

Learning to Influence Vehicles' Routing in Mixed-Autonomy Networks by Dynamically Controlling the Headway of Autonomous Cars

Xiaoyu Ma¹ and Negar Mehr²

Abstract—It is known that autonomous cars can increase road capacities by maintaining a smaller headway through vehicle platooning. Recent works have shown that these capacity increases can influence vehicles' route choices in unexpected ways similar to the well-known Braess's paradox, such that the network congestion might increase. In this paper, we propose that in mixed-autonomy networks, i.e., networks where roads are shared between human-driven and autonomous cars, the headway of autonomous cars can be directly controlled to influence vehicles' routing and reduce congestion. We argue that the headway of autonomous cars — and consequently the capacity of link segments — is not just a fixed design choice; but rather, it can be leveraged as an infrastructure control strategy to dynamically regulate capacities. Imagine that similar to variable speed limits which regulate the maximum speed of vehicles on a road segment, a control policy regulates the headway of autonomous cars along each road segment. We seek to influence vehicles' route choices by directly controlling the headway of autonomous cars to prevent Braess-like unexpected outcomes and increase network efficiency. We model the dynamics of mixed-autonomy traffic networks while accounting for the vehicles' route choice dynamics. We train an RL policy that learns to regulate the headway of autonomous cars such that the total travel time in the network is minimized. We will show empirically that our trained policy can not only prevent Braess-like inefficiencies but also decrease total travel time¹.

I. INTRODUCTION

Road networks will soon be shared between human-driven and autonomous cars, i.e., they will operate under mixed vehicle autonomy. Unlike human-driven cars which cannot be externally controlled, we can assume a level of control over autonomous cars, which provides us with an additional control input to affect traffic networks. Several recent works have considered controlling autonomous cars' actions at a vehicle level for improving traffic conditions [1], [2], [3], [4], [5]. In this work, we focus on the system-level performance of mixed-autonomy networks. A particularly attractive feature of autonomous cars for influencing system-level mobility is that they can facilitate vehicle platooning. Vehicle platoons are groups of vehicles that can maintain a shorter headway using adaptive cruise control technologies. As a result, forming vehicle platoons can result in increased capacity of road segments [6]. Such capacity increases can be up to threefold if all vehicles are autonomous [6].

*This work is supported by the National Science Foundation, under CAREER Award ECCS-2145134.

¹Xiaoyu Ma is with the Department of Electrical and Computer Engineering, UIUC, 306 N Wright St, Urbana, IL 61801, USA xiaoyum2@illinois.edu

²Negar Mehr is with the Department of Aerospace Engineering, UIUC, 104 S Wright St, Urbana, IL 61801, USA negar@illinois.edu

¹The code is available at:

<https://github.com/labicon/RL-Traffic-Dynamics>




Fig. 1: We propose that we can dynamically control the headway of autonomous cars β_1^a and β_2^a on road segments to influence the vehicles' routing and reduce congestion. We assume that the headway of human-driven cars (β_1^h and β_2^h) remains uncontrolled and constant.

While such capacity increases have the potential to increase traffic throughput and reduce congestion, it was shown in recent works that if all vehicles select their routes *selfishly*, the capacity increases of autonomous cars might in fact *worsen* congestion [7], [8]. It is well known that traffic networks tend to operate at equilibria, where vehicular flows are routed along the network paths such that no vehicle can gain any savings in its travel time by unilaterally changing its route [9]. Under such selfish route choices, in our previous work [7], [8], we showed that the capacity increases of autonomous cars can influence vehicles' route choices such that the overall network delay may increase. This unexpected behavior is similar to the well-known Braess's paradox [10] where it was shown that increasing a link capacity by adding lanes to the link can increase the overall network delay.

Several recent works have focused on control schemes that alleviate the inefficiencies that result from selfish route choices of vehicles in mixed-autonomy networks. For example, [11], [12], [13] have considered tolling mixed-autonomy traffic networks such that network inefficiencies are minimized. In [14], [15], [16], the altruistic routing of autonomous cars was studied for controlling mixed-autonomy networks, where the routing of autonomous cars was controlled such that when human-driven cars react selfishly to the routing of autonomous cars, the overall network performance is improved. In all these works, the assumption is that autonomous cars maintain platoons of cars with a pre-specified headway, and control mechanisms are sought to alleviate the inefficiencies that arise from selfish routing.

