

UC Berkeley

Other Recent Work

Title

ACC, queue storage, and worrisome news for cities

Permalink

<https://escholarship.org/uc/item/5w46m77f>

Authors

Lapardhaja, Servet

Doig Godier, Jean

Cassidy, Michael J

et al.

Publication Date

2024-10-01

DOI

10.1016/j.trc.2024.104809

Peer reviewed

ACC, Queue Storage, and Worrisome News for Cities

Servet Lapardhaja, Department of Civil and Environmental Engineering, University of California, Berkeley

Jean Doig Godier, Department of Civil and Environmental Engineering, University of California, Berkeley

Michael J. Cassidy, Department of Civil and Environmental Engineering, University of California, Berkeley

Xingan (David) Kan, Department of Civil, Environmental, and Geomatics Engineering, Florida Atlantic University

*Published in Transportation Research Part C: Emerging Technologies, v. 167, October 2024,
<https://doi.org/10.1016/j.trc.2024.104809>*

ACC, Queue Storage, and Worrisome News for Cities

Servet Lapardhaja, Jean Doig Godier and Michael J. Cassidy*

Department of Civil and Environmental Engineering, University of California, Berkeley
and

Xingan (David) Kan

Department of Civil, Environmental, and Geomatics Engineering, Florida Atlantic University.

*Corresponding author; Email: cassidy@ce.berkeley.edu

ABSTRACT

Rush-period traffic conditions in two idealized settings are forecast into the future, when most drivers will presumably rely on adaptive cruise control (ACC) while operating their cars. Field experiments emulating the full range of congested conditions confirm that, for a given traffic speed, the spacings for ACC-vehicles tend to be larger than those in present-day congestion, where vehicles are fully human-controlled. These larger spacings mean smaller densities, which mean, in turn, that queues will be less compacted than at present. The queues will therefore expand over greater distances in the future, as more ACC-controlled vehicles enter the scene. These wider-spread, uncompacted queues spell trouble for cities, where queue storage during a rush is often a problem already.

Simulations calibrated to the field-measured data were used to explore this unintended consequence of ACC for various foreseeable futures. Assumptions favorable to ACC were adopted throughout, to produce what are likely lower-bound estimates of future queue-storage problems. These lower bounds served as simple means to address forecast uncertainties. This is because our best-case outcomes for all futures examined are still far worse than the glowing predictions from elsewhere of how ACC may someday eliminate congestion. The first idealized setting was inspired by Downtown Los Angeles, where moderately high congestion already persists during each rush, but where physically long street links help with queue storage. We predict that, owing to ACC alone, rush-period vehicle hours traveled (VHT) on this first network will grow from present-day levels by as much as 12%. In the second setting, inspired by Midtown Manhattan where congestion is already severe and link lengths are short, VHT is predicted to grow by as much as 87%. Higher bottleneck capacities often promised of ACC are shown to be of little value when spillover queues constrain bottleneck flows from reaching those capacities. Adjusting onboard ACC controllers to produce smaller jam spacings was tested through simulation. The tests show how looming problems might be averted by this intervention, and futures thus improved.

Keywords: adaptive cruise control; vehicle automation; urban congestion; queue storage

1. INTRODUCTION

A vehicle equipped with adaptive cruise control, or ACC, comes with onboard sensors that monitor movements of the vehicle ahead, and a control system that automatically adjusts its own vehicle's speed in response to those downstream movements. When activated, ACC technology automatically maintains its vehicle's speed-dependent car-following distances. These are generated based upon the driver's choice of preferred following-distance. Preferences are selected from a menu of short-, medium-, or long settings, with extra-long as an optional setting on some vehicle models. The technology is available on most new cars sold today. It is a component in present-day vehicle automation (Gunter et al., 2020; Makridis et al., 2021; Shang & Stern, 2021), and is an anticipated component for self-driving vehicles of the future (Tesla, 2023).

Calls against activating ACC when traveling in congested, slow-moving traffic have come from select quarters on grounds of safety (GMC, 2023; Honda, 2022; Sparks, 2022; Volvo, 2020). Yet, louder voices seemingly champion more indiscriminate use of this technology. The US Department of Transportation, for example, heralds ACC as means of enhancing "throughput" and various other related measures of traffic performance (USDOT, 2019). The claims suggest that part of ACC's value lies in using it when in congested rush-period traffic, when good performance matters most. Drivers themselves seem amenable to this more unfettered use of the technology: most participants in a recent federal study (USDOT, 2022) report feeling comfortable with all following-distances automatically generated by ACC, even when traveling at speeds as low as 40 km/h, and after having selected the short setting as the preferred following-distance. Moreover, it stands to reason that, having paid for ACC technology, drivers will be eager to use it. And unrestricted use of ACC seems part of a march toward fully automated driving (BMDV, 2021; Walz, 2023). The trend thus seems to be one of broad ACC-engagement, whether in congested or uncongested traffic.

We therefore proceed by assuming that drivers of ACC-equipped vehicles will tend to use the technology, whatever the traffic conditions might be, both in nearer term- and more distant futures. The study then follows a logic that we view as elegantly simple. Responses of a single make and model of ACC-equipped cars powered by internal combustion engines (ICE), and of a single ACC-equipped battery-powered electric car (EV), were measured in field experiments. Measurements were taken of all available settings for preferred following distance and were conducted over a full range of congested vehicle speeds, including the jammed (zero) speed, and for various free-flow speeds. These data were then used to calibrate a well-known car-following model, to replicate field-measured ACC-responses under each setting for following-distance. Parameter values for fully-human-controlled cars were taken from the literature. The calibrated models were used to simulate various future scenarios for idealized versions of two distinct US cities. Assumptions favorable to ACC were adopted throughout. Yet even these best-case outcomes illustrate why favorable forecasts regarding ACC's future impacts on "throughput" or any other measure of traffic performance are not only exaggerated, but that in the absence of intervention, one can expect ACC to add to urban congestion considerably. We demonstrate that degradation of this sort will likely be an unintended consequence of the vehicle spacings that ACC selects in heavily congested and jammed traffic. We further demonstrate that this consequence is likely to occur for most any foreseeable future, and is especially likely in older, denser cities that are already highly congested during a rush and that have more limited queue-storage space. An intervention to avert these looming problems is explored by following the same logic.

1.1 Background

Many of the early studies of ACC's impacts on traffic were theoretical in nature, in that they relied upon car-following models like those in Gipps (1981) and Treiber et al. (2000), which were formulated for fully-human-controlled driving. Parameter values for these models were in some instances calibrated to ACC responses measured in field experiments but were often selected solely on speculation. The models were invariably applied to individual (isolated) bottlenecks. The focus was typically on how ACC impacted the capacity (i.e., the queue discharge flow) of said bottlenecks. Consensus on this matter has remained elusive, as described below.

Works in Vander Werf et al. (2001; 2002) used parameter values that emulated idealized, highly responsive ACC technologies that would prevent the formation of stop-and-go waves in congestion. The literature has since reported that ACC will generate these traffic instabilities (Darbha & Rajagopal, 1999; Gunter et al., 2019, 2020; Milanés & Shladover, 2014; Shang & Stern, 2021). However, Vander Werf et al. (2001; 2002) sought best-case estimates of ACC's effect on bottleneck capacity, which stop-and-go driving has long been thought to diminish. Even for these best-case conditions, the latter-cited works report that the future emergence of ACC will produce little, if any, capacity gains over present-day rates. By instead using model parameters that generate stop-and-go, more recent work in Shang and Stern (2021) predicts that a future in which all vehicles operate with ACC will see bottleneck capacity drop by a staggering 35%.

In contrast, other studies (e.g., Goñi-Ros et al., 2019; Kesting et al., 2008; Papacharalampous et al., 2015; Talebpour & Mahmassani, 2016) report that bottleneck capacities will grow as ACC-vehicles increase in number. Optimistic estimates collectively range from 12% growth when one in four vehicles function under ACC (Kesting et al., 2008) to more than 35% growth when all vehicles operate with ACC (Talebpour & Mahmassani, 2016).

The lack of agreement across studies is notable. One might wonder if it stems from treating model predictions of capacity as ground truth. Doing so came with advantages, e.g., it enabled studies of how ACC's effects may be altered by vehicle lane-changing maneuvers, stop-and-go instabilities, and other traffic features not observable in small-scale field experiments of the kind conducted in Bose and Ioannou (2003), Ntousakis et al. (2015), and Shang and Stern (2021). It also enabled analysts to vary parametrically the mix of ACC- and human-controlled vehicles in traffic, as was done in James et al. (2019) and in most of the other studies cited above.

Of concern, parametric tests in James et al. (2019) reveal to us problems with the car-following models used in many of the previous studies. For example, we see that the calibrated model of Milanés and Shladover (2014) inexplicably predicts that higher bottleneck discharge flows are produced by selecting a longer (not shorter-) preferred following-distance. Another of the models (Xiao et al., 2017) is shown, logically, to predict that higher discharge flows come by selecting the short following-distance, but that these flows steadily diminish as more and more ACC-controlled cars with short distance settings enter the mix, as if too much of a good thing somehow becomes bad. And this same model predicts that ACC will generate impossibly large jam densities.

These predictions of large density would have been good news for cities, had the predictions been realistic. This seems not to be the case, however, as field-measured data from two more recent studies (Shi & Li, 2021; Li et al., 2022) indicate.

These two latter-cited works used field-measured data as ground truth, and not the predictions from car-following models. In both cases, the field data were of ACC-responses measured in small 2- and 3-car platoons operating on real roads. Bivariate plots were constructed of densities and flows, the inverses of the measured car spacings and headways, respectively; and

plots in Li et al. (2022) were augmented with field data collected in Makridis et al. (2021). The plots make clear that driver selections of short- and medium following-distances engendered queue discharge flows that were higher than those observed in present-day traffic. Selecting long-length following-distances led to smaller than present-day discharge flows.

By combining these latter-cited empirical findings with those reflecting driver choices of preferred flowing-distance (Nowakowski et al., 2010),¹ one might conclude that ACC will have favorable impacts on the future discharge flows through isolated bottlenecks. Alas, a practical problem emerges from this: most cities are not plagued by single, isolated bottlenecks, but rather by multiple bottlenecks that can interact in flow-damaging ways when queues grow long (Carlson et al., 2010, 2014; Cassidy et al., 2002; Daganzo, 1998; Daganzo et al., 1999). Any promise of higher bottleneck capacity from ACC might therefore become meaningless when queues that spillover from one street link to the next constrain a bottleneck's flow from reaching the higher capacity. This suggests that a greater determinant of future urban congestion will be ACC's impact on queue expansion, rather than on bottleneck capacity. The news on this front turns out to be bad. Clues that point to this bad news are evident in Shi and Li (2021) and Li et al. (2022).

