NIST Internal Report
NIST IR 8442

Workshop Report: Standards and Performance Metrics for On-Road Autonomous Vehicles

Craig Schlenoff
Suzanne Lightman
Vinh Nguyen
Prem Rachakonda
Tao Zhang
Omar Aboul-Enein
Nicholas Barbosa
Cameron Miller
Daniel Sawyer
Ann Virts

NIST | NATIONAL INSTITUTE OF
STANDARDS AND TECHNOLOGY
U.S. DEPARTMENT OF COMMERCE

NIST Internal Report
NIST IR 8442

Workshop Report: Standards and Performance Metrics for On-Road Autonomous Vehicles

Craig Schlenoff
Suzanne Lightman
Vinh Nguyen
Prem Rachakonda
Tao Zhang
Omar Aboul-Enein
Nicholas Barbosa
Cameron Miller
Daniel Sawyer
Ann Virts
National Institute of Standards and Technology
Gaithersburg, MD 20899

October 2022

NIST Technical Series Policies
[Copyright, Fair Use, and Licensing Statements](#)
[NIST Technical Series Publication Identifier Syntax](#)

Publication History
Approved by the NIST Editorial Review Board on 2022-09-26

NIST Author ORCID iDs
Vinh Nguyen: 0000-0003-4617-0892
Tao Zhang: 0000-0002-0684-4111
C Cameron Miller: 0000-0002-2960-0996

Abstract

On-road autonomous vehicles are expected to significantly influence key aspects of everyday life. However, these complex systems can pose a safety risk in the event of unexpected system performance. Therefore, the National Institute of Standards and Technology (NIST) held the Standards and Performance Metrics for On-Road Autonomous Vehicles Workshop to solicit stakeholder feedback with respect to challenges and opportunities in developing standards and performance metrics for this complex interdisciplinary field. In addition, this workshop aimed to foster a community consisting of stakeholders from a variety of disciplines and domains in autonomous vehicles. This workshop was conducted virtually over two days (March 8th-9th, 2022) consisting of introductory remarks, panels, breakout sessions, and concluding remarks. This report is a summary of the technical topics discussed by participants during the workshop.

Key words

Artificial Intelligence; Autonomous Vehicles; Cybersecurity; Intelligent Vehicles; Vehicle Communications; Vehicle Sensing.

Table of Contents

1. Introduction

This section provides background context behind the National Institute of Standards and Technology (NIST) effort to advance the state of on-road autonomous vehicles in addition to detailing the purpose for the Performance Metrics for On-Road Autonomous Vehicles Workshop held virtually on March 8[th] and 9[th], 2022. Workshop recordings can be accessed through the following website https://www.nist.gov/news-events/events/2022/03/standards-and-performance-metrics-road-autonomous-vehicles.

1.1. Background

On-road autonomous vehicles are projected to influence key aspects of everyday life including transportation, goods delivery, manufacturing, public safety, and security. Morgan Stanley indicated that self-driving cars could save the US economy $646 billion in annual savings[1]. In addition, a McKinsey & Company study found that this technology has the potential to reduce motor vehicle deaths by 90%[2]. Furthermore, the US Department of Commerce estimated that 15.5 million U.S. workers could be affected by this technology[3] through new job opportunities and retraining.

However, on-road autonomous vehicles can pose a safety risk in the event of unexpected system performance. In addition, the autonomous vehicle sector is a vastly inter-disciplinary field with subject areas including perception, planning, control, safety, cybersecurity, and communications. NIST is a nonregulatory government agency under the Department of Commerce that is equipped to develop test methods, metrics, and standards to characterize the performance of these complex systems. Therefore, industry and government agencies are looking to NIST to characterize the performance of these complex systems to mitigate risk to both the manufacturers and the consumers.

1.2. Workshop Goals

This workshop was aimed at: 1) soliciting stakeholder feedback to identify key areas where standards and performance metrics are needed to help advance the on-road autonomous vehicle field, 2) engaging a multidisciplinary community consisting of stakeholders from a variety of disciplines and domains in autonomous vehicles, and 3) to convene industry experts in the on-road autonomous vehicle space and reduce barriers to entry and innovation within the market.

To achieve those goals, NIST invited keynote speakers to introduce the important topics of the workshop and invited multidisciplinary panelists from government, academia, and the private sector to participate in focused panel discussions. Finally, the workshop facilitated breakout sessions to allow participants to offer feedback in the key topic areas. This report details the discussion by workshop participants.

---

[1] Shared Mobility on the Road of the Future, Morgan Stanley, 2015
[2] Ten ways autonomous driving could redefine the automotive world, McKinsey & Company, 2015
[3] The Employment Impact of Autonomous Vehicles, US Department of Commerce, 2017

2. Introductory Remarks and Keynote Speeches

At the beginning of each day, introductory remarks and keynote speeches were provided from prominent government officials and domain experts. These presentations were given to level set and inspire discussion for follow-on workshop elements.

