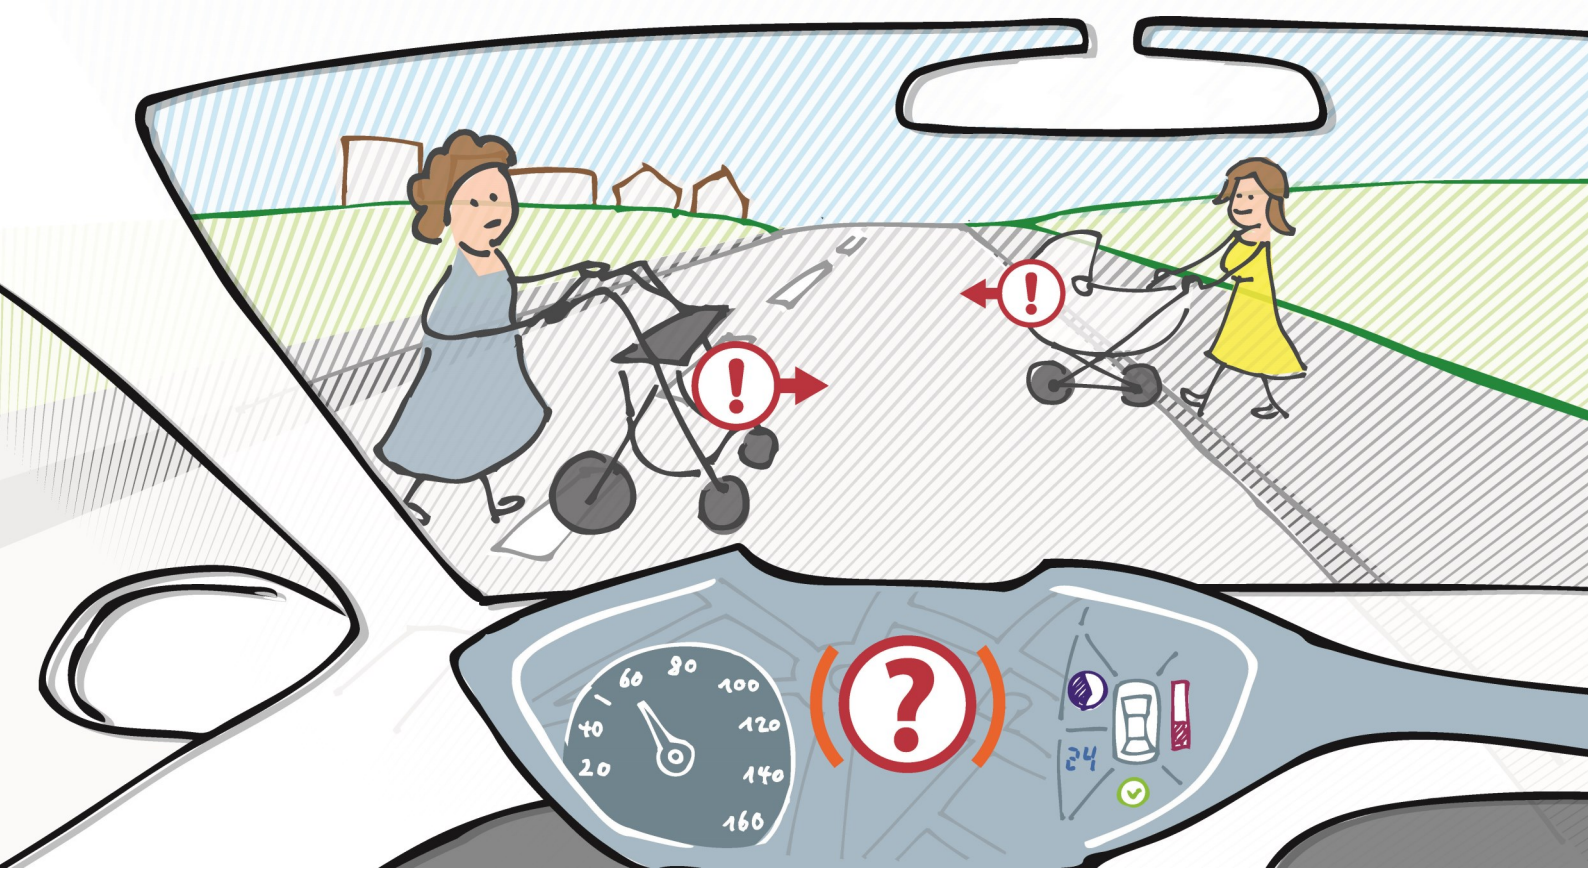

Autonomous Driving:

Cars in a moral dilemma?

In the context of the long-term goal of fully autonomous vehicles, numerous challenges confront the vehicle behavior. How should autonomous vehicles behave in critical situations? Can the desired behavior of an autonomous vehicle be determined in the first place? Do we need machine ethics for autonomous vehicles and if so, how do we achieve it?

This white paper highlights fundamental decision problems and rule-based ("top down") as well as learning based ("bottom up") approaches. It turns out that despite diverse challenges the way for fully autonomous vehicles is not only possible but even necessary from a social perspective.

Stuttgart, July 2015

mm1

Autonomous vehicles as a moral challenge

Increasing autonomy of digital and connected systems

Digitalization and automation are playing ever greater roles in our lives. Products and services are connected with each other and are finding their way into the virtual world of the Internet. Complex algorithms make automated decisions and control processes. The people's ability to take self-determined action based on their own principles, i.e. human autonomy (cf. Plauen & Welzer, 2014), is augmented by autonomous artificial systems. We transfer tasks to these systems that we do not want to or cannot carry out (better) ourselves.

The autonomous vehicle as a guiding principle for the automobile industry

The automobile industry, too, is characterized by digitalization and automation. Modern vehicles are equipped with numerous assistance systems which actively and passively support drivers in their driving task. Ultimately, when automation is fully perfected, the autonomous vehicle won't need a driver at all.[1] Advocates of autonomous vehicles list numerous arguments in favor of this technology: fewer accidents, improved traffic flow, more efficient use of resources, access to individual mobility for groups of persons previously excluded as well as increased comfort and productivity thanks to additional activities that are made possible while driving.[2]

