SAE INTERNATIONAL

SAE EDGE™
RESEARCH REPORT

Unsettled Topics Concerning Human and Autonomous Vehicle Interaction

Rahul Razdan, PhD

SAE EDGE™
RESEARCH REPORT

Unsettled Topics Concerning Human and Autonomous Vehicle Interaction

Rahul Razdan, PhD
Raztech Research Institute

EDGE DEVELOPMENT TEAM

Ken Ford, PhD, *Florida Institute for Human and Machine Cognition*

Justin Mason, PhD, *University of Florida*

Sanjay Ranka, PhD, *University of Florida*

Brian Reimer, PhD, *New England University Transportation Center at MIT*

Raivo Sell, PhD, *Tallinn University of Technology*

Edward Straub, PhD, *Automated Vehicle Safety Consortium*

Kevin Tang, PhD, *University of Florida*

SAE INTERNATIONAL®

Warrendale, Pennsylvania, USA

About the Publisher

SAE International® is a global association of more than 128,000 engineers and related technical experts in the aerospace, automotive and commercial-vehicle industries. Our core competencies are life-long learning and voluntary consensus standards development. Visit sae.org

SAE EDGE™ Research Report Disclaimer

SAE EDGE™ Research Reports focus on topics that are dynamic, in which knowledge is incomplete, and which have yet to be standardized. They represent the collective wisdom of a group of experts and serve as a practical guide to the reader in understanding unsettled subject matter. They are not meant to provide a recommended practice or protocol. The experts have assembled as a community of practitioners to contribute and collectivize their thoughts and points of view; these are not the positions of the institutions or businesses with which they are affiliated, nor is one contributor's perspective advanced over others. SAE EDGE™ Research Reports are the property of SAE International and SAE alone is responsible for their content.

About This Publication

SAE EDGE™ Research Reports provide state-of-the-art and state-of-industry examinations of the most significant topics in mobility engineering. SAE EDGE™ contributors are experts from research, academia, and industry who have come together to explore and define the most critical advancements, challenges, and future direction in areas such as vehicle automation, unmanned aircraft, cybersecurity, advanced propulsion, advanced manufacturing, Internet of Things, and connectivity.

Related Resources

SAE MOBILUS® Automated & Connected Knowledge Hub
https://saemobilus.sae.org/automated-connected/

SAE Team

Frank Menchaca, Chief Growth Officer
Michael Thompson, Director of Standards, Information and Research Publications
Monica Nogueira, Director of Content Acquisition and Development
Beth Ellen Dibeler, Product Manager
William Kucinski, Managing Technical Editor

Printed in USA

EPR2020025
ISSN 2640-3536
e-ISSN 2640-3544
ISBN 978-1-4686-0279-1

To purchase bulk quantities, please contact: SAE Customer Service

E-mail: CustomerService@sae.org
Phone: 877-606-7323 (*inside USA and Canada*)
+1-724-776-4970 (*outside USA*)
Fax: +1-724-776-0790

https://www.sae.org/publications/edge-research-reports

About the Editor

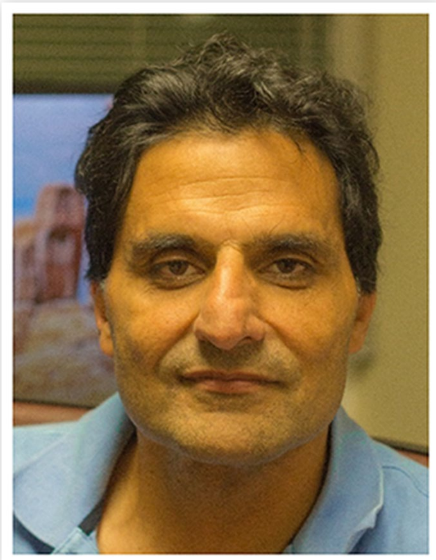

Dr. Rahul Razdan is a seasoned scientist and business executive, who has held significant roles in the world of academia, startups, and Fortune 500 companies. In corporate roles, he was the Senior Vice President of Strategy of Flextronics, General Manager of the System and Functional Verification Business at Cadence Design Systems, and Alpha CPU architect at Digital Equipment Corporation. He has built successful startups in areas such wireless power, machine learning, and low-power electronics design.

In academia, beyond a large range of worldwide collaborations, he maintains academic appointments at the University of Florida, Florida Polytechnic University, and Institute for Human Machine Cognition. Today, he is CEO of Raztech Research Institute, a technology-based research institute which incubates ideas at the intersection of technology and society. Active projects for the institute include the Autonomous Vehicle Verification Consortium, NextGen Edu, and Anew Design Automation. He has a bachelor's (1984) and master's degree (1985) from Carnegie Mellon University and a doctorate from Harvard University (1994).

contents

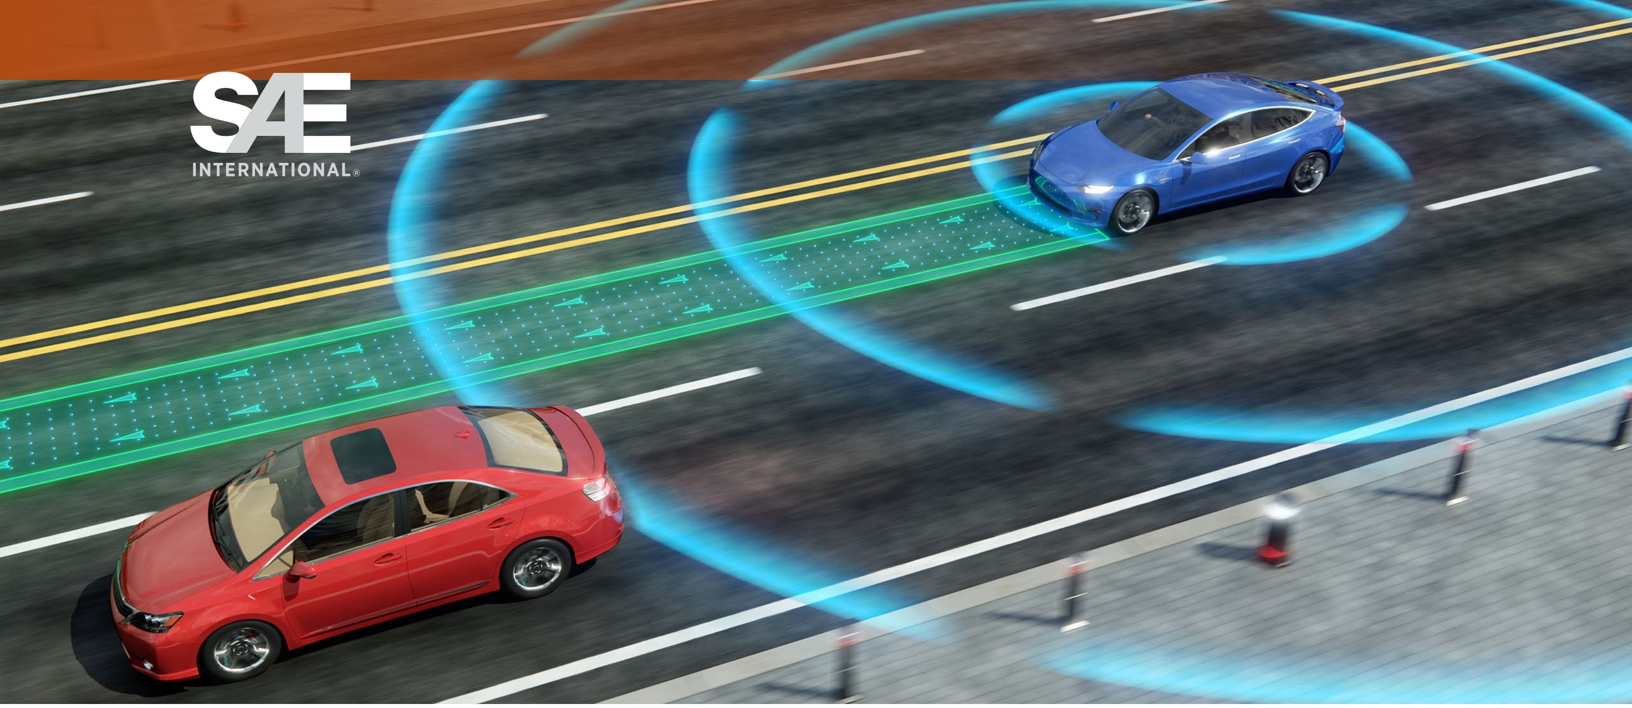

Unsettled Topics Concerning Human and Autonomous Vehicle Interaction

Abstract

Autonomous technology has the potential to greatly benefit personal transportation, last-mile delivery, logistics, and many other mobility applications. In many of these applications, the mobility infrastructure is a shared resource in which all the players must cooperate. In fact, the driving task has been described as a "tango" where we—as humans—cooperate to enable a robust transportation system. Can autonomous systems participate in this tango? Does that even make sense?

