

Chapter 10

Autonomous Vehicles



This chapter describes what autonomous vehicles are and how they work. We then discuss the associated risks and benefits not only from an ethical perspective, but also from a legal and environmental view. Core issues discussed are liability and safety of autonomous vehicles.

Some object to the term “autonomous vehicles.” In philosophy and ethics, “autonomy” is a concept stating that human beings are to make decisions on their own. They should not be coerced by anyone, and they should be free to give themselves their own rules within reason (Greek: “autos” = self + “nomos” = law). This is not what AI or autonomous vehicles (AVs) could do. Some think a better term would be “highly automated vehicles” or “fully automated vehicles”. In the context of AVs “autonomous” is used in the more standard robotic sense, which is simply the ability to operate for a protracted period of time without a human operator. However, since the term AV is well-established in the discussion, we will continue using it here. In the AV context, we need to distinguish between different levels or autonomy.

10.1 Levels of Autonomous Driving

There are several systems of classifying the levels of autonomous driving, in particular, one by the US Society of Automotive Engineers (SAE) and another by the German Association of the Automotive Industry (VDA). They differ, however, only in the details. The general idea behind these classifications goes as follows:

- **Level 0: No driving automation** This is a traditional vehicle without any automated functionality.

- **Level 1: Driver assistance** The vehicle has one type of automated functionality. For example, braking automatically when encountering an obstacle.
- **Level 2: Partial driving automation** The vehicle can perform both braking and accelerating functions as well as changing lanes functions. However, the driver has to monitor the system at all times and be ready to take control whenever necessary. For example, all Tesla vehicles are officially considered Level 2 automation.
- **Level 3: Conditional driving automation** The driver does not need to monitor the system at all times. Under certain circumstances the system can work autonomously. The system gives the driver time (10 seconds for example) before handing back control. In 2018, the Audi A8 claimed to be the first car capable of Level 3 automation.
- **Level 4: High driving automation** The vehicle can perform all driving functions under standard circumstances. The driver is not required to take control under standard circumstances. Non-standard conditions would include inclement weather.
- **Level 5: Full driving automation** The vehicle can perform all driving functions in all circumstances. In the German classification, this is labelled as “no driver”, making the car completely autonomous.

10.2 Current Situation

As of 2019 there are no commercially available autonomous vehicles beyond level 2 or level 3 automation. It is interesting to note that a number of new companies are challenging the traditional manufacturers, with some of them being on the market already (Tesla), and some of them are testing and collecting data while preparing for market entry at a later time (Waymo see Fig. 10.1). Test driving is, on the other hand, widespread already. Nevada gave Google a permit to test AVs in 2009. In Europe, the Netherlands, Germany and the UK permit AV testing, A number of US federal states have passed or changed regulations to allow autonomous test driving under certain conditions.

Testing of AVs in the US started in 1995 when an experimental autonomous vehicle developed by Carnegie Mellon University drove from Pittsburgh to San Diego. Rideshare platforms like Uber or Lyft are developing and testing autonomous taxis. Singapore completed a test on autonomous taxis in 2018. Japan is planning to introduce autonomous taxis on a large scale for the Tokyo 2020 Olympics. Moreover, in a number of places, autonomous buses are employed, usually in controlled environments.

10.3 Ethical Benefits of AVs

According to the available studies, AVs could bring about a number of safety benefits, especially by avoiding accidents and fatal crashes. It is estimated that around 90–95% of all automobile accidents are chiefly caused by human mistakes (Crew 2015).



Fig. 10.1 Waymo's fully self-driving Chrysler Pacifica Hybrid minivan on public roads (Source Waymo)

This raises the question of whether or not a society might be ethically obligated to legally permit autonomous vehicles.

Still, some argue that the presence of AVs may indirectly induce certain types of accidents when the AVs strictly follow the law: For example, following the Atlanta highway speed limit of 55 mph, when the speed of the average human driver might rather be 70 mph, will likely cause people to navigate around the autonomous vehicle, possibly resulting in a greater number of “human-caused” accidents. However, the autonomous cars cannot be blamed for these type of accidents since human drivers did not observe the rules of traffic in the first place.

10.4 Accidents with AVs

There have already been some accidents with autonomous vehicles. Even though the number is small, they have generated a lot of media attention. One of the first was an accident with a Tesla in 2016, where the vehicle did not recognise a truck, killing the driver. Sensory equipment has been improved since then (Hawkins 2018).

In March 2018, an Uber Volvo fatally hit a woman crossing the road in Tempe (Arizona), an accident which probably no human driver could have avoided. The car's emergency braking system was however deactivated on purpose to avoid false positives, and the driver was not acting appropriately (O'Kane 2018).