The density-flow data plotted in both those studies include measurements made at slower speeds commensurate with light to moderate congestion, e.g., see Fig. 5a in Shi and Li (2021) and Fig. 5b in Li et al. (2022). Best-fit lines through those congested data were extrapolated to estimate jam densities. Resulting estimates ran only as high as 90 cars/lane-km in Shi and Li (2021), and as low as 50 cars/lane-km in Li et al. (2022). The estimates characterize ACC-dominated queues to be conspicuously uncompacted, e.g., the lower estimate in Li et al. (2022) would mean that each stopped car occupies, on average, a whopping 20m of longitudinal road space. Estimation errors may have colored these outcomes, perhaps due to the heroic extrapolations made in the absence of severely congested data, and perhaps because some of the lower-speed data included from Makridis et al. (2021) were sampled in short-lived, non-steady-state conditions.

A modestly higher jam density of 101 cars/lane-km was separately estimated in Li et al. (2022) from an unspecified number of jam spacings sampled of an EV, a Tesla Autopilot. Findings from our own field experiments also show that an EV operating under ACC can exhibit larger jammed densities than do ICE-powered ACC-enabled cars. We attribute these larger densities to the EV's regenerative braking system, as will be discussed in sec. 2.3.

Though the above estimates of jam density often seem suspiciously small, they point to something both important and troubling: for a given vehicle speed in congestion, most of the ACC-generated densities appear to be lower than those that occur when vehicles are entirely controlled by human drivers. (This finding was evidently overlooked in earlier studies, with their focus on discharge flows through isolated bottlenecks.) In a prescient response to their finding on density, the authors of Li et al. (2022) both: warn of potential queue storage problems in the future, when ACC-equipped vehicles will come in greater numbers; and call for further study on the matter. The warning is justified, for the reasons discussed below in sec 1.2; and the present paper answers the call for further study and does so in the manner described in sec. 1.3.

1.2 The Concern

Problems are known to occur when vehicle queues grow long during a rush and spillover to multiple links upstream, e.g., see Daganzo (1998). Spillover queues on freeways, for example, can block off-ramps and starve them of exit flows, thus adding to system delay (Cassidy et al., 2002; Newell, 1993). Still greater problems can occur on city streets where long queues wrap themselves

¹ The reference reports that 70% of ACC-equipped drivers select either the small- or medium following distances.

around what are often short city blocks, to become like a snake that has swallowed its tail. The queues spread to other links and trigger a vicious cycle of worsening congestion over time, with ever-diminishing trip-completion rates and skyrocketing delays, a process called gridlock (Daganzo, 1996, 2007; Mahmassani et al., 2013).

As alluded to in Li et al. (2022), the lower estimates of density (i.e., higher vehicle spacings) in congestion mean that ACC-dominated queues in the future will be less compacted than at present and will thus expand over greater distances. All else equal, these greater queue expansions will mean more problems of the kind described above.² With this unintended consequence in mind, we now answer the call in Li et al. (2022) for further study on the matter.

1.3 Study Scope

The present paper responds to the call, in effect, by starting over, from “square one.” New field experiments were performed, since the very low estimates for severely congested and jammed densities from previous experiments appear suspicious to us. As will be described in the following section, the present experiments were separately conducted in two sets. The first entailed a 3-car platoon of ACC-equipped cars of a single make and model, each one ICE-powered. The second set featured a 2-car platoon of ACC-equipped, battery-powered EVs. The tests were akin to earlier field experiments like those of Shi and Li (2021) and Li et al. (2022), but this time the field data included measurements at lower speeds commensurate with more severe congestion, along with numerous measurements of jammed conditions; and all data were extracted from steady-state conditions. The resulting estimates of queued densities are mostly larger than those reported in Shi and Li (2021) and Li et al. (2022). Still, our new density estimates are typically smaller than those from present-day, fully-human-controlled traffic. The new densities are evidently low enough to create damaging spillover queues as ACC market shares grow in the future.

Future damage will be illustrated for city streets. Their shorter link lengths (relative to freeways) render them especially vulnerable to spillover problems caused by less-compacted queues. Rather than relying upon models that appear suspect in James et al. (2019), assessments were performed using the car-following model formulated in Gipps (1981) and resident in the Aimsun microscopic traffic simulation platform (Aimsun, 2017). As will be described in sec. 3, the parameters of this model were separately calibrated to replicate the fundamental diagrams for seven distinct ACC classes, these being: each of three settings for preferred following-distance available on the ICE-powered cars used in the present field experiments; and the four settings available on our lone EV from which measurements were extracted. The same resident car-following model with suitably calibrated parameters was used to emulate fully-human-controlled cars. When possible, the calibrated models for each car class were constrained, to generate queue discharge flows and queued densities that were very close to the corresponding values estimated from the field data. Failing this, the models were calibrated to generate values favorable to ACC, to engender conservative predictions of future problems.

The seven ACC classes were simulated in sec. 4 for two distinct case-study sites: one an idealized version of Downtown Los Angeles, the other of Midtown Manhattan. The simulations were performed for each site over a wide range of plausible future patterns of technology adoption.

² Metaphorically speaking, one might wish to compact vehicle queues, to store more of them on a street link, much as one might compact trash, to store more of it in a waste bin. Overflows cause problems, whether from waste bins or congested street links.

To be conservative, assumptions favorable to ACC were adopted wherever possible.³ The impacts thus predicted reflect best-case outcomes, but still signal impending problems for wide-ranging cases. An intervention to avert these problems is evaluated in sec. 4 as well. Caveats to the present study and future research needs are discussed in sec. 5.

2. FIELD STUDIES

The experiments measured how the ACC-controlled test cars adjusted their spacings in response to their leaders, while traveling on real roads and a highway.⁴ The tests were conducted using the platoons previously described, each consisting of an ACC-equipped leader and either the two ICE-powered ACC-followers or the single EV, also ACC-equipped. Experiments took place only during off-peak times when traffic was light. In this way, platooned cars could safely alter their speeds in unison, to emulate entries into, and exits from, wide-ranging congestion states.

In each experiment, the platoon leader initially traveled at a high speed that was pre-selected from a menu of free-flow rates, each to emulate what might occur on an uncongested facility of a particular type. The leader then decelerated to a lower speed that had also been selected *a priori*, this time from a menu reflecting a full range of congestion levels. The leader eventually returned to the initial high speed and the entire sequence of prescribed speeds was regulated by the leader's ACC controller onboard. The term "cycle" will henceforth denote one sequence of high-low-high speeds.

Each prescribed combination of high and low speeds was executed under each of the available options for preferred following-distance; and each set of combined speeds and preferred following-distance was performed for at least 12 cycles. The steps described below were taken to ensure that estimates were of steady-state conditions.

Each high- and low speed was maintained for 10s or more.⁵ Measurements to be described in secs. 2.1 and 2.2 were recorded at short intervals (of 0.04s or 0.2s), from which vehicle trajectories were constructed. Average vehicle spacing and headway were jointly determined for each period when the speed of a leader and its follower was approximately the same.

A total of 1,264 joint measurements of average spacing and headway (and thus of their inverses, density and flow) were obtained in the above fashion from across all experiments. Certain details for the experiments involving the ICE-platoon were different from those involving the platooned EV. These details are described below, starting with the ICE experiments.

³ Favorable assumptions include: (a) no future growth in travel demands via cars that would on their own add to future congestion; (b) optimistic EV adoption rates, which by our observations would be less damaging to traffic flows for reasons demonstrated in due course; (c) no short-length turn lanes that would have been especially susceptible to queue spillovers; (d) the smallest value for present-day, human-generated jam density found in the literature, to minimize the problematic differences with still smaller, ACC-induced densities; and (e) still other ACC-favoring assumptions to be described in due course. In further deference to ACC, the estimates of congested densities in Shi and Li (2021) and Li et al. (2022), which are even smaller, and thus more worrisome than ours, were excluded from present analyses.

⁴ Conducting the field experiments on distinct facilities was not viewed as a concern, since ACC responses are reportedly independent of road type (Xiao & Gao, 2010).

⁵ Slightly shorter durations were often used when entire platoons were brought to a stop, since conditions were thereafter not changing with time and were thus fully in steady state.

2.1 ICE-Platoon (Toyota Corollas)

This set of experiments involved the 3-car platoon and was comprised entirely of ACC-equipped, ICE-powered Toyota Corollas of model year 2020. Measurements were extracted from the platoon's second and third test cars. Two facilities in South Florida were used.⁶

Experiments began on the shoulder of a facility, where the second and third Corollas were each physically separated from its respective leader by 1.4 m, as measured from leader's rear to follower's front.⁷ Once in a travel lane, the leader then: accelerated to the speed of 88 km/h; decelerated to one of the lower speeds from the menu of {72, 56, 40, 24, 0} (km/h); and returned to the initial speed of 88 km/h, so that a new cycle might ensure shortly thereafter.⁸ Speedometer readings for each of the three platooned cars were automatically recorded at 0.2-s intervals. Resulting estimates of the distances traveled over each time step were recorded using Onboard Diagnostics II data loggers.

A re-initialization to re-separate platooned vehicles (by 1.4 m) was performed following every two cycles. This, to limit systematic measurement errors in each vehicle's speedometer, which might have accumulated with distance traveled.

2.2 EV-Platoon (Hyundai IONIQ 5)

The second set of experiments featured a 2-car platoon of ACC-equipped EVs; the leader a Toyota Camry (model year 2021) and the follower a Hyundai IONIQ 5 (model year 2022). Measurements were extracted from the latter. It came with an additional setting for preferred following-distance, that being: extra-long. Isolated, low-trafficked stretches of three public roads in Northern California were used this time around for measuring ACC responses.⁹

The lead vehicle decelerated from a speed that had been pre-selected from a menu of {95, 88, 72, 56} (km/h) to one of the lower speeds from the menu of {72, 56, 40, 24, 0} (km/h), and eventually returned to its earlier high speed. Lower-cost GPS devices produced by Racebox had by this time entered the marketplace and were now used to measure each vehicle's geographic coordinates at 0.04-s intervals. These more precise devices eliminated the need for periodically re-initializing vehicle separations. Since platoons were now composed of 2 cars rather than 3, measurements were extracted over additional cycles, so that data points were of comparable number in each set of (ICE and EV) experiments.

