TRR
JOURNAL OF THE TRANSPORTATION RESEARCH BOARD

Research Article

Transportation Research Record
2025, Vol. 2679(1) 1768–1784

DOI: 10.1177/03611981241257257
journals.sagepub.com/home/trr
Sage

Introduction to the Third Generation Simulation Dataset: Data Collection and Trajectory Extraction

Rami Ammourah[1], Pedram Beigi[2], Bingyi Fan[3], Samer H. Hamdar[2], John Hourdos[4], Chun-Chien Hsiao[1], Rachel James[5], Mohammdreza Khajeh-Hosseini[1], Hani S. Mahmassani[3], Dana Monzer[3], Tina Radvand[1], Alireza Talebpour[1], Mahdi Yousefi[1], and Yanlin Zhang[1]

Abstract

This study aims to provide accurate trajectory datasets capable of characterizing human–automated vehicle interactions under a diverse set of scenarios in diverse highway environments. Distinct methods were utilized to collect data from Level 1, Level 2, and Level 3 automated vehicles: (1) fixed location aerial videography (a helicopter hovers over a segment of interest); (2) moving aerial videography (a helicopter follows the automated vehicles as they move in a much longer segment than in the first method); and (3) infrastructure-based videography (multiple overlapping cameras located on overpasses creating a comprehensive image of the study area). Utilizing the fixed location aerial videography approach, trajectories were extracted on I-90/I-94 in Chicago, IL. The moving aerial videography approach was adopted to extract four datasets on I-90/I-94 and I-294 in Chicago, IL. Finally, two datasets were collected on I-395 and George Washington University Campus in Washington, D.C., using the infrastructure-based videography approach. Extracting multiple complete and accurate vehicle trajectories raises a set of methodological and practical challenges that vary across the three data measurement approaches. The methodological details to extract these trajectories are presented in this paper along with the lessons learned with respect to data collection setup, instrumentation, and experimental design efforts.

Keywords

data and data science, data analysis, general, probe vehicle data, operations, automated/autonomous vehicles

The past few years have seen breakthroughs in connectivity and automation technologies, enabling the development of highly connected automated vehicles (CAVs). The introduction of such technologies to the transportation system has introduced a new set of interactions among roadway users (e.g., human–CAV interactions) that did not exist a few years ago. These interactions can 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.

CAVs' decision logic (at any automation level) relies on a broad set of information that they receive from onboard sensors, as well as (potentially) vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communications. At higher automation levels (Level 4 automation and above), fortunately, sophisticated design approaches have resulted in human-like behavior in most driving instances (although the response of these vehicles to high-risk and/or rare driving instances can be still governed by robotics-based and possibly rule-based

[1]Department of Civil and Environmental Engineering, University of Illinois at Urbana-Champaign, Urbana, IL
[2]Department of Civil and Environmental Engineering, George Washington University, Washington, DC
[3]Northwestern University Transportation Center, Evanston, IL
[4]Turner-Fairbank Highway Research Center, McLean, VA
[5]Office of Transportation Policy Studies, Washington, DC

Corresponding Author:
Hani S. Mahmassani, masmah@northwestern.edu

approaches). Unfortunately, the market penetration rate of higher levels of automation is significantly limited (at the time of this paper, fully automated fleet operations are offered in only two major cities in the U.S.A., Phoenix and San Francisco, both still in geofenced areas), and lower levels of automation (Levels 1–3) still utilize design approaches that can result in decision-making algorithms that are fundamentally different from humans' decision-making logic. These differences are rooted in the utilized state-of-the-art perception, motion planning, and control algorithms and are extremely hard to change (it is often required to do a complete design overhaul to achieve human-like behavior). Our understanding of the impacts of such differences in decision-making logic on the transportation system and its users is still in its infancy. In fact, despite some limited evidence (e.g., Rahmati et al. [*1*] found that human drivers choose a smaller headway when following adaptive cruise control [ACC]-driven vehicles as opposed to following other human-driven vehicles), it is not clear whether the underlying mechanisms of human–CAV interactions are fundamentally different from human–human interactions. Understanding the differences and similarities in this regard is, however, critical for managing the transportation system in the next several years, as Level 1 and Level 2 automation will dominate the new vehicle market and the market penetration of these vehicles will likely increase exponentially in the transportation system.

The first step toward addressing this fundamental question is to capture and investigate human–CAV interactions in a real-world setting. 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–CAV interactions (*2*, *3*). Moreover, while extensive data have been available from human–human interactions (e.g., NGSIM [*4*], HighD [*5*], pNeuma [*6*], etc.), data from human–CAV interactions have been quite limited. The existing datasets that contain CAV data either mostly focus on the behavior of CAVs only and do not focus on human–CAV interactions (e.g., the OpenACC [*7*] dataset is focused on the behavior of ACC-driven vehicles) or 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 also shared datasets from operating Level 3 and Level 4 autonomous vehicles on the road; e.g., Waymo Open Dataset [*8*], Ford Autonomous Vehicle Dataset [*9*], nuScenes [*10*], BDD100K [*11*], ApolloScape [*12*], Level 5 Open Data [*13*], etc.). It is noteworthy that there are some datasets from operating ACC systems in traffic (*14*); however, these datasets mostly focus 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 CAV datasets and their characteristics was offered by Yurtsever et al. (*15*) and Yin and Berger (*16*). Accordingly, none of the aforementioned datasets can be utilized to study the impacts of CAVs on human drivers in their vicinity. The only available trajectory dataset from CAV operations in a real-world setting for the purpose of understanding and characterizing human–CAV interactions was collected by Khajeh-Hosseini et al. (*17*) in Austin, TX. They collected vehicle trajectory data based on aerial videography while having three ACC-driven vehicles 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.

Accordingly, to address the above limitations, this study focuses on collecting data from human–CAV interactions in a real-world setting. Multiple trajectory datasets were collected with various levels of automation (Levels 1–3) 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–CAV as well as human–human interactions in CAVs' vicinity. These datasets can be utilized for various traffic flow and safety analyses, including but not limited to (1) investigating the impacts of Level 1 and Level 2 CAVs on human behavior; (2) the string stability of ACC-driven vehicles in real-world settings; (3) the dynamics and impacts of CAV–heavy vehicle interactions on traffic flow and safety; (4) the underlying dynamics of mandatory and discretionary lane-changing maneuvers; (5) car-following behavior under various traffic flow regimes; and (6) microscopic and macroscopic traffic flow analyses. In addition to the datasets, this study offers the following key contributions. (1) 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–CAV interactions). (2) A detailed experimental design approach is presented based on extensive experimentation with different design approaches to accurately capture data from vehicle interactions. Note that a new data collection approach is also developed as part of this study utilizing a moving helicopter. (3) A detailed discussion on lessons learned from multiple data collection efforts is also offered.

The remainder of this paper is organized as follows: the next section presents the details of the data collection process (including location, time, etc.). This section is

Table 1. Data Collection Summary

Site location	Duration	Time of day	Segment length	Number of lanes	Type of data	Number of CAVs	Data collection method	Altitude
I-90/I-94	2 h	3–5 p.m.	2.5 mi	4	Isolated Level 2	3	Vehicle following	700 ft
I-90/I-94	2 h	4–6 p.m.	0.8 mi	3–6	Isolated Level 2	2	Hovering	1000 ft
I-294	2 h	3–5 p.m.	3.0 mi	4	ACC platoon	3	Vehicle following	1000 ft
I-294	2 h	3–5 p.m.	3.0 mi	4	Isolated Level 2	2	Vehicle following	1000 ft
I-395	2 h	8–10 a.m.	0.34 mi	3	Isolated Level 2	2	Infrastructure-based	Ground level
George Washington University Campus	2 h	3–5 p.m.	4 blocks	1–2	Isolated Level 3	1	Infrastructure-based	The roof of multiple buildings

Note: CAVs = connected automated vehicles; ACC = adaptive cruise control.

followed by a discussion on the trajectory extraction process. The details of the collected trajectory data are presented next along with some analysis of the trajectory datasets. Finally, the paper is concluded with summary remarks and lessons learned from the data collection and trajectory extraction processes.

Data Collection Locations

Four locations were selected for data collection: (1) I-90/I-94 in Chicago, IL, (2) I-294 near Hinsdale, IL, (3) I-395 in Washington, D.C., and (4) George Washington University Campus in Washington D.C. In addition to including ACC-driven vehicles, each of these locations offers a unique driving environment that was not captured in any of the existing datasets. Table 1 summarizes the key characteristics of each data collection location.

I-90/I-94 in Chicago, IL

This data collection effort focused on Level 2 automated vehicles and contained a major weaving section on I-90/I-94 in Chicago, IL. Two sets of data were collected in this location. Figure 1, *a* and *b*, shows the overview of the segments that were used for data collection. The first set covered about 2.5 mi 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 between three and six lanes (four lanes for the majority of the 2.5 mi).

The second set offers 0.8 mi of highway driving (0.4 mi in each direction, as shown in Figure 1*b*) and a major weaving section in each direction (where I-90 and I-94 are separated in the northbound direction and join together in the southbound direction). The segment has two off-ramps and 2 on-ramps in the northbound direction and has between three and six lanes in each direction.

I-294 near Hinsdale, IL

This data collection effort focused on Level 1 and Level 2 automated vehicles on I-294 to the southwest of Chicago, IL. Figure 1, *a* and *c*, shows an overview of the segment that was used for data collection. This highway segment is controlled and operated by Illinois Tollway. This section offers 3 mi 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.

I-395 in Washington, D.C

This data collection effort focused on Level 2 automated 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 *e*, shows an overview of the segment that was used for data collection. This section is 0.34 mi 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-hand side. In addition to this on-ramp, the section covers an off-ramp on the right-hand side.

