

Trusted Autonomous Vehicles: an Interactive Exhibit

Hugo L. S. Araujo
Carlos Diego N. Damasceno
Rayna Dimitrova

Genovefa Kefalidou
Mehdi Mehtarizadeh
Mohammad Reza Mousavi

Jemima Onime
Jan Oliver Ringert
José Miguel Rojas

Nervo Xavier Verdezoto
Syed Wali

*School of Informatics
University of Leicester (UK)*

Abstract—Recent surveys about autonomous vehicles show that the public is concerned about the safety consequences of system or equipment failures and the vehicles’ reactions to unexpected situations. We believe that informing about the technology and quality, e.g., safety and reliability, of autonomous vehicles is paramount to improving public expectations, perception and acceptance. In this paper, we report on the design of an interactive exhibit to illustrate (1) basic technologies employed in autonomous vehicles, i.e., sensors and object classification; and (2) basic principles for ensuring their quality, i.e., employing software testing and simulations. We subsequently report on the delivery of this exhibit titled “Trusted Autonomous Vehicles” at the Royal Society Summer Science Exhibition 2019. We describe the process of designing and developing the artefacts used in our exhibit, the theoretical background associated to them, the design of our stand, and the lessons learned. The activities and findings of this study can be used by other educators and researchers interested in promoting trust in autonomous vehicles among the general public.

Index Terms—Autonomous vehicles, Royal society, Science fair, Software quality, Science communication

I. INTRODUCTION

The automotive industry is facing a drastic change of paradigm: computer systems (and in particular AI-enabled software) are taking the central stage in an industry which was traditionally centred around mechanics and later mechatronics. It is predicted that by 2025 (and possibly much sooner than that) computer systems will account for more than half of the value of a modern car [1] and with the sharp decrease in the price of hardware, software will be the dominant cost factor [2]. Moreover, the absolute majority of the innovation in modern cars is embodied by software [2]–[4].

This new paradigm is a great enabler for more automated and autonomous functions. The Society of Automotive Engineers (SAE) classifies the levels of autonomy in 6 levels [5]: from no automation (level 0) to full automation under all circumstances without any need for a handover of control (level 5). Advanced driver-assistance systems (ADAS), which take over much of the driving tasks (under the assumption of the possibility of rapid handover to a human driver), are already present in modern high-end cars. There are commercial

level 2 and 3 autonomous vehicles, which take over all driving functions under predefined circumstances, yet these systems still rely on the possibility of a rapid handover of control to human drivers. Several pilots of level 3 and 4 autonomy are ongoing across the world, such as the ones by the American company Waymo¹ or the British company Oxbotica², to name a few.

Drawing a precise timeline for the wide adoption of higher levels of automation still remains challenging. However, there is a steady and consistent trend towards more autonomy. It is well known from the available body of literature, regarding autonomous vehicles [6]–[8] and other autonomous technologies (e.g., in the healthcare domain), that raising public understanding of these technologies and establishing a level of trust is a key factor in their widespread adoption. There are clear indications that we still lack effective means of establishing a basic understanding about the safety and trust in autonomous vehicles and many doubts and misconceptions prevail in the general public [9], [10]. Hence, coming up with an effective model of public engagement in this domain is a very relevant and timely challenge, which is addressed in this paper.

In this paper, we report on a set of activities organized to raise public understanding and awareness about connected and autonomous vehicles. The activities were organized at a prestigious 7-day event at the Royal Society and built upon several years of research at our group at the University of Leicester. The exhibition attracted about 12,000 visitors and our statistics show that some 4,000-5,000 people engaged in our various activities. We report on the social and theoretical background of our activities and the details of their technical design. We also report on our experience in public engagement and the most prominent concerns put forward by the general public. The rest of the paper is organized as follows. In Section II, we review the literature regarding public trust in autonomous systems and vehicles, the technical design, and safety analysis of connected and autonomous vehicles. In Section III, we describe the technical design of our activities aimed to raise public awareness and understanding

This work is supported by the Leicester AI Network, through a Research Network grant awarded by the Leicester Institute for Advanced Studies, the Royal Society Summer Science Exhibition, and Zenzie

¹<https://waymo.com>

²<https://www.oxbotica.com>

about autonomous vehicles. In Section IV, we report on the exhibition in which these activities were deployed. Section V is dedicated to observations from the exhibition and lessons learned. Finally, we conclude in Section VI by reviewing our main achievements and proposing some avenues for future research.

II. BACKGROUND

A. Trust in autonomous vehicles

Trust in a system is defined as the belief or attitude that the system is helpful in achieving the user’s goal, particularly in uncertain situations [11]. Trust is studied in a myriad of contexts using different methodologies, but the present paper focuses on the issue of trust in human-vehicle interaction, particularly in the case of connected and autonomous vehicles.

