

Autonomous vehicles need social awareness to find optima in multi-agent reinforcement learning routing games.

Anastasia Psarou

Jagiellonian University, Faculty of Mathematics and Informatics

Krakow, Poland

anastasia.psarou@doctoral.uj.edu.pl

Dominik Gawel

University of Warsaw, Faculty of Mathematics, Informatics, and Mechanics

Warsaw, Poland

dg448617@students.mimuw.edu.pl

Łukasz Gorczyca

Jagiellonian University, Faculty of Mathematics and Informatics

Krakow, Poland

lukasz.gorczyca@student.uj.edu.pl

Rafał Kucharski

Jagiellonian University, Faculty of Mathematics and Informatics

Krakow, Poland

rafal.kucharski@uj.edu.pl

ABSTRACT

Previous work has shown that when multiple selfish Autonomous Vehicles (AVs) are introduced to future cities and start learning optimal routing strategies using Multi-Agent Reinforcement Learning (MARL), they may destabilize traffic systems, as they would require a significant amount of time to converge to the optimal solution, equivalent to years of real-world commuting. We demonstrate that moving beyond the selfish component in the reward significantly relieves this issue. If each AV, apart from minimizing its own travel time, aims to reduce its impact on the system, this will be beneficial not only for the system-wide performance but also for each individual player in this routing game. By introducing an intrinsic reward signal based on the marginal cost matrix, we significantly reduce training time and achieve convergence more reliably. Marginal cost quantifies the impact of each individual action (route-choice) on the system (total travel time). Including it as one of the components of the reward can reduce the degree of non-stationarity by aligning agents' objectives. Notably, the proposed counterfactual formulation preserves the system's equilibria and avoids oscillations. Our experiments show that training MARL algorithms with our novel reward formulation enables the agents to converge to the optimal solution, whereas the baseline algorithms fail to do so. We show these effects in both a toy network and the real-world network of Saint-Arnoult. Our results optimistically indicate that social awareness (i.e., including marginal costs in routing decisions) improves both the system-wide and individual performance of future urban systems with AVs.

KEYWORDS

Multi-agent reinforcement learning, autonomous vehicles, urban route choice, marginal cost matrix

ACM Reference Format:

Anastasia Psarou, Łukasz Gorczyca, Dominik Gawel, and Rafał Kucharski. 2026. Autonomous vehicles need social awareness to find optima in multi-agent reinforcement learning routing games..

In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAA-MAS, 10 pages.

1 INTRODUCTION

Every day, selfish human drivers [41] make routing decisions to travel from an origin to a destination through a traffic network. For instance, they commute from home to work by selecting among multiple possible routes, typically the one with the lowest expected travel time [7]. The anticipated integration of Autonomous Vehicles (AVs) into future transportation systems will introduce new dynamics, as AVs may not only make independent routing decisions but also possess capabilities for coordination, learning, or communication [13]. Unlike human drivers, the objectives of AVs may not necessarily be selfish. To model AV routing decisions, Multi-Agent Reinforcement Learning (MARL) has emerged as a promising method, with previous research already exploring this application [4, 22, 38, 39]. Yet despite its promise, MARL presents challenges when applied to real-world traffic environments. Prior research has shown that when multiple *selfish* AVs simultaneously learn routing strategies using state-of-the-art MARL algorithms, they may destabilize traffic networks [30]. This can happen because the algorithms used in that work require long training iterations, equivalent to several years of real-world commuting, to converge to the optimal solution (system optimal or individually optimal), or in some cases, fail to converge at all.

To address this instability and enable faster convergence to the optimal solution, we propose an intrinsic reward signal that complements the standard *selfish* travel-time-based reward. In intrinsically motivated Reinforcement Learning (RL) [10], the reward signal encourages agents to explore or learn behaviors beyond what is explicitly defined by selfish (external) rewards, like social behaviors. Jaques et al. [21] demonstrated that rewarding agents for having causal influence over other agents' actions, using counterfactual reasoning, can promote coordination and communication between them. **In this work, we demonstrate that incorporating the impact an AV agent's presence has on the other agents of the system into its reward value enables the system to reach the optimum solution faster.** Specifically, we show that:

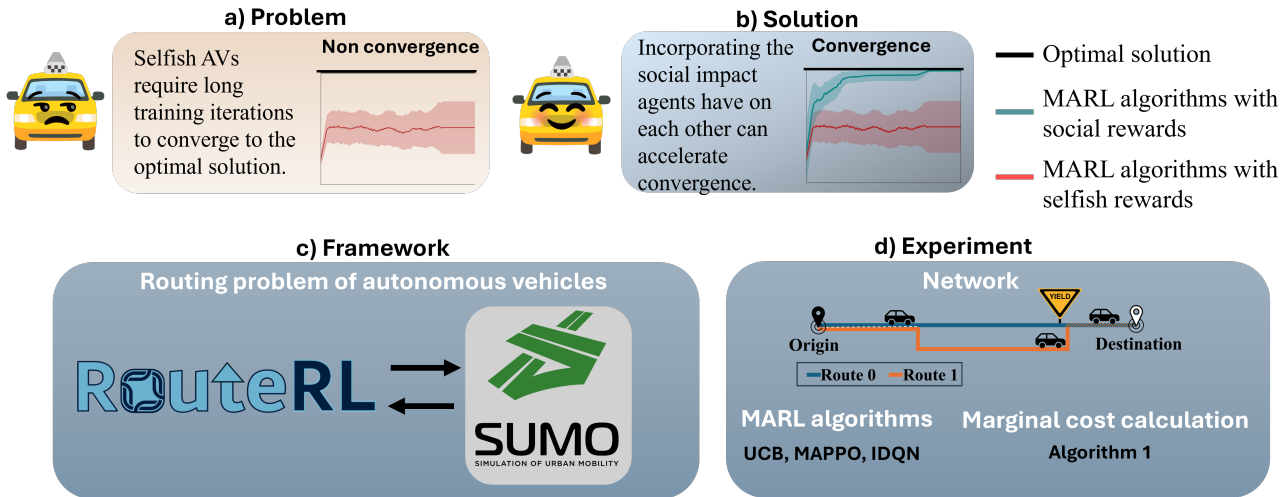


Figure 1: Overview: When multiple selfish AVs are introduced into future cities and have to simultaneously learn optimal routing strategies using MARL they need long training iterations to converge to the optimal solution (a). However, if we incorporate the impact an AV’s presence has on the other agents of the system, into its travel-time based reward value, can improve the convergence of MARL algorithms applied to the route choice problem of AVs (b). This enables faster convergence to the system-optimal and is can also be more beneficial for individual AV agents. We show this with experiments using RouteRL [5], a Multi-Agent Reinforcement Learning (MARL) framework that models the routing decisions of AVs and human drivers (c), and demonstrate our results using the Two-Route (Yield) network (TRY), d), where agents choose between *Route 0*, which is the shortest route without priority, and *Route 1*, a slightly longer alternative with priority.