This report will examine the current interaction points between humans and autonomous systems, the shortcomings of the current state of these systems with a particular focus on advanced driver assistance systems, the requirements for human-machine interfaces as imposed by human perception, and finally, the progress being made to close the gap.

NOTE: SAE EDGE™ Research Reports are intended to identify and illuminate key issues in emerging, but still unsettled, technologies of interest to the mobility industry. The goal of SAE EDGE™ Research Reports is to stimulate discussion and work in the hope of promoting and speeding resolution of identified issues. SAE EDGE™ Research Reports are not intended to resolve the challenges they identify or close any topic to further scrutiny.

RAHUL RAZDAN, PhD
Raztech Research Institute

EDGE Development Team

Ken Ford, PhD, *Florida Institute for Human and Machine Cognition*
Justin Mason, PhD, *University of Florida*
Sanjay Ranka, PhD, *University of Florida*
Brian Reimer, PhD, *New England University Transportation Center at MIT*
Raivo Sell, PhD, *Tallinn University of Technology*
Edward Straub, PhD, *Automated Vehicle Safety Consortium*
Kevin Tang, PhD, *University of Florida*

ISSN 2640-3536

Introduction

State of the Industry

Autonomous technology can provide, accessibility, and new-use models. But, to provide this utility, there must be an interaction between humans and autonomous systems. Many human-machine reports focus on the question of "Will humans accept autonomous technology?" and typically involve polling perceptions of the general population. However, the results tend to be fickle because initial perceptions are based on third-party (media) accounts and true acceptance cannot be tested until there are actual engagements with the technology. As a recent example, early in its lifecycle, e-commerce technology was perceived as risky; yet over a period of time, the perception of risk faded. With that point of view in mind, the focus of this report is to explore the deeper interaction between humans and automated vehicles (AVs), as it pertains to the driving task: the "tango" of driving between humans and its implications for AVs [1]. In this dance, the participants know the rules, but may never have met each other (Figure 1). Thus, the task is not quite teaming (e.g., human-machine manufacturing) but more dynamic ad hoc cooperation.

This report is split into three sections. The first explores the current state-of-art for autonomous mobility solutions with a focus on human-machine interaction, its current limitations, and the consequences of these limitations. The second section examines human perception and communication as it relates to the driving task. Finally, the third section outlines current efforts to improve human-machine interaction processes.

Autonomous Vehicle Perception Mechanisms

The general structure of automated driving systems (ADS) consists of sensors—radar, light detection and ranging (LiDAR), and camera—which feed perception systems that build an accurate view of the external environment (Figure 2). Based on this view and goals for movement, an AV (equipped with an ADS) navigates itself through the transportation infrastructure. At a high level, LiDAR provides accurate

Francesco Abrignani/Shutterstock.com

FIGURE 1. **Driving is a cooperative task akin to a tango.**

FIGURE 2. **An AV using radar to perceive its surroundings.**

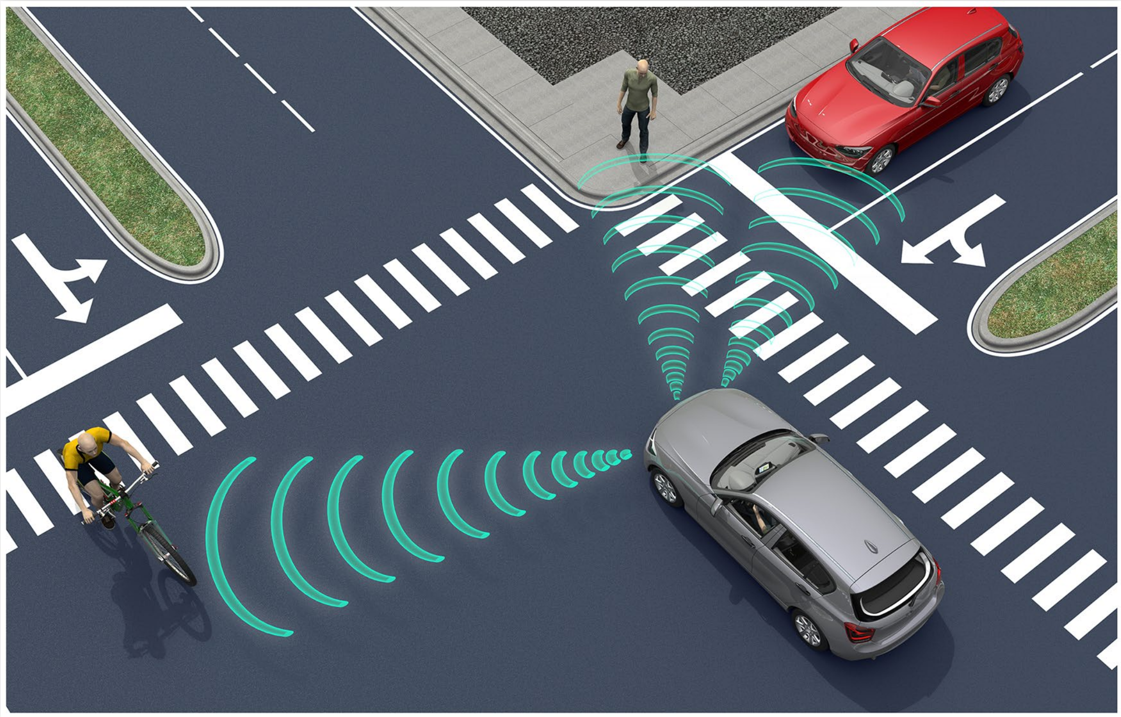

distance calculation, radar provides accurate velocity information, and cameras provide the pixels, which are at the center of object recognition.

The industry has aligned itself around SAE International's Levels of Driving Automation framework [2]. Currently, there are millions of vehicles equipped with advanced driver assistance systems (ADAS). These systems offer many features, but three of particular interest are collision avoidance, lane following, and automated parking, as they can take active control of the automobile. Beyond ADAS, there are several use models in various stages of testing, including "robo-taxis" in geo-fenced areas, truck convoying, and public transportation shuttles.

Today, the open question in the AV community is "What level of validation is sufficient to be acceptable to the public?" Some would more coarsely say, "Humans kill over 40,000 people yearly, but one AV accident seems to be a show-stopper." This group would advocate a more aggressive deployment approach with the notion that the result *must* be better than the human alternative, to which the counterargument is "**What accidents will AVs get into which humans could have easily avoided?**"

This discussion is a subset of a larger discussion in the broader artificial intelligence (AI) community relative to the appropriate role of AI systems. As pointed out in "No AI Is an Island: The Case for Teaming Intelligence," the effectiveness of AI systems is a function of their ability to add value in the context of higher-level human-machine teaming [3]. Further, AI must do so with a type of intelligence, which is fundamentally different to that of humans. Thus, human-machine interaction models are the key. Specifically, AVs must set appropriate expectations for their behavior and part of the behavior definition must be the communication paradigm in regard to humans. Let us examine both topics for transportation.

Expectation and Communication

Today, AVs (Level 2 and above) are placed into the public road system with a footprint, which is exactly that of a human. **With this use model, AVs inherit all the expectation attributes of a human driver**. Is this reasonable? "On computational wings: The Prospects and Putative Perils of AI" points out that imbuing AI systems with human characteristics is similar to asking planes to fly by flapping their wings [4]. Humans and machines each have their own strengths and weaknesses.

Following this line of thinking, AVs must establish expectations of their behavior on the roads, which are distinct from those of humans. Further, this persona must set expectations clearly through a well-defined operational design domain (ODD). Based on this ODD, one can build a clear validation and verification framework. Today, clear ODDs do not exist for any SAE-defined level of automation. Even ADAS, the lowest level of AV automation has no clear ODD definitions.

