Third Generation Simulation (TGSIM) Data: A Closer Look at The Impacts of Automated Driving Systems on Human Behavior

www.its.dot.gov/index.htm

Final Report – May 2024
FHWA-JPO-24-133

U.S. Department of Transportation

Produced by University of Illinois at Urbana-Champaign
U.S. Department of Transportation
Office of the Assistant Secretary for Research and Technology
Federal Highway Administration

Notice

Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
FHWA-JPO- 24-133		

4. Title and Subtitle	5. Report Date
Third Generation Simulation Data (TGSIM): A Closer Look at the Impacts of Automated Driving Systems on Human Behavior	May 2024
	6. Performing Organization Code

7. Author(s)	8. Performing Organization Report No.
Alireza Talebpour (0000-0002-5412-5592) Hani S. Mahmassani (0000-0002-8443-8928) Samer H. Hamdar (0000-0001-6896-367X)	

9. Performing Organization Name and Address	10. Work Unit No. (TRAIS)
University of Illinois at Urbana-Champaign Urbana, Illinois 61801	11. Contract or Grant No.

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
Federal Highway Administration Office of Safety Research and Development 6300 Georgetown Pike McLean, VA 22101	14. Sponsoring Agency Code HOIT-1

15. Supplementary Notes

The project CORs were Rachel James and John Hourdos. The JPO Program Manager was Hyungjun Park

16. Abstract

This study collected and processed accurate trajectory datasets capable of characterizing human-automated vehicle interactions under a diverse set of scenarios in highway and city environments. Multiple methods were utilized to collect data: fixed location aerial videography, moving aerial videography, and infrastructure-based videography. Fixed location aerial videography was used to collect data on I–90/I–94 in Chicago, IL; this methodology collected data using a stationary helicopter hovering over a segment of interest at a higher altitude to capture more roadway in a single image. Moving aerial videography was used to collect four datasets on I–90/I–94 and I–294 in Chicago, IL; this methodology collected data using a helicopter following a string of automated vehicles at a lower altitude as they drive a longer roadway segment. Infrastructure-based videography was utilized to collect data on I–395 in Washington D.C. and on the George Washington University campus (located in the Foggy Bottom neighborhood); this methodology utilized multiple overlapping cameras located on overpasses and buildings stitched together and synchronized, creating a comprehensive field of view of the study area. Extracting multiple vehicle trajectories from video data raises a set of methodological and practical challenges that vary across the three data collection approaches. This report details the methodologies used to collect data and extract vehicle trajectories. Additionally, this report provides a sample analysis showing how the collected data can be utilized in modeling automated vehicle driving behavior. Finally, this report offers a discussion on lessons learned during this project, including data collection set-up, instrumentation, and experimental design.

17. Keywords	18. Distribution Statement
Vehicle Trajectory Data, Automated Vehicles, Human-Automated Vehicle Interactions	No restrictions.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
N/A	N/A	79	

Acknowledgements

The authors express their sincere appreciation to John Hourdos, Rachel James, and Gene McHale, along with the other esteemed researchers and staff at the FHWA Turner Fairbanks Research Laboratory, for their support and expert guidance throughout the entirety of this project. Their invaluable contributions have significantly enriched this work. We are deeply grateful for their unwavering commitment and support.

U.S. Department of Transportation
Office of the Assistant Secretary for Research and Technology
Intelligent Transportation Systems Joint Program Office

Table of Contents

List of Tables

List of Figures

Executive Summary

In an era marked by rapid advancements in automation technologies, the growing usage of Automated Driving Systems (ADS) and Advanced Driver Assistance Systems (ADAS) is reshaping the dynamics of human-vehicle interactions. This study focuses into these interactions through an extensive data collection effort. The collected dataset aims to capture the interaction between humans and automated vehicles across different scenarios, environments, and levels of automation.

This project employed three videography methods to capture different driving scenarios and levels of autonomy. (1) Fixed location aerial videography was used to collect data with mixed SAE Level 2 automated vehicles with active automated lane changing functions mixed in the traffic flow on I–90/I–94 in Chicago, IL; this methodology collected data using a stationary helicopter hovering over a segment of interest at a higher altitude to capture more roadway in a single image. (2) Moving aerial videography was used to collect four datasets on I–90/I–94 and I–294 in Chicago, IL, including SAE Level 1 and Level 2 automated vehicle fleets mixed in the traffic; this methodology collected data using a helicopter following a string of automated vehicles at a lower altitude as they drive a longer roadway segment. (3) Infrastructure-based videography was utilized to collect data on I–395 in Washington D.C. and on the George Washington University campus (located in the Foggy Bottom neighborhood); this methodology utilized multiple overlapping cameras located on overpasses and buildings stitched together and synchronized, creating a comprehensive field of view of the study area.

This report also presents the details of the corresponding data processing framework. The data processing framework was designed to handle the complexity and volume of the collected video data, ensuring the extraction of accurate, meaningful insights into vehicle trajectories. The steps include data preprocessing, object detection, object tracking, image stabilization, trajectory construction, data cleaning, and data validation. Data preprocessing is the initial phase involved in extracting raw images from video frames, followed by the generation of reference images to serve as a basis for coordinate transformation. Then, the object detection and tracking steps utilized deep-learning-based algorithms to detect and track vehicles across frames. Followed by the image stabilization and trajectory construction steps for stabilizing and aligning video data with a fixed global coordinate. Next, the data cleaning was performed to remove the anomalies and ensure consistency in the final datasets. The collected data are manually evaluated, and any potential issues are mitigated through an accurate error correction process to ensure the validity of the extracted trajectories.

A sample analysis based on the collected datasets is also performed. The results suggested that human drivers have different behaviors in the presence of ADS-equipped vehicles, such as changes in car-following patterns, which can impact traffic flow dynamics. These findings highlight the need for an improvement in both traffic flow models and safety analysis methods to accommodate the presence of ADS-equipped vehicles.

This report concludes with a reflective discussion on the lessons learned throughout the project, covering aspects such as data collection setup, instrumentation, and experimental design. These insights aim to

support future data collection and trajectory extraction efforts and offer guidance for the continuous improvement in future data collection efforts.

Chapter 1. Introduction

The past few years have seen breakthroughs in automation technologies, enabling the development of advanced driver assistance system (ADAS)-equipped vehicles and automated driving system (ADS)-equipped vehicles[1]. ADAS are intelligent systems that reside inside the vehicle and assist the main driver in a variety of ways (*Kala, 2016*). While there are various implementations of ADAS and these features can change from one manufacturer to another, SAE Level 1™-equipped vehicles are generally capable of providing longitudinal **or** lateral support to the driver, which includes adaptive cruise control (ACC) that automatically adjusts the vehicle's speed to keep a pre-set distance from the vehicle in front of it, or lane keeping assist (LKA) that automatically and gently steers to prevent the vehicle from departing the lane (*NHTSA, 2021*). SAE Level 2™-equipped vehicles can provide the driver with both longitudinal and lateral support simultaneously. SAE Level 3™ and 4 are often referred to as highly automated vehicles. Both SAE Level 3 and 4 vehicles can complete a trip from an origin to a destination as long as they operate withing their certified operational design domain (ODD) (*Czarnecki, 2018*). The key difference is that SAE Level 3 vehicles still require driver/human oversight in case the system fails, while SAE Level 4™ vehicles do not need any human oversight or intervention when they operate in conditions they have been certified for. SAE Level 5 vehicles do not need any human oversight under any operating condition.

As of May 2018, at least one ADAS feature is available on 92.7 percent of new vehicles available in the U.S. (*AAA, 2019*) According to the U.S. Autonomous Vehicle Tracker, in the first half of 2022, more ADAS-equipped vehicles were sold than Level 0 nonautomated vehicles with:

- Level 2+ ADAS-equipped vehicles accounting for 3.9 percent of total U.S. car sales.
- Level 2 ADAS-equipped vehicles accounting for 46.5 percent of total U.S. car sales.
- Level 1 ADAS-equipped vehicles accounting for 19.8 percent of total U.S. car sales.
- Level 0 nonautomated vehicles accounting for 29.8 percent of total U.S. car sales (R&A Electronics, 2023).

Unlike nonautomated vehicles[2], ADAS- and ADS-equipped vehicles' decision logic relies on a broad set of data they receive from onboard sensors. The algorithms used for motion control in ADAS-equipped vehicles utilize design approaches that can result in decision-making algorithms that are fundamentally different from humans' decision-making logic (*Lu et al., 2022*). Thus, the introduction of ADAS-equipped vehicles to the transportation system has resulted in a new set of interactions among roadway users (e.g., human-AV interactions[3]) that did not exist a few years ago. As the number of vehicles equipped with

[1] This report uses automated vehicle (AV) as an umbrella term inclusive of ADAS- and ADS-equipped vehicles.
[2] Nonautomated vehicle refers to SAE Level 0™-equipped vehicles with no automated driving support features.
[3] In this report, human-AV refers to SAE Level 0 and SAE Level 1–2 interactions, although humans are still responsible for the driving task in SAE Level 1–2 vehicles, even when the vehicles are using the automated support future.

these technologies increases on our roadways, more and more interactions between these technologies and humans are expected.

These interactions will shape the future of traffic flow theory and fundamentally change the way we model traffic flow dynamics. Unfortunately, our understanding of these interactions remains limited and there have not been many efforts to characterize them (*Mahmassani et al., 2018*). Capturing the interactions between human drivers and AVs is essential in addressing many future research questions, including, but not limited to:

How will AVs impact traffic? Past research assumed AVs will have an instant positive impact on traffic flow (*Yu et al., 2021*), however, such claims need to be verified based on the real design limitations of these technologies (as opposed to the hypothetical potential of these technologies).

What opportunities exist for new congestion management strategies with various AV market penetration rates? A deeper understanding of human-AV interactions can be leveraged to develop novel strategies that have not previously been possible.

How can AV data be used to improve traffic and planning models? The sensors on AVs may increase the availability of trajectory level vehicle data, which can be utilized to calibrate car-following models and lane-changing models in mixed traffic and increase the accuracy of microscopic simulation tools.

- How do ADAS-equipped vehicles and nonautomated vehicles interact? Some limited evidence exists that humans exhibit different behavior around AVs compared to their interactions with other non-automated vehicles (e.g., Rahmati et al.(*2019*) found that human drivers choose a smaller headway when following ACC-driven vehicles as opposed to following other human-driven vehicles). However, it is not yet clear whether the underlying mechanisms of human-AV interactions are fundamentally different from human-human interactions.

Research suggests that although ADAS features were only available on 10 percent of vehicles in use worldwide in 2020, the penetration rate is anticipated to increase to 30 percent in 2025 and 50 percent by 2030 (*Canalys, 2021*). Given that ADAS-equipped vehicles are rapidly gaining market penetration, understanding the human-ADAS interaction is critical for managing the transportation system in the mixed traffic environment.

The first step toward addressing the aforementioned limitations is to capture and investigate human-AV interactions in a real-world setting (*Mahmassani et al., 2018*). Unfortunately, most of the existing studies focus on navigating automated vehicles safely and on predicting the movement of human-driven vehicles, without explicitly focusing on the potential differences between human-human and human-AV interactions (*Deo et al., 2018, Sadigh et al., 2016*). Moreover, while extensive data have been available from human-human interactions (e.g., NGSIM (*Punzo et al., 2011*), HighD (*Krajewski et al., 2018*), pNeuma *(Barmpounakis et al., 2020*), etc.), data from human-AV interactions have been quite limited. The existing datasets that contain AV data usually focus on one of two areas:

- Primary focus on the behavior of AVs and disregard human-AV interactions (e.g., OpenACC (*Makridis et al., 2021*) dataset which is focused on the behavior of ACC-driven vehicles).
- Primary focus on training perception systems and do not contain the necessary elements (e.g., accurate distance to human-driven vehicles or vehicle speed) required to understand driver behavior in the vicinity of these vehicles (e.g., various autonomous driving companies have shared datasets from operating SAE Level 3 and SAE Level 4 autonomous vehicles on the road; e.g. Waymo Open

Dataset (*Sun et al., 2020*), Ford Autonomous Vehicle Dataset (*Agarwal et al., 2020*), nuScenes (*Caesar et al., 2020*), BDD100k (*Yu et al., 2020*), ApolloScape (*Huang et al., 2019*), Level 5 Open Data (*Kesten et al., 2019*) etc.).

It is also noteworthy that while some datasets were gathered from operating ACC systems in traffic (*Makridis et al., 2020*), they were mostly collected in a closed course environment and focused on the ACC vehicles and the stability of the ACC system rather than explicitly looking into the impacts of ACC-driven vehicles on human-driven vehicles in their vicinity.

A detailed list of available AV datasets and their characteristics was offered by Yurtsever et al. (*2020*) and Yin and Berger (*2017*). Accordingly, none of the aforementioned datasets can be directly utilized to study the impacts of AVs on human drivers in their vicinity.

To the best of the authors' knowledge, at the time of the writing of this report, the only available trajectory dataset from AV operations in a real-world setting for understanding and characterizing human-AV interactions was collected by Khajeh-Hosseini et al. (*2022*) in Austin, TX. The authors collected vehicle trajectory data based on aerial videography while three ADAS-equipped vehicles with ACC actively operating in the traffic stream. Their dataset, however, contained only five instances of ACC-driven vehicles, each for less than a minute, and did not result in reliable conclusions about the impacts of ACC-driven vehicles on human drivers.

In parallel to the work presented in this document, a second project was awarded by FHWA with a similar objective. The "Advanced Driver Assistance System-Equipped Vehicle Datasets Collected in Central Ohio" project developed a methodology to collect and process the raw vehicle sensor data to extract trajectory data from SAE Level 2 instrumented as well as from all adjacent vehicles in the field of view of the cameras carried by the instrumented vehicle. This data collection effort used both a discreet production ADAS-equipped vehicle and a readily identifiable ADAS-equipped vehicles where the sensor stack was visible to adjacent vehicles. The project collected the dataset under diverse road and traffic conditions with a broad spectrum of roadway types, roadway conditions, traffic densities, and speed limits at locations around Columbus, OH. Although the project final report is still in the publishing stage, the produced datasets are available at the ITS DataHub[4].

Project Objectives

Accordingly, to address the above research gaps, this study focuses on collecting data from human-AV interactions in a real-world setting. Multiple trajectory datasets were collected with various SAE Levels of automation (SAE Level 1 and SAE Level 2) in Chicago, IL, and Washington D.C. The datasets cover various traffic flow regimes (from free-flow to fully congested traffic) and offer insight into human-AV as well as human-human interactions in the vicinity of ADAS-equipped vehicles. These datasets can be utilized for various traffic flow and safety analyses, including but not limited to:

- Investigating the impacts of SAE Level 1 and SAE Level 2 ADAS-equipped vehicles on human behavior.

[4] https://data.transportation.gov/Automobiles/Advanced-Driver-Assistance-System-ADAS-Equipped-Tw/vhz2-exyi/about_data

- Investigating the string stability of ACC-driven vehicles in real-world settings.
- Investigating the dynamics and impacts of ADAS-heavy vehicle interactions on traffic flow and safety.
- Investigating the underlying dynamics of mandatory and discretionary lane-changing maneuvers.
- Investigating the car-following behavior under various traffic flow regimes.
- Investigating the microscopic and macroscopic traffic flow analyses.

In addition to collecting datasets in a variety of naturistic environments, this study offers the following key contributions:

- A robust methodology is proposed to extract accurate vehicle trajectory data based on high-altitude videography in a highway environment (with a focus on capturing human-AV interactions).
- A detailed experimental design approach is presented based on extensive experimentation with different design approaches to accurately capture data from vehicle interactions. Note that new data collection approaches are also developed as part of this study utilizing a moving helicopter and a multitude of fixed cameras working collectively.
- A detailed discussion on lessons learned from multiple data collection efforts is offered.
- An example of how this data can be utilized to improve the quality of microscopic simulation tools by introducing a new hybrid ACC model that captures disengagement instances is provided to better illustrate the value of the collected data.

Table 1: Data Collection Summary.

Site Location	Duration	Time of Day	Segment Length	Number of Lanes	Vehicle Type	Number of AVs	Data Collection Method	Altitude
I–90/I–94	2h	3 p.m.–5 p.m.	4.0 km	4	SAE Level 2	3	Vehicle Following	~200m
I–90/I–94	2h	4 p.m.–6 p.m.	1.3 km	3 to 6	SAE Level 2	2	Hovering	~335m
I–294	2h	3 p.m.–5 p.m.	4.8 km	4	ACC String	3	Vehicle Following	~300m
I–294	2h	3 p.m.–5 p.m.	4.8 km	4	SAE Level 2	2	Vehicle Following	~300m
Foggy Bottom	3h	2 p.m.–5 p.m.	4 Blocks	Varies	SAE Level 3	1	Infrastructure-Based	Varies
I–395	3h	8:30 a.m.-10:30 a.m.	0.55 km	3	SAE Level 2	2	Infrastructure-Based	Overpass

Data Collection Locations

The team collected data at two different sites in both Chicago and Washington, D.C.:

- I–90/ I–94 in Chicago, IL (see figure 1.a and 1.b).
- I–294 near Hinsdale, IL (see figure 1a. and 1.c).
- I–395 in Washington D.C. (see figure 1.d).
- George Washington University Campus in Washington D.C. (see figure 1.e).

Figures 1.a and 1.d show the geographic relationship between Chicago and Washington D.C. data collection sites, respectively. Table 1 summarizes the key characteristics of each data collection location. These locations were chosen because they include a plethora of rich interactions between vehicles and because of the different geometric road features present such as merging and diverging lanes.

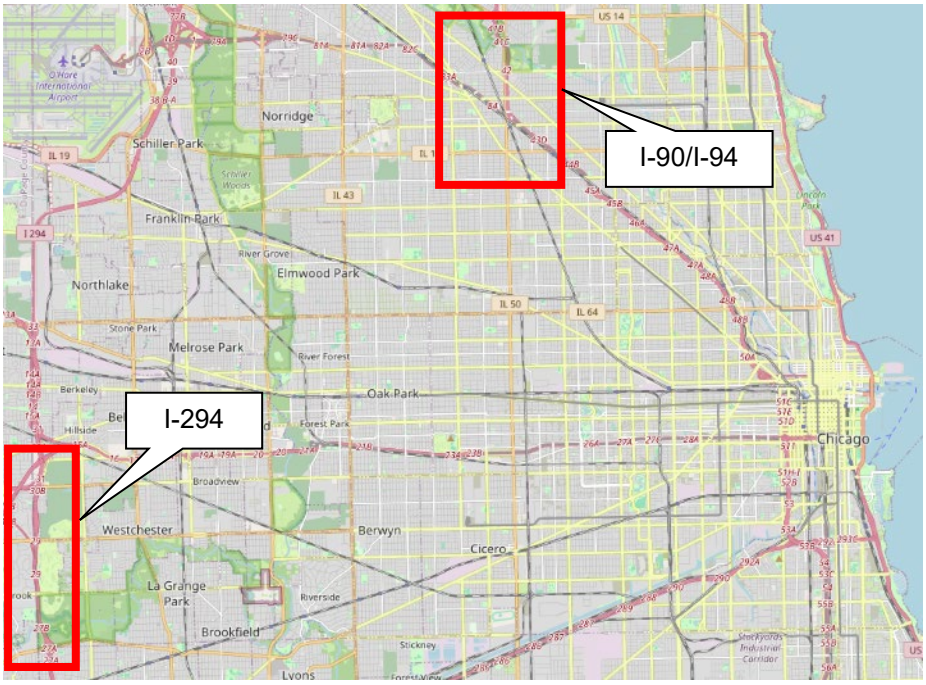

(a) Chicago Map and Location of Two Data Collection Sites.

(b) Location of Data Collection on I-90/I-94, Chicago, IL.

(c) Location of Data Collection on I-294, Chicago, IL.

(d) Location of Data Collection on I-395, Washington, D.C.

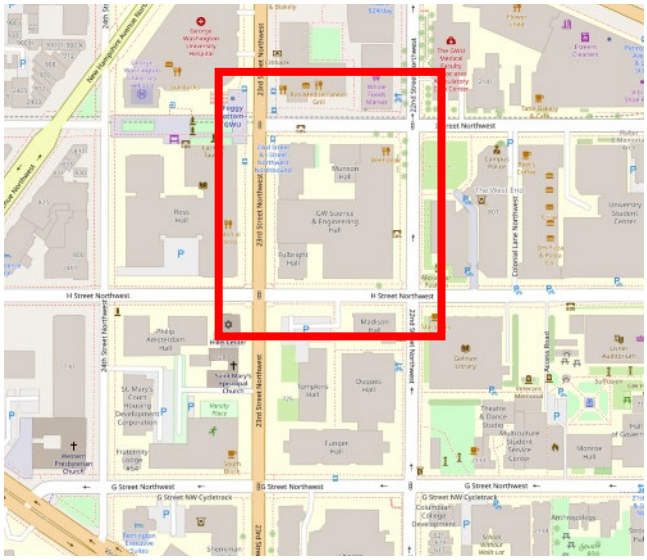

(e) One Block was Utilized for Data Collection in Foggy Bottom, Washington, D.C.

©OpenStreetMap[5] with additions by FHWA.