2.3 Data Inspection

A first order of business entailed looking for accumulative errors in our measurements of the ICE-platoon. Concerns were dispelled by Fig. 1. It presents median values of the densities and flows estimated from all first cycles (shown with unshaded data points) and second cycles (darkened points). Note that these data are further partitioned by both steady-state vehicle speeds (shown via light lines radiating from the origin) and preferred following-distance. Best-fit polynomial curves are included in the figure to aid in visual inspection.

The figure makes clear that paired unshaded and shaded (i.e., first- and second cycle) data often exhibit differences, both in their densities and flows. But these differences are not large

⁶ These facilities were: (1) an 8-km stretch in both directions of US Highway 441 (between Atlantic and Boynton Beach Blvd.) in Delray Beach; and (2) a 4-km stretch in both directions of Flying Cows Rd. (between US Highway 98 and Rustic Rd.) in Wellington.

⁷ The 1.4m distance was chosen arbitrarily, as a matter of convenience: that distance could be easily measured using the rear cameras on our test cars.

⁸ The platoon decelerated to zero only when traveling on Flying Cows Rd., but not while on US 441.

⁹ These were roughly 10-km stretches of Pendrick, Robben, and Sikes Rds. in the rural community of Dixon.

compared against the displacements between the data points and their best-fit curves. In short, the ACC responses display random variation, such that the visible differences between first- and second cycle estimates shown in Fig. 1 are to be expected and do not suggest systematically accumulating errors in measurement. We therefore used all the data collected in the field experiments, including those from the ICE-platoon extracted in second cycles.

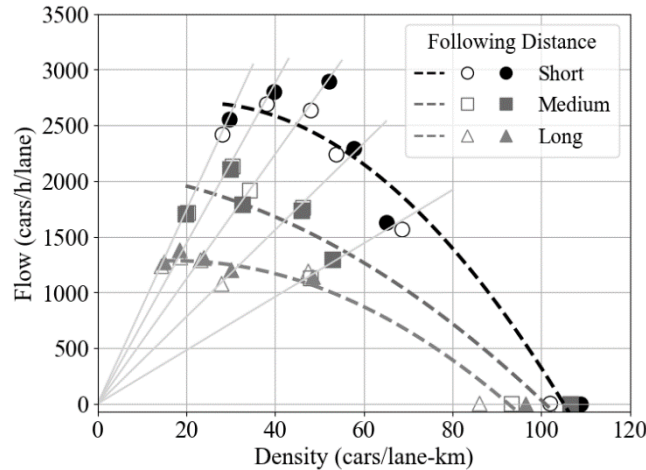


Figure 1. Median measurements for Cycle 1 (unshaded) and Cycle 2 (shaded) ICE Experiments with Toyota Corollas

All the field-measured data are displayed in Figs. 2a and 2b for the ICE- and EV-platoon, respectively. The slightly enlarged, shaded data points are the median values for each cluster of data of the same vehicle speed and preferred following-distance. Best-fit curves through these medians are shown.

Visual inspection of Fig. 2a reveals that the selection of the short following-distance produced an average queue discharge flow for our ICE-powered vehicle model (the apex of the figure’s darkest concave-shaped relation) that is greater than those reported of present-day traffic, e.g., Seeherman and Skabardonis (2013); TRB (2016). This is consistent with findings in Shi and Li (2021) and Li et al. (2022). Figure 2b shows that the EV also displayed the highest discharge flow when the short following-distance was selected. The EV’s discharge flow was also high when the following-distance was set to medium.

Turning attention now to jam densities, Fig. 2a reveals that estimates for the ICE-platoon are 106, 97, and 88 cars/lane-km, depending upon the preferred following-distance selected. Most of these estimates are larger than those in the two earlier-cited works but still smaller than what is commonly reported for traffic comprised solely of fully-human-controlled vehicles (Chiabaut et al., 2009; Hoogendoorn et al., 2015; Knoop & Daamen, 2017; Lárraga & Alvarez-Icaza, 2010; Li et al., 2022; Rakha et al., 2008). Note how the present estimates of jam density get smaller in Fig. 2a as the settings for following distance get larger. Hence for the ICE-powered ACC-car model tested here, settings that produce lower queue discharge flows also produce lower congested densities, and thus have greater potential for queue spillovers.

The pattern is somewhat reversed for the EV presently tested, i.e., the selection of shorter following-distances produced lower jam densities, as is evident in Fig. 2b. Selection of the short- and medium distances produced jam densities akin to those of the ICE-cars when under the same

distance settings.¹⁰ But setting the EV at long- and extra-long following-distances respectively produced jam densities of 130- and 133 cars/lane-km, values nearly on par with what is observed in human-controlled traffic. This fortunate outcome seems linked to the EV’s regenerative braking system, which is known to govern and EV’s braking when performed gradually (Chau, 2014).

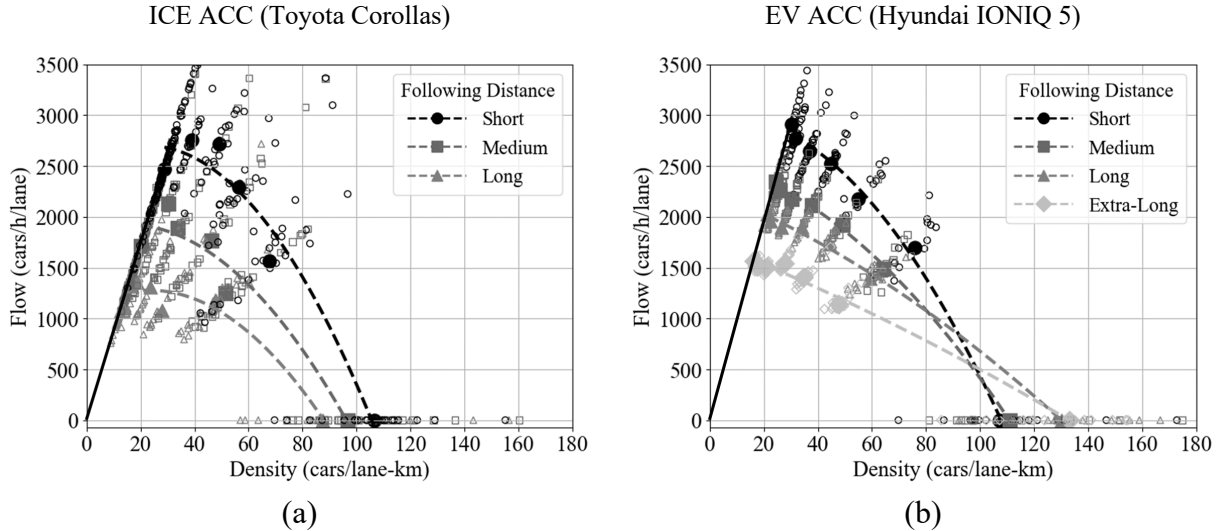


Figure 2. Fundamental diagrams and field data: (a) ICE Toyota Corollas, (b) EV Hyundai IONIQ 5.

The mechanism and outcome are evident in Figs. 3a and b. The paired trajectories in each figure trace the EV platoon, as both vehicles decelerated to full stops during the field experiments. Figure 3a depicts the platoon when the follower operated under the short setting for preferred following distance. Note from the figure how the follower came to a complete stop only 0.9s after its leader had done so. The jam spacing thus produced was 10.3m, as annotated in the figure.

In contrast, Fig 3b shows what occurred when the follower EV operated under the extra-long following-distance, and thus had the luxury of decelerating more gradually, and as governed by the vehicle’s regenerative braking. Notice this time the follower’s more prolonged deceleration: it rolled to a full stop a much lengthier 6.3s after that of its leader, and produced a more compacted jam spacing, 7.4m, relative to what was observed when the short following-distance was selected.

We cannot say whether ACC controllers on all EV makes and models exhibit this beneficial feature. Our inspections of the previously existing data sources cited in sec. 1.2 shed little light on the matter. Still, the present finding shows that it is possible to control EVs (and perhaps also ICE-powered, ACC-enabled cars) to bring about more compacted queues. We will exploit this finding in analyses to come.

¹⁰ In contrast to what is observed in Fig. 2a, the EV exhibits a jam density that is slightly larger under the medium following distance than under the short one. We attribute this (modest) difference to the (perhaps equally modest) idiosyncrasies of ACC technology. Moreover, the relatively small difference between the EV’s jam spacings under short and medium settings might be attributed to higher granularity: the EV came with a fourth (i.e., extra-long) setting not available on the ICE-powered cars.

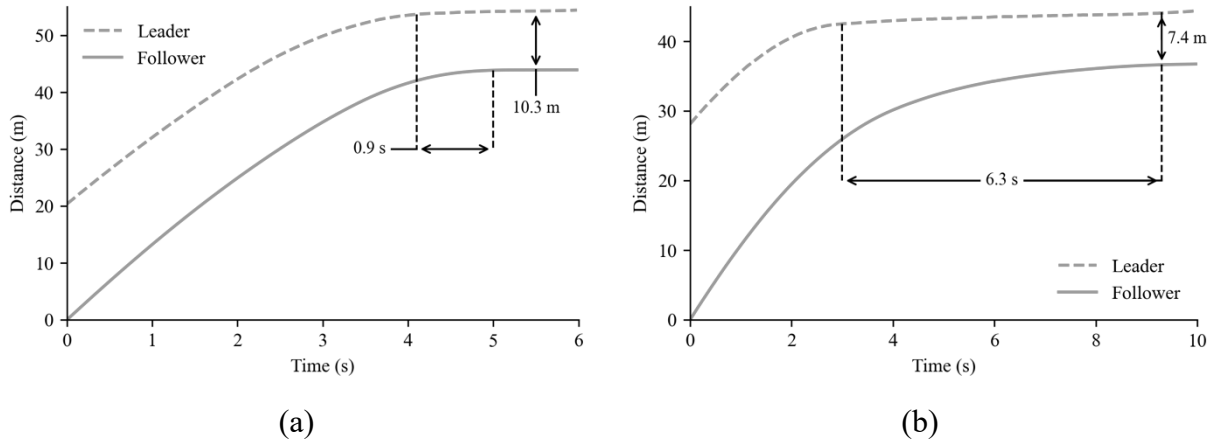


Figure 3. EV Trajectories (Hyundai IONIQ 5): (a) short following-distance, (b) extra-long distance.

3. MODEL CALIBRATION

Parameters in the Aimsun microscopic simulation platform (Aimsun, 2017) were calibrated to match the field-measured data described in sec. 2. The Toyota Corollas introduced in sec. 2.1 were henceforth used to represent ICE-powered cars in general and the Hyundai IONIC 5 (sec. 2.2) to represent EVs. To emulate human-controlled cars, for which Aimsun is highly reputed (Ahmed et al., 2021; Panwai & Dia, 2005), parameters were carefully selected from the literature, as will be described momentarily. Hence, our choice of model platform enabled simulations of both present-day conditions and of futures when ACC (and EVs) presumably will claim steadily greater market shares (Calvert et al., 2017; Litman, 2020; Tillema et al., 2017). The platform enabled this in a relatively simple fashion, and with good calibration results.