- **Incorporating a social component in the reward of the AV agents accelerates convergence to the optimal solution.** Specifically, including the marginal cost in the immediate reward of the AVs brings the system faster to the System Optimal (SO) solution (as demonstrated in Figs. 1b, 2), both in deterministic and non-deterministic traffic settings. This leads to more efficient solutions, as both AVs and human agents experience shorter average travel times, once the training process is complete, as shown in Table 1.
- **The equilibria of the system remain unchanged as the coefficient of the counterfactual intrinsic reward increases.** This result demonstrates that, as the coefficient of the intrinsic reward increases, convergence is accelerated towards the unchanged system optimal equilibria, as depicted in Figs. 3, 4.
- **Our results are demonstrated in a small and in a bigger real-world network.** Initially, we show that AVs converge faster in the Two-Route (Yield) network (TRY), depicted in Fig. 1d that is deliberately designed for studying this problem. In section 3.5, we extend our method in the real-world Saint-Arnoult traffic network and show that the system and individual travel times are reduced when social AVs are introduced in the system (Table 3).
- **Our method, specifically the way we compute the marginal cost matrix and incorporate it in the reward signal of an agent, is general and can be incorporated in any MARL algorithm for the routing of AVs.**

1.1 Related work

In recent years, AVs have attracted significant research interest [23, 35, 43, 45]. Studies suggest that their integration may improve road safety and reduce traffic congestion [1, 46]. Other studies highlight potential negative effects, such as increased emissions [27]. However, limited research has examined the route choice decisions of AVs and their potential effects, both positive and negative, in urban environments.

Jamróz et al. [20] investigated the route choice problem for Connected Autonomous Vehicles (CAVs) and Human Driven Vehicles (HDVs) and demonstrated that CAVs, by optimizing different goals, can either benefit or systematically disadvantage human drivers. In their analysis, CAVs differ from HDVs in their capabilities to make collective decisions. Several works have explored the AV route choice problem with the use of MARL [3, 4, 22, 38, 39]. For instance, Akman et al. [5] developed RouteRL, a MARL framework that simulates the routing decisions of AVs in different traffic networks shared with human drivers. RouteRL uses SUMO [24], a continuous, microscopic traffic simulator. Building on this work, Psarou et al. [30] demonstrated that when multiple selfish AVs simultaneously learn optimal routing strategies using MARL, they may destabilize traffic networks by increasing travel times for both AVs and human drivers.

MARL encompasses several learning paradigms. Independent learning (IL) is one of them, where each agent learns its own policy while treating the other agents as part of the environment [36]. A widely used RL algorithm is Deep Q-learning [26], whose IL

MARL variant is the Independent Deep Q-learning (IDQN). Another MARL learning paradigm is the Centralized Training Decentralized Execution (CTDE), where agents learn decentralized policies in a centralized way that are used independently in the execution phase. The Multi-Agent Proximal Policy Optimization (MAPPO) algorithm [44] is a representative CTDE approach and is regarded as a baseline for cooperative MARL tasks [28]. Finally, in certain scenarios, bandit settings, each agent selects a single action per episode. A classic multi-armed bandit algorithm is the Upper Confidence Bound (UCB) algorithm [8].

In congested road networks, each additional traveler imposes an external cost on others by increasing their travel times [40]. Peeta et al. [29] defined the Marginal Travel Time (MTT) of a time-dependent path as the impact that adding one vehicle to a given path at a specific timestep has on the overall system travel time in a macroscopic setting. Sheffi et al. [34] described the MTT of a link as the "marginal contribution of an additional traveler on the link to the total travel time on this link."

Marginal cost pricing (MCT) has been studied as a means of improving traffic efficiency. The core idea is to charge each traveler a toll equal to the additional delay other drivers in the system experience because of the traveler's presence in the system [11]. Building on this concept, Ramos et al. [31] proposed a reinforcement learning algorithm that realigns heterogeneous preferences of drivers over travel time with the marginal cost imposed on the user as a toll. This tolling algorithm leads to a system-efficient equilibrium. Additionally, Sharon et al. [33] studied the impact of the MCT and demonstrated that underestimating its value can result in system performance comparable to scenarios without any tolling.

In this work, we incorporate the marginal travel time into the agent's reward function; similar reward-shaping strategies have been adopted in prior studies addressing the routing problem of vehicles. For example, Grunitzki et al. [16] studied a setting where drivers learn their route choices, using difference rewards. In the case of difference rewards, a default action is chosen, and the agent's reward is calculated as the difference between the travel time of the system and the travel time of the system when the agent chooses the default action. In their method, the difference reward of an agent is defined as the global system reward minus the global system reward when the specific agent is eliminated from the system. This approach is inspired by Tumer et al. [39] that demonstrated that using difference utility rewards can lead agents to achieve outcomes that are not only beneficial for themselves but also more advantageous for the system as a whole. In both of these approaches, the agent's reward is computed by approximating the system's travel time and excluding the states where the specific agent has no impact.

Our approach is related to counterfactual techniques. In COMA-style counterfactual baselines, the credit term marginalizes over the agent's actions while holding others fixed, which can be costly in large action spaces [14]. In contrast, our method requires only a single counterfactual evaluation (presence versus absence), which makes it substantially more tractable in large action spaces. Additionally, our counterfactual preserves the set of equilibria in the system and improves stability [12, 17], as we show in section 3.4.

To the best of our knowledge, no prior work has addressed the problem identified in [30]: selfish AV agents, when learning

their route choices using MARL, converge to suboptimal solutions that destabilize traffic networks. In particular, it remains unclear whether introducing social behavior to the AVs could mitigate this problem. By addressing this problem, we not only resolve the convergence issue but also enhance the societal impact, making systems with limited capacity more efficient and effective.

2 METHOD

We begin by introducing the simulation environment in section 2.1 within which, in section 2.2, we present the problem formulation. The immediate reward function of each agent is shown in section 2.3, followed with formulation of the marginal travel time (MTT) and explanation how it is incorporated into the intrinsic reward structure of each AV agent in section 2.4.

2.1 Urban routing problem

To simulate the day-to-day routing decisions of human drivers and AVs, we utilize RouteRL [5], a MARL framework integrated with SUMO [24], a continuous microscopic traffic simulator that enables modeling complex urban traffic networks and agent interactions. Within RouteRL, drivers everyday (iteration) choose a route to travel from their origin to their destination; for example, traveling from home to work at a specific time. Feasible routes are computed by Janux [2]. Set of routes do not change during the simulation or from one day to the next, and both departure times and demand remain constant across days of simulation.

Within RouteRL, human drivers are represented as self-interested utility maximizers who iteratively adapt their route choices based on their experienced travel times [9, 15]. Their learning process is assumed to guide the system toward a Wardrop equilibrium, a state in which no individual can improve their outcome by unilaterally changing routes [41]. Once this equilibrium is reached, AVs are introduced into the network and use MARL algorithms to learn optimal routing policies through repeated interactions with the environment.

Both AVs and human drivers interact with the environment by selecting a route from a predefined set of discrete alternatives to reach their destination. Agents' decisions, due to limited capacity, contribute to system travel costs (congestion). This setting appears non-stationary from the perspective of a single agent, as the actions of the agents jointly influence the environment. AVs receive as observation the current state of the network, that is, the route choices of the vehicles that have already departed. Additionally, selfish AV agents receive a reward equal to the negative of their travel time upon arrival, and their objective is to maximize this reward value (see details in [5]).