As pointed out by "Tomorrow's Human-Machine Design Tools: From Levels of Automation to Interdependencies," automation-level definitions defined independently of human-machine teaming behavior limit the utility of these systems [5].

Since we are working with humans, a long ODD manual, which exhaustively outlines activity, is unlikely to suffice. Engineering products such as the Intel Responsibility-Sensitive Safety model or UL safety standards provide a great deal of utility for engineering safety but do not pass the human-machine interaction model test [6]. Rather, one needs clear, short, and easily understood concepts such that the broader public can absorb them quickly and efficiently. The world of information technology (IT) learned to use this approach with reuse of concepts such as "file" and "window." What is the equivalent for AVs? This is an open question right now and cause of much of angst. An example like truck convoying is an interesting start [7]. Convoying is a high-level behavior concept, which is immediately and easily understood by humans.

Further—to tango—observability, predictability, and detectability are the necessary minimum properties of teamwork, whether for humans or machines [8]. In performing the driving task, humans communicate direction, intention, and risk (as in the frightened face of a young driver) constantly with each other. This communication, combined with contextual understanding of the situation, allows the transportation system to function efficiently.

One might reasonably ask "How well do current AVs employ these properties?" Well, AVs do not communicate any of these attributes in a visible way to third parties. The lack of these communication mechanisms is a significant drawback for AVs. This brings us back to the main question: **When will AVs be accepted from a safety point of view?**

The answer: This will occur when there are clear expectations around AV behavior with high-level nonhuman connected ODDs; only then will the key elements of cooperative behavior (observability, predictability, and detectability) be respected. What is the result of this lack of expectations? The results can be seen in the early accident data from the California Department of Motor Vehicles' "Requirements for the Next-Generation Autonomous Vehicle Ecosystem," where the vast majority of accidents are low-speed rear-end collisions of humans hitting AVs [9]. It appears that they were expecting a driver with human characteristics.

Operational Design Disciplines and Measurable Safety

Assuming one has a clear ODD, the next step is obtaining clear, measurable metrics for adherence within the ODD definition (Figure 3).

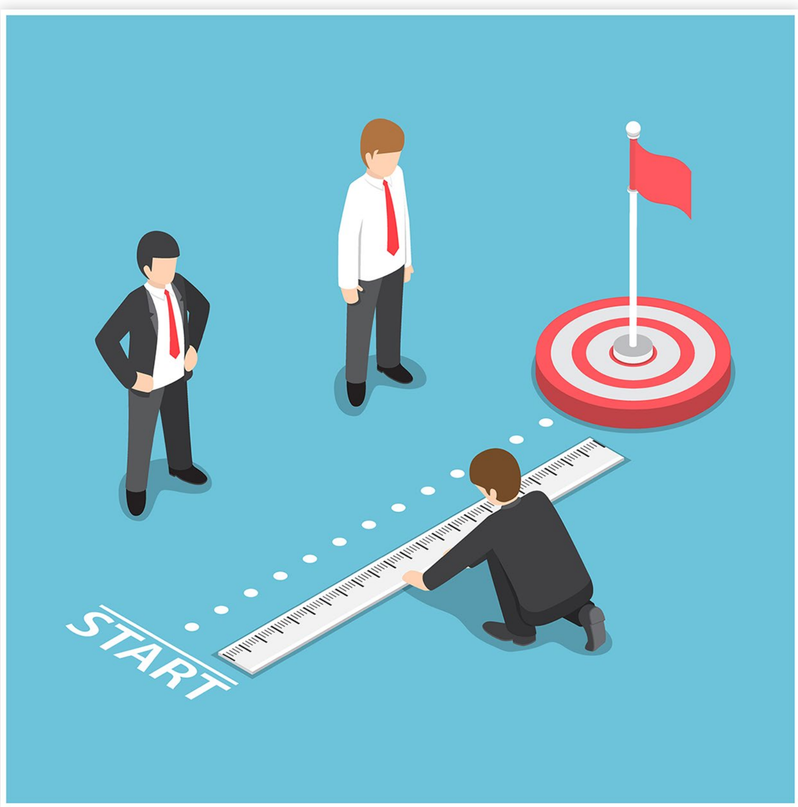

Jiw Ingka/Shutterstock

FIGURE 3. **The importance of measurement and metrics.**

"If you can't measure it, you can't improve it," said renowned management guru Peter Drucker, meaning that one cannot define success unless there are metrics, which are defined and tracked [10]. With metrics, one can track progress and adjust development to produce positive movement. Without them, one is always in a haze of doubt. Interestingly, the lack of clear metrics not only hurts the producer (i.e., What is progress?), the haze of uncertainty causes the buyer to impose a discount on the product's value.

Hard metrics are not new to the automotive industry. Fuel efficiency and performance are just two of a host of hard metrics published and tracked by the industry. However, ADAS safety systems are one place where there is a remarkable lack of clear metrics. Compounding that is public confusion surrounding ADAS features, which has led to consumer advocacy organizations (i.e., AAA, JD Power, Consumer Reports, and the National Safety Council) to publish a common naming structure consisting of four major categories (driving control assistance, collision warnings, collision intervention, and parking assistance) [11].

Where Are the Regulators?

To date, regulators have taken a laissez-faire approach. The US Department of Transportation's "Ensuring American Leadership in Automated Vehicle Technologies: Automated Vehicles 4.0" or "AV 4.0" report does not offer any guidance on ADAS despite the fact that systems, such as driving control assistance and collision intervention, can actively engage in the driving task [12].

What about Congress?

Indeed, Congress has been working on a bipartisan bill on autonomous vehicles. The bill contains many fine focus points, especially on the topic of cybersecurity, but it does not address ADAS regulation [13]. Perhaps the reason for this lack of concern is that these systems just work. But, how well do they work? The Insurance Highway Safety Institute has studied the effectiveness of ADAS and found that tasks such as "active lane-keeping" were a challenge for the commercially available vehicles, and conflict-avoidance systems are getting better [14]. Thus, it appears that there is some cause for concern.

In summary, what is the result of the lack of hard metrics?

1. **Customer Value:** The first predictable result is that the customer does not value the capability and will not pay much incrementally for it [15]. It is hard to pay for something where the value is "We will try to help you in an undefinable manner, but ultimately you are on your own."
2. **Insurance Risk:** Insurers are refusing to provide safety discounts for ADAS [16]. The original users of data mining claim that data does not yet support such discounts.
3. **Maintenance Cost:** ADAS consists of sophisticated sensors, which require care in terms of maintenance, and from a consumer point of view, this can be a "hidden" cost. This situation has become sufficiently problematic that Consumer Reports put out a bulletin warning customers of this issue [17].

Where Does This Leave ADAS?

Consumers do not highly value ADAS; it has hidden costs, and insurers are not sure it adds value to safety. Finally, the ultimate liability for ADAS is that it has not been litigated through the court system yet. The conventional wisdom is that the driver carries the liability but just one accident of premature breaking in collision intervention can remarkably change the liability picture.

Thus, it seems there is a dire need for some metrics, which define functions and expectations clearly. What might this look like? Let's consider the case of radar-based conflict avoidance systems:

1. **Curvature, Hills, and Valleys:** The expectation is that the conflict avoidance should work independent of the curvature of the road or whether you are on a hill or valley. Is this true? Where is the evidence?
2. **Road Debris or Potholes:** Should conflict avoidance be triggered based on road debris or potholes? If so, what exactly would trigger it? What should not trigger it?
3. **Driver Attentiveness:** Should conflict avoidance be triggered based on driver attentiveness? If so, what exactly constitutes driver attentiveness?
4. **Physical Environment:** The expectation is that conflict avoidance will work independent of the surrounding environment. It should not matter if you are in a tunnel or under a bridge. Yet, we know radar-based solutions are susceptible to the "ghost" problem due to reflection and interference [18]. Where are the metrics that show the ability to address these problems?

Some of the underlying technology and standardization efforts are forming to address this issue, such as the UN's effort with automated lane-keeping systems [19]. Also, standardization bodies such as the Association for Standardization of Automation and Measuring Systems are building a framework called OpenODD, which can reasonably model the interesting scenarios [20]. These two pieces form underlying components upon which to build reasonable engineering metrics.

