CAN Coach: Continuous CAN-based Feedback to Change Driver Behavior

Matthew Nice (corresponding author)
Department of Civil and Environmental Engineering
Institute for Software Integrated Systems
Vanderbilt University
1025 16th Avenue South
Nashville, TN 37212
matthew.nice@vanderbilt.edu

Daniel B. Work
Department of Civil and Environmental Engineering
Institute for Software Integrated Systems
Vanderbilt University
1025 16th Avenue South
Nashville, TN 37212
dan.work@vanderbilt.edu

6149 words + 2 tables = 6649 words

ABSTRACT

In this work, we address the question of whether driver behavior can be modified to improve the ability of a driver to maintain a constant time gap using only data from stock sensors on a commercial car reported on the *controller area network* (CAN). We introduce the *CAN Coach*, which is a system that continuously feeds time gap sensor information from the CAN bus back to the driver in real time. Three sets of preliminary experiments are conducted in which the study vehicle follows a lead vehicle driving a specified driving profile to assess the potential of the CAN Coach to modify driver behavior. The experiments consider *Normal driving* (the driver is given no prompt and no feedback), *Instructed driving* (driver is given a prompt to drive at a two second time gap, but is not given any realtime feedback from the CAN Coach), and *Coached driving* (two second prompt and CAN Coach feedback). The mean time gap errors from the 2 second target are 0.39 s (Normal Driving), 0.09 s (Instructed Driving), and 0.01 s (Coached Driving). The standard deviation of the time gap error with the CAN Coach reduced by 72% and 68% from Normal Driving and Instructed Driving respectively. Given this reduction of mean and standard deviation of the time gap error, we conclude that it is possible to "coach" drivers using only data from the CAN.

INTRODUCTION

Modern cars have *controller area networks* (CANs) that are in place to facilitate communication between *electronic control units* (ECUs) on the vehicle. CAN uses a multi-master data bus where all messages are multicast. There are dozens of systems on the CAN, featuring hundreds of sensors, that support everything in the car from optimal engine function to climate control to driver safety. Year by year, more electronic features are added (e.g., adaptive cruise control, blind spot monitoring, lane keep assist). The amount of data produced on the CAN will continue to accelerate as cars become increasingly automated. The data that is multicast on the CAN bus provides information on the car, the driver, and the external environment. CAN messages are often sent in the order of hundreds of Hertz, and the thousands of data signals are typically a standard eight bytes in size. This means gigabytes of data per hour are sent on a vehicle's CAN bus.

Considering the ubiquity of CAN technology and the holistic nature of the data it provides, there is substantial potential for finding insights about drivers under various levels of vehicle automation [1, 2]. Moreover, considering the real-time nature of the CAN and the increased number of sensors reporting data on it, it is now possible to implement CAN applications that augment safety, improve eco-driving, or adjust car following behavior. Since CAN data is available widely within today's vehicle fleet, these applications could scale broadly and at low cost.

The demonstrable value of CAN data is already established in commercial aftermarket applications. Insurance companies offer discounts on rates in exchange for a relatively reduced set of data on a driver available through the *on board diagnostics* (OBD-II) interface (a subset of the data available on the CAN) to characterize driver behavior and ultimately risk. Major car manufacturers now collect data on their vehicles, and collaborate on CAN data research [3, 4, 5]. Like the rapid introduction of mobility applications that were created during the transition from traditional cell phones to smartphones, we expect a similar transition to occur on vehicles using data available on the CAN bus. The goal of this work is to illustrate such an application, that is able to change driver car following behavior using only CAN data.

Though there has been discussion of the potential of CAN data to support application development [1], there have not yet been studies published showing real-time driver feedback to achieve a precise human-in-the-loop control objective. The literature features machine learning work characterizing data traces from CAN to categorize emergent driving behavior like changing lanes, making turns, and identifying drivers [1, 2]. These analyses were all done in post-processing. Because the data from the CAN has potential to be processed in real time, it can also be used in real time to modify driver behavior.

In this work, we address the question of whether driver behavior can be modified to improve the ability of a driver to maintain a constant *time gap* (i.e., the time separation between a lead vehicle's rear bumper and the following vehicle's front bumper when the follower travels at a constant speed) using only data from the CAN[1]. We introduce CAN Coach, a system that provides real-time time gap feedback to the driver based on sensor data already available on the CAN bus. We illustrate through preliminary experiments the potential for feedback using CAN data to change driver behavior.

[1]This work was conducted under IRB approval #200343

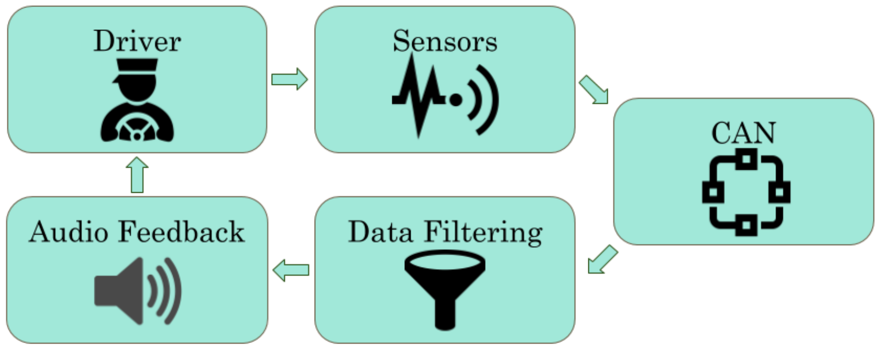

FIGURE 1 : CAN Coach system. Stock sensors on the vehicle report measurements on the CAN, which are used to determine the error between a desired time gap (to the car ahead) control objective and the current time gap. The magnitude and direction of the error is fed back to the driver, who adjusts the vehicle. The changes are measured by the sensors and the feedback loop is closed.

Related Work

This section discusses related work on time gap control, driver behavior intervention, and audio feedback to drivers. The work [6] provides continuous auditory feedback to the driver whenever the time gap is below one second. Because of the feedback, drivers spend less time driving at time gaps below or near one second. Cautious drivers do not change their following behavior as much, because the feedback only occurs at the one second threshold which is not often reached by these drivers. In contrast to [6], [7] gives drivers discrete feedback on their time headway. Human subjects placed in a driving simulator change their behavior in the correct direction, with error of 0.3 s to 0.8 s [7].