Lee and Moray [12] propose a general theory for trust in automation; their theory assumes the persistence of natural laws and it describes trust in terms of three dimensions: performance (what the system does), process (how the system is built), and purpose (why the system does something). We address all three dimensions of trust in our activities, but our primary focus is on performance (particularly, safety) and process (particularly, design). Inspired by earlier studies [6]–[8], [13], [14], our hypothesis is that ensuring safety and understanding the design of autonomous vehicles contribute significantly to establishing trust.

Hoff and Bashir [15] perform a secondary study on the literature on trust and come up with an extensive model of trust classification in terms of three general categories: dispositional trust (personal and cultural characteristics of the user), situational trust (contextual factors regarding the system and the user for example, in terms of workload, attention, and personal mood) and learned trust (e.g., pre-existing knowledge about the system and observed system performance). Our activities address learned trust, particularly in the context of safety.

Körber, Baseler and Bengler [14] report that trust can be directly influenced by a prior training and also show that elevated trust leads to lower number of intervention in the autonomous behavior by human users. Our activity builds upon their observation regarding the possibility of elevating trust and its positive effect on the use of autonomous functions.

Recent results [6]–[8] suggest that trust is one of the key factors in the public adoption of connected and autonomous vehicles (CAVs) and technical transparency is one of the pillars of establishing trust. That is why trust [16], explainability, and perspicuity [17] have become the centre of research attention in this domain in the recent years.

B. Autonomous vehicle software design

Large research projects on autonomous driving go back to the U.S. Automated Highway System project in the nineties [18]. More recently, some experiments have explored autonomous vehicles based on a single machine learning component [19], [20]. However, most current approaches – including industrial-strength ones – rely on a decomposed

software architecture design. A high-level architecture of an autonomous vehicle is often composed of four elements: sensing, perception, planning and control [21]–[23]³. While holistic approaches aim to develop optimal system architecture designs, each of these elements rely on a collection of well-established technologies, such as object detection [24] and classification [25], obstacle detection [26] and motion planning and control [27].

Several automotive software platforms have achieved various levels of maturity in recent years, e.g., AutoSAR [28], Autoware [29] with ROS [30], and Android Automotive [31]. These platforms often prescribe architectural styles, e.g., components with well-defined interfaces in AutoSAR or publish-subscribe architectures in ROS.

Our activities are informed by our ongoing research in model-based design, design-space exploration, and analysis of cyber-physical and automotive systems; in particular, we have been using techniques such as model-based testing [32], [33] and search-based testing [34], [35] to perform safety analysis of automotive systems and architectural transformation and synthesis [36] for their design. We used the outcomes of these lines of research to inform the design of our activities described in Section III.

III. DESIGNED ACTIVITIES

A. Autonomous vehicle technology demonstrator

Rationale. In a recent study in the UK on machine learning, self-driving cars were cited as “having the greatest risk to society” [10]. We believe that a better understanding of the technology employed in CAVs and a demonstration of their safety features is a necessary step for the public to gain a proper understanding and reasonable trust in CAVs. Towards this aim, we have designed an activity featuring an autonomous car (scale 1:8 Audi Q2 shown in Fig. 1) driving on a 3m x 5m track. This activity demonstrates how sensors, here a camera and a LIDAR, help CAVs perceive the surrounding world. Visitors can press a buzzer that moves a pedestrian in front of the car right onto the middle of the track. Visitors can then observe how the car automatically stops within a safe distance from the pedestrian.

System implementation. Our system implementation consists of three parts:

a) *Car model and software:* Our demonstrator is based a scale 1:8 Audi Q2 model car. The car contains an Intel Core i3 processor, 16 GB RAM, a 128 GB M.2 SSD hard drive and an NVIDIA GeForce GTX1050Ti graphics card for computing. It is equipped with a LIDAR (RPLIDAR A2, 180° field of view, detection range <6m, update rate 10Hz), a 130° mono video camera (front), an 80° mono video camera (back), wheel speed sensors, and a 9-axis motion tracking sensor.

The car runs an Ubuntu OS with Tensorflow⁴, OpenCV⁵, and ADTF (Automotive Data and Time-Triggered Frame-

³alternative decomposition strategies exist, e.g., perception, integration, control and actuation

⁴<https://www.tensorflow.org/>

⁵<https://opencv.org/>




Fig. 1. An autonomous vehicle (Audi Q2, scale 1:8) and a model pedestrian controlled by a buzzer to enter the driving path of the vehicle

work)⁶. ADTF is a framework for the model-based configuration and scheduling of automotive software components. Software components are written in C++11 and pre-processed for scheduling and execution in an ADTF session.

The ADTF component and connector architecture of our AV is shown in Fig. 2. Components on the left side provide sensor readings to the software, e.g., component `LaserScanner` provides a set of points (angle and distance) detected by the LIDAR. Components on the right provide access to actuators and virtual displays of the car. As an example, the component `MotorActuator` takes as an input power, a percentage of the maximal voltage for the motor, and `steering`, the desired steering angle. Component `WheelSpeedController` takes as input both the current and the desired speed and computes the target power for the motor actuator (computed by a proportional derivative controller [37]).

