



Ethical and Social Aspects of Self-Driving Cars

Tobias Holstein, Gordana Dodig Crnkovic and Patrizio Pelliccione

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

February 9, 2018

Ethical and Social Aspects of Self-Driving Cars

Tobias Holstein
Mälardalen University
Västerås, Sweden
tobias.holstein@mdh.se

Gordana Dodig-Crnkovic, Patrizio Pelliccione
Chalmers University of Technology | University of
Gothenburg
Gothenburg, Sweden
[gordana.dodig-crnkovic,patrizio]@chalmers.se

ABSTRACT

As an envisaged future of transportation, self-driving cars are being discussed from various perspectives, including social, economical, engineering, computer science, design, and ethics. On the one hand, self-driving cars present new engineering problems that are being gradually successfully solved. On the other hand, social and ethical problems are typically being presented in the form of an idealized unsolvable decision-making problem, the so-called trolley problem, which is grossly misleading. We argue that an applied engineering ethical approach for the development of new technology is what is needed; the approach should be applied, meaning that it should focus on the analysis of complex real-world engineering problems. Software plays a crucial role for the control of self-driving cars; therefore, software engineering solutions should seriously handle ethical and social considerations. In this paper we take a closer look at the regulative instruments, standards, design, and implementations of components, systems, and services and we present practical social and ethical challenges that have to be met, as well as novel expectations for software engineering.

CCS CONCEPTS

• **Computer systems organization** → **Embedded systems**; *Redundancy*; Robotics; • **Networks** → Network reliability;

KEYWORDS

Self-Driving Cars, Autonomous Cars, Trolley Problem, Decision Making, Ethics, Social Aspects, Software Engineering, Challenges

1 INTRODUCTION

Increasingly, prototypical self-driving vehicles are participating in public traffic [48] and are planned to be sold starting in 2020 [56, 61]. Public awareness and media coverage contribute to a manifold of discussions about self-driving vehicles. This is currently amplified through recent accidents with autonomous vehicles [24, 58].

Software is playing a key role in modern vehicles and in self-driving vehicles. Gigabytes of software run inside the Electronic Control Units (ECUs), which are small computers embedded in the vehicle. The number of ECUs has grown in the last 20 years from 20 to more than 100. Software in cars is growing by a factor of 10 every 5 to 7 years, and in some sense car manufacturers are becoming software companies [47]. These novelties ask for a change on how the software is engineered and produced and for a disruptive renovation of the electrical and software architecture of the car, as testified by the effort of Volvo Cars [47].

Moreover, self-driving vehicles will be connected with other vehicles, with the manufacturer cloud, e.g., for software upgrades,

with Intelligent Transport Systems (ITS), Smart Cities, and Internet of Things (IoT). Self-driving vehicles will combine data from inside vehicle with external data coming from the environment (other vehicles, the road, signs, and the cloud). In such a scenario, different applications will be possible: smart traffic control, better platooning coordination, and enhanced safety in general. However, the basic assumption is that future self-driving connected cars must be socially sustainable. A typical discussion about ethical aspects of self-driving cars starts with ethical thought experiment, so called “trolley problem” described in [29] and [66], that has been discussed in number of articles in IEEE [7, 9, 33], ACM [30, 40, 43], Scientific American [16, 37, 41], Science [11, 36], other high-profile journals [14, 32, 34], conference workshops [8, 50] and other sources [2, 6, 44, 54]. Here is the general scenario being discussed:

A self-driving vehicle drives on a street with a high speed. In front of the vehicle a group of people suddenly blocks the street. The vehicle is too fast to stop before it reaches the group. If the vehicle does not react immediately, the whole group will be killed. The car could however evade the group by entering the pedestrian way and consequently killing a previously not involved pedestrian. The following alternations of the problem exist: (A) Replacing the pedestrian with a concrete wall, which in consequence will kill the passenger of the self-driving car; (B) Varying the personas of people in the group, the single pedestrian or the passenger. The use of personas allows including an emotional perspective [10], e.g., stating that the single pedestrian is a child, a relative, a very old or a very sick human, or a brutal dictator, who killed thousands of people.

Even though the scenarios are similar, the responses of humans, when asked how they would decide, differ [11]. The problem is that the question asked has limited number of possible answers, which are all ethically questionable and perceived as bad or wrong. Therefore, a typical approach to this problem is to analyze the scenarios by following ethical theories, such as utilitarianism, other forms of consequentialism or deontological ethics [42]. For example, utilitarianism would aim to minimize casualties, even if it means to kill the passenger, by following the principle: the moral action is the one that maximizes utility (or in this case minimizes the damage). Depending on the ethics framework, different arguments can be used to justify the decision.

Applying ethical doctrines to analyze a given dilemma and possible answers can presently only be done by humans. How would self-driving cars solve such dilemmas? There are a number of publications that suggest to implement moral principles into algorithms of self-driving cars [17, 18, 33]. We find that this does not solve the problem, but it reassures that the solution is calculated based on a given set of rules or other mechanisms, moving the problem to engineering, where it is implemented.

It is worth to notice that the engineering problem is substantially different from the hypothetical ethical dilemma. While an ethical dilemma is an idealized constructed state that has no good solution, an engineering problem is always by construction such that it can differentiate between better and worse solutions. A decision making process that has to be implemented in a self-driving car can be summarized as follows. It starts with an awareness of the environment: Detecting obstacles, such as a group of humans, animals or buildings, and also the current context/situation of the car using external systems (GPS, maps, street signs, etc.) or locally available information (speed, direction, etc.). Various sensors have to be used to collect all required information. Gaining detailed information about obstacles would be a necessary step before a decision can be made that maximizes utility and/or minimizes damage. A computer program calculates solutions and chooses the solution with the optimal outcome. The self-driving car executes the calculated action and the process repeats itself.

The process itself can be used to identify concrete ethical challenges within the decision making by considering the current state of the art of technology and its development. In a concrete car both the parts of this complex system and the way in which it is created have a critical impact on the decision-making. This includes for instance the quality of sensors, code, and testing. We also see ethical challenges in design decisions, such as whether a certain technology is used because of its lower price, even though the quality of information for the decision making would be substantially increased if more expensive technology (such as sensors) would be used.

