
SAINT-ACC: Safety-Aware Intelligent Adaptive Cruise Control for Autonomous Vehicles Using Deep Reinforcement Learning

Lokesh Das¹ Myounggyu Won¹

Abstract

We present a novel adaptive cruise control (ACC) system namely SAINT-ACC: Safety-Aware Intelligent ACC system (SAINT-ACC) that is designed to achieve simultaneous optimization of traffic efficiency, driving safety, and driving comfort through dynamic adaptation of the inter-vehicle gap based on deep reinforcement learning (RL). A novel dual RL agent-based approach is developed to seek and adapt the optimal balance between traffic efficiency and driving safety/comfort by effectively controlling the driving safety model parameters and inter-vehicle gap based on macroscopic and microscopic traffic information collected from dynamically changing and complex traffic environments. Results obtained through over 12,000 simulation runs with varying traffic scenarios and penetration rates demonstrate that SAINT-ACC significantly enhances traffic flow, driving safety and comfort compared with a state-of-the-art approach.

1. Introduction

An adaptive cruise control (ACC) system enables vehicles to automatically maintain the desired headway distance to the preceding vehicle by using various kinds of sensors such as radar, LiDAR, and cameras (Vahidi & Eskandarian, 2003). It is one of the critical components of autonomous vehicles (Wang et al., 2020). Recent research demonstrates that it can also be used to enhance traffic flow (Ntousakis et al., 2015; Delis et al., 2016; Nikolos et al., 2015) by adaptively adjusting the inter-vehicle gap in response to dynamically changing traffic conditions (Bekiaris-Liberis & Delis, 2019; Goñi-Ros et al., 2019; Das & Won, 2020).

Latest intelligent ACC systems, however, focus mostly on

¹Department of Computer Science, University of Memphis, TN, United States. Correspondence to: Myounggyu Won <mwon@memphis.edu>.

enhancing traffic flow, overlooking the impact of adaptive adjustment of inter-vehicle gap on driving safety and comfort (Bekiaris-Liberis & Delis, 2019; Goñi-Ros et al., 2019). As such, developing an intelligent ACC system that optimizes not only traffic flow but also driving safety and comfort is an important research problem. Recently, a surrogate safety measure (SSM) (Nadimi et al., 2020) such as the time-to-collision (TTC) model (Gettman & Head, 2003) is incorporated in the design of an intelligent ACC system to account for driving safety and comfort in making a control decision on the inter-vehicle gap (Zhu et al., 2020; Alizadeh et al., 2019). However, we demonstrate that state-of-the-art safety-aware ACC systems (Zhu et al., 2020; Alizadeh et al., 2019) do not perform driving safety assessment effectively under dynamically changing and complex traffic environments (Zhu et al., 2020) because the implications of dynamic adaptation of the critical parameters of the driving safety model on balancing traffic efficiency and safety are not fully considered.

In this paper, we develop the Safety-Aware Intelligent ACC system (SAINT-ACC) that performs driving safety assessment more effectively through dynamically updating the safety model parameters in response to varying traffic conditions. A novel dual reinforcement learning (RL) agent approach is designed aiming to maximize traffic efficiency, driving safety, and driving comfort simultaneously. Specifically, a separate RL agent is designed to find and adapt the optimal TTC threshold (Gettman & Head, 2003) based on rich traffic information including both macroscopic and microscopic traffic data obtained from the surrounding environment. This RL agent, by providing the optimal TTC threshold as feedback, interacts with the main RL agent designed to derive an optimal inter-vehicle gap that maximizes traffic flow, driving safety, and comfort at the same time.

The dual RL agents are trained under various highway scenarios including highways with on-ramps and off-ramps. A simulation platform based on the combination of a traffic simulator, SUMO (Krajzewicz et al., 2012), and a vehicular network simulator, Veins (Sommer et al., 2019), is used to perform extensive simulations over 12,000 runs with varying penetration rates and merging/exiting traffic density to evaluate the performance. Results demonstrate the superior

performance of the first ‘dynamic’ driving safety model based on RL that assesses driving safety more effectively under complex traffic conditions compared to the existing ‘static’ TTC model of the state-of-the-art intelligent ACC systems (Zhu et al., 2020) and the applicability of SAINT-ACC for various transportation applications such as truck platooning and traffic control.

2. Related Work

The methods for assessing driving safety for autonomous vehicles are categorized largely into two groups: rule-based and machine learning-based approaches. The rule-based methods are based on a safety model represented as a function to gauge the level of driving safety. These methods are easy to implement and are capable of assessing driving safety very quickly, thereby being adequate for fast-moving vehicles that should make a control decision quickly. Nilsson *et al.* design a utility function to determine whether or not it is safe for an autonomous vehicle to change lanes (Nilsson et al., 2016). The utility function computes a utility value based on the average vehicle speed, traffic density on a lane, and the remaining time until the vehicle reaches the end of the lane; If the utility value of a target lane is greater than that of the current lane, changing lanes is considered safe. Zheng *et al.* develop a safety model based on a threshold for the deceleration of a lane-changing vehicle and the immediately following vehicle (Zheng et al., 2019). If the deceleration of a vehicle is greater than the threshold, it is considered unsafe to change lanes. Numerous works utilize the surrogate safety measures (SSM) to evaluate the safety impacts of autonomous vehicle technologies (Morando et al., 2018; Viridi et al., 2019; Rahman & Abdel-Aty, 2018). However, these rule-based approaches do not cope well with dynamically changing traffic conditions (Ye et al., 2020).

