2019 Forum on the Impact of Vehicle Technologies and Automation on Users – Design and Safety Implications: A Summary Report

April 2020

607 14th Street, NW, Suite 201 | Washington, DC 20005 | 202-638-5944

AAA Foundation for Traffic Safety

Title

2019 Forum on the Impact of Vehicle Technologies and Automation on Users – Design and Safety Implications: A Summary Report

(April 2020)

Foreword

2019 marked the third Forum on the Impact of Vehicle Technologies and Automation on Users. The evolution and advancement of vehicle technologies continues at a rapid pace, along with our appreciation of the many important research questions still needing answers. Such answers will help to realize the many potential safety and mobility benefits of automated technology.

This report summarizes presentations and discussion from the 2019 Forum held in the University of California San Diego. Stakeholders from academia, industry and government gathered to discuss and exchange information and ideas about the impact that emerging transportation technologies are having on road users, with a focus on the design and safety implications of these systems. This report should be of interest to researchers and practitioners who are involved with work related to vehicle technologies and automation.

C. Y. David Yang, Ph.D.

Executive Director
AAA Foundation for Traffic Safety

About the Sponsor

AAA Foundation for Traffic Safety
607 14th Street, NW, Suite 201
Washington, DC 20005
202-638-5944
www.aaafoundation.org

Founded in 1947, the AAA Foundation for Traffic Safety is a nonprofit, publicly supported charitable research and education organization dedicated to saving lives by preventing traffic crashes and reducing injuries when crashes occur. Funding for this report was provided by voluntary contributions from AAA/CAA and their affiliated motor clubs, individual members, AAA-affiliated insurance companies and other organizations or sources.

Table of Contents

Abbreviated Terms

ADAS – Advanced Driver Assistance Systems

AAA – American Automobile Association

AV – Automated Vehicles

CAN – Controller Area Network

EDR – Electronic Data Recorder

eHMI – External Human-Machine Interface

LIDAR – Light Detection and Ranging

NDS – Naturalistic Driving Study

ODD – Operational Design Domain

OEM – Original Equipment Manufacturer

OTA – Over-the-Air

PDO – Property Damage Only

VIN – *V*ehicle Identification Number

Introduction

On Nov. 4 and 5, 2019, the AAA Foundation for Traffic Safety and the University of California San Diego hosted a forum to discuss issues and identify future research needs on the impact of vehicle technologies and automation on users, focusing on the design and safety implications of these systems.

The forum in San Diego, the third in the series, was attended by academics, industry representatives, government agencies, advocacy groups, and other research organizations (see Appendix A for a list of participating organizations). The forum was co-sponsored by the AAA Life Insurance Company, the AAA Club Alliance, and AAA Northeast.

The main objectives of this forum were to: (a) gather representatives/experts from the research community, government, and industry to discuss issues and identify research needs and critical considerations regarding the design and safety of vehicle technologies and automation, and (b) develop a summary report documenting the outcomes of the panel and group discussions and share it with other stakeholders to improve coordination and encourage collaboration.

On Day 1, three expert panels were convened to discuss a variety of topics related to vehicle technology, automation, and their interaction with transportation system users. The panels were: The State of Vehicle Technology and Automation (Panel 1), Design Recommendations and Considerations (Panel 2), and Implications to Future Safety and Operations (Panel 3). Each panel was followed by an extended question-and-answer period. On Day 2, all forum attendees engaged in a world café exercise, a structured conversational process in which small groups discuss and share knowledge on a topic or question. The panel presentations and discussions, world café exercise, and outcomes are described in the sections below. The forum agenda can be found in Appendix B.

Day 1: Introductions and Panel Presentations

Drs. David Yang (*AAA Foundation for Traffic Safety*) and Albert Pisano (*University of California San Diego*) opened the forum with welcoming remarks.

Panel 1: State of Vehicle Technologies & Automation

(Panelists: Dr. Henrik Christensen, *University of California San Diego* and Tom Alkim, *European Commission;* Facilitated by Dr. Tara Kelley-Baker, *AAA Foundation for Traffic Safety*)

Dr. Christensen discussed the importance of considering vehicle technologies and automation in the context of the road eco-system, including changes in infrastructure and potential security issues. He noted several challenges related to sensor limitations (e.g., effective range of LIDAR) and data processing and management (e.g., onboard processing of sensor data; transferring data from the vehicle; updating and maintaining high-density maps). Current machine learning algorithms are capable of identifying and tracking vehicles and motorcycles with fairly high accuracy. However, the same level of classifying and predicting pedestrian movement has not yet been attained due to their complex behaviors and movements (though classifiers are improving and are now capable of identifying objects carried by pedestrians). Dr. Christensen noted that, overall, the progress of AV implementation has been slower than anticipated and there are many issues needing resolution, including regulations and liability issues.