Behavioral intervention is also done for goals other than time gap control. For example, eco-driving experiments are done with behavioral training as the method of intervention [8]. The work [8] measures driver's eco-driving in simulated environments, then does offline training on eco-driving, and puts the driver back in the simulator to measure the effects of behavioral intervention.

Choosing the correct sound is an important design consideration in auditory warning systems (e.g., as is already used in lane departure warnings and blind spot monitoring systems). The work [4] examines which sounds help in discrete intervention (i.e., a single warning when a minimum time gap threshold is violated) to stop dangerously close following. It investigates which sounds are informative but least irritating and therefore effective as a feature in production systems designed to increase safety for long-haul truckers without annoyance. Because [3] uses continuous ambient sounds to give truckers complete information on the state of the vehicle, e.g., lane position, it was determined through truck driver surveys that the lack of sound is the preferred mechanism to give feedback to drivers. In contrast to [4] which only aims to avoid very low time gaps, our work aims to coach drivers affirmatively to follow a specific time gap. Our more ambitious aim comes at the cost that it would make sounds too often to be usable in a production-level system.

The time gap is a quantity that can be readily computed from sensor data on a follower vehicle. The time gap is the space gap (distance between lead vehicle rear bumper and follower vehicle front bumper) divided by the velocity of the follower car. Time gap is a surrogate safety metric for cars [9, 10, 11], in which larger time gaps are associated with increased safety. In freeway driving, drivers drive at nearly half the recommended safe time gap [12]. It is also relevant to note that the time gap and total throughput of the roadway are inversely related, so it is important to prevent the time gap from becoming too large. Accordingly, feedback to promote a safe and overall throughput-efficient time gap is desired.

The work in [5] informs drivers in a simulator about the time headway of the car to promote appropriate reliance on automation and support effective transitions between manual and ACC control. Drivers who were given continuous information about time headway are better at controlling time headway and managing braking interventions. They argue that continuous feedback is more effective than discrete driver warnings.

Contribution and Outline

The main contribution of this article is a preliminary demonstration of a CAN-based controller for time gap. In other words, using on-board stock vehicle sensors to give feedback in real-time is effective in changing the time gap driving behavior. A real-time feedback system, CAN Coach, is built to read CAN data and compute the error between the desired time gap and the current time gap. A set of experiments are conducted to compare the time gap distribution of a driver under *i*) *Normal Driving* with no coaching and no prompting; *ii*) *Instructed Driving* with a prompt to follow a two second time gap but no coaching; and *iii*) *Coached Driving* with CAN Coach that provides real time feedback to assist the driver in achieving a two second time gap. Compared to normal and instructed driving, The mean time gap error using CAN Coach is reduced from 0.39 s (Normal) and 0.09 s (Instructed) to 0.01 s. The standard deviation of time gap error is also reduced using CAN Coach by 72% (compared to Normal) and 68% (compared to Instructed).

We caution that the results presented here are from a preliminary study on a single driver, so the generalizablity across multiple drivers remains an open question we are exploring in our future work. This work provides a starting point to motivate the need for a comprehensive study in the next phase of our work.

The remainder of this article is organized as follows. In the Section *CAN Data Background*, we provide an overview of CAN and a description of the data obtained from the vehicle platform. In the Section *Methodology*, we describe the data processing required to transform multiple raw CAN messages into relevant time gap data that can be used as feedback to the driver. The implementation of the auditory feedback is also discussed, and the description of the driving experiments is provided. The Section *Results* describes how the validity of the CAN data is assessed, and provides the analysis of the experiments conducted to determine the impact of CAN Coach on driver behavior. The Section *Conclusions and Future Work* outlines the planned extensions from our preliminary work presented here.

BACKGROUND ON CAN

This section provides a brief background on CAN data, the hardware used to collect it for this study, and details on the relevant CAN signals used. For this study the precision, accuracy, and recency of the data provided for feedback contribute directly to the performance of the CAN Coach. The ceiling of the potential performance is set by the quality of the real-time CAN data that are the basis of the driver feedback (i.e., bad sensor data or incorrect processing will provide inaccurate feedback to the driver).

To build the time gap controller in CAN Coach, we need measurements of the velocity of the following vehicle and the distance from the front of the following vehicle to the tail of the leading vehicle. By accessing the CAN data in the following vehicle, this information can be found

TABLE 1 : Example of CAN Messages that must be deciphered and translated to build CAN Coach.

CAN Messages				
Time (s)	***Bus***	***Address***	***Data Signal***	***Data Length (B)***
3.048295	0	945	00000105000800ca	8
3.048318	1	865	8053460056005c7c	8
3.048447	1	384	a8fff80000000028	8
...	...	...	...	...

via wheel encoder and radar sensors. Radar sensors are increasingly common on stock vehicles today that directly measure information needed to compute the space gap and relative velocity, so the relative trajectories of the follower and the leader are evident in the data. To have relevant feedback in real-time, there are two related challenges with radar data. First, incoming radar data tracks many objects in a wide field of view and needs to be processed to extract a high quality estimate of the location and relative velocity of the lead vehicle. Second, any processing on the radar data must be done quickly so that the feedback to the driver remains relevant. Latent data could substantially degrade the quality of control achievable using CAN Coach.

Collecting CAN Data

Each message recorded from the CAN includes a timestamp of record, the bus the message was sent on, the data signal(s) contained within the message, and the length of the data in bytes, which is critical for proper decoding (i.e., translating into human understandable values).

A subset of this data is shown in Table 1. From this data, we can understand what the vehicle's sensors are reporting after translation. The address gives the context to decipher the information embedded in the hexadecimal message. A separate *CAN database* (referred to as a DBC) provides information and instruction on how to read CAN signals. The DBC has stringent formatting rules that define the messages, the signals subset of the message, the start, size, scale, offset, min, max, units, endian-ness, and signed-ness of the message. Some messages are multiplexed too. The DBC is critical for successful decoding of the CAN data, since any errors in the DBC result in errors or complete failure of the ability to interpret the information contained in the CAN message.