Since building and engineering of self-driving vehicle involve various stakeholders, such as software/hardware engineers, sales people, management, etc., we can also pose the following questions: does the actual self-driving car have a moral on its own or is it the moral of its creators? And who is to blame for the decision making of a self-driving car? In [22] the argument is put forward that the systemic view must be used in case of socio-technological systems. Thus the problems in the system can originate or be a combination coming from inadequate solutions in various steps from requirements specification to implementation, testing, deployment maintenance, safety regulation and other normative support etc.

Besides the self-driving vehicle itself, it is also important to address yet another complex system: self-driving vehicles participating in public traffic among cars with human drivers. Therefore, it is important to investigate how self-driving vehicles are actually built, how ethical challenges are addressed in their design, production, and use and how certain decisions are justified. Discussing this before self-driving vehicles are officially introduced into the market, allows taking part in the setting and definition of ethical ground rules. McBride states that "Issues concerning safety, ethical decision making and the setting of boundaries cannot be addressed without transparency" [43]. We think that transparency is only one factor, as it is necessary to start further investigations and discussions.

In order to give a more detailed perspective on the complex decision making process, we propose to create a conceptual ethical model that connects the different components, systems and stakeholders. It shows inter-dependencies and allows pinpointing ethical challenges that will be presented in the concluding recommendations.

Focusing on important ethical challenges that should currently be addressed and solved is an important step before ethical aspects of self-driving cars can actually be meaningfully discussed from the point of view of societal and individual stakeholders as well as designers and producers. It is important to focus not on abstract thought experiments but on concrete conditions that influence the behavior of self-driving cars and their safety as well as our expectations from them.

The paper is structured as follows. A short introduction to self-driving cars and their current state of the art is provided in Section 2, with the emphasis on the description of the decision making principles given in Section 2.1 and the role of software in Section 2.2. Ethical and social challenges are addressed in Section 3 regarding technical aspects, and Section 4 addressing social aspects. Section 5 describes the current state of norms and standards, while conclusions and final remarks are presented together with recommendations in Section 6.

2 SELF-DRIVING CARS BASICS

The term "autonomous" could be ambiguous to some readers. It can be used to describe certain autonomous features or functions, such as advanced driver assistance systems, that for example assist the driver in keeping the lane or adjust to the speed of vehicles ahead. Those systems are designed to assist, but the driver is always responsible and has to intervene if critical situations occur.

We use the term "self-driving" cars to avoid wrong interpretations of the terms "fully autonomous" or "driverless". Self-driving cars refer to cars that may operate self-driving without human help or even without a presence of human being. This means that the unoccupied car can drive from place A to B to pick up someone. This is the highest level of autonomy for cars and corresponds to the last level of five as defined by the Society of Automotive Engineers [51] and United States National Highway Traffic Safety Administration (NHTSA), who, since September 2016, adopted SAE's classification with level 1 (no automation), level 2 (drive assistance), level 3 (partial automation), level 4 (conditional automation), and level 5 (full automation) [45, p.9].

A concrete example is the self-driving Waymo car [65], former known as the Google car [35], a fully autonomous and self-driving vehicle.

2.1 Decision Making Process in Self-Driving Cars

Developing self-driving cars that act without a driver means to replace a human, who today is performing the complex tasks of driving, with a computer system executing the same tasks. Figure 1 shows both variants and allows a comparison.

There is an important difference in the feedback loop. While humans continuously learn, for example from their mistakes or misbehaviour, automotive software might be confined to slow updates. Approaches with self-adaptive software, such as machine learning approaches, which learns and reacts immediately, aim to overcome this constraint. Extraordinary road signs for example, which are new to the self-driving car's software, present a risk as they can pass unnoticed/uninterpreted, while they could be understood by a human through context/interpretation. Also

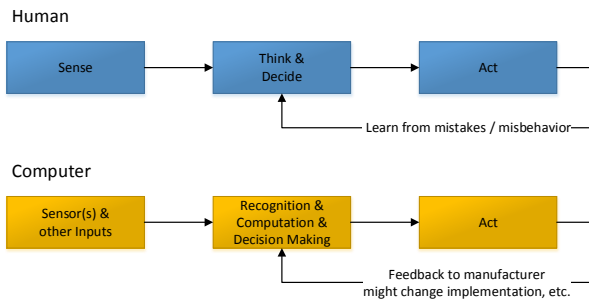


Figure 1: Comparison of human and computers sense, think and act process (cf. [31]) which we extended by adding a feedback loop

unexpected and dangerous situations, like an attack or threat near or even against the vehicle might not be correctly interpreted by a self-driving car compared to a human.

Depending on the technology and the amount of sensors, the type and quality of information that is gathered differ. However, this extremely complex process might be difficult to imagine and in order to give an idea of what self-driving cars “see” we refer to the visualization depicted in Figure 2. It shows a rendered point cloud, based on the data gathered by a laser radar (LIDAR) mounted on the top of the vehicle.




Figure 2: Point cloud image of a vehicle approaching an intersection illustrates the complexity [25]

2.2 Complexity of Decision Making and the Role of Software

The amount of sensors used to detect objects around the vehicle and its surrounding environment differs among car manufacturers. Figure 3 shows an abstraction made to discuss the types of information used and how they relate to each other.

Most of the functionality in the automotive domain is based on software [12]. Software is written by software engineers and at least for important components extensively tested to ensure their

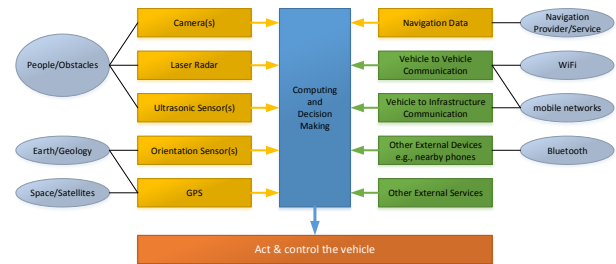


Figure 3: Abstract representation of decision making in autonomous vehicles composed from various sources (cf. [25, 59, 63, 64])

correct functioning. In self-driving cars software relies on different disciplines, such as computer vision, machine learning, and parallel computing, but also on various external services. It is a complex process to calculate a decision, and it is also difficult to test those against all possible real world scenarios [63].