Figure 1. Map. Data Collection Locations.

I–90/I–94 in Chicago, IL

I–90/I–94 was selected to capture instances of mandatory and discretionary lane changing, as well as last-minute forced lane-changing maneuvers.

This data collection effort operated SAE Level 2 driving automation system-equipped vehicles with partial driver automation capabilities along a segment containing a major weaving section in Chicago, IL.

[5] OpenStreetMap License can be found here: https://www.openstreetmap.org/copyright

I-294 near Hinsdale, IL

I–294 was selected to test human drivers' responses to vehicle strings and automated lane-changing maneuvers. I–294 also contains a high percentage of heavy vehicles.

This data collection effort operated SAE Level 1 and SAE Level 2 driving automation system-equipped vehicles on I–294 to the southwest of Chicago, IL. Figure 1.a and Figure 1.c shows the overview of the segment that was used for data collection.

I–395 in Washington D.C.

I–395 was selected to capture several human-AV interactions, including during automated lane-changing maneuvers. Heavy vehicles' responses to L2 ADAS-equipped vehicles can be also seen on I–395.

This data collection effort operated SAE Level 2 driving automation system-equipped vehicles in a controlled access roadway segment (i.e., urban expressway) covering a complex freeway setting in the eastbound direction of I–395. Figure 1.d and Figure 1.e show the overview of the segment that was used for data collection.

George Washington University Campus (Foggy Bottom) in Washington D.C.

The George Washington University campus was selected because it offers rich interactions among AVs and non-motorized modes of transportation in a city environment.

This data collection effort operated SAE Level 3 driving automation system-equipped vehicles within a city environment in realistic urban intersections. The vehicles were specifically observed performing a series of complex maneuvers at intersections controlled by stop signs and traffic signals, including both protected and permitted left turns. The data collection segment is depicted in the figure 1.e,

Report Organization

The remainder of this report contains details about data collection, processing, and management throughout the life cycle of the project. The remaining chapters are organized as follows:

Chapter 2. Data Collection: This chapter details the data collection procedure followed in this work. It discusses the vehicles used, the methodology, further details on the collection locations, and data processing.

Chapter 3. Data Overview: This chapter presents an in-depth view of the collected data.

Chapter 4. Sample Analysis: This chapter provides a preliminary analysis and results of some of the collected data.

Chapter 5. Lessons Learned: This chapter discusses lessons learned from life cycle of the data collection project, which aims to make the process easier for future researchers.

Chapter 6. Conclusions: This chapter provides a conclusion and summary of the work.

Chapter 2. Data Collection

This chapter discusses the elements and stages of the data collection process, including the data collection vehicles, data collection sites, and the data collection and trajectory extraction methodologies.

Data Collection Vehicles

The first step toward successful data collection is to utilize appropriate vehicles. Unfortunately, this project coincided with the COVID-19 pandemic, which caused several material shortages and other supply chain issues (*The White House, 2021*) and prevented the full development of test vehicles in time to conduct the data collection. Accordingly, this project utilized production vehicles available through commercial vehicle rentals. Different vehicles were utilized for each vehicle automation SAE Level.

SAE Level 1 ADAS-equipped Vehicles with driver assistance features

SAE Level 1 ADAS-equipped vehicles provide continuous assistance with throttle control or steering. Two examples of such driver assistance features ACC and LKA. **Error! Reference source not found.**a presents a sample group of SAE Level 1 vehicles utilized in this study. The vehicles utilized in this study were manufactured after 2020. During the experiments, these vehicles formed a string of ACC-driven vehicles throughout the study segment.

SAE Level 2 ADAS-equipped Vehicles with partial driver automation features

SAE Level 2 ADAS-equipped vehicles provide continuous assistance with throttle control and steering. These vehicles, unlike SAE Level 1 vehicles, can perform both longitudinal and lateral movements between lanes (i.e., lane-changing maneuvers). Figure 2.b presents a sample group of SAE Level 2 vehicles utilized in this study. At the time of data collection, all the vehicles had the 2020.48 software version or newer.

SAE Level 3 ADS-equipped Vehicles with conditional driver automation features

SAE Level 3 ADS-equipped vehicles handle all driving instances within their ODD; however, the driver is still responsible to ensure the safety of the operations and serve as the fallback user if the system makes a request. A similar make and model as SAE Level 2 vehicles were utilized in the related data collection efforts. The key difference, however, is the software version. At the time of data collection, all the vehicles had the 2021.4 software version or newer.

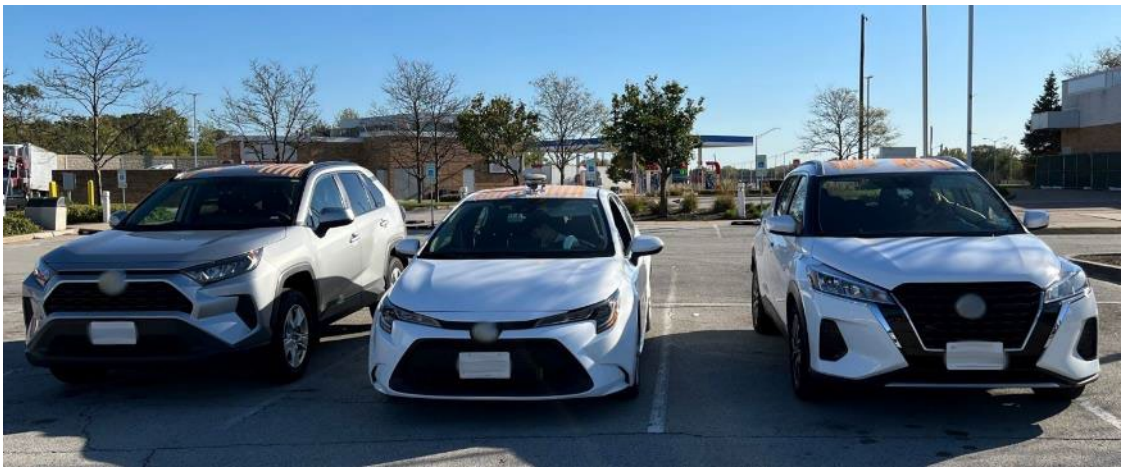

Source: FHWA

(a) SAE Level 1 vehicles

Source: FHWA

(b) SAE Level 2 and SAE Level 3 vehicles

Figure 2. Photos. Sample vehicles used in the data collection.

Video Data Collection Methodologies

This section presents the experimental design and data collection details. It is important to note that data collection methodology and experimental design varied from one location to another to capture comprehensive datasets from the perspective of human-AV and human-human interactions. Accordingly, various approaches have been tested and only the final design is presented in this section. However, some of the lessons learned throughout the experimental design process are also presented in Chapter 5 to help interested readers and researchers with future experimental designs. Table 1 summarizes the key characteristics of the data collection design. All data collection occurred on weekdays.

Due to differences in the data collection locations, the research team utilized three distinct approaches for collecting the video data: fixed location aerial videography, moving aerial videography, and infrastructure-based videography. The following sections discuss the details of these methods.

Fixed Location Aerial Videography

Fixed location aerial videography refers to a hovering helicopter at a relatively stationary location covering a fixed length of a roadway. This approach is probably the most common approach among recent trajectory data collection efforts (e.g., HighD (*Krajewski et al., 2018*) and Austin (*Kesten et al., 2019*) datasets). The trajectory extraction process for this method, discussed in the next section, is also the simplest among the three approaches. This method bypasses issues related to the movement of the recording mechanism, making post-processing steps much less complex and more accurate. However, the round-trip time to get to the start of the segment of interest in the fixed location aerial videography can significantly limit the number of automated vehicle trajectories in the dataset (depending on the number of automated vehicles used in the data collection). This was one of the main issues in the dataset collected by Khajeh-Hosseini et al. (*2022*) in Austin, TX. Moreover, the length of the trajectories can be quite limited considering the lens distortion (avoiding fish-eye lenses increases the detection accuracy, but limits the field of view), flight altitude, and image/video quality. Limited trajectory lengths can negatively impact any study on human-AV interactions due to the limited number of data points (another problem identified by Khajeh-Hosseini et al. (*2022*) in their dataset).

Source: FHWA

Figure 3. Photo. Example of fixed location aerial video field of view

Moving Aerial Videography

In the moving aerial videography approach, a helicopter follows one or more target vehicles as they move through the segment of interest. This approach offers a significant advantage when collecting data from a small set of automated vehicles. Limited trajectory lengths can negatively impact any study on human-AV interactions due to the limited number of data points (a problem identified by Khajeh-Hosseini et al. (*2022*) in their dataset) as well as the limited number of other vehicles interacting with the AVs.

The moving aerial videography approach addresses both limitations and offers long trajectories from AVs while following them as they move through various traffic flow regimes. However, this approach has never been implemented to collect vehicle trajectory data at scale. Thus, the approach becomes much more complex than the fixed location aerial videography method in later post-processing steps and requires

high-performing detection and tracking algorithms, as well as image transformation and stabilization methods.

Source: FHWA

Figure 4. Photo. Example of moving aerial video field of view

Infrastructure Based Videography

This method refers to cameras attached to existing infrastructure such as buildings and roadway overpasses. The cameras are stationary and capture the same location.

The infrastructure-based videography approach was utilized in Washington D.C. due to the restricted airspace of the Washington D.C. region. A simplified version of this approach was utilized in the NGSIM (*Punzo et al., 2011*) dataset. For I–395 data collection, this study utilizes several cameras positioned on overpasses along I–395. For the Foggy Bottom data collection, the cameras were positioned on the roof of multiple buildings on the George Washington University Foggy Bottom campus. There is a lot of overlap between cameras to reduce blind spots. For the remaining blind spots trajectory extraction required interpolation. The advantage of this method is that multiple cameras can be used in coordination to capture a more holistic view of a certain study area; however, fusing and processing data from different input sources is a significant challenge.

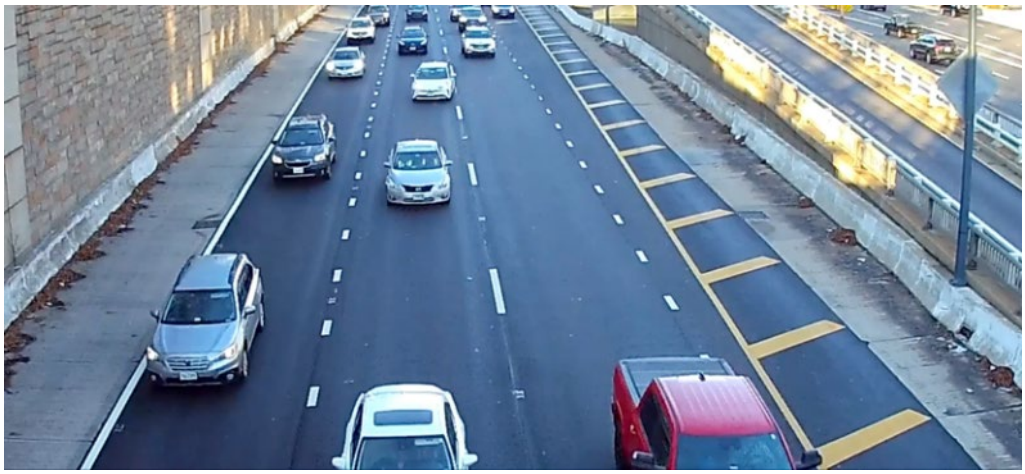

Source: FHWA

Figure 5. Photo. Example of infrastructure-based video field of view

Data Collection Sites

The team collected data at two different sites in both Chicago and Washington, D.C.:

- I–90/I–94 in Chicago, IL (see Figure 1.a and Figure 1.b).
- I–294 near Hinsdale, IL (see Figure 1.a and Figure 1.c).
- I–395 in Washington D.C. (see Figure 1.d and Figure 1.e).
- George Washington University Campus in Washington D.C. (see Figure 1.e., and Figure 1.f).

The data collection sites were selected because they offer various traffic flow conditions. They also allow for strings of ADAS-equipped vehicles and include rich vehicle interactions (discretionary and mandatory lane changes, weaving, etc.) There are also interesting road features, such as complex weaving segments, that result in interesting traffic maneuvers.

Figure 1.a and Figure 1.d show the geographic relationship between the two different Chicago and Washington D.C. data collection sites, respectively. Table 1 summarizes the key characteristics of each data collection location. The following subsections discuss key features of each of the data collection sites.

I–90/I–94 in Chicago, IL

Two datasets were collected with SAE Level 2 ADAS-equipped vehicles on I–90/I–94, using both the fixed location aerial and moving aerial videography approaches.

I–90/I–94 Fixed Location Aerial Videography Data Collection

One dataset was collected using the fixed location aerial videography approach on a short segment of I–90/I–94 focusing on the merge and diverge points (see Figure 1.b). The helicopter was hovering at an

altitude of approximately 335 m. The altitude of the helicopter enabled the camera to capture 1.3 km of highway driving (0.6 km in each direction as shown in Figure 1.b) and a major weaving section in each direction (where I–90 and I–94 diverge in the northbound direction and merge in the southbound direction). The segment has two off-ramps and two on-ramps in the northbound direction, and between 3–6 lanes in each direction. All roads have 88kph (55mph) speed limits.

Two SAE Level 2 ADAS-equipped vehicles drove through this segment between 4 p.m. and 6 p.m. They entered the northbound direction upstream of the target section and exited the target section on the right through I–94[6]. This exit location was selected for safety reasons to limit the number of automated lane-changing maneuvers (a maximum of three automated lane-changing maneuvers were needed to complete a run; conducting more lane-changing maneuvers within 0.40 km from the entry point to the highway was infeasible and unsafe). The Level 2 technology required the lane-changing requests to be initiated by the driver but were executed by the system. The first lane-changing request was initiated at the entry point (after merging to the highway from the on-ramp). The subsequent two lane-changing requests were initiated right after the previous lane-changing maneuver was completed.

The data was collected via a professional high-resolution camera at 30 frames per second (fps) at 8K resolution. The data collection altitude was set to 213 m (due to cloud ceiling on the day of the data collection) and the data was collected on a cloudy day.

I–90/I–94 Moving Aerial Videography Data Collection

Unfortunately, I–90/I–94 was highly congested on the day of data collection, which made it challenging to circulate the Level 2-equipped vehicles as originally planned. In fact, during the two hours of data collection, the two vehicles of interest were only captured 9 times in the aerially recorded videos. This motivated a shift to moving aerial videography for data collection.

The previous sections on Video Data Collection Methodologies have extensively discussed the limitations associated with data collection through Fixed Location Aerial Videography. The primary constraints include the limited length of observable trajectories and the quantity of automated vehicles' trajectories captured. Therefore, the second dataset collected on I–90/I–94 was focused on addressing the limitations of the first dataset and was collected using the moving aerial videography approach. Throughout the data collection (from 3 p.m. to 5 p.m.), the helicopter followed an SAE Level 2 ADAS-equipped vehicle in the northbound direction at an altitude of approximately 200 m above the highway. The data collection effort covered about 4.0 km in the northbound direction between N Kimball Ave. and N Wilson Ave. This section includes five off-ramps and three on-ramps on the right, and one on-ramp and one off-ramp on the left. The entire segment has 3–6 lanes (4 lanes for the majority of the 4.0 km). The vehicle entered the highway from an on-ramp on the right, drove for about 4.0 km, continued to I–90, and exited the highway using an right side off-ramp. Since the data was collected in the northbound direction only, three SAE Level 2 ADAS-equipped vehicles were used in this study to ensure a continuous movement of vehicles along the highway during the data collection period. Accordingly, once a vehicle reached the end of the segment on I–90, the helicopter came back to the starting point on I–90/I–94 and started following another vehicle. In the meantime, the first vehicle returned to the starting point to wait for another turn.

[6] Note that while I–90 offered more interesting traffic flow dynamics (I–90 towards the O'Hare Airport was significantly more congested and many drivers changed lanes at the last minute to exit the weaving section through I–90), exiting through I–90 required at least four lane-changing maneuvers.

The data was collected via a high-resolution camera at 30 fps at 8K resolution. The data collection altitude was set to 300 m and the data was collected on a cloudy day.

Note that both experiments aimed at collecting data from car-following and mandatory lane-changing maneuvers of SAE Level 2 ADAS-equipped vehicles. All the vehicles were equipped with GPS+IMU units and they were all from the same manufacturer that offer SAE Level 2 ADAS-equipped vehicles.

I–294 near Hinsdale, IL

Two datasets were collected in October 2021 (from three data collection efforts) from SAE Level 1 and SAE Level 2 automated vehicles on I–294 (see Figure 1) using the moving aerial videography approach. The aerial-based data collection efforts were intended to capture the impact of SAE Level 1- and SAE Level 2-equipped vehicles on adjacent traffic. This highway segment is controlled and operated by the Illinois Tollway. This section offers 4.8 km of highway driving in each direction and contains a high percentage of heavy vehicles. The segment has mostly four lanes in each direction and covers a major on-ramp in the southbound direction. In addition to this major on-ramp, this section includes one off-ramp in the southbound direction, and two off-ramps and an on-ramp in the northbound direction.

Each dataset covered about two hours of data from 3 p.m. to 5 p.m. One of the two datasets focused on a string of three full-range ACC-driven vehicles. The structure of the ACC string was similar to the one utilized by Khajeh-Hosseini et al. (*2022*) in Austin, TX. During the experiment, the three vehicles manually changed lanes to reach the second lane from the left. Once they reached the second lane, all vehicles initiated their ACC with a minimum distance setting (exact headway values are not provided by the manufacturer). The vehicles stayed in this mode throughout the segment (unless a safety concern was recognized by the driver). Note that due to large following distances, the string often broke into two or three individual ACC-driven vehicles. Drivers did not try to form the string in any inorganic way (i.e., the string was formed as soon as the vehicle breaking the string changed lanes again).

As discussed previously, the Austin dataset faced a key challenge: the number of platooning runs captured during two hours of data collection was extremely limited. This was due to using only three vehicles and considerable turnaround time. Instead, this study followed the string with a helicopter throughout the data collection segment, capturing a continuous trajectory for each ACC-driven vehicle throughout the segment. Note that in addition to the aerial-based data collection, all the vehicles in the string were equipped with a GPS+IMU. The collected information formed the ground truth when extracting data from aerial videos and was utilized in the data cleaning process. The second dataset from this location focused on collecting data from SAE Level 2 automated vehicles. The two key differences between SAE Level 1 and SAE Level 2 datasets are:

- Number of vehicles: Two vehicles were used in the SAE Level 2 experiment, whereas three vehicles were utilized in the SAE Level 1 experiment.
- Lane changing: SAE Level 1 vehicles executed lane changes manually and entered an ACC string after reaching the desired lane. In contrast, the SAE Level 2 vehicles started automated lane-changing maneuver the moment they merged into the highway from the on-ramp. The lane-changing request was initiated by the driver as soon as the manual lane-changing from the on-ramp maneuver was complete.

All datasets were collected via a professional high-resolution camera at 30 fps at 8K resolution. The data collection altitude was set to 300 m to capture approximately 400 m upstream and downstream of the subject vehicles at any given time. The SAE Level 2 dataset was collected on a cloudy day, while the SAE Level 1 dataset was collected on a sunny day.

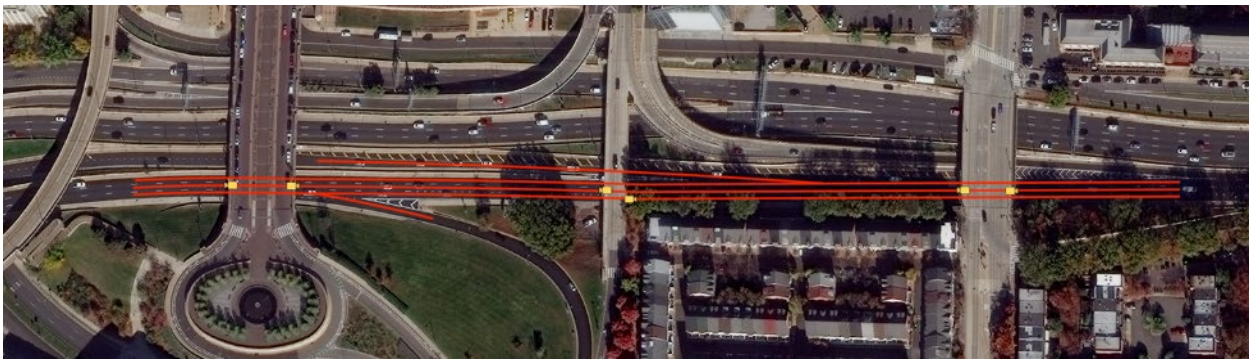

©Google Earth with additions by FHWA

Figure 6. Illustration. Camera setup for I–395.

I–395 in Washington D.C.

