Towards a Framework for Ethical Decision Making in Automated Vehicles

Leon René Sützelfeld¹, Peter König^{1,2}, Gordon Pipa¹

¹*Institute of Cognitive Science, Osnabrück University, Wachsbleiche 27, 49090 Osnabrück, Germany*

²*Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg-Eppendorf, Martinistr. 52, 20246, Hamburg, Germany*

Abstract

Automated driving technology advances quickly, and self-driving vehicles will soon no longer need human supervision. The ethical questions that the technology brings with it, however, are diverse and not always easily solvable. In particular, the question of morally right behavior in dilemma situations presents an unsolved issue to date, a solution to which has the potential to make our roads significantly safer and fairer. The main contribution of this paper is a discussion of central issues with regard to the regulation and implementation of ethical decision making systems in automated vehicles. We review existing guidelines with respect to their theoretical and practical feasibility, and highlight a number of unsolved issues that need to be addressed in the pursuit of a framework for automated ethics. We conclude with a list of 10 central demands and open issues that should be discussed on a societal level and need to be addressed by regulators and manufacturers.

Keywords: Automated Vehicles, Ethical Decision Making, Dilemma Situations, Framework, Open Issues

1. Introduction

Driven by fast progress in artificial intelligence, automated driving technology has in recent years become a key area of research and development in the automobile industry. However, getting to the point where automated vehicles (AVs) can be trusted to drive on public roads without human supervision isn't just a question of technology. It brings with it a number of

ethical and regulatory challenges that need to be addressed in order to facilitate a smooth launch of the technology, avoid costly legal disputes and public backlashes, and ultimately increase road safety [1, 2].

Many of the open ethical questions concerning automated driving revolve around route choices, handling of user data, or biases in AI-driven perception systems [3, 4]. Most prominently debated, and at the heart of this paper, are decisions that AVs may have to make in response to dangerous situations, or imminent collisions. Take, for example, a situation where a driver turns left at an intersection and overlooks an oncoming AV. The AV can only avoid a collision by swerving into oncoming traffic itself, and has to make a decision without knowing exactly how the events will play out. Staying in its lane may result with near certainty in a severe collision, while swerving into oncoming traffic could lead to a collision with another car, or that car may be able to avoid the collision with an emergency maneuver. This example alone raises a number of questions: How should probabilistic predictions be handled? Can the AV be at all allowed to swerve into the opposite lane, committing a traffic violation? And to what extent is it permissible to put otherwise uninvolved parties at risk? A number of trolley-like studies have shed a light on human behavior in such situations, which create a frame of reference as to what we deem morally right behavior [5, 6, 7, 8, 9]. But many of the biases we show as humans in our personal judgment of specific situations are at odds with the moral foundations of our law and society. So while these studies can inform the debate, the observed behavior can't simply be copied over to AVs [10, 1].

Another aspect of ethical behavior in road traffic are trade-offs between utility and safety. They include the adaptation of the driving speed to the immediate environment of the car, determinations of when it is safe enough to overtake a cyclist or another car, and even the precise positioning of an AV within its lane. Arguably, both common and critical situations would benefit from, or even require, an ad-hoc ethical evaluation of a given situation, and a set of behavioral principles that govern the AV's actions based on this evaluation. What these evaluations should look like, and what behavioral principles should govern the self-driving vehicles' behavior, however, is an unsolved problem to date.

2. The Case for Ethical Decision Making in AVs

Before we discuss the central issues for the design of ethical decision making systems in AVs, we will motivate this debate by laying out when and

why such systems become a necessity.

2.1. Levels of Automation

At what point ethical decision making becomes relevant for AVs can be deduced from the Society of Automotive Engineers’s (SAE) categorization in six levels of automation, ranging from no automation (level 0) to fully automated (level 5) [11]. On the first three levels, 0-2, the task of monitoring the environment and responding to events whenever necessary lies with the human driver. Level 1 and 2 systems are consequently categorized as driver support features. While these can include limited response to critical situations, i.e., emergency braking, they are not designed to be relied on, and the driver is supposed to intervene whenever necessary to guarantee safe operation of the vehicle. The responsibility to make ethical decisions, therefore, still lies with the driver. From level 3 on, the task of monitoring the environment, as well as performing collision avoidance, lies with the system. Such systems are classified as Automated Driving Systems (ADS), and no longer require the driver, or user, to be attentive to the road. While on level 3 the driver needs to be fallback-ready, “sufficient time for a typical person to respond appropriately”[11] has to be provided when issuing a request to intervene. Since this will take at least 5 seconds for an average driver [12], it is not an option for any immediately dangerous situation. The task of any required ethical assessments and decision making is thus passed on to the ADS for levels 3 to 5.

The case can be made that a level 3 or level 4 system would not necessitate ad-hoc ethical assessments, if automated driving is restricted to specific driving domains in which critical situations that require thorough ethical evaluation are exceedingly rare. This is a valid point, provided sufficient statistical support that the driving domains present highly predictable low risk scenarios. For most country roads, and in particular for all suburban and urban areas, however, such exceptions seem out of the question. Finally, level 5 is characterized by an ability to operate the car under all circumstances and in all domains in which a human driver could operate it, and consequently can not be exempt based on the available driving domains.

2.2. Need for Ethical Decision Making in AVs

A common argument against the need for ethical decision making systems in AVs is that these will rarely cause any accidents, with hopes of up to a 1000-fold reduction in accident frequency [13]. To see if this claim has merit,

we can adduce the safety levels of both AV prototypes and currently available level-2 systems as a frame of reference.

Favarò et al. [14] report on accident statistics by the California Department of Motor Vehicles, showing that AV prototypes between 2014 and 2017 caused accidents at a similar rate as human drivers. Notably, unexpected braking by the prototypes lead to a very high number of rear-end collisions, driving up the total rate of accident involvement 10-fold. With respect to safety in commercial releases, Tesla set a benchmark with their late 2015 release of Autopilot, a level-2 ADAS restricted to highway use. Unfortunately, reliable data on the standard of safety of the Tesla Autopilot system is difficult to obtain. Statistics provided by the US Insurance Institute for Highway Safety were analyzed by Marshall [15] and Thomas [16], who estimate that Teslas equipped with Autopilot show a slight reduction in accident rates of up to 35% (~ 1.5 -fold). However, the data did not differentiate between manual driving and Autopilot mode, and the numbers may be confounded with other safety-improvements.

