Basic numbers needed to understand the traffic safety effect of Automated Cars

Magdalena Lindman, Irene Isaksson-Hellman, Johan Strandroth

Abstract In order to simplify the development process, most automated driving technologies employ limitations in their operational scope. Comparing automated driving performance in these specific contexts to manual driving performance in general then provide an unfair comparison. Instead, the reference should be made as specific and as precise as possible.

The objective of the present study was to identify basic and relevant numbers required to evaluate the safety performance of automated driving by identifying conflict situations for targeted traffic environments and establishing crash rates for these conflict situations. Results in terms of crash rate differences confirmed that relevant references must be specified according to the automated driving technology limitations, i.e. traffic environment, vehicle type and conflict situation. The need to distinguish between crash outcomes (fatal, severe or minor injuries, or property damage only crashes) is obvious. The number of kilometres necessary to establish future significant differences in these crash rates indicate that retrospective traffic safety evaluations will be far from realistic for many years to come.

This report provides an overview on basic numbers required to evaluate traffic safety performance of automated driving cars, a discussion on the methodology to achieve these numbers, and insights into expectations on future traffic safety evaluations of automated driving traffic safety performance.

Keywords Automated driving, Crash rate, Safety performance evaluation, Traffic accident data,

I. INTRODUCTION

As cars with automated driving (AD) are being developed, this range of technologies is often referred to as the ultimate solution to all traffic safety issues. Thus, questions are being raised on their safety potential, e.g. how safe these cars must be and whether they bring new traffic safety issues that need to be addressed. One qualifier in any attempt to answer these kinds of questions is the knowledge of relevant traffic safety reference numbers that can be used when evaluating the performance of AD technology.

Traffic safety can be evaluated using a limited number of indicators, among which the number of crashes and personal injuries are the most obvious. Crash rate represents a meaningful and comparative measure of these indicators since it reflects crash numbers in relation to a measure of exposure. Thus, crash rate can be used as a reference when evaluating effects from both injury- and crash-preventing arrangements in the traffic environment. Typical examples include crash and injury rate per kilometers travelled, per capita, per registered vehicles or per insured vehicle years, and they have proven to be useful for ranking the performance of countries, comparing safety levels of road surface conditions, or evaluating vehicle safety performance [1-6], to mention just a few examples.

In a recent study, autonomous cars crashed 3.2 times per million miles, compared with 4.2 times for human drivers [7]. In [8], a rate of 77 reported injuries and 1.09 fatalities per 100 million miles driven by human drivers in 2013 was found, based on estimates of crash rates calculated from official US data.

For more practical applications, as in the case of ongoing and planned AD research projects, a higher level of detail of such an analysis is needed. First, it is important to delimit the analysis according to the specific AD research project scope, if there are such limitations. As an example, in the DriveMe project [9], automated driving will be directed to specific roads in Sweden where no vulnerable road users (VRU´s) are allowed and where lanes with oncoming traffic are separated, which will narrow down the number of traffic situations to include in the analysis. Secondly, different crash severity levels have to be investigated to ensure that all desired

M. Lindman (e-mail: magdalena.lindman@volvocars.com; tel: +46 31 3256257) is Technical Expert at Volvo Cars Safety Centre. I. Isaksson-Hellman is Research Leader for Traffic Safety at If P&C Insurance. Johan Strandroth is Traffic Safety Analyst at the Swedish Transport Administration.

indicators for the application can be studied. For the AD application in DriveMe, this means that beyond fatality and injury crashes, property damage only crashes should also be considered when examining the potential and performance of the technology within the selected traffic situations.

Fig. 1 displays VCC's general arrangement for evaluating the performance of crash-avoidance technologies. The set-up requires basic numbers describing the current traffic situation to achieve a clear reference to be used in the traffic safety evaluations. Thus, in Step I, the delimitations in the study are provided and an estimate of the current crash situation provides a reference for the evaluation (Fig. 1). The actual change in traffic situations from cars with crash-avoidance technologies is estimated either in a predictive way by performing simulations of reconstructed crashes (II), or retrospectively, by studying the outcome in real world data (III). These estimations are then compared to the reference from Step I for the performance evaluation in Step IV.

Fig. 1. The VCC general set-up for evaluating crash-avoidance technology benefits. (Details of the Predictive Estimation Methodology, the baseline generation process and the simulation setup, are not considered in this report and are hence omitted in this figure.)

The objective of the present study was to specify basic numbers needed to evaluate the safety performance of automated cars, related to Step I. in the process illustrated in Fig. 1. The means to achieve this objective are to identify safety-relevant conflict situations in the DriveMe targeted environments, and to analyse different Swedish datasets to establish crash rates for the identified conflict situations. Based on this, we can then examine the extent of data required in future research on automated car crash rates.

II. METHODS

To establish target & specification for AD traffic safety evaluation in the DriveMe project, real world traffic data was analysed in order to identify relevant conflict situations. Accordingly, crash rates were defined for different levels of crash severity, and the amount of data mileage needed to confirm a lower crash rate, was estimated based on statistical assumptions.