The main functionality of our CAV demonstrator is to stay within the driving lane, i.e., to calculate a desired speed and steering angle. There are multiple options for calculating these parameters. A popular approach for steering is to rely on visual lane detection in combination with high-resolution maps (see [38] for a recent survey). For our exhibit with uncertain light conditions and possible reflections on the track, we chose not to rely on processing of the camera image. Our solution is based on the 180° scan of the LIDAR mounted in the front of the car. Roughly, component `LaserDrive` takes as input the points where the LIDAR has detected obstacles and it calculates the speed and steering to maintain a safety distance from the outer road boundary (see the red and white tape in Fig. 1).

Another important functionality of our CAV technology demonstrator is to detect and stop for obstacles. Component `ObstacleDetection` takes as input the current steering

angle, speed, and points detected by the LIDAR. The output of the component is an adjusted speed value. In case an obstacle is in the path that the car is driving (calculated from the speed and steering angle) the adjusted speed value will be 0, i.e., the car will come to a stop. To explain how the LIDAR works and to illustrate the emergency stop feature, we chose to visualize the points detected by the LIDAR. A visualization is shown inside component `VideoDisplay` in Fig. 2. The visualization shows an upcoming curve. The partial ellipse at the bottom of the image is the area in front of the car tilted by the current steering angle. If the LIDAR detects an obstacle within the area in front of the car, the visualization changes the color of the dots and component `ObstacleDetection` will reduce any positive speed to 0 to prevent a collision.

Finally, our exhibit displays the camera image and the LIDAR visualization on a large screen that is connected to a Raspberry Pi 3B⁷ running the Raspian OS. The Raspberry Pi connects to the AV via NoMachine⁸ to show the visualizations produced by the ADTF components shown in Fig. 2, right.

b) Object detection and classification: For the general public in the UK, artificial intelligence (AI) is a controversial topic that is not thoroughly understood but often carries negative perceptions [39]. We aimed to explain one of the basic AI components of CAVs: object detection and classification from camera images [24], [25]. On the one hand, we have employed this activity to show that AI in CAVs usually performs well-defined, small-scoped tasks that it is good at – the interactivity of our exhibit allowed visitors to try it out themselves. On the other hand, the sensor visualization of the LIDAR clearly motivated the need for understanding *what* it is that is picked up by the sensor – a CAV’s reaction should differ whether it is a pedestrian or some other object, e.g., a floating plastic bag.

Technically, we have used a pre-trained neural network to detect and classify objects in images that were recorded by a camera facing our visitors. Specifically, we have used OpenCV to capture images, analyzed these using a You Only Look Once (YOLO) [40] neural network (YOLOv3)⁹, and displayed the results of the detection on the screen in real-time. The software was executed on regular desktop computer with an Intel Core i5 processor and an NVIDIA GeForce GTX1050 graphics card running an Ubuntu OS.

c) Interactive pedestrian: To demonstrate a safety feature of CAVs we have decided to automate a model pedestrian to enter the road and force the car to perform an emergency stop (as shown in Fig. 1). The pedestrian is interactively controlled by the user on the pressing of a buzzer. The pedestrian stays on the road until the buzzer is pressed again.

The mechanism to move the pedestrian is based on Lego EV3¹⁰. The EV3 computer has an ARM 300MHz CPU and we installed EV3dev¹¹, a Debian Linux-based OS with support for Python 3. We have added a WiFi dongle to the EV3 to