One of the problems is that all calculations are based on an abstraction of the real world. This abstraction is an approximate representation of a real world situation and thus the decision making will create decisions for an imperfect world. This is a twofold problem, because the more information is available the better the decisions might be, but at the same time more interpretation and filtering might have to be used to get the data that actually is useful for the decision making.

Engineers have to decide what kind of data to use, how reliable or trustworthy the data are and how to balance the different sources of information in their algorithms. Also different sensors have their specific limitations and to overcome those, a combination of multiple sensors might be used. The overall problem is usually referred to as sensor fusion. This problem is exacerbated in the case of connected vehicles since data will come not only from the sensors of the car, but also from other vehicles, street infrastructure, etc. In this case other factors should be taken into account since it is not possible to have a perfect knowledge about the devices that are used to sense information and about their status.

Imagine heavy weather conditions, the navigation reports a street ahead, the radar is reporting a clear street, but the visual camera reports an obstacle straight ahead. How will this “equation” be solved and what will be the result? The wrong decision might lead to an accident, when important information of some sensors is disregarded and other sensors do not detect the obstacle or hazard in front of the vehicle [58]. Car manufacturers are constantly improving and testing the recognition capabilities of their systems [59]. It is a multi-factor optimization task, which aims to find an optimal solution under consideration of costs, quality, and potential risk factors.

Some manufacturers are thinking to count miles covered without any accident, however this might be infeasible since a vehicle should cover around 11 billion of miles to demonstrate with 95% of confidence and 80% power that autonomous vehicle failure rate is lower than the human driver failure rate [39]. Moreover, this calculation holds if the software within the car does not change

over time. Nowadays, manufacturers are increasingly interested in continuous integration and deployment techniques that promise to update the software even after the vehicle has been sold and is on the street, like a common smart-phone. However, changing even a single line of code might require to start counting from 0 the number of covered miles.

3 ETHICAL ASPECTS OF THE TECHNICAL CHALLENGES IN SELF-DRIVING CARS

In the following, we will discuss ethical deliberations surrounding the autonomous vehicle, including involved stakeholders, technologies, social environments, and costs vs. quality. The multifaceted and complex nature of reality emphasizes again the importance to look broader instead of focusing on single ethical dilemma like the trolley problem.

3.1 Safety

Safety is the most fundamental requirement of autonomous cars. The central question is then: how should a self-driving car be tested? What guidelines should be fulfilled to ensure that it is safe to use? There are several standards, such as the ISO 26262, that specify the safety standard for road vehicles. For self-driving cars standards are under development, based on experiences being made. Google Car tests show one million kilometres without any accident, is this a measurement to certify its software? As we discussed above, this should not be a reasonable assurance for safety. Should a self-driving car make a driver licence as suggested in [43]? How would that work?

The source code of autonomous cars are typically commercial and not publicly available. One possibility to assure code correctness via independent control. Should there be an independent organization to check those? But could it actually be checked? Who else than the developers at a car manufacturer or supplier will understand such a complex system?

An alternative route seems to be preferred by legislators - instead of control of the software which is in the domain of the producers, legislation focus on behaviour that is being tested, based on the "Proven in Use" Argument.

Testing of present-day cars should demonstrate the compliance of their behaviour with legislative norms [20]. Disengagements, accidents and reaction times based on data released in 2016 from the California trials are discussed in [21].

In the case of the software of the car will evolve even when the vehicle is already on the street, testing should account for this new challenge.

When it comes to hardware and hardware-software systems, there have been discussions about the prices of laser radars compared to cameras or ultra-sonic sensors. Laser radars are very expensive, but deliver high quality data in diverse weather conditions. Ultra-sonic sensors or cameras are less accurate and sensitive under weather conditions like rain. Should a car manufacturer choose a cheap over an expensive sensor, even if this raises the likelihood of errors/faults/accidents? In advanced driving assistance systems, the driver would take over, if a critical situation could not be handled by the system. What happens in self-driving cars? Will the car just stop and wait until the rain is over? Will passengers be able

and allowed to intervene? Under which conditions? Would it be required to have a driving licence for a self-driving car? Or would the police have a possibility to intervene, and in what way, when a car behaves inadequately or even dangerously? Also would the police even have the possibility to stop a self-driving car that is behaving correctly, with the sole purpose of checking the passengers?

The economic aspects might be seen as the highest priority. Using cheap equipment might lead to wrong decision-making and in a self-driving car, it would be impossible to interfere with the decisions made. Assuming that wrong decision may lead to a loss of human lives or property, having chosen a cheap component could therefore be ethically unacceptable.

Learning from experience is the most important basis for improvement of safety in self-driving cars. This is for instance envisioned by the CEO of Tesla, Elon Musk, in the Tesla's second 10 year master plan "part deux", where the third element of the four major elements is: develop a vehicle self-driving capability that is 10x safer than manual via massive fleet learning¹.

3.2 Security

For autonomous cars, security is of paramount importance, and software security is a fundamental requirement. As an indication of the development we mention that in August 2017 UK's Department for Transport, published the document "Key principles of vehicle cyber security for connected and automated vehicles" [19]. It is built on the following eight basic principles:

- (1) Organizational security is owned, governed, and promoted at board level;
- (2) Security risks are assessed and managed appropriately and proportionately, including those specific to the supply chain;
- (3) Organizations need product aftercare and incident response to ensure systems are secure over their lifetime;
- (4) All organizations, including sub-contractors, suppliers, and potential 3rd parties, work together to enhance the security of the system;
- (5) Systems are designed using a defence-in-depth approach;
- (6) The security of the software is managed throughout its lifetime;
- (7) The storage and transmission of data is secure and can be controlled;
- (8) The system is designed to be resilient to attacks and respond appropriately when its defences or sensors fail.

Similar documents are mentioned, such as Microsoft Security Development Lifecycle (SDL), SAFE Code best practices, OWASP Comprehensive, lightweight application security process (CLASP), and HMG Security policy framework [19].

There have been a number of attacks at car systems and sensors (e.g., LIDAR and GPS) that were used to manipulate the cars behaviour. Attacks might be inevitable, but should there be a minimum security threshold to allow a self-driving car to be used? This leads to another question: How secure must the systems and the connections be?

In aircrafts "black boxes" are used to determine what happened after a crash. Should this be also a part of a self-driving car?

¹<https://www.tesla.com/blog/master-plan-part-deux>

What about security issues and software updates? Should a self-driving car be allowed to drive, when it does not have the latest software version running? What about bugs in the new software?