Using SUMO, we can simulate deterministic and non-deterministic traffic dynamics. In deterministic scenarios, traffic dynamics remain fixed, and identical joint actions result in identical travel times. We model non-determinism by varying SUMO's random seed, allowing identical route choices to produce different travel times for each agent, as in real-world traffic networks.

2.2 Problem formulation

We abstract the daily (repeated) route choices made by humans (and AVs in the future) in capacitated networks (limited resources)

as a repeated congestion game [19]. This can be formalized as a one-cycle Agent Environment Cycle (AEC) problem [37], defined in [3, 30] as a tuple:

$$\langle \mathcal{S}, \mathcal{I}, \{\mathcal{U}_i\}_{i \in \mathcal{I}}, \{r_i\}_{i \in \mathcal{I}}, \{\Omega_i\}_{i \in \mathcal{I}}, v \rangle$$

We consider that in each episode (day or iteration) \mathcal{I} AV agents in order v of departure time, choose a route from their action space. We consider:

- A finite set of states \mathcal{S} .
- A finite set of agents \mathcal{I} such that $N = |\mathcal{I}|$ (\mathcal{H} is the set of human agents)
- For each agent $i \in \mathcal{I}$:
 - The agent selects an action $u_i \in U_i$, forming a joint action $\mathbf{u} \in \mathbf{U} \equiv U^N$.
 - r_i is the immediate reward agent i aims to maximize, as defined in [3, 30].
 - Ω_i is the set of possible observations for agent i .
- v is the agent selection mechanism. Agents act sequentially in the order of their departure times, so that those with earlier departure times act first.

2.3 Intrinsic reward signal

In the proposed method, we demonstrate that incorporating the marginal travel time of other agents into the intrinsic reward of each AV, in the context of the route choice problem, enables the system to converge faster towards user equilibrium. Following the work of [21], the reward of an agent j is formulated as:

$$r_j(\mathbf{u}) = \alpha e_j(\mathbf{u}) + \beta m_j(\mathbf{u}), j \in \mathcal{I} \quad (1)$$

where e_j is the extrinsic environmental reward and m_j is the influence reward of an agent j and α, β are the weighting coefficients of each of these components.

2.4 Marginal travel time

In this study, we followed the counterfactual definitions of the marginal travel time discussed in section 1.1 and formulated it as the change in the travel time of agent i that arises when agent j is removed from the system, compared to when agent j remains present in the simulation. We define the marginal cost matrix as $M_{i,j}(\mathbf{u})$, a square matrix where rows represent an AV agent i and columns represent an AV agent j . The entry at position (i, j) denotes the travel time difference experienced by agent i when agent j is present compared to when agent j is removed from the system. The travel time of agent i when agent j is removed from the system is indicated as $e_i(\mathbf{u}, \mathbf{u}^{-j})$. The intrinsic reward for agent j , denoted as $m_j(\mathbf{u})$, is defined as the sum of values in column j of the marginal cost matrix, representing the total impact agent j has on the travel times of others. This impact is then transformed using the hyperbolic tangent function to produce a bounded, sign-preserving score that limits the influence of extreme values while preserving the direction of the effect.

$$m_j(\mathbf{u}) = \sum_{\substack{i=1 \\ i \neq j}}^N \tanh(e_i(\mathbf{u}) - e_i(\mathbf{u}^{-j})) \quad (2)$$

Algorithm 1 Marginal cost matrix calculation for joint action \mathbf{u}_T

Require: Simulation run T , parameters Θ , agent set \mathcal{I} , joint action \mathbf{u}_T , travel times t_T

Ensure: Marginal cost matrix M

```

1:  $M_{i,j}(\mathbf{u}) \leftarrow \mathbf{0}^{N \times N}$ 
2: for  $j \in \mathcal{I}$  do
3:    $\mathcal{I}^{(-j)} \leftarrow \mathcal{I} \setminus \{j\}$  ▷ Remove agent  $j$ 
4:    $\mathbf{u}_T^{(-j)} \leftarrow \mathbf{u}_T \setminus \{u_j\}$ 
5:   env  $\leftarrow$  TrafficEnvironment( $\mathcal{I}^{(-j)}, \Theta$ )
6:   Run simulation  $\mathcal{E}$  with env under actions  $\mathbf{u}_T^{(-j)}$ 
7:   for  $i \in \{1, \dots, A_m\}$  do
8:      $t_{\mathcal{E}}(i) \leftarrow$  travel time of agent  $i$  in run  $\mathcal{E}$ 
9:      $M_{i,j}(\mathbf{u}) \leftarrow t_{\mathcal{E}}(i) - t_T(i)$ 
10:  end for
11: end for
12: return  $M_{i,j}(\mathbf{u})$ 

```

$$m_j(\mathbf{u}) = \sum_{i=1}^N M_{i,j}(\mathbf{u}), \quad N = |\mathcal{I}| \text{ or } N = |\mathcal{H} \cup \mathcal{I}| \quad (3)$$

To construct the marginal cost matrix $M_{i,j}(\mathbf{u})$, we rerun (for each joint action \mathbf{u}) simulation N times, once for each removed AV agent, while keeping the joint action and SUMO random seed fixed. In each run, a single agent j is removed from the environment, and the resulting changes in travel times for all other agents are recorded. These values populate column j of the marginal cost matrix, $M_{i,j}$, as in Algorithm 1. As a result, constructing the full matrix for one episode (one joint action) requires generating N additional and separate environment instances for each joint action. For example, in the TRY network depicted in Fig. 1d, which has 2 discrete actions (routes), a scenario with 10 agents yields 1024 joint actions. If each joint action is encountered during training, an additional 1024×10 environment runs will be needed.

To better understand the observed phenomenon, we consider two scenarios for calculating marginal cost. In the first scenario, AVs incorporate in their intrinsic reward the impact they have on the other AVs of the system, referred to as *AV group marginal* (eq. 3). To better understand how AVs can affect the total system, we incorporate the *system marginal* (eq. 3), which is the impact of an AV on all the drivers in the system, including both humans and AVs, denoted as *system marginal*.

Notably, when the intrinsic part of the reward (eq. 1) is computed as above, we obtain the interplay between the two often contradicting possible states in the urban mobility systems. The first part represents the natural, selfish inclination of utility-maximizing travelers to arrive at the destination as quickly as possible, while the latter represents the ideal cooperation to make the system as efficient as possible (see more in [18]). Mixing the two in one formula may lead to interesting trade-off solutions. Here, we exploit this property in a system where user equilibrium (UE) is also system-optimal (SO) (which is rarely the case in complex systems), we sketch a proof suggesting that the resulting solutions will be stable somewhere between SO and UE in section 3.4.

Table 1: Average travel times during the testing phase, last 100 episodes of the experiment, (in seconds) for each subgroup (AVs, humans), with standard deviations in parentheses.