⁶<https://www.elektrobit.com/products/automated-driving/eb-assist/adtf/>

⁷<https://www.raspberrypi.org/products/raspberry-pi-3-model-b/>

⁸<https://www.nomachine.com/>

⁹<https://pjreddie.com/darknet/yolo/>

¹⁰<https://www.lego.com/en-gb/mindstorms/about-ev3>

¹¹<https://www.ev3dev.org/>

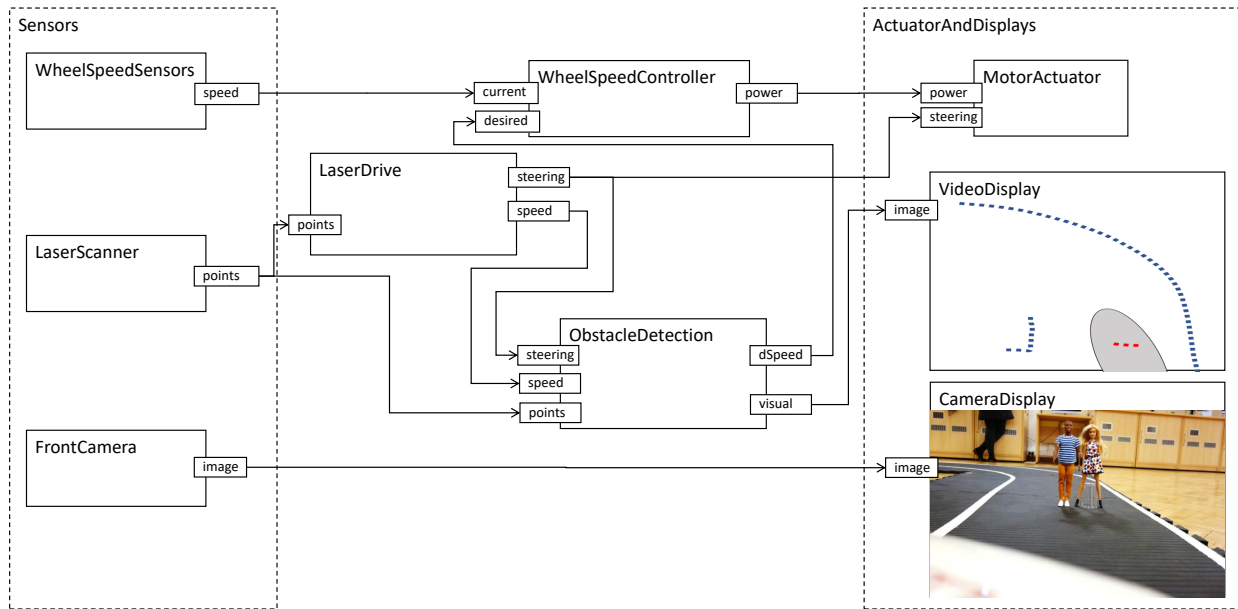


Fig. 2. A simplified Component and Connector model of the main software components of the autonomous vehicle as implemented in ADF and C++

communicate with an ESP8266 Arduino chip (NodeMCU v3). The ESP8266 is connected to a buzzer.

We use the ESP8266 chip as a WiFi access point. The EV3 connects to it and starts a Python web server based on the class `SimpleHTTPServer`. Once the button is pressed, the ESP8266 sends an HTTP request to its connected EV3 to operate the motors and move the pedestrian.

As a safety measure to prevent damage to our equipment, we have added an ultrasonic sensor to the EV3 to detect the car. If the car is detected by the sensor and it is too close, the pedestrian would hit the side of the car if it was moved. In these cases we prevent the pedestrian from entering the road.

B. Experiencing validation and verification in a driving simulator

Rationale. One of the main technical challenges [41] and public concerns [42] relates to how CAVs manage and respond to rare events. To handle rare events, such as birds flying over the road or cyclists suddenly crossing the road, engineers often rely on virtual testing environments such as driving simulators [43]. Moreover, ensuring safety against the virtually infinite possible scenarios of use and reasonable misuse of CAVs poses a formidable challenge for validation and verification. Hence, using intelligent techniques, such as evolutionary algorithms, to generate scenarios [34] and deriving quantitative measures and guarantees of safety are among the essential steps towards establishing trust in CAVs.

To illustrate how CAVs can be tested and how their safety analysis can be automated, we have designed an activity where visitors could experience the actual process of testing and safety assurance of autonomous vehicles. To achieve this, we have designed scenarios and populated them with obstacles (e.g., pedestrians, cyclists, and cars running street junctions) in

order to create potentially challenging situations that can allow for quantitatively measuring the safety of driving behaviors. To demonstrate this to the general public, we let the public drive these scenarios and compared their behavior to our calculated optimal behaviors. To this end, we have used the CARLA driving simulator [44] as a test environment.

System implementation. CARLA (Car Learning to Act)¹² is an open-source simulator for autonomous driving research. It has been developed using the Unreal engine to support training, prototyping and validation of autonomous driving models [44]. The CARLA project consists of two modules: the simulator module that controls the logic, physics, and rendering of actors and sensors in the scenarios; and the Python API that provides an interface to control and retrieve data from the simulator. The source code of the CARLA simulator, their Python API library and other side projects, e.g., end-to-end driving via conditional imitation learning [20], reinforcement learning [44], and test scenario configuration; are freely available through GitHub [45].

The `scenario_runner` project [46] includes an execution engine to configure test scenarios in the CARLA simulator. For our work, we have extended the project with additional driving scenarios and a user interface for the purpose of this exhibition. Visitors could control an ego car using a steering wheel and pedals. The goal is to let the exhibit’s visitors experience scenarios that can be used to assess the safety of self-driving AIs, and to explain how our research leads to automatically generating such challenging scenarios. Figure 3 shows a snapshot from the CARLA simulator.

Each new scenario comprises a 250 meters route populated with challenging events and weather conditions. Up to four events were designed to be triggered in each driving session:

¹²<http://carla.org/>




Fig. 3. Snapshot from the CARLA driving simulator



Fig. 4. Snapshot of the final score

- 1) a cyclist coming out of an alleyway, building or enclosure,
- 2) a car backing out of a driveway,
- 3) a static accident on the road and
- 4) a car running a red light.