These data suggest that, at least for the first generation of cars equipped with level 3 systems or higher, safety levels will be far from the goal of a 1000-fold reduction compared to human drivers. Moreover, there is an incentive not wait for the technology to become orders of magnitude safer than human drivers before it is introduced to the market: The RAND Corporation [17] estimates that even an introduction with safety levels on par with that of humans could save thousands of lives in the long run, due to a faster advancement of the technology when large amounts of data are available sooner. We thus find it reasonable to assume that AVs' safety levels will be comparable to, or only moderately better than human drivers, at the time of their introduction.

Importantly, even a large safety-improvement in AVs would not keep them from being involved in accidents caused by other road users. Without an ad-hoc evaluation of the risk factors involved in a critical situation, any AV would have to be programmed to strictly stay in its lane and to use emergency braking as the only means to react to a dangerous situation. A balancing of multiple sources of risk, i.e., risk management [18], and an ethical evaluation of the involved factors will often be necessary to sensibly avoid or mitigate collisions, regardless of who is at fault for the situation. We thus argue in line with Goodall [18] that ethical assessments and appropriate risk management algorithms should be regarded a necessity for AVs, and qualify to be mandated by legislation.

2.3. Safety Potential of Ethical Decision Making Systems

This leads us to the question how often AVs can be expected to end up in a situation where the described mechanisms would noticeably improve the outcome of a critical situation. The German Federal Office of Statistics annually publishes detailed accounts of crash statistics for Germany [19]. In 2017, a total of 2.6 million accident reports were filed, with roughly 300.000 (12%) of them entailing at least minor injuries to one of the involved persons. The latter are assessed and analyzed in more detail in the report. Based on these numbers, we make a projection with respect to safety-benefits of ethical decision making systems for two different future scenarios. The first scenario is set in the medium term future and assumes 5% of all kilometers to be driven by AVs, and AV accident rates 5 times lower than humans¹. The second scenario is set further in the future, with 50% of all kilometers driven by AVs, accident rates among human drivers cut in half, and AV accident rates 1000 times lower than those of human drivers in 2017. Across all kinds of accidents, we assumed that in any accident with AV involvement there is a 10% chance that an ethical assessment of the situation would allow the AV to change its trajectory to achieve a “better” outcome of the situation. This could mean a less severe collision with reduced bodily harm to those involved, or a fairer outcome from an ethical point of view. We further assume that in a quarter of these cases, a collision can be prevented altogether.

The results are summarized in table 1, and all estimations are laid out in detail in appendix 7. Most importantly, the mid-term future scenario already yields an annual 12,920 accidents with human injuries in which AVs are involved. Among these are an estimated 1,163 cases in which the outcome could at least be improved, and 293 accidents that could be avoided altogether. In the far future scenario, these numbers increase by a factor of 2.7. Keeping in mind that these are estimates for Germany alone, this outlines an enormous potential for improved safety and more ethical, i.e., “fairer” behavior, when AVs are equipped with systems performing real time ethical evaluations as part of their decision making process.

¹These numbers were contrived exemplarily to provide conservative estimates for the expected number of accidents with AV involvement.

Scenario	Total	AV Involved	Caused by AV	Outcome Improved	Prevented
2017	302,656	-	-	-	-
near future	291,247	12,920	2,756	1,163	293
far future	85,473	34,009	143	3,163	791

Tabell 1: Annual number of accidents with personal injuries (total, with AV involvement, and caused by an AV) in the estimated near and far future scenarios (5% and 50% of all kilometers driven by AVs, respectively.) *Outcome Improved* and *Prevented* refer to the assumed safety potential of ethical decision making systems. Table constitutes summary of the tables found in appendix 7.

3. Towards a Framework for Ethical Decision Making in AVs

Despite the need for ethical decision making systems in AVs, there is no suitable framework to date that could serve as the basis for a practical solution. We will, therefore, review existing guidelines and demands with regard to their theoretical and practical feasibility, and based on this compile a list of 10 requirements and open questions for the implementation and regulation of ethics in AVs.

A precedent with respect to fundamental moral rules is set by the report of the ethics commission for automated and connected driving, convened by the German Federal Ministry of Transport[1]. As it is the first document of its kind issued by a federal ministry, it stands as a point of orientation for other lawmakers. Although their views aren’t legally binding, we are not aware of any contending ethical statements from official sides at the time of writing, and, therefore, treat these demands as a starting point in this discussion.

3.1. Need for Explicit Regulation

Existing regulations touching upon ethical aspects of certain traffic situations, which were conceived to regulate the traffic behavior of human drivers, cannot be applied in the same way to AVs. As a case in point, the German ethics commission states that “[a]lthough a human driver would act illegally if they killed a person in an emergency to save one or more other people, he would not necessarily be held guilty. Such judgments, employed in retrospect for special circumstances cannot be converted without further ado into abstract-general ex-ante judgments and thus also not into appropriate programming”[1]. In other words, while a human driver can be forgiven after making a questionable decision in an emergency situation, the same reasoning cannot be applied to AVs. An AV makes trajectory decisions based on information about its own state and its immediate environment, collected and

combined into a world map by its perception module [20]. It applies a pre-defined set of behavioral principles, specifying the car’s actions as a function of the perceived state, the desired route, and other factors. These behavioral principles also define how the AV would react in an emergency situation. Since they are pre-defined in the programming, any morally reprehensible decision would have to be construed as premeditated. Moreover, since these behavioral principles are known at the time of the AV’s certification, it would be difficult to argue any wrongdoing on the side of the manufacturers after the fact for behavior that arises from these principles. By virtue of the precautionary principle, regulatory bodies thus have an obligation to verify that the implemented behavioral principles are morally acceptable. Consequently, in-place regulations conceived for human drivers do not suffice and the regulations for automated driving need to specify explicitly what constitutes morally acceptable behavior² (§1).

3.2. *Transparency*

An important requirement of the ethics commission’s report is for the public to be adequately informed about the decision making systems [1] (§2). Birnbacher and Birnbacher [2] take the same line, arguing that ethical decisions should under no circumstances be made by black box algorithms, and that the implemented moral norms should ideally be understood and shared by everybody. Although the demand for transparency appears natural, its implications should not be underestimated. The precise decision making process of human drivers in critical situations is concealed by severe time constraints, and their ethical assessment of the situation is usually not precisely explicable, nor put under scrutiny. As a consequence, human drivers can proclaim to generally act according to specific ethical principles (e.g., non-discrimination with respect to personal features) and still break these principles in a critical situation. A transparent decision making process, however, requires and enforces congruence of principles and actions. The need for an explicit algorithmic implementation of ethical principles may thus point out a number of inconvenient questions.