To address the limitation of the rule-based methods, machine learning techniques are utilized for assessing driving safety (Zhu et al., 2020). Numerous research has demonstrated that machine learning-based approaches outperform rule-based methods in evaluating driving safety (Ye et al., 2020). Mirchevska *et al.* adopt RL integrated with a safety verification model to assure that the vehicle action is guaranteed to be safe at any time (Mirchevska et al., 2018). However, their method is designed to make only a binary decision, *i.e.*, whether taking an action is safe or not. Hoel *et al.* adopt a deep Q-Network to allow vehicles to make a lane-changing decision safely (Hoel et al., 2018). However, a simple reward function is used considering only extreme cases, *i.e.*, providing a penalizing reward only when a crash occurs; As a result, it does not evaluate driving safety effectively. Similar RL-based approaches are proposed, *e.g.*, Lin *et al.* utilize deep RL for driving safety in training a merging vehicle to minimize collisions (Lin et al., 2019); Baheri *et al.*

incorporate a rule-based model with a dynamically-learned safety module based on RL (Baheri et al., 2019). However, the reward function of these RL-based solutions takes into account only particular dangerous scenarios in giving penalizing rewards, *e.g.*, when merging vehicles stop suddenly or collide. To address this problem, Ye *et al.* develop a solution based on proximal policy optimization (PPO) using deep RL (Ye et al., 2020). Specifically, the reward function incorporates the ‘near-collision’ penalty instead of relying only on the situations where a collision actually occurs. Zhu *et al.* design a better reward function based on human driving data and various other performance criteria such as safety, efficiency, and driving comfort (Zhu et al., 2020). In particular, to evaluate driving safety, a surrogate safety measure called the time to collision (TTC) is incorporated with the reward function. However, a fixed threshold for TTC is used in their model (*i.e.*, 4 seconds), constraining the performance under dynamically changing traffic conditions. Similarly, Alizadeh *et al.* develop a deep RL agent based on TTC to enable autonomous vehicles to change lanes safely (Alizadeh et al., 2019). However, it also relies on a fixed TTC threshold (1.8 seconds). In contrast to these state-of-the-art research works, we develop a novel dual RL agent approach where the optimal TTC threshold is determined adaptively depending on dynamically changing traffic conditions, which is provided as feedback to the main RL agent for controlling the inter-vehicle gap more effectively.

3. Motivation

A surrogate safety measure (SSM) is a widely used indicator of driving safety (Nadimi et al., 2020). There are different kinds of SSMs such as time to collision (TTC), unsafe density (UD), proportion of stopping distance (PSD), comprehensive time-based measure (CTM), *etc.* (Guido et al., 2011). The most widely used one is TTC that represents the amount of time remaining until a collision with the front vehicle will occur if the current vehicle keeps its trajectory and speed (Gettman & Head, 2003). More specifically, TTC is based on a threshold denoted by TTC^* defined to determine whether or not a collision will occur; Formally, the TTC of vehicle F at time t denoted by $TTC_F(t)$ with respect to the preceding vehicle L is defined as follows: $TTC_F(t) = \frac{X_L(t) - X_F(t) - l_L}{v_F(t) - v_L(t)}$, $\forall v_F(t) > v_L(t)$, where X , v , and l are the position, speed, and body length of a vehicle, respectively. The TTC model is only valid when the speed of vehicle F is greater than that of the leading vehicle L (*i.e.*, $v_F(t) > v_L(t)$). A dangerous situation is detected at time t if the measured TTC of a following vehicle is smaller than the threshold TTC^* at time t (*i.e.*, $TTC_F(t) < TTC^*$), which indicates that the effectiveness of the TTC model depends on how the threshold TTC^* is determined.

An interesting observation is that state-of-the-art intelligent ACC systems designed based on the TTC model for evaluating driving safety rely on a fixed TTC threshold (TTC*) (Zhu et al., 2020; Alizadeh et al., 2019). In this section, we perform a simulation study to understand the effect of TTC* on the performance of an intelligent ACC system aiming to motivate the needs for a novel method to adapt TTC* dynamically in response to the current traffic conditions.

We implement a state-of-the-art intelligent ACC system based on RL integrated with the TTC model (Zhu et al., 2020). More specifically, in contrast to (Zhu et al., 2020) built upon a fixed TTC* (i.e., 4 seconds), in this study, we attempt to vary TTC* to demonstrate the effect of the threshold on the performance of the ACC system in terms of driving safety and traffic efficiency. A highway segment with an on-ramp is created for this simulation study. The details of the simulation settings are described in Section 5. The driving safety is represented as the number of ‘near’ collisions. Here a near collision is defined as a situation where there is a vehicle with the inter-vehicle gap smaller than a minimum gap (default 2.5m) (Krajzewicz et al., 2012). Additionally, it is ensured that if the speed of a vehicle is smaller than or equal to that of the leading vehicle, a near collision is not registered even if the inter-vehicle gap is smaller than the minimum gap. The traffic efficiency is represented by the average vehicle speed of all vehicles on the given highway segment.




Figure 1. The effect of TTC* on driving safety.




Figure 2. The effect of TTC* on traffic efficiency.

Fig. 1 exhibits the measured driving safety as a function of TTC*. We observe that there exists a tradeoff between driving safety and traffic efficiency depending on TTC*. More specifically, as TTC* decreases, driving safety degrades; the reason is that some dangerous situations are not detected because TTC* is too small - recall that a dangerous situation is detected when the TTC value for a vehicle is smaller than TTC*; in contrast, as depicted in Fig. 2, traffic efficiency increases as TTC* becomes smaller. Another interesting observation is that when TTC* is too small, the performance gain in terms of driving safety becomes marginal. The results indicate that there exists a TTC* that strikes the optimal balance between driving safety and traffic efficiency and that traffic efficiency can be improved further while sustaining

a similar level of driving safety. The results motivate us to develop an adaptive intelligent ACC system that adapts TTC* dynamically to achieve the optimal balance between traffic efficiency and driving safety.