The challenge of social acceptance

Societal norms determine the acceptance of new technologies.

Do autonomous vehicles really have the potential to succeed on the market? This requires a product to offer outstanding features across four dimensions: technology, customer benefit, business model and social acceptance (Scholz, 2015). Social acceptance is understood to mean the absence of resistance to, or even a benevolent tolerance of the technology in question or its implementation (cf. Schäfer & Keppler, 2013). It depends on whether the implementation of the technology is compatible with existing social and legal standards.

In many cases, ethical arguments are not only taken into account in an overall consideration of costs/benefits, but they also serve as constraints that need to be

[1] Cf. Gasser et al. (2012) for the levels of automation of driving functions in vehicles.

[2] Additional effects could include a fundamental change in our relationship to automobiles. Depending on the development of this technology, the traditional business model pursued by the automobile industry might change radically (e.g. usage instead of ownership, fully autonomous cab services, etc.).

met. No matter how great the benefits for society as a whole may be (e.g. treatment of neurological diseases through stem cell research), unethical use of technology is simply unacceptable to a majority of society.

The moral dilemma of autonomous vehicles

The development of autonomous vehicles is also accompanied by an ethically motivated debate. Introduction of such vehicles is preceded by the promise that damage to persons will be significantly reduced. However, the risk of accidents will remain despite the increasing degree of automation and maturity of the technology, e.g. due to implementation errors, known system limitations or uncontrollable traffic situations. This alone poses a considerable obstacle for social acceptance, as the demands on security and reliability of autonomous vehicles by society will be far higher than those on human drivers.