Furthermore, after any incident, the decision making system must allow to objectively trace and explicate why a specific behavior was triggered. The

²A more in-depth look into options for the regulation and certification of ethical decision making systems is given in appendix 7.

demand for transparency, therefore, not only concerns the public communication of the used approach, but requires the approach itself to be principally comprehensible by laypersons.

3.3. Public Consensus and Personalized Ethics Settings

Getting the ethical programming of an AV to reflect a public consensus may be difficult to achieve, as many people’s intuitive values are at odds with existing legislation, or established legal precedence. A potential victim’s age, for example, plays a major role to many people across all cultural backgrounds [5, 8, 9, 6, 7], but a differentiation by age or other individual factors is prohibited in most jurisdictions. For example, article 3 section 3 of the German constitution states “No person shall be favoured or disfavoured because of sex, parentage, race, language, homeland and origin, faith or religious or political opinions. No person shall be disfavoured because of disability” [21]. Consequently, any attempt to differentiate or discriminate is strongly discouraged as well by the German ethics commission [1] (§3). Obviously, the programming of AVs must adhere to the constitutional rights of the respective jurisdiction.

Further complicating the goal of reflecting a public consensus in the AVs’ programming is a significant variability in moral values even between individuals of the same culture [5, 7] (§4). One way to address this issue could be to use a parametric approach, aiming to express a person’s moral values numerically. Such numerical representations could be assessed in forced choice decision tasks [5, 6], and the population’s central tendency could be used in AVs to represent the most common ethical preferences of a society.

Alternatively, personalized ethics settings could allow the AV’s owners to set up their cars to match their own ethical preferences, but this option is not without controversy [22, 23, 24, 2]. While much of this debate is focused on an optional over-protection of the AV’s occupants, the option of allowing other ethical settings to be personalized has not gained much attention so far. A potential argument in favor of personalized ethics settings comes from research on forecasting algorithms: “Algorithm aversion” describes the tendency of people to choose human forecasters over algorithmic ones in all kinds of domains, even if they saw the algorithm perform significantly better than the human forecaster [25]. Dietvorst et al. [26] showed that algorithm aversion can be mitigated or avoided when people are given a chance to modify the algorithm’s prediction, even if the changes they can make are very limited. Whether or not such findings translate to the field of automated ethics is

an open, yet very interesting question. Allowing for limited personalization of certain ethics settings could potentially increase people’s trust in, and satisfaction with algorithmic moral decisions, and might improve the adoption rates of AVs. While we are agnostic to the debate of personalized ethics settings vs. mandatory ethics settings at this point, it should be noted that the option of personalizing ethical preferences could be viewed as an advantage or even requirement for ethical decision making algorithms in the future.

4. Behavioral Principles in Dilemma Situations

The German ethics commission’s report states important ethical rules for the behavior of automated vehicles in emergency situations, such as the protection of humans before all other considerations of utility, legal interests, and safety of animals and property [1]. It further states that “[a] set-off of victims is prohibited. A general programming to reduce the number of personal injuries can be justifiable. Those involved in the generation of mobility risks must not sacrifice uninvolved persons.” The ethics commission thus appears to allow some form of utilitarian decision making, while insisting on a number of strict prohibitions stemming from a deontologically shaped understanding of morality. Unfortunately, most of these demands don’t allow for a straightforward derivation of behavioral principles and are therefore unsuitable for a framework of automated ethical decision making. In the following, we will establish a set of important principles of ethical decision making in AVs, outline based on these where the aforementioned demands by the ethics commission are inadequate, and discuss alternatives.

4.1. Probabilistic Data and Robustness

The state of the vehicle and its immediate environment, as well as predicted future events will always be probabilistic in nature. The current state is reconstructed from sensor data, which always introduce some inaccuracies. The exact position and speed, for example of other cars or pedestrians, can only be estimated and will always carry some uncertainty. Even more so, future events can never be foreseen with absolute precision, and critical situations can be particularly difficult to predict. The behavior of other actors is often volatile, and the magnitude of a collision in terms of injuries can often only be coarsely approximated. Thus, ethical decision making systems must be able to operate on probabilistic data (§5).

As a corollary of the probabilistic nature of the available data for decision making, the robustness of the ethical evaluations must be ensured. By robustness we mean that small changes in the situation should not lead to fundamental changes in the ethical evaluation. This is important for two reasons: First, ethical evaluations have to be repeated in short intervals. In an emergency situation, it is important that the ethical assessment doesn't jump back and forth between fundamentally different states, as this could cause volatile and unpredictable behavior by the vehicle. Second, if the ethical evaluations change fundamentally based on minimal differences between near identical situations, we lose transparency in the decision making process, violating the previously set demand. Thus, similar situations should lead to similar ethical evaluations (§6.1).

4.2. Categorical Distinctions, Thresholding, and Reasonable Decisions

The probabilistic nature of predictions implies that we often cannot know with certainty whether a person would die in a collision or only be injured. As a consequence, there is no straightforward solution to a categorical distinction between injuries and fatalities in the decision making logic. In the same way, we often cannot know whether or not a person would be injured in an accident, calling a categorical treatment of this variable into question as well. This becomes problematic when regulations or moral rules presume precisely these categorical distinctions, as is the case in the prohibition of all sacrifices of persons not involved for the greater good. Leaving aside for now the problematic conception of who is involved and who isn't, we have two choices: Either we refrain from all actions that pose the slightest risk of death to those not involved, or we define a threshold for the amount of risk above which the decision would be defined as a sacrifice – both of which are problematic.

Refraining from even the slightest risk is often incompatible with a notion of reasonableness. Let's say we need to swerve onto the sidewalk in order to save a pedestrian's life, and there is a 99% probability that no-one will be harmed at all, but a 1% chance that a second pedestrian on the sidewalk will be hit and killed in the process. We believe that it would be unreasonable not to take this small chance if it means saving another person with certainty. Yet if we deem even miniscule risks to bystanders unacceptable, this option is precluded from the start. The exact point at which it becomes unreasonable not to save the person at risk would need to be debated, but most will agree that at *some* point the absolute protection of bystanders from any risk will

become unreasonable. We believe that a notion of reasonableness should be upheld whenever possible (§6.2).

Introducing thresholds that define at which level of risk a decision is considered a sacrifice could recover some notion of reasonableness, but only at the cost of robustness and transparency of the decisions. When the probability of a critical aspect is perceived near a defined threshold, smallest changes in the perception can fundamentally change the system’s ethical evaluation of the situation, leading to the previously discussed lack of robustness. It would also be difficult to justify any precise threshold value. Most ethically relevant aspects of emergency situations exist on a continuous spectrum, and an artificial dichotomization would often not do them justice. Any precise threshold defining who counts as involved in a critical situation, or what constitutes a sacrifice, would be entirely arbitrary.