Mr. Alkim discussed many initiatives under way in Europe related to AV, safety, mobility, and infrastructure readiness. He reiterated several core yet recurring themes that have been discussed at international venues, including safety validation, issues related to trust, ethical issues, data protection/cybersecurity, and cooperation. He also discussed what would be acceptable behaviors in the AV driving environment, how we define them, and how they would differ from the human-operated driving environment. He tracked the position of automated vehicles along the Gartner Hype Cycle over the past decade, noting that in the most recent years, AVs were broken out according to different levels of technology. He also described the conceptualization of operational design domains (ODD) within the Strategic Transport Research and Innovation Agenda (STRIA) Roadmap. Namely, that beyond the first and last mile, ODD would likely be characterized by smaller, local barriers (or gaps) within larger boundaries. Moreover, he classified different driving environments by the level of challenges for AV based on the dimensions of complexity and velocity. Based upon this definition, rural roads are the most challenging, while business and campus parking lots are the least.

Panel 2: Design Recommendations & Considerations

(Panelists: Dr. Linda Boyle, *University of Washington*, Dr. John Campbell, *Exponent*, and Dr. James Jenness, *Westat*; Facilitated by Dr. William Horrey, *AAA Foundation for Traffic Safety*)

In considering the design of systems, Dr. Boyle emphasized the importance of understanding the motivations and/or constraints of the controller (i.e., the agent who is controlling the vehicle, whether driver, automated system or the joint product). System-driver handoffs or takeovers will be influenced by this understanding and there is potential confusion arising from many sources. She noted that a key question is how to design for these limitations and how best to support communication between system and driver in order to reduce mismatches in what drivers believe to be the mode of operation and the actual mode (articulated in terms of a signal detection framework). Dr. Boyle also addressed how vehicles could become a platform for work — especially for knowledge workers. Along these lines, she discussed safety, satisfaction, and productivity issues related to AVs. To increase people's satisfaction, AVs need to minimize motion sickness, communicate the systems' awareness and intentions, account for non-driving activities, and make other societal considerations.

Dr. Campbell addressed that the basic designs of in-vehicle interfaces and controls (e.g., location, color, message, text, luminance, etc.) have been well understood and codified through past human factors research. Moreover, studies of alerts and warning systems have also provided insight for design. In the context of ADAS and AV failure scenarios, the concepts of trust (including under-trust and over-trust), system understanding (i.e., mental models), and driver take-over readiness are critical. Dr. Campbell also discussed several important discrepancies between what the system promises (i.e., myths) and what happens in reality, noting some of the unintended consequences that are possible in the context of human-automation interactions. Echoing Dr. Boyle, he underscored the importance of facilitating more meaningful communication, beyond status indicators and warnings, to better calibrate driver trust and mental models. This includes, but is not limited to, conveying information concerning the amount and type of support or capability provided by the automation to the driver. He instilled playbooks and waypoint visualizations as useful design metaphors. He also highlighted the role of attention management strategies and behavioral incentives as means of promoting better drivers' visual scanning (looking at the right thing at the right time), thereby improving driver's readiness. He also noted that better measures of driver engagement should be sought (cf. eyes on road).

Dr. Jenness discussed the important interactions between automation-equipped vehicles and other road users. Under manual operation, the vehicles' appearance, motion, and intentional behavior are all used to predict its actions and how other users should interact with it. Drivers also provide implicit and explicit cues. In the context of AVs, however, many existing assumptions or expectations regarding these cues could change. External human-machine interfaces (eHMI), where the vehicle itself could be used as an interface for other road users, provide one possible avenue to ameliorating communication between vehicles and other road users. There are, however, many challenges: communicating to the appropriate users (e.g., an individual on a crowded street); short-range versus long-range communication and the appropriate timing; the traffic context and environment (e.g., different traffic volumes, localities, weather); and other unintended consequences and liability concerns. Designers also need to consider how much information is too much, the importance of standardization, and the needs of a full range of road users (e.g., blind pedestrians). There is potential to design the system using dynamic vehicle cues. Cues, in different forms, can foster and reinforce traffic culture and AV can use predictable and repeatable gestures. Dr. Jenness also raised the potential of leveraging or repurposing existing signals or elements to facilitate communication between road users.