Should the vehicle be connected or should the vehicle be completely disconnected? On one side, the most secure system is the one that is disconnected from the network. On the side, it would be unethical to do not deploy immediately new software or a new version of the software on the car if there is evidence that the new update will fix important problems that might endanger human lives. In order to enable the massive fleet learning and to do the software update, connectivity is needed. Moreover, connected vehicles might receive information from other systems that will enhance the understanding of the reality thus opening new and promising safety scenarios. Imagine, for instance, a pedestrian on a side of a building, totally invisible to the instrumentations of the car, that is approaching a cross and that will most probably have an impact with the vehicle².

3.3 Privacy

The more information taken into consideration for the decision making, the more it might interfere with data and privacy protection. For example, a sensor that detects obstacles, such as human beings in front of the car is based on visual information. Even the use of a single sensor could invade privacy, if the data is recorded/reported and/or distributed without the consent of the involved people. The general question is: How much data is the car supposed to collect for the decision making? Who will access those data? When will these data be destroyed?

What about using active signals by devices people carry around to detect moving obstacles in front or near the car? What about people who do not carry such devices? Would they more likely be hit by the self-driving car, because they were not “present enough” in the data?

And how much data is actually used for evaluation? Is it anonymous? Does it contain more data than “just” the position of a human? Can it be connected to other types of data like the phone number, the bank account, the credit cards, personal details, or health data?

Those and similar questions are met by legislation such as Regulation (EU) 2016/679 of the European Parliament and of the Council (the General Data Protection Regulation) setting a legal framework to protect personal data [28], and discussed in [62].

3.4 Trust

Trust is an issue that appears in various forms in autonomous cars e.g. in production (when assembled, trust is the requirement for both hardware and software components) as well as in use of the car. A human might define where the car has to go, but the self-driving car will make the decisions how to get there, following the given rules and laws. However the self-driving car might already distribute data like the target location to a number of external services, such as traffic information or navigation data, which are used in the calculation of the route. But how trustworthy are those data sources (e.g., GPS, map data, external devices, other vehicles)?

In regard of the used sensors and hardware, the question is, how trustworthy are those? How can trust be implemented, when so many different systems are involved?

3.5 Transparency

The transparency is of central importance for many of the previously introduced challenges. Without transparency none of them could be analyzed, because the important information would be missing. “Transparency is a prerequisite for ethical engagement in the development of autonomous cars. There can be nothing hidden, no cover-ups, no withholding of information” [43]. It is a multi-disciplinary challenge to ensure transparency, while respecting e.g., copyright, corporate secrets, security concerns and many other related topics.

How much should be disclosed, and disclosed to whom? The car development ecosystem includes many other companies acting as suppliers that produce both software and hardware components. Should the entire ecosystem be transparent? Also to whom should it be transparent? How to manage the intellectual property rights? Some initial formulations are already present in the current policy documents and initial legislative that will be discussed later on.

Declaration of Amsterdam [4] lists among the objectives “to adopt a ‘learning by experience’ approach, including, where possible, cross-border cooperation, sharing and expanding knowledge on connected and automated driving and to develop practical guidelines to ensure interoperability of systems and services”.

Goodman and Flaxman in [34] present EU regulations on algorithmic decision-making and a “right to explanation” that is the right for user to ask for an explanation of an algorithmic (machine) decision that was made about them. The Department of Motor Vehicles provides the law requirements [20] “Under the testing regulations, manufacturers are required to provide DMV with a Report of Traffic Accident Involving an Autonomous Vehicle (form OL 316) within 10 business days of the incident”. The list of all incidents can be found in [5].

3.6 Reliability

One of the basic questions is: How reliable is the cell network? What if there is no mobile network available? What if sensor(s) fail? Should there be redundancy for everything? Is there a threshold that determines when the car is reliable, e.g., when two out of four sensors fail?

In connected vehicles there are different levels that should be considered for reliability purposes. First the diagnostic of the vehicle that might be subject to failures. Then, the vehicle sensors that enable the vehicle to sense the surrounding environment of the vehicle. Finally, the data coming from external entities, like other vehicles and road infrastructures. Reliability approaches should consider all these levels.

3.7 Responsibility and Accountability

In the case of autonomous cars responsibility will obviously be redefined. The question is how will responsibility be defined in case of incidents and accidents. Regarding ethical aspects of responsibility, a lot can be learned from the existing Roboethics and the debate about responsibility in autonomous robots, e.g., [23]. This is still

²<https://www.youtube.com/watch?v=w0rPQpjZhXg>

an open problem even though important steps forward are being made by legislators, such as mentioned “Key principles of vehicle cyber security for connected and automated vehicles” [19].

3.8 Quality Assurance Process

Detailed Quality assurance programs covering all relevant steps must be developed in order to ensure high quality components. The question is also how is the decision making to be implemented. How to ensure overall quality of the product? What about the lifetime of components? How will maintenance be organized and quality assured? When car manufacturers follow a non-transparent process of software engineering, how could anyone make sure that the car follows a certain ethical guideline? Whose responsibility will it be that car software follows ethical principles?

One part of the Quality Assurance (QA) process regards assembling of components. All parts of a vehicle are designed, fabricated and then assembled to the overall car. A standard non-autonomous premium vehicle today has more than 100 electronic control units that are responsible for the control of e.g., the engine, the wipers, the navigation system or the dashboard [47]. We assume that for self-driving the number will be increased. Parts are usually built by not one, but a multitude of suppliers. This requires an extensive design and development process, which again involves various disciplines, such as requirements engineering, software engineering or project management. It is an overall extensive process which holds ethical questions and challenges. Thus, it is necessary to include ethical deliberations in the overall process but also in all sub processes. As it is stated in [52]: “value-based ethical aspects, which today are implicit, should be made visible in the course of design and development of technical systems, and thus a subject of scrutiny”.

Including ethics-aware decision making in all processes will help to make ethically justified decisions. This is important when it comes to questions: Which parts/components are used for a vehicle? Can we choose a cheaper component with less accuracy instead? Is the reliability of this part high enough for a self-driving car?

