

Incorporating Driver Relaxation into Factory Adaptive Cruise Control and its Implications on Traffic Operation

Hao Zhou^a, Jorge Laval^{a,*}

^a*School of Civil and Environmental Engineering, Georgia Institute of Technology, Atlanta, United States*

Abstract

Current adaptive cruise control (ACC) systems adopt fixed desired time headway, which leads to an abrupt speed reduction after being cut-in by a lane changer in front or when changing lanes too close to the new leader. In contrast, human drivers behave differently and feature a variable spacing within 20 or 30 seconds right after a cut-in or lane change. Motivated by the smooth transition found in driver relaxation, the paper aims to incorporate relaxation into ACC systems. Based on the open-source ACC platform, Openpilot, Comma.ai, the paper proposes a feasible relaxation model compatible with current factory ACCs, which has also been tested using a market car with stock ACC hardware.

The study further investigates the impact of relaxation ACC on traffic operation. Numerical simulation suggests that incorporating relaxation into ACC can help: i) reduce the magnitude of speed perturbations in both cut-in vehicles and followers; ii) stabilize the lane-changing traffic by reducing the speed variance and prevent the lateral propagation of congestion, and iii) increase the average vehicle speed and capacity in merging traffic.

Keywords: factory ACC, Openpilot, relaxation, merge traffic

1. Introduction

With adaptive cruise control (ACC) becoming a popular package on recent car models, lots of research attention has been drawn from the traffic community to investigate its impact on traffic operation. Notably, we see a large body of literature Li et al. (2021), Makridis et al. (2020), Shi and Li (2021), Montanino and Punzo (2021), Zhou et al. (2021a,b) working on the string stability of the longitudinal control in ACCs under the regular car following (CF) state, which mainly focuses on the propagation of speed perturbations from the leader to followers in the platoon.

Unlike the behaviors under the normal CF state, another key component of the longitudinal control, the ACC's reaction model before or after lane changes, has received little attention. To see why the behaviors before or after lane changes can be different, many works of literature in the traffic community has investigated the analogous behaviors in human driving, referred to as the 'relaxation' phenomenon Cohen (2004), Laval and Leclercq (2008), Zheng et al. (2013), which means human drivers are willing to accept smaller time headways within 20 or 30 seconds after the lane change near the merge area. While 'relaxation' in human driving has been discovered and studied for many years, so far we haven't seen relevant works investigating the behaviors in ACC's control under the same circumstances. Since no breakthrough design has been reported from the automakers to claim a flexible time headway in different scenarios, we conjecture that relaxation does not exist in current factory ACCs. Such guess is also consistent with daily riding experiences from many ACC users when their cars are being cut and decelerate abruptly to yield.

The relaxation in human driving indeed brings a lot of benefits. On the single vehicle/lane level, accepting a small time headway prevents the vehicle from over-reacting to a sudden new leader. It also avoids the abrupt speed reduction in the ACC and the congestion propagation towards the upstream followers Milanés and Shladover (2016).

*Corresponding author

Email address: jorge.laval@ce.gatech.edu (Jorge Laval)

¹790 Atlantic Dr NW, Atlanta, GA, 30313

On the aggregated level, simulation studies Kim and Coifman (2013) found that driver relaxation can help increase the capacity of the merge area.

Motivated by the benefits of relaxation found in human driving, the paper aims to transfer it to ACCs, especially for the factory ACC units on new market car models. Despite many relaxation modeling studies Laval and Leclercq (2008), Leclercq et al. (2007), to the best of our knowledge, the authors only found one work Milanés and Shladover (2016) that implemented the linearly relaxed desired time headways in cooperative ACC platoons. In this study, the on-road experiments were conducted to show their model can effectively handle lane changes without causing major perturbations. However, their model was only applied to the in-house vehicles that need special modifications and hardware to meet the requirement of the CACC environment. To incorporate relaxation into factory ACC units on recent car models need more investigation and possibly different algorithms.

A recent open-source ACC platform, Openpilot, Comma.ai com has gained growing popularity in the AV industry Stumpf (2020). It uses an after-market device to connect to the stock interfaces of the onboard ACC units, allowing us to obtain the origin sensor data and write commands to the stock car control unit as the stock ACC does. To the best of our knowledge, this is by far the closest way for us to understand the design of the "black-box" factory ACCs units on recent market vehicles. The study Zhou et al. (2021b) presents the two-level control pipeline as shown in Fig.1. We believe, due to the same ACC hardware, the "blackbox" stock ACC shares a similar design to what's in Openpilot.

With that being said, the paper tries to incorporate relaxation into factory ACC systems based on Openpilot. In section 2 we will first present the control algorithm of ACCs in Openpilot before and after a cut-in or lane change. In section 3 we propose a feasible design to incorporate relaxation into current factory ACC systems. In section 4 we implement and validate the design on a regular market car model. Section 5 investigates the impact of relaxation ACC on traffic operation and section 6 concludes the paper.

2. State of the practice

2.1. Factory linear ACC

A recent study Zhou et al. (2021b) presents the linear planner in an open-sourced factory ACC, where the desired spacing s_{ego} and the target speed v_{target} for the ego vehicle are computed as follows:

$$v_{target} = v_{lead} + k(s_{des} - s_{ego}) \quad (1)$$

$$s_{des} = \tau \cdot v_{lead} + \delta \quad (2)$$

This linear planner is simple but somewhat surprising because it adopts the lead vehicle speed v_{lead} , rather than the ego speed v_{ego} that is commonly used in literature to calculate the desired spacing and the target speed.

We argue that using v_{lead} in the planner algorithm is safer because the radar detections of both the speed and spacing can be leveraged for the next step planner target. For example, when the leader decelerates abruptly, both the shortened spacing and the reduced lead speed contribute to moving the v_{target} towards the correct direction. In contrast, if the v_{ego} is used in (1) and (2), the planner only responds to the distance error, which might take a longer time to accumulate to a significant value and lead the ego vehicle to brake in time.

From the controller perspective, (1) is a speed controller that relies on two terms, in which v_{lead} is a feedforward term and $k(s_{des} - s_{ego})$ is a feedback term proportional to the distance error. Comparing the feedback controller which needs to wait for the error to accumulate first, adding the feedforward term is expected to make the speed controller more responsive and provides extra safety.