One dataset was collected using SAE Level 2-equipped vehicles on I–395 (see Figure 1) using the infrastructure-based videography approach. This section is 0.5 km long and covers a major weaving/mandatory lane-changing section between L'Enfant Plaza and 4th Street SW. This location has three lanes in the eastbound direction and a major on-ramp on the left side. In addition to this on-ramp, the section covers an off-ramp on the right side. The expressway includes one diverging lane at the beginning of the section on the right side and one merging lane in the middle of the section on the left side. For the purposes of data extraction, the shoulder of the merging lane is also considered a travel lane since some vehicles illegally use it as an extended on-ramp to pass other drivers. The camera setup comprised of six 4K cameras mounted on tripods, positioned on three overpasses along I-395 (see Figure 6). These cameras captured distinct segments of the highway, and their combined overlapping and non-overlapping footage resulted in a continuous trajectory for the entire section covering 0.5 km. The data collection involved multiple cameras recording at 1080P resolution and 30 fps. The data were collected on a sunny day. The two SAE Level 2-equipped vehicles traveled through the designated section between 8:30 a.m. and 10:30 a.m.

The objective of this data collection was to evaluate the impact of SAE Level 2-equipped vehicles on adjacent vehicles and their behavior in congested areas, particularly in complex merging sections. Accordingly, the focus of the experiment was on performing automated lane-changing maneuvers throughout the segment, like the SAE Level 2 experiment on I–294.The dataset was collected during the morning rush hour, aiming to capture the congestion patterns in that specific area.

George Washington University Campus (Foggy Bottom) in Washington D.C.

Data was collected in the Foggy Bottom neighborhood of Washington, D.C. using the infrastructure-based videography approach. The data collection segment, depicted in Figure 1.f, provide an overview of the specific area where data was collected. This segment encompasses one block within the George Washington University campus, with test vehicles moving mainly in a counterclockwise direction. This

area includes four intersections as well as a single block of I Street NW, H Street NW, 22nd Street NW, and 23rd Street NW. The camera installation covered four intersections within a single block, utilizing five buildings (see Figure 7). A total of 12 4K cameras were employed, with varying degrees of overlap and non-overlapping coverage. The cameras utilized in this setup captured footage at a resolution of 1080P and 30 fps (see Figure 8 for two sample images). The data was collected during the evening rush hour on a sunny day, between 4 p.m. to 6 p.m.

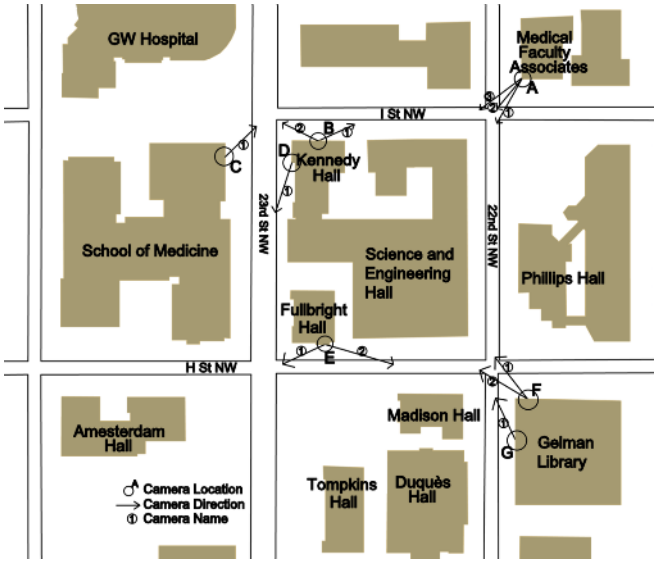

Source: FHWA

Figure 7. Map. Camera installation setup at the George Washington University Campus (Foggy Bottom) in Washington D.C.

Source: FHWA

Figure 8. Photos. Two sample images from the camera setup at the George Washington University Campus (Foggy Bottom) in Washington D.C.

The area presented a highly intricate environment, with a significant volume of pedestrians traversing the intersections. The coverage area encompassed pedestrian crossings adjacent to the Foggy Bottom Metro Station, as well as other crossings located between university buildings. The research vehicle (SAE Level 3) is circulating in a counterclockwise direction around the block, actively engaging with other vehicles, pedestrians, signal control systems, and stop signs. The intent of the experiment was to test the operation of a SAE Level 3 vehicle but unfortunately the complex environment of the area generated very frequent disengagement events during which the safety driver took control.

Data Processing Methodology

Vehicle trajectory data can be extracted from the video frames recorded from a bird's-eye view (via the helicopter) or a side view considering the perspective (via infrastructure-mounted cameras). This study utilizes the methodology proposed by Khajeh-Hosseini et al. (*2022*) as the core of their trajectory extraction method. There are, however, several key modifications to their approach that were made to meet the requirements of this study. This section discusses the details of those modifications, while briefly discussing the fundamentals of the method proposed by Khajeh-Hosseini et al. (*2022*). The trajectory extraction process consists of six steps:

1. Data preprocessing: extracting raw images and generating reference images.
2. Vehicle detection (see Figure 9b): bounding boxes around each detected object.
3. Vehicle tracking (see Figure 9c): linking the new detections to previous observations.
4. Image stabilization: converting the location of every vehicle in an image to the fixed coordinate system picked on the ground.

5. Trajectory construction: tracking changes in the location of objects over time.
6. Data cleaning: identifying and addressing problematic trajectories.

After the pre-processing step, the vehicles are detected and tracked as indicated in Figure 9 with a sample image. Then in the image stabilization step, all the images are transformed to match a reference field of view. Then the vehicles are detected in every image and tracked over the sequence of images. Finally, the vehicles' locations and trajectories are constructed by converting the image coordinates to the adopted reference coordinates on the ground. The underlying processes in each step are presented below.

It is also worth mentioning that in the case of moving aerial videography, the data was divided into 'runs'. Each run refers to the target vehicles traversing the full length of the study area either through the northbound or the southbound side of the highway. This was done to reduce the complexity of the data postprocessing steps and to eliminate the unusable data relating to the turning of the helicopter and camera until the following of the target vehicles was established again.

Preprocessing

The first step in the extraction process is preprocessing the collected videos. Data preprocessing requires two steps: (1) raw image extraction and (2) reference image generation.

Raw Image Extraction

Every video recording is converted to a sequence of images (i.e., frames) separated at a constant rate over time (e.g., 10 fps). This is done using the publicly available RED software[7], the software of the camera which was used for video recording. Vehicles will be detected and tracked in these raw images to generate the trajectory data. Accordingly, the extraction rate depends on the frequency of data in the trajectory dataset.

Reference Image Generation

The key to an effective and accurate trajectory extraction process is that in every video frame, the vehicles' location should be estimated for a fixed coordinate system and reference point on the ground. While this is a fairly straightforward process in the fixed location aerial videography approach, it was very challenging for the moving aerial videography and infrastructure-based videography approaches. This study developed a consistent methodology to convert each image into a fixed coordinate system (i.e., reference image) for each of the data collection sites. For each experiment run, a separate reference image is created even for the same data collection sites to ensure the effectiveness in subsequent transportation steps. Figure 10 shows a sample reference image developed for I–294 data.

[7] https://www.red.com/downloads

(a) Pre-Processing: Raw Image Extraction

(b) Vehicle Detection

(c) Vehicle Tracking

Source: FHWA

Figure 9. Illustration. Vehicle detection and tracking in aerial images.

Source:OpenstreetMap with FHWA additions

Figure 10. Map and corresponding reference image for one run of I–294 data.

Moving Aerial Videography Approach and Fixed Location Aerial Videography Approach

This study utilizes satellite images to create the reference images. The reference image for the Chicago datasets was created by first carefully selecting a representative set of high-resolution images from the original videos. The images cover the entire study area and have a preset overlap between them. Those images were then superimposed on a satellite image manually as well as automatically using distinct features on the images utilizing Photoshop software (*Adobe*). Manual matching is done by accurately stretching the superimposed images to match features on the satellite image, such as lane markings, on and off-ramps, buildings, and other distinct features. The satellite image is finally removed, and the matched images are then stitched together to form the final reference image. For higher accuracy, multiple reference images were created for the same location at different times to account for the changes in weather, sun, and shadows.

Infrastructure-Based Videography Approach

The process of generating the reference image was different in Washington D.C. Semi-automatic camera calibration was conducted based on Perspective-n-Point (PnP) (*Tang et al., 2018*) since cameras are fixed and angles are close to the horizontal angle concerning the ground. PnP originates from camera calibration, in which using a set of n 3D points in the world reference frame and their corresponding 2D image projections as well as the calibrated intrinsic camera parameters, determine the six degrees of freedom (DOF) pose of the camera. These DOFs include rotation (roll, pitch, and yaw) and 3D translation of the camera with respect to the fixed cartesian coordinate system in x and y predefined by the reference image. Based on this information, the homography matrix can be computed by mapping each pixel in the image to its real-world location. Finally, radial distortions can be corrected manually, resulting in the accurate location of objects in the 3D world based on 2D image data.

Object Detection

Object detection techniques in computer vision have evolved considerably in the past few years, mainly due to the recent advancements in deep neural networks (DNN). Most of the existing approaches can be used to identify and locate the vehicles in the aerial images with a little additional training. There are multiple popular convolutional neural network (CNN)-based object detectors such as R-CNN (*Girshick et al., 2014*), RetinaNet (*Lin et al., 2017*) and YOLOv5 (*Redmon et al., 2016*). The weights and parameters of a pre-trained CNN-based object detector can be fine-tuned by training on a dataset of aerial images with known vehicle annotations. The trained model can be used in the vehicle detection process to identify and locate vehicles on the aerial images. This study utilizes RetinaNet (*Lin et al., 2017*) in the fixed location aerial videography and moving aerial videography approaches, and uses YOLOv5 (*Jocher et al., 2022*) in the infrastructure-based videography approach. Note that YOLOv5 is utilized for the infrastructure-based videography approach (instead of RetinaNet) due to its exceptional performance dealing with images that are not perfect birds-eye view (due to collecting data from overpasses and buildings) as well as its ability to accurately detect multiple classes of objects with very distinct characteristics. The input of this step is the extracted raw images, and the output is annotated images (i.e., the images with bounding boxes over the detected vehicles shown in Figure 9b), and text files detailing the location of each bounding box in each image.

Object Tracking

Tracking is the process of linking the new detections to previous observations. The tracking methodology proposed by Khajeh-Hosseini et al. (*2022*) includes data association and track maintenance. The inputs of this step are the text files from the object detection step, and the outputs are tracking text files.

Data association is associating the detected vehicles in the current image frame with the vehicles identified in the previous ones. This is a critical step considering that not all the transformations are perfect and some frames must be removed from the data. Accordingly, the process should detect vehicles with some gap in distance between them in two frames.

The track maintenance initiates new trajectories (i.e., the list of bounding boxes assigned to the same ID), maintains the trajectories, and deletes them. The track maintenance initiates trajectories with unique IDs for all the vehicles detected in the first image frame. After that, for every image, all the newly detected vehicles are compared with the existing trajectories using the data association. The trajectories are updated as new detection is associated with them.

A new trajectories is constructed for any new observation that is not associated with the current trajectories. Moreover, if a trajectories is not updated in multiple frames, the track maintenance deletes that trajectories. A tracked object maintained by the track maintenance contains both the unique ID of the trajectories and the coordinates of the bounding box of its last observation.

Image Stabilization

The location of every vehicle in an image is estimated by converting its position on the image to the fixed coordinate system picked on the ground. Consequently, it is essential to find the mapping function between the image coordinates and the adopted ground coordinate. Image stabilization is the process of converting the field of view of all the images (i.e., frames) to a reference image for which the mapping function to the ground coordinate is known. Accordingly, the location of every pixel in the image is mapped to the known ground coordinates. The input to this step is the tracking text files, and the output is the same tracking text files but with their coordinates transformed to a predefined global coordinate system (*Rublee et al., 2011*).

The image stabilization is performed in three steps:

1. Detecting the key features in both reference and input images.
2. Finding the matching features between the two images.
3. Estimating transformation between them.

There exist different algorithms for good key features' detection in images such as Harris corner detector (*Derpanis, 2004*), Scale-Invariant Feature Transform (SIFT) (*Lowe, 2004*), Speeded up Robust Feature (SURF) (*Bay et al., 2006*), and Oriented FAST and Rotated BRIEF (ORB) (*Rublee et al., 2011*). Following the recommendation by Khajeh-Hosseini et al. (*2022*), this study adopts SIFT (*Lowe, 2004*) as the feature detection algorithm.

The second step is matching the features between the reference image and the input image. One naive approach is to compare every feature in the reference image with every feature in the input image to find the best matching pairs; however, this would be very time-consuming and impractical for video data at a

high frame rate. Accordingly, similar to Khajeh-Hosseini et al. (*2022*), this study utilizes the Fast Library for Approximate Nearest Neighbors (FLANN) matcher to match features between the images (*Muja et al., 2013*).

The final step is finding the perspective transformation, specifically the homography, between the reference and input images considering the best matching features. Khajeh-Hosseini et al. (*2022*) realized that despite utilizing state-of-the-art algorithms, some of the feature matchings are incorrect. They used Random Sample Consensus (RANSAC) (an algorithm to find the model parameters from a dataset with many outliers through an iterative process (Derpanis et al., 2010)) to estimate the homography transformation between two images considering the matched key features. To further improve the transformations (considering that the sizes of the reference images in this study are significantly larger than that of Khajeh-Hosseini et al. (*2022*)), this study also removed the matchings that correspond to the vehicles. This is a logical step since the vehicles in the reference image and current image cannot be at the same ground locations and feature matching using vehicles can result in undesirable transformations.

Trajectory Construction

Detection and tracking steps are based on bounding boxes that represent the vehicle's location in image coordinates (i.e., row and column of pixels). These coordinates need to be converted to a fixed ground coordinate system (e.g., meters or feet) for trajectory extraction. Every pixel is identified by its row and column number in the image map. The pixel coordinate can be transformed into a cartesian coordinate system by taking axes parallel to the columns and rows of the image map and knowing the pixel size on the ground. Note that the pixel size on the ground depends on the flight elevation and image quality and is the key to the mapping function between the two coordinate systems. Khajeh-Hosseini et al. (*2022*) took the front bumper to indicate the vehicles' location, and the trajectory of the vehicle is the list of its location over space and time. Moreover, based on the recommendation by Khajeh-Hosseini et al. (*2022*), a Kalman filter based smoothing was applied to reduce the noise in the vehicle's location estimates (caused by noisy bounding boxes from image stabilization, errors in image transformations, and issues associated with vehicle detection and tracking processes). The input of this step is the transformed tracking text files, and the output is text files including the details of each vehicle's trajectory.

The trajectory extraction process is slightly more complex for the infrastructure-based videography approach. There are multiple cameras (some overlapping) and detection and tracking results should be combined to generate a complete trajectory. The process shares certain similarities with the aerial-based methods (collecting video data for each camera and developing trajectories for each segment). However, semi-automatic camera calibration is conducted based on PnP (*Tang et al., 2018*) since cameras are fixed and angles are close to the horizontal angle to the ground. A homography matrix can be computed, and radial distortions can be corrected manually. In this manner, vehicular trajectory data is developed using a trajectory data procedure over the six continuous segments. However, due to multiple cameras recording individually, the identity of the subject vehicle may vary in this case since the vehicle would switch from segment to segment during the recording. Thus, the development of an extended trajectory for such cases is considered a major research gap, especially under conditions of heterogeneous traffic. If a vehicle leaves the first segment, it will be detected in the second segment with minor differences in the time stamp and have information such as longitudinal position, lateral position, and matching vehicle category, but it might have been tracked with a different ID. Thus, manual ID-matching is used based on the sequence of vehicles passing in each lane from each segment to the next segment (covered by another camera).

In addition, to ensure continuity over the entire road space (discontinuity in trajectory data happens when vehicles go under overpasses and they cannot be observed for several seconds), missing data points are predicted using linear interpolation, which is used to predict the longitudinal and lateral position of each vehicle as a function of time. The Kalman filter based smoothing is then applied to eliminate the noise in trajectory datasets. This stitching process is used to develop extended trajectory data throughout the entire segment.

Data Cleaning

Although all the previous steps are automated (except for ID matching before and after overpasses) and result in fairly accurate trajectories, there are still some rare cases that cannot be captured and tagged via the presented process (including parked vehicles on the side of a highway, vehicles on the overpasses, etc.). Accordingly, every single trajectory in the dataset is inspected and verified. This process was conducted through the following steps:

1. Problematic trajectories are identified using several filters such as trajectories that have zero speed values for long durations, trajectory lengths below a minimum threshold, and trajectories that traverse the opposite direction of traffic flow.
2. The IDs of those trajectories are then noted and checked against the original image frames with plotted tracking IDs to verify that they are indeed problematic.
3. Finally, the erroneous trajectories and corresponding IDs are removed from the trajectory datasets if they do not impact the integrity of the datasets.

Data Validation

During the data collection, vehicle location is recorded using a phone GPS. Although the locations are not very accurate, they generally have a systematic error. Accordingly, the extracted trajectories from the data should follow the same trend and pattern as the GPS location. Unfortunately, since the camera feed collected does not offer any time stamp (e.g., helicopter GPS and the camera system were not connected directly and were operating on two different systems), it is not possible to create a one-on-one match between the vehicle GPS data and extracted trajectories. However, since the same trend is expected in the extracted data as in the vehicle GPS data, the collected GPS data is utilized to tune the Kalman filter parameters to ensure trajectory data remains close to the observed data after smoothing.

In addition to utilizing the GPS data for validation and trajectory extraction, this study conducts an extensive evaluation of the outcome of each step to ensure the accuracy of the extracted trajectories. Errors in detection, tracking, transformation, etc. are all evaluated and mitigated through algorithm re-training, smoothing, and manual editing. This step is necessary to ensure the validity of the extracted data.

Chapter 3. Data Overview

This chapter provides an overview of the collected datasets, along with data visualizations and quality control examples. The following tables summarize the three main data collection effort results in terms of duration and number of vehicle trajectories extracted. The data collection efforts took place on the following days:

1. I-294 L1: October 18, 2021 between 3pm and 5pm
2. I-294 L2: October 5, 2021 between 3pm and 5pm
3. I-90/94 Stationary: October 6, 2021 between 4pm and 6pm
4. I-90/94 Moving: May 27, 2022 between 3pm and 5pm
5. I-395: January 26, 2023 between 8:30am and 10:30am
6. Foggy Bottom: May 4, 2023 between 2pm and 5pm

Table 2. Aerial Videography efforts in Chicago, IL

Datasets	Duration (hr)	Date	Time	Distance traveled (m)		Trajectory counting	
				Total	AVs	Small vehicles	Large vehicles
I-294 L1	2	18-Oct-21	3 p.m. - 5 p.m.	1888767	124523	960	205
I-294 L2	2	5-Oct-21	3 p.m. - 5 p.m.	378735	10348	592	120
I-90/94 Stationary	2	6-Oct-21	4 p.m. - 6 p.m.	2932819	2794	9693	325
I-90/94 Moving	2	27-May-22	3 p.m. - 5 p.m.	992175	7324	2232	290

Table 3. Infrastrucutre Videography Effort in Wshington D.C.

Dataset	Duration (hr)	Date	Time	Distance traveled (m)		Trajectory counting		
				Total	Avs	Passenger Cars	Trucks	Buses
I-395	2	26-Jan-23	8:30 am - 10:30am	4329066	9568	10361	275	8

Table 4. George Washington University Arterial Videography

Dataset	Duration (hr)	Distance traveled (m)		Trajectory counting					
		Total	AVs	Persons	Bicycles	Scooters	Passenger cars	Motocycles	Buses
Foggy Bottom	3	636629	7539	15307	214	43	4547	28	220

The presented approach was applied to the collected data in all four locations. Figure 11a–d illustrates sample trajectory data from I–294 near Hinsdale, IL in time-space diagrams for each of the four lanes. The colors are simply used to differentiate between vehicles. In this sample, there are significant shockwaves and congestion occurring in most of the lanes. However, the figure indicates that the level of congestion decreases significantly from Lane 1 to Lane 4, as the slopes of the trajectories become steeper, indicating higher speeds. A significant portion of vehicles in Lanes 1 and 2 are heavy vehicles, resulting in slower-moving traffic (comparing the slopes of lines in the traffic).