In this paper, we take a different approach and propose that in mixed-autonomy networks, the capacity increases of autonomous cars can be *directly regulated* to reduce network inefficiencies and improve performance. We argue

that the headway of autonomous cars— and consequently the capacity increases of autonomous cars— is not just a fixed design choice; but rather, it can be leveraged as an *infrastructure control* strategy to *dynamically* regulate road capacities (see Fig. 1). Similar to variable speed limits, which regulate the maximum speed of vehicles on a road segment, a control policy can regulate the headway of autonomous cars along each road segment. This is in fact closely related to variable speed limits which have been widely studied for networks of only human-driven cars, where link capacities are regulated by dynamically regulating the maximum speed along network links [17], [18], [19], [20]. In this work, we seek to influence vehicles’ route choices by directly controlling the headway of autonomous cars to prevent Braess-like unexpected outcomes and increase network efficiency.

To investigate the impact of our proposed controller, we model the dynamics of mixed-autonomy traffic networks while accounting for the vehicles’ route choice dynamics. Since traffic dynamics and vehicles’ choice dynamics are highly nonlinear, finding an appropriate controller is extremely challenging. To tackle this, we propose to train an RL policy that learns to regulate the headway of autonomous cars such that the total travel time in the network is minimized. We will show through empirical analysis that not only can our trained policy prevent Braess-like inefficiencies but also it can decrease total travel time and improve network performance. Our evaluations reveal that dynamically regulating the headway of autonomous cars has the potential to reduce the inefficiencies of selfish routing in mixed-autonomy traffic networks. In summary, our main contributions are:

- We propose to dynamically regulate the headway of autonomous cars to influence vehicles’ route choices.
- We use traffic dynamics and vehicles’ route choice dynamics to train an RL policy for dynamically controlling the headway of autonomous cars.
- We demonstrate empirically through our experiments that such an RL policy can reduce the inefficiencies that are inherent in selfish route choices and can improve the performance of mixed-autonomy traffic networks.

II. NETWORK MODEL

We model the traffic network as a directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{L})$ where \mathcal{N} is the set of nodes and \mathcal{L} is the set of links. Each link $l \in \mathcal{L}$ in the network is a road segment which is a directed edge from a node $u \in \mathcal{N}$ to a node $v \in \mathcal{N}$. We assume that we are given a set of origin-destination (O/D) pairs denoted by \mathcal{W} , which contains k O/D pairs. Each O/D pair $w \in \mathcal{W}$ is denoted as (o_w, d_w) , where $o_w, d_w \in \mathcal{N}$ ($o_w \neq d_w$) represent the origin and destination nodes of O/D pair w respectively. We define the vector of network demand to be $r(t) = (r_w(t))_{w \in \mathcal{W}}$ where $r_w(t)$ is the potentially time-varying demand profile of OD pair w . Note that $r_w(t)$ specifies the total amount of vehicular flow (including both human-driven and autonomous vehicles) that needs to be routed along O/D pair w at time t .

For each O/D pair $w \in \mathcal{W}$, we assume that there exists at least one path from its origin to its destination. We




Fig. 2: The fundamental diagram of mixed-autonomous traffic for varying autonomy fractions α_l along link l . As the fraction of autonomous vehicles $\alpha_l(t)$ increases, the link capacity $C_l(\alpha_l)$ also increases while the jam density \bar{n}_l is constant.

use \mathcal{P}_w to denote the set of all paths that connect O/D pair w . Additionally, let $\mathcal{P} = \cup_{w \in \mathcal{W}} \mathcal{P}_w$ denote the set of all network paths. For each $p \in \mathcal{P}$, we denote the time-dependent flow of autonomous and human-driven cars along p by $f_p^a(t)$ and $f_p^h(t)$, respectively. We further define the vectors of autonomous and human-driven flows as $f(t) := (f_p^h(t), f_p^a(t))_{p \in \mathcal{P}}$, where $f^a(t) = (f_p^a(t))_{p \in \mathcal{P}}$ and $f^h(t) = (f_p^h(t))_{p \in \mathcal{P}}$. We define the total flow on a link $l \in \mathcal{L}$ at time t as $f_l(t) = f_l^a(t) + f_l^h(t)$. We also define the autonomy fraction along a link $l \in \mathcal{L}$ to be the fraction of cars along the link that are autonomous, i.e. $\alpha_l(t) := \frac{f_l^a(t)}{f_l^a(t) + f_l^h(t)}$.