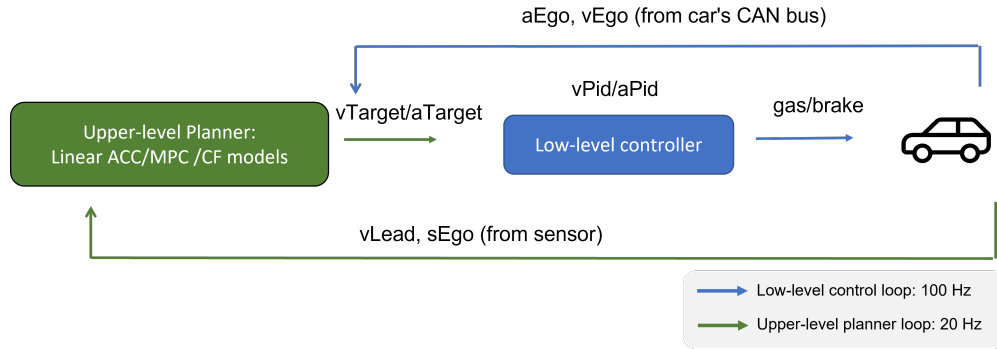


Figure 1: Control pipeline of Openpilot, an open-source factory ACC platform

Note that the target speed from (1) is only a speed reference for the vehicle's low-level controller; see the control pipeline in Fig.1. Overall, the ego vehicle speed v_{ego} tracks the target speed v_{target} but the speed changing rate is constrained to the acceleration and deceleration bounds. More details of the low-level control can be found in Zhou et al. (2021b).

2.2. ACC's control algorithm in lane changes

Now we introduce the factory ACC reaction model when it detects two different leaders during a lane change of itself or from a cut-in leader.

The algorithm used in Openpilot is shown in **Algorithm 1**. In short, the control paradigm always picks the leader that makes the car goes slower and after the lane change, it treats the new leader as a regular leader who has the speed and spacing that has suddenly changed. More specifically, the new spacing, which is often far smaller than the equilibrium spacing, is immediately used in (1) for planning a target speed. The large error between the short new spacing and the equilibrium spacing produces an abrupt decrease in the v_{target} , causing the large discrepancy with the ego speed. In this case, the ego vehicle will use its deceleration limit to move towards the v_{target} , until the speed discrepancy caused by the distance error is reduced. Note that the ego vehicle decelerates at its maximum until the equilibrium spacing is achieved, which is certainly too conservative compared to an average human driver.

Algorithm 1 Reaction model to lane changes with a linear ACC in Openpilot

Input: regular leader l_1 ; new leader l_2

Output: v_{target}

```

1: if  $l_1$  is not None then
2:    $d_{des}$  is derived by plugging  $l_1.v_{lead}$  into (2)
3:   Calculate  $v_{target}$  based on  $d_{des}$  using (1)
4:   if  $l_2$  is not None then
5:     Plug  $l_2.v_{lead}$  in (2) to compute the  $d_{des}^2$ 
6:     Plug  $d_{des}^2$  to compute the  $v_{target}^2$ 
7:     if  $v_{target}^2 < v_{target}^2$  then
8:        $l_1 = l_2$  # Now treat the cut-in vehicle as the main leader
9:        $d_{des} = d_{des}^2, v_{target} = v_{target}^2$ 
10:    end if
11:  end if
12: end if

```

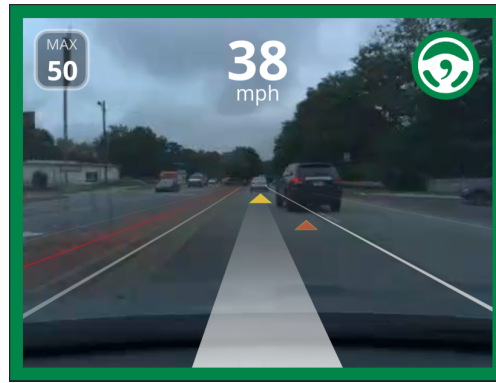
In the case where the ego vehicle is cut, the current ACC control algorithm suggests the ego vehicle will decelerate as much as possible to yield, starting from the point when the radar detects the lane changer or a vision model predicts the cut-in behavior. For example, Fig.2 (a) shows the detection of both the regular leader and the new cut-in leader in Openpilot, where the spacing and speed of the two leaders are output by the radar in Fig.2 (b) and (c). In around

step 2800, the radar detected the cut-in leader which has a spacing less than 10 m, consequently, the ego vehicle was forced to decelerate harshly, see Fig.2 (c) and (d) for the ego speed and acceleration, respectively.

Note that the reaction model to a cut-in vehicle also indicates no competition from the ACC follower since the ACC always treats the cut-in as a new leader and yields it. This finding explains why the ACC vehicles are frequently 'cut' in real-world traffic as reported by many ACC users.

If the ego vehicle is the lane changer, the ACC will respond immediately with its new leader according to (1). The target speed v_{target} given by the new leader is often much smaller than its current speed due to the suddenly shortened spacing.

In both cases, we argue that the ACC behaviors are 'unnatural' compared to human driving because it decelerates abruptly due to the suddenly dropped spacing with the new leader. The characteristics and model for driver relaxation will be discussed shortly.



(a) Front camera view: at step 2800, the orange triangle indicates a cut-in new leader.

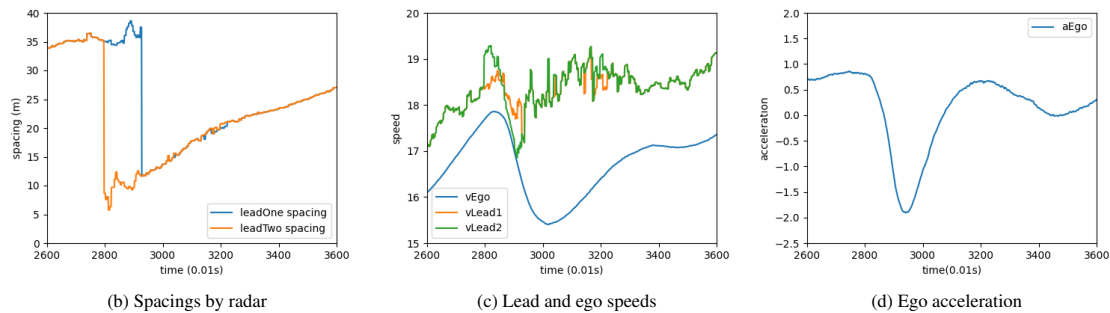


Figure 2: Radar detection and ACC reactions to a cut-in leader

3. Incorporating relaxation into factory ACCs

Cohen (2004) formulated the first model to capture human driver's relaxation behaviors where the time headway is linearly relaxed from an initial time headway τ_{init} to the equilibrium value τ within a relaxation period that lasts for T_m seconds; see Fig.3. This similar model has been re-invented or adopted by many follow-up studies Schakel et al. (2012), Milanés and Shladover (2016).