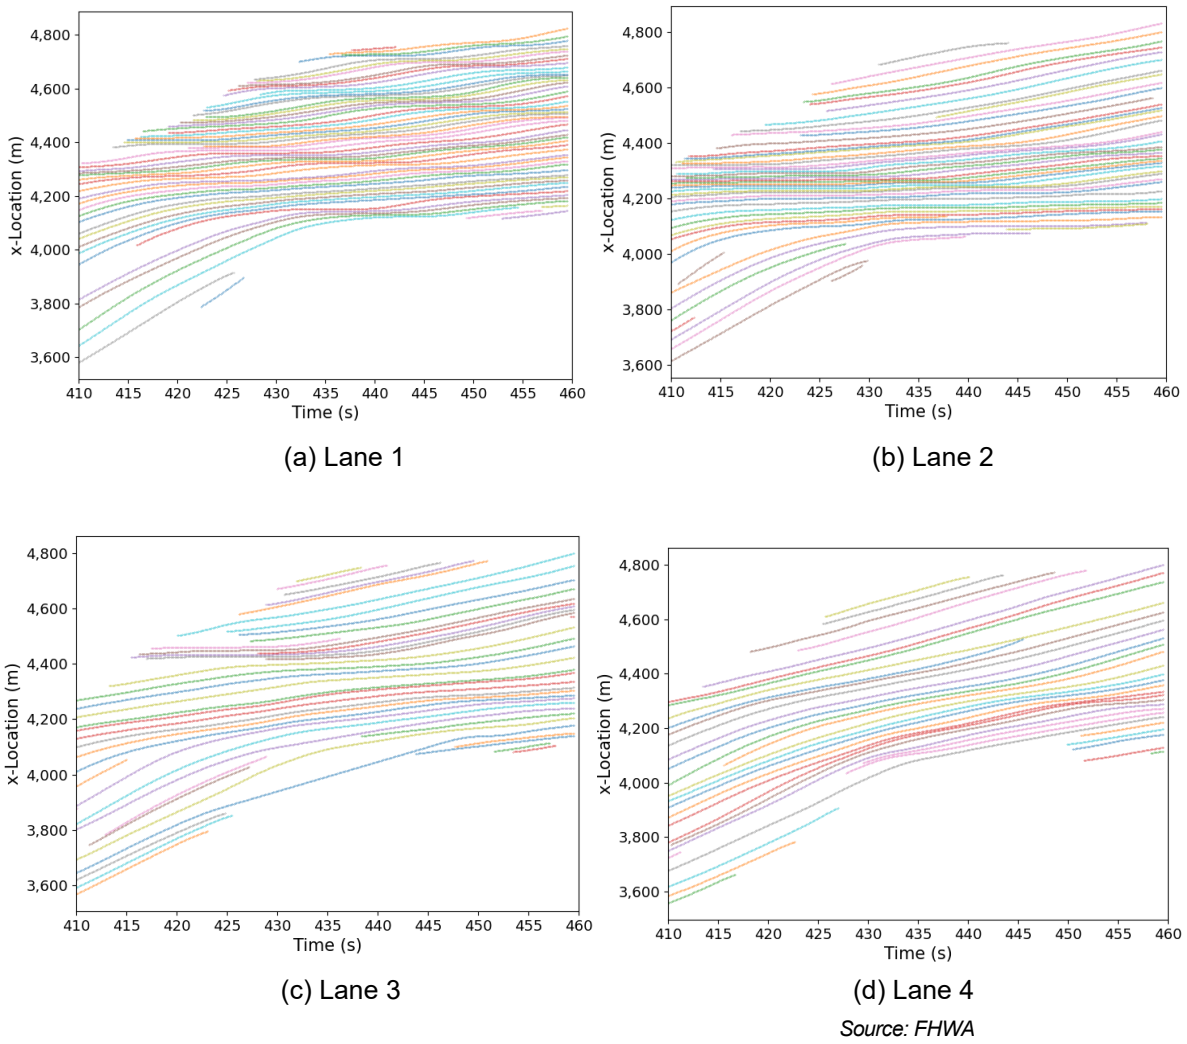

(a) Lane 1

(b) Lane 2

(c) Lane 3

(d) Lane 4

Source: FHWA

Figure 11. Graph. Sample trajectory data from a single run on I–294 near Hinsdale, IL.

Figure 12a and b presents the acceleration and speed profiles of the leader (i.e., first vehicle), Figure 12c and d presents the acceleration and speed profiles of the follower (i.e., third vehicle) in the ACC string as

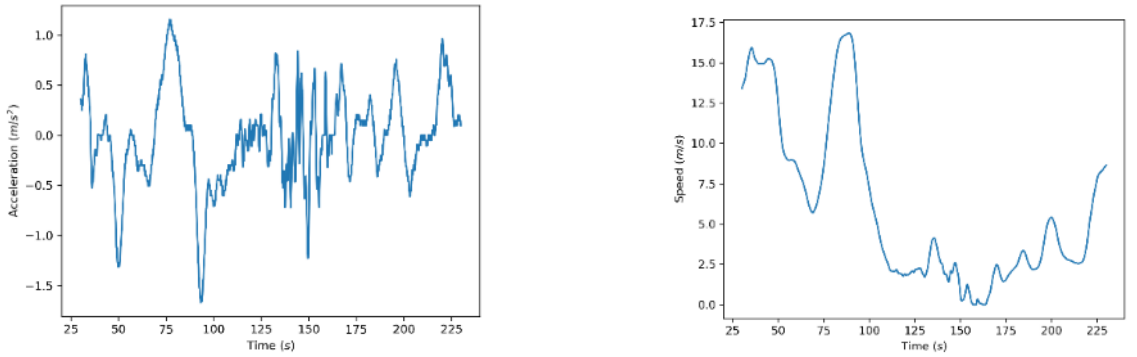

(a) Acceleration Profile: First ACC-Driven Vehicle

(b) Speed Profile First ACC-Driven Vehicle

(c) Acceleration Profile: Third ACC-Driven Vehicle

(d) Speed Profile: Third ACC-Driven Vehicle

(e) Acceleration Profile: String Followers

(f) Speed Profile: String Followers

Source: FHWA

Figure 12. Graph. Acceleration and Speed Profiles of ACC-Driven Vehicles in the String and the String Followers.

well as the string's immediate followers (Figure 12e and f). The patterns in the acceleration and speed profiles are consistent across the string members. The string shows signs of string stability as disturbances in the first vehicles lose their magnitude as they propagate upstream toward the third ACC-

driven vehicle. Interestingly, similar to the ACC string, the immediate follower follows the driving patternsof the last ACC-driven vehicle to a great extent. This is especially true in the speed profile, however, in the acceleration profile, there is a slight difference despite the similarity, which could indicate the better capability of the ACC-driven vehicles to follow in the pattern of their leaders compared to a nonautomated vehicle.

The string follower changed frequently (followers are labeled f1, f2, and f3, where different colors and line formats show different followers in Figure 12.e and Figure 12.f). This was mainly because the ACC string was traveling slower than the rest of the traffic (i.e., traveling around the speed limit in a free-flow traffic regime). Note that the acceleration and speed profiles are smooth within an acceptable level of fluctuations (between -1.5 and 1.0 m/s^2). A similar analysis has been conducted for the entire dataset (ADAS-equipped vehicles and non-automated vehicles) to ensure the quality of the trajectory data.

Figure 13 shows the overall speed profile of the vehicles in the helicopter's view (all lanes) as the ADAS-equipped vehicles (and the helicopter) move along the highway for an uncongested run (Figure 13.a) and a congeste (Figure 13.b). Two patterns are visible in Figure 13.a:

- The speed of the vehicle moving forward is fairly constant (free flow traffic regime). This is mainly due to the fact that the vehicles are moving in a lane with free-flow traffic regime.

- There are, however, multiple smaller shockwaves that can be seen (color changes in the figure, also circled in the figure). These smaller-scale changes in average speed are mainly due to the mismatch in speeds between the two rightmost lanes (slow-moving lanes) and the two leftmost lanes (free-flow moving lanes).

On average, between 400 and 500 vehicle trajectories are captured in each run with about 15–20 percent of the trajectories belonging to heavy vehicles (depending on the run and time of the data collection).

Many similar patterns can be observed in Figure 13.b. The key difference, however, is that major shockwaves are observed, and ADAS-equipped vehicle slowdown is also obvious (reaching the speed of close to zero around 1,500 m from the start of the segment). Moreover, the number of vehicles captured is 3 to 4 times more compared with Figure 13.a. Regardless of the congestion level, the unique structure of the data collection roadway segment results in almost all the heavy vehicle trajectories located in the two rightmost lanes and no heavy vehicle trajectories in the leftmost lane. Accordingly, the dataset presents many challenging lane-changing maneuvers involving passenger cars moving into much slower lanes with a high percentage of heavy vehicles.

Figure 14 presents sample trajectory data from I–90/ I–94 in Chicago, IL collected using the fixed location aerial videography approach. Because the helicopter is stationary (and not moving with traffic), the video captures a significantly higher number of vehicles (compared with the moving aerial videography approach) due to many vehicles entering and leaving the segment in a short time window. The number of heavy vehicles, however, is significantly lower compared with I–294 data.[8] Figure 14 shows how the level of congestion decreases as we move from three leftmost lanes to three rightmost lanes. Stop-and-go and

[8] I–294 is considered a major truck route and heavy vehicles mostly avoid downtown Chicago congestion on I–90/ I–94 by using I–294.

slow-and-go traffic can be observed in the leftmost lanes (circlet in the figures), while the traffic moves at almost free-flow in the rightmost lanes. This is especially visible in lane 10 (leftmost) where the slope of the trajectories almost becomes flat at one point. This dynamic creates interesting scenarios as some vehicles use the rightmost lanes and move into the leftmost lanes at the last-minute forcing complex lane-changing maneuvers. Such behavior—despite being a common behavior—has not been captured in any other trajectory dataset that we know of). Note that in this figure, the SAE Level 2 ADAS-equipped vehicles are in the second and third lane from the right and are mostly traveling at free-flow speed. They, however, get involved with the complex lane-changing maneuvers that are captured in the dataset collected through the moving aerial videography approach. Figure 15 illustrates the overall speed profile for a single run of SAE Level 2 ADAS-equipped vehicles through the segment.

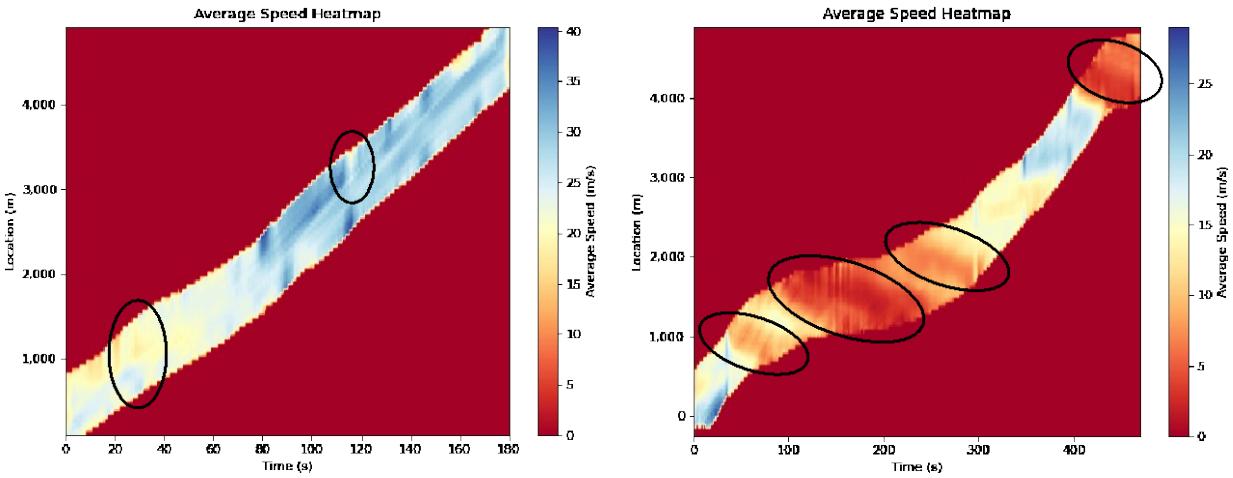

(a) an uncongested sample run (b) a congested sample run

Note: The helicopter moves with the ACC-driven vehicles in both runs.

Source: FHWA

Figure 13. Heat map. Average Speed on I–294 near Hinsdale, IL.

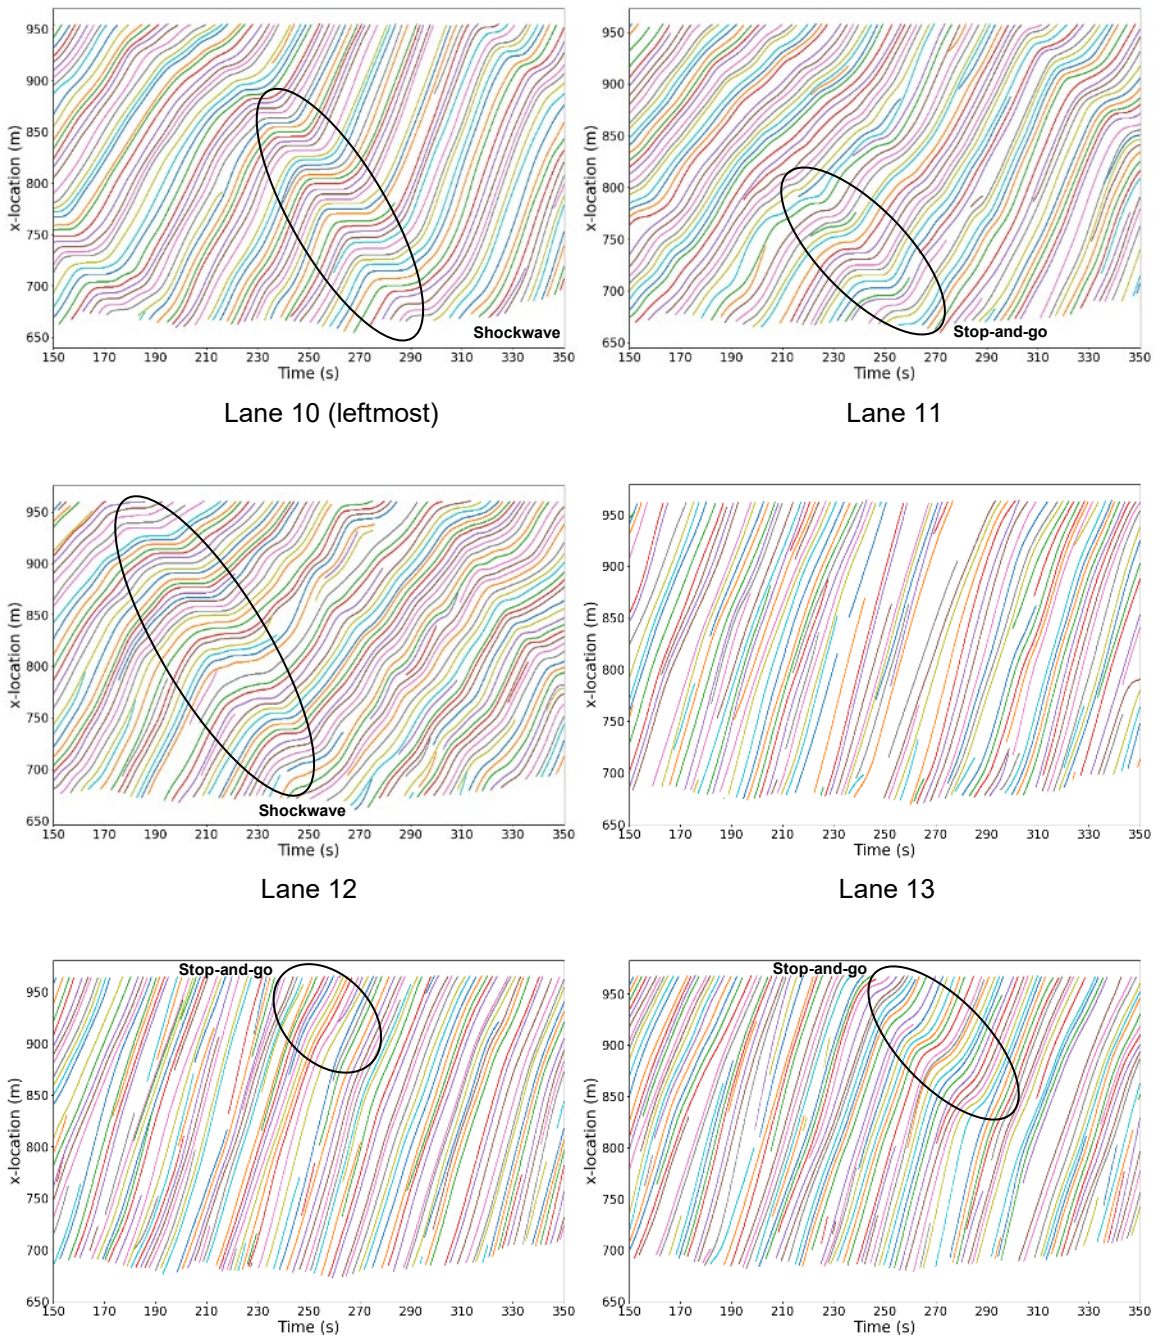

Source: FHWA

Figure 14. Time space diagrams. Sample trajectory data on I–90/I–94 in Chicago, IL. Shockwaves are indicated with black lines.

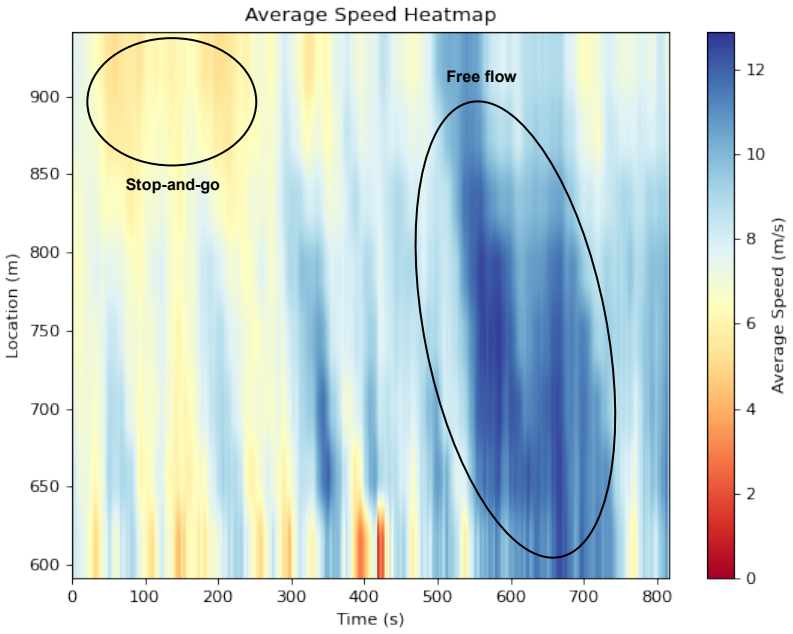

Source: FHWA

Figure 15. Heat map. Average Speed on I–90/I–94 in Chicago, IL for a sample run.

Figure 16 presents sample trajectory data from I–395 in Washington D.C. This section of I–395 is significantly congested and offers interesting dynamics between an on-ramp and the main segment lanes. There is a significant amount of lane changing occurring on the highway, either to access the off-ramp or to avoid congestion near the on-ramp. The segment also contains a significant percentage of heavy vehicles, ranging between 4–8 percent, (the largest percentage of heavy vehicles were observed on the leftmost lane, contrary to I–294). Figure 16 shows the clear shockwave in Lane 1 (rightmost lane) and multiple small shockwaves in the merging lane.

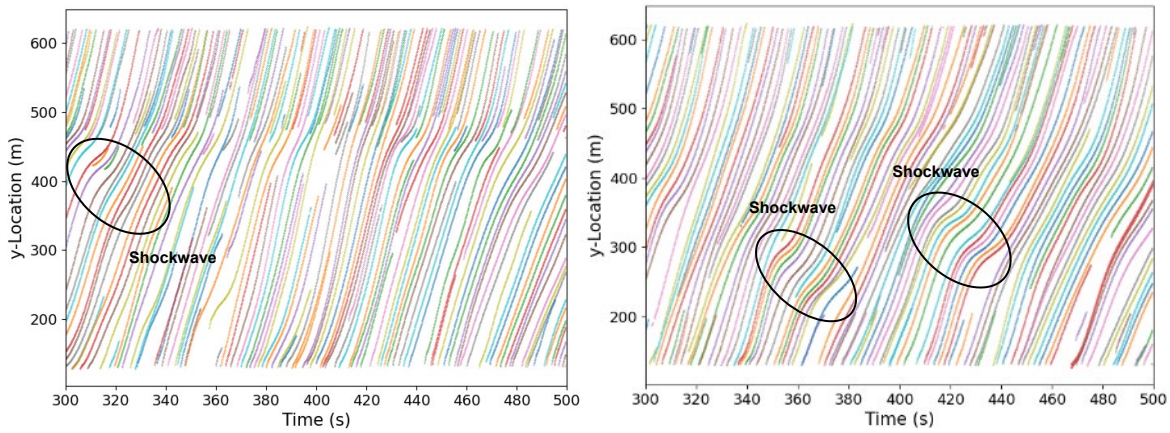

Lane 1 Lane 2 and Merging lane

Source: FHWA

Figure 16. Time space diagram. Sample trajectory data from I–395 in Washington D.C.

Figure 17 shows the average speed heatmap for the I–395 dataset/run. It is evident that there is significant congestion during rush hour, particularly in the middle section where merging lanes join the highway. As time progresses, speeds increase, and congestion gradually dissipates.

