorporates self-attention to learn the phase correlation and competition in contrast to FRAP that uses human knowledge. This method has shown to outperform FRAP based MPLight and the CoLight over numerous traffic datasets.

5.3.2 Analysis. Our results show that SocialLight in Table 1 outperforms all existing baselines over a wide variety of traffic flows in terms of reducing average travel time. In particular, we observe the highest performance gains over the New York traffic flow sets that comprises 196 traffic agents. We report more modest performance gains over the Hangzhou and Jinan traffic flow datasets.

Over traffic flow datasets compiled for Jinan and Hangzhou urban networks, SocialLight shows modest improvements compared to the state-of-the-art AttentionLight. Despite large gains on NewYork, the relatively modest performance gains of SocialLight over the Hangzhou and Jinan can be attributed to the saturation. This saturation results from their relatively small scale in terms of the number of agents in the training datasets. This saturation is evident from the performance classical baselines such as Max-Pressure over these traffic flow. For instance, the average trip time performance is actually observed to degrade with Co-Light and MP-Light as compared to the MaxPressure baseline on the Hangzhou and Jinan datasets. Attention-light is found to have very marginal performance gains as compared to the MaxPressure baseline (except for Hangzhou-2). This observations indicate that any potential performance gains from these datasets would be minimal due to

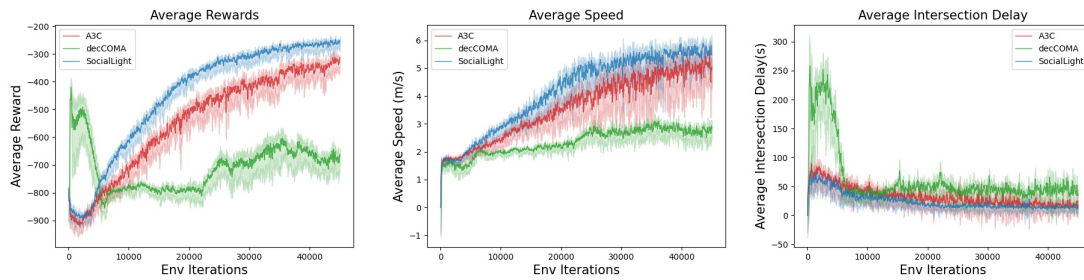


Figure 4: Training Plots of methods over average rewards, vehicle speed and intersection delay on the SUMO Manhattan synthetic data-set. SocialLight is shown in Blue, A3C with Neighborhood Rewards is in red and SocialLight with original COMA advantages which we refer to as decCOMA is in green. Observe that decCOMA fails to improve over cumulative returns.

saturation. Hence, SocialLight performs better than AttentionLight and the MaxPressure baselines on most traffic datasets.

We then turn our attention to the New York traffic flow sets, which are harder to learn over due to their sheer scale. Of the two flow sets there, New York-2 is the hardest traffic dataset [37] in terms of vehicle arrival rate. Most RL methods (CoLight, MPLight and AttentionLight) produce modest performance gains compared to the classical FixedTime and MaxPressure baselines on both the New York traffic flowsets. Co-light improves over the New York-2 dataset compared to the classical methods, but fails to do so over the New York-1 dataset. In contrast, AttentionLight fails to improve over the classic MaxPressure controller in the NewYork-2 dataset. We believe that these improvements indicate the need for true cooperation, which Co-Light and AttentionLight only achieve via their network designs.

SocialLight pushes the state-of-the-art in terms of traffic performance on both these New York datasets, where our performance gains over the current state-of-the-art AttentionLight and CoLight are even more pronounced ($\geq 20\%$). We believe that this is due to the impact of agents marginalizing individual contributions to the neighborhood reward, which encourages better cooperation as each agent understands its role in the local traffic performance. By overlapping the agents’ neighborhoods (i.e., local areas of enhanced cooperation), our approach effectively leads to improved network-wide cooperation and thus large performance gains.

5.4 Impact of Individual Contribution Marginalization on Learning and Final Performance

This ablation study aims to identify the impact of contribution marginalization on both the learning process and the final performance of trained policies. We compare SocialLight with vanilla Asynchronous Actor Critic (A3C) with reward sharing (rewards of neighbors summed up). As shown in Fig. 4, SocialLight exhibits improved sample efficiency and improved returns as compared to A3C even though both algorithms are on-policy and have the same architectures and augmented state inputs for their policy networks.

Furthermore, we note that decentralized contribution marginalization introduced in SocialLight is also shown to have improved the stability of training. While both algorithms converge, Fig 4 indicates larger variances in the returns, average speed and intersection delays during training for A3C with neighborhood reward sharing.

We further compare both the networks over a validation set of synthetic traffic flows generated from separate seeds that generate traffic flows during training. The improvement over the total trip time and cumulative delay shown in Table 2 highlight how the marginalization of individual contributions promote cooperation between traffic agents for large performance gains.