In Cohen's model, the relaxation rate, or the slope in Fig.3 varies a lot which depends on the initial time headway. We argue that this is not necessary and instead we keep the same relaxation rate of the desired time headway.

Now we introduce a linear relaxed model with small modification to work with current linear factory ACC model; see Fig.3. At time t , the most recent planning step is denoted as \hat{t} , then the relaxed time headway $\tau(t)$ is computed as follows:

$$\tau(t) = \max\{\tau_{init}, \tau_{min} + \frac{e(\hat{t})}{T_m}(\tau - \tau_{min})\} \quad (3)$$

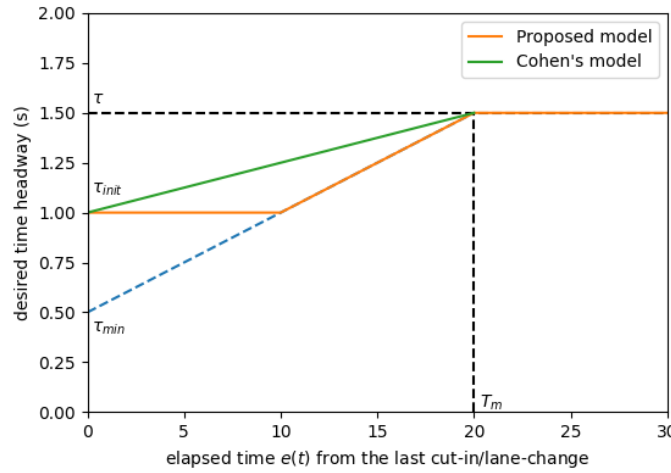


Figure 3: Proposed relaxation model for factory ACCs

Unless the initial time headway τ_{init} is smaller than the minimum safety value τ_{min} , in the first stage of the relaxation process no spacing adjustment is required. According to the target speed planner (1), the ego vehicle is essentially moving the ego vehicle speed towards the lead vehicle speed v_{lead} , since the distance error is zero. This period remains until the elapsed time increases and the desired spacing grows to a larger value than the current spacing.

The detailed algorithm to incorporate this model into the factory ACC planner is provided in **Algorithm 2**, where a key new parameter, the elapsed time $e(t)$ from the last cut-in or lane change is introduced to the ACC planner.

Algorithm 2 Incorporating relaxation into ACC

Input: Elapsed time e from last cut-in/lane-change; Maximum relaxation time T_m ; Minimum time headway τ_{min}

```

1: for every planning step  $\hat{t}$  : do
2:   if  $l_2.dRel \neq l_1.dRel$  then
3:      $e(\hat{t}) = 0$ 
4:   else
5:      $e(\hat{t}) = e(\hat{t}) + d\hat{t}$ 
6:      $e(\hat{t}) = \min(e(\hat{t}), T_m)$ 
7:   end if
8:   Current time headway  $\tau_{cur}(\hat{t}) = s_{ego}(\hat{t})/v_{lead}$ 
9:   Desired distance  $d_{des}(t) = v_{lead} \cdot \tau(t)$  from (3)
10:  Using  $d_{des}(t)$  to calculate  $v_{target}$  and execute the low-level control
11: end for

```

4. Validation on real cars

To justify the proposed model and algorithm, in this section we implement the relaxation ACC on a regular market vehicle, a 2019 Honda Civic. Then the relaxation ACC will be tested on open roads, where both the cut-in and lane change scenarios are included.

4.1. Control details

Now we first show the control details of the relaxation ACC when a vehicle cuts in. In Fig.4 (a), we exhibit the true spacings detected by the radar, where the cut-in vehicle suddenly decreases the spacing of the second leader of the ACC. Also, the relaxed and non-relaxed spacing are compared against each other, where the relaxed spacing is only slightly larger because it is already close to the equilibrium spacing. Fig.4 (b) shows the abrupt decrease in the target speed. Notice that the sharp drop is mainly due to the change of lead spacing, where the origin spacing is much larger than the equilibrium spacing and the new spacing with the cut-in leader falls short. It indicates that the relaxed desired headway does not predominate the linear ACC planner, which helps safety.