A possible solution to unite robustness of the decision making logic and reasonableness of the resulting decision would be to conceptualize ethically relevant properties on a continuous scale, and treat moral rules as soft constraints to the car’s behavior. This would allow for a compromise between deontological and utilitarian considerations. The system would principally base its decisions on a comparison of the stakes involved for different parties in a situation, but could additionally disincentivise against the violation of important moral rules, as well as traffic violations. We, therefore, believe that an insistence on categorical distinctions and absolute prohibitions is not justified at the cost of reasonableness, robustness, and transparency in the decisions made. We do concede, however, that this is a possible point of contention and thus encourage the debate of this issue. Ultimately, a fundamental decision has to be made between insisting on strict categorizations and absolute prohibitions on the one hand, and making reasonable and robust decisions on the other (§7).

4.3. A Single Scale of Ethical Cost

The notion of reasonableness and the issues relating to categorical distinctions highlight another aspect that needs to be considered: Given that small or miniscule risks of human injury are often present in emergency situations and collision avoidance, it seems unreasonable to treat these as categorically superseding any risk to animals or material damages. For example, if a pedestrian steps in front of the car unexpectedly, and the car estimates a 10% chance of having a light collision with that pedestrian (speed at impact not exceeding 5 kph), should it take all risk from the pedestrian by swerving and

crashing into a lamp post instead? Such behavior may be perceived by many as disproportionate. To prevent it, the well-being of humans would have to be set off with material and other ethical cost, such as the well-being of animals, at least so long as the risk of human injury is low. To assert that human well-being is still prioritized over all else, its valuation in the decision making logic would have to clearly exceed that of all other ethically relevant factors. The value of animals and material damages could also be capped to ensure that any severe or fatal injuries to humans supersede any amount of entirely materialistic damage or risk to animals. So long as severe injuries are very unlikely or precluded altogether, careful balancing of the involved cost factors could provide a solution that avoids unreasonable or disproportionate decisions. We thus believe it is justified to consider and debate the option of offsetting human well-being with other considerations on a single scale of ethical cost (§8).

4.4. Involvement, Responsibility, and Protection of AV Occupants

To address the question who should be considered involved in an emergency situation, and who would be protected as an innocent bystander, Birnbacher and Birnbacher [2] contend that from a moral point of view it is not relevant whether or not those in danger are in the movement direction of the car, since, unless otherwise specified, everyone in danger is equally innocent. The authors further argue that even if we allow a distinction between active “harming” and passive “letting harm occur”, this distinction does not apply to AVs, since any behavior by the car would be active behavior, as pre-defined by the programmer. We agree with the authors in that an AV has no logically defined mode of “not behaving” while in motion, and all of its behavior can be regarded active. Those potentially in danger, however, arguably differ with respect to the amount of risk to themselves that they consent to, by means of their location and behavior. A pedestrian jaywalking generally consents to a larger amount of risk than a pedestrian crossing the street in an orderly manner, who, in turn, still consents to a greater level of risk than a person standing on the sidewalk. We suggest that a distinction be made with respect to the location of those who are potentially at risk (§9), i.e., the drivable surface could be divided into zones of differential protection to those within them. For example, sidewalks could be declared protected zones, drastically disincentivizing AVs from swerving in the direction of pedestrians.

With respect to assessing the responsibility of any involved party for causing a dangerous situation, there is an approach that lends itself to this

possibility. The Responsibility-Sensitive Safety model (RSS) [13] is based on common-sense notions of accountability in traffic, and essentially formalizes a blame metric. While it was conceived to guide an AV’s behavior such that it is never to blame for an accident, its assignments of blame could in principal be used to bring a notion of fairness into an AV’s ethical evaluation of critical situations. While the technical and legal feasibility of such an approach would have to be examined in detail, and potential ethical issues would still have to be discussed, considering the amount of responsibility for a dangerous situation in the decision making logic could help to increase the perceived fairness of ethical decision making systems in AVs.

Either way, the general idea of differentiating by location and accountability of the involved parties could offer an elegant solution to another fiercely debated question [23, 22] – whether or not an AV should be allowed to provide special protection to its occupants at the cost of other road users. When other road users are getting protective advantages based on their location and possibly their (lack of) liability, one can discuss an appropriate protective classification of AV passengers. These are passive and thus never directly to blame for causing a dangerous situation. Yet by using the AV, they consent to some amount of mobility risk linked to the high speeds of the vehicle in comparison to pedestrians. Their protective classification would, therefore, sensibly lie between that of those at fault for causing the dangerous situation and that of bystanders on the sidewalk.

4.5. Breaking the Rules

Many scenarios in which smart collision avoidance and crash behavior can significantly improve the outcome of critical situations would require an AV to violate traffic regulations. Jumping a red light may be necessary to avoid being rear-ended, and swerving off the road, onto a sidewalk or into the opposite lane could sometimes prevent a head-on collision. We believe that it is unreasonable and even morally reprehensible not to avoid a collision if we know the risk to third parties to be minimal or non-existent.

Allowing for traffic violations would probably already be in line with the ethics commission’s demands, which state that the protection of humans has to take precedence over other legal interests [1]. Beyond this, it could be viewed as sensible to allow for traffic violations also to prevent material damages, if no considerable risk to others is created in the process. Similar to moral rules, traffic regulations could be implemented as soft constraints, assigning

a cost to certain violations that can be outweighed by defined amounts of risk-minimization.

Ultimately, traffic regulations serve the purpose of improving safety, and not standing in its way. We thus argue that responsibly breaking traffic regulations to benefit overall safety is not only acceptable, but a moral imperative when the circumstances demand it. AVs should thus be allowed or even encouraged by regulation to break certain rules if the car is equipped to make the respective decisions *sensibly* (§10).

4.6. Synopsis

In summary, the above discussion shows a number of severe limitations of the guidelines by the German ethics commission [1], and gives rise to the following list of demands and open questions that need to be addressed in the pursuit of a solid framework for algorithmic ethical decision making in AVs:

1. An AVs ethical behavior needs to be regulated. A precise definition of acceptable behavioral principles for AVs in critical situations is necessary, and a verification that these principles are upheld by the software needs to be part of the certification procedure for AVs.
2. The behavioral principles implemented in decision making systems of AVs must be transparent, traceable, publicly communicated, and be principally comprehensible by laypersons, precluding black-box algorithms.
3. Personal characteristics of potentially involved persons, such as their age or gender, must not be considered in the decision making procedure.
4. It needs to be determined, to which degree behavioral principles in dilemma situations can and should reflect a public consensus, and whether or not personalized ethics settings, possibly restricted to a defined range, are an option.
5. All behavioral rules and systems that implement these must do justice to uncertainties in perception and the probabilistic nature of predictions.