Databases

Police reported data

The Swedish Transport Accident Data Acquisition (STRADA) is managed by the Swedish Transport Agency (STA) and includes police records and hospital data. In Sweden, crashes on public roads that feature at least one injured person are recorded by the police. Injury levels (fatal, serious, slight) are also assigned by the officer attending the crash scene. The hospital data section of STRADA gathers information from emergency hospitals in Sweden. As there are underreporting of the hospital data and the extent of this is not known, only the police reported data was used for crash rate estimations in the present study.

Insurance claims data

Insurance data from the insurance company If/Volvia were used. A motor insurance for cars covers damages to vehicles and property, as well as injuries to people involved in crashes. Collision Damage insurance pays for vehicle damage to the policyholder's own car, while Third Party Liability insurance covers personal injuries and damage to other vehicles and property. In addition, a unique Swedish concept is the Car Damage Warranty (CDW), valid for the car’s first three years and covering damage to the policyholder’s own car. For all new Volvo cars this CDW is handled by If /Volvia. Different coverages may apply depending on who is at fault. Exposure is measured in insured vehicle years.

Vehicle Inspection Data

All vehicles in Sweden are inspected on a regular basis by the Motor Vehicle Inspection Company. Starting in 1998, all records of driven distances registered at the time vehicles were inspected were saved. Based on this data, Statistics Sweden is able to estimate the traffic mileage for groups of vehicles during a certain year. The estimates cover vehicle mileage of Swedish vehicles irrespective of where they are used.

Traffic flow data

The National Road Database (NVDB) describes the properties of the entire Swedish road network where motor vehicles are allowed to drive. Traffic flow of passenger cars, powered two-wheelers and heavy vehicles is also measured on the road network, and data on traffic flow and road properties from NVDB can be requested from the Swedish Transport Administration. A traffic measurement system for the estimation of annual average daily traffic (AADT) is based on a large number of sample-measurements at randomly selected locations in the State-administered road network. The measurement of AADT and the driving distances provided by Statistics Sweden serve different purposes. While the driving distances show the total mileage driven by the Swedish vehicle fleet per vehicle type, the AADT should be seen as a representative measurement aiming to show the AADT distribution divided by road type. Thus they cannot be merged into a consistent dataset to display vehicle mileage by road and vehicle type.

Data selection and definitions

Conflict situations

Situations in real world traffic safety data are, in general, classified according to *conflict situations*, which describe (only) the movement of the involved road users in relation to each other before the crash or near-crash. A *conflict situation* do not include any information about why the crash occurred or the crash circumstances (e.g. traction lost, driver’s distraction, light conditions, etc.). For illustrative examples, see Table 1.

With the purpose of investigating conflict situations relevant for the DriveMe project, STRADA was analysed in two steps: 1. for the specific route where AD technologies will be restricted to operate; and 2. for DriveMe route-alike traffic environments.

The STRADA web-based data client provides a graphical interface where the road segments in the DriveMe route (Fig. 2) are selected. From this web-service, reports were plotted and exported as aggregated data. For all crashes reported by the police and/or emergency hospitals in the period 2012–2016, injury level and crash situation were investigated. The pre-defined classification of injury level in this STRADA web application based on the Injury Severity Scale (ISS) was used.

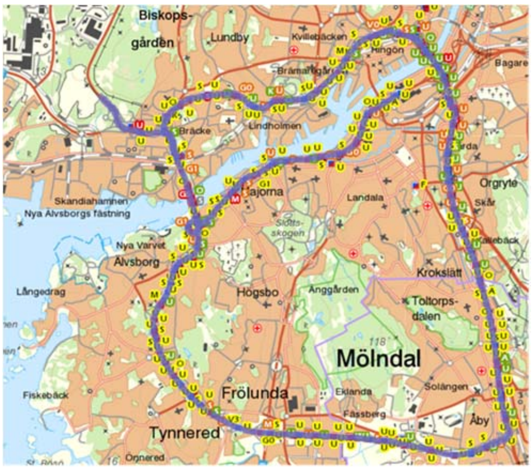

Fig. 2. Road segments in the DriveMe route selected in the STRADA web-based data client.

In addition, DriveMe-route-alike traffic environments were studied in order to examine a larger dataset that possibly would capture other traffic situations for consideration. This was done using a STRADA database extraction retrieved from STA in November 2016. A selection of passenger cars in the period 2010–2015 was used, and the road type ‘Motorway’ was selected to filter out traffic environments according to the DriveMe delimitations, where lanes with oncoming traffic are separated.

As it was not possible to classify conflict situations automatically in the dataset, a review of the text description of the crashes was performed. A large proportion of the crashes was preceded by conflict situations that correspond to the traffic safety evaluation target for the DriveMe project (see Table I).

Crash rates

For the estimation of crash rates, different subsets of data relating to four aspects of the DriveMe AD application were considered: traffic environment; vehicle type; injury level; and conflict situation. In order to do this, different crash and exposure datasets were combined.

Crash datasets: police reported crashes in STRADA for all traffic elements, for passenger cars and for a specific car make and model, the Volvo XC60, were filtered out. Again, road type ‘Motorway’ was used to define AD environments. The insurance claims data provided crashes with the Volvo XC60, MY 2010–2015, in the period 2010–2015 in rear-end and lane-change conflict situations.