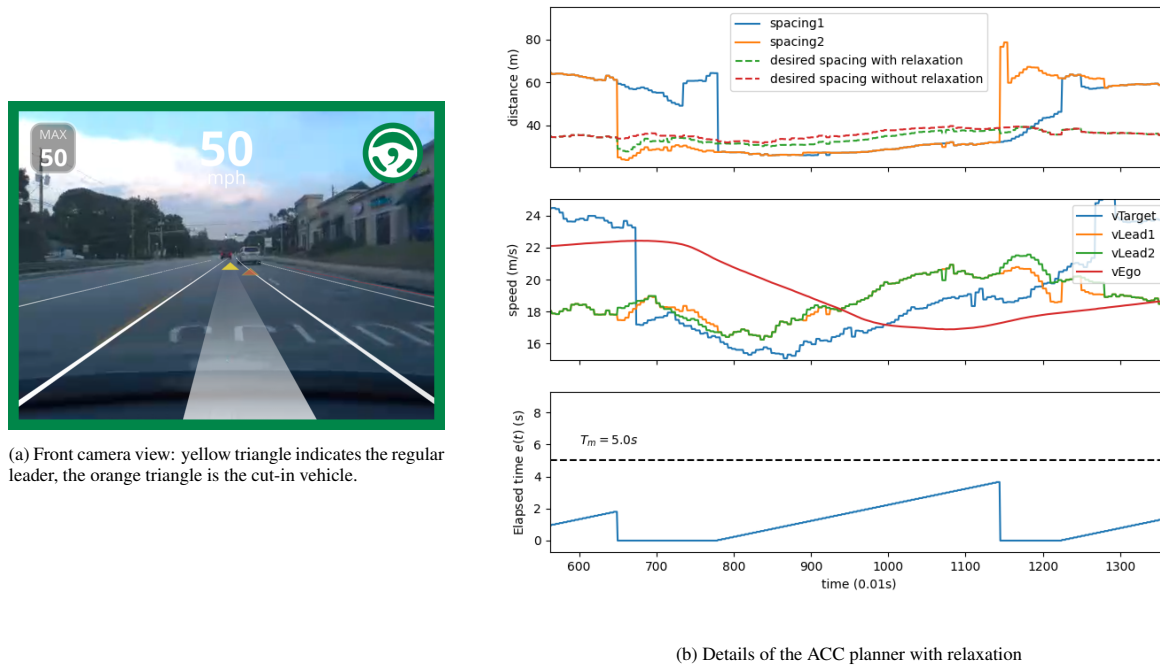


Figure 4: Implementation of relaxation ACC on a real car for the cut-in case

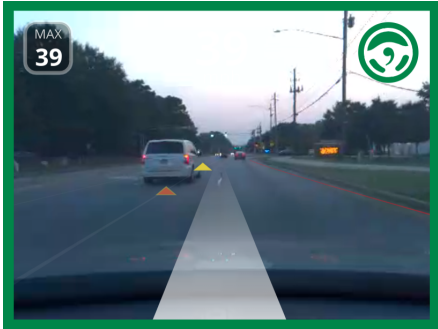
We further show the implementation of the relaxation ACC when the ego vehicle makes a lane change and follows a new target leader. Fig.5 (a) shows the small spacing with the new target leader, and the relaxation ACC produces a much smaller desired spacing in this scenario because the initial time headway (τ_{init}) is far from the equilibrium time headway τ . In Fig.5 (b), one can see the ego vehicle speed does not change too much. This is because the target lane lead speed is larger, which cancels out the deceleration demand caused by a smaller spacing. The overall effect is desirable since very little speed perturbation is caused and there still exists enough safety distance.

In both the lane-change and cut-in cases, the introduced time tracker $e(t)$; see Fig.4 (c) and Fig.5 (c); runs perfectly, which verified our design in Algorithm 2.

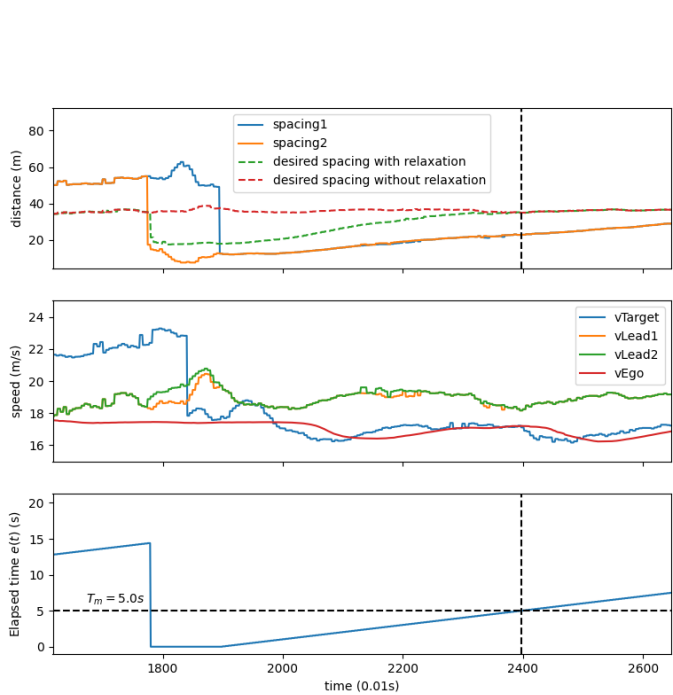
4.2. Safety concerns

The relaxation ACC shares the spirit of human drivers who are willing to accept temporary smaller spacings after lane changes. While we have shown two successful implementations above, on the other hand, it is certainly worth investigating whether it is still safe in some corner cases such as a minimal initial spacing takes place.

Now we show, it is feasible to tune the low-level controllers and the deceleration bounds in the ACC planner to obtain sufficient deceleration capabilities, which helps alleviate the safety concerns caused by small gaps. In the example; see Fig.6, the results of our filed tests where a large deceleration bound can guarantee safety even if the initial spacing is less than 10 m.



(a) Front camera view: yellow triangle indicates the regular leader (lead 1), the ego vehicle is the lane change vehicle, the orange vehicle is the new target lane leader (lead 2).



(b) Details of the ACC planner with relaxation

Figure 5: Implementation of relaxation ACC on a real car for the lane change case

We argue that more safety design can be involved in the lane-changing decision process before the relaxation model causes any safety issues. For example, the lead vehicle speed advantage or the lead vehicle acceleration can both be used to guarantee there is safety room for the lane change. Meanwhile, more advanced machine learning models can help future AVs better predict the lead vehicle trajectory, which can be incorporated into the decision-making process before a lane change is executed. Then whether to activate the relaxation model can be more intelligent and safe as well.

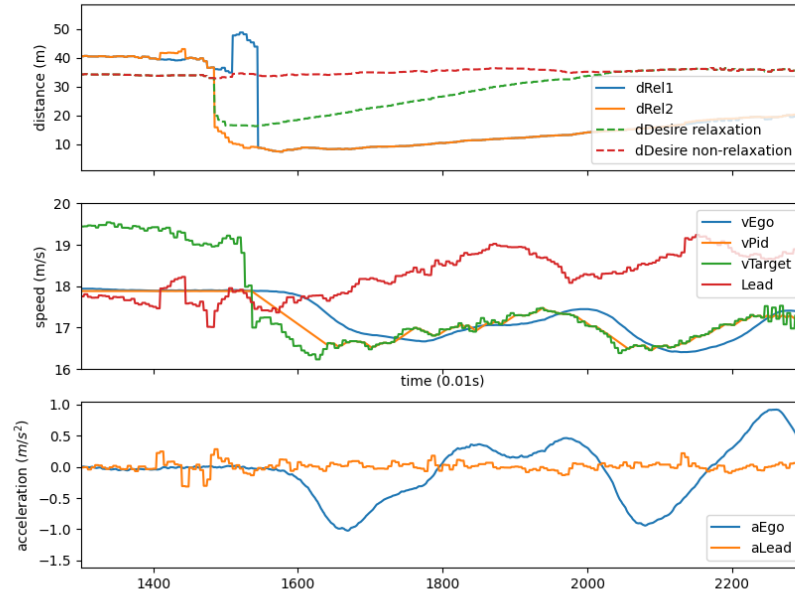


Figure 6: Safety test against the relaxation after a lane change very close to its leader

5. Implications of the relaxation ACC

Having introduced and implemented the relaxation model for factory ACCs, now we are interested to investigate their implications on traffic operation, especially the merge area featuring frequent lane changes.

