Characterizing Level 2 Automation in a Naturalistic Driving Fleet

January 2024 | Final Report

SAFE-D
SAFETY THROUGH DISRUPTION

VIRGINIA TECH
TRANSPORTATION INSTITUTE
VIRGINIA TECH

Texas A&M Transportation Institute

San Diego State University
Leadership Starts Here

Disclaimer

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No. VTTI-00-024	2. Government Accession No.	3. Recipient's Catalog No.

4. Title and Subtitle Characterizing Level 2 Automation in a Naturalistic Driving Fleet	5. Report Date January 2024 6. Performing Organization Code:
7. Author(s) Miguel A. Perez, Paolo Terranova, Mariette Metrey, Haden Bragg, and Nicholas Britten	8. Performing Organization Report No.
9. Performing Organization Name and Address: Safe-D National UTC Virginia Tech Transportation Institute 3500 Transportation Research Plaza Blacksburg, VA 24061	10. Work Unit No. 11. Contract or Grant No. 69A3551747115/Project VTTI-00-024
12. Sponsoring Agency Name and Address Office of the Secretary of Transportation (OST) U.S. Department of Transportation (US DOT)	13. Type of Report and Period Final Research Report 08/2019-01/2024 14. Sponsoring Agency Code

15. Supplementary Notes

This project was funded by the Safety through Disruption (Safe-D) National University Transportation Center, a grant from the U.S. Department of Transportation – Office of the Assistant Secretary for Research and Technology, University Transportation Centers Program.

16. Abstract

Introducing automation into the vehicle fleet disrupts how vehicles operate and potentially affects what drivers do with these features and expect from vehicle performance. Therefore, it is imperative to study driver adaptations in response to these innovations. This investigation leveraged 47 vehicles from the Virginia Tech Transportation Institute Level 2 (L2) Naturalistic Driving Study to analyze driver behavior with L2 automation features. Results showed no sizeable differences between periods of L2 feature usage and general driving periods with respect to time-of-day and calendar-related metrics. Most L2 feature usage occurred on motorways, following design expectations. L2 features were activated for 7.2 minutes in trips lasting an average of 22.8 minutes, or about 32% of the L2 trip duration. Driver-initiated overrides were predominantly done by braking or accelerating the vehicle, with steering-based overrides being minimal and likely involving lane changes without using a turn signal. Intervention requests were the most common takeover request, followed by requests due to insufficient driver hand contact with the steering wheel. Findings suggest that as L2 features penetrate the U.S. fleet in non-luxury consumer vehicles, system usage will be common and comparable with previous findings for luxury offerings. While evidence of potential system misuse was observed, future work may further operationalize system misuse and assess the prevalence of such behaviors.

17. Key Words Level 2 automation, naturalistic driving, driver behavior	18. Distribution Statement No restrictions. This document is available to the public through the Safe-D National UTC website, as well as the following repositories: VTechWorks, The National Transportation Library, The Transportation Library, Volpe National Transportation Systems Center, Federal Highway Administration Research Library, and the National Technical Reports Library.

19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 25	22. Price $0

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

Abstract

Introducing automation into the vehicle fleet disrupts how vehicles operate and potentially affects what drivers do with these features and expect from vehicle performance. Therefore, it is imperative to study driver adaptations in response to these innovations. This investigation leveraged 47 vehicles from the Virginia Tech Transportation Institute Level 2 (L2) Naturalistic Driving Study to analyze driver behavior with L2 automation features. Results showed no sizeable differences between periods of L2 feature usage and general driving periods with respect to time-of-day and calendar-related metrics. Most L2 feature usage occurred on motorways, following design expectations. L2 features were activated for 7.2 minutes in trips lasting an average of 22.8 minutes, or about 32% of the L2 trip duration. Driver-initiated overrides were predominantly done by braking or accelerating the vehicle, with steering-based overrides being minimal and likely involving lane changes without using a turn signal. Intervention requests were the most common takeover request, followed by requests due to insufficient driver hand contact with the steering wheel. Findings suggest that as L2 features penetrate the U.S. fleet in non-luxury consumer vehicles, system usage will be common and comparable with previous findings for luxury offerings. While evidence of potential system misuse was observed, future work may further operationalize system misuse and assess the prevalence of such behaviors.

Acknowledgements

The authors would like to acknowledge Martha Gizaw for her assistance with some video processing during the project, along with Charles Layman for his creation of video analysis tools and Jacobo Antona-Makoshi for his review of earlier drafts of this report. In addition, many thanks to all the VTTI personnel involved in the collection of the VTTI L2 NDS data that was leveraged in this work.

This project was funded by the Safety through Disruption (Safe-D) National University Transportation Center, a grant from the U.S. Department of Transportation – Office of the Assistant Secretary for Research and Technology, University Transportation Centers Program.

Table of Contents

List of Figures

List of Tables

Introduction

The introduction of automation features into the vehicle fleet is disrupting the way vehicles operate. Likewise, the introduction of more vehicles with automated features of increasing ability into the fleet can potentially affect what drivers do when the features are active and the expectations that drivers have related to vehicle functionality and characteristics. For these reasons, it is imperative to study driver adaptations in response to these innovations.

The study of vehicle and driver adaptations as SAE International (2021) Level 2 (L2) features are introduced into the fleet requires the collection of relevant data. The Virginia Tech Transportation Institute (VTTI) L2 Naturalistic Driving Study (NDS), which includes over 200 vehicles equipped with L2 automation features, was leveraged in this investigation to support analyses of driver behavior with these systems.

The work focused on isolating and characterizing the VTTI L2 NDS participants' use of adaptive cruise control (ACC) and lane keeping assistance systems (LKAS)/lane centering (LC) in tandem. At present, we have only limited knowledge of the circumstances under which L2 features are activated in the real world and the extent to which those circumstances are compatible with the operating envelopes envisioned by designers. Most previous work in this area, described in the next section, has focused on the resumption of control and prevalence of non-driving behaviors when L2 features are active or has focused on high-end vehicles that are not typical on U.S. roadways.

The primary research questions addressed by this project were:

- Driver usage
    - Under what environmental conditions are L2 automation features activated or deactivated?
    - On what roadway types are L2 automation features activated or deactivated?
    - How long are L2 features activated, and are there particular trip characteristics that modulate these activations?
- Takeover requests
    - What situations (e.g., environmental change, roadway change) lead to system takeover requests?
        - How frequently do these conditions occur?

Background

Advancements in automotive technology promise to transform driving and the role of human drivers by introducing Automated Driving Systems (ADS) to our roadways. The long-term goal of ADS is to replace human drivers' performance of all the driving subtasks across all environments or operational design domains (ODDs). While ADS technology has advanced considerably, no such system exists today. Instead, currently commercially available vehicles are equipped with Advanced Driver Assistance Systems (ADAS) capable of SAE L2 autonomy.

L2 automation, which is also called partial driving automation, performs the control-level driving, while the driver is responsible for supervising the ADAS features and attending to the roadway (SAE, 2021). When L2 autonomy features are engaged, the driver must be prepared to resume control of the system at all times with little or no warning. Additionally, L2 features have manufacturer-specific system performance requirements that only allow the features to be used within a pre-set speed range and/or ODD.

The monitoring role of the human driver during partially automated driving reveals one of the ironies of automation, whereby a human operator is asked to monitor the automated system to ensure it is working correctly despite the promise of the automated system to perform the task better than a human operator (Bainbridge, 1983). The shift from the role of an active controller of the vehicle to a monitoring role has potential safety implications. It is well-established that humans are ineffective at monitoring automated systems for extended periods of time, which can lead to safety decrements (Fisher et al., 2020; Mackworth, 1948). Indeed, research examining the use of automated systems in the field of aviation has found that these safety decrements are due to distraction, over-trust in automated system capabilities, user complacency, skill loss, and reduced situation awareness (Fisher et al., 2020).

Literature Review

L2 Use in Real-World Driving

To date, several NDSs (Dunn et al., 2019; Fridman, Brown, Glazer, et al., 2019; Russell et al., 2018; Stapel et al., 2022) and field operational tests (FOTs; Gaspar & Carney, 2019; Orlovska, Novakazi, et al., 2020; Orlovska, Wickman, & Söderberg, 2020; Reagan et al., 2019) have investigated and published results on drivers’ use of L2 automation features in commercially available vehicles. Of these studies, five reported investigating the drivers’ patterns of usage (e.g., when and/or where the drivers elected to activate and deactivate L2 features), the prevalence of and circumstances leading to takeover requests, or both. Table 1 summarizes the different characteristics (i.e., sample size, location, fleet characteristics, and study duration) for each study.