Exposure datasets: traffic flow data by road type were selected from a digital platform provided by the Swedish Transport Administration (STA). Roads with median barrier were assumed to represent a DriveMe-alike road environment. This include highways with grade separation in rural and urban areas as well as 2+1 roads [10] with median barriers, and thus does not exactly correspond to the STRADA selection.

Mileages from vehicle inspection data were provided as average numbers of kilometres travelled per car make, model and year. These were aggregated to vehicle-years in traffic for the period 2010–2015; see Fig. 3 for an example.

Insurance claims data hold exposure data in terms of insured vehicle years calculated for the cars involved in selected conflict situations.

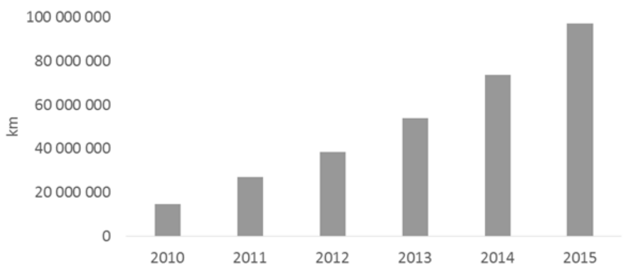

Fig. 3. Vehicle-years in traffic for the Volvo XC60 between 2010 and 2015.

The crash rate *P* was estimated for the different combinations of crash and exposure datasets and was defined as:

$$P = \frac{n}{E} \qquad (1)$$

where *n* is the number of crashes and *E* is either the number of kilometres travelled or insured vehicle years, depending on the dataset used. Using normal approximation to a Poisson distribution, two-sided confidence intervals for the rate were calculated:

$$P \pm z_{1-\frac{\alpha}{2}} \sqrt{\frac{P}{E}} \qquad (2)$$

To define confidence intervals for rate differences, again Poisson distribution and test-based methods were considered [11]:

$$\chi^2 = \frac{\left(n_a - \frac{(n_a + n_b)E_a}{(E_a + E_b)}\right)^2}{\frac{(n_a + n_b)E_a E_b}{(E_a + E_b)^2}}$$

Then, RD confidence limits were calculated by:

$$RD = RD \pm z_{1-\frac{\alpha}{2}} \sqrt{\frac{RD^2}{\chi^2}} \tag{3}$$

To investigate if a hypothetical crash rate P_{hyp} is lower than an actual crash rate with significance level α, the upper confidence bound can be estimated:

$$\frac{n + z_{1-\alpha}\sqrt{n}}{E} < P_{hyp}$$

Then, to determine kilometres needed to be driven in order to confirm this, we can solve it with:

$$E = P_{hyp}\left(\frac{z_{1-\alpha}}{P - P_{hyp}}\right)^2 \tag{4}$$

III. RESULTS

Conflict situations

In total, 802 crashes during 2012–2016 were available in the AD route in DriveMe, with crashes reported by the police and/or emergency hospitals. Two of the crashes involved fatally injured persons, though these are not included in the Swedish official statistics. (Since 2010, official roadway crash suicide figures in Sweden are handled separately. This approach was an adaptation of the definition of accidents to what applies for other types of traffic (aviation, railroads, maritime) as well as to the practice in most other European countries.) Crashes involving serious injuries (ISS9+) were reported in 1% of cases, moderate injuries (ISS4-8) in 5%, and minor injuries (ISS1-3) in 87%. Passenger cars were present in 528 of the 802 crashes. The majority of crashes (74%) were with another passenger car, followed by single vehicle crashes (18%), crashes with motorcycles (4%), and crashes with heavy vehicles (3%). Notable among passenger car collision partners was the presence of animals, mopeds, bicyclists and pedestrians, totaling 6% of all crashes in the DriveMe route.

For the DriveMe-route-alike traffic environments, 5% of all crashes including at least a minor injury and 4% of all fatal crashes involving a passenger car took place on motorways. The distribution of crash types according to the STRADA typology is shown in Fig. 4a)–c). The results differ for different selections of crash injury outcome. When considering all injury levels, car following and single vehicle crashes hold the majority; when considering fatal crashes, more than 20% are pedestrian- and oncoming crashes. Further, when looking only at Volvo car models in fatal crashes, the share of pedestrian and oncoming crashes is even more pronounced by 36%. As it was not possible to classify conflict situations automatically in the dataset, a review of the text description of the crashes was performed. A large proportion of the crashes was preceded by conflict situations that correspond to the traffic safety evaluation target for the DriveMe project (see Table I).