These events were intended and designed to maximize collision potential; obstacles were placed in areas of the map where accidents are most likely to happen, e.g., roundabouts, street junctions, and places with high visual clutter. The placement also considered the trajectory and velocity of the ego car to project a point where to trigger events, i.e., where cyclists and cars move towards the ego vehicle with an appropriate velocity.

At the end of the driving session a score is assigned. Figure 4 shows the final screen showed to visitors with their respective final score. For each scenario, there is a time limit of 100 seconds; the faster a scenario is completed the higher the final score given to the visitor. Once the track is completed, the remaining time is multiplied by 100 to compute a partial score. However, each misconduct results in a weighted penalty. Thus, the score range is $]-\infty, timeLimit * 100]$. A positive score means that the ego car has completed the route fast with very few misconducts. Table I shows the penalty associated with each misconduct.

The score is designed so that any collision with drivers, pedestrians or cyclists should result in a negative score. Furthermore, to discourage shortcuts, which could result in a higher score, the penalty for driving off-route is high. A future plan for adjusting the score is to weight the collision based on its intensity and to set speed ranges for each part of

Misconducts	Penalty points
Collision with static objects	500 per collision
Collision with cars, cyclist and, pedestrians	5000 per collision
Driving away from the predefined path	100 per meter
Route not completed	100 per missing percent

TABLE I
SCORE PENALTY

the track and penalize cars driving outside of these ranges.

IV. VENUE AND STAND

The Royal Society Summer Science Exhibition¹³ is a UK, high-profile science event that took place in July 2019. In seven days the event attracted more than 12,000 visitors from schools and the general public. Our exhibit with the title *Trusted Autonomous Vehicles* was selected as one of 21 exhibits from more than 70 proposals by UK research institutes and universities.

An overview of our stand is shown in Fig. 5. The CAV demonstrator is shown at the top. It consists of a 3m x 5m track on an elevated platform to provide a good view of the car. Visitors have access to a buzzer (shown below the track) to control the model pedestrian entering the track. A TV screen, mounted on the back wall above the track, allows visualizing the LIDAR scan and camera of the car.

Visitors do not have direct physical access to the track but only to the 2m x 4m area shown at the bottom of Fig. 5. The visitor area has a wall with printed information to the right. The driving simulator is next to the wall on the right and faces the track to allow other visitors to see the screen from the back. A terminal is mounted on a plinth at the left outer corner of the stand. This terminal uses a camera to demonstrate real-time object detection and classification based on AI.

The arrangement allowed us to maximize the number of people who can view the CAV demonstrator, and to manage the queue of people who wait for using the driving simulator. The terminal and camera facing outwards of the stand made it possible to discuss with additional visitors and helped attract passersby.

Our team consisted of 15 exhibition staff members (academics and students). The team received visitors at the stand for up to 11h per day. At every point in time we had scheduled 6 staff members to interact with visitors. Our statistics indicate approximately 5,000 visitors to our stand throughout the exhibition week.

V. LESSONS LEARNED

Developing and delivering the exhibit was a valuable experience from various perspectives: from designing activities for large public to learning about the concerns and the attitude of the general public towards CAV, about which we briefly report below.

¹³<https://royalsociety.org/science-events-and-lectures/2019/summer-science-exhibition/>




Fig. 5. Layout of the exhibition stand. The 3m x 5m track is elevated 60cm above the 2m x 4m stand area

Regarding the design of the stand, an important observation was that hosting thousands of visitors through a one-week period essentially relied on having a range of activities: the object detection activity at the entrance of our exhibition served as an activity that could engaged a large number of people at once and serve as an introduction to our exhibition without requiring one-to-one interaction between visitors and exhibitors. After a brief capturing interaction with a visitor or group of visitors, the exhibitor was able to lead them to an introduction to other sensors, to pushing the pedestrian in front of the car and discuss its behavior and visualization elements, and finally, to driving in the simulator. Had all activities required a one-on-one personal interaction with visitors, we would not have had the capacity to deal with the large volume of visitors expected. Through the activities with personal interaction, we had a chance of providing a more in-depth information and discussion with the more engaged visitors.

a) Autonomous vehicle demonstrator: We were able to use the CAV demonstrator for various purposes. On the one hand, the simple setup of the demonstrator and the sensor visualizations made it possible for visitors to understand how the system operates and what sensors it relies on. On the other hand, we could demonstrate the emergency break safety feature that is already available in commercial cars today. This seemed to convince many visitors of the value of research towards CAVs for increased road safety.

Many visitors, old and young, were trying to get the pedestrian hit by the car by sending it onto the road when the car was very close. A sensor of the pedestrian avoided moving onto the road at the very last minute to prevent physical damage of the equipment. This behavior is not very realistic: there will be driving situations where no vehicle could avoid a collision.