A. Capacity Model

We assume that the flow along each link is governed by the fundamental diagram of the link [21], i.e. the maximum flow along a link is the link’s capacity, and the maximum density of the vehicles along a link is the jam density. Unlike networks with only human-driven cars where the link capacity is constant, when autonomous cars with variable headways are deployed, a link capacity changes as a function of both the headway of autonomous cars and the autonomy fraction along the link. Let the headway distance maintained by autonomous cars and human-driven cars on link $l \in \mathcal{L}$ at time t be $\beta_l^a(t)$ and $\beta_l^h(t)$ respectively. In the following, we discuss how the capacity changes of a link can be modeled.

We let $\tilde{n}_l(\alpha_l(t))$ denote the critical density of link l at time t . Critical density determines when the traffic on the link changes from free flow to congestion. Similar to link capacity, the critical density $\tilde{n}_l(\alpha_l(t))$ is a function of the autonomy fraction along the link α_l . The critical density is the maximum density of vehicles that can travel on link l in free flow (see Fig. 2). It was shown in [15], [13] that in mixed-autonomy networks, the critical density of a link is equal to

$$\tilde{n}_l(\alpha_l(t)) = \frac{b_l}{\alpha_l(t) \cdot \beta_l^a(t) + (1 - \alpha_l(t)) \cdot \beta_l^h(t)}, \quad (1)$$

where b_l is the number of lanes on link l . The weighted sum of headways in the denominator of Equation (1) represents

the average headway kept between one car and its preceding car when all vehicles are running in free flow. Let the capacity of a link l at time t be denoted by $C_l(\alpha_l(t))$. It can be then shown that the capacity of link l is $C_l(\alpha_l(t)) = v_l \cdot \tilde{n}_l(\alpha_l(t))$, where v_l is the free-flow velocity on link l . The jam density, denoted as \tilde{n} , is the density on a link when the link gets totally jammed and the speed of vehicles running in the link reduces to zero, and is assumed to be a constant.

The density of a link l at time t is defined as $n_l(t) = n_l^a(t) + n_l^h(t)$, where $n_l^a(t)$ and $n_l^h(t)$ are the density of autonomous and human-driven cars respectively. Following flow conservation, for any link $l \in \mathcal{L}$, at every time step t , we have:

$$n_l(t+1) = n_l(t) + f_{l,in}(t) - f_{l,out}(t). \quad (2)$$

where $f_{l,in}(t)$ and $f_{l,out}(t)$ are the input and output flows of link l at time t .

Following the fundametnal diagram, the flow of each link $l \in \mathcal{L}$ at time t can be calculated as

$$f_l(n_l(t)) = \begin{cases} v_l \cdot n_l(t), & \text{if } n_l(t) \leq \tilde{n}_l(\alpha_l(t)) \\ \frac{v_l \cdot \tilde{n}_l(\alpha_l(t)) \cdot (\tilde{n}_l - n_l(t))}{\tilde{n}_l - \tilde{n}_l(\alpha_l(t))}, & \text{if } \tilde{n}_l(\alpha_l(t)) \leq n_l(t) \leq \tilde{n}_l \\ 0, & \text{else} \end{cases} \quad (3)$$

Equation (3) indicates that when the density $n_l(t)$ is smaller than the critical density, vehicles move at the free-flow speed. When $n_l(t)$ is larger than the critical density, the vehicles are in the congested regime as shown in Fig. 2.

B. Latency Function

For each link l in \mathcal{L} , we define a binary parameter $s_l(t)$ indicating whether the link is in free flow at time step t ($s_l(t) = 0$) or congested ($s_l(t) = 1$). Then, using the results from [13], the latency along a link $l \in \mathcal{L}$ can be found by:

$$e_l(f_l(t), s_l) = \begin{cases} \frac{d_l}{v_l}, & \text{if } s_l(t) = 0 \\ d_l \left(\frac{\tilde{n}_l}{f_l(t)} + \frac{\tilde{n}_l(\alpha_l(t)) - \tilde{n}_l}{v_l \tilde{n}_l(\alpha_l(t))} \right), & \text{if } s_l(t) = 1 \end{cases} \quad (4)$$

where d_l is the length of link l . Eq. (4) captures the time that it takes for one vehicle to travel along the link. Consequently, for each path $p \in \mathcal{P}$, the path latency e_p is the sum of link latency along the links that constitute path p :