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

The focus of this data collection effort was on Level 3 automated vehicles operating 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-hand turns. Figure 1, *d*, *f*, and *g*, provides an overview of the specific area where data was collected. This segment

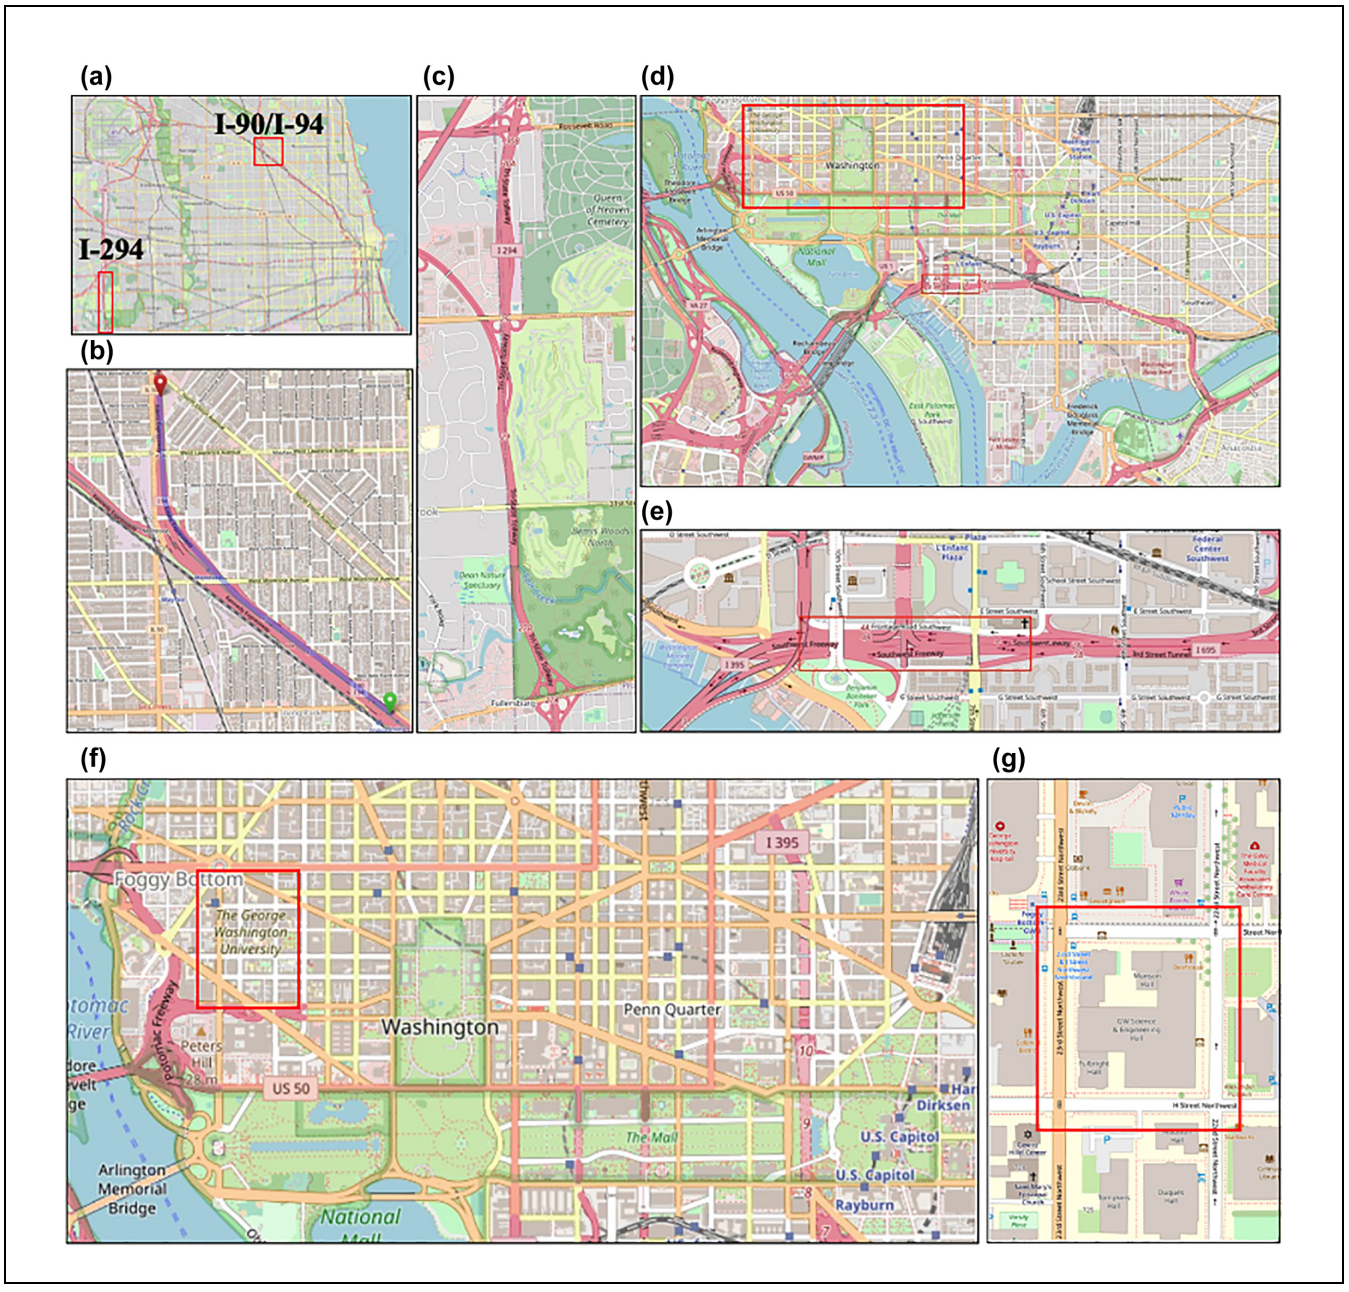

Figure 1. Data collection locations: (*a*) Chicago map and location of two data collection sites, (*b*) location of data collection on I-90/I-94, (*c*) location of data collection on I-294, (*d*) Washington, D.C., map and locations of the data collection site, (e) location of data collection on I-395, (*f*) location of data collection in Foggy Bottom, and (*g*) one block was utilized for data collection in Foggy Bottom.

encompasses one block within the George Washington University campus, with test vehicles moving mainly in a counter-clockwise direction. This area includes four intersections as well as a single block of I Street, H Street, 22nd Street, and 23rd Street.

Data Collection Setup

This section presents the experimental design and data collection details. It is important to note that data collection methodology and experiment design varied from one location to another to ensure capturing the most comprehensive datasets from human–CAV 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 presented toward the end of the paper to help interested readers with future experimental design. Table 1 summarizes the key characteristics of the data collection design.

Three distinct approaches were utilized for data collection in this study: (1) fixed location aerial videography; (2) moving aerial videography, and (3) infrastructure-based 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., the HighD [*5*] and Austin [*17*] datasets). The trajectory extraction process, discussed in the next section, is also the simplest among the three approaches. In the moving aerial videography approach, a helicopter follows one or more target vehicles as they move through the segment of interest. While this approach might not yield a significant benefit when dealing with human-driven vehicles only, it offers a significant advantage when collecting data from a small set of automated target vehicles. The round-trip time to get to the start of the segment of interest in the fixed location aerial videography approach can significantly limit the number of automated vehicle trajectories in the dataset. This was one of the main issues of the dataset collected by Khajeh-Hosseini et al. (*17*) in Austin, TX. Moreover, the length of the trajectories can be quite limited, resulting in capturing the impacts of automation in very limited traffic conditions (another problem identified by Khajeh-Hosseini et al. [*17*] in their dataset). The moving aerial videography approach addresses both of these limitations and offers long trajectories from automated vehicles while following them as they move through various traffic flow regimes. Note that this approach has never been implemented to collect vehicle trajectory data at scale. The infrastructure-based videography approach was utilized in Washington, D.C., particularly because of the limitations associated with flying a helicopter in that area. A simplified version of this approach was utilized in the NGSIM (*4*) dataset. The main difference is that this study utilizes several cameras positioned on overpasses along I-395, many of which are complimentary and overlapping, while others require interpolation.

I-90/I-94 in Chicago, IL

Two datasets were collected from Level 2 automated vehicles on I-90/I-94. 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*). Two Level 2 automated vehicles circulated and drove through this segment between 4 and 6 p.m. They entered the northbound direction upstream of the target section and exited the targe section on the right through I-94. 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 was needed to complete a run; conducting more lane-changing maneuvers within 0.25 mi from the entry point to the highway was infeasible). Note that while I-90 offered more interesting traffic flow dynamics (I-90 toward 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. In the southbound direction, the vehicles entered I-94 upstream of the segment of interest and merged with I-90. The data was collected via a RED camera at 30 fps at 8K resolution. The data collection altitude was set to 700 ft and the data was collected on a cloudy day.

The second dataset 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 to 5 p.m.) the helicopter followed a Level 2 automated vehicle in the northbound direction. The vehicle entered the highway from an on-ramp on the right, drove for about 2.5 mi, continued to I-90, and exited the highway through an off-ramp. Since the data was collected in the northbound direction only, three Level 2 vehicles were used in this study. 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 waiting for another turn. The data was collected via a RED camera at 30 fps at 8K resolution. The data collection altitude was set to 1000 ft and the data was collected on a cloudy day.

Note that both of these experiments aimed at collecting data from car-following and mandatory lane-changing maneuvers of Level 2 automated vehicles. All the vehicles were equipped with Global Positioning System (GPS) + inertial measurement units (IMUs) and they were all from the same manufacturer that offers Level 2 automated driving.