	Scenario	Reward type	Algorithms	Travel times	
				AVs	Humans
Deterministic	Selfish AVs		UCB	65.7 (15.6)	65.1 (24.3)
			MAPPO	62.0 (14.8)	64.4 (27.3)
			IDQN	64.1 (14.5)	70.3 (28.6)
	AV group marginal		UCB	57.9 (12.0)	51.7 (14.4)
			MAPPO	57.8 (12.4)	51.9 (14.8)
			IDQN	56.8 (11.5)	50.9 (13.6)
	System marginal		UCB	57.9 (12.6)	51.2 (14.5)
			MAPPO	57.7 (12.2)	51.1 (14.3)
			IDQN	59.2 (12.0)	51.7 (13.4)
Non Deterministic	Selfish AVs		UCB	62.9 (14.4)	61.5 (19.8)
			MAPPO	62.0 (15.1)	64.3 (27.5)
			IDQN	63.7 (14.6)	67.7 (27.8)
	AV group marginal		UCB	56.9 (11.7)	51.0 (13.8)
			MAPPO	57.0 (11.7)	51.0 (13.9)
			IDQN	59.1 (12.4)	52.1 (13.8)
	System marginal		UCB	57.0 (12.0)	50.9 (13.8)
			MAPPO	57.0 (11.7)	51.0 (13.9)
			IDQN	59.7 (12.1)	52.0 (13.9)

3 RESULTS

We begin by introducing the minimal experimental setting employed in this work in section 3.1. Next, in section 3.2, we present the main finding of this study: *our proposed reward formulation achieves convergence to the optimal solution*. Afterwards, we examine the impact of different β values on the convergence of the MARL algorithms in section 3.3. In section 3.4, we show that under certain conditions the equilibria of the system remain unchanged despite increasing in the β coefficient. Finally, in section 3.5, We replicate our findings for the real-world Saint-Arnoult network, demonstrating that, after a short training, 50% of the AV agents achieve shorter travel times, as well as shorter system travel times, when the marginal cost is introduced.

3.1 Experimental setting

To demonstrate the impact of the intrinsic reward on the solution and training, we select the carefully tailored minimal toy network example, used in all the experiments of this work, except for section 3.5 where we experiment with more complex and realistic setting. In the TRY network agents everyday have the option to choose between two routes: Route 0, which is faster but does not have priority (yield sign), and Route 1, which is slightly longer but has priority. These experiments involve 22 agents. Initially, all are human drivers who learn their route-choice behavior for 200 days (iterations). During this learning phase, they converge to consistently selecting Route 0 (which, by design, is both Wardrop equilibrium and System Optimum).

Following the initial 200 episodes (days), 10 of the human drivers are assumed to switch to using AVs for their daily commute. We consider that AVs start learning their routing strategies from

episode 200 using different MARL algorithms. We select most popular, promising and fitting state-of-the-art RL algorithms, building on findings from URB benchmark [3] we use here: MAPPO [44], IDQN [26], and UCB [8]. After episode 200, we assume human drivers’ strategies remain fixed, reflecting the realistic assumption that human drivers typically settle on a perceived optimal route after repeated experience [41]. For simplicity, we assume that remaining human drivers will not adapt to the changes induced by AVs and their training.

3.2 Performance of the marginal cost reward formulation

In this section, we demonstrate that when running different MARL algorithms with our new reward formulation under both deterministic and non-deterministic traffic dynamics, the AV agents converge faster to the optimal solution, as depicted in Fig. 2. Specifically, in Fig. 2(a, b, d, e), after as little as 100 iterations, the agents reach a joint action that is close to the optimal one.

Additionally, we consider two scenarios of marginal cost calculation discussed in Section 2.4 into our demonstration in Fig. 2. One where AVs account for their impact on the other AVs in the system, referred to as *AV group marginal*, and the another where the intrinsic reward reflects the impact on all drivers in the system, denoted as *system marginal*. As shown in Fig. 2. We did not observe significant differences between the two formulations, both the *AV group marginal* and the *system marginal* require a similar amount of time to converge to the optimal solution.

Table 2: Sample from the marginal cost matrix for the AV group marginal case, computed at joint action [1, 0, 1, 0, 0, 0, 1, 0, 1, 1], under deterministic traffic dynamics, for first three AVs (values in seconds).

ID	AV1	AV3	AV5
AV 1	0	0	0
AV 3	-21	0	-9.6
AV 5	0	-1.2	0

Table 1 shows the average travel times of AVs and human agents during the testing phase, where AVs utilize the policy they learned without including any exploration. In both deterministic and non-deterministic traffic settings, the average travel times of AVs and human drivers are shorter in both the *AV group marginal* and the *System marginal* across all the tested algorithms. By introducing marginal cost, travel times are reduced from over 70s (IDQN) to 57 seconds for AVs and to 51 for human drivers. Notably, AV agents achieve similar travel times in both marginal formulations (*AV group marginal*, *System marginal*), which are shorter compared to when the marginal cost is not introduced. These results support our claim that, when the marginal cost is introduced into the reward of AV agents, they converge faster to the system and individually optimal solutions.

Table 2, samples values of the marginal cost matrix. Notably, the matrix is often not lower triangular, which indicates that the system does not satisfy the First-In-First-Out (FIFO) property. This occurs because the priority scheme of the yield sign allows later-departing