Table 1. Summary of NDS & FOT L2 Automation Studies

Study	Sample Size	Location	Vehicle Models	Study Duration
L2 MFA NDS (Dunn et al., 2019; Russell et al., 2018)	120 participants 10 vehicles	Washington, DC, metro area	2017 Audi Q7 (*n* = 2) 2015 Infiniti Q50 (*n* = 2) 2016 Mercedes-Benz E350 (*n* = 2) 2015 Tesla Model S (*n* = 2) 2016 Volvo XC90 (*n* = 2)	4 weeks
MIT-AVT (Fridman, Brown, Glazer, et al., 2019; Fridman, Brown, Kindelsberger, et al., 2019)	122 participants 29 vehicles	Boston, MA, metro area	Cadillac CT6 (*n* = 2) Range Rover Evoque (*n* = 2) Tesla Model S (*n* = 16) Tesla Model X (*n* = 7) Volvo S90 (*n* = 2)	37 months
VCC L2 NDS (Dunn et al., 2019)	50 participants 50 vehicles	Washington, DC, metro area	2015 Tesla Model S (*n* = 3) 2016 Tesla Model S (*n* = 8) 2017 Tesla Model X (*n* = 1) 2014 Acura MDX (*n* = 1) 2015 Acura TLX (*n* = 2) 2016 Acura RDX (*n* = 2) 2015 Ford Fusion (*n* = 1) 2017 Ford Fusion (*n* = 1) 2016 Honda Accord (*n* = 2) 2017 Honda Accord (*n* = 2) 2015 Hyundai Genesis (*n* = 2) 2016 Hyundai Genesis (*n* = 1) 2017 Chrysler Pacifica (*n* = 1) 2014 Jeep Cherokee (*n* = 1) 2016 Hyundai Sonata (*n* = 2)	12 months
MIT-AVT FOT (Reagan et al., 2019)	39 participants 4 vehicles	Boston, MA, metro area	2016 Range Rover Evoque (*n* = 2) 2017 Volvo S90 (*n* = 2)	4 weeks
Gaspar and Carney (2019)	10 participants 1 vehicle	Iowa City, IA	2017 Tesla Model S	1 week
Ovrlovska, Novakazi, et al. (2020)	132 participants 132 vehicles	Gothenburg, Sweden	6 undisclosed Volvo models	7 months
Orlovska, Wickman, and Söderberg (2020)	218 participants 218 vehicles	China Sweden United States	7 undisclosed Volvo models	7 months
Stapel et al. (2022)	10 participants 2 vehicles	The Netherlands	7 BMW 540i 3 Tesla Model S	3 months

Driver Usage of L2 Features

Investigations on the usage of novel features are important since insight into use patterns of these technologies may allow for better design and implementation. These investigations may also reveal safety concerns, particularly in L2 automation, as the driver is still responsible for constant supervision. Often, these investigations also identify system misuse, as specific driving conditions and environments may not be supported by the technology, but drivers may engage the systems, nonetheless. Previous work exists in this domain but has focused on luxury vehicles and systems (e.g., Fridman, Brown, Glazer, et al., 2019; Fridman, Brown, Kindelsberger, et al., 2019; Stapel et al., 2022). Therefore, such previous work provides a baseline understanding and investigation into these issues, but a gap still exists in understanding the use of L2 in mainstream vehicles that more directly represent the current U.S. vehicle fleet.

First and foremost, L2 system use has been shown to vary across drivers. A survey of 386 owners of Autopilot-equipped Teslas (Hardman et al., 2019) found that driver self-reported L2 system use could be used to define them into very frequent (15%), frequent (39%), semi-frequent (36%), and infrequent (10%) users. These different user types exhibited diverse use patterns. Very frequent users reported using the automated system on 92% of their trips and for 67% of their commutes. In contrast, infrequent users reported using the automated system for 25% of their trips and for 9% of their commute. While all groups reported being more likely to use the system on freeways free from traffic during clear weather, the very frequent and frequent users also reported using the system on other road types (e.g., urban roads), traffic, and weather conditions.

With these differences across drivers in mind, the following sections summarize current knowledge pertaining to L2 system use in the context of areas related to the research questions.

Research Question: Under what environmental conditions are L2 automation features activated or deactivated?

Time of Day

No study sufficiently reported the times of day that participants elect to activate L2 features. Moreover, the sampling methodology of some previous work lacks the ability to answer this question. An exception is Dunn et al. (2019), who noted that most sampled epochs[1] (selected based on detecting L2 activations across the available data) occurred during daylight hours (indicating a higher prevalence of system use at those times). In addition, Stapel et al. (2022) tried to observe time of day in which the L2 features were active, but night driving was omitted from analysis due to a limited number of activation events. For some of the drivers in Stapel et al., however, there was significantly less use of ACC and LKAS during evening drives.

Weather

Research in this area suggests that drivers generally engage L2 features in clear weather, although adverse weather activation is observed in rare instances. Specifically, Russell et al. (2018) found that drivers were more likely (odds ratio = 1.88) to activate L2 automation features during clear weather compared to when the weather was overcast, raining, misting/lightly raining, foggy, or rainy and foggy. Of the sampled epochs with both features engaged, the weather was clear during 78.8% of the epochs. Similarly, Dunn et al. (2019) found that the weather conditions were largely

---

[1] Epochs refer to specific periods of time selected from the overall driving data due to specific characteristics of interest. The epochs are identified based on the trip that they are selected from and their start and end times within the data stream.

(~80%) clear or partly cloudy during times when drivers had activated the L2 features. Conversely, the researchers rarely observed times when there were rainy conditions during L2 activation.

The activation of L2 systems in weather is interesting from several perspectives. First, weather has known effects in increasing crash risk, so it is possible that there are safety risks associated with L2 system use during inclement weather. Second, the functionality of the systems that support L2 use may be affected by weather. Sensors in these systems must be able to see and perceive their surroundings for timely and appropriate system actions. Yoneda et al. (2019) summarized the impacts of adverse weather conditions, including sun glare, rain, fog, and snow, on common sensors used in automation technologies.

Traffic

Drivers in previous research generally engaged L2 features when traffic was free-flowing or stable. More specifically, Russell et al. (2018) found that traffic density and flow were significant predictors of L2 features engagement. Drivers were more likely to activate L2 features during stable maneuvering traffic (odds ratio = 1.22) than during unstable maneuvering traffic. Similarly, Fridman, Brown, Glazer, et al. (2019) reported that, during periods of speed-restricted traffic, particularly when the travel speed was 10 mph or less, Tesla drivers spent 45.5% of the time under manual control, 15.87% of the time with Traffic-Aware Cruise Control (TACC) engaged, and 19.3% of the time with Autopilot engaged. These researchers also noted that most of the time and distance that drivers traveled with Autopilot engaged were in free-flowing traffic. Similar findings were reported by Stapel et al. (2022), who found that manual driving in their data was more frequent at speeds below 70 kph. Additionally, the L2 features of ACC and LKAS were least used during moving congested traffic conditions, defined by speed ranging between 30 and 80 kph. Finally, in follow-up participant interviews, Orlovska, Novakazi, et al. (2020) reported that the most widely mentioned preferred traffic condition for L2 system usage was on open roads or highways “when traffic was more in a flow.”

Research Question: On what roadway types are L2 automation features activated or deactivated?

Russell et al. (2018) found that drivers were more likely to engage L2 automation features on controlled highways (odds ratio = 2.14) than on business/industrial roads and on two-way divided roads (odds ratio = 4.00) compared to one-way roads. Specifically, of the epochs sampled with both features engaged, 85.6% (*n* = 1,108) occurred during driving on controlled highways and 96.1% (*n* = 1,231) occurred on two-way divided roads. Similarly, Dunn et al. (2019) found that ~80% and 90% of the epochs sampled with L2 features activated occurred on controlled highways and two-way divided roads, respectively.

Fridman, Brown, Kindelsberger, et al. (2019), in turn, observed that drivers tended to primarily choose to activate L2 features on roads with speed limits of 55 mph and above, noting that the U.S. Department of Transportation generally categorizes roads with speed limits of 55 mph and above as interstate, freeway, multilane highway, and other types of arterial roads. Nevertheless, their research also indicated a significant proportion of L2 feature engagements occurred on local roads. Similarly, Reagan et al. (2019) found that nearly 40% of the mileage driven with L2 features active was on interstates, other freeways, and expressways. Orlovska, Novakazi, et al. (2020) complemented these findings with the additional perspective of use frequency. Their findings suggested that high-use drivers (who used L2 automation for about 33% of their trips, on average), used L2 features at lower speeds (0 to 18.6 mph) and on lower volume roads compared to low-use drivers (L2 use ~1% of their trips, on average). Conversely, low-use drivers used L2 features on highways more often than high-use drivers. Participants in this study considered country roads unsuitable for L2 system use and subsequently did not generally elect to engage the features there. Similarly, Morando et al. (2020) showed that Tesla Autopilot system use accounted for ~70% of all the highway miles driven. Finally, Stapel et al. (2022) concluded that the L2 automation features in BMWs and Teslas in their study were used 57% and 63% of the time with highway driving.