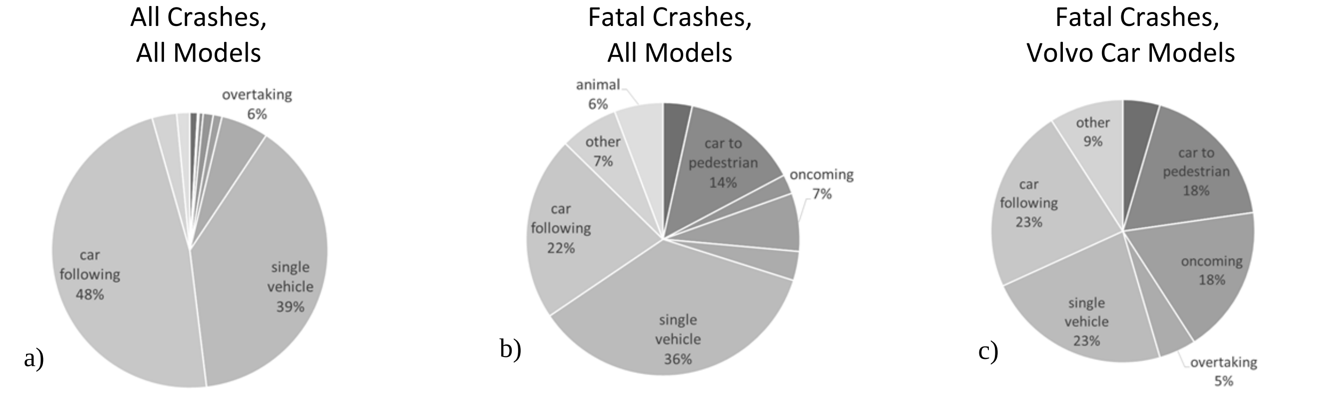

Fig. 4. a)–c). Distribution of crashes with at least one passenger car involved types in motorways, according to the STRADA typology: a) all injury severity level crashes, b) fatal crashes and c) fatal crashes with Volvo car models on motorways. Legends for shares <= 3% were omitted in the figures.

TABLE I
CONFLICT SITUATIONS FOR TRAFFIC SAFETY EVALUATIONS IN THE DRIVEME PROJECT

Road_departure$_{R}$	host vehicle leaves the road to the Right
Road_departure$_{L}$	host vehicle leaves the road to the Left
Oncoming$_{hv_in_opponents_lane}$	host and opponent vehicle travel in Opposite directions (no turning intentions, such as in crossing-related conflict situations), host vehicle enters the opponent vehicles lane
Oncoming$_{hv_in_lane}$	host and opponent vehicle travel in Oncoming directions, opponent vehicle enters host vehicle's lane
Oncoming	host and opponent vehicle travel in Oncoming directions, only one lane/car-width is available in the road layout
SDhv$_{B}$ (RE-F)	host and opponent vehicle travel in the Same Direction in the same lane, host vehicle is Behind the opponent vehicle (Rear-End Frontal)
SDhv$_{F}$ (RE-R)	host and opponent vehicle travel in the Same Direction in the same lane, host vehicle is in Front of the opponent vehicle (Rear-End Rear)
ovSD_lanechange$_{RT}$	host and opponent vehicle travel in the Same Direction in different lanes, opponent vehicle changes to the lane on its Right side
ovSD_lanechange$_{LT}$	host and opponent vehicle travel in the Same Direction in different lanes, opponent vehicle changes to the lane on its Left side
hvSD_lanechange$_{RT}$	host and opponent vehicle travel in the Same Direction in different lanes, host vehicle changes to the lane on its Right side
hvSD_lanechange$_{LT}$	host and opponent vehicle travel in the Same Direction in different lanes, host vehicle changes to the lane on its Left side
Object on Road	host vehicle is moving forward, an object is lying on the road
Falling object	host vehicle is moving forward, an object is falling down onto the road
Car-to-Cyclist, Car-to-Pedestrian, and Car-to-Animal conflict situations:	
all SCP situations	Straight Crossing Path, cyclist/pedestrian/animal from right
SD	host vehicle going Straight, cyclist/pedestrian/animal in Same Direction
Oncoming	host vehicle going straight, cyclist/pedestrian/animal from Opposite direction

Crash rates

Three different combinations of crash and exposure datasets were examined, providing 19 different levels of crash rate estimates (Table II). Rates for different levels of crash consequences were possible to distinguish in STRADA, although property damage only crashes are not available in this dataset, while the opposite is the case for insurance data. The latter includes exposure data for each car make and model, while STRADA is combined with two separate datasets. Data from vehicle inspections allow for information on car makes and models, while data from traffic flow measurements differentiate between road-types. Exemplifying the data, a Volvo XC60 experiences one crash with fatal injury (to any person inside or outside the vehicles involved) per 436 million kilometres driven, based on data from 169,780 vehicle-years in traffic. The corresponding travel distance for a crash causing severe or fatal injury is 85 million kilometres, and for a minor-severe or fatal injury it is 10 million kilometres. Passenger car driving in DriveMe-specific traffic environments represents ~40% of all kilometres travelled, according to traffic flow data. Crash rates based on insurance data were specified for four of the relevant conflict situations identified. This allows traffic safety performance evaluation studies of critical situations to be sited within an overall population of crashes.