Relevant CAN Signals

This section details the CAN signals that are used in this work, summarized in Table 2. Many types of velocity signals are reported on the CAN. Each wheel reports an associated velocity, the speedometer has an associated velocity, and the driver support unit reports a velocity. In this work we use the velocity corresponding to the speedometer because it reflects the movement for the vehicle as a whole, and is reported for reference by other vehicle systems on multiple busses. The brakes message has a binary Brake Pressed signal, and an 8-bit sized proportional signal. The proportional signal is helpful for testing feedback to the driver, because it is sensitive to small changes in the brake pressure. The gas pedal signal contains a binary Gas Pressed signal, and a proportional signal analogous to the brakes. Acceleration is reported in x,y, and z directions. The x-direction corresponds to longitudinal accelerations, which is informative to validate velocity, gas pedal, and brake signals.

The radar data is perhaps the most informative data to gain context about the conditions ahead of the vehicle, but is also the most challenging to interpret. The data is reported on a set of *tracks*, each of which may contain an object in the field of view of the sensor. Each track reports the Cartesian (x,y) coordinates of the object, the relative velocity of the object with respect to the vehicle, the relative acceleration, a binary validity signal, and a score ranging from 0-100. Processing this collection of data to determine the distance to the vehicle immediately ahead, critical for the CAN Coach, is explained later in the Section *Methodology*.

Hardware

This section describes the experimental hardware platform used. The base vehicle used in this work is a stock 2020 Toyota Rav4 Hybrid vehicle. The vehicle has as standard equipment an *Adaptive Cruise Control* (ACC) system and *Lane Tracing Assist* (LTA) system. The ACC system depends on a stock forward looking radar unit that transmits relevant data on the CAN. To access this data, we use a Gray Panda manufactured by Comma.ai as a data logging device. The Gray Panda is connected to a standard laptop via USB, where messages are decoded, processed, and used to generate an auditory feedback for the driver. The laptop is used for convenience; in our ongoing work we intend to replace it with a low cost Raspberry Pi 4.

METHODOLOGY

This section details the real time processing approach used to build the CAN Coach, and the experimental design. The hardware was chosen with minimized cost in mind for future scalability. As much as possible, stock vehicle features were utilized to maximize the value from the vehicle CAN. The select data signals were processed in real time to keep computing costs low and throughput quick. Off-the-shelf algorithms are leveraged wherever possible to ensure the flow of real-time data was reliable and accurate. The CAN Coach controller itself was kept as minimal as possible. As a human-in-the-loop system, digital precision and subtleties get easily lost. There are essentially only three discrete options for a human driver following another vehicle: speed up, slow down, do nothing. The controller was designed to guide the driver to do one of those three actions.

Real Time Processing

This section discusses the decisions made to process CAN data for driver feedback. The general approach is summarized as follows. The radar and velocity signals are read and processed to compute the current time gap. The current time gap is then compared to the desired time gap (two seconds in our experiments). The difference (or time gap error) is then used to generate sounds to the driver that indicate if the time gap needs to be increased or decreased. The time gap is calculated at 20 Hz, with sounds output to the driver through the vehicle audio system at a reduced rate of 3 Hz.

The most challenging measurement to obtain is the a reliable space gap measurement to the vehicle immediately ahead. The CAN includes a pre-filtered space gap signal reporting at 5 HZ, which is rounded down to the integer meter. This signal cannot be directly used because it was observed that it does not always report when a vehicle is immediately ahead, due to undetermined radar processing mechanisms that are not exposed. To improve performance of CAN Coach, we elect to instead process the raw radar tracks and estimate the lead vehicle directly from the raw data.

TABLE 2 : Summary of CAN Messages Used for CAN Coach and experimental analyses.

Message Name	**Signals**	**Frequency (Hz)**	**Description**
Raw Radar Track 1	Latitude, Longitude, Relative Velocity	20	Tracks objects in radar field, their position and velocity. Also meta-reports data valid and score 0-100.
Raw Radar Track 2	Latitude, Longitude, Relative Velocity	20	Tracks objects in radar field, their position and velocity. Also meta-reports data valid and score 0-100.
...	...	...	...
Raw Radar Track 16	Latitude, Longitude, Relative Velocity	20	Tracks objects in radar field, their position and velocity. Also meta-reports data valid and score 0-100.
Pre-filtered Radar	Space Gap	5	Filtered radar data reporting the space gap of the lead vehicle.
Velocity	Speed	40	Speed of the car from the speedometer.
Acceleration	x-direction,y-direction,z-direction	40	Acceleration of the car from an accelerometer.
Brakes	Brake Pressed, Brake Proportion	76	A binary brake pressed signal and an 8-bit (0-256) brake pressure signal.
Gas Pedal	Gas Pressed, Gas Proportion	62	A binary gas pressed signal and proportional gas pedal pressure signal ranging from 0-1.

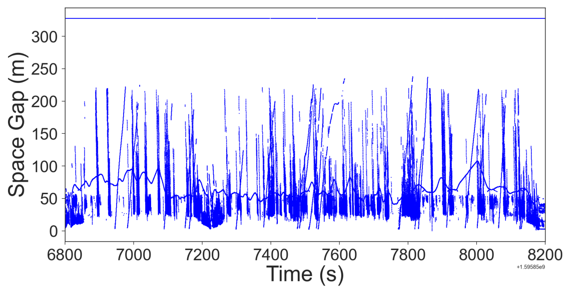

(a) All raw radar space gap data.

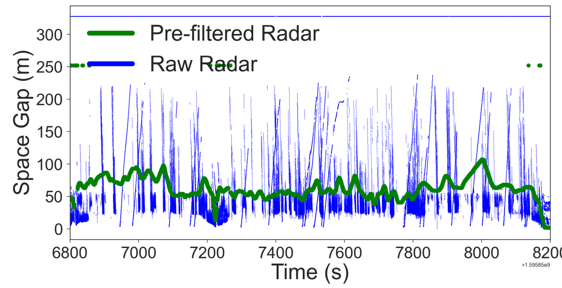

(b) All raw radar space gap data with pre-processed radar overlaid.

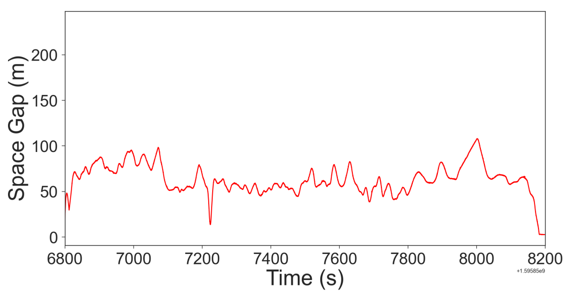

(c) Filtered raw radar space gap data.