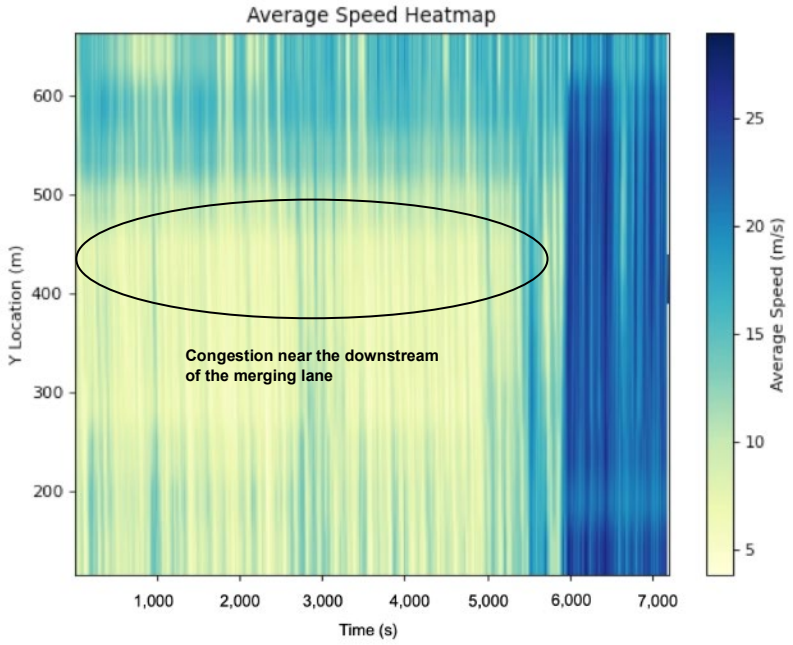

Source: FHWA

Figure 17. Heat map. Average Speed on I-395 in Washington D.C.

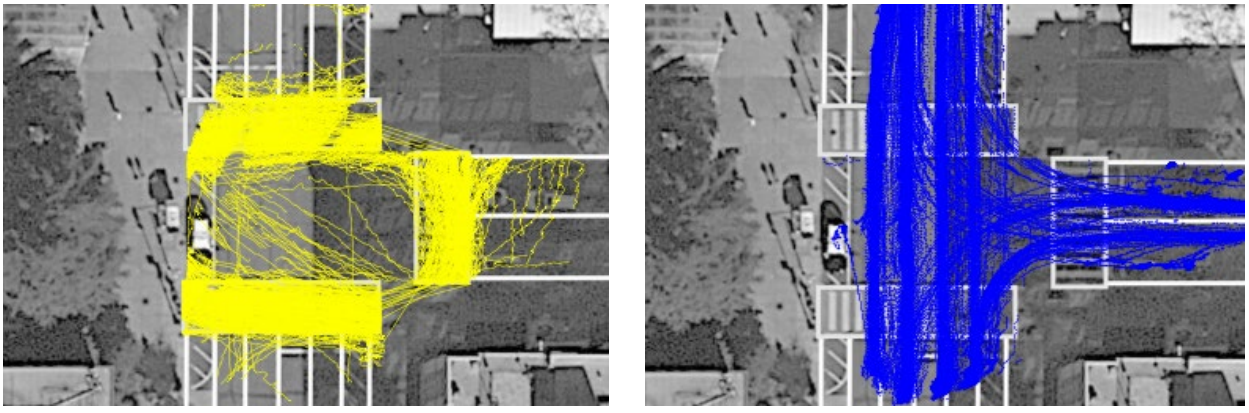

(a) Pedestrian trajectories (b) Vehicle trajectories

Source: FHWA

Figure 18. Illustration. Trajectories in George Washington University Campus in Washington D.C.

Figure 18 displays the trajectories depicted on the satellite image, offering supplementary context and spatial comprehension for the data. It illustrates the dynamics between vehicles and vulnerable road users (such as pedestrians and cyclists) within a highly populated urban environment at the Foggy Bottom campus of George Washington University, in Washington, D.C. The pedestrians are observed within the four intersections, while the vehicles are observed throughout the entire city block.

Data Dictionary

Table 5 shows the dictionary of the five datasets collected from I-294 and I-90/94, along with a detailed explanation of the lane indexing for each dataset:

Table 5: Data Dictionary for the I-294 and I-90/94 datasets

ID	A unique ID is given to each detected vehicle (i.e. numerical values)
time	Specific time at which current vehicle information is measured relative to the start of the run (in seconds)
xloc-kf	x coordinate of the vehicle relative to the global coordinates of the reference image after Kalman Filtering (in meters)
yloc-kf	y coordinate of the vehicle relative to the global coordinates of the reference image after Kalman Filtering (in meters)
lane-kf	Lane ID as defined on the reference image after Kalman Filtering (refer to the following paragraphs)
speed-kf	Measured speed after Kalman Filtering (m/s)
acceleration-kf	Measured acceleration after Kalman Filtering (m/s^2)
length-smoothed	The smoothed length of the bounding box corresponding to the detected vehicle (a proxy for vehicle length)
width-smoothed	The smoothed width of the bounding box corresponding to the detected vehicle (a proxy for vehicle width)
type-most-common	Type of the vehicle detected (e.g. small vehicle, large vehicle)
ACC	Labeled 'Yes' if the current vehicle detected is an ACC vehicle, and 'No' otherwise.

For the two datasets collected from I-294, lanes are labeled as -2, -3, -4, and -5 for the southbound and 2, 3, 4, and 5 for the northbound. Lane -5 is the leftmost lane from the perspective of the reference image, and lane 5 is the rightmost lane.

The I-294 site included three overpasses. Vehicle trajectories on these locations were interpolated to fill the gap from the loss of view of the vehicle from the helicopter. The coordinates of the overpasses for each run are as follows:

- For run-index 1: The xloc-kf ranges for the three overpasses are as follows:
  - Overpass 1: 2130 to 2170
  - Overpass 2: 3330 to 3370
  - Overpass 3: 3730 to 3770

- For run-index values 3, 7, 9, 11, 19, 21: The xloc-kf ranges for the three overpasses are:
  - Overpass 1: 1530 to 1570
  - Overpass 2: 2730 to 2770
  - Overpass 3: 3130 to 3170

- For run-index values 8, 18, 20: The xloc-kf ranges for the three overpasses are:
  - Overpass 1: 1760 to 1800
  - Overpass 2: 2160 to 2200
  - Overpass 3: 3360 to 3400

For the dataset collected with a moving helicopter on I-90/94, all vehicles are traversing northbound with the lanes labeled as 1, 2, 3, 4, 5, and 6, where lane 1 is the innermost lane, and express lanes are not included.

For the dataset collected with a stationary helicopter on I-90/94 at the diverging area, lanes are categorized based on the run index. Odd-index runs (1, 3, 5) use lane IDs 1 through 6, with lane 1 being the innermost. Even-index runs (2, 4, 6) use lane IDs 10 through 15, with lane 10 being the innermost. In both cases, express lanes are not included.

Table 6 provides an overview of the type of data collected from I-395, and the detailed lane indexing, and vehicle type notation are as follows:

- For lane-kf, -5 is the rightmost lane, -4 is the middle lane, and -3 is the leftmost lane of the highway. -2 is the merging from the left side of the highway. -1 is the shoulder of the merging lane. -6 is the diverging lane from the right side of the highway.
- For type-most-common, 1 is the Passenger Cars, 2 is the Trucks, 3 is the Buses, and 4 is the AV Vehicles.

Similarly to the I-294 site, the I-394 site also includes overpasses where the cameras were stationed. These overpasses create blind spots for which the vehicle trajectories had to be interpolated. The locations of the I-394 overpasses is the following:

- Overpass (1) yloc-kf from 200m, to 235m
- Overpass (2) yloc-kf from 380m, to 395m
- Overpass (3) yloc-kf from 555m, to 585m

Table 6: Data Dictionary for the I-395 dataset

ID	A unique ID is given to each detected vehicle (i.e. numerical values)
time	Specific time at which current vehicle information is measured relative to the start of the run (in seconds)
xloc-kf	x coordinate of the vehicle relative to the global coordinates of the reference image after Kalman Filtering (in meters)
yloc-kf	y coordinate of the vehicle relative to the global coordinates of the reference image after Kalman Filtering (in meters)
lane-kf	Lane ID as defined on the reference image after Kalman Filtering (refer to following paragraphs)
speed-kf	Measured speed after Kalman Filtering (m/s)
acceleration-kf	Measured acceleration after Kalman Filtering (m/s^2)
length-smoothed	The smoothed length of the bounding box corresponding to the detected vehicle (a proxy for vehicle length)
width-smoothed	The smoothed width of the bounding box corresponding to the detected vehicle (a proxy for vehicle width)
type-most-common	Type of the vehicle detected (refer to following paragraphs)

Table 7 provides an overview of the type of data collected from the Foggy Bottom Campus of George Washington University. Different from the datasets from the highway environment, this dataset is collected in a grid road network. Therefore, the speed and acceleration are reported separately on x and y dimensions.

The lane-kf attribute ranges from 1 to 49, denoting the crosswalks and branches of four intersections in the data collection area, and the legend in the dataset has a clear annotation of the locations of each area.

The type-most-common designates road users into 8 types as the following: 0 is persons, 1 is bicycles, 2 is scooters, 3 is passenger cars, 4 is automated vehicles, 5 is motorcycles, 6 is buses, and 7 is trucks.

Table 7: Data Dictionary for the Foggy Bottom dataset

ID	A unique ID is given to each detected vehicle (i.e. numerical values)
time	Specific time at which current vehicle information is measured relative to the start of the run (in seconds)
xloc-kf	x coordinate of the vehicle relative to the global coordinates of the reference image after Kalman Filtering (in meters)
yloc-kf	y coordinate of the vehicle relative to the global coordinates of the reference image after Kalman Filtering (in meters)
lane-kf	Lane ID as defined on the reference image after Kalman Filtering
speed-kf-x	Measured speed in x dimension after Kalman Filtering (m/s)
speed-kf-y	Measured speed in y dimension after Kalman Filtering (m/s)
acceleration-kf-x	Measured acceleration in x dimension after Kalman Filtering (m/s^2)
acceleration-kf-y	Measured acceleration in y dimension after Kalman Filtering (m/s^2)
length-smoothed	The smoothed length of the bounding box corresponding to the detected vehicle (a proxy for vehicle length)
width-smoothed	The smoothed width of the bounding box corresponding to the detected vehicle (a proxy for vehicle width)
type-most-common	Type of the vehicle detected

Chapter 4. Sample Analysis (I-294 L1)

Overview

ACC technology has been around for over two decades. It has been developed to improve driving comfort, reduce driver errors, improve road safety, as well as increase traffic flow, and reduce fuel consumption (*Xiao et al., 2010*). There has been a significant number of studies that focused on developing robust, string stable, efficient, and safe ACC systems *(Makridis et al., 2020, Ioannou et al., 1993, Yu et al., 2022*). In the majority of these studies, the following vehicle reacts to the leading vehicle(s) and adjusts its speed to maintain a safe gap between them (*Nowakowski et al., 2011*). One of the earliest comprehensive studies on ACC systems is by Swaroop et al. (*1994*), who proposed a comparison between various spacing and control policies and identified the scenarios where an ACC string can be string stable. Based on their findings, a constant time-headway policy resulted in string stable ACC strings. Note that this study also utilizes a variation of their constant time-headway policy to model ACC-based behavior. More recently, using a different approach, Zhu et al. (*2020*) extracted car-following events from the NGSIM dataset to train and test a reinforcement learning model for safe, efficient, and comfortable vehicle velocity control. Their approach aimed to develop an ACC system that behaves similarly to human drivers. Ahmed et al. (*2021*) detail the different types of car-following models and present a thorough review of models used in microsimulation tools for automated vehicles.

One of the key missing aspects in the majority of ACC modeling studies is the fact that the current ACC systems need human supervision and drivers often disengage the system for various reasons. While the mechanical properties of the control of ACC vehicles and underlying dynamics of disengagements have been investigated extensively (*Varotto et al., 2015*), there is not a comprehensive framework to model the movement of ACC-capable vehicles through these cycles of engage-disengage. This is a particularly critical task since when the ACC system is not active since the vehicle is controlled by a human driver, and its movements are governed by human decision-making logic. The data collected from ACC-driven vehicles on I–294 provides the opportunity to develop car-following models to capture the impact of the engage-disengage phenomenon on vehicles' following behavior. It is important to note that the majority of trajectory datasets do not have (or do not identify) ACC-driven vehicles. Accordingly, the main focus of this chapter is to introduce a comprehensive modeling framework to capture the behavior of ACC-capable vehicles going through cycles of engage-disengage and to illustrate the advantages of the collected dataset over existing datasets.

Fortunately, similar to the ACC models, the car-following behaviors of human drivers have been investigated extensively. The intelligent driver model (IDM) is one of the models that have been adopted to model human drivers' behavior (*Treiber et al., 2000*). Talebpour et al. (*2011*) proposed a multi-regime car-following model based on the model by Hamdar et al. (*2008*) that considers different car-following behavior for congested and uncongested traffic regimes. While there is no shortage of analytical and physics-based approaches to modeling car-following behavior, several data-driven approaches have been proposed recently. Ozkan and Ma (*2021*) used inverse reinforcement learning to model car-following behaviors of different human drivers when interacting with a mixture of connected-automated vehicles and human-driven vehicles. In another study, Kamjoo et al. (*2023*) introduced multiple data-driven car-

following models to capture car-following behavior in adverse weather conditions. Accordingly, the key challenge that this modeling effort aims to tackle is to develop a probabilistic framework to capture when drivers switch between ACC-controlled and human-controlled instances.

Model Formulation

Car-following models typically focus on modeling the acceleration of the target vehicle in response to the behavior of its leader(s). When implemented in a discrete-time simulation framework, the model takes the input values observed at the end of the previous time step and outputs an acceleration value that is assumed to be uniformly applied on the target vehicle over the duration of the current time step.

In a typical ACC-based drive, the limitations of the ACC would likely prevent the drivers from taking a completely hands-off approach in actual highway traffic. They are often required to intervene. Some ACC systems will not engage below a certain speed threshold, and some refuse to re-engage after the vehicle has stopped in a stop-and-go traffic pattern. In fact, SAE J3016 makes it clear that humans are completely responsible for vehicle decisions even if the ADAS is controlling the longitudinal and lateral behavior. Thus, a naturalistic trajectory of an ACC-activated vehicle will likely contain periods where the human driver is in full control of the vehicle. To better capture such instances, this modeling effort proposes a hybrid framework that has a high-SAE Level discrete choice component to model the probability of human intervention. The hybrid model has three components: two underlying models for human and ACC car-following, and a logit model to select the car-following model in effect at each time step. The specific subcomponents including the logit model can be replaced if other researchers are interested in adopting this framework for their needs. This study uses the Intelligent Driver Model (IDM) (*Treiber et al., 2000*) and a Constant Time Headway (CTH) policy *(Nowakowski et al., 2011*) for the underlying human and ACC models, respectively. The logit model for model selection has the following setup:

$$f = \beta_0 + \beta_1 v + \beta_2 \Delta v + \beta_3 s$$

Figure 19. Equation. Logit function.

$$p = \frac{e^f}{1 + e^f}$$

Figure 20. Equation. Logit Probability.

$$a_{hybrid} = p.a_{IDM} + (1 - p).a_{CTH}$$

Figure 21. Equation. Hybrid Acceleration Equation.

Where:

f is the logit function.

p is the logit probability.

β's are parameters to be calibrated.

a_{IDM} is the output acceleration using the IDM model.

a_{CTH} is the output acceleration using the CTH model.

a_{hybrid} is the joint acceleration given the probability p.

Table 8. Recommended or Initial Values of Model Parameters

Group	Symbol	Physical Meaning	Recommended/Initial Value (Treiber et al., 2000)
IDM	T	Reaction time	1.6 m
	a	Maximum acceleration	0.73 m/s^2
	b	Maximum deceleration	1.67 m/s^2
	δ	Exponent	4
	s_0	Minimum space headway	2 m
	v_0	Targeted speed	33.3 m/s
CTH	h_w	Targeted time headway	1.0 s
	α	Coefficient	0.1
Logit	$\beta_{0,1,\ldots}$	Coefficients	0.001

Note that more state variables, such as the density, average speed, and variance of the speed of the traffic flow, can be added to the logit model, and the number of β's will increase accordingly. However, the chosen parameters were suitable and sufficient for the requirements of this study. The symbols and recommended or initial values of all parameters to be estimated under this section are summarized in Table 5. Note that the original CTH paper (*Swaroop et al., 2001*) model did not include recommended values, accordingly, the values listed in the table are initial guesses.

Intelligent Driver Model (IDM)

The IDM is one of the most popular models for human drivers' car-following behavior. It outputs an acceleration of the target vehicle based on the input traffic states. To calculate the acceleration, the desired space headway is first calculated based on:

$$s^* = \max\left(0, s_0 + vT + \frac{v\Delta v}{2\sqrt{ab}}\right)$$

Figure 22. Equation. Space headway between vehicles

Then, the acceleration is calculated by:

$$a_{IDM} = \mathrm{a}\left(1 - \left(\frac{v}{v_0}\right)^{\delta} - \left(\frac{s^*(v, \Delta v)}{s}\right)^2\right)$$

Figure 23. Equation. IDM model acceleration

Where:

v is the speed of the target vehicle.

v_0 is the target speed

Δv is the speed difference between the target vehicle and the leader.

s is the space headway between the target vehicle and its leader.

s^* is the desired space headway between the target vehicle and its leader.

s_0 is the minimum space headway

T is the reaction time

a and b are coefficients

α_{IDM} is the output acceleration using the IDM model

δ is the exponent

Constant time headway (CTH) model

CTH models were proposed to capture an ACC vehicle's car-following behaviors. They assume that an ACC controller is designed to maintain a preset time headway, which in turn is proportional to the difference in the speeds of the target and lead vehicles. While the internal ACC algorithms and low-SAE Level of automation controllers vary among the manufacturers and are unknown to the research team, a constant headway model can capture the general logic behind most ACC systems and the maneuvers as a result of low-SAE Level of automation controlling. The research team uses the CTH policy formulation given in Swaroop and Rajagopal (*Swaroop et al., 2001*):

$$a_{CTH} = \frac{1}{h_w}\Delta v - \alpha\left(v - \frac{\Delta v}{h_w}\right)$$

Figure 24. Equation. CTH model acceleration

where h_w and α are parameters to be calibrated with typical values presented in Table 5.

Calibration Process

Any car-following model should be calibrated to ensure its ability to replicate real-world observations. Car-following models, whether of human drivers or robotic drivers, contain tunable parameters that can be tweaked depending on various factors, including local conditions, type of roadway, driver behavior, etc. Existing car-following models can be calibrated to fit existing empirical datasets. This can bring insight into the ability of the original car-following model to capture traffic flow dynamics, as well as its capability of modeling unseen events (to be utilized in simulation tools and studies). Various calibration approaches have been developed in the literature. For instance, Chen et al. (*2010*) used the trajectory datasets provided by the NGSIM program to calibrate the IDM and Microscopic Traffic Simulator (MITSIM) car-following models using a genetic algorithm. They achieved tracking gap errors for both models that are generally below 30 percent. Vasconcelos et al. (*2014*) presented a procedure to collect trajectory data and calibrate the Gipps car-following model (*Gipps, 1981*) using two estimation methods. Other model calibration efforts are abundant in the literature where IDM, ACC, as well as other car-following models, are calibrated using empirical and simulated data *(Kesting et al., 2008, Cattin et al., 2019, Rakha et al., 2007*). The remainder of this section presents the calibration approaches developed for the proposed hybrid car-following model.

Along with the new framework, this study proposes two calibration methods, each with its own rationale. The first one takes a modal choice perspective and explicitly recognizes mode choice (ACC vs. human-driven) as a separate step from acceleration. The second formulates the problem as an optimization problem and optimizes all the parameters simultaneously using the genetic algorithm.

Two-loop iterative method

The three components of the hybrid model are calibrated in an iterative two-stage process. The process has an outer and an inner loop. Each iteration of the outer loop starts with two inner-loop operations, in which the IDM and CTH models are calibrated with the genetic algorithm. Their outputs are then used to fit the logit model and to prepare data for the next iteration of the outer loop. The following is a description of the two-loop procedure, as illustrated in Figure 25.

- In the initial iteration of the outer loop, all data points (v, s, and Δv of the target vehicle of all time steps) are used to calibrate the IDM and CTH models.
- The IDM and CTH are calibrated separately with the genetic algorithm in two inner loops.
- Errors in the IDM and CTH models are calculated with the optimal parameters from the previous inner loops. Error values are calculated separately for each data point.
- The errors of the two driving modes are compared. For each data point, the mode with the smaller error is considered the preferred car-following mode of that time step.
- v, s, and Δv data (X) and the modal preference (y) of each time step are fed into a Scikit-learn Logistic Regression module (*Pedregosa et al., 2011*) to fit the logit model.
- The fitted logit model predicts the more likely modal choices (y^{pred}) of all data points.
- Separate subsets of data are created to further calibrate the IDM and CTH models. For each data point, if $y^{pred} > 0.5$ (indicating a preference for IDM), it is assigned to the subset for IDM, otherwise to the set for CTH. An exploration factor is added to recognize the fact that it may not be possible to have

a clear-cut determination of CTH or IDM for some data points. For a 20 percent possibility, a data point is assigned to a randomly chosen model, and the predicted preference is ignored.

- For the inner loops in the next iteration, half of the initial population is inherited from the best chromosomes of the previous iteration, while the other half is randomly initialized.
- The procedure goes back to Step 2. Another iteration of the outer loop begins. The outer loop repeats until the maximum number of iterations is reached.