4. Safety-Aware Intelligent ACC System (SAINT-ACC)

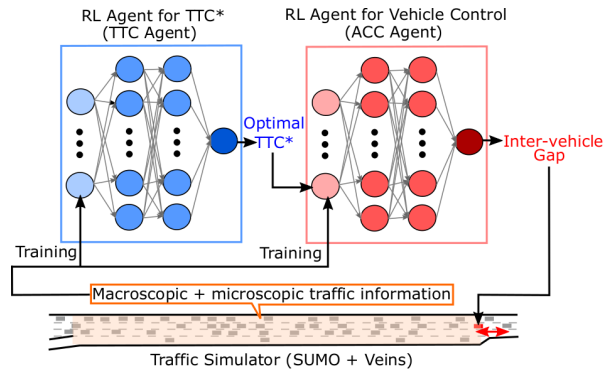


Figure 3. An overview of SAINT-ACC.

An overview of SAINT-ACC is depicted in Fig. 3. It is built upon a feedback loop comprising of dual RL agents that is used to adapt TTC* in response to dynamically changing traffic conditions and in turn utilizes the obtained optimal TTC* to control the inter-vehicle gap to maximize performance. The TTC agent is designed to obtain the optimal value of TTC*, and the ACC agent is used to control the inter-vehicle gap. Both RL agents are trained based on the macroscopic traffic information including the vehicle density and average vehicle speed as well as the microscopic traffic information such as the vehicle length and vehicle acceleration collected from the surrounding traffic environment. More specifically, the TTC agent interacts with the ACC agent continually by providing the optimal value of TTC* so that the ACC agent can control the inter-vehicle gap based on the latest and accurate driving safety assessment to improve traffic efficiency and driving safety simultaneously. The dual RL design of SAINT-ACC allows for more effective exploration and improved learning with a smaller action set (Zahavy et al., 2018), i.e., we observe that our state space increases by 25X when it is designed with a single consolidated RL agent, and the design of the reward function becomes more complicated with reduced efficiency as it should evaluate heterogeneous performance criteria.

4.1. Designing the TTC Agent

The TTC agent is used to adjust TTC* dynamically in response to the current traffic conditions. This section explains the details of the RL model for the TTC agent.

State Space: The state space is designed to represent the current traffic conditions. It incorporates both the macroscopic and microscopic traffic information including most of the widely used traffic parameters in the literature, *i.e.*, vehicle acceleration/deceleration, time headway, driving imperfection (according to SUMO (Krajewicz et al., 2012)), minimum inter-vehicle gap, vehicle length, mainline vehicle density, mainline average vehicle speed, non-mainline vehicle density, non-mainline average vehicle speed, and the length of ramp (if applicable). A salient aspect of the state space model is that it is designed to consider the traffic parameters of the mainline and non-mainline traffic separately so that SAINT-ACC can be applied to more diverse highway types. For example, while SAINT-ACC is primarily designed for highways with on/off ramps, we demonstrate that it can be applied to other road types such as a roundabout. We create a roundabout scenario and apply SAINT-ACC by modeling the traffic in terms of the mainline traffic and merging traffic. SAINT-ACC allows vehicles in the roundabout to create a gap to enable merging vehicles to enter the roundabout more smoothly without reducing the speed significantly, thereby improving the average vehicles speed compared with a baseline algorithm (Fig. 4).




Figure 4. The state space of SAINT-ACC modeling the traffic of a roundabout.

Action Space: Since the goal of the TTC agent is to find the optimal TTC^* and adapt it dynamically depending on real-time traffic conditions, the action space is primarily designed to control TTC^* . A total of 21 actions are defined with respect to a wide range of values for TTC^* between 0 and 10 with an interval of 0.5. We note that TTC^* greater than 10 is not useful as 10 seconds is enough time to avoid a collision in most highway scenarios (Vogel, 2003).




Figure 5. Convergence of the reward function of the TTC agent.

Reward Function: The reward function denoted by R_{TTC} is a core part of the TTC agent. It is used to evaluate the effectiveness of the selected TTC^* . The reward function is defined based on the concept of a ‘near’ collision (as explained in Section 3) and actual collisions in simulation. More specifically, the reward function determines the reward values based on the number of false positives (FP), number of false negatives (FN), and number of actual collisions (AC). A false positive is recorded in the event that

the TTC model with a given TTC^* detects a dangerous situation (*i.e.*, $TTC < TTC^*$) and a near-collision is not registered in simulation. On the other hand, a false negative is recorded when the TTC model with a given TTC^* does not detect a dangerous situation (*i.e.*, $TTC \geq TTC^*$) and a near-collision is registered. An actual collision is recorded in the event that two cars actually crash in simulation. The reward function R_{TTC} is then calculated as follows: $R_{TTC} = -\alpha_1(\text{Total Number of FPs}) - \alpha_2(\text{Total Number of FNs}) - \alpha_3(\text{Total Number of ACs})$. We determine those parameters as $\alpha_1 = 1$, $\alpha_2 = 2$, and $\alpha_3 = 10$ through numerous simulations. Note that a highest priority is placed on the number of actual collisions (α_3).

Table 1. Parameters used for the neural network of the TTC Agent.

Neural network architecture	Input size 13 (state size), output size 21 (action size), 1 hidden layer with 30 neurons in each layer, “random_normal” for kernel initialization
Activation functions	ReLU for hidden layer, “linear” for output layer
Replay buffer size	100k samples
$(\gamma, \epsilon_0, \epsilon_{min}, \lambda_{decay})$	(0.95, 1.0, 0.01, 0.99985)
Batch size	64
Loss function	Mean square error (MSE)
Optimization method	Adam with learning rate 0.0001
Target network update frequency	5 episodes (each episode runs for 7 minute of simulation time)
Input Normalization	Batch normalization

Neural Network: The neural network used to train the TTC agent consists of one hidden layer with 30 neurons. The linear activation function is used for the output layer to allow the deep Q network to produce the Q value with no specific bounds, not necessarily meaning no activation function as the weights are involved. The parameters for the neural network (summarized in Table 1) are carefully selected based on the extensive trial and error process. Especially, we adopt the ϵ -greedy policy (Wunder et al., 2010) to balance between exploration and exploitation in training our neural network. An action is randomly chosen from the action space with probability ϵ , while an action is selected using the greedy method with probability $1 - \epsilon$. It is ensured that ϵ is decreased gradually as the greedy algorithm iterates, *i.e.*, $\epsilon = \max(\epsilon_0 \cdot \lambda_{decay}, \epsilon_{min})$. The values selected for these parameters ϵ_0 , λ_{decay} , ϵ_{min} for training our neural network are summarized in Table 1. Fig. 5 displays the rewards for the TTC agent measured over a large number of episodes, demonstrating the convergence of the reward function.