I-294 near Hinsdale, IL

Three datasets are collected from Level 1 and Level 2 automated vehicles on I-294 (see Figure 1) using the moving aerial videography approach. The aerial-based data collection efforts are intended to capture the impact of CAVs on the entire traffic stream and traffic flow dynamics. Each dataset covers about 2 h of data from 3 to 5 p.m. Out of the three datasets, one dataset focuses on a platoon of three full-range ACC-driven vehicles. The structure of the ACC platoon is similar to that utilized by Khajeh-Hosseini et al. (*17*) in Austin, TX. As discussed previously, the Austin dataset faced a key challenge: the number of platooning runs captured

during 2 h of data collection was extremely limited. This was because of using only three vehicles and a considerable turnaround time. Instead, this study follows the platoon 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 platoon are equipped with a GPS + IMU. The collected information forms the ground truth when extracting data from aerial videos and is utilized in the data-cleaning process. The remaining two datasets focus on collecting data from Level 2 automated vehicles. The only difference between the Level 1 and Level 2 datasets is that two Level 2 vehicles are followed throughout the 3-mi segment. All these datasets were collected via a RED camera at 30 fps at 8K resolution. The data collection altitude was set to 1000 ft. One dataset (Level 2) was collected on a cloudy day and the other two datasets (one Level 1 and one Level 2) were collected on a sunny day.

I-395 in Washington, D.C

One dataset was collected from Level 2 automated vehicles on I-395 (see Figure 1) using the infrastructure-based videography approach. The setup of the experiment is very similar to the first dataset collected on I-90/I-94. Level 2 vehicles were circulating and moving through the segment of interest. All the data were collected using multiple cameras at 1080P resolution and 30 fps. The dataset was collected during the morning peak hour and aims to capture the congestion in that area.

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

Data was collected in Foggy Bottom, D.C., using the infrastructure-based videography approach. The installation covered four intersections within a single block, utilizing five buildings. A total of 12 cameras were employed, with varying degrees of overlapping and non-overlapping coverage. The cameras utilized in this setup captured footage at a resolution of 1080P and 30 fps. The data was collected during the evening peak hour on a sunny day. 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 autonomous vehicle being tested is circulating mainly in a counterclockwise direction around the block, actively engaging with other vehicles, pedestrians, signal control systems, and stop signs.

Methodology

Vehicle trajectory data can be extracted from the video frames recorded in the bird's-eye view (via the helicopter) or in a side view considering the perspective (via infrastructure-mounted cameras) (see Figure 2*a*). This study utilizes the methodology proposed by Khajeh-Hosseini et al. (*17*) as the core of their trajectory extraction method. There are, however, several key modifications to their approach 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. (*17*). The trajectory extraction process consists of six steps: (1) pre-processing, (2) vehicle detection, (3) vehicle tracking, (4) image stabilization, (5) trajectory construction, and (6) data cleaning. Figure 2 shows these steps for a sample image. As indicated in this figure, after the pre-processing step, 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.

Pre-Processing

The first step in the extraction process is pre-processing the collected videos. Two steps should be taken in the pre-processing.

***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). 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 can be very challenging for the moving aerial videography and infrastructure-based videography approaches. To facilitate the process for all of the approaches, this study developed a consistent methodology to convert each image into a fixed coordinate system (i.e., reference image). Accordingly, this study utilizes satellite images to create the reference image. The reference images for Chicago datasets were created by first carefully selecting a representative set of

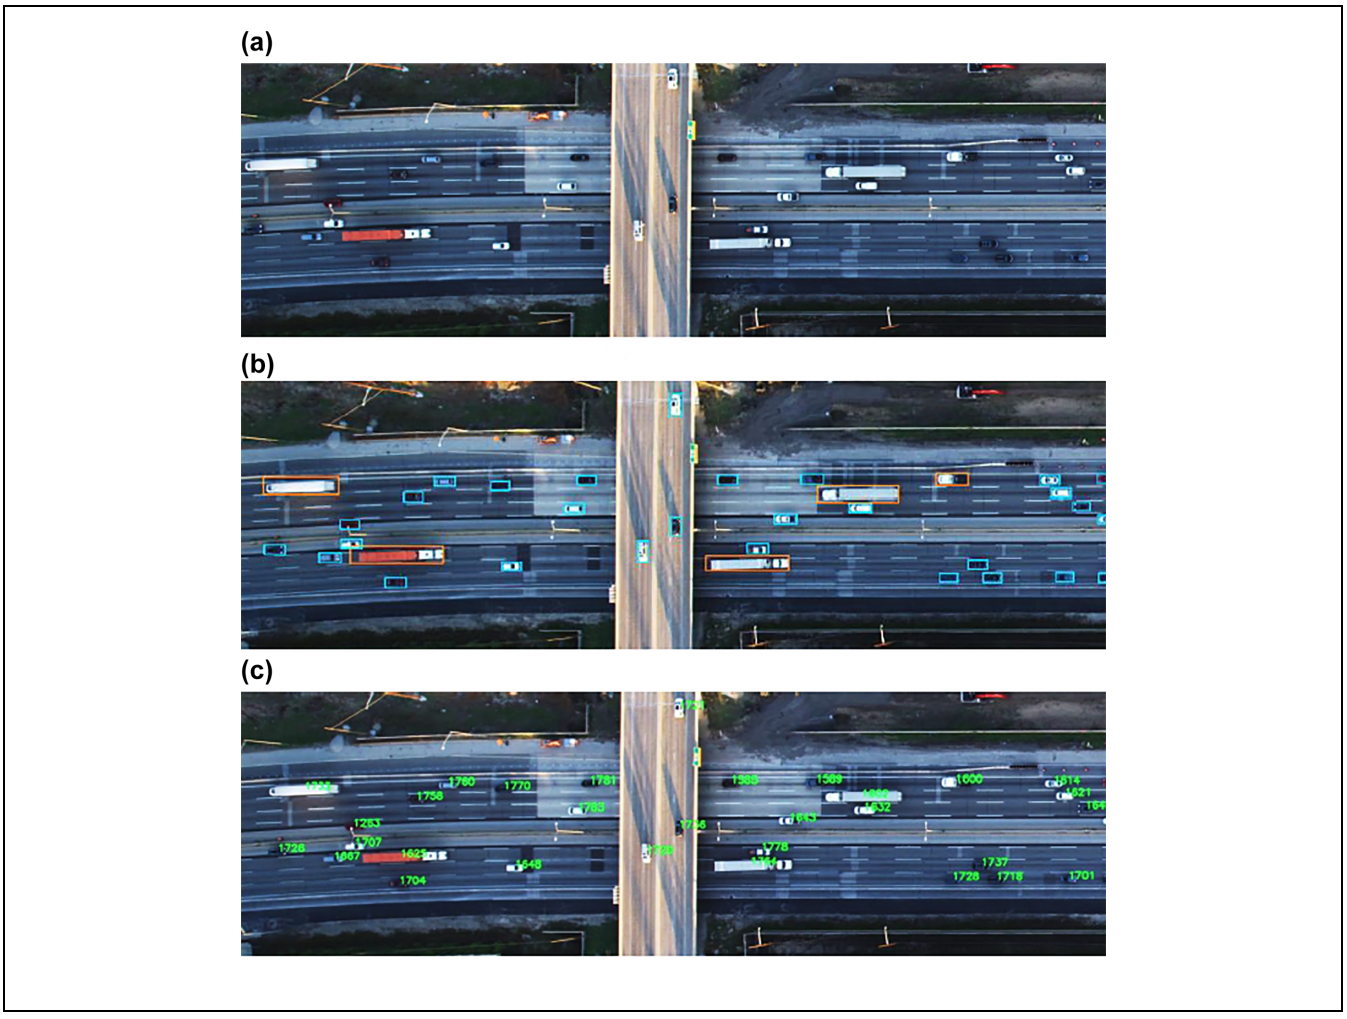

Figure 2. Vehicle detection and tracking in aerial images: (*a*) birds-eye view, (*b*) vehicle detection, and (*c*) vehicle tracking.

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 the satellite image manually as well as automatically using distinct features on the images utilizing Photoshop software. 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. The process of generating the reference image was, however, different in Washington, D.C. Semi-automatic camera calibration was conducted based on perspective-n-point (PnP) (*18*), since cameras are fixed, and angles are close to the horizontal angle with respect to the ground. PnP originates from camera calibration, in which using a set of n three-dimensional (3D) points in the world reference frame and their corresponding two-dimensional (2D) image projections, as well as the calibrated intrinsic camera parameters, determine the six degree 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 world. Based on this information, the homography matrix can be computed and 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 because of the recent advancements in deep neural networks (DNNs). Most of the existing approaches can be used to identify and locate the vehicles in the aerial images with a small additional training process. There are multiple popular convolutional neural network

(CNN)-based object detectors, such as R-CNN (*19*), RetinaNet (*20*), and YOLOv5 (*21*). 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 in the aerial images. This study utilizes RetinaNet in the fixed location aerial videography and moving aerial videography approaches and uses YOLOv5 in the infrastructure-based videography approach. Note that YOLOv5 is utilized for the infrastructure-based videography approach (instead of RetinaNet) because of its exceptional performance dealing with images that are not perfect bird's-eye views (caused by collecting data from overpasses).

Object Tracking

Tracking is the process of linking the new detections to previous observations. The tracking methodology proposed by Khajeh-Hosseini et al. (*17*) includes data association and track maintenance. Data association is associating the detected vehicles in the current image frame to the vehicles identified in the previous ones. This is a critical step considering that not all the transformation can be perfect and some has to be removed from the data. Accordingly, the process should detect vehicles with some gap in distance between them in two frames. The track maintenance is in charge of initiating new tracks, maintaining the tracks, and deleting them. The track maintenance initiates tracks 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 tracks using the data association. The tracks are updated as new detections are associated with them.