5.5 Impact of Modified Advantages on Training Stability and Convergence

Further analysis shows the impact of the modified advantages that marginalizes individual contributions. Compared to simply applying COMA advantages [9], our modifications are shown to improve training stability as shown in Fig 4 and overall returns. We believe that original COMA advantages converge to sup-optimal policies due to the highly distributed nature of our training process, where high initial biases in the critic inhibit monotonic policy improvement during the initial phase of training. As both the policy and networks learn in conjunction, the policy network learned via the original COMA advantages converge to sub-optimal traffic control policies, differently from SocialLight.

6 CONCLUSION

This paper presents SocialLight, a fully decentralized training framework that learns cooperative traffic light control policies via distributedly marginalizing individual contributions to each agent’s local neighborhood reward. We show that our method improves over the scalability for cooperative learning, thereby improving final traffic performance (average trip time) especially over large traffic networks such as the New York grid with 196 traffic intersections. These performance gains suggest that our method could improve the overall quality of learned policies on real-life citywide networks.

Our method leverages the fixed spatial structure of traffic systems to define overlapping neighborhoods over which agents can marginalize their contributions, to scale cooperative learning without the need for a centralized critic. Future work will focus on extending this idea to general mixed competitive-cooperative games, where such distributed spatial structures are more difficult to identify and leverage. There, we will aim to develop methods to learn suitable neighborhoods over which individual agents can marginalize their contributions over to improve the scalability of cooperative learning.

ACKNOWLEDGMENTS

This work was partly supported by A*STAR, CISCO Systems (USA) Pte. Ltd and National University of Singapore under its CISCO Accelerated Digital Economy Corporate Laboratory (Award I21001E0002).