Panel 3: Implications to Future Safety & Operations

(Panelists: Christopher Hart, *Hart Solutions, LLC*, Dr. Jessica Cicchino, *Insurance Institute for Highway Safety*, and Dr. Bryant Walker Smith, *University of South Carolina*; Facilitated by Brian Tefft, *AAA Foundation for Traffic Safety*)

Mr. Hart discussed some of the important lessons learned in aviation automation, including human-centric approaches, allocation of function, and the implementation of alerts. In describing lessons learned from accident investigations in other domains, he highlighted the interplay between automation system failures, the communication between the system and operator, and the operator state. While acknowledging the great potential of these systems, he discussed many new challenges that have not been encountered or prominent in aviation automation, including the role of deep learning of artificial intelligence, the low level of training in drivers (compared to pilots), the necessity of testing on actual roads, the need for graceful exits (cf. stopping the car in its lane or simply disengaging automation), the time constraints faced by drivers versus pilots when struggling to recover from a system error, the complexities of mixed fleets, and of the traffic environment in general.

Dr. Cicchino described the outcomes of several recent safety evaluations of ADAS and AV technology, noting that current generations of crash avoidance technologies have been shown to be effective in crash mitigation and that, across OEMs, vehicles are showing improvements over time. For example, more are achieving superior crash ratings. In evaluations of Level 2 automation, some OEMs are showing advantages with respect to collision or property damage-only liability. In discussing the utility and accessibility of crash and vehicle data, Dr. Cicchino also described several important challenges, such as understanding which vehicle has what features. The database they have developed also identifies system features by matching up with the vehicle identification number (VIN). However, this is not always comprehensive. She also noted that crash data recordings do not yet need to include information about the automation state (i.e., whether the system is engaged or not) and, currently, it is sometimes difficult to account for system software updates, whether over-the-air or dealer-based.

Dr. Walker Smith discussed safety in a more holistic view of the AV environment, including the public policy perspective. This also includes the role and impact of mass systemic failures, such as coordinated large-scale cyberattacks, network outages, and other unintended consequences. He underscored that safety is a lifetime commitment, not the product of a single event or outcome. That is, safety transcends the lifecycle of the system, from design through implementation and beyond. Thus, it is not necessarily about the technology itself but rather trustworthiness in brands. Dr. Walker Smith also discussed how the advent of AV technology also raises new questions about who drives an automated vehicle; even current legal definitions of ‘driver’ confound the matter. Fundamentally, he argued that companies will drive AVs, not drivers and computers and, because companies are legal subjects that act through their humans and machine agents, AVs might be considered only as safe as the companies behind them. Therefore, he suggested that we need to regulate the company not the technology, focusing instead on processes and systems (e.g., Automated Operation of Vehicles Act) and penalize the company for breaching trust.

Day 2: World Café Discussion and Outcomes

Following a similar model as in the previous years, the main charge on Day 2 was to break into small groups and discuss a series of questions. In 2019, a World Café approach was adopted, which is a structured conversational process for sharing knowledge where small groups discuss a topic at a table, with groups switching tables after a set time (here, 20-minutes) and getting introduced to the previous discussion at their new table by a "table host." After rotation, the groups seek to build upon the discussions carried out by earlier groups. In the current exercise, eight small groups were exposed to four different questions:

1. Given current states of vehicle technology, what are the most pressing research needs concerning user interactions with automated systems? How can these research needs be addressed?
2. What are the most important measures of safety in the context of vehicle automation, and what are appropriate targets or thresholds? (Consider both pre-deployment and post-deployment.)
3. How can we improve our understanding of driver interaction with vehicle automation? What methods can we use?
4. What are the key focus areas related to safety in the next few years to continue to advance vehicle automation?

Information from the group presentations, the notes from group interactions, and the feedback gleaned from individuals have been distilled and synthesized in the sections below. While this captures many of the main themes, it does not do justice to the rich, dynamic, and interconnected threads comprising the group discussions. Further, there was some overlap in the scope of the questions as well as in the ensuing discussion. This is reflected in the sections below.