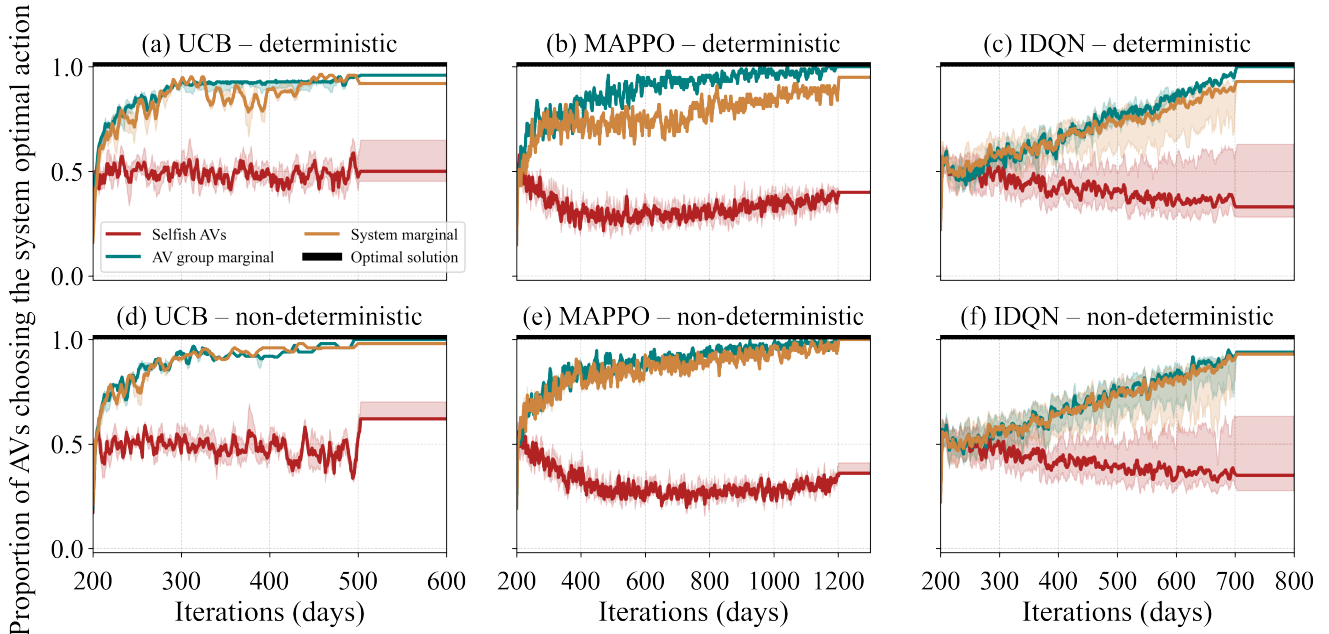


Figure 2: Introducing the marginal travel time in the reward of the AV agents accelerates convergence of the AVs to the optimal solution. We demonstrate this by incorporating two types of marginal travel times (MTTs): one that considers the impact on AV agents, *AV group marginal*, and the second that considers the impact on all drivers, *system marginal* (eq. 3). We also demonstrate that our method enhances convergence in non-deterministic traffic dynamics, which more accurately represent real-world traffic conditions. When the proportion of AVs choosing the optimal action is close to 1, it indicates that nearly all agents selected the optimal solution. A proportion of 0.5 indicates that half of the agents chose the optimal option. The last 100 iterations of the plots depict the evaluation mode, where agents use the learned policy without exploration.

agents to influence earlier-departing ones, thereby amplifying non-stationarity. For example, in Table 2, AV agent 3 selects route 0, which does not have priority, and AV agents 1 and 5 choose route 1. Agent 1 departs at timestep 1, agent 3 at timestep 3 and agent 5 at timestep 5. Consequently, agent 3 must yield to agents 1 and 5 in the intersection. Therefore, removing these two agents from the system would benefit agent 3 - which is reflected in the values of matrix M .

3.3 Impact of the social component weight β

Here, we investigate how incorporating agents' marginal impact into the reward value influences the convergence of the learning algorithms. Fig. 3 shows the convergence of the UCB [8], MAPPO [44], and IDQN [26] algorithms when the coefficient of the intrinsic reward, β takes different values. For UCB, β equal to 10 results in convergence to the optimal solution after approximately 300 training episodes. Suggesting that higher values of β (100 and 200) enable the agents to reach the optimal solution faster.

For the MAPPO algorithm, however, increasing β from 10 to 100 does not noticeably improve the convergence rate. For IDQN, all the experiments with the tested beta values converge to the optimal solution after approximately 400 iterations (except from the ones with $\beta = 0.3$). However, IDQN requires more iterations compared to UCB and MAPPO, where the system- and individually optimal

solutions are nearly reached after roughly one-third of the total training, as shown in Fig. 3. Based on these findings, and in line with the experimental results presented in Fig. 3, we select $\beta = 200$ for all the experiments shown in this paper, except for this section and section 3.5. The varying number of selected training iterations for each algorithm indicates that, depending on the selected number of training iterations, tuning the β value can facilitate convergence to the optimal solution.

3.4 Equilibria

Now we analyze the equilibria that arise in the system when the intrinsic rewards described in eq. 3 are introduced into the rewards of the AV agents. Equilibrium solutions, and specifically the Nash equilibria, have been adopted as the standard solution concept in MARL [6], we aim to understand how these equilibria can be affected by changes in the reward structure. Since any change in the equilibria can alter the agents' optimal policies, and consequently their observed behaviors [32]. Specifically, we investigate how varying the coefficients of the intrinsic and extrinsic rewards (α, β), as defined in eq. 1, influences the resulting equilibria.

To formally determine whether a joint action is a Nash equilibrium, we test the 2^{10} (1024) possible joint actions. For each of these possible joint actions $\mathbf{u} \in U^{10}$, we test unilateral deviations. A joint action is a Nash equilibrium if no agent can increase its




Figure 3: Effect of different values of the shaping coefficient β on the convergence of the UCB (a), MAPPO (b), and IDQN (c) algorithms. Higher values of β lead to faster convergence to the optimal solution. When the proportion of agents choosing the system optimal action is close to 1, it indicates that nearly all agents selected the optimal solution.

reward, sum of the negative travel time of the agent and the MTT, by changing its own action while the decisions of all other agents remain fixed. When assessing a unilateral deviation in the Nash equilibrium calculation, an agent j compares whether it is better for him to choose the alternative action (route 1 if the initial action was route 0 or route 0 if the initial action is route 1 in the TRY network).

We first analyze the case where agents receive no intrinsic reward and aim to minimize their own travel times ($\alpha = 1, \beta = 0$). In this setting, the system has a unique equilibrium in which all agents choose route 0. This corresponds to a *Wardrop Equilibrium* [42], the analogue of the Nash equilibrium in the route choice context, and in our specific setting, a *System Optimum*, as the total travel time across the network is minimized in that state [25].

We next explore how the introduction of intrinsic rewards and the variation of β affect the equilibrium structure (Fig. 4). Starting with $\alpha = 1, \beta = 0.3$, we observe that the equilibrium remains unchanged, compared to the state where $\alpha = 1, \beta = 0$. **Increasing β up to 100 leads to the same outcome: the equilibrium state is unaffected and it remains unique.**

Proof. One of our claims is that introducing the counterfactual marginal cost-based intrinsic reward (eq. 2) improves convergence without altering the underlying game. If the intrinsic reward were to shift the equilibrium, then our improvement would be dismissed as optimizing a different problem. Therefore, it is crucial to show that increasing the coefficient β beyond a certain threshold does not affect the equilibrium state. The proof described in this section concerns the TRY network but, can easily be extended to bigger networks.

The counterfactual intrinsic reward for an agent j , as described in eq. 2 is obtained by re-simulating the system, once with the agent j present and once where agent j is removed. Since agent j is removed from the simulation in this counterfactual no matter which route it would have chosen, the travel times of the remaining agents are the same whether j chose route 0 or route 1:

$$\sum_{\substack{i=1 \\ i \neq j}}^N e_i(\mathbf{u}^{-j}|_{u_j=0}) = \sum_{\substack{i=1 \\ i \neq j}}^N e_i(\mathbf{u}^{-j}|_{u_j=1}), \quad N = |I|$$




Figure 4: Number of equilibria in the system for different α, β coefficients. When $\beta > 0$ and $\alpha = 1$, the system has one unique equilibrium, the same as the equilibrium when the AV agents are selfish ($\beta = 0$ and $\alpha = 1$). This suggests that with β values above a certain threshold and $\alpha = 1$, the equilibrium state remains unchanged, as discussed in section 3.4. Additionally, for hypothetical values of $\beta < 0$ (malicious behavior, [20]), multiple equilibria become possible in the system.

Unilateral deviation $u_j : 0 \rightarrow 1$ (with \mathbf{u}^{-j} fixed). Let's consider the deviation of agent j from its chosen action while keeping the actions of all the other agents fixed. The only component of the intrinsic reward that is affected by agent j deviating is the travel times of all the individual agents when agent j chooses action 0 or 1. Hence, adding the marginal-cost term contributes an action-invariant constant and rescales the same $\sum_{i=1, j \neq i}^N (e_i(\mathbf{u}^{-j}))$ by a scale β , leaving the ordering of routes, and therefore best responses and the Nash equilibrium unchanged as β increases.