Of course, conditions in which L2 system use is possible can be constrained by system limitations, such as geofencing and the ability of the sensors to reliably detect necessary roadway features. This is the case, for example, for the Cadillac Super Cruise system. Gershon et al. (2021) specifically researched this system and observed that Super Cruise was activated for 25% of the ~22,000 miles of data collected, which were primarily driven on limited access highways.

Research Question: How long are L2 features activated, and are there particular trip characteristics that modulate these activations?

Drivers in the Massachusetts Institute of Technology Advanced Vehicle Technology study (MIT-AVT) had mean Autopilot engagement durations of 4.8 minutes during trips where they did not encounter “tricky situations.” These tricky situations were events that the researchers judged could have led to property damage, injury, or death if the driver did not attend to the situation. During trips where drivers encountered tricky situations, however, the mean Autopilot engagement duration was 3.7 minutes before and 5.2 minutes after a tricky situation. These differences, however, were not statistically significant. System availability also seems to play a role in activation duration, as illustrated by the shorter system use durations observed in the MIT-AVT by Gershon et al. (2021): about 2.8 minutes for a Super Cruise system whose operation was ODD-constrained.

In terms of trip duration and L2 activations, the maximum trip duration in the MIT-AVT was 126.6 minutes. As a comparison, in the Virginia Connected Corridor (VCC) L2 NDS, the mean duration of trips with either L2, ACC, or LKAS activation was 42.3 minutes (Dunn et al., 2019). Similarly, across the trips sampled by Russell et al. (2018). the mean trip duration was 53.0 minutes. When trips were clustered into a shorter trip cluster (mean trip duration: 32.8 minutes), middle-trip cluster (71.0 minutes), and a longer trip cluster (159.4 minutes), participants were more likely to activate at least one L2 feature during longer trips. As a further comparison, while no information about activation durations was provided, Orlovska, Novakazi, et al. (2020) reported that drivers had higher usage of L2 features during longer trips (i.e., greater than ~31 miles; 33.9% of trips) compared to shorter (i.e., less than ~9 miles; 4.7%) and medium-length (i.e., greater than ~9 miles and less than ~31 miles; 19.2%) trips.

Takeover Requests

The use of L2 features is inextricably connected to transitions in control of the vehicle. These transitions are sometimes initiated by a request from the vehicle that requires the driver to resume control of the dynamic driving task (DDT). These requests are known as takeover requests and can range in terms of the urgency of the request and the severity of the condition, resulting in the need for a takeover. Given the importance of these situations, some past research has examined the conditions in which they occur and the responses they elicit.

To complicate their study, takeover requests are not always defined uniformly. Dunn et al. (2019) examined three types of takeover requests in the VCC L2 NDS data: forward collision warning (FCW), immediate takeover, and hands-on-wheel (HOW) alerts. FCWs alert the driver when they are too close or are closing too quickly to the lead vehicle. Immediate takeover alerts relate to the lateral control of the vehicle and alert the driver when they need to resume control of steering the vehicle. The HOW alert occurs when the system does not detect the driver's hands on the steering wheel after a period of time and issues visual and/or auditory prompts, sometimes in increasing urgency stages, for the driver to place their hands back on the steering wheel. In total, Dunn et al. found 63 FCW, 61 immediate takeover, and 391 HOW alerts during their study. In contrast, the analysis of the L2 Mixed Function Automation (MFA) NDS (Dunn et al., 2019; Russell et al., 2018) only investigated takeover requests generated by the lateral control features, identifying 449 of these takeover requests within that dataset.

Unlike these previous studies, Fridman, Brown, Glazer, et al. (2019) also measured driver-initiated disengagements in addition to system-initiated takeovers. The system-initiated takeover was the same immediate takeover alert investigated by Dunn et al. (2019). FCW and HOW system-initiated disengagements were not investigated. The three types of driver-initiated disengagement represented the three methods that a driver could use to disengage Autopilot: braking, steering, or toggling the Autopilot stalk. In total, nearly 18,938 disengagements were found in the data, the vast majority (18,800) of which were human initiated. Of the system-initiated requests, 13 were HOW alerts and 115 were immediate takeover alerts.

Additional classification schemes for takeover requests have emerged from the California Department of Motor Vehicles (DMV) disengagement data. It is relevant to note that reported disengagements in these data come from autonomous vehicle testing (generally operating at L3 or higher) and not from drivers using personal vehicles equipped with similar L2 features. The California DMV rule defines a reportable disengagement as a deactivation of the autonomous mode when a failure of the autonomous technology is detected, or when the safe operation of the vehicle requires that the autonomous vehicle test driver disengages autonomous mode and takes immediate manual control of the vehicle. Chen et al. (2018), for example, categorized disengagements into (a) passive, defined as when the automation system recognizes a problem and requires the driver to take over; and (b) active, defined as when the automation system does not recognize a problem and the driver has to manually take over. These two types of disengagements were each further categorized into subcategories. The subcategories for passive disengagement accounted for hardware issues, software issues, weather conditions, and road surface condition. The subcategories for active disengagements accounted for software limitations, hardware issues, emergency situations, and precautionary interventions. Overall, Chen et al. found that software failures and limitations accounted for the vast majority of the passive and active disengagements, accounting for 79.6% and 87.7% of active and passive disengagements, respectively.

Research Question: What situations (e.g., environmental change, roadway change) lead to system takeover requests?

Dunn et al. (2019) did not investigate the situations that lead to the system takeover requests in either the VCC L2 NDS or the L2 MFA NDS datasets. Dixit et al. (2016) examined causes of reported disengagements and found six main categories of reasons for the disengagement: system failure (56.1%), driver-initiated (26.6%), road infrastructure (10.0%), other road users (5.0%), construction zones (1.6%), and weather (0.8%). Fridman, Brown, Glazer, et al. (2019), in contrast, classified disengagements, including takeover requests, based on whether they occurred during epochs with or without the previously mentioned tricky situations. Disengagements with a tricky situation encounter were further categorized as occurring before or after the encounter. In their data, a total of 81 takeover requests occurred without a tricky situation present, none occurred before a tricky situation was encountered, and 47 occurred after a tricky situation was encountered. In Gershon et al. (2021), 14% of system-initiated takeover requests from L2 state to L1 state were due to the system exiting the system's ODD. Takeover requests leading to a manual driving state were generally due to driver inattentiveness or exiting the system's ODD.

Research Question: How frequently do takeover requests occur?

Gershon et al. (2021) found that drivers transitioned control an average of nearly 10 times per trip when they had access to L2 functionality. Fridman, Brown, Glazer, et al. (2019) reported one average tricky situation engagement every 9.2 miles of Autopilot driving; however, there was no quantification of the rate at which system-initiated takeovers occurred.

Method

Data Sample

The final VTTI L2 NDS collection included 249 vehicles. Of these, a subset of 47 vehicles were used for this study. These 47 vehicles were selected from all those that had accessible vehicle network information, permitting the automated detection of L2 feature activation state from the time series data. Alternative efforts were made under this study to use the instrument cluster video view and machine-assisted processing to gather L2 feature activation state for the remaining vehicles, but these efforts did not generate data for the current investigation. Thus, this investigation was completed on the 47-vehicle sample, which included Subaru models (i.e., Forester, Impreza, Crosstek, Outback, Ascent, Legacy) sold between 2017 and 2021. These 47 vehicles accounted for 49,655 trips and 191,328 miles that were analyzed as part of this work.

The analysis periods consisted of all the instances in the vehicles of interest when both the ACC and LKAS were engaged. For simplicity, these events are referred to as “L2 activation epochs” in the rest of this document.

L2 Activation Epochs Selection

The beginning of the L2 activation epoch was defined as the first timestamp when both ACC and LKAS were active. The epoch end was considered to be the first timestamp when either the ACC, LKAS, or both were disengaged. Activation epochs had to last a minimum of 5 seconds without interruption to be considered for analysis. Interruptions between consecutive activation epochs that were shorter than 5 seconds were disregarded. Thus, adjacent activation epochs with less than a 5-second gap were combined prior to analysis.