2.1. James K. Olthoff, PhD

At the time of the workshop, Dr. James K. Olthoff was performing the Non-Exclusive Functions and Duties of the Under Secretary of Commerce for Standards and Technology & Director at NIST. Currently, Dr. Olthoff is the Associate Director for Laboratory Programs at NIST. Dr. Olthoff described NIST's role in establishing standards and emerging technologies. He recognized the vast interdisciplinary field of the autonomous vehicles sector, and the many factors that the sector relates to including perception, planning, control, safety, cybersecurity, and communications. He noted the difficulty of identifying another technology that touches so many elements of NIST's portfolio. Autonomous vehicles bring great benefits but also risks, so their system performance needs to be sound and secure. He noted that developing the required level of trust for such vehicles requires standards and measurements and is therefore, well-suited to NIST's mission. He described NIST's previous history in the field since at least 2004 with the Defense Applied Research Projects Agency (DARPA) Learning Applied to Ground Vehicles (LAGR) program and remarked on the progress in the years since. Finally, Dr. Olthoff thanked the attendees for their participation and for bringing their perspectives and experience. Together, he challenged the workshop to help build confidence that autonomous vehicles will behave safely, securely, and appropriately, thus developing the trust in that technology that the public needs.

2.2. Nell Abernathy

Mrs. Nell Abernathy is currently the Director of the Office of Policy and Strategic Planning of the Department of Commerce. Mrs. Abernathy shared her pleasure in representing the Department of Commerce and shared some of the agency's priorities of broadly shared prosperity, domestic innovation, investments in the industrial base, and international collaboration. The mission is to support these advancements and innovation, and she reminded attendees that we will be unable to succeed in the 21st century unless we are able to harness the talents of all Americans and all communities. As we facilitate and expedite the development and commercialization of emerging technologies, we recognize the need to promote collaboration among federal, state, local, tribal, territorial, business, not-for-profit, and academic institutions. As a champion of advancing standards and developing measurement science solutions to challenging technical problems, NIST is well positioned to represent the federal government, drawing upon NIST's experience in areas such as Artificial Intelligence (AI), cybersecurity, and materials- and physics-based technologies. She described the focus areas for the workshop as we aim to continue U.S. leadership in promoting an open dialogue in promoting on-road autonomous vehicle development. On behalf of the Secretary of Commerce, she thanked attendees for their participation and for their past, current, and future progress.

2.3.   Trent Victor, PhD

Dr. Trent is currently the Director of Safety Research and Best Practices at Waymo, which is currently developing robot-taxis. His keynote speech was subject to proprietary information protection during this presentation.

2.4.   Tim Kurth

Mr. Tim Kurth is currently the Chief Counsel (R), United States House of Representatives of the Energy and Commerce Committee Consumer Protection and Commerce Subcommittee. Mr. Kurth welcomed attendees to Day 2. He described the importance of standards and performance metrics that will be supported by NIST and the National Highway Traffic Safety Administration (NHTSA) as on-road autonomous vehicles move to commercialization. He described the benefits that on-road autonomous vehicles will provide as a new form of mobility. He shared the risks of inaction, for example, noting that Department of Transportation 2020 figures reflected more than 38,000 fatalities in crashes nationwide cause by human error – a statistic that the U.S. hopes will be greatly diminished by on-road autonomous vehicles. Mr. Kurth described legislation supporting NIST work on national strategies, through the Department of Commerce, in the areas of artificial intelligence, Internet of Things, advanced manufacturing, quantum computing, blockchain technology, new and advanced materials, unmanned delivery services, and three-dimensional printing. He advocated for progress on a national privacy and data security framework.

2.5.   Adam Campbell, PhD

Dr. Adam Campbell is currently the Senior Manager Safety Innovation and Impact at Gatik. Dr. Campbell introduced Gatik as an organization that currently uses on-road autonomous vehicles for logistics and deliveries. He opened the discussion with a hypothetical question about how one should react when an autonomous vehicle is involved in a collision that a human driver would have avoided. He proposed that an autonomous vehicle should perform at least as well as a human, thus establishing human driver performance as a baseline. He reminded attendees that innovative changes have occurred in the past, illustrated by the difference between 1900 and 1913 where horse-drawn carriages gave way to automobiles. He noted that technology often precedes policy, but that good teams working together help to progress to safer, more reliable solutions. Drawing on discussions from Day 1, he reflected on some of the considerations that will go into that benchmarking and the need to look at ways to address challenges, including the difficulty of establishing metrics that will prove safety. Research demonstrates that it is difficult to perform a statistical comparison of existing drivers due to their sheer number and performance variation. A human benchmark will be beneficial to supporting system design and testing. Therefore, confidence and metrics can be provided if our resulting design and performance match those of a human being. He pointed out that industry can take advantage of research, such as studies performed by NHTSA, and that comparing our data and observations of our autonomous vehicles to these types of data will help set some benchmarks about what to expect and explain vehicle performance. He also shared that, if industry doesn't develop metrics to at least set minimum performance expectations, then the courts will, and the consequences of lagging could be significant. He

thanked participants again for their work and encouraged the attendees to continue their progress.