To do so, we run the same factory ACC code base in simulations, where the radar spacing and speeds are replaced with values from simulation, and a simple function simulates the vehicle actuator according to Zhou et al. (2021b).

We first start investigating the impact on a single vehicle and its followers. Then we extend to the multi-lane traffic where a moving bottleneck causes frequent lane changes.

5.1. Impact on a single vehicle/lane

To show the impact of relaxation on a single vehicle, we compare the responses of a cut-in vehicle that changes to a new lane with the same speed and position; see Fig.7 (a) and (b). Due to the new short spacing after the lane change, the lane changer needs to adjust its speed and regain the equilibrium spacing. The magnitude of the speed perturbation and the time it takes to regain the equilibrium can largely depend on whether the ACC has relaxation or not.

Clearly, we see the non-relaxation experiences larger speed drop in the cut-in vehicle and the followers of the cut-in vehicle. The relaxation ACC shows a smaller speed adjustment but a longer time to regain equilibrium. Also, thanks to the less speed reduction, the relaxation ACC has a shorter final spacing after equilibrium, which prevents more capacity drop than the non-relaxation case.

Note that the congestion wave in Fig.7 (a) and (b) results from the ACC reaction model that takes effect twice, where the cut-in vehicle first responds to the new leader with or without relaxation. Meanwhile, the new follower responds to the cut-in vehicle using the same model. More details in the dynamics and control are shown in Fig.8 and Fig.9 for the cut-in and the new follower.

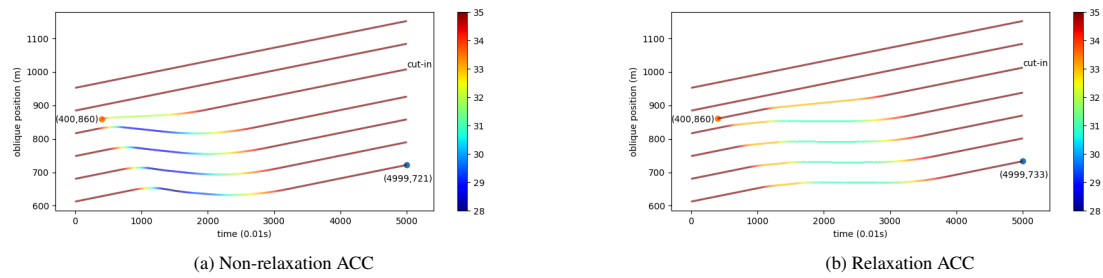


Figure 7: Comparing the trajectory of the cut-in with/without relaxation

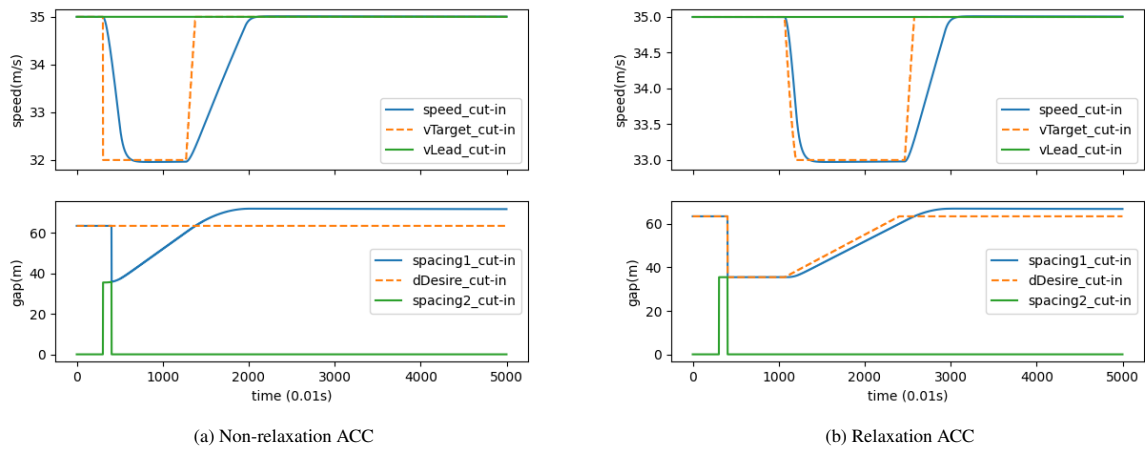


Figure 8: Comparing the control detail of the cut-in vehicle with/without relaxation

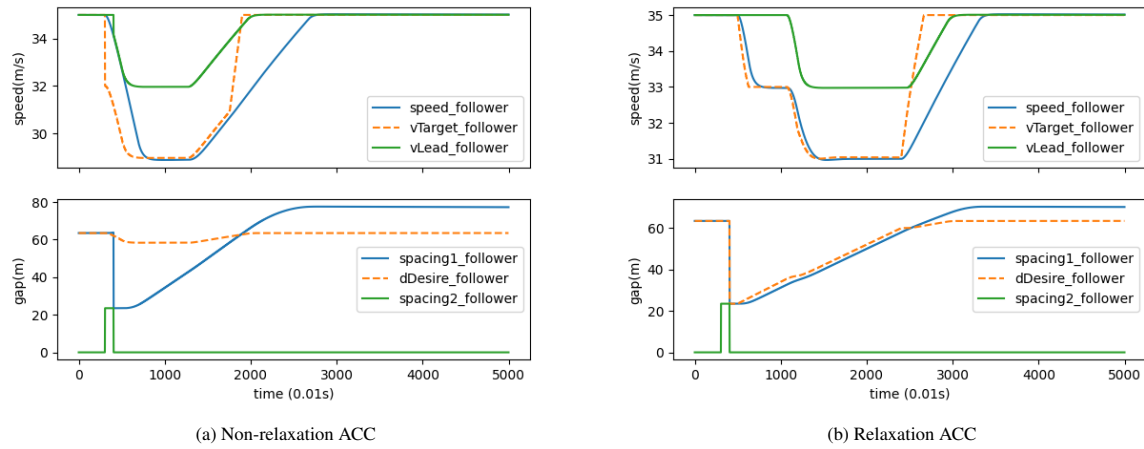


Figure 9: Comparing the control detail of the target lane follower with/without relaxation

5.2. Impact on the merging traffic

5.2.1. Multilane simulation settings

The multi-lane highway consists of 4 lanes (lane 0, 2, 4 and 6) with a free flow speed of 35 m/s, which is also the cruise speed for all ACC vehicles except that a moving bottleneck, the first vehicle on the rightmost lane travels at a lower cruise speed of 30 m/s, which induces lane changes.