Compared to human drivers, the demands on safety and reliability of autonomous vehicles will be higher.

But even if one assumes that the systems are functioning perfectly, the question arises of how the vehicle's behavior should be designed. The vehicle must be able to cope with any traffic situation, and do so in a way that satisfies the social demands on the "right" behavior. A certain level of consensus on fundamental ethical principles is a prerequisite for developing suitable algorithms. The following though experiment, however, illustrates that it is unclear whether such a consensus can actually be reached:[3]

An autonomous vehicle is traveling on a single-lane road that is flanked by a guardrail on the left and a sidewalk on the right. Suddenly, two senior citizens step onto the street. Under normal circumstances, the vehicle would initiate an emergency stop or significantly reduce its speed in order to perform an evasive maneuver along the (deserted!) sidewalk. But let us assume that the remaining braking distance is too short for an emergency stop and that a mother and her son are walking along the sidewalk. Hence, the vehicle must make a decision: Should it (A) stay on the road and hit the senior citizens or (B) swerve onto the sidewalk and hit the mother and her child?

There is disagreement, in moral terms, on what would be the better option. On

[3] Dilemmas with this basic structure are known as "Sophie's Choice", named after the movie by Alan J. Pakula (1982).

the one hand, it should be taken into account that the senior citizens have already attained a considerable age, and that the mother – and her child in particular – still have a long life ahead of them (therefore option A). On the other hand, younger people have a considerably higher chance of surviving severe injuries than older people (Goodall, 2014; p. 62) (option B). But that aside, one may also argue that any weighing of human life is immoral, which is why the vehicle should accept its fate and act as passively as possible; i.e. braking and keeping the wheels straight (synonymous with option A).

The divergence of moral views within society confronts developers of autonomous vehicles with a seemingly impossible task.

In essence, this means that developers of autonomous vehicles are faced with a seemingly unresolvable task. No matter what moral intuitions they integrate: a large proportion of the population will consider this implementation to be morally reprehensible. As a consequence, this would endanger social acceptance. The marketability of autonomous vehicles can be called into question based on this finding. However, the following will show that concentrating on moral dilemmas is not a useful approach to normative discussions on this subject and that these dilemmas only have limited relevance.

Solution approaches from an ethical perspective

The exemplary arguments presented for and against options A and B correspond to the two most significant ethics theories: utilitarianism and Kantianism. Both theories are difficult to apply in practice.

Utilitarianism: the end justifies the means

Utilitarianism determines the moral value of an action based on its consequences. The actions that are to be carried out are supposed to maximize utility for all people. However, taking into account the consequences of all courses of action is not possible; at least not without drawing arbitrary spacial and temporal boundaries.[4] And no system will ever be able to process infinite amounts of data (Yampolsky, 2011; p. 4).

Furthermore, the implications for action in the context of utilitarianism often contradict moral intuition. Just imagine an autonomous vehicle that will inevitably

[4] In the context of the automobile industry, this issue can be illustrated excellently by an unauthorized commercial created by film students from Ludwigsburg. Cf. "Filmstudenten benutzen Mercedes zum Tyrannenmord", in: HORIZONT, August 23, 2013

collide with one of two motorcyclists (A and B). Unlike B, A is not wearing a helmet and has provoked this dangerous situation through his reckless riding style. B, on the other hand, is a cautious motorcyclist and also the brother of the vehicle passenger. From a utilitarian standpoint, utility would be maximized if the car steered towards B, as the helmet would provide him with a higher chance of survival. However, many people's sense of justice would be infringed upon by this course of action.

A utilitarian decision logic does not lead to consensual vehicle behavior.

Kantianism: Act in accordance with duty, no matter what happens

Kantianism demands that actions follow certain rules (e.g. "do not run over people"), regardless of the consequences. Kant's categorical imperative[5] prohibits using other people merely as a means to an end (e.g. to save the life of a third person) (Johnson, 2014). An autonomous vehicle programmed according to Kantianism would thus be unable to select the smaller of two evils in cases where an extreme appreciation of values is required. But it is easy to imagine a situation in which the number of people that could be saved is so high that our moral intuition would allow harming one individual (e.g. fully occupied bus vs. single driver).