The calibration process used to model ACC-engaged cars is described in sec. 3.1, along with discussion regarding our parameter choices for human-controlled cars. The full set of specially calibrated parameters used in our simulations is given in sec. 3.2. Comparisons between Aimsun’s resulting predictions and the field data are furnished in sec. 3.3.

3.1 Process

Calibrating Aimsun’s car-following parameters to emulate our ACC-engaged test cars was painstakingly performed for cases of homogeneous traffic, meaning that: model parameters were selected via separate simulation runs, always of traffic comprised entirely of one car type (ICE-powered or EVs), and for which all simulated cars were set to a single, same preferred following-distance (short, medium, long, or extra-long in the case of EVs). Runs were performed for each feasible combination of car type and following distance, such that parameters were separately calibrated for 7 cases of homogeneous, ACC-controlled traffic.

In each case, parameters were fine-tuned via trial-and-error over many runs. In addition to its advantage of simplicity, our trial-and-error approach facilitated delicate, manual adjustments to parameter values, so that model predictions typically matched field estimates of queue discharge flows and queued densities. Failing this, the approach enabled us to calibrate models that produced conservative (ACC-favoring) predictions for these key traffic variables; see Fig. 4 and its

discussion in sec. 3.3.¹¹ Each run’s simulated outcomes were generated for a single travel lane of 1-km length. Cars entered at the upstream end, at a rate approaching the lane’s capacity. A speed limit was posted at the lane’s midpoint, to emulate slowdowns that could have been caused by congestion from a bottleneck downstream. The speed limit varied parametrically at rates of {3, 6, 15, 25, 40, 50} (km/h), to mimic varying congestion levels. Congested densities and flows were jointly measured via a simulated vehicle detector placed just upstream of the posted speed limit.

Four parameters were calibrated via this process since these affect the shapes of the Aimsun-generated fundamental diagrams. Three of these parameters were: the “clearance” parameter (in meters), which dictates jam density and is defined as the average jammed distance between a leader’s rear and its follower’s front; and two other unitless parameters called “sensitivity” and “headway aggression,” which dictate the shape of a fundamental diagram’s congested branch. The fourth parameter was driver reaction time (in secs). For all simulations, this latter parameter was altered from the default value of 0.8s to 0.66s, so that Aimsun could replicate the higher queue discharge flows observed by ACC-engaged test cars when smaller following-distances had been selected.

Human-controlled traffic was constrained to match nearly piecewise-linear, triangular-shaped fundamental diagrams with a free flow speed of 50 km/h, distinct average queue discharge flows of 1,900- or 1,650 cars/h/lane, to describe each of our two distinct urban settings (TRB, 2016), and a jam density of 138 cars/lane-km. The latter was the smallest value found in the literature (Chiabaut et al., 2009). Larger values found in (Hoogendoorn et al., 2015; Knoop & Daamen, 2017; Lárraga & Alvarez-Icaza, 2010; Li et al., 2022; Rakha et al., 2008) were not used, so that our predictions of ACC’s deleterious effect on jam density would be the smallest within reason. Our choice for the human-controlled jam density was therefore a conservative one, in that it limited any added queue storage problems predicted under ACC-dominated futures.

3.2 Parameter Estimates

The numeric values ultimately selected for the three case-varying parameters are displayed for each homogenous-traffic case in Table 1. Also shown for reference in the table’s final row are Aimsun’s default parameter values. Bold italics are used in the table to highlight values that are different from their respective defaults. The final column in Table 1 presents the maximum acceleration rates used for each car type since these are known to differ across ICE and EV technologies. A value of 4.5 m/s² was selected for EVs based on Hyundai (2022), and commensurate with their higher acceleration capabilities. The default value of 3 m/s² was used for ICE-powered cars.

3.3 Comparisons

The fundamental diagrams (FDs) produced via simulations are included in Figs. 4a – d. Those for the ACC-controlled EV are shown across two figures (4b and 4c) to avoid clutter. The shaded data points in the figures are the averages of 10 simulations. Note from Figs. 4a – c that simulated FDs (shown with solid curves) tend to match rather nicely their empirically-estimated counterparts (dashed curves) that are reproduced from Figs. 2a and b. The largest discrepancies occur between simulated and empirical FDs for the EV when set at the extra-long following-distance; see Fig. 4c. Still, the congested branch for this simulated FD lies on or above the empirical branch, indicating that our analyses to come will estimate ACC’s negative impacts conservatively. (For a given

¹¹ In contrast, more systematic approaches to model calibration have often produced poor results; see Cheng et al. (2021); Kesting and Treiber (2008).

congested density, simulated flow will never be smaller than the empirical rate.) Moreover, Fig. 4c unveils that good matches were obtained between field-estimated and simulated jam densities.¹²

TABLE 1 Input Parameters (ICE: Toyota Corollas; EV: Hyundai IONIQ 5)

Vehicle Type (Following - Dist. Setting)	Clearance (m)	Sensitivity Factor	Headway Aggressiveness	Acceleration (m/s ²)
ICE ACC (Short)	5.2	0.98	1	3
ICE ACC (Medium)	5.5	1.16	0	3
ICE ACC (Long)	6.4	1.62	-0.75	3
EV ACC (Short)	4.5	0.96	1	4.5
EV ACC (Medium)	4.2	1.10	0	4.5
EV ACC (Long)	3.0	1.20	-0.75	4.5
EV ACC (Extra-Long)	2.9	1.50	-1	4.5
Human Driver (high discharge flow)	2.3	1.10	-1	3
Human Driver (low discharge flow)	2.3	1.35	-1	3
Aimsun Default	1	1	0	3

As per our input to Aimsun, the jam density for the human-controlled FD in Fig. 4d is larger than those for the ACC cases in Figs. 4a – c. To highlight the differences for the reader, the lightly shaded area underneath the FD (with 1,900 cars/h/lane as the queue discharge) in Fig. 4d is reproduced in the three other figures. From visual inspection of Figs. 4a and b, note the smallness of the ACC-generated jam densities relative to the estimate for fully-human-controlled traffic. Note too how these differences become small only in Fig. 4c. This is a harbinger of problems likely to come, should future traffic include ACC-vehicles that are not EVs set at long- and extra-long following-distances. Examples are offered next.

4. ILLUSTRATIONS

This section offers a flavor of the impacts that one might expect of ACC in nearer term- and more-distant futures. Recall that we do so by assuming that ACC-engaged ICE-powered cars and EVs function identically to our field-tested Toyota Corollas and the Hyundai IONIQ 5, respectively.

Further recall that our illustrations are loosely inspired by two urban settings. The one inspired by Downtown Los Angeles comes with a street network comprised of comparatively longer block lengths built with car travel in mind. The other setting inspired by Midtown Manhattan has shorter blocks characteristic of older cities. Street geometry for both idealized settings was kept simple, again in the pursuit of conservative predictions regarding ACC-induced future problems. Included among these simplifications is the absence of any turn pockets, within which less-compacted, ACC-queues of turning-traffic might have readily spilled-over to block through-moving flows in adjacent lanes.

¹² Readers might note using Table 1 how higher values for Aimsun’s “clearance” parameter produced higher jam spacings and thus lower jam densities.

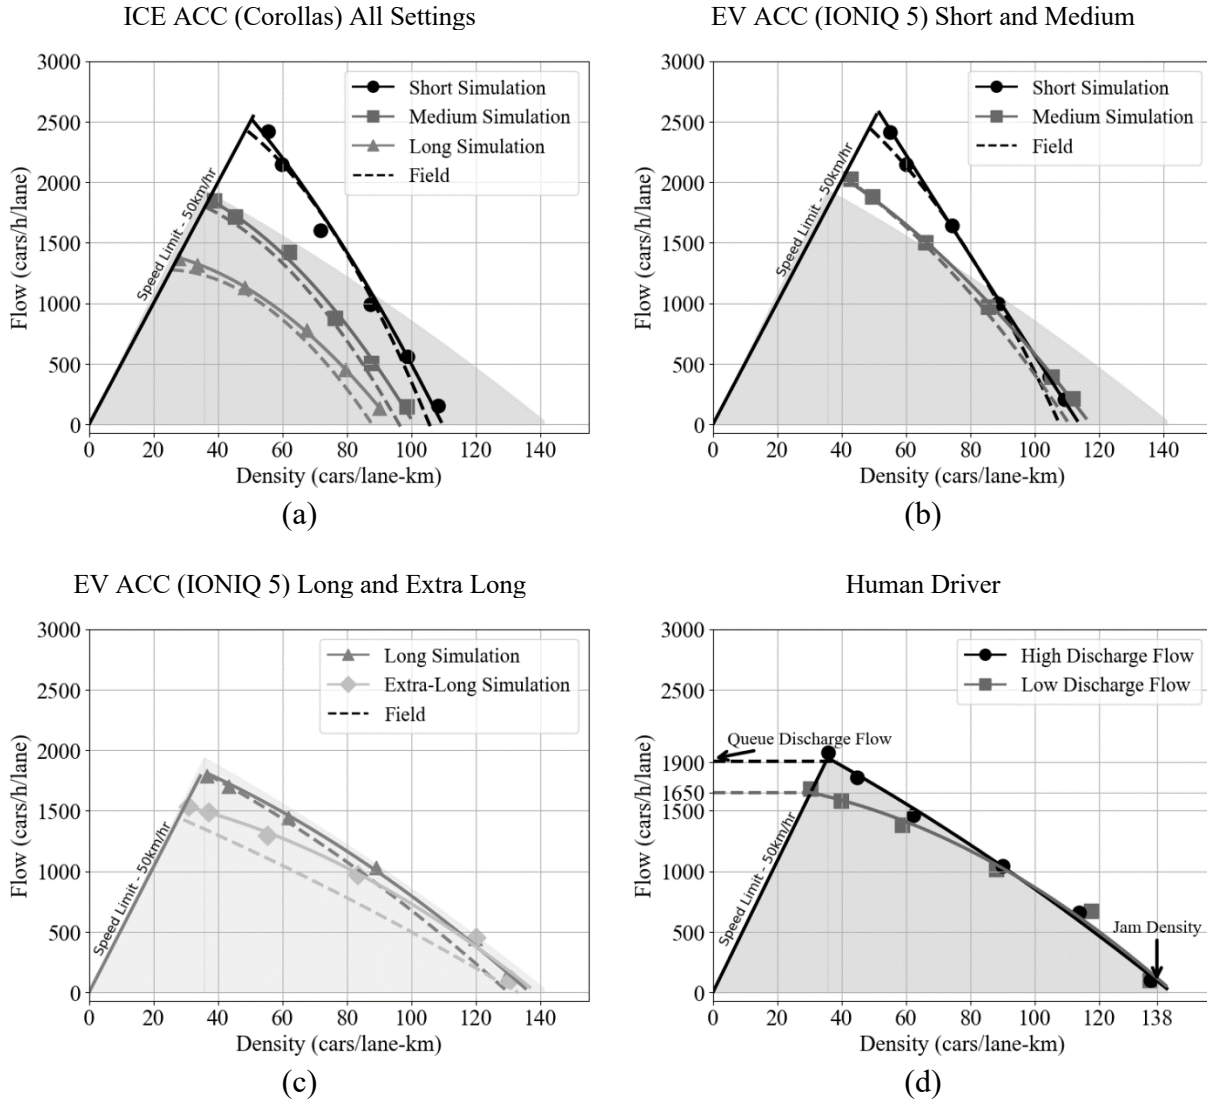


Figure 4. Simulated and empirical FDs. (a) Toyota Corolla ICE-ACC, (b): Hyundai IONIQ 5 EV-ACC, short and medium following-distances, (c) Hyundai IONIQ 5 EV-ACC, long and extra-long distances, (d) Human-controlled traffic.