$$e_p = \sum_{l' \in p} e_{l'}(f_{l'}(t)) \quad (5)$$

C. Route Choice Dynamics

To capture the impact of the headway of autonomous cars on the vehicles' route choices, we need to model the choice dynamics of vehicles in the network. We model this through evolutionary dynamics introduced in [15] to model how a population of vehicles along a path $p \in \mathcal{P}$ choose their route. More specifically, for a path p connecting O/D pair w , we have:

$$f_p^h(t+1) = \frac{f_p^h(t) \exp(-\mu^h e_p(t))}{\sum_{p' \in \mathcal{P}_w} f_{p'}^h(t) \exp(-\mu^h e_{p'}(t))}, \quad (6)$$

$$f_p^a(t+1) = \frac{f_p^a(t) \exp(-\mu^a e_p(t))}{\sum_{p' \in \mathcal{P}_w} f_{p'}^a(t) \exp(-\mu^a e_{p'}(t))}, \quad (7)$$

where μ^h and μ^a are the rationality factors for human-driven and autonomous vehicles that capture the selfishness of vehicles when selecting their route. In (6) and (7), the fraction of vehicles running on path p at time step $t+1$ is considered to be inversely proportional to the exponential of the delay experienced by users of that road. The parameters μ^h and μ^a were introduced in previous work as learning rates in the context of humans' routing choices and simulate a congestion game [22]. Intuitively, when the value of μ^h or μ^a is large, it is more likely for the group of corresponding vehicles to actively change their route due to a change in the latency of their previously selected route.

III. LEARNING TO ADJUST THE HEADWAY OF AUTONOMOUS CARS

Equipped with a model that can simulate the impact of autonomous cars' headway on the route choice of vehicles, we can now train a policy that learns to control the headway of autonomous cars such that when all cars select their routes greedily, the overall performance of the network is improved. We discuss how we train a policy to achieve this goal.

A. Reward function

An intuitive reward function for training the policy is the negative of the Total Travel Time (TTT) of the network. The network TTT is defined as the sum of vehicle densities across all network links over a finite horizon of time T .

$$r(T) = - \sum_{t=1}^T \sum_{l \in \mathcal{L}} n_l(t). \quad (8)$$

Since we are trying to alleviate congestion, maximizing (8) can be viewed as minimizing the total densities of all vehicles in the network. This choice is motivated by the common use of TTT in the transportation literature to capture network performance (see [23] for a list of references).

B. Policy Training

We would like to find a policy π that outputs the headway of autonomous cars along each link β_l^a such that the reward (8) is maximized when vehicles choose their routes according to (6) and (7). We assume that we cannot control the headway of human-driven cars, and we have control only over the headway of autonomous cars along each link β_l^a . However, this is an extremely challenging control problem since the traffic and route choice dynamics of the vehicles are highly nonlinear and complex. To tackle this challenge, we propose to train a policy that learns to adjust the autonomous cars' headway for minimizing TTT. We train our policy in a model-free setup where we leverage our traffic dynamics model to develop a traffic simulator capable of capturing the dependence of vehicles' route choices on vehicles' headway. We used this as a simulation platform for training our RL agent that learns to adjust the headway of autonomous cars over a fixed time horizon T .




Fig. 3: The geometry of the classic Braess network (a) and the extended 8-link network with embedded Braess geometry (b).

The problem can be incorporated as a finite-horizon discounted Markov decision process (MDP), defined by the tuple $(\mathcal{S}, \mathcal{A}, P, c, \rho_0, \gamma)$, where \mathcal{S} is the set of states, \mathcal{A} is the set of actions, $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is the transition probability distribution, $r : \mathcal{S} \rightarrow \mathbb{R}$ is the reward function, $\rho_0 : \mathcal{S} \rightarrow \mathbb{R}$ is the distribution of the initial state s_0 which is equal to the initial density of the vehicles, and $\gamma \in (0, 1)$ is the discount factor. The observation space is the same as the state space $S \in \mathcal{S}$, which includes the density of vehicles on the road segments, while the action space $A \in \mathcal{A}$ is the headway of autonomous cars along the links at each time step $(\beta_l^\alpha(t))_{l \in \mathcal{L}}$. The transition function is determined by the system dynamics (2), (6), and (7).

Our implemented simulator can be used in conjunction with state-of-the-art RL algorithms. We selected Proximal Policy Optimization (PPO) [24] to train our RL agent. Since our state and action spaces are continuous, we chose PPO which is a policy-gradient-based method. PPO dynamically adjusts the policy π according to the gradient of the reward with respect to the policy parameters. The clipping mechanism in PPO can protect the policy from over-updating itself with respect to large and steep gradients.