The lane-changing model in the simulation is from Laval and Daganzo (2006), where the lanes are separated into road cells. A Poisson process generates lane change particles with acceleration bounds using lane-changing desire proportional to the speed difference between road cells.

Parameter	Value	Description
Desired time headway τ	1.5s	regular following state
Sensitivity factor for distance error k	$2/\tau$	guarantees the marginal SS Zhou et al. (2021b)
Minimum desired time headway τ_{min}	0.5s	relaxation parameter
Maximum relaxation time T_m	20s	relaxation parameter
Acceleration bound	$0.5 + 0.02(35-v) \text{ m/s}^2$	intercept for linear acceleration bound
Deceleration bound d_0	$-1.5 - 0.03(35-v) \text{ m/s}^2$	intercept for linear deceleration bound
Safe rear gap for cut-in	10 m	threshold value for changing lane
Safe lead gap for cut-in	10 m	threshold value for changing lane

Table 1: ACC parameters used in the simulation

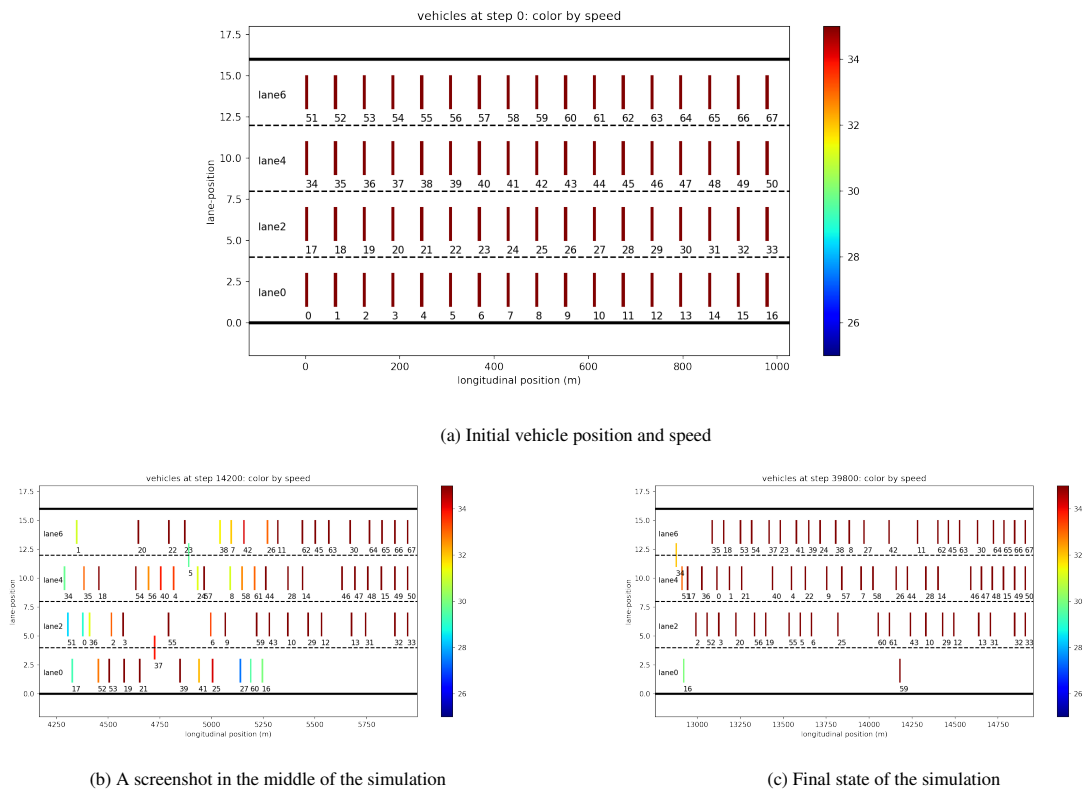


Figure 10: Simulation of multi-lane traffic with lane changes: the lane changes are initialized by a moving bottleneck on the right-most lane (veh16), which reduces its speed from 35 m/s to 30 m/s at the beginning and maintain it throughout the simulation. All vehicles are colored by speeds. Lane 0,2,4 and 6 are main lanes for travel. All vehicles on inter-lanes are temporarily waiting for the safe gap to execute the lane changes.

5.2.2. Impact on lane changes

In this simulation, a moving bottleneck forces the right lane vehicles to decelerate, causing a speed difference that induces lane changes towards the left. In this process, the congestion propagates laterally from the shoulder lane to the inner lanes.

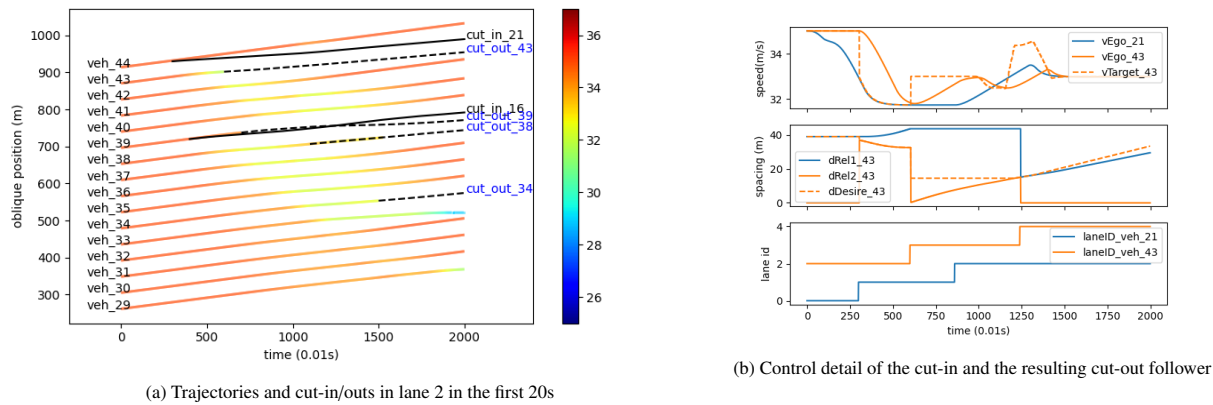


Figure 11: Lateral propagation of lane changes: cut-in vehicles can easily force their followers to change lanes towards faster lanes, leading to the lateral propagation of congestion across lanes.