One might develop an extension of the scenario – where collisions happen – to demonstrate these limitations. Visitors experimented with the demonstrator in additional ways beyond what we had envisioned. As an example, visitors placed mobile phones and other types of obstacles on the track to test the CAV. This behavior suggests not only enhanced engagement levels but also strong interest from visitors on technology aspects and safety. These interactions also led visitors to engage in many interesting discussions with the exhibitors.

b) Test scenarios in driving simulator: The driving simulator was the most popular of our three activities. Table II displays minimal, median, and maximal driving scores across all participants. We believe that the simulator activity demonstrated to the public that humans do not always have the skills or reaction time to handle unexpected events and that our driving scenarios are able to uncover similar limitations.

	Human collisions	Static collisions	Distance off-track	Time left	Route completed	Final score
Min	0	0	0 m	0 s	0 %	-582,400
Median	2	3	47 m	12 s	100 %	-20,810
Max	25	81	5,356 m	79 s	100 %	5,844

TABLE II
MINIMAL, MEDIAN, AND MAXIMAL DRIVING SCORES

Regarding the driving scores, we have noticed that only 11% of the participants (i.e., 128 out of 1114) obtained a positive score, with the highest score being 5,844 while the lowest is -582,400 points. Figure 6 depicts an histogram which shows the number of people and their respective score.

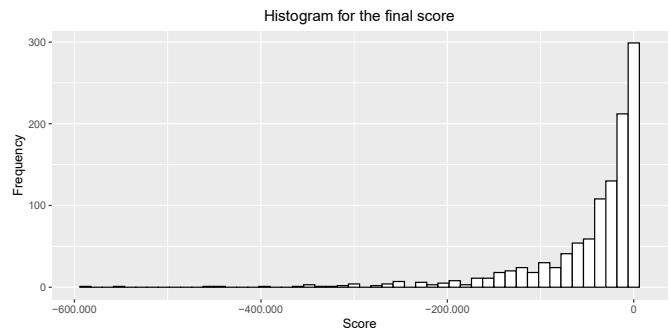


Fig. 6. Histogram for the final score

A. Frequent questions and concerns

We now list some of the questions and concerns that were frequently brought forward to us by visitors.

Most visitors asked us about a timeline or date for CAVs to become prevalent on public roads. Many visitors were curious about liability issues and legal implications in case of accidents. A few visitors asked about decisions in ethical dilemmas [47], where an accident cannot be avoided and the CAV has to decide whom to harm.

Visitors also shared several concerns with us. Some feared that jobs will be lost by automation; others were afraid that cars will become more vulnerable to cyber-attacks. Some

expressed concerns that due to CAVs the demand for mobility will increase and hence, roads will become more congested. Finally, some were concerned about losing autonomy because they will not be in charge of driving themselves anymore.

In anticipation of some of these controversial questions, we had incorporated different pieces of information and statistics to our stand design, e.g., about the different levels of autonomy, projected statistics about road safety, and projected figures about the jobs created in the CAV market. For some others questions and concerns, we gradually produced educational material, particularly in terms of interviews with the press and made them available to the public.

VI. CONCLUSIONS

Based on our research activities and using open-source, free, and free-for-academic-use software, we designed three activities to inform the general public about the principles behind the design, validation, and verification of connected and autonomous vehicles (CAVs). Our activities comprised: 1) a prototype implementation of an autonomous vehicle on a model car equipped with a camera, a LIDAR and USS sensors, with which the public could interact by pushing a pedestrian in front of the car, 2) an object detection software with which the general public could interact using various objects, and 3) a simulator where challenging driving scenarios are populated with rare and hazardous events, in which the visitors could drive and compare their behavior with safe behavior using our quantitative measure.

Our research in this domain continues both on the technical front and on the human-machine interaction side. On the technical front, we are developing a toolset for generating challenging scenarios using various evolutionary algorithms [34]. On the human-machine interaction side, we are measuring public perception of trust and our influence by training the public in the basic concepts of safety in CAVs.