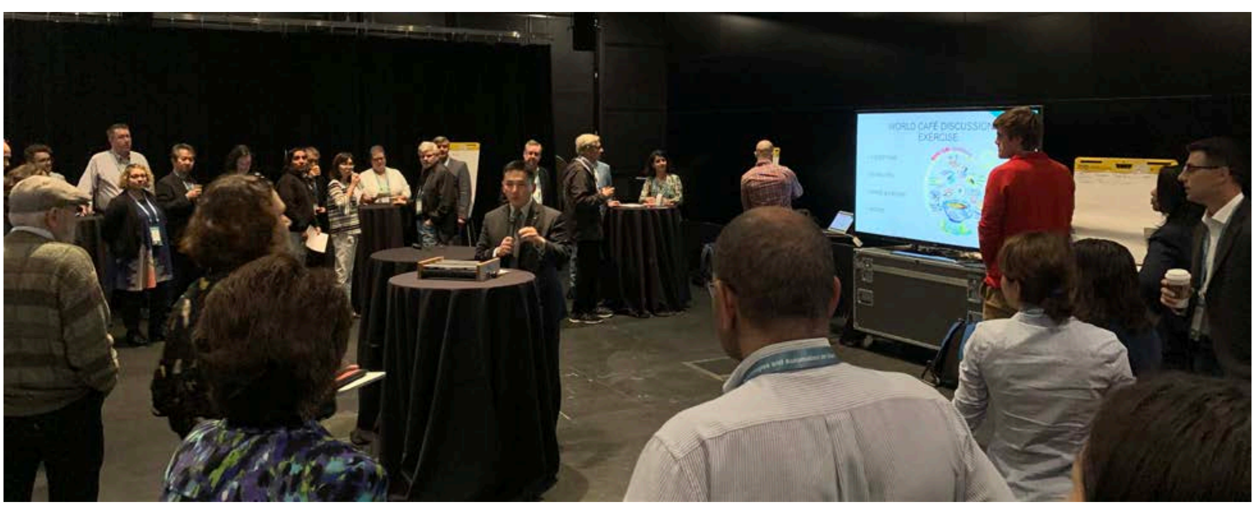

Question 1. Given current states of vehicle technology, what are the most pressing research needs concerning user interactions with automated systems? How can these research needs be addressed?

Following from previous years, the 2019 forum brought to light many important research questions related to driver interactions with and education concerning automated technology as well as the impact on the broader road system, including other road users. The list below attempts to capture the most prominent questions, though it is noted that similar issues arose in the discussion of the other world café questions and so viable and pressing research needs can also be found in subsequent descriptions. With respect to the research needs articulated below, the groups expressed that many merited broader consideration in the context of different levels of automation, different users (individual differences), and contexts (e.g., urban vs. rural), etc. That is, there are numerous interactions that need to be explored and/or accounted for under these topics (e.g., looking at training and education across a wide range of driving populations and experiences). The following list has been group by similar topics although is not presented in order of importance:

<u>Education, Understanding, and System Feedback</u>

1) Given the general lack of knowledge and education (mental models) concerning ADAS and AVs:
    a. What are the training requirements for these systems? What are critical thresholds along the learning curve (i.e., those that can translate to safe and appropriate use)?
    b. What is the role and impact of system branding and nomenclature?
    c. How does advertising impact driver perceptions and understanding of systems?
2) How does the driver's mental model of system operations change or vary over time?
3) What are the most promising approaches for 2-way communication between drivers and systems, including appropriate and transparent feedback (e.g., driver state monitoring, passive and active feedback, etc.)?
4) What is the role and impact of customizable interfaces? What can vehicles learn about their own drivers (e.g., driver history, style, preference) and how can this information be used to tune the driving algorithms?
5) What system properties help establish user trust (calibrated trust)?
6) What technologies are most challenging to drivers? What mistakes are made?

<u>System Use and Takeovers</u>

7) Concerning disuse of automated systems: Why do drivers turn systems on or off? What is the difference between and impact of intentional (opt-in/opt-out) versus unintentional disuse? How does this vary by context? What properties of the system map onto these decisions?

8) When automation technology fails, what does the driver do (what behaviors, inputs)? What vehicle and contextual data can be gathered to support finger grained analysis of these scenarios?
9) What factors impact driver takeover quality and timing? Can a human intervene safely when they have mind-wandered? How to quantify driver engagement?
10) With respect to graceful system exits, what is a safe-state maneuver and is it safe enough? What are the necessary preconditions or state of readiness of the fallback user?
11) For system takeover, what are the characteristics of appropriate and safe transitions? Do drivers know when to take over? How does the driver's level of trust impact their taking over the system?
12) Under what conditions are systems being used inappropriately (misuse) and what factors are implicated (e.g., who is driving, where are they driving, what are they doing, how often are they doing it)?
13) To what degree is AV-related risk tolerated by different driving populations? What factors serve to encourage risk tolerance or offset perceived risks (e.g., fuel/time savings)?