4.2. Designing the ACC Agent

The ACC agent is trained to control the optimal inter-vehicle gap to maximize traffic flow, driving safety, and driving comfort. The TTC^* value provided by the TTC agent is used in the reward function of the ACC agent to make the control decision.

State Space: The state space of the ACC agent is designed similarly as that of the TCC agent, *i.e.*, the state space represents the current traffic conditions based on macroscopic and microscopic traffic information collected from the surrounding traffic environment.

Action Space: The action space of the ACC agent is designed to represent a set of control actions performed by a vehicle to adjust the inter-vehicle gap. Specifically, the control action is to set the inter-vehicle gap to a certain value selected from a predefined set of available values for the inter-vehicle gap. A total of 25 different actions corresponding to the headway gaps between 1m and 25m with an interval of 1m are used. A vehicle takes an action of adjusting the inter-vehicle gap every second in accordance with the current traffic conditions. This dynamic adaptation of the inter-vehicle gap is particularly useful for highways with merging/exiting traffic because creating an appropriate inter-vehicle gap is crucial for merging and exiting vehicles to change lanes smoothly without causing traffic perturbation that can lead to formation and propagation of a traffic jam. Such a positive effect of adjusting the inter-vehicle gap on preventing a traffic jam has been well studied in the literature (Goñi-Ros et al., 2019).

Reward Function: The reward function of the ACC agent denoted by R_{ACC} consists of three subfunctions representing driving safety R_s , traffic efficiency R_e , and driving comfort R_c , respectively. In designing the subfunction for driving safety, we employ a TTC-based approach of the state-of-the-art intelligent ACC system (Zhu et al., 2020) which measures driving safety based on the TTC model as follows: $F_{TTC} = \begin{cases} \log(\frac{TTC}{4}) & \text{if } 0 \leq TTC \leq 4 \\ 0 & \text{otherwise} \end{cases}$. An

interesting aspect is that the TTC threshold TTC^* is fixed to 4 in this model. In contrast, we make two critical modifications to integrate with the TTC agent: (1) we incorporate the optimal TTC^* which is dynamically adjusted by the TTC agent instead of relying on a fixed value to perform assessment of driving safety more effectively, and (2) we consider not only the driving safety for the ego vehicle but also the driving safety for surrounding vehicles in a given highway segment as a whole. Consequently, our subfunction for driving safety is calculated as follows: $R_s = \begin{cases} \sum_{i \in V} \log(\frac{TTC_i}{TTC^*}) & \text{if } 0 \leq TTC \leq TTC^* \\ 0 & \text{otherwise} \end{cases}$,

where TTC_i is the TTC of an individual vehicle i , and V is

a set of vehicles in the given highway segment.

The subfunction for traffic efficiency (R_e) is calculated based on the degree of traffic flow which is measured using the average delay to pass a given highway segment. More precisely, if the average delay is greater than the expected delay for congested traffic conditions (which is known priori), a penalizing reward is imposed. On the other hand, if the average delay is smaller than the delay for congested traffic conditions, a positive reward is provided. Consequently, the subfunction for traffic efficiency is defined as follows:

$$R_e = \begin{cases} +1 & \text{if avg delay} \leq \frac{\text{highway length}}{\text{avg speed under congestion}} \\ -1 & \text{if avg delay} > \frac{\text{highway length}}{\text{avg speed under congestion}} \end{cases}$$

Lastly, the subfunction for driving comfort (R_c) is defined based on an existing driving comfort model (Zhu et al., 2020; Jacobson et al., 1980). In this driving comfort model, the level of driving comfort is quantified using the change rate of acceleration called Jerk as follows: $R_c = -\frac{\text{jerk}^2}{27.04}$. Here, since our sampling interval is one second (*i.e.*, vehicles make a control decision on the inter-vehicle gap every second), and the acceleration is bounded between -2.6 and 2.6 m/s^2 , the greatest possible jerk is $\frac{2.6 - (-2.6)}{1.0} = 5.2$. As such, to bound the result of R_c between 0 and 1, we divide by 27.04 which is the square of the greatest possible jerk.

Now combining all three subfunctions, the reward function R_{ACC} is defined as follows: $R_{ACC} = \beta_1 R_e + \beta_2 R_s + \beta_3 R_c$, where, β_1, β_2 , and β_3 are weighting parameters which can be used to allow the user to vary the weight of a certain performance criteria. In our simulation, we use $\beta_1 = \beta_2 = \beta_3 = 1$ to optimize traffic efficiency, driving safety, and driving comfort with the same weight.

Neural Network: The neural network of the ACC agent is designed similarly as that of the TTC agent. However, through extensive trial and error, we observe that the most effective neural network configuration for the ACC agent is slightly different as follows: We use one hidden

layer with 30 neurons; the replay buffer size is 100k samples; and the batch size is 64. We also adopt the ϵ -greedy policy (Wunder et al., 2010) for training the neural network for the ACC agent with a slightly smaller λ_{delay} . Fig. 6 shows the reward values as a function of episodes demonstrating the convergence of rewards for the ACC agent.




Figure 6. Convergence of the reward function of the ACC agent.