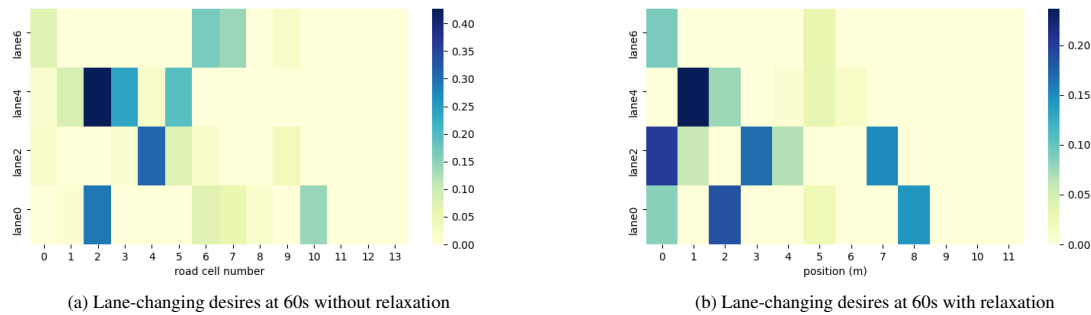


Figure 12: Comparison of lane-changing desires with and without relaxation: screenshot taken at 60s after the beginning of the simulation

Fig.13 compares the number of lane changes throughout the simulation with and without relaxation. The results suggest that relaxation can help suppress the number of discretionary lane changes caused by speed differences. Notice that the slopes of the cumulative curves in Fig.13 indicate the average speed variation between adjacent lanes. In the non-relaxation case, we see a larger slope at the beginning period of the simulation, which indicates the moving bottleneck causes more speed variances than the relaxation case. To provide more evidence, Fig.14 (a) and (b) compare the aggregated speeds at the beginning state of the simulation (60s) in all road cells. The speeds with relaxation are more uniform and higher than the case without the relaxation. This explains the smaller lane-changing desires seen in Fig.12 and the lower number of lane change counts in Fig.13.

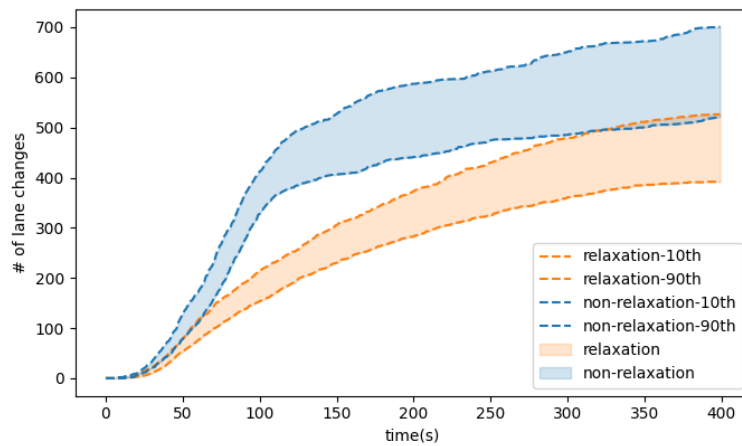


Figure 13: Comparison of cumulative lane changes with and without relaxation: results from 100 iterations for each case

5.2.3. Effects on speed and capacity

Now we further investigate the impact of relaxation ACCs on the speed and capacity of the merge traffic.

First, we compare the speed pattern with and without the relaxation model simultaneously, i.e. 60s after the moving bottleneck takes effect; see Fig.14, where the average speeds of road cells are displayed. The results in Fig.14 (a) and (b) suggest that relaxation improves the average speed level and reduces the variance in speeds, which helps suppress the number of lane changes as shown in Fig.13.

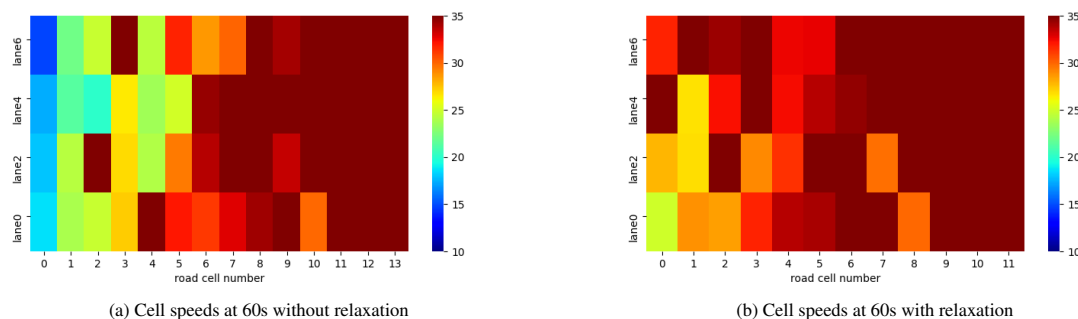


Figure 14: Comparison of lane cell speeds with and without relaxation: screenshot taken at 60s after the beginning of the simulation

We further show the aggregated results to compare the average speed of all vehicles in the merge traffic, and the capacity, measured by the size of voids in the final equilibrium state when the simulation ends.

Due to the randomness in the lane change process, we ran 100 iterations for both the relaxation and non-relaxation cases, where the results are summarized in Fig.15. For capacity, the distributions of void lengths in Fig.15 (a) indicate that the relaxation helps reduces the bubbles in the merge traffic even if the acceleration capabilities are the same. For the speed, Fig.15 (b) shows an obvious edge in speed of the relaxation model over the non-relaxation ACC.