REFERENCES

- [1] M. Roemer and A. Kramer, *The Intelligent Car*. A.T. Kearney, 2010.
- [2] J. Mössinger, "Software in automotive systems," *IEEE Software*, vol. 27, no. 2, pp. 92–94, 2010.
- [3] H. Aerts and H. Schaminée, "How software is changing the automotive landscape," *IEEE Software*, vol. 34, no. 6, pp. 7–12, 2017.
- [4] A. Pretschner, M. Broy, I. H. Kruger, and T. Stauner, "Software engineering for automotive systems: A roadmap," in *Future of Software Engineering (FOSE '07)*. IEEE, 2007, pp. 55–71.
- [5] Society of Automobile Engineers, "Taxonomy and definitions for terms related to on-road motor vehicle automated driving systems," available from: https://www.sae.org/standards/content/j3016_201806/, Accessed August 2019.
- [6] J. K. Choi and Y. G. Ji, "Investigating the importance of trust on adopting an autonomous vehicle," *International Journal of Human-Computer Interaction*, vol. 31, p. 697–702, 2015.
- [7] S. Nordhoff, J. de Winter, M. Kyriakidis, B. van Arem, and R. Happee, "Acceptance of driverless vehicles: Results from a large cross-national questionnaire study," *Journal of Advanced Transportation*, 2018, article ID 5382192.
- [8] T. Zhang, D. Tao, X. Qu, X. Zhang, R. Lin, and W. Zhang, "The roles of initial trust and perceived risk in public's acceptance of automated vehicles," *Transportation Research Part C: Emerging Technologies*, vol. 98, pp. 207 – 220, 2019.
- [9] B. Schoettle and M. Sivak, "A survey of public opinion about autonomous and self-driving vehicles in the us, the uk, and australia," University of Michigan, Ann Arbor, Transportation Research Institute, Tech. Rep., 2014. [Online]. Available: <https://deepblue.lib.umich.edu/handle/2027.42/108384>
- [10] Ipsos MORI, "Public views of machine learning," Royal Society, Tech. Rep., 2017, findings from public research and engagement conducted on behalf of the Royal Society.
- [11] J. D. Lee and K. A. See, "Trust in automation: Designing for appropriate reliance," *The Journal of the Human Factors and Ergonomics Society*, vol. 46, pp. 50–80, 2004.
- [12] J. D. Lee and N. Moray, "Trust, control strategies and allocation of function in human-machine systems," *Ergonomics*, vol. 35, pp. 1243–1270, 1992.
- [13] M. Hengstler, E. Enkel, and S. Duelli, "Applied artificial intelligence and trust – the case of autonomous vehicles and medical assistance devices," *Technological Forecasting and Social Change*, pp. 105–120, 2016.
- [14] M. Körber, E. Baseler, and K. Bengler, "Introduction matters: Manipulating trust in automation and reliance in automated driving," *Applied Ergonomics*, vol. 66, pp. 18 – 31, 2018.
- [15] K. Hoff and M. Bashir, "Trust in automation: Integrating empirical evidence on factors that influence trust," *The Journal of the Human Factors and Ergonomics Society*, vol. 57, p. 407–434, 2015.
- [16] M. Kwiatkowska, "When to trust a self-driving car..." 2019, Milner Award Lecture at the Royal Society.
- [17] H. Hermans, "Towards a science of perspicuous computing - lessons learnt from the analysis of automotive emissions control systems," 2019, keynote address at the 4th Workshop on Formal Reasoning about Causation, Responsibility, and Explanations in Science and Technology.
- [18] Transportation Research Board, *National Automated Highway System Research Program: A Review: A Review—Special Report 253*. Transportation Research Board, 1998, vol. 253.
- [19] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, X. Zhang, J. Zhao, and K. Zieba, "End to End Learning for Self-Driving Cars," *arXiv e-prints*, Apr 2016.
- [20] F. Codevilla, M. Müller, A. López, V. Koltun, and A. Dosovitskiy, "End-to-end driving via conditional imitation learning," in *International Conference on Robotics and Automation (ICRA)*, 2018.
- [21] M. Campbell, M. Egerstedt, J. P. How, and R. M. Murray, "Autonomous driving in urban environments: approaches, lessons and challenges," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 368, no. 1928, pp. 4649–4672, 2010. [Online]. Available: <https://doi.org/10.1098/rsta.2010.0110>
- [22] J. Leonard, J. How, S. Teller, M. Berger, S. Campbell, G. Fiore, L. Fletcher, E. Frazzoli, A. Huang, S. Karaman *et al.*, "A perception-driven autonomous urban vehicle," *Journal of Field Robotics*, vol. 25, no. 10, pp. 727–774, 2008.
- [23] S. Behere and M. Torngren, "A functional architecture for autonomous driving," in *2015 First International Workshop on Automotive Software Architecture (WASA)*, May 2015, pp. 3–10.
- [24] Y. Amit and P. Felzenszwalb, *Object Detection*. Boston, MA: Springer US, 2014, pp. 537–542. [Online]. Available: https://doi.org/10.1007/978-0-387-31439-6_660
- [25] S. Lazebnik, *Object Class Recognition (Categorization)*. Boston, MA: Springer US, 2014, pp. 533–536. [Online]. Available: https://doi.org/10.1007/978-0-387-31439-6_337
- [26] L. Matthies, *Obstacle Detection*. Boston, MA: Springer US, 2014, pp. 543–549. [Online]. Available: https://doi.org/10.1007/978-0-387-31439-6_52
- [27] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 1, no. 1, pp. 33–55, March 2016.
- [28] S. Furst and M. Bechter, "Autosar for connected and autonomous vehicles: The autosar adaptive platform," in *2016 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshop (DSN-W)*. IEEE, 2016, pp. 215–217.
- [29] S. Kato, E. Takeuchi, Y. Ishiguro, Y. Ninomiya, K. Takeda, and T. Hamada, "An open approach to autonomous vehicles," *IEEE Micro*, vol. 35, no. 6, pp. 60–68, Nov 2015.
- [30] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "Ros: an open-source robot operating system,"