A new track is constructed for any new observation that is not associated with the current tracks. Moreover, if a track is not updated in multiple frames, the track maintenance deletes that track. A tracked object maintained by the track maintenance contains both the unique ID of the track 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 map to the fixed coordinate system picked on the ground. Consequently, it is essential to find the mapping function between the image coordinate to 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. The image stabilization is performed in three steps: the first is detecting the key features in both reference and input images; the second is finding the matching features between the two images; and the third is estimating the transformation between them. There exist different algorithms for good key features detection in images such as the Harris corner detector (*22*), scale-invariant feature transform (SIFT) (*23*), speeded up robust feature (SURF) (*24*), and oriented FAST and rotated BRIEF (ORB) (*25*). Following the recommendation by Khajeh-Hosseini et al. (*17*), this study adopts SIFT as the feature detecting 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 collection at a high frame rate. Accordingly, similar to Khajeh-Hosseini et al. (*17*), this study utilizes the fast library for approximate nearest neighbors (FLANN) matcher to match features between the images (*26*). 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. (*17*) realized that despite utilizing state-of-the-art algorithms, some of the feature matchings are incorrect. They used a combination of random sample consensus (RANSAC) (an algorithm to find the model parameters from a dataset with many outliers through an iterative process [*27*]) to estimate the homography transformation between two images considering the matched key features. To further improve the transformations (considering that the size of the reference image can be significantly larger than that of Khajeh-Hosseini et al. [*17*]), 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 transformation.

Trajectory Construction

The detection and tracking steps are based on bounding boxes that represent the vehicle's location in the 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 located by its row and column number in the image map. The pixel coordinate can be transformed into a cartesian coordinate system by taking the 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 is

the key to the mapping function between the two coordinate systems. Khajeh-Hosseini et al. (*17*) took the front bumper to indicate the vehicle's location, and the trajectory of the vehicle is the list of its location over space and time. Moreover, based on the recommendation of Khajeh-Hosseini et al. (*17*), a Kalman filter was applied to reduce the noise in the vehicle's location estimates (caused by noisy bounding boxes from image stabilization and vehicle detection and tracking processes).

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. It will be similar to the previous procedure to collect video data for each camera and develop trajectories for each segment. However, semi-automatic camera calibration will be conducted based on PnP (*18*), 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, because of multiple cameras recording individually, the identity of the subject vehicle may vary in this case since the vehicle may switch from segment to segment during the recording. Thus, the development of an extended trajectory is always characterized as a 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 below the overpasses. In addition, to have continuity over the entire road space, 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 is applied to eliminate the trajectory dataset's noise. This stitching process is used to develop extended trajectory data for the entire 0.34 mi.

Data Cleaning

Although all the previous steps are automated 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 overpasses, etc.). Accordingly, every single trajectory in the dataset should be inspected and verified. This process was conducted through the following steps.

Problematic trajectories, such as trajectories that have zero speed values, trajectory lengths below a minimum threshold, and trajectories that traverse the opposite direction of traffic flow, are first identified using several filters. 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. Finally, the erroneous trajectories and corresponding IDs are removed from the trajectory datasets.

Data Overview

The presented approach was applied to the collected data in all three locations. Figure 3 illustrates sample trajectory data from I-294 near Hinsdale, IL. In this particular snapshot, despite the overall trend in the dataset, there are only minor shockwaves. However, the figure indicates that the level of congestion decreases significantly from Lane 1 to Lane 4. In fact, a significant portion of vehicles in Lanes 1 and 2 are heavy vehicles, resulting in slower-moving traffic (this can be verified by comparing the slopes of the trajectories). Figure 4 presents the acceleration and speed profiles of the three vehicles in the ACC platoon as well as the platoon's immediate followers. The patterns in the acceleration and speed profiles are consistent across the three platoon members. The platoon shows signs of string stability, as any disturbance in the first vehicles loses its magnitude as it propagates upstream toward the second and third ACC-driven vehicles. Unlike the ACC platoon, the immediate follower did not follow the speed profile of the last ACC-driven vehicle. In fact, the follower changed frequently (different colors show different followers in Figure 4, *g* and *h*). This was mainly because the ACC platoon was traveling more slowly 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 with an acceptable level of fluctuation. A similar analysis has been conducted for the entire dataset.

Figure 5 shows the overall speed profile of the vehicles in the helicopter view (all lanes) as the ACC-driven vehicles (and the helicopter) move along the highway. Two patterns are visible in this figure. The speed of the vehicle moving forward is fairly constant (free-flow traffic regime). This is mainly because the vehicles are moving in a lane with a free-flow traffic regime. There are, however, multiple shockwaves that can be seen (color changes in the figure). These smaller-scale changes in average speed are mainly because of 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% of the trajectories

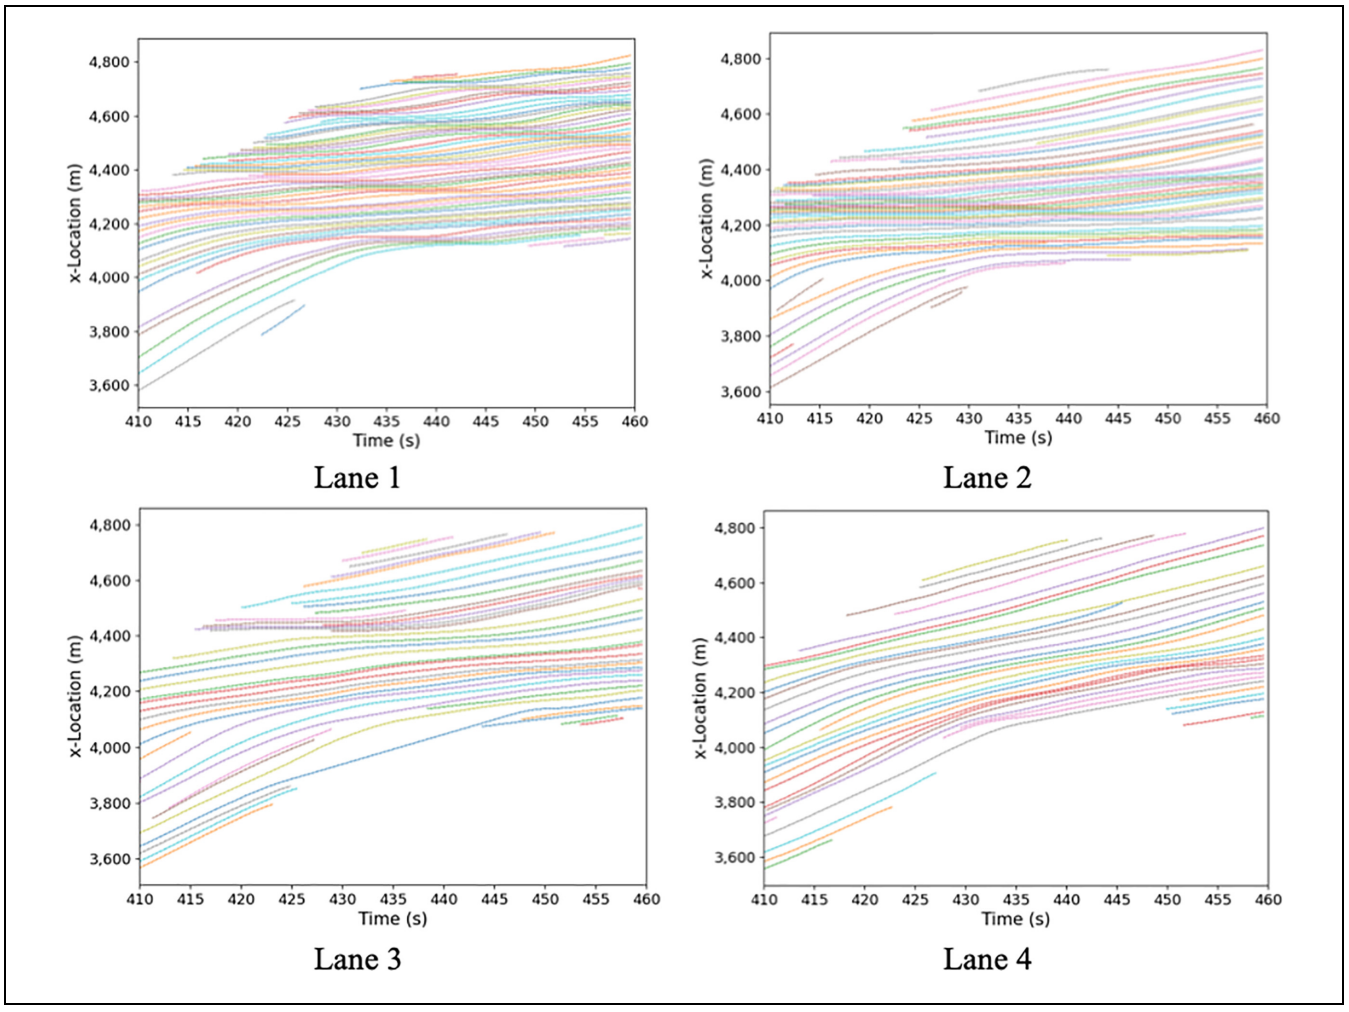

Figure 3. Sample trajectory data from a single run on I-294 near Hinsdale, IL.

belonging to heavy vehicles (depending on the run and time of the data collection). The unique structure of the data collection segment results in almost all the heavy vehicle trajectories being located in the two rightmost lanes and no heavy vehicle trajectory 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 6 presents sample trajectory data from I-90/I-94 in Chicago, IL, collected using the fixed location aerial videography approach. The number of trajectories captured is significantly higher in this case (compared with the moving aerial videography approach) because of many vehicles entering and leaving the segment in a short time window. The number of heavy vehicles, however, is significantly lower compared with the I-294 data (I-294 is considered a major truck route and heavy vehicles mostly avoid the I-90/I-94 congestion by using I-294). Figure 6 shows how the level of congestion decreases as we move from the three leftmost lanes to the three rightmost lanes. Stop-and-go and slow-and-go traffic can be observed in the leftmost lanes, while the traffic moves at almost free-flow in the rightmost lanes. 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 (which was never captured in any other trajectory dataset). Note that in this figure ACC-driven vehicles are in the second and third lane from the right and mostly traveling at free-flow. They, however, get involved with the complex lane-changing maneuvers in the dataset collected using the moving aerial videography approach. Figure 7 illustrates the overall speed profile through the I-90/I-94 segment. Multiple major shockwaves can be observed in this figure, resulting in stop-and-go waves with speeds close to 0 m/s.