TABLE II

CRASH RATE ESTIMATES BASED ON VARIOUS SWEDISH DATASETS

Sample	Data	Exposure	Crash outcome	RateRate [cl]
Passenger cars Accident year 2010–2015	Crashes from STRADA, Exposure based on data from vehicle inspections	380,649,033,034 km by 30,739,089 car-years in traffic	fatal injury	2.87E-09 [2.70E-09, 3.04E-09]
			severe or fatal injury	2.98E-08 [2.92E-08, 3.03E-08]
			minor, severe or fatal injury	2.01E-07 [2,00E-07, 2.03E-07]
Volvo XC60 Accident year 2010–2015		3,052,073,344 km by 169,780 car-years in traffic	fatal injury	2.30E-09 [5.94E-10, 3.99E-09]
			severe or fatal injury	1.18E-09 [7.94E-09, 1.56E-08]
			minor, severe or fatal injury	9.80E-08 [8.69E-08, 1.09E-07]
All motor vehicles Accident year 2014	Crashes from STRADA, exposure based on data traffic flow measurements	79,215,000 000 km	fatal injury	3.22E-09[2.82E-09, 3.61E-09]
			severe or fatal injury	2.48E-08[2.37E-08, 2.59E-08]
			minor, severe or fatal injury	1.36E-07[1.34E-07, 1.39E-07]
Passenger cars Accident year 2014		63,812,684 470 km	fatal injury	2.59E-09[2.19E-09, 2.98E-09]
			severe or fatal injury	2.31E-08[2.19E-08, 2.43E-08]
			minor, severe or fatal injury	1.43E-07[1.40E-07, 1.45E-07]
Passenger cars in AD road type Accident year 2014		25,337,000,000 km	fatal injury	7.50E-10[4.14E-10, 1.09E-09]
			severe or fatal injury	5.88E-09[4.94E-09, 6.82E-09]
			minor, severe or fatal injury	4.34E-08[4.08E-08, 4.60E-08]
Volvo XC60 MY2010–2015 *in rear-end frontal crashes* Accident year 2010–2015	Crashes and exposure from insurance data	71,062 insured vehicle years	all property damage and injury crashes	4.40E-03 [3.92E-03, 4.89E-03]
Volvo XC60 MY2010–2015 *in rear-end rear crashes* Accident year 2010–2015		114,589 insured vehicle years		8.74E-03 [8.19E-03, 9.28E-03]
Volvo XC60 MY2010–2015 *in lane change crashes where host vehicle changes lane* Accident year 2010–2015		71,062 insured vehicle years		1.94E-03 [1.62E-03, 2.27E-03]
Volvo XC60 MY2010–2015 *in lane change crashes where opponent changes lane* Accident year 2010–2015		114,589 insured vehicle years		2.13E-03 [1.86E-03, 2.40E-03]

Crash rates were compared for different samples (Table II). Significant differences, indicating the need for separate references in future safety performance evaluations in these situations, were found when comparing:

- all motor vehicle crashes vs passenger car crashes, for all injury levels;
- passenger car crashes in all traffic environments vs in DriveMe AD road types, for all injury levels;
- passenger car crashes vs Volvo XC60 crashes, for all injury levels except fatal-only crashes;
- the striking vs struck vehicle in both lane-change and rear-end crashes.

The low number of crashes involving fatalities (n=7) in or around the Volvo XC60 is partly related to the relatively small amount of kilometres driven during the period studied (3,052 million km), thus no difference was found for passenger car fatal crashes vs Volvo XC60 fatal crashes. However, the trend indicates that this comparison will also connect to the other significant differences once more data are available.

TABLE III
CRASH RATE DIFFERENCE 95% CONFIDENCE LIMITS

	CI 95%
all motor vehicles **vs** passenger cars, all traffic environments, fatal injury	[6.84E-11, 1.20E-09]
all motor vehicles **vs** passenger cars, all traffic environments, severe+fatal injury	[1.28E-10, 3.36E-09]
all motor vehicles **vs** passenger cars, all traffic environments, minor+severe+fatal injury	[2.40E-09, 1.02E-08]
all passenger cars, in all traffic environments **vs** in DriveMe AD road types, fatal injury	[6,67E-10, 3.00E-09]
all passenger cars, in all traffic environments **vs** in DriveMe AD road types, severe+fatal injury	[1.37E-08, 2.07E-08]
all passenger cars, in all traffic environments **vs** in DriveMe AD road types, minor+severe+fatal injury	[1.08E-07, 9.04E-08]
all passenger cars **vs** Volvo XC60, fatal injury	[-1.33E-09, 2.48E-09]
all passenger cars **vs** Volvo XC60, severe+fatal injury	[1.19E-08, 2.41E-08]
all passenger cars **vs** Volvo XC60, minor+severe+fatal injury	[8.76E-08, 1.20E-07]
Volvo XC60, all traffic environments, all crash outcomes incl. property damage only crashes, rear-end frontal crashes **vs** in in rear-end rear crashes	[3.54, 5.12]
Volvo XC60, all traffic environments, all crash outcomes incl. property damage only crashes lane change crashes where the host vehicle is changing lane **vs** where the opponent is changing lane	[0.18741, 0.18735]

The kilometres needed to be driven in order to confirm a lower crash rate can be estimated using the upper confidence level for a selected significance level. This was done for some of the rate estimates given in Table I, and Fig. 5 shows the number of kilometres needed, with 95% and 90% confidence, to launch a 50% lower rate for crashes involving at least minor injuries. For example, 7 billion passenger car kilometres in DriveMe AD road types are required to confirm reduced severe to fatal crashes with 95% confidence. The only ways to decrease these driving distances to achieve traffic safety performance results in terms of reduced crash rates, is to either choose another safety indicator, e.g. all crashes involving minor or no injuries, or lower the estimate precision. Estimating a 50% lower crash rate with 90% confidence for passenger cars in DriveMe AD road types requires 76 million kilometres.