Activation Metrics

Activation Timing

Each activation epoch was described as a function of the time of day, the distance and time driven with L2 features active when the activation occurred, and the time and distance proximity to the closest crash or near-crash (in days and kilometers, respectively). The activation timing was further classified into the following categories, which served as a surrogate for traffic and environmental conditions:

- Morning: 5 a.m. to 12 p.m.
- Afternoon: 12 to 5 p.m.
- Evening: 5 to 9 p.m.
- Night: 9 p.m. to 4 a.m.

Activation Location

Each activation epoch was characterized as a function of the roadway type(s) on and the speed limits in which the activation occurred. These data elements were identified via GPS location map-matching that leveraged an in-house OpenStreetMap-data-driven version of the Valhalla Map Matching service. The proportion of miles with L2 features engaged was calculated as a function of speed limit exceedance and roadway type (i.e., motorway, trunk, primary, secondary, tertiary, residential, unclassified, or service).

Deactivation Timing

Data at the end of the activation were examined to differentiate between human-initiated and system-initiated deactivations and to identify any challenging scenarios (e.g., presence of curves).

Driver-initiated deactivations were assumed to occur when the driver applied the brakes or steered more than 3.5° in either direction within 3 seconds prior to the deactivation. Driver overrides were assumed to occur when the driver accelerated or steered (more than 3.5° in either direction) within

2 seconds before the deactivation. System-initiated deactivations were assumed to occur when there was absence of braking, steering, or accelerating, and (1) there was an L2 warning signal within 5 seconds of the deactivation, or (2) when the vehicle was stationary within 2 seconds prior to the deactivation.

Furthermore, to identify the presence of curves, the radius of curvature was calculated throughout the trip using the formula:

$$R = \frac{v}{\omega}$$

where v is the vehicle's speed and ω is the yaw rate. If the radius of curvature was smaller than 1,000 m within 10 seconds after the deactivation, the presence of a curve was assumed.

Takeover Request Characterization

It was assumed that the vehicle required the driver to resume vehicle control if an L2 warning signal was present within 5 seconds prior to the deactivation. In these instances, the hands-off-wheel time was summed over the previous 30 seconds to quantify the level of driver disengagement present in the situation.

System Use Propensity

Duration of system usage with respect to the duration and the length of the entire trip was calculated.

Analysis

The research questions of interest in this investigation were examined through descriptive statistics and analysis. The distributions of L2 activations and takeover requests across different driver, vehicle, trip, and environmental characteristics were assessed from the available data.

Results

L2 activation epochs were observed from 3,588 trips encompassing 60,327 miles (7.2% and 31.5% of the total trips and miles in the sample, respectively). These 3,588 trips included 20,043 L2 activation epochs encompassing 28,048 miles (Table 2).

Table 2. Total Number of L2 Trips and Activations

	Frequency	Distance [miles]
L2 Trips	3,588	51,978
L2 Activations	20,043	28,048

Activation Timing

Approximately 69% of all activations and 72% of the trips happened during the daytime (morning or afternoon; Figure 1).

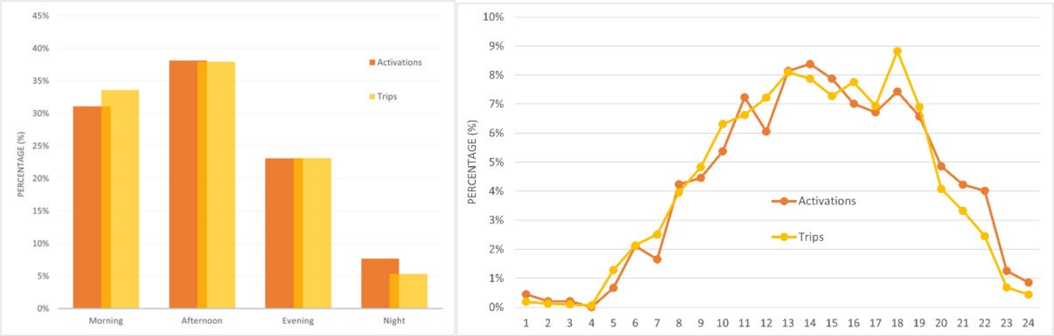

Figure 1. Graphs. Percentage of activations and trips as a function of time of day (left) and hour of day (right).

Their activation distributions as a function of day of the month did not exhibit any notable pattern of difference against the trip distributions (Figure 2, right). Similarly, the activations and trip distributions as a function of the day of the week were similar, except for Sunday, where a comparatively higher percentage of activations were observed (Figure 2, left).

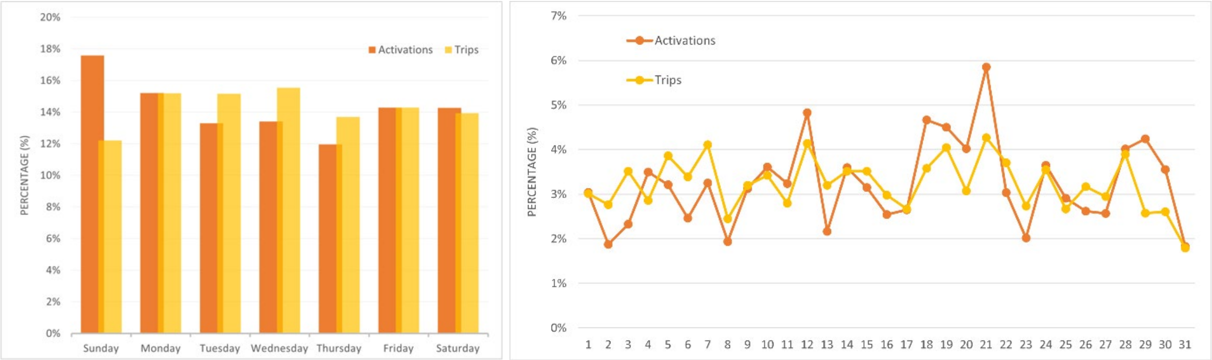

Figure 2. Graphs. Percentage of activations and trips as a function of day of the week (left) and day of the month (right).

Activation Proximity to a Crash or Near-crash

There were no crashes recorded in the dataset. Only eight of the 47 vehicles in the analysis had a near-crash. These eight vehicles had a total of nine near-crash events. These eight vehicles also experienced 5,205 L2 activation epochs. There were four near-crashes in which L2 system activation was observed in the same trip. However, even in the case of closest activation temporal proximity to the near-crash, the deactivation occurred at least 7 seconds before the event. In line with this finding, none of the L2 activations were deemed to be associated with a near-crash event.

Activation Location

Information about the road types traversed was available for approximately 88% of the total distance traveled while utilizing L2 features. Within this distance, approximately 74% of the traveled distance occurred on motorways, 12% on trunk roads, 8% on primary roads, and 5% on secondary roads. System use in tertiary, residential, unclassified, and service roads only accounted for about 1% of overall miles traveled with L2 features engaged.

Given the roads generally traveled while L2 features were engaged, speed limit information availability was high (~92% of the distance traveled while L2 features were engaged). The speed limit was exceeded by 10 mph or more on 28% of the distance traveled with known speed limits (Table 3).

Table 3. Percentage of Miles Traveled with L2 Features Active While Exceeding the Speed Limit

Speed limit exceedance	≥ +0 mph	≥ +5 mph	≥ +10 mph	≥ +15 mph
Percentage of miles traveled exceeding speed limits	88.9%	63.1%	28.1%	8.7%

Deactivation Timing

Driver-initiated Deactivations

In approximately 17% of the activation epochs, the brakes were engaged within the 3 seconds preceding the deactivation of the L2 features. No sharp steering maneuvers, however, were observed within the same interval across the activation epochs (Table 4).

Table 4. Driver-initiated Deactivation Frequencies

	Yes	No	Unknown	Total
Brake application [3 sec]	3,316	16,727	0	20,043
Steering [3 sec]	0	20,043	0	20,043

Driver Overrides

Drivers commonly accelerated to override the L2 features. Acceleration was observed within the 2 seconds prior to deactivation in around 70% of the activation epochs. Like for driver-initiated deactivations, however, no instances of sharp steering maneuvers were observed in that same interval across the activation epochs (Table 5).

Table 5. Driver Override Frequencies

	Yes	No	Unknown	Total
Acceleration [2 sec]	14,027	6,007	9	20,043
Steering [2 sec]	0	20,043	0	20,043

System-initiated Deactivations

Considering all the activations, in most cases (61%) the systems did not trigger any warnings to the driver in the 5 seconds prior to the deactivations (Table 6). Intervention requests (11%), hands-off warnings (3%), and lane departure warnings (2%), however, were also present in the data. Notably, there were no instances of system malfunctions or timeouts. The warning data were not available for 23% of the cases.