This algorithm assumes that the dynamic choice between the IDM and the CTH model through the logit discrete choice model will improve overall performance in each iteration of the outer loop. The logit model operates as a classifier of the data. It is assumed that the human driver was more likely to have been in control of the target vehicle at a given time step if the error of the IDM was smaller than that of the CTH model, and vice versa. This information should fit a logit model that would recognize the general conditions of traffic under which the human or ACC should be in control. Note that the predicted choice of the logit model may be different from the actual choice due to relatively small errors/noise in trajectory measurements. This is expected, considering that the data contains noise and the IDM and CTH models are not perfect. After the initial iteration, the IDM and CTH models are calibrated with their preferred data points. The purpose is to improve their performance for the specific traffic conditions under which they were assumed to be in effect and to have the improved car-following models to reinforce the classification capability of the logit model. Finally, randomization is introduced in Steps 7 and 8 of the process to mitigate the bias, if any, of insufficiently trained models in the early iterations from perpetuating into later iterations.

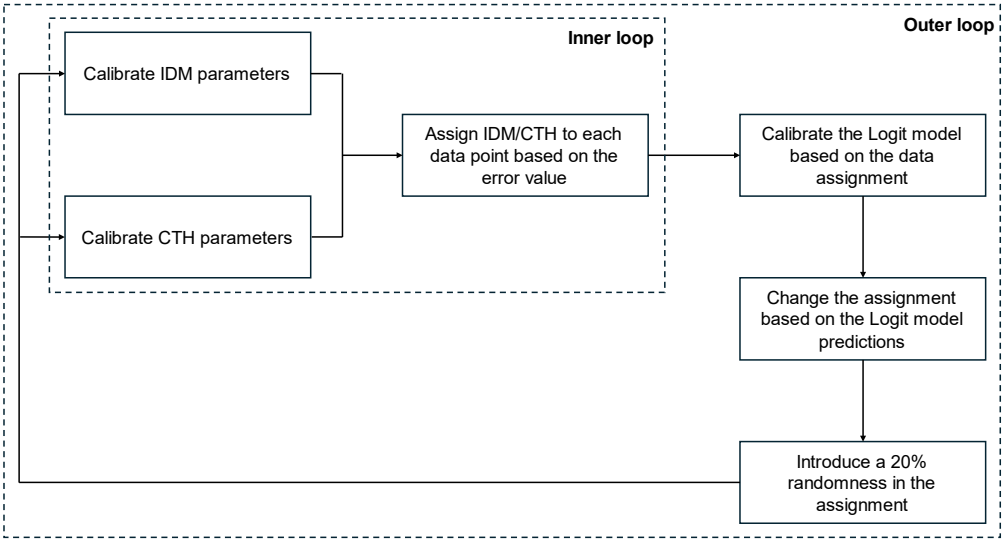

Source: FHWA

Figure 25. Illustration. Flowchart of the two-loop calibration method.

Inner loop: Genetic algorithm-based calibration

The genetic algorithm is used to calibrate the parameters of the IDM and CTH models in the inner loop. The goal is to find a set of parameters that minimize an error term, which is defined as the difference between the outputs of the car following the model and the observed values. The genetic algorithm is a heuristic iterative method that is ideal for solving this non-linear, intractable optimization problem

(*Sivanandam et al., 2008*). It starts from a search space of randomly generated initial sets of parameters, which are called “chromosomes” in a “population”. In each iteration, the fitness of each chromosome of the current population is evaluated by simulation, and a certain number of best fit chromosomes, called “parents”, are recombined in the “crossover” step to breed a new generation of population. Random changes are then introduced in the “mutation” step before the new generation goes into the next iteration. The algorithm stops when the fitness hits a predetermined level or when the maximum number of iterations has been reached. The best-fit chromosome of the final iteration is regarded to be the optimal set of parameters.

This study implemented the genetic algorithm as described in Hamdar et al. (*2015*). The size of the population is fixed at 100. The initial ranges of the IDM parameters are set around the recommended values in Kesting and Treiber (*2008*) (no recommended values were given for the CTH model.) In each iteration, ten parents participate in the crossover, breeding 90 new chromosomes. One parameter in each chromosome is randomly chosen for mutation. Its value will have a random change not exceeding 10 percent.

The error is based on the absolute difference between the predicted, $x_t^{predicted}$, and observed, $x_t^{observed}$, values of positions. In each iteration of the inner loop, the target vehicle’s movement is simulated with each chromosome (i.e., set of parameters) under consideration. The simulation is reset to the observed position and speed of the target vehicle at each time step. This would prevent the errors from accumulating across the time steps and distorting the fitting of the logit model. For this study, the "fitness" in the genetic algorithm is simply defined as the negative average of the errors.

One-loop joint estimation

As an alternative to the two-loop procedure described above, this procedure estimates all 14 parameters jointly in the genetic algorithm loop. It does not explicitly fit the logit model either. For this method, a different approach based on simulation was utilized since there is no need for time step wise error to fit the logit model. The process starts with the initial position and speed of the vehicle and predicts the movement under the hybrid model. For each subsequent time step, the predicted position and speed of the target vehicle and the actual values of the leader as inputs to predict new movement (headway and speed difference are re-calculated accordingly). The simulation is run through the entire length of the available data points without resetting the predicted position and speed, except when there is a sudden drop in the vehicle’s headway between the leading vehicle. The error is the average difference between the cumulative predicted position and the observed location of all time steps. It is believed that this provides a more rigorous evaluation of the model’s performance, which is the ability to reconstruct a long trajectory on its own.

The genetic algorithm procedure used here is nearly identical to the inner loop described above. The main difference is in the crossover and mutation steps. Instead of mutating one parameter at each iteration, three parameters, each chosen from the IDM, CTH, and logit components, are randomly chosen to mutate. Crossover also happens within each component. Half of the parameters of a group come from one parent, while the other half comes from the other.

Experiments

The dataset used for this study contains 5,650 data points, consisting of three runs (which means a north-south or south-north trip through the highway segment), representing about 10.5 minutes of traffic. Two ADAS-equipped vehicles that were driven by the research team were identified from the dataset as the target vehicles, for which the proposed car-following model was individually calibrated. For the two-loop estimation, the outer loop iterated for a total of 10 times, and the inner loop (i.e., simulation, genetic algorithm) ran for 1,000 iterations each time. In total, they took approximately 300 seconds. For the one-loop joint estimation, the process was terminated after 150 iterations, which took approximately 800 seconds. The latter was significantly slower because the sequential setup of the simulation does not allow for vectorized computation. All calibrations were performed in a Python 3.8 (*Van Rossum and Drake, 2009*) environment on a personal computer. In this section, we illustrate some of the results and key findings of the calibration process. Note that, while the model was calibrated with the entire set of 5,650 data points, in this section we only showcase graphs plotted with the data points from Run 18 due to space constraints. It is the longest run and contains the most diverse set of traffic conditions.

Calibration Results

The calibration results are summarized in Table 6. "Leader" means the first car in the ACC string, which followed the general flow of traffic; "Follower" means the second car, which followed the Leader. All sets of parameters are withing physical limits, except for some extreme values, for example, the desired deceleration (b) in the Joint Leader Follower set. Such a value, 22 m/s^2, is unrealistically large for a normal passenger vehicle. For comparison, the AASHTO recommended value is 3.4 m/s^2 (*AASHTO, 2011*). Extreme values, however, can be avoided by setting limits on the model parameters (we did not impose any limitation on the parameter values in this study because we were able to conduct successful simulation experiments even with some irregular parameter values).

Table 9: Calibration Results under Both Methods

Group	Symbol	Two-loop Leader	Follower	Joint Estimation Leader (One-Loop)	Follower
IDM	T (s)	1.303	0.055	0.777	0.580
	a (m/s^2)	0.311	1.094	0.691	1.563
	b (m/s^2)	1.143	5.558	5.005	22.956
	δ	8.840	5.429	3.340	6.038
	s_0 (m)	0.000	13.688	3.766	10.062
	v_0 (m/s)	28.080	27.870	25.894	58.053
CTH	h_w (s)	12.785	1.349	0.580	4.150
	α	0.008	0.076	0.045	0.016
Logit	β_0 offset	-0.242	-0.636	0.001	-0.002
	β_1 for v	0.012	0.015	-0.004	-0.002
	β_2 for s	-0.000	0.008	-0.005	-0.003
	β_3 for Δv	0.021	0.038	0.009	0.015

Analysis and Discussion

To capture the actual performance of the model, we look at the predicted changes in speed. This is done because examining the actual position or speed values, which are large in magnitude compared to the error values, does not always give much insight into the performance of the model. Figure 26 illustrates the performance of the models using the two-loop iterative procedure explained under the model formulation and calibration section. On the other hand, Figure 27 showcases the performance of the models calibrated under the joint estimation procedure. "Speed difference" in the figures means the difference between the speed at the current time step and that of the previous time step.

Figure 26.a and Figure 27.a show the difference in speed values between every two consecutive time steps for the CTH model versus the ground truth (i.e., actual) values. It can be seen that the model is generally capable of following the general trend of the actual behavior of the driver, but it does go a little astray at times. By definition of the Constant Time Headway ACC model, it would be impossible to match the actual oscillations that occur in natural driving at the current complexity of the CTH model. On the other hand, Figure 26.b and Figure 27.b show the performance of the IDM model. Similarly, it can be inferred that the IDM model is able to capture some of the general trends of the data but at certain locations fails to do so. Upon further examination of the figures, we can see that, at certain times, the IDM model outperforms the CTH model and follows the actual data more closely. This can be seen between 120 and 200 seconds, where the IDM model is better at capturing the ground truth. On the other hand, we can also see that the CTH model is better at capturing the actual trends of driving at other times. Both models output extreme values (i.e., spikes in the figures) at certain times. The predicted speed difference under the joint estimation method is particularly volatile for the Leader. Consequently, we make use of the hybrid model that is able to capture the best of both models and mitigate the extreme values, and produce results that are more closely related to the actual values. The results of the hybrid model can be seen in Figure 26.c and Figure 27.c. The hybrid model minimizes the errors of both the CTH and IDM models and, as illustrated, is able to best represent the actual data.

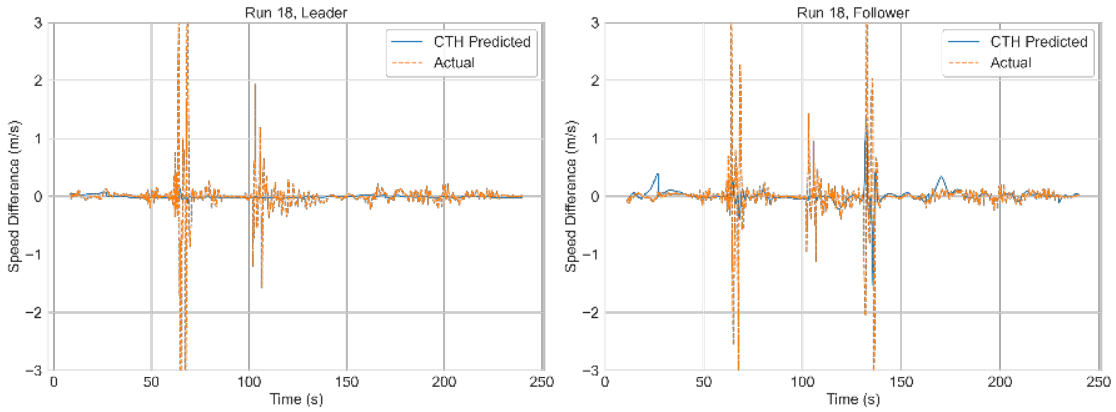

(a) CTH Model Speed Difference (Left: Leader, Right: Follower)

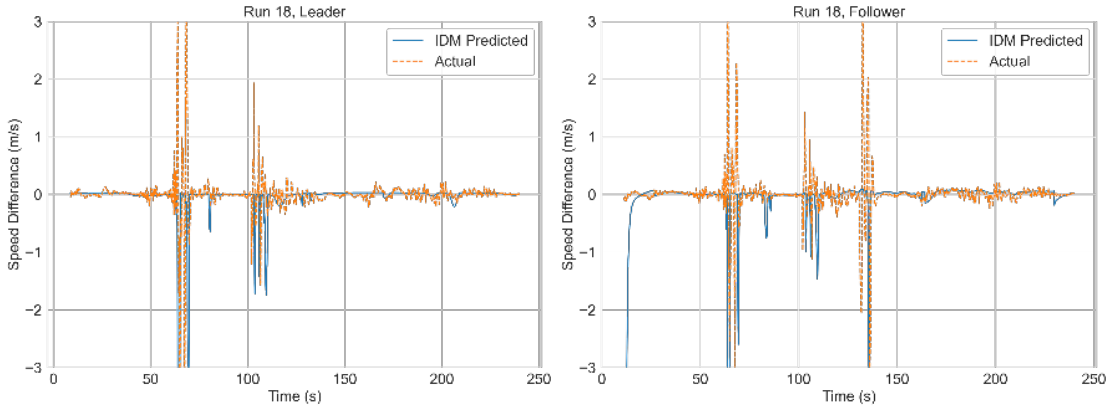

(b) IDM Speed Difference (Left: Leader, Right: Follower)

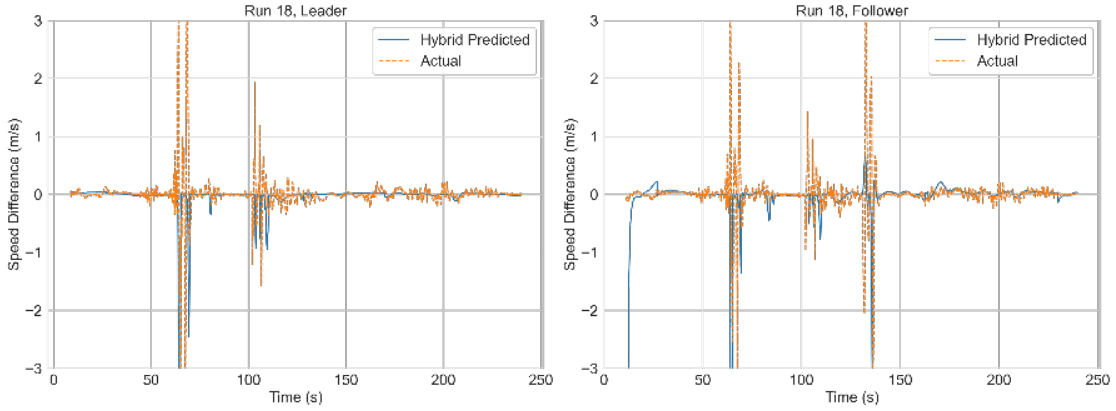

(c) Hybrid Model Speed Difference (Left: Leader, Right: Follower)

Source: FHWA

Figure 26.Graphs. Step-wise change in speed predicted by the models calibrated under the two-loop procedure.

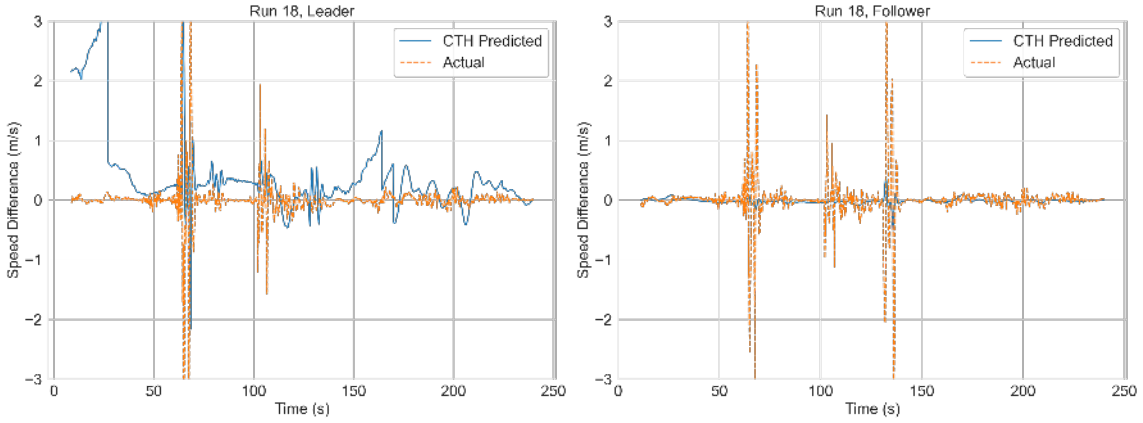

(a) CTH Model Speed Difference (Left: Leader, Right: Follower)

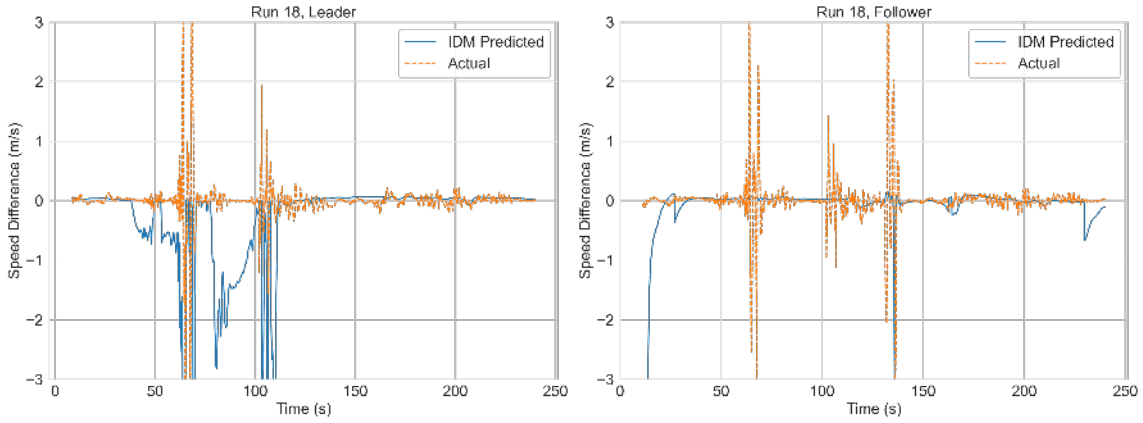

(b) IDM Speed Difference (Left: Leader, Right: Follower)

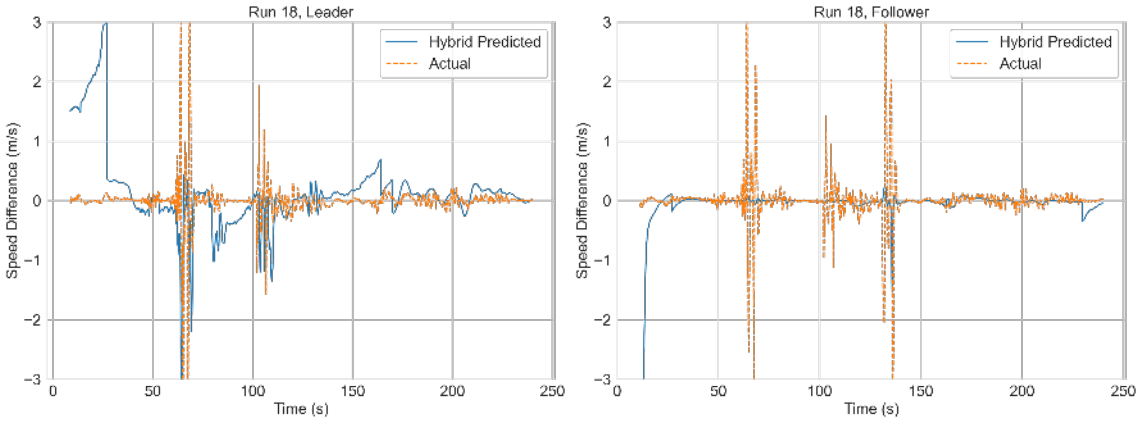

(c) Hybrid Model Speed Difference(Left: Leader, Right: Follower)

Source: FHWA

Figure 27: Stepwise change in speed predicted by the models calibrated under the joint estimation procedure.

Simulation Results

In order to evaluate the performance of the hybrid model, a continuous simulation is performed of the target vehicles, uninterrupted from its initial position and speed at the earliest point of the available data, except when there was a sudden drop in the vehicle's headway between the leading vehicle (Points of discontinuity in the "Predicted" curves in Figure 28.a and Figure 29.a correspond to when the position and speed were reset.) This is similar to the simulation approach used in the joint estimation procedure. The simulation results are plotted in Figure 28 and Figure 29. They show that, while the models calibrated under both methods are able to reconstruct the general trends of the trajectory, the one that was calibrated under the joint estimation procedure has better performance. While both sets of parameters almost perfectly replicate the position trajectory of the Follower and have small errors in space headway, the parameters under the joint estimation method were able to reconstruct significantly better trajectories for the Leader. It shows the robustness of the calibrated model. Interestingly enough, the two-loop method showed better progress during calibration but did not perform as well in simulation validation.

The result shows that good stepwise fitness in calibration does not guarantee global accuracy in a continuous simulation.