<u>Other Road Users and Impacts</u>

14) How will manual drivers interact with automation-equipped vehicles?
15) What kinds of interactions will take place between AVs, drones, scooters, bikes, pedestrians and which ones are safety-critical? What is the future of intersection management (e.g., no lights/stopping)?
16) How can AV technology be used to support first responders?
17) How will AVs affect economic development? (e.g., truck stops, fuel, public transit, transport goods)

Addressing the Research Needs

The question of how these research needs might be addressed was primarily answered in terms of the different methodological approaches or appropriate measures that could be applied. These were often characterized in terms of whether the approaches shed insight into "what" a driver was doing with a given system versus "why" they were doing it. Groups also called for the careful definition of different measures.

In terms of approaches, a broad range was noted:

- On road-data (including field operational tests and data from fleets)
- On-board data monitoring, including CAN bus
- Naturalistic driving studies (new or leveraging existing data sets)
- Driving simulation

- Ethnographic and other observational approaches (i.e., what driver is doing with system)
- Survey approaches or other subjective measures (e.g., preferences, trust, workload)
- Focus groups
- Interviews (i.e., why driver is (not) doing something)
- Verbal (think aloud) protocols

The specific method notwithstanding, the groups also emphasized the importance of developing a better understanding of how interactions, attitudes, trust, performance, and safety change over time, calling for more longitudinal approaches in data collection. Such methods can also help evaluate the process by which drivers learn how to use the systems and develop their mental models. Studies should also be more deliberate in accounting for different driver populations as well as differences in driver abilities and other individual differences. Moreover, it is important to reiterate that the different types or levels of automation warrant appropriate consideration of the approach as well as the associated measures.

Although not an exhaustive list, the following measures were brought up in the discussion of Question 1 (and see subsequent questions for a parallel treatment):

- Accounting of basic interactions between driver and the system – when systems are engaged or disengaged by drivers and how effectively; what happens when a system issues a takeover request or disengages?
- Behaviors and interactions with and by other road users, including pedestrians, cyclists, and drivers of non-AVs.
- More expansive physiological measures of the driver, including brain activity (e.g., electroencephalography, EEG; stress/arousal, eye tracking).
- Subjective and objective indices of trust (e.g., self-ratings, compliance with system, response time).

Although not a specific measure per se, the groups noted that more thought and effort should be devoted to determining how different measures can be consolidated or linked to support modeling and other algorithmic approaches. Streamlining data collection was also touted as an important area for development. The groups also noted that the use of crash-likely locations or scenarios could help shape the scope of some data collection efforts. Lastly, the groups discussed some of the challenges of data sharing (e.g., by OEMs) and standardization of “open source” data. They also speculated on the utility and prospects of a reporting system for near misses (akin to the Aviation Safety Reporting System used in aviation).

Question 2. What are the most important measures of safety in the context of vehicle automation, and what are appropriate targets or thresholds? (Consider both pre-deployment and post-deployment.)

The groups considered this question in the context of some guiding principles: do no harm, arrive well, and feel good (comfort). They suggested that safety needed to consider a wide array of road users (drivers, pedestrians, cyclists, etc.) and that consideration be given to crashes that might have been caused or contributed to by automated vehicles, even if these vehicles were ultimately not involved. Ideally, candidate measures would embody traditional qualities: reliable, diagnostic, repeatable, and accurate. The groups also appreciated that safety and its related measurement would be multidimensional; it would be impossible to distill this down to a single value. For many measures, the importance of the denominator was cited; that is, a raw count absent any information about exposure (e.g., distance travelled) would not be informative. The groups touted the need to look for convergence across multiple measures, while acknowledging the strengths and drawbacks of different approaches in isolation.

Important metrics were thought of as being classified as either pre- or post-deployment. Pre-deployment metrics can be gathered from numerous settings/approaches: Test track, simulations (microscopic and mesoscopic traffic simulations; human-in-the-loop driving simulations), field tests, mixed reality, NDS data, and information gleaned through the design cycle. Post-deployment approaches might entail vehicle data (what vehicle sees, what vehicle interprets, vehicle's 'mental model,' data regarding the surrounding environment), on-road data, black box data (electronic data recorder, EDR), telematics (including vehicle kinematics, driver behavior) as well as crashes (from on-road as well as from relevant databases). Measures were further distinguished by the target or source of the measurement, for example, whether the driver or user of the automated system, the vehicle or system itself, or from external sources such as the infrastructure, environment or other road users.