Table 6. Warning Types Observed Within 5 sec of Deactivations

Warning types [5 sec]	Frequency
No warning	12,319
Request to intervene	2,237
Hands-off warning	603
Lane departure warning	310

Warning types [5 sec]	Frequency
System malfunctioning	0
System timeout	0
No data available	4,574
Total	20,043

Excluding cases where the drivers braked, steered, or accelerated in the 2 seconds before the deactivation, the systems did not trigger any warnings to the driver in the 5 seconds prior to the deactivations in 61% of the cases. Intervention requests (20%), hands-off warnings (3%), and lane departure warnings (1%), however, were present in the data (Table 7). The vehicle was stationary in only 48 cases (0.24%).

Table 7. Warning Types Observed Within 5 sec of Deactivations When Braking, Steering, and Acceleration Did Not Occur in the 2 sec Before the Deactivations

Warning types [5 sec]	Frequency
No warning	2,771
Request to intervene	898
Hands-off warning	158
Lane departure warning	52
System malfunctioning	0
System timeout	0
No data available	696
Total	4,575

Finally, in 51% of the cases, the radius of curvature was smaller than 1,000 m in the 10 sec after the deactivation, indicating the probable presence of a curve (Table 8).

Table 8. Presence of Curves Following Deactivations

	Yes	No	Unknown	Total
Presence of a curve	10,156	9,885	2	20,043

Takeover Request Characterization

Several different alerts were presented to drivers in the 5 seconds preceding the L2 feature deactivation. The most frequent alert was a request to intervene (11.2% of cases), followed by hands-off-wheel warning (3.0%) and lane departure warning (1.5%). Requests to intervene in particular were likely to have resulted from events where hands-off-wheel time was longer (Table 9).

Table 9. Distribution of Alert Frequency and Associated Hands-off-Wheel Time

Alert [5 sec prior]	Frequency	No hands in contact for [sec]				Always in contact
		$0 < t < 1$	$1 < t < 5$	$5 < t < 10$	$10 < t < 30$	
Hands-off-wheel	603	165	330	50	30	28

Alert [5 sec prior]	Frequency	No hands in contact for [sec]				Always in contact
		0 < *t* < 1	1 < *t* < 5	5 < *t* < 10	10 < *t* < 30	
Request to intervene	2237	445	1,390	249	74	79
Lane departure	310	39	219	37	15	0

System Use Propensity

Time Analysis

The mean trip duration for trips with L2 usage was 22.8 minutes. In these trips, the L2 features were active a mean time of 7.21 minutes (Table 10), about 32% of the L2 trip duration.

Table 10. Descriptive Statistics of L2 Feature Activations and Associated Trip Duration

	Activation Duration [min]	L2 Trip Duration [min]	Trip Duration, All Vehicle Trips [min]
Mean	7.21	22.81	10.05
Median	2.84	16.16	4.96
St. dev.	15.93	28.50	18.56
Max	213.62	380.39	1,434.98
Min	0.08	0.27	0.00
Total time	25,854.83	81,832.33	470,989.81

A majority of the L2 feature activations (~92%) lasted less than 20 minutes, with over 80% lasting less than 10 minutes. This tendency towards a preponderance of shorter durations was also observed, albeit much less saliently, for the overall duration in trips where L2 features were activated. For those trips, about 32% and 29% of the trips lasted between 0 to 10 and 10 to 20 minutes, respectively (Figure 3).

Individual L2 feature activations tended to last longer as trip duration was longer (Figure 4, left). As may be expected based on the mean duration of trips with L2 activations (22.8 min), however, most activations (~65%) tended to occur in trips with durations shorter than 60 minutes (Figure 4, right).

The comparison of active and inactive time proportions for L2 features across various trip duration intervals revealed a relatively consistent pattern wherein the features remain active for approximately 30% of the time (Figure 5).

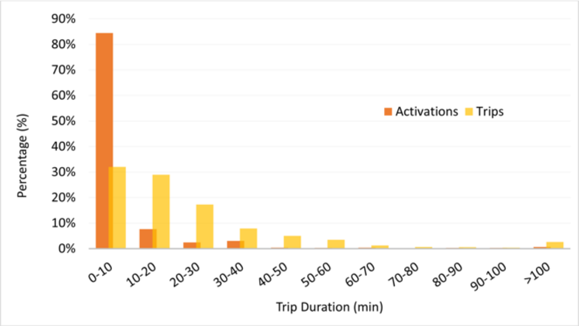

Figure 3. Graph. Distributions of time duration of L2 feature activations and of the trips in which those activations occurred.

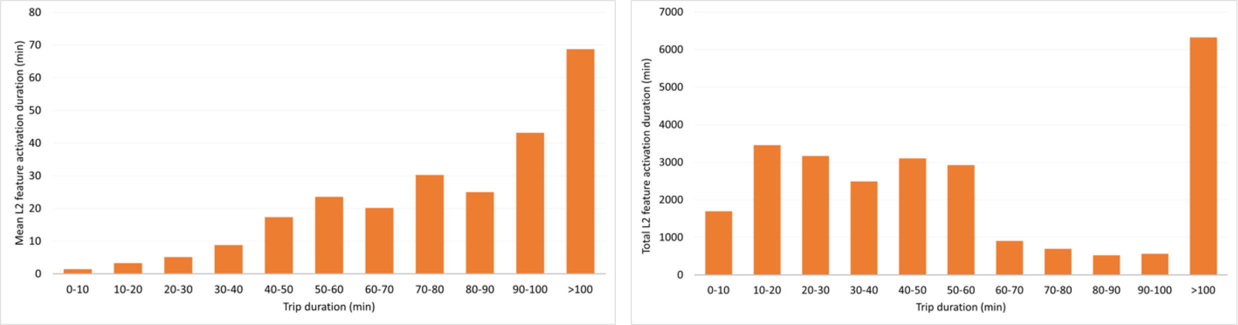

Figure 4. Graphs. Mean (left) and total (right) L2 feature activation duration as a function of trip duration.

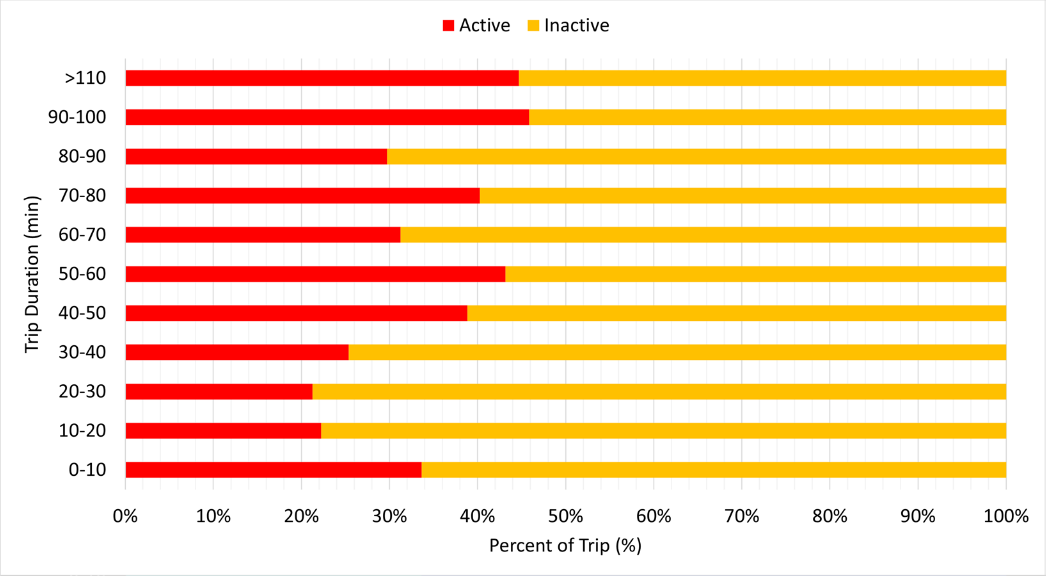

Figure 5. Graph. Comparison of active and inactive time proportions for L2 features as a function of trip duration.

Mileage Analysis

The mean trip distance for trips with L2 usage was 15.50 miles. In these trips, the L2 features were active a mean distance of 7.20 miles (Table 11), about 46.5% of the L2 trip distance.

Table 11. Descriptive Statistics of L2 Feature Activations and Associated Trip Distance

	Activation Distance [miles]	L2 Trip Distance [miles]	Trip Distance, All Vehicle Trips [miles]
Mean	7.20	15.50	4.08
Median	2.17	8.40	0.46
St. dev.	18.35	27.12	13.56
Max	260.93	327.06	338.96
Min*	0	0	0
Total time	28,047.89	60,327.25	191,328.17

* - Sensor malfunctions generated files with movement, but with zero mileage, this in turn generated these minimums of zero miles traveled.