The result of applying the categorical imperative collides with moral intuition.

Ex ante, a dilemma remains a dilemma

As shown above, neither utilitarianism nor Kantianism is able to provide solutions for the described dilemma that would be compatible with the moral intuition of the majority of people. With the desired social acceptance in mind, these moral theories are therefore unsuitable for top-down implementation in autonomous vehicles (Hars, 2014; p. 32 f.). Programmers of such vehicles should thus not be expected to come up with solutions for moral dilemmas, given that it has not even been possible to reach consensus within the moral philosophy community.

We should not expect programmers of autonomous vehicles to resolve moral dilemmas, if there is no consensus even in moral philosophy.

Ex post, moral assessment is possible

Retrospective assessment of the dilemma, on the other hand, is undisputed. In line with the widely accepted concept of moral responsibility, a person is only morally responsible when three prerequisites were met (Braham & van Hees, 2010 & 2011). For example, one must have been able to not bring about the ac-

[5] Kant has formulated the main principle as a categorical imperative: "Act only according to that maxim by which you can at the same time will that it should become universal law." – Immanuel Kant, Groundwork of the Metaphysics of Morals, 1785

Whether an unavoidable personal injury has been caused by man or machine should not have an impact on how we evaluate the situation.

tual result; hence, there had to be a freedom of choice between multiple reasonable options. In the situation described above, every possible action alternative would result in serious personal injury and is thus not reasonable. Consequently, the driver in the dilemma – man or machine – is not morally responsible for the personal injuries caused.[6]

Solution approaches from an empirical perspective

In the discussion of moral theory, the dilemma described has been assessed from a moral perspective. However, the question remains: How should an autonomous vehicle actually react in such a situation?

A pragmatic perspective is now taken to answer this question. There are four possible solutions: firstly, by taking into account introductory scenarios; secondly, by taking into account dynamics and coordination; thirdly, the primacy of caution and damage mitigation; and fourthly, by implementing non-deterministic decision-making systems.

Avoiding the dilemma thanks to introductory scenarios

Even if the desired goal is a vehicle that acts fully autonomously in every possible traffic situation, it is unrealistic to assume a "Big Bang" to occur on the introductory date of automated driving functions. Instead, the introduction of highly, and then fully, automated driving will take place step by step[7]; thus, it will at first be restricted to specific scenarios. These scenarios will only permit autonomous driving in defined traffic areas. In addition, it must still be possible for vehicle passengers to manually take over driving tasks at any time.

In introduction scenarios for highly automated driving, the moral dilemma does not arise.

The most likely introductory scenario for autonomous vehicles is driving on defined highway sections (Fraunhofer IAO et al., 2015). Clearance of these sections is based on the evaluation of a wide range of parameters, and only granted for defined periods of time. Additionally, it must be assumed that driving onto or leaving a highway will continue to be done manually.

The gradual introduction of autonomous driving actually helps to answer the

[6] It is certainly possible that one is not morally responsible for an action and its consequences; despite simultaneous existence of causal or criminal responsibility.

[7] With regard to automation levels, cf. the definitions of the degree of automation of vehicles provided by the German Federal Highway Research Institute (BASt) or the Society of Automobile Engineers (SAE).

question of social acceptance in combination with the initial dilemma. After all, there are no women with baby strollers on a highway, nor do senior citizens cross the traffic lanes. Of course, critical situations – in which the autonomous vehicle must make a decision on how to minimize damage – may also occur on highways. While it is not possible to find a rule-based algorithm for the proper weighing between "collision with child" and "collision with adults", it is certainly possible to define a rule-based control mechanism for the highway scenario. Simply put, this could be as follows: 1) Search for a collision-free path by adapting longitudinal and lateral control; 2) if it is not possible to find such a path, choose the path with the least kinetic impact energy (and reduce kinetic energy as quickly as possible by braking).