Many of the safety metrics had clear relevance to fatalities, injuries, and property damage:

- Crashes
- Near misses, including measures of time-to-contact and evasive maneuvers (system kinematics)
- Safety critical events (broadly defined)
- Crashes that were avoided by the system (not necessarily equated with near misses or safety critical events)
- Hazard detection/perception and response (e.g., system/software situation awareness)
- During takeover, time to stop or stabilize the vehicle
- Safety and other impacts of graceful (or not so graceful) system exits

Although not mutually exclusive to safety metrics, a number of performance measures were also discussed:

- Number and rate of takeover requests and critical disengagements
- Takeover time and quality
- Driver engagement and disengagement
- Driver situation awareness, eyes on the road
- Driver knowledge or mental model of the system, mode, operational design domain (ODD)
- Knowledge or expectations of other road users regarding AV
- Attitudes and behaviors of other road users
- Driver errors
- Level or currency of software updates or certification
- Ease of use, understandability
- Unintended consequences, negative behavioral adaptation

Additionally, there were a variety of measures that related to the overall system use and system status in different situations. Again, many of these relate to measures listed above or could be considered important covariates or classifiers in data analysis:

- Overall use pattern, time or proportion of drive with system engaged
- Usage inside and outside of the system's ODD
- Miles or trips incident-free
- Driver state (e.g., through driver state monitoring systems)
- System status, e.g., automation availability, mode/state, warning states
- Contextual information, including weather, traffic, and road conditions, etc.

Safety Thresholds and Benchmarking

With respect to benchmarking, the discussions were understandably high level. The groups acknowledged several challenges related to the balance of safety, efficiency, economics, public acceptance, and other factors. There was little debate that, ultimately, a reduction in the number and severity of crashes over time was a laudable aim (e.g., Vision Zero), but that the entire road system should be considered (to account for new types of crashes and other road users). Moreover, different types of crashes might be weighed differently in decisions about benchmarks (e.g., injury versus non-injury crashes). The groups also noted that careful appraisal should be made of the differential impact across crashes involving diverse causal factors, such as speed, fatigue, alcohol and other drugs, and distraction. Generally speaking, improvements to the status quo were seen favorably, although less substantial gains in crash reduction in the near term could be undermined by public acceptance and trust of the technology. Public acceptance could be gleaned directly through opinion polling, perhaps with a sensitivity to a potential "tipping point," and indirectly through sales and market penetration. Trust was considered and could easily be assessed against several potential benchmarks, including taxi and rideshare drives, school buses, and other professional drivers.

The groups also raised the question as to who would be the ultimate authority on safety thresholds—whether governments, manufacturers, or the public (through demand). They also noted that a boarder consideration of safety metrics and benchmarks would also need to factor in the nature of future deployments. Changes to the ownership and fleet model could have a profound impact.

Question 3. How can we improve our understanding of driver interaction with vehicle automation? What methods can we use?

The discussions of this question tended to intersect with the research needs noted in Question 1 as well as the methods and measures cited in Question 2. Despite some overlap in coverage, some important nuance emerged as well.

In order to improve our understanding of driver interactions with vehicle automation, the groups noted several areas begging future research or approaches that could enhance our current understanding. These are presented in no particular order of importance:

- Determine what about driver understanding can be handled through design so that it is intuitive for drivers. In doing so, articulate what material is appropriate and/or necessary for different forms of driver training or education and seek related examples or parables to help teach or reinforce. Determine when training should be administered, whether at the point of sale, as part of licensing/renewal (or technology/vehicle certification), or at other points in time and if and when training should be revisited (in light of innovation in AV over time and changes in driver comfort/trust). Leverage an understanding of an individual's deficits for training purposes (mapping mental model).
- Determine the accuracy or adequacy of driver mental models for AV systems relative to the system functionality and limitations. Devise appropriate means to reduce the gaps in driver mental models through system feedback, training, branding, etc.
- Embed social scientists on teams with engineers who research and/or develop AV systems. In doing so, try to bridge gap between system designers and the end users.
- Develop a better mapping for driver engagement while using AV technology. Discover why drivers choose to engage or disengage in particular moments or situations; and understand how to improve driver readiness for takeover. Understand what properties of the system design that map onto those decisions to engage or disengage. Consider how the external environment influences the driver's choice to use automation or to disengage it.
- Consider standardization of how systems function across vehicle, and how they are mapped through HMI design, to promote more consistently across users in terms of expectations concerning functions and limitations.

- Promote more longitudinal research on drivers of different backgrounds, experience levels, income, age, demographics, etc. Need to understand not only drivers but riders (e.g., with higher levels of automation). Through research, encourage multitude of approaches but with emphasis on real world settings and situations.
- Further explore the interaction between AV-equipped vehicles and other non-AV drivers.
- Invest in the psychology of relationships that real users have with their AV, including some of the emotional dimensions. Understand the emotions experienced and how they impact consumer choice to engage or disengage automation.
- Explore approaches to incentivize people to use and buy AV-equipped vehicles.