A vast majority of the L2 feature activations (92%) occurred over less than 20 miles, with over 85% lasting less than 10 miles. This tendency towards a preponderance of shorter trip distances was also observed, albeit much less saliently, for the overall distance traveled in trips where L2 features were activated. For those trips, about 35% and 25% of the trips lasted between 0 to 5 and 5 to 10 miles, respectively (Figure 6).

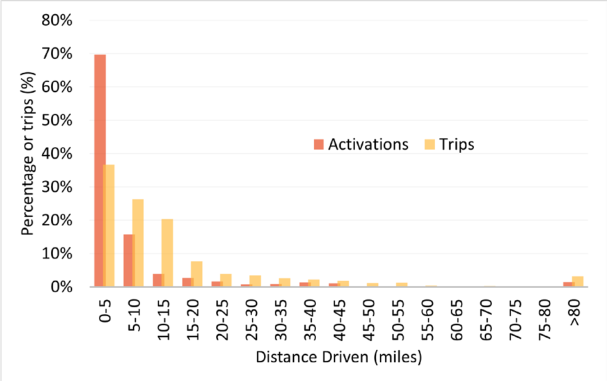

Figure 6. Graph. Distributions of distance driven during L2 feature activations and during the trips in which those activations occurred.

As was the case for time duration of trips, longer activations showed a consistent pattern of occurring in longer trips (Figure 7, left). The total distribution of activations across total trip mileage (Figure 7, right) was similar to the distribution observed for the trip time durations (Figure 4, right). Longer distance trips, however, accounted for a sizeable proportion of the overall L2 feature activation miles.

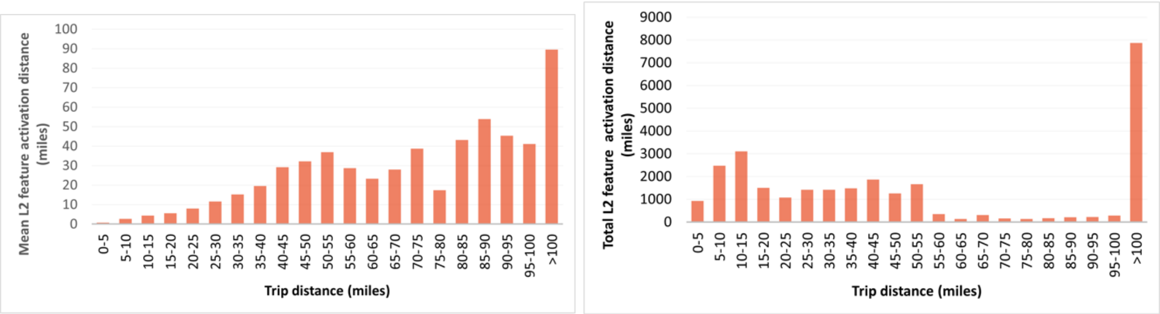

Figure 7. Graphs. Mean (left) and total (right) L2 feature activation distance as a function of trip distance.

Also similar to the time duration of trips, the comparison of active and inactive mileage proportions for L2 features across various trip distance intervals revealed a relatively consistent pattern. For the case of mileage, however, the observed proportion of distance that L2 features were active was around 40% (compared to ~30% for the time duration metric; Figure 8).

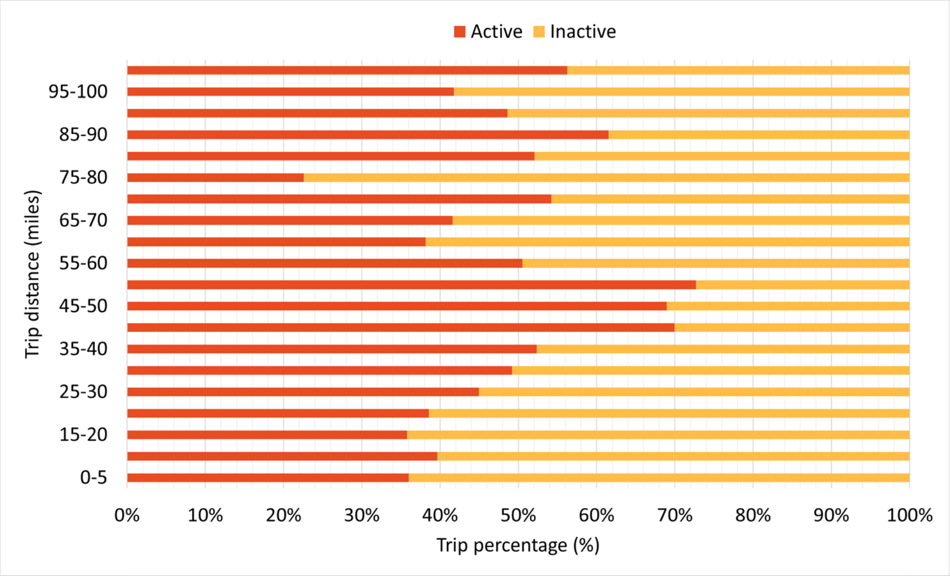

Figure 8. Graph. Comparison of active and inactive time proportions for L2 features as a function of trip distance.

Discussion

The results of this investigation showed similarity with previous efforts to describe L2 automation feature use, but with an increased level of confidence arising from a large number of vehicles that were participant-owned. The investigation also extends previous research to include non-luxury vehicles, which represent a large proportion of the overall vehicle fleet in the United States. The discussion of these results is structured around the research questions motivating the work.

Driver Usage

Under what environmental conditions are L2 automation features activated or deactivated?

While constraints in the study approach prevented examination of weather and traffic patterns, assessment of the L2 automation feature usage times shows no sizeable differences between periods of L2 feature usage and general driving periods. For time of day, usage is more frequent during daytime periods, which matches general driving patterns. This observation matches the findings of Dunn et al. (2019). Reduced evening use also aligns with the findings of Stapel et al. (2022). More generally, distributions of L2 feature usage across a given month and the days of the week are also relatively uniform, also matching general driving.

On what roadway types are L2 automation features activated or deactivated?

The vast majority of L2 feature usage occurred in motorways, with trunk, primary, and secondary roads also registering some usage. The 70% usage observed here compares favorably with Dunn et al. (2019), but the odds of motorway usage over primary or secondary road usage appear to be larger than observed in Russell et al. (2018). The latter observation may be partly due to distinct roadway classification schemes used in this study and Russell et al.'s work. A distinct classification scheme may also have contributed to an observed motorway usage that was larger than reported by Reagan et al. (2019) for interstates (~40%).

Interestingly, speeding while L2 features were being used seems to be fairly prevalent. The prevalence estimate was about 30% when using a 10-mph speeding threshold. This proportion is slightly larger than Perez et al. (2021) generally observed for time-based proportions in non-L2 vehicles (<25%). Although the mileage-based estimates in this investigation may not be fully comparable to the time-based estimates in Perez et al., the finding nonetheless suggests a fairly strong tendency for drivers to set their ACC speed settings above the speed limit.

How long are L2 features activated, and are there particular trip characteristics that modulate these activations?

On average, L2 features were activated for 7.2 minutes in trips lasting an average of 22.8 minutes, or about 32% of the duration of trips featuring L2 use. Trips where L2 feature usage was observed, however, were over 10 minutes longer than the average trip. Mileage-based use patterns were similar, but the representation of L2 feature usage rises from 30% to 40% when a mileage basis is used. These observations generally suggest longer periods of use than observed in previous literature by between 2 and 5 minutes (Gershon et al., 2021). Contrary to Dunn et al. (2019) and Orlovska, Novakazi, et al. (2020), however, there were no clear trends indicating disproportionately longer L2 feature use in longer trips.

Takeover Requests

What situations (e.g., environmental change, roadway change) lead to system takeover requests? How frequently do these conditions occur?

The analysis related to this question focused on both L2 feature deactivations and associated takeover requests. Driver-initiated overrides were predominantly done by braking or accelerating the vehicle, with steering-based overrides being minimal or non-existent and likely involving lane changes without using a turn signal. Generally, these lane changes would result in lane departure warnings, which were indeed observed in the data. Intervention requests, where the system asks the driver to take over, were the most common takeover requests, followed by requests due to insufficient driver hand contact with the steering wheel. Notably, some of the intervention requests would have resulted from extended hands-off-wheel time. The system failures noted as a large proportion of the disengagements in Dixit et al. (2016) were not present at all in this sample, which may be partly due to increased system reliability but may also be due to differences in what was considered a system failure across both studies. Interestingly, since the systems in this study were not map-based, it is difficult to assess reliably the extent to which ODD changes led to deactivations, as noted in Gershon et al. (2021).