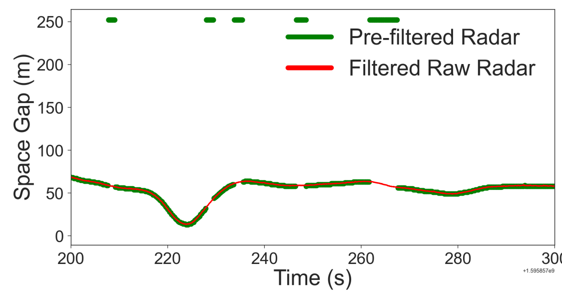

(d) Filtered raw radar sensor data overlaid with pre-processed radar.

FIGURE 2 : Progression of radar data processing for the CAN Coach.

Raw radar data contains the space gap data for the lead vehicle, but also space gap data for up to 15 other objects in the field of view of the radar unit. Figures 2a and 2b give an illustration of the challenge. In Figure 2a, numerous diagonal lines correspond to objects moving towards or away from the vehicle over time that are not directly in the same lane as the vehicle. The solid line above 300 m corresponds to when a track does not contain an object. Figure 2b contains raw data from Figure 2a and also the pre-filtered (but incomplete) signal corresponding to the lead vehicle.

To clean the raw radar data and improve upon the pre-filtered space gap signal, a Kalman filtering problem [13] is solved in real time. The Kalman filter tracks as state the position of the lead vehicle. To determine which measurements are associated with the lead vehicle, and thus available to perform an update to the tracked position in the filter based on the measurement, a data association problem must be solved. The data association problem determines which measurement corresponds to the lead vehicle, amongst the many raw radar measurements reported on the CAN. The data association problem is solved via the agglomerative clustering algorithm using a standard implementation found in [14].

Figure 2c shows the result of running the Kalman filter and agglomerative clustering. Compared to the raw data in Figure 2a, it is apparent a clean signal has been obtained. In Figure 2d, we show the pre-filtered space gap data (containing missing data), with the result of our filtering methods. Here we can verify that the filtering is consistent with the pre-filtered signal when data is present, but also fills in the missing entries and thus gives a reliable space gap signal.

Controller Design

We briefly discuss the decisions used to create driver feedback within the CAN Coach to change driver behavior. The CAN Coach has simple feedback, because complexity can be a burden to the driver [3]. Specifically, when the driver is a threshold amount away from a two second time gap, a signal is output from the controller to alert the driver of the car's state. Since auditory or tactile driving warnings correspond with better driver reaction times than visual warnings or no warning at all [15], we select an auditory feedback approach. When the time gap is outside of 2.0s +/- 0.05s, a sound is output to increase or shorten the time gap. There are only three discrete options for a human driver controlling a car following another vehicle: speed up, slow down, or do nothing. High pitch feedback is given to indicate the time gap is high and the driver should speed up to shorten the time gap. Low pitch feedback is given to the driver to indicate the time gap is low and the driver should slow down to increase the time gap.

Experimental Design

This section covers the experimental design to assess the CAN Coach's ability to change driver behavior with respect to following a two second time gap. In each experiment there are two vehicles, a leader and a follower; the follower drives in one of three experimental conditions, while the lead vehicle drives under one of two protocols, described next.

Experimental Conditions

- *Normal Driving*: Regular driving with no prompt for time gap. The driver is given instructions to follow the vehicle ahead without changing lanes. As a safety precaution, the driver is alerted that the vehicle ahead may change speeds during the test. This establishes a baseline in absence of instruction and feedback. Some evidence shows that a mild cognitive task (such as driving for an experiment) induces longer time gap [16], so the time gap behavior shown experimentally likely represents a longer time gap than one observed candidly.

- *Instructed Driving*: The driver is given all information provided during Normal driving. Additionally, they are asked to maintain a two second time gap to the vehicle ahead. This instruction establishes a baseline in absence of feedback with which to compare the CAN feedback.

- *Coached Driving*: The driver is given the same instructions as Instructed driving. Additionally, the CAN Coach is activated when driving. The driver is informed that high pitches indicate the time gap is too large, and low pitches indicate the time gap is too small. No sound indicates no adjustment is necessary.

Experimental Protocol

The lead vehicle driving protocol is a subset of the tests designed in [17]. Oscillatory tests are designed to collect transient data to understand how the CAN coach performs under non-constant headway and lead vehicle speeds. Two speed oscillation magnitudes are considered, namely 2.2 m/s (5 mph) and 4.5 m/s (10 mph). The two protocols are described as follows.

- *Small Dips*: The speed is fluctuated between 29.0 m/s (65 mph) and 26.8 m/s(60 mph), with each speed being held for at least 30 seconds.

- *Large Dips*: The speed is fluctuated between 26.8 m/s (60 mph) and 31.3 m/s (70 mph) with each speed being held for at least 30 seconds.

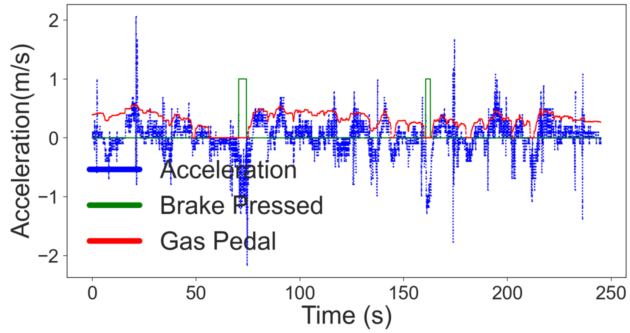

FIGURE 3 : Acceleration, gas, and brakes signals during the Normal Driving Low Dips test. Brake Pressed is a binary signal indicating if the brake is pressed. Gas Pedal is a proportional signal from 0-1.

RESULTS

This section discusses the tests to assess the CAN data validity, and the primary results from driving under Normal Driving, Instructed Driving, Coached Driving. It compares across experimental conditions and protocols.

Data Validity

The data required for this experiment needed to be precise, accurate, and recent. Without an official manufacturer's DBC, one must always question the identity, accuracy, and precision of the measurements found on the CAN. Here we assess the data validity.