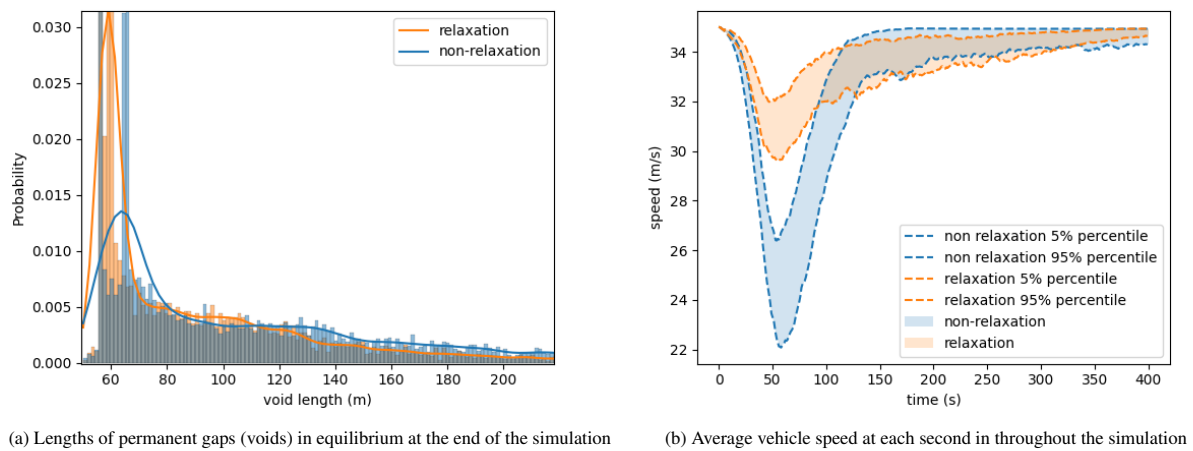


Figure 15: Effects of relaxation ACC on stabilizing the multi-lane traffic

6. Conclusion

The paper investigated the current ACC behaviors before and after a cut-in or lane change happens in comparison with the well-known driver relaxation phenomenon in human driving. Based on the open-source algorithms in Openpilot, we conjecture that current factory ACCs do not have a relaxation design which explains the abrupt speed reductions observed in real life when ACCs are cut in or change lanes.

The study proposed a feasible design to incorporate driver relaxation into the factory ACCs. The design adopts the two leaders detected by the stock radar/cameras in factory ACC units and introduces a relaxed desired time headway. The algorithm is ready for practice and has been tested on a regular market car with the stock ACC hardware.

The study further investigated the implications of the relaxation ACC on traffic operation. Through numerical simulations, the paper finds that: i) incorporating relaxation into ACC help mitigate the speed perturbations in both of the cut-in vehicles and their followers, reducing the magnitude and propagation of the congestion wave in a single lane; ii) relaxation also helps reduce the total number of lane changes in the merging traffic, thanks to the smaller speed differences across lanes; and iii) the relaxation ACC improves both the capacity and the average speed of the merging traffic.

The paper proposes a ready-for-practice relaxation ACC model and showcases its potential in stabilizing lane-changing traffic. While the relaxation ACC model has been tested on real cars, we acknowledge the results from the multilane simulation need more justification. Although we use the same code base as the factory ACC algorithm in Openpilot for the longitudinal control, the lane-changing decision model is subject to error which needs more investigation for future work.

Acknowledgment

This research was funded by NSF Awards #1932451 and #1826162.

References

- , . Introducing comma three. URL: <https://comma.ai/>.
- Cohen, S., 2004. Application of relaxation procedure for lane changing in microscopic simulation models. *Transportation research record* 1883, 50–58.
- Kim, S., Coifman, B., 2013. Driver relaxation impacts on bottleneck activation, capacity, and the fundamental relationship. *Transportation Research Part C: Emerging Technologies* 36, 564–580.
- Laval, J.A., Daganzo, C.F., 2006. Lane-changing in traffic streams. *Transportation Research Part B: Methodological* 40, 251–264.
- Laval, J.A., Leclercq, L., 2008. Microscopic modeling of the relaxation phenomenon using a macroscopic lane-changing model. *Transportation Research Part B: Methodological* 42, 511–522.

- Leclercq, L., Chiabaut, N., Laval, J., Buisson, C., 2007. Relaxation phenomenon after lane changing: Experimental validation with ngsim data set. *Transportation Research Record* 1999, 79–85.
- Li, T., Chen, D., Zhou, H., Laval, J., Xie, Y., 2021. Car-following behavior characteristics of adaptive cruise control vehicles based on empirical experiments. *Transportation research part B: methodological* 147, 67–91.
- Makridis, M., Mattas, K., Ciuffo, B., Re, F., Kriston, A., Minarini, F., Rognelund, G., 2020. Empirical study on the properties of adaptive cruise control systems and their impact on traffic flow and string stability. *Transportation research record* 2674, 471–484.
- Milanés, V., Shladover, S.E., 2016. Handling cut-in vehicles in strings of cooperative adaptive cruise control vehicles. *Journal of Intelligent Transportation Systems* 20, 178–191.
- Montanino, M., Punzo, V., 2021. On string stability of a mixed and heterogeneous traffic flow: A unifying modelling framework. *Transportation Research Part B: Methodological* 144, 133–154.
- Schakel, W.J., Knoop, V.L., van Arem, B., 2012. Integrated lane change model with relaxation and synchronization. *Transportation Research Record* 2316, 47–57.
- Shi, X., Li, X., 2021. Empirical study on car-following characteristics of commercial automated vehicles with different headway settings. *Transportation Research Part C: Emerging Technologies* 128, 103134.
- Stumpf, R., 2020. Consumer reports ranks this aftermarket driver assistance kit above tesla autopilot, cadillac super cruise. URL: <https://www.thedrive.com/news/37833/consumer-reports-ranks-this-aftermarket-driver-assistance-kit-above-tesla-autopilot-cadillac-super-cruise>
- Zheng, Z., Ahn, S., Chen, D., Laval, J., 2013. The effects of lane-changing on the immediate follower: Anticipation, relaxation, and change in driver characteristics. *Transportation research part C: emerging technologies* 26, 367–379.
- Zhou, H., Laval, J.A., Zhou, A., Wang, Y., Wu, W., Qing, Z., Peeta, S., 2021a. Review of learning-based longitudinal motion planning for autonomous vehicles: Implications on traffic congestion. *Transportation Research Board 100th Annual Meeting* Transportation Research Board .
- Zhou, H., Zhou, A., Li, T., Chen, D., Peeta, S., Laval, J., 2021b. Impact of the low-level controller on string stability of adaptive cruise control system. *arXiv preprint arXiv:2104.07726* .