IV. EXPERIMENT

A. Experiment Design

In this section, we introduce the details of our experimental setup. We will start with a Braess network geometry shown in Fig. 3 (a). We will then consider a larger network with 8 links that has the Braess geometry embedded in it as shown in Fig. 3 (b). Note that the Braess geometry is a very complicated network geometry that was shown in our previous work that can result in counter-intuitive vehicles' route choices [8]. For example, in [7], it was shown that in a Braess network geometry, the capacity increases that result from vehicle platooning on the middle link (Link 4) can indeed increase congestion levels when vehicles select their routes selfishly. We chose the Braess network to verify if our control policy can prevent such unexpected behaviors and even more, improve the performance.

We considered a network where Link 0, 1, 2, and 3 are designed to be 240 kilometers long while the middle link is 60 kilometers long. The links are designed to have a free-flow velocity of 30 meters per second. Each time step lasts for one minute. We assume that the headway of autonomous




Fig. 4: The time-dependent demand injected at node O for both the classic Braess network and the 8-link extended Braess network.

cars can be changed every 10 minutes, and the total length of the time sequence is 200 minutes. We chose each time step to last for 10 minutes accordingly. Although higher frequency headway changes are feasible, we wanted to see if a significant improvement in traffic performance can be sought even with less frequent updates on headways. We set the number of lanes on Link 0 and 3 to be two times larger than those on Link 1 and 2. The number of lanes on Link 4 is designed to be especially large so that we can better capture the impact of capacity increases. For the initial conditions, we put more vehicles onto Links 0 and 2 at the beginning of the time horizon so that it is more crowded on the top route of the network. For the 8-link Braess network, we keep Links 0 and 4 to be the same as the classic Braess network, while Link 5 and 7 are designed to have the same characteristics as Links 2 and 4.

We considered a time-varying vehicular demand entering the network through node O . We chose the demand profile such that it resembles the peak-hour demand profiles where the demand increases to a peak demand rapidly and then decreases (see Fig. 4). Such a choice allows us to verify our controller under realistic demand profiles which may vary significantly over time. We also consider a cool-down period in our demand profile similar to [23] to evaluate the capability of our control policy in discharging vehicles that remained on the network. The simulator will reset itself to initial values and start a new training sequence after 20 consecutive episodes. The autonomous cars' headway is constrained to remain between 1 meter to 10 meters.

B. Training Details

Our experiments were carried out on Google Colab. The simulated environment is scripted in a Python environment connected with the baseline RL agent available in the Stable-Baselines library [25]. We chose a learning rate of 2×10^{-4} , and set the number of steps to run for each environment per update to be 2048, the batch size to be 64, and the clipping range to be 0.2. Training the RL agent on the 5-link classic Braess Network typically takes at least 4 million time steps to reach convergence, and it will take approximately 5 hours of training on CPU. The 8-link extended Braess Network




Fig. 5: Heat maps of vehicle densities that result from two policies on the Braess network (a) and extended 8-link Braess network (b). The top rows are the heat maps of the network in which autonomous cars use the same constant headway coefficient as human-driven cars at all times. In contrast, the bottom rows are heat maps of the networks in which the RL agent regulates the headway of autonomous cars. By changing the headway of autonomous cars dynamically, at least one route in the network remains in free flow, i.e. congestion is not propagated throughout the entire network.

takes around 10 million time steps to converge, and it will take approximately 9 hours to finish. Note that the training time depends on the hardware available. The CPU that we used while training our policy on Google Colab was AMD EPYC 7B12.

V. RESULTS AND ANALYSIS

In this section, we report the performance of our trained policy. We compare the network TTT under our trained policy as opposed to two baselines which treat the headway of autonomous cars as constant. 1) Uniform constant headway: the baseline where the headways of both human-driven and autonomous cars are constant and equal. 2) Minimum Headway: the baseline where the headway of autonomous cars is set to be the minimum possible feasible headway. The Minimum Headway baseline was selected to investigate how the minimum headway and consequently maximum capacity increase of autonomous cars will affect vehicles' route choices. In comparison against these baselines, we also discuss some of the important factors that can affect the system's performance including the fraction of autonomous cars in the demand and the vehicles' rationality level captured in vehicles' rationality factors μ^a and μ^h .