An additional argument that defuses the moral dilemma of the autonomous vehicle in introductory scenarios is the option – and legal obligation – of the passenger in the driver's seat to assume control of the vehicles at any time. Legal provisions ultimately exempt the developers from moral responsibility. For it is people who have the moral responsibility, as they have the final say in the decision-making process.

Avoiding the dilemma by means of coordinative systems

It is a well-known fact that finding a solution for conflicts that arise in traffic situations is based on the complex and dynamic behavior of all those involved. In addition to making it possible to reach a consensus, communication also enables subsequent coordination. This idea, for example, is incorporated in concept vehicles, which communicate their intentions to other road users through simple signals.[8] This includes coordination between autonomous vehicles, but also between autonomous vehicles and pedestrians, conventional vehicles or traffic infrastructure.[9] Thus, coordinative systems provide an additional technical solution approach when it comes to avoiding the dilemma.

The solution is not only a result of one's own behavior, but the result of the behavior of all those involved in a complex and dynamic system.

The dynamics in solving problems jointly shows that the perception of the auton-

[8] The visual and acoustic signals generated by the Mercedes-Benz F 015 concept vehicle are one example of this approach.

[9] One problem, which is not explained in detail here, is centered around the question of when solitary decision-making should take place and when coordination with other players should be pursued. For more information on this issue, see, for example, Tranni et al.: From Solitary to Collective Behaviours: Decision Making and Cooperation, in: Advances in Artificial Life, Springer Verlag 2007, p. 575-584.

omous vehicle as a solitary and reactively acting subject – which is implicit in the dilemma scenario – is misleading. Rather, we are in a process of joint decision-making and coordinated actions. The range of possible solutions in a given situation is thus broadened considerably. In this case, this means: collaboration between vehicles and stop lights could proactively ensure that all persons are in a safe zone. Additionally, a vehicle's warning signals could ensure that people become aware of the immediate danger and move to a safe location.

Mitigating the dilemma with systems acting cautiously

The moral urgency of the decision in a dilemma situation is toned down by the reduction of potential damage.

Reducing potentially negative consequences can be seen as an additional solution approach ("mitigating the dilemma"). In human perception, the moral implications of a decision are directly tied to the severity of the consequences.

In relation to the dilemma in question, the result is the moral imperative of autonomous vehicles adopting a defensive driving style, characterized by prudence and not insisting on the formal entitlement to specific rights. Speed marks the decisive influencing factor in urban areas. After all, in a collision, impact speed is the most important determining factor of injury severity for those involved (Kühn et al., 2007), (Richter et al., 2001; p. 733). A reduction in speed results in a significant decrease in potentially negative consequences and considerably mitigates the decision-making dilemma described above. Furthermore, the probability of accidents occurring is itself reduced when vehicles travel at slower speeds.[10] In the context of the debate, the authors therefore make a case for reducing the maximum speed of autonomous vehicles to no more than 30 km/h in urban areas.

With regard to the social acceptance of autonomous vehicles, the question arises whether such a limitation will have a positive or negative effect. The authors assume that such a step will result in an overall positive perception[11], as maximum speed is less important in city traffic than average speed and consistent traffic flow. The greater the penetration of autonomous vehicles, the better the traffic

[10] It is assumed that the slower speed will provide more vulnerable traffic participants and the vehicle itself with more time to react in potentially critical situations.

[11] Nevertheless, it is also worth mentioning that fully autonomous vehicles must comply with all applicable traffic laws. In road traffic, this can actually lead to a lack of acceptance by vehicle operators or even result in provoking high-risk driving maneuvers (e.g. overtaking other vehicles).

flow can be organized (Au et al., 2014).