3. Panel Discussions

Five one-hour panel discussions were conducted to inform workshop participants about the current state of standards, test methods, and performance metrics of on-road autonomous vehicles. Panelists and moderators were respected experts in the vast and interdisciplinary field of automated and connected vehicle technology. Panels were held in the following focus areas: safety, communications, cybersecurity & privacy, artificial intelligence, and sensor perception.

3.1. Safety

Moderator: Edward Griffor, NIST

Panelists: Chan Lieu, Aurora; Edward Straub, SAE; Jack Weast, Intel; Bert Kaufman, Zoox; Robert Dingli, Plus; Robyn Robertson, Traffic Injury Research Foundation

The panelists began by discussing the biggest challenges facing on-road autonomous vehicle safety. Challenges were both non-technical and technical. A significant non-technical challenge is convincing the public that vehicles are safe while a significant technical challenge is interacting with other road users. Predicting human behavior is difficult and will rely on advancement of computation. Within these non-technical and technical challenges, the panelists discussed the notion of how to create a "safe system" while accommodating various use cases.

Panelists agreed, there is no such thing as perfect safety since the act of driving is inherently risky. Therefore, the panelists posed questions to the audience:

- How do we know the vehicle is safe enough to deploy?
- What is an acceptable level of safety?
- Who sets the acceptable level of risk?
- How do we go about incorporating different factors in our requirements?

The panelists then turned to the question of defining standards and performance metrics. The discussion continued to define safety in given scenarios. Panelists highlighted the unbounded number of potential scenarios that could be created and tested. The challenge with a scenario-based approach is the constantly changing landscape and additions of new scenarios. A set of known scenarios would be a place to start but would need to be updated and challenged continuously. A finite set of use-cases could be utilized during the development process. The aim would be that the case set is sufficiently large to demonstrate reasonable safety while recognizing there are edge cases which are untested. If the operational design is enterprise scale, the data set would be very large and a barrier to entry for smaller organizations.

Another approach to defining safety would be to utilize a metrics-based approach. The panelists discussed the challenges of defining the right level of abstraction for metrics. Some metrics may be too general to be testable, some metrics may be too specific to be applicable to all situations. The panelists discussed the feasibility of separating testing metrics from scenarios to maintain safety standards while not inhibiting the development of the technology.

The panelists moved on to discuss the baseline environment of the on-road autonomous vehicle and the effects of the environment on safety considerations. Panelists highlighted current considerations in their organizations including emergency vehicles, road conditions, pedestrians, infrastructure data, and other non-autonomous vehicles. While these considerations are important to the on-road autonomous vehicle safety, the panelists suggested a need to describe baseline expectations of safety from these environmental variables. There does not currently exist a standard vocabulary for discussing these variables and the corresponding expectations of safety. The panelists agreed that infrastructure changes (e.g., consistent roads, data on road work, road marking) would be helpful, and that those changes take a long time to implement and deliver value.

The panelists concluded by discussing human-machine interaction. Panelists highlighted the need for intuitive and responsive interface with the vehicle. While there are some use-cases where the vehicle is empty, the human remains a large part of the system. Therefore, the human inside and outside of the vehicle will affect the safety of the vehicle. There currently is no baseline of what to expect or make predictions of what the humans inside and outside the vehicle will do. With the number of variables inherent in the human-machine interaction, simulation was discussed as a tool to reduce errors after deployment. Panelists agreed that a feedback loop of acquired knowledge was needed after deployment to continually raise the safety bar for autonomous vehicles.

3.2. Communications

Moderator: Scott McCormick, Teleoperation Consortium

Panelists: John McNicol, British Standards Institute; Stan Schneider, Real-Time Innovations; Maarten Sierhuis, Nissan North America; Fan Bai, General Motors; Bettina Erdem, Continental Corporation

The panel first discussed the challenges and solutions to teleoperation. Generally, the panel agreed that teleoperation had many challenges that were focused on the interaction between the human and the machine. Being able to handle corner cases (e.g., construction zones, police interaction, rescue missions) is essential to the mass deployment of driverless cars. The vehicles already possess a tremendous amount of processing power; however, connectivity of the remote operator to the vehicle is a large hurdle. Panelists discussed the need for an ontology of teleoperation modes (time-sensitive and non-time-sensitive teleoperation). The panel suggested standards around the training and certification of the remote operators and remote operator stations as potentially useful, as most autonomous vehicles have a remote operator that can step in during unexpected situations.

The panel discussed the broader topic of communication technology. From a conceptual perspective, the panelists discussed the communication of a vehicle with another vehicle, the vehicle to the infrastructure, and the infrastructure to the vehicle. There are currently many different approaches to the communication of the system components. There is currently no standard data format, channel, or presentation layer for users. There is a need to directly share communication-focused data sets and not worry about language barriers.

A challenge to any communication approach will be the channel over which the communication is transmitted. The panelists discussed the use of many bandwidths for

communication. Similarly, the panelists highlighted the challenge of 5G roll out in rural areas. Latency will be a challenge no matter which channels are used. The availability of timely and relevant data to any given vehicle will be important to maintaining safety.