Methods

A variety of methods, approaches and measures were discussed in terms of potential for improving our understanding of driver interactions with automation. Many such techniques shed insight into what a driver is doing in situ while others help address the question of why drivers are doing these things. With respect to understanding what drivers are doing with automated systems, observational approaches including naturalistic driving approaches can help characterize the activities, behaviors and their relative frequency and context. These could also capture the types and frequencies of user interactions with the system, including the level of involvement/input. In contrast, other qualitative or subjective approaches might be helpful in understanding the motivations behind certain behaviors (e.g., interviews, focus groups, surveys).

With respect to observational approaches, ideally these can be augmented with vehicle-based data (e.g., CanBus, LIDAR, radar, audio, video of areas around the vehicle—not just limited to driver/vehicle). Some approaches might bridge observational and qualitative approaches by later asking drivers about specific occurrences. For example, using the pre-recorded video data to show their own events or behaviors (e.g., “what were you thinking then?” or “why did you choose to engage the system then?”). Talk aloud paradigms carried out in situ might also be a means of enhancing naturalistic data.

Other laboratory settings or approaches can offer insight on both dimensions (what and why), affording more control in attempts to isolate effects or examine more specific use case scenarios, while also provided opportunity to supplement with subjective and questionnaire data or interviews. These include simulation, computer-based testing, and Wizard of Oz approaches. Lab-based studies are more likely to be amenable to an array of physiological data and runtime assessment of workload (e.g., NASA-TLX) and perceptions (frustration, annoyance, etc.), which can greatly augment the output—especially for articulating impacts that could be sub-conscious.

Cognitive task analyses (or related approaches) can be useful in detailing the nature and sequence of interactions involved in a given driving or AV-related task. These can also serve to emphasize potential error scenarios.

Perhaps more importantly is the potential to capitalize on certain elements such as driver state monitoring systems, such that the actual fleet becomes a potential platform for research. The abundant data that can be gleaned from a fleet raises new potential for data mining and machine/deep learning approaches. These approaches, in turn, can help define surrogates for safety critical incidents for research purposes and help develop tools for filtering these data efficiently. It was also noted that vehicles themselves could become data collection platforms, where they learn about their drivers: driving habits, demos, literacy, and numeracy and thereby customize the interface and system function accordingly.

In isolation, each approach has its benefits and drawbacks, so ideally, a multitude of approaches can be deployed. Then convergence or divergence across methods can be examined. Within the AV space, one important caveat that the groups noted is the degree of experience or exposure that users have with the technology a priori. It is not uncommon for participants or respondents in a study to have zero familiarity with the systems being addressed. This calls for certain care in drawing firm conclusions.

Regardless of the setting or experimental approach, the discussions called for a more comprehensive assessment or treatment of factors that can influence behavior or outcomes, including a driver's knowledge and trust in the system as well as a more in-depth accounting of individual differences. Importantly, there were broad calls for more longitudinal studies to better understand the learning curve, explore how people progress/digress in knowledge, trust (subjective and objective) and appropriate use with increased exposure.

Question 4. What are the key focus areas related to safety in the next few years to continue to advance vehicle automation?

Three broader themes emerged from the discussion, related to the operationalization of safety, other important concepts, and a variety of use case scenarios. Again, there is clear overlap with some of the topics that emerged in discussions of the previous questions.

Definitions, Benchmarks, and Metrics
At the forefront, discussants recognized the importance of defining "safety" and identifying measures appropriate for this definition. It was acknowledged that safety and safety-related measures associated with conventional vehicles and drivers might not necessarily apply for AVs.

To date, much of the views of safety have been based on the fatality and crash rates of AV to human drivers. The groups offered up a broader conceptual framework that could be further refined and developed to guide the operationalization of safety (and safety benefits) within the following spaces: driver-vehicle interactions, vehicle-environment interactions, technical capabilities, socio-cultural factors, and oversight and regulations.

The next logical challenge related to setting thresholds or benchmarks for “what is safe enough” (see also Question 2 above). The general consensus was that AV safety will asymptote at some point, such that 100% safety will be an unlikely achievement — barring a long time frame and significant expenditure of resources. For example, in any road system, there will be instances where there is no or very little predictability as to when an accident will happen (someone randomly jumping out in front of AV). Such crashes will be very difficult to prevent due to randomness and the allowance of time and space. As such, the focus should be on determining what is safe enough for public use such that near term gains can be reaped.