6. Ethical evaluations should be robust, i.e., similar situations should result in similar ethical evaluations, and a notion of reasonableness in the decisions should be upheld whenever possible.
7. It needs to be debated, to what extent an insistence on categorical distinctions with absolute prohibitions can be reconciled with robust, transparent, and reasonable decisions in a probabilistic framework.
8. The protection of human life needs to take priority over that of animals or property. It should be discussed if and to what extent a set-off of human well-being with other ethically relevant factors is acceptable when it serves to avoid disproportionate decisions.
9. In the decision making logic, a distinction needs to be made with respect to the location of a person at risk. It should also be debated, whether the amount of responsibility for the dangerous situation can or should be factored in as well.
10. AVs must be allowed to violate certain traffic regulations if the car is equipped to make the respective decisions sensibly.

5. Approaches Towards a Solution

With respect to concrete solutions, some existing approaches generally comply with the stated demands. van Otterlo [27] proposes a utility function accumulating all costs associated with each trajectory option, and selecting the one with the lowest cost. The costs per trajectory option are a sum of the “ethical costs” per killed obstacle and other consequences of choosing a given trajectory. However, this approach does not explicitly consider the probabilistic basis for the decisions. Addressing this aspect, Goodall [18] suggests that AVs could perform risk analysis by calculating expectation values, i.e., the product of the probability of an unwanted event, and its magnitude. An event with a 10% chance of occurring, which would cost two lives would thus have an expectation value of 0.2. The process of balancing multiple sources of risk, or multiple expectation values, called risk management, would then be the basis for the decision making process. We believe that a combination of utility functions and risk management could be conceived that fulfills all of the previously established framework conditions. However, any concrete approach would have to pick a side on the open questions outlined above.

5.1. Examples of Dilemma Situations

To provide a more illustrative account of some of the issues outlined above, we will exemplify them with two hypothetical scenarios.

Scenario 1

An AV is driving in city traffic with two passengers on board, doing the allowed 50 kph. The road has a single lane in each direction and sidewalk to the right of the car. Without checking for traffic, a pedestrian (pedestrian 1) starts crossing the road a few meters in front of the car, and the AV is left with three options:

1. Initiate an emergency breaking maneuver without swerving. The car's perception module estimates an 80% chance of a collision, and the collision would likely kill the pedestrian. The AV's passenger would not be harmed.
2. Swerve on the sidewalk to the right and initiate emergency breaking. A second pedestrian (pedestrian 2) is on the sidewalk a little further away. The perception module estimates a 70% chance of a collision with the second pedestrian. Due to the increased distance, and thus reduced speed at impact, the collision is estimated to cause severe injuries, but unlikely to kill them. The AV's passenger would not be harmed.
3. Swerve to the left into oncoming traffic and initiate emergency breaking. A car with a single driver and no further passengers would be hit with near certainty (95%), likely resulting in minor injuries for all passengers in the two cars.

Scenario 2

An AV with one passenger on board is stopped at a red light on a suburban intersection. The car behind the AV with one person on board initiates a late emergency braking maneuver, and a rear-end collision is probable. The intersection is empty with one car approaching from the left (two occupants). This car has a green light, but would probably be able to swerve or come to a stop in time if necessary. The AV has the following options:

1. Stay put with a 70% chance of being rear-ended. The AV predicts minor injuries to passengers of both vehicles if the collision happens, while the car approaching from the left would be unaffected.

2. Jump the red light to give the car behind enough space to come to a stop. The AV calculates a 10% chance of colliding with the car approaching from the left, which would likely result in minor injuries for the other car’s passengers, and moderate injuries to the AV’s passenger.

These examples illustrate the probabilistic nature of the situations in which the AV needs to make a decision, and show that a large number of different factors need to be evaluated at the same time. These include the well-being of multiple parties, material damages, traffic regulations, and ethical considerations. It is easy to further elaborate these scenarios, e.g., involving animals, persons of different age, or unknown numbers of passengers in a car or bus. This shows that the complexity of real world decisions more often than not touches upon a majority of the above ten requirements and open questions.

6. Trajectory Choices

Finding appropriate solutions to these issues would go a long way in creating a framework for the regulation and implementation of ethical decision making systems. However, two important practical considerations have not yet been addressed: (1) Trajectory-selection in continuous space instead of a low number of distinct options, and (2) the option of optimizing trajectories in regular driving.

6.1. *Evaluations of Continuous Space*

In most examples of traffic dilemmas, the scenarios are presented in a trolley-like fashion with typically two or three very distinct behavioral options for the driver or vehicle. While this makes it easier for humans to get a grasp of the situation and discuss it, this view may be rather detached from how path planning and decision making works in reality. Between every two distinct behavioral options, there is technically an infinite number of trajectories the car could take, and it is unclear how it should decide, which of these to evaluate in detail. Thus, the problem is not only how to decide, but already which options to investigate for a later decision.

Instead of evaluating a low number of trajectories representing fundamentally different choices, one could integrate the ethical evaluation with the car’s internal map of its immediate surrounding to achieve an ethical evaluation of continuous space. This map, in Gruyer et al. [20] referred to as

the local dynamic perception map (LDPM), represents the vehicle’s position and dynamic state, lanes and surfaces, as well as possible obstacles and their predicted trajectories³. Based on the LDPM, a map of the car’s environment could be created, in which each point in space (or pixel) is assigned a value that answers the question "how good or bad would it be to choose a trajectory that leads the car over this point in space?". Figure 1 gives an intuition of how such an approach could work: (A) is a depiction of the scene in the real world, with the blue car being the AV. In its perception module it detects a number of features of its surrounding scene, such as relevant zones (lanes, walkways, etc.), depicted in (B), and relevant objects (pedestrians, other cars, etc.), depicted in (C). Since many of the relevant objects in (C) are moving, a prediction of their future locations is computed and indicated as probability densities. In the ethical evaluation step, the detected zones and locations of relevant objects would be transformed into disincentives to move into a given space. (D) shows a small disincentive to move into the opposite lane, and moderate disincentives to move onto the sidewalk. (E) shows strong disincentives to move into spaces predicted to be occupied by other objects, in particular by pedestrians on the sidewalk. Finally, (F) is a superimposition of D, and E, and could be used to precisely plan the vehicle’s trajectory. Note that the illustration doesn’t cover all contents of an LDPM, nor all aspects relevant to the ethical evaluation of the scene, and disregards the passing of time for simplicity.