From the perspective of system deactivation frequency, the systems were deactivated an average of 5.6 times per trip where L2 feature usage was observed, and 0.33 times per mile driven in those trips. The observed frequency per trip is lower than noted in Gershon et al. (2021; ~10 times per trip). The observed miles driven per deactivation (~3) are also lower than the 9.2 miles driven between tricky situations reported in Fridman, Brown, Glazer, et al. (2019). These figures may, however, not be directly comparable due to different system capabilities.

Conclusions and Recommendations

Altogether, the findings in this investigation suggest that as L2 features penetrate the U.S. fleet in non-luxury consumer vehicles, usage of the systems is quite common and comparable with previous findings for luxury offerings. While some evidence of potential system misuse was observed, future work can strive to further operationalize system misuse and further assess the prevalence of such behaviors. Drivers were observed to mostly use these systems on roadways for which they were designed to operate. Nevertheless, given that over 3,000 of the 20,000 system activations ended with some sort of warning being provided to the driver, it seems like transitions in system state are not always fluid for these systems, and certainly not always driver initiated. Future research should be dedicated to understanding the reasons for these deactivations more clearly and specifically. Future research should also be devoted to the operation of newer L2 features, as these systems continue to rapidly evolve and become increasingly mainstream in new vehicle offerings.

Additional Products

Unless otherwise noted, the products listed in this section are available on the [project page] on the Safe-D website. The dataset used for the project is available within the Safe-D Dataverse.

Education and Workforce Development Products

This project partially supported the work of five graduate students: Mr. Nicholas Britten, Ms. Martha Gizaw, Mr. Haden Bragg, Ms. Mariette Metrey, and Mr. Paolo Terranova. Throughout the project, the students developed firsthand experience with project management and task completion, as well as experience completing literature reviews, developing data analysis plans, and conducting naturalistic driving data analysis. The students also increased their exposure and understanding of different L2 automation features and systems, noting differences in implementation and operation. The project also supported one exhibit focused on vehicle instrumentation and related data, presented to over 6,000 attendees at the Virginia Tech Science Festival in 2019. An undergraduate-level educational module using the study dataset was also developed and is available within the Safe-D website. The module provides some background in this area of research and guides students through relevant analysis approaches: https://safed.vtti.vt.edu/projects/characterizing-level-2-automation-in-a-naturalistic-driving-fleet/

Technology Transfer Products

This project was completed with supporting funding from the National Surface Transportation Safety Center of Excellence (NSTSCE). The NSTSCE board includes personnel from General Motors Corporation, State Farm Insurance, the Virginia Department of Transportation, and the Federal Motor Carrier Safety Administration. These members have been periodically briefed about the progress for this project and passed relevant information along to their organizations. In addition, outputs from this project have supported the Automated Mobility Partnership, a consortium of private organizations interested in pre-competitive collaboration in the development of data sources and analytics approaches that support continued development of ADS and ADAS.

Beyond this collaboration, the following technology transfer products were created:

- Dictionary to standardize the system states of ADAS technologies, particularly those related to L2 automation features. The dictionary is available on the SAGE advance TransportRxiv as a publicly available preprint. [Link]
- A human-in-the-loop machine-assisted approach for inferring L2 feature activation state from video of a vehicle's instrument panel (see the Appendix).

Data Products

A spreadsheet containing the data used in this investigation is available from the Safe-D Dataverse. [Link]

References

Bainbridge, L. (1983). Ironies of automation. *Automatica, 19*(6), 775-779.

Chen, L., Dongpu, C., Yifan, Z., Auger, D. J., Sullman, M., Huaji, W., Dutka, L. M., Skrypchuk, L. & Mouzakitis, A. (2018). Analysis of autopilot disengagements occurring during autonomous vehicle testing. *IEEE/CAA Journal of Automatica Sinica, 5*(1), 58-68.

Dixit, V. V., Chand, S., & Nair, D. J. (2016). Autonomous vehicles: Disengagements, accidents and reaction times. *PLoS One, 11*(12), e0168054.

Dunn, N., Dingus, T., & Soccolich, S. (2019). *Understanding the impact of technology: Do advanced driver assistance and semi-automated vehicle systems lead to improper driving behavior?* (Technical Report). AAA Foundation for Traffic Safety.

Fisher, D. L., Horrey, W. J., Lee, J. D., & Regan, M. A. (2020). *Handbook of human factors for automated, connected, and intelligent vehicles.* https://www.taylorfrancis.com/books/9781315269689

Fridman, L., Brown, D. E., Glazer, M., Angell, W., Dodd, S., Jenik, B., . . . Reimer, B. (2019). MIT advanced vehicle technology study: Large-scale naturalistic driving study of driver behavior and interaction with automation. *IEEE Access, 7,* 102021-102038.

Fridman, L., Brown, D. E., Kindelsberger, J., Angell, L., Mehler, B., & Reimer, B. (2019). *Human side of Tesla Autopilot: Exploration of functional vigilance in real-world human-machine collaboration*.

Gaspar, J., & Carney, C. (2019). The effect of partial automation on driver attention: A naturalistic driving study. *Human Factors, 61*(8), 1261-1276.

Gershon, P., Seaman, S., Mehler, B., Reimer, B., & Coughlin, J. (2021). Driver behavior and the use of automation in real-world driving. *Accident Analysis and Prevention, 158*, 106217.

Hardman, S., Lee, J. H., & Tal, G. (2019). How do drivers use automation? Insights from a survey of partially automated vehicle owners in the United States. *Transportation Research Part A: Policy and Practice, 129*, 246-256.

Mackworth, N. H. (1948). The breakdown of vigilance during prolonged visual search. *Quarterly Journal of Experimental Psychology, 1*(1), 6-21.

Morando A., Gershon P., Mehler B., & Reimer B. (2020). Driver-initiated Tesla Autopilot disengagements in naturalistic driving. *Proceedings of the 12th International Conference on Automotive User Interfaces and Interactive Vehicular Applications, Washington, DC, Association for Computing Machinery*, 57–65.

Orlovska, J., Novakazi, F., Lars-Ola, B., Karlsson, M., Wickman, C., & Söderberg, R. (2020). Effects of the driving context on the usage of Automated Driver Assistance Systems

(ADAS): Naturalistic driving study for ADAS evaluation. *Transportation Research Interdisciplinary Perspectives, 4,* 100093.

Orlovska, J., Wickman, C., & Söderberg, R. (2020). Naturalistic driving study for Automated Driver Assistance Systems (ADAS) evaluation in the Chinese, Swedish and American markets. *Procedia CIRP, 93,* 1286-1291.

Perez, M.A., Sears, S., Valente, J.T., Huang, W., & Sudweeks, J. (2021) Factors modifying the likelihood of speeding behaviors based on naturalistic driving data. *Accident Analysis and Prevention, 159*, 106267.

Reagan, I. J., Hu, W., Cicchino, J. B., Seppelt, B., Fridman, L., & Glazer, M. (2019). Measuring adult drivers' use of Level 1 and 2 driving automation by roadway functional class. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting, 63*(1), 2093–2097.

Russell, S. M., Blanco, M., Atwood, J., Schaudt, W. A., Fitchett, V., & Tidwell, S. (2018). *Naturalistic study of Level 2 driving automation functions* (Report No. DOT HS 812 642). National Highway Traffic Safety Administration.

SAE International. (2021). J3016_202104 Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles.

Stapel, J., Happee, R., Christoph, M., van Nes, N., & Martens, M. (2022). Exploring the usage of supervised driving automation in naturalistic conditions. *Transportation Research Part F: Traffic Psychology and Behavior, 90*, 397-411.

Yoneda, K., Suganuma, N., Yanase, R., & Aldibaja, M. (2019). Automated driving recognition technologies for adverse weather conditions. *IATSS Research, 43*(4), 253-262.

Appendix. Human-in-the-Loop Machine-Assisted Approach for Inferring L2 Feature Activation State from Instrument Panel Video

A human-in-the-loop machine-assisted approach for inferring L2 feature activation state from instrument panel (IP) video was developed to support the need to summarize IP video based on the presence of icons that indicate L2 feature status. The approach, and associated tool, was developed by Charles Layman, a Data Engineer within the Division of Data and Analytics at Virginia Tech, under the auspices of this project. The tool allows machine vision processes to mine IP video data efficiently, while also providing an interface for human validation of the L2 feature states to occur. The tool was developed due to limited success in previous efforts to fully automate detection of L2 feature states using full IP video. Those efforts were particularly hindered by the high severity, frequency, and variability of glare and other reflections on the IP cover as the vehicle was in motion, which could overload the camera sensors due to their high brightness in comparison to the IP light sources of interest. The goal of the tool was to streamline a human-in-the-loop process to support the automated detection of L2 feature states.