REFERENCES

- [1] Mohammad Aslani, Mohammad Saadi Mesgari, and Marco Wiering. 2017. 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* 85 (2017), 732–752.
- [2] Mohammad Aslani, Saadi Mesgari, and Marco Wiering. 2017. 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* 85 (12 2017), 732–752. <https://doi.org/10.1016/j.trc.2017.09.020>
- [3] Noe Casas. 2017. Deep deterministic policy gradient for urban traffic light control. *arXiv preprint arXiv:1703.09035* (2017).
- [4] Chi-Chun Chao, Jun-Wei Hsieh, and Bor-Shiun Wang. 2022. Cooperative Reinforcement Learning on Traffic Signal Control. *arXiv preprint arXiv:2205.11291* (2022).
- [5] Chacha Chen, Hua Wei, Nan Xu, Guanjie Zheng, Ming Yang, Yuanhao Xiong, Kai Xu, and Zhenhui Li. 2020. Toward a thousand lights: Decentralized deep reinforcement learning for large-scale traffic signal control. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 3414–3421.
- [6] Tianshu Chu, Jie Wang, Lara Codecà, and Zhaojian Li. 2019. Multi-agent deep reinforcement learning for large-scale traffic signal control. *IEEE Transactions on Intelligent Transportation Systems* 21, 3 (2019), 1086–1095.
- [7] Tianshu Chu, Jie Wang, Lara Codecà, and Zhaojian Li. 2019. Multi-Agent Deep Reinforcement Learning for Large-scale Traffic Signal Control. <https://doi.org/10.48550/ARXIV.1903.04527>
- [8] Samah El-Tantawy and Baher Abdulhai. 2012. Multi-Agent Reinforcement Learning for Integrated Network of Adaptive Traffic Signal Controllers (MARLIN-ATSC). *Conference Record - IEEE Conference on Intelligent Transportation Systems*, 319–326. <https://doi.org/10.1109/ITSC.2012.6338707>
- [9] Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. 2017. Counterfactual Multi-Agent Policy Gradients. <https://doi.org/10.48550/ARXIV.1705.08926>
- [10] Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. 2018. Counterfactual multi-agent policy gradients. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 32.
- [11] Wade Genders and Saideh Razavi. 2016. Using a deep reinforcement learning agent for traffic signal control. *arXiv preprint arXiv:1611.01142* (2016).
- [12] Jayesh K Gupta, Maxim Egorov, and Mykel Kochenderfer. 2017. Cooperative multi-agent control using deep reinforcement learning. In *International conference on autonomous agents and multiagent systems*. Springer, 66–83.
- [13] Peter Koonce and Lee Rodegerdts. 2008. *Traffic signal timing manual*. Technical Report. United States. Federal Highway Administration.
- [14] Daniel Krajzewicz. 2010. Traffic simulation with SUMO—simulation of urban mobility. In *Fundamentals of traffic simulation*. Springer, 269–293.
- [15] Li Li, Yisheng Lv, and Fei-Yue Wang. 2016. Traffic signal timing via deep reinforcement learning. *IEEE/CAA Journal of Automatica Sinica* 3, 3 (2016), 247–254.
- [16] Pablo Alvarez Lopez, Michael Behrisch, Laura Bieker-Walz, Jakob Erdmann, Yun-Pang Flötteröd, Robert Hilbrich, Leonhard Lücken, Johannes Rummel, Peter Wagner, and Evamarie Wießner. 2018. Microscopic Traffic Simulation using SUMO. In *The 21st IEEE International Conference on Intelligent Transportation Systems. IEEE Intelligent Transportation Systems Conference (ITSC)*. <https://elib.dlr.de/124092/>
- [17] Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. *Advances in neural information processing systems* 30 (2017).
- [18] P Lowrie. 1990. Scats—a traffic responsive method of controlling urban traffic. *Sales information brochure published by Roads & Traffic Authority, Sydney, Australia* (1990).
- [19] Volodymyr Mnih, Adria Puigdomenech Badia, Mehdi Mirza, Alex Graves, Timothy Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*. PMLR, 1928–1937.
- [20] Seyed Sajad Mousavi, Michael Schukat, and Enda Howley. 2017. Traffic light control using deep policy-gradient and value-function-based reinforcement learning. *IET Intelligent Transport Systems* 11, 7 (2017), 417–423.
- [21] Tomoki Nishi, Keisuke Otaki, Keiichiro Hayakawa, and Takayoshi Yoshimura. 2018. Traffic signal control based on reinforcement learning with graph convolutional neural nets. In *2018 21st International conference on intelligent transportation systems (ITSC)*. IEEE, 877–883.
- [22] Afshin Oroojlooy, Mohammadreza Nazari, Davood Hajnejzad, and Jorge Silva. 2020. Attendlight: Universal attention-based reinforcement learning model for traffic signal control. *Advances in Neural Information Processing Systems* 33 (2020), 4079–4090.
- [23] LA Prashanth and Shalabh Bhatnagar. 2011. Reinforcement learning with average cost for adaptive control of traffic lights at intersections. In *2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 1640–1645.
- [24] Roger P Roess, Elena S Prassas, and William R McShane. 2004. *Traffic engineering*. Pearson/Prentice Hall.
- [25] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
- [26] Zheng Tang, Milind Naphade, Ming-Yu Liu, Xiaodong Yang, Stan Birchfield, Shuo Wang, Ratnesh Kumar, David Anastasiu, and Jenq-Neng Hwang. 2019. Cityflow: A city-scale benchmark for multi-target multi-camera vehicle tracking and re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 8797–8806.
- [27] Elise Van der Pol and Frans A Oliehoek. 2016. Coordinated deep reinforcement learners for traffic light control. *Proceedings of learning, inference and control of multi-agent systems (at NIPS 2016)* (2016).
- [28] Pravin Varaiya. 2013. Max pressure control of a network of signalized intersections. *Transportation Research Part C: Emerging Technologies* 36 (2013), 177–195.
- [29] Hua Wei, Chacha Chen, Guanjie Zheng, Kan Wu, Vikash Gayah, Kai Xu, and Zhenhui Li. 2019. Presslight: Learning max pressure control to coordinate traffic signals in arterial network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1290–1298.
- [30] Hua Wei, Nan Xu, Huichu Zhang, Guanjie Zheng, Xinshi Zang, Chacha Chen, Weinan Zhang, Yanmin Zhu, Kai Xu, and Zhenhui Li. 2019. Colight: Learning network-level cooperation for traffic signal control. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 1913–1922.
- [31] Hua Wei, Guanjie Zheng, Vikash Gayah, and Zhenhui Li. 2019. A Survey on Traffic Signal Control Methods. *arXiv preprint arXiv:1904.08117* (2019).
- [32] Hua Wei, Guanjie Zheng, Huaxiu Yao, and Zhenhui Li. 2018. Intellilight: A reinforcement learning approach for intelligent traffic light control. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2496–2505.
- [33] MA Wiering, J van Veenen, Jilles Vreeken, and Arne Koopman. 2004. Intelligent traffic light control.
- [34] Donghan Xie, Zhi Wang, Chunlin Chen, and Daoyi Dong. 2020. Iedqn: Information exchange dqn with a centralized coordinator for traffic signal control. In *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.
- [35] Chao Yu, Akash Velu, Eugene Vinitzky, Yu Wang, Alexandre Bayen, and Yi Wu. 2021. The Surprising Effectiveness of PPO in Cooperative, Multi-Agent Games. <https://doi.org/10.48550/ARXIV.2103.01955>
- [36] Chao Yu, Akash Velu, Eugene Vinitzky, Yu Wang, Alexandre Bayen, and Yi Wu. 2021. The Surprising Effectiveness of PPO in Cooperative, Multi-Agent Games. <https://doi.org/10.48550/ARXIV.2103.01955>
- [37] Liang Zhang, Qiang Wu, and Jianming Deng. 2022. AttentionLight: Rethinking queue length and attention mechanism for traffic signal control. <https://doi.org/10.48550/ARXIV.2201.00006>
- [38] Guanjie Zheng, Yuanhao Xiong, Xinshi Zang, Jie Feng, Hua Wei, Huichu Zhang, Yong Li, Kai Xu, and Zhenhui Li. 2019. Learning phase competition for traffic signal control. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 1963–1972.