6.2. Trajectory Optimization in Regular Driving

An ethical evaluation of continuous spaces would pave the way for yet another option: Optimizing the vehicle’s trajectory not just in critical situations, but permanently. Since there are no precise boundaries between careful maneuvering, collision avoidance, and crash behavior, defining a precise switch point at which the ethical decision making system takes over would appear arbitrary and might leave safety potential unused. Instead, a constant ethical assessments for the space around them would allow for a sensible adaptation of their driving speed, let them assess when it is safe enough to overtake a slower vehicle or a cyclist, and position themselves in their lane in a way that always minimizes the risk of a collision.

³A more detailed account of the perception and control architecture in AVs is given in appendix 7.

Figure 1: Sketch: Ethical evaluations in continuous space. A: Depiction of a scene in the real world. B: Detected surfaces (road, lanes, sidewalks, etc.) as part of the LDPM. C: Detected objects, color-coded by type, and corresponding movement prediction (probability densities) as part of the LDPM. D: Ethical evaluation of drivable surfaces: Deeper reds indicate higher ethical cost (stronger disincentives) to drive in that region. E: Ethical evaluation of objects based on their predicted locations, i.e., probability densities weighted by vulnerability and location. F: Full ethical evaluation of the scene (superimposition of D and E).

Even in the absence of any immediate and recognized hazards, occluded areas, such as the space behind parked vehicles on the side of the road, could be considered individual sources of risk, creating small disincentives against driving past them at a very close distance. Under certain circumstances, their risk cost could be treated as elevated, for example when other parts of the sidewalk are visibly crowded with pedestrians. In a narrow street with parked cars on the side but no other traffic, the car would then move closer to the center in order to leave more room for possibly occluded pedestrians, and more time for itself to avoid a collision.

However, since very small amounts of risk are present in almost any situation, the decision making system would have to go beyond a mere minimization of risk to a process of counterbalancing risk and utility. Without

accounting for utility, the system would often be likely to stop the car immediately, since not moving at all is the safest option in most cases. The aspects of utility that would be considered could include economy of time – creating an incentive to keep the car moving –, passenger comfort, energy consumption and conformity with traffic culture. The full integration of these aspects of utility could lead to ethical decision making systems covering every aspect of driving.

7. Conclusion

In this work, we analyzed when and why ethical decision making systems become necessary in AVs, we reviewed existing guidelines for the behavior of AVs in dilemma situations, and we compiled a set of 10 demands and open questions that need to be addressed in the pursuit of a framework for ethical decision making in AVs. The principal issues outlined in this work show that neither the regulations for algorithmic ethical decisions in AVs, nor a practical implementation of such a system can be conceived independently of one another. Any tangible solution for automated ethics needs to reconcile fundamental law, technical feasibility, and the moral values of the society. The goal must be to not only allow the cars to be as fair and safe as possible, and guard the manufacturers from legal liability with respect to the ethical programming of their cars, but also to foster public trust in automated driving technology, and facilitate its adoption. It is, therefore, important that regulatory bodies and the industry engage in a dialogue to find an appropriate solution. At the same time, the central ethical issues need to be discussed publicly, and we emphatically encourage this debate.

Acknowledgments

The authors would like to thank Ane Dalsnes Storsæter for her valuable input and collaboration in early stages of this project.

Appendix

Projected Accident Numbers Involving AVs in Germany

Estimations in detail: Tables 2 and 3.

Options for Regulation and Certification

Most ADSs currently in development use a perception and control architecture referred to as the sense-plan-act control methodology [13]. This architecture consists of three main stages, as detailed in Gruyer et al. [20]:

1. The *perception stage* collects, filters and processes data from all sensors and systems, as well as information from external sources such as other vehicles, road infrastructure, and maps. Functionally, the tasks of the perception system are to detect the road, lanes, relevant features of the road and obstacles, detect or sense relevant information about the environment, such as weather, road signs, etc., and to correctly apprehend the vehicle’s position and dynamic state. This also includes behavior identification and trajectory prediction for moving objects in the scene. The information derived from various sources is then combined in a local map of the car’s immediate surrounding, in Gruyer et al. [20] referred to as the local dynamic perception map (LDPM).
2. The *planning and decision-making stage* uses the LDPM and route information to select the car’s precise trajectory and make tactical decisions, such as the desired speed, acceleration or deceleration rates. Ad-hoc risk analysis, risk management and ethical assessments, and the subsequent decision making would happen within this module.
3. Finally, the *control stage* issues precise control orders to the steering system, engine or motors, and brakes, that ensure that the trajectory and other tactical decisions, decided on in the planning stage, are effectuated by the vehicle.

This three-tiered structure follows a logical perception-to-action sequence, and provides a leverage point for regulation, allowing the individual sub-tasks, i.e., perception, planning, and action, to be verified and regulated separately. With respect to ethical aspects of the car’s behavior, it allows for the isolation of different factors from one another that could no longer be disentangled when only verifying the system’s behavior as a whole. For example, finding the car to behave more cautious around smaller people could mean that this group is more easily detected in the perception stage, or it could hint towards a bias favoring women and children (who are smaller on average), hard-coded in the car’s value system. Separate regulation and

certification of the three stages would thus allow for more precise oversight of the individual modules' capabilities and programming.

Birnbacher and Birnbacher [2] plead for unified decision making algorithms across manufacturers, and modular certification of an AV's perception, planning, and control stages. In particular, the authors suggest unified interfaces between these modules, which would facilitate the certification and verification of these parts of the software. As part of the certification procedure, the planning and decision making module could be given a large number of different scenarios in the form of LDPMs, or be put in a simulation environment producing LDPMs, and would return trajectory commands in a fixed format. Regulators could then test if the desired outcomes are produced in any number of test cases, without the procedure being complicated by inaccuracies or mistakes originating in the perception and control modules. On the downside, using unified interfaces could mean using LDPMs with a fixed feature set. As technology improves and the perception modules get better, this could constitute a bottleneck, preventing information to be passed on between the individual modules.