4 ETHICAL ASPECTS OF SOCIAL CHALLENGES

Self-driving cars will influence job markets, as for example for taxi drivers, chauffeurs or truck drivers. The perception of cars will change and cars might be seen as a service that is used for transportation. The idea of having a vehicle that is specialized for the specific use, e.g. off road, city road, long travels might become attractive. This might impact the business model of car manufacturer and their market.

This in itself poses ethical problems: what strategy should be applied for people loosing jobs because of the transition to self-driving cars? It is expected that the accident frequency will decrease rapidly, so car insurances may become less important. This may affect insurance companies in terms of jobs and the business. There is a historical parallel with process of industrialization and automatization, and there are experiences that may help anticipate and better plan for the process of transition.

4.1 Stakeholders - General Public Interests

Humans concerns must be taken into account in the decision making of self-driving cars. Should there be an emergency button to allow the human to interfere with the decision making of the self-driving car? Putting the human back in the loop of decision making also inflicts with the autonomy of the system. Is it then truly self-driving? Giving passengers a choice to interfere with the decisions of the self-driving car puts the passenger back in charge, who would be responsible to press or *not* press the button in all situations. In the context of the self-driving car the computer decision might be better, but it might also be worse than human, because of possible errors [26].

Another perspective on the human interest is the granularity of the settings or configurations given to the user. How for example will a route be planned?

In an extreme scenario self-driving cars might even avoid or reject to drive to a certain region or position. Would that be an interference with the freedom of choice, will passengers be informed about the reasons for such decisions? It is important to determine how much control the human should have, that will be taken into account when making design choices for a self-driving car.

4.2 Possible New Selling Points

The automotive industry has a highly competitive market. What will be the difference between buying a self-driving car of brand A compared to brand B?

Taking away the primary and secondary tasks of driving, i.e. the driving controls, safety features, assistance etc., leaves only entertainment and comfort functions in control of the passengers, the former drivers. The interior becomes more important and factors that cannot be controlled less in the focus of the user. What will be the main buying criteria? Will it be the interior/exterior, speed (as often with traditional cars) or other new services? Will it be possible for the users of the car to choose the priorities in its decision-making? The latter is difficult, since decision choices [57] supporting the survival of passengers over other traffic participants by car manufacturers would have legal implications in most countries [15]. The question is also who will own the cars. Will they become a service for individual users, and owned by companies? Buying criteria will be different depending on the ownership.

Surveys based on hypothetical trolley problem scenarios show that people feel less attracted to buy a car that would sacrifice the passengers in order to save more human lives [11]. Would that decision be left to car manufacturers? Existing policy documents do not seem to leave possibilities open for anti-social cars to be developed [3, 13, 27, 46, 49].

Table 1 and 2 present the summary of ethical and social challenges with recommendations (action points) grouped by requirement to be taken into account in policy-making as well as software design and development for self-driving cars.

5 LEGISLATION, STANDARDS, AND GUIDELINES

Present-day regulatory instruments for transportation systems are based on the assumption of human-driven vehicles. As the development and introduction of increasingly automated and connected

Table 1: Summary of the technical challenges and recommendations grouped by requirement

Requirements	Challenges	Recommendations
Safety	<ul style="list-style-type: none"> • Trade-off between safety and other aspects like economic aspects • Boundaries of autonomy of self-driving cars and human (passenger) interactions • Police control and possibility of intervention with self-driving cars • Systemic solutions to guarantee safety in organizations (regulations, authorities, safety culture) 	<ul style="list-style-type: none"> • Assure means to guarantee that safety is not sacrificed because of other aspects (this is not so different from today. This can happen also for the braking system of non-autonomous cars) • It should be specified how a self-driving car will behave in cases that the car is not able to deal autonomously. In the future passengers will be able to drive a car • There is the need of clarifying the relationship of police and self-driving cars • New techniques and standards are necessary to guarantee safety in self-driving cars that will continuously update their software
Security	<ul style="list-style-type: none"> • Identification and declaration of minimal necessary security requirements, that work as a threshold for deployment of self-driving cars • Security in systems and connections • Deployment of software updates • Storing and using received and generated data in a secure way 	<ul style="list-style-type: none"> • Provide technical solutions that will guarantee minimum security under all foreseeable circumstances • Anticipate and prevent worst case scenarios regarding security breaches • Continuous learning process must be in place to provide active security • Assure accessibility of the data even in the case of accidents so that it could be analyzed and lessons learned.
Privacy	<ul style="list-style-type: none"> • Trade-off between privacy and data collection/recording • Use of technology that detects humans near/around the car, even if those humans do not carry any kind of electronics 	<ul style="list-style-type: none"> • Following/applying legal frameworks to protect personal data, such as Regulation (EU) 2016/679 of the European Parliament [28] (discussed in [62]) • Justify the use of collected data through a transparent decision making process
Trust	<ul style="list-style-type: none"> • How trust between both software and hardware components of complex systems can be implemented is not clear 	<ul style="list-style-type: none"> • Further research on how to implement trust across multiple systems • Provide trusted connections between components as well as external services
Transparency	<ul style="list-style-type: none"> • Information disclosure, what and to whom • Transparency of the ecosystem • Management of intellectual property rights 	<ul style="list-style-type: none"> • Ensure transparency and provide insight into decision making • Actively share knowledge, gained by "learning through experience", to ensure the interoperability of systems and services. • Transparency is a prerequisite for the herein introduced challenges, since it is the key to potentially undisclosed background information
Reliability	<ul style="list-style-type: none"> • Reliability of required networks and solution for the case when the network is unavailable • Reliability of sensors, and need for redundancy • Way to determine when a car is not reliable anymore 	<ul style="list-style-type: none"> • Define different levels for reliability (diagnostics, vehicle input sensors, external services) • Determine reliability for components and the overall car
Responsibility and Accountability	<ul style="list-style-type: none"> • Responsibility and accountability in case of incidents and accidents • Responsibility that car software follows ethical principles 	<ul style="list-style-type: none"> • Consider research and learn from robotics, i.e. Roboethics [23] • Support development of solutions, e.g. by contributing to existing approaches [19]
Quality Assurance (QA) Process	<ul style="list-style-type: none"> • Quality of components • Quality of decision making • Lifetime and maintenance • Trade-offs between non-transparent processes and external QA control of adherence of ethical principles/guidelines 	<ul style="list-style-type: none"> • Ethical deliberations must be included in the process of design and development of self-driving cars • Ethics-aware decision making must be part of the process to ensure ethically justified decisions