Current ADAS Incentive Structures Building even just a "simple" robust conflict avoidance system is a nontrivial task. ADAS designers face the following difficult challenges:

1. **Radar:** In aerospace applications, radar is very effective because every object encountered is of high

interest. However, in the automotive context, radar signals reflect off the environment in such a manner to create blind spots and ghost objects.

2. **Camera:** It is easy enough to get a picture, but interpreting it is quite challenging. Today's machine learning (ML) algorithms are the equivalent of a Google Images search, so if the object has not been seen and recognized, the ADAS does not know what to do with it.
3. **Tracking:** Both radar and cameras are fixed on the car, which makes them susceptible to situations of road curvature, hills, and valleys from a field-of-view point of view.
4. **Driver intent:** Is the driver accelerating toward a car with the intent of a passing maneuver or is the driver distracted? How does the system decide?

With some understanding of the problem, one must sympathize for the ADAS designer's challenges. The response to these recognized challenges has been a permissive regulatory environment that encourages innovation. For ADAS, this permissiveness is provided by the statement that the driver is ultimately responsible. However, is this "crutch" impeding innovation?

Let's consider the behavior it seems to engender:

1. **Vague and unclear functionality:** What exactly can the customer count on when buying an ADAS-equipped vehicle?
2. **Lazy validation:** Partial validation and partial coverage are the norm in the industry today. How do we know that the next revision of the software is actually better than the last one from a safety point of view?

Both issues are covered by the loophole that the driver carries ultimate liability. Further, the current structure creates a predisposition against braking, since premature braking attaches liability to ADAS while a more permissive (aggressive) approach shifts liability to the driver.

Recommendations

1. **AV Persona Definitions:** Without setting clear expectations distinct from humans, autonomous systems will have difficulty making the safety argument. It is advisable that there is a standard that definitively and visibly separates AV operation from human operation.
2. **ODD Definitions:** With an AV persona, one must define an ODD. Since this ODD is to be leveraged by human drivers, it must be clear enough to be easily understood.
3. **Regulatory and Incentive Structures:** Regulations and associated incentive structures that do not push for ODD clarity and associated measurable safety are doing the industry a disservice. Either through self-regulation or governmental intervention, one needs clarity on measurable safety to build capability, which customers actually value.

Human Perception and Driving

For a tango to work, AVs must understand how their human partners perceive the world. This understanding is the basis of communication. Of course, learning the techniques that humans use for perception can potentially lead to better ways to build autonomous perception (Figure 4). Similar to AVs, humans use senses (vision, hearing, touch, taste, and smell) to navigate through the world. However, the most powerful organ in this perception task is the brain. In this context, while our focus is humans, it should be noted that this level of perception also extends to animals, and animals are part of the ecosystem that an AV must handle gracefully.

Research in animal communications conjures up images of Jane Goodall studying the great apes in Tanzania. A great deal of research in that field focuses on higher-level functions such as social interactions, animal cognition, or even emotional life. However, underneath those functions, there is a much more basic elemental aspect of intelligence consisting of simply operating in the physical world (Figure 5). This capability seems to be so innate that there is little research on its basic functionality. What are the key characteristics of this capability?

First, there is a large class of intelligence, which is connected to the insight of physical future behavior of other actors:

1. **Focus:** Through observation of eyes and sometimes ears, animals can interpret the direction of the other actors' attention.
2. **Body Positioning:** Sitting, walking, and running postures are all interpreted. Positions that can lead to a high degree of acceleration are of special interest.
3. **Gestures and Intent**: Facial as well as full body gestures are recognized for intent.
4. **Movement:** As opposed to recognition based on static images, humans use recognition based on movement. This capability is what allows a human to recognize a familiar person based on their gait from a distance.

FIGURE 4. **Basic senses of human perception.**

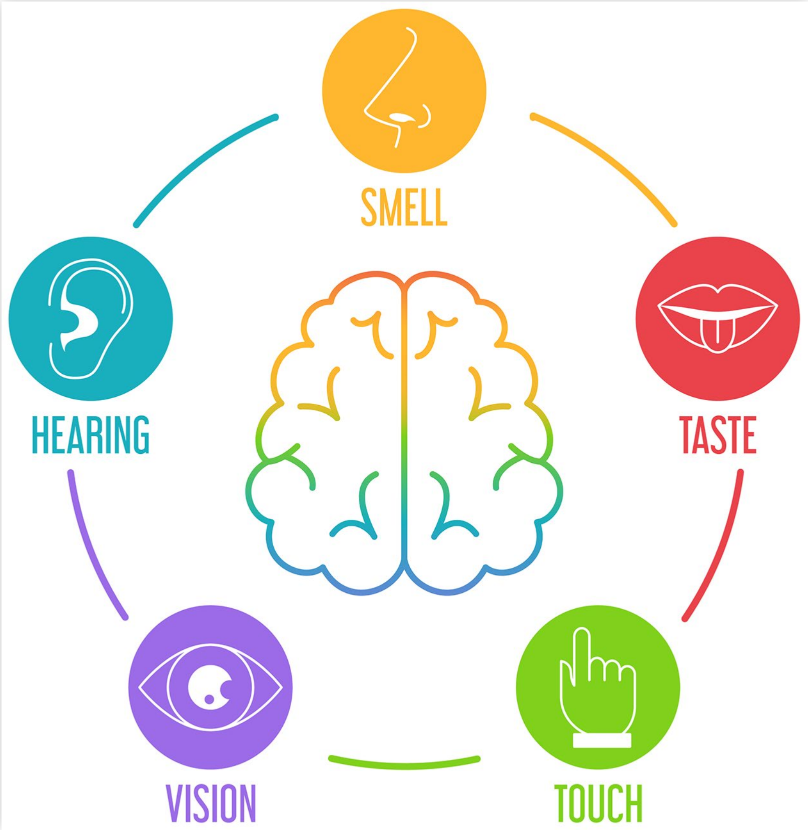

ImLucky/Shutterstock.com

FIGURE 5. **Animal perception, decision-making, and movement.**

master24/Shutterstock.com

Second, there is a basic calculation of physics. No, animals are not solving Newton's equations of motion, but animals innately maintain balance, calculate interception trajectories, and manage threats. In fact, the power of the understanding of physics in perception can be seen in illusions such as Adelson's Checker shadow or even in parachute landing training [21, 22].

Third, animals maintain a basic virtual mental model of their surrounding environment, and this model seems to generate an expectation, which drives perception [23]. The difference between perception and expectation combined with absolute distance seems to drive behavior. For example, a "surprise" awareness of an unknown object in close proximity drives a highly visceral response.

With this range of perception, mental modeling, and nonverbal cross-species communication, animals perform "threat assessment" and "path planning." Of course, humans have the same capabilities in the lower levels of the brain.

What Does All of This Have to Do with Autonomous Vehicles in a Well-Regulated Transportation Network?

Even the simplest AV system must understand human perception to be useful. As an example, ADAS collision-avoidance systems must assess driver attentiveness. More advanced systems must understand the future intent of pedestrians at intersections. At a transportation system level, it is the nature of human beings to anthropomorphize and we also do so with vehicles. Human beings interpret micro-breaking, micro-acceleration, lane drifting, and other factors in our own threat assessment of the situation. All these nonverbal movements are a source of active communication for human beings. Layered on top of this interpretation is more explicit nonverbal communication through eye contact or hand gestures. Overall, this creates a nonverbal "language-of-driving," which effectively makes cooperative transportation work.

Interestingly, this language-of-driving has all the characteristics of spoken language. It varies significantly by region. As an example, the communication provided by the horn in India and the US is quite different. In fact, the variation is so high that it is not uncommon for licensed drivers (typically western) to hire local drivers on their travels. What is the lesson?

When autonomous vehicles do not participate in this communication, they create danger in the overall system. To be effective, it is likely that AVs will have to be able to interpret a broader language and this analysis may well extend to the behavior of animals (e.g., deer, cats, dogs, etc.). Also, it may be useful for AVs to be clearly marked (e.g., color) so that humans can suitably adapt to the fact their behavior is going to be different than human drivers.