Figure 8 presents sample trajectory data in I-395 in Washington, D.C. This section of I-395 is significantly congested and offers interesting dynamics between the

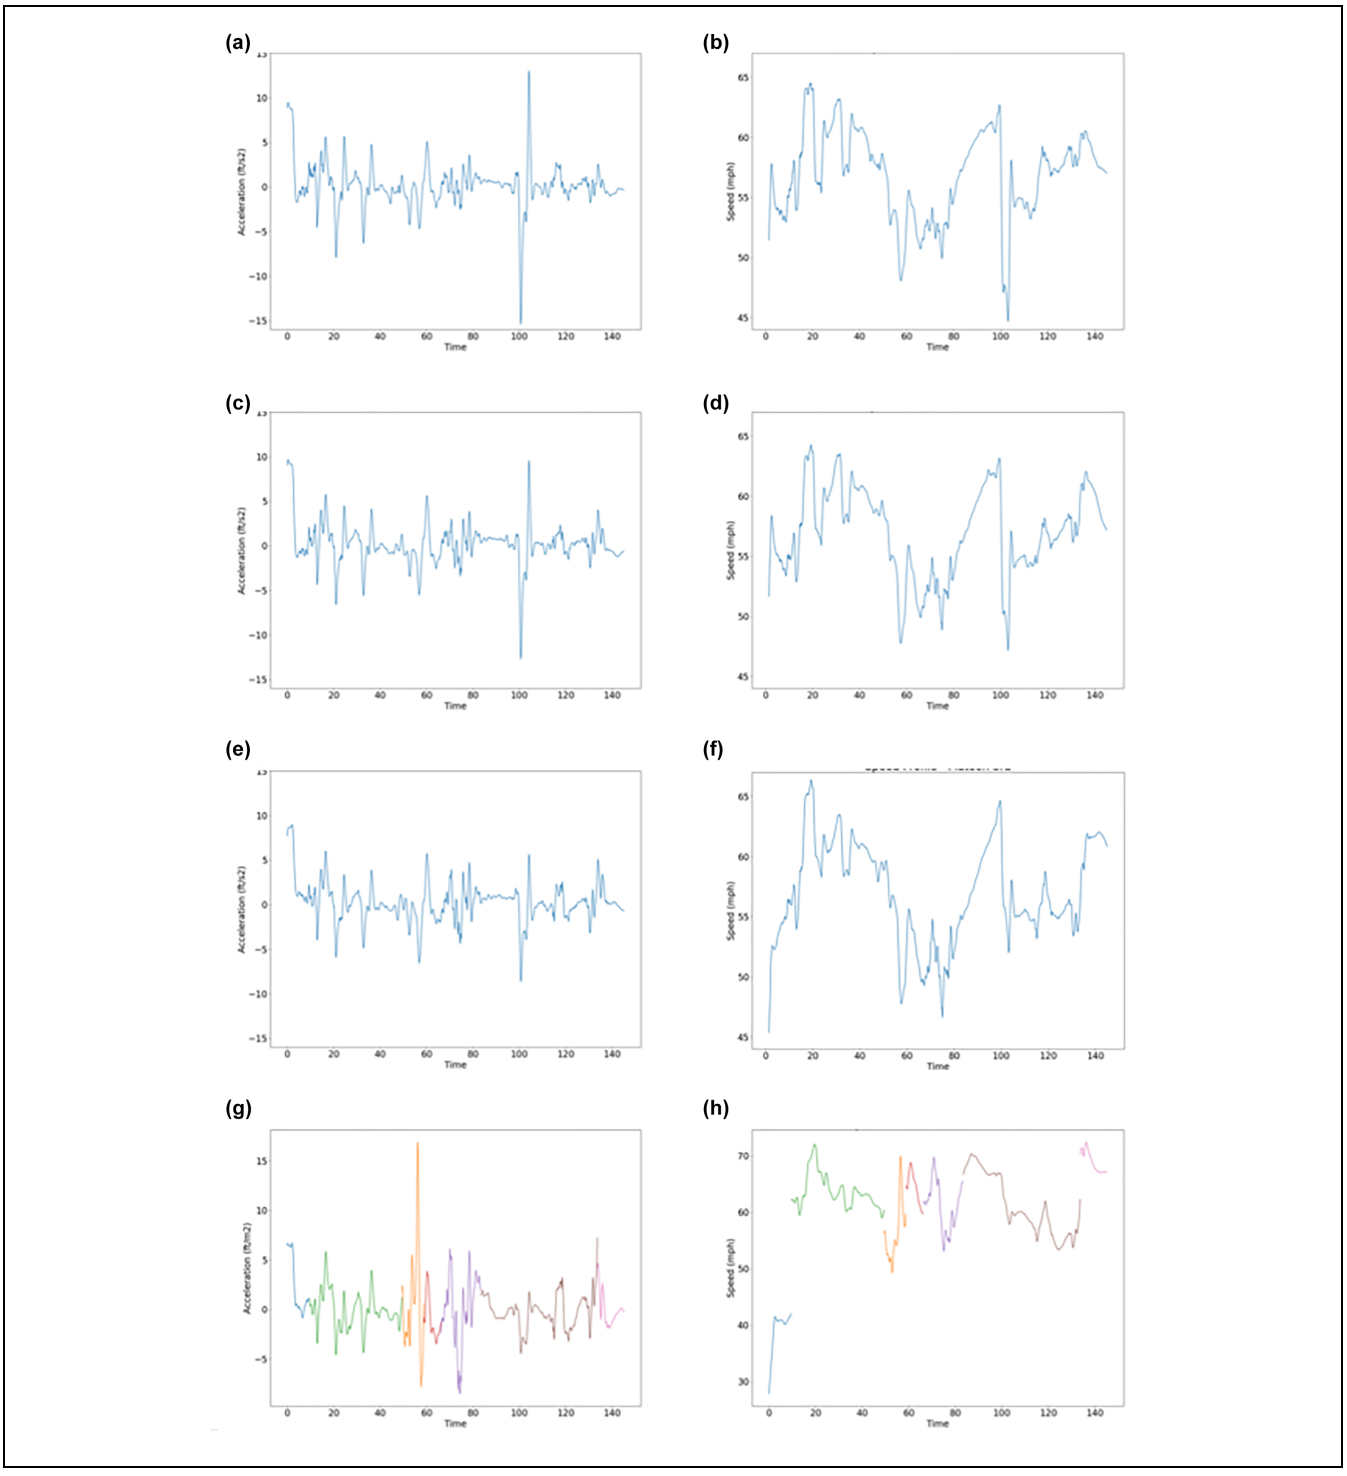

Figure 4. Acceleration and speed profiles of adaptive cruise control (ACC)-driven vehicles in the platoon and the platoon followers: (*a*) acceleration profile: first ACC-driven vehicle, (*b*) speed profile: first ACC-driven vehicle, (*c*) acceleration profile: second ACC-driven vehicle, (*d*) speed profile: second ACC-driven vehicle, (*e*) acceleration profile: third ACC-driven vehicle, (*f*) speed profile: third ACC-driven vehicle, (*g*) acceleration profile: platoon followers, and (*h*) speed profile: platoon followers.
Note: Color online only.