5. Results

We implement SAINT-ACC and a state-of-the-art intelligent ACC system (referred to as ‘Zhu’ hereafter) (Zhu et al., 2020) on a traffic simulation framework, SUMO (Krajzewicz et al., 2012), which is integrated with a vehicular network simulator, Veins (Sommer et al., 2019). The state-of-the-art intelligent ACC system is based on RL that aims to optimize traffic efficiency, driving safety, and driving comfort simultaneously by designing a reward function that incorporates these three performance criteria. Compared to SAINT-ACC, however, Zhu does not account for dynamically changing traffic conditions effectively as the traffic efficiency is determined according to a static probability density function, and the driving safety depends on a fixed TTC*. The RL framework for SAINT-ACC and Zhu are implemented using Python based on Keras and Tensorflow (Abadi et al., 2016); It is interfaced with SUMO via Traffic Control Interface (TraCI) (Wegener et al., 2008). Simulations are performed with a PC running MacOS (Majove 10.14.5) with a 1.4GHz Intel Core i5 CPU and 8GB of RAM.

We consider two most common highway types, *i.e.*, a highway segment with an on-ramp and that with an off-ramp. We create a 1.5km highway segment with a on/off ramp with the length of 360m. There is a 180m-long acceleration lane used by vehicles to accelerate before they reach the on-ramp. The traffic is generated and injected into the main road at a rate of 1800 veh/h/lane. A vehicle used for this simulation has the body length of 4~5m and changes lanes based on a widely adopted lane-changing model (Erdmann, 2015). SAINT-ACC allows vehicles to take an action of controlling the inter-vehicle gap every one second in response to dynamically changing traffic conditions. For non-ACC vehicles, the default car-following model (Krauss) with driver imperfection parameters is used. In particular, the minimum headway distance for each vehicle is intentionally set to 0m to permit collisions in simulation for the purpose of evaluating driving safety. Also, both the maximum acceleration and deceleration of a vehicle are set to sufficiently high to evaluate driving comfort considering the fact that the maximum acceleration and deceleration to ensure driving comfort should be within the range between 2 m/s^2 and 3 m/s^2 (Hoberock, 1976). The performance of SAINT-ACC is compared with Zhu and the baseline method (referred to as ‘Base’ hereafter). In the baseline method, vehicles are not equipped with any intelligent ACC technology. We measure driving safety, traffic efficiency, and driving comfort by varying the technology penetration rate, vehicle density on the on-ramp as well as off-ramp. Here the penetration means the percentage of vehicles that are equipped with an intelligent ACC system, and the vehicle density refers to the number of vehicles on a particular road segment. Specifically, driving safety is measured based on

the number of near-collisions as defined in Section 3; Traffic efficiency is measured using the average speed of vehicles in the highway segment; And driving comfort is measured based on the acceleration and deceleration of vehicles.

5.1. Highway with an On-Ramp




Figure 7. An illustration of a highway segment with an on-ramp.

We evaluate the performance on a highway segment with an on-ramp. A top view of the highway segment is depicted in Fig. 7 that illustrates blue-colored vehicles adjust the inter-vehicle gap to allow red-colored merging vehicles to change lanes. In this simulation study, we measure the traffic efficiency, driving safety, and driving comfort of SAINT-ACC, Zhu, and Base by varying the penetration rate. We run 200 random simulations to represent a single data point.




Figure 8. Driving safety for highway with on-ramp.




Figure 9. Traffic efficiency for highway with on-ramp.

Table 2. Performance gain (driving safety) compared with Zhu for varying penetration rates (PR).

PR	5%	10%	20%	40%	80%
Gain	2.6%	8.5%	7.1%	27.7%	42.4%

Fig. 8 shows the results on driving safety. No significant pattern is observed for the driving safety of Base as vehicles are not equipped with an intelligent ACC system. In contrast, SAINT-ACC and Zhu achieve significantly better driving safety compared with Base. Specifically, SAINT-ACC achieves by up to 99.2% higher driving safety than that of Base when the penetration rate is 80%. It is also observed that SAINT-ACC outperforms Zhu by up to 42.4% at the penetration rate of 80%. Table 2 summarizes the performance gain compared with Zhu for different penetration rates to supplement the graph (Fig. 8) not clearly showing the performance gain due to the too large y-axis

range. On average, the driving safety is enhanced by 17.5% which translates into 7.2 less near collisions. Given a short period of time for each run, 7.2 less near collisions is notable as it can accumulate into a very large number over time. These results demonstrate the efficacy of adaptively adjusting TTC* and using it to control the inter-vehicle gap.




Figure 10. A space-time diagram.

higher traffic efficiency compared with Base because they adapt the inter-vehicle gap so that merging vehicles change the lane to the mainline smoothly with reduced interruption to traffic flow of the mainline. The average speed of SAINT-ACC is higher by up to 36.2% compared with Base when the penetration rate is 80%. Fig. 10 shows a time-space diagram that more clearly represents the qualitative behavior of SAINT-ACC compared with Base, suppressing formation of a traffic jam caused by merging vehicles.




Figure 11. Driving comfort for highway with on-ramp.




Figure 12. Driving safety for highway with off-ramp.

Fig. 11 shows the results on driving comfort. The literature shows that the maximum acceleration and deceleration to assure comfortable driving is in the range between 2 m/s² and 3 m/s² (Hoberock, 1976). Considering this range, we observe that driving comfort is not significantly impacted by all approaches regardless of the penetration rate. An interesting observation is that when the penetration rate is small (*i.e.*, smaller than 40%), both SAINT-ACC and Zhu exhibit lower driving comfort level compared with that of Base. The reason is possibly that when vehicles adjust the inter-vehicle gap to allow merging vehicles to change lanes, the following vehicles that are not equipped with an intelligent ACC system need to decrease their speed significantly, causing degradation of driving comfort. We also observe that driving comfort for both SAINT-ACC and Zhu keep improving with higher penetration rates. When most vehicles are equipped with an intelligent ACC system, we observe that both approaches provide much more comfortable driving experience. In particular, SAINT-ACC has 18.2% higher

driving comfort level than that of Zhu when the penetration rate is 80%.