Another (fatal) accident happened with Tesla cars in March 2018 in Mountain View/California, and in Laguna Beach/California in May 2018. In both cases, it turned out that the respective driver ignored several warnings by the “autopilot” system.

These (and similar) cases show we need clear rules about how AVs are to be used by customers. Drivers should not sit behind the wheel playing games or reading their email. Companies should be required to communicate these rules in adequate ways.

10.5 Ethical Guidelines for AVs

Besides detailed regulations that have been passed in a number of countries, a broader ethics code for AVs has been adopted. In July 2016, the German Federal Minister of Transport and Digital Infrastructure appointed a national ethics committee for automated and connected driving. The committee was composed of 14 members, among them professors of law, ethics and engineering, as well as representatives from automotive companies and consumer organisations. The chairman was a former judge of the German Federal Constitutional Court. In addition, hearings with additional experts from technical, legal and ethical disciplines were conducted, as well as driving tests with several AVs. In June 2017, the committee presented 20 ethical guidelines for AVs (Federal Ministry of Transportation and Digital Infrastructure 2017). Many of these will have a significant direct or indirect bearing on the car industry and their ethics policies, both for Germany as well as the EU (Luetge 2017).

There are some researchers that feel that caution is necessary when discussing the ethical benefits of autonomous vehicles (Brooks 2017; Marshall 2018) and risks associated with autonomous vehicles (Cummings 2017). One issue is that the data of miles driven in autonomous mode might be skewed in terms of ideal driving conditions, that the behaviour of autonomous vehicles differs from certain ‘unwritten’ norms of human drivers (Surden and Williams 2016), and that people may place themselves at greater risk because of the presence of autonomous vehicles (Rothenbücher et al. 2016). The RAND corporation suggested in 2016 that an AV must drive 275 million miles without a fatality to prove that they are safe (Kalra and Paddock 2016). RAND researchers later however also stated that AVs should be deployed as soon as possible in order to save lives (Marshall 2017).

10.6 Ethical Questions in AVs

At least some of the accidents indicate that it may be ethically questionable to call a certain mode “autopilot” (which Tesla does), suggesting false connotations, since the driver of a Tesla car (which is officially Level 2) is required to keep the hands on the wheel at all times. Just stating that this is a matter of one’s own responsibility is insufficient—as with many other cases from different industries. In the 1970s, for

example, a massive scandal erupted for the Nestle company, as people believed it did not give adequate rules for mothers how to use their baby milk products. In the not-so-distant future, AVs might be introduced on a massive scale. Companies have a responsibility lay out drivers' responsibilities very clearly.

In addition, there are a number of other ethical questions usually associated with AVs.

10.6.1 Accountability and Liability

The question of liability will be critical. According to the 1968 Vienna Convention on Road Traffic (Vienna Convention on Road Traffic 1968, updated 2014, effective since 2016 (United Nations 1968), the driver is the one responsible for their car. However, in cases where an autonomous system is in control, this responsibility does not make much sense. The German Ethics Code therefore states that in these cases, the liability has to be turned over to the car manufacturer and the company operating or developing the software. Liability in autonomous driving will become a case of product liability and it will also require monitoring devices to be built into AVs. Future AV driving recorders might be similar to the flight recorders used on contemporary aircraft.

10.6.2 Situations of Unavoidable Accidents

Much of the literature on ethics of AVs revolves around situations similar to the "trolley problem" where an accident is unavoidable and an AV has to choose between two evils (Bonneton et al. 2016; Lin 2016). Should the car choose to swerve to avoid hitting four people and hit one person instead? Does it matter if the people are children, or older adults? Should it hit someone crossing the street at a red light rather than someone observing the traffic rules (see Fig. 10.2)?

How frequent these situations will be in practice is controversial. Some have argued that these situations will be quite rare, given that the conditions for them are extremely narrow. But if we accept their possibility for the moment, then one issue is to avoid creating a machine or an algorithm that selects targets according to personal characteristics (such as age or gender, which the German Ethics Code prohibits, or Social Credit Score), but still allowing a company to program code that reduces the overall number of fatalities or injuries. This is a complex problem, which is usually not as simple as selecting one of several targets to be killed with a guarantee. For example, there will probably be different probabilities for injuries or casualties of different targets. This may result in complicated situations in which it would not be ethical to forego the opportunity to reduce the overall damage to persons. We should make it clear that current autonomous vehicles do not engage in trolley problem