Speed

The CAN velocity measurements corresponding to the speedometer are precise to 4+ significant figures. The speed measurements from the CAN bus are shown to be valid by comparing the integral of the velocity data with a cumulative distance travelled between data points from a high-precision GPS sensor used for validation. There was a 0.01% difference in the integrated velocity from the GPS distance travelled measurement.

Gas, Brakes, and Acceleration

This section confirms the accuracy of the gas pedal and brake data, which are validated using acceleration data also available on the CAN. While stationary, the brake pedal measurements are tested, where it is confirmed that the brake value linearly scales from 0 when not pressed to the largest value when the pedal is maximally pressed. The gas pedal data was validated in relation to the acceleration data. In Figure 3 we show how the gas pedal and brake usage are corroborated with acceleration data. When the proportion of the gas pedal is higher, the acceleration goes up. When the gas pedal is released, the acceleration goes down gradually. When the brakes are engaged, the acceleration goes down more sharply. The acceleration was independently validated for accuracy by matching its integral to the velocity signal. The maximum value of the gas pedal is not confirmed due to the desire to avoid extreme engine RPMs.

Space Gap
CAN radar messages are used to measure space gap accurately in [18], which are also confirmed with high-precision GPS measurements. Since these experiments were performed with a different year-model vehicle, here we take additional steps to confirm the space gap measurements are accurate. Ten measurements are made in approximately one meter intervals from 0-10 meters when the vehicle is at rest and an object is ahead at the set distance. All of the CAN measurements match the true object distance empirically measured with a tape measure to within +/- 2 centimeters. This confirms that the measurements were accurate and precise.

Normal Driving

In this section, we show the results of the tests run under Normal Driving (see the Section *Methods*). Figure 4a and Figure 4b show the timeseries of the velocity of the following car (blue), the lead car (green), and the space gap between the vehicles for the Low Dips and High Dips tests respectively. It is clear in Figure 4b that when the lead car (green) jumps up from 26.8 m/s (60 mph) and 31.3m/s (70 mph), the following car (blue) is delayed in adjusting speed, overshoots, and then settles in matching the speed of the lead vehicle.

In Figure 4c, a Velocity vs. Space Gap scatter plot for the Normal Driving condition under both the High and Low Dips tests. Since the time gap is defined as the space gap divided by the velocity, the slope of a line on this plot is equivalent to the time gap. To aid interpretability, two lines in the figure are added, corresponding to time gaps of 2.2 s (lower line) and 1.8 s (upper line). The points on the scatter plot are plotted with the color assigned to their time on the colorbar. Normal Driving is primarily outside the time gap lines, showing it is primarily not close to two seconds time gap at varying speeds and space gaps.

Figure 4d is a kernel density estimate of the time gap for the Normal Driving experiments. *Kernel density estimation* (KDE) is a non-parametric statistical method to estimate the *probability density function* (PDF) of a random variable. The time gap for Normal Driving has a mean of 1.61 s, a standard deviation of 0.27 s, and an interquartile range of 0.32 s. We conclude the driver does not follow a two second time gap under normal driving.

The roughly unimodal shape of the KDE shows that the driver's following distance is dictated by time gap, not a fixed distance. This fact is supported in Figure 4c, where the range of velocity (23.3-33.9 m/s) is generally coupled with the range of space gap (30.0-65.6 m).

Instructed Driving

In this section, we report the results of the experiments when the driver operates in under the Instructed Driving condition.

Under Instructed Driving, the driver must use rough estimation techniques to both figure out their current time gap and then actuate based on that estimate. For example, this can be done by finding a fixed point ahead just passed by the vehicle ahead, and counting the time until passing that point. Although it is possible for skilled drivers to perform these calculations reliably, it requires the reliable estimation of two seconds passing, and is more difficult to adjust to time gaps at sub-second resolution (e.g., a time gap of 1.8 s).

Figure 5 summarizes the results of the experiments. Figure 5a shows the speed-space gap plot, and shows that the time gaps under instructed driving are generally near the instructed two second time gap. Figure 5b shows the corresponding KDE. The time gap mean is 2.09 s, interquartile range 0.28 s, and standard deviation 0.23 s. Using an average speed of the follower,

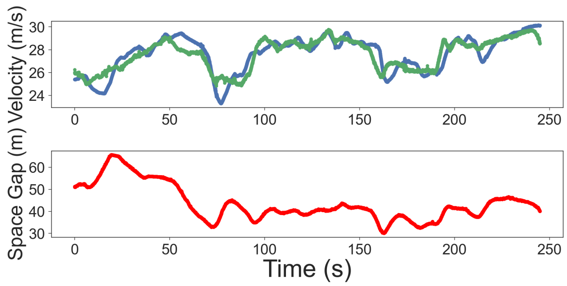

(a) Velocity of lead car (Green) and following car (Blue), and space gap for Normal Driving Low Dips test.

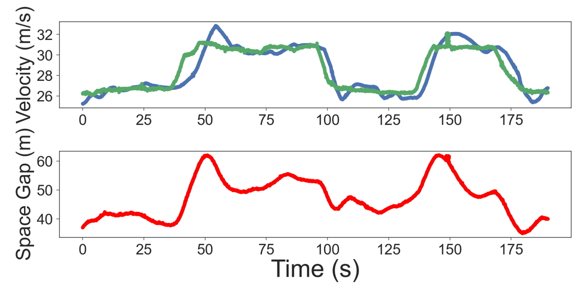

(b) Velocity of lead car (Green) and following car (Blue), and space gap for Normal Driving High Dips test.

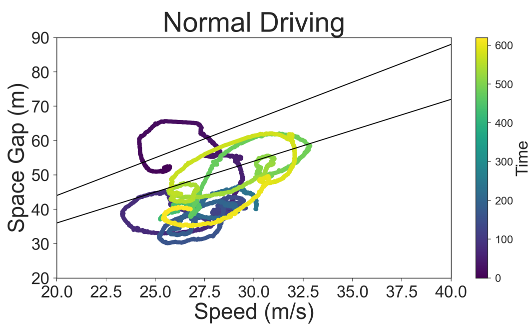

(c) Speed versus space gap for Normal Driving.

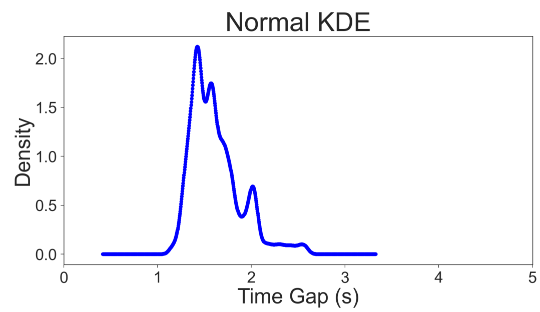

(d) Kernel density estimate for Normal Driving tests.