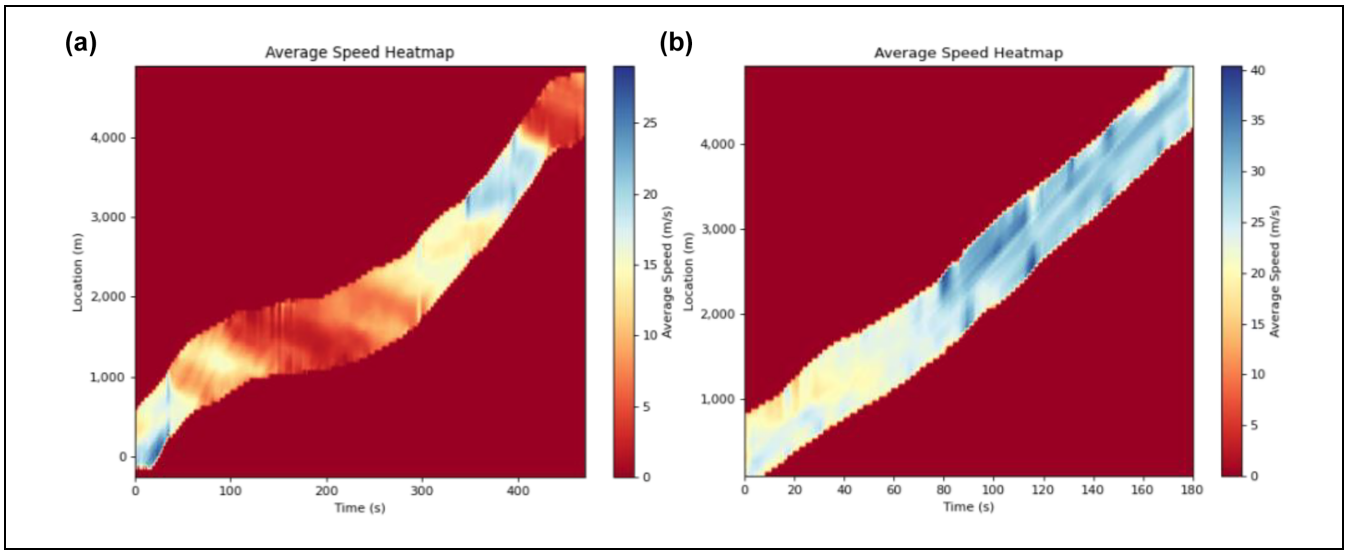

Figure 5. Average speed on I-294 near Hinsdale, IL, for (*a*) a congested sample run and (*b*) an uncongested sample run. The helicopter moves with the adaptive cruise control-driven vehicles in both runs.

on-ramp and the main segments. The segment also contains a significant percentage of heavy vehicles, ranging between 4% and 8% (the largest percentage of heavy vehicles was observed on the leftmost lane, in contrast to I-294). Figure 8 shows the clear shockwaves and a very interesting dynamic of speeding up and slowing down as vehicles approach the merging area. A significant number of mergers are occurring in the leftmost lane. Figure 9 illustrates the overall speed profile through the I-395 segment. A considerable number of shockwaves can be observed in this figure.

Figure 10 presents sample trajectory data collected at George Washington University Campus in Washington, D.C. To better present the data, the trajectories of vulnerable road users (e.g., pedestrians, cyclists, and scooter riders) and vehicles (e.g., passenger cars, motorcycles, buses, and trucks) are presented separately (while in the data, they overlap throughout time and space). A considerable number of interactions among multiple modes of transportation can be observed in the data across the entire 3 h of data collection and at all four intersections.

Lessons Learned

The data collection process has gone through multiple iterations and the authors 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 from automated vehicles.

One of the main considerations for data collection is the location. Clearly, it is critical to identify locations that offer unique insight into traffic flow dynamics. In this study, the I-294 location was selected as a representation of a typical highway environment. The segment was selected to offer various traffic flow regimes (i.e., free-flow, slow-and-go, stop-and-go, and fully congested). The length of the segment is also suitable for forming and maintaining platoons of ACC-driven vehicles. The southbound also has a major on-ramp that results in interesting mandatory lane-changing instances. Moreover, I-294 has a very high heavy vehicle percentage and offers unique insight into CAV–truck interactions. The I-90/I-94 location was selected because of the existence of a major weaving section with a significant number of mandatory lane-changing instances, including last-minute forced lane-changing maneuvers. The ability of automated lane-changing of Level 2 CAVs was tested in this location. Finally, the I-395 location was selected because of its location and its importance to the Washington, D.C., transportation network. It offers a diverse traffic flow regime and mandatory lane-changing instances from an on-ramp located on the left-hand side of the highway. This section also offers considerable heavy vehicle traffic that creates interesting CAV–truck interactions. Other considerations also play an important role and can determine the most efficient data collection approach (e.g., we utilized the infrastructure-based videography approach in Washington, D.C., since flying a helicopter in Washington, D.C., was challenging). Such a consideration also limited the location of the data collection and reduced the segment length (compared with the original data collection plan). The day of the data collection also plays an important role and various daily and seasonal traffic variations should be considered. For instance, we delayed the data collection on I-294 until

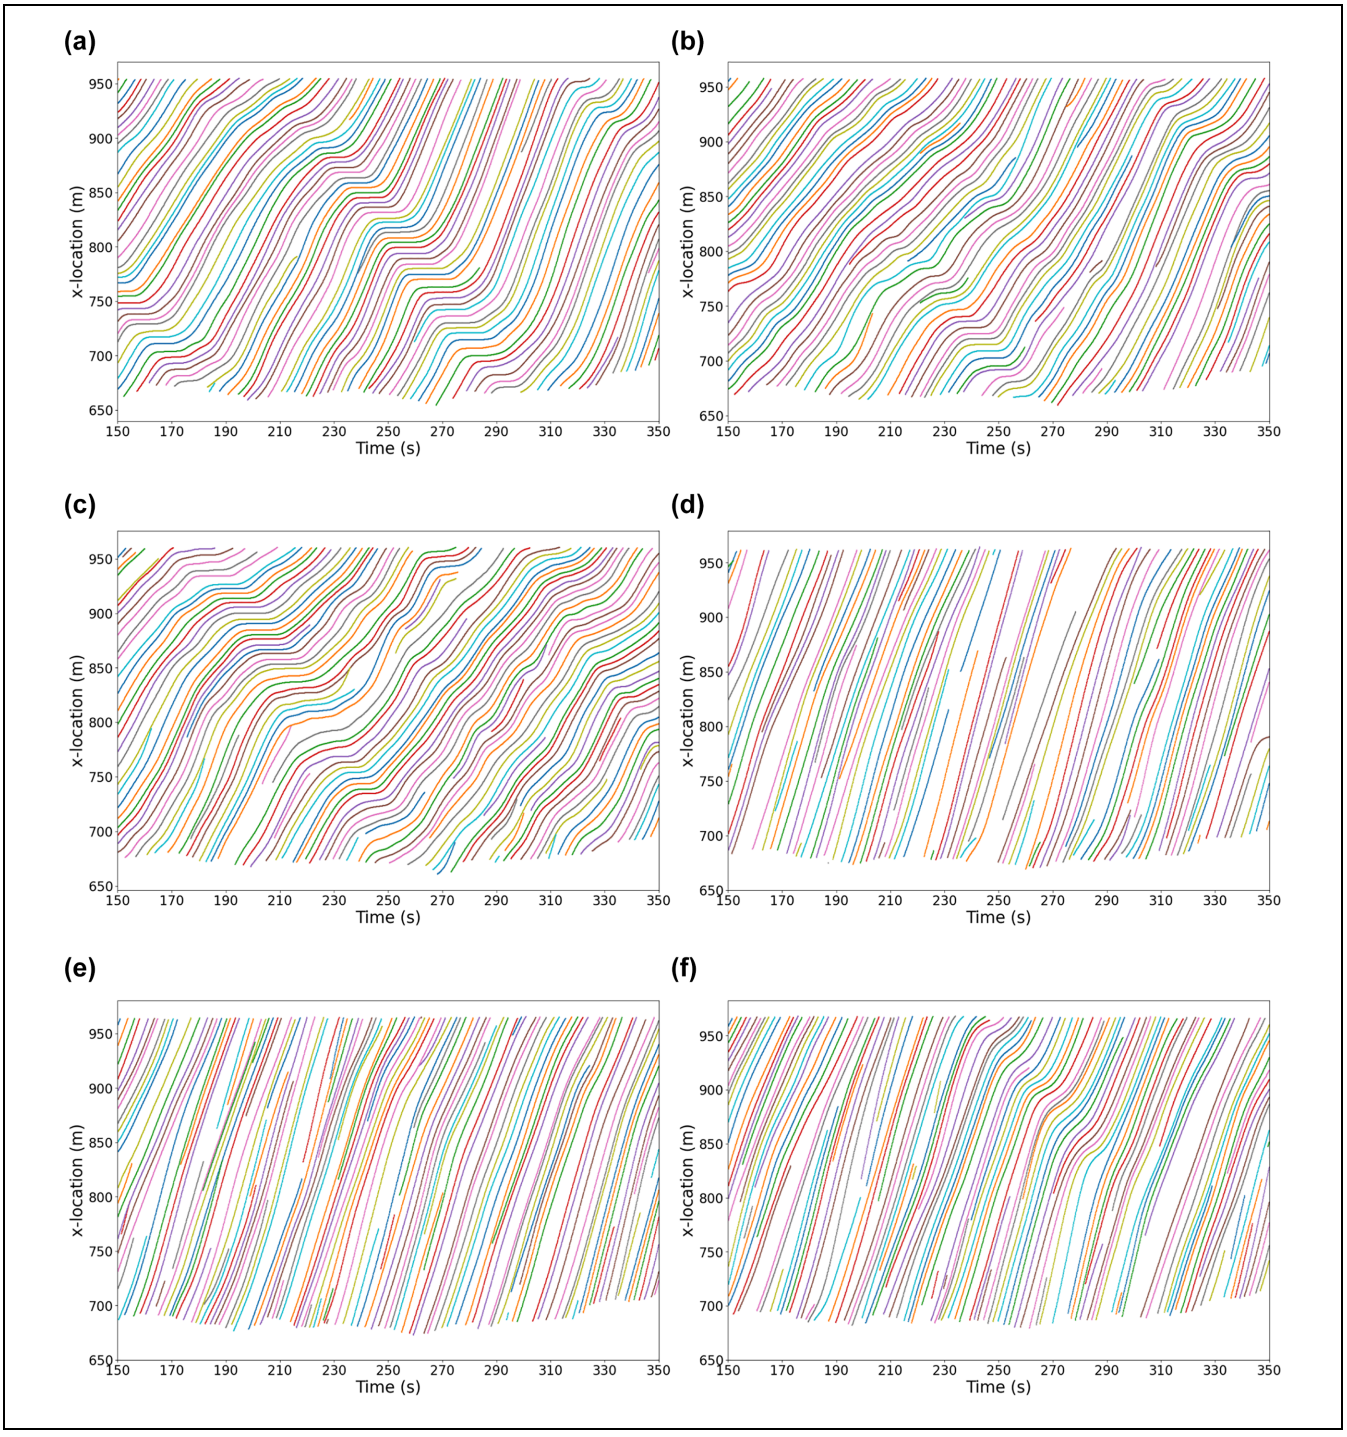

Figure 6. Sample trajectory data on I-90/I-94 in Chicago, IL: (*a*) Lane 10, (*b*) Lane 11, (*c*) Lane 12, (*d*) Lane 13, (e) Lane 14, and (*f*) Lane 15.
Note: Lane numbering starts from the left with Lane 10 as the leftmost lane and increases to Lane 15 as the rightmost lane.