Forecasted travel demands are similarly conservative. For example, the demands were set to approximately reproduce present-day congestion levels in Los Angeles and Manhattan during the morning rush and were assumed not to grow in future years. Our forecasts regarding ACC adoption behavior reflect guesswork. Yet, we shall see why reasonable variations in our guesses would not likely skirt the sorts of problems that we predict for nearer-term and more-distant futures.

Inputs to our illustrations are presented in secs. 4.1 – 4.3. Disquieting predictions for intervention-free futures are in sec. 4.4. Discussion on why problems should still be expected for other forecasted (guessed-at) inputs is offered in sec. 4.5. Findings that point to a promising intervention are in sec. 4.6.

4.1 Site Geometries and Infrastructure

The street networks for both idealized settings were assumed to form a square grid consisting of two perpendicular sets of 20 parallel, bidirectional arterials with two travel lanes in each direction; see Fig. 5 and its inset. Turning maneuvers were permitted from each approach. Centerline spacing for all street links was 210m and 170m for the Los Angeles- and Manhattan-inspired sites, respectively. Each link had a posted speed limit of 50 km/h. Recall that queue discharge flows for fully-human-controlled cars were taken to be 1,900 cars/h/lane for links on the Los Angeles-like grid and 1,650 cars/h/lane for those on the Manhattan-inspired setting. Discharge flows for ACC-cars were as estimated; see again Figs. 4a – c.

Pretimed, two-phase traffic signals (serving N-S and E-W movements) were assumed to control every intersection on the grid. The common cycle length was long, 90s, which is good practice when addressing congested traffic (see Koshi, 1989). Equal green splits were used (i.e., green- and amber times were 45s in each phase), to accommodate each network’s symmetry. These greens were synchronized with zero offsets, which is known to work well in congested settings (Daganzo & Lehe, 2016).

4.2 Travel Demands

To roughly recreate real-world conditions presently observed during the morning rush, each of our idealized city’s congestion level was defined as its maximal rise in average vehicle pace (the maximum reduction in rush-hour speed), normalized by the average free flow pace. These were measured with Google Maps as in (Doig et al., 2024) for a large sample of real morning commute trips on city streets in Los Angeles and Manhattan. The real-world congestion level thus estimated for Los Angeles was 2.7 and the estimate for Manhattan was 5.0.

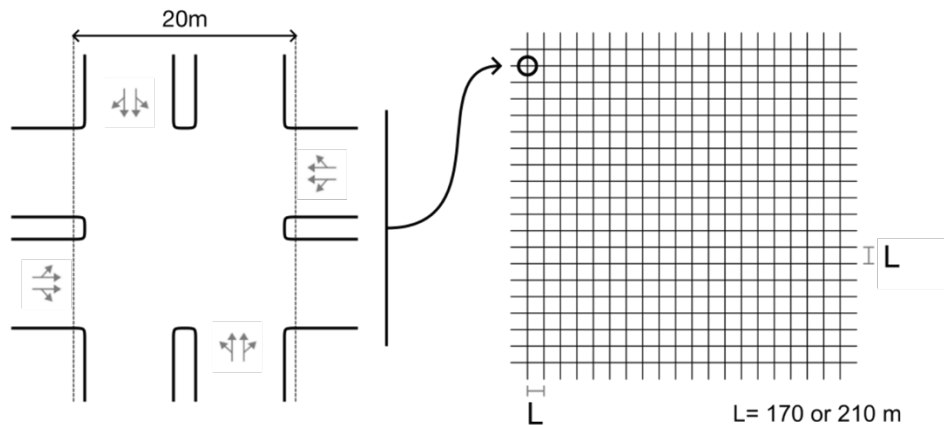


Figure 5. Network grid with intersection geometry shown in inset.

Aggregate demands for the idealized versions of both cities were assumed to follow the time-varying pattern displayed by the cumulative count curves in Fig. 6. Note how the rates changed: from lower to higher at time 20 mins from the simulation’s start; back to the lower rate at time 60 mins; and to zero at 80 mins. The termination of demand at time 80 mins allowed all cars to be served by the end of each simulation. This, in turn, enabled us to assign finite travel times to all cars and thus draw comparisons across scenarios more effectively. The full cessation of demand also truncated lingering congestion at the end of a rush and is yet another feature of our

analyses that works in ACC’s favor. The time-varying demand rates were scaled for each setting so that, when coupled with a time-invariant OD matrix of trip fractions to describe the spatial pattern of demand, the simulated congestion level for each idealized city matched its corresponding real-world estimate from Google Maps.

All simulated origins and destinations occurred at intersections on the idealized grids. Trip fractions were the same for every OD in the matrix (i.e., OD pairs were uniformly and independently distributed over space), and trips between all OD pairs followed a Poisson process of equal rate. Each simulated vehicle entered the network on one of the four departing links of its origin intersection, when safe entry was possible and entry delays are included in all analyses to come. Upon reaching their destination intersections, vehicles were immediately removed from the network, with no added time spent looking for parking.

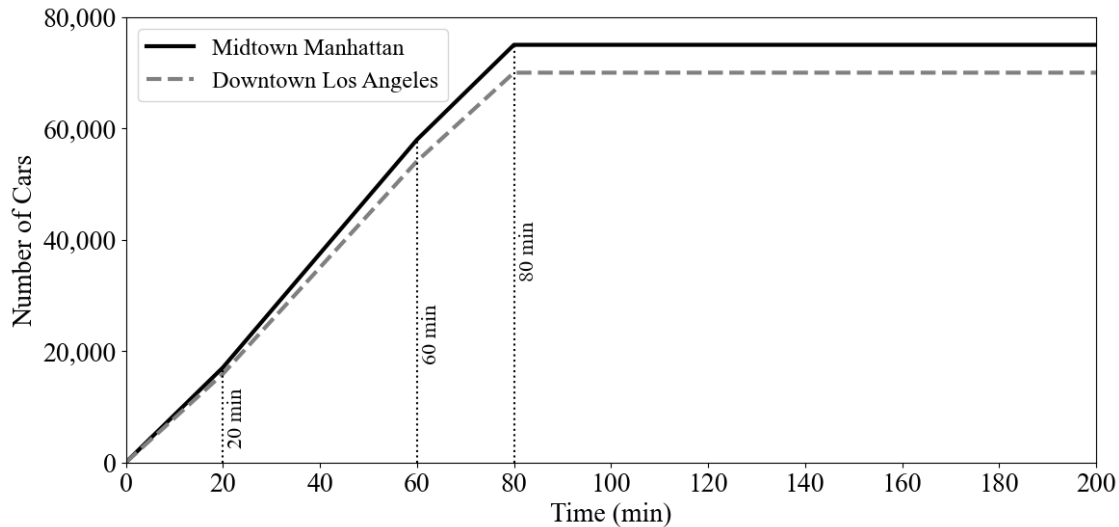


Figure 6. Time-varying aggregate demands

4.3 Adoption Forecasts

We reasonably assume that adoptions will grow with time, not only for ACC (Calvert et al., 2017; Litman, 2020; Tillema et al., 2017) but for EVs as well, such that ICE-powered cars will presumably decline in number (CARB, 2022; Choi, 2021; Lambert, 2021). Our illustrations will follow the guessed-at projections displayed in Fig. 7. We view these as optimistic, in the sense that other guesses, involving greater numbers of ICE-powered cars, for example, produced bleaker prognostications, though not shown here for brevity. Our optimistic guesses at various futures thus further facilitate conservative forecasts of ACC’s damages.

Note from Fig. 7 that our projections entail (i) a nearer-term future in which half of all cars are equipped with ACC and half of these are EVs; (ii) an especially optimistic, more distant future in which 90% of all cars are EVs equipped with ACC, and (iii) a more modestly optimistic, longer-run future in which 80% of cars are of this latter class. Simulated outcomes for all these futures are compared against present-day conditions for which all cars are assumed to be fully human controlled.

Recall that drivers of ACC-equipped cars are assumed always to engage the technology. These drivers are further assumed to select a preferred following-distance that does not change during a trip. The distribution of preferred distances for our ACC-driver populations was taken from the empirical study in Nowakowski et al. (2010), which to our knowledge is the only one of

its kind. In it, 50.4% of ACC-drivers reportedly select the short following-distance when driving their cars, 18.5% the medium distance, and 31.1% the long setting. It was assumed for the present analyses that drivers of ACC-controlled EVs select the long- and extra-long settings with equal propensity and total 31.1% of all ACC-equipped EVs.