FIGURE 4 : Summary of time gap for the Normal Driving condition.

0.09 s mean time gap error corresponds to about 2.5 m error in position, or about a half car length. Similarly, the standard deviation of 0.23 s corresponds to about 1.5 car lengths (6.5 m) positional error at the tested following speeds. Overall the performance by the driver with respect to the mean time gap is very good, but the standard deviation (0.23 s) and interquartile range (0.28 s) are fairly wide, similar to the Normal Driving condition where standard deviation is 0.27 s and interquartile range is 0.32 s.

Coached Driving

In this section we examine the Coached Driving test. In Figure 6 we can see a summary of the Coached Driving condition. The driver follows the two second time gap very closely. The points in the scatter plot Figure 6a are almost totally inside the lines, i.e., very close to two seconds time gap at varying speed and spacing. The KDE of the time gap is centered and peaked tightly over the targeted 2.0 seconds. The mean time gap is 1.99 s, with a standard deviation of 0.13 s. The interquartile range is 0.09 s. Again, using an average speed for conversion, the 0.01 s mean time gap error corresponds to a spacing error of less than 28 centimeters. The standard deviation is more than two times smaller and interquartile range is more than three times smaller than Normal Driving and Instructed Driving, respectively. An interquartile range of 0.09s translates to about 2.5 m at the testing speeds, so the middle 50% of the data is less than +/- 1.25 m away from 1.99 s of time gap.

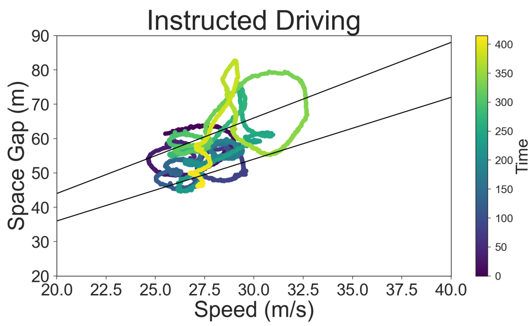

(a) Scatter plot of Instructed Driving.

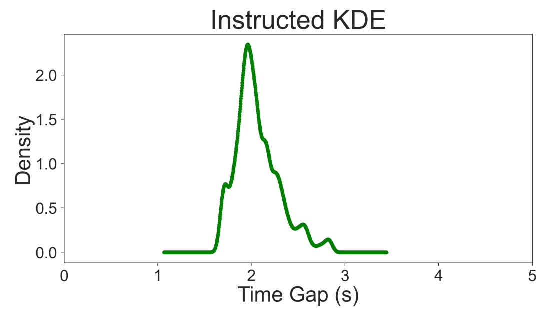

(b) KDE of Instructed Driving tests.

FIGURE 5 : Summary of time gap for Instructed Driving (no feedback) tests.

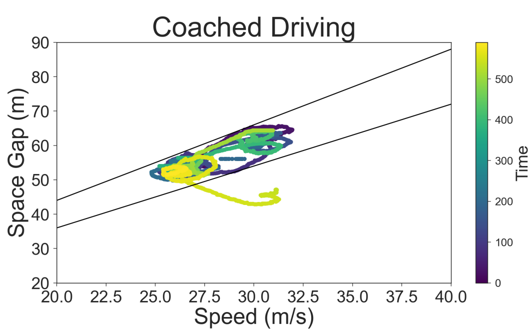

(a) Scatter plot of Coached Driving.

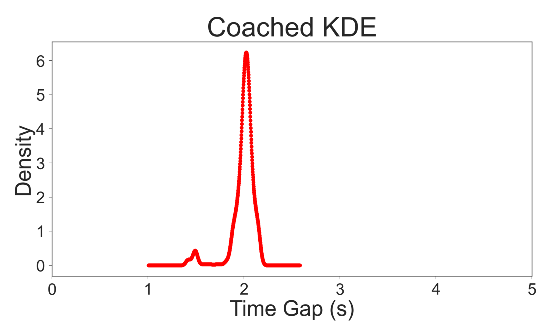

(b) KDE of Coached Driving tests.

FIGURE 6 : Summary of time gap for Coached Driving tests.

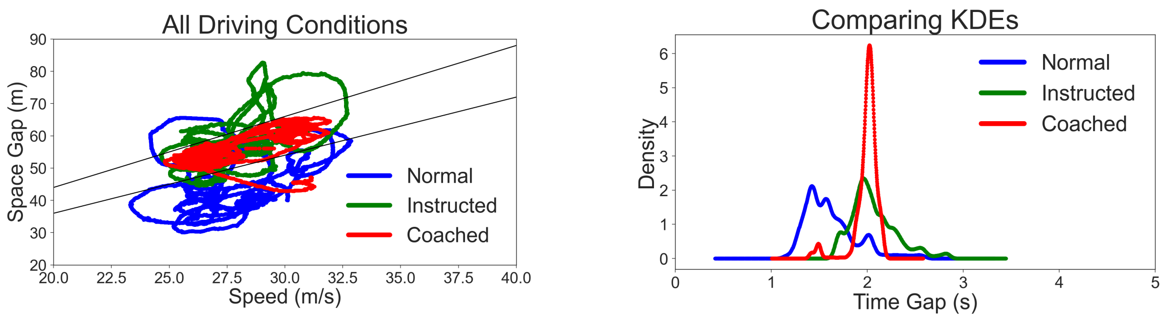

(a) Scatter plot of all three driving conditions.

(b) KDEs for all three driving conditions.

FIGURE 7 : Summary plots comparing the time gap behavior in all three conditions.

Comparing Driving Conditions

In this section we compare different conditions to one another. We will examine the three conditions' time gaps, and acceleration and brake use.

Comparing Time Gaps

Figure 7 shows the data from the three driving conditions overlaid on top of each other. This figure allows us to clearly see the trends from condition to condition.