Agent j minimizes their reward: $r_j(\mathbf{u}) = e_j(\mathbf{u}) + \beta m_j(\mathbf{u})$. To verify whether the equilibrium is preserved, it is sufficient to examine the effect of a single agent unilaterally switching its route while all the other agents keep their choices fixed. We then define the own travel-time change of agent j and the corresponding system-wide change as follows:

$$\delta_j \triangleq e_j(\mathbf{u}^{-j}|_{u_j=1}) - e_j(\mathbf{u}^{-j}|_{u_j=0}), \quad \Delta C_j \triangleq \sum_{\substack{i=1 \\ i \neq j}}^N \tanh(\delta_i).$$

Agent's $j \in I$ reward change when switching from 0 to 1 is

$$\Delta r_j(\beta) \triangleq r_j(\mathbf{u}^{-j}|_{u_j=1}) - r_j(\mathbf{u}^{-j}|_{u_j=0}) = \underbrace{\delta_j}_{\text{own change}} + \beta \underbrace{\Delta C_j}_{\text{system change}}.$$

Invariance of best response (route 0 better for all $\beta \geq 0$).

Because \tanh is odd and strictly increasing, we have $\text{sign}(\tanh(\delta_i)) = \text{sign}(\delta_i)$. If in addition, we assume aligned externality, i.e., if $\text{sign}(\Delta C_j) = \text{sign}(\delta_j)$ then for every $\beta \geq 0$ holds $\text{sign}(\Delta r_j(\beta)) = \text{sign}(\delta_j)$. Indeed, if $\delta_j > 0$ and $\Delta C_j \geq 0$, then for every $\beta \geq 0$ we have $\Delta r_j(\beta) = \delta_j + \beta \Delta C_j \geq \delta_j > 0$, for $\delta_j < 0$ and $\Delta C_j \leq 0$, we get for every $\beta \geq 0$ that $\Delta r_j(\beta) = \delta_j + \beta \Delta C_j \leq \delta_j < 0$ and if

$\delta_j = 0 = \Delta C_j$, hence $\Delta r_j(\beta) = 0$ for all β . More generally, on a finite range $\beta \in [0, \beta_{\max}]$, where $\beta_{\max} = \sup\{\beta \geq 0 : \Delta r_j(\beta) > 0\}$

$$\Delta r_j(\beta) > 0 \forall \beta \in [0, \beta_{\max}] \iff \Delta C_j > -\frac{\delta_j}{\beta_{\max}}. \quad (4)$$

Fix $\beta_{\max} > 0$. Because $r_j(\beta)$ is affine in $\beta \in [0, \beta_{\max}]$, its minimum occurs at an endpoint. If $\Delta C_j \geq 0$, Δr_j is nondecreasing so its minimum is at $\beta = 0$. Thus

$$\Delta r_j(\beta) > 0 \forall \beta \in [0, \beta_{\max}] \iff \Delta r_j(0) = \delta_j > 0.$$

If $\Delta C_j < 0$, Δr_j is strictly decreasing so its minimum is at $\beta = \beta_{\max}$. Thus

$$\Delta r_j(\beta) > 0 \forall \beta \in [0, \beta_{\max}] \iff \Delta r_j(\beta_{\max}) = \delta_j + \beta_{\max} \Delta C_j > 0.$$

Therefore, when $\beta \in [0, \beta_{\max}]$, the equilibria of the system remain intact.

3.5 Real-world network

To test our claim that incorporating the marginal cost into the reward of the AV agents enhances system-wide performance and can reduce the travel times of individual AV agents, we present results from a bigger real-world network where the SO and UE solutions do not coincide. Specifically, we apply our new reward formulation to the real-world network of Saint-Arnoult from the URB benchmark [3], which includes 1289 nodes and 2011 edges. In the experiment, we consider 111 traveling agents, with each agent selecting from an action space of three available routes (leading to a huge joint action space 3^{110}). We consider that after period of human training, 10 of the human agents switch to traveling with AVs, learning the optimal route choice for 300 iterations using the UCB algorithm, and subsequently evaluating the learned policy for an additional 10 iterations. Notably, now (unlike in TRY network) the UE and SO solutions do not coincide.

Table 3: Travel times (in seconds) of 10 AVs in Saint-Arnoult with selfish objective and our proposed marginal intrinsic reward. Even in this short training experiment, for 6 out of 10 AVs using marginal cost in their reward formulation ($\beta = 0.3$) allowed to reduce their individual travel times (selfish rewards). Both the system-wide and AV group travel times were improved.

AV ID	Selfish AVs	AV group marginal	System marginal
2	125.17 (0.44)	126.95 (0.11)	127.01 (0.18)
37	116.56 (1.14)	117.59 (0.84)	117.55 (0.04)
80	112.43 (0.36)	113.24 (1.55)	112.40 (0.16)
52	132.54 (1.93)	132.36 (0.02)	132.36 (0.7)
107	165.10 (0.33)	163.04 (0.13)	163.41 (0.02)
4	181.86 (0.53)	183.57 (0.14)	180.12 (0.44)
95	192.98 (0.74)	193.06 (0.23)	193.86 (1.27)
74	215.28 (0.82)	213.25 (0.71)	212.51 (0.22)
102	213.56 (0.35)	214.88 (0.24)	215.20 (0.57)
34	400.33 (1.08)	393.61 (0.05)	393.24 (0.13)
System	27493.92 (15.55)	27494.81 (2.96)	27488.84 (0.74)
AV group	1855.05 (2.55)	1851.31 (1.58)	1847.85 (3.37)

In Table 3, we compare the average travel times (among replications) of each individual AV agent, of the overall system, and of the

AV group when the training is complete. The comparison is made under three reward schemes for AVs: the selfish travel-time-based reward, and the selfish travel-time-based reward including the AV group marginal, and the System marginal (eq. 1, 3). We observe that with a relatively short training (300 iterations, i.e. as much as in the simple TRY network, whereas we are now in a much more complex setting), more than 50% of the agents achieve shorter individual travel times when the marginal cost is introduced, as shown in Table 3. Additionally, the system and the AV group average travel times are reduced when we introduce the social component in the reward of the AVs. This result emphasizes that incorporating socially oriented behavior in AVs can be efficient in improving both individual and system-level performance of future transportation systems.

4 DISCUSSION AND CONCLUSIONS

In this work, we demonstrated that introducing socially aware MARL-based AV agents into future traffic networks can help relief the convergence challenges that arise when AVs act selfishly. Specifically, we complement the selfish travel-time-based reward with a counterfactual, equilibrium-preserving intrinsic reward based on the marginal cost that each AV imposes on the other drivers of the system. We demonstrate that our novel reward formulation not only improves the overall system performance but, surprisingly, is beneficial for the individual AV agents. Additionally, we show that the proposed reward formulation preserves the system’s equilibria while accelerating convergence toward the system-optimal solution.