- [1] C. Lütge, The german ethics code for automated and connected driving, *Philosophy & Technology* (2017).
- [2] D. Birnbacher, W. Birnbacher, Automatisiertes fahren. ethische fragen an der schnittstelle von technik und gesellschaft, *Information Philosophie* 4 (2016) 8–15.
- [3] T. Holstein, G. Dodig-Crnkovic, Avoiding the intrinsic unfairness of the trolley problem (2018).
- [4] B. Wilson, J. Hoffman, J. Morgenstern, Predictive inequity in object detection, *arXiv preprint arXiv:1902.11097* (2019).
- [5] E. Awad, S. Dsouza, R. Kim, J. Schulz, J. Henrich, A. Shariff, J.-F. Bonnefon, I. Rahwan, The moral machine experiment, *Nature* 563 (2018) 59.
- [6] L. R. Sütfeld, R. Gast, P. König, G. Pipa, Using virtual reality to assess ethical decisions in road traffic scenarios: applicability of value-of-life-based models and influences of time pressure, *Frontiers in behavioral neuroscience* 11 (2017) 122.

- [7] L. R. Sütfeld, B. V. Ehinger, P. König, G. Pipa, How does the method change what we measure? comparing virtual reality and text-based surveys for the assessment of moral decisions in traffic dilemmas, 2019.
- [8] L. T. Bergmann, L. Schlicht, C. Meixner, P. König, G. Pipa, S. Boshammer, A. Stephan, Autonomous vehicles require socio-political acceptance—an empirical and philosophical perspective on the problem of moral decision making, *Frontiers in behavioral neuroscience* 12 (2018) 31.
- [9] A. K. Faulhaber, A. Dittmer, F. Blind, M. A. Wächter, S. Timm, L. R. Sütfeld, A. Stephan, G. Pipa, P. König, Human decisions in moral dilemmas are largely described by utilitarianism: Virtual car driving study provides guidelines for autonomous driving vehicles, *Science and engineering ethics* (2018) 1–20.
- [10] L. R. Sütfeld, R. Gast, P. König, G. Pipa, Response: Commentary: Using virtual reality to assess ethical decisions in road traffic scenarios: Applicability of value-of-life-based models and influences of time pressure, *Frontiers in behavioral neuroscience* 12 (2018) 128.
- [11] S. of Automotive Engineers, SAE Standard J3016: Taxonomy and Definitions for Terms Related to On-Road Motor Vehicle Automated Driving Systems, SAE International, 2018.
- [12] A. Eriksson, N. A. Stanton, Takeover time in highly automated vehicles: noncritical transitions to and from manual control, *Human factors* 59 (2017) 689–705.
- [13] S. Shalev-Shwartz, S. Shammah, A. Shashua, On a formal model of safe and scalable self-driving cars, arXiv preprint arXiv:1708.06374 (2017).
- [14] F. M. Favarò, N. Nader, S. O. Eurich, M. Tripp, N. Varadaraju, Examining accident reports involving autonomous vehicles in california, *PLOS ONE* 12 (2017) 1–20.
- [15] A. Marshall, Tesla’s favorite autopilot safety stat just doesn’t hold up, *wired* (2018).
- [16] B. A. Thomas, A closer inspection of tesla’s autopilot safety statistics, <https://medium.com/@mc2maven/>

a-closer-inspection-of-teslas-autopilot-safety-statistics-533eebe0869d, 2018.

- [17] RAND Corporation, Why waiting for perfect autonomous vehicles may cost lives, <https://www.rand.org/blog/articles/2017/11/why-waiting-for-perfect-autonomous-vehicles-may-cost-lives.html>, 2017.
- [18] N. J. Goodall, Away from trolley problems and toward risk management, *Applied Artificial Intelligence* 30 (2016) 810–821.
- [19] S. Bundesamt, Verkehr verkehrsunfälle 2017, 2018.
- [20] D. Gruyer, V. Magniera, K. Hamdia, L. Claussmann, O. Orfila, A. Rakotonirainy, Perception, information processing and modeling: Critical stages for autonomous driving applications, *Annual Reviews in Control* 44 (2017) 323–341.
- [21] Basic law for the federal republic of germany, <https://www.btg-bestellservice.de/pdf/80201000.pdf>, ??? Accessed: 2019-06-11.
- [22] J.-F. Bonnefon, A. Shariff, I. Rahwan, The social dilemma of autonomous vehicles, *Science* 352 (2016) 1573–1576.
- [23] J. Gogoll, J. F. Müller, Autonomous cars: in favor of a mandatory ethics setting, *Science and engineering ethics* 23 (2017) 681–700.
- [24] P. Lin, Here’s a terrible idea: robot cars with adjustable ethics settings, *Wired. com*. Available via <http://www.wired.com/2014/08/heres-a-terrible-idea-robot-cars-with-adjustable-ethics-settings> (2014).
- [25] B. J. Dietvorst, J. P. Simmons, C. Massey, Algorithm aversion: People erroneously avoid algorithms after seeing them err., *Journal of Experimental Psychology: General* 144 (2015) 114.
- [26] B. J. Dietvorst, J. P. Simmons, C. Massey, Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them, *Management Science* 64 (2016) 1155–1170.

- [27] M. van Otterlo, From algorithmic black boxes to adaptive white boxes: Declarative decision-theoretic ethical programs as codes of ethics, arXiv preprint arXiv:1711.06035 (2017).