Another interesting experiment is to investigate the behavior of SAINT-ACC without the TTC agent. We run the same set of simulations with and without the TTC agent. The results regarding traffic efficiency are shown in Fig. 13, demonstrating that without the dynamic adjustment of the TTC threshold, the ACC agent does not improve the performance much. Similar results are observed for driving safety and comfort.




Figure 13. The behavior of the ACC agent without the TTC agent.

5.2. Highway with Off-ramp




Figure 14. An illustration of a highway segment with off-ramp.

The performance of SAINT-ACC is measured on a highway with an off-ramp. A top view of the highway segment is depicted in Fig. 14. In this simulation scenario, vehicles exiting the highway are randomly selected from the three lanes of the main road. Fig. 12 shows the results on driving safety. Similar to the results obtained from the highway segment with an on-ramp, no significant pattern of driving safety is observed for Base since vehicles are not equipped with an intelligent ACC system. In contrast, as the penetration rate increases, driving safety is significantly increased for both SAINT-ACC and Zhu. Specifically, driving safety for SAINT-ACC is significantly higher than that of Base by up to 98.7%. It is observed that dynamic adjustment of the inter-vehicle gap allows exiting vehicles to change lanes smoothly to exit the highway. In particular, in comparison with Zhu, we observe that SAINT-ACC improves driving safety by up to 13.8% when the penetration rate is 5%.

Fig. 15 shows the results on traffic efficiency. Both SAINT-ACC and Zhu have significantly higher traffic efficiency than Base, especially as the penetration rate increases. In particular, SAINT-ACC has higher traffic efficiency by up to 134.2% compared with Base. This result demonstrates that SAINT-ACC enables exiting vehicles to change lanes to exit the highway without causing much interruption to surrounding vehicles since the inter-vehicle gap is appropriately adjusted. We also observe that as the penetration rate increases, the traffic efficiency of both SAINT-ACC and

Zhu increases. Interestingly, the performance gap between the two approaches becomes greater when the penetration is small; It is observed that SAINT-ACC achieves higher traffic efficiency by up to 6.82% compared with Zhu. Although the performance gap for traffic efficiency seems small, the results demonstrate the capability of SAINT-ACC achieving higher traffic efficiency while keeping superior driving safety than that of Zhu.




Figure 15. Traffic efficiency for highway with off-ramp.




Figure 16. Driving comfort for highway with off-ramp.

Fig. 16 shows the results on driving comfort. Similar results are observed as those obtained from the highway with an on-ramp. Overall, the impact of exiting vehicles on driving comfort is insignificant in our simulation setting. Another interesting observation is that driving comfort for both SAINT-ACC and Zhu keeps increasing as the penetration rate increases. Compared with Base, SAINT-ACC improves the level of driving comfort by up to 40.2% when the penetration is as high as 80%. However, when the penetration rate is small (*e.g.*, < 30%), driving comfort for SAINT-ACC becomes even worse than that for Base. We think that the reason is because the following vehicles without an intelligent ACC system have to accelerate/decelerate as the front vehicle equipped with SAINT-ACC or Zhu adjusts the inter-vehicle gap.

5.3. Effect of Merging Traffic




Figure 17. Effect of merging traffic on driving safety.




Figure 18. Effect of merging traffic on traffic efficiency.

Merging traffic is one of the main causes of traffic congestion (Milanés et al., 2010). We evaluate the performance of SAINT-ACC by varying the merging traffic density defined as the number of vehicles on a merging lane per unit time (veh/h/km). Fig. 17 displays the driving safety of all three

approaches as a function of the merging vehicle density. The results show that the driving safety of Base degrades when there are more merging vehicles. It is also observed that the traffic becomes increasingly unstable as the merging vehicle density increases as indicated by the large standard deviation. In contrast, both SAINT-ACC and Zhu successfully suppress near-collisions even when the merging vehicle density is very high. Specifically, SAINT-ACC improves driving safety by up to 93.6% compared with that of Base, indicating that SAINT-ACC manages effectively complex traffic conditions with very high merging vehicle density. In comparison with Zhu, SAINT-ACC achieves higher driving safety by up to 8.7% demonstrating the effectiveness of dynamically adjusting TTC*.

Fig. 18 shows the results on traffic efficiency. It is observed that when no intelligent ACC solution is applied, traffic efficiency degrades as the merging vehicle density increases. In contrast, both SAINT-ACC and Zhu effectively sustain high average vehicle speed despite the increasing merging vehicle density. In particular, SAINT-ACC achieves higher average speed by up to 83.9% compared with that of Base, demonstrating the effectiveness of dynamic adaptation of the inter-vehicle gap in response to varying merging vehicle density. We also observe that SAINT-ACC achieves higher traffic efficiency compared with Zhu.




Figure 19. Effect of merging traffic on driving comfort.




Figure 20. Effect of exiting traffic on driving safety.

It is observed that the level of driving comfort decreases for all three approaches as the merging vehicle density increases as depicted in Fig. 19. It is also observed that both SAINT-ACC and Zhu improve the driving comfort level compared with that of Base. However, it should be mentioned that the results are insignificant as the average acceleration of the vehicles for all approaches is within the range of the acceleration for comfortable driving (Hoberock, 1976) regardless of the merging vehicle density.

5.4. Effect of Exiting Traffic

Exiting traffic (vehicles exiting the highway) is another important cause for traffic congestion (Günther et al., 2012). We evaluate the performance of SAINT-ACC by varying the exiting traffic density defined as the number of vehicles exiting the highway per unit time (veh/h/km). In particular,




Figure 21. Effect of exiting traffic on traffic efficiency.