- in *ICRA workshop on open source software*, vol. 3, no. 3.2. Kobe, Japan, 2009, p. 5.
- [31] "Android automotive," available from: <https://source.android.com/devices/automotive>, Accessed July 2019.
- [32] H. Araujo, G. Carvalho, M. Mohaqeqi, M. R. Mousavi, and A. Sampaio, "Sound conformance testing for cyber-physical systems: Theory and implementation," *Sci. Comput. Program.*, vol. 162, pp. 35–54, 2018.
- [33] A. Aerts, M. R. Mousavi, and M. A. Reniers, "Model-based testing of cyber-physical systems," in *Cyber-Physical Systems Foundations, Principles and Applications*. Elsevier, 2016, ch. 19, pp. 287–304.
- [34] H. Araujo, G. Carvalho, M. R. Mousavi, and A. Sampaio, "Multi-objective search for effective testing of cyber-physical systems," in *Proceedings of the 17th International Conference on Software Engineering and Formal Methods (SEFM 2019)*. Springer, 2019.
- [35] J. M. Rojas, M. Vivanti, A. Arcuri, and G. Fraser, "A detailed investigation of the effectiveness of whole test suite generation," *Empirical Software Engineering*, vol. 22, no. 2, pp. 852–893, 2017.
- [36] V. Bertram, S. Maoz, J. O. Ringert, B. Rumpe, and M. von Wenckstern, "Component and connector views in practice: An experience report," in *20th ACM/IEEE International Conference on Model Driven Engineering Languages and Systems, MODELS 2017, Austin, TX, USA, September 17-22, 2017*. IEEE Computer Society, 2017, pp. 167–177. [Online]. Available: <https://doi.org/10.1109/MODELS.2017.29>
- [37] K. J. Åström and T. Hägglund, *PID controllers: theory, design, and tuning*. Instrument society of America Research Triangle Park, NC, 1995, vol. 2.
- [38] A. Bar-Hillel, R. Lerner, D. Levi, and G. Raz, "Recent progress in road and lane detection: a survey," *Mach. Vis. Appl.*, vol. 25, no. 3, pp. 727–745, 2014. [Online]. Available: <https://doi.org/10.1007/s00138-011-0404-2>
- [39] S. Cave, K. Coughlan, and K. Dihal, "'scary robots': Examining public responses to AI," in *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society, AIES 2019, Honolulu, HI, USA, January 27-28, 2019*, V. Conitzer, G. K. Hadfield, and S. Vallor, Eds. ACM, 2019, pp. 331–337. [Online]. Available: <https://doi.org/10.1145/3306618.3314232>
- [40] J. Redmon, S. K. Divvala, R. B. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*. IEEE Computer Society, 2016, pp. 779–788. [Online]. Available: <https://doi.org/10.1109/CVPR.2016.91>
- [41] M. M. Waldrop, "Autonomous vehicles: No drivers required," *Nature*, vol. 518, no. 7537, p. 20–23, February 2015. [Online]. Available: <https://doi.org/10.1038/518020a>
- [42] B. Schoettle and M. Sivak, "Public opinion about self-driving vehicles in china, india, japan, the u.s., the u.k., and australia," University of Michigan, Ann Arbor, Transportation Research Institute, Tech. Rep., 2014. [Online]. Available: <https://deepblue.lib.umich.edu/handle/2027.42/109433>
- [43] Y. Kang, H. Yin, and C. Berger, "Test your self-driving algorithm: An overview of publicly available driving datasets and virtual testing environments," *IEEE Transactions on Intelligent Vehicles*, vol. 4, no. 2, pp. 171–185, June 2019. [Online]. Available: <https://ieeexplore.ieee.org/document/8667012>
- [44] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in *Proceedings of the 1st Annual Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, S. Levine, V. Vanhoucke, and K. Goldberg, Eds., vol. 78. PMLR, 13–15 Nov 2017, pp. 1–16. [Online]. Available: <http://proceedings.mlr.press/v78/dosovitskiy17a.html>
- [45] Carla, "CARLA GitHub project - Open-source simulator for autonomous driving research.," <https://github.com/carla-simulator/>, 2019, [Online; accessed 25-Jun-19].
- [46] —, "CARLA Scenario_Runner - Traffic scenario definition and execution engine," https://github.com/carla-simulator/scenario_runner, 2019, [Online; accessed 25-Jun-19].
- [47] E. Awad, S. Dsouza, R. Kim, J. Schulz, J. Henrich, A. Shariff, J.-F. Bonnefon, and I. Rahwan, "The moral machine experiment," *Nature*, vol. 563, no. 7729, p. 59, 2018.