The panel discussed long and short-range communication platforms and the need to create a holistic view of the world on the vehicle platform. The topic of sensor and perception fusion was discussed as a need for achieving the single world map view. The coordination, cooperation, and collaboration of multiple vehicles with different sensor outputs (i.e., sensor fusion) was generally agreed as a goal of this holistic view. From this single view, a "democratization of safety" (i.e., accessible to everyone) could be achieved on the road.

The panel also discussed the need for standards in the vehicle communication space. Panelists agreed that timing of standards is important. If a standard is too early, it stifles innovations, if it is too late, there is no market cohesion and progress. For autonomous vehicles, the panelists suggested a taxonomy of different environmental variables (speeds, road conditions, traffic data, and meta-data) that could be communicated between vehicles. There are communication standards for non-autonomous vehicles which could be extrapolated to on-road autonomous vehicles.

3.3. Cybersecurity & Privacy

Moderator: Katherine McClaskey, Cybersecurity and Infrastructure Security Agency

Panelists: Adam Mistick, Ford; Anuja Sonalkar, Steer; Christos Papadopoulos, University of Memphis; Amie Stepanovich, Future of Privacy Forum

The panelists began by discussing the unique challenges to cybersecurity in the on-road autonomous vehicle domain. While confidentiality and integrity are important, panelists agreed that availability is the critical security principle to focus on. The panelists agreed that ensuring the vehicle is working as intended and available to the driver are the biggest concerns with cybersecurity.

As on-road autonomous vehicles become more sophisticated and add more features, the complexity of those systems increases. This complexity accelerates cybersecurity concerns as the communication between features and components on the vehicle are vulnerabilities for malicious attackers to target. The panel discussed threat scenarios in which the attackers sent fraudulent signals to the vehicle or caused a denial of service of a component. These scenarios could render the vehicle unsafe to drive and thus pose a significant risk to the driver and others on the road.

A common topic of discussion was supply chain considerations. The current disruption in microchip processing is affecting on-road autonomous vehicles the same way that it is affecting personal computers. These disruptions cause the same types of threats to on-road autonomous vehicles and thus must be addressed. Alternatively, the panelists suggested that the on-road autonomous vehicles are similar to other systems that exist today while also having custom hardware and software technology. There are inputs and outputs, therefore, some best practices and concepts can be leveraged or translated to the on-road autonomous vehicle space.

Finally, the panelists discussed privacy and trust. The end user ultimately decides whether to trust the on-road autonomous vehicle and its capabilities. As the data collection of the on-road

autonomous vehicle platform increases, the ability for the system to maintain privacy becomes more difficult. Giving users a way to make informed choices about their data while also maintaining their safety is a key balance to strike when constructing cybersecurity and privacy for on-road autonomous vehicles.

3.4.   Artificial Intelligence

Moderator: Elham Tabassi, NIST

Panelists: Daniel Carruth, Mississippi State; Alberto Lacaze, Robotic Research; Hussein Mehanna, Cruise; Alfred Chen, University of California Irvine; Vijay Patnaik, Applied Intuition; Kevin Zaseck, Toyota Research Institute

Panelists highlighted important challenges facing the AI community in on-road autonomous vehicles. Challenges ranged from trustworthiness, security, safety, and complexity. Currently, the public is having an open debate on the trustworthiness of on-road autonomous vehicle AI. This debate stems from issues surrounding the safety of the vehicles whether from errors in object detection, data processing, or malicious actors. The development of safety measures remains difficult due to the complexity of breaking those metrics down into component level measures and tests. The panelists suggested that more conversation is needed around these topics while best practices are being developed.

The panel then discussed the topic of testing AI systems. Panelists suggested the use of simulators has shown benefits before using human testing to raise acceptance rates. The panelists also highlighted the balance of coverage and resource utilization. While AI is a new tool, there are previous examples of safety critical industries (such as the airline industry) that have best practices surrounding tooling. The community could leverage those tools and techniques while not duplicating previous efforts. Concurrently, panelists suggested that testing AI in on-road autonomous vehicles be conducted continuously and transparently because sensors (and their corresponding output) can decay, conditions can change, and algorithms can be updated. Therefore, all tests, along with their associated measures, should be monitored and communicated clearly.

The panelists turned their attention to the topic of AI safety. As with previous panels, the discussion centered around the dependency of safety on operational design domains (ODD). While a given scenario may be safe in one ODD, the same levels of safety will not suffice in another ODD. Flexibility in determining the context will be critical for AI safety. As a potential lens to view the challenge, the panel suggested a risk-based approach. Determining frequency and severity of a given incident or failure can be a way to assess risk and thus the mitigations and tests needed. While the panel agreed there was no such thing as a zero-risk solution, they also suggested that the mitigation must be less harmful than the original risk. Therefore, the panel suggested that AI systems should be able to determine when they are operating outside of risk parameters and take corrective actions based on a continuous feedback loop.