Resolving the dilemma through non-deterministic decision-making systems

Despite the approaches presented for avoiding and minimizing the dilemma, it cannot be ignored that it is definitely realistic in the context of city traffic. It is therefore absolutely imperative to implement an action competence in the autonomous vehicle that is able to cope with the decision-making problem in a dilemma situation.

In the end we still need to deal with the dilemma.

This is where the rule-based thought process reaches a dead end. As already explained in the moral theory consideration, the initial dilemma is primarily characterized by the fact that it is not possible to find a generally accepted rule for the "right actions". Hence, no such rule can be translated into algorithms. Yet, an alternative to the rule-based solution process can be derived from the following thought: a fully autonomous vehicle should behave like a perfect driver, thus imitating the ideal human being. So, how must a system be designed so that its mode of operation in principle follows human role models?

Rule-based approaches lead to a dead end.

Artificial neural networks represent a technical approach to replicating the human brain. Neural networks connect a variety of neurons (functional units) that receive signals, process them, and, depending on the strength of the impulse received, convert them into a separate impulse which is in turn to be transmitted to other neurons.[12] The architecture of artificial neural networks results in fundamental characteristics that correspond to those of the human brain. This provides them with the ability to learn and adapt, they work in parallel to a large extent and are resistant to disruptions. In addition, artificial neural networks are not exclusively deterministic, i.e. not strictly rule-based. They rather work stochastically, meaning with probabilities and a system-imminent vagueness in their behavior. Furthermore, they can generalize problems, hence, resolve specific problems based on information already stored in form of network connections without ever having solved this exact problem in the past. It is this characteristic that ultimately brings the moral dilemma to mind: Although we must solve it, we ultimately do not want to – at least not on the basis of rigid rules.

[12] A useful introduction to neural networks can be found at http://www.dkriesel.com/science/neural_networks

Some of the characteristics described can now be applied to the specific problem: Firstly, a neural network can be trained thanks to its ability to adapt. This training is, for example, possible by confronting the network with environmental situations in which consensus exists on the "best" moral solution. The result of this learning process encompasses correspondingly modified connection characteristics between the individual neurons (Klüver, 2012; p. 127). Figuratively speaking, the neural network implicitly learns the moral values on which the target solutions are based ("you shall not harm any person").

Thanks to the ability of generalization, neural networks are able to bring forth a concrete behavior

Secondly, due to its ability to generalize, a problem can be presented to a neural network for which we ourselves do not have an answer. And this is exactly the situation in a moral dilemma. The neural network is able to comprehend the situation and, thanks to its ability to generalize, derive instructions on the course of action. However, the way a neural network derives its solution for a specific situation cannot be replicated in detail. This is similar to human decisions: they may be comprehensible in their entirety, but the cytological processes leading to a decision will never be fully understood, let alone due to their complexity.

That is why a trick is used in the implementation of an artificial neural network for resolving the moral dilemma: a system is developed which is able to resolve the problem – without a predefined solution. The responsibility for resolving the dilemma is transferred to the autonomous vehicle, without the possibility of a specific moral recourse on its human "programmer" – because a specific answer to the specific decision-making problem has never been provided, nor has a specific action instruction (rule) been implemented in form of an algorithm. The vehicle "makes the best" out of the situation within the constraints of its abilities.

Will society accept such an implementation approach? From the authors' perspectives, the chances for success do not depend on the specific handling of one specific environmental situation, but rather on the functioning in its entirety, i.e. the general ability to perform in daily life. Acceptance should not be problem as long as the fully automatic control works well in the majority of cases (in the sense of a "parts per million" claim). Dealing with other cases has already been detailed in the context of the moral assessment.

The perspective for society as a whole

The discussion of the ethical-normative implications of introducing autonomous vehicles must include an assessment of the status quo with respect to normatively relevant categories such as driving safety. This does not mean that finding the best possible solution for the above dilemma can simply be dismissed as being irrelevant. However, the excessive focus on specific and very rare situations obstructs the view of the global advantages offered by the technology.