What about Cultural or Social Interactions? Are They Important for the Tango of Driving?

For an answer, we can look toward human behavior examples, such as the autism spectrum. One of the challenges for autistic individuals is the interpretation of social cues. In a similar period of time as the AV revolution, an increasing number of people with autism have wanted to join the driving public, and this has prompted research studies on their effectiveness in the driving task. "The Challenge of Driving With Asperger's" provides insights into this, and several comments in the article are directly applicable relative to AVs [24]:

- **Common sense:** "Obeying rules is generally a good thing but can be taken too far if rules are applied inflexibly or without taking context into account. For example, does a "Stop at White Line" sign mean that the line is where you should stop only if you need to stop—or that you should stop every time you come to it?"
- **Implied communication:** "And cooperating with other drivers involves perhaps the hardest task for people with Asperger's: reading nonverbal social cues. On the road, that happens through the "gestures" drivers make through the motion of their cars—by changing lanes boldly or hesitantly, for instance. Those motions amount to signals flashed from driver to driver so routinely that most people are hardly aware of the messages being sent about intention or mood."
- **Readiness to drive:** "According to a survey conducted by Cecilia Feeley, a project manager at the Center for Advanced Infrastructure and Transportation at Rutgers University, only 24 percent of adults with autism—many of whom described themselves as 'higher functioning'—said they were independent drivers, compared with 75 percent of the population as a whole."

In response to the desire of autistic individuals to drive, there are no laws against driving with autism, but safety is key [25]. To assess driving ability, it is suggested that there are some important factors and skills to master first, including motor coordination, pre-planning, flexibility to change, social judgment, and the ability to focus, multitask, and prioritize.

Thus, social interactions are critical to driving tasks. One of the key interactions is the communication of unease or discomfort. When faced with this signal, humans become more alert. Currently, in the ADAS use models, there is an expectation for humans to be alert at a machine-like level and take over the driving task when needed. In this sort of human-machine teaming, it is

FIGURE 6. **Facial gestures form the foundation of communication.**

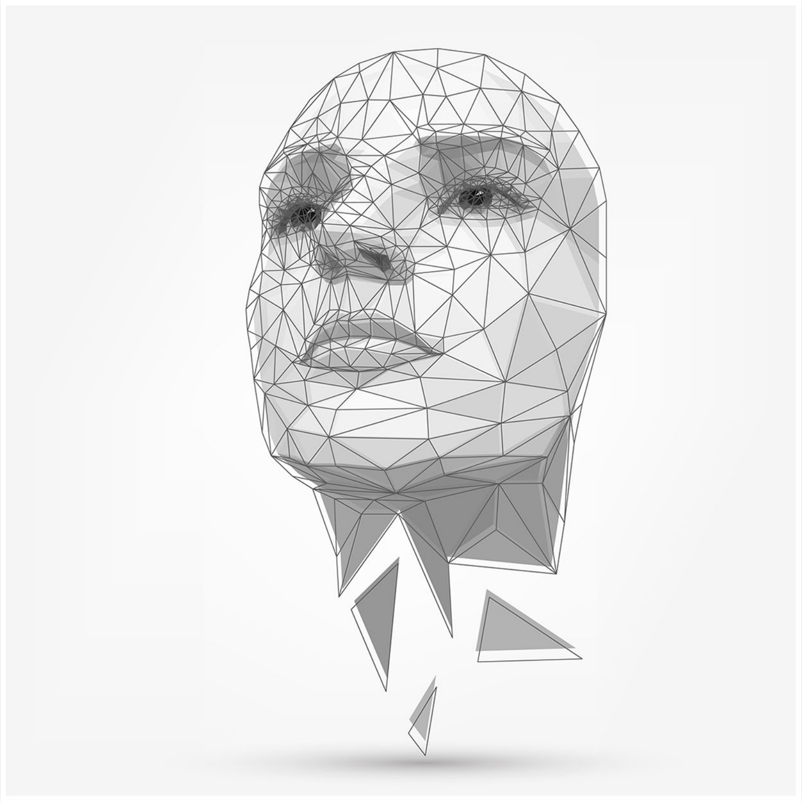

important for the AV to communicate risk in a manner understandable to humans.

Finally, there is a basic question about the exact manner of this communication. If IT history is any guide, machine interfaces that mimic humans world become critical to the success of the technology. Examples include Windows, Alexa, and a mouse. In the world of AVs, the simplest mechanism to communicate ideas of focus, fear, and gestures is the human face (Figure 6). Maybe a humanoid robot interface, like in Philip K. Dick's short science fiction story *We Can Remember It for You Wholesale* may be the easiest way for automated systems to team within a human world [26].

Recommendations

1. **Human Perception Systems**: To interact with humans, AVs must interpret and understand human perception to some degree. Clarity in defining and communicating this understanding is important for both parties.
2. **Language of Driving:** Human communities build languages for cooperative teaming. To participate in the act of cooperative transportation, AVs will have to understand this language. Depending on the level of expectation communicated by the AV, this language may extend into social interaction models.
3. **Communicating risk:** One of the most important communication paradigms consists of communicating risk. This is especially true in the situation where AVs expect humans to dynamically take over the driving task.

Progress in Human and Machine Interaction

In the context of AV technology, the area of human-machine teaming is relatively new and the issues are significant. However, both industry and academia are making steady progress. In this section, we will discuss four such developments:

1. **Passenger Communication:** The Automated Vehicle Safety Consortium (AVSC) led by SAE Industry Technologies Consortia (SAE ITC) has defined the behavior and interaction models between a passenger and the AV.
2. **Pedestrian Communication:** Researchers at Tallinn University of Technology are experimenting with visible AV intent communication for the benefit of pedestrians.
3. **Improved Perception Systems:** Industry and academia are looking at making perception more

robust. One approach from Massachusetts Institute of Technology (MIT) focuses on using movement of abstract objects as a signature for object recognition.

4. **Formal Language-of-Driving Definition:** A group at Florida Polytechnic University and University of Florida is launching a project to define the language-of-driving more formally using lessons from natural language processing and linguistics.

Passenger Communication

AVSC consists of a large representative set of AV manufacturers and members actively involved in AV testing and on-road pilot programs (Figure 7). AVSC has worked with SAE ITC to build group standards around the safe deployment of SAE Level 4 and Level 5 ADS and has recently released an "AVSC Best Practice for Passenger-Initiated Emergency Trip Interruption." [27]

The document outlines the interactions for a passenger-initiated emergency stop (PES) and passenger-initiated emergency call (PEC):

- Interfaces should be easily recognizable, reasonably accessible, and always available; PES and PEC should be easily identifiable but clearly distinguishable from one another based on visual, tactile, or audio cues.
- The stop must ALWAYS be controlled and provide an immediate and predictable response with humans always available to field the call for PEC. Overall system should be designed to assume an emergency may be taking place.
- Once the feature has been initiated, the interior should be illuminated and controllable ambient sound reduced (e.g., radio/entertainment), so fleet operations can assess the situation (we call this "Enhanced Diagnosis"). Further, other road users should be alerted that an emergency is taking place and an emergency maneuver may take place.
- Post-stop actions apply to PES and PEC if the PEC results in a stop:
    - Alert passengers and fleet operations once the stop maneuver is complete; the vehicle should not restart or resume driving until authorized by fleet operations personnel
    - Remind passengers to observe their surroundings and be aware of potential hazards in the environment if they leave the vehicle

If all ADS-equipped vehicles are equipped with a PES function, a PEC function, or both, passenger-initiated emergency interfaces may provide a sense of agency to passengers and increase their willingness to trust and use the system.

From a human-machine point of view, these specifications progress the communication paradigm, but there are still unresolved issues:

1. Standard Form Factor for the PEC and PES function
2. Well-defined ODDs for the PES maneuver

Lastly, a key requirement of an effective system is to have in-built fail-safe mechanisms based on the environment. Essentially, just like a helicopter is grounded on a

SB7/Shutterstock.com

FIGURE 7. **Emergency action and communication.**

foggy day, the system should safely ground itself. This is by no means an easy task, but it may be a requirement for AVs.