Table 2: Summary of social challenges and recommendations grouped by requirement

Requirements	Challenges	Recommendations
Social challenges of disruptive technology	<ul style="list-style-type: none"> • Handling job losses (e.g., taxi/truck drivers, traditional mechanics, insurance agents, etc.) • Change of related markets and business models (e.g., car insurances, car manufacturers, etc.) 	<ul style="list-style-type: none"> • Prepare strategic solutions for people loosing jobs • Take advice/learn from historic parallels to industrialization and automatization
Stakeholders - general public interests	<ul style="list-style-type: none"> • Human concerns in the decision making of self-driving cars, e.g. possibility of interference with the decision making • Freedom of choice hindered by the system (e.g. it may not allow to drive into a certain area) • Values and priorities: Ensure that general public values will be embodied in the technology, with interests of minorities taken into account 	<ul style="list-style-type: none"> • Determine and communicate the amount of control a human has in context of the self-driving car • The freedom of choice determined by regulations • Active involvement of stakeholders in the process of design and requirements specification
Selling points	<ul style="list-style-type: none"> • Self-driving cars will have different buying criteria, depending on who will own the cars, big companies, social institutions such as municipalities or individual users, as they all have different preferences. Among those preferences environmental and sustainability criteria can be expected to play central role • Existing policy documents don't seem to leave possibilities open for anti-social self-driving cars to be developed [13, 27, 46, 49]. 	<ul style="list-style-type: none"> • Priorities and choices for the self-driving cars will result from the dialog between producers and future users • Ensure that existing and future policies and standards prevent the possibility of developing "anti-social" self-driving cars.
Legislation, norms, policies and standards	<ul style="list-style-type: none"> • Keeping legislation up-to-date with current level of automated driving, and emergence of self-driving cars • Creating and defining global legislation frameworks for the implementation of interoperable and development of increasingly automated vehicles • Defining the guidelines that will be adopted by society for building self-driving cars. • Including ethical guidelines in design and development processes 	<ul style="list-style-type: none"> • Car producers supporting and collaborating with legislators in their task to keep up-to-date with the current level of automated driving • Legislative support and contribution to global frameworks to ensure a smooth enrollment of the emerging technology • Include ethics in the overall process of design, development and implementation of self-driving cars. Ensure Ethics training for involved engineers [52, 55] • Establish and maintain a functioning socio-technological system in addition to functional safety standards.

cars proceed, from level 1 towards level 5 of automation, legislation needs constant updates [3, 27, 46, 49]. It has been recognized that present state regulatory instruments for human-controlled vehicles will not be adequate for self-driving cars: "existing NHTSA authority is likely insufficient to meet the needs of the time and reap the full safety benefits of automation technology. Through these processes, NHTSA will determine whether its authorities need to be updated to recognize the challenges autonomous vehicles pose" [46].

On 14 April 2016 EU member states endorsed the Declaration of Amsterdam [4] that addresses legislation frameworks, use of data, liability, exchange of knowledge and cross-border testing for the emerging technology. It prepares a European framework for the implementation of interoperable connected and automated vehicles by 2019 [27]. It also considers roles of stakeholders:

Agreement by all stakeholders on the desired deployment of the new technologies will provide developers with the certainty they need for investments. For an effective communication between the technological and political spheres, categorization and terminology are

being developed which define different levels of vehicle automation. [49]

The question is thus how to ensure that self-driving cars will be built upon ethical guidelines, which will be adopted by society. The strategy is to rely on rigorously monitoring the behaviour of cars, while the details of implementation are within the responsibility of producers. That means among others that design and implementation of software should follow ethical guidelines. An example of ethical guidelines trying to think one step further is described in Sarah Spiekermann's book *Ethical IT innovation* [55].

The approach based on "learning by experience" and "Proven in use" argument [1, 4, 53] presupposes a functioning socio-technological assurance system that has strong coupling among legislation, guidelines, standards and use, and promptly adapts to lessons learned. Ethical analysis in [22, 38, 60] addresses this problem of establishing and maintaining a functioning learning socio-technological system, while [38] discusses why functional safety standards are not enough.

6 CONCLUSIONS AND FINAL REMARKS

Self-driving vehicles have been recognized as the future of transportation systems and will be successively introduced into the transport systems globally [3, 46, 49]. It is now the right time to start an investigation into the manifold of ethical challenges surrounding self-driving and connected vehicles [27]. As this new technology is being tested and gradually allowed on the roads under controlled conditions, the focus should be on the practical technological solutions and their social consequences, rather than on idealized unsolvable problems such as much discussed trolley problem. Conclusions reached from idealized problem discussions would be that it has no general solution under all circumstances. We can compare this situation with the development and introduction of first cars. If the developers of traditional driver-controlled cars asked about general responsibility of a human driver for traffic accidents before allowing them to enter traffic, they would never be accepted, as safety in general and under all circumstances cannot be guaranteed and indeed human factor is the major safety concern. This does not mean that we should not take care of the basic requirements like security, safety, privacy, trust etc. and social challenges in general including legislation and stakeholders interests. On the contrary, those real-world techno-social problems must be taken seriously.

Focusing on unsolvable idealized ethical dilemmas such as the trolley problem obfuscates true ethical challenges, starting with characteristics of the whole techno-social system supporting new technology, with the emphasis on maximizing learning, on machine-, individual-, and social-level [13, 22]. The decision-making process and its implementation, which is central for the behaviour of a car, might internally use unreliable or insecure technology. Emerging technology of self-driving cars should follow ethical guidelines that stakeholders agree upon and should not be an autonomous black box with unknown performance. This poses new expectations, which affect software engineering that is involved in all its stages - from its regulatory infrastructure, to the requirements engineering, development, implementation, testing and verification [2, 7, 13, 36, 44]. As software is integral part of a complex software-hardware-human-society system, we presented different types of issues that we anticipate will affect software engineering in the near future.