The panel also discussed the importance of data to AI solutions in autonomous vehicles. Panelists highlighted three broad issues regarding data curation and data traceability, which included the following. First, while the data within testing environments has produced positive results, the ability to trace issues back to the initial, incomplete data set remains a challenge. Second, there was discussion regarding approaches for data acquisition. It was noted that there

are currently no standard approaches to acquiring, generating, or mining relevant data for most use cases. Third, there was concurrent discussion that, when dealing with large data sets, the underlying distribution of the data may appear to be normal; however, the algorithm is biased against a certain outcome or demographic. Understanding bias in the data, and thus the AI, is important to consider when deploying autonomous vehicle systems.

3.5. Sensor Perception

Moderator: Kamel Saidi, NIST

Panelists: William Buller, Michigan Tech Research Institute; Rajeev Thakur, Velodyne; Brent Fisher, Auto Lidar Expert; Ryan Lamm, Southwest Research Institute; Zeb Barber, Aurora

Panelists discussed the challenges facing the sensor market in on-road autonomous vehicles. They highlighted the fact that there are many uncoordinated vendors. Each sensor defines their own metrics for reporting targets and physical measures. This lack of standard vocabulary on metrics leads to point-to-point decisions between vendors and original equipment manufacturers (OEMs) which increases the resourcing needed to compete in the market. Similarly, overall vehicle system requirements are not standardized. This lack of definition results in scattered sensor requirements and increases complexity for developing automotive solutions.

The panel also discussed the importance of the ODD in developing sensors. Currently a bottom-up approach is used where sensors are built and combined to create an ODD. Panelists suggested creating a standard definition of an ODD to provide a top-down approach. This approach would create valid use-cases for sensor fusion to achieve a specific mission. Panelists also commented that a common definition would allow for the industry to address gaps in current sensor technology by identifying the capabilities that are not present instead of assuming more of the same will solve the problem.

The panelists discussed benchmarks for sensors. Current practice is for OEMs to solicit a Request for Quote (RFQ) and attempt to find the best solution from the responses. This type of activity is driving development around specific challenges and the language of a given organization and not challenges that may be present across the market. The panelists suggested that creating benchmarks for both sensor level metrics (e.g., target size, target detection, number of pixels on target, frame rate) and algorithm level, sensor-agnostic metrics could be beneficial to standardizing the space. Additionally, the panel suggested that standardizing the vocabulary of the processing space would be helpful (e.g., perception, computation, actuation).

The panel also discussed the different issues of operating within the test range versus the real world. There are many variables that are present in deployment which are not accounted for in the lab, sometimes referred to as 'confusers' to a perception system. The panel provided many examples of these variables (e.g., road conditions, signal lights, signs, environmental conditions). These variables, even when considered, do not have standard definitions nor the ability to determine suboptimal values. The panel suggested that defining these variables and determining ranges for those variables would be helpful to increasing repeatability and reliability of sensor systems.

The panel also discussed sensor fusion and highlighted the balance between approaching sensor fusion from a hardware or software perspective. In the past, hardware was more

expensive and computing power was difficult to acquire; this led to more simple and elegant solutions in the software space. Conversely, hardware and computing power is more readily available to organizations and thus the constraints on software have relaxed. As the volume of data increases, the panelists suggested creating latency metrics and sensor fusion metrics to appropriately incentivize sensor processing hardware and software development.

4. Breakout Room Discussions

Breakout room discussions were conducted to solicit feedback on focus areas with respect to performance testing and standards in on-road autonomous vehicles through active audience participation. Each breakout room was facilitated by at least one professional moderator and was conducted through two, 90-minute sessions.

4.1. Safety

During the discussion, participants highlighted that safety was the top concern when considering on-road autonomous vehicles. Given its importance, participants suggested that it would be helpful for NIST to work with other federal agencies to develop appropriate safety risks. In addition, NIST could help define an appropriate level of risk which is a baseline for all autonomous vehicles. Note that NIST is a non-regulatory agency that cannot set policy, but can supply experimental data, information, and measurement techniques that can be used by others to help formulate policy. Participants also provided that acceptable risk is dependent on a given context and that a challenge will be to determine a common set of risks.

Participants also focused on the concept of testing safety requirements. Participants discussed the distinction between the relative challenges for test requirements for at-fault crashes versus not-at-fault crashes. The discussion did not reach any concrete conclusions. Participants discussed the potential to focus the safety question to examine the context of a situation to determine risk requirements.

While there is not sufficient knowledge to define a test that would determine that a vehicle is safe or not, any tests that are defined should include parameters for the random events that can occur within the environment of an autonomous vehicle. Given the wide range of parameters which could be defined, some participants suggested there was no credible test that can be designed to quantity safety in any meaningful way. To accommodate this gap, participants highlighted a need to think about the use-case of the technology when thinking about the test to deploy. This approach could produce reference scenarios, and metrics, which are tied to use-cases. If this approach is taken, use-cases could be selected by organizations, thus rendering the one-size-fits-all standard as negligible. While this outcome is not ideal, participants agreed, having a baseline standard would be beneficial as compared to having no standards.