Figure 22. Effect of exiting traffic on driving comfort.

exiting vehicles are randomly selected from all three lanes of the main road in our simulation. Fig. 20 shows the effect of exiting traffic density on driving safety. It can be seen that the driving safety of Base significantly degrades as there are more vehicles exiting the highway. In contrast, both SAINT-ACC and Zhu keep the number of near-collisions fairly small despite the increasing number of exiting vehicles. Especially, SAINT-ACC enhances driving safety by up to 91.5% compared with that of Base. Furthermore, SAINT-ACC achieves better driving safety in comparison with that of Zhu as well by up to 24% when the exiting vehicle density is 2,000 veh/h/km. The results clearly demonstrate the advantages of adjusting the inter-vehicle gap adaptively using the optimal TTC* in response to dynamically changing traffic conditions.

Fig. 21 depicts the results on traffic efficiency with varying exiting traffic density. We observe that the traffic efficiency of all three approaches decreases as the exiting vehicle density increases. Interestingly, it is observed that traffic efficiency is less significantly affected by exiting traffic compared to how merging traffic influences on traffic efficiency. A possible explanation is that vehicles exiting the highway are randomly distributed on all three lanes; as such, the traffic flow of all three lanes are affected by the exiting vehicles; In contrast, all merging vehicles attempt to change lanes from a single lane (*i.e.*, a ramp), thereby affecting only the traffic flow on the lane adjacent to the ramp, creating greater traffic perturbation. Overall, compared with Base and Zhu, SAINT-ACC improves the traffic flow by up to 11.8% and 4.1%, respectively. Fig. 22 displays the results on driving comfort which demonstrate that SAINT-ACC improves the level of driving comfort by up to 41.3% and 7.6% in comparison with that of Base and Zhu, respectively.




Figure 23. Computation delay.

5.5. Computational Delay

We measure the computational delay required to make a control decision on the inter-vehicle gap. The computational delay is measured 100 times per epoch for SAINT-ACC and Zhu. Fig. 23 depicts the cumulative distribution function (CDF) graph of the computational delay. The average computational delay for Zhu and SAINT-ACC is 22.3ms and 26.9ms, respectively. It can be seen that the computational delay for SAINT-ACC is 17% higher on average than that of Zhu. The reason is SAINT-ACC has a more complex RL model trained with a larger number of traffic parameters; additionally, SAINT-ACC is designed with dual RL models. However, the results demonstrate that the average computation delay for SAINT-ACC is still sufficiently small for vehicles to make a control decision every one-second update interval.

5.6. Update Interval

SAINT-ACC makes a control decision of the inter-vehicle distance every second, allowing a vehicle to cope with dynamically changing traffic conditions more effectively. Additionally, SAINT-ACC allows that the update interval can be easily changed depending on traffic situations. We evaluate the effect of the update interval on the performance of SAINT-ACC focusing on traffic efficiency. Results are depicted in Fig. 24. As shown, as the update interval increases, the traffic efficiency gradually degrades. The reason is straightforward that vehicles by updating the inter-vehicle distance more frequently cope better with the dynamically changing traffic conditions. Another interesting observation is that when the update interval is sufficiently small (*i.e.*, < 3 seconds), the performance gain is marginal for a smaller update interval, allowing us to choose an update interval between 1 second and 3 seconds.




Figure 24. Effect of the update interval of the inter-vehicle distance.

6. Conclusion

We have presented a safety-aware intelligent ACC system that maximizes traffic efficiency, driving safety, and driving comfort simultaneously by finding and adjusting the optimal TTC threshold used to control the inter-vehicle gap adaptively in response to dynamically changing traffic conditions. Our future work is to implement a proof-of-concept SAINT-ACC system on a real vehicle platform and perform experiments to validate the performance under real-world traffic environments using open-source semi-autonomous driving systems such as Comma.AI (Santana & Hotz, 2016).