The process entails parsing the IP video and a small subset of the time series data, subsequently generating a static web page from HTML, CSS, jQuery, and additional required files (Figure 9). Display regions of interest for the various indicators on the web page were predetermined and provided as inputs to the process. Above the dashed line, the web page displays the view in a given region of interest at any desired point in time, and additional interaction with that portion of the page initiates a video feed that allows the user to watch the change in the relevant regions of interest over time. Below the dashed line, the web page displays a spatial dimension summary of the region of interest, which allows the whole trip to be scrolled through and viewed as a single image. The spatial summary was computed through a weighted mean of the vertical pixels in the image, effectively collapsing each frame of video for the region of interest into a line of pixels with varying intensity. This allows identification of changes in the image to be done relatively easily by human data annotators (Figure 10). While the presence of glare in the video can still represent an issue, the underlying signal is usually still visible, if it is present. Glare also tends to be more transient in nature than the signals of interest, which creates patterns in the data easily visible to the human eye.

This tool will continue to be leveraged in future extended analyses of the VTTI L2 NDS dataset used for this study.

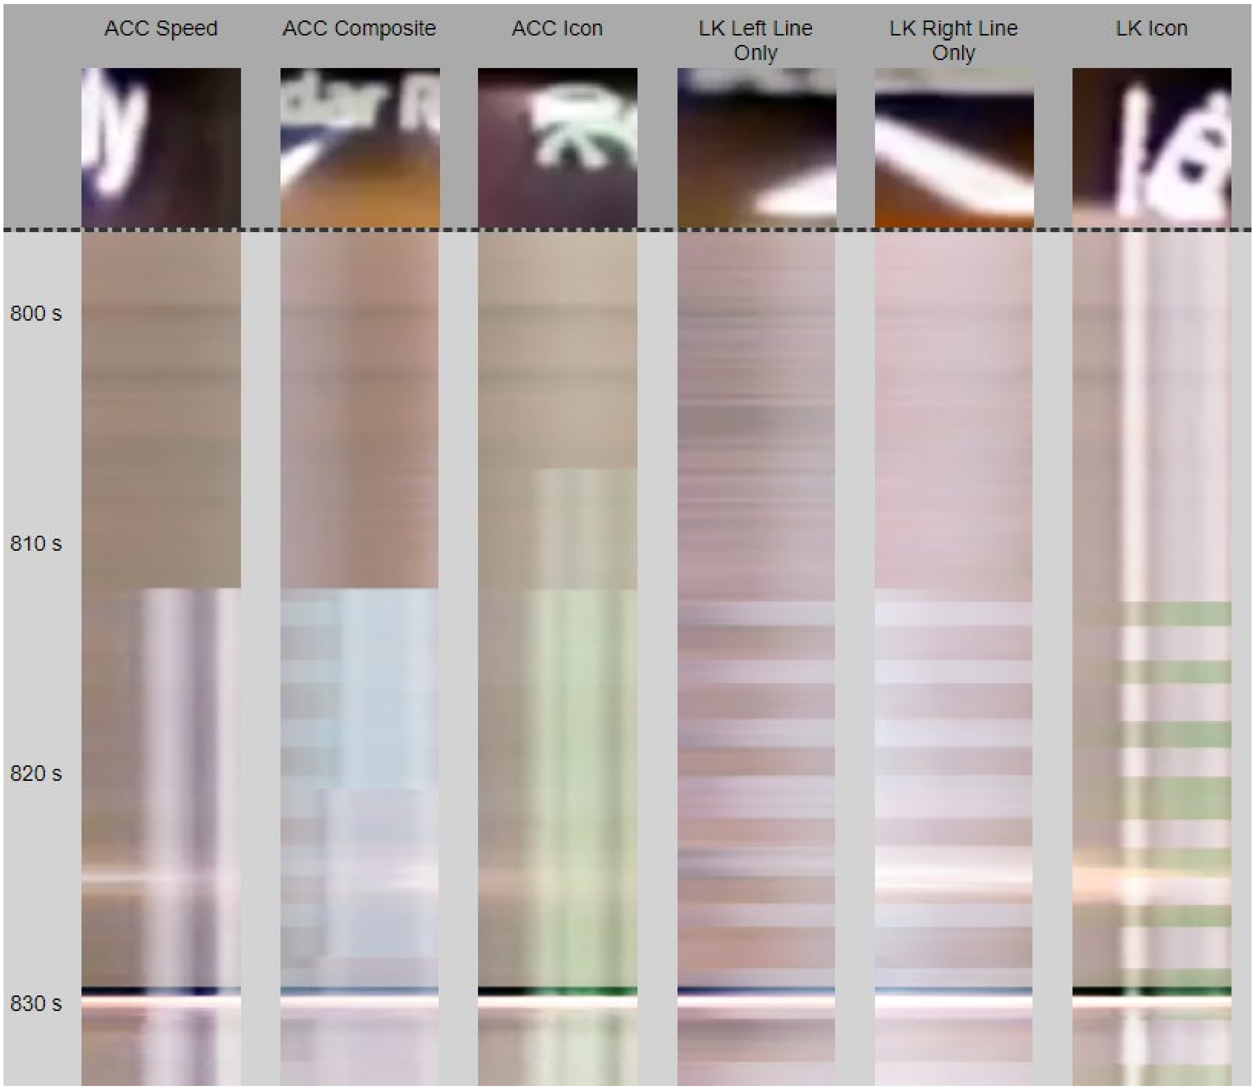

Figure 9. Image. Sample view of the static web page resulting from video parsing.

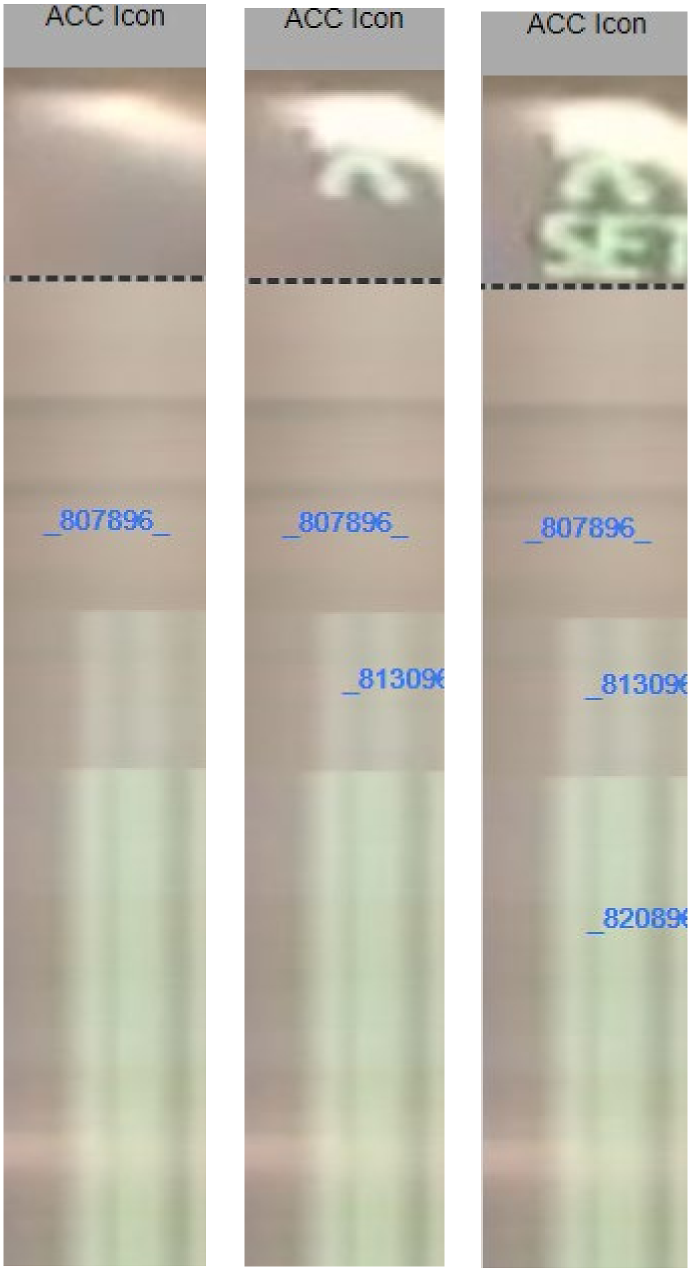

Figure 10. Image. Sequence of interactions with the web page showing an ACC activation. Note the progression in pattern within the region below the dashed line as the icon appears and ACC speed is set. Numbers in blue indicate the timestamp for the frame being shown above the dashed line.