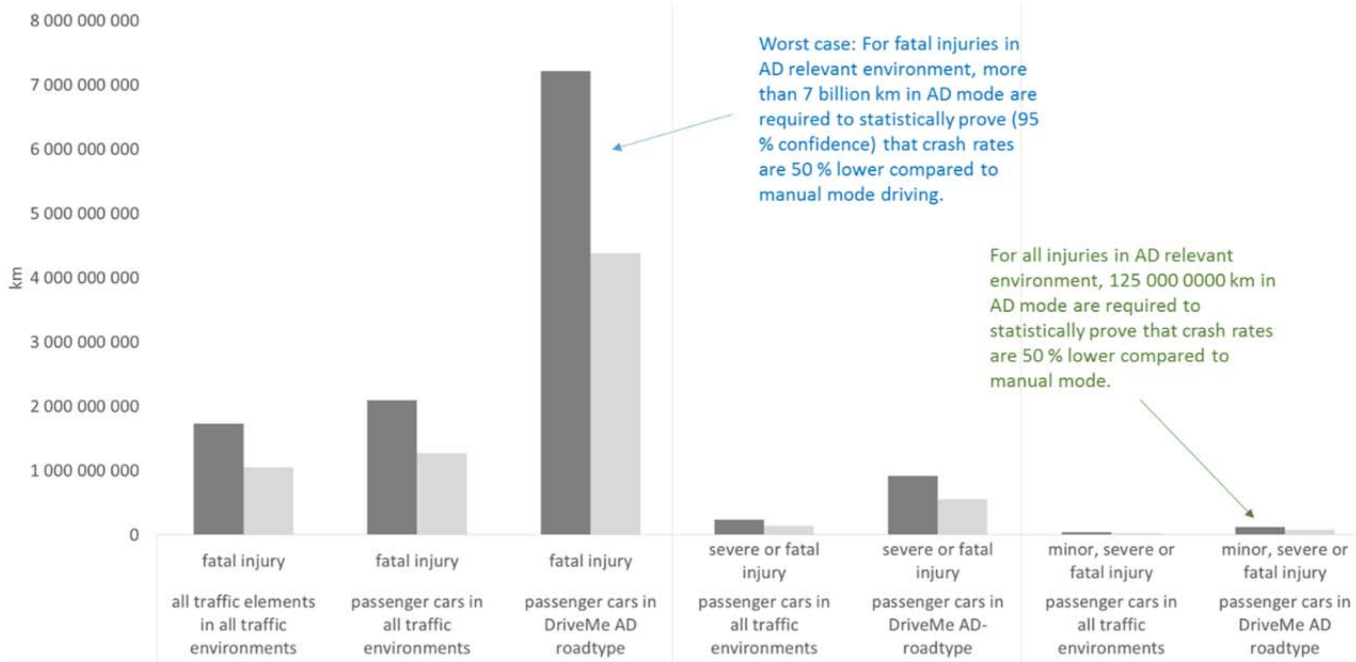

	fatal injury all traffic elements in all traffic environments	fatal injury passenger cars in all traffic environments	fatal injury passenger cars in DriveMe AD roadtype	severe or fatal injury passenger cars in all traffic environments	severe or fatal injury passenger cars in DriveMe AD-roadtype	minor, severe or fatal injury passenger cars in all traffic environments	minor, severe or fatal injury passenger cars in DriveMe AD roadtype
95% confidence	1 735 690 448	2 093 075 388	7 217 111 097	234 299 484	920 302 757	37 976 406	124 659 192
90% confidence	1 054 184 240	1 271 244 587	4 383 365 009	142 303 499	558 952 585	23 065 247	75 712 668

Fig. 5. Number of kilometres needed, with 95% and 90% confidence, to achieve a 50% lower rate for crashes.

IV. Discussion

From reports on ongoing efforts in the field of Automated Driving, it can be deduced that current AD approaches employ specific limitations in their operational scope (regarding road types, driving conditions, etc.), usually in order to simplify the complex technology development process. Comparing AD performance in these very specific driving contexts to manual driving performance in general should, then, provide a skewed comparison. Instead, the reference against which AD performance is compared, should be made as specific and as precise as possible, to ensure results are generalisable. To take a rather extreme but illustrative example, if the AD application addresses parking situations, the reference should be something very different in terms of both conflict situations considered and population size needed for real world safety performance evaluation, compared to an AD application for highway driving.

The aim of the current study was to take a step towards greater specificity. By adding probable crash rates for the very specific traffic environment, car make and model, which are in the scope of the DriveMe project [9], it is clear that the comparatively generic analysis in [7-8] should be further refined and specified.