accidents 2017	party at fault	other party involved	fraction of cases	impr. factor	accidents in scenario	% impr. outcome	# impr. outcome	% prev. collision	# prev. collision
(1) Collision with another vehicle which starts, stops or is stationary ("hitting a parked car")									
18,972	human driver	human driver	0.9025	1	17,122	0	0	0	0
	human driver	AV	0.0475	1	901	0	0	0	0
	AV	human driver	0.0475	5	180	0.1	18	0.25	5
	AV	AV	0.0025	5	9	0.1	1	0.25	0
(2) Collision with another vehicle moving ahead or waiting ("rear-ending a car in traffic")									
54,505	human driver	human driver	0.9025	1	49,191	0	0	0	0
	human driver	AV	0.0475	1	2,589	0.1	259	0.25	65
	AV	human driver	0.0475	5	518	0.1	52	0.25	13
	AV	AV	0.0025	5	27	0.1	3	0.25	1
(3) Collision with another vehicle moving laterally in the same direction ("side-swiping")									
15,698	human driver	human driver	0.9025	1	14,167	0	0	0	0
	human driver	AV	0.0475	1	746	0.1	75	0.25	19
	AV	human driver	0.0475	5	149	0.1	15	0.25	4
	AV	AV	0.0025	5	8	0.1	1	0.25	0
(4) Collision with another oncoming vehicle ("head-on collision")									
21,893	human driver	human driver	0.9025	1	19,758	0	0	0	0
	human driver	AV	0.0475	1	1,040	0.1	104	0.25	26
	AV	human driver	0.0475	5	208	0.1	21	0.25	5
	AV	AV	0.0025	5	11	0.1	1	0.25	0
(5) Collision with another vehicle which turns into or crosses a road ("collision at intersection")									
82,542	human driver	human driver	0.9025	1	74,494	0	0	0	0
	human driver	AV	0.0475	1	3,921	0.1	392	0.25	98
	AV	human driver	0.0475	5	784	0.1	78	0.25	20
	AV	AV	0.0025	5	41	0.1	4	0.25	1
(6) Collision between vehicle and pedestrian ("hitting a pedestrian")									
28,324	human driver	pedestrian	0.475	1	13,454	0	0	0	0
	pedestrian	human driver	0.475	1	13,454	0	0	0	0
	AV	pedestrian	0.025	5	142	0.1	14	0.25	4
	pedestrian	AV	0.025	1	708	0.1	71	0.25	18
(7) Collision with an obstacle in the carriageway ("hitting an obstacle")									
2,738	human driver	obstacle	0.475	1	1,301	0	0	0	0
	obstacle	AV	0.475	1	1,301	0	0	0	0
	AV	obstacle	0.025	5	14	0.1	1	0.25	0
	obstacle	AV	0.025	1	68	0.1	7	0.25	2
(8+9) Leaving the carriageway to the right or left ("driving off the road")									
40,020	human driver	-	0.95	1	38,019	0	0	0	0
	AV	-	0.05	5	400	0	0	0	0
(10) Accident of another kind ("all others")									
37,964	human driver	other party	0.9025	1	34,263	0	0	0	0
	other party	AV	0.0475	1	1,803	0	0	0	0
	AV	other party	0.0475	5	361	0.1	36	0.25	9
	other party	AV	0.0025	1	95	0.1	10	0.25	3
302,656	total				291,247		-		-
	AV involved				12,920		1,163		293
	AV at fault				2,756		241		62

Tabell 2: Overview of estimations for scenario 1 (5% of all kms driven by AVs). *Accidents 2017* refers to the number of accidents by kind of accident (numbers from Germany, 2017). In this scenario, we assume a random selection of 5% of the involved cars to be AVs instead of conventional vehicles, creating (mostly) four constellations of *party at fault* vs. *other party involved*, each making up a certain *fraction of cases* within the specified kind of accident. The *improvement factor* is set to 5-fold reduction in accident rates for all accidents caused by AVs, and to 1 (no change) for all other parties at fault. The number of *accidents in scenario* is calculated as *accidents 2017* times *fraction of cases* divided by *improvement factor*. We make the assumption that in 10% of the remaining cases with AV involvement, an ethical assessment of the situation would allow the AV to change the trajectory to avoid the crash or mitigate the damages or bodily harm, or achieve a better outcome in terms of fairness. We further make the assumption that in a quarter of these cases, a collision can be prevented altogether.

accidents 2017	party at fault	other party involved	fraction of cases	impr. factor	accidents in scenario	% impr. outcome	# impr. outcome	% prev. collision	# prev. collision
(1) Collision with another vehicle which starts, stops or is stationary ("hitting a parked car")									
18,972	human driver	human driver	0.25	2	2,372	0	0	0	0
	human driver	AV	0.25	2	2,372	0	0	0	0
	AV	human driver	0.25	1000	5	0.1	1	0.25	0
	AV	AV	0.25	1000	5	0.1	1	0.25	0
(2) Collision with another vehicle moving ahead or waiting ("rear-ending a car in traffic")									
54,505	human driver	human driver	0.25	2	6,813	0	0	0	0
	human driver	AV	0.25	2	6,813	0.1	681	0.25	170
	AV	human driver	0.25	1000	14	0.1	1	0.25	0
	AV	AV	0.25	1000	14	0.1	1	0.25	0
(3) Collision with another vehicle moving laterally in the same direction ("side-swiping")									
15,698	human driver	human driver	0.25	2	1,962	0	0	0	0
	human driver	AV	0.25	2	1,962	0.1	196	0.25	49
	AV	human driver	0.25	1000	4	0.1	0	0.25	0
	AV	AV	0.25	1000	4	0.1	0	0.25	0
(4) Collision with another oncoming vehicle ("head-on collision")									
21,893	human driver	human driver	0.25	2	2,737	0	0	0	0
	human driver	AV	0.25	2	2,737	0.1	274	0.25	69
	AV	human driver	0.25	1000	5	0.1	1	0.25	0
	AV	AV	0.25	1000	5	0.1	1	0.25	0
(5) Collision with another vehicle which turns into or crosses a road ("collision at intersection")									
82,542	human driver	human driver	0.25	2	10,318	0	0	0	0
	human driver	AV	0.25	2	10,318	0.1	1,032	0.25	258
	AV	human driver	0.25	1000	21	0.1	2	0.25	1
	AV	AV	0.25	1000	21	0.1	2	0.25	1
(6) Collision between vehicle and pedestrian ("hitting a pedestrian")									
28,324	human driver	pedestrian	0.25	2	3,541	0	0	0	0
	pedestrian	human driver	0.25	2	3,541	0	0	0	0
	AV	pedestrian	0.25	1000	7	0.1	1	0.25	0
	pedestrian	AV	0.25	1	7,081	0.1	708	0.25	177
(7) Collision with an obstacle in the carriageway ("hitting an obstacle")									
2,738	human driver	obstacle	0.25	2	342	0	0	0	0
	obstacle	AV	0.25	2	342	0	0	0	0
	AV	obstacle	0.25	1000	1	0.1	0	0.25	0
	obstacle	AV	0.25	1	685	0.1	69	0.25	17
(8+9) Leaving the carriageway to the right or left ("driving off the road")									
40,020	human driver	-	0.5	2	10,005	0	0	0	0
	AV	-	0.5	1000	20	0	0	0	0
(10) Accident of another kind ("all others")									
37,964	human driver	other party	0.45	2	8,542	0	0	0	0
	other party	AV	0.05	2	949	0	0	0	0
	AV	other party	0.45	1000	17	0.1	2	0.25	1
	other party	AV	0.05	2	949	0.1	95	0.25	24
302,656	total				85,473		-		-
	AV involved				34,009		3,163		791
	AV at fault				143		13		3

Tabell 3: Overview of estimations for scenario 2 (50% of all kms driven by AVs). In this scenario, we assume a random selection of 50% of the involved cars to be AVs instead of conventional vehicles. The *improvement factor* is set to 1000-fold reduction in accident rates for all accidents caused by AVs, to a 2-fold reduction for human drivers, and to 1 (no change) for pedestrians at fault and obstacle collisions. Everything else is the same as in table 2.