Ultimately, participants highlighted that the guiding metric of safety is consumer trust. There could be tremendous value in the use-cases, tests, reference documentation, and metrics and that these approaches would be wasted work if the results were not communicated clearly to the public. Participants suggested that any approach taken by NIST should include a public outreach and communications component.

4.2. Communication

Participants began by discussing the different type of communication networks that exist (e.g., Wi-Fi, satellite, dedicated short-range communications, cellular). This discussion centered around the requirements of those networks for communication between the vehicle and the environment around it. Many participants mentioned performance requirements for networked communication (e.g., teleoperation, latency, processing time). Examples of work done both domestically and internationally with integrating different types of networks into a cohesive system were discussed.

The conversation turned towards the overall ecosystem of communications that will need to occur to ensure that on-road autonomous vehicles are providing as much functionality as possible. Participants mentioned the need for standards to normalize communication patterns within the vehicle, between the vehicle and infrastructure, and between the vehicle and the environment. Similarly, participants suggested providing ways for technology providers to integrate their systems together using communication protocols. As more stakeholders are added to the growing ecosystem of on-road autonomous vehicle communication, the need for a common language grows. Participants highlighted the need for business- and use-cases to be developed so organizations can see value in the standards.

Participants then shifted the discussion to data fusion. Many participants highlighted the processing location challenges. There are multiple approaches to data fusion (e.g., centralized, de-centralized, distributed, hierarchical, hybrid) and selecting one or standardizing on such an approach is difficult due to the varied nature of use-cases. Participants highlighted the need for a common language around the confidence and quality of data coming from a given sensor. It was agreed that a common world view for each vehicle is desired and therefore, sensor and data fusion require standardized metadata at multiple levels of aggregation (e.g., sensor, vehicle, environment, infrastructure).

Participants moved on to discussing the security of the communications of an on-road autonomous vehicle. Given the connected nature of on-road autonomous vehicles, participants raised concerns over malicious actors taking over legitimate teleoperation of the vehicle. Additionally, participants discussed unintended behavior within normal operating conditions, such as platooning and lemming effects. It was agreed that security of communications was an open question and warranted further conversation.

4.3. Cybersecurity & Privacy

The discussions identified numerous areas where NIST can assist with on-road autonomous vehicle considerations in both the research side and with the development of standards and best practices. The attendees recognized that there are cybersecurity elements in all aspects of on-road autonomous vehicles. The community can build upon and adapt many existing models, such as vulnerability assessments, software/code evaluation, and advances in integrated DevSecOps (development, security, and operations) (i.e., to introduce security earlier in the software development) methods. The complexity of on-road autonomous vehicles may prevent many of the current solutions from being directly applicable, so much work remains to adapt existing guidance for this emerging technology.

One element that needs further research is the risk imposed by after-market changes. Even if detailed and stringent security models are imposed at the manufacturing stage, there are post-delivery factors (unauthorized modifications, physical threats to street signs and other markers) that must be considered. As with more general cybersecurity considerations, supply chain provisions must be included in the on-road autonomous vehicles research and standards. The integrity of software and components is a critical factor.

Attendees pointed out that the vehicle life cycle is quite different from many other types of technology. Product cycles in traditional IT are short, while a vehicle may be expected to be operation for many years. Therefore, the longer cycle heightens the need for stringent and visionary planning. Furthermore, cryptography for securing internal and extra-vehicular communications will be a key consideration, possibly to include innovations in quantum computing. Processes for security software patching and updates will be a vital concern that must be consistently and reliably addressed. The community will also need to consider security implications for critical traffic infrastructure such as signage and intra-vehicle communications.

Breakout participants agreed that standards and guidance might include cybersecurity elements of roadworthiness criteria, possibly to include software currency elements, resilience considerations (i.e., how to address a loss of communications or a failed or attacked component), and domain-specific considerations (i.e., criteria for a passenger vehicle will be different from those for a truck hauling hazardous materials). These standards are likely to be regionally/geographically influenced.

A final element of the breakout session concerned privacy considerations. Participants agreed that a significant amount of data will be exchanged through the operation of on-road autonomous vehicles, and there will be concerns about what personal and private information is collected and monitored. This consideration represents an area where NIST may be able to provide some guidance, such as by recommending methods for opting in and what tiers of privacy data might be included. That guidance may consider privacy data ownership and discussion about who owns the data that describes where drivers have travelled and how the vehicle has been operated.

4.4. Artificial Intelligence

In the AI discussion, participants discussed both the general needs of the community for addressing risks to/from on-road autonomous vehicles and specific activities that NIST may be able to support. The initial conversation related to testing methods, including those for AI algorithms supporting vehicles. Because there are many differing contexts for different types of on-road autonomous vehicles, there is not a single set of tests that would fit all scenarios. Attendees felt that NIST (and the standards community, in general) may be able to develop test scenarios for specific ODDs. Attendees also noted that, because test conditions change, such as when sensors' performance degrade, any algorithms or test criteria developed must be adaptive enough to compensate for real-world conditions.