Currently, nearly 3,400 people are killed on German roads every year; approximately 400,000 are injured (German Federal Statistical Office, 2015). About 90% of all car accidents can be traced back to human error (Hars, 2015; p. 27; Lin, 2014), and one-third of these accidents are related to drinking and driving (Thierer & Hagemann, 2014; p. 15). If autonomous vehicles are able to significantly reduce the number and/or severity of personal traffic-related injuries, then this would be of considerable significance for the normative assessment. Thanks to numerous sensors and sufficient computing power, autonomous vehicles will be able to react to traffic situations more quickly than human beings ever could. And – other than humans – they are never drunk, tired, distracted or erratic in their behavior. Indeed, technology-related accidents that result in injuries to other traffic participants may occur with autonomous vehicles, but nevertheless, it is not to be expected that such accidents, in terms of their frequency and severity, will come anywhere close to the scope of accidents that are caused by human error today. It is thus expected that in Germany alone, a 10% market penetration of autonomous vehicles would save 120 lives per year; a penetration of 50% would save 1,200 lives (cf.: Thierer & Hagemann, p. 16-17).

It is of great importance for the normative evaluation that autonomous vehicles significantly reduce personal injury in road traffic.

mm1

Conclusion

Social acceptance is an important factor for the introduction of autonomous vehicles. It is thus essential to establish a serious discourse on how such vehicles are to behave and where limitations to their use are to be defined. Such a discourse should not be limited to automated vehicles, but rather it should also generally deal with the increasing transfer of decision-making and operational autonomy to digital and networked systems.[13]

In future, the question of whether autonomous vehicles are socially accepted will be obsolete. The question will rather be: how long will society accept human drivers?

From the authors' perspective, the fear that the moral dilemma described above could influence social acceptance and thus prevent the introduction of autonomous vehicles is unfounded. Even if it is not possible to resolve the dilemma itself a priori, we are in a position to understand technical limitations the same way as human ones. We also have the possibility of preventing a dilemma from occurring, or mitigating its consequences. Ultimately, it is even possible to circumvent the dilemma itself in form of non-deterministic systems, without ignoring it. And, generally, autonomous vehicles come with the promise of making all traffic safer.

In future, the question of whether autonomous vehicles are socially accepted will be obsolete. The question will rather be: how long will society accept human drivers?

[13] Fundamental questions on the future relationship between man and machine arise in the context of using artificial neural networks in safety-critical applications. The photographer Max Aguilera-Hellweg puts it in a nutshell: "I am not certain whether I like the idea of having a machine decide on life and death. This will inevitably result in collateral damage. [...]. Reality is never black or white, it is made up of all possible shades of gray. That's what I like about it. Every time you begin to understand something, new questions arise. Incessantly. You always encounter a new layer which you were unaware of before. Doing research on robots means opening Pandora's Box." (quote from: Roboter – Noch Maschine oder schon Mensch? Marc Felix & Bruno Victor-Pujebet, Arte 2014)

mm1

About the authors:

Volker Scholz is co-founder and Managing Partner of mm1, the Consultancy for Connected Business. His clients benefit from two decades of consulting experience in the conception of innovative products as well as the development of new business fields in the context of the connected world. He is responsible for mm1's activities in the areas of Connected Car, Sustainable Innovation and Agile & Lean Startup. He has been focusing on issues involving the success factors of autonomous vehicles for several years.

Marius Kempf is a consultant at mm1 who focuses on technology assessment in digitalization and automation processes. During his studies in Philosophy & Economics, he worked on numerous issues that can be applied to Connected Business – for example to the design of autonomous vehicles.

mm1 Consulting & Management PartG
Bolzstraße 6
70173 Stuttgart
Tel.: +49 711 184210-0
office@mm1.de

Bibliography

Au, Tsz, Zhang, Shun & Stone, Peter (2014). Autonomous Intersection Management for Semi-Autonomous Vehicles. In Handbook of Transportation, Routledge, Taylor & Francis Group.