Pedestrian Communication

Communication between the car and pedestrians at a crosswalk is a difficult and important problem for automation. A research team from the Tallinn University of Technology in Estonia took on this challenge, and the result was the ISEAUTO shuttle [28]. For the shuttle, they designed a single LED panel illuminated by 128 lights where three patterns are communicated to pedestrians (Figure 8). Green indicates an invitation to cross, stripes indicate that "I see you," and a red cross that indicates one should not cross. A questionnaire was created to collect feedback and personal data from those who interacted with robot platform (Figure 9). Overall, the comfort level with the shuttle increased due to the communication paradigm [29].

Improved Perception Systems

Conventional AI algorithms in AVs such as Waymo and Tesla work toward a divergent model. They opt for a model more akin to a "knowledge oracle," where the AV observes its whole environment constantly and with full detail. In fact, this is a key part of the safety value statement offered by AV manufacturers for addressing the distracted-driving problem. The conventional methodology for AVs is to train the AI engines to recognize labeled objects with ever-growing databases. Of course, recognizing an object in a pixel map (or LiDAR point cloud or imaging radar) in all of the potential orientations is a very difficult problem. Invariably, there is confusion caused by objects such as a van that has a person's image wrapped around the exterior, or by a pedestrian walking with a bicycle (as in the Phoenix Uber fatality) [30]. Like Google image searches, this "data-up" method of solving the problem has serious robustness issues because one is always missing the next interesting training set.

In response, AV manufacturers have pushed out the training to the physical world. Tesla touts the use of its

FIGURE 8. **Symbols used by ISEAUTO for communication.**

Pattern	Zebra Line	Arrows	Cross
Action	Pass	Pass	Stop
Visualization			

Tallinn University of Technology

FIGURE 9. **ISEAUTO vehicle communicating with pedestrian.**

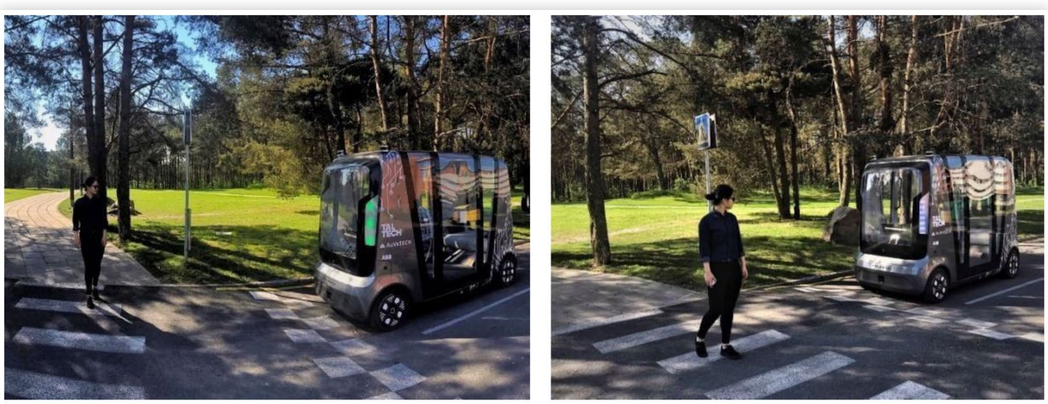

Tallinn University of Technology

fleet—nearly one million vehicles—with Autopilot as a mechanism for reaching closure on this task. There are three reasons that this approach may not be wise and ultimately may not work as currently constructed:

1. **Reactive Methodology:** Fundamentally, the process is reactive in nature. That is, one records the world, builds a process to analyze the data, performs root cause analysis, and solves the problem. Each step has significant challenges and expenses.
2. **Execution Velocity:** A real-world physical test bed has its advantages, but the real world moves very slowly. Most cars are parked most of the time (less than 5% utilization is the norm), so the number of driven miles for even a million cars is not very high.
3. **Sampling Bias:** Perhaps the most important shortcoming is that of sampling bias. The Tesla road fleet is highly correlated to specific areas, so whatever validation is done is limited to those situations. What about the next interesting situation? How does it define a clear expectation for the next customer for where the AV will work?

Overall, compositional reasoning over objects and their relationships is a challenging ML problem and solutions will enable ML models to act intelligently in complex and unforeseen environments. Most of the effective deep learning approaches developed are supervised. These tend to not work well when unforeseen objects (or their unforeseen actions) are in play. The work in unsupervised learning is not yet mature but some strides are being made by combining object representation and dynamics for inferencing future behavior based on video streams.

What Are Some of the Solution Vectors? Given the robustness challenges of the current approach, it would seem that using the ideas of focus and abstraction from the human world would be useful. In fact, Bluespace, a Silicon Valley startup, claims that there is utility in such an approach for AVs [31]. Will this approach work? Time will tell. In the world of chess, the initial solutions were all based on raw data and computation, but eventually, the winning solutions used a combination of human insight and computing power [32]. One gets the feeling that AV technology is on a similar technological arc. That is, it seems reasonable that one does not need to know the ultimate details of all objects to enable autonomous operations in an vehicle. Combining higher-level insights can add robustness while lowering power and cost. Afterall, the human brain can drive while multitasking, spending only about 20-30 Watts of power (akin to the power demand of a low-end conventional laptop) [33].

Another active thread is the use of movement for AI learning. Similar to animals, humans use motion as a part of object recognition, but this has not been a big part of AV object recognition. Rather, AV object recognition has traditionally focused on training AI engines on individual images.

Recently, MIT released a video data set called DriveSeg. The data consists of two minutes and 47 seconds of high-resolution video captured during a daytime trip around the busy streets of Cambridge, Massachusetts. The video's 5,000 frames are densely annotated manually with per-pixel human labels of 12 classes of road objects. Then, a set of videos captured from a range of scenes drawn from MIT Advanced Vehicle Technology Consortium data coarsely annotated through a novel semiautomatic annotation approach developed by MIT [34]. This dataset is meant to enable deeper object recognition algorithmic development based on movement.

In conclusion, as machine perception moves toward the fundamental forms of human perception, there is a greater ability for AV systems to communicate with humans.

Formal Language Definitions for the Tango In the world of natural language processing and understanding (e.g., Alexa, etc.), researchers have built a reasonable basic understanding for spoken and written language. However, this much more basic form of communication and perception is just at its beginning stages. It seems the capability is so innate that we did not know we did it, until we had to recreate it in AVs. Is it possible to formally capture this language-of-driving? Is it possible to build a methodology for collecting the information in a structured fashion? A group at Florida Polytechnic University, Florida Institute for Human and Machine Cognition, and University of Florida is starting to look at this fundamental problem.

In a classic modular system, human-spoken communication has the following modules: phonetics (mechanism behind speech production and perception), phonology (the grammar of speech sounds), morphology (the grammar of meaning bearing units), syntax (the grammar of meaningful forms for deriving unique messages), semantics (the mechanism of how the meaning of morphemes, words, phrases, and sentences are derived), pragmatics (how contexts can contribute to meaning), and a lexicon (a dictionary that maps meanings with forms). Such modular systems also exist with human sign languages. While the precise nature of these modules derived from our understanding of spoken/sign languages are still under active research and debate, it is clear that the language of driving must consist of some, if not all, of elements of these modules. To illustrate this, let us turn to the highest level of the communicative modules—pragmatics—and specifically the theory of mind (ToM).

ToM is a much-needed component of pragmatics. There is evidence that language development correlates with the development of ToM in humans. It is the ability to attribute mental states (such as beliefs, intents, desires, knowledge) to others and understand how they might differ from those of our own. It is known to be a deficit with people having autism spectrum disorders. ToM is crucial for any human social interaction. Pragmatic theories of human communication assume an understanding of beliefs and mental states of others as enabled by ToM to infer the communicative content of the interlocutors. More than often, the intended meaning of a signal (e.g., an utterance) cannot be understood literally from the signal, which is underspecified and has many different meanings—instead, they must rely on the actual context. Recent work in AI has started to consider ToM in two distinct directions:

1. Teaching an AI to understand/develop a ToM of a human [35]
2. Understanding how humans understand the ToM of an AI (Theory of AI's mind) [36]

There are already tasks that are used to test ToM such as the "Sally-Anne" test, which tests one's ability to recognize that others can hold false beliefs about the world [37]. A formal definition of the language-of-driving must therefore incorporate a component of ToM of humans and/or of AVs. Existing evaluative tasks of ToM can be adapted for evaluating AV's language-of-driving.