Participants expressed that, as with other areas such as cybersecurity and privacy, NIST is well-suited to provide measurement standards, but may not be an appropriate party to apply those measures nor to define acceptability criteria. NIST can help define the concept of conducting testing, including the units to be measured, and algorithms, methods, and models

that may be applicable for a given domain. Acceptability criteria might eventually be produced by legislators, insurance companies, and risk assessors.

The group agreed that a valuable contribution from NIST would be to provide consistent and standard terminology for the on-road autonomous vehicle industry. Many terms (e.g., transparency, explainability, predictability) are used in the industry but need more consistent definition. A common language is needed to support tools, techniques, and other criteria. Attendees felt that system designers and developers need actionable guidance. While the community has discussed many points of focus for decades, it is most helpful when guidance can be clearly understood and describe specific technical actions. Notably, there was a comment that safety and security are not the same, and even definition of those terms (and the difference between them) needs to be developed for this sector.

As discussion continued, attendees recognized that testing and information sharing may be hampered by the fact that systems, algorithms, and even previous results include proprietary data. There are also many layers of data – the actual vehicle engineering components may have an operational context, but there are AI aspects that may be distinct. Even if a product is certified as "secure" or "safe", there would need to be mechanisms for ensuring that no changes to the system might result in degraded safety or reliability.

At the close of the discussion, it was agreed that there is value in NIST assisting with definitions, glossaries, and test cases for common scenarios (e.g., for components, sub-components, and individual elements) that could provide better consistency in secure development and operation of on-road autonomous vehicles.

4.5. Sensor Perception

The breakout session began with a discussion on the challenges facing the space. These challenges included the volume, velocity, and veracity of the data being generated by the sensors. Specifically, participants highlighted the need for standards in the sensor data space to develop metrics on data quality. Additionally, participants noted that these types of standards would aid in the ability for upstream processing to make better decisions. Furthermore, participants noted that current traffic infrastructure is not sufficient to support on-road autonomous vehicles.

Participants considered the topic of LIDAR in on-road autonomous vehicles. Challenges were presented including no baseline of data, multiple sources of data, and no standardized return rate. Participants were split as to the necessity of LIDAR in every on-road autonomous vehicle. If a standard is desired in the LIDAR space, participants suggested collaborating with industry to align standards across the entire space, as well as determining definitions and performance requirements.

Participants discussed the distributed nature of computing within an on-road autonomous vehicle. There are multiple sensors on a vehicle computing their sensor specific data and feeding the answers to compartmentalized processes to make decisions. This approach results in a lack of transparency when a failure occurs. Some participants noted that a lack of transparency also exists in specifications for sensors. As a solution, participants suggested creating sensor standards, complete with definitions and performance metrics, to alleviate the transparency issue.

From a market perspective, participants discussed the current sensor integration process. Currently, testing and integration onto a vehicle is done in a custom build for each project, with little to no commonality between projects. There was disagreement between participants if interchangeability of sensors, data systems, or processing was possible. Some participants suggested common testing mechanisms to provide common language on object definition, object detection, environmental variables, and functionality of sensors and perception.

4.6. Overarching Themes

During the panels and breakout sessions of the workshop, the NIST team identified the following prevalent themes.

While the concept of on-road autonomous vehicle safety had its own track, many participants highlighted the notion that safety requirements would drive the other track's outcomes. This top-down conceptual organization requires that safety requirements be considered first and used as input into the other concept's workstreams.

Many participants across the sessions highlighted the need for performance requirements. These requirements were intended to be as specific as possible such as within the sensors and perception session, defining wavelengths of light and luminosity thresholds and within the safety session, defining metrics which could be used to measure safety other than miles driven. These performance requirements were not necessarily the same for each track but rather consistent in the desire for them to exist.

Within a few tracks, the concept of fusion was discussed. Whether it was sensor or data, participants highlighted the need to tackle the challenge of aggregating, aligning, and determining if multiple data points represent the same real object. Participants provided multiple approaches to this challenge (e.g., distributed, centralized, mesh) No matter the approach, creating reliable sensor fusion aids in achieving a holistic, "one world view", of the environment.

Finally, participants recognized that any efforts in the on-road autonomous vehicle space should be collaborative. NIST is committed to a public-private partnership model that seeks to achieve common understanding of the challenges ahead and the work to be done.

5. Next Steps

NIST would like to thank all participants, panelists, moderators, facilitators, and scribes for their valuable contributions to the workshop. Based on these contributions, NIST is drafting a program plan for work to be done in the future. Please check back in with the program at our website   https://www.nist.gov/programs-projects/nist-and-autonomous-vehicles   to   receive updates on next steps.

NIST is always open to feedback. Please email questions, comments, or suggestions to the following email address: autonomousvehicles@nist.gov.