In our experimental evaluation, we compare AV agents trained using MARL algorithms and optimizing different objectives: our proposed reward formulation versus the selfish travel-time based reward. On the simple TRY network, we demonstrate that socially aware AVs converge faster to the optimal solution, which in this case coincides with the individually optimal solution. In the real-world Saint-Arnoult network, after a short training period, more than 50% of the socially aware AVs achieve shorter individual travel times, and the overall system travel times are also improved.

A limitation of this work includes the computational cost of calculating the marginal cost matrices, as they will grow exponentially with size of the joint action space. Moreover, we deliberately focused on the rare case where SO and UE coincide. Future studies shall better exploit how including the social component may be both effective in reducing Price of Anarchy and remaining attractive and fair for AV users.

This work can contribute towards demonstrating the importance of introducing some form of socially-aware collaboration between AV agents in future traffic systems. Understanding the potential impact of introducing AVs into future cities is of paramount importance, particularly when it comes to identifying potential negative effects and examining whether some form of collaboration between the vehicles could help maintain overall system stability, equity and performance.

REFERENCES

- [1] Mohamed Abdel-Aty and Shengxuan Ding. 2024. A matched case-control analysis of autonomous vs human-driven vehicle accidents. *Nature Communications* 15 (2024), 4931. <https://doi.org/10.1038/s41467-024-48526-4>
- [2] Ahmet Onur Akman. 2025. JanuX. <https://github.com/COEXISTENCE-PROJECT/JanuX>
- [3] Ahmet Onur Akman, Anastasia Psarou, Michał Hoffmann, Łukasz Gorczyca, Łukasz Kowalski, Paweł Gora, Grzegorz Jamróz, and Rafał Kucharski. 2025. URB – Urban Routing Benchmark for RL-equipped Connected Autonomous Vehicles. In *Advances in Neural Information Processing Systems*.
- [4] Ahmet Onur Akman, Anastasia Psarou, Zoltán György Varga, Grzegorz Jamróz, and Rafał Kucharski. 2024. Impact of Collective Behaviors of Autonomous Vehicles on Urban Traffic Dynamics: A Multi-Agent Reinforcement Learning Approach. In *Seventeenth European Workshop on Reinforcement Learning*. <https://openreview.net/forum?id=88zP8xh5D2>
- [5] Ahmet Onur Akman, Anastasia Psarou, Łukasz Gorczyca, Zoltán György Varga, Grzegorz Jamróz, and Rafał Kucharski. 2025. RouteRL: Multi-agent reinforcement learning framework for urban route choice with autonomous vehicles. *SoftwareX* 31 (2025), 102279. <https://doi.org/10.1016/j.softx.2025.102279>
- [6] Stefano V. Albrecht, Filippos Christianos, and Lukas Schäfer. 2024. *Multi-Agent Reinforcement Learning: Foundations and Modern Approaches*. MIT Press. <https://www.marl-book.com>
- [7] W. Brian Arthur. 1994. Inductive Reasoning and Bounded Rationality. *The American Economic Review* 84, 2 (1994), 406–411. <http://www.jstor.org/stable/2117868>
- [8] Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. Finite-time Analysis of the Multiarmed Bandit Problem. *Machine Learning* 47, 2-3 (2002), 235–256. <https://doi.org/10.1023/A:1013689704352>
- [9] Ennio Cascetta. 2009. *Transportation Systems Analysis: Models and Applications*. Springer Optimization and Its Applications, Vol. 29. Springer. <https://doi.org/10.1007/978-0-387-75857-2>
- [10] Nuttapon Chentanez, Andrew Barto, and Satinder Singh. 2004. Intrinsically Motivated Reinforcement Learning. In *Advances in Neural Information Processing Systems*, L. Saul, Y. Weiss, and L. Bottou (Eds.), Vol. 17. MIT Press. https://proceedings.neurips.cc/paper_files/paper/2004/file/4be5a36baca8ab9d2066debf4e65c1-Paper.pdf
- [11] Richard Cole, Yevgeniy Dodis, and Tim Roughgarden. 2006. How much can taxes help selfish routing? *J. Comput. System Sci.* 72, 3 (2006), 444–467. <https://doi.org/10.1016/j.jcss.2005.09.010> Network Algorithms 2005.
- [12] Sam Devlin and Daniel Kudenko. 2011. Theoretical considerations of potential-based reward shaping for multi-agent systems. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1 (Taipei, Taiwan) (AAMAS '11)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 225–232.
- [13] Daniel J. Fagnant and Kara Kockelman. 2015. Preparing a nation for autonomous vehicles: opportunities, barriers and policy recommendations. *Transportation Research Part A: Policy and Practice* 77 (2015), 167–181. <https://doi.org/10.1016/j.tra.2015.04.003>
- [14] Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. 2024. Counterfactual Multi-Agent Policy Gradients. arXiv:1705.08926 [cs.AI] <https://arxiv.org/abs/1705.08926>
- [15] Christian Gawron. 1998. *Simulation-Based Traffic Assignment: Computing User Equilibria in Large Street Networks*. Ph.D. Dissertation. Universität zu Köln. <https://kups.uni-koeln.de/OpenAccess>.
- [16] Rafael Grunitzki, Guilherme de Oliveira Ramos, and Ana L C Bazzan. 2014. Individual versus difference rewards on reinforcement learning for route choice. In *2014 Brazilian Conference on Intelligent Systems*. IEEE, 179–184.
- [17] Marek Grześ. 2017. Reward Shaping in Episodic Reinforcement Learning. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems (São Paulo, Brazil) (AAMAS '17)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 565–573.
- [18] Michał Hoffmann, Michał Bujak, Grzegorz Jamróz, and Rafał Kucharski. 2025. Wardropian Cycles make traffic assignment both optimal and fair by eliminating price-of-anarchy with Cyclical User Equilibrium for compliant connected autonomous vehicles. arXiv:2507.19675 [eess.SY] <https://arxiv.org/abs/2507.19675>
- [19] Ron Holzman and Nissan Law-Yone. 1997. Strong equilibrium in congestion games. *Games and economic behavior* 21, 1-2 (1997), 85–101.
- [20] Grzegorz Jamróz, Ahmet Onur Akman, Anastasia Psarou, Zoltán György Varga, and Rafał Kucharski. 2025. Social implications of coexistence of CAVs and human drivers in the context of route choice. *Scientific Reports* 15, 1 (2025), 6768.
- [21] Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Caglar Gulcehre, Pedro A. Ortega, DJ Strouse, Joel Z. Leibo, and Nando de Freitas. 2019. Social Influence as Intrinsic Motivation for Multi-Agent Deep Reinforcement Learning. arXiv:1810.08647 [cs.LG] <https://arxiv.org/abs/1810.08647>