With regards to benchmarking, there was some suggestion that the systems should be at least as safe as the human driver. An analogy was offered: “When do you take your keys away from elderly grandparents?” One possible answer is “When you are uncomfortable with them driving your kids.” Stemming from this, engineers might view safety using the same or similar lens: would I let my kids use this (AV) system? Of course, beneath the surface of these broader questions and analogies is how do we actually compare AV safety compared to the safety of human drivers?

From a system’s perspective, the AV technology should also be at least better than the last crash, therefore there is constant improvement. Importantly, the groups acknowledged that forecasting and benchmarking of automated systems needed to consider the different levels of automation: where do the different systems or features fall on the “safety” graph? Benchmarking should also take into consideration how safe or successful systems need to be in order to be accepted and used by the driving public.

Other Key Factors

While safety was the primary consideration in the discussions, the importance of other factors was also acknowledged, especially in as much as these factors could indirectly impact safety, e.g., through uptake and market penetration. The groups highlighted the role of other desirable outcomes related to AVs, including efficiency, mobility, public acceptance, and willingness to buy, as well as the pleasure, comfort, and satisfaction of drivers and other road users. All of these should also be considered in addition to safety, although it is clear that there will be tradeoffs along many dimensions (e.g., safety and efficiency/throughput). The groups also noted the importance of other upstream factors that could impact acceptance and uptake, such as inclusivity, social equitableness, education, privacy, and cybersecurity.

Closing Remarks

This forum, the third in the series, was convened with the ultimate aim of promoting engagement and discussion amongst key stakeholders from research, industry, government and other entities. As technology progresses, it is important to keep sight on those research questions, design considerations and other core safety issues that will affect the success of these systems. Although we are constantly apprised of new and important research on the topic, the discussion and outcomes from this forum help to reinforce that still much more work is needed. Let us stay engaged.

Appendix A: List of Organizations That Participated in the 2019 Forum

AAA Allied Group
AAA Foundation for Traffic Safety
AAA National
AAA Oregon/Idaho
AAA Western & Central New York
Adventure Cycling Association
Alliance of Automobile Manufacturers
Arizona Commerce Authority
Auto Club of Southern California
Dynamic Research Inc.
Exponent
European Commission
Google
Hart Solutions, LLP
Insurance Institute for Highway Safety
International Association of Traffic and Safety Sciences
National Highway Traffic Safety Administration
National Institutes of Health / National Institute on Aging
National Safety Council
Osaka University
Purdue University
Red Scientific
San Diego State University
Texas A&M Transportation Institute
Transport Canada
University of Alabama Center for Advanced Public Safety
University of California Berkeley
University of California Berkeley Partners for Advanced Transportation Technology
University of California Davis Institute of Transportation Studies
University of California San Diego
University of California San Diego Design Lab
University of Iowa National Advanced Driving Simulator
University of Kansas
University of Massachusetts Amherst
University of South Carolina
University of Utah
University of Washington
Virginia Tech Transportation Institute
Westat
WSP

Appendix B: 2019 Forum Agenda

AAA Foundation for Traffic Safety — UC San Diego

2019 Forum: Impact of Vehicle Technologies and Automation on Users – Design & Safety Implications

MONDAY, NOVEMBER 4	
8:30 a.m.	Registration and Continental Breakfast
9:00 a.m.	Welcome and Forum Objectives
9:30 a.m.	Panel 1: State of Vehicle Technologies and Automation
10:45 a.m.	Refreshment Break
11:00 a.m.	Panel 2: Design Recommendations and Considerations
12:00 p.m.	Lunch and Networking
1:00 p.m.	Facilitated Panel 2 Discussion
1:45 p.m.	Panel 3: Implications to Future Safety and Operations
2:45 p.m.	Refreshment Break
3:05 p.m.	Facilitated Panel 3 Discussion
3:50 p.m.	Plan for Day Two Discussion
4:00 p.m.	End of Day One
5:30 p.m.	Transportation Departs from Hyatt Regency La Jolla for Forum Dinner
6:00 p.m. – 8:30 p.m.	Forum Dinner – Viewpoint Brewing Co.

TUESDAY, NOVEMBER 5	
8:30 a.m.	Continental Breakfast
9:00 a.m.	World Café Discussion Exercise
10:30 a.m.	Refreshment Break
11:00 a.m.	Small Group Reports and Feedback
11:50 a.m.	Closing Comments

Please refer to the Forum Mobile App for more detailed information.

FPO