Overall, humans use nonverbal communication to clarify our intentions when negotiations are needed, especially when there are ambiguities. For example, while crossing the intersection, we (e.g., pedestrians) rely on cues from drivers' behaviors such as eye contact, postures, and gestures to negotiate right of way. This is a relatively simple form of communication that can be enhanced by using an external human-machine interface (HMI). It is evident that we must define this language before we start developing software or interventions to fix current issues arising between users and AVs. The current social norms for the language of driving will need to adapt to accommodate AVs. Similar to when automobiles were introduced in the early twentieth century, we will need to adapt our vehicle design, infrastructure, and social norms.

Recommendations

1. **Improved HMI Conventions:** Building on the excellent work of AVSC, industry must continue to clearly define AV interactions with humans in their various roles as driver, pedestrian, third-party driver, and passenger.
2. **Improved Perception:** The gap between AI and human perception systems are broad. As AI perception improves, one of the key vectors of success will be successful communication to humans of the AI state.
3. **Formalized Language of Driving:** Cooperative teaming is a key aspect of using the transportation system. Formalization of the language-of-driving and building mechanisms for capture, validation, and communication of this language will be critical to AV success.

Summary

Driving is a cooperative task, which has many similarities to a dance where the participants figure out the rules in an ad hoc manner. This dance has been well honed by human beings over the last hundred years. There are instances of direct communication (even language) through gestures, but a great deal of the communication is implicit. Much like a dance, drivers notice the subtle behavior (speed of acceleration, stability on road, etc.) and react to these motions with their own set of motions.

For AVs to add value, they must deal with the reality of implicit and explicit human communication. This communication paradigm is bidirectional. That is, AVs must understand human communication, but just as importantly, communicate their internal state to humans in a naturalistic manner. Finally, a very important concept is the connection of communication with expectation.

Expected behavior is at the center of critical concepts such as legal liability. As an example, the original AVs were horses, and since everyone had a clear understanding and associated expectation of a horse, legal liability attached naturally. There were concepts of “spooking” a horse. Is there a concept of spooking an AV? If so, those concepts must be explicitly set and communicated.

Unfortunately, the current state of the industry is that there are no clear expectations set on AV behavior. Rather, AVs are placed in the middle of the driving dance. Today, the assumption by all parties is that they will react in the same manner as humans. From a human-machine interaction point of view, this situation has an incredible number of challenges. This is the case even for ADAS, the lowest level of autonomous capability. The result of these fuzzy expectations are fuzzy ODDs, and the resulting confusion of consumers. Regulators, insurance companies, and even vehicle makers have not yet reached coherence on this issue, and it is likely the coherence may be driven by the courts.

Beyond expectations, the underlying challenges of autonomous perception are daunting, and the current architectures have significant robustness issues. Promising techniques for improving autonomous perception are investigating concepts such as focus, abstraction, and movement—often gaining insights from human perception. Even more importantly, the alignment of these perception and communication systems must be done to allow the two parties to reasonably communicate with each other.

In terms of the state-of-art, early research in human-machine interaction has focused on dynamic signage, but more complex naturalistic communication may in fact go in the direction of a manipulating a human face. Further, AI research is now starting to explore concepts of movement and abstraction to build more robust object recognition systems. Communication standards for specific circumstances such as emergency stops are slowly being developed by SAE ITC. Overall, human-machine cooperation is very much in its early days.

SAE EDGE™ Research Reports

SAE EDGE™ Research Reports, like the present report on "Unsettled Topics Concerning Human and Autonomous Vehicle Interaction," are intended to push further out into still unsettled areas of technology of interest to the mobility industry. SAE launches these reports before attempting to form a joint working group, let alone a cooperative research program or a standards committee.

SAE EDGE™ reports are intended to be quick, concise overviews of major unsettled areas where vital new technologies are emerging. An unsettled area is characterized more by confusion and controversy than established order. Early practitioners must confront an absence of agreement. Their challenge is often not to seize the high ground but to find common ground. These scouting reports from the frontiers of investigation are intended merely to begin the process of sorting through critical issues, contributing to a better understanding of key problems, and providing helpful suggestions about possible next steps and avenues of investigation.

SAE EDGE™ Research Reports, therefore, are fundamentally distinct from the more formal working groups approach and far removed from the more mature research program and standard's development process.

Next Steps for Human-Machine Interaction

This publication should be considered only as a first step toward clarifying the issues around human-machine interaction. The intention behind this and other SAE EDGE™ Research Reports is to start a dialogue among interested parties on important industry-wide topics that require further attention. The expectation is that these explorations of unsettled areas of technology will lead to the formation of working groups and, ultimately, committees that can address and resolve the issues they raise, producing a framework for developing a common vocabulary of definitions, best practices, protocols, and standards needed to support continued progress toward safer and more innovative products.

The experts' collaboration that gave rise to this publication demonstrated a great willingness on the part of the industry to define the terminology, procedures, and eventually the standards needed to enable human-machine interaction technology to move ahead as quickly and efficiently as possible. SAE International has demonstrated its lead in this and closely related areas by the AVSC group.

This SAE EDGE™ Research Report on human-machine interaction identifies the following key topics for further pursuit, both through continued informal discussions among industry practitioners and through more formal working groups:

- ADAS ODD: This must be defined with associated performance metrics.
- Autonomy Communication: Standardization of autonomy communication to humans must continue beyond limited emergency situations
- Perception Research: We must collaborate to accelerate the pace of perception research. It is the critical cog to autonomy progress.

Recommendations

The overall recommendations of this SAE EDGE™ Research Report can be summarized as follows:

1. **AV Expectations:** It is critical to define and set clear measurable behaviors and metrics for AV functionality. These expectations become the bedrock for ODD definitions for regulators, customers, and engineers. Eventually, these expectations become the best way to manage liability.
2. **AV Communication:** AVs must find ways to naturalistically communicate their capabilities and internal state to third parties. Without this communication, third parties may well misinterpret capabilities.
3. **AV Understanding:** Depending on expectations and communication, AVs must be able to understand and react to third-party explicit and implicit communication signals.

Definitions

ADAS - Advanced Driver Assistance Systems

ADS - Automated Driving System

AI - Artificial Intelligence

AV - Automated Vehicle

IT - Information Technology

ITC - [SAE] Industry Technologies Consortia

HMI - Human-Machine Interface

LiDAR - Light Detection and Ranging

MIT - Massachusetts Institute of Technology

ML - Machine Learning

ODD - Operational Design Domain

PEC - Passenger-initiated Emergency Call

PES - Passenger-initiated Emergency Stop

US - United States

Acknowledgments

Recognition should go first to all the participants, many of who also provided feedback on the draft version of this publication. Without their input and initiative, this SAE EDGE™ Research Report would not have been possible.

Ken Ford, PhD, Florida Institute for Human and Machine Cognition
Justin Mason, PhD, University of Florida
Sanjay Ranka, PhD, University of Florida
Brian Reimer, PhD, New England University Transportation Center at MIT
Raivo Sell, PhD, Tallinn University of Technology
Edward Straub, PhD, Automated Vehicle Safety Consortium
Kevin Tang, PhD, University of Florida

The authors of this document together with the SAE Team responsible for its creation join in expressing their deepest appreciation to all the individuals mentioned above.