A. Performance Comparison

We apply our policy to the 5-link network for a horizon of 200 minutes. We measure TTT and compare it with TTT over the same time horizon under our baseline methods. Table I shows the percentage improvements in performance in terms of TTT when $\mu = 0.1$ and the fraction of autonomous cars in the overall demand is equal to $\alpha_{O-D} = 0.8$. As Table I demonstrates, not only can our policy avoid the unexpected increases in TTT due to the capacity increases of autonomous cars when using minimum headway, but also it improves the performance compared to the case when no vehicles' headway is controlled and are all equal to the headway of human-driven cars. Using the headway coefficients produced by our trained policy, the TTT is 4.0% better than the network using uniform constant headway and is 11.1% better than the network using minimum headway. This indicates that as demonstrated in our prior work [7], simply requiring

autonomous cars to maintain platoons of vehicles is not necessarily going to improve network performance, it may actually worsen the performance when vehicles select their routes greedily as too many vehicles may be incentivized to change their routes in response to the capacity increases of autonomous cars. We ran a similar experiment with the 8-link network and found that our trained policy resulted in 2.8% of improvement in TTT compared to the policy using uniform constant headway, and 0.2% better than the policy using minimum headway. The percentage improvement is smaller in the 8-link network mainly due to the insufficient discharge of vehicles, as the time required for vehicles to exit the network is inherently longer than the classic Braess network.

	Improvement w.r.t. Baseline 1	Improvement w.r.t. Baseline 2
Braess Network	4.03%	11.16%
8-link Network	2.80%	0.20%

TABLE I: The performance improvement of our policy compared to two constant headway policies: 1) uniform constant headway and 2) constant minimum headway.

To further show the difference in the system trajectories, Fig. 5 provides a comparison of the vehicles' distribution in the form of a heat map where a relatively congested link is painted in red, while a link in free flow is painted in green. We can see from the top row of these figures that with no headway control, both Link 0 and Link 1, end up congested. However, by changing the headway of autonomous cars dynamically (bottom row), at least one route in the network remains in free flow, i.e., the congestion is not propagated throughout the entire network. This allows the vehicles to have at least one route available to exit the network, thus improving the performance in terms of TTT.

In our experiments, we found that two critical factors play essential roles in the effectiveness of the trained policy. One is the vehicles' rationality parameter μ , representing the level of vehicles' selfishness in their route choice dynamics. The other is the total fraction of autonomous cars in the network. In the following, we will discuss the impact of these parameters.

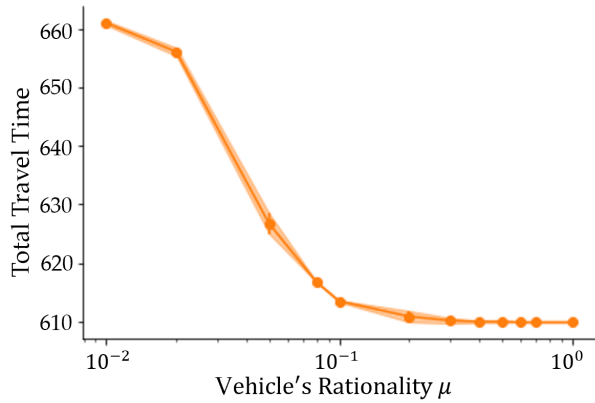


Fig. 6: Plots of TTT as a function of the vehicles’ rationality. We assume that the rationality coefficient of human-driven and autonomous cars are equal, i.e. $\mu^a = \mu^h = \mu$. The improvement in TTT is more significant with larger μ . Note that the x-axis is plotted in a log scale.

B. Vehicles’ Rationality

The rationality parameters μ^a and μ^h are coefficients that appear in Eq. (6) and (7), and they affect the portion of vehicles that change their route due to a change in the latency of their selected routes. Generally speaking, the larger μ^h and μ^a are, the more the vehicles will be likely to change their path to take a path with lower latency. Consequently, μ^a and μ^h will have a significant impact on how much one can influence vehicle’ routing by controlling the headway of autonomous cars.

To study this impact, we measured TTT under our trained policy for different levels of vehicles’ rationality. We tried 5 random seeds for each rationality parameter, and the experiment is carried out on the classic Braess network. The resulting TTT for various rationality parameters is shown in Fig. 6. As the figure demonstrates, as the vehicles’ rationality increases, TTT decreases, i.e. the impact of autonomous vehicles’ headway on the route choices of vehicles become more significant. Interestingly, the profile of links’ headway generated by the RL agent differs for different levels of rationality. For example, the RL agent is more likely to generate smaller headways when using smaller μ^a and μ^h , i.e. less selfish drivers, so that the critical capacity of the links can be increased more dramatically, which will help the policy provide larger incentives for vehicles to change their routes. When dealing with larger μ^a and μ^h , the RL agent will assign larger headways to some links so that these links seem more crowded to the vehicles away from those links. In other words, the less selfish and greedy the vehicles are, the more the RL policy leverages the capacity increases of autonomous cars. This indicates that we need to carefully infer and estimate vehicles’ selfishness when deriving a control strategy, as it can cause a significant difference in the resulting policies.