It is also the right time to discuss the border between what is technically possible in relation to what is ethically justifiable. Even if this might limit the possibilities, it will set the necessary ground for further developments. The discussion should cover different dimensions, namely business, technical, process, and organization. First of all, there is the need to open a serious trade off analysis between business needs and ethics. As discussed above we should certainly avoid to compromise safety because of business priorities, e.g. equipping the car with cheaper but unreliable sensors. For what concerns technical aspects, it is of key importance to include ethical thinking and reasoning into the design and development process of autonomous and self-driving vehicles. Ethical aspects should be considered in every phase of a software development process, from requirements, till testing, maintenance, and evolution. Architectural and design decisions should be taken through a process that includes ethics as first-class actor and by involving stakeholders that are relevant to this concern. These architectural and design decisions should then be embedded into the code that

will run the self-driving vehicles and ensure its ethical aspects are taken care of. It is also necessary to enforce the transparency on those processes, so that independent evaluations become possible. Proper development processes, supported by suitable organization structure should promote and enable a serious discussion of ethics, and should emphasize the human interests, to make sure that the freedom of choice does not disappear in the new era of fully autonomous and self-driving vehicles.

REFERENCES

- [1] What is the iso 26262 functional safety standard? Technical report, National Instruments, 2014.
- [2] Moral Machine. <http://moralmachine.mit.edu>, 2016.
- [3] Ethics commission on automated driving presents report: First guidelines in the world for self-driving computers. Technical report, Federal Ministry of Transport and Digital Infrastructure, 2017.
- [4] On our way towards connected and automated driving in Europe. Technical report, Government of the Netherlands, 2017.
- [5] Report of traffic accident involving an autonomous vehicle (ol 316). https://www.dmv.ca.gov/portal/dmv/detail/vr/autonomous/autonomousveh_ol316+, 2017.
- [6] J. Achenbach. Driverless cars are colliding with the creepy Trolley Problem. <https://www.washingtonpost.com/news/innovations/wp/2015/12/29/will-self-driving-cars-ever-solve-the-famous-and-creepy-trolley-problem/>, December 2015.
- [7] E. Ackerman. People Want Driverless Cars with Utilitarian Ethics, Unless They're a Passenger. <https://spectrum.ieee.org/cars-that-think/transportation/self-driving/people-want-driverless-cars-with-utilitarian-ethics-unless-theyre-a-passenger>, June 2016.
- [8] H. S. Alavi, F. Bahrami, H. Verma, and D. Lalanne. Is driverless car another weiserian mistake? In *Proceedings of the 2017 ACM Conference Companion Publication on Designing Interactive Systems*, DIS '17 Companion, pages 249–253, New York, NY, USA, 2017. ACM.
- [9] S. Applin. Autonomous vehicle ethics: Stock or custom? *IEEE Consumer Electronics Magazine*, 6(3):108–110, July 2017.
- [10] A. Bleske-Rechek, L. Nelson, J. P. Baker, M. Remiker, and S. J. Brandt. Evolution and the trolley problem: People save five over one unless the one is young, genetically related, or a romantic partner. 4:115–127, 01 2010.
- [11] J.-F. Bonnefon, A. Shariff, and I. Rahwan. The social dilemma of autonomous vehicles. *Science*, 352(6293):1573–1576, 2016.
- [12] M. Broy, I. H. Kruger, A. Pretschner, and C. Salzmann. Engineering Automotive Software. *Proceedings of the IEEE*, 95(2):356–373, feb 2007.
- [13] V. Charisi, L. A. Dennis, M. Fisher, R. Lieck, A. Matthias, M. Slavkovik, J. Sombetzki, A. F. T. Winfield, and R. Yampolskiy. Towards moral autonomous systems. *CoRR*, abs/1703.04741, 2017.
- [14] I. Coca-Vila. Self-driving cars in dilemmatic situations: An approach based on the theory of justification in criminal law. *Criminal Law and Philosophy*, Jan 2017.
- [15] Daimler. Daimler clarifies: Neither programmers nor automated systems are entitled to weigh the value of human lives - Daimler Global Media Site, 2016.
- [16] K. Deamer. What the First Driverless Car Fatality Means for Self-Driving Tech. <https://www.scientificamerican.com/article/what-the-first-driverless-car-fatality-means-for-self-driving-tech/>, 2016.
- [17] L. Dennis, M. Fisher, M. Slavkovik, and M. Webster. *Ethical Choice in Unforeseen Circumstances*, pages 433–445. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.
- [18] L. Dennis, M. Fisher, M. Slavkovik, and M. Webster. Formal verification of ethical choices in autonomous systems. *Robotics and Autonomous Systems*, 77(Supplement C):1 – 14, 2016.
- [19] Department for Transport (DfT) and Centre for the Protection of National Infrastructure (CPNI). The key principles of cyber security for connected and automated vehicles. Technical report, 2017.
- [20] Department of Motor Vehicles (State of California). Testing of Autonomous Vehicles. <https://www.dmv.ca.gov/portal/dmv/detail/vr/autonomous/testing>.
- [21] V. V. Dixit, S. Chand, and D. J. Nair. Autonomous vehicles: Disengagements, accidents and reaction times. *PLOS ONE*, 11(12):1–14, 12 2016.
- [22] G. Dodig Crnkovic and B. Çürüklü. Robots: ethical by design. *Ethics and Information Technology*, 14(1):61–71, Mar 2012.
- [23] G. Dodig-Crnkovic and D. Persson. Sharing moral responsibility with robots: A pragmatic approach. In *Proceedings of the 2008 Conference on Tenth Scandinavian Conference on Artificial Intelligence: SCAI 2008*, pages 165–168, Amsterdam, The Netherlands, The Netherlands, 2008. IOS Press.
- [24] D. Dolgov. Google self-driving car project - monthly report - september 2016 - on the road. Technical report, Google, 2016.