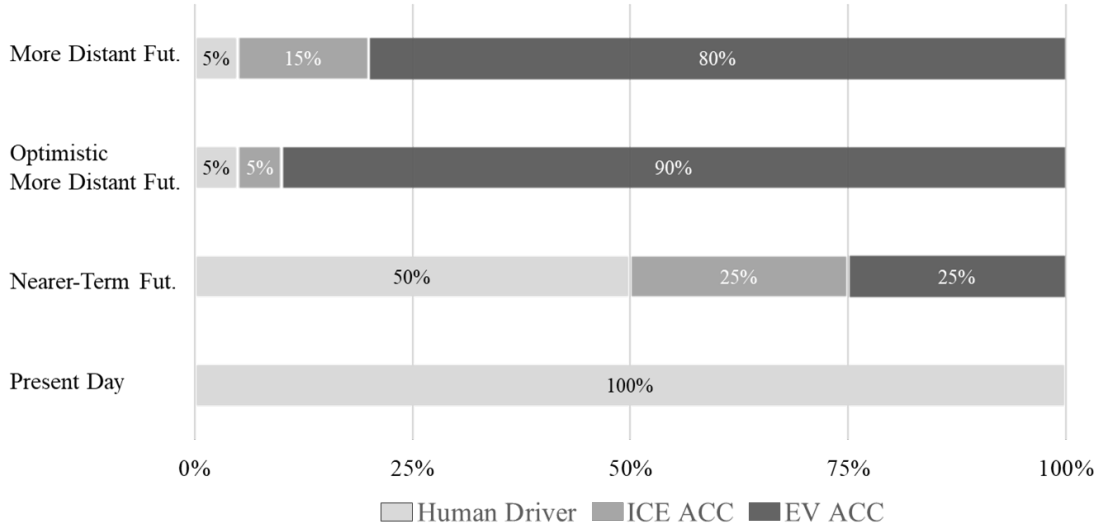


Figure 7. Adoption Forecasts

4.4 Performance Predictions

Outcomes are presented for the Los Angeles- and Manhattan-inspired settings in Figs. 8a and 8b, respectively. The solid, bold curves are the aggregate demands for the respective sites. The thin solid curves display cumulative exit counts for the present day and the dashed and dotted curves are exit counts in the forecasted futures. Each exit curve is the average of 10 simulations with distinct random seeds.¹³ The curves are plotted in such ways that the area between a demand curve and an exit curve is the networkwide vehicle-hours traveled (VHT) for the respective case, and the vertical displacements between said curves are the time-varying vehicle accumulations, also network-wide. Visual inspection of the figures reveals that future traffic conditions are mostly predicted to be worse than at present – for both idealized cities.

Further visual inspection indicates that the future is less bleak for the Los Angeles-like setting. This is thanks in part to its longer link lengths that bring greater queue-storage space. Still, Fig. 8a shows that networkwide VHT and accumulation (i.e., congestion) on the Los Angeles-like site are clearly worsened in our nearer-term future, when ICE-powered cars are expected to still comprise a large share of traffic. The VHT for this epoch is predicted to rise 12% from the present-day mark. Delay (not shown by the figure) is predicted to rise by 25%.

Closer inspection of the dotted, nearer-term exit curve in Fig. 8a reveals why problems occur. Use of a straightedge would confirm that the slopes of that curve (i.e., network exit flows) began to diminish at around time 60 mins from the simulations’ start, when accumulation on the Los Angeles-inspired network reached about 14,000 cars. This sort of reduction in exit flow is a telltale sign of the onset of gridlock, e.g., see Daganzo (1996, 2007) and Mahmassani et al. (2013). The gridlock process was eventually halted – artificially – thanks to the abrupt cessation of demand

¹³ All outcomes to follow were averaged in this same fashion.

that occurred at the 80-minute mark. The gridlock process would have been prolonged, and its pernicious effects exacerbated, had demand diminished in more gradual and realistic fashion near the end of the rush.

The more distant futures look brighter for the Los Angeles-like site. The optimistic, longer-run exit curve is nearly superimposed atop the present-day one and is therefore scarcely visible in Fig. 8a. Thanks again to the longer link lengths and lower present-day congestion level, simulations of this epoch produced sizable periods over which the Los Angeles-like network was not fully engulfed by residual queues. The higher queue discharge flows brought by ACC, and the higher accelerations brought by EVs, thus boosted exit flows during these periods. This favorable effect largely compensated for the damage done by ACC's less-compacted queues. Thus, networkwide VHT predicted for the optimistic distant future is roughly equal that of the present day.

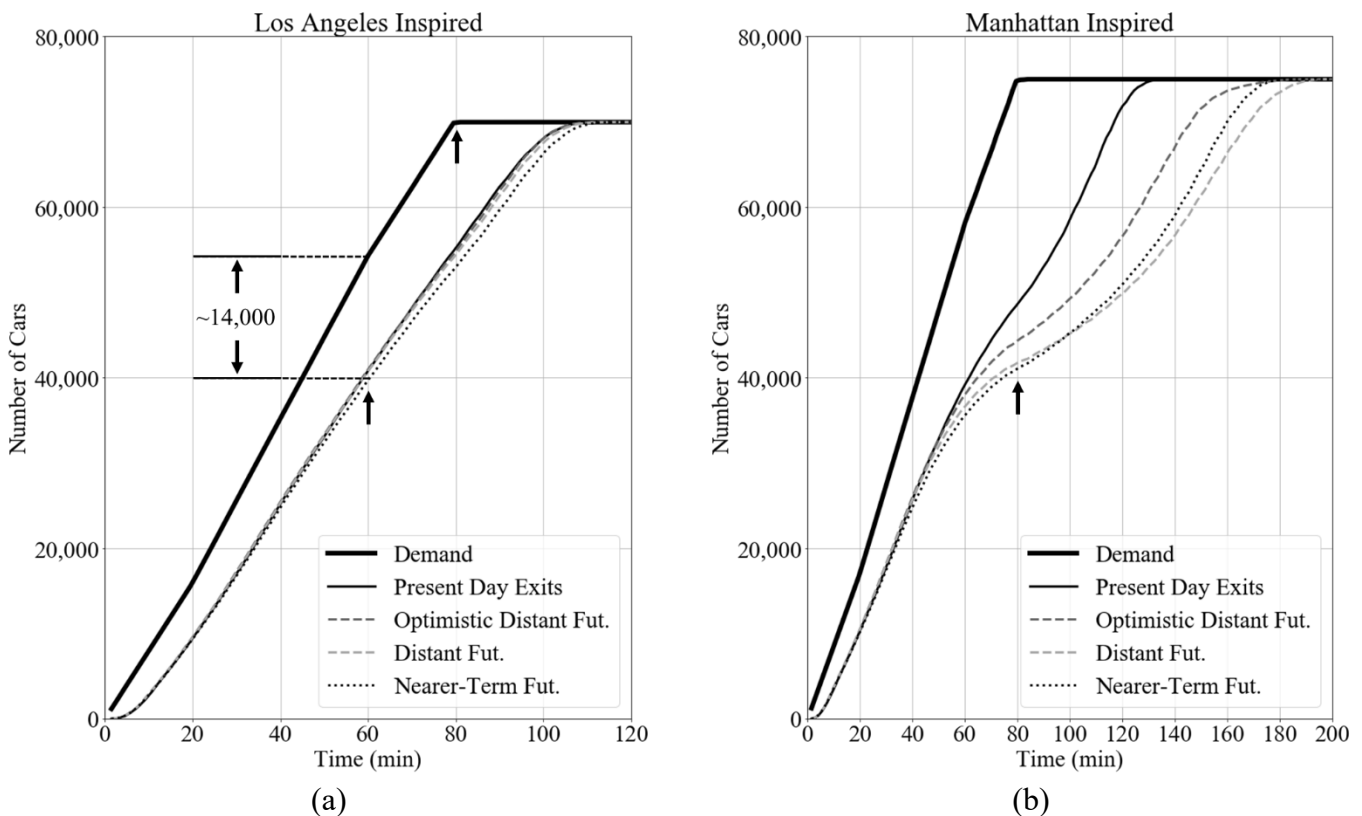


Figure 8. Cumulative Curve Diagrams for (a) Los Angeles-like and (b) Manhattan-like settings.

The exit curve for the less-optimistic distant future also lies close, though not as close to its present-day counterpart. This less optimistic future produced the least-compacted queues of all cases, but the damage done relative to the present day is small, thanks again to the higher average queue discharge flows of ACC-cars and the higher vehicle accelerations of EVs.

Our predictions of less-damaging distant futures for the Los Angeles-inspired setting may be consoling. But these are based on a good many favorable assumptions. Moreover, these predictions are still a far cry from the sorts of benefits promised of ACC by its advocates.

Things look far worse for the Manhattan-like setting, as shown in Fig. 8b. This is because the shorter block lengths with their reduced queue storage space helped make the setting already very congested in the present day and amplified the problems with ACC predicted for the various futures. Note in Fig. 8b that until around the 80-min mark, the damage predicted for the three distinct futures ranked identically to those of the Los Angeles-like site, such that the worst damage was predicted for the nearer term. But by time 80 mins, street links of the Manhattan-like site were so swamped by queues that the benefits of higher discharge flows and accelerations came to an end. The less-optimistic distant future (with its least-compacted queues) therefore ultimately emerged as the most troubling case for the Manhattan-like site. The more-clearly evident onset of gridlock in this case (i.e., the visible reduction in exit flow prior to the 60-min mark) caused networkwide VHT to grow by a whopping 87% over present-day conditions. Things were better for the optimistic distant future, but even then VHT grew relative to the present by 46%.

4.5 Further Notes on the Future

One can take issue with our guesses at ACC and EV adoption patterns for the future. Yet, moderate variations on these guessed-at forecasts would not likely change predicted outcomes much. This seems clear in that our nearer-term adoption forecasts are substantially different from our more distant ones; see again Fig. 7. And yet Figs. 8a and 8b each reveal that horizontal and vertical displacements among a city's exit curves for all three futures are small relative to their displacements from their respective demand curves drawn in bold. It follows that denser, more congested cities like New York have reason for concern, whatever the errors in our adoption forecasts might be. Newer, more sprawling cities like Los Angeles may have grounds for concern as well, given that our predictions of future congestion are probably very loose lower bounds.

One can also take issue with our reliance on a dated study, Nowakowski et al. (2010), of ACC-drivers' collective choices for following-distance. Indeed, the inherent uncertainty underlying these future choices adds to the uncertainties in our outcomes presented in Figs. 8a and b. What now seems certain, however, is that whatever the distribution of preferred following-distances might someday be, the effects on queued densities, and not on queue discharge flows, will matter most in many cities. Analysis now proceeds with this in mind.

4.6 Are Different Futures Possible?

We address the question using Fig. 9. It includes the demand and present-day exit curves, along with two additional exit curves, all for the Manhattan-inspired site. The dotted exit curve in that figure describes what turns out to be a dystopian future in which all cars are EVs and ACC-engaged, and all drivers select the short setting for their preferred following-distance. In this case, the average queue discharge flow is the highest possible (and considerably higher than present-day rates), and jam densities are relatively large (though smaller than at present); see again Figs. 4a and d. Yet, we see from the dotted exit curve in Fig. 9 how conditions devolve in the future: VHT is predicted to grow by a staggering 94%.