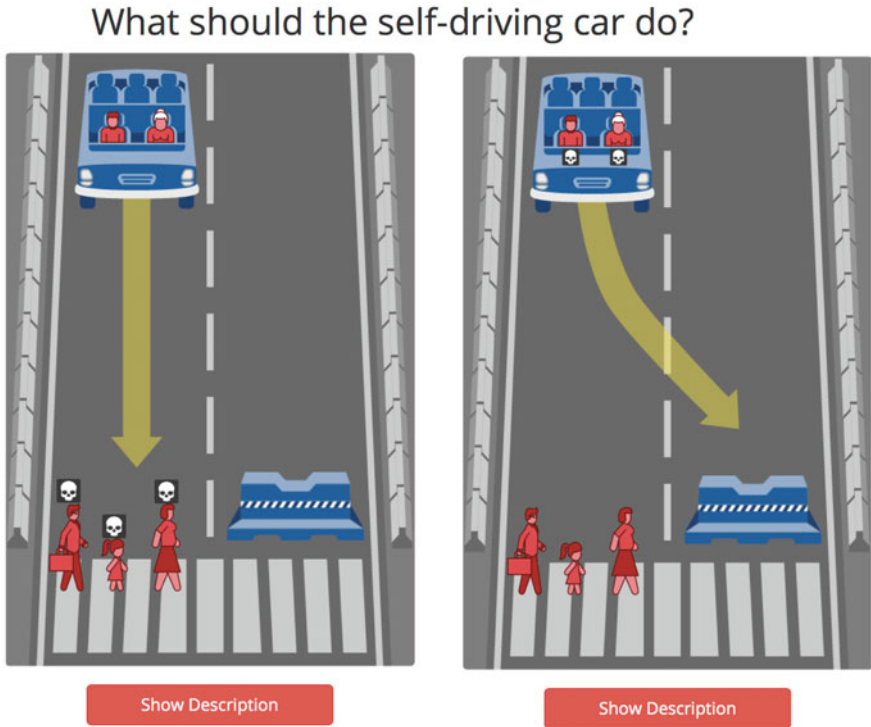


Fig. 10.2 Example question from the Moral Machine experiment that confronted people with trolley problems (Source MIT)

calculations. These vehicles attempt to avoid all obstacles and make no attempt to identify the obstacle, although in some cases they avoid obstacles based on size.

Cars do already know the number of passengers since they feature a seat belt warning system. Car to car wireless communication is also already available. In the case of an imminent crash, it would be technically feasible for the two cars to negotiate driving behaviour, similar to the Traffic Collision Avoidance System used already today in air planes. The car with more passengers might be given priority. While this technical possibility exists, it will be up to society to decide if such a system would be desirable.

However, what the German code explicitly does not say is that individual victims in different scenarios are allowed to be offset against each other. To some extent, this is the lesson of the German “Luftsicherheitsgesetz” (Aviation Security Act) being ruled unconstitutional by the German Federal Constitutional Court in 2006. The Aviation Security Act would have allowed to shooting down of hijacked aircraft that were thought to be used as weapons. In that case, individually known subjects would have been sacrificed for the sake of others. In the case of an anonymous program, however, no victims are known individually in advance. Rather, it is an abstract

guideline, the exact consequences of which cannot be foreseen, and which reduces the overall risk for all people affected by it. In this way, the risk could be regarded as similar to the risk that comes with vaccination.

Finally, parties not involved in the traffic situation with an AV should not be sacrificed. This implies that an algorithm should not unconditionally save the driver of an AV. However, as the German Ethics Code states, the driver should not come last either: after all, who would want to buy such a car?

10.6.3 Privacy Issues

The problem of privacy, which has come to the forefront especially with developments like the European General Data Protection Regulation, effective 2018 (see Sect. 8), is also relevant for AVs. These cars collect massive amounts of data every second they move, which is essential for evaluating their performance and improving their safety. On the other hand, such data may be very personal, tracking the passengers' locations and the behaviour of the driver. Such data might be used for purposes for which they were not intended. Ultimately, starting the engine of an AV will probably imply that the driver (and/or the passengers) accept the terms of conditions of the car manufacturer who will require access to such data.

It should also be noted that this problem is seen from different perspectives in different regions around the globe, especially considering the US or China, where data collection is seen as being less problematic than, for example, in Germany or other EU countries.

Autonomous vehicles generate data as they move from one location to another. This data may include observations of the individuals inside the car. Moreover, observations can be used to evaluate whether or not a person should be allowed to drive the car. Technology could be used to prevent people with suspended license or drinking problems from driving. For an autonomous vehicle, the vehicle itself may determine if the person's behaviour indicates intoxication and decide whether or not to allow the person to drive. Perhaps more of a concern, is when an autonomous vehicle uses data about the person's emotional state, mood, or personality to determine if their driving privileges should be curtailed.