The study highlights several important aspects that need to be taken into account in future, related work. First, it is clear that while some of the factors required to precisely define the reference dataset are quite accessible, others are less so. For example, traffic environment (roads where lanes with oncoming traffic are separated) was possible to isolate in one combination of datasets. Unfortunately, basic classifications of vehicle type, make and model was not accessible the same data.

Similar problems as found when evaluating the effect from optionally mounted crash avoidance systems will remain for AD safety estimations. In the case of AD cars, are the characteristics of car buyers that are attracted to purchase these vehicles comparable drivers in the general crash population? And will factors as driver complacency and vigilance change over time and develop with use of the technology - and may this introduce further factors needed to consider in in the safety evaluation? Likewise, temporary driver related pre-crash factors such as fatigue and distraction were not available for combination of crash and exposure data.

Other factors, however, represent new problems that AD performance assessment bring to the table. One is whether the reference with which the AD car performance is compared should include obviously illegal driver behaviour in manual driving, such as drunk driving or speeding. The police often report if a driver is suspected of being under the influence of alcohol, so these cases could probably be accounted for, but for speeding the information is generally missing. Excluding the latter because the data is not there would put the AD cars at a disadvantage, since they are proposed to be completely law-abiding. As the data, in any case, is not available in any exposure dataset, how should this be handled? This is a topic for further discussion in the safety performance estimation community.

Another issue is the presence of driver support systems in the reference crash group. Should AD performance be compared against manual driving only (with no driver assistance systems engaged), or should crashes which occur under assisted driving also be included? The datasets used in this study include cars with various fitment levels of safety and comfort systems, such as AEB (Auto Emergency Braking) and ACC (Automatic Cruise Control). While fitment rates probably can be controlled for by using the insurance data [3-4], and given that access to vehicle identification numbers in police reported data is provided, it seems much more difficult to ascertain whether, for example, ACC or Cruise Control was engaged when the crash occurred. This also needs to be discussed further.

As for all performance estimations, the current study also highlights the fact that as analysis precision increases, the data amounts required for statistical certainty increase along with it. Moreover, that increase is a factor that affects the existing performance evaluation activity in a significant way. As the results above clearly illustrate, AD cars would have to accumulate extreme mileages in real world driving if the intent is to verify their level of safety compared to manual driving. Under Swedish conditions, the population size required for a corresponding investigation is even larger, i.e. 7,220 million kilometres would be necessary to establish whether a similar decrease in fatal crashes involving passenger cars on motorways can be achieved with AD technology, although factors such as car make and model were not accounted for in this example.

Even if lower severity level crashes are included, this problem persists. A relevant reference for the DriveMe project would be crashes that also include minor injuries. Based on data from police reports and traffic flow measurements, the current study estimates that 125 million passenger car kilometres on motorways are required to establish whether a 50% lower rate for crashes involving at least minor injuries can be achieved in AD mode. To set this in context, this mileage exceeds the total mileage – in all traffic environments – travelled by all XC60 cars in Sweden during 2015 (Fig. 3). In practice, such extreme mileages are unrealistic to collect for any vehicle manufacturer for experimental testing purposes. It follows that the evaluation of AD safety performance against manual driving calls for alternative methods that can provide status indications during function development.

The only way to do this, is to perform predictive studies based on real world traffic situations by using simulations (Fig. 1). In this approach, a clear target for the study is first agreed (such as "all crashes including property damage only crashes, in roads where no VRUs are allowed and where lanes with oncoming traffic are separated"). This target defines the relevant conflict situations, as was the aim of the present study. Next, a baseline consisting of reconstructed real world traffic situations is established and then used in simulations with and without the crash-preventing technology. The output describes the, possibly altered, traffic situation outcome with the vehicle under test.

An advantage of using simulations is that near-crashes can be made part of the baseline in addition to crashes. Near-crashes, where a driver's manoeuvre contributed to avoiding the crash, are interesting when compared with the performance of an AD car equipped with other perception, situation interpretation and action capabilities than a human driver. Near-crash baselines could, for example, be generated based on real world traffic situations in naturalistic driving study (NDS) or field operational test (FOT) datasets.

It should also be emphasised that while this type of predictive, simulation-based analysis is quite established as a methodology [12], the application to AD functions generates some new challenges not previously encountered. Most importantly, since the AD car continuously evaluates and adapts to the traffic situation, both the number, the detail level and the length of the reconstructed traffic situations need to increase, as well as the complexity of the function representation. This in turn puts higher demands on the simulation tools and methods used.

Additionally, new classes of situations need to be added to the simulations that target the interaction between the AD technology and the driver. The first one is related to driver capabilities. It has been speculated both that the transition from automated to manual control can have negative safety implications and that driver skill degradation may occur if AD mode is engaged a significant portion of the total driving time. The latter is related to injury causation: when the driver is not supposed to be manoeuvring the car, he might want to sit in a more relaxed position. These possible new seating positions have to be accounted for, as they may or may not challenge the performance of the passive safety systems. Thus, it should be noted that simulation can never be a complete replacement for real world performance estimation. Rather, good simulations depend heavily on having a strong empirical background for the data on which the simulations are performed. The need for accruing as much mileage of real world driving in AD mode still stands, as that will be the only way to anchor parametrisations (in the simulations) to new real world situations with AD.