Nonetheless, there are several anomalies in the graphs. In Figure 28.b, Figure 28.c, Figure 29.b and Figure 29.c, both the speed and the space headway had unnatural shapes. The speed dipped to 0 before rebounding, and the space headway had large oscillations. This happened around when we reset the position and speed because there was a large reduction in the Follower's space headway in the original data, from 35 m to 12 m. The reduction likely caused a near-collision condition and forced the car-following models to respond with a very large deceleration. It is believed to be a case of another vehicle cutting in front of the target vehicle. The simulation was reset because neither of the car-following models was able to normally cope with such a near-collision situation. The drivers of the project vehicles do remember performing hard break maneuvers several times during the data collection because of other cars cutting in. Besides this anomaly in the original data, there are several irregular values in the observed speed and headway of the target vehicles, which can be seen from the "Actual" curves in the figures. Note that some of these anomalies may represent actual anomalies in the highway traffic, while the others are likely caused by errors in data collection and extraction. Those extreme values may have distorted the calibrated model parameters. This is especially true for the Leader, whose traffic states were very volatile, and the calibration was unreliable at times.

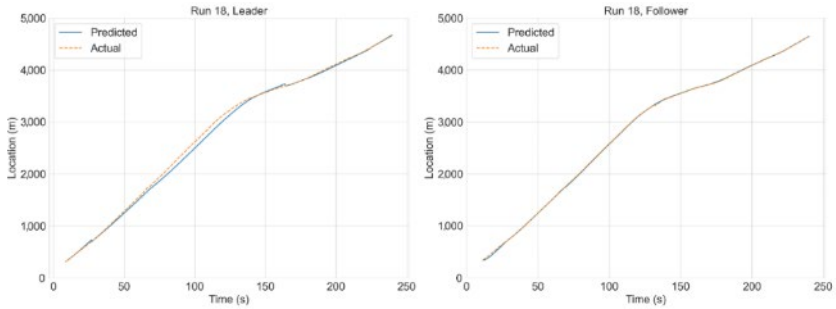

(a) Cumulative Location (Left: Leader, Right: Follower)

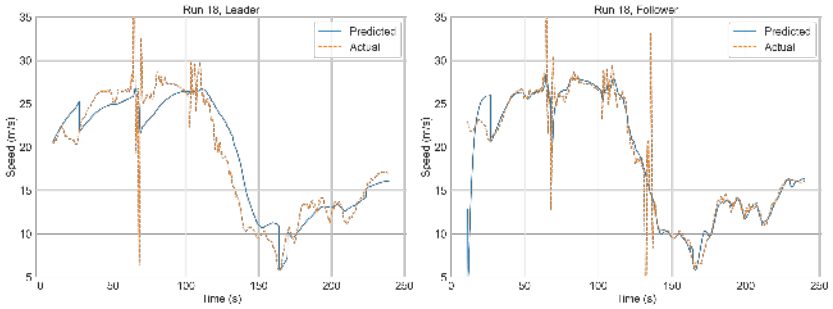

(b) Cumulative Speed (Left: Leader, Right: Follower)

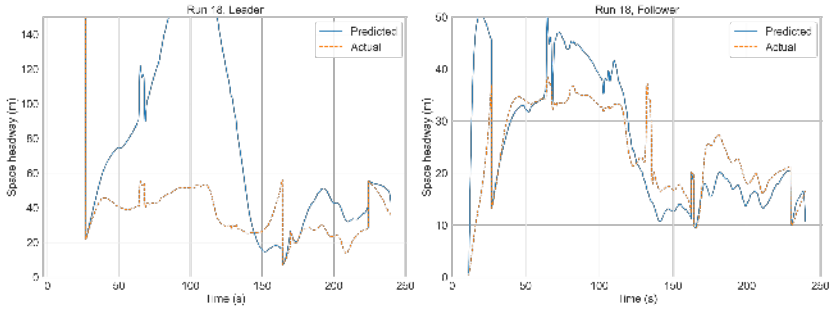

(c) Cumulative Headway (Left: Leader, Right: Follower)

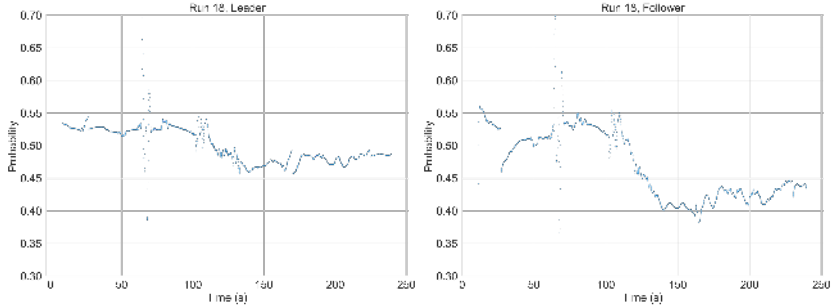

(d) Probability Distribution. 1 is IDM and 0 is CTH (Left: Leader, Right: Follower)

Source: FHWA

Figure 28: Continuous simulation results of the hybrid model calibrated under the two-loop calibration procedure.

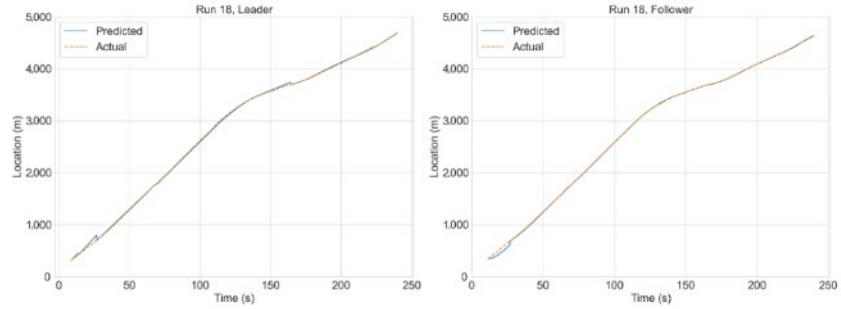

(a) Cumulative Location (Left: Leader, Right: Follower)

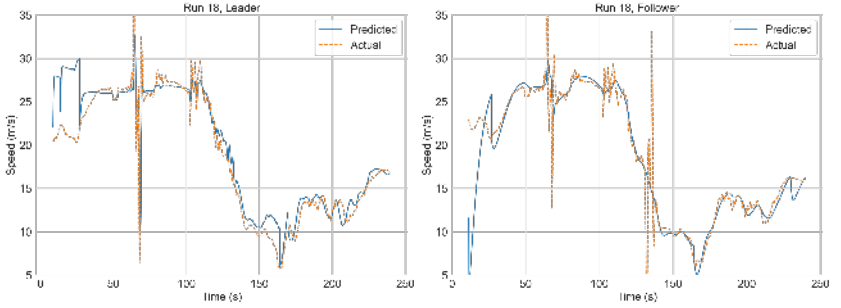

(b) Cumulative Speed (Left: Leader, Right: Follower)

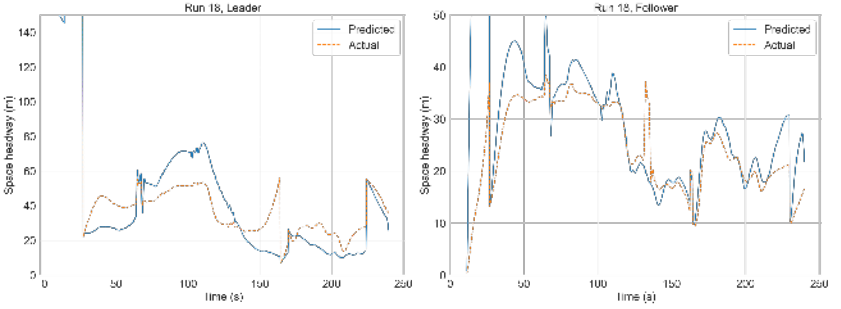

(c) Cumulative Headway (Left: Leader, Right: Follower)

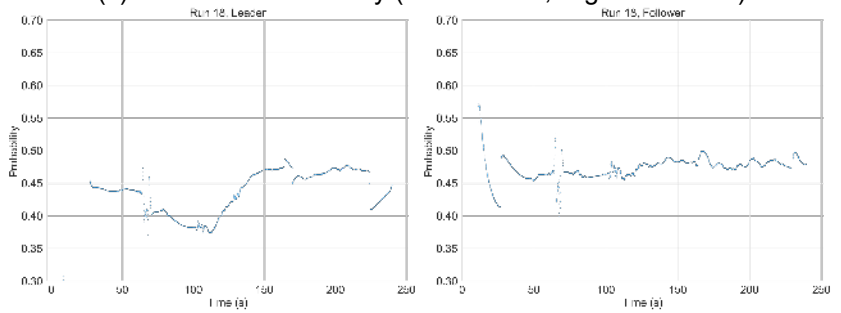

(d) Probability Distribution. 1 is IDM and 0 is CTH (Left: Leader, Right: Follower)

Source: FHWA

Figure 29: Continuous simulation results of the hybrid model calibrated under the joint estimation procedure.

Chapter 5. Lessons Learned

Developing the data collection plan was highly iterative and the project team conducted multiple test runs at several locations to identify the best combinations of time and location, camera setup, and vehicle setup to collect the most comprehensive trajectory dataset. The following section presents some of the lessons learned throughout this process.

Location Considerations and Data Collection objectives

One of the main considerations for data collection is the location. It is critical to identify locations that offer unique insight into traffic flow dynamics. In this study, four data collection sites were selected: I–294, I–90/I–94, I–395, and Foggy Bottom, D.C.

The I–294 location was selected as a representation of a typical highway environment:

- Segment offered various traffic flow regimes (i.e., free-flow, slow-and-go, stop-and-go, and fully congested).
- Segment length was adequate for forming and maintaining strings of ADAS-equipped vehicles.
- Southbound direction also has a major on-ramp that results in interesting mandatory lane-changing instances.
- Segment has a very high heavy vehicle percentage and offers unique insight into passenger car AV interactions with trucks.

The I–90/I–94 was selected for its ability to provide more insight into lane-changing behavior:

- Location has a major weaving section with a significant number of mandatory lane-changing instances, including last-minute forced lane-changing maneuvers.
- Location provides insight into the automated lane-changing ability of SAE Level 2 vehicles as well as human lane-changing decision making.

The I–395 location was selected to capture complex urban expressways.

- Segment selected due to its location and its importance to the Washington D.C. transportation network.
- Segment offered a diverse traffic flow regime and mandatory lane-changing instances from an on-ramp located on the left side of the highway.
- Segment offered considerable heavy vehicle traffic that creates interesting ADAS-truck interactions.

Finally, the Foggy Bottom location (George Washington University campus) was selected as the only arterial environment for data collection:

- Location enabled capturing vehicle interactions with various motorized and non-motorized modes of transportation (including pedestrians, bicyclists, and scooterists).

- Location provided low-speed interactions at both stopped-controlled and signalized intersections provide a unique data collection opportunity from interactions of these modes with automated vehicles.

It is critical to note that these locations were selected based on the project objectives. Any interested researcher is advised to start by defining the goals and objectives and select the roadway segments that address the requirements associated with their objectives. The presented approaches in this study can be applied to almost any location throughout the US.

In addition to the location requirements and data collection objectives, other considerations play an important role and can determine the most efficient data collection approach. For example, the project team utilized the infrastructure-based videography approach in Washington D.C. since it has the most restricted airspace in the country. Such a consideration also limited the location of the data collection and reduced the segment length that could feasibly be used for data collection as compared to other videography methods. The day of the data collection also plays an important role and various daily and seasonal traffic variations should be considered. In addition, any usual situations or special events that may impact traffic should be considered. For instance, the data collection on I–294 had to be delayed until the COVID-19 pandemic restrictions were lifted and traffic returned to normal (this was particularly important for the truck traffic).

Object Detection Considerations

When analyzing the data, accurate and precise detections of vehicles play a key role and prevent a lot of challenges in later data extraction steps. Accurately detecting large articulated heavy trucks can be particularly challenging. It is recommended to use ‘polygons’ for truck detection rather than rectangles since rectangles become very inaccurate, especially around curved road segments, due to the length of the heavy vehicles. It is also preferable to train a rotated bounding box model to detect vehicles. This is especially helpful when the study area consists of significant curvatures or turns at intersections. Such an approach eliminates an additional level of inaccuracy caused by right-angle bounding boxes at curved locations when the bounding boxes do not only cover the vehicle but also some of the surrounding environment.

Moreover, it is important to note that despite the existence of state-of-the-art object detection algorithms and pre-trained models, the addition of a manual annotation step to each dataset and then using transfer learning and retraining to fine-tune each model to each dataset reduces the number of missed detections. This is critical in future steps of trajectory extraction and data cleaning. This is particularly important for helicopter-based data collection efforts.

Unfortunately, the existing calibrated models rarely have similar images as part of their training process (e.g., satellite images are not utilized to train for vehicle detection) and all the algorithms tested by the project team failed to accurately capture vehicles in images captured by a helicopter. This limitation was remedied by re-calibrating these models based on manually annotated data.

Reference Image and Image Transformation Considerations

Creating a unique reference image for each dataset (and sometimes for every run in the same dataset), regardless of the existence of the satellite image, is advisable. This is because satellite images, despite being less distorted than the reference images, are harder in practice to match to the collected frames and can result in lower accuracy transformations. This is, in part, because the satellite images have no perspective, while the collected images can have some perspective (e.g., the helicopter is not always directly above the road and can be slightly to the side). Overall, the project team found that spending extra time to ensure that the reference image is accurate can potentially save several hours in future data post-processing and cleaning.

When constructing a reference image based on a satellite image, ensuring that every feature matches perfectly with the satellite image is difficult because the captured images (as part of the data collection) inevitably form a certain angle with the road. In this case, one should try to ensure that the road features (e.g., road markings and damages to the pavement) and road infrastructure features (e.g., road signs and toll plazas) match the satellite image. At the same time, ensuring that the captured images are smoothly stitched together to form the overall reference picture is essential to obtain a better transformation afterward.

After creating the reference image, each lane line must be depicted to facilitate the subsequent lane assignment and trajectory extraction steps. The lane lines should be fitted as accurately as possible to the reference image using distinct color (e.g., white - R:255, G:255, B:255) lines to mark the boundary between each lane in the reference image. The process can be done manually, however, that may result in very few white pixels that are discontinuous. Accordingly, linear interpolation should be performed based on the coordinates of the neighboring white pixel points to ensure continuity and avoid bugs in the process of lane assignment.

Depending on the size of the dataset, minimizing the size of the reference image— specifically the feature search area—can significantly increase the speed of the transformation process (which is one of the most time and resource intensive steps). In other words, looking for features in a small subset of the reference image is recommended to speed up the process.

Finally, producing low-resolution (faster playback) transformed images for subsequent manual checking is a key step in quality control and ensuring correct transformation is happening. It is important to note that transformation can result in negative coordinates (outside of the limits of the reference image). Such cases happen when there is an issue with the feature selection process. Plotting, visualizing, and actively looking for such issues help with producing higher-quality datasets and can reduce the data cleaning process significantly.

Trajectory Extraction Considerations

Accounting for distinct objects and features in the study area before beginning data collection and processing could also save a lot of time and improve the tracking process (e.g., accounting for overpass locations can help remove inaccurate detections, such as vehicles on the overpasses, and avoid gaps in vehicle trajectories). It is important to mention that visualizing and inspecting trajectory data after extraction is an essential step to catch and fix errors.

The data association step of the tracking process is not foolproof. Post-processing and manual checking are required to remove false positive associations. An example of a false positive association is where a vehicle stops being detected (due to leaving the view of the helicopter, becoming occluded, etc.) and the association step wrongly associates the frame of the next closest object to that same vehicle. This is time-consuming—but necessary—since there are cases that even the most sophisticated algorithms cannot capture automatically (e.g., a vehicle at the edge of the image leaving the helicopter frame and coming back after several seconds).

Removing vehicles before feature extraction in automated image transformation can improve the transformation outcome. A similar conclusion can be made for any temporary object in the image (e.g., removing trains from the images on I–90/ I–94 can improve the transformations).

ID matching is one of the most significant parts of data cleaning. The same vehicle can have different IDs due to shortly becoming occluded (e.g., going under an overpass), which results in gaps in the constructed trajectories. It is highly recommended that each trajectory be inspected through trajectory plotting, annotating original images with vehicle tracking IDs, and cross-validation with the extracted data. The research team attempted to implement multiple automated ID-matching processes. While some processes achieved high accuracies, the methodology still required manual ID matching. No algorithms can guarantee 100 percent accuracy, and thus manual matching is critical and cannot be avoided.

An accurate lane classification process is needed to ensure correct lane assignment and to avoid subsequent issues with crossing trajectories and manual data cleaning. However, this cannot always be achieved with an automated process. In particular, when the camera is not directly above the road, the center of the bounding box can fall on the adjacent lane (while the vehicle has not crossed the line) due to the size of the vehicle. Such a challenge can be fixed either by utilizing 3D bounding boxes (which can introduce additional issues during the detection and transformation, especially in birds-eye view videos) or by manual checking of the instances. Fortunately, such instances put the vehicle in the adjacent lane for only a short period of time, making finding such issues fairly straightforward by searching for two lane-changing maneuvers within a short period of time. This study selected 10 seconds as the period of time to mark such lane-changing instances to ensure that all of these instances can be captured (although most of the cases only last for a fraction of a second).

It is important to identify and differentiate missed detections (missed trajectories) from lane-changing maneuvers, as they both can look the same when trajectories are inspected on a lane-by-lane basis. Moreover, often, vehicles may leave the camera frame and show up again in the frame after several seconds (this happens if a vehicle is traveling slower or faster than the helicopter), which results in trajectories with a large gap. Fixing such instances, as discussed, requires manual inspection of the trajectory data.

Finally, it is critical to smooth the resulting trajectories, as the final trajectories are not always smooth due to several factors including detection and transformation challenges. Smoothing, however, should be carried out in a way to ensure the integrity of the data. Considering vehicle dynamics (and type of the vehicle) along with a tunable smoothing approach (e.g., Kalman filter based smoothing) is essential to ensure that the final trajectories are smooth and represent the actual observations.

Other Considerations

While all the aforementioned factors are important and can influence what data can represent, the time of the day directly impacts the quality of the extracted trajectories. Collecting data during morning and evening peak hours results in the most congestion, but these times of day typically have long shadows, as illustrated in Figure 30. These shadows can significantly degrade the quality of vehicle detection, leading to increased errors in trajectory data. It is crucial, therefore, that the data extraction process incorporates strategies to mitigate the impact of these conditions, ensuring the reliability of the collected data. Note that any data collection should be performed under safe weather and road surface conditions. However, interestingly enough, the research team found cloudy days were the best weather conditions for data collection due to the lack of shadows (as long as the cloud ceiling is high enough for the helicopter to fly at a desired altitude).

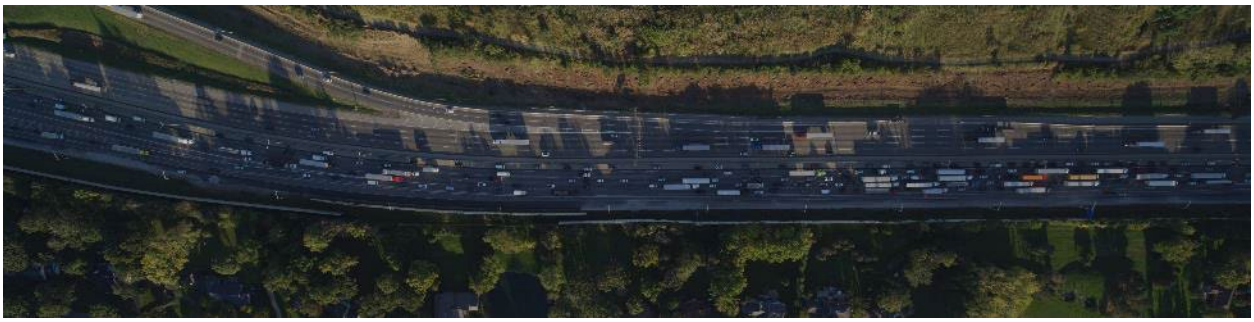

Figure 30: A sample frame from data collection on I-294 with long shadows.

Camera and helicopter setup also play critical roles in the quality of the collected data. The video quality and frame rate should be high enough to provide accurate detection and the required point frequency in the trajectory data. This study utilized 8K videos in the helicopter, but even with that quality, the flight altitude should have been limited to 300 m. Assuming a lens with no distortion, the camera could capture about 0.8 km of the highway at that altitude. Flying at a higher altitude (as was done at some of the collection sites, mostly during test flights), however, can significantly reduce vehicle detection accuracy (and trajectory accuracy consequently). The video resolution was significantly reduced in Washington D.C. since the cameras are much closer to the highway and vehicles can be detected accurately at a 4K resolution.

Depending on the nature of the data collection and its location, many of the aforementioned manual processes can become automated, significantly reducing the need for many of the cleaning and post-processing steps. For instance, many of the above challenges can be avoided when the helicopter is hovering and detection and transformation can be completed with much higher accuracy (i.e., fewer missed detections, vehicles going out of frame, ID matching, etc.) compared with a moving helicopter case.

Finally, it is important to note that every study is different and requires considerations that are specific to the study area, method of data collection, type of data, required outcomes, and deliverables. Due to the study requirements and needed accuracy, the trajectory extraction process should be involves several iterations/repetitions until reaching satisfactory results.