the COVID-19 pandemic restrictions were lifted and traffic got back to normal. While all the aforementioned factors are important and can influence what data can represent, the time of day directly affects the quality of the extracted trajectories. Collecting data during morning and evening peak hours results in the most congestion, but the data extraction process should deal with long shadows (which lowers the quality of vehicle detection). Note that any data collection should be performed under safe weather and road surface conditions. However, interestingly enough, because of the lack of shadows, cloudy days were found to be the best

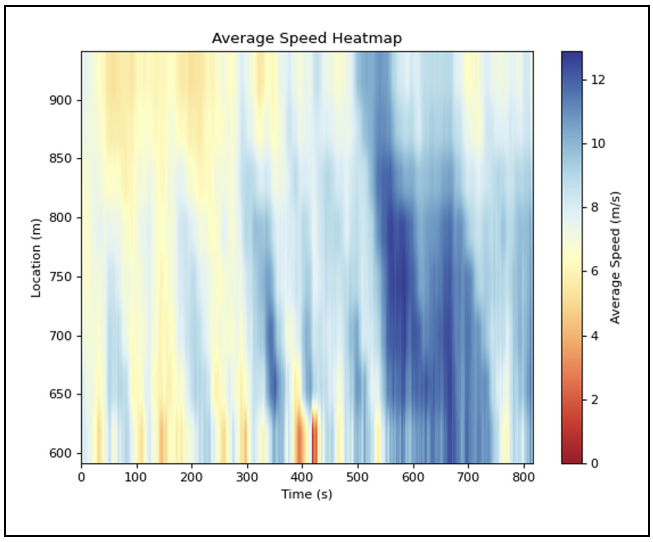

Figure 7. Average speed on I-90/I-94 in Chicago, IL.

environment for data collection (when the cloud ceiling is high enough).

The camera and helicopter setup also plays a critical role 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 1000 ft. Using a lens with no distortion, the camera could capture about 0.5 mi of the highway at that altitude. Flying at a higher altitude, however, significantly reduces the vehicle detection accuracy. 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 that resolution.

When analyzing the data, accurate and precise detections of vehicles play a key role, especially for heavy

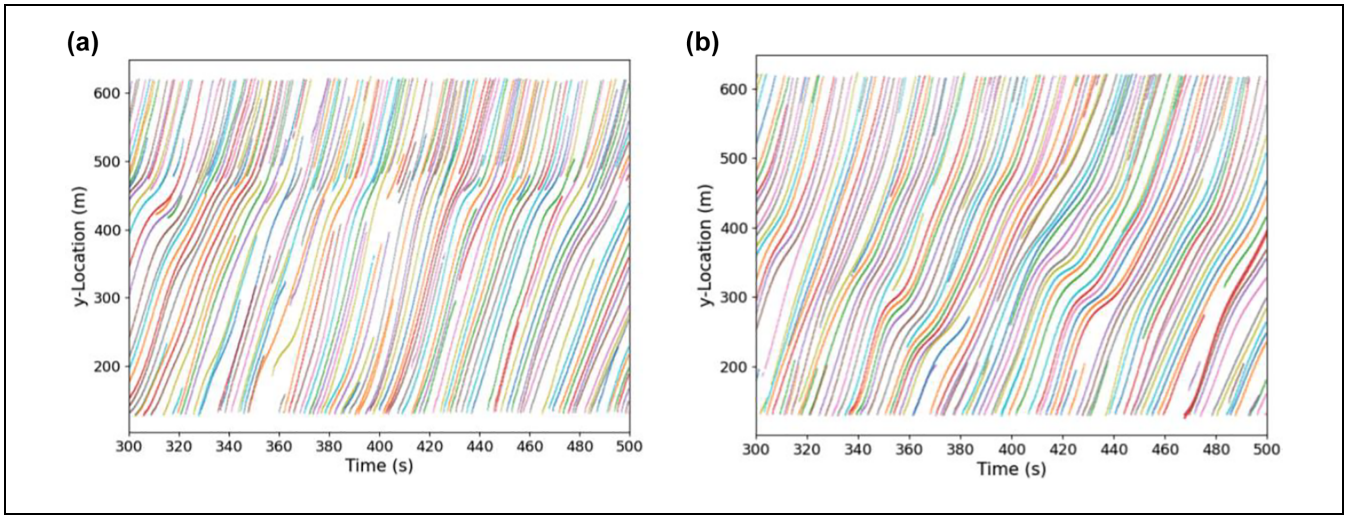

Figure 8. Sample trajectory data from I-395 in Washington, D.C.: (*a*) leftmost lane and (*b*) rightmost lane.

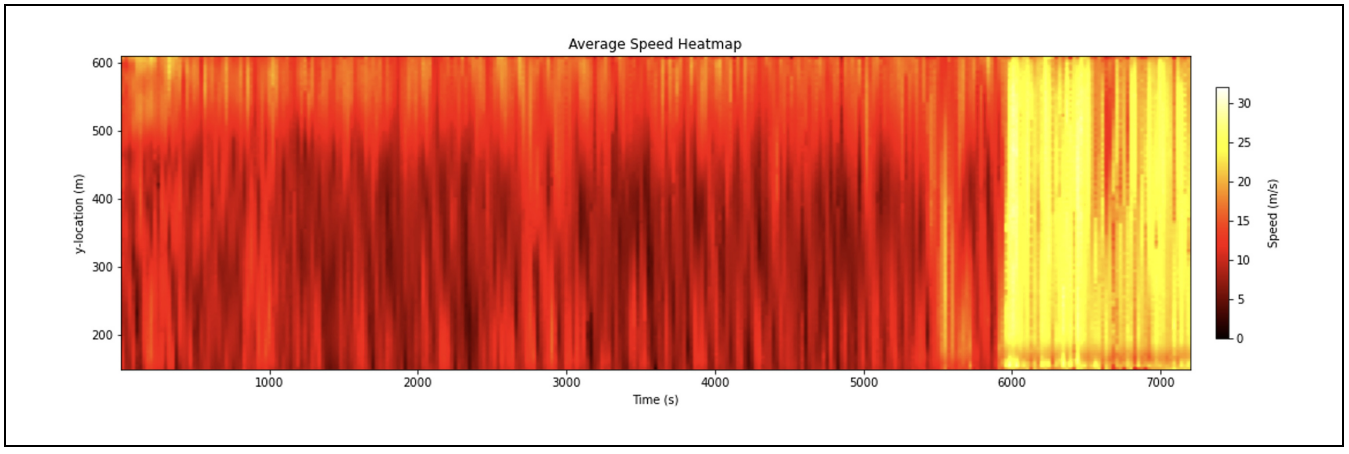

Figure 9. Average speed for the middle lane on I-395 in Washington, D.C.

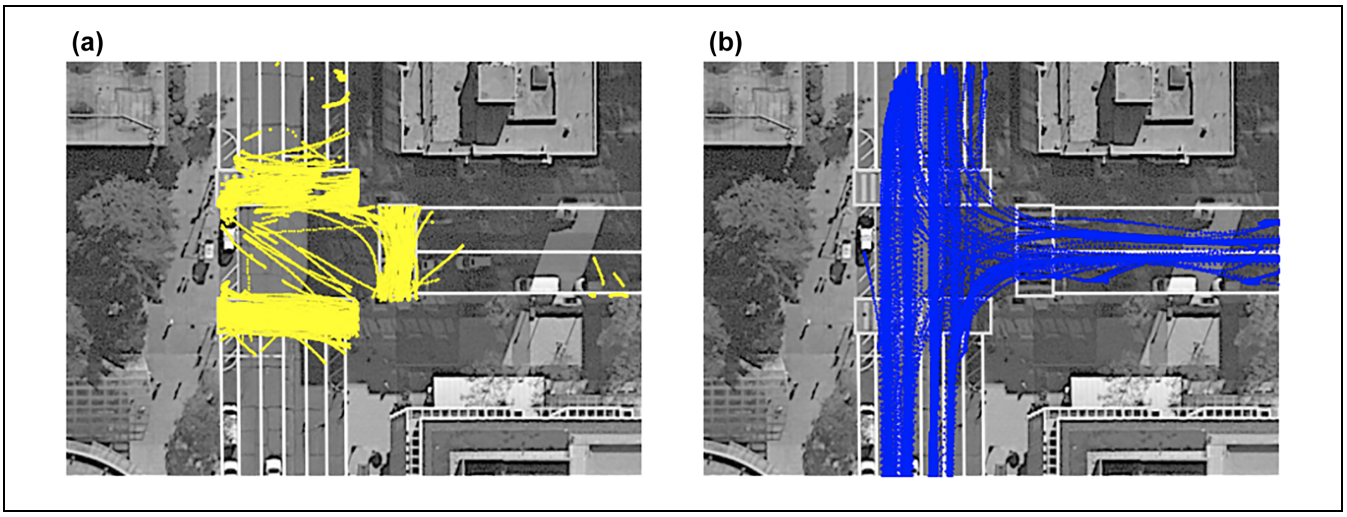

Figure 10. Sample trajectory data from George Washington University Campus in Washington, D.C.: (*a*) vulnerable road users and (*b*) vehicles.

vehicles, and would save many challenges in the other data extraction steps. In particular, we recommend using "polygons" for truck detection rather than rectangles since rectangles become very inaccurate, especially around curved road segments (because of the length of heavy vehicles). Moreover, it is recommended to train the detection algorithm using a small sample of the data that was collected to improve detection accuracy. Accounting for distinct objects and features in the study area before beginning data collection and processing could also save a great deal of time and improve the tracking process (e.g., accounting for overpass locations can help remove inaccurate detections—vehicles on the overpasses—and avoid gaps in vehicle trajectories). Finally, visualizing and inspecting trajectory data after extraction is an essential step to catch and fix errors.