Although being a common conclusion in most safety performance real world assessments, a key question is: in what way should the evaluation methods be improved? Details needed in both crash and exposure data have already been discussed. Also, different analysis approaches could be considered. E.g., the induced exposure method [13] might be used to establish a reference for relative risk difference estimations for AD cars and human driven cars. Future attempts to apply this, or other methods, may show whether it is suitable for this purpose, considering factors like the characteristics of AD car buyers and drivers in the general crash population.

Another important question is: given the introduction of assisted driving in general and automated driving in particular, what safety indicators will be relevant to assess crash-avoidance technology performance? Normally, it is stated that crashes and personal injuries are relevant traffic safety indicators, but even these indisputable metrics need further discussion. In [7], the correlation of incidents in the SHRP2 data to police reported crashes is discussed. Crashes reported to an insurance company is a much more interesting alignment, since these cover all crashes leading to at least property damage. Insurance claims also covers crashes with vulnerable road users, but does not consider vehicle impacts, such as curb strikes, where the car is neither damaged nor comes to a stop. Given recent data-collection procedures, where all levels of incidents are possible to record, a highly relevant question is: what constitutes a relevant crash?

The results of the current study indicate that crash rates to be used as a future reference in traffic safety evaluations must be based on data with comparable vehicle makes and models in relevant traffic environments. However, no single combination of datasets was available to provide necessary input with information on all crashes (including property damage only crashes), on crash rates in specific traffic environments, different injury severity levels and for specific conflict situations. It follows that both crash-and exposure data collection need to be improved if they are to be able to provide input to future traffic safety analysis of self-driving cars.

V. CONCLUSIONS

By identifying conflict situations targeted for the DriveMe project and establishing corresponding crash rates for different segments of crashes, the results of this study provide necessary reference data for traffic safety predictions and follow-up studies on AD cars. It was determined that various classifications of data are needed to provide relevant references, and that no single dataset was available to provide all the variables for this purpose.

In all, the present study demonstrates a fresh set-up on the basic numbers needed to evaluate traffic safety performance of AD cars, as well as a discussion on the methodology to achieve these numbers. The results also contribute essential insights to AD stakeholders on the expectations on future traffic safety evaluations of AD safety performance.

VI. ACKNOWLEDGEMENT

We thank Lotta Jakobsson, Alix Edwards, Mikael Ljung Aust, Trent Victor, Emma Tivesten, Anders Eugensson, Marcus Rothoff and Bo Svanberg and for comments that greatly improved the manuscript.

VII. REFERENCES

[1] OECD/ITF (2016) Road Safety Annual Report 2016. *OECD Publishing*, Paris.

[2] Eriksson, O. (2014) Accident risk at local unevenness or other road surface defects. VTI rapport 811.

[3] Isaksson-Hellman, I., Lindman, M. (2016) Evaluation of the crash mitigation effect of low-speed automated emergency braking systems based on insurance claims data. *Traffic Injury Prevention*, 17: pp. 42–7.

[4] Isaksson-Hellman, I., Lindman, M. (2015) Evaluation of Rear-End Collision Avoidance Technologies based on Real World Crash Data. *Proceedings of the 3rd International Symposium on Future Active Safety Technology Toward zero traffic accidents*, 2015, Gothenburg.

[5] HLDI Bulletin (2015) Volvo City Safety loss experience — a long term update. *HLDI Bulletin*, 32(1).

[6] Doyle, M., Edwards, A., Avery, M. (2015) AEB real world validation using UK motor insurance claims data. *24th International Technical Conference on the Enhanced Safety of Vehicles (ESV)*, 2015, Gothenburg.

[7] Blanco, M., *et al*. (2016) Automated vehicle crash rate comparison using naturalistic data. VTTI report.

[8] Kalra, N., Paddock, S. M. (2016) Driving to Safety: How Many Miles of Driving Would It Take to Demonstrate Autonomous Vehicle Reliability? www.rand.org/t/RR1478

[9] Victor, T., *et al.* (2017) When Autonomous Vehicles Are Introduced on a Larger Scale in the Road Transport System: The Drive Me Project. In: *Automated Driving*. *Springer International Publishing*, Switzerland, pp.541–6.

[10] Vadeby, A. (2016) Traffic Safety Effects of Narrow 2+1 Roads with Median Barrier in Sewden. *17th International Conference Road Safety On Five Continents.* Rio de Janeiro.

[11] Sahai, H., Khurshid, A. (1996) Statistics in epidemiology: methods, techniques, and applications. *CRC Press*, Boca Raton.

[12] Page, Y., *et al*. (2015) A Comprehensive and Harmonized Method for Assessing the Effectiveness of advanced driver assistance systems by virtual simulation: the P.E.A.R.S. initiative. *24th International Technical Conference on the Enhanced Safety of Vehicles (ESV)*, 2015, Gothenburg, pp.1–12.

[13] Evans, L. (1998) Antilock Brake Systems and Risk of Different Types of Crashes in Traffic. *ESV Conf paper no 98-S2-O-12.* Windsor.