Rahul Razdan, PhD
Raztech Research Institute

References

1. Straub, E.R. and Schaefer, K.E., “It Takes Two to Tango: Automated Vehicles and Human Beings Do the Dance of Driving—Four Social Considerations for Policy,” *Transportation Research Part A: Policy and Practice* 122(C):173-183, 2019.
2. SAE International, “Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles J3016_201806,” 2018, https://doi.org/10.4271/J3016_201806.
3. Johnson, M. and Vera, A., “No AI Is an Island: The Case for Teaming Intelligence,” *AI Magazine* 40(1):16-28, 2019, https://doi.org/10.1609/aimag.v40i1.2842.
4. Ford, K., “On computational Wings: The Prospects & Putative Perils of AI,” 2018, https://www.youtube.com/watch?v=OsfZ8R2K6ZM.
5. Johnson, M., Bradshaw, J.M., and Feltovich, P.J., “Tomorrow’s Human-Machine Design Tools: From Levels of Automation to Interdependencies,” *Journal of Cognitive Engineering and Decision Making* 12(1):77-82, 2017.
6. Responsibility-Sensitive Safety (RSS), “A Mathematical Model for Automated Vehicle Safety,” https://www.mobileye.com/responsibility-sensitive-safety/, accessed Dec. 3, 2020.
7. Razdan, R., “Will Truck Convoying Be the First Viable Commercial Application for AV Technology?,” *Forbes*, Mar. 21, 2020, https://www.forbes.com/sites/rahulrazdan/2020/03/21/will-truck-convoying-be-the-first-viable-commercial-application-for-av-technology-/#5468dd8f1454, accessed Sept. 25, 2020.
8. Razdan, R. et al., “Requirements for the Next-Generation Autonomous Vehicle Ecosystem,” in *IEEE Southeastcon*, Huntsville, AL, 2019.
9. Sahawneh, S., Alnaser, A.J., Akbaş, M.İ., Sargolzaei, A. et al., “Requirements for the Next-Generation Autonomous Vehicle Ecosystem,” in *2019 SoutheastCon*, Huntsville, AL, 2019, 1-6, https://doi.org/10.1109/SoutheastCon42311.2019.9020400.
10. Drucker, P., *Management: Tasks, Responsibilities, Practices* (New York: HarperCollins, 1974), 84-85. ISBN:978-0-7506-4389-4.
11. AAA, “Advanced Driver Assistance Technology Names,” Jan. 2019, https://www.aaa.com/AAA/common/AAR/files/ADAS-Technology-Names-Research-Report.pdf, accessed Dec. 3, 2020.
12. National Science & Technology Council and US Department of Transportation, “Ensuring American Leadership in Automated Vehicle Technologies: Automated Vehicles 4.0,” Jan. 2020, https://www.transportation.gov/sites/dot.gov/files/2020-02/EnsuringAmericanLeadershipAVTech4.pdf, accessed Dec. 3, 2020.
13. Razdan, R., “New AV Bill, Its Bipartisan, Is It Better?,” *Forbes* Feb. 24, 2020, https://www.forbes.com/sites/rahulrazdan/2020/02/24/reacting-to-new-bipartisan-av-bill---good-start-but-missing-adas-urgency/#cb0a61475449, accessed Sept. 25, 2020.
14. IIHS, “IIHS Examines Driver Assistance Features in Road, Track Tests,” Aug. 2, 2018, https://www.iihs.org/news/detail/iihs-examines-driver-assistance-features-in-road-track-tests, accessed Sept. 25, 2020.
15. Preston, B., “The Hidden Cost of Car Safety Features,” n.d., Retrieved Sept. 26, 2020, from https://www.consumerreports.org/car-repair/the-hidden-cost-of-car-safety-features/, accessed Sept. 25, 2020
16. Global Fleet, “Insurers Say No Discounts for ADAS,” July 31, 2019, https://www.globalfleet.com/en/technology-and-innovation/north-america/features/insurers-say-no-discounts-adas?a=API07&t%5B0%5D=ADAS&t%5B1%5D=Insurance&curl=1, accessed: Sept. 25, 2020.
17. Preston, B., “The Hidden Cost of Car Safety Features,” Consumer Reports, Jan. 30, 2020, https://www.consumerreports.org/car-repair/the-hidden-cost-of-car-safety-features/, accessed Sept. 25, 2020.
18. Field, K., “As Radar Technology on cars proliferates, so do interference concerns,” Fierce Electronics, Dec. 30, 2019, https://www.fierceelectronics.com/electronics/as-radar-technology-cars-proliferates-so-do-interference-concerns, accessed Sept. 25, 2020.

19. Proposal for a New UN Regulation on Uniform Provisions Concerning the Approval of Vehicles with Regard to Automated Lane Keeping Systems, https://undocs.org/ECE/TRANS/WP.29/2020/81, accessed Dec. 3, 2020.
20. ASAM OpenODD, "ASAM OpenODD Standard," https://www.asam.net/project-detail/asam-openodd/, accessed Sept. 25, 2020.
21. Adelson, E.H., "Checkershadow Illusion," http://persci.mit.edu/gallery/checkershadow, accessed Sept. 25, 2020.
22. Razdan, R., "Skydiving Lessons for Autonomous Driving Perception Systems," *Forbes*, 2020, https://www.forbes.com/sites/rahulrazdan/2020/07/12/skydiving-lessons-for-autonomous-driving-perception-systems/#4ebb62d56684, accessed Sept. 26, 2020.
23. Avgar, T., Deardon, R., and Fryxell, J.M., "An Empirically Parameterized Individual Based Model of Animal Movement, Perception, and Memory," *Ecological Modelling* 251:158-172, 2013, ISSN:0304-3800.
24. Razdan, R., "Temple Grandin, Elon Musk and the Interesting Parallels between Autonomous Vehicles and Autism," *Forbes*, 2020, https://www.forbes.com/sites/rahulrazdan/2020/06/07/temple-grandin-elon-musk-and-the-interesting-parallels-between-autonomous-vehicles-and-autism/#6a90604c53cb, accessed Sept. 26, 2020.
25. Autism Society, "Driving—Autism Society," 2020, https://www.autism-society.org/living-with-autism/autism-through-the-lifespan/adulthood/living-with-autism-autism-through-the-lifespan-adulthood-driving/, accessed Sept. 26, 2020.
26. Dick, P.K., "We Can Remember It for You Wholesale," 1990.
27. "AVSC Best Practice for Passenger-Initiated Emergency Trip Interruption," AVSC00003202006, SAE Industrial Technologies Consortia, June 30, 2020.
28. Sell, R., "Estonia's First Self-Driving Car Iseauto," Sept. 21, 2018, https://www.youtube.com/watch?v=0URhTZ2L4F4, accessed Sept. 25, 2020.
29. Wang, R., Sell, R., Rassolkin, A., Otto, T. et al., "Intelligent Functions Development on Autonomous Electric Vehicle Platform," *Journal of Machine Engineering* 20(2):114-125, 2020, https://doi.org/10.36897/jme/117787.
30. Vlasic, B. and Boudette, N.E., "Self-Driving Tesla Was Involved in Fatal Crash, U.S. Says," *The New York Times*, June 30, 2016, https://www.nytimes.com/2016/07/01/business/self-driving-tesla-fatal-crash-investigation.html, Retrieved Dec. 11, 2018.
31. Razdan, R., "Skydiving Lessons for Autonomous Driving Perception Systems," *Forbes*, 2020, https://www.forbes.com/sites/rahulrazdan/2020/07/12/skydiving-lessons-for-autonomous-driving-perception-systems/#4ebb62d56684, accessed Sept. 26, 2020.
32. Praxtime by Nathan Taylor, "What Chess and Moore's Law Teach Us about the Progress of Technology," 2020, https://praxtime.com/2014/03/24/chess-technology-progress/, accessed Sept. 26, 2020.
33. Hypertextbook.com, "Power of a Human Brain—The Physics Factbook," 2020, https://hypertextbook.com/facts/2001/JacquelineLing.shtml#:~:text=The%20brain%20consumes%20energy%20at,of%20the%20brain%2020%20W, accessed Sept. 26, 2020.
34. Razdan, R., "MIT and Toyota Announce an Important Dataset for Improving Perception Research," *Forbes*, 2020, https://www.forbes.com/sites/rahulrazdan/2020/06/24/mit-and-toyota-announce-an-important-dataset-for-improving-perception-research/#2ed337d672e0, accessed Sept. 26, 2020.
35. Rabinowitz, N.C., Perbet, F., Song, H.F., Zhang, C. et al., "Machine Theory of Mind," 2018, https://arxiv.org/abs/1802.07740, accessed Sept. 25, 2020.
36. Chandrasekaran, A., Yadav, D., Chattopadhyay, P., Prabhu, V. et al., "It Takes Two to Tango: Towards Theory of AI's Mind," CoRR. Abs/1704.00717, 2017, http://arxiv.org/abs/1704.00717, accessed Sept. 25, 2020.
37. Baron-Cohen, S., Leslie, A., and Frith, U., "Does the Autistic Child Have a Theory of Mind?" *Cognition* 21:37-46, 1985, https://doi.org/10.1016/0010-0277(85)90022-8.

Contact Information

EDGEresearch@sae.org