10.6.4 Security

Security of AVs is an important issue, since hackers might breach the security of cars or their systems and cause substantial harm. In 2015, for example, hackers were able to remotely control a Jeep Cherokee while driving (Greenberg 2015). Cars might be reprogrammed to deliberately crash, even on a massive scale. Therefore, state-of-the-art security technology will have to be involved in order to make AVs as secure as possible.

10.6.5 Appropriate Design of Human-Machine Interface

The interface between a human driver and the AV is very important, since in a traffic situation, time is highly critical (Carsten and Martens 2018). It must be clear at all times who is in control, the handover procedure must be clearly defined, and the logging of behaviour and car control must be programmed in an appropriate way.

10.6.6 Machine Learning

AV software will certainly use machine learning in an offline manner. The systems will learn from traffic scenarios they encounter, from accidents or from the data generated by other vehicles. Vehicles, however, are not dynamically learning who and what to avoid while they are driving. Moreover, learning typically occurs on the level of an entire fleet of cars. And even here, learning must be robust to an extent where small mistakes (or external manipulations) cannot cause large-scale negative effects. Vehicles that learn on their own, might become quite erratic and unpredictable.

10.6.7 Manually Overruling the System?

Finally, one unresolved problem is whether a level 4 or 5 AV should be programmed in such a way as to allow the driver to overrule the autonomous system at any point. Engineers regularly reject this possibility stating that this feature would make a system much more vulnerable than it could be. Waymo even argued that its car would have avoided a collision in 2018 with a motorcyclist if the driver had not taken back control (Laris 2018).

Others argue that to deny overruling would jeopardise acceptance of AVs. At least in the current situation where people are not yet familiar with the technology, it might be better to allow for overruling, making a driver feel more comfortable. But this may be reviewed at some point in the future.

10.6.8 Possible Ethical Questions in Future Scenarios

While the ethical questions mentioned above are relevant today, there are some ethical aspects of AVs which are frequently discussed in public, even though they may be only relevant in the distant future. One of these is whether the introduction of AVs on a massive scale might lead to a critical centralisation of power in the control centres of this technology, which in turn might lead to a total surveillance of citizens and/or a deliberate or accidental misuse. Certainly, there are some scenarios here that should be kept in mind, even if a Big Brother scenario is unrealistic since it neglects the competition among different companies.

Another dystopian outcome that is frequently mentioned is whether AVs might become mandatory at some point in the future, forbidding the use of non-autonomous vehicles at all (Sparrow and Howard 2017). We believe that at this point, it is much too early to decide on this. There should certainly be a public discussion about whether such a step should be taken. But, right now, we are technically far from its implementation, and will rather, at least for the next decades, have to deal with a situation with mixed traffic in which autonomous vehicles will mix with traditional ones. Studies predict that this situation will already substantially prevent many accidents and save a lot of lives (Crew 2015).

Discussion Questions:

- On what level of “autonomy” (according to SAE and VDA) would you be comfortable to drive alone? What if you were with your child? Explain your reasoning.
- According to the Vienna Convention, who is responsible for a car? Is this law appropriate for AVs? Why? Why not?
- Do you think that young and old people should be treated differently by an AV’s algorithm? Why? Why not?

Further Reading:

- Jean-Francois Bonnefon, Azim Shariff, and Iyad Rahwan. The social dilemma of autonomous vehicles. *Science*, 352(6293):1573–1576, 2016. ISSN 0036-8075. Doi: 10.1126/science.aaf2654. URL <https://doi.org/10.1126/science.aaf2654>
- Noah J. Goodall. Can you program ethics into a self-driving car? *IEEE Spectrum*, 53(6):28–58, 2016. ISSN 0018-9235. Doi: 10.1109/MSPEC.2016.7473149. URL <https://doi.org/10.1109/MSPEC.2016.7473149>
- Patrick Lin. Why ethics matters for autonomous cars. In Markus Maurer, J Christian Gerdes, Barbara Lenz, Hermann Winner, et al., editors, *Autonomous driving*, pages 70–85. Springer, Berlin, Heidelberg, 2016. ISBN 978-3-662-48845-4. Doi: 10.1007/978-3-662-48847-8_4. URL https://doi.org/10.1007/978-3-662-48847-8_4
- Christoph Luetge. The German ethics code for automated and connected driving. *Philosophy & Technology*, 30(4):547–558, Dec 2017. ISSN 2210-5441. Doi: 10.1007/s13347-017-0284-0. URL <https://doi.org/10.1007/s13347-017-0284-0>.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