A better future lies in maximizing jam density, to compact vehicle queues more tightly, rather than in maximizing bottleneck discharge flows. To illustrate, the dashed curve in Fig. 9 describes exit flows when all cars are EVs, ACC-engaged, and set for the extra-long following-distance. This setting reduces queue discharge flow from the short-setting case, but produces a jam density closer to present-day traffic; see Figs. 4c and d. Indeed, Fig. 9's dashed exit curve lies much closer to the solid, present-day one. Networkwide VHT is in this case predicted to grow by 14%, which is still considerable, but much better than 94%.

We note for emphasis that the dashed exit curve in Fig. 9 describes the very best outcome that ACC (alone) can achieve, subject to all our assumptions thus far adopted. Under these assumptions, VHT cannot diminish beyond what is dictated by said curve, i.e., this lower bound cannot be further diminished by altering our forecasts of ACC or EV adoption, nor by varying the distribution of preferred following-distances. The finding can be a bellwether for further research on the topic, and another promising example of this is offered below.




Figure 9. Cumulative Curve Diagrams for Manhattan-like setting (EVs modeled as the Hyundai IONIQ 5).

5. CONCLUSIONS

Field measurements of steady-state conditions with numerous observations of jammed states confirmed that in severe, low-speed congestion, queues of ACC-engaged cars are less-compacted than those of fully-human-operated vehicles. Though present estimates of severely congested and jammed vehicle spacings are less alarming than those from previous studies, our estimates raise concerns over queue spillovers. Simulations featuring car-following models calibrated to the new data and tested in two idealized settings illustrate that these lower-density, ACC-dominated queues spell trouble for many cities in the future, as ACC’s market share grows over time. Though the street networks in some cities may have queue storage space to spare, which might avert the problems presently predicted, our forecasts raise red flags even for car-centric, moderately congested cities, like our Los Angeles-inspired setting. Not surprisingly, queue storage problems were forecast to do greatest harm in our Manhattan-like setting, which featured short block lengths and already-high levels of city-street congestion. Perhaps what is surprising – and disquieting – is the extent of the damage predicted, especially since so many of our assumptions were conservative in this respect.

The present study comes with caveats. These include the idealizations used for both the Los Angeles- and Manhattan-like settings. These idealizations should probably not lessen the concerns presently raised, however. It follows from Daganzo (1998) that the problems unveiled via our idealized settings will likely be amplified, not diminished, in more complex, real-world cities.

Caveats also include the uncertainties inherent in the inputs to our forecasts, in terms of commuters' future adoption behaviors toward ACC and EVs; and owing to the distribution of preferred following-distances that we used in many instances. Obtaining definitive inputs would be difficult to say the least, and we have argued why altering these data would likely do little to dampen concerns presently raised. Moreover, the inputs used in the present work are sufficient to unveil a possible way forward to a better future.

This way forward is also linked to another caveat of the present study, that being: our field measurements, which were extracted for only one make and model of ICE-powered cars (Toyota Corollas of model year 2020) and from only one EV (a 2022 Hyundai IONIQ 5). The measurements from these cars were alone used to calibrate the traffic models used in our forecasts. Although Makridis et al. (2021) minimizes the concerns here by suggesting that ACC systems function similarly across manufacturers, our field measurements suggest something different, i.e., that ACC-engaged EVs are controlled differently from ACC-engaged ICE-powered cars, and that the former can produce smaller jam spacings and thus a better future. It is unknown to us, however, whether this favorable feature is common to other EV makes and models when operating under ACC. Still, the evidence that this sort of beneficial control is possible is encouraging on its own.

5.1 Moving Forward

Given what was found regarding our lone EV, fine-tuning ACC controllers to select smaller spacings in congestion (particularly when jammed at red lights) may be a way forward for cities with limited queue storage. In addition to the favorable outcome presented in sec. 4.6, other simulations (not shown for brevity) were performed after recalibrating ACC-cars to adopt jam spacings akin to present-day values, no matter the preferred following-distance selected by drivers. For our most optimistic longer-run future, the Manhattan-like setting saw networkwide VHT diminish by 39% from the present day. This reduction may be less impressive than what has been promised by ACC's advocates (USDOT, 2019), yet seems a hopeful outcome given present findings. Moreover, what had been our conservative assumptions (e.g., concerning link geometry and future demand levels) might in this new case be causing us to underpredict the benefits of the fine-tuned ACC controllers.

Mandating shorter vehicle separations in congested traffic may, however, introduce genuine safety concerns, e.g., risks of rear-end collisions. Questions of liability might thus take center stage. And automobile manufacturers may worry about the comfort perceived by consumers when confronted with shorter following-distances. Adapting controllers as described above might therefore encounter resistance.

Another option might be to prohibit use of ACC when driving in congested traffic, especially when on city streets. Such a ban would be akin to the calls in (GMC, 2023; Honda, 2022; Sparks, 2022; Volvo, 2020), and could now be justified not only for safety reasons, but on societal grounds too, i.e., to curb congestion and its externalities. However, limiting the use of a technology feature for which drivers have paid money might prove unpopular and difficult to enforce. It would, moreover, be a step backward in the march toward greater vehicle automation and would end (for a time) any pretense of self-driving technologies as panacea for congestion. We therefore imagine a reluctance to intervene in this way.

The above uncertainties call for further research and deliberation. Time is of the essence: market share for ACC is growing steadily (Calvert et al., 2017; Litman, 2020; Tillema et al., 2017), and fixing problems *post facto* may be more difficult than addressing them proactively. Care and

objectivity will be important as well, both for the researchers who conduct future studies and for the decision-makers who might eventually implement the findings.

As regards the latter group, US DOT, for example, might better serve the cause by playing less the role of promoter and more that of careful evaluator, and perhaps that of regulator. The likelihood that this institutional transformation will occur is not clear to us. A recent US DOT publication leaves room for doubt; see USDOT (2022), p. 20. It seemingly assigns responsibility for determining suitable control standards to the manufacturers of ACC-equipped vehicles. This transfer of authority strikes us as a troubling abdication. Perhaps leadership will come instead from state and local governments throughout the US, or perhaps from elsewhere in the world. Wherever it might come, leadership is sorely needed to confront the unintended consequence of ACC evidently now looming on the horizon.

ACKNOWLEDGMENTS

The field tests described in this paper were sponsored by the National Science Foundation (NSF) under the Civil Infrastructure Systems (CIS) program, award number 2301446. The data analysis was partially funded by the National Institute of Congestion Reduction (NICR) a National University Transportation Center sponsored by US DOT.

REFERENCES

- Ahmed, H. U., Huang, Y., & Lu, P. (2021). A Review of Car-Following Models and Modeling Tools for Human and Autonomous-Ready Driving Behaviors in Micro-Simulation. *Smart Cities*, 4(1), Article 1. <https://doi.org/10.3390/smartcities4010019>
- Aimsun. (2017). *Aimsun Next*. <https://www.aimsun.com/>
- BMDV. (2021, July 27). *Germany will be the world leader in autonomous driving*. German Federal Ministry for Digital and Transport. <https://bmdv.bund.de/SharedDocs/EN/Articles/DG/act-on-autonomous-driving.html>
- Bose, A., & Ioannou, P. A. (2003). Analysis of traffic flow with mixed manual and semiautomated vehicles. *IEEE Transactions on Intelligent Transportation Systems*, 4(4), 173–188. <https://doi.org/10.1109/TITS.2003.821340>
- Calvert, S., Schakel, W., & Van Lint, J. (2017). Will automated vehicles negatively impact traffic flow? *Journal of Advanced Transportation*, 2017.
- CARB. (2022, August 25). *California moves to accelerate to 100% new zero-emission vehicle sales by 2035* [California Air Resources Board]. <https://ww2.arb.ca.gov/news/california-moves-accelerate-100-new-zero-emission-vehicle-sales-2035>
- Carlson, R. C., Papamichail, I., & Papageorgiou, M. (2014). Integrated feedback ramp metering and mainstream traffic flow control on motorways using variable speed limits. *Transportation Research Part C: Emerging Technologies*, 46, 209–221. <https://doi.org/10.1016/j.trc.2014.05.017>
- Carlson, R. C., Papamichail, I., Papageorgiou, M., & Messmer, A. (2010). Optimal mainstream traffic flow control of large-scale motorway networks. *Transportation Research Part C: Emerging Technologies*, 18(2), 193–212. <https://doi.org/10.1016/j.trc.2009.05.014>

- Cassidy, M. J., Anani, S. B., & Haigwood, J. M. (2002). Study of freeway traffic near an off-ramp. *Transportation Research Part A: Policy and Practice*, 36(6), 563–572.
- Chau, K. T. (2014). 21—Pure electric vehicles. In R. Folkson (Ed.), *Alternative Fuels and Advanced Vehicle Technologies for Improved Environmental Performance* (pp. 655–684). Woodhead Publishing. <https://doi.org/10.1533/9780857097422.3.655>
- Cheng, Q., Liu, Z., Lin, Y., & Zhou, X. (Simon). (2021). An s-shaped three-parameter (S3) traffic stream model with consistent car following relationship. *Transportation Research Part B: Methodological*, 153, 246–271. <https://doi.org/10.1016/j.trb.2021.09.004>
- Chiabaut, N., Buisson, C., & Leclercq, L. (2009). Fundamental diagram estimation through passing rate measurements in congestion. *IEEE Transactions on Intelligent Transportation Systems*, 10(2), 355–359.
- Choi, J. (2021, January 5). *Massachusetts to require 100 percent of car sales to be electric by 2035*. The Hill. <https://thehill.com/policy/energy-environment/532684-massachusetts-to-require-100-percent-of-car-sales-to-be/>
- Daganzo, C. F. (1996). The nature of freeway gridlock and how to prevent it. *Transportation and Traffic Theory*, 629–646.
- Daganzo, C. F. (1998). Queue spillovers in transportation networks with a route choice. *Transportation Science*, 32(1), 3–11.
- Daganzo, C. F. (2007). Urban gridlock: Macroscopic modeling and mitigation approaches. *Transportation Research Part B: Methodological*, 41(1), 49–62.
- Daganzo, C. F., Cassidy, M. J., & Bertini, R. L. (1999). Possible explanations of phase transitions in highway traffic. *Transportation Research Part A: Policy and Practice*, 33(5), 365–379. [https://doi.org/10.1016/S0965-8564\(98\)00034-2](https://doi.org/10.1016/S0965-8564(98)00034-2)
- Daganzo, C. F., & Lehe, L. J. (2016). Traffic flow on signalized streets. *Transportation Research Part B: Methodological*, 90, 56–69.
- Darbha, S., & Rajagopal, K. R. (1999). Intelligent cruise control systems and traffic flow stability. *Transportation Research Part C: Emerging Technologies*, 7(6), 329–352. [https://doi.org/10.1016/S0968-090X\(99\)00024-8](https://doi.org/10.1016/S0968-090X(99)00024-8)
- Doig, J., Daganzo, C. F., & Cassidy, M. J. (2024). How and when cordon metering can reduce travel times. *Transportation Research Part C: Emerging Technologies*, 104581. <https://doi.org/10.1016/j.trc.2024.104581>
- Gipps, P. G. (1981). A behavioural car-following model for computer simulation. *Transportation Research Part B: Methodological*, 15(2), 105–111.
- GMC. (2023). *About Adaptive Cruise Control*. <https://www.gmc.com/support/vehicle/driving-safety/driver-assistance/adaptive-cruise-control>
- Goñi-Ros, B., Schakel, W. J., Papacharalampous, A. E., Wang, M., Knoop, V. L., Sakata, I., van Arem, B., & Hoogendoorn, S. P. (2019). 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.