It is important to note that every study is different and requires specific considerations that are specific to the study area, method of data collection, type of data, required outcomes, and deliverables. Because of the study requirements and needed accuracy, the trajectory extraction process should be often repeated until reaching satisfactory results.

Conclusion

This study introduces a set of accurate trajectory datasets focusing on characterizing human–automated vehicle interactions under a diverse set of scenarios in diverse highway environments. This study also introduces a detailed introduction to the methodologies utilized for converting video frames to vehicle trajectory data under three distinct data collection methods: (1) fixed location aerial videography (a helicopter hovers over a segment of interest); (2) moving aerial videography (a helicopter follows the automated vehicles as they move in a much longer segment than in the first method), and (3) infrastructure-based videography (multiple overlapping cameras located on overpasses create a comprehensive image of the highway). Utilizing the fixed location aerial videography approach, a dataset was collected on I-90/I-94 in Chicago, IL, from Level 2 vehicles, and utilizing the moving aerial videography approach, another dataset was collected on I-90/I-94 from Level 2 vehicles. The two datasets were collected on I-294 near Hinsdale, IL, from Level 1 and Level 2 vehicles. One dataset was collected on I-395 in Washington, D.C., from Level 2 vehicles using the infrastructure-based videography approach. Finally, one dataset was collected near the George Washington University Campus in Washington, D.C., from Level 3 vehicles using the infrastructure-based videography approach. This paper also presents an overview of the datasets and provides some basic statistics about each collected dataset. Finally, a discussion on the lessons learned throughout the data collection efforts in different locations is presented.

Author Contributions

The authors confirm contribution to the paper as follows: study conception and design: R. Ammourah, P. Beigi, B. Fan, C.-C. Hsiao, S.H. Hamdar, J. Hourdos, R. James, M. Khajeh-Hosseini, H.S. Mahmassani, D. Monzer, T. Radvand, A. Talebpour, M. Yousefi, Y. Zhang; data collection: R. Ammourah, P. Beigi, B. Fan, C.-C. Hsiao, S.H. Hamdar, M. Khajeh-Hosseini, H.S. Mahmassani, D. Monzer, T. Radvand, A. Talebpour, M. Yousefi, Y. Zhang; analysis and interpretation of results: R. Ammourah, P. Beigi, B. Fan, C.-C. Hsiao, S.H. Hamdar, J. Hourdos, R. James, M. Khajeh-Hosseini, H.S. Mahmassani, D. Monzer, T. Radvand, A. Talebpour, M. Yousefi, Y. Zhang; draft manuscript preparation: R.

Ammourah, P. Beigi, B. Fan, C.-C. Hsiao, S.H. Hamdar, J. Hourdos, R. James, M. Khajeh-Hosseini, H.S. Mahmassani, D. Monzer, T. Radvand, A. Talebpour, M. Yousefi, Y. Zhang. All authors reviewed the results and approved the final version of the manuscript.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This study was supported by the U.S. Department of Transportation (U.S. DOT) Federal Highway Administration (FHWA): Third Generation Simulation (TGSIM) Data: A Closer Look at The Impacts of Automated Driving Systems on Human Behavior.

ORCID iDs

Rami Ammourah https://orcid.org/0000-0003-0912-8265
Pedram Beigi https://orcid.org/0000-0002-3020-4085
Chun-Chien Hsiao https://orcid.org/0009-0006-9402-176X
Samer H. Hamdar https://orcid.org/0000-0001-6896-367X
John Hourdos https://orcid.org/0000-0001-9888-0445
Rachel James https://orcid.org/0000-0001-9138-510X
Mohammdreza Khajeh-Hosseini https://orcid.org/0000-0002-1657-9198
Hani S. Mahmassani https://orcid.org/0000-0002-8443-8928
Dana Monzer https://orcid.org/0000-0002-4688-6359
Tina Radvand https://orcid.org/0000-0002-4454-8567
Alireza Talebpour https://orcid.org/0000-0002-5412-5592
Yanlin Zhang https://orcid.org/0000-0002-5342-2755

References

1. 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: Journal of the Transportation Research Board*, 2019. 2673: 367–379.
2. Deo, N., and M. M. Trivedi. Convolutional Social Pooling for Vehicle Trajectory Prediction. *Proc., IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE Xplore, 2018, pp. 1468–1476.
3. Sadigh, D., S. Sastry, S. A. Seshia, and A. D. Dragan. Planning for Autonomous Cars That Leverage Effects on Human Actions. *Robotics: Science and Systems*, Vol. 2, 2016, pp. 1–9.
4. 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.
5. Krajewski, R., J. Bock, L. Kloeker, and L. Eckstein. The HighD Dataset: A Drone Dataset of Naturalistic Vehicle Trajectories on German Highways for Validation of Highly Automated Driving Systems. *Proc., 2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, Maui, HI, 2018.
6. 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.
7. 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.
8. Sun, P., H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, et al. Scalability in Perception for Autonomous Driving: Waymo Open Dataset. *Proc., IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE Xplore, 2020, pp. 2446–2454.
9. Agarwal, S., A. Vora, G. Pandey, W. Williams, H. Kourous, and J. McBride. Ford Multi-AV Seasonal Dataset. *The International Journal of Robotics Research*, Vol. 39, No. 12, 2020, pp. 1367–1376.
10. Caesar, H., V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom. nuScenes: A Multimodal Dataset for Autonomous Driving. *Proc., IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE Xplore, 2020, pp. 11621–11631.
11. Yu, F., H. Chen, X. Wang, W. Xian, Y. Chen, F. Liu, V. Madhavan, and T. Darrell. Bdd100k: A Diverse Driving Dataset for Heterogeneous Multitask Learning. *Proc., IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE Xplore, 2020, pp. 2636–2645.
12. Huang, X., X. Cheng, Q. Geng, B. Cao, D. Zhou, P. Wang, Y. Lin, and R. Yang. The Apolloscape Dataset for Autonomous Driving. *Proc., IEEE Conference on Computer Vision and Pattern Recognition Workshops*, IEEE Xplore, 2018, pp. 954–960.
13. Kesten, R., M. Usman, J. Houston, T. Pandya, K. Nadhamuni, A. Ferreira, M. Yuan, et al. *Lyft Level 5 Av Dataset 2019*. Vol. 1, 2019, p. 3. https://level5.lyft.com/dataset.
14. 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: Journal of the Transportation Research Board*, 2020. 2674: 471–484.
15. 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.
16. Yin, H., and C. Berger. When to Use What Data Set for Your Self-Driving Car Algorithm: An Overview of Publicly Available Driving Datasets. *Proc., 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, Yokohama, Japan, 2017.

17. 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*, Vol. 208, 2022, p. 118060.
18. Tang, Z., G. Wang, H. Xiao, A. Zheng, and J.-N. Hwang. Single-Camera and Inter-Camera Vehicle Tracking and 3D Speed Estimation Based on Fusion of Visual and Semantic Features. *Proc., IEEE Conference on Computer Vision and Pattern Recognition Workshops*, IEEE Xplore, 2018, pp. 108–115.
19. Girshick, R., J. Donahue, T. Darrell, and J. Malik. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. *Proc., IEEE Conference on Computer Vision and Pattern Recognition*, IEEE Xplore, 2014, pp. 580–587.
20. Lin, T.-Y., P. Goyal, R. Girshick, K. He, and P. Dollár. Focal Loss for Dense Object Detection. *Proc., IEEE International Conference on Computer Vision*, IEEE Xplore, 2017, pp. 2980–2988.
21. Redmon, J., S. Divvala, R. Girshick, and A. Farhadi. You Only Look Once: Unified, Real-Time Object Detection. *Proc., IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, New York, 2016, pp. 779–788.
22. Derpanis, K. G. *The Harris Corner Detector*, Vol. 2. York University, Toronto, ON, Canada, 2004, pp. 1–2.
23. Lowe, D. G. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision*, Vol. 60, No. 2, 2004, pp. 91–110.
24. Bay, H., T. Tuytelaars, and L. Van Gool. Surf: Speeded up Robust Features. *Proc., Computer Vision–ECCV 2006: 9th European Conference on Computer Vision*, Graz, Austria, Springer, Berlin, Heidelberg, May 7–13, 2006, pp. 404–417.
25. Rublee, E., V. Rabaud, K. Konolige, and G. Bradski. ORB: An Efficient Alternative to SIFT or SURF. *Proc., 2011 International Conference on Computer Vision*, Barcelona, Spain, 2011.
26. Muja, M., and D. Lowe. Fast Library for Approximate Nearest Neighbors (FLANN). 2013. git://github. com/mariusmuja/flann.git. http://www.cs.ubc.ca/research/flann.
27. Derpanis, K. G. Overview of the RANSAC Algorithm. *Image Rochester NY*, Vol. 4, No. 1, 2010, pp. 2–3.