Figure 7b allows us to easily compare the time gap behavior from all three conditions. Normal driving establishes the natural driver following behavior, which has a mean time gap lower than two seconds with a wide variance. Instructed Driving noticeably improved the accuracy of the time headway, but the variance remains large. Coached Driving demonstrates that the feedback is successful in changing the driving behavior by giving the driver continuous information about the state of the vehicle. The Coached Driving mean is almost exactly two seconds time gap (1.99) with standard deviation and interquartile range three times smaller than Normal Driving and Instructed Driving.

Gas Pedal and Brake Usage

In this section we examine the use of the pedals by the driver in all three experimental conditions. For example, a potential negative finding could be that the driver uses the throttle and brake significantly more under Instructed or Coached conditions in order to meet the control objective. This would be unlikely to be maintainable for extended periods of time and could have implications on vehicle wear and tear. We observe that Instructed and Coached driving result in insignificant changes to the actuation amount. There are 6, 3, and 4 brake events recorded under Normal, Instructed, and Coached Driving respectively. The small number of Brake events is likely due to the high speed freeway driving setting under which the tests are conducted.

Consequently time gap adjustment primarily occurs through the gas pedal. Figure 8a shows a normalized histogram of the gas pedal proportion. The distribution of gas pedal use is approximately the same. For Normal, Instructed, and Coached Driving gas pedal position the mean was 0.29, 0.27, and 0.29, and the standard deviation was 0.14, 0.13, and 0.15 respectively. These preliminary results suggest that the driving condition has no effect on the actuation behavior of the driver.

Figure 8b compares gas pedal and time gap data from the Coached Driving High Dips test.

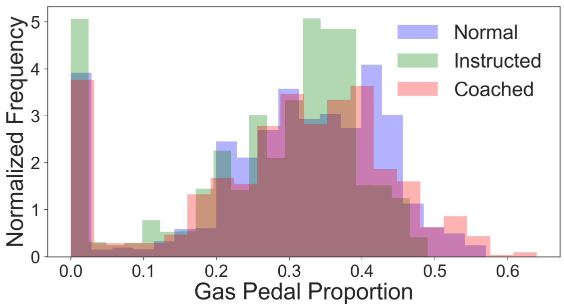

(a) Gas pedal proportion in each condition.

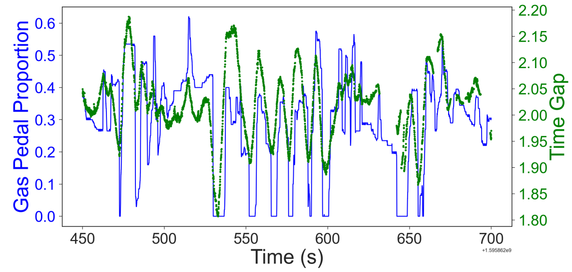

(b) Gas pedal proportion plotted with time gap.

FIGURE 8 : Figure 8a is an histogram of the gas pedal use in the three experimental conditions. The gas pedal proportion ranges from 0-1. The frequency of the values is normalized to account for the slightly different lengths of time for each condition. Figure 8b is a comparison of the gas pedal proportion and the time gap. It shows how the two signals are associated with each other.

The clear pattern shows that the gas pedal use dictated the adjustments to time gap.

Comparing Driving Tests

This section compares the driving tests in each condition and under each protocol. Normal Driving (Figure 9a) under high and low dips show two similar shaped KDEs, although a shift is observed in the mean time gap to a longer time gap under the higher speed dips. This is intuitive that the repeated larger speed variations resulted in the driver shifting away from the low time headway in the low dips test. Under Instructed driving, the time gap means are similar, but the time gap variance is notably larger for the high dips test (Figure 9b). In Figure 9c, the driving under CAN coach shows negligible change in the mean and standard deviation of the time gap. The consistency of the CAN Coach under both experimental protocols is distinct compared to the other testing condition without feedback. Figure 9d compares the duration of all the tests in all of the experimental conditions. They all have similar durations.

CONCLUSIONS AND FUTURE WORK

This article illustrates an application that is able to change driver car following behavior using only data on the CAN. It is a proof of the concept that CAN data can be processed in near real-time and effect driver behavior to a desired outcome, in this case a time gap. The CAN Coach was introduced as a system to provide real time auditory feedback to achieve the desired time gap using data from the CAN. Using the feedback, a single driver is shown to maintain a time gap with smaller error and variance compared to normal driving and driving under instructions to maintain a time gap (but without feedback).

The main limitation of this work is the generalizeability of the results, which require testing on more drivers. Based on the preliminary evidence shown in this work, we are planning additional experiments to test across a range of drivers. We are also interested to test other control objectives, including ones designed to improve the so-called string stability of the driver. String stable drivers may reduce the presence of phantom traffic jams, which can provide additional aggregate safety and and traffic flow benefits.

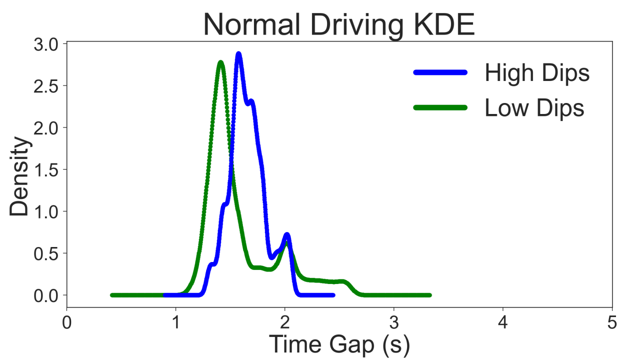

(a) Normal Driving tests KDEs.

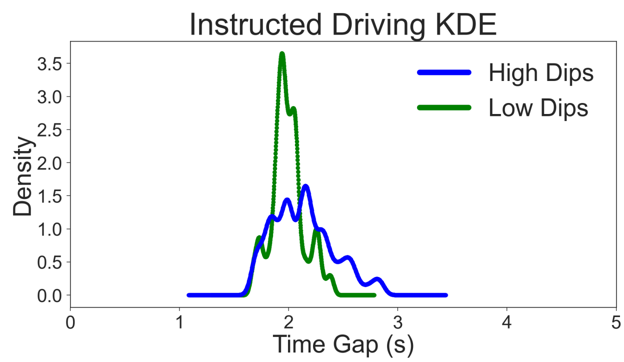

(b) Instructed Driving (no feedback) tests KDEs.