C. The total fraction of autonomous cars

The total fraction of autonomous vehicles in the network will also inarguably affect the performance of our policy. The

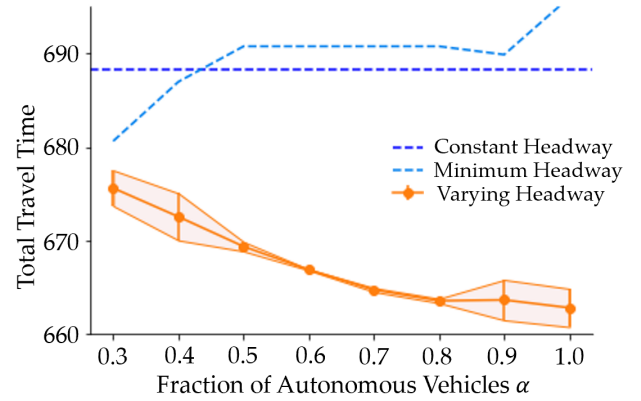


Fig. 7: Plot of TTT as a function of the demand autonomy fraction α_{O-D} , acquired by using the RL agent trained for each corresponding α_{O-D} separately. As the plot shows, the higher the fraction of autonomy is, the larger the savings in TTT are. Baselines using uniform constant headway and minimum headway are shown in dashed lines.

larger the fraction of autonomy is, the more control authority we will have for influencing vehicles’ route choices. To study this, we keep the total number of vehicles constant and only vary the fraction of vehicles that are autonomous along O/D pair O-D denoted by α_{O-D} . We varied the value of α_{O-D} and trained our policy to learn a policy tailored to the particular choice of autonomy fraction α_{O-D} . Improvements in TTT against baselines using uniform constant or minimum headway are plotted in Fig. 7.

To examine the policies learned for each α_{O-D} , we trained our policy on each α_{O-D} at least 5 times to compute TTT. Overall, the performance of our policy keeps elevating when increasing the fraction of autonomous vehicles in the total demand. This is intuitive since effectively, the portion of the flow that we can control keeps growing. We see that for lower values of α_{O-D} , the variance in TTT is larger. This is because, in the low α_{O-D} regime, the change in the headway maintained by autonomous vehicles is not having a significant effect on the traffic dynamics.

VI. CONCLUSION AND FUTURE WORK

In this work, we examined the feasibility of relieving traffic congestion by dynamically controlling the headway of autonomous vehicles. We modeled traffic networks with time-varying route choice dynamics. We used this model to train a policy for dynamically controlling the headway of autonomous cars on the network links to improve the efficiency of the network. We measured the performance of our policy and demonstrated that our trained policy reduces total travel time. Our results indicate that controlling the headway of autonomous cars can be utilized to influence vehicles’ route choices and reduce overall delays. In this work, we experimented on a relatively small-scale network due to an underpowered training platform. For our future work, we aim to apply our framework to a larger-scale network and further investigate the application of multi-agent RL when facing multiple O/D pairs in the network.