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
- Alizadeh, A., Moghadam, M., Bicer, Y., Ure, N. K., Yavas, U., and Kurtulus, C. Automated lane change decision making using deep reinforcement learning in dynamic and uncertain highway environment. In *Intelligent Transportation Systems Conference (ITSC)*, pp. 1399–1404, 2019.
- Baheri, A., Nagesh Rao, S., Tseng, H. E., Kolmanovsky, I., Girard, A., and Filev, D. Deep reinforcement learning with enhanced safety for autonomous highway driving. *arXiv preprint arXiv:1910.12905*, 2019.
- Bekiaris-Liberis, N. and Delis, A. Feedback control of freeway traffic flow via time-gap manipulation of ACC-equipped vehicles: A PDE-based approach. *IFAC-PapersOnLine*, 52(6):1–6, 2019.
- Das, L. and Won, M. D-ACC: dynamic adaptive cruise control for highways with on-ramps based on deep q-learning. *arXiv preprint arXiv:2006.01411*, 2020.
- Delis, A. I., Nikolos, I. K., and Papageorgiou, M. Simulation of the penetration rate effects of acc and cacc on macroscopic traffic dynamics. In *International Conference on Intelligent Transportation Systems (ITSC)*, pp. 336–341, 2016.
- Erdmann, J. Sumo’s lane-changing model. In *Modeling Mobility with Open Data*, pp. 105–123, 2015.
- Gettman, D. and Head, L. Surrogate safety measures from traffic simulation models. *Transportation Research Record*, 1840(1):104–115, 2003.
- Goñi-Ros, B., Schakel, W. J., Papacharalampous, A. E., Wang, M., Knoop, V. L., Sakata, I., van Arem, B., and Hoogendoorn, S. P. Using advanced adaptive cruise control systems to reduce congestion at sags: An evaluation based on microscopic traffic simulation. *Transportation Research Part C: Emerging Technologies*, 102:411–426, 2019.
- Guido, G., Saccomanno, F., Vitale, A., Astarita, V., and Festa, D. Comparing safety performance measures obtained from video capture data. *Journal of Transportation Engineering*, 137(7):481–491, 2011.
- Günther, G., Coeymans, J. E., Muñoz, J. C., and Herrera, J. C. Mitigating freeway off-ramp congestion: A surface streets coordinated approach. *Transportation Research Part C: Emerging Technologies*, 20(1):112–125, 2012.
- Hoberock, L. L. A survey of longitudinal acceleration comfort studies in ground transportation vehicles. Technical report, Council for Advanced Transportation Studies, 1976.
- Hoel, C.-J., Wolff, K., and Laine, L. Automated speed and lane change decision making using deep reinforcement learning. In *International Conference on Intelligent Transportation Systems (ITSC)*, pp. 2148–2155, 2018.
- Jacobson, I. D., Richards, L. G., and Kuhlthau, A. R. Models of human comfort in vehicle environments. *Human Factors in Transport Research*, 2, 1980.
- Krajzewicz, D., Erdmann, J., Behrisch, M., and Bieker, L. Recent development and applications of SUMO-Simulation of Urban MObility. *International Journal on Advances in Systems and Measurements*, 5(3&4), 2012.
- Lin, Y., McPhee, J., and Azad, N. L. Anti-jerk on-ramp merging using deep reinforcement learning. *arXiv preprint arXiv:1909.12967*, 2019.
- Milanés, V., Godoy, J., Villagrà, J., and Pérez, J. Automated on-ramp merging system for congested traffic situations. *IEEE Transactions on Intelligent Transportation Systems*, 12(2):500–508, 2010.
- Mirchevska, B., Pek, C., Werling, M., Althoff, M., and Boedecker, J. High-level decision making for safe and reasonable autonomous lane changing using reinforcement learning. In *International Conference on Intelligent Transportation Systems (ITSC)*, pp. 2156–2162, 2018.
- Morando, M. M., Tian, Q., Truong, L. T., and Vu, H. L. Studying the safety impact of autonomous vehicles using simulation-based surrogate safety measures. *Journal of Advanced Transportation*, 2018, 2018.
- Nadimi, N., NaserAlavi, S. S., and Asadamraji, M. Calculating dynamic thresholds for critical time to collision as a safety measure. In *Proceedings of the Institution of Civil Engineers-Transport*, pp. 1–21, 2020.
- Nikolos, I. K., Delis, A. I., and Papageorgiou, M. Macroscopic modelling and simulation of ACC and CACC traffic. In *International Conference on Intelligent Transportation Systems (ITSC)*, pp. 2129–2134, 2015.
- Nilsson, J., Silvlín, J., Brannstrom, M., Coelingh, E., and Fredriksson, J. If, when, and how to perform lane change maneuvers on highways. *IEEE Intelligent Transportation Systems Magazine*, 8(4):68–78, 2016.
- Ntousakis, I. A., Nikolos, I. K., and Papageorgiou, M. On microscopic modelling of adaptive cruise control systems. *Transportation Research Procedia*, 6:111–127, 2015.

- Rahman, M. S. and Abdel-Aty, M. Longitudinal safety evaluation of connected vehicles' platooning on expressways. *Accident Analysis & Prevention*, 117:381–391, 2018.
- Santana, E. and Hotz, G. Learning a driving simulator. *arXiv preprint arXiv:1608.01230*, 2016.
- Sommer, C., Eckhoff, D., Brummer, A., Buse, D. S., Hagenauer, F., Joerer, S., and Segata, M. Veins: The open source vehicular network simulation framework. In *Recent Advances in Network Simulation*, pp. 215–252. Springer, 2019.
- Vahidi, A. and Eskandarian, A. Research advances in intelligent collision avoidance and adaptive cruise control. *IEEE Transactions on Intelligent Transportation Systems*, 4(3):143–153, 2003.
- Virdi, N., Grzybowska, H., Waller, S. T., and Dixit, V. A safety assessment of mixed fleets with connected and autonomous vehicles using the surrogate safety assessment module. *Accident Analysis & Prevention*, 131:95–111, 2019.
- Vogel, K. A comparison of headway and time to collision as safety indicators. *Accident Analysis & Prevention*, 35(3):427–433, 2003.
- Wang, C., Gong, S., Zhou, A., Li, T., and Peeta, S. Cooperative adaptive cruise control for connected autonomous vehicles by factoring communication-related constraints. *Transportation Research Part C: Emerging Technologies*, 113:124–145, 2020.
- Wegener, A., Piórkowski, M., Raya, M., Hellbrück, H., Fischer, S., and Hubaux, J.-P. TraCI: an interface for coupling road traffic and network simulators. In *Communications and Networking Simulation Symposium*, pp. 155–163, 2008.
- Wunder, M., Littman, M. L., and Babes, M. Classes of multiagent q-learning dynamics with epsilon-greedy exploration. In Fürnkranz, J. and Joachims, T. (eds.), *Proceedings of the 27th International Conference on Machine Learning (ICML-10), June 21-24, 2010, Haifa, Israel*, pp. 1167–1174. Omnipress, 2010.
- Ye, F., Cheng, X., Wang, P., and Chan, C.-Y. Automated lane change strategy using proximal policy optimization-based deep reinforcement learning. *arXiv preprint arXiv:2002.02667*, 2020.
- Zahavy, T., Haroush, M., Merlis, N., Mankowitz, D. J., and Mannor, S. Learn what not to learn: Action elimination with deep reinforcement learning. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 3566–3577, 2018.
- Zheng, Y., Ran, B., Qu, X., Zhang, J., and Lin, Y. Cooperative lane changing strategies to improve traffic operation and safety nearby freeway off-ramps in a connected and automated vehicles environment. *IEEE Transactions on Intelligent Transportation Systems*, 2019.
- Zhu, M., Wang, Y., Pu, Z., Hu, J., Wang, X., and Ke, R. Safe, efficient, and comfortable velocity control based on reinforcement learning for autonomous driving. *Transportation Research Part C: Emerging Technologies*, 117:102662, 2020.