Chapter 6. Conclusions

This study produced a set of trajectory datasets focusing on capturing human-AV vehicle interactions under a diverse set of scenarios in highway and urban street environments. This study also includes a detailed introduction to the methodologies utilized for converting video frames to vehicle trajectory data recorded by three distinct data collection methods:

- Fixed location aerial videography (a helicopter hovers over a segment of interest). Utilizing the fixed location aerial videography approach, a dataset was collected on I–90/I–94 in Chicago, IL capturing automated lane-changing maneuvers by SAE Level 2 ADAS-equipped vehicles.
- Moving aerial videography (a helicopter follows the subject vehicles as they navigate a much longer segment than the first method). Utilizing the moving aerial videography approach, another dataset was collected on I–90/I–94 in Chicago, IL, and two datasets were collected on I–294 near Hinsdale, IL. The dataset on I–90/I–94 and one of the datasets on I–294 focused on automated lane-changing maneuvers by SAE Level 2 ADAS-equipped vehicles. The other dataset on I–294 focused on SAE Level 1 ACC-driven automated vehicles.
- Infrastructure-based videography (multiple overlapping cameras located on overpasses or on top of buildings to create a comprehensive image of the study environment). Utilizing the infrastructure-based videography approach, two datasets were collected in Washington D.C. The dataset on I–395 focused on SAE Level 2 ADAS-equipped vehicles (the automated lane-changing opportunities were, however, significantly more limited compared with the Chicago datasets due to heavier congestion and a shorter data collection segment). The dataset on the George Washington University campus (located in Foggy Bottom) focused on SAE Level 3 automated driving system-equipped vehicles in a city environment.

Overall, these datasets can facilitate research aimed towards a comprehensive understanding of human-automated vehicle interactions in both congested and uncongested traffic regimes as well as both highway and city environments. Covering different SAE Levels of automation also provides the opportunity to conduct a deep analysis of the impacts of different automation SAE Levels on traffic flow dynamics, safety, and human behavior.

It is also important to note that despite the diverse set of datasets collected as part of this project, the collected data has some limitations:

- Not all the data points in the trajectory data were observed. Each data collection site had some blind spots that were not covered in the collected footage (e.g., each highway had multiple overpasses blocking the view of the camera). To ensure the continuity of the trajectories, the unobserved data was generated through simple interpolation. Such data points (despite being a very small subset of the data) are marked in the dataset and depending on the application, researchers are encouraged to choose the entire trajectory data or the observed portion.
- Despite the existence of long trajectories for SAE Level 2 ADAS-equipped vehicles in the dataset, the number of automated lane-changing maneuvers is still limited. This is mainly because the data collection avoided unnecessary (and excessive) lane-changing maneuvers (to ensure the safety of the drivers). Accordingly, there is a significant mismatch between the number of automated lane-changing

maneuvers and the number of lane-changing maneuvers performed by other vehicles (that were not the test vehicles). It is noteworthy that each dataset from SAE Level 2 vehicles should have enough automated lane-changing maneuvers to analyze the impacts of this maneuver on the traffic stream and to study human interactions with this maneuver.

- The dataset does not clearly mark the ACC and Lane-Keeping disengagement instances.
- Due to heavy and irregular activity in and around the arterial road, data collection from SAE Level 3 automated driving proved to be challenging and only a small portion of the Foggy Bottom dataset covers SAE Level 3 ADS-equipped vehicles.
- The data collection utilized a small number of commercially available passenger vehicles for each level of autonomy. The performance of automated driving systems can greatly vary for different vehicles from different manufacturers (and even different makes and models from the same manufacturer). A more focused study is necessary if one needs to capture the impacts of different vehicles from different manufacturers on traffic and human-automated vehicle interactions.

To illustrate the capabilities of the collected datasets, this study also presented the formulation of a hybrid car-following model that is intended to capture the behavior of an ACC-capable commercial vehicle in real-world settings. ACC is assumed to be in control of the vehicle most of the time, while the human driver is expected to occasionally intervene. The model has three components: two underlying car-following models that are assumed to represent the behaviors of a pure human or ACC and a logit model for the probability of choosing one of the models under specific traffic states. This study also presented two calibration methods that were designed to calibrate the model. The genetic algorithm is an important part of both methods because of its excellent performance in calibrating non-linear models with a large number of parameters. The car-following model was calibrated using both of the proposed methods. The calibration was conducted using the SAE Level 1 dataset collected on I–294. The results showed that the proposed model is capable of capturing the general trends of vehicular dynamics, and the joint estimation method produced a set of model parameters that result in better performance at reconstructing the trajectory in the simulation.

Finally, this report is concluded with a discussion on the lessons learned from this data collection effort. The limitations of the presented trajectory extraction method are discussed and detailed suggestions are made to help interested researchers avoid unnecessary trial-and-error to collect similar datasets.

Potential Research Utilizing the Collected Data

While this report is not focused on providing a comprehensive list of potential research problems that can be facilitated by the collected data, a short list of research topics that have not been studied extensively in the literature and can be investigated more thoroughly using the collected data is presented below.

- Lane-changing behavior and its impact on traffic flow dynamics is a topic worthy of investigation. This area of research has not been studied extensively mainly due to the lack of comprehensive and reliable data. The data collected as part of this study provides an unparalleled opportunity to conduct a thorough analysis of the lane-changing behavior and to develop models that accurately capture the lane-changing behavior of non-automated and SAE Level 2 vehicles.
- Modeling the behavior of SAE Level 1, SAE Level 2, and SAE Level 3 in a naturalistic environment is a critical step towards capturing the impacts of these technologies on traffic flow dynamics. While many of the studies in these areas focus on data from test tracks, the collected data in this study offers the opportunity to investigate and model the behavior of these vehicles in real-world settings.

- Data-driven modeling and simulation of traffic flow have the potential to fundamentally change the way we simulate different modes of transportation. Data availability has always been the key challenge preventing the development of robust data-driven models. The collected data offers a considerable amount of data enabling the development and training of data-driven traffic flow models.

Future Data collection Needs

While this study offers comprehensive trajectory data from vehicles, pedestrians, and other non-motorized modes of transportation, the research community can still benefit from the following additional data:

- This study offers comprehensive trajectory data from SAE Level 1 and SAE Level 2 vehicles in the highway environment (the environment that they are designed to operate in) as well as data from SAE Level 3 vehicles in urban settings. With the introduction of commercial SAE Level 4 vehicles and their superior performance compared with SAE Level 3 vehicles in urban environments, it is critical to collect data from SAE Level 4 automation.
- This study offers comprehensive trajectory data from vehicles and pedestrians. However, data from other non-motorized modes of transportation (e.g., bicyclists and scooterists) are limited. There is a critical need to collect data from these modes interacting with other modes (including non-automated vehicles) as the market penetration of these modes increases in urban communities.
- Vehicle technologies have been changing (e.g., the SAE Level 3 vehicles utilized in this study have received multiple updates since the data collection was conducted with each update increasing the performance and safety of their operation). There is a critical need to create methodologies based on new, limited data collection efforts, to quickly adjust the existing datasets for the evolving changes in vehicle technologies.

Chapter 7. References

AAA. 2019. Advanced Driver Assitance Technology Names. Last accessed Oct. 5, 2023. https://www.aaa.com/AAA/common/AAR/files/ADAS-Technology-Names-Research-Report.pdf

AASHTO. A Policy on Geometric Design of Highways and Streets. American Association of State Highway and Transportation Officials, Washington, D.C., 2011

Adobe Photoshop, C. Adobe https://www.adobe.com/products/photoshop.html . RRID: SCR_014199.

Ahmed, H. U., Y. Huang, and P. Lu. A Review of Car-Following Models and Modeling Tools for Human and Autonomous-Ready Driving Behaviors in Micro-Simulation. Smart Cities, Vol. 4, No. 1, 2021, pp. 314–335.

Agarwal, S., Vora, A., Pandey, G., Williams, W., Kourous, H., and McBride J. "Ford Multi-AV Seasonal Dataset", *The International Journal of Robotics Research*, Volume: 39 issue: 12 (2020), page(s): 1367-1376 https://arxiv.org/abs/2003.07969

Barmpounakis, E., and N. Geroliminis. On the New Era of Urban Traffic Monitoring with Massive Drone Data: The pNEUMA Large-Scale Field Experiment. Transportation research part C: emerging technologies, Vol. 111, 2020, pp. 50–71.

Bay, H., Tuytelaars, T. and Van Gool, L., 2006. Surf: Speeded up robust features. In Computer Vision–ECCV 2006: 9th European Conference on Computer Vision, Graz, Austria, May 7-13, 2006. Proceedings, Part I 9 (pp. 404-417). Springer Berlin Heidelberg.

Canalys. 2021. Huge opportunity as only 10% of the 1 billion cars in use have ADAS features. Last accessed Oct. 5, 2023. Canalys Newsroom - Huge opportunity as only 10% of the 1 billion cars in use have ADAS features

Cattin, J., L. Leclercq, F. Pereyron, and N. El Faouzi. Calibration of Gipps' Car-following Model for Trucks and the Impacts on Fuel Consumption Estimation. IET Intelligent Transport Systems, Vol. 13, No. 2, 2019, pp. 367–375.

Chen, C., L. Li, J. Hu, and C. Geng. Calibration of MITSIM and IDM Car-Following Model Based on NGSIM Trajectory Datasets. Presented at the Proceedings of 2010 IEEE International Conference on Vehicular Electronics and Safety, 2010.

Czarnecki, K., 2018. Operational design domain for automated driving systems. Taxonomy of Basic Terms “, Waterloo Intelligent Systems Engineering (WISE) Lab, University of Waterloo, Canada.

Deo, N. and Trivedi, M.M., 2018. Convolutional social pooling for vehicle trajectory prediction. In Proceedings of the IEEE conference on computer vision and pattern recognition workshops (pp. 1468-1476).

Derpanis, K. G. The Harris Corner Detector. York University, Vol. 2, 2004.

Derpanis, K. G. Overview of the RANSAC Algorithm. Image Rochester NY, Vol. 4, No. 1, 2010, pp. 2–3.

Gipps, P. G. A Behavioural Car-Following Model for Computer Simulation. Transportation Research Part B: Methodological, Vol. 15, No. 2, 1981, pp. 105–111.

Girshick, R., Donahue, J., Darrell, T. and Malik, J., 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 580-587).

Hamdar, S. H., M. Treiber, H. S. Mahmassani, and A. Kesting. Modeling Driver Behavior as Sequential Risk-Taking Task. Transportation research record, Vol. 2088, No. 1, 2008, pp. 208–217.

Hamdar, S. H., H. S. Mahmassani, and M. Treiber. From Behavioral Psychology to Acceleration Modeling: Calibration, Validation, and Exploration of Drivers’ Cognitive and Safety Parameters in a Risk-Taking Environment. Transportation Research Part B: Methodological, Vol. 78, 2015, pp. 32–53.

Huang, Xinyu, Peng Wang, Xinjing Cheng, Dingfu Zhou, Qichuan Geng, and Ruigang Yang. "The apolloscape open dataset for autonomous driving and its application." IEEE transactions on pattern analysis and machine intelligence 42, no. 10 (2019): 2702-2719.

Ioannou, P. A., and C.-C. Chien. Autonomous Intelligent Cruise Control. IEEE Transactions on Vehicular technology, Vol. 42, No. 4, 1993, pp. 657–672.

Jocher, G., A. Chaurasia, A. Stoken, J. Borovec, Y. Kwon, K. Michael, J. Fang, Z. Yifu, C. Wong, and D. Montes. Ultralytics/Yolov5: V7. 0-Yolov5 Sota Realtime Instance Segmentation. Zenodo, 2022.

Kala, Rahul. 2016. On-road intelligent vehicles: Motion planning for intelligent transportation systems. Butterworth-Heinemann.

Kamjoo, E., R. Saedi, A. Zockaie, M. Ghamami, T. Gates, and A. Talebpour. Developing Car-Following Models for Winter Maintenance Operations Incorporating Machine Learning Methods. Transportation research record, Vol. 2677, No. 2, 2023, pp. 519–540.

Kesten, R., M. Usman, J. Houston, T. Pandya, K. Nadhamuni, A. Ferreira, M. Yuan, B. Low, A. Jain, and P. Ondruska. Lyft Level 5 Av Dataset 2019. *urlhttps://level5. lyft. com/dataset*, Vol. 1, 2019, p. 3.

Kesting, A., and M. Treiber. Calibrating Car-Following Models by Using Trajectory Data: Methodological Study. Transportation Research Record, Vol. 2088, No. 1, 2008, pp. 148–156.

Khajeh-Hosseini, M., A. Talebpour, S. Devunuri, and S. H. Hamdar. An Unsupervised Learning Framework for Detecting Adaptive Cruise Control Operated Vehicles in a Vehicle Trajectory Data. Expert Systems with Applications, 2022, p. 118060.

Krajewski, R., Bock, J., Kloeker, L. and Eckstein, L., 2018, November. The highd dataset: A drone dataset of naturalistic vehicle trajectories on german highways for validation of highly automated driving systems. In 2018 21st international conference on intelligent transportation systems (ITSC) (pp. 2118-2125). IEEE.

Lu, Y., J. Fu, G. Tucker, X. Pan, E. Bronstein, B. Roelofs, B. Sapp, B. White, A. Faust, and S. Whiteson. Imitation Is Not Enough: Robustifying Imitation with Reinforcement Learning for Challenging Driving Scenarios. arXiv preprint arXiv:2212.11419, 2022.

Lin, T.Y., Goyal, P., Girshick, R., He, K. and Dollár, P., 2017. Focal loss for dense object detection. In Proceedings of the IEEE international conference on computer vision (pp. 2980-2988).

Lowe, D. G. Distinctive Image Features from Scale-Invariant Keypoints. International journal of computer vision, Vol. 60, No. 2, 2004, pp. 91–110.

Mahmassani, H. S., A. Elfar, S. E. Shladover, and Z. Huang. Development of an Analysis/Modeling/Simulation (AMS) Framework for V2I and Connected/Automated Vehicle Environment. United States. Department of Transportation. Intelligent Transportation, 2018

Makridis, M., K. Mattas, B. Ciuffo, F. Re, A. Kriston, F. Minarini, and G. Rognelund. Empirical Study on the Properties of Adaptive Cruise Control Systems and Their Impact on Traffic Flow and String Stability. Transportation research record, Vol. 2674, No. 4, 2020, pp. 471–484.

Makridis, M., K. Mattas, A. Anesiadou, and B. Ciuffo. OpenACC. An Open Database of Car-Following Experiments to Study the Properties of Commercial ACC Systems. Transportation research part C: emerging technologies, Vol. 125, 2021, p. 103047.

Muja, M., and D. Lowe. Fast Library for Approximate Nearest Neighbors (FLANN). git://github. com/mariusmuja/flann. git. url: http://www. cs. ubc. ca/research/flann, 2013.

NHTSA, 2021 https://www.nhtsa.gov/vehicle-safety/driver-assistance-technologies?gad_source=1#the-topic-technology-saves-lives

Nowakowski, C., Shladover, S.E., Cody, D., Bu, F., O'Connell, J., Spring, J., Dickey, S. and Nelson, D., 2011. Cooperative adaptive cruise control: Testing drivers' choices of following distances (No. UCB-ITS-PRR-2011-01).

Ozkan, M. F., and Y. Ma. Modeling Driver Behavior in Car-Following Interactions with Automated and Human-Driven Vehicles and Energy Efficiency Evaluation. IEEE Access, Vol. 9, 2021, pp. 64696–64707.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, and V. Dubourg. Scikit-Learn: Machine Learning in Python. the Journal of machine Learning research, Vol. 12, 2011, pp. 2825–2830.

Punzo, V., M. T. Borzacchiello, and B. Ciuffo. On the Assessment of Vehicle Trajectory Data Accuracy and Application to the Next Generation SIMulation (NGSIM) Program Data. Transportation Research Part C: Emerging Technologies, Vol. 19, No. 6, 2011, pp. 1243–1262.

Rahmati, Y., M. Khajeh Hosseini, A. Talebpour, B. Swain, and C. Nelson. Influence of Autonomous Vehicles on Car-Following Behavior of Human Drivers. Transportation research record, Vol. 2673, No. 12, 2019, pp. 367–379.

Rakha, H., C. C. Pecker, and H. B. B. Cybis. Calibration Procedure for Gipps Car-Following Model. Transportation Research Record, Vol. 1999, No. 1, 2007, pp. 115–127.

Redmon, J., Divvala, S., Girshick, R. and Farhadi, A., 2016. You only look once: Unified, real-time object detection. In Proceedings of the IEEE conference on computer vision and pattern recognition (pp. 779-788).

Rublee, E., Rabaud, V., Konolige, K. and Bradski, G., 2011, November. ORB: An efficient alternative to SIFT or SURF. In 2011 International conference on computer vision (pp. 2564-2571). IEEE.

Sadigh, D., Sastry, S., Seshia, S.A. and Dragan, A.D., 2016, June. Planning for autonomous cars that leverage effects on human actions. In Robotics: Science and systems (Vol. 2, pp. 1-9).

Sivanandam, S., S. Deepa, S. Sivanandam, and S. Deepa. Genetic Algorithm Optimization in C/C++. Introduction to Genetic Algorithms, 2008, pp. 263–316.

Sun, P., Kretzschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B. and Vasudevan, V., 2020. Scalability in perception for autonomous driving: Waymo open dataset. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 2446-2454).

Swaroop, D., J. K. Hedrick, C. Chien, and P. Ioannou. A Comparision of Spacing and Headway Control Laws for Automatically Controlled Vehicles1. Vehicle system dynamics, Vol. 23, No. 1, 1994, pp. 597–625.

Swaroop, D., and K. Rajagopal. A Review of Constant Time Headway Policy for Automatic Vehicle Following. ITSC 2001. 2001 IEEE Intelligent Transportation Systems. Proceedings (Cat. No. 01TH8585), 2001, pp. 65–69.

Tang, Z., Wang, G., Xiao, H., Zheng, A. and Hwang, J.N., 2018. Single-camera and inter-camera vehicle tracking and 3D speed estimation based on fusion of visual and semantic features. In Proceedings of the IEEE conference on computer vision and pattern recognition workshops (pp. 108-115).

Talebpour, A., H. S. Mahmassani, and S. H. Hamdar. Multiregime Sequential Risk-Taking Model of Car-Following Behavior: Specification, Calibration, and Sensitivity Analysis. Transportation research record, Vol. 2260, No. 1, 2011, pp. 60–66.

Treiber, M., A. Hennecke, and D. Helbing. Congested Traffic States in Empirical Observations and Microscopic Simulations. Physical review E, Vol. 62, No. 2, 2000, p. 1805.

Van Rossum, G., and F. L. Drake. Introduction to Python 3: Python Documentation Manual Part 1. CreateSpace, 2009.

Varotto, S. F., R. G. Hoogendoorn, B. van Arem, and S. P. Hoogendoorn. Empirical Longitudinal Driving Behavior in Authority Transitions between Adaptive Cruise Control and Manual Driving. Transportation Research Record, Vol. 2489, No. 1, 2015, pp. 105–114.

Vasconcelos, L., L. Neto, S. Santos, A. B. Silva, and Á. Seco. Calibration of the Gipps Car-Following Model Using Trajectory Data. Transportation research procedia, Vol. 3, 2014, pp. 952–961.

“Why the Pandemic Has Disrupted Supply Chains”. The White house (2021) https://www.whitehouse.gov/cea/written-materials/2021/06/17/why-the-pandemic-has-disrupted-supply-chains/

Xiao, L., and F. Gao. A Comprehensive Review of the Development of Adaptive Cruise Control Systems. Vehicle system dynamics, Vol. 48, No. 10, 2010, pp. 1167–1192.

Yin, H., and C. Berger. When to Use What Data Set for Your Self-Driving Car Algorithm: An Overview of Publicly Available Driving Datasets. Presented at the 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC), 2017.

Yu, L., and R. Wang. Researches on Adaptive Cruise Control System: A State of the Art Review. Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering, Vol. 236, No. 2–3, 2022, pp. 211–240.

Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., Darrell, T. “BDD100K: A Diverse Driving Dataset for Heterogeneous Multitask Learning”. Published at IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2020. https://arxiv.org/abs/1805.04687

Yu, H., Jiang, R., He, Z., Zheng, Z., Li, L., Liu, R., Chen, X.: Automated vehicle-involved traffic flow studies: A survey of assumptions, models, speculations, and perspectives. Transp. Res. Part C: Emerging Technol. 127, 103101 (2021). https://doi.org/10.1016/j.trc.2021.103101

Yurtsever, E., J. Lambert, A. Carballo, and K. Takeda. A Survey of Autonomous Driving: Common Practices and Emerging Technologies. IEEE access, Vol. 8, 2020, pp. 58443–58469.

Zhu, M., Y. Wang, Z. Pu, J. Hu, X. Wang, and R. Ke. Safe, Efficient, and Comfortable Velocity Control Based on Reinforcement Learning for Autonomous Driving. Transportation Research Part C: Emerging Technologies, Vol. 117, 2020, p. 102662.

U.S. Department of Transportation
ITS Joint Program Office – HOIT
1200 New Jersey Avenue, SE
Washington, DC 20590

Toll-Free “Help Line” 866-367-7487

www.its.dot.gov

FHWA-JPO-24-133

U.S. Department of Transportation