- [22] Daniel A. Lazar, Erdem Biyik, Dorsa Sadigh, and Ramtin Pedarsani. 2021. Learning How to Dynamically Route Autonomous Vehicles on Shared Roads. arXiv:1909.03664 [math.OC] <https://arxiv.org/abs/1909.03664>
- [23] Pingfei Li, Xinyu Zhu, Yao Ren, Zhengping Tan, Wenhao Hu, You Zhang, and Cheng Xu. 2024. Generalization of cut-in pre-crash scenarios for autonomous vehicles based on accident data. *Scientific Reports* 14 (2024), 17664. <https://doi.org/10.1038/s41598-024-68263-4>
- [24] Pablo Alvarez Lopez, Michael Behrisch, Laura Bieker-Walz, Jakob Erdmann, Yun-Pang Flötteröd, Robert Hilbrich, Leonhard Lüken, 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/>
- [25] Deepak K Merchant and George L Nemhauser. 1978. Optimality conditions for a dynamic traffic assignment model. *Transportation Science* 12, 3 (1978), 200–207.
- [26] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. 2013. Playing Atari with Deep Reinforcement Learning. arXiv:1312.5602 [cs.LG] <https://arxiv.org/abs/1312.5602>
- [27] Nuri C. Onat, Jafar Mandouri, Murat Kucukvar, Burak Sen, Saddam A. Abbasi, Wael Alhajyaseen, Adeeb A. Kutty, Rateb Jabbar, Marcello Contestabile, and Abdel Magid Hamouda. 2023. Rebound effects undermine carbon footprint reduction potential of autonomous electric vehicles. *Nature Communications* 14 (2023), 6258. <https://doi.org/10.1038/s41467-023-41992-2>
- [28] Georgios Papoudakis, Filippos Christianos, Lukas Schäfer, and Stefano V. Albrecht. 2021. Benchmarking Multi-Agent Deep Reinforcement Learning Algorithms in Cooperative Tasks. arXiv:2006.07869 [cs.LG] <https://arxiv.org/abs/2006.07869>
- [29] Srinivas Peeta and Hani S. Mahmassani. 1995. System optimal and user equilibrium time-dependent traffic assignment in congested networks. *Annals of Operations Research* 60 (1995), 81–113. <https://doi.org/10.1007/BF02031941>
- [30] Anastasia Psarou, Ahmet Onur Akman, Łukasz Gorczyca, Michał Hoffmann, Grzegorz Jamróz, and Rafał Kucharski. 2025. Collaboration Between the City and Machine Learning Community is Crucial to Efficient Autonomous Vehicles Routing. arXiv:2502.13188 [cs.MA] <https://arxiv.org/abs/2502.13188>
- [31] Gabriel de O. Ramos, Roxana Rădulescu, Ann Nowé, and Anderson R. Tavares. 2020. Toll-Based Learning for Minimising Congestion under Heterogeneous Preferences. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (Auckland, New Zealand) (AAMAS '20)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1098–1106.
- [32] Jette Randløv and Preben Alstrøm. 1998. Learning to Drive a Bicycle Using Reinforcement Learning and Shaping. In *Proceedings of the Fifteenth International Conference on Machine Learning (ICML 1998)*, Jude W. Shavlik (Ed.). Morgan Kaufmann, San Francisco, CA, USA, 463–471. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.52.3038&rep=rep1&type=pdf>
- [33] Guni Sharon, Stephen D. Boyles, Benjamin Alkoby, and Peter Stone. 2019. Marginal Cost Pricing with a Fixed Error Factor in Traffic Networks. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (Montreal QC, Canada) (AAMAS '19)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1539–1546.
- [34] Yosef Sheffi. 1985. *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods*. Prentice-Hall, Englewood Cliffs.
- [35] Raphael E. Stern, Shumo Cui, Maria Laura Delle Monache, Rahul Bhadani, Matt Bunting, Miles Churchill, Nathaniel Hamilton, R'mani Haulcy, Hannah Pohlmann, Fangyu Wu, Benedetto Piccoli, Benjamin Seibold, Jonathan Sprinkle, and Daniel B. Work. 2018. Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments. *Transportation Research Part C: Emerging Technologies* 89 (2018), 205–221. <https://doi.org/10.1016/j.trc.2018.02.005>
- [36] Ming Tan. 1997. *Multi-agent reinforcement learning: independent vs. cooperative agents*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 487–494.
- [37] J. K. Terry, Benjamin Black, Nathaniel Grammel, Mario Jayakumar, Ananth Hari, Ryan Sullivan, Luis Santos, Rodrigo Perez, Caroline Horsch, Clemens Dieffendahl, Niall L. Williams, Yashas Lokesh, and Praveen Ravi. 2021. PettingZoo: Gym for Multi-Agent Reinforcement Learning. arXiv:2009.14471 [cs.LG] <https://arxiv.org/abs/2009.14471>
- [38] Luiz A. Thomasini, Lucas N. Alegre, Gabriel O. Ramos, and Ana L. C. Bazzan. 2023. RouteChoiceEnv: a Route Choice Library for Multiagent Reinforcement Learning. In *Adaptive and Learning Agents Workshop at AAMAS*.
- [39] Kagan Tumer and Adrian Agogino. 2006. Agent Reward Shaping for Alleviating Traffic Congestion. (01 2006).
- [40] William S. Vickrey. 1969. Congestion Theory and Transport Investment. *The American Economic Review* 59, 2 (1969), 251–260. <http://www.jstor.org/stable/1823678>
- [41] J G WARDROP. 1952. ROAD PAPER. SOME THEORETICAL ASPECTS OF ROAD TRAFFIC RESEARCH. *Proceedings of the Institution of Civil Engineers* 1, 3 (May 1952), 325–362. <https://doi.org/10.1680/ipeds.1952.11259>
- [42] J. G. Wardrop. 1952. ROAD PAPER. SOME THEORETICAL ASPECTS OF ROAD TRAFFIC RESEARCH. <https://api.semanticscholar.org/CorpusID:131127018>
- [43] Yichuan Yang and Zongyu Song. 2025. The public perception and adaptability of laws and regulations of autonomous driving vehicles. *Humanities and Social Sciences Communications* 12 (2025), 1224. <https://doi.org/10.1057/s41599-025-05600-6>

- [44] Chao Yu, Akash Velu, Eugene Vinitsky, Yu Wang, Alexandre M. Bayen, and Yi Wu. 2021. The Surprising Effectiveness of MAPPO in Cooperative, Multi-Agent Games. *CoRR* abs/2103.01955 (2021). arXiv:2103.01955 <https://arxiv.org/abs/2103.01955>
- [45] Qiyuan Zhang, Christopher D. Wallbridge, Dylan M. Jones, and Phillip L. Morgan. 2024. Public perception of autonomous vehicle capability determines judgment of blame and trust in road traffic accidents. *Transportation Research Part A: Policy and Practice* 179 (2024), 103887. <https://doi.org/10.1016/j.tra.2023.103887>
- [46] Shi-Teng Zheng, Rui Jiang, Xiqun Chen, Junfang Tian, Xiao Han, Ruidong Yan, Bin Jia, Xiaobo Qu, Zhen-Hua Li, Lan-Da Gao, Fang Zhang, De-Zhao Zhang, and Ziyou Gao. 2025. Communicating with two vehicles immediately ahead boosts traffic capacity sixfold in connected autonomous vehicle platoons. *Communications Engineering* 4 (2025), 160. <https://doi.org/10.1038/s44172-025-00500-8>