- [25] S. I. Earth Imaging Journal (EIJ): Remote Sensing, Satellite Images. Lidar boosts brain power for self-driving cars, 2012.
- [26] L. Eckstein and M. Schwalm. Wahrnehmung: Auge oder Kamera - wer sieht besser? - ZF Friedrichshafen AG, 2016.
- [27] Ethics Commission. Automated and connected driving. Technical report, Federal Ministry of Transport and Digital Infrastructure, 2017.
- [28] European Union. Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/ec (general data protection regulation). Technical report, European Union, 2016.
- [29] P. Foot. The problem of abortion and the doctrine of double effect. *Oxford Review*, 5, 1967.
- [30] A.-K. Frison, P. Wintersberger, and A. Riener. First person trolley problem: Evaluation of drivers' ethical decisions in a driving simulator. In *Adjunct Proceedings of the 8th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, AutomotiveUI '16 Adjunct, pages 117–122, New York, NY, USA, 2016. ACM.
- [31] G. Ghisio. Challenges for the Automotive Platform of the Future, 2016.
- [32] N. J. Goodall. Vehicle automation and the duty to act. In *Proceedings of the 21st world congress on intelligent transport systems*, pages 7–11, 2014.
- [33] N. J. Goodall. Can you program ethics into a self-driving car? *IEEE Spectrum*, 53(6):28–58, June 2016.
- [34] B. Goodman and S. Flaxman. European Union regulations on algorithmic decision-making and a "right to explanation". *ArXiv e-prints*, June 2016.
- [35] Google. Google self-driving car project, 2016.
- [36] J. D. Greene. Our driverless dilemma. *Science*, 352(6293):1514–1515, 2016.
- [37] L. Greenemeier. Driverless Cars Will Face Moral Dilemmas. <https://www.scientificamerican.com/article/driverless-cars-will-face-moral-dilemmas/>, 2016.
- [38] A. Johnsen, G. D. Crnkovic, K. Lundqvist, K. Häröinen, and P. Pettersson. Risk-based decision-making fallacies: Why present functional safety standards are not enough. In *2017 IEEE International Conference on Software Architecture Workshops (ICSAW)*, pages 153–160, April 2017.
- [39] N. Kalra and S. M. Paddock. Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability? *Transportation Research Part A: Policy and Practice*, 94(Supplement C):182 – 193, 2016.
- [40] K. Kirkpatrick. The moral challenges of driverless cars. *Commun. ACM*, 58(8):19–20, July 2015.
- [41] S. Kuchinskas. Crash Course: Training the Brain of a Driverless Car. <https://www.scientificamerican.com/article/autonomous-driverless-car-brain/>, 2013.
- [42] B. MacKinnon. *Ethics: Theory and Contemporary Issues, Concise Edition*. Cengage Learning, 2012.
- [43] N. McBride. The ethics of driverless cars. *SIGCAS Comput. Soc.*, 45(3):179–184, Jan. 2016.
- [44] C. Mooney. Save the driver or save the crowd? Scientists wonder how driverless cars will 'choose'. <https://www.washingtonpost.com/news/energy-environment/wp/2016/06/23/save-the-driver-or-save-the-crowd-scientists-wonder-how-driverless-cars-will-choose/>, 2016.
- [45] National Highway Traffic Safety Administration (NHTSA). Federal automated vehicles policy - accelerating the next revolution in roadway safety. Technical report, U.S. Department of Transportation, 2016.
- [46] N. H. T. S. A. (NHTSA). "dot/nhtsa policy statement concerning automated vehicles" 2016 update to "preliminary statement of policy concerning automated vehicles". Technical report, National Highway Traffic Safety Administration (NHTSA).
- [47] P. Pelliccione, E. Knauss, R. Haldal, S. M. Ågren, P. Mallozzi, A. Almingier, and D. Borgentun. Automotive architecture framework: The experience of volvo cars. *Journal of Systems Architecture*, 77(Supplement C):83 – 100, 2017.
- [48] M. Persson and S. Elfström. Volvo Car Group's first self-driving Autopilot cars test on public roads around Gothenburg, 2014.
- [49] S. Pillath. Briefing: Automated vehicles in the EU. *European Parliamentary Research Service (EPRS)*, (January):12, 2016.
- [50] A. Riener, M. P. Jeon, I. Alvarez, B. Pfleging, A. Mirnig, M. Tscheligi, and L. Chuang. 1st workshop on ethically inspired user interfaces for automated driving. In *Adjunct Proceedings of the 8th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*, AutomotiveUI '16 Adjunct, pages 217–220, New York, NY, USA, 2016. ACM.
- [51] SAE. Taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles. *Global Ground Vehicle Standards*, (J3016):30, 2016.
- [52] G. Sapienza, G. Dodig-Crnkovic, and I. Crnkovic. Inclusion of ethical aspects in multi-criteria decision analysis. In *2016 1st International Workshop on Decision Making in Software ARCHitecture (MARCH)*, pages 1–8, April 2016.
- [53] H. Schäbe and J. Braband. Basic requirements for proven-in-use arguments. *CoRR*, abs/1511.01839, 2015.
- [54] A. Shashkevich. Stanford professors discuss ethics involving driverless cars. <https://news.stanford.edu/2017/05/22/stanford-scholars-researchers-discuss-key-ethical-questions-self-driving-cars-present/>, may 2017.
- [55] S. Spiekermann. *Ethical IT Innovation: A Value-Based System Design Approach*. Taylor & Francis, 2015.
- [56] J. D. Stoll. Gm executive credits silicon valley for accelerating development of self-driving cars, 2016.
- [57] M. Taylor. Self-Driving Mercedes-Benzes Will Prioritize Occupant Safety over Pedestrians. <https://blog.caranddriver.com/self-driving-mercedes-will-prioritize-occupant-safety-over-pedestrians/>, October 2016.
- [58] Tesla. A tragic loss | tesla deutschland, 2016.
- [59] Tesla. Upgrading Autopilot: Seeing the World in Radar | Tesla Deutschland, 2016.
- [60] A. Thekkilakattil and G. Dodig-Crnkovic. Ethics aspects of embedded and cyber-physical systems. In *2015 IEEE 39th Annual Computer Software and Applications Conference*, volume 2, pages 39–44, July 2015.
- [61] Toyota. New toyota test vehicle paves the way for commercialization of automated highway driving technologies | toyota global newsroom, 2015.
- [62] S. Wachter, B. Mittelstadt, and L. Floridi. Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *International Data Privacy Law*, 7(2):76–99, 2017.
- [63] M. M. Waldrop. Autonomous vehicles: No drivers required. *Nature*, 518:20–3, 2015.
- [64] Waymo. Technology - Waymo, 2017. <https://waymo.com/tech/>.
- [65] Waymo. Waymo, September 2017. <https://waymo.com>.
- [66] P. Wintersberger, A. K. Prison, A. Riener, and S. Hasrioglu. The experience of ethics: Evaluation of self harm risks in automated vehicles. In *2017 IEEE Intelligent Vehicles Symposium (IV)*, pages 385–391, June 2017.