REFERENCES

- [1] C. Wu, A. Kreidieh, E. Vinitsky, and A. M. Bayen, "Emergent behaviors in mixed-autonomy traffic," in *Conference on Robot Learning*, pp. 398–407, 2017.
- [2] S. Darbha and K. Rajagopal, "Intelligent cruise control systems and traffic flow stability," *Transportation Research Part C: Emerging Technologies*, vol. 7, no. 6, pp. 329–352, 1999.
- [3] J. Yi and R. Horowitz, "Macroscopic traffic flow propagation stability for adaptive cruise controlled vehicles," *Transportation Research Part C: Emerging Technologies*, vol. 14, no. 2, pp. 81–95, 2006.
- [4] N. Mehr, R. Li, and R. Horowitz, "A game theoretic model for aggregate bypassing behavior of vehicles at traffic diverges," *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, 2018.
- [5] N. Mehr, R. Li, and R. Horowitz, "A game theoretic macroscopic model of lane choices at traffic diverges with applications to mixed-autonomy networks," *Transportation Research Part B: Methodological*, vol. 144, pp. 45–59, 2021.
- [6] J. Lioris, R. Pedarsani, F. Y. Tascikaraoglu, and P. Varaiya, "Platoons of connected vehicles can double throughput in urban roads," *Transportation Research Part C: Emerging Technologies*, vol. 77, pp. 292–305, 2017.
- [7] N. Mehr and R. Horowitz, "Can the presence of autonomous vehicles worsen the equilibrium state of traffic networks?," in *2018 IEEE Conference on Decision and Control (CDC)*, pp. 1788–1793, IEEE, 2018.
- [8] N. Mehr and R. Horowitz, "How will the presence of autonomous vehicles affect the equilibrium state of traffic networks?," *IEEE Transactions on Control of Network Systems*, vol. 7, no. 1, pp. 96–105, 2019.
- [9] J. G. Wardrop, "Some theoretical aspects of road traffic research," in *Inst Civil Engineers Proc London/UK*, 1952.
- [10] D. Braess and G. Koch, "On the existence of equilibria in asymmetrical multiclass-user transportation networks," *Transportation Science*, vol. 13, no. 1, pp. 56–63, 1979.
- [11] N. Mehr and R. Horowitz, "Pricing traffic networks with mixed vehicle autonomy," in *2019 American Control Conference (ACC)*, pp. 2676–2682, IEEE, 2019.
- [12] D. A. Lazar, S. Coogan, and R. Pedarsani, "Optimal tolling for heterogeneous traffic networks with mixed autonomy," in *2019 IEEE 58th Conference on Decision and Control (CDC)*, pp. 4103–4108, IEEE, 2019.
- [13] E. Bıyık, D. A. Lazar, R. Pedarsani, and D. Sadigh, "Incentivizing efficient equilibria in traffic networks with mixed autonomy," *IEEE Transactions on Control of Network Systems*, vol. 8, no. 4, pp. 1717–1729, 2021.
- [14] E. Bıyık, D. A. Lazar, R. Pedarsani, and D. Sadigh, "Altruistic autonomy: Beating congestion on shared roads," in *International Workshop on the Algorithmic Foundations of Robotics*, pp. 887–904, Springer, 2018.
- [15] D. A. Lazar, E. Bıyık, D. Sadigh, and R. Pedarsani, "Learning how to dynamically route autonomous vehicles on shared roads," *Transportation research part C: emerging technologies*, vol. 130, p. 103258, 2021.
- [16] M. Kolarich and N. Mehr, "Stackelberg routing of autonomous cars in mixed-autonomy traffic networks," in *2022 American Control Conference (ACC)*, pp. 4654–4661, IEEE, 2022.
- [17] P.-W. Lin, K.-P. Kang, and G.-L. Chang, "Exploring the effectiveness of variable speed limit controls on highway work-zone operations," in *Intelligent transportation systems*, vol. 8, pp. 155–168, Taylor & Francis, 2004.
- [18] B. Khondaker and L. Kattan, "Variable speed limit: an overview," *Transportation Letters*, vol. 7, no. 5, pp. 264–278, 2015.
- [19] A. Muralidharan and R. Horowitz, "Optimal control of freeway networks based on the link node cell transmission model," in *American Control Conference (ACC)*, 2012, pp. 5769–5774, IEEE, 2012.
- [20] R. L. Bertini, S. Boice, and K. Bogenberger, "Dynamics of variable speed limit system surrounding bottleneck on german autobahn," *Transportation Research Record*, vol. 1978, no. 1, pp. 149–159, 2006.
- [21] M. J. Cassidy, K. Jang, and C. F. Daganzo, "Macroscopic fundamental diagrams for freeway networks: Theory and observation," *Transportation Research Record*, vol. 2260, no. 1, pp. 8–15, 2011.
- [22] K. Lam, W. Krichene, and A. Bayen, "On learning how players learn: Estimation of learning dynamics in the routing game," *2016 ACM/IEEE 7th International Conference on Cyber-Physical Systems (ICCPS)*, 2016.
- [23] G. Gomes and R. Horowitz, "Optimal freeway ramp metering using the asymmetric cell transmission model," *Transportation Research Part C: Emerging Technologies*, vol. 14, no. 4, pp. 244–262, 2006.
- [24] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," Aug 2017.
- [25] "Baselines3 docs - reliable reinforcement learning implementations¶."