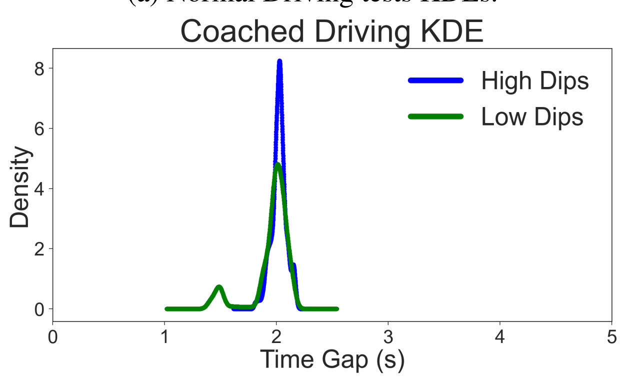

(c) Coached Driving tests KDEs.

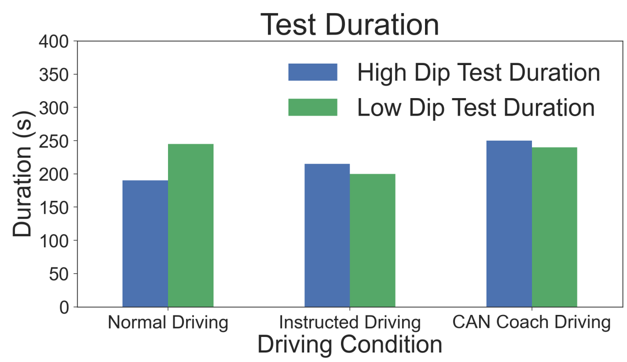

(d) Duration of each test for each condition.

FIGURE 9 : Comparing the differences between the two driving tests. 9a, 9b, and 9c show KDEs for each of the driving tests. 9d compares the two tests by duration of the tests in each condition.

ACKNOWLEDGEMENTS

The authors are grateful to the Compositional Systems Lab at the University of Arizona for their collaboration on interpreting and using CAN data in transportation cyber physical systems. The authors are grateful to Will Barbour, Caroline Janssen, and Yanbing Wang for their assistance in executing driving experiments.

This material is based upon work supported by the National Science Foundation under Grant No. CNS-1837652 (Work). This material is based upon work supported by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) under the Vehicle Technologies Office award number CID DE-EE0008872 (Nice). The views expressed herein do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

References

[1] U. Fugiglando, E. Massaro, P. Santi, S. Milardo, K. Abida, R. Stahlmann, F. Netter, and C. Ratti, "Driving behavior analysis through can bus data in an uncontrolled environment," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 2, pp. 737–748, 2018.

[2] T. Wakita, K. Ozawa, C. Miyajima, K. Igarashi, K. Itou, K. Takeda, and F. Itakura, "Driver identification using driving behavior signals," *IEICE Transactions on Information and Systems*, vol. 89, no. 3, pp. 1188–1194, 2006.

[3] P. Bazilinskyy, J. Stapel, C. de Koning, H. Lingmont, T. de Lint, T. van der Sijs, F. van den Ouden, F. Anema, and J. de Winter, "Graded auditory feedback based on headway: An on-road pilot study," *Varieties of interaction: from User Experience to Neuroergonomics*, 2018.

[4] P. Bazilinskyy, P. Larsson, E. Johansson, and J. C. de Winter, "Continuous auditory feedback on the status of adaptive cruise control, lane deviation, and time headway: An acceptable support for truck drivers?," *Acoustical Science and Technology*, vol. 40, no. 6, pp. 382–390, 2019.

[5] B. D. Seppelt and J. D. Lee, "Making adaptive cruise control (acc) limits visible," *International journal of human-computer studies*, vol. 65, no. 3, pp. 192–205, 2007.

[6] S. H. Fairclough, A. J. May, and C. Carter, "The effect of time headway feedback on following behaviour," *Accident Analysis & Prevention*, vol. 29, no. 3, pp. 387–397, 1997.

[7] M. Risto and M. H. Martens, "Supporting driver headway choice: The effects of discrete headway feedback when following headway instructions," *Applied ergonomics*, vol. 45, no. 4, pp. 1167–1173, 2014.

[8] M. J. Sullman, L. Dorn, and P. Niemi, "Eco-driving training of professional bus drivers–does it work?," *Transportation Research Part C: Emerging Technologies*, vol. 58, pp. 749–759, 2015.

[9] C. Johnsson, A. Laureshyn, and T. De Ceunynck, "In search of surrogate safety indicators for vulnerable road users: a review of surrogate safety indicators," *Transport Reviews*, vol. 38, no. 6, pp. 765–785, 2018.

[10] A. Laureshyn, Å. Svensson, and C. Hydén, "Evaluation of traffic safety, based on micro-level behavioural data: Theoretical framework and first implementation," *Accident Analysis & Prevention*, vol. 42, no. 6, pp. 1637–1646, 2010.

[11] V. Astarita, D. C. Festa, V. P. Giofrè, and G. Guido, "Surrogate safety measures from traffic simulation models a comparison of different models for intersection safety evaluation," *Transportation research procedia*, vol. 37, pp. 219–226, 2019.

[12] J.-J. Martinez and C. Canudas-de Wit, "A safe longitudinal control for adaptive cruise control and stop-and-go scenarios," *IEEE Transactions on control systems technology*, vol. 15, no. 2, pp. 246–258, 2007.

[13] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Transactions of the ASME–Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.

[14] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[15] J. Scott and R. Gray, "A comparison of tactile, visual, and auditory warnings for rear-end collision prevention in simulated driving," *Human factors*, vol. 50, no. 2, pp. 264–275, 2008.

[16] S. M. Pampel, S. L. Jamson, D. L. Hibberd, and Y. Barnard, "Old habits die hard? the fragility of eco-driving mental models and why green driving behaviour is difficult to sustain," *Transportation research part F: traffic psychology and behaviour*, vol. 57, pp. 139–150, 2018.

[17] G. Gunter, D. Gloudemans, R. E. Stern, S. McQuade, R. Bhadani, M. Bunting, M. L. Delle Monache, R. Lysecky, B. Seibold, J. Sprinkle, *et al.*, "Are commercially implemented adaptive cruise control systems string stable?," *IEEE Transactions on Intelligent Transportation Systems*, 2020.

[18] Y. Wang, G. Gunter, M. Nice, and D. B. Work, "Estimating adaptive cruise control model parameters from on-board radar units," 2019.