- Gunter, G., Gloudemans, D., Stern, R. E., McQuade, S., Bhadani, R., Bunting, M., Delle Monache, M. L., Lysecky, R., Seibold, B., & Sprinkle, J. (2020). Are commercially implemented adaptive cruise control systems string stable? *IEEE Transactions on Intelligent Transportation Systems*, 22(11), 6992–7003.
- Gunter, G., Janssen, C., Barbour, W., Stern, R. E., & Work, D. B. (2019). Model-based string stability of adaptive cruise control systems using field data. *IEEE Transactions on Intelligent Vehicles*, 5(1), 90–99.
- Honda. (2022). *Adaptive Cruise Control (ACC) with Low Speed Follow*. <https://techinfo.honda.com/rjanisis/pubs/OM/AH/AT202222IOM/enu/details/131229047-67924.html>
- Hoogendoorn, S. P., Knoop, V. L., van Lint, H., & Vu, H. L. (2015). Applications of the Generalized Macroscopic Fundamental Diagram. In M. Chraïbi, M. Boltes, A. Schadschneider, & A. Seyfried (Eds.), *Traffic and Granular Flow '13* (pp. 577–583). Springer International Publishing. https://doi.org/10.1007/978-3-319-10629-8_65
- Hyundai. (2022). *IONIQ 5 Technology*. HYUNDAI MOTORS. <https://www.hyundai.com/worldwide/en/eco/ioniq5/technology>
- James, R. M., Melson, C., Hu, J., & Bared, J. (2019). Characterizing the impact of production adaptive cruise control on traffic flow: An investigation. *Transportmetrica B: Transport Dynamics*, 7(1), 992–1012.
- Kesting, A., & Treiber, M. (2008). Calibrating Car-Following Models by Using Trajectory Data: Methodological Study. *Transportation Research Record*, 2088(1), 148–156. <https://doi.org/10.3141/2088-16>
- Kesting, A., Treiber, M., Schönhof, M., & Helbing, D. (2008). Adaptive cruise control design for active congestion avoidance. *Transportation Research Part C: Emerging Technologies*, 16(6), 668–683.
- Kim, K., & Cassidy, M. J. (2012). A capacity-increasing mechanism in freeway traffic. *Transportation Research Part B: Methodological*, 46(9), 1260–1272.
- Knoop, V. L., & Daamen, W. (2017). Automatic fitting procedure for the fundamental diagram. *Transportmetrica B: Transport Dynamics*, 5(2), 129–144. <https://doi.org/10.1080/21680566.2016.1256239>
- Koshi, M. (1989). Cycle time optimization in traffic signal coordination. *Transportation Research Part A: General*, 23(1), 29–34. [https://doi.org/10.1016/0191-2607\(89\)90137-4](https://doi.org/10.1016/0191-2607(89)90137-4)
- Lambert, F. (2021, November 10). *Countries and automakers agree to go all-electric by 2040 in weak new goal set at COP26*. Electrek. <https://electrek.co/2021/11/10/countries-automakers-agree-go-all-electric-by-2040-weak-new-goal-cop26/>
- Lárraga, M. E., & Alvarez-Icaza, L. (2010). Cellular automaton model for traffic flow based on safe driving policies and human reactions. *Physica A: Statistical Mechanics and Its Applications*, 389(23), 5425–5438. <https://doi.org/10.1016/j.physa.2010.08.020>
- Li, T., Chen, D., Zhou, H., Xie, Y., & Laval, J. (2022). Fundamental diagrams of commercial adaptive cruise control: Worldwide experimental evidence. *Transportation Research Part C: Emerging Technologies*, 134, 103458.

- Litman, T. (2020). *Autonomous vehicle implementation predictions: Implications for transport planning*.
- Mahmassani, H. S., Saberi, M., & Zockaie, A. (2013). Urban network gridlock: Theory, characteristics, and dynamics. *Procedia-Social and Behavioral Sciences*, 80, 79–98.
- Makridis, M., Mattas, K., Anesiadou, A., & Ciuffo, B. (2021). OpenACC. An open database of car-following experiments to study the properties of commercial ACC systems. *Transportation Research Part C: Emerging Technologies*, 125, 103047.
- Milanés, V., & Shladover, S. E. (2014). Modeling cooperative and autonomous adaptive cruise control dynamic responses using experimental data. *Transportation Research Part C: Emerging Technologies*, 48, 285–300.
- Newell, G. F. (1993). A simplified theory of kinematic waves in highway traffic, part II: Queueing at freeway bottlenecks. *Transportation Research Part B: Methodological*, 27(4), 289–303.
- Nowakowski, C., O’Connell, J., Shladover, S. E., & Cody, D. (2010). *Cooperative adaptive cruise control: Driver acceptance of following gap settings less than one second*. 54(24), 2033–2037.
- Ntousakis, I. A., Nikolos, I. K., & Papageorgiou, M. (2015). On Microscopic Modelling of Adaptive Cruise Control Systems. *Transportation Research Procedia*, 6, 111–127. <https://doi.org/10.1016/j.trpro.2015.03.010>
- Panwai, S., & Dia, H. (2005). Comparative evaluation of microscopic car-following behavior. *IEEE Transactions on Intelligent Transportation Systems*, 6(3), 314–325. <https://doi.org/10.1109/TITS.2005.853705>
- Papacharalampous, A. E., Wang, M., Knoop, V. L., Ros, B. G., Takahashi, T., Sakata, I., van Arem, B., & Hoogendoorn, S. P. (2015). *Mitigating congestion at sags with adaptive cruise control systems*. 2451–2457.
- Rakha, H., Farzaneh, M., Arafteh, M., & Sterzin, E. (2008). Inclement Weather Impacts on Freeway Traffic Stream Behavior. *Transportation Research Record*, 2071(1), 8–18. <https://doi.org/10.3141/2071-02>
- Seherman, J., & Skabardonis, A. (2013). Rethinking the Driver Population Factor: Examination of Freeways in California. *Transportation Research Record: Journal of the Transportation Research Board*, 2395(1). <https://doi.org/10.3141/2395-12>
- Shang, M., & Stern, R. E. (2021). Impacts of commercially available adaptive cruise control vehicles on highway stability and throughput. *Transportation Research Part C: Emerging Technologies*, 122, 102897.
- Shi, X., & Li, X. (2021). Constructing a fundamental diagram for traffic flow with automated vehicles: Methodology and demonstration. *Transportation Research Part B: Methodological*, 150, 279–292.
- Sparks, E. (2022, May 24). *Is Adaptive Cruise Control a Safe Feature to Use?* Eckell Sparks Attorneys at Law. <https://www.eckellsparks.com/2022/05/24/adaptive-cruise-control-safe-feature-use/>

- Talebpour, A., & Mahmassani, H. S. (2016). Influence of connected and autonomous vehicles on traffic flow stability and throughput. *Transportation Research Part C: Emerging Technologies*, 71, 143–163.
- Tesla. (2023). *Autopilot and Full Self-Driving Capability*. <https://www.tesla.com/support/autopilot>
- Tillema, T., Gelauff, G., van der Waard, J., Baveling, J., & Moorman, S. (2017). *Paths to a self-driving future: Five transition steps identified*.
- Transportation Research Board. (2016). *Highway Capacity Manual 6th Edition: A Guide for Multimodal Mobility Analysis*. The National Academies Press. <https://doi.org/10.17226/24798>
- Treiber, M., Hennecke, A., & Helbing, D. (2000). Congested traffic states in empirical observations and microscopic simulations. *Physical Review E*, 62(2), 1805–1824. <https://doi.org/10.1103/PhysRevE.62.1805>
- USDOT. (2019). *Advanced Driver Assistance Systems (ADAS): Cooperative Adaptive Cruise Control (CACC)*. https://www.itskrs.its.dot.gov/sites/default/files/executive-briefings/2019/EB03-%20ACC%20DRIVER%20ASSISTANCE_%2005_16_19_FINAL.pdf
- USDOT. (2022). *Preferred Following Distance as a Function of Speed—Function-Specific Automation (Level 1) Applications*. <https://highways.dot.gov/research/publications/safety/FHWA-HRT-22-107>
- Vander Werf, J., Shladover, S. E., Miller, M. A., & Kourjanskaia, N. (2002). Effects of adaptive cruise control systems on highway traffic flow capacity. *Transportation Research Record*, 1800(1), 78–84.
- VanderWerf, J., Shladover, S., Kourjanskaia, N., Miller, M., & Krishnan, H. (2001). Modeling Effects of Driver Control Assistance Systems on Traffic. *Transportation Research Record*, 1748(1), 167–174. <https://doi.org/10.3141/1748-21>
- Volvo. (2020, March 19). *Limitations for adaptive cruise control*. <https://www.volvocars.com/en-om/support/car/xc40/19w17/article/53495f6e8386b6b9c0a8015150beb4c1>
- Walz, E. (2023, June 26). *China supports commercialization of Level-3 autonomous driving technology*. Automotive Dive. <https://www.automotivedive.com/news/china-commercialization-level-3-autonomous-driving/653872/>
- Xiao, L., & Gao, F. (2010). A comprehensive review of the development of adaptive cruise control systems. *Vehicle System Dynamics*, 48(10), 1167–1192. <https://doi.org/10.1080/00423110903365910>
- Xiao, L., Wang, M., & van Arem, B. (2017). Realistic Car-Following Models for Microscopic Simulation of Adaptive and Cooperative Adaptive Cruise Control Vehicles. *Transportation Research Record*, 2623(1), 1–9. <https://doi.org/10.3141/2623-01>