Appendix A: Workshop Agenda

(All times are Eastern Standard Time, UTC-5)

Day 1: Tuesday, March 8, 2022

Time	Topic
9:00 AM - 9:30 AM	**Welcome to NIST** **James K. Olthoff,** Performing the Non-Exclusive Functions and Duties of the Under Secretary of Commerce for Standards and Technology & Director, National Institute of Standards and Technology **Opening Remarks** **Nellie Abernathy,** Department of Commerce
9:30 AM - 10:00 AM	**Keynote Speech** **Trent Victor,** Waymo
10:00 AM - 11:00 AM	**Panel 1: Safety** For autonomous vehicles to be accepted, they must be trusted and perceived as safe systems. This panel discusses the requirements for testing conditions and evaluating safety of autonomous vehicles and the current state of Operation Design Domains (ODD). • **Ed Griffor,** NIST **(moderator)** • **Chan Lieu,** Aurora • **Ed Straub,** SAE • **Jack Weast,** Intel • **Bert Kaufman,** Zoox • **Robert Dingli,** Plus • **Robyn Robertson,** Traffic Injury Research Foundation **Panel Organizer:** Craig Schlenoff, craig.schlenoff@nist.gov
11:00 AM - 12:00 PM	**Lunch Break**
12:00 PM - 1:00 PM	**Panel 2: Communication** Quality communication protocols are critical in monitoring, messaging, and teleoperation of connected and autonomous vehicles. This panel discusses testing methods and messaging requirements for quality communication (both peer-to-peer and cellular) in on-road autonomous vehicles. • **Scott McCormick,** Teleoperation Consortium **(moderator)**

Time	Topic
	• **John McNicol,** British Standards Institute • **Stan Schneider,** Real-Time Innovations • **Maarten Sierhuis,** Nissan North America • **Fan Bai,** General Motors • **Bettina Erdem,** Continental Corporation **Panel Organizer:** Tao Zhang, tao.zhang@nist.gov
1:00 PM - 2:00 PM	**Panel 3: Cybersecurity and Privacy** As vehicles become more connected and more automated functionalities are introduced, these complex systems must be secure and protect privacy while adhering to strict real-time quality constraints. This panel discusses standards and taxonomy in defining cybersecurity and privacy for on-road autonomous vehicles. • **Katherine McClaskey,** Cybersecurity and Infrastructure Security Agency **(moderator)** • **Adam Mistick,** Ford • **Anuja Sonalkar,** Steer • **Christos Papadopoulos,** University of Memphis • **Amie Stepanovich,** Future of Privacy Forum **Panel Organizer:** Suzanne Lightman, suzanne.lightman@nist.gov
2:00 PM - 2:30 PM	**Break**
2:30 PM - 3:30 PM	**Panel 4: Artificial Intelligence** Autonomous vehicles conduct a variety of automated cognition functions including lane following, path planning, and classification. This panel discusses metrics, datasets, and simulated/physical testing methods for evaluating artificial intelligence in on-road automated vehicles. • **Elham Tabassi,** NIST **(moderator)** • **Daniel Carruth,** Mississippi State • **Alberto Lacaze,** Robotic Research • **Hussein Mehanna,** Cruise • **Alfred Chen,** University of California Irvine • **Vijay Patnaik,** Applied Intuition • **Kevin Zaseck,** Toyota Research Institute **Panel Organizer:** Vinh Nguyen, vinh.t.nguyen@nist.gov
3:30 PM - 4:30 PM	**Panel 5: Sensor Perception**

Time	Topic
	The driving environment consists of multiple obstacles to perception (both manual and automated) including occlusion, variable reflectivity, and edge cases. This panel discusses testing conditions and metrics for evaluating sensor data and fusion techniques in autonomous vehicles. • **Kamel Saidi,** NIST **(moderator)** • **William Buller,** Michigan Tech Research Institute • **Rajeev Thakur,** Velodyne • **Brent Fisher,** Auto Lidar Expert • **Ryan Lamm,** Southwest Research Institute • **Zeb Barber,** Aurora **Panel Organizer:** Prem Rachakonda, prem.rachakonda@nist.gov
4:30 PM - 5:00 PM	**Closing Remarks**
5:00 PM	**Adjourn for the day.**

Day 2: Wednesday, March 9, 2022

Time	Topic
10:00 AM - 10:30 AM	**Day 2 Opening** **Tim Kurth,** Chief Counsel (R), United States House of Representatives, Energy and Commerce Committee Consumer Protection and Commerce Subcommittee **Recap of Day 1**
10:30 AM- 11:00 AM	**Keynote Speech: Adam Campbell,** Gatik
11:00 AM - 12:30 PM	**Breakout Session 1** **Breakout Rooms/Topic Areas:** • Safety • Communication • Artificial Intelligence • Cybersecurity • Sensor Perception
12:30 PM - 1:30 PM	**Lunch Break**

Time	Topic
1:30 PM - 3:00 PM	**Breakout Session 2** **Breakout Rooms/Topic Areas:** • Safety • Communication • Artificial Intelligence • Cybersecurity • Sensor Perception
3:00 PM - 3:30 PM	**Break**
3:30 PM - 5:00 PM	**Breakout Debrief** Breakout facilitators will host a panel to discuss findings. **Closing Remarks**
5:00 PM	**Adjourn**