Braham, Matthew & Hees, Martin van (2011). Responsibility Voids. *Philosophical Quarterly* 61: 6-15.

Braham, Matthew & Hees, Martin van (2012). An Anatomy of Moral Responsibility. *Mind* 121: 601-634.

Fraunhofer IAO, Fraunhofer FOKUS, Institute for Climate Protection, Energy and Mobility—Law, Economics and Policy & mm1 Consulting & Management (vorauss. 2015). Hochautomatisiertes Fahren auf Autobahnen – Industriepolitische Schlussfolgerungen. Study on behalf of the Federal Ministriy for Economic Affairs and Energy

Gasser, Tom M. et al. (2012). Rechtsfolgen zunehmender Fahrzeugautomatisierung. *Bundesanstalt für Straßenwesen: Fahrzeug Kompakt* 11/12.

Goodall, Noah J. (2014). Ethical Decision Making during Automated Vehicle Crashes. *Transportation Research Record: Journal of the Transportation Research Board* 2424: 58-65.

Hars, Alexander (2015). Selbstfahrende Fahrzeuge: neun Thesen. *Zeitschrift für die gesamte Wertschöpfungskette Automobilwirtschaft* 1/2015: 27-34

Johnson, Robert (2014). Kant's Moral Philosophy. *Stanford Encyclopedia of Philosophy* (Summer 2014 Edition), Edward N. Zalta (Hrsg.).

Klüver, Christina et al. (2012). Modellierung komplexer Prozesse durch naturanaloge Verfahren: Soft Computing und verwandte Techniken. Springer-Vieweg (Wiesbaden).

Kühn, Matthias et al. (2007). Fussgängerschutz: Unfallgeschehen, Fahrzeuggestaltung, Testverfahren. *Springer Verlag* (Berlin).

Kriesel, David (2007). Ein kleiner Überblick über Neuronale Netze, available on http://www.dkriesel.com

Lin, Patrick (2014). The Robot Car of Tomorrow May Just be Programmed to Hit you. *Wired.com*. URL = http://www.wired.com/2014/05/the-robot-car-of-tomorrow-might-just-be-programmed-to-hit-you/

Plauen, Michael & Welzer, Harald (2014). Autonomie: Eine Verteidigung. S. Fischer Verlag (Berlin).

Richter, Martinius et al. (2001). Problematik der Verletzungen von Kindern und Jugendlichen im Straßenverkehr. Unfallchirurg 8/2001. Springer Verlag (Berlin)

Schäfer, Martina & Keppler, Dorothee (2013). *Modelle der technikorientierten Akzeptanzforschung: Überblick und Reflexion am Beispiel eines Forschungsprojekts zur*

Implementierung innovativer technischer Energieeffizienz-Maßnahmen. Technische Universität Berlin.

Scholz, Volker (2015) Autonomes Fahren - Herausforderung in vier Dimensionen. *Zeitschrift für die gesamte Wertschöpfungskette Automobilwirtschaft* 1/2015: 6-11

German Federal Statistical Office (2015). Zahl der Verkehrstoten im Jahr 2014 um 0,9% gestiegen. *Pressemitteilung* Nr.065.

German Federal Statistical Office (2013). Unfallentwicklung auf Deutschen Straßen 2012.

Thierer, Adam & Hagemann, Ryan (2014). Removing Roadblocks to Intelligent Vehicles and Driverless Cars. *Mercatus Working Paper*.

Vanderbilt, Tom (2012). Five Reasons the Robo-Car Haters are Wrong. *Wired.com*. URL = < http://www.wired.com/2012/02/robo-car-haters-are-wrong/>

Yampolskiy, Roman V. & Fox, Joshua (2012). Safety Engineering for Artificial General Intelligence. *Topoi* 32 (2): 217-226.

mm1

mm1 Consulting & Management PartG

Bolzstraße 6
D-70173 Stuttgart
Tel.: +49 711 184210-